(Almost) Optimal Constructions of UOWHFs from 1-to-1, Regular One-way Functions and Beyond

Abstract

We revisit the problem of black-box constructions of universal one-way hash functions (UOWHFs) from several (from specific to more general) classes of one-way functions (OWFs), and give constructions accordingly that either improve the best previously known or generalize to a broader class of one-way functions. In addition, the parameters we achieve are either optimal or almost optimal simultaneously up to small factors, e.g., $O(\log n)$ or arbitrarily small $\omega(1)$.

- For any 1-to-1 one-way function (and not necessarily a one-way permutation), we give an optimal construction of UOWHFs with key and output length $\Theta(n)$ by making a single call to the underlying OWF. This improves the constructions of Naor and Yung (STOC 1989) and De Santis and Yung (Eurocrypt 1990) that need key length $O(n \cdot \omega(\log n))$.
- For any known-(almost-)regular one-way function with known hardness, we give another optimal construction of UOWHFs with key and output length $\Theta(n)$ and a single call to the one-way function.
- For any known-(almost-)regular one-way function, we give a construction of UOWHFs with key and output length $O(n{\cdot}\omega(1))$ and by making $\omega(1)$ non-adaptive calls to the one-way function. This improves the construction of Barhum and Maurer (Latincrypt 2012) that requires key and output length $O(n{\cdot}\omega(\log n))$ and $\omega(\log n)$ calls.
- For any one-way function f that is weakly unknown-regular (i.e., the set of x's with maximal number of siblings is of an n^{-c}-fraction for some constant c), we give a construction of UOWHFs with key length $O(n{\cdot}\log n)$ and output length $\Theta(n)$. This generalizes the construction of Ames et al. (Asiacrypt 2012) which requires an unknown-regular one-way function (i.e., $c = 0$).

Along the way, we introduce several technical tools and techniques that might be of independent interest. The first tool is a technical lemma about universal hashing with nice symmetry to the leftover hash lemma. Secondly, we show that almost 1-to-1 (except for a negligible fraction) one-way functions and known (almost-)regular one-way functions are equivalent in the known-hardness (or non-uniform) setting, by giving an optimal construction of the former from the latter. Thirdly, we show how to transform any one-way function that is far from regular (but only weakly regular on a noticeable fraction of domain) into an unknown-almost-regular one-way function.

Keywords: Foundations, One-way Functions, Universal One-way Hash Functions, Target Collision Resistance.

1 Introduction

Informally, a family of compressing hash functions, denoted by $\mathcal{G}$, is called *universal one-way*, if given a random function $g \in \mathcal{G}$ and a random (or equivalently, any pre-fixed) input x, it is infeasible for any efficient algorithm to find any $x' \neq x$ satisfying $g(x) = g(x')$. The seminal result that one-way functions (OWFs) imply universal one-way hash functions (UOWHFs) [21] is one of the central results upon which modern cryptography is successfully founded. It further implies that digital signature (as defined in [10]) can be based on any one-way function [19]. Other important applications of UOWHFs include constructions of Cramer-Shoup encryption scheme [4], statistically hiding commitment scheme [13, 14], etc.

UOWHFs FROM ANY OWFs. The principle possibility result that UOWHFs can be based on any OWF was established by Rompel [21] (with some corrections given in [22, 17]). However, Rompel's construction was quite complicated and extremely unpractical. In particular, for any one-way function on n-bit inputs it requires key length $\tilde{O}(n^{12})$ and output length $\tilde{O}(n^8)$. Haitner et al. [12] improved the construction via the notion of inaccessible entropy [14], and reduced key and output length to $\tilde{O}(n^7)$. We mention also recent development by Gennaro and Venkitasubramaniam [6] that further reduces the key and output lengths. Despite of all these improvements, the constructions are mainly of theoretical interest and are too inefficient to be of any practical use.

UOWHFs FROM SPECIAL OWFs. Another line of research focuses on more efficient (and nearly practical) constructions of UOWHFs from special structured OWFs. Naor and Yung gave an elegant construction of UOWHFs with key and output length $\Theta(n)$ which does a single call to any one-way permutation. More specifically, let $f : \{0,1\}^n \to \{0,1\}^n$ be any one-way permutation, let h be a random permutation (over n bits) from a pairwise-independent hash permutation family $\mathcal{H}$, and let $\mathsf{trunc} : \{0,1\}^n \to \{0,1\}^{n-1}$ be a truncating function that outputs the first $n-1$ bits of input, then the following

$$\mathcal{G}_{owp} \stackrel{\mathsf{def}}{=} \{\ (\mathsf{trunc} \circ h \circ f) : \{0,1\}^n \to \{0,1\}^{n-1}\ ,\ h \in \mathcal{H}\ \}$$

is a family of UOWHFs with 1 bit of shrinkage (i.e., compress by 1 bit), where "$\circ$" denotes function composition. However, for a slightly weaker primitive, namely, 1-to-1 one-way functions, the authors of [19] only gave a rather complicated construction. De Santis and Yung [23] gave an improved construction (from any 1-to-1 OWF $f : \{0,1\}^n \to \{0,1\}^l$) as below:

$$\mathcal{G}_{1-to-1} \stackrel{\mathsf{def}}{=} \{\ (h^n_{n-1} \circ \ldots \circ h^{l-1}_{l-2} \circ h^l_{l-1} \circ f) : \{0,1\}^n \to \{0,1\}^{n-1}\ ,\ h^n_{n-1} \in \mathcal{H}^n_{n-1}, \ldots, h^l_{l-1} \in \mathcal{H}^l_{l-1}\ \}\ ,$$

where each $\mathcal{H}^i_{i-1}$ denotes a family of pairwise-independent hash functions that compress i-bit strings into $(i-1)$ bits. Although $\mathcal{G}_{1-to-1}$ enjoys linear output length and a single function call, it requires[1] key length $O(\omega(\log n) \cdot n)$ for describing all the hash functions. In addition, the work of [23] also introduced a construction from any known-regular[2] one-way function with key and output length $O(\omega(\log^2 n) \cdot n)$ and $O(\omega(1) \cdot \log n)$ adaptive calls, which was recently improved by Barhum and Maurer [3] to key and output length $O(\omega(\log n) \cdot n)$ and $O(\omega(1) \cdot \log n)$ non-adaptive calls. Based on unknown-regular one-way functions, Ames et al. [1] presented a more general construction with output length $\Theta(n)$, key length $O(\log n \cdot n)$ and $\tilde{O}(n)$ adaptive calls. In summary, as tabulated in Table 1, the best known construction

[1] A back-of-the-envelope calculation suggests that $\mathcal{G}_{1-to-1}$ needs key length $O(l \cdot (l-n))$, and we know (see Fact 3.1) that every 1-to-1 one-way function implies another one-way function $f' : \{0,1\}^{n' \in \Theta(n)} \to \{0,1\}^{n' + \omega(\log n)}$ that is 1-to-1 except on a negligible fraction of inputs, which implies that the key length of [19, 23] can be pushed to $O(\omega(\log n) \cdot n)$.

[2] A function f is regular if 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. More generally (as introduced in [27]), f is weakly unknown-regular if the fraction of x's with maximal $|f^{-1}(f(x))|$ (which is not necessarily efficiently computable) is noticeable. We stress that here "weakly" is used to describe "regularity" (rather than "one-wayness" as in "weakly one-way functions").

Table 1: A summary of the parameters from existing constructions [19, 23, 3, 1] and our work, where KR-OWF and UR-OWF are the shorthands for known-regular and unknown-regular one-way functions respectively, ε-hard KR-OWF additionally assumes that the hardness parameter ε of KR-OWF is known, and n^{-c}-WUR-OWF is the shorthand for weakly unknown-regular one-way functions (see Footnote 2 and formally Definition 2.6).

	Assumption	Output Length	Key Length	# of Calls	Type
[19]	OWP	$\Theta(n)$	$\Theta(n)$	1	non-adaptive
[23, 19]	1-to-1 OWF	$\Theta(n)$	$O(\omega(\log n)\cdot n)$	1	non-adaptive
[23]	KR-OWF	$O(\omega(\log^2 n)\cdot n)$	$O(\omega(\log^2 n)\cdot n)$	$O(\omega(\log n))$	adaptive
[3]	KR-OWF	$O(\omega(\log n)\cdot n)$	$O(\omega(\log n)\cdot n)$	$O(\omega(\log n))$	non-adaptive
[1]	UR-OWF	$\Theta(n)$	$O(\log n\cdot n)$	$\tilde{O}(n)$	adaptive
ours	1-to-1 OWF	$\Theta(n)$	$\Theta(n)$	1	non-adaptive
ours	ε-hard KR-OWF	$\Theta(n)$	$\Theta(n)$	1	non-adaptive
ours	KR-OWF	$O(\omega(1)\cdot n)$	$O(\omega(1)\cdot n)$	$O(\omega(1))$	non-adaptive
ours	n^{-c}-WUR-OWF	$\Theta(n)$	$O(\log n\cdot n)$	$\tilde{O}(n^{2c+1})$	adaptive

requires key length $O(n\cdot\log n)$ even for a 1-to-1 one-way function, and needs to make $O(\omega(1)\cdot\log n)$ calls (or $\tilde{O}(n)$ adaptive calls if one wants linear output length at the same time) to a regular one-way function.

Summary of our constructions. In the paper, we give the following constructions from the respective aforementioned one-way functions. The first two constructions achieve optimal parameters simultaneously, the third is almost optimal up to an arbitrarily small super-constant factor $\omega(1)$ (e.g., $\log\log\log n$ or even less), and the fourth has optimal output length $\Theta(n)$ and key length $O(n\cdot\log n)$. We remark that further improvement on key length $O(n\cdot\log n)$ (of the fourth construction) requires more key-length efficient domain extender (for Merkle-Damgård construction of UOWHFs) than Shoup's [24], which seems far beyond reach of conventional techniques[3]. Finally, the first three constructions have optimal shrinkages (per invocation of OWF) by matching the upper bound of Gennaro et al. [5].

1. For any 1-to-1 one-way function, we give an optimal construction of UOWHFs with key and output length $\Theta(n)$ and a single OWF call. This improves the constructions of Naor and Yung [19] and De Santis and Yung [23] that require key length $O(n\cdot\omega(\log n))$.

2. For any known-regular one-way function with known hardness, we give another optimal construction of UOWHFs with key and output length $\Theta(n)$ and a single call.

3. For any known-regular one-way function, we give a construction of UOWHFs with key and output length $O(\omega(1)\cdot n)$ and $\omega(1)$ non-adaptive calls. This improves the construction of Barhum and Maurer [3] that requires key and output length $O(n\cdot\omega(\log n))$ and $\omega(\log n)$ calls.

4. For any one-way function f that is weakly unknown-regular on a noticeable fraction (e.g., n^{-c} for constant c) of domain, we give a construction of UOWHFs with key length $O(n\cdot\log n)$ and output length $\Theta(n)$. This generalizes the construction of Ames et al. [1], where an unknown-regular one-way function (i.e., $c=0$) is required.

On the (a)symmetry to PRGs. Our results further improve the understanding about the inherent "black-box duality" [5, 14, 12] between one-way functions and pseudorandom generators. Firstly, we

[3] Asymmetrically, the case of range extension for a PRG g is much easier by just composing g with itself iteratively.

introduce a technical lemma (see Lemma 3.1) which is dual to the leftover hash lemma and might be of independent interest. Informally, it says that when applying a universal hash function h to any "flat" random variable X of entropy (**no more than**) $\mathbf{a}$ to produce an $(\mathbf{a}+\mathbf{d})$-bit output, h will be injective on X except for a $\mathbf{2^{-d}}$ fraction. In contrast, the leftover hash lemma states that when hashing any X of Rényi entropy (**no less than**) $\mathbf{a}$ into $(\mathbf{a}-\mathbf{d})$-bit strings, the resulting output distribution will be $\mathbf{2^{-d/2}}$-close (in terms of statistical distance) to uniform, where the symmetry is highlighted in bold. Secondly, constructions #2 and #3 above match the best known results about constructions of PRGs from known-regular OWFs (see [28]), namely, seed length $O(\omega(1)\cdot n)$ or even $\Theta(n)$ if the hardness of the underlying OWF is known. Thirdly, construction #1 is symmetric to the PRG construction [27] based on the same class of one-way functions, where succinct key/seed length $O(n\cdot\log n)$ is achieved via bounded space generators. Finally (and perhaps more interestingly), construction #1 is asymmetric to the case of PRGs, where we do not know how to construct a linear seed length PRG from an arbitrary 1-to-1 one-way function in general[4].

On the efficiency, feasibility and limits. Constructions #1, #2 and #3 are practically relevant as most one-way function candidates turn out to be known-almost-regular or even 1-to-1. Goldreich, Levin and Nisan [9] showed how to base 1-to-1 one-way functions on concrete intractable problems such as RSA and DLP. We further prove (as a byproduct of construction #2) the equivalence of almost 1-to-1 (i.e., 1-to-1 except for a negligible fraction) one-way functions and known-(almost-)regular one-way functions in the known-hardness setting, by giving an optimal construction of the former from the latter. Moreover, unknown regular one-way functions further reduce the knowledge required about the underlying one-way functions, and the problem of basing cryptographic primitives (PRGs, UOWHFs, etc.) on weaker assumptions is of theoretic interests. It improves our understanding about the feasibility and limits of black-box reductions. In particular, Holenstein and Sinha [15], Barhum and Holenstein [2] showed that $\Omega(n/\log n)$ black-box calls to an arbitrary (including unknown-regular) one-way function is necessary to construct PRGs and UOWHFs, and the lower bound is matched by explicit constructions of PRGs [11] and UOWHFs [1] respectively. We carry on this study even further by considering a more general class of one-way functions called weakly unknown-regular one-way functions (as introduced in [27]), namely, the underlying one-way function can have an arbitrary structure as long as the set of x's with maximal number of siblings (i.e., x and x' are siblings of each other if $f(x)=f(x')$) is of noticeable fraction. The authors of [27] gave a construction of PRG with seed length $O(n\cdot\log n)$ from weakly unknown-regular OWFs. However, their analysis is quite ad-hoc (see Remark 5.1), and doesn't seem to generalize to the case of UOWHFs. As an intermediate step of construction #4, we prove a statement that "iterating such a one-way function (that is weakly regular on only a noticeable fraction) sufficiently many times yields a one-way function that is regular on an overwhelming fraction" and thus unifies the approach to the two closely related objects (i.e., PRGs and UOWHFs). We mention an (arguable) analogue to this problem, namely, hardness amplification of one-way functions [26], where a function that is weakly one-way (for which every efficient algorithm has a noticeable fraction to fail upon) can be turned into strongly one-way (hard to invert on an overwhelming portion) by parallel repetition [26] or even sequential composition (assuming additionally that the underlying function is regular) [11].

The roadmap. We outline below the steps to build UOWHFs from the respective one-way function $f:\{0,1\}^n\to\{0,1\}^l$ introduced above. To simplify the presentation, we assume without loss of generality that $l\in O(n)$ for 1-to-1 one-way functions and even length-preserving (i.e., $l=n$) for

[4]Given a 1-to-1 one-way function f, one might think of getting a PRG by hashing $f(U_n)$ into $n-s$ bits concatenated with $s+1$ hard-core bits of f, where $s\in\omega(\log n)$ is the necessary entropy loss due to the leftover hash lemma. This is in general not possible without knowing the hardness of the underlying f. See more discussions and the relaxed solutions to this problem by Goldreich [7, Section 3.5.1.3]. For example, we get a linear seed-length PRG of the following weaker form, i.e., for every $\varepsilon=1/\mathsf{poly}(n)$ there exists a weak PRG of seed length $\Theta(n)$ whose output distribution is ε-indistinguishable from uniform to all PPT distinguishers. Alternatively, we use parallel repetition to obtain a standard PRG with seed length $O(\omega(1)\cdot n)$ [28].

arbitrary one-way functions. We state this as Fact 3.1 with a full proof given in Appendix A. In fact, Haitner et al. [11] showed that any one-way function implies a length-preserving one-way function, and our Fact 3.1 proves an even stronger version that (1) any 1-to-1 one-way function $f : \{0,1\}^n \to \{0,1\}^l$ implies a one-way function $f' : \{0,1\}^{n' \in \Theta(n)} \to \{0,1\}^{l' \in \Theta(n)}$ that is 1-to-1 except for a negligible fraction; (2) any one-way function f with $\alpha \le |f^{-1}(y)| \le \alpha \cdot \beta$ implies another length-preserving one-way function $f' : \{0,1\}^{n' \in \Theta(n)} \to \{0,1\}^{n'}$ with $\alpha' \le |f'^{-1}(y)| \le \alpha' \cdot \beta$ except for a negligible fraction, where the size of range β is preserved, and α' is efficiently computable if α is.

- Based on 1-to-1 OWFs. We adapt the classic Naor-Yung construction (for one-way permutation) to any 1-to-1 one-way function as follows:

$$\mathcal{G}_1 \stackrel{\text{def}}{=} \{\ (\mathsf{trunc} \circ h \circ f) : \{0,1\}^n \to \{0,1\}^{n-s}\ ,\ h \in \mathcal{H}\ \}\ ,$$

  where $\mathcal{H}$ is a family of universal[5] hash permutations on l bits, and $\mathsf{trunc} : \{0,1\}^l \to \{0,1\}^{n-s}$ is a truncating function that outputs the first $n-s$ bits of input. We give a proof that if f is a (t,ε)- 1-to-1 OWF f then the resulting $\mathcal{G}_1$ is a $(t - n^{O(1)},\ 2^s \cdot \varepsilon)$-UOWHF family with key and output length $\Theta(n)$ and shrinkage s (see Definition 2.2 and Definition 2.7 for formal definitions). The construction enjoys optimal parameters and somewhat counter-intuitively the security bound drops only by factor 2^s (which is optimal by [5]) rather than by 2^{l-n+s} (i.e., exponential in the number of bits truncated). We refer to the proof of Theorem 3.1 and Remark 3.1 for more technical details and further discussions.

