

EFFICIENCY IMPROVEMENTS IN CONSTRUCTING PSEUDORANDOM GENERATORS FROM ONE-WAY FUNCTIONS*

IFTACH HAITNER[†], OMER REINGOLD[‡], AND SALIL VADHAN[§]

Abstract. We give a new construction of pseudorandom generators from any one-way function. The construction achieves better parameters and is simpler than that given in the seminal work of Håstad, Impagliazzo, Levin, and Luby [*SIAM J. Comput.*, 28 (1999), pp. 1364–1396]. The key to our construction is a new notion of *next-block pseudoentropy*, which is inspired by the notion of “inaccessible entropy” recently introduced in [I. Haitner, O. Reingold, S. Vadhan, and H. Wee, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009, pp. 611–620]. An additional advantage over previous constructions is that our pseudorandom generators are parallelizable and invoke the one-way function in a nonadaptive manner. Using [B. Applebaum, Y. Ishai, and E. Kushilevitz, *SIAM J. Comput.*, 36 (2006), pp. 845–888], this implies the existence of pseudorandom generators in NC^0 based on the existence of one-way functions in NC^1 .

Key words. one-way functions, pseudorandom generators, pseudoentropy

AMS subject classifications. 94A60, 68Q01

DOI. 10.1137/100814421

1. Introduction. The result of Håstad et al. [18] that one-way functions imply pseudorandom generators is one of the centerpieces of the foundations of cryptography and the theory of pseudorandomness.

From the perspective of cryptography, it shows that a very powerful and useful cryptographic primitive (namely, pseudorandom generators) can be constructed from the minimal assumption for complexity-based cryptography (namely, one-way functions). With this starting point, numerous other cryptographic primitives can also be constructed from one-way functions, such as private-key cryptography [6, 27], bit-commitment schemes [28], zero-knowledge proofs for NP [9], and identification schemes [4].

From the perspective of pseudorandomness, their result provides strong evidence that pseudorandom bits can be generated very efficiently, with smaller computational resources than the “distinguishers” to whom the bits should look random. Such kinds of pseudorandom generators are needed, for example, for hardness results in learning [35] and the natural proofs barrier for circuit lower bounds [29]. Moreover, the paper of Håstad et al. [18] introduced concepts and techniques that now permeate the theory of pseudorandomness, such as pseudoentropy and the leftover hash lemma.

A drawback of the construction of Håstad et al. [18], however, is that it is quite complicated. While it utilizes many elegant ideas and notions, the final construction

*Received by the editors November 10, 2010; accepted for publication (in revised form) August 1, 2012; published electronically June 25, 2013. Versions of this paper have previously appeared in STOC 2010 [16] and on ECCO [15].

<http://www.siam.org/journals/sicomp/42-3/81442.html>

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (iftachh@cs.tau.ac.il). This author’s research was supported by US-Israel BSF grant 2010196. Part of this research was conducted while this author was at Microsoft Research, New England Campus.

[‡]Microsoft Research, Silicon Valley Campus, Mountain View, CA 94043, and Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel (omreing@microsoft.com). This author’s research was supported by US-Israel BSF grant 2006060.

[§]School of Engineering and Applied Sciences and Center for Research on Computation and Society, Harvard University, Cambridge, MA 02138 (salil@seas.harvard.edu). This author’s research was supported by NSF grant CNS-0831289 and by US-Israel BSF grants 2006060, 2010196.

combines these in a rather ad hoc and indirect fashion due to various technical issues. In addition to being less satisfactory from an aesthetic and pedagogical perspective, the complexity of the construction also has a significant impact on its efficiency. Indeed, it is too inefficient to be implemented even for very modest settings of parameters.

In the last few years, progress has been made on simplifying the construction of Håstad et al. (see [20]) and improving its efficiency (see [13]). These constructions, however, still retain the overall structure of the Håstad et al. construction and thus retain some of the complex and ad hoc elements.

In this paper, we present a significantly more direct and efficient construction of pseudorandom generators from one-way functions. The key to our construction is a new notion of *next-block pseudoentropy*, which is inspired by the recently introduced notion of “inaccessible entropy” [17].

1.1. Håstad–Impagliazzo–Levin–Luby construction. Informally, a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a *one-way function* (OWF) if it is easy to compute (in polynomial time) and hard to invert even on random inputs. (See section 2 for formal definitions.) A polynomial-time computable function $G : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ is a *pseudorandom generator* (PRG) if it is stretching (i.e., $m(n) > n$) and its output distribution is pseudorandom (i.e., $G(U_n)$ is computationally indistinguishable from $U_{m(n)}$). The theorem of Håstad et al. [18] relates these notions.

THEOREM 1.1. *If there exists a one-way function, then there exists a pseudorandom generator.*

The key notion underlying their construction is the following generalization of pseudorandomness.

DEFINITION 1.2 (pseudoentropy, informal). *A random variable X has pseudoentropy at least k if there exists a random variable Y such that*

1. X is computationally indistinguishable from Y ;
2. $H(Y) \geq k$, where $H(\cdot)$ denotes Shannon entropy.¹

A pseudoentropy generator (PEG)² is a polynomial-time computable function $G : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ such that $X = G(U_n)$ has pseudoentropy at least $H(G(U_n)) + \Delta(n)$ for some $\Delta(n) \geq 1/\text{poly}(n)$. We refer to $\Delta(n)$ as the entropy gap of G .

That every pseudorandom generator $G : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ is a pseudoentropy generator can be seen by taking $Y = U_{m(n)}$ and noting that $H(Y) = m(n)$, but $H(G(U_n)) \leq H(U_n) = n$. Pseudoentropy generators are weaker in that Y may be very far from uniform and may even have $H(Y) < n$ (as long as $H(G(U_n))$ is even smaller).

The construction of pseudorandom generators from one-way functions proceeds roughly in the following steps:

OWF to PEG: Given a one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, Håstad et al. [18] define $\text{PEG}(x, h, i) = (f(x), h, h(x)_{1,\dots,i})$, where h is an appropriate hash function and $h(x)_{1,\dots,i}$ denotes the first i bits of $h(x)$. PEG can be shown to be a pseudoentropy generator with an entropy gap of roughly $\Delta = \log n/n$: Whenever $i = \log |f^{-1}(x)| + \Theta(\log n)$ (which happens with probability $\Theta(\log n/n)$), the first $\approx \log |f^{-1}(x)|$ bits of $h(x)$ extract all the entropy

¹The Shannon entropy of a random variable X is defined to be $E_{x \leftarrow X}[\log(1/\Pr[X = x])]$.

²Håstad et al. [18] refer to such a generator as a *false entropy generator* and require a pseudoentropy generator to have output pseudoentropy (at least) $n + \Delta(n)$, rather than just $H(G(U_n)) + \Delta(n)$. For the informal discussion here, however, we prefer not to introduce the additional term “false entropy.”

of x , and then we get $\Theta(\log n)$ bits of pseudoentropy by the Goldreich–Levin hardcore-bit theorem [8].

Converting Shannon entropy to min-entropy and amplifying the gap:

Next, Håstad et al. use a direct product construction $\text{PEG}'(x_1, \dots, x_t) = (\text{PEG}(x_1), \dots, \text{PEG}(x_t))$ to convert pseudoentropy into pseudo-*min*-entropy and increase the entropy gap to be $\omega(\log n)$. This turns out to require taking $t = \tilde{O}(n/\Delta)^2$ copies.

Randomness extraction: By hashing, Håstad et al. extract pseudorandom bits from the pseudo-min-entropy achieved so far. By also hashing the seed x to extract any remaining entropy, they obtain a pseudorandom generator. Specifically, they show that $G(x, h_1, h_2) = (h_1, h_2, h_1(\text{PEG}'(x)), h_2(x))$ is a pseudorandom generator if the output lengths of h_1 and h_2 are chosen appropriately. The choice of output lengths depends on the amount of min-entropy in the output of PEG' , which in turn depends on the amount of entropy in the output of PEG . Unfortunately, these quantities may be infeasible to compute; this is handled by the next step.

Enumeration: Håstad et al. enumerate over all $u = O(n/\Delta)$ possible values k for the output entropy of PEG (up to an accuracy of $\Delta/2$), construct a pseudorandom generator G_k for each, use composition to make each G_k stretch its seed by a factor greater than u , and then take $G(x_1, \dots, x_u) = G_1(x_1) \oplus \dots \oplus G_u(x_u)$ as their final pseudorandom generator.

The total seed length in this informal description is $n \cdot t \cdot u = \tilde{O}(n^4/\Delta^3) = \tilde{O}(n^7)$. In fact, we have been cheating a bit in order to present the construction in a more modular way than in [18]. (The issues we ignored have to do with handling uniform adversaries, for which the (non)samplability of source Y in Definition 1.2 is an issue.) The actual seed length in the main construction presented in [18] is of $O(n^{10})$ (and the construction involves additional complications). A construction of seed length $O(n^8)$ is outlined in [18] and has been formalized and proved in [20].

Above we see three main sources of inefficiency in the construction: (1) the entropy gap Δ being fairly small, (2) the conversion of Shannon entropy to min-entropy, and (3) enumerating guesses for the output entropy of the initial pseudoentropy generator. Haitner, Harnik, and Reingold [13] show how to save a factor of n in the enumeration step (by constructing a pseudoentropy generator in which more is known about how the entropy is distributed) to obtain a seed length of $O(n^7)$, but still all of the steps remain.

A further complication in the construction of Håstad et al. [18] is that the reductions demonstrating the correctness of the construction are much more complex for uniform adversaries. This aspect of the proof has recently been simplified and made much more modular via Holenstein’s uniform hardcore lemma [19, 20].

In case the one-way function is secure against exponential running time ($2^{\Omega(n)}$) adversaries, Holenstein [20] showed how to reduce the seed length to $\tilde{O}(n^4)$ (or $O(n^5)$ to obtain a pseudorandom generator with exponential security), which was then improved by Haitner, Harnik, and Reingold [11] to $\tilde{O}(n)$ (or $O(n^2)$ to obtain a pseudorandom generator with exponential security).³

³In more detail, Holenstein’s construction generalizes [18] for one-way functions of “any hardness,” while Haitner, Harnik, and Reingold [11] take a totally different route (based on the “randomized iterate” of a function introduced by Goldreich, Krawczyk, and Luby [7]) and obtain constructions based on exponentially hard one-way functions, as well as on (unknown-)regular one-way functions.

1.2. Our approach. Our construction is based on a generalization of the notion of a pseudoentropy generator. It is similar in spirit to the Blum–Micali notion of next-bit unpredictability [2], which was shown by Yao [36] to be equivalent to his (now standard) definition of pseudorandomness. In one form, this equivalence says that the pseudorandomness of a random variable X is equivalent to each bit of X being indistinguishable from uniform given the previous ones. That is, $X = (X_1, \dots, X_n)$ is computationally indistinguishable from $U_n = (Y_1, \dots, Y_n)$ if and only if for every i , $(X_1, X_2, \dots, X_{i-1}, X_i)$ is computationally indistinguishable from $(X_1, X_2, \dots, X_{i-1}, Y_i)$. It is thus natural to consider what happens if we require not that X_i be pseudorandom given the previous bits, but only that X_i have *pseudoentropy* given the previous bits. More generally, we can allow the X_i 's to be blocks instead of bits.

DEFINITION 1.3 (next-block pseudoentropy, informal). *A random variable $X = (X_1, \dots, X_m)$ has next-block pseudoentropy at least k if there exists a set of random variables $Y = \{Y_1, \dots, Y_m\}$, each jointly distributed with X , such that the following hold:*

1. *For every i , $(X_1, X_2, \dots, X_{i-1}, X_i)$ is computationally indistinguishable from $(X_1, X_2, \dots, X_{i-1}, Y_i)$.*
2. $\sum_i H(Y_i | X_1, \dots, X_{i-1}) \geq k$.

A next-block pseudoentropy generator (*NBPEG*) is a polynomial-time computable function $G : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell m}$ such that $(X_1, \dots, X_m) = G(U_n)$ has next-block pseudoentropy at least $H(G(U_n)) + \Delta(n)$, where again $\Delta(n)$ is called the entropy gap.

That is, in total, the bits of X “look like” they have k bits of entropy given the previous ones. Note that the case of one block ($m = 1$) amounts to the definition of a pseudoentropy generator. Also note that, when $m > 1$, allowing Y to be correlated with X in this definition is essential: For example, if all the blocks of X are always equal to each other (and have noticeable entropy), then there is no way to define Y that is independent of X and satisfies the first condition.

With this notion, our construction proceeds as follows.

OWF to NBPEG: Given a one-way function f , we define $G(x, h) = (f(x), h, h(x)_1, h(x)_2, \dots, h(x)_n)$, where $h : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is an appropriate hash function and $h(x)_i$ is the i th bit of $h(x)$. Note that this is the same as the construction of Håstad et al. [18], except that we do not randomly truncate the output. At first, this seems problematic; by revealing all of $h(x)$, it becomes easy for an adversary to compute x , and thus the pseudoentropy of output equals its real entropy (i.e., we have *zero* entropy gap). We show, however, that it does indeed have *next-block pseudoentropy at least $n + \log n$* , which is even larger than the seed length of G . We have gained in two ways here. First, the entropy gap is now $\Delta = \log n$ instead of $\Delta = \log n/n$. Second, we know the total amount of entropy in the output (though not the amount contributed by the individual blocks). These two advantages improve the complexity and security of the rest of the construction. Furthermore, the fact that the next-block pseudoentropy is larger than the seed length simplifies the construction, as we do not need to extract any additional entropy from the seed.

