

Pseudorandom Generators from Regular One-Way Functions: New Constructions with Improved Parameters

Yu Yu^{1,2}, Xiangxue Li², and Jian Weng³

¹ Institute for Theoretical Computer Science, Institute for Interdisciplinary Information Sciences, Tsinghua University, P.R. China

² Department of Computer Science, East China Normal University, P.R. China

³ Department of Computer Science, Jinan University, Guangzhou 510632, P.R. China

Abstract. We revisit the problem of basing pseudorandom generators on regular one-way functions, and present the following constructions:

- For any known-regular one-way function (on n -bit inputs) that is known to be ε -hard to invert, we give a neat (and tighter) proof for the folklore construction of pseudorandom generator of seed length $\Theta(n)$ by making a single call to the underlying one-way function.
- For any unknown-regular one-way function with known ε -hardness, we give a new construction with seed length $\Theta(n)$ and $O(n/\log(1/\varepsilon))$ calls. Here the number of calls is also optimal by matching the lower bounds of Holenstein and Sinha (FOCS 2012).

Both constructions require the knowledge about ε , but the dependency can be removed while keeping nearly the same parameters. In the latter case, we get a construction of pseudo-random generator from any unknown-regular one-way function using seed length $\tilde{O}(n)$ and $\tilde{O}(n/\log n)$ calls, where \tilde{O} omits a factor that can be made arbitrarily close to constant (e.g. $\log \log \log n$ or even less). This improves the *randomized iterate* approach by Haitner, Harnik and Reingold (CRYPTO 2006) which requires seed length $O(n \cdot \log n)$ and $O(n/\log n)$ calls.

1 Introduction

The seminal work of Håstad, Impagliazzo, Levin and Luby (HILL) [14] that one-way functions (OWFs) imply pseudorandom generators (PRGs) constitutes one of the centerpieces of modern cryptography. Technical tools and concepts (e.g. pseudo-entropy, leftover hash lemma) developed and introduced in [14] were found useful in many other contexts (such as leakage-resilient cryptography). Nevertheless, a major drawback of [14] is that the construction is quite involved and too inefficient to be of any practical use, namely, to obtain a PRG with comparable security to the underlying OWF on security parameter n , one needs a seed of length $O(n^8)$ ¹. Research efforts (see [15,13,23], just to name a few) have been followed up towards simplifying and improving the constructions, and

¹ More precisely, the main construction of [14] requires seed length $O(n^{10})$, but [14] also sketches another construction of seed length $O(n^8)$, which was proven in [15].

the current state-of-the-art construction [23] requires seed length $O(n^3)$. Let us mention all aforementioned approaches are characterized by a parallel construction, namely, they run sufficiently many independent copies of the underlying OWFs (rather than running a single trail and feeding its output back to the input iteratively) and there seems an inherent lower bound on the number of copies needed. This is recently formalized by Holenstein and Sinha [16], in particular, they showed that any black-box construction of a PRG from an arbitrary OWF f requires $\Omega(n/\log n)$ calls to f in general.²

PRGS FROM SPECIAL OWFS. Another line of research focuses on OWFs with special structures that give rise to more efficient PRGs. Blum, Micali [3] and Yao [26] independently introduced the notion of PRGs, and observed that PRGs can be efficiently constructed from one-way permutations (OWPs). That is, given a OWP f on input x and its hardcore function h_c (e.g. by Goldreich and Levin [10]), a single invocation of f already implies a PRG $g(x) = (f(x), h_c(x))$ with a stretch³ of $\Omega(\log n)$ bits and it extends to arbitrary stretch by repeated iterations (seen by a hybrid argument):

$$g^\ell(x) = (h_c(x), h_c(f^1(x)), \dots, h_c(f^\ell(x)), \dots)$$

where $f^i(x) \stackrel{\text{def}}{=} f(f^{i-1}(x))$ and $f^1(x) \stackrel{\text{def}}{=} f(x)$. The above PRG, often referred to as the BMY generator, enjoys many advantages such as simplicity, optimal seed length, and minimal number of calls. Levin [19] observed that f is not necessarily a OWP, but it suffices to be one-way on its own iterate. Unfortunately, an arbitrary OWF doesn't have this property. Goldreich, Krawczyk, and Luby [9] assumed known-regular⁴ OWFs and gave a construction of seed length $O(n^3)$ by iterating the underlying OWFs and applying k -wise independent hashing in between every two iterations. Later Goldreich showed a more efficient (and nearly optimal) construction from known-regular OWFs in his textbook [7], where in the concrete security setting the construction does only a single call to the underlying OWF (or $\omega(1)$ calls in general). The construction was also implicit in many HILL-style constructions (e.g. [15,13]). Haitner, Harnik and Reingold [12] refined the technique used in [9] (which they called the *randomized iterate*) and adapted the construction to unknown regular OWFs with reduced seed length $O(n \cdot \log n)$. Informally, the randomized iterate follows the route of [9] and applies a random pairwise independent hash function h_i in between every two applications of f , i.e.

$$f^1(x) \stackrel{\text{def}}{=} f(x); \text{ for } i \geq 2 \text{ let } f^i(x; h_1, \dots, h_{i-1}) \stackrel{\text{def}}{=} f(h_{i-1}(f^{i-1}(x; h_1, \dots, h_{i-2})))$$

The key observation is “*the last iterate is hard-to-invert*” [11], more precisely, function f , when applied to $h_{i-1}(f^{i-1}; h_1, \dots, h_{i-2})$, is hard-to-invert even if

² The lower bound of [16] also holds in the concrete security setting, namely, $\Omega(n/\log(1/\varepsilon))$ calls from any ε -hard OWF.

³ The stretch of a PRG refers to the difference between output and input lengths.

⁴ A function $f(x)$ is regular if the every image has the same number (say α) of preimages, and it is known- (resp., unknown-) regular if α is efficiently computable (resp., inefficient to approximate) from the security parameter.

h_1, \dots, h_{i-1} are made public. The generator follows by running the iterate $O(n/\log n)$ times, and outputting $\Omega(\log n)$ hardcore bits per iteration, which requires seed length $O(n^2/\log n)$ and can be further pushed to $O(n \cdot \log n)$ using derandomization techniques (e.g., Nisan’s bounded-space generator [20]). The randomized iterate matches the lower bound on the number of OWF calls⁵, but it remains open if any efficient construction can achieve linear seed length and $O(n/\log n)$ OWF calls simultaneously.

SUMMARY OF CONTRIBUTIONS. We contribute an alternative proof for the folklore construction of PRGs from known-regular OWFs via the notion of unpredictability pseudo-entropy, which significantly simplifies and tightens the proofs in [7]. We also give a new construction from any unknown-regular one-way function using seed length $\tilde{O}(n)$ and making $\tilde{O}(n/\log n)$ calls, where both parameters are optimal in the concrete security setting and nearly optimal in general (up to an arbitrarily close to constant factor), and this improves the randomized iterate [11]. We sketch both constructions as follows.

ENTROPY OBSERVATION. We start by assuming a (t, ε) -OWF f (see Definition 2) with known regularity 2^k (i.e., every image has 2^k preimages under f). The key observation is that for uniform X (over $\{0, 1\}^n$) we have X given $f(X)$ has $k + \log(1/\varepsilon)$ bits of pseudo-entropy (defined by the game below and formally in Definition 5). That is, no adversary A of running time t can win the following game against the challenger C with probability greater than $(2^{-k} \cdot \varepsilon)$. The rationale is that conditioned on any $f(X) = y$ random variable X is uni-




Fig. 1. The interactive game between A and C that defines unpredictability pseudo-entropy, where $x \leftarrow U_n$ denotes sampling a random $x \in \{0, 1\}^n$

formly distributed on set $f^{-1}(y) \stackrel{\text{def}}{=} \{x : f(x) = y\}$ of size 2^k , and thus even if any deterministic (or probabilistic) A recovers a $x' \in f^{-1}(y)$, the probability that $X = x'$ is only 2^{-k} .

PRGS FROM KNOWN-REGULAR OWFS. Given the above observation, we immediately obtain the following folklore construction using three extractions along with a three-line proof.

⁵ As explicitly stated in [16], the lower bound of $\Omega(n/\log n)$ calls also applies to unknown regular OWFs.