- Based on known-(almost-)regular ε-hard OWFs. Given an almost-regular f (see Definition 2.5) which is known to be (t,ε)-one-way, i.e., ε is efficiently computable, we define the following function family

$$\mathcal{G}_2 \stackrel{\text{def}}{=} \{\ g : \{0,1\}^n \to \{0,1\}^{n-s}\ ,\ g(x) = (g_1(x), h_1(x)),\ g_1 = \mathsf{trunc} \circ h \circ f,\ h \in \mathcal{H}\ , h_1 \in \mathcal{H}_1\ \}$$

  where $\mathcal{H}$ is a family of universal hash permutations, and let $\mathcal{H}_1$ and trunc be a family of universal hash functions and the truncating function (both with appropriate output sizes) respectively. We show that $\mathcal{G}_2$ is a UOWHF family with key and output length $\Theta(n)$ and shrinkage s. The rationale is that for any[6] $x \neq x'$ colliding on $g \in \mathcal{G}_2$ it either satisfies "$f(x) = f(x') \wedge h_1(x) = h_1(x')$" or "$f(x) \neq f(x') \wedge \mathsf{trunc}(h(f(x))) = \mathsf{trunc}(h(f(x')))$". The former is bounded information-theoretically by our hashing lemma, and the latter is computationally bounded (and reducible to the one-way-ness of f). Interestingly, by abstracting out function $f'(x, h_1) \stackrel{\text{def}}{=} (f(x), h_1(x), h_1)$ from the above construction, we further show that f' is a one-way function that is 1-to-1 except for a negligible fraction. We refer to Theorem 4.1, Lemma 4.1 and Theorem 4.2 for the details.

- Based on known-(almost-)regular OWFs. Next, we consider any known-(almost)-regular OWF f whose hardness parameter is ε unknown (i.e., ε is negligible but may not be efficiently computable). In this case, we run q independent copies of f, and we get a construction by making q non-adaptive calls with shrinkage $q \log n$, key and output length $O(q \cdot n)$, where $q \in \omega(1)$ can be any efficiently computable super-constant. The parallel repetition technique was also used in similar contexts (e.g., the construction of PRG from any known regular OWF [28]). We refer to Theorem 4.3 for the detailed construction and proof.

[5] Many existing UOWHF constructions use pairwise (or even 3-wise) independent hashing to facilitate the analysis, but in fact universal hashing suffices here.

[6] More precisely, x is sampled at random and x' can be any efficient function of x such that $x \neq x'$.

- Based on weakly unknown-regular OWFs. Finally, we proceed to the more general assumption that the one-way function is only weakly unknown-regular on a noticeable fraction of the domain. We show iterating this f sufficiently many times yields a one-way function f' that is unknown-almost-regular, and thus plugging this f' into the construction of Ames et al.[1] yields a construction of UOWHFs with output length $\Theta(n)$ and key length $O(n \cdot \log n)$.

2 Preliminaries

Notations and definitions. We use $[n]$ to denote set $\{1, \ldots, n\}$. We use capital letters (e.g., X, Y) for random variables, standard letters (e.g., x, y) for values, and calligraphic letters (e.g. $\mathcal{X}$, $\mathcal{Y}$) for sets. The support of a random variable X, denoted by $\mathsf{Supp}(X)$, refers to the set of values on which X takes with non-zero probability, i.e., $\{x : \Pr[X = x] > 0\}$. For a binary string $x = x_1 \ldots x_n$, denote by $x_{[t]}$ the first t bits of x, i.e., $x_1 \ldots x_t$. $x \| y$ refers the concatenation of x and y. We denote by $\mathsf{trunc} : \{0,1\}^n \to \{0,1\}^t$ a truncating function that outputs the first t bits of input, i.e., $\mathsf{trunc}(x) = x_{[t]}$. $|\mathcal{S}|$ denotes the cardinality of set $\mathcal{S}$. For function $f : \{0,1\}^n \to \{0,1\}^{l(n)}$, we use shorthand $f(\{0,1\}^n) \stackrel{\mathsf{def}}{=} \{f(x) : x \in \{0,1\}^n\}$, and denote by $f^{-1}(y)$ the set of y's preimages under f, i.e., $f^{-1}(y) \stackrel{\mathsf{def}}{=} \{x : f(x) = y\}$. We say f is length-preserving if $l(n) = n$. We use $s \leftarrow S$ to denote sampling an element s according to distribution S, and let $s \stackrel{\$}{\leftarrow} \mathcal{S}$ denote sampling s uniformly from set $\mathcal{S}$, and $y := f(x)$ denote value assignment. We use U_n and $U_{\mathcal{X}}$ to denote uniform distributions over $\{0,1\}^n$ and $\mathcal{X}$ respectively, and let $f(U_n)$ be the distribution induced by applying function f to U_n. For probabilistic algorithm A, we use $\mathsf{A}(x; r)$ to denote the output of A on input x and internal coin r.

Collision probability. We use $\mathsf{CP}(X)$ to denote the collision probability of X, i.e., $\mathsf{CP}(X) \stackrel{\text{def}}{=} \sum_x \Pr[X = x]^2$, and denote by $\mathsf{CP}(X|Z)$ the average collision probability of X conditioned on another (possibly correlated) random variable Z by

$$\mathsf{CP}(X|Z) \stackrel{\text{def}}{=} \mathbb{E}_{z \leftarrow Z} \left[\sum_x \Pr[X = x|\ Z = z]^2 \right] \quad .$$

Simplifying Notations. To simplify the presentation, we use the following simplified notations. Throughout, most parameters are functions of the security parameter n (e.g., $t(n)$, $\varepsilon(n)$, $r(n)$) and we often omit n when clear from the context (e.g., t, ε, r). Parameters (e.g., ε, r) are said to be known if they are polynomial-time computable from n. By notation $f : \{0,1\}^n \to \{0,1\}^l$ we refer to the ensemble of functions $\{f : \{0,1\}^n \to \{0,1\}^{l(n)}\}_{n \in \mathbb{N}}$. 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 context. For example, in $h(y) \stackrel{\mathsf{def}}{=} h \cdot y$ the first h denotes a function, the second h refers to a string (a finite field element) that describes the function, and '$\cdot$' denotes multiplication between field elements.

Definition 2.1 (ρ-almost universal hashing) *A family of functions $\mathcal{H} = \{h : \{0,1\}^l \to \{0,1\}^t\}$ is ρ-almost universal if for any distinct $x_1, x_2 \in \{0,1\}^l$, it holds that*

$$\Pr_{h \stackrel{\$}{\leftarrow} \mathcal{H}} [h(x_1) = h(x_2)] \leq \rho \ .$$

In the special case $\rho = 2^{-t}$, we say that $\mathcal{H}$ is universal.

It is folklore that almost universal families of hash functions can be efficiently constructed.

Fact 2.1 (efficient constructions of almost universal hashing) *For any integers $t \leq l$, there exists a family of $O(l/t) \cdot 2^{-t}$-almost universal hash functions $\mathcal{H} = \{h : \{0,1\}^l \to \{0,1\}^t\}$ such that $\mathcal{H}$ has description length $O(t)$ and every $h \in \mathcal{H}$ is computable in time $\mathsf{poly}(l)$.*

A CONCRETE EXAMPLE. Assume without loss of generality that t divides l , i.e., $l = k{\cdot}t$ for some $k \in \mathbb{N}$ (otherwise use $l' = \lceil (l/t) \rceil \cdot t$ instead of l), and parse x as a sequence of t-bit strings $(x_1, \ldots, x_k)$. Then, we have that $\mathcal{H} = \{h_a : h_a(x) \stackrel{\text{def}}{=} \sum_{i=1}^{k} a^i{\cdot}x_i,\ a, x_i \in GF(2^t)\}$ is a family of $k{\cdot}2^{-t}$-almost universal hash functions of description length t. In fact, we might be able to use explicit almost pairwise independent hash functions [18, 25] to achieve even smaller ρ (e.g., $\rho = O(2^{-t})$ for any $l \in \mathsf{poly}(t)$), but the above construction already suffices for our applications.

Definition 2.2 (one-way functions) *A sequence of functions $\{f : \{0,1\}^n \to \{0,1\}^{l(n)}\}_{n\in\mathbb{N}}$ is $(t(n),\varepsilon(n))$-one-way if f is polynomial-time computable and for any probabilistic algorithm A of running time $t(n)$*

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n} [\mathsf{A}(1^n, f(x)) \in f^{-1}(f(x))] \ \leq \ \varepsilon(n).$$

Hereafter we use simplified notation $f : \{0,1\}^n \to \{0,1\}^{l(n)}$ for the above one-way function, where $t(\cdot)$ and $1/\varepsilon(\cdot)$ are super-polynomial.

Definition 2.3 (a family of one-way functions) *A sequence of function family $\mathcal{F} = \{\mathcal{F}_n\}_{n\in\mathbb{N}}$, where $\mathcal{F}_n = \{f_u : \{0,1\}^n \to \{0,1\}^{l(n)}, u \in \{0,1\}^{q(n)}\}$, is $(t(n),\varepsilon(n))$-one-way if for any $n \in \mathbb{N}$, $u \in \{0,1\}^{q(n)}$ and $x \in \{0,1\}^n$, the value $f_u(x)$ can be computed in polynomial time, and for any probabilistic algorithm A of running time $t(n)$, we have that*

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n;\ u \stackrel{\$}{\leftarrow} \{0,1\}^{q(n)}} [\ \mathsf{A}(1^n, u, f_u(x)) \in f_u^{-1}(f_u(x))\ ] \ \leq \ \varepsilon(n) \quad .$$

Likewise, we use simplified notation $\mathcal{F} = \{f_u : \{0,1\}^n \to \{0,1\}^{l(n)}, u \in \{0,1\}^{q(n)}\}$ for $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$.

Definition 2.4 (almost 1-to-1 functions) *A function $f : \{0,1\}^n \to \{0,1\}^{l(n)}$ is $\epsilon(n)$-almost 1-to-1 if there exists a negligible function $\epsilon(n)$, such that for every $n \in \mathbb{N}$ we have*

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n} [\ \exists x' :\ x' \neq x \wedge f(x) = f(x')\ ] \leq \epsilon(n).$$

In particular, f is 1-to-1 for $\epsilon(n) \equiv 0$.

Definition 2.5 (almost regular functions) *A function $f : \{0,1\}^n \to \{0,1\}^{l(n)}$ is $\alpha(n)$-regular if there exists an integer function $\alpha(n)$, called the regularity function, such that for every $n \in \mathbb{N}$ and $x \in \{0,1\}^n$ we have*

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

For integer functions $\alpha(n)$ and $\beta(n)$, f is $(\alpha(n),\ \alpha(n){\cdot}\beta(n))$-almost regular if for every $n \in \mathbb{N}$ and $x \in \{0,1\}^n$ we have

$$\alpha(n) \ \leq \ |f^{-1}(f(x))| \ \leq \ \alpha(n) \cdot \beta(n).$$

In particular, f is known-(almost)-regular if α is polynomial-time computable, or otherwise it is called unknown-(almost)-regular. In case that f is also $(t(n),\varepsilon(t))$-one-way, standard "almost-regularity" refers to that f is $(\alpha(n),\ \alpha(n){\cdot}\beta(n))$-almost-regular for $\beta(n) \leq \mathsf{poly}(n)$ or at most $\beta(n) \in (1/\varepsilon(n))^{O(1)}$ for certain small constant $O(1)$.

Definition 2.6 (weakly unknown-regular OWFs [27]) *Let $f : \{0,1\}^n \to \{0,1\}^{l(n)}$ be a one-way function, and for every $n \in \mathbb{N}$, divide domain $\{0,1\}^n$ into sets $\mathcal{X}_1, \ldots, \mathcal{X}_n$ (i.e., $\mathcal{X}_1 \cup \ldots \cup \mathcal{X}_n = \{0,1\}^n$) such that $\mathcal{X}_j \stackrel{\mathsf{def}}{=} \{x : 2^{j-1} \leq |f^{-1}(f(x))| < 2^j\}$, and define function $\max(n) \stackrel{def}{=} \max\{j : |\mathcal{X}_j| > 0\}$, i.e.,*

$|\mathcal{X}_{\max(n)}| > 0$ *and* $|\mathcal{X}_{\max(n)+1} \cup \ldots \cup \mathcal{X}_n| = 0$. *We say that* f *is* ***weakly unknown-regular*** *if there exists a constant* c *such that for all sufficiently large* n*'s :*

$$\Pr[U_n \in \mathcal{X}_{\max(n)}] \geq n^{-c} \quad . \tag{1}$$

Note that $\max(\cdot)$ *can be arbitrary (not necessarily efficient) functions and thus unknown-regular one-way functions fall into a special case*[7] *for* $c = 0$.

Definition 2.7 (UOWHFs [19]) *A sequence of function family* $\mathcal{G} = \{\mathcal{G}_n\}_{n\in\mathbb{N}}$, *where* $\mathcal{G}_n = \{g_u : \{0,1\}^{\ell(n)} \to \{0,1\}^{\ell(n)-s(n)}, u \in \{0,1\}^{q(n)}, \ell \in \mathsf{poly}\}$, *is a family of* $(t(n),\varepsilon(n))$*-universal one-way hash functions if for every* $n \in \mathbb{N}$, $u \in \{0,1\}^{q(n)}$ *and* $x \in \{0,1\}^{\ell(n)}$, *the value* $g_u(x)$ *can be computed in polynomial time, and for every probabilistic algorithm* A *of running time* $t(n)$, *it holds that*

$$\Pr_{x\xleftarrow{\$}\{0,1\}^{\ell(n)};\ u\xleftarrow{\$}\{0,1\}^{q(n)};\ x'\leftarrow\mathsf{A}(1^n,x,u)}[\ x \neq x' \wedge g_u(x) = g_u(x')\ ] \ \leq \ \varepsilon(n) \quad .$$

The difference between input and output lengths (i.e., $s(n)$*) is called* ***shrinkage****. Standard asymptotic security requires* $t(\cdot)$ *and* $1/\varepsilon(\cdot)$ *to be super-polynomial. For succinctness, hereafter we will use shorthand* $\mathcal{G} = \{g_u : \{0,1\}^{\ell(n)} \to \{0,1\}^{\ell(n)-s(n)}, u \in \{0,1\}^{q(n)}\}$ *for* $\{\mathcal{G}_n\}_{n\in\mathbb{N}}$ *defined above.*

3 UOWHFs from 1-to-1 One-way Functions

3.1 A Technical Lemma

We introduce Lemma 3.1 below with nice duality to the leftover hash lemma which will be useful in our constructions and might be of independent interest. We mention that the lemma actually generalizes to almost flat sources (stated as Lemma A.1 in Appendix A), which is used in the proof of Fact 3.1.

Lemma 3.1 (The injective hash lemma) *For any integers* a, d, k *and* l *satisfying* $a \leq l$, *let* Y *be any uniform distribution over some set* $\mathcal{Y} \subseteq \{0,1\}^l$ *of size* 2^a, *and let* $\mathcal{H} \stackrel{\mathsf{def}}{=} \{h : \{0,1\}^l \to \{0,1\}^{a+d}\}$ *be a family of* $(k\cdot 2^{-(a+d)})$*-almost universal hash functions. Then, we have that*

$$\Pr_{y\leftarrow Y,\ h\xleftarrow{\$}\mathcal{H}}[\ \exists \tilde{y} \in \mathcal{Y} :\ \tilde{y} \neq y \wedge h(\tilde{y}) = h(y)\ ] \leq k\cdot 2^{-d} \quad .$$

Recall that $k = 1$ *corresponds to the special case that* $\mathcal{H}$ *is universal.*

Proof. The (almost) universality of $\mathcal{H}$ implies an upper bound on $\mathsf{CP}(H(Y)|H)$, i.e.,

$$\mathsf{CP}(\ H(Y) \mid H\ ) \leq \mathsf{CP}(Y) + \max_{y_1\neq y_2}\{\ \Pr_{h\xleftarrow{\$}\mathcal{H}}[\ h(y_1) = h(y_2)\ ]\ \} \ = \ 2^{-a}(1 + k\cdot 2^{-d}) \quad .$$

where we consider the random experiment of sampling y_1 and y_2 i.i.d. to Y and thus the collision probability of $H(Y)$ given H is bounded by the sum of $\Pr[Y_1 = Y_2]$ and $\Pr[H(y_1) = H(y_2)]$ for any $y_1 \neq y_2$.

[7] In fact, our construction #4 only requires that f is **weakly unknown almost-regular**, i.e., $\Pr[U_n \in \mathcal{X}_{\max(n)-O(\log n)} \cup \ldots \cup \mathcal{X}_{\max(n)}] \geq n^{-c}$ instead of (1), where unknown-almost-regular one-way functions become a special case for $c = 0$.