Entropy equalization: Here we use a technique from [17] to convert our knowledge about the total entropy (summed over all blocks) into knowledge about the entropy in the individual blocks. We evaluate G on $u = O(n/\Delta)$ independent seeds and concatenate the outputs, but randomly shifted by $i \stackrel{R}{\leftarrow} [n]$ coordinates. This increases our seed length and our entropy by a multiplicative factor of approximately u , but now almost all the blocks have pseudoentropy

at least the average pseudoentropy of the blocks of G .

Converting Shannon entropy to min-entropy and amplifying the gap: This works the same as in [18]. Again we take $t = \tilde{O}(n/\Delta)^2$ copies but concatenate them within each block to obtain an m -block generator G' . Now each of the m blocks is indistinguishable from having high *min-entropy* conditioned on the previous ones. Thus, what we have is a computational analogue of a *block source* [3], which is a random source in which each block has high min-entropy conditioned on the previous ones.

Randomness extraction: For this step, we use a known method for block-source extraction [3, 37] and define $G(x, h) = (h, h(G'(x)_1), \dots, h(G'(x)_m))$, where h is a universal hash function. Since we know how much pseudo-min-entropy is in each block, there is no difficulty in choosing the output length of h .

In total, our seed length is $O(n \cdot u \cdot t) = \tilde{O}(n^4)$. For the case of exponentially hard one-way functions, we can obtain $\Delta = \Omega(n)$ and thus achieve seed $\tilde{O}(n)$ matching [11] (but, unlike in [11], our construction uses nonadaptive calls to the one-way function).

In addition, our pseudorandom generator makes $q = u \cdot t$ invocations of the one-way functions and achieves additive stretch $s = \Omega(u \cdot t \cdot \Delta) = \Omega(q \cdot \log n)$. The ratio $s/q = \Omega(\log n)$ is optimal for black-box constructions, as shown by Gennaro et al. [5], and was not achieved by any previous construction from general one-way functions.

Note that our construction involves no “guessing” of entropies, neither in the construction of the initial NBPEG G , nor in an enumeration step at the end. While the entropy equalization “costs” the same (namely, $u = O(n/\Delta)$) as enumeration did, it is actually doing more for us. Enumeration handled our lack of knowledge of a single entropy value (for which there were only $O(n/\Delta)$ choices), but here equalization is handling lack of knowledge for approximately n entropy values (one for each block), for which there are exponentially many choices. Moreover, enumeration required composing the pseudorandom generators to increase their stretch, resulting in a construction that is highly sequential and makes adaptive use of the one-way function. Our pseudorandom generators make nonadaptive use of the one-way function and are parallelizable (e.g., in NC^1) for getting pseudorandom generators with small stretch. Using Applebaum, Ishai, and Kushilevitz [1], this implies the existence of pseudorandom generators in NC^0 based on the existence of one-way functions in NC^1 .

1.3. Relation to inaccessible entropy. The notion of next-block pseudoentropy generators was inspired by the notion of *inaccessible entropy generators* in [17]. These are generators G that also produce m blocks (x_1, \dots, x_m) with the property that it is infeasible for an adversary to *generate* a sequence of blocks (x_1, \dots, x_m) that are consistent with the output of G in such a way that entropy of the individual blocks x_i is high (conditioned on the state of the adversary after generating the previous blocks). Thus, in a computational sense, the output of G has *low entropy*. For this reason, the notions of next-block pseudoentropy generators and inaccessible entropy generators seem to be dual to each other.

The initial construction of an inaccessible entropy generator in [17] is $G(x) = (f(x)_1, \dots, f(x)_n, x)$, which is very similar to our construction of a next-block pseudoentropy generator except that there is no hashing and the bits of $f(x)$ instead of $h(x)$ are treated as separate blocks. This initial step is followed by entropy equalization and gap amplification steps that are exactly the same as those we use (but are analyzed with respect to dual notions). The final hashing step in [17] (to construct statistically hiding commitment schemes) is more complex than ours and is necessarily interactive.

Interestingly, the notion of an inaccessible entropy generator was introduced in an attempt to make the construction of statistically hiding commitment schemes from one-way functions “as simple” as the construction of pseudorandom generators from one-way functions, via manipulating notions of computational entropy. (The previous construction, from [14], was extremely complex.) In return, that effort has now inspired our simplifications and improvements to the construction of pseudorandom generators.

1.4. Subsequent work. A subsequent paper by Vadhan and Zheng [34] improves our results in two ways. First, they remove the hashing from our construction of a next-bit pseudoentropy generator. Specifically, they show that if $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a one-way function, then $G(x) = (f(x), x_1, \dots, x_n)$ is already a next-block pseudoentropy generator (as we conjectured in oral presentations of our work). Second, they show how to avoid the seed-length blow-up due to entropy equalization and thereby reduce the seed length from $\tilde{O}(n^4)$ to $\tilde{O}(n^3)$ (at the price of making adaptive calls to the one-way function).

A new paper by Holenstein and Sinha [23] shows that any black-box construction of a pseudorandom generator from a one-way function on n -bit inputs must invoke the one-way function $\Omega(n/\log n)$ times. Their lower bound also applies to regular one-way functions (of unknown regularity) and is tight in this case (due to the constructions of [7, 13]). Our construction from general one-way functions, as well as the improved version in [34], invokes the one-way function $\tilde{O}(n^3)$ times. It remains an open question whether or not the superlinear number of invocations or the superlinear seed length is necessary, or the constructions can be further improved.

1.5. Paper organization. Notation and definitions used through this paper are given in section 2, while the new notion of a next-block pseudoentropy generator is formally defined in section 3. In section 4 we present our construction of a next-block pseudoentropy generator from one-way functions, while in section 5 we show how to use next-block pseudoentropy generators to get a pseudorandom generator. Finally, in section 6 we use the above reductions to prove the main result of this paper.

2. Preliminaries.

2.1. Random variables. Let X and Y be random variables taking values in a discrete universe \mathcal{U} . We adopt the convention that when the same random variable appears multiple times in an expression, all occurrences refer to the same instantiation. For example, $\Pr[X = X]$ is 1. The *support* of a random variable X is $\text{Supp}(X) := \{x : \Pr[X = x] > 0\}$. We write $x \stackrel{R}{\leftarrow} X$ to indicate that x is selected according to X . Similarly, given a finite set S , we let $s \stackrel{R}{\leftarrow} S$ denote that s is selected according to the uniform distribution on S .

We write $\Delta(X, Y)$ to denote the *statistical difference* (also known as variation distance) between X and Y ; i.e.,

$$\Delta(X, Y) = \max_{T \subseteq \mathcal{U}} |\Pr[X \in T] - \Pr[Y \in T]|.$$

If $\Delta(X, Y) \leq \varepsilon$ (resp., $\Delta(X, Y) > \varepsilon$), we say that X and Y are ε -close (resp., ε -far).

2.2. Entropy measures. We refer to several measures of entropy in this work. The relation and motivation of these measures are best understood by considering a notion that we refer to as the *sample-entropy*: For a random variable X and $x \in$

Supp(X), we define the sample-entropy of x with respect to X to be the quantity

$$H_X(x) := \log(1/\Pr[X = x]).$$

The sample-entropy measures the amount of “randomness” or “surprise” in the specific sample x , assuming that x has been generated according to X . Using this notion, we can define the *Shannon entropy* $H(X)$ and *min-entropy* $H_\infty(X)$ as follows:

$$H(X) := \mathbb{E}_{x \stackrel{R}{\leftarrow} X} [H_X(x)],$$

$$H_\infty(X) := \min_{x \in \text{Supp}(X)} H_X(x).$$

Flattening Shannon entropy. It is well known that the Shannon entropy of a random variable can be converted to min-entropy (up to a small statistical distance) by taking independent copies of this variable.

LEMMA 2.1.

1. Let X be a random variable taking values in a universe \mathcal{U} , let $t \in \mathbb{N}$, and let $\varepsilon > 0$. Then with probability at least $1 - \varepsilon - 2^{-\Omega(t)}$ over $x \stackrel{R}{\leftarrow} X^t$,

$$|H_{X^t}(x) - t \cdot H(X)| \leq O(\sqrt{t \cdot \log(1/\varepsilon)} \cdot \log(|\mathcal{U}| \cdot t)).$$

2. Let X, Y be jointly distributed random variables, where X takes values in a universe \mathcal{U} , let $t \in \mathbb{N}$, and let $\varepsilon > 0$. Then with probability at least $1 - \varepsilon - 2^{-\Omega(t)}$ over $(x, y) \stackrel{R}{\leftarrow} (X^t, Y^t) := (X, Y)^t$,

$$|H_{X^t|Y^t}(x|y) - t \cdot H(X|Y)| \leq O(\sqrt{t \cdot \log(1/\varepsilon)} \cdot \log(|\mathcal{U}| \cdot t)).$$

Proof.

1. For $x = (x_1, \dots, x_t)$, we have $H_{X^t}(x) = \sum_{i=1}^t H_X(x_i)$. Thus, when $x \stackrel{R}{\leftarrow} X^t$, $H_{X^t}(x)$ is the sum of t independent random variables $H_X(x_i)$, and thus we can obtain concentration around the expectation (which is $t \cdot H(X)$) via Chernoff–Hoeffding bounds. These random variables $H_X(x_i)$ are not bounded (as is required to apply the standard Chernoff–Hoeffding bound), but they are unlikely to be much larger than $O(\log |\mathcal{U}|)$. Specifically, for every $\tau > 0$ we have

$$\begin{aligned} \Pr_{x_i \stackrel{R}{\leftarrow} X} [H_X(x_i) \geq \log(|\mathcal{U}|/\tau)] &\leq \sum_{x_i \in \mathcal{U}: H_X(x_i) \geq \log(|\mathcal{U}|/\tau)} \Pr[X = x_i] \\ &\leq |\mathcal{U}| \cdot 2^{-\log(|\mathcal{U}|/\tau)} \\ &= \tau. \end{aligned}$$

A Chernoff bound for random variables with such exponentially vanishing tails follows from [32], and it says that the probability that the sum deviates from the expectation by at least $\Delta \cdot (\log(|\mathcal{U}|/\tau)) + 2\tau t$ is at most $\exp(-\Omega(\Delta^2/t)) + \exp(-\Omega(\tau t))$, provided that $\tau \in [0, 1]$. An appropriate choice of $\Delta = O(\sqrt{t \log(1/\varepsilon)})$ and $\tau = \min\{1, O(\log(1/\varepsilon)/t)\}$ completes the proof.

2. The proof of part 2 follows by similar reasons, noting that $H_{X^t|Y^t}(x|y) = \sum_{i=1}^t H_{X|Y}(x_i|y_i)$. \square

2.3. One-way functions and pseudorandom generators. We recall the standard definitions of one-way functions and pseudorandom generators.

DEFINITION 2.2 (one-way functions). *Let $f: \{0, 1\}^n \mapsto \{0, 1\}^m$ be a polynomial-time computable function, where n is a security parameter and $m = m(n)$. For $T = T(n)$ and $\varepsilon = \varepsilon(n)$, we say that f is a (T, ε) -one-way function if for every probabilistic algorithm A running in time T and all sufficiently large n , we have*

$$\Pr[A(Y) \in f^{-1}(Y)] \leq \varepsilon,$$

where the probability is taken over $Y = f(U_n)$ and the coin tosses of A . We say that f is a one-way function if it is a $(p(n), 1/p(n))$ -one-way function for every polynomial p .

DEFINITION 2.3 (pseudorandom generators). *Let X be a random variable, depending on a security parameter n and taking values in $\{0, 1\}^m$ for $m = m(n)$. For $T = T(n)$ and $\varepsilon = \varepsilon(n)$, we say that X is (T, ε) -pseudorandom if for every probabilistic distinguisher D running in time T and all sufficiently large n , we have*

$$|\Pr[D(X) = 1] - \Pr[D(U_m) = 1]| \leq \varepsilon.$$

A polynomial-time computable function $G: \{0, 1\}^n \mapsto \{0, 1\}^m$ with $m = m(n) > n$ is a (T, ε) -pseudorandom generator if $G(U_n)$ is (T, ε) -pseudorandom.

We say that X is pseudorandom if it is $(p(n), 1/p(n))$ -pseudorandom for every polynomial p . Similarly, G is a pseudorandom generator if $G(U_n)$ is pseudorandom.

3. Next-block pseudoentropy. In this section we formally define the new notion of next-block pseudoentropy for the cases of both Shannon entropy and min-entropy. The definitions differ from the informal definition given in the introduction (Definition 1.3) in the following two ways, both of which are important for the treatment of uniform adversaries:

- We require indistinguishability even against algorithms that have an oracle for sampling from the joint distribution (X, Y_i) . (This enables us to show, using a hybrid argument, that pseudoentropy increases when we are taking many independent copies of X . In the case of nonuniform adversaries, no oracle for sampling from (X, Y_i) is needed, as the samples can be nonuniformly hardwired into the adversary.)
- In order to achieve the first item, we allow the random variables Y_i to depend on the distinguisher.