- RANDOMNESS EXTRACTION FROM $F(X)$. $f(X)$ has min-entropy $n - k$, and thus we can extract nearly $n - k$ statistically random bits.
- RANDOMNESS EXTRACTION FROM X . X has min-entropy k given any $y = f(X)$, so we can extract another k statistically random bits.
- PSEUDORANDOMNESS EXTRACTION FROM X . The second extraction reduces the unpredictability pseudo-entropy of X given $f(X)$ by no more than k (i.e., $\log(1/\varepsilon)$ bits remaining by the entropy chain rule), and hence we use Goldreich-Levin hardcore functions [10] to extract another $O(\log(1/\varepsilon))$ bits.

NOVELTY OF OUR ANALYSIS. While the construction was already known in literature (explicit in [7] and implicit in HILL-style generators [14,13,23]), its full proof was only seen in [7] and we further simplify the analysis via the use of unpredictability pseudo-entropy. In addition to simplicity, our technique can also be used to refine and tighten the proof given in [7] (see Section 3.2 and Remark 2 for details). We mention that our proofs sketched above are not implied by the recent work of [13,23]. In particular, the construction of [13] applies a dedicated universal hash function h of description length $O(n^2)$ to $f(X)$ such that the first k output bits are statistically random and the next $O(\log n)$ bits are computationally random, and this holds even if k is unknown (which is desired for a general OWF f whose preimage size may vary for different images). However, in our context k is known and it is crucial that the description length of the hash functions is linear, for which we do two extractions from $f(X)$ using h_2 and h_c respectively. We also stress that our observation that “ X given $f(X)$ has unpredictability pseudo-entropy $k + \log(1/\varepsilon)$ ” is incomparable with the counterpart in [23, Thm 1.5], which was informally stated as “ $(f(X), X)$ has next-bit pseudo-entropy $n + \Omega(\log n)$ ”. Firstly, our proof enjoys simplicity and tightness whereas theirs employs the uniform version of Min-Max Theorem which is much involved and interesting in its own right. Secondly, next-bit pseudo-entropy was a newly introduced notion [13] and whether it implies (or is implied by) unpredictability pseudo-entropy is unknown to our knowledge, and the ways of extraction from these two sources are different. See Section 3.2 and Remark 3 for details.

CONCRETE VS. ASYMPTOTIC SECURITY. The above construction is optimal (in seed length and the number of OWF calls), but requires the knowledge about parameter ε , more precisely, we need ε to decide entropy loss d such that the first extraction outputs $n - k - d$ bits with statistical error bounded by $2^{-d/2}$ (by the Leftover Hash Lemma [14]) and let the third extraction output more than d bits to achieve a positive stretch. It is unknown how to remove the dependency on ε for free (see also the discussions in [7]). Fortunately, there is a known repetition trick to solve the problem using seed length $\tilde{O}(n)$ and $\tilde{O}(1)$ OWF calls, where notation \tilde{O} omits a factor of $q \in \omega(1)$ (i.e. q can be any factor arbitrarily close to constant such as $\log \log n$).

PRGs FROM UNKNOWN-REGULAR OWFS. We also give a new construction oblivious of the regularity of f . The construction follows the steps below.

- CONVERT TO KNOWN REGULARITY. The key idea is to transform any unknown regular OWF into another known regular OWF (over a special domain). That is, for a (length-preserving) unknown-regular (t, ε) -OWF $f : \{0, 1\}^n \rightarrow \mathcal{Y}$ where $\mathcal{Y} \subseteq \{0, 1\}^n$ is the range of f , define function $\bar{f} : \mathcal{Y} \times \{0, 1\}^n \rightarrow \mathcal{Y}$ as $\bar{f}(y, r) \stackrel{\text{def}}{=} f(y \oplus r)$ where “ \oplus ” denotes bitwise XOR. It is not hard to see that \bar{f} has regularity 2^n (regardless of the regularity of f) and it preserves the hardness of f .
- CONSTRUCT A Z -SEEDED PRG. Similar to that observed in the 1st construction, $\bar{f}(Y, R)$ hides $n + \log(1/\varepsilon)$ bits of pseudo-entropy about (Y, R) , and thus we can extract $n + O(\log(1/\varepsilon))$ pseudorandom bits, namely, we get a special PRG \bar{g} that maps random elements over $\mathcal{Y} \times \{0, 1\}^n$ to pseudorandom ones over $\mathcal{Y} \times \{0, 1\}^{n+O(\log(1/\varepsilon))}$. This PRG is known as the “ Z -seeded PRG” [23], one that given input distribution Z outputs (Z', U_s) which is computationally indistinguishable from (Z, U_s) , where in the above case $Z = (Y, R)$ and stretch $s = O(\log(1/\varepsilon))$. Note that if Z were U_{2n} then this would be a standard PRG.
- ITERATIVE COMPOSITION OF Z -SEEDED PRG. Nevertheless, to use the above Z -seeded PRG \bar{g} we need to efficiently sample from $Y = f(U_n)$ (i.e. uniform distribution over \mathcal{Y}), which costs n random bits despite that the entropy of Y may be far less than n . Quite naturally (following [23,3]), the construction invests n bits (to sample a random $y \leftarrow f(U_n)$) at initialization, runs \bar{g} in iterations, and outputs $O(\log(1/\varepsilon))$ bits per iteration. The stretch becomes positive after $O(n/\log(1/\varepsilon))$ iterations, which matches the lower bounds of [16]. The seed length remains of order $\Theta(n)$ by reusing the coins for universal hash and G-L functions at every iteration, thanks to the hybrid argument.
- REMOVE DEPENDENCY ON ε (OPTIONAL). Similarly, in case that ε is unknown, we pay a penalty factor $\tilde{O}(1)$ for using the repetition trick. That is, we construct a PRG from any unknown-regular OWF using seed length $\tilde{O}(n)$ and $\tilde{O}(n/\log n)$ OWF calls.

2 Preliminaries

NOTATIONS AND DEFINITIONS. We use capital letters (e.g. X, Y, A) for random variables, standard letters (e.g. x, y, a) for values, and calligraphic letters (e.g. $\mathcal{X}, \mathcal{Y}, \mathcal{S}$) for sets. $|\mathcal{S}|$ denotes the cardinality of set \mathcal{S} . For function f , we let $f(\mathcal{X}) \stackrel{\text{def}}{=} \{f(x) : x \in \mathcal{X}\}$ be the set of images that are mapped from \mathcal{X} under f , and denote by $f^{-1}(y)$ the set of y 's preimages under f , i.e. $f^{-1}(y) \stackrel{\text{def}}{=} \{x : f(x) = y\}$. We say that distribution X is flat if it is uniformly distributed over some set \mathcal{X} . We use $s \leftarrow S$ to denote sampling an element s according to distribution S , and let $s \leftarrow \mathcal{S}$ denote sampling s uniformly from set \mathcal{S} , and $y := f(x)$ denote value assignment. We use U_n to denote the flat distribution over $\{0, 1\}^n$ independent of the rest random variables in consideration, and let

$f(U_n)$ be the distribution induced by applying function f to U_n . We use $\text{CP}(X)$ to denote the collision probability of X , i.e., $\text{CP}(X) \stackrel{\text{def}}{=} \sum_x \Pr[X = x]^2$, and collision entropy $\mathbf{H}_2(X) \stackrel{\text{def}}{=} -\log \text{CP}(X) \geq \mathbf{H}_\infty(X)$. We also define average (aka conditional) collision entropy and average min-entropy of a random variable X conditioned on another random variable Z by

$$\begin{aligned} \mathbf{H}_2(X|Z) &\stackrel{\text{def}}{=} -\log \left(\mathbb{E}_{z \leftarrow Z} \left[\sum_x \Pr[X = x|Z = z]^2 \right] \right) \\ \mathbf{H}_\infty(X|Z) &\stackrel{\text{def}}{=} -\log \left(\mathbb{E}_{z \leftarrow Z} \left[\max_x \Pr[X = x|Z = z] \right] \right) \end{aligned}$$

An entropy source refers to a random variable that has some non-trivial amount of entropy. A function $\mu : \mathbb{N} \rightarrow [0, 1]$ is negligible if for every polynomial poly we have $\mu(n) < 1/\text{poly}(n)$ holds for all sufficiently large n 's. We define the *computational distance* between distribution ensembles $X \stackrel{\text{def}}{=} \{X_n\}_{n \in \mathbb{N}}$ and $Y \stackrel{\text{def}}{=} \{Y_n\}_{n \in \mathbb{N}}$ as follows: we say that X and Y are $(t(n), \varepsilon(n))$ -close, denoted by $\text{CD}_{t(n)}(X, Y) \leq \varepsilon(n)$, if for every probabilistic distinguisher D of running time up to $t(n)$ it holds that

$$| \Pr[D(1^n, X_n) = 1] - \Pr[D(1^n, Y_n) = 1] | \leq \varepsilon(n) .$$

The *statistical distance* between X and Y , denoted by $\text{SD}(X, Y)$, is defined by

$$\text{SD}(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_x |\Pr[X = x] - \Pr[Y = x]| = \text{CD}_\infty(X, Y)$$

We use $\text{SD}(X, Y|Z)$ (resp. $\text{CD}_t(X, Y|Z)$) as shorthand for $\text{SD}((X, Z), (Y, Z))$ (resp. $\text{CD}_t((X, Z), (Y, Z))$).