Further, denote $\mathcal{S}_1 \stackrel{\mathsf{def}}{=} \{(z,h) : |\{\tilde{y} \in \mathcal{Y} : h(\tilde{y}) = z\}| = 1\}$ and $\mathcal{S}_2 \stackrel{\mathsf{def}}{=} \{(z,h) : |\{\tilde{y} \in \mathcal{Y} : h(\tilde{y}) = z\}| \geq 2\}$, and we have the following lower bound

$$\begin{aligned}
&\mathsf{CP} \quad (\ H(Y) \mid H\ ) \\
&= \quad \sum_h \Pr[H=h] \left(\sum_{z:(z,h)\in\mathcal{S}_1} \Pr[H(Y)=z|H=h]^2 + \sum_{z:(z,h)\in\mathcal{S}_2} \Pr[H(Y)=z|H=h]^2 \right) \\
&\geq \quad 2^{-a} \cdot \sum_{(z,h)\in\mathcal{S}_1} \Pr[H=h, H(Y)=z] \ + \ \min_{(z,h)\in\mathcal{S}_2} \{\ \Pr[h(Y)=z]\ \} \cdot \sum_{(z,h)\in\mathcal{S}_2} \Pr[H=h, H(Y)=z] \\
&= \quad 2^{-a} \cdot \Pr[\ (H(Y),H) \in \mathcal{S}_1\ ] \ + \ 2^{-a+1} \cdot \Pr[\ (H(Y),H) \in \mathcal{S}_2\ ] \\
&= \quad 2^{-a}(1 + \Pr[\ (H(Y),H) \in \mathcal{S}_2]) \quad ,
\end{aligned}$$

where the inequality is due to that any $(z,h) \in \mathcal{S}_1$ satisfies $\Pr[h(Y)=z] = 2^{-a}$, and for any $(z,h) \in \mathcal{S}_2$ we have $\Pr[h(Y)=z] \geq 2^{-a+1}$ (recall that Y is uniform over $\mathcal{Y}$ by assumption). Taking into account both the lower and upper bounds on $\mathsf{CP}(H(Y)|H)$, we get $\Pr[\ (H(Y),H) \in \mathcal{S}_2\ ] \leq k \cdot 2^{-d}$ and thus complete the proof. □

3.2 Simplifying Assumption about Output Length

We argue that the input and output lengths of a 1-to-1 one-way function $f : \{0,1\}^n \to \{0,1\}^{l(n)}$ can be assumed to be linearly related (i.e., $l(n) \in O(n)$) without loss of generality. For almost regular one-way functions, we can even assume that they are length-preserving (i.e., $l(n) = n$). We state it as Fact 3.1 with proof given in Appendix A.

Fact 3.1 (two folklore facts) *For any $r_1 = r_1(n) \leq r_2 = r_2(n)$ and any efficiently computable $\kappa = \kappa(n) \in O(n)$, we have*

1. *Any 1-to-1 (t,ε)-one-way function $f : \{0,1\}^n \to \{0,1\}^l$ implies a $(t - n^{O(1)}, \varepsilon + \mathsf{poly}(n) \cdot 2^{-\kappa})$-one-way function $f' : \{0,1\}^{n' \in \Theta(n)} \to \{0,1\}^{(n'+\kappa) \in \Theta(n)}$ which is 1-to-1 except on a $(\mathsf{poly}(n) \cdot 2^{-\kappa})$-fraction of inputs, i.e.,*

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^{n'}} [\ \exists x' \in \{0,1\}^{n'} : x' \neq x \ \wedge \ f'(x) = f'(x')\ ] \ \leq \ \mathsf{poly}(n) \cdot 2^{-\kappa}$$

2. *Any $(2^{r_1}, 2^{r_2})$-almost regular (t,ε)-one-way function $f : \{0,1\}^n \to \{0,1\}^l$ implies a length-preserving $(t - n^{O(1)}, \varepsilon + \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa)})$-one-way function $\bar{f} : \{0,1\}^{n' \in \Theta(n)} \to \{0,1\}^{n'}$ which is $(2^{n+\kappa+r_1}, 2^{n+\kappa+r_2})$-almost regular except on a $(\mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa)})$-fraction of inputs, i.e.,*

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^{n'}} [\ 2^{n+\kappa+r_1} \ \leq \ |\bar{f}^{-1}(\bar{f}(x))| \ \leq \ 2^{n+\kappa+r_2}\ ] \ \geq \ 1 - \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa)} \quad .$$

Note that $2^{r_2-r_1}$ can be arbitrarily large (i.e., not necessarily bounded by $\mathsf{poly}(n)$), and thus the second statement above applies to any one-way function f. It suffices to set $\kappa = \omega(\log n)$ to have a negligible error bound, and in case $\kappa = \Theta(n)$ the bound will be exponentially small.

3.3 UOWHFs from 1-to-1 OWFs

We will assume in the remainder of the paper that the underlying 1-to-1 one-way function has linear output length (i.e., $l(n) \in O(n)$) and that the almost-regular and weakly unknown-regular one-way functions in consideration are length-preserving (i.e., $l(n) = n$). We first state a fact about the hard-to-invertness of f, and then adapts the Naor-Yung construction to any 1-to-1 one-way functions.

Fact 3.2 *For any 1-to-1 $(t,\ \varepsilon)$-one-way function $f:\{0,1\}^n\to\{0,1\}^l$ and any probabilistic algorithm Inv of running time t, it holds that*

$$\Pr_{y^*\xleftarrow{\$}\{0,1\}^l}\left[\ f(\mathsf{Inv}(y^*))=y^*\ \right]\ \le\ 2^{-(l-n)}\cdot\varepsilon\quad .$$

Proof.

$$\Pr_{y^*\xleftarrow{\$}\{0,1\}^l}\left[\ f(\mathsf{Inv}(y^*))=y^*\ \right]\le\ \Pr_{y^*\xleftarrow{\$}\{0,1\}^l}\left[y^*\in f(\{0,1\}^n)\right]\cdot\ \Pr_{y^*\xleftarrow{\$}f(\{0,1\}^n)}\left[\ f(\mathsf{Inv}(y^*))=y^*\ \right]\le\ 2^{-(l-n)}\cdot\varepsilon\quad .$$

□

Theorem 3.1 (UOWHFs from 1-to-1 OWFs) *Let $f:\{0,1\}^n\to\{0,1\}^{l\in O(n)}$ be any 1-to-1 $(t,\ \varepsilon)$-one-way function, let $\mathcal{H}$ be a family of permutations*[8] *over $\{0,1\}^l$ as follows:*

$$\mathcal{H}=\{h:\{0,1\}^l\to\{0,1\}^l\ ,\ h(y)\overset{\mathsf{def}}{=}h\cdot y,\quad \textit{where}\ \ y\in GF(2^l),\ \vec{0}\neq h\in GF(2^l)\ \}\ ,$$

let $\mathsf{trunc}:\{0,1\}^l\to\{0,1\}^{n-s}$ be a truncating function, where $s=s(n)$ is efficiently computable. Then, we have that

$$\mathcal{G}_1\overset{\mathsf{def}}{=}\{\ (\mathsf{trunc}\circ h\circ f\ ):\{0,1\}^n\to\{0,1\}^{n-s(n)}\ ,\ h\in\mathcal{H}\ \}$$

is a family of $(t-n^{O(1)},\ 2^s\cdot\varepsilon)$-UOWHFs with shrinkage s, key and output length $\Theta(n)$.

Proof. Suppose for contradiction that there exists a $\mathcal{G}_1$-collision finder A of running time t' that on input (x,h), breaks the target collision resistance with some non-negligible probability ε', i.e.,

$$\Pr_{x\xleftarrow{\$}\{0,1\}^n,h\xleftarrow{\$}\mathcal{H}}\left[\ x'\leftarrow\mathsf{A}(x,h):\ x\neq x'\wedge h(f(x))_{[n-s]}=h(f(x'))_{[n-s]}\ \right]\ =\ \varepsilon'\ >\ 2^s\cdot\varepsilon$$

We define algorithm $\mathsf{Inv}^{\mathsf{A}}$ (that inverts f on input $y^*\xleftarrow{\$}\{0,1\}^l$ by invoking A) as in Algorithm 1. By Claim 3.1, conditioned on $f(x)\neq y^*$ it is equivalent to consider that $\mathsf{Inv}^{\mathsf{A}}$ samples (x,h,v) from $\{0,1\}^n\times\mathcal{H}\times\mathcal{V}$ uniformly and independently, and then determines the value of y^*. We argue that $\mathsf{Inv}^{\mathsf{A}}$ inverts f with the following probability (see the rationale below)

$$\begin{aligned}
&\Pr_{y^*\xleftarrow{\$}\{0,1\}^l,\ x\xleftarrow{\$}\{0,1\}^n,\ v\xleftarrow{\$}\mathcal{V}}\left[\ f(\mathsf{Inv}^{\mathsf{A}}(y^*))=y^*\ \right]\\
\ge\ &\Pr_{x\xleftarrow{\$}\{0,1\}^n,y^*\xleftarrow{\$}\{0,1\}^l}\left[\ f(x)=y^*\ \right]\ +\ \Pr_{x\xleftarrow{\$}\{0,1\}^n,y^*\xleftarrow{\$}\{0,1\}^l}\left[\ f(x)\neq y^*\ \right]\\
&\times\ \Pr_{x\xleftarrow{\$}\{0,1\}^n,h\xleftarrow{\$}\mathcal{H},x\neq(x'\leftarrow\mathsf{A}(x,h)),v\xleftarrow{\$}\mathcal{V}}\left[\ h(f(x))_{[n-s]}=h(f(x'))_{[n-s]}\ \wedge\ y^*=f(x')\mid f(x)\neq y^*\ \right]\\
\ge\ &2^{-l}\ +\ (1-2^{-l})\cdot\varepsilon'\cdot\ \Pr_{v\xleftarrow{\$}\mathcal{V}}\left[y^*=f(x')\mid f(x)\neq y^*\ \wedge\ f(x)\neq f(x')\wedge\ h(f(x))_{[n-s]}=h(f(x'))_{[n-s]}\ \right]\\
=\ &2^{-l}\ +\ (1-2^{-l})\cdot\varepsilon'\cdot\frac{1}{|\mathcal{V}|}\ =\ 2^{-l}\ +\ (1-2^{-l})\cdot\varepsilon'\cdot\frac{1}{2^{l-n+s}-1}\ >\ \varepsilon'\cdot 2^{-(l-n+s)}\ >\ \varepsilon\cdot 2^{-(l-n)}\ ,
\end{aligned}$$

where A takes only x and h as input (i.e., independent of v), and thus conditioned on that A produces a valid $x'\neq x$ satisfying $h(f(x'))_{[n-s]}=h(f(x))_{[n-s]}$, we have by Claim 3.1 that string y^* is uniformly distributed over set $\mathcal{Y}^*\overset{\mathsf{def}}{=}\{y^*:y^*=f(x)-v\cdot h^{-1},v\in\mathcal{V}\}$. Note that the already fixed $f(x')$ is also an element of $\mathcal{Y}^*$ and thus y^* hits $f(x')$ with probability $1/|\mathcal{Y}^*|=1/|\mathcal{V}|=1/(2^{l-n+s}-1)$. We thus complete the proof by reaching a contradiction to Fact 3.2. □

[8] In fact, $\mathcal{H}$ constitutes a family of universal hash permutations. However, our proofs will only use the concrete construction of $\mathcal{H}$ and benefit from its algebraic property over finite fields, rather than assuming a general universal $\mathcal{H}$ plus a constructible property [14] (i.e., given any x and y there exists a PPT that outputs $h\xleftarrow{\$}\{h\in\mathcal{H}:h(x)=y\}$) .

Algorithm 1 $\mathsf{Inv}^{\mathsf{A}}$ that inverts f on input y^* using random coins (x, v).

Input: $y^* \xleftarrow{\$} \{0,1\}^l$

Sample $x \xleftarrow{\$} \{0,1\}^n$
if $f(x) = y^*$ **then**
  Output x and **terminate**.
end if
sample $h := (f(x) - y^*)^{-1} \cdot v$, where $v \xleftarrow{\$} \mathcal{V} = \{v \in \{0,1\}^l \setminus \{\vec{0}\} : v_{[n-s]} = \overbrace{0 \ldots 0}^{n-s}\}$
{*note*: The above implies $h \xleftarrow{\$} \{h \in \mathcal{H} : \ h(f(x))_{[n-s]} = h(y^*)_{[n-s]}\}$ by the $GF(2^l)$ arithmetics. }
$x' \leftarrow \mathsf{A}(x, h)$
if $f(x')=y^*$ **then**
  Output x'
else
  Output $\perp$
end if
Terminate

Claim 3.1 (equivalent sampling) *Let the values h, v, x, y^* be sampled as in Algorithm 1 (or as in Algorithm 3), and conditioned on the event $y^* \neq f(x)$, it is equivalent to sample $(x,\ h,\ v) \xleftarrow{\$} \{0,1\}^n \times \mathcal{H} \times \mathcal{V}$ uniformly and independently and then determine $y^* := f(x) - v \cdot h^{-1}$.*

Proof of Claim 3.1. We known that (x, v) is uniformly sampled from $\{0,1\}^n \times \mathcal{V}$ by definition, and thus it suffices to show that "fix any (x, v), and conditioned on $y^* \neq f(x)$ (i.e., Y^* is uniform distributed over $\{0,1\}^l \setminus \{f(x)\}$), it holds that h is uniform over $\mathcal{H}$". As $v \neq \vec{0}$ ($\mathcal{V}$ excludes $\vec{0}$ by definition), it follows that $h = (f(x) - Y^*)^{-1} \cdot v$ is uniform over $\{0,1\}^l \setminus \{\vec{0}\}$. Finally, for any given $(x,\ h,\ v)$, one efficiently determines the value $y^* = f(x) - v \cdot h^{-1}$ due to the arithmetics over the finite field. □

Remark 3.1 (on the optimal security bounds.) *Theorem 3.1 enjoys optimal security degradations, in particular, the collision resistance deteriorates exponentially only with respect to shrinkage s (which is optimal by [5]), i.e., not to the number of bits truncated (i.e., $l-n+s$). This is due to the fact that we reduce the collision-finding problem to that of inverting a random y^* over $\{0,1\}^l$, where the probability that y^* is valid image (i.e., over $f(\{0,1\}^n)$) is $2^{-(l-n)}$ and thus cancels the factor $(l-n)$.*

We also state a simple corollary of Theorem 3.1 below where the underlying one-way function f is 1-to-1 except for a negligible fraction. See its proof in Appendix A.

Corollary 3.1 (UOWHFs from almost 1-to-1 OWFs) *Let f, $\mathcal{H}$, trunc and $\mathcal{G}$ be the same as assumed (or defined) in Theorem 3.1 except that f is $\delta(n)$-almost 1-to-1 (instead of perfectly 1-to-1), where $\delta(n) \leq 1/2$. Then, $\mathcal{G}_1$ is a family of $(t - n^{O(1)},\ 2^{s+1} \cdot \varepsilon + \delta)$-universal one-way hash functions with shrinkage $s(n)$, key and output length $\Theta(n)$.*

4 UOWHFs from Known Regular OWFs

We proceed to the more general case that f is a known almost-regular function. Recall that by Fact 3.1 we can assume WLOG that the underlying almost regular one-way function is length-preserving. We first show an optimal construction where the hardness parameter ε is known.

4.1 Compressing the Output is Necessary but Not Sufficient

We attempt to generalize the Naor-Yung approach for one-way permutations (and 1-to-1 one-way functions) to almost regular one-way functions by compressing (using $\mathsf{trunc} \circ h$) the output $Y = f(X)$ into $\mathbf{H}_\infty(Y) - s'$ bits, where $\mathbf{H}_\infty(Y)$ denotes the min-entropy of Y and $s' \in O(\log(1/\varepsilon))$. However, this only gives a weak form of guarantee, as stated in Lemma 4.1 below, that given a random x it is infeasible for efficient algorithms to find any $f(x') \neq f(x)$ such that $\mathsf{trunc}(h(f(x'))) = \mathsf{trunc}(h(f(x)))$. Otherwise said, it does not rule out the possibility that one may easily find $x' \neq x$ satisfying $f(x') = f(x)$. Hence, compressing the output is only a useful intermediate step to obtain UOWHFs. Lemma 4.1 below further generalizes Theorem 3.1 to known-(almost-)regular functions, whose proof is similar to that of Theorem 3.1 and thus we defer it to Appendix A to avoid redundancy.

Lemma 4.1 *For any constant c, and any efficiently computable $r = r(n)$ and $s' = s'(n)$, let $f : \{0,1\}^n \to \{0,1\}^n$ be any $(2^r, 2^r n^c)$-almost regular (length-preserving) (t, ε)-one-way function, let $\mathcal{H}$ be a family of universal hash permutations over $\{0,1\}^n$, i.e.,*

$$\mathcal{H} = \{h : \{0,1\}^n \to \{0,1\}^n \ , \ h(y) \stackrel{\mathsf{def}}{=} h \cdot y, \quad \textit{where} \quad y \in GF(2^n), \ \vec{0} \neq h \in GF(2^n) \ \} \ ,$$

let $\mathsf{trunc} : \{0,1\}^n \to \{0,1\}^{n-r-c\cdot\log n-s'}$ be a truncating function. Then, for any $\tilde{\mathsf{A}}$ of running time $t - n^{O(1)}$ (for some universal constant $O(1)$) we have that

$$\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n, \ h \stackrel{\$}{\leftarrow} \mathcal{H}} [\ x' \leftarrow \tilde{\mathsf{A}}(x,h) \ \wedge \ f(x) \neq f(x') \ \wedge \ \mathsf{trunc}(h(f(x))) = \mathsf{trunc}(h(f(x'))) \ ] \ \leq \ n^c \cdot 2^{s'} \cdot \varepsilon \ .$$

4.2 UOWHFs from Known (Almost-)Regular OWFs with Known Hardness

We first give an optimal construction assuming that the inversion probability upper bound ε is known. Note that in addition to hashing the output $f(x)$ (as we did in Lemma 4.1), we also hash the input x to ensure that no distinct x' collides with x with respect to the resulting function.

Theorem 4.1 (UOWHFs from known almost-regular OWFs with known ε) *For constant c, let $f : \{0,1\}^n \to \{0,1\}^n$ be any $(2^r, 2^r n^c)$-almost regular (length-preserving) (t,ε)-one-way function, where $r = r(n)$ and $\varepsilon = \varepsilon(n)$ are any efficiently computable functions. Let shrinkage $s = s(n)$ be any efficiently computable function, and let $\mathcal{H}$ and trunc be as defined in Lemma 4.1 with $s' = (s + \log(1/\varepsilon) - c\log n)/2$, and let $\mathcal{H}_1 = \{h_1 : \{0,1\}^n \to \{0,1\}^{r+c\log n+s'-s}\}$ be a family of universal hash functions. Then, we have that*

$$\mathcal{G}_2 \stackrel{\mathsf{def}}{=} \{ \ g : \{0,1\}^n \to \{0,1\}^{n-s} \ , \ g(x) \stackrel{\mathsf{def}}{=} (g_1(x), h_1(x)), g_1 \stackrel{\mathsf{def}}{=} (\mathsf{trunc} \circ h \circ f), \ h \in \mathcal{H} \ , h_1 \in \mathcal{H}_1 \ \}$$

is a $(t - n^{O(1)}, O(\sqrt{2^s \cdot n^c \cdot \varepsilon}))$-universal one-way hash function family with key and output length $\Theta(n)$.