Similar issues arise for treating uniform adversaries with standard pseudoentropy.

DEFINITION 3.1 (next-block (Shannon) pseudoentropy). *Let X be a random variable taking values in \mathcal{U}^m , where X , \mathcal{U} , and m may all depend on a security parameter n . For $T = T(n)$, $k = k(n)$, and $\varepsilon = \varepsilon(n)$, we say that X has (T, ε) next-block pseudoentropy at least k if for every oracle-aided distinguisher $D^{(\cdot)}$ of running time at most T , there exists a set of random variables $\{Y_1, \dots, Y_m\}$ over \mathcal{U} such that*

1. $\sum_{i=1}^m H(Y_i | X_1, \dots, X_{i-1}) \geq k$, and
2. $\mathbb{E}_{i \leftarrow [m]} [\Pr[D^{O_{X,Y}}(X_1, \dots, X_i) = 1] - \Pr[D^{O_{X,Y}}(X_1, \dots, X_{i-1}, Y_i) = 1]] \leq L\varepsilon$, where $O_{X,Y}(i)$, for $i \in [m]$, samples according to the joint distribution (X, Y_i) , and $L = L(n)$ is a bound on the number of calls made by D to $O_{X,Y}$ (including the challenge itself).

We say that X has next-block pseudoentropy at least k if it has $(p(n), 1/p(n))$ -next-block pseudoentropy at least k for every polynomial p . We say that every block of X has (T, ε) -next-block pseudoentropy at least $\alpha = \alpha(n)$ if condition 1 above is replaced with $H(Y_i | X_1, \dots, X_{i-1}) \geq \alpha$ for every $i \in [m]$.

Note that in the informal description given in the introduction we required that (X_1, \dots, X_i) be indistinguishable from $(X_1, \dots, X_{i-1}, Y_i)$ for every i . This would be

equivalent to putting absolute values inside the indistinguishability condition (condition 2) above (and below). This is purely for technical convenience, as D can use $O((m/\varepsilon)^2)$ random samples from its oracle to test whether the (signed) advantages inside the expectation are positive or negative to within an accuracy of $\pm\varepsilon/2m$ and negate itself for some values of i in order to ensure a positive advantage of at least $L\varepsilon/2$.

DEFINITION 3.2 (next-block pseudo-min-entropy). *Let X be a random variable taking values in \mathcal{U}^m , where X , \mathcal{U} , and m may all depend on a security parameter n . For $T = T(n)$, $\alpha = \alpha(n)$, and $\varepsilon = \varepsilon(n)$, we say that every block of X has (T, ε) -next-block pseudo-min-entropy α if for every oracle-aided distinguisher $D^{(\cdot)}$ running in time at most $T(n)$, there exists a set of random variables $\{Y_1, \dots, Y_m\}$ over \mathcal{U} such that*

1. $H_\infty(Y_i \mid X_{1, \dots, i-1}) \geq \alpha$, and
2. $E_{i \leftarrow [m]} [\Pr[D^{O_{X,Y}}(X_1, \dots, X_i) = 1] - \Pr[D^{O_{X,Y}}(X_1, \dots, X_{i-1}, Y_i) = 1]] \leq L \cdot \varepsilon$, where $O_{X,Y}$ and L are as in Definition 3.1.

We say that every block of X has next-block pseudo-min-entropy α if every block of X has $(p(n), 1/p(n))$ -next-block pseudo-min-entropy α for every polynomial p .

Unless explicitly stated otherwise, in the following sections we view a distribution over $\{0, 1\}^t$ as a t -block distribution. When we refer to the next-block pseudoentropy properties of a function $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$, these refer to the random variable $G(\mathcal{U}_n)$.

4. One-way functions to next-block pseudoentropy generator. This section shows how to construct a next-block pseudoentropy generator G_{nb}^f out of a one-way function $f : \{0, 1\}^n \mapsto \{0, 1\}^n$.

THEOREM 4.1 (next-block pseudoentropy generator from one-way functions). *Let n be a security parameter and $f : \{0, 1\}^n \mapsto \{0, 1\}^n$ be a polynomial-time computable function. Then there exists a polynomial-time computable generator $G_{nb} : \{0, 1\}^d \mapsto \{0, 1\}^m$, with $d = d(n) \in O(n)$ and $m = m(n) \in O(n)$, such that the following hold:*

Security. Assume that f is a (T, ε) -one-way function for some $T = T(n)$ and $\varepsilon = \varepsilon(n)$. Then for any $\text{poly}(n)$ -time computable value of $\varepsilon' = \varepsilon'(n) > 2^{-n/4}$, G_{nb} has $(T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon')$ -next-block-pseudoentropy $k = d + \log(1/\varepsilon) - O(\log n)$. Moreover, the reduction from the security of G_{nb} to that of f is fully black-box.⁴

Complexity. G_{nb} is computable in NC^1 with one (uniformly random⁵) oracle call to f .

When f is a standard one-way function, we can take $T = 1/\varepsilon = n^{c'}$ for an arbitrarily large constant c' and set $\varepsilon' = 1/n^{\gamma c'}$ for a small universal constant γ , to deduce that G_{nb} has $(n^{\Omega(c')}, 1/n^{\Omega(c')})$ -next-block pseudoentropy at least $k = d + c' \log n - O(\log n)$. In particular, G_{nb} has next-block pseudoentropy at least $k = d + \log n$.

Our construction employs a family of hash functions $\mathcal{Q} = \{\mathcal{Q}_n = \{q : \{0, 1\}^n \mapsto \{0, 1\}^n\}\}$. We will shortly discuss the properties needed from \mathcal{Q} . Given an appropriate family \mathcal{Q} , we can define G_{nb}^f quite easily.

CONSTRUCTION 4.2. *On security parameter n , define the algorithm G_{nb} on domain $\{0, 1\}^n \times \mathcal{Q}_n$, for $\mathcal{Q}_n = \{q : \{0, 1\}^n \mapsto \{0, 1\}^n\}$, and oracle $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$*

⁴That is, the proof of security treats both f and a possible adversary as black boxes (i.e., as oracles). See [30] for more details.

⁵Here and similarly throughout the paper, we mean that when G_{nb} is evaluated on a uniformly random seed, its oracle call to f is uniformly random.

as follows:

$$G_{nb}^f(s, q) := (f(s), q, q(s)_1, \dots, q(s)_n),$$

where s is n bits long and $q(s)_i$ denotes the i th bit of $q(s)$. (Note that we abuse notation and write q for both the function and its description.)

The following properties regarding the efficiency of G_{nb}^f are immediate.

LEMMA 4.3. *If \mathcal{Q} is in NC^1 , then G_{nb}^f is in NC^1 with one oracle call to f . If the description length of \mathcal{Q}_n is $O(n)$, then the input length of $G_{nb}^{(\cdot)}$ is linear in its first argument (as $|q| = O(|s|)$). Finally, G_{nb}^f invokes f exactly once (and thus is nonadaptive with respect to f).*

On a uniformly random seed for G_{nb} , the single invocation of f is on a uniformly random input as well.

Indeed, we define \mathcal{Q} that is both efficient and has a short description. The main requirement from \mathcal{Q} , however, has to do with ensuring that G_{nb}^f is a next-block pseudoentropy generator. Let us start by presenting the following strategy showing that the entropy gap (i.e., $k-d$) is at least $\log n$ when f is a standard one-way function. Let $D_f(y) := \log |\{x \in \{0, 1\}^n : f(x) = y\}|$, and let S be uniformly distributed over $\{0, 1\}^n$. Then the distribution of S conditioned on $y = f(S)$ still has $D_f(y)$ bits of entropy. We would like \mathcal{Q} to extract these bits and in addition to extract $\log n$ bits of pseudoentropy. More concretely, we ask that the first $D_f(y) + \log n$ bits of $q(S)$ be pseudorandom even given $y = f(S)$ and q (to simplify notation, we ignore round-off errors and treat $D_f(y)$ as an integer). Given such a \mathcal{Q} , we are essentially done (at least when considering nonuniform security⁶). Consider the distributions $X = G_{nb}^f(S, Q)$ and the distribution $(Y_1, \dots, Y_n) := (f(S), Q, R_1, \dots, R_k, Q(S)_{K+1}, \dots, Q(S)_n)$, where S and Q are uniformly distributed, $K := D_f(f(S)) + \log n$, and the R_i 's are uniformly random bits. By the above discussion, X and Y (more formally, X and $\{Y_i\}$) are next-block indistinguishable. In addition, we have

$$\begin{aligned} \sum_{i=1}^m \mathbb{H}(Y_i \mid X_{1, \dots, i-1}) &\geq \mathbb{H}(f(S)) + \mathbb{H}(Q) + \mathbb{H}(R_1, \dots, R_K \mid f(S)) \\ &= \mathbb{H}(f(S)) + \mathbb{H}(Q) + \mathbb{E}[D_f(f(S)) + \log n] \\ &= n + \log |\mathcal{Q}_n| + \log n, \end{aligned}$$

and therefore G_{nb}^f is indeed a next-block pseudoentropy generator.

The first remaining challenge is to construct such a family \mathcal{Q} . As we discuss below, it is easy to obtain all the above properties with hash functions that have description length n^2 . For better efficiency, we settle on \mathcal{Q} with slightly weaker properties (where the pseudorandom bits extracted by $q \in \mathcal{Q}$ are pseudorandom up to advantage $1/n$ rather than an arbitrary inverse polynomial advantage). An additional challenge is achieving next-block pseudoentropy in the (more standard) uniform setting. The difficulty is that we need X and Y to be next-block indistinguishable even given oracles that sample these distributions. While X is efficiently samplable (and thus an oracle that samples X is easy to implement), Y may not be (as $D_f(y)$ may be hard to compute). To overcome this difficulty we employ Holenstein's uniform hardcore lemma [19]. Employing the hardcore lemma also closes the gap between the properties of \mathcal{Q} we obtain and the stronger properties in the discussion.

⁶That is, the distinguisher is nonuniform and does not get oracle access to $O_{X,Y}$.

4.1. The family \mathcal{Q} and unpredictability. A family \mathcal{Q} with the above properties, but with description length n^2 , can be obtained by defining $q(s)$ to be As , where A is a uniformly chosen $n \times n$ matrix over $\text{GF}(2)$. For a random $y = f(S)$, the leftover hash lemma [25] yields that the first $D_f(y) - c \log n$ bits of $Q(S)$ are statistically close to uniform up to statistical distance $1/n^{\Omega(c)}$. An additional $(c + 1) \cdot \log n$ bits are pseudorandom by reduction to the Goldreich–Levin hardcore predicate [8]. An interesting open problem is to come up with a family \mathcal{Q} that has similar properties and in addition has description length n . Instead, in this paper we relax the requirements from \mathcal{Q} .

Defining $q(s) = As$ is equivalent to selecting each one of the output bits of $q(s)$ to be a uniformly selected location of the Hadamard encoding of s . If instead we let each bit be a location in a polynomially long encoding of s , we get description length $n \log n$. As long as this encoding possesses good list-decoding properties, such a construction still suffices for our purposes. To save the final $\log n$ factor, we look at an encoding of x into logarithmically long symbols (and thus need only $n/\log n$ symbols as the output of $q(s)$). The following lemma formalizes the properties of the encoding we use. As with the Hadamard code, the code we use satisfies both the role of extracting $D_f(y) - O(\log n)$ truly random bits via the leftover hash lemma and the role of extracting $O(\log n)$ pseudorandom bits similarly to a hardcore function.

LEMMA 4.4. *There exists an NC^1 algorithm Enc such that, on input $s \in \{0, 1\}^n$, the algorithm Enc produces $t = \text{poly}(n)$ symbols $\text{Enc}(s)_1, \text{Enc}(s)_2, \dots, \text{Enc}(s)_t$ with each $\text{Enc}(s)_i \in \{0, 1\}^\ell$ for $\ell = \lceil \log n \rceil$ and such that the following properties hold:*

Almost 2-universal. *For every two distinct n -bit strings $s \neq s'$ it holds that*

$$\Pr_{i \in [t]} [\text{Enc}(s)_i = \text{Enc}(s')_i] \leq 2^{-\ell} \cdot (1 + 1/(2n^5)).$$

List decoding. *There exists a polynomial-time algorithm Dec that, on input 1^n and given oracle access to a function $\tilde{A} : [t] \times \{0, 1\}^\ell \rightarrow \{0, 1\}$, outputs a list of $\text{poly}(n)$ strings that includes every $s \in \{0, 1\}^n$ satisfying the following:*

$$\Pr_{i \xleftarrow{R} [t]} [\tilde{A}(i, \text{Enc}(s)_i) = 1] - \Pr_{i \xleftarrow{R} [t], z \xleftarrow{R} \{0, 1\}^\ell} [\tilde{A}(i, z) = 1] > 1/5n^2.$$

Note that the oracle \tilde{A} has a domain of size $t \cdot 2^\ell = \text{poly}(n)$, so Dec has enough time to query it on all inputs (i.e., “local decoding” is not needed).