SIMPLIFYING ASSUMPTIONS AND NOTATIONS. To simplify the presentation, we make the following assumptions without loss of generality. It is folklore that one-way functions can be assumed to be length-preserving (see [12] for formal proofs). Throughout, most parameters are functions of the security parameter n (e.g., $t(n)$, $\varepsilon(n)$, $\alpha(n)$) and we often omit n when clear from the context (e.g., t , ε , α). Parameters (e.g. ε , α) are said to be known if they are known to be polynomial-time computable from n . By notation $f : \{0, 1\}^n \rightarrow \{0, 1\}^l$ we refer to the ensemble of functions $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^{l(n)}\}_{n \in \mathbb{N}}$. As slight abuse of notion, *poly* might be referring to the set of all polynomials or a certain polynomial, and h might be either a function or its description, which will be clear from the context.

Definition 1 (universal hash functions [4]). A family of functions $\mathcal{H} \stackrel{\text{def}}{=} \{h : \{0, 1\}^n \rightarrow \{0, 1\}^l\}$ is called a universal hash family, if for any $x_1 \neq x_2 \in \{0, 1\}^n$ we have $\Pr_{h \leftarrow \mathcal{H}}[h(x_1) = h(x_2)] \leq 2^{-l}$.

Definition 2 (one-way functions). A function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{l(n)}$ is $(t(n), \varepsilon(n))$ -one-way if f is polynomial-time computable and for any probabilistic algorithm A of running time $t(n)$

$$\Pr_{y \leftarrow f(U_n)} [A(1^n, y) \in f^{-1}(y)] \leq \varepsilon(n).$$

For $\varepsilon(n) = 1/t(n)$, we simply say that f is $\varepsilon(n)$ -hard. f is a one-way function if it is $\varepsilon(n)$ -hard for some negligible function $\varepsilon(n)$.

Definition 3 (regular functions). A function f is α -regular if there exists an integer function α , called the regularity function, such that for every $n \in \mathbb{N}$ and $x \in \{0, 1\}^n$ we have

$$|f^{-1}(f(x))| = \alpha(n).$$

In particular, f is known-regular if α is polynomial-time computable, or called unknown-regular otherwise. Further, f is a (known-/unknown-) regular OWF if f is a OWF with (known/unknown) regularity.

Definition 4 (pseudorandom generators [3,26]). A function $g : \{0, 1\}^n \rightarrow \{0, 1\}^{l(n)}$ ($l(n) > n$) is a $(t(n), \varepsilon(n))$ -secure PRG if g is polynomial-time computable and

$$CD_{t(n)}(g(1^n, U_n), U_{l(n)}) \leq \varepsilon(n).$$

where $(l(n) - n)$ is the stretch of g , and we often omit 1^n (security parameter in unary) from g 's parameter list. We say that g is a pseudorandom generator if both $1/t(n)$ and $\varepsilon(n)$ are negligible.

Definition 5 (unpredictability pseudo-entropy[2,17]). For distribution ensemble $(X, Z) \stackrel{\text{def}}{=} \{(X_n, Z_n)\}_{n \in \mathbb{N}}$, we say that X has $k(n)$ bits of pseudo-entropy conditioned on Z for all $t(n)$ -time adversaries, denoted by $\mathbf{H}_{t(n)}(X|Z) \geq k(n)$, if for any $n \in \mathbb{N}$ and any probabilistic adversary A of running time $t(n)$

$$\Pr_{(x,z) \leftarrow (X_n, Z_n)} [A(1^n, z) = x] \leq 2^{-k(n)}$$

Alternatively, we say that X is $2^{-k(n)}$ -hard to predict given Z for all $t(n)$ -time adversaries.

Unpredictability pseudo-entropy can be seen as a relaxed form of min-entropy by weakening adversary's running time from unbounded to parameter $t(n)$, which (presumably) characterizes the class of practical adversaries we care about. Note that the notion seems only meaningful in its conditional form as otherwise (when Z is empty) non-uniform attackers can simply hardwire the best guess about X , and thus $\mathbf{H}_{t(n)}$ collapses to \mathbf{H}_∞ . Let us mention the unpredictability pseudo-entropy is different from (and in fact, strictly weaker than [2,17]) the HILL pseudo-entropy [14], which is another relaxed notion of min-entropy by considering its computationally indistinguishable analogues.

3 Pseudorandom Generators from Regular One-Way Functions

3.1 Technical Tools

The first technical tool we use is the leftover hash lemma. Informally, it states that when applying a random universal hash function to min-entropy (or Rényi entropy) source, one obtain random strings that are statistical close to uniform even

conditioned on the description of hash function. The objects were later formalized as randomness extractors [21]. Universal hash functions are also good condensers (whose outputs have nearly maximal entropy) for a wider range of parameters than extractors.

Lemma 1 (leftover hash lemma [14]). *For any integers $d < k \leq n$, there exists a (efficiently computable) universal hash function family $\mathcal{H} \stackrel{\text{def}}{=} \{h : \{0, 1\}^n \rightarrow \{0, 1\}^{k-d}\}$ such that for any joint distribution (X, Z) where $X \in \{0, 1\}^n$ and $\mathbf{H}_2(X|Z) \geq k$, we have*

$$\text{SD}(H(X), U_{k-d} \mid H, Z) \leq 2^{-\frac{d}{2}}$$

where H is uniformly distributed over the members of \mathcal{H} , the description size of H is called seed length, and d is called entropy loss, i.e., the difference between the entropy of X (given Z) and the number of bits that were extracted from X .

Lemma 2 (condensers from hash functions). *Let $\mathcal{H} \stackrel{\text{def}}{=} \{h : \{0, 1\}^n \rightarrow \{0, 1\}^k\}$ be any universal hash function family and let (X, Z) be any random variable with $X \in \{0, 1\}^n$ and $\mathbf{H}_2(X|Z) \geq k$. Then, for H uniform distributed over \mathcal{H} we have $\mathbf{H}_2(H(X) \mid H, Z) \geq k - 1$.*

Proof. Let X_1 and X_2 be i.i.d. to $X \mid Z = z$ (i.e. X conditioned on $Z = z$).

$$\begin{aligned} 2^{-\mathbf{H}_2(H(X)|H,Z)} &= \mathbb{E}_{h \leftarrow H, z \leftarrow Z} \left[\Pr_{x_1 \leftarrow X_1, x_2 \leftarrow X_2} [H(x_1) = H(x_2) \mid H = h, Z = z] \right] \\ &\leq \mathbb{E}_{z \leftarrow Z} \left[\Pr_{x_1 \leftarrow X_1, x_2 \leftarrow X_2} [x_1 = x_2 \mid Z = z] \right] \\ &+ \mathbb{E}_{z \leftarrow Z} \left[\Pr_{h \leftarrow H} [h(x_1) = h(x_2) \mid x_1 \neq x_2, Z = z] \right] \\ &\leq 2^{-k} + 2^{-k} = 2^{-(k-1)}. \end{aligned}$$

We refer to [22,6,18] for extremely efficient constructions of universal hash functions with short description (of length $\Theta(n)$), such as multiplications between matrices and vectors, or over finite fields.