Proof. Define shorthands $\mathcal{E}_1 \stackrel{\mathsf{def}}{=} (x \neq x' \wedge f(x) = f(x') \ \wedge h_1(x) = h_1(x'))$ and $\mathcal{E}_2 \stackrel{\mathsf{def}}{=} (f(x) \neq f(x') \ \wedge g_1(x) = g_1(x'))$. For any $\mathcal{G}_2$-collision finder A, we have (with explanations below)

$$\begin{aligned}
&\Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n, \ (h,h_1) \stackrel{\$}{\leftarrow} (\mathcal{H},\mathcal{H}_1), \ x' \leftarrow \mathsf{A}(x,h,h_1)} [\ x \neq x' \ \wedge \ g(x) = g(x') \ ] \\
&\leq \Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n, \ (h,h_1) \stackrel{\$}{\leftarrow} (\mathcal{H},\mathcal{H}_1), \ x' \leftarrow \mathsf{A}(x,h,h_1)} [\ \mathcal{E}_1 \ \vee \ \mathcal{E}_2 \ ] \\
&\leq \Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n, \ h_1 \stackrel{\$}{\leftarrow} \mathcal{H}_1} [\ \exists \ x' \neq x \ \wedge \ f(x) = f(x') \ \wedge \ h_1(x) = h_1(x') \ ] \\
&\quad + \Pr_{x \stackrel{\$}{\leftarrow} \{0,1\}^n, \ (h,h_1) \stackrel{\$}{\leftarrow} (\mathcal{H},\mathcal{H}_1), \ x \neq (x' \leftarrow \mathsf{A}(x,h,h_1))} [\ f(x) \neq f(x') \ \wedge \ g_1(x) = g_1(x') \ ] \\
&\leq 2^{-(s'-s)} \ + \ n^c \cdot 2^{s'} \cdot \varepsilon \ = \ \sqrt{2^s \cdot n^c \cdot \varepsilon} \ + \ \sqrt{2^s \cdot n^c \cdot \varepsilon} \ = \ 2\sqrt{2^s \cdot n^c \cdot \varepsilon} \ ,
\end{aligned}$$

where the second inequality follows by a union bound, namely, for a random x, if there is some $x' \neq x$ colliding on $g \in \mathcal{G}_2$ then it must satisfy either $\mathcal{E}_1$ or $\mathcal{E}_2$. We already know by Lemma 4.1 that the second term is bounded by $n^c \cdot 2^{s'}\varepsilon$, and thus it remains to show that the first term is bounded by $2^{-(s'-s)}$. Conditioned on any $y = f(X)$ random variable X is a flat distribution on a set of size at most $2^r \cdot n^c$, so we apply Lemma 3.1 (setting $a \leq r + c \cdot \log n$, $d \geq s' - s$ and $k = 1$) to get

$$\begin{aligned}
&\Pr_{x \xleftarrow{\$} \{0,1\}^n,\ h_1 \xleftarrow{\$} \mathcal{H}_1} [\ \exists\ x' \neq x\ \wedge\ f(x) = f(x')\ \wedge\ h_1(x) = h_1(x')\ ] \\
&= \mathbb{E}_{y \leftarrow f(U_n)} [\ \Pr_{x \xleftarrow{\$} f^{-1}(y),\ h_1 \xleftarrow{\$} \mathcal{H}_1} [\ \exists\ x' \neq x\ \wedge\ f(x) = f(x')\ \wedge\ h_1(x) = h_1(x')\ ]\ ] \\
&\leq \mathbb{E}_{y \leftarrow f(U_n)} [\ 2^{-(s'-s)}\ ] = 2^{-(s'-s)}\ .
\end{aligned}$$

which completes the proof. □

4.3 An Alternative Approach to Section 4.2

A neater (and perhaps more intuitive) approach is to construct an almost 1-to-1 one-way function f' (with input and output lengths $\Theta(n)$) based on f (stated as Theorem 4.2) and then plug f' into Corollary 3.1 (using f' in place of f). This statement is interesting in its own right as it implies that almost 1-to-1 one-way functions and known-(almost-)regular one-way functions (with known hardness) are equivalent. Taking a closer look at Theorem 4.2 we find that this almost 1-to-1 f' is also present (as an intermediate function) in construction $\mathcal{G}_2$ of Theorem 4.1 (except with slightly different length parameters). Lemma 4.2 and Lemma 4.3 state the almost injectiveness and one-way-ness of f' respectively, for which we determine a judicious value for d (based on ε) in Theorem 4.2 to achieve injectiveness and one-way-ness simultaneously. However, we could not adapt this approach to the case when ε is unknown, and thus we mainly focus on the former construction (as in Theorem 4.1) and extend it to any known(-almost)-regular one-way functions in Section 4.4.

Theorem 4.2 (almost 1-to-1 OWF f' from almost regular OWF f with known ε) *For any constant c, and any efficiently computable $r = r(n)$ and $d = d(n) \in O(n)$, let $f : \{0,1\}^n \to \{0,1\}^n$ be any $(2^r, 2^r n^c)$-almost regular (length-preserving) (t, ε)-one-way function, define*

$$f' : \{0,1\}^n \times \mathcal{H}_1 \to \{0,1\}^n \times \{0,1\}^{r + c\cdot\log n + d} \times \mathcal{H}_1$$

$$f'(x, h_1) \stackrel{def}{=} (f(x), h_1(x), h_1)$$

where $\mathcal{H}_1$ is a family of universal hash functions from n bits to $r + c\cdot\log n + d$ bits. Then, for $d = \frac{\log(1/\varepsilon) - c\cdot\log n - 3}{3}$ we have that f' is $2\sqrt[3]{\varepsilon \cdot n^c}$-almost 1-to-1 and $(t - O(n),\ 2\sqrt[3]{\varepsilon \cdot n^c})$-one-way with input and output lengths $\Theta(n)$.

Proof. The almost 1-to-1-ness and one-way-ness of f' follow from Lemma 4.2 and Lemma 4.3 respectively by setting parameter $d = \frac{\log(1/\varepsilon) - c\cdot\log n - 3}{3}$. □

The proofs of Lemma 4.2 and Lemma 4.3 are given in Appendix A due to lack of space.

Lemma 4.2 (f' is almost 1-to-1) *f' defined in Theorem 4.2 is 2^{-d}-almost 1-to-1.*

Lemma 4.3 (f' is one-way) *f' defined in Theorem 4.2 is a $(t - O(n),\ \sqrt{2^{d+3} \cdot n^c \cdot \varepsilon})$-one-way function.*

4.4 UOWHFs from any Known (Almost-)Regular OWFs

REMOVING THE DEPENDENCY ON ε. Unfortunately, Theorem 4.1 doesn't immediately apply to an arbitrary regular function as in general we assume no knowledge about the hardness parameter ε (other than that ε is negligible). To see the difficulty, consider the proof of Theorem 4.1 where the security of the resulting UOWHF is bounded by the sum of two terms, i.e., $2^{-(s'-s)} + n^c \cdot 2^{s'} \cdot \varepsilon$. Without knowing ε, one may end up setting some super-polynomial $2^{s'}$ (to make the first term negligible) which kills the second term $n^c \cdot 2^{s'} \cdot \varepsilon$. Same problems arise in similar situations (e.g., construction of PRGs from regular OWFs [28]). A remedy for this is parallel repetition: for any efficiently computable $q \in \omega(1)$, run q copies of f, apply hashing and truncating functions (setting $s' = 2\log n$) to every $f(x)$ (to get a bound $O(\varepsilon \cdot n^{c+2})$), which shrinks the entropies by $2q \log n$ bits, and finally apply a single hashing (to the q inputs of f jointly) that expands $q \cdot \log n$ bits (to yield another negligible term n^{-q}). This gives a family of UOWHFs with shrinkage $2q \log n - q \log n = q \log n$, and key and output length $O(q \cdot n)$ for any (efficiently computable) super-constant q.

Definition 4.1 (parallel repetition) *For any function $g : \mathcal{X} \to \mathcal{Y}$, we define its q-fold parallel repetition $g^q : \mathcal{X}^q \to \mathcal{Y}^q$ as*

$$g^q(x_1, ..., x_q) = (\ g(x_1)\ , ...,\ g(x_q)\ ) \quad .$$

For simplicity, we will use shorthand $\vec{x} \stackrel{def}{=} (x_1, \ldots, x_q)$ and thus $g^q(\vec{x}) = g^q(x_1, \ldots, x_q)$.

Theorem 4.3 (UOWHFs from any known almost-regular OWFs) *Let $f : \{0,1\}^n \to \{0,1\}^n$ be any $(2^r, 2^r n^c)$-almost regular (length-preserving) (t,ε)-one-way function, where c is any constant, and $r = r(n)$ is any efficiently computable function. Then, for any efficiently computable $q = q(n) \in \omega(1)$, let $\mathcal{H}$ and trunc be as defined in Lemma 4.1 with $s' = 2\log n$, and let $\mathcal{H}_1 = \{h_1 : \{0,1\}^{q \cdot n} \to \{0,1\}^{q(r+(c+1)\log n)}\}$ be a family of universal hash functions, we have that*

$$\mathcal{G}_3 \stackrel{\text{def}}{=} \{\ g : \{0,1\}^{qn} \to \{0,1\}^{qn - q\log n}\ ,\ g(\vec{x}) \stackrel{\text{def}}{=} (g_1(\vec{x}), h_1(\vec{x})), g_1 \stackrel{\text{def}}{=} (\mathsf{trunc} \circ h \circ f)^q,\ h \in \mathcal{H}\ , h_1 \in \mathcal{H}_1\ \}$$

is a $(t - n^{O(1)}, n^{-q} + q \cdot n^{c+2} \cdot \varepsilon)$-universal one-way hash function family with key and output length $O(q \cdot n)$, and shrinkage $q \cdot \log n$.

Proof. Similar to the proof of Theorem 4.1, define $\mathcal{E}_1 \stackrel{\text{def}}{=} \left(\vec{x} \neq \vec{x'} \wedge f^q(\vec{x}) = f^q(\vec{x'})\ \wedge h_1(\vec{x}) = h_1(\vec{x'})\right)$ and $\mathcal{E}_2 \stackrel{\text{def}}{=} \left(f^q(\vec{x}) \neq f^q(\vec{x'})\ \wedge g_1(\vec{x}) = g_1(\vec{x'})\right)$, we have

$$\begin{aligned}
&\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ (h,h_1) \xleftarrow{\$} (\mathcal{H},\mathcal{H}_1),\ \vec{x'} \leftarrow \mathsf{A}(\vec{x},h,h_1)} [\ \vec{x} \neq \vec{x'}\ \wedge\ g(\vec{x}) = g(\vec{x'})] \\
\leq\ &\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ (h,h_1) \xleftarrow{\$} (\mathcal{H},\mathcal{H}_1),\ \vec{x'} \leftarrow \mathsf{A}(\vec{x},h,h_1)} [\ \mathcal{E}_1\ \vee \mathcal{E}_2\ ] \\
\leq\ &\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ h_1 \xleftarrow{\$} \mathcal{H}_1} [\ \exists\ \vec{x'} \neq \vec{x}\ \wedge\ f^q(\vec{x}) = f^q(\vec{x'})\ \wedge\ h_1(\vec{x}) = h_1(\vec{x'})\ ] \\
&+ \Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ (h,h_1) \xleftarrow{\$} (\mathcal{H},\mathcal{H}_1)} [\ \vec{x'} \leftarrow \mathsf{A}(\vec{x},h,h_1)\ \wedge\ f^q(\vec{x}) \neq f^q(\vec{x'})\ \wedge\ g_1(\vec{x}) = g_1(\vec{x'})\ ] \\
\leq\ &2^{-q\log n} + q \cdot n^{c+2} \cdot \varepsilon\ =\ n^{-q} + q \cdot n^{c+2} \cdot \varepsilon\ ,
\end{aligned}$$

where the second inequality follows by a union bound, and the first term of the third inequality is due to that conditioned on any $\vec{y} = f^q(\vec{X})$ random variable $\vec{X}$ is uniform over some set of size at most

$(2^r \cdot n^c)^q$, so we apply Lemma 3.1 (setting $a \leq q(r + c \cdot \log n)$, $d \geq q \log n$ and $k = 1$) to get

$$\begin{aligned}
&\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ h_1 \xleftarrow{\$} \mathcal{H}_1}\left[\ \exists\ \vec{x'} \neq \vec{x}\ \wedge\ f^q(\vec{x}) = f^q(\vec{x'})\ \wedge\ h_1(\vec{x}) = h_1(\vec{x'})\ \right] \\
&= \mathbb{E}_{\vec{y} \leftarrow f^q(U_{qn})}\Big[\ \Pr_{\vec{x} \xleftarrow{\$} (f^q)^{-1}(\vec{y}),\ h_1 \xleftarrow{\$} \mathcal{H}_1}\left[\ \exists\ \vec{x'} \neq \vec{x}\ \wedge\ f^q(\vec{x}) = f^q(\vec{x'})\ \wedge\ h_1(\vec{x}) = h_1(\vec{x'})\ \right]\ \Big] \\
&\leq \mathbb{E}_{\vec{y} \leftarrow f^q(U_{qn})}\left[\ 2^{-q \log n}\ \right] \ =\ n^{-q}\ .
\end{aligned}$$

We proceed to the proof of bounding the second term. Suppose for contradiction that there exists A_{g_1} of running time $t - n^{O(1)}$ such that

$$\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ (h,h_1) \xleftarrow{\$} (\mathcal{H},\mathcal{H}_1)}\left[\ \vec{x'} \leftarrow \mathsf{A}_{g_1}(\vec{x}, h, h_1)\ \wedge\ f^q(\vec{x}) \neq f^q(\vec{x'})\ \wedge\ g_1(\vec{x}) = g_1(\vec{x'})\ \right]\ >\ q \cdot n^{c+2} \cdot \varepsilon$$

Then, define $\tilde{\mathsf{A}}$ as in Algorithm 2. Conditioned on A_{g_1} finds a collision, i.e., $f^q(\vec{x}) \neq f^q(\vec{x'})$ and $g_1(\vec{x}) = g_1(\vec{x'})$, there exists at least one $i^* \in [q]$ satisfying $f(x_{i^*}) \neq f(x'_{i^*})$ and $\mathsf{trunc}(h(f(x_{i^*}))) = \mathsf{trunc}(h(f(x'_{i^*})))$. We have

Algorithm 2 ($\mathsf{trunc} \circ h$)-collision finder $\tilde{\mathsf{A}}$ on input (x, h).

Input: $(x, h) \xleftarrow{\$} \{0,1\}^n \times \mathcal{H}$

Sample $\vec{x} = (x_1, \ldots, x_q) \xleftarrow{\$} \{0,1\}^{qn}$, $h_1 \xleftarrow{\$} \mathcal{H}_1$, $i \xleftarrow{\$} [q]$
$\vec{x'} = (x'_1, \ldots, x'_q) \leftarrow \mathsf{A}_{g_1}(\ (x_1, \ldots, x_{i-1}, x, x_{i+1}, \ldots, x_q)\ , h, h_1)$ {i.e., replace x_i with x }
return x'_i

$$\begin{aligned}
&\Pr_{x \xleftarrow{\$} \{0,1\}^n,\ h \xleftarrow{\$} \mathcal{H}}\left[\ x' \leftarrow \tilde{\mathsf{A}}(x, h)\ \wedge\ f(x) \neq f(x')\ \wedge\ \mathsf{trunc}(h(f(x))) = \mathsf{trunc}(h(f(x')))\ \right] \\
\geq\ &\Pr_{\vec{x} \xleftarrow{\$} \{0,1\}^{qn},\ (h,h_1) \xleftarrow{\$} (\mathcal{H},\mathcal{H}_1)}\left[\ \vec{x'} \leftarrow \mathsf{A}_{g_1}(\vec{x}, h, h_1)\ \wedge\ f^q(\vec{x}) \neq f^q(\vec{x'})\ \wedge\ g_1(\vec{x}) = g_1(\vec{x'})\ \wedge\ i = i^*\ \right] \\
>\ &q \cdot n^{c+2} \varepsilon \cdot (1/q)\ =\ n^{c+2} \varepsilon\ ,
\end{aligned}$$

which is a contradiction to Lemma 4.1 (recall that $s' = 2 \log n$) and thus completes the proof. □

5 UOWHFs from Regular OWFs and Beyond

5.1 UOWHFs from Any (Almost-)Regular OWFs

Ames et al. [1] presented an elegant construction based on any almost-regular OWFs[9], where no knowledge is required about the regularity of the OWF. Furthermore, their construction enjoys output length $\Theta(n)$ and key length $O(n \cdot \log n)$ and makes $O(n/\log n)$ calls to the underlying OWF. To see this, we set $s = \Omega(\log n)$ in Theorem 5.1 and thus get a construction of UOWHFs by making $\kappa = O(n/\log n)$ calls to any $(\alpha, \alpha \cdot \beta)$-almost regular (t, ε)-OWF, where α and β need not to be efficiently computable, and the construction tolerates regularity slackness for any $\beta = n^{O(1)}$ or even certain $\beta = (1/\varepsilon)^{O(1)}$. We note that the number of calls $O(n/\log n)$ is optimal (for black-box constructions) in general by matching the lower bound of [2].

[9] The authors of [1] mainly stated the neat case, i.e., for $\beta = 1$ and $s = 1$, and similar to Theorem 4.3 it (implicitly in their proof) generalizes to Theorem 5.1, where almost regularity and logarithmic shrinkage are considered.