Proof. Enc can be taken to be the concatenation of the Reed–Solomon code of degree $d = n - 1$ over $\mathbb{F} = \text{GF}(2^{c\ell})$ with the Hadamard code over $\mathbb{F}' = \text{GF}(2^\ell)$, where $\ell = \lceil \log n \rceil$ and $c \in \mathbb{N}$ is a sufficiently large constant to be determined below. That is, codewords are of length $t = |\mathbb{F}|^2 = \text{poly}(n)$, and for $s = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$ and $(a, b) \in \mathbb{F}^2$, we set $\text{Enc}(s)_{a,b} = (\sum_{i=0}^{n-1} s_i a^i) \odot b \in \mathbb{F}'$, where \odot denotes dot product (viewing elements of \mathbb{F} as c -dimensional vectors over \mathbb{F}').

The almost-universality property follows from a standard argument: the probability that $\text{Enc}(s)_{a,b} = \text{Enc}(s')_{a,b}$ is bounded by the probability that two corresponding symbols in (the same) random location of two distinct Reed–Solomon codewords agree (which is at most $d/|\mathbb{F}| < n/n^c$) plus the probability that two corresponding symbols in (the same) random location of two distinct Hadamard codewords agree (which is $1/|\mathbb{F}'| = 1/2^\ell$). By taking c sufficiently large, we get the desired property.

For the list-decoding property, given $\tilde{A} : [t] \times \{0, 1\}^\ell \rightarrow \{0, 1\}$ that can distinguish a random symbol of some (unknown) codeword $\text{Enc}(s)$ from uniform with advantage greater than $1/5n^2$, we show how to obtain from \tilde{A} a list $k = \text{poly}(n)$ of “received

words” $r_1, \dots, r_k \in (\{0, 1\}^\ell)^t$ such that $\text{Enc}(s)$ agrees with one of the r_i ’s in more than an $\alpha = (1 + 1/5n^2)/2^\ell$ fraction of positions. By applying the list-decoding algorithm of [10, Thm. 7] to each of the r_i ’s, we are guaranteed to recover each such s , provided that

$$\alpha \geq \frac{1}{|\mathbb{F}'|} + \sqrt{\frac{d}{|\mathbb{F}'|}} + O\left(\frac{1}{t}\right).$$

Since $\alpha = 1/|\mathbb{F}'| + \Omega(1/n^3)$ and $\sqrt{d/|\mathbb{F}'|} + O(1/t) = O(\sqrt{1/n^{c-1}})$, the inequality is satisfied for a sufficiently large choice of the constant c .

We still need to show how to obtain the received words r_1, \dots, r_k from \tilde{A} . For this, we use the standard transformation from distinguishers to predictors [36]. For each pair $u, v \in \mathbb{F}'$, we define

$$r_{u,v}[i] = \begin{cases} u & \text{if } \tilde{A}(i, u) = 1, \\ v & \text{otherwise.} \end{cases}$$

Now consider any s such that

$$\Pr_{i \stackrel{\mathbb{R}}{\leftarrow} [t]} [\tilde{A}(i, \text{Enc}(s)_i) = 1] - \Pr_{i \stackrel{\mathbb{R}}{\leftarrow} [t], z \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^\ell} [\tilde{A}(i, z) = 1] > 1/5n^2.$$

Let’s analyze the expected agreement between s and $r_{u,v}$, when u and v are uniformly random distinct elements of $\{0, 1\}^\ell$:

$$\begin{aligned} & \Pr_{u,v,i} [r_{u,v}[i] = \text{Enc}(s)_i] \\ &= \Pr_{u,v,i} [u = \text{Enc}(s)_i \wedge \tilde{A}(i, \text{Enc}(s)_i) = 1] + \Pr_{u,v,i} [v = \text{Enc}(s)_i \wedge \tilde{A}(i, u) = 0] \\ &= \frac{1}{2^\ell} \cdot \Pr_i [\tilde{A}(i, \text{Enc}(s)_i) = 1] + \frac{1}{2^\ell} \cdot \Pr_{i, u \neq \text{Enc}(s)_i} [\tilde{A}(i, u) = 0] \\ &= \frac{1}{2^\ell} \cdot \left(\Pr_i [\tilde{A}(i, \text{Enc}(s)_i) = 1] + 1 - \Pr_{i, u \neq \text{Enc}(s)_i} [\tilde{A}(i, u) = 1] \right) \\ &> \frac{1}{2^\ell} \cdot \left(1 + \Pr_i [\tilde{A}(i, \text{Enc}(s)_i) = 1] - \Pr_{i,u} [\tilde{A}(i, u) = 1] \right) \\ &> \frac{1 + 1/5n^2}{2^\ell}. \quad \square \end{aligned}$$

CONSTRUCTION 4.5. Let n , Enc , t , and ℓ be as in Lemma 4.4. The description of a random hash function $q \in \mathcal{Q}_n$ is composed of $\lceil n/\ell \rceil$ random indices $i_1, \dots, i_{\lceil n/\ell \rceil} \in [t]$. On input s , define $q(s) = \text{Enc}(s)_{i_1}, \dots, \text{Enc}(s)_{i_{\lceil n/\ell \rceil}}$ (which for simplicity is assumed to be exactly n bits long).

LEMMA 4.6. Let n be a security parameter, let \mathcal{Q} be as in Construction 4.5, let G_{nb} be the oracle-aided algorithm for Construction 4.2 (with respect to \mathcal{Q}), and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a (T, ε) -one-way function for $T = T(n) \geq n, \varepsilon = \varepsilon(n)$. Then there exists a constant $c > 0$ such that

$$\sum_{i=0}^{n-1} \Pr[P(f(S), Q, Q(S)_1, \dots, Q(S)_i) \neq Q(S)_{i+1}] \geq \frac{n - H(f(S)) + \log(1/\varepsilon) - c \log n}{2},$$

even when P is allowed to run in time T/n^c , where S and Q are uniformly distributed over $\{0, 1\}^n$ and \mathcal{Q}_n , respectively.

Note that above (and below) $Q(S)_1, \dots, Q(S)_{i+1}$ refer to individual *bits* of $Q(S)$, not ℓ -bit blocks.

Proof. Consider the following two distributions:

- $X := (f(S), Q, Q(S)_1, \dots, Q(S)_K)$, where K is the random variable $D_f(f(S)) + \log(1/\varepsilon) - c' \log n$, and
- $Y := (f(S), Q, R_1, \dots, R_K)$, where R_1, \dots, R_K are uniformly random and independent bits.

We show that, for a sufficiently large choice of the constant c' , no algorithm with running time $T' = T/n^{c'}$ can distinguish X from Y with advantage better than $1/n$. The lemma then follows by observing that the sum of the error probabilities in predicting the bits R_1, \dots, R_K is at least $E[K]/2$. In more detail, first note that K is always smaller than n , because if $D_f(t) \geq n - \log(1/\varepsilon)$ for some t , then $\Pr[U_n \in f^{-1}(t)] \geq \varepsilon$, which gives rise to a trivial inversion algorithm that succeeds with probability at least ε . Thus,

$$\begin{aligned} & \sum_{i=0}^{n-1} \Pr[P(f(S), Q, Q(S)_1, \dots, Q(S)_i) \neq Q(S)_{i+1}] \\ & \geq \sum_{i=0}^{n-1} \Pr[i < K \wedge (P(f(S), Q, Q(S)_1, \dots, Q(S)_i) \neq Q(S)_{i+1})] \\ & \geq \sum_{i=0}^{n-1} (\Pr[i < K \wedge (P(f(S), Q, R_1, \dots, R_i) \neq R_{i+1})] - 1/n) \\ & \geq \sum_{i=0}^{n-1} (\Pr[i < K]/2 - 1/n) \\ & = E[K]/2 - 1 \\ & = (n - H(f(S)) + \log(1/\varepsilon) - c' \log n - 2)/2. \end{aligned}$$

We assume for simplicity that ℓ always divides both $D_f(t)$ and $\log(1/\varepsilon)$. Let $A = D_f(f(S))/\ell$ and $a' = \log(1/\varepsilon)/\ell$. Then recall that $q(s)$ is defined to be $(\text{Enc}(s)_{i_1}, \dots, \text{Enc}(s)_{i_{\lceil n/\ell \rceil}})$. Therefore, we need to prove that $(\text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{A+a'-c'}})$ is $1/n$ -pseudorandom even given $f(S)$ and $Q = (I_1, \dots, I_{\lceil n/\ell \rceil})$. To prove this, we introduce a hybrid distribution Z , defined as follows:

- $Z := (f(S), Q, \text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{A-5}}, Z_{A-4}, \dots, Z_{A+a'-c'})$, where $Z_1, \dots, Z_{A+a'-c'}$ are uniformly random and independent elements of $\{0, 1\}^\ell$.

First we argue that Z has statistical distance at most $1/10n^2$ from $Y = (f(S), Q, Z_1, \dots, Z_{A+a'-c'})$. Let's condition on any fixed value $f(S) = t$. This also determines that $A = D_f(t)/\ell =: a$. The distribution of S conditioned on $f(S) = t$ still has min-entropy $D_f(t) = a \cdot \ell$. As Enc is almost universal (see the first property in Lemma 4.4), for every two distinct n -bit strings $s \neq s'$ we have that

$$\Pr[\text{Enc}(s)_{I_1}, \dots, \text{Enc}(s)_{I_{a-5}} = \text{Enc}(s')_{I_1}, \dots, \text{Enc}(s')_{I_{a-5}}] \leq 2^{-(a-5)\ell} \cdot (1 + 1/100n^4).$$

Therefore, the leftover hash lemma [25, 26] yields that (conditioned on $f(S) = t$) the prefix $(Q, \text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{A-5}})$ is δ -close to uniform for $\delta = 1/2 \cdot \sqrt{2^{-5\ell} + 1/100n^4} \leq 1/10n^2$.

Now we argue that there is no algorithm B that runs in time T' and distinguishes $X = (f(S), Q, \text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{A+a'-c'}})$ from Z with advantage better than $9/10n$. Assume that there is such an algorithm A . By a hybrid argument,

we get that for a random index J between $A - 4$ and $A + a' - c'$, the algorithm B distinguishes with advantage at least $9/10n^2$ the distributions H^J and H^{J+1} , where

- $H^j = (f(S), Q, \text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{j-1}}, Z_j, \dots, Z_{A+a'-c'})$.

Now, by a Markov argument, it holds that, with probability at least $1/2n^2$ over fixing B 's coin tosses and all components of H^J and H^{J+1} other than $I_J, \text{Enc}(S)_{I_J}$, and Z_J , B distinguishes $(I_J, \text{Enc}(S)_{I_J})$ from (I_J, Z_J) with advantage at least $2/5n^2$. Thus, if we can generate those components given a value of $f(S)$, we can plug B into the decoder Dec guaranteed by Lemma 4.4 to recover S and hence invert f . We will not generate the components exactly but closely enough so that the inversion still succeeds with nonnegligible probability.

The inverter gets as input $t = f(S)$ and operates as follows:

1. Guess the value of $a \stackrel{R}{\leftarrow} [n/\ell]$.
2. Sample uniformly at random coin tosses for B .
3. Sample uniformly at random $i_1, \dots, i_{a-5} \stackrel{R}{\leftarrow} [t]$ and $z_1, \dots, z_{i_{a-5}} \stackrel{R}{\leftarrow} \{0, 1\}^\ell$.
4. Select at random $j \stackrel{R}{\leftarrow} \{a - 4, \dots, a + a' - c'\}$.
5. Sample uniformly at random $i_{a-4}, \dots, i_{j-1}, i_{j+1}, \dots, i_{a+a'-c'} \stackrel{R}{\leftarrow} [t]$ and $z_{a-4}, \dots, z_{j-1}, z_{j+1}, \dots, z_{a+a'-c'} \stackrel{R}{\leftarrow} \{0, 1\}^\ell$.
6. We have fixed all of the inputs of B aside from i_j and z_{i_j} . Denote the resulting algorithm after all of these fixings as \tilde{B} .
7. Invoke $\text{Dec}^{\tilde{B}}(1^n)$, where Dec is as in Lemma 4.4.
8. Select at random an element in the list returned by $\text{Dec}^{\tilde{B}}$.

Let's examine how the distribution of B 's inputs that are fixed by the inverter differs from that in H^J and H^{J+1} :

- The inverter chooses a uniformly at random. This value is correct (i.e., equal to $\ell \cdot D_f(t)$) with probability at least $1/n$.
- The inverter chooses $z_1, \dots, z_{i_{a-5}}$ uniformly at random, whereas in H^J these are distributed according to $(\text{Enc}(S)_{I_1}, \dots, \text{Enc}(S)_{I_{a-5}})$. However, by the indistinguishability of Z and Y , the latter distribution is $(1/10n^2)$ -close to uniform (given $f(S)$ and Q). Thus, this reduces the probability of a good fixing by at most $1/10n^2$, from $1/2n^2$ to $2/5n^2$.
- The inverter chooses z_{a-4}, \dots, z_{j-1} uniformly at random, whereas in H^J these are distributed according to $(\text{Enc}(S)_{I_{a-4}}, \dots, \text{Enc}(S)_{I_{j-1}})$. However, this amounts to guessing at most $u = (a' + 5 - c') \cdot \ell$ bits, and hence is correct with probability at least $2^{-u} \geq \varepsilon \cdot n^{c'}/n^5$.