RECONSTRUCTIVE EXTRACTORS. We will also need objects that extract pseudorandomness from unpredictability pseudo-entropy sources. Unfortunately, the leftover hash lemma (and randomness extractors [21] in general) does not serve the purpose. Goldreich and Levin [10] showed that the inner product function is a reconstructive bit-extractor for unpredictability pseudo-entropy sources. Further, there are two ways to extend the inner product to multiple-bit extractors: (1) multiplication with a random matrix of length $O(n^2)$ and extracts almost all entropy (by a hybrid argument); (2) multiplication with a random Toeplitz matrix of length $\Theta(n)$ and extracts $O(\log(1/\epsilon))$ bits (due to Vazirani’s XOR lemma [25,10]). We will use the latter multi-bit variant (as stated below) to keep the seed length linear. Interestingly, the Toeplitz matrix based functions also constitute pairwise independent and universal hash function families.

Theorem 1 (Goldreich-Levin [10]). For distribution ensemble $(X, Y) \in \{0, 1\}^n \times \{0, 1\}^*$, and for any integer $m \leq n$, there exists⁶ a function family $\mathcal{H}_C \stackrel{\text{def}}{=} \{h_c : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ of description size $\Theta(n)$, such that

– If $Y = f(X)$ for any (t, ε) -OWF f and X uniform over $\{0, 1\}^n$, then we have

$$CD_{t'}(H_C(X), U_m \mid Y, H_C) \in O(2^m \cdot \varepsilon) . \tag{1}$$

– If X is ε -hard to predict given Y for all t -time adversaries, namely, entropy condition satisfies $\mathbf{H}_t(X|Y) \geq \log(1/\varepsilon)$, then we have

$$CD_{t'}(H_C(X), U_m \mid Y, H_C) \in O(2^m \cdot (n \cdot \varepsilon)^{\frac{1}{3}}) . \tag{2}$$

where $t' = t \cdot (\varepsilon/n)^{O(1)}$ and function H_C is uniformly distributed over the members of \mathcal{H}_C .

Remark 1 (unpredictability vs. one-wayness). To see the difference between the two versions above, consider the interactive game in Figure 1, where by unpredictability A’s prediction is successful only if $x = x'$, but in contrast A inverts OWF f as long as he finds any x' satisfying $f(x') = y$. Recall that the proof of the theorem can be seen as an efficient local list decoding procedure for the Hadamard code, where in the former case the decoder returns a random member from the candidate list while in the latter case it goes through all candidates and outputs the one x' satisfying $f(x') = y$ (if exists). We refer to Goldreich’s exposition [8] for further details.

We recall two folklore facts below, namely the chain rule of unpredictability (pseudo-)entropy and the replacement inequality. Intuitively, any leakage $Y \in \{0, 1\}^l$ decreases the unpredictability about secret X by a factor of no more than 2^l , which can be seen by a simple reduction (e.g., by replacing Y with a random string). The replacement inequality states that any information that is (efficiently) computable from the knowledge of the adversary does not help further reduce the unpredictability (pseudo-)entropy of the secret in consideration.

Fact 1 (chain rule of entropies). For any joint distribution (X, Y, Z) where $Y \in \{0, 1\}^l$, we have

$$\begin{aligned} \mathbf{H}_\infty(X|Y, Z) &\geq \mathbf{H}_\infty(X|Z) - l , \\ \mathbf{H}_{t'}(X|Y, Z) &\geq \mathbf{H}_t(X|Z) - l , \end{aligned}$$

where $t' \approx t$.

Fact 2 (replacement inequalities). For any joint distribution (X, Y, Z) and any t_h -time computable function $h : \mathcal{Y} \rightarrow \{0, 1\}^*$, we have

$$\begin{aligned} \mathbf{H}_\infty(X|h(Y), h, Z) &\geq \mathbf{H}_\infty(X|Y, Z) , \\ \mathbf{H}_{t-t_h}(X|h(Y), h, Z) &\geq \mathbf{H}_t(X|Y, Z) . \end{aligned}$$

⁶ For example (see [10]), we can use an $m \times n$ Toeplitz matrix $a_{m,n}$ to describe the family of functions, i.e., $\mathcal{H}_C \stackrel{\text{def}}{=} \{h_c(x) \stackrel{\text{def}}{=} a_{m,n} \cdot x, \text{ where } x \in \{0, 1\}^n, a_{m,n} \in \{0, 1\}^{m+n-1}\}$.

3.2 PRGs from OWFs with Known Regularity and Hardness

We state our motivating observation as the lemma below.

Lemma 3 (regular OWFs imply unpredictability pseudo-entropy). *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a 2^k -regular (t, ε) -OWF. Then, we have*

$$\mathbf{H}_t(X \mid f(X)) \geq k + \log(1/\varepsilon) \quad , \tag{3}$$

where X is uniform over \mathcal{X} .

Proof. The (t, ε) -one-wayness of f guarantees that for any deterministic adversary A of running time t

$$\Pr_{x \leftarrow \mathcal{X}, y:=f(x)} [A(y) \in f^{-1}(y)] \leq \varepsilon$$

which in turn implies (as conditioned on $f(X) = y$, X is uniform over $f^{-1}(y)$ of size 2^k):

$$\Pr_{x \leftarrow \mathcal{X}, y:=f(x)} [A(y) = x] \leq 2^{-k} \cdot \varepsilon$$

which is essentially Equation (3) by taking a negative logarithm. Note that the above argument extends to probabilistic t -time A as well, by considering $A(y; r)$ on every fixing of his random coin r .

THE CONSTRUCTION FOR KNOWN α AND ε . As sketched in introduction, our first construction essentially extracts from the joint distribution $(X, f(X))$ three times, namely, use universal hash function h_1 to extract nearly (up to entropy loss) $n - k$ bits from $f(X)$, and then apply h_2 and h_c to extract k statistical random bits and another $\Theta(\log(1/\varepsilon n))$ pseudo-random bits from X respectively. For convenience, we assume without loss of generality that the regularity is a power of two, i.e., $\alpha = 2^k$.

Theorem 2 (preliminary construction based on known regularity and hardness). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a known 2^k -regular length-preserving (t, ε) -OWF, let d, s be any integer functions satisfying $9d + 6s = 2 \log(1/\varepsilon n)$, let $\mathcal{H}_1 \stackrel{\text{def}}{=} \{h_1 : \{0, 1\}^n \rightarrow \{0, 1\}^{n-k-d}\}$, $\mathcal{H}_2 \stackrel{\text{def}}{=} \{h_2 : \{0, 1\}^n \rightarrow \{0, 1\}^k\}$ be universal hash function families, let $\mathcal{H}_C \stackrel{\text{def}}{=} \{h_c : \{0, 1\}^n \rightarrow \{0, 1\}^{d+s}\}$ be a Goldreich-Levin function family, and let g be*

$$g : \{0, 1\}^n \times \mathcal{H}_1 \times \mathcal{H}_2 \times \mathcal{H}_C \rightarrow \{0, 1\}^{(n-k-d)+k+(d+s)} \times \mathcal{H}_1 \times \mathcal{H}_2 \times \mathcal{H}_C$$

$$(x, h_1, h_2, h_c) \mapsto (h_1(f(x)), h_2(x), h_c(x), h_1, h_2, h_c)$$

where $x \in \{0, 1\}^n$, $h_1 \in \mathcal{H}_1$, $h_2 \in \mathcal{H}_2$, $h_c \in \mathcal{H}_C$. Then, g is a $(t \cdot (\varepsilon/n)^{O(1)}, O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{3}}))$ -secure PRG with stretch s .

We deal with the situation where $n - k - d \leq 0$ by letting h_1 output nothing. Another special case $k = 0$ (i.e., f is a OWP) is handled by letting h_1 and h_2 output the identity and empty strings respectively.

