THEORY OF COMPUTING, Volume 16 (8), 2020, pp. 1–55
www.theoryofcomputing.org

Inaccessible Entropy II: IE Functions and Universal One-Way Hashing

Iftach Haitner* Thomas Holenstein Omer Reingold† Salil Vadhan‡
Hoeteck Wee§

Received January 22, 2015; Revised September 6, 2020; Published October 13, 2020

Abstract: This paper uses a variant of the notion of *inaccessible entropy* (Haitner, Reingold, Vadhan and Wee, STOC 2009), to give an alternative construction and proof for the fundamental result, first proved by Rompel (STOC 1990), that *Universal One-Way Hash Functions (UOWHFs)* can be based on any one-way functions. We observe that a small tweak of any one-way function f is already a weak form of a UOWHF: consider the function $F(x,i)$ that returns the i-bit-long prefix of $f(x)$. If F were a UOWHF then given a random x and i it would be hard to come up with $x' \neq x$ such that $F(x,i) = F(x',i)$. While this may not be the

A preliminary version of this paper appeared as "Universal One-Way Hash Functions via Inaccessible Entropy" in EUROCRYPT 2010 [9].

*Supported by ERC starting grant 638121, ISF grant 1076/11 and US-Israel BSF grant 2010196.

†Initial research and conference version supported by US-Israel BSF grant 2006060. Preparation of this journal version supported by NSF grant CCF-1763311.

‡Initial research and conference version supported by NSF grants CNS-0831289 and US-Israel BSF grants 2006060 and 2010196. Preparation of this journal version supported by NSF grant CCF-1763299 and a Simons Investigator Award.

§Most of the work was done while at Queens College, CUNY. Supported in part by PSC-CUNY Award #6014939 40 and NSF CAREER Award CNS-0953626.

ACM Classification: F.0

AMS Classification: 03D15

Key words and phrases: cryptography, complexity theory, one-way functions, computational entropy

DOI: 10.4086/toc.2020.v016a008

case, we show (rather easily) that it is hard to sample x' with almost full entropy among all the possible such values of x'. The rest of our construction simply amplifies and exploits this basic property. Combined with other recent work, the construction of three fundamental cryptographic primitives (Pseudorandom Generators, Statistically Hiding Commitments and UOWHFs) out of one-way functions is now to a large extent unified. In particular, all three constructions rely on and manipulate computational notions of entropy in similar ways. Pseudorandom Generators rely on the well-established notion of pseudoentropy, whereas Statistically Hiding Commitments and UOWHFs rely on the newer notion of inaccessible entropy.

In an additional result we reprove the seminal result of Impagliazzo and Levin (FOCS 1989): a reduction from "uniform distribution" average-case complexity problems to ones with arbitrary (polynomial-time samplable) distributions. We do that using techniques similar to those we use to construct UOWHFs from one-way functions, where the source of this similarity is the use of a notion similar to inaccessible entropy. This draws an interesting connection between two seemingly separate lines of research: average-case complexity and universal one-way hash-functions.

1 Introduction

Unlike the more common notions of computational entropy, e. g., *pseudoentropy* [15], that are only useful as a *lower* bound on the "computational entropy" of a distribution, *accessible entropy* is an *upper* bound on computational entropy. In particular, it measures the entropy (Shannon, or other types) of the distribution computed by a resource-bounded machine.

The inaccessible entropy of the distribution is the gap between its (Shannon) entropy and its accessible entropy. Inaccessible entropy was introduced by Haitner, Reingold, Vadhan and Wee [12] as a means to give a simpler construction and proof of statistically hiding commitment from one-way functions (reproving the result of [10]), and to construct constant-round statistically hiding commitment from constant-round zero-knowledge proof for NP. In this article introduce simpler variant of their notion to give an alternative construction and proof for the fundamental result, first proved by Rompel [26], that *Universal One-Way Hash Functions (UOWHFs)* can be based on one-way functions. In an additional result, we reprove the seminal result of Impagliazzo and Levin [18]: a reduction from "uniform distribution" average-case complexity problems to ones with arbitrary (though polynomial samplable one) distributions. The latter is proved using similar techniques to the ones we use to construct UOWHFs from one-way functions, where the source of this similarity is the use of a similar notion of inaccessible entropy. This draws an interesting connection between two seemingly separate lines of research: average-case complexity and universal one-way hash-functions.