In total, the list-decoding algorithm successfully recovers S (among a list of $\text{poly}(n)$ possibilities) with probability at least $(1/n) \cdot (2/5n^2) \cdot (\varepsilon \cdot n^{c'}/n^5) = \varepsilon \cdot n^{c'}/\text{poly}(n)$. Choosing at random from the list of decodings reduces the inversion probability by another factor of $\text{poly}(n)$. Taking c' sufficiently large, the inverter succeeds with probability at least ε . We note that the inverter runs in time $\text{poly}(n)$ making at most $2^\ell \cdot t = \text{poly}(n)$ queries to B . Thus if B runs in time $T(n)/n^c$ for a sufficiently large constant c , the inverter runs in time at most $T(n)$, contradicting the security of f . \square

4.2. Proving next-block pseudoentropy via hardcore lemma. Lemma 4.6 shows that the output of G_{nb} (after $f(S)$ and Q) satisfies a weak next-bit unpredictability. In this section, we use Holenstein's uniform variant of Impagliazzo's hardcore lemma [24] to translate this weak next-bit unpredictability into next-bit pseudoentropy, thereby proving Theorem 4.1.

The hardcore lemma translates weak unpredictability (of a single bit) into strong unpredictability on a noticeable fraction of inputs, as shown next.

PROPOSITION 4.7 (see [21]). *Let n be a security parameter, and let $h: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ and $V: \{0, 1\}^n \mapsto \{0, 1\}$ be polynomial-time computable functions. Let $\delta_0 = \delta_0(n) \in [0, 1]$, $\delta = \delta(n) \in [\delta_0, 1]$, and $\gamma = \gamma(n) \in [0, 1] > 2^{-n/3}$. Assume that*

$$\Pr[M(\delta_0, \gamma, h(U_n)) = V(U_n)] \leq 1 - \delta/2$$

for every probabilistic algorithm M running in time $T = T(n)$ and large enough n . Then for every oracle-aided predictor P running in time $T \cdot (\gamma\delta_0/n)^{O(1)}$ and all sufficiently large n , there exists a set $L \subseteq \{0, 1\}^n$ of density at least δ such that

$$\Pr_{W \stackrel{R}{\leftarrow} L} [P^{\chi_L}(h(W)) = V(W)] < (1 + \gamma)/2,$$

where χ_L is the characteristic function of L , provided that all the queries of P to χ_L are computed independently of the input $h(W)$. Furthermore, the reduction is fully black-box.

The above differs from Holenstein’s hardcore lemma in that it does not require that δ be known to M , but requires only a lower bound δ_0 on δ . However, the above can be readily deduced from Holenstein’s version; see the appendix.

We now reinterpret Holenstein’s hardcore lemma in terms of conditional pseudoentropy, similarly to the reinterpretation of Impagliazzo’s hardcore lemma in [31].

PROPOSITION 4.8. *Let n be a security parameter, $\delta_0 = \delta_0(n) \in [0, 1]$, $\delta = \delta(n) \in [\delta_0, 1]$, and $\gamma = \gamma(n) \in [0, 1] > 2^{-n/3}$. Let (A, B) be a poly(n)-time samplable random variable over $\{0, 1\}^{\ell(n)} \times \{0, 1\}$ such that*

$$\Pr[M(\delta_0, \gamma, A) = B] \leq 1 - \delta/2$$

for every probabilistic algorithm M running in time $T = T(n)$ and large enough n . Then for every oracle-aided distinguisher D running in time $T' = T \cdot (\delta_0\gamma/n)^{O(1)}$ and all sufficiently large n , there is a random variable C , jointly distributed with (A, B) , such that

1. $H(C|A) \geq \delta$;
2. $\Pr[D^{O_{A,B,C}}(A, B) = 1] - \Pr[D^{O_{A,B,C}}(A, C) = 1] \leq \gamma$,

where $O_{A,B,C}$ is an oracle that samples according to the joint distribution (A, B, C) .

Proof. Let $(h, V) : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)} \times \{0, 1\}$ be the poly-time sampling algorithms for (A, B) , i.e., $(h(U_n), V(U_n)) = (A, B)$. (By renaming the security parameter n , we may assume that the sampling algorithms use n coin tosses.) Thus we may apply Proposition 4.7 to the pair (h, V) . For any given subset $L \subseteq \{0, 1\}^n$ of density δ , we define a probabilistic function $V_L : \{0, 1\}^n \rightarrow \{0, 1\}$, where

$$V_L(r) = \begin{cases} V(r) & \text{if } r \notin L, \\ \text{a random bit} & \text{if } r \in L. \end{cases}$$

From this we get a random variable C_L jointly distributed with (A, B) , defined by $(A, C_L) = (h(U_n), V_L(U_n))$. Notice that $H(C_L|A)$ is at least the density of L , namely δ . We show that taking $C = C_L$ for some L suffices. Suppose for contradiction that we have an oracle-aided distinguisher D running in time T' such that, for every L of density δ , $\Pr[D^{O_{A,B,C_L}}(A, B) = 1] - \Pr[D^{O_{A,B,C_L}}(A, C_L) = 1] > \gamma$. Since B and C_L are identical when $U_n \notin L$, we have

$$\Pr[D^{O_{A,B,C_L}}(A, B) = 1|U_n \in L] - \Pr[D^{O_{A,B,C_L}}(A, C_L) = 1|U_n \in L] > \gamma.$$

Since C_L is a uniformly random bit when $U_n \in L$, we can apply the standard reduction from distinguishing to predicting to obtain an oracle-aided predictor P , running in time $T' + O(1)$ such that

$$(1) \quad \Pr[P^{X_L}(A) = B | U_n \in L] > (1 + \gamma)/2.$$

Specifically, on input x , P generates a random bit $b \stackrel{R}{\leftarrow} \{0, 1\}$, runs $D^{O_{A,B,C_L}}(x, b)$, outputs b if D outputs 1, and outputs $-b$ if D outputs 0. P can simulate random samples from the oracle O_{A,B,C_L} by choosing $r \stackrel{R}{\leftarrow} \{0, 1\}^n$ and outputting $(h(r), V(r), V_L(r))$, which can be efficiently computed using P 's oracle access to χ_L . Equation (1) can be rewritten as

$$\Pr_{W \stackrel{R}{\leftarrow} L} [P^{X_L}(h(W)) = V(W)] > (1 + \gamma)/2.$$

This contradicts Proposition 4.7. \square

We now use this form of the hardcore lemma to deduce Theorem 4.1 from Lemma 4.6.

THEOREM 4.9 (Theorem 4.1, restated). *Let n be a security parameter, and let $f : \{0, 1\}^n \mapsto \{0, 1\}^n$ be a polynomial-time computable function. Then there exists a polynomial-time computable generator $G_{nb} : \{0, 1\}^d \mapsto \{0, 1\}^m$, with $d = d(n) \in O(n)$ and $m = m(n) \in O(n)$, such that the following hold:*

Security. Assume that f is a (T, ε) -one-way function for some $T = T(n)$ and $\varepsilon = \varepsilon(n)$. Then for any poly(n)-time computable value of $\varepsilon' = \varepsilon'(n) > 2^{-n/4}$, G_{nb} has $(T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon')$ -next-block-pseudoentropy $k = d + \log(1/\varepsilon) - O(\log n)$. Moreover, the reduction from the security of G_{nb} to that of f is fully black-box.

Complexity. G_{nb} is computable in NC^1 with one (uniformly random⁷) oracle call to f .

Proof of Theorem 4.9. Let $W = (Q, S)$ be uniformly distributed over the domain of G_{nb} (i.e., $\{0, 1\}^n \times \mathcal{Q}_n$), let $X = G_{nb}(W)$, and denote the length of X by $m = m(n) = 2n + \log |\mathcal{Q}_n|$. Assume that the theorem does not hold. In particular, there exists an oracle-aided algorithm D whose running time is bounded by $T' = T'(n)$, and the following holds for every set of distributions $\{Y_i\}_{i \in [m]}$ over $\{0, 1\}$, jointly distributed with X , with $\sum_{i=1}^m \mathbb{H}(Y_i | X_1, \dots, X_{i-1}) \geq k = d + \log(1/\varepsilon) - c \cdot \log n$:

$$(2) \quad \mathbb{E}_{i \stackrel{R}{\leftarrow} [m]} [\Pr[D^{O_{X,Y}}(X_1, \dots, X_i) = 1] - \Pr[D^{O_{X,Y}}(X_1, \dots, X_{i-1}, Y_i) = 1]] \geq L \cdot \varepsilon' \geq \varepsilon',$$

where L is the number of oracle queries made by D and c is the constant of Lemma 4.6.

Recall that $X = (f(S), Q, Q(S)_1, \dots, Q(S)_n)$. By Lemma 4.6, if we take $A = (f(S), Q, Q(S)_1, \dots, Q(S)_{I-1})$ and $B = Q(S)_I$ for I uniformly random over $[n]$, we have

$$(3) \quad \Pr[M(\varepsilon', A) = B] \leq 1 - \delta/2$$

for

$$\delta = \frac{n - \mathbb{H}(f(S)) + \log(1/\varepsilon) - c \log n}{n}$$

⁷Here and throughout the paper, when we say that G_{nb} makes a uniformly random call to f , we refer to the distribution of the call when we run G_{nb} on a uniformly random seed.

and any algorithm M running in time $T/n^c - p(n) \leq T/2 \cdot n^c$, where $p \in \text{poly}$ is the computing time of ε' , and for the inequality we assume without loss of generality that $T > 2p(n) \cdot n^c$ (otherwise the conclusion of the theorem holds trivially). In the following we use Proposition 4.8 for showing that (3) is in contradiction to (2).

Let C be any random variable, jointly distributed with A , such that $H(C|A) \geq \delta$. We can use any such C to define a set $\{Y_i\}$ of random variables jointly distributed with X :

$$Y_i|_{X=(x_1, \dots, x_m)} = \begin{cases} x_i & \text{if } i \leq n + \log |\mathcal{Q}|, \\ C|_{I=i, X=(x_1, \dots, x_n)} & \text{otherwise.} \end{cases}$$

Then

$$\begin{aligned} \sum_{i=1}^m H(Y_i|X_1, \dots, X_{i-1}) &= \sum_{i=1}^{n+\log |\mathcal{Q}|} H(X_i|X_1, \dots, X_{i-1}) + \sum_{i=1}^n H(C|A, I=i) \\ &= H(f(S), Q) + n \cdot H(C|A) \\ &\geq H(f(S)) + H(Q) + n - H(f(S)) + \log(1/\varepsilon) - c \log n \\ &= d + \log(1/\varepsilon) - c \log n = k, \end{aligned}$$

where d is the seed length of G_{nb} . Therefore (2) holds; i.e., D is a next-block distinguisher for X and $\{Y_i\}$ with advantage ε' . Observe that D has zero advantage when $i \leq n + \log |\mathcal{Q}|$ (because then $Y_i = X_i$) and thus must gain its entire advantage when $i > n + \log |\mathcal{Q}|$. In the latter case, D is distinguishing (A, B) from (A, C) with advantage ε' . Moreover, each of D 's oracle queries to $O_{X,Y}(i)$ can be simulated with $O(n)$ queries to $O_{A,B,C}$. (It may take an expected $O(n)$ trials to get a sample in which I has the desired value.) Thus, we obtain an algorithm D_0 running in time $T_0 = O(T' \cdot n)$ such that $D_0^{O_{A,B,C}}$ distinguishes (A, B) from (A, C) with advantage greater than ε' for any C such that $H(C|A) \geq \delta$. Proposition 4.8 yields the existence of an algorithm M of running time $T_0 \cdot (n/\varepsilon')^{O(1)}$ such that $\Pr[M(\varepsilon', A) = B] > 1 - \delta/2$ for infinitely many n 's. (We can set $\delta_0 = 1/n$, since if $\delta < 1/n$, the conclusion of the theorem holds trivially.) Setting the constant in the definition of T' to a large enough value (recall that $T' = T \cdot (\varepsilon'/n)^{O(1)}$) yields that the running time of M is bounded by $T/2 \cdot n^c$, in contradiction to (3). \square

5. From next-block pseudoentropy to pseudorandom generators. In this section we show how to transform a next-block pseudoentropy generator into a pseudorandom generator.

THEOREM 5.1 (next-block pseudoentropy generator to pseudorandom generator). *Let n be a security parameter, and let $m = m(n)$, $\Delta = \Delta(n) \in [1/\text{poly}(n), n]$, and $\kappa = \kappa(n) \in \{1, \dots, n\}$ be $\text{poly}(n)$ -time computable. For every polynomial-time computable, m -block generator $G_{nb}: \{0, 1\}^n \mapsto \{0, 1\}^m$, there exists a polynomial-time computable generator $G: \{0, 1\}^d \mapsto \{0, 1\}^{d \cdot (1 + \Omega(\Delta/n))}$ with seed length*

$$d = d(n) = O\left(\frac{n^2 \cdot m^2 \cdot \kappa \cdot \log^2 n}{\Delta^3}\right)$$

such that the following hold:

Security. *If G_{nb} has (T, ε) -next-block pseudoentropy at least $n + \Delta$, for $T = T(n)$, $\varepsilon = \varepsilon(n)$, then G is a $(T - n^{O(1)}, n^{O(1)} \cdot (\varepsilon + 2^{-\kappa}))$ -pseudorandom generator. Moreover, the reduction is fully black-box.*

Complexity. G is computable in NC^1 with $O(d/n)$ (uniformly random) oracle calls to G_{nb} .