Proof. The entropy conditions for the (pseudo)-randomness extractions are guaranteed by Lemma 4. We have by Equation (4), Equation (5) and the leftover hash lemma that the first $n - d$ bits extracted are statistically random, namely,

$$\begin{aligned} & \text{SD}((H_1(f(X)), H_2(X)), U_{n-d} \mid H_1, H_2) \\ & \leq \text{SD}(H_1(f(X)), U_{n-k-d} \mid H_1) + \text{SD}(H_2(X), U_k \mid H_1(f(X)), H_1, H_2) \\ & \leq 2 \cdot 2^{-\frac{d}{2}} = 2 \cdot 2^{\frac{d}{3} + \frac{1}{9} \log(\varepsilon n)} = O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{9}}) \end{aligned}$$

Next, as stated in Equation (6), conditioned on the prefix of $n - d$ random bits (and the seeds used), X remains $(t - n^{O(1)}, \varepsilon)$ -hard to predict, and thus by Goldreich-Levin (Theorem 1)

$$\begin{aligned} \text{CD}_{t'}(H_C(X), U_{d+s} \mid H_1(f(X)), H_2(X), H_1, H_2, H_C) &= O(2^{d+s} \cdot (n \cdot \varepsilon)^{\frac{1}{3}}) \\ &= O(2^{-\frac{d}{2}}) = O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{9}}) \end{aligned}$$

holds for $t' = t \cdot (\varepsilon/n)^{O(1)}$. The conclusion follows by a triangle inequality.

Lemma 4 (entropy conditions). *Let $f, \mathcal{H}_1, \mathcal{H}_2$ be defined as in Theorem 2, we have*

$$\mathbf{H}_\infty(f(X)) = n - k \quad , \tag{4}$$

$$\mathbf{H}_\infty(X \mid h_1(f(X)), h_1) \geq \mathbf{H}_\infty(X) - (n - k - d) = k + d \quad , \tag{5}$$

$$\mathbf{H}_{t-n^{O(1)}}(X \mid h_1(f(X)), h_2(X), h_1, h_2) \geq \mathbf{H}_t(X \mid f(X), h_2(X), h_2) \geq \log(1/\varepsilon) \tag{6}$$

hold for every $h_1 \in \mathcal{H}_1, h_2 \in \mathcal{H}_2$, and X uniform over $\{0, 1\}^n$.

Proof. Equation (4) follows from the regularity of f , i.e., every $y = f(x)$ has 2^k preimages, and thus $f(X)$ is uniformly distributed over a set of size 2^{n-k} . Equation (5) is due to the chain rule of min-entropy (see Fact 1). The first inequality of Equation (6) is the replacement inequality (see Fact 2), and the second one is obtained by applying the chain rule of unpredictability entropy to Equation (3), i.e., $\mathbf{H}_t(X \mid f(X), h_2(X), h_2) \geq \mathbf{H}_t(X \mid f(X)) - k = \log(1/\varepsilon)$.

Therefore, we already complete the proof for the PRG with linear seed length by doing a single call to any 2^k -regular ε -hard OWF provided that ε and k are known. We provide an alternative (and simpler) proof to that given by Goldreich [7] for essentially the same construction via unpredictability pseudo-entropy.

ON TIGHTENING SECURITY BOUNDS. Concretely, if the underlying OWF is $n^{-\log n}$ - (resp., $2^{-\frac{n}{3}}$ -) hard, then the outputs of the resulting PRG will be nearly $n^{-\frac{\log n}{9}}$ - (resp., $2^{-\frac{n}{27}}$ -) close to uniform (with respect to reasonably weakened adversaries than counterparts of the OWF). The main lossy step in the reduction is that we considered function $f'(x, h_2) \stackrel{\text{def}}{=} (f(x), h_2(x), h_2)$, where by Equation (6) X is (ε, t) -hard to predict given $f'(X)$ and thus we directly applied Equation (2) to get the inferior bounds. However, a closer look at f' suggests that it is almost 1-to-1, which implies that f' is a OWF (stated as in Lemma 5), which allows

us to use the tight version of Goldreich-Levin Theorem (see Equation (1)). This is actually the approach taken by [7], where however f' was only shown to be roughly $\varepsilon^{1/5}$ -hard (by checking the proof of [7, Prop 3.5.9]). We give a refined analysis below to get the tighter $\sqrt{\varepsilon}$ -hardness of f' , and this eventually leads to the improved construction as in Theorem 3.

Lemma 5 (unpredictability and almost 1-to-1 imply one-wayness). *Let f and \mathcal{H}_2 be as defined in Theorem 2, then function $f'(x, h_2) \stackrel{\text{def}}{=} (f(x), h_2(x), h_2)$ is a $(t, 3\sqrt{\varepsilon})$ - one-way function.*

Proof. Suppose for contradiction there exists A of running time t such that

$$\Pr [A(f'(X, H_2)) \in f'^{-1}(f'(X, H_2))] > 3\sqrt{\varepsilon}$$

Recall that $f(X)$ has min-entropy $n - k$ and conditioned on any $y = f(X)$ X has min-entropy k , and thus by the condensing property of universal hashing (see Lemma 2, setting $Z = f(X)$ and $H = H_2$) $\mathbf{H}_2(H_2(X)|H_2, f(X)) \geq k - 1$, which implies that $\text{CP}(f(X), H_2(X) | H_2) \leq 2^{-(n-k)} \cdot 2^{-(k-1)} = 2^{-(n-1)}$. It follows from Lemma 6 (setting $a = 2^{-n}/\sqrt{\varepsilon}$, $X_1 = (f(X), H_2(X))$, $Z_1 = H_2$) that $f'(X, H_2)$ hits set \mathcal{S} (defined below) with negligible probability, i.e., $\Pr[f'(X, H_2) \in \mathcal{S}] \leq 2\sqrt{\varepsilon}$ where

$$\begin{aligned} \mathcal{S} &\stackrel{\text{def}}{=} \{(y, w, h_2) : \Pr[(f(X), h_2(X)) = (y, w) | H_2 = h_2] \geq 2^{-n}/\sqrt{\varepsilon}\} \\ &= \{(y, w, h_2) : |f'^{-1}(y, w, h_2)| \geq 1/\sqrt{\varepsilon}\} . \end{aligned}$$

Then, let \mathcal{E} be the event that A inverts f' on any image whose preimage size is bounded by $1/\sqrt{\varepsilon}$, i.e., $\mathcal{E} \stackrel{\text{def}}{=} A(f'(X, H_2)) \in f'^{-1}(f'(X, H_2)) \wedge f'(X, H_2) \notin \mathcal{S}$

$$\begin{aligned} \Pr [A(f'(X, H_2)) = X] &\geq \Pr [\mathcal{E}] \cdot \Pr[A(f'(X, H_2)) = X | \mathcal{E}] \\ &> (3\sqrt{\varepsilon} - 2\sqrt{\varepsilon}) \cdot \left(\frac{1}{1/\sqrt{\varepsilon}}\right) = \varepsilon , \end{aligned}$$

where the probability of hard-to-invertness is related to unpredictability by the maximal preimage size. The conclusion follows by reaching a contradiction to the (t, ε) -unpredictability of X given $f'(X, H_2)$ (as stated in Equation (6)).

Lemma 6 (\mathbf{H}_2 implies \mathbf{H}_∞ with small slackness). *Let (X_1, Z_1) be a random variable, for $a > 0$ define $\mathcal{S}_a \stackrel{\text{def}}{=} \{(x, z) : \Pr[X_1 = x | Z_1 = z] \geq a\}$, it holds that $\Pr[(X_1, Z_1) \in \mathcal{S}_a] \leq \text{CP}(X_1 | Z_1)/a$.*

Proof. The proof is a typical Markov type argument.

$$\begin{aligned} \text{CP}(X_1 | Z_1) &= \mathbb{E}_{z \leftarrow Z_1} \left[\sum_x \Pr[X_1 = x | Z_1 = z]^2 \right] \\ &= \sum_{(x,z)} \Pr[(X_1, Z_1) = (x, z)] \cdot \Pr[X_1 = x | Z_1 = z] \\ &\geq \sum_{(x,z) \in \mathcal{S}_a} \Pr[(X_1, Z_1) = (x, z)] \cdot \Pr[X_1 = x | Z_1 = z] \\ &\geq a \cdot \Pr[(X_1, Z_1) \in \mathcal{S}_a] . \end{aligned}$$

Theorem 3 (improved construction based on known regularity and hardness). *For the same $f, g, \mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_C$ as assumed in Theorem 2 except that d and s satisfy $3d + 2s = \log(1/\varepsilon)$, we have that g is a $(t \cdot (\varepsilon/n)^{O(1)}, O((2^{2s} \cdot \varepsilon)^{1/6})$ -secure PRG with stretch s .*

