Theoretical Computer Science 569 (2015) 58–69

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Pseudorandom generators from regular one-way functions: New constructions with improved parameters ☆

Yu Yu [a], Xiangxue Li [b,*], Jian Weng [c]

[a] Department of Computer Science and Engineering, Shanghai Jiao Tong University, China
[b] Department of Computer Science and Technology, East China Normal University, China
[c] Department of Computer Science, Jinan University, China

ARTICLE INFO

Article history:
Received 18 March 2014
Received in revised form 6 December 2014
Accepted 19 December 2014
Available online 30 December 2014
Communicated by G. Persiano

Keywords:
Foundations
Pseudorandom generators
One-way functions
Randomized iterate

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 (2012) [6].

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 et al. (2006) [4] 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) [2] 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 [2] were found useful in many other contexts (such as leakage-resilient cryptography). Nevertheless, a major drawback of [2] 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 [3–5], just to name a few) have been followed up towards sim-

☆ A preliminary version [1] of this paper was presented at the conference Asiacrypt 2013.

* Corresponding author at: Department of Computer Science and Technology, East China Normal University, Dongchuan RD 500#, Shanghai 200241, China.
E-mail addresses: yuyuathk@gmail.com (Y. Yu), xxli@cs.ecnu.edu.cn (X. Li).

[1] More precisely, the main construction of [2] requires seed length $O(n^{10})$, but [2] also sketches another construction of seed length $O(n^8)$, which was formalized and proven in [3].

http://dx.doi.org/10.1016/j.tcs.2014.12.013
0304-3975/

plifying and improving the constructions, and the current state-of-the-art construction [5] requires seed length $O(n^3)$. Let us mention that 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 [6], 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 [7] and Yao [8] independently introduced the notion of PRGs, and observed that PRGs can be efficiently constructed from one-way permutations (OWPs). That is, given an OWP f on input x and its hardcore function h_c (e.g. by Goldreich and Levin [9]), a single invocation of f already implies a PRG $g(x) = (f(x), h_c(x))$ with a stretch[3] of $\Omega(\log n)$ bits and it extends to arbitrary stretch by repeated iterations (seen by a hybrid argument):

$$g^{\ell}(x) = \big(h_c(x), h_c\big(f^1(x)\big), \ldots, h_c\big(f^{\ell}(x)\big), \ldots\big)$$

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 [10] observed that f is not necessarily an 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 [11] assumed known-regular[4] 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 [12], 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. [3,4]). Haitner, Harnik and Reingold [13] refined the technique used in [11] (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 [11] and applies a random pairwise independent hash function h_i in between every two applications of f, i.e.

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

The key observation is "*the last iterate is hard-to-invert*" [14], more precisely, function f, when applied to $h_{i-1}(f^{i-1}; h_1, \ldots, h_{i-2})$, is hard-to-invert even if $h_1, \ldots, h_{i-1}$ 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 [15]). The randomized iterate matches the lower bound on the number of OWF calls,[5] 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 [12]. 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 [14]. We sketch both constructions as follows.

Entropy observation. We start by assuming a (t, ε)-OWF f (see Definition 2.2) 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 2.5). 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 uniformly distributed on the 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 an $x' \in f^{-1}(y)$, 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.

- 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.

[2] The lower bound of [6] 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 (see Definition 3.2).

[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.

[5] As explicitly stated in [6], the lower bound of $\Omega(n/\log n)$ calls also applies to unknown regular OWFs.