In Theorem 5.1, it may be convenient to view $\kappa(n)$ as the security parameter of the construction. In particular, when $\kappa(n)$ is logarithmic in $1/\varepsilon(n)$, we get that $(T(n), \varepsilon(n))$ -next-block pseudoentropy turns into a $(T(n)/\text{poly}(n), \text{poly}(n) \cdot \varepsilon(n))$ -pseudorandom generator.

We prove Theorem 5.1 via the following sequence of reductions:

1. In section 5.1 we show how to get a better handle on the output distribution of the G_{nb} —specifically, we apply a generic transformation on G_{nb} to get a generator for which the (conditional) pseudoentropy of each of its output blocks is the same (i.e., $(n + \Delta)/m$).
2. In section 5.2 we consider a direct product of the latter next-block pseudoentropy generator and show that this action both increases the absolute gap between the next-block pseudoentropy of the generator and its real entropy (i.e., its input length) and transforms its next-block pseudoentropy into next-block pseudo-min-entropy.
3. In section 5.3 we show how to extract pseudorandomness from the output of the latter type of generators.
4. In section 5.4, we put the above parts together to prove Theorem 5.1.

To simplify notation, we prove the first three steps with respect to arbitrary next-block pseudoentropy distributions. Given a distribution X over \mathcal{U}^m , a set of distributions $Y = \{Y_i\}_{i \in [m]}$ over \mathcal{U} , and an oracle-aided algorithm $D^{(\cdot)}$, we let $\delta_{X,Y}^D := \mathbb{E}_{i \in [m]} [\delta_{X,Y,i}^D := \Pr[D^{O_{X,Y}}(X_1, \dots, X_i) = 1] - \Pr[D^{O_{X,Y}}(X_1, \dots, X_{i-1}, Y_i) = 1]]$, where $O_{X,Y}(i)$ samples according to the joint distribution (X, Y_i) (see Definition 3.1). Finally, in all of the following claims we assume that the description of the “universe” \mathcal{U} is polynomial in n .

5.1. Entropy equalization. In this section we show how to manipulate a given distribution to gain a better characterization of its next-block pseudoentropy, without losing “too much” pseudoentropy. The following transformation is closely related to a similar reduction from [17]. The idea is the following: Consider an m -block random variable X over \mathcal{U}^m with next-block pseudoentropy k . Now generate $m \cdot \ell$ blocks by concatenating ℓ independent copies one after the other. Finally, for a random $j \in [m]$, erase the first j blocks and the last $m - j$ blocks. We now have a new variable \tilde{X} with $m \cdot (\ell - 1)$ blocks, and for every location i the block in the i th location of \tilde{X} is a block of X in a random location. It is not hard to prove (as we do below) that the next-block pseudoentropy of each block is at least k/m . On the other hand, the real entropy of \tilde{X} is at most ℓ times that of X . Taking large enough ℓ , we get that the (relative) difference between next-block pseudoentropy and real entropy has not significantly decreased.

For $j \in [m]$ and $z^{(1)}, \dots, z^{(\ell)} \in \mathcal{U}^m$, we let

$$\text{Equalizer}(j, z^{(1)}, \dots, z^{(\ell)}) := (z_j^{(1)}, \dots, z_m^{(1)}, \dots, z_1^{(\ell)}, \dots, z_{j-1}^{(\ell)}).$$

LEMMA 5.2. *Let n be a security parameter, and let $m = m(n) = \text{poly}(n)$ and $\ell = \ell(n) = \text{poly}(n)$ be $\text{poly}(n)$ -time computable integer functions, where $\ell(n) > 1$. Let X be a random variable over \mathcal{U}^m with (T, ε) -next-block pseudoentropy at least k for $T = T(n)$, $\varepsilon = \varepsilon(n)$, and $k = k(n)$. Let J be uniformly distributed over $[m]$, and let $\tilde{X} = \text{Equalizer}(J, X^{(1)}, \dots, X^{(\ell)})$, where the $X^{(i)}$'s are independent and identically distributed (i.i.d.) copies of X . Then every block of \tilde{X} has $(T - O(\ell \cdot m \cdot \log |\mathcal{U}|), \ell \cdot \varepsilon)$ -*

next-block pseudoentropy at least k/m . Moreover, the reduction between the security of \tilde{X} and X is fully black-box.

Proof. Let $m' = (\ell - 1) \cdot m$, and let $Y = \{Y_1, \dots, Y_m\}$ be a set of random variables jointly distributed with X . In the following we think of Y as a single random variable $Y = (Y_1, \dots, Y_m)$ jointly distributed with X , though we sample only a single entry Y_i per instance of Y . Let $Y^{(1)}, \dots, Y^{(\ell)}$ be i.i.d. copies of Y , and let $\tilde{Y} = \text{Equalizer}(J, Y^{(1)}, \dots, Y^{(\ell)})$ be jointly distributed with \tilde{X} in the natural way— J takes the same value as in \tilde{X} , and for every $j \in [\ell]$, $Y^{(j)}$ is jointly distributed with $X^{(j)}$ according to the joint distribution (X, Y) . Notice that $\tilde{Y}_i = Y_{J+i-1 \bmod m}$ (where we define $m \bmod m$ to equal m rather than 0) and that $J + i - 1$ is uniformly distributed in $[m]$.

Thus, for every $i \in [m']$ we have that

$$(4) \quad \begin{aligned} \mathbb{H}(\tilde{Y}_i \mid \tilde{X}_{1, \dots, i-1}) &\geq \mathbb{H}(Y_{J+i-1 \bmod m} \mid X_1, \dots, X_{(J+i-1 \bmod m)-1}) \\ &= \mathbb{E}_{i' \stackrel{\mathbb{R}}{\leftarrow} [m]} [\mathbb{H}(Y_{i'} \mid X_1, \dots, X_{i'-1})]. \end{aligned}$$

Let \tilde{D} be an adversary that violates the next-block pseudoentropy of \tilde{X} . We define D for breaking the next-block pseudoentropy of X as follows: On input (x_1, \dots, x_{i-1}, z) , D generates a random sample $x' = (x'_1, \dots, x'_{m'})$ from \tilde{X} (using $O_{X,Y}$). It then selects $i' \in [m']$ uniformly at random such that $i' = j + i - 1 \bmod m$, where j is the value of J in the generation of x' , and returns $\tilde{D}^{O_{\tilde{X}, \tilde{Y}}}(x'_1, \dots, x'_{i'-i}, x_1, \dots, x_{i-1}, z)$, while answering \tilde{D} 's queries to $O_{\tilde{X}, \tilde{Y}}$ using $O_{X,Y}$.

We note that D makes at most ℓ times more oracle queries than \tilde{D} , and that D can be implemented in the running time of \tilde{D} plus $O(\ell \cdot m \cdot \log |\mathcal{U}|)$.

For every Y as above with $\sum_{i \in [m]} \mathbb{H}(Y_i \mid X_{1, \dots, i-1}) \geq k$, (4) yields that \tilde{Y} is a random variable that \tilde{D} should be able to next-block distinguish from \tilde{X} . Since the query to \tilde{D} done by D is distributed identically to a random challenge to \tilde{D} with respect to the joint distribution (\tilde{X}, \tilde{Y}) , it follows that

$$\delta_{\tilde{X}, \tilde{Y}}^{\tilde{D}} = \delta_{X, Y}^D \leq L \cdot \varepsilon = \tilde{L} \cdot (\ell \varepsilon),$$

where L and \tilde{L} are the number of oracle calls made by D and \tilde{D} , respectively. This is in contradiction to the assumption about the next-block pseudoentropy of X . \square

5.2. Next-block pseudoentropy converts to pseudo-min-entropy. In this section we show how to transform next-block (Shannon) pseudoentropy to next-block pseudo-min-entropy, while increasing the overall entropy gap. The transformation of X is simply a t -fold parallel repetition of X (i.e., every block of the new random variable X^t is composed of t corresponding blocks of t independent copies of X). This generalizes an analogous transformation that was used by Håstad et al. [18] in the context of standard (i.e., single-block) pseudoentropy.

Given an m -block random variable V taking values in \mathcal{U}^m and an integer $t > 0$, we let $V^t = ((V_1^{(1)}, \dots, V_1^{(t)}), \dots, (V_m^{(1)}, \dots, V_m^{(t)})) \in (\mathcal{U}^t)^m$, where the $V^{(i)}$'s are i.i.d. copies of V .

LEMMA 5.3. *Let n be a security parameter, let $m = m(n) = \text{poly}(n)$ and $t = t(n) = \text{poly}(n)$ be $\text{poly}(n)$ -time computable functions, and let X be a random variable over \mathcal{U}^m , where every block of X has (T, ε) -next-block pseudoentropy at least α for $T = T(n)$, $\varepsilon = \varepsilon(n)$, $\alpha = \alpha(n)$. Then for every $\kappa = \kappa(n) > 0$, it holds that every block of X^t has (T', ε') -next-block pseudo-min-entropy α' , where*

- $T' = T'(n) = T - O(m \cdot t \cdot \log |\mathcal{U}|)$,
- $\varepsilon' = \varepsilon'(n) = t^2 \cdot (\varepsilon + 2^{-\kappa} + 2^{-c \cdot t})$ for a universal constant $c > 0$, and
- $\alpha' = \alpha'(n) = t \cdot \alpha - \Gamma(t, \kappa, |\mathcal{U}|)$ for $\Gamma(t, \kappa, |\mathcal{U}|) \in O(\sqrt{t \cdot \kappa} \cdot \log(|\mathcal{U}| \cdot t))$.

Furthermore, the reduction between the security of X^t and X is fully black-box.

Note that the $t \cdot \alpha$ term is the largest we could hope for the pseudoentropy—getting α bits of pseudoentropy per copy. However, since we wish to move from a pseudoform of Shannon entropy (measuring randomness on average) to a pseudoform of min-entropy (measuring randomness with high probability), we may have a deviation that grows like \sqrt{t} . By taking t large enough, this deviation becomes insignificant.

In more detail, consider the case that X has next-block pseudoentropy at least α with polynomial security; i.e., T and $1/\varepsilon$ can be taken to be arbitrarily large polynomials in n . We would like to deduce that X^t has next-block pseudo-min-entropy α' with polynomial security. Moreover, assume $\mathcal{U} = \{0, 1\}$. Then we should take κ and t to be arbitrarily large multiples of $\log n$, and we have $\alpha' = t \cdot (\alpha - O(\sqrt{(\log n)/t}) \cdot \log t)$. So if we would like to have pseudo-min-entropy at least $t \cdot (\alpha - \delta)$, we should take t to be $\text{polylog}(n)/\delta^2$. In our application, we have $\delta = \Theta(\Delta/n) = \Theta(\log n/n)$, so we take $t = \tilde{O}(n^2)$ copies.

Proof. For a random variable Y_i over \mathcal{U} jointly distributed with X , we define the (jointly distributed with X^t) variable $(Y_i)^t = ((Y_i)^{(1)}, \dots, (Y_i)^{(t)})$, where for each $j \in [t]$ the entry $(Y_i)^{(j)}$ is jointly distributed with the entry $X^{(j)}$ in X^t according to the joint distribution (X, Y_i) . Given an adversary D_t that violates the next-block pseudo-min-entropy of X^t , we define D for breaking the next-block pseudoentropy of X as follows: Given an input (x_1, \dots, x_{i-1}, z) and oracle access to $O_{X,Y}$, D first samples $j \stackrel{R}{\leftarrow} [t]$ and (x^t, y^t) from $(X^t, (Y_i)^t)$ (using $O_{X,Y}$). Then it replaces $(x_{j,1}^t, \dots, x_{j,i-1}^t)$ (the $(i-1)$ prefix of the j th column of x^t) with (x_1, \dots, x_{i-1}) , sets $z^{[j]} = (x_{i,1}^t, \dots, x_{i,j-1}^t, z, y_{j+1}^t, \dots, y_i^t)$, and returns $D_t^{O_{X^t, (Y_i)^t}}(x_{1, \dots, i-1}^t, z^{[j]})$, while answering D_t 's queries to $O_{X^t, (Y_i)^t}$ using $O_{X,Y}$. That is, D queries D_t on the j th hybrid between $(X_{1, \dots, i-1}^t, (Y_i)^t)$ and $X_{1, \dots, i}^t$.

We note that D makes at most t times more oracle calls than D_t , and that it can be implemented in the same time as D plus $O(t \cdot m \cdot \log |\mathcal{U}|)$.

Assuming that $H(Y_i | X_{1, \dots, i-1}) \geq \alpha$ for $i \in [m]$, Lemma 2.1 yields that there exists a random variable W over \mathcal{U}^t , jointly distributed with X^t , such that the following hold:

1. $\Delta((X_{1, \dots, i-1}^t, (Y_i)^t), (X_{1, \dots, i-1}^t, W)) \leq 2^{-\kappa} + 2^{-c \cdot t}$ for a universal constant $c > 0$, and
2. $H_\infty(W | x_{1, \dots, i-1}) \geq \alpha - \Gamma(t, \kappa, |\mathcal{U}|)$ for every $x \in \text{Supp}(X^t)$.