Proof sketch. The proof is similar to Theorem 2. The first $n - d$ bits extracted are $2^{-d/2}$ -statistically random, conditioned on which the next $d + s$ bits are $O(2^{d+s} \sqrt{\varepsilon})$ -computationally random. It follows that the bound is $2^{-d/2} + O(2^{d+s} \sqrt{\varepsilon}) = O(2^{-d/2}) = O((2^{2s} \cdot \varepsilon)^{1/6})$. □

Remark 2 (a comparison with [7]). We provide an alternative (and much simplified) proof to the counterpart in [7]. Both approaches start by considering function $f'(x, h_2) = (f(x), h_2(x), h_2)$ and observing that distribution $f'(X, H_2)$ is nearly of full entropy (i.e. the amount of random bits used to sample $f'(X, H_2)$). The analysis of [7] then gets somewhat involved to show that f' is a $(t, O(\varepsilon^{1/5}))$ -OWF, and we simply apply the chain rule to get that the unpredictability pseudo-entropy about X given $f'(X, H)$ is at least $\log(1/\varepsilon)$ (see Lemma 4). Therefore, by Goldreich-Levin one can extract more bits from X to make a PRG. Combined with another idea that f' is nearly 1-to-1 and thus unpredictability implies one-wayness, our proof also implies a tighter version of [7], namely, f' is a $(t, 3\sqrt{\varepsilon})$ -OWF.

Remark 3 (next-bit vs. unpredictability pseudo-entropy). We mention that our observation that “ X given $f(X)$ has unpredictability pseudo-entropy $k + \log(1/\varepsilon)$ ” is incomparable with the counterpart⁷ in [23] that “ $(f(X), X)$ has next-bit pseudo-entropy $n + \Omega(\log n)$ ”. First, the proof of [23] is fundamentally different via the uniform version of Min-Max Theorem which is technically involved and useful in much broader contexts [24]. Secondly, there are no known reductions in relating unpredictability pseudo-entropy to next-bit pseudo-entropy from either directions, and in the former case one needs special extractors (that support reconstruction) while for the latter one needs to concatenate many copies of next-bit entropy sources and to extract many times (see [23, Figure 1]).

THREE EXTRACTIONS ARE NECESSARY. We argue that three extractions (using h_1, h_2 and h_c) seem necessary. One might think that the first two extractions (using h_1 and h_2) can be merged using a single universal hash function (that applies to the source $(X, f(X))$ and outputs $n - d$ bits). However, by doing so we cannot ensure the entropy condition (see Equation (6)) for the third extraction (using h_c). From another perspective, the merge would remove the dependency on the regularity and thus result in a generic construction that does a single call to any unknown regular OWFs, which is a contradiction to [16]. Furthermore, it seems necessary to extract from X at least twice, namely, using h_2 and h_c to get statistically and computationally random bits respectively.

⁷ In fact, this was first observed in [13] via the application of a special universal hash function of description length $O(n^2)$, and the work of [23] shows that the use of the hash function is not necessary.

3.3 PRGs from Any Known Regular OWFs: Removing the Dependency on ϵ

The parameterization of the aforementioned construction depends on ϵ , but sometimes ϵ is unknown or not polynomial-time computable. It is thus more desirable to have a construction based on any known-regular OWF regardless of parameter ϵ (as long as it is negligible). We observe that by setting entropy loss to zero (in which case hash functions are condensers) and letting G-L functions extract $O(\log n)$ bits the resulting generator is a generic (i.e. without relying on ϵ) pseudo-entropy generator (PEG) with a (collision) entropy stretch of $O(\log n)$ bits. Note however the output of the PEG is not indistinguishable from uniform but from some high collision entropy sources (with small constant entropy deficiency), which implies a PRG by running $q \in \omega(1)$ copies of the PEG and doing a single extraction from the concatenated outputs.

Definition 6 (pseudo-entropy generators). *Function $g : \{0, 1\}^n \rightarrow \{0, 1\}^{l+e}$ ($l > n$) is a (t, ϵ) \mathbf{H}_2 -pseudo-entropy generator (PEG) if g is polynomial-time computable and there exists a random variable $Y \in \{0, 1\}^{l+e}$ with $\mathbf{H}_2(Y) \geq l$*

$$\text{CD}_t(g(U_n), Y) \leq \epsilon.$$

where $(l - n)$ is the stretch of g , and e is the entropy deficiency. We say that g is an \mathbf{H}_2 -pseudo-entropy generator if $1/\epsilon$ and t are both super-polynomial.

Theorem 4 (PEGs from any known-regular OWFs). *For the same $f, g, \mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_C$ as assumed in Theorem 2 except that $d = 0$ and $s = 2 \log n + 2$, we have that if f is a known-regular one-way function then g is a \mathbf{H}_2 -pseudo-entropy generator with stretch $2 \log n$ and entropy deficiency 2.*

Proof sketch. It is not hard to see (using Lemma 2) that for $d = 0$ we have

$$\begin{aligned} & 2^{-\mathbf{H}_2(H_1(f(X)), H_2(X) \mid H_1, H_2)} = \text{CP}(H_1(f(X)), H_2(X) \mid H_1, H_2) \\ & \leq \Pr_{x_1, x_2 \leftarrow U_n, h_1 \leftarrow H_1} [h_1(f(x_1)) = h_1(f(x_2))] \\ & \quad \times \Pr_{x_1 \leftarrow X_1, x_2 \leftarrow X_2, h_2 \leftarrow H_2} [h_2(x_1) = h_2(x_2) \mid f(X_1) = f(X_2)] \\ & \leq 2^{-(n-k-1)} \cdot 2^{-(k-1)} = 2^{-(n-2)}. \end{aligned}$$

And we have by Lemma 5 and Goldreich-Levin the $2 \log n + 2$ hardcore bits are pseudo-random given $H_1(f(X))$ and $H_2(X)$, which completes the proof. \square

Theorem 5 (PRGs from any known-regular OWFs). *For any known k , there exists a generic construction of pseudo-random generator with seed length $\tilde{O}(n)$ by making $\tilde{O}(1)$ calls to any (length-preserving) 2^k -regular one-way function.*

Proof sketch. The idea is to run $q \in \omega(1)$ independent copies of the PEGs as in Theorem 4 to get an entropy stretch of $2q \log n$ followed by a single randomness extraction with entropy loss $q \log n$. This yields a PRG with stretch $q \log n$ that is roughly $O(q \cdot n^2 \sqrt{\varepsilon} + n^{-q})$ computationally indistinguishable from uniform randomness, where n^{-q} is negligible for any $q \in \omega(1)$. \square

3.4 PRGs from Any Unknown Regular OWFs

THE FIRST ATTEMPT: A PARALLEL CONSTRUCTION. A straightforward way to adapt the construction to unknown regular OWFs is to pay a factor of $n/\log n$. That is, it is not hard to see the construction for known regularity $\alpha = 2^k$ remains secure even by using an approximated value $\tilde{\alpha} = 2^{\tilde{k}}$ with accuracy $|\tilde{k} - k| \leq \log n$. This immediately implies a parallel construction by running $n/\log n$ independent copies of our aforementioned construction, where each i^{th} copy assumes regularity $2^{i \cdot \log n}$. Therefore, at least one (unknown) copy will be a PRG and thus we simply XOR the outputs of all copies and produce it as the output. Unfortunately, similar to the HILL approach, the parallelism turns out an inherent barrier to linear seed length. We will avoid this route by giving a sequential construction.

STEP 1: CONVERT TO KNOWN REGULARITY. Now we present the construction from any (length-preserving) unknown-regular OWF. We first transform it into a hardness-preserving equivalent with known regularity 2^n , as stated in Lemma 7.

Lemma 7 (unknown to known regularity). *For any length-preserving unknown regular (t, ε) -OWF $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, define*

$$\begin{aligned} \bar{f} : \mathcal{Y} \times \{0, 1\}^n &\rightarrow \mathcal{Y} \\ \bar{f}(y, r) &\stackrel{\text{def}}{=} f(y \oplus r) \end{aligned} \tag{7}$$

where $\mathcal{Y} \stackrel{\text{def}}{=} f(\{0, 1\}^n) \subseteq \{0, 1\}^n$, “ \oplus ” denotes bit-wise XOR. Then, \bar{f} is a 2^n -regular $(t - O(n), \varepsilon)$ -OWF.