Definition 5.1 (the generalized iterate [1]) *Let $n \in \mathbb{N}$, function $f : \{0,1\}^n \to \{0,1\}^n$, and let $\mathcal{H}$ be a family of pairwise-independent hash functions from $\{0,1\}^{n+s}$ to $\{0,1\}^n$. For $i \leq \kappa \in \mathbb{N}$, $x_1 \in \{0,1\}^n$, $v_1, \ldots, v_\kappa \in \{0,1\}^s$ and vector $\vec{h}^\kappa = (h_1, \ldots, h_\kappa) \in \mathcal{H}^\kappa$, recursively define the i^{th} randomized iterate by:*

$$x_1 \xrightarrow{f} y_1 \xrightarrow[v_1]{h_1} x_2 \xrightarrow{f} y_2 \xrightarrow[v_2]{h_2} \cdots \quad x_\kappa \xrightarrow{f} y_\kappa \xrightarrow[v_\kappa]{h_\kappa} x_{\kappa+1} \xrightarrow{f} y_{\kappa+1}$$

$$y_i = f(x_i),\ x_{i+1} = h_i(y_i \| v_i)$$

We denote the κ^{th} iterate by function g_f^κ, i.e., $y_{k+1} = g_f^\kappa(v_1 \| \ldots \| v_\kappa, x_1, \vec{h}^\kappa)$, where $x_1 \xleftarrow{\$} \{0,1\}^n$, $v_1, \ldots, v_\kappa \xleftarrow{\$} \{0,1\}^s$, $\vec{h}^\kappa \leftarrow Shoup(U_{O(n \cdot \log n)})$ and $Shoup : \{0,1\}^{O(n \cdot \log n)} \to \mathcal{H}^\kappa$ is Shoup's generator [24].

Theorem 5.1 (UOWHFs from unknown almost-regular OWFs [1]) *For security parameter $n \in \mathbb{N}$, any (not necessarily efficient) $\alpha = \alpha(n)$, $\beta = \beta(n) \geq 1$ and any efficiently computable $s = s(n)$, $\kappa = \kappa(n)$ such that $s(n) \cdot \kappa(n) \geq n + s(n)$, let $f : \{0,1\}^n \to \{0,1\}^n$ be any $(\alpha, \alpha \cdot \beta)$-almost regular (length-preserving) (t, ε)-one-way function, let g_f^κ, $\mathcal{H}$ and $Shoup : \{0,1\}^{q \in O(n \cdot \log n)} \to \mathcal{H}^\kappa$ be defined as in Definition 5.1. Then, we have that*

$$\mathcal{G} = \{ g_f : \{0,1\}^{s \cdot \kappa} \to \{0,1\}^n \ ,\ g_f(z) = g_f^\kappa(z, x_1, Shoup(u)), z \stackrel{def}{=} v_1 \| \ldots \| v_\kappa, x_1 \in \{0,1\}^n, u \in \{0,1\}^q \ \}$$

is a family of $(t - n^{O(1)},\ \mathsf{poly}(\beta, 2^s, \kappa) \cdot \varepsilon^{\Theta(1)})$-UOWHFs with key length $O(n \cdot \log n)$, output length n and at least s bits of shrinkage.
Notice that x_1 is not input to hash function g_f but the part (together with u) of the description of g_f.

5.2 UOWHFs from Weakly Unknown-Regular OWFs

Next, we introduce the construction where the underlying OWF can be far from regular, as long as the fraction of x's that have maximal number of siblings is noticeable. The proof of the theorem below is deferred to Section 5.5, where we put together all the necessary technical ingredients.

Theorem 5.2 *Assume that f is a weakly unknown-regular one-way function on a noticeable fraction (i.e., n^{-c} for constant c) of domain. Then, there exists an explicit construction of UOWHF family (stated as Construction 5.1) with output length $\Theta(n)$, key length $O(n \cdot \log n)$ by making $n^{2c+1} \cdot \omega(1)$ black-box calls to f.*

5.3 An Explicit Construction

The main idea is to transform any weakly unknown-regular one-way function f into a family of functions $\mathcal{F} = \{f_u : u \in \{0,1\}^{O(n \log n)}\}$ such that $\mathcal{F}$ is almost regular and that it preserves the one-way-ness of f. $\mathcal{F}$ is constructed with a succinct description u based on (the derandomized version of) the randomized iterate. Finally, we sample a random $f_u \xleftarrow{\$} \mathcal{F}$ and plug it into Theorem 5.1 to get the UOWHFs as desired. We refer to Construction 5.1 for more details.

Definition 5.2 (the randomized iterate [11, 8]) *Let $n \in \mathbb{N}$, function $f : \{0,1\}^n \to \{0,1\}^n$, and let $\mathcal{H}$ be a family of pairwise-independent length-preserving hash functions over $\{0,1\}^n$. For $k \in \mathbb{N}$, $x_1 \in \{0,1\}^n$ and vector $\vec{h}^k = (h_1, \ldots, h_k) \in \mathcal{H}^k$, recursively define the i^{th} randomized iterate by:*

$$x_1 \xrightarrow{f} y_1 \xrightarrow{h_1} x_2 \xrightarrow{f} y_2 \xrightarrow{h_2} \cdots \quad x_k \xrightarrow{f} y_k \xrightarrow{h_k}$$

$$y_i = f(x_i),\ x_{i+1} = h_i(y_i)\ .$$

We denote the i^{th} iterate by function f^i, i.e., $y_i = f^i(x_1, \vec{h}^k)$, where $\vec{h}^k$ is possibly redundant as for $i \le k+1$ y_i only depends on $\vec{h}^{i-1}$.
*The **randomized version** refers to the case where $x_1 \xleftarrow{\$} \{0,1\}^n$ and $\vec{h}^k \xleftarrow{\$} \mathcal{H}^k$.*
*The **derandomized version** refers to that $x_1 \xleftarrow{\$} \{0,1\}^n$, $u \xleftarrow{\$} \{0,1\}^{q \in O(n\cdot\log n)}$, $\vec{h}^k := BSG(u)$, where $BSG : \{0,1\}^q \to \{0,1\}^{k\cdot\log|\mathcal{H}|}$ is a bounded-space generator that 2^{-2n}-fools every $(2n+1, k, \log|\mathcal{H}|)$-LBP (see Definition B.1), and $\log|\mathcal{H}|$ is the description length of $\mathcal{H}$ (e.g., $2n$ bits for concreteness).*

We refer to Definition B.1 for the definition of bounded-width layered branching program (LBP). We note that the aforementioned bounded-space generators exist due to Theorem B.1 (see Appendix B) by setting $s(n) = 2n+1$, $v(n) = \log|\mathcal{H}| = \Theta(n)$, $k(n) = \mathsf{poly}(n)$ and $\varepsilon(n) = 2^{-2n}$ and thus $q(n)=O(n\cdot\log n)$.

Remark 5.1 (on what is proven in [27].) *In [27], the authors introduced weakly unknown-regular one-way functions and then constructed a pseudorandom generator with seed length $O(n\cdot\log n)$ based on the randomized iterate. They showed that "every $k = n^{2c}\cdot\log n\cdot\omega(1)$ iterations are hard-to-invert", i.e., for any j it is hard to predict x_j given $y_{j+k} = f^{j+k}(x_1, BSG(u))$ and u. A PRG thus follows by outputting $\log n$ hardcore bits for every k iterations. In this paper, we first adapt their findings to show that $f_u(\cdot) = f^k(\cdot, BSG(u))$ constitutes a family of one-way functions, i.e., given $y_k = f_u(x_1)$ and u it is infeasible to find any x_1' such that $y_k = f^k(x_1', BSG(u))$. This is stated as Lemma 5.2 with proof given in Appendix B. However, it is still insufficient to construct UOWHFs with the one-way-ness of f_u. We further show in Lemma 5.3 that a random $f_u \xleftarrow{\$} \mathcal{F}$ is almost regular (in a slightly weaker sense than Definition 2.5 but already suffices for our needs).*

Construction 5.1 (an explicit construction) *For constant c and $k = n^{2c}\cdot\log n\cdot\omega(1)$, let f be as defined in Definition 2.6, let f^k and $BSG(\cdot)$ be as defined in Definition 5.2, define function family*

$$\mathcal{F} \stackrel{def}{=} \{\ f_u : \{0,1\}^n \to \{0,1\}^n, f_u(x)=f^k(x, BSG(u)), u \in \{0,1\}^{O(n\cdot\log n)}\ \} \tag{2}$$

and further define $\mathcal{G}$ as in Theorem 5.1, i.e.,

$$\mathcal{G} = \{g_{f_u} : \{0,1\}^{\kappa\cdot s} \to \{0,1\}^n\ ,\ g_{f_u}(z) = g^{\kappa}_{f_u}(z, x_1, Shoup(u')), x_1 \in \{0,1\}^n, u, u' \in \{0,1\}^{O(n\cdot\log n)}\ \}$$

where $g^{\kappa}_{f_u}$ and $Shoup(\cdot)$ are as defined in Definition 5.1, and $\kappa\cdot s \ge n+s$ (e.g., set $s = \log n$, $\kappa = \Omega(n/\log n)$).

WHY A FAMILY OF OWFS? Note that the UOWHF g_{f_u} operates on input z and enjoys output length $\Theta(n)$, and it is described by key $(x_1, u, u') \in \{0,1\}^{O(n\cdot\log n)}$. An alternative is to view f_u as a single one-way function rather than a family of OWFs, i.e., $\tilde{f}(x,u) \stackrel{\text{def}}{=} (f_u(x), u)$ and plug $\tilde{f}$ into the UOWHF construction as in Theorem 5.1. However, in this case, the output and key lengths of the UOWHF are $O(n\cdot\log n)$ and $O(n\cdot\log^2 n)$ respectively since $\tilde{f}$ now has input and output length $O(n\cdot\log n)$.

Definition 5.3 (events) *For any $n, j \le k \in \mathbb{N}$, define events*

$$\mathcal{E}_j \stackrel{\mathsf{def}}{=} \left((X_1, \vec{H}^k) \in \{\ (x_1, \vec{h}^k) : y_j = f^j(x_1, \vec{h}^k) \in \mathcal{Y}_{\max}\ \}\right)$$

$$\mathcal{E}'_j \stackrel{\mathsf{def}}{=} \left((X_1, U_q) \in \{\ (x_1, u) : y_j = f^j(x_1, BSG(u)) \in \mathcal{Y}_{\max}\ \}\right)$$

where $\mathcal{Y}_{\max} \stackrel{def}{=} \{y : 2^{\max-1} \le |f^{-1}(y)| < 2^{\max}\}$, $(X_1, \vec{H}^k)$ and (X_1, U_q) are uniform over $\{0,1\}^n \times \mathcal{H}^k$ and $\{0,1\}^n \times \{0,1\}^q$ respectively.
Note that by definition $\mathcal{Y}_{\max} = f(\mathcal{X}_{\max})$ (see Definition 2.6) and thus $\Pr[f(U_n) \in \mathcal{Y}_{\max}] \ge n^{-c}$.

We will use the following inequalities from [27] and reproduce their proofs in Appendix B. It is not hard to see that (3), (5) and (7) hold for the randomized version. For example, we have by the pairwise independence of $\mathcal{H}$ that all $x_1, \ldots, x_k$ are i.i.d. to U_n so that (5) immediately follows and (7) follows by a Chernoff bound. Then, for every inequality (4), (6) and (8), we define an LBP (see Definition B.1) and argue that the advantage of the LBP on $\vec{H}^k$ and $BSG(U_q)$ is bounded by 2^{-2n} and thus (4), (6) and (8) follow from their respective counterparts (3), (5) and (7) by adding an additive term 2^{-2n}.

Lemma 5.1 (some inequalities from [27]) *For any $n, k \in \mathbb{N}$, it holds that*

$$\mathsf{CP}(\ Y_k \mid \vec{H}^k) \ \leq \ k{\cdot}2^{\max-n+1}\ , \tag{3}$$

$$\mathsf{CP}(\ Y'_k \mid U_q) \ \leq \ k{\cdot}2^{\max-n+1} + 2^{-2n}, \tag{4}$$

$$\forall j \in [k]: \ \Pr[\mathcal{E}_j] \ \geq \ n^{-c}\ , \tag{5}$$

$$\forall j \in [k]: \ \Pr[\mathcal{E}'_j] \ \geq \ n^{-c} - 2^{-2n}\ , \tag{6}$$

$$\Pr[\mathcal{E}_1 \vee \mathcal{E}_2 \vee \ldots \vee \mathcal{E}_k] \ \geq \ 1 - \exp^{-k/n^{2c}}\ , \tag{7}$$

$$\Pr[\mathcal{E}'_1 \vee \mathcal{E}'_2 \vee \ldots \vee \mathcal{E}'_k] \ \geq \ 1 - \exp^{-k/n^{2c}} - 2^{-2n}\ , \tag{8}$$

where $Y_k = f^k(X_1, \vec{H}^k)$ and $Y'_k = f^k(X_1, BSG(U_q))$.

Lemma 5.2 ($\mathcal{F}$ is one-way) *Assume that f is a (t,ε)-OWF that is weakly unknown-regular on an n^{-c} fraction of domain, let $\mathcal{F} = \{f_u\}$ be as defined in (2). Then, for any PPT A of running time $t - n^{O(1)}$ it holds that*

$$\Pr_{u \xleftarrow{\$} \{0,1\}^q,\ x \xleftarrow{\$} \{0,1\}^n} [\ \mathsf{A}(u, f_u(x)) \in f_u^{-1}(\ f_u(x)\ )\ ] \ \leq \ \sqrt{2^9{\cdot}k^4 \cdot n^{3c} \cdot \varepsilon}\ + \exp^{-k/n^{2c}}\ +\ 2^{-2n}\ . \tag{9}$$

where $q \in \Theta(n{\cdot}\log n)$, $u \in \{0,1\}^q$ and $f_u(x){=}f^k(x, BSG(u))$ as defined in Definition 5.2.

Although non-trivial, the above lemma is mainly attributed to and adapted from a related statement in [27] (see Remark 5.1). We refer the readers to Appendix B for its adapted proof.

5.4 $\mathcal{F}$ is Almost-Regular

Lemma 5.3 ($\mathcal{F}$ is almost-regular) *For $n, k \in \mathbb{N}$, let c and f be as defined in Definition 2.6, let $\mathcal{H}$,f^k and $BSG : \{0,1\}^{q \in O(n \cdot \log n)} \to \{0,1\}^{k \cdot \log |\mathcal{H}|}$ be as defined in Definition 5.2. Then, for any $a \geq 0$ it holds that*

$$\Pr_{u \xleftarrow{\$} \{0,1\}^q,\ x \xleftarrow{\$} \{0,1\}^n} [\ 2^{\max-a-1} \ \leq \ |f_u^{-1}\ (\ f_u(x)\ )| \ \leq \ 2^{\max+a+1}\ ] \ \geq \ 1 \ - \ k{\cdot}2^{-a+2} \ - \ \exp^{-k/n^{2c}} \ - \ 2^{-2n}, \tag{10}$$

where $u \in \{0,1\}^q$ and $f_u(x){=}f^k(x, BSG(u))$.

Proof. We define $\mathcal{S}_{low} \stackrel{\text{def}}{=} \Big((X_1, U_q) \in \{(x,u) : 0 < |f_u^{-1}(f_u(x))| < 2^{\max-a-1}\}\Big)$ and $\mathcal{S}_{up} \stackrel{\text{def}}{=} \Big((X_1, U_q) \in \{(x,u) : |f_u^{-1}(f_u(x))| > 2^{\max+a+1}\}\Big)$, where X_1 is uniform over $\{0,1\}^n$. Clearly, the left-hand of (10) is

lower bounded by $1-\Pr[\mathcal{S}_{low}]-\Pr[\mathcal{S}_{up}]$ and thus it suffices to upper bound both $\Pr[\mathcal{S}_{low}]$ and $\Pr[\mathcal{S}_{up}]$. We first have

$$\begin{aligned}
\Pr[\mathcal{S}_{low}] &= \Pr[\mathcal{S}_{low}\wedge(\mathcal{E}'_1\vee\mathcal{E}'_2\vee\ldots\vee\mathcal{E}'_k)]+\Pr[\mathcal{S}_{low}\wedge\neg(\mathcal{E}'_1\vee\mathcal{E}'_2\vee\ldots\vee\mathcal{E}'_k)]\\
&\leq \Pr[\bigvee_{j=1}^{k}(\mathcal{S}_{low}\wedge\mathcal{E}'_j)] \;+\; \Pr[\neg(\mathcal{E}'_1\vee\mathcal{E}'_2\vee\ldots\vee\mathcal{E}'_k)]\\
&\leq \sum_{j=1}^{k}\Pr[\mathcal{S}_{low}\wedge\mathcal{E}'_j] \;+\; (\exp^{-k/n^{2c}} \;+\; 2^{-2n})\\
&\leq k{\cdot}2^{-a} \;+\; \exp^{-k/n^{2c}} \;+\; 2^{-2n}
\end{aligned}$$

where the first inequality is trivial, the second is by the union bound and (8), and the third is due to that for every $j\in[k]$ with shorthand $f_{u,j}(x)\stackrel{\text{def}}{=}f^j(x,BSG(u))$ it holds that

$$\begin{aligned}
\Pr[\mathcal{S}_{low}\wedge\mathcal{E}'_j] &= \sum_u \Pr[U_q=u]\cdot \sum_{x:\ f_{u,j}(x)\in\mathcal{Y}_{\max}\ \wedge\ 0<|f_u^{-1}(f_u(x))|<2^{\max-a-1}} \Pr[X_1=x|U_q=u]\\
&\leq \sum_u \Pr[U_q=u]\cdot \sum_{x:\ f_{u,j}(x)\in\mathcal{Y}_{\max}\ \wedge\ 0<|f_{u,j}^{-1}(f_{u,j}(x))|<2^{\max-a-1}} \Pr[X_1=x\mid U_q=u]\\
&\leq \sum_u \Pr[U_q=u]\cdot|\mathcal{Y}_{\max}|{\cdot}2^{\max-a-1}{\cdot}2^{-n}\\
&\leq 2^{n+1-\max}\cdot 2^{-n+\max-a-1} \;=\; 2^{-a}
\end{aligned}$$

where the first inequality is due to Fact 5.1 (setting $f_1{=}f_{u,j}$, $f_2=f{\circ}h_{k-1}{\circ}\ldots{\circ}f\circ h_j$ and thus $\bar{f}=f_u$), the second follows from the fact that there are $|\mathcal{Y}_{\max}|$ possible values for $f_{u,j}(x)\in\mathcal{Y}_{\max}$ and every $f_{u,j}(x)$ has less than $2^{\max-a-1}$ preimages (by definition of $\mathcal{S}_{low}$), and the third is due to $|\mathcal{Y}_{\max}|{\leq}2^{n+1-\max}$. Next we proceed to bounding the second term, i.e., $\Pr[\mathcal{S}_{up}]\;\leq\;k{\cdot}2^{-a+1}$. We have:

$$\begin{aligned}
k{\cdot}2^{\max-n+1}+2^{-2n} &\geq \mathsf{CP}(\,Y'_k\mid U_q) \;=\mathbb{E}_{u\leftarrow U_q}[\,\sum_y \Pr[\,f_u(X_1)=y\mid U_q=u]^2\,]\\
&> 2^{\max+a-n+1}\cdot\mathbb{E}_{u\leftarrow U_q}[\sum_{y:\ |f_u^{-1}(y)|>2^{\max+a+1}\}}\Pr[\,f_u(X_1)=y\mid U_q=u]\,]\\
&= 2^{\max+a-n+1}\cdot\Pr[\mathcal{S}_{up}]\;,
\end{aligned}$$

where the first inequality is by (4), and the second is due to that for any (y,u) satisfying $|f_u^{-1}(y)|>2^{\max+a+1}$ it holds that $\Pr[\,f_u(X_1)=y\mid U_q=u]>2^{\max+a-n+1}$. It follows that $\Pr[\mathcal{S}_{up}]\leq(k{\cdot}2^{\max-n+1}+2^{-2n})/2^{\max+a-n+1}{\leq}2^{-a+1}$ and hence completes the proof. □