For $j \in [t]$, let $Z^{[j]} = ((X_i^t)_{1, \dots, j}, (Y_i^t)_{j+1, \dots, t})$ (i.e., the j th hybrid between X_i^t and $(Y_i)^t$). It follows that for each $i \in [m]$, we have

$$\begin{aligned} & \delta_{X,Y,i}^D \\ & := \frac{1}{t} \cdot \sum_{j \in [t]} (\Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, Z^{[j]}) = 1] - \Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, Z^{[j-1]}) = 1]) \\ & = \frac{1}{t} \cdot (\Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, Z^{[t]}) = 1] - \Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, Z^{[0]}) = 1]) \\ & = \frac{1}{t} \cdot (\Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, X_i^t) = 1] - \Pr[D_t^{O_{X^t, (Y_i)^t}}(X_{1, \dots, i-1}^t, (Y_i)^t) = 1]) \\ & \geq \frac{1}{t} \cdot (\Pr[D_t^{O_{X^t, W}}(X_{1, \dots, i-1}^t, X_i^t) = 1] - \Pr[D_t^{O_{X^t, W}}(X_{1, \dots, i-1}^t, W) = 1]) \\ & \quad - L' \cdot (2^{-\kappa} - 2^{-c \cdot t}) \end{aligned}$$

$$=: \frac{1}{t} \cdot (\delta_{X^t, W, i}^{D_t} - L' \cdot (2^{-\kappa} - 2^{-c \cdot t})),$$

where L' is a bound on the number of oracles calls made by D_t . Taking expectations over $i \stackrel{R}{\leftarrow} [m]$, we have

$$\begin{aligned} \delta_{X, Y}^D &\geq \frac{1}{t} \cdot (\delta_{X^t, W}^{D_t} - L' \cdot (2^{-\kappa} + 2^{-c \cdot t})) \\ &\geq \frac{1}{t} \cdot (t^2 \cdot L' \cdot (\varepsilon + 2^{-\kappa} + 2^{-c \cdot t}) - L' \cdot (2^{-\kappa} + 2^{-c \cdot t})) \\ &\geq L' \cdot t \cdot \varepsilon = L\varepsilon, \end{aligned}$$

where $L = t \cdot L'$ is a bound on the number of oracle calls made by D . Since this holds for all $\{Y_i\}$ such that $H(Y_i \mid X_{1, \dots, i-1}) \geq \alpha$, this contradicts the assumption about the next-block pseudoentropy of X . \square

5.3. Next-block pseudo-min-entropy to pseudorandomness. For our final step, we assume that X is such that each of the m blocks of X has large next-block pseudo-min-entropy α . Using a two-universal independent hash function S , we extract almost all its pseudo-min-entropy of each block individually. The result is a sufficiently long pseudorandom bit sequence. This is a computational analogue of block-source extraction in the literature on randomness extractors [3, 37].

LEMMA 5.4. *Let n be a security parameter, and let $m = m(n) = \text{poly}(n)$, $t = t(n) = \text{poly}(n)$, $\alpha = \alpha(n) \in [t(n)]$, and $\kappa = \kappa(n) \in [\alpha(n)]$ be $\text{poly}(n)$ -time computable integer functions. There exists an efficient procedure $\text{Ext} \in \text{NC}^1$ that, on input $x \in (\{0, 1\}^t)^m$ and $s \in \{0, 1\}^t$, outputs a string $y \in \{0, 1\}^{t+m \cdot (\alpha - \kappa)}$ such that the following holds.*

Let X be a random variable over $(\{0, 1\}^t)^m$ such that every block of X has (T, ε) -next-block pseudo-min-entropy α for $T = T(n)$ and $\varepsilon = \varepsilon(n)$; then $\text{Ext}(X, U_t)$ is $(T - m \cdot t^{O(1)}, m \cdot (\varepsilon + 2^{-\kappa/2}))$ pseudorandom. Moreover, the reduction between the pseudorandomness of $\text{Ext}(X, U_t)$ and the security X is fully black-box.

Proof. Let $\text{Ext}(x, s) := (s, s(x_1), \dots, s(x_m))$, where s is interpreted as a member of a family of two-universal hash functions from t bits to $\alpha - \kappa$ bits in NC^1 (such as $s(x) := s \cdot x$ over $\text{GF}(2^t)$ truncated to $\alpha - \kappa$ bits). Let D_{PRG} be an adversary for the pseudorandomness of $\text{Ext}(X, U_t)$, and let δ_{PRG} be its distinguishing advantage. We define D for breaking the next-block pseudoentropy of X as follows: On input (x_1, \dots, x_{i-1}, z) , D returns $D_{\text{PRG}}(s, s(x_1), \dots, s(x_{i-1}), s(z), U_{(\alpha - \kappa) \cdot (m - i)})$, where s is uniformly chosen from $\{0, 1\}^t$.

We note that D makes no oracle calls (and thus we count its query complexity as one) and that its running time is at most that of D_{PRG} plus $m \cdot \text{poly}(t)$.

Let $Z^{[i]}(W) := (S, S(X_1), \dots, S(X_{i-1}), S(W), U_{(\alpha - \kappa) \cdot (m - i)})$ for a uniformly distributed hash function S . Let $Y = \{Y_1, \dots, Y_m\}$ be a set of random variables over \mathcal{U} jointly distributed with X , with $H_\infty(Y_i \mid X_{1, \dots, i-1} = x_{1, \dots, i-1}) \geq \alpha$ for every $x \in \text{Supp}(X)$ and $i \in [m]$. The leftover hash lemma [25, 26] yields that $Z^{[i]}(Y_i)$ has statistical difference at most $2^{-\kappa/2}$ from $Z^{[i-1]}(X_{i-1})$. Thus

$$\begin{aligned} \delta_{\text{PRG}} &= \Pr[D_{\text{PRG}}(Z^{[m]}(X_m)) = 1] - \Pr[D_{\text{PRG}}(Z^{[0]}(X_0)) = 1] \\ &= \sum_{i=1}^m \left(\Pr[D_{\text{PRG}}(Z^{[i]}(X_i)) = 1] - \Pr[D_{\text{PRG}}(Z^{[i-1]}(X_{i-1})) = 1] \right) \\ &\leq \sum_{i=1}^m \left(\Pr[D_{\text{PRG}}(Z^{[i]}(X_i)) = 1] - \Pr[D_{\text{PRG}}(Z^{[i]}(Y_i)) = 1] + 2^{-\kappa/2} \right) \end{aligned}$$

$$\begin{aligned}
 &= m \cdot (\delta_{X,Y}^D + 2^{-\kappa/2}) \\
 &\leq m \cdot (\varepsilon + 2^{-\kappa/2}). \quad \square
 \end{aligned}$$

5.4. Putting it together. We are now ready to prove Theorem 5.1.

THEOREM 5.5 (Theorem 5.1, restated). *Let n be a security parameter, and let $m = m(n)$, $\Delta = \Delta(n) \in [1/\text{poly}(n), n]$, and $\kappa = \kappa(n) \in \{1, \dots, n\}$ be $\text{poly}(n)$ -time computable. For every polynomial-time computable, m -block generator $G_{nb}: \{0, 1\}^n \mapsto \{0, 1\}^m$, there exists a polynomial-time computable generator $G: \{0, 1\}^d \mapsto \{0, 1\}^{d \cdot (1 + \Omega(\Delta/n))}$ with seed length*

$$d = d(n) = O\left(\frac{n^2 \cdot m^2 \cdot \kappa \cdot \log^2 n}{\Delta^3}\right)$$

such that the following hold:

Security. *If G_{nb} has (T, ε) -next-block pseudoentropy at least $n + \Delta$, for $T = T(n)$, $\varepsilon = \varepsilon(n)$, then G is a $(T - n^{O(1)}, n^{O(1)} \cdot (\varepsilon + 2^{-\kappa}))$ -pseudorandom generator. Moreover, the reduction is fully black-box.*

Complexity. *G is computable in NC^1 with $O(d/n)$ (uniformly random) oracle calls to G_{nb} .*

Proof. Let $X = G_{nb}(U_n)$. Without loss of generality, we may assume that the number m of output blocks (=bits) of G_{nb} is a power of 2 (by padding with zeros if necessary). X can be generated using n random bits and has next-block pseudoentropy $k = n + \Delta$.

We now set $\ell = \lceil 2(n + \Delta + \log m)/\Delta \rceil = O(n/\Delta)$ and apply entropy equalization (Lemma 5.2) to obtain $\tilde{X} := \text{Equalizer}(J, X^{(1)}, \dots, X^{(\ell)})$, where J is uniformly distributed over $[m]$ and the $X^{(i)}$'s are i.i.d. copies of X . \tilde{X} can be generated using $d_\ell = \log m + \ell \cdot n$ random bits and has $m' = (\ell - 1) \cdot m$ blocks. Every block of \tilde{X} has $(T_\ell = T - O(\ell \cdot m), \varepsilon_\ell = \ell \cdot \varepsilon)$ -next-block pseudoentropy at least $\alpha_\ell = k/m$. Thus the total next-block pseudoentropy in \tilde{X} is at least

$$\begin{aligned}
 m' \cdot \alpha_\ell &= m \cdot (\ell - 1) \cdot k/m \\
 &= (n + \Delta) \cdot (\ell - 1) \\
 &\geq n \cdot \ell + \log m + \Delta \ell / 2 \\
 &= d_\ell + \Delta \ell / 2,
 \end{aligned}$$

where the inequality follows from the setting of ℓ .

Next we apply t -fold parallel repetition (Lemma 5.3) to obtain $(\tilde{X})^t$ for a parameter $t = \text{poly}(n)$ to be set below. $(\tilde{X})^t$ can be generated using $d_t = t \cdot d_\ell$ random bits and has m' blocks of t bits each. Every block of $(\tilde{X})^t$ has $(T_t = T_\ell - O(m' \cdot t), \varepsilon_t = t^2 \cdot (\varepsilon_\ell + 2^{-\kappa} + 2^{-\Omega(t)}))$ -next-block pseudo-min-entropy $\alpha_t = t \cdot \alpha_\ell - \Gamma(t, \kappa)$.

Finally, we apply Lemma 5.4, extracting $\alpha_t - 2\kappa$ bits from each block of $(\tilde{X})^t$ using a seed of length t . This yields our final output $\text{Ext}((\tilde{X})^t, U_t)$, which is (T', ε') -pseudorandom, for

$$\begin{aligned}
 T' &= T_t - m' \cdot t^{O(1)} = T - \text{poly}(n), \\
 \varepsilon' &= m' \cdot (\varepsilon_t + 2^{-\kappa}) = \text{poly}(n) \cdot (\varepsilon + 2^{-\kappa} + 2^{-\Omega(t)}).
 \end{aligned}$$

$\text{Ext}((\tilde{X})^t, U_t)$ can be generated using a seed of length

$$d = d_t + t = t \cdot d_\ell + t = O(t \cdot \ell \cdot n),$$

and it has an output length of

$$\begin{aligned} d' &= m' \cdot (\alpha_t - \kappa) + t \\ &= m' \cdot (t \cdot \alpha_\ell - \Gamma(t, \kappa) - \kappa) + t \\ &\geq t \cdot d_\ell + t\Delta\ell/2 - m' \cdot (\Gamma(t, \kappa) + \kappa) + t \\ &= d + t\Delta\ell/2 - m' \cdot O(\sqrt{t\kappa} \cdot \log t + \kappa) \\ &\geq d + t\Delta\ell/4 = (1 + \Omega(\Delta/n)) \cdot d, \end{aligned}$$

where the last inequality follows from an appropriate setting of

$$t = O\left(\left(\frac{m'}{\Delta\ell}\right)^2 \cdot \kappa \cdot \log^2\left(\frac{m'\kappa}{\Delta\ell}\right)\right) = O\left(\frac{m^2 \cdot \kappa \cdot \log^2 n}{\Delta^2}\right).$$

Finally, we can bound the seed length d by

$$d = O(t \cdot n \cdot \ell) = O\left(\frac{n^2 \cdot m^2 \cdot \kappa \cdot \log^2 n}{\Delta^3}\right). \quad \square$$

6. Deducing the main results. We are now ready to prove the main result of the paper.

THEOREM 6.1 (pseudorandom generators from one-way functions). *Let n be a security parameter and $f : \{0, 1\}^n \mapsto \{0, 1\}^n$ a polynomial-time computable function. For all poly(n)-time computable functions $\varepsilon = \varepsilon(n) \leq 1/n^c$ (where c is a universal constant) and $\kappa = \kappa(n) \in [n/4]$, there exists an efficient generator G from strings of length $d = d(n) = O(n^4 \cdot \kappa \cdot \log^2 n / \log^3(1/\varepsilon))$ to strings of length $d \cdot (1 + \Omega(\log(1/\varepsilon))/n)$ such that the following hold:*

Security. Assume that f is a (T, ε) -one-way function for $T = T(n)$. Then for every poly(n)-time computable function $\varepsilon' = \varepsilon'(n) \geq 2^{-\kappa}$, G is a $(T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon' \cdot n^{O(1)})$ -pseudorandom generator. Moreover, the reduction is fully black-box.

Complexity. G is computable in NC^1 with $O(d/n)$ (uniformly random) oracle calls to f .

Proof. We let c be the constant of Theorem 4.1. We start by applying Theorem 4.1 on f to get a generator $G_{nb} : \{0, 1\}^{d_1} \mapsto \{0, 1\}^{m_1}$ that has $(T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon')$ -next-block-pseudoentropy $k = d_1 + \log(1/\varepsilon)/2$, where $d_1, m_1 \in O(n)$.