Proof. On uniform (y, r) over $\mathcal{Y} \times \{0, 1\}^n$, $y \oplus r$ is uniform over $\{0, 1\}^n$. Thus, any algorithm inverts \bar{f} to produce (y, r) with probability ε implies another algorithm that inverts f with the same probability by outputting $y \oplus r$. Let us assume that f is α -regular. Then, for any $y_1 = \bar{f}(y, r) = f(y \oplus r)$ we have $|f^{-1}(y_1)| = \alpha$, and for any $x \in f^{-1}(y_1)$ we have $|\{(y, r) \in \mathcal{Y} \times \{0, 1\}^n : y \oplus r = x\}| = |\mathcal{Y}| = 2^n/\alpha$, which implies $|\bar{f}^{-1}(y_1)| = \alpha \cdot (2^n/\alpha) = 2^n$.

STEP 2: Z-SEEDED PRG. Similarly to the known regular case, we first assume ε is known and then eliminate the dependency. Intuitively, the output of \bar{f} hides n bits of min-entropy about its input (by the 2^n -regularity) plus another $\log(1/\varepsilon)$ bits of pseudo-entropy (due to the one-wayness), and thus one can extract $n + O(\log(1/\varepsilon))$ pseudorandom bits. This is formalized in Lemma 8, where we build a Z-seeded PRG \bar{g} that expands random elements over $\mathcal{Y} \times \{0, 1\}^n$ into pseudorandom ones over $\mathcal{Y} \times \{0, 1\}^{n+O(\log(1/\varepsilon))}$. The proof of Lemma 8 is similar to that of Theorem 2, and we defer it to the appendix.

Definition 7 (*Z*-seeded PRG [23]). A function $g^z : \mathcal{Z} \rightarrow \mathcal{Z} \times \{0, 1\}^s$ is a (t, ε) -secure *Z*-seeded PRG with stretch s if g^z is polynomial-time computable and $\text{CD}_t(g^z(\mathcal{Z}), (\mathcal{Z}, U_s)) \leq \varepsilon$.

Lemma 8 (construct *Z*-seeded PRG). Let f, \bar{f} be defined as in Lemma 7, for any integers d, s satisfying $7d + 6s = 2 \log(1/\varepsilon n)$, let $\mathcal{H} \stackrel{\text{def}}{=} \{h : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{n-d}\}$ be a universal hash function family, let $\mathcal{H}_C \stackrel{\text{def}}{=} \{h_c : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{d+s}\}$ be a *G-L* function family, define \bar{g} as

$$\begin{aligned} \bar{g} : \mathcal{Y} \times \{0, 1\}^n \times \mathcal{H} \times \mathcal{H}_C &\rightarrow \mathcal{Y} \times \{0, 1\}^{n+s} \times \mathcal{H} \times \mathcal{H}_C \\ \bar{g}(y, r, h, h_c) &\stackrel{\text{def}}{=} (\bar{f}(y, r), (h(y, r), h_c(y, r)), h, h_c) \end{aligned} \tag{8}$$

Then, we have that \bar{g} is a $(t \cdot (\varepsilon/n)^{O(1)}, O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{7}}))$ -secure *Z*-seeded PRG for $\mathcal{Z} = (Y, R, H, H_C)$, where (Y, R) is identically distributed to $(f(U_n^1), U_n^2)$, and H, H_C are uniform over $\mathcal{H}, \mathcal{H}_C$ respectively.

STEP 3: SEQUENTIAL COMPOSITION. Notice, however, \bar{g} is NOT a standard PRG with positive stretch as the only black-box way to sample distribution Y is to compute $f(U_n)$, which costs n random bits (despite that $\mathbf{H}_\infty(Y)$ might be far less than n). Quite naturally and thanks to the sequential composition, the construction simply iterates \bar{g} , reuses the random seeds (in each iteration), and outputs $s = O(\log(1/\varepsilon))$ bits per iteration.

Lemma 9 (sequential composition [23,3]). Let $g^z : \mathcal{Z} \rightarrow \mathcal{Z} \times \{0, 1\}^s$ be a (t, ε) -secure *Z*-seeded PRG, for $1 \leq i \leq \ell$ iteratively compute $(z_i, w_i) := g^z(z_{i-1})$, and define $g^{z, \ell}(z_0) \stackrel{\text{def}}{=} (z_\ell, w_1, \dots, w_\ell)$. Then, we have that $g^{z, \ell}$ is a $(t - \ell \cdot n^{O(1)}, \ell \cdot \varepsilon)$ -secure *Z*-seeded PRG with stretch $\ell \cdot s$.

Proof. The proof is seen by a hybrid argument.

Theorem 6 (PRGs from any unknown-regular OWFs with known hardness). Let function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be any (possibly unknown) regular length-preserving (t, ε) -OWF, define $\bar{f}, \bar{g}, \mathcal{H}, \mathcal{H}_C, s$ as in Lemma 8, and define

$$\begin{aligned} g : \{0, 1\}^n \times \{0, 1\}^n \times \mathcal{H} \times \mathcal{H}_C &\rightarrow (\{0, 1\}^s)^\ell \times \{0, 1\}^n \times \mathcal{H} \times \mathcal{H}_C \\ g(x, r_0, h, h_c) &\stackrel{\text{def}}{=} (w_1, w_2, \dots, w_\ell, r_\ell, h, h_c) \end{aligned}$$

where g computes $y = f(x)$, and sequentially composes (as in Lemma 9) the *Z*-seeded PRG \bar{g} (on input $z_0 = (y, r_0, h, h_c)$) ℓ times to produce output $w_1, w_2, \dots, w_\ell, r_\ell, h, h_c$. Then, for any $s \leq \log(1/\varepsilon n)/3$, we have that function g is a $(t \cdot (\varepsilon/n)^{O(1)} - \ell \cdot n^{O(1)}, O(\ell \cdot (2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{7}}))$ -secure PRG with stretch $\ell \cdot s - n$.

Proof. We can almost complete the proof by Lemma 8 and Lemma 9 except that the stretch of g (as a standard PRG) is $\ell \cdot s - n$ instead of $\ell \cdot s$. This is because we need to take into account that n bits are used to sample y at initialization.

CONCRETE PARAMETERS. Therefore, for any unknown-regular OWF with known hardness, we obtain a PRG with linear seed length, and by letting $s \in \Theta(\log(\frac{1}{\varepsilon n}))$ the number of calls $\ell \in \Theta(n/s) = \Theta(n/\log(1/\varepsilon n))$ matches the lower bound of [16]. This extends to the general case (where the hardness parameter is unknown) by repetition.

Theorem 7 (PRGs from any unknown-regular OWFs). *There exists a generic construction of pseudo-random generator with seed length $\tilde{O}(n)$ by making $\tilde{O}(n/\log n)$ calls to any unknown-regular one-way function.*

Proof sketch. For any unknown-regular OWF f , define \bar{g} as in Lemma 8 except setting $d = 0$ and $s = 2 \log n + 1$. It is not hard to see that the resulting \bar{g} is a \mathbf{H}_2 -pseudo-entropy generator with stretch $2 \log n$ and entropy deficiency 1 (proof similar to that in Theorem 4). We then use the repetition trick (similar to Theorem 5), namely, for any $q \in \omega(1)$ run q independent copies of \bar{g} and do a single extraction on the concatenated output with entropy loss set to $q \log n$. This gives us a Z' -seeded PRG \bar{g}' for $Z' = (Y, U_n, H, H_c)^q$ with stretch $q \cdot \log n$. Again, sequential composing \bar{g}' for $\ell' = \lceil (qn+1)/q \log n \rceil \in O(n/\log n)$ iterations yields a standard PRG

$$g' : \{0, 1\}^{2qn} \times \mathcal{H}^q \times \mathcal{H}_C^q \rightarrow \{0, 1\}^{2qn+s'} \times \mathcal{H}^q \times \mathcal{H}_C^q$$

where the stretch $s' = (q \cdot \log n) \cdot \ell' - q \cdot n \geq 1$. This completes the proof. \square

Acknowledgements. We thank Guang Yang, Colin Jia Zheng and Yunlei Zhao for useful comments. We are also grateful to Thomas Holenstein for clarifying his lower bound results [16] at the very early stage of this work. We thank the anonymous reviewers of ASIACRYPT 2013 for very helpful comments and suggestions that significantly improve the presentations of the PRG constructions. Finally, Yu Yu thanks Leonid Reyzin for interesting discussions about “saving private randomness [5]” at the Warsaw workshop [1], and thank Stefan Dziembowski for making the meeting possible. This research work was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61172085, 61061130540, 61073174, 61103221, 61070249, 60703031, 11061130539, 61021004 and 61133014. This work was also supported by the National Science Foundation of China under Grant Nos. 61272413 and 61133014, the Fok Ying Tung Education Foundation under Grant No. 131066, and the Program for New Century Excellent Talents in University under Grant No. NCET-12-0680.