Fact 5.1 *Let $f_1:\mathcal{X}\to\mathcal{Y}$ and $f_2:\mathcal{Y}\to\mathcal{Z}$ be any functions, and let X be any random variable over $\mathcal{X}$. Then, for any integer $t>0$ and any set $\mathcal{X}_a\subseteq\mathcal{X}$ it holds that*

$$\sum_{x:\ x\in\mathcal{X}_a\ \wedge\ 0<|\bar{f}^{-1}(\bar{f}(x))|<t}\Pr[X=x]\;\leq\sum_{x:\ x\in\mathcal{X}_a\ \wedge\ 0<|f_1^{-1}(f_1(x))|<t}\Pr[X=x]$$

where $\bar{f}=f_2{\circ}f_1$

Proof. We use shorthands $\mathcal{X}_1\stackrel{\text{def}}{=}\{x:0<|\bar{f}^{-1}(\bar{f}(x))|<t\}$ and $\mathcal{X}_2\stackrel{\text{def}}{=}\{x:0<|f_1^{-1}(f_1(x))|<t\}$. It suffices to show that $\mathcal{X}_1\subseteq\mathcal{X}_2$. This is not hard to see since any x satisfying $0<|\bar{f}^{-1}(\bar{f}(x))|<t$ implies $0<|f_1^{-1}(f_1(x))|<t$. □

5.5 Putting Things Together

Proof sketch of Theorem 5.2. Consider a (t,ε)-OWF f as defined in Definition 2.6. Although f is far from regular, iterating it (as defined in Construction 5.1) sufficiently many, say $k = n^{2c} \cdot \log n \cdot \omega(1)$, times yields a family of one-way functions $\mathcal{F}$ with description size $O(n \cdot \log n)$, as stated in Lemma 5.2. Furthermore, Lemma 5.3 states that, for $\alpha = 2^{\max - a - 1}$ and any $\beta = 2^{2a+2} \geq 4$, a random function $f_u \xleftarrow{\$} \mathcal{F}$ is $(\alpha,\ \alpha \cdot \beta)$-almost regular except for a $(O(k/\sqrt{\beta}) + \mathsf{negl}(n))$-fraction. Therefore, plug f_u into Theorem 5.1 and set $s = \log n$, $\beta = (1/\varepsilon^{O(1)})$ for some small enough constant $O(1)$ so that $\mathsf{poly}(\beta, 2^s, \kappa) \cdot \varepsilon^{\Theta(1)}$ remains negligible, we obtain a family of UOWHFs with output length $\Theta(n)$ and key length $O(n \cdot \log n)$. In total, it makes $\kappa = O(n/\log n)$ calls to f^k for $k = n^{2c} \cdot \log n \cdot \omega(1)$ and thus $O(n^{2c+1} \cdot \omega(1))$ calls to f. □

References

[1] Scott Ames, Rosario Gennaro, and Muthuramakrishnan Venkitasubramaniam. The generalized randomized iterate and its application to new efficient constructions of UOWHFs from regular one-way functions. In *ASIACRYPT*, pages 154–171, 2012.

[2] Kfir Barhum and Thomas Holenstein. A cookbook for black-box separations and a recipe for uowhfs. In *Proceedings of the 5th Theory of Cryptography Conference (TCC 2013)*, pages 662–679, 2013.

[3] Kfir Barhum and Ueli Maurer. UOWHFs from OWFs: Trading regularity for efficiency. In *LATINCRYPT*, pages 234–253, 2012.

[4] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003.

[5] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM Journal on Computing*, 35(1):217–246, 2005.

[6] Rosario Gennaro and Muthuramakrishnan Venkitasubramaniam. Can you compress where you expand? new and efficient hash functions. announced at the rump session of ASIACRYPT 2013, 2013. http://asiacrypt.2013.rump.cr.yp.to/.

[7] Oded Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.

[8] Oded Goldreich, Hugo Krawczyk, and Michael Luby. On the existence of pseudorandom generators. *SIAM Journal on Computing*, 22(6):1163–1175, 1993.

[9] Oded Goldreich, Leonid A. Levin, and Noam Nisan. On constructing 1-1 one-way functions. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 13–25. 2011.

[10] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.

[11] Iftach Haitner, Danny Harnik, and Omer Reingold. On the power of the randomized iterate. In *Proceedings of the 26th International Cryptology Conference (CRYPTO 2006)*, pages 22–40, 2006.

[12] Iftach Haitner, Thomas Holenstein, Omer Reingold, Salil P. Vadhan, and Hoeteck Wee. Universal one-way hash functions via inaccessible entropy. In *EUROCRYPT*, pages 616–637, 2010.

[13] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM Journal on Computing*, 39(3):1153–1218, 2009.

[14] Iftach Haitner, Omer Reingold, Salil P. Vadhan, and Hoeteck Wee. Inaccessible entropy. In *Proceedings of the 41st ACM Symposium on the Theory of Computing*, pages 611–620, 2009.

[15] Thomas Holenstein and Makrand Sinha. Constructing a pseudorandom generator requires an almost linear Number of calls. In *Proceedings of the 53rd IEEE Symposium on Foundation of Computer Science*, pages 698–707, 2012.

[16] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the 26th ACM Symposium on the Theory of Computing*, pages 356–364, 1994.

[17] Jonathan Katz and Chiu-Yuen Koo. On constructing universal one-way hash functions from arbitrary one-way functions. *IACR Cryptology ePrint Archive*, 2005. http://eprint.iacr.org/2005/328.

[18] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.

[19] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In D. S. Johnson, editor, *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 33–43, Seattle, Washington, 15–17 May 1989.

[20] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

[21] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 387–394, Baltimore, Maryland, 14–16 May 1990.

[22] John Rompel. *Techniques for computing with low-independence randomness.* PhD thesis, Massachusetts Institute of Technology, 1990. http://dspace.mit.edu/handle/1721.1/7582.

[23] Alfredo De Santis and Moti Yung. On the design of provably secure cryptographic hash functions. In *EUROCRYPT*, pages 412–431, 1990.

[24] Victor Shoup. A composition theorem for universal one-way hash functions. In Bart Preneel, editor, *Advances in Cryptology—EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 445–452. Springer-Verlag, 2000.

[25] D. R. Stinson. Universal hashing and authentication codes. *Designs, Codes, and Cryptography*, 4(4):369–380, 1994.

[26] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *Proceedings of the 23rd IEEE Symposium on Foundation of Computer Science*, pages 80–91, 1982.

[27] Yu Yu, Dawu Gu, and Xiangxue Li. The randomized iterate revisited - almost linear seed length PRGs from a broader class of one-way functions. *Electronic Colloquium on Computational Complexity (ECCC)*, TR14-82, 2014.

[28] Yu Yu, Xiangxue Li, and Jian Weng. Pseudorandom generators from regular one-way functions: New constructions with improved parameters. In *ASIACRYPT*, pages 261–279, 2013.

A Lemmata and Proofs Omitted

Lemma A.1 (The injective hash lemma for almost flat sources) *For any integers a, d, k and l satisfying $a \leq l$ and any real value $\beta \geq 1$, let Y be any random variable over a support set $\mathcal{Y} \subseteq \{0,1\}^l$ such that every $y \in \mathcal{Y}$ satisfies $2^{-a} \leq \Pr[Y=y] \leq 2^{-a}\beta$, and let $\mathcal{H} \stackrel{\mathsf{def}}{=} \{h : \{0,1\}^l \to \{0,1\}^{a+d}\}$ be a family of $(k \cdot 2^{-(a+d)})$-almost universal hash functions. Then, we have that*

$$\Pr_{y \leftarrow Y,\ h \xleftarrow{\$} \mathcal{H}}[\ \exists \tilde{y} \in \mathcal{Y} :\ \tilde{y} \neq y \wedge h(\tilde{y}) = h(y)\ ] \leq k \cdot \beta \cdot 2^{-d}\ .$$

Proof. It follows from $|\mathcal{Y}| \cdot 2^{-a} \leq \sum_{y \in \mathcal{Y}} \Pr[Y=y] = 1$ that $|\mathcal{Y}| \leq 2^a$. We denote by $U_{\mathcal{Y}}$ the uniform distribution over $\mathcal{Y}$. Note that $U_{\mathcal{Y}}$ and Y are of the same support set. We define

$$\mathcal{S}^* \stackrel{\mathsf{def}}{=} \{(y,h) \in (\mathcal{Y}, \mathcal{H}) : \exists \tilde{y} \in \mathcal{Y} \ \wedge\ \tilde{y} \neq y \ \wedge\ h(\tilde{y}) = h(y)\}$$

and thus by the standard injective hash lemma (i.e., Lemma 3.1) it holds that

$$\Pr[(U_{\mathcal{Y}}, H) \in \mathcal{S}^*] \ \leq\ k \cdot 2^{-(a+d-\log|\mathcal{Y}|)} \ \leq\ k \cdot 2^{-d}\ .$$

On the other hand, we have

$$\begin{aligned}
\Pr[(U_{\mathcal{Y}}, H) \in \mathcal{S}^*] &= \sum_h \Pr[H=h] \sum_{y:(y,h) \in \mathcal{S}^*} \frac{1}{|\mathcal{Y}|} \\
&\geq \sum_h \Pr[H=h] \sum_{y:(y,h) \in \mathcal{S}^*} 2^{-a} \\
&\geq \sum_h \Pr[H=h] \sum_{y:(y,h) \in \mathcal{S}^*} \Pr[Y=y]/\beta \ =\ \Pr[(Y,H) \in \mathcal{S}^*]/\beta\ .
\end{aligned}$$

It immediately follows that $\Pr[(Y,H) \in \mathcal{S}^*] \leq k \cdot \beta \cdot 2^{-d}$. □

Lemma A.2 (regularity-preserving OWF) *For any $r_1 = r_1(n) \leq r_2 = r_2(n)$, and any efficiently computable $\kappa = \kappa(n) \in O(n)$ and $\delta = \delta(n) \in O(n)$, let $f : \{0,1\}^n \to \{0,1\}^l$ be any $(2^{r_1}, 2^{r_2})$-almost regular (t,ε)-one-way function, let $\mathcal{H} = \{h : \{0,1\}^l \to \{0,1\}^{n+\delta+\kappa}\}$ be a family of $(\mathsf{poly}(n) \cdot 2^{-(n+\delta+\kappa)})$-almost universal hash functions with description length[10] $O(n)$, define function $f' : \{0,1\}^n \times \mathcal{H} \to \{0,1\}^{n+\delta+\kappa} \times \mathcal{H}$ as*

$$f'(x,h) = (h(f(x)), h)\quad . \tag{11}$$

Then, we have

1. Regularity-preserving. *f' is $(2^{r_1}, 2^{r_2})$-regular except on a $\mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa+\delta-r_2)}$-fraction of inputs, i.e.,*

$$\Pr_{x \xleftarrow{\$} \{0,1\}^n,\ h \xleftarrow{\$} \mathcal{H}}[\ 2^{r_1} \ \leq\ |f'^{-1}(f'(x,h))| \leq 2^{r_2}\ ] \ \geq\ 1 - \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa+\delta-r_2)}\quad .$$

2. Hardness-preserving. *f' is a $(t - n^{O(1)},\ \varepsilon + \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa+\delta-r_2)})$-one-way function.*

[10] Such efficient $\mathcal{H}$ exists for any efficiently computable $l = l(n) \in \mathsf{poly}(n)$ and $\kappa = \kappa(n), \delta = \delta(n) \in O(n)$ by Fact 2.1.

Proof of Lemma A.2. As for every $y = f(x)$ we have $2^{r_1} \leq |f^{-1}(y)| \leq 2^{r_2}$ it suffices to show that the fraction of y's (drawn from $Y = f(U_n)$) on which h is 1-to-1 is overwhelming, i.e.,

$$\begin{aligned}
&\Pr_{x \xleftarrow{\$} \{0,1\}^n,\ h \xleftarrow{\$} \mathcal{H}} [\ 2^{r_1} \ \leq\ |f'^{-1}(f'(x,h))| \ \leq\ 2^{r_2}\ ] \\
&\geq \Pr_{y \leftarrow f(U_n),\ h \xleftarrow{\$} \mathcal{H}} [\ \neg \exists \tilde{y} \in f(\{0,1\}^n) : h(\tilde{y}) = h(y) \ \wedge\ y \neq \tilde{y}\ ] \\
&= 1 \ - \Pr_{y \leftarrow f(U_n),\ h \xleftarrow{\$} \mathcal{H}} [\ \exists \tilde{y} \in f(\{0,1\}^n) : h(\tilde{y}) = h(y) \ \wedge\ y \neq \tilde{y}\ ] \\
&\geq 1 - \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa+\delta-r_2)}\ ,
\end{aligned}$$

where the second inequality is by Lemma A.1 (setting $Y = f(U_n)$, $a = n - r_1$, $\beta = 2^{r_2 - r_1}$, $d = r_1 + \delta + \kappa$, $k = \mathsf{poly}(n)$). Further, it is not hard to see that any f'-inverting algorithm $\mathsf{A}_{f'}$ implies an f-inverting algorithm A_f. That is, on input $f(x)$, A_f applies random h to $f(x)$, and then invokes $\mathsf{A}_{f'}$ on $(h(f(x)), h)$ to recover x. The inversion probability of A_f is

$$\begin{aligned}
&\Pr_{y \leftarrow f(U_n)}[\mathsf{A}_f(y) \in f^{-1}(y)] \\
\geq\ &\Pr_{y \leftarrow f(U_n), h \xleftarrow{\$} \mathcal{H}} [\mathsf{A}_{f'}(h(y), h) \in f'^{-1}(h(y), h) \ \wedge \neg \left(\exists \tilde{y} \in f(\{0,1\}^n) : h(\tilde{y}) = h(y) \ \wedge\ y \neq \tilde{y} \right)\ ] \\
=\ &1 \ - \Pr_{y \leftarrow f(U_n), h \xleftarrow{\$} \mathcal{H}} [\mathsf{A}_{f'}(h(y), h) \notin f'^{-1}(h(y), h) \ \vee \left(\exists \tilde{y} \in f(\{0,1\}^n) : h(\tilde{y}) = h(y) \ \wedge\ y \neq \tilde{y} \right)\ ] \\
\geq\ &1 \ - \Pr_{y \leftarrow f(U_n), h \xleftarrow{\$} \mathcal{H}} [\mathsf{A}_{f'}(h(y), h) \notin f'^{-1}(h(y), h)] \ - \Pr_{y \leftarrow f(U_n), h \xleftarrow{\$} \mathcal{H}} [\exists \tilde{y} \in f(\{0,1\}^n) : h(\tilde{y}) = h(y) \ \wedge\ y \neq \tilde{y}\ ] \\
\geq\ &\Pr_{y \leftarrow f(U_n), h \xleftarrow{\$} \mathcal{H}} [\mathsf{A}_{f'}(h(y), h) \in f'^{-1}(h(y), h)] \ - \ \mathsf{poly}(n) \cdot 2^{-(2r_1+\kappa+\delta-r_2)}\ ,
\end{aligned}$$

where the first inequality refers to that A_f inverts f if $\mathsf{A}_{f'}$ inverts f' on those $(h(y), h)$ for which there exists no distinct $\tilde{y}$ satisfying $h(\tilde{y}) = h(y)$, the second inequality is the union bound, and the third is due to the probability that h is not injective on Y as given above. This completes the proof. □