In the next step we apply Theorem 5.1 with respect to the above G_{nb} , κ , and $\Delta := \log(1/\varepsilon)/2$ to get an efficient generator from strings of length

$$d = O\left(\frac{d_1^2 \cdot m_1^2 \cdot \kappa \cdot \log^2 n}{\Delta^3}\right) = O\left(\frac{n^4 \cdot \kappa \cdot \log^2 n}{\log^3(1/\varepsilon)}\right)$$

to strings of length $d \cdot (1 + \Omega(\log(1/\varepsilon))/n)$ that is $(T \cdot (\varepsilon'/n)^{O(1)} - n^{O(1)}, (\varepsilon_1 + 2^{-\kappa}) \cdot n^{O(1)}) = (T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon' \cdot n^{O(1)})$ -pseudorandom. \square

The above theorem yields the following important corollaries.

COROLLARY 6.2 (pseudorandom generator from one-way functions—polynomial security case). *Let n be a security parameter and $f : \{0, 1\}^n \mapsto \{0, 1\}^n$ a one-way function. Then there exists a pseudorandom generator G from strings of length $d = d(n) \in \tilde{O}(n^4)$ to strings of length $d \cdot (1 + \Omega((\log n)/n))$.*

Furthermore, the reduction is fully black-box. G is computable in NC^1 with $O(d/n)$ (uniformly random) oracle calls to f .

Proof. Applying Theorem 6.1 on f , $\varepsilon = 1/n^c$, and $\kappa = \log^2 n$, we get an efficient generator G of the stated input and output lengths. Assume now that G is not a pseudorandom generator. Namely, there exists $p \in \text{poly}$ such that G is not $(p(n), 1/p(n))$ -pseudorandom. Therefore, G is not $(T \cdot (\varepsilon'/n)^{O(1)}, \varepsilon' \cdot \text{poly}(n))$ -pseudorandom for $\varepsilon' := 1/p(n) \cdot \text{poly}(n) > 2^{-\kappa}$ and $T = \text{poly}(n) \cdot p(n)$. It follows that f is not a (T, ε) -one-way function, in contradiction. \square

COROLLARY 6.3 (pseudorandom generator from one-way functions—exponential hardness case). *Let $f: \{0, 1\}^n \mapsto \{0, 1\}^n$ be a $(2^{\Omega(n)}, 2^{-\Omega(n)})$ -one-way function. Then*

1. *there exists a $(2^{\Omega(n)}, 2^{-\Omega(\log^2 n)})$ -pseudorandom generator G from strings of length $d = d(n) \in \tilde{O}(n)$ to strings of length $d \cdot (1 + \Omega(1))$, and*
2. *there exists a $(2^{\Omega(n)}, 2^{-\Omega(n)})$ -pseudorandom generator G from strings of length $d = d(n) \in \tilde{O}(n^2)$ to strings of length $d \cdot (1 + \Omega(1))$.*

Furthermore, in both cases it holds that the reduction is fully black-box, and G is computable in NC^1 with $O(d/n)$ (uniformly random) oracle calls to f .

Proof. The proof is immediate from Theorem 6.1, taking $\kappa = \log^2(n)$ and $\kappa \in \Omega(n)$ in the first and second cases, respectively, and $\varepsilon' = 2^{-\kappa}$. \square

Appendix. The uniform hardcore lemma. In this section, we show how to deduce our version of the uniform hardcore lemma (Proposition 4.7) from what Holenstein proves. Holenstein’s statement of the hardcore lemma requires the functions $\gamma(n)$ and $\delta(n)$ to be computable in time $\text{poly}(n)$ and to be at least $1/\text{poly}(n)$. However, the same proof yields the following, where we give δ and γ as inputs to the algorithm M and allow a loss of $(\delta\gamma)^{O(1)}$ in the running time.

PROPOSITION A.1 (see [21, Thm. 6.8]). *Let n be a security parameter, and let $h: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ and $V: \{0, 1\}^n \mapsto \{0, 1\}$ be polynomial-time computable functions. Let $\delta = \delta(n) \in [\delta_0, 1]$ and $\gamma = \gamma(n) \in [0, 1] > 2^{-n/3}$. Assume that*

$$\Pr[M(\delta, \gamma, h(U_n)) = V(U_n)] \leq 1 - \delta/2 + \gamma^2\delta^5/8192$$

for every probabilistic algorithm M running in time $T = T(n)$ and large enough n . Then for every oracle-aided predictor P running in time $T \cdot (\gamma\delta/n)^{O(1)}$ and all sufficiently large n , there exists a set $L \subseteq \{0, 1\}^n$ of density at least δ such that

$$\Pr_{W \stackrel{R}{\leftarrow} L} [P^{\chi_L}(h(W)) = V(W)] < (1 + \gamma)/2,$$

where χ_L is the characteristic function of L , provided that all the queries of P to χ_L are computed independently of the input $h(W)$. Furthermore, the reduction is fully black-box.

Note that the hypothesis of Holenstein’s lemma is slightly stronger in that it works even if the success probability of M is slightly higher than $1 - \delta/2$; this slackness is convenient for deducing Proposition 4.7, which we do now.

Proof of Proposition 4.7. Assume that

$$(5) \quad \Pr[M(\delta_0, \gamma, h(U_n)) = V(U_n)] \leq 1 - \delta/2$$

for every probabilistic algorithm M running in time $T = T(n)$ and large enough n . Then we will argue that

$$(6) \quad \Pr[M'(\delta', \gamma, h(U_n)) = V(U_n)] \leq 1 - \delta'/2 + \alpha$$

for every probabilistic algorithm M' running in time $T' = T(n) \cdot (\delta_0\gamma)^{O(1)}$ and large enough n , where δ' is the smallest multiple of $\alpha = \delta_0^5\gamma^2/8192$ that is at least as

large as δ (so $\delta' \leq \delta + \alpha$). Suppose for contradiction that there is an M' violating inequality (6). Then we can obtain an M violating inequality (5) by trying $\delta' = i \cdot \alpha$ for each $i = 0, \dots, \lfloor 1/\alpha \rfloor$, estimating the success probability to within accuracy $\alpha/3$ (by random sampling), and finally running M' with the value of δ' that achieves maximum estimated success probability.

Since $\alpha \leq (\delta')^5 \gamma^2 / 8192$, we can apply Proposition A.1 to obtain, for every predictor P running in time $T' \cdot (\gamma \delta' / n)^{O(1)} = T \cdot (\gamma \delta_0 / n)^{O(1)}$, a hardcore set L of density $\delta' \geq \delta$. \square

Acknowledgments. We thank Benny Applebaum, Oded Goldreich, Thomas Holenstein, and Emanuele Viola for very helpful conversations.

REFERENCES

- [1] B. APPLEBAUM, Y. ISHAI, AND E. KUSHILEVITZ, *Cryptography in NC⁰*, SIAM J. Comput., 36 (2006), pp. 845–888.
- [2] M. BLUM AND S. MICALI, *How to generate cryptographically strong sequences of pseudo-random bits*, SIAM J. Comput., 13 (1984), pp. 850–864.
- [3] B. CHOR AND O. GOLDBREICH, *Unbiased bits from sources of weak randomness and probabilistic communication complexity*, SIAM J. Comput., 17 (1988), pp. 230–261.
- [4] U. FEIGE, A. FIAT, AND A. SHAMIR, *Zero-knowledge proofs of identity*, J. Cryptology, 1 (1988), pp. 77–94.
- [5] R. GENNARO, Y. GERTNER, J. KATZ, AND L. TREVISAN, *Bounds on the efficiency of generic cryptographic constructions*, SIAM J. Comput., 35 (2005), pp. 217–246.
- [6] O. GOLDBREICH, S. GOLDWASSER, AND S. MICALI, *How to construct random functions*, J. ACM, 33 (1986), pp. 792–807.
- [7] O. GOLDBREICH, H. KRAWCZYK, AND M. LUBY, *On the existence of pseudorandom generators*, SIAM J. Comput., 22 (1993), pp. 1163–1175.
- [8] O. GOLDBREICH AND L. A. LEVIN, *A hard-core predicate for all one-way functions*, in Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC), 1989, pp. 25–32.
- [9] O. GOLDBREICH, S. MICALI, AND A. WIGDERSON, *Proofs that yield nothing but their validity, or all languages in NP have zero-knowledge proof systems*, J. ACM, 38 (1991), pp. 691–729.
- [10] V. GURUSWAMI AND M. SUDAN, *List decoding algorithms for certain concatenated codes*, in Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC), 2000, pp. 181–190.
- [11] I. HAITNER, D. HARNIK, AND O. REINGOLD, *Efficient pseudorandom generators from exponentially hard one-way functions*, in Automata, Languages and Programming, Part II, Lecture Notes in Comput. Sci. 4052, Springer, Berlin, 2006, pp. 228–239.
- [12] I. HAITNER, D. HARNIK, AND O. REINGOLD, *On the power of the randomized iterate*, in Advances in Cryptology—CRYPTO 2006, Springer, Berlin, 2006, pp. 22–40.
- [13] I. HAITNER, D. HARNIK, AND O. REINGOLD, *On the power of the randomized iterate*, SIAM J. Comput., 40 (2011), pp. 1486–1528.
- [14] I. HAITNER, M.-H. NGUYEN, S. J. ONG, O. REINGOLD, AND S. VADHAN, *Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function*, SIAM J. Comput., 39 (2009), pp. 1153–1218.
- [15] I. HAITNER, O. REINGOLD, AND S. VADHAN, *Efficiency Improvements in Constructing Pseudorandom Generators from One-Way Functions*, Technical Report TR10-089, Electronic Colloquium on Computational Complexity, 2010.
- [16] I. HAITNER, O. REINGOLD, AND S. VADHAN, *Efficiency improvements in constructing pseudorandom generators from one-way functions*, in Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), 2010, pp. 437–446.
- [17] I. HAITNER, O. REINGOLD, S. VADHAN, AND H. WEE, *Inaccessible entropy*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), 2009, pp. 611–620.
- [18] J. HÅSTAD, R. IMPAGLIAZZO, L. A. LEVIN, AND M. LUBY, *A pseudorandom generator from any one-way function*, SIAM J. Comput., 28 (1999), pp. 1364–1396.
- [19] T. HOLENSTEIN, *Key agreement from weak bit agreement*, in Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), 2005, pp. 664–673.
- [20] T. HOLENSTEIN, *Pseudorandom generators from one-way functions: A simple construction for any hardness*, in Theory of Cryptography, Lecture Notes in Comput. Sci. 3876, Springer, Berlin, 2006, pp. 443–461.

- [21] T. HOLENSTEIN, *Strengthening Key Agreement Using Hard-Core Sets*, Ph.D. thesis, Department of Computer Science, ETH Zürich, Zürich, Switzerland, 2006.
- [22] T. HOLENSTEIN AND M. SINHA, *Constructing a Pseudorandom Generator Requires an Almost Linear Number of Calls*, preprint, arXiv:1205.4576v1 [cs.CR], 2012.
- [23] T. HOLENSTEIN AND M. SINHA, *Constructing a pseudorandom generator requires an almost linear number of calls*, in Proceedings of the 53rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 2012, pp. 698–707.
- [24] R. IMPAGLIAZZO, *Hard-core distributions for somewhat hard problems*, in Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1995, pp. 538–545.
- [25] 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 (STOC), 1989, pp. 12–24.
- [26] R. IMPAGLIAZZO AND D. ZUCKERMAN, *How to recycle random bits*, in Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1989, pp. 248–253.
- [27] M. LUBY AND C. RACKOFF, *How to construct pseudorandom permutations from pseudorandom functions*, SIAM J. Comput., 17 (1988), pp. 373–386.
- [28] M. NAOR, *Bit commitment using pseudorandomness*, J. Cryptology, 4 (1991), pp. 151–158.
- [29] A. A. RAZBOROV AND S. RUDICH, *Natural proofs*, J. Comput. System Sci., 55 (1997), pp. 24–35.
- [30] O. REINGOLD, L. TREVISAN, AND S. P. VADHAN, *Notions of reducibility between cryptographic primitives*, in Theory of Cryptography, Lecture Notes in Comput. Sci. 2951, Springer, Berlin, 2004, pp. 1–20.
- [31] M. SUDAN, L. TREVISAN, AND S. VADHAN, *Pseudorandom generators without the XOR lemma*, J. Comput. System Sci., 62 (2001), pp. 236–266.
- [32] S. P. VADHAN, *Constructing locally computable extractors and cryptosystems in the bounded-storage model*, J. Cryptology, 17 (2004), pp. 43–77.
- [33] S. VADHAN AND C. J. ZHENG, *Characterizing Pseudoentropy and Simplifying Pseudorandom Generator Constructions*, Technical Report TR11-141, ECCO, 2011; available online from <http://eccc.hpi-web.de/report/2011/141/>.
- [34] S. VADHAN AND C. J. ZHENG, *Characterizing pseudoentropy and simplifying pseudorandom generator constructions*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012, pp. 817–836.
- [35] L. G. VALIANT, *A theory of the learnable*, Commun. ACM, 27 (1984), pp. 1134–1142.
- [36] A. C. YAO, *Theory and applications of trapdoor functions*, in Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1982, pp. 80–91.
- [37] D. ZUCKERMAN, *Simulating BPP using a general weak random source*, Algorithmica, 16 (1996), pp. 367–391.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.