References

1. Workshop on leakage, tampering and viruses (June 2013), <http://www.crypto.edu.pl/events/workshop2013>
2. Barak, B., Shaltiel, R., Wigderson, A.: Computational analogues of entropy. In: Arora, S., Jansen, K., Rolim, J.D.P., Sahai, A. (eds.) RANDOM 2003 and APPROX 2003. LNCS, vol. 2764, pp. 200–215. Springer, Heidelberg (2003)

3. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo random bits. In: Proceedings of the 23rd IEEE Symposium on Foundation of Computer Science, pp. 112–117 (1982)
4. Carter, J.L., Wegman, M.N.: Universal classes of hash functions. *Journal of Computer and System Sciences* 18, 143–154 (1979)
5. Dedić, N., Harnik, D., Reyzin, L.: Saving private randomness in one-way functions and pseudorandom generators. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 607–625. Springer, Heidelberg (2008)
6. Dodis, Y., Elbaz, A., Oliveira, R., Raz, R.: Improved randomness extraction from two independent sources. In: APPROX-RANDOM, pp. 334–344 (2005)
7. Goldreich, O.: *Foundations of Cryptography: Basic Tools*. Cambridge University Press (2001)
8. Goldreich, O.: Three XOR-lemmas — an exposition. In: Goldreich, O. (ed.) *Studies in Complexity and Cryptography*. LNCS, vol. 6650, pp. 248–272. Springer, Heidelberg (2011)
9. Goldreich, O., Krawczyk, H., Luby, M.: On the existence of pseudorandom generators. *SIAM Journal on Computing* 22(6), 1163–1175 (1993)
10. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: Johnson, D.S. (ed.) *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, Seattle, Washington, May 15–17, pp. 25–32 (1989)
11. Haitner, I., Harnik, D., Reingold, O.: On the power of the randomized iterate. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 22–40. Springer, Heidelberg (2006)
12. Haitner, I., Harnik, D., Reingold, O.: On the power of the randomized iterate. *SIAM Journal on Computing* 40(6), 1486–1528 (2011)
13. Haitner, I., Reingold, O., Vadhan, S.P.: Efficiency improvements in constructing pseudorandom generators from one-way functions. In: *Proceedings of the 42nd ACM Symposium on the Theory of Computing*, pp. 437–446 (2010)
14. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: Construction of pseudorandom generator from any one-way function. *SIAM Journal on Computing* 28(4), 1364–1396 (1999)
15. Holenstein, T.: Pseudorandom generators from one-way functions: A simple construction for any hardness. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 443–461. Springer, Heidelberg (2006)
16. Holenstein, T., Sinha, M.: Constructing a pseudorandom generator requires an almost linear Number of calls. In: *Proceedings of the 53rd IEEE Symposium on Foundation of Computer Science*, pp. 698–707 (2012)
17. Hsiao, C.-Y., Lu, C.-J., Reyzin, L.: Conditional computational entropy, or toward separating pseudoentropy from compressibility. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 169–186. Springer, Heidelberg (2007)
18. Lee, C.-J., Lu, C.-J., Tsai, S.-C., Tzeng, W.-G.: Extracting randomness from multiple independent sources. *IEEE Transactions on Information Theory* 51(6), 2224–2227 (2005)
19. Levin, L.A.: One-way functions and pseudorandom generators. *Combinatorica* 7(4), 357–363 (1987)
20. Nisan, N.: Pseudorandom generators for space-bounded computation. *Combinatorica* 12(4), 449–461 (1992)
21. Nisan, N., Zuckerman, D.: Randomness is linear in space. *Journal of Computer and System Sciences* 52(1), 43–53 (1996)

22. Stinson, D.R.: Universal hash families and the leftover hash lemma, and applications to cryptography and computing. *Journal of Combinatorial Mathematics and Combinatorial Computing* 42, 3–31 (2002), <http://www.cacr.math.uwaterloo.ca/~dstinson/publist.html>
23. Vadhan, S.P., Zheng, C.J.: Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In: *Proceedings of the 44th ACM Symposium on the Theory of Computing*, pp. 817–836 (2012)
24. Vadhan, S.P., Zheng, C.J.: A uniform min-max theorem with applications in cryptography. In: Canetti, R., Garay, J.A. (eds.) *CRYPTO 2013, Part I. LNCS*, vol. 8042, pp. 93–110. Springer, Heidelberg (2013)
25. Vazirani, U.V., Vazirani, V.V.: Efficient and secure pseudo-random number generation (extended abstract). In: *Proceedings of the 25th IEEE Symposium on Foundation of Computer Science*, pp. 458–463 (1984)
26. Yao, A.C.-C.: Theory and applications of trapdoor functions (extended abstract). In: *Proceedings of the 23rd IEEE Symposium on Foundation of Computer Science*, pp. 80–91 (1982)

A Proofs Omitted

Proof of Lemma 8. Note that $\bar{f}(Y, R)$ is identically distributed to Y , so it is equivalent to show

$$CD_{t,(\varepsilon/n)^{O(1)}}((H(Y, R), H_C(Y, R)), U_{n+s} \mid \bar{f}(Y, R), H, H_C) = O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{7}}) .$$

It follows from the $(t - O(n), \varepsilon)$ -one-way-ness of \bar{f} (see Lemma 7) and Lemma 3 that

$$\mathbf{H}_{t-O(n)}((Y, R) \mid \bar{f}(Y, R)) \geq n + \log(1/\varepsilon) . \tag{9}$$

Then, similar to Lemma 4, we have the following entropy conditions

$$\mathbf{H}_{\infty}((Y, R) \mid \bar{f}(Y, R)) = n ,$$

$$\begin{aligned} \mathbf{H}_{t-O(n)}((Y, R) \mid \bar{f}(Y, R), h(Y, R), h) &\geq \mathbf{H}_{t-O(n)}((Y, R) \mid \bar{f}(Y, R)) - (n - d) \\ &\geq d + \log(1/\varepsilon) , \end{aligned}$$

hold for any $h \in \mathcal{H}$, where the second inequality is by applying the chain rule to Equation (9). Therefore,

$$\begin{aligned} &CD_{t,(\varepsilon/n)^{O(1)}}((H(Y, R), H_C(Y, R)), U_{n+s} \mid \bar{f}(Y, R), H, H_C) \\ &\leq SD(H(Y, R), U_{n-d} \mid \bar{f}(Y, R), H) \\ &\quad + CD_{t,(\varepsilon/n)^{O(1)}}(H_C(Y, R), U_{d+s} \mid \bar{f}(Y, R), H(Y, R), H, H_C) \\ &\leq 2^{-\frac{d}{2}} + O(2^{d+s} \cdot (n \cdot \varepsilon \cdot 2^{-d})^{\frac{1}{3}}) = 2^{-\frac{d}{2}} + O(2^{d+s} \cdot (2^{\frac{-(7d+6s)}{2}} \cdot 2^{-d})^{\frac{1}{3}}) \\ &= O(2^{-\frac{d}{2}}) = O(2^{\frac{3s+\log(\varepsilon \cdot n)}{7}}) = O((2^{3s} \cdot \varepsilon \cdot n)^{\frac{1}{7}}) \end{aligned}$$

where the first inequality is triangle, the statistical distance is due to the leftover hash lemma and the computational distance of the second inequality is by the Goldreich-Levin Theorem. □