Proof of Fact 3.1. The first statement immediately follows from Lemma A.2 by setting $r_1 = r_2 = \delta = 0$. As for the second statement, let f' be as defined in (11) from Lemma A.2 with $\delta = n$, we further define a padded function $\bar{f} : \{0,1\}^{2n+\kappa} \times \mathcal{H} \to \{0,1\}^{2n+\kappa} \times \mathcal{H}$ as

$$\bar{f}(x, \mathsf{dummy}, h) \stackrel{\text{def}}{=} f'(x, h)\ ,$$

where $x \in \{0,1\}^n$, $\mathsf{dummy} \in \{0,1\}^{n+\kappa}$, and $h \in \mathcal{H}$ (which is of size $O(n)$). Note that for every $(h(f(x)), h)$ the preimage-size of $\bar{f}$ is multiplied by a factor of $2^{n+\kappa}$ than that of f' due to the $(n+\kappa)$-bit padding dummy. Finally, with $\delta = n$ the negligible term $\mathsf{poly} \cdot 2^{-(2r_1+\kappa+\delta-r_2)}$ is further bounded by $\mathsf{poly} \cdot 2^{-(2r_1+\kappa)}$. This concludes the second statement. □

Fact A.1 *For any $(2^r, 2^r n^c)$-almost regular (length-preserving) (t, ε)-one-way function $f : \{0,1\}^n \to \{0,1\}^n$ and any inverter Inv of running t, it holds that*

$$\Pr_{y^* \xleftarrow{\$} \{0,1\}^l} [\ f(\mathsf{Inv}(y^*)) = y^*\ ] \ \leq\ 2^{-r} \cdot \varepsilon\ .$$

Proof.

$$\begin{aligned}
&\Pr_{y^* \xleftarrow{\$} \{0,1\}^n}[\ f(\mathsf{Inv}(y^*)) = y^*\ ] \\
&= \sum_{y \in f(\{0,1\}^n)} \Pr_{y^* \xleftarrow{\$} \{0,1\}^n}[y = y^*] \cdot \Pr[f(\mathsf{Inv}(y)) = y] \\
&\leq 2^{-r} \cdot \sum_{y \in f(\{0,1\}^n)} \Pr[f(U_n) = y] \cdot \Pr[f(\mathsf{Inv}(y)) = y] \\
&\leq 2^{-r} \cdot \varepsilon
\end{aligned}$$

where the first inequality is due to $\Pr[f(U_n) = y] \geq 2^{r-n}$ and the second is due to f's one-way-ness. □

Proof of Lemma 4.1. Suppose for contradiction that there exists an efficient $\tilde{\mathsf{A}}$ of running time t' such that

$$\Pr_{x \xleftarrow{\$} \{0,1\}^n, h \xleftarrow{\$} \mathcal{H}}[\ x' \leftarrow \tilde{\mathsf{A}}(x,h) :\ f(x) \neq f(x') \wedge h(f(x))_{[u]} = h(f(x'))_{[u]}\ ] \ =\ \varepsilon' \ >\ n^c \cdot 2^{s'} \cdot \varepsilon$$

where $u = n - r - c\log n - s'$. We proceed to the definition of algorithm $\mathsf{Inv}^{\tilde{\mathsf{A}}}$ (that inverts f by invoking $\tilde{\mathsf{A}}$) as in Algorithm 3. By Claim 3.1, conditioned on $f(x) \neq y^*$ it is equivalent to consider that $\mathsf{Inv}^{\tilde{\mathsf{A}}}$

Algorithm 3 $\mathsf{Inv}^{\tilde{\mathsf{A}}}$ that inverts f on input y^* using random coins (x, v).

Input: $y^* \xleftarrow{\$} \{0,1\}^n$

Sample $x \xleftarrow{\$} \{0,1\}^n$
if $f(x) = y^*$ **then**
  Output x and **terminate**.
end if
sample $h := (f(x) - y^*)^{-1} \cdot v$, where $v \xleftarrow{\$} \mathcal{V} = \{v \in \{0,1\}^n \setminus \{\vec{0}\} : v_{[u]} = \overbrace{0 \ldots 0}^{u}\}$
{*note*: The above implies $h \xleftarrow{\$} \{h \in \mathcal{H} :\ h(f(x))_{[u]} = h(y^*)_{[u]}\}$ by the algebraic structure of h. }
$x' \leftarrow \tilde{\mathsf{A}}(x,h)$
if $f(x')=y^*$ **then**
  Output x'
else
  Output $\perp$
end if
Terminate

samples (x,h,v) from $\{0,1\}^n \times \mathcal{H} \times \mathcal{V}$ uniformly and independently, from which y^* can be determined.

Then, we argue that $\mathsf{Inv}^{\tilde{\mathsf{A}}}$ inverts f with the following probability (see the rationale below)

$$\begin{aligned}
&\Pr_{y^*\xleftarrow{\$}\{0,1\}^n,\ x\xleftarrow{\$}\{0,1\}^n,\ v\xleftarrow{\$}\mathcal{V}}[\ f(\mathsf{Inv}^{\tilde{\mathsf{A}}}(y^*))=y^*\ ]\\
\geq\ &\Pr_{x\xleftarrow{\$}\{0,1\}^n,y^*\xleftarrow{\$}\{0,1\}^n}[\ f(x)=y^*\ ]\ +\ \Pr_{x\xleftarrow{\$}\{0,1\}^n,y^*\xleftarrow{\$}\{0,1\}^n}[\ f(x)\neq y^*\ ]\\
&\times\ \Pr_{x\xleftarrow{\$}\{0,1\}^n,h\xleftarrow{\$}\mathcal{H},x\neq(x'\leftarrow\tilde{\mathsf{A}}(x,h)),v\xleftarrow{\$}\mathcal{V}}[\ h(f(x))_{[u]}=h(f(x'))_{[u]}\ \wedge\ y^*=f(x')\mid f(x)\neq y^*\ ]\\
\geq\ &2^{-n}\ +\ (1-2^{-n})\cdot\varepsilon'\cdot\ \Pr_{v\xleftarrow{\$}\mathcal{V}}[\ y^*=f(x')\mid f(x)\neq y^*\ \wedge\ f(x)\neq f(x')\ \wedge\ h(f(x))_{[u]}=h(f(x'))_{[u]}\ ]\\
=\ &2^{-n}\ +\ (1-2^{-n})\cdot\varepsilon'\cdot\frac{1}{|\mathcal{V}|}\\
=\ &2^{-n}\ +\ (1-2^{-n})\cdot\varepsilon'\cdot\frac{1}{2^{r+c\log n+s'}-1}\ >\ n^c\cdot 2^{s'}\cdot\varepsilon\cdot 2^{-(r+c\log n+s')}\ =\ \varepsilon\cdot 2^{-r}\quad,
\end{aligned}$$

where $\tilde{\mathsf{A}}$ takes only x and h as input (i.e., independent of v), and thus conditioned on that $\tilde{\mathsf{A}}$ produces a valid $f(x')\neq f(x)$ satisfying $h(f(x'))_{[u]}=h(f(x))_{[u]}$, we have by Claim 3.1 that string y^* is uniformly distributed over set $\mathcal{Y}^*\overset{\mathsf{def}}{=}\{y^*: y^*=f(x)-v\cdot h^{-1}, v\in\mathcal{V}\}$. Note that the already fixed $f(x')$ is also an element of $\mathcal{Y}^*$ and thus y^* hits $f(x')$ with probability $1/|\mathcal{Y}^*|=1/|\mathcal{V}|=1/(2^{n-u}-1)$. We thus complete the proof by reaching a contradiction to Fact A.1. □

Proof of Corollary 3.1. Denote by $\mathcal{E}$ the event that f is 1-to-1 on uniformly random x, i.e., $\mathcal{E}\overset{\text{def}}{=}\ X\in\{x:|f^{-1}(f(x))|=1\}$. It follows from Theorem 3.1 that conditioned on $\mathcal{E}$ (with probability $\Pr[\mathcal{E}]\geq 1-\delta$) that $\mathcal{G}_1$ is a family of $(t-n^{O(1)},\ 2^{s'}\cdot\varepsilon)$-universal one-way hash functions with actual shrinkage $s'=\log(2^n\cdot\Pr[\mathcal{E}])-(n-s)=s-\log\Pr[\mathcal{E}]\leq s+1$. Therefore, overall the resulting $\mathcal{G}$ is a family of $(t-n^{O(1)},\ 2^{s+1}\cdot\varepsilon+\delta)$-universal one-way hash functions with shrinkage s. □

Proof of Lemma 4.2.

$$\begin{aligned}
&\Pr_{x\xleftarrow{\$}\{0,1\}^n,\ h_1\xleftarrow{\$}\mathcal{H}_1}[\ \exists x':x'\neq x\ \wedge\ f'(x,h_1)=f'(x',h_1)]\\
=\ &\mathbb{E}_{y\leftarrow f(U_n)}\left[\ \Pr_{x\xleftarrow{\$}f^{-1}(y),\ h_1\xleftarrow{\$}\mathcal{H}_1}[\exists x'\in f^{-1}(y):x'\neq x\ \wedge\ h_1(x)=h_1(x')]\ \right]\\
\leq\ &\mathbb{E}_{y\leftarrow f(U_n)}[\ 2^{-d}\ ]\ =\ 2^{-d}\quad,
\end{aligned}$$

where the inequality is due to that conditioned on any $y=f(X)$ random variable X is a flat distribution on a set of size at most $2^r\cdot n^c$, so we apply Lemma 3.1 (setting $\mathcal{Y}=f^{-1}(y)$, $a\leq r+c\cdot\log n$, and $k=1$). □

Proof of Lemma 4.3. Suppose that there exists some inverter A' for f' of running time $t-O(n)$ such that

$$\Pr[\ \mathsf{A}'(f'(X,H_1))\in\ f'^{-1}\ (\ f'(X,H_1)\ )\ ]\ >\ \sqrt{2^{d+3}\cdot n^c\cdot\varepsilon}\ .$$

First we consider the collision probability of $H_1(X)$ given $f(X)$ and H_1, i.e.,

$$\begin{aligned}
\mathsf{CP}(\ H_1(X)\mid f(X),H_1\ )\ &=\ \mathbb{E}_{y\leftarrow f(U_n)}[\ \mathsf{CP}(\ H_1(X)\mid f(X)=y,\ H_1\ )\ ]\\
&\leq\ \mathbb{E}_{y\leftarrow f(U_n)}[\ \mathsf{CP}(\ X\mid f(X)=y\ )\ +\ \max_{x_1\neq x_2,f(x_1)=f(x_2)}\{\ \Pr[H_1(x_1)=H_1(x_2)]\ \}\ ]\\
&\leq\ \mathbb{E}_{y\leftarrow f(U_n)}[\ 2^{-r}\ +\ 2^{-r-c\log n-d}\ ]\\
&\leq\ 2^{-r+1}
\end{aligned}$$

where we note that the argument is similar to the proof of the injective hash lemma. To apply Lemma B.1, let $\mathcal{W} = \{0,1\}^{m=r+c\log n+d}$, $\mathcal{Z} = f(\{0,1\}^n) \times \mathcal{H}_1$ and thus $e = d + c \cdot \log n + 1$, and define Adv as the success probability of A' on the corresponding input, i.e.,

$$\mathsf{Adv}(w, z = (y, h_1)) \stackrel{\text{def}}{=} \Pr[\ \mathsf{A}'(y, w, h_1) \in f'^{-1}(y, w, h_1)\ ]\ ,$$

where the probability is taken over the internal coins of A'. Thus,

$$\begin{aligned} & \Pr[\ \mathsf{A}'(f(X), U_{r+c\log n+d}, H_1) \in f'^{-1}(f(X), U_{r+c\log n+d}, H_1)\ ] \\ = \ & \mathbb{E}[\mathsf{Adv}(U_{\mathcal{W}}, Z)] \ \geq \ \mathbb{E}[\mathsf{Adv}(W, Z)]^2/2^{e+2} \\ = \ & \Pr[\ \mathsf{A}'(f(X), H_1(X), H_1) \in f'^{-1}(f(X), H_1(X), H_1)\ ]^2/2^{d+c\cdot\log n+3} \\ > \ & (\sqrt{2^{d+3} \cdot n^c \cdot \varepsilon})^2/2^{d+c\cdot\log n+3} \ = \ \varepsilon \end{aligned}$$

where the first inequality is due to Lemma B.1 and the second is by the assumption. This immediately implies another inverter A for f that on input y, it samples $h_1 \stackrel{\$}{\leftarrow} \mathcal{H}_1$, $w \leftarrow U_{r+c\log n+d}$, invokes $(x', h_1') \leftarrow \mathsf{A}'(y, w, h_1)$ and produces x' as output. In particular, A inverts f with the following probability

$$\Pr[\mathsf{A}(f(X)) \in f^{-1}(f(X))] \ \geq \ \Pr[\ \mathsf{A}'(f(X), U_{r+c\log n+d}, H_1) \in f'^{-1}(f(X), U_{r+c\log n+d}, H_1)\ ] \ > \ \varepsilon$$

which is a contradiction to the one-way-ness of f and thus completes the proof. □

B Proofs Reproduced and Adapted from [27]

A technical Lemma. To prove Lemma 4.3 and Lemma 5.2, we will need the following lemma that was implicit in [11] (and folklore in leakage-resilient cryptography) and was abstracted out in [27]. Informally, it states that "if any algorithm wins a one-sided game (e.g., inverting a OWF) on uniformly sampled challenges only with some negligible probability, then it cannot do much better (beyond a negligible advantage) in case that the challenges are sampled from any distribution of logarithmic Rényi entropy deficiency".

Lemma B.1 (one-sided game on imperfect randomness [27]) *For any $e \leq m \in \mathbb{N}$, let $\mathcal{W}$ and $\mathcal{Z}$ be any sets with $|\mathcal{W}| = 2^m$, let $\mathsf{Adv} : \mathcal{W} \times \mathcal{Z} \to [0,1]$ be any (deterministic) real-valued function, let (W, Z) be any random variable over set $\mathcal{W} \times \mathcal{Z}$ with $\mathsf{CP}(W|Z) \leq 2^{e-m}$, we have*

$$\mathbb{E}[\mathsf{Adv}(W, Z)] \ \leq \ \sqrt{2^{e+2} \cdot \mathbb{E}[\mathsf{Adv}(U_{\mathcal{W}}, Z)]} \tag{12}$$

where $U_{\mathcal{W}}$ denotes uniform distribution over $\mathcal{W}$ and independent of Z.

Proof of Lemma B.1. For any given δ define $\mathcal{S}_\delta \stackrel{\text{def}}{=} \{(w,z) : \Pr[W = w|Z = z] \geq 2^{e-m}/\delta\}$

$$\begin{aligned} 2^{e-m} \ & \geq \sum_z \Pr[Z = z] \sum_w \Pr[W = w|Z = z]^2 \\ & \geq \ \sum_z \Pr[Z = z] \sum_{w:(w,z)\in\mathcal{S}_\delta} \Pr[W = w|Z = z] \cdot 2^{-(m-e)}/\delta \\ & \geq \ (2^{e-m}/\delta) \cdot \Pr[(W, Z) \in \mathcal{S}_\delta]\ , \end{aligned}$$

and thus $\Pr[(W, Z) \in \mathcal{S}_\delta] \leq \delta$. It follows that

$$\begin{aligned} \mathbb{E}[\mathsf{Adv}(W, Z)] \ & = \ \sum_{(w,z)\in\mathcal{S}_\delta} \Pr[(W, Z) = (w, z)] \cdot \mathsf{Adv}(w, z) \ + \ \sum_{(w,z)\notin\mathcal{S}_\delta} \Pr[Z = z] \cdot \Pr[W = w|Z = z] \cdot \mathsf{Adv}(w, z) \\ & \leq \ \sum_{(w,z)\in\mathcal{S}_\delta} \Pr[(W, Z) = (w, z)] \ + \ (2^e/\delta) \cdot \sum_{(w,z)\notin\mathcal{S}_\delta} \Pr[Z = z] \cdot 2^{-m} \cdot \mathsf{Adv}(w, z) \\ & \leq \ \delta \ + \ (2^e/\delta) \cdot \mathbb{E}[\mathsf{Adv}(U_{\mathcal{W}}, Z)]\ , \end{aligned}$$

and we complete the proof by setting $\delta = \sqrt{2^e \cdot \mathbb{E}[\mathsf{Adv}(U_{\mathcal{W}}, Z)]}$. $\square$

Proof of (3),(5) and (7). We have that x_1, $x_2 = h_1(y_1)$, …, $x_k = h_{k-1}(y_{k-1})$ are all i.i.d. to U_n due to the universality of $\mathcal{H}$, which that $\mathcal{E}_1$, … and $\mathcal{E}_k$ are i.i.d. events with probability at least n^{-c}. For every $j \in [k]$, define $\zeta_j = 1$ iff $\mathcal{E}_j$ occurs (and $\zeta_j = 0$ otherwise). It follows by a Chernoff-Hoeffding bound that

$$\Pr[\ (\neg\mathcal{E}_1) \wedge \ldots \wedge (\neg\mathcal{E}_k)\ ] \ = \ \Pr[\ \sum_{j=1}^{k} \zeta_j = 0\ ] \ \leq \ \exp^{-k/n^{2c}}$$

which yields (7) by taking a negation. Finally, Regarding (3), consider two instances of the random iterate seeded with independent x_1 and x'_1 and a common random $\vec{h}^k$, the collision probability is upper bounded by the sum of events that the first collision occurs on points y_1, y_2, …, $y_k \in \mathcal{Y}_{[\max]}$ respectively. We thus have by the pairwise independence of $\mathcal{H}$ that

$$\begin{aligned}
&\mathsf{CP}(\ Y_k\ \mid \vec{H}^k) \\
&\leq \ \Pr_{x_1, x'_1 \xleftarrow{\$} \{0,1\}^n}[f(x_1) = f(x'_1)] \ + \ \sum_{j=2}^{k} \left(\Pr_{y_{j-1} \neq y'_{j-1},\ h_{j-1} \xleftarrow{\$} \mathcal{H}}[f(x_j) = f(x'_j)] \right) \\
&\leq \ k \cdot \mathsf{CP}(f(U_n)) \ \leq \ k \sum_{i=1}^{\max} \sum_{y \in \mathcal{Y}_i} \Pr[f(U_n) = y] \cdot 2^{i-n} \ = \ k \sum_{i=1}^{\max} \Pr[f(U_n) \in \mathcal{Y}_i] \cdot 2^{i-n} \\
&\leq \ k \cdot 2^{\max - n}(1 + 2^{-1} + \ldots + 2^{-(\max - 1)}) \ \leq \ k \cdot 2^{\max - n + 1}\ .
\end{aligned}$$

$\square$

Definition B.1 (bounded-width layered branching program - LBP) *An (s, k, v)-LBP M is a finite directed acyclic graph whose nodes are partitioned into $k + 1$ layers indexed by $\{1, \ldots, k+1\}$. The first layer has a single node (the source), the last layer has two nodes (sinks) labeled with 0 and 1, and each of the intermediate layers has up to 2^s nodes. Each node in the $i \in [k]$ layer has exactly 2^v outgoing labeled edges to the $(i+1)^{th}$ layer, one for every possible string $h_i \in \{0,1\}^v$.*

Theorem B.1 (bounded-space generator [20, 16]) *Let $s(n), k(n), v(n) \in \mathbb{N}$ and $\varepsilon(n) \in (0,1)$ be polynomial-time computable functions. Then, there exist a polynomial-time computable function $q(n) \in \Theta(v(n) + (s(n) + \log(k(n)/\varepsilon(n))) \log k(n))$ and a generator $BSG : \{0,1\}^{q(n)} \to \{0,1\}^{k(n) \cdot v(n)}$ that runs in time $\mathsf{poly}(s(n), k(n), v(n), \log(1/\varepsilon(n)))$, and $\varepsilon(n)$-fools every $(s(n), k(n), v(n))$- LBP M, i.e.,*

$$\mid \Pr[M(U_{k(n) \cdot v(n)}) = 1] \ - \ \Pr[M(BSG(U_n)) = 1] \mid \ \leq \ \varepsilon(n)\ .$$

Proof of (4). For any $k \in \mathbb{N}$, consider the following $(2n, k, \log|\mathcal{H}|)$-LBP M_1: on source node input $(y_1 = f(x_1), y'_1 = f(x'_1))$. For $1 \leq i \leq k$, at each i^{th} layer M_1 computes $y_i := f(h_{i-1}(y_{i-1}))$ and $y'_i := f(h_{i-1}(y'_{i-1}))$. Finally, at the $(k+1)^{th}$ layer M_1 outputs 1 iff $y_k = y'_k \in \mathcal{Y}_{\max}$. Imagine running two iterates with random x_1, x'_1 and seeded by a common hash function from distribution either $\vec{H}^k$ or $BSG(U_q)$, we have

$$\mathsf{CP}(\ Y_k \mid \vec{H}^k\ ) \ = \ \Pr_{(x_1, x'_1) \leftarrow U_{2n},\ \vec{h}^k \leftarrow \vec{H}^k}[M_1(x_1, x'_1, \vec{h}^k) = 1]$$

$$\mathsf{CP}(\ Y'_k \mid BSG(U_q)\ ) \ = \ \Pr_{(x_1, x'_1) \leftarrow U_{2n},\ \vec{h}^k \leftarrow BSG(U_q)}[M_1(x_1, x'_1, \vec{h}^k) = 1]$$

and thus

$$\begin{aligned}
&| \ \mathsf{CP}(Y_k \mid \vec{H}^k \ ) \ - \ \mathsf{CP}(\ Y'_k \mid BSG(U_q) \ ) \ | \\
&\leq \ \mathbb{E}_{(x_1,x'_1)\leftarrow U_{2n}} \left[\ | \ \Pr[M_1(x_1,x'_1,\vec{H}^k)=1] \ - \ \Pr[M_1(x_1,x'_1,BSG(U_q))=1] | \ \right] \\
&\leq \ 2^{-2n} \quad .
\end{aligned}$$

It follows by (3) that

$$\mathsf{CP}(Y'_k \mid BSG(U_q) \ ) \ \leq \ \mathsf{CP}(\ Y_k \mid \vec{H}^k \ ) \ + \ 2^{-2n} \ \leq \ k{\cdot}2^{\max-n+1} + 2^{-2n} \quad .$$

Note that for any $\vec{h}^k$ and any $u_1, u_2 \in BSG^{-1}(\vec{h}^k)$,

$$\mathsf{CP}(\ Y'_k \mid U_q = u_1 \ ) \ = \ \mathsf{CP}(\ Y'_k \mid U_q = u_2) \ = \ \mathsf{CP}(\ Y'_k \mid BSG(U_q) = \vec{h}^k \ ) \ .$$

Therefore,

$$\mathsf{CP}(Y'_k \mid U_q \ ) \ = \ \mathsf{CP}(Y'_k \mid BSG(U_q) \ ) \ \leq \ k{\cdot}2^{\max-n+1} + 2^{-2n} \quad .$$

□

Proof of (6). Similar to that of (4), we define another $(n+1, k, \log|\mathcal{H}|)$-LBP M_2 that on source node input $(x_1, \mathsf{tag}_1 = 0)$, it computes $y_i := f(x_i)$, $x_{i+1} := h_i(y_i)$, for every $i \leq k$ and sets $\mathsf{tag}_i = 1$ if $i = j$ and $y_i \in \mathcal{Y}_{\max}$ (or otherwise $\mathsf{tag}_i := \mathsf{tag}_{i-1}$). Finally, it outputs tag_k. Thus,

$$\Pr[\mathcal{E}'_j] \geq \Pr[\mathcal{E}_j] - 2^{-2n} \geq n^{-c} - 2^{-2n} \quad .$$

□

Proof of (8). Consider the following $(n+1, k, \log|\mathcal{H}|)$-LBP M_3: on source node input (x_1, tag_1) and layered input vector $\vec{h}^k$, it computes $y_i := f(x_i)$, $x_{i+1} := h_i(y_i)$, at each i^{th} layer, and sets $\mathsf{tag}_i = 1$ iff either $\mathsf{tag}_{i-1} = 1$ or $y_i \in \mathcal{Y}_{\max}$. Finally, M_3 outputs tag_k. By the bounded space generator we have

$$| \ \Pr[M_3(X_1,\vec{H}^k)=1] \ - \ \Pr[M_3(X_1,BSG(U_q)=1] \ | = \ | \Pr[\ \bigvee_{i=1}^{k} \mathcal{E}_i \ ] - \Pr[\ \bigvee_{i=1}^{k} \mathcal{E}'_i \ ] \ | \ \leq \ 2^{-2n} \ ,$$

and thus by (7)

$$\Pr[\ \bigvee_{i=1}^{k} \mathcal{E}'_i \ ] \ \geq \ \Pr[\ \bigvee_{i=1}^{k} \mathcal{E}_i \ ] \ - \ 2^{-2n} \ \geq \ 1 - \exp^{-k/n^{2c}} - 2^{-2n} \quad .$$

□

B.1 $\mathcal{F}$ Is a Family of One-way Functions

Proof of Lemma 5.2. Assume for contradiction that there exists A (of running time $t - n^{O(1)}$) that inverts f_u with some non-negligible ε_{A}, i.e.,

$$\Pr_{u\leftarrow U_q, \ x\xleftarrow{\$}\{0,1\}^n} [\ \mathsf{A}(u, f_u(x)) \in f_u^{-1}(\ f_u(x) \ ) \ ] \ \geq \varepsilon_{\mathsf{A}} \quad .$$

We use shorthand $\mathcal{C}$ for the event that A inverts f_u, i.e.,

$$\mathcal{C} \ \stackrel{\mathsf{def}}{=} \ \left((X_1, U_q) \in \{ \ (x,u) : \mathsf{A}(u, f_u(x)) \in f_u^{-1}(f_u(x)) \ \} \right)$$

and thus

$$\begin{aligned}
\varepsilon_{\mathsf{A}} &\leq \Pr[\mathcal{C}] \\
&= \Pr[\mathcal{C} \wedge (\mathcal{E}'_1 \vee \mathcal{E}'_2 \vee \ldots \vee \mathcal{E}'_k)] + \Pr[\mathcal{C} \wedge \neg(\mathcal{E}'_1 \vee \mathcal{E}'_2 \vee \ldots \vee \mathcal{E}'_k)] \\
&\leq \Pr[\bigvee_{j=1}^{k} (\mathcal{C} \wedge \mathcal{E}'_j)] \;+\; \Pr[\neg(\mathcal{E}'_1 \vee \mathcal{E}'_2 \vee \ldots \vee \mathcal{E}'_k)] \\
&\leq \sum_{j=1}^{k} \Pr[\mathcal{C} \wedge \mathcal{E}'_j] \;+\; (\exp^{-k/n^{2c}} + 2^{-2n})
\end{aligned}$$

where the third inequality follows from the union bound and (8). We have by an averaging argument that there exists $j^* \in [k]$ such that $\Pr[\mathcal{C} \wedge \mathcal{E}'_{j^*}] \geq (\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})/k$. That is, conditioned on event $\mathcal{E}'_{j^*}$, algorithm A inverts $f_u(x) = f^k(x, BSG(u))$ to produce $x' \in {f_u}^{-1}(f_u(x))$ with probability

$$\Pr[\; \mathcal{C} \mid \mathcal{E}'_{j^*} \;] \;=\; \frac{\Pr[\mathcal{C} \wedge \mathcal{E}'_{j^*}]}{\Pr[\; \mathcal{E}'_{j^*} \;]} \;\geq\; \Pr[\mathcal{C} \wedge \mathcal{E}'_{j^*}] \;\geq\; (\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})/k \;\; .$$

The above implies an algorithm M^{A} (as given in Algorithm 4) that inverts $y_{j^*} = f^{j^*}(x, BSG(u))$ (with respect to f) to get $x_{j^*} \in f^{-1}(y_{j^*})$ with almost the same probability. Loosely speaking, on input y_{j^*}, the algorithm M^{A} evaluates the iterate to obtain y_k, invokes A on y_k to get x_1, and produces x_{j^*} (again by evaluating the iterate on x_1) as a candidate preimage of y_{j^*} under function f. The only issue is that j^* is unknown, so it simply makes a random guess $j \xleftarrow{\$} [k]$, which hits j^* with probability $1/k$. Therefore, it holds that

Algorithm 4 M^{A}.

Input: $u \in \{0,1\}^q$, $y \in \{0,1\}^n$

Sample $j \xleftarrow{\$} [k]$;
$\big(\vec{h}^k = (h_1, \ldots, h_k)\big) := BSG(u)$;
Let $\tilde{y}_j := y$;
FOR $i = j+1$ TO k
    Compute $\tilde{x}_i := h_{i-1}(\tilde{y}_{i-1})$, $\tilde{y}_i := f(\tilde{x}_i)$;
$\tilde{x}_1 \leftarrow \mathsf{A}(u, \tilde{y}_k)$;
FOR $i = 1$ TO $j-1$
    Compute $\tilde{y}_i := f(\tilde{x}_i)$, $\tilde{x}_{i+1} := h_i(\tilde{y}_i)$;

Output: $\tilde{x}_j$

$$\Pr[\; \mathsf{M}^{\mathsf{A}}(U_q, Y'_{j^*};\; j \;) \in f^{-1}(Y'_{j^*}) \;\mid\; j = j^* \;\wedge\; \mathcal{E}'_{j^*} \;] \;\geq\; (\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})/k \;\; , \tag{13}$$

where we recall that $Y'_{j^*} = f^{j^*}(X_1, BSG(U_q))$. We state in Claim B.1 that replacing the above Y'_{j^*} (which correlated to U_q) with $f(U_n)$ (which is independent of U_q) the inverting probability weakens only by a $1/\mathsf{poly}(n)$ factor and thus M^{A} becomes an inverter for f.

$$\begin{aligned}
\varepsilon \;&\geq\; \Pr[\; \mathsf{M}^{\mathsf{A}}(\; U_q, f(U_n);\; j \;) \in f^{-1}(f(U_n)) \;] \\
&=\; \Pr[j = j^*] \cdot \Pr[f(U_n) \in \mathcal{Y}_{\max} \;] \cdot \Pr[\; \mathsf{M}^{\mathsf{A}}(\; U_q, f(U_n);\; j \;) \in f^{-1}(f(U_n)) \;\mid\; j = j^* \;\wedge\; f(U_n) \in \mathcal{Y}_{\max} \;] \\
&\geq\; \frac{(\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})^2}{2^9 \cdot k^4 \cdot n^{3c}} \;\; ,
\end{aligned}$$

where the first inequality is due to the one-way-ness of f. This yields an upper bound on ε_{A} (by taking a square root) as desired and thus completes the proof. □

THE RATIONALE FOR CLAIM B.1. By Lemma B.1, the collision probability of (U_q, Y'_{j^*}) conditioned on $\mathcal{E}'_{j^*}$ is small enough and close to that of the uniform distribution of $(U_q, U_{\mathcal{Y}_{\max}})$. Thus, any algorithm that inverts the former (i.e., Y'_{j^*} given U_q) with a non-negligible probability will invert the the latter (i.e., $U_{\mathcal{Y}_{\max}}$ given uncorrelated U_q) with also a non-negligible probability.

Claim B.1

$$\Pr[\ \mathsf{M}^{\mathsf{A}}(\ U_q, f(U_n);\ j\ ) \in f^{-1}(f(U_n))\ \mid\ j = j^*\ \wedge\ f(U_n) \in \mathcal{Y}_{\max}\ ]\ \geq\ \frac{(\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})^2}{2^9 \cdot k^3 \cdot n^{2c}}\ . \quad (14)$$

Proof of Claim B.1. To apply Lemma B.1, let $\mathcal{W} = \{0,1\}^q \times \mathcal{Y}_{\max}$, let Z be empty set, W be the distribution of (U_q, Y'_{j^*}) conditioned on $\mathcal{E}'_{j^*}$ (i.e., $Y'_{j^*} \in \mathcal{Y}_{\max}$), and define

$$\mathsf{Adv}(u, y)\ \stackrel{\text{def}}{=}\ \Pr[\ \mathsf{M}^{\mathsf{A}}(u, y;\ j\ ) \in f^{-1}(y)\ \mid\ j = j^*\ ]\ .$$

where the probability is taken over the internal coins of M^{A}. Thus, we have

$$\mathsf{Adv}(W)\ =\ \Pr[\ \mathsf{M}^{\mathsf{A}}(U_q, Y'_{j^*};\ j\ ) \in f^{-1}(Y'_{j^*})\ \mid\ j = j^*\ \wedge\ \mathcal{E}'_{j^*}\ ]\ .$$

and

$$\begin{aligned}
\mathsf{CP}(W)\ &=\ \mathsf{CP}(\ (U_q, Y'_{j^*}\ )\ \mid\ \mathcal{E}'_{j^*}\ ) = \frac{\mathsf{CP}(\ (U_q, Y'_{j^*}\ )\ \wedge\ \mathcal{E}'_{j^*})}{\Pr[\ \mathcal{E}'_{j^*}\ ]^2} \\
&\leq\ \frac{\mathsf{CP}(\ (U_q, Y'_{j^*})\ )}{\Pr[\ \mathcal{E}'_{j^*}\ ]^2} \\
&\leq\ \frac{2^{-q} \cdot \mathsf{CP}(\ Y'_{j^*}\ |U_q\ )}{(n^{-c} - 2^{-2n})^2} \\
&\leq\ \frac{j^* \cdot 2^{\max - n + 1} + 2^{-2n}}{(n^{-2c}/4) \cdot 2^q}\ \leq\ \frac{16k \cdot n^{2c}}{2^{n - \max + q}}\ \leq\ \underbrace{(32k \cdot n^{2c})}_{2^e} \cdot 2^{-m}\ ,
\end{aligned}$$

where the fourth inequality is due to $2^{-2n} \leq j^* \cdot 2^{\max - n + 1}$ and $j^* \leq k$ and the fifth inequality is by $2^{m-q} = |\mathcal{Y}_{\max}| \leq 1/2^{\max - 1 - n}$. We thus have

$$\begin{aligned}
&\Pr[\ \mathsf{M}^{\mathsf{A}}(\ U_q, f(U_n);\ j\ ) \in f^{-1}(f(U_n))\ \mid\ j = j^*\ \wedge\ f(U_n) \in \mathcal{Y}_{\max}\ ] \\
=\ &\frac{\sum_{(u,y) \in \{0,1\}^q \times \mathcal{Y}_{\max}} 2^{-q} \cdot \Pr[f(U_n) = y] \cdot \mathsf{Adv}(u, y)}{\sum_{y \in \mathcal{Y}_{\max}} \Pr[f(U_n) = y]} \\
\geq\ &\frac{\sum_{(u,y) \in \{0,1\}^q \times \mathcal{Y}_{\max}} 2^{-q} \cdot \frac{1}{2|\mathcal{Y}_{\max}|} \cdot \mathsf{Adv}(u, y)}{\sum_{y \in \mathcal{Y}_{\max}} \frac{2}{|\mathcal{Y}_{\max}|}} \\
=\ &\frac{\mathbb{E}[\mathsf{Adv}(U_{\mathcal{W}})]}{4}\ \geq\ \frac{\mathbb{E}[\mathsf{Adv}(W)]^2}{2^{e+4}} \\
\geq\ &\frac{(\varepsilon_{\mathsf{A}} - \exp^{-k/n^{2c}} - 2^{-2n})^2}{2^9 \cdot k^3 \cdot n^{2c}}\ ,
\end{aligned}$$

where the first inequality is because for any $y \in \mathcal{Y}_{\max}$ we have $1/2|\mathcal{Y}_{\max}| \leq \Pr[f(U_n) = y] \leq 2/|\mathcal{Y}_{\max}|$, the second inequality follows from (12) and the third is due to (13) and $2^e = 32k \cdot n^{2c}$. □