We start by discussing our construction of universal one-way hash functions, where the result about average-case complexity is described in Section 1.3. *Universal one-way hash functions* (UOWHFs), as introduced by Naor and Yung [23], are a weaker form of *collision-resistant hash functions*; a function family $\mathcal{F}$ is collision resistant if given a randomly chosen function $f \in \mathcal{F}$, it is infeasible to find any pair of distinct inputs x, x' such that $f(x) = f(x')$. UOWHFs only require *target collision resistance*, where the adversary must specify one of the inputs x before seeing the description of the function f. We give a formal definition.

Definition 1.1. A family of functions

$$\mathcal{F}_k = \left\{ F_z \colon \{0,1\}^{n(k)} \to \{0,1\}^{m(k)} \right\}_{z \in \{0,1\}^k},$$

for $n(k), m(k) \in \mathsf{poly}(k)$, is a family of universal one-way hash functions (UOWHFs) if it satisfies:

Efficiency: given $z \in \{0,1\}^k$ and $x \in \{0,1\}^{n(k)}$, $F_z(x)$ can be evaluated in time $\mathrm{poly}(k)$.

Shrinking: $m(k) < n(k)$.

Target Collision Resistance: the probability that a PPT (probabilistic polynomial-time) adversary A succeeds in the following game is negligible in k:

1. Let $(x, \mathsf{state}) \leftarrow \mathsf{A}(1^k)$.
2. Let $z \leftarrow \{0,1\}^k$.
3. Let $x' \leftarrow \mathsf{A}(\mathsf{state}, z)$.
4. A *succeeds* if $x, x' \in \{0,1\}^{n(k)}$, $x \neq x'$, and $F_z(x) = F_z(x')$. [1]

It turns out that this weaker security property suffices for many applications. The most immediate application given in [23] is *secure fingerprinting*, whereby the pair $(f, f(x))$ can be taken as a compact "fingerprint" of a large file x, such that it is infeasible for an adversary, seeing the fingerprint, to change the file x to x' without being detected. More dramatically, [23] also showed that UOWHFs can be used to construct secure digital signature schemes, whereas all previous constructions (with proofs of security in the standard model) were based on trapdoor functions (as might have been expected to be necessary due to the public-key nature of signature schemes). More recently, UOWHFs have been used in the Cramer–Shoup encryption scheme [8] and in the construction, from one-way functions, of statistically hiding commitment schemes [10, 12].

Naor and Yung [23] gave a simple and elegant construction of UOWHFs from any one-way *permutation*. [28] generalized the construction of [23] to get UOWHFs from *regular* one-way functions. Rompel [26] gave a more involved construction to prove that UOWHFs can be constructed from an arbitrary one-way function, thereby resolving the complexity of UOWHFs (as one-way functions are the minimal complexity assumption for complexity-based cryptography, and are easily implied by UOWHFs); this remains the state of the art for arbitrary one-way functions.[2] While complications may be expected for constructions from arbitrary one-way functions (due to their lack of structure), Rompel's analysis also feels quite ad hoc. In contrast, the construction of pseudorandom generators from one-way functions in [15], while also somewhat complex, involves natural abstractions (e. g., pseudoentropy) that allow for modularity and measure for what is being achieved at each stage of the construction.

[1] The $\leftarrow$ notation is explained in Sec. 2.1. Namely, on security parameter 1^k, algorithm A first samples an element x in the function family input domain. Then given (the description of) a function F_z uniformly drawn from the family, algorithm A has to find a collision with x: an element $x' \neq x$ that F_z maps to the same output value. To avoid discussing *stateful* algorithms, we do not allow A to keep a "state" between the game stages. Rather, we enable it to transfer information between the stages using the auxiliary string state.

[2] More details of the proof given in [26] are worked out, with some corrections, in [27, 20].

In this paper, we give simpler constructions of UOWHFs from one-way functions, based on (a variant of) the recently introduced notion of *inaccessible entropy* [12]. In addition, one of the constructions obtains slightly better efficiency and security than Rompel's original construction.

1.1 Inaccessible entropy

For describing our construction, it will be cleaner to work with a variant of UOWHFs where there is a *single* shrinking function $F : \{0,1\}^n \to \{0,1\}^m$ (for each setting of the security parameter k) such that it is infeasible to find collisions with *random inputs*. So in our model an adversary A is given a uniformly random $x \leftarrow \{0,1\}^n$, outputs an x' such that $F(x') = F(x)$, and succeeds[3] if $x' \neq x$. Note that we can assume without loss of generality that $x' = \mathsf{A}(x)$ is always a preimage of $F(x)$ (A has the option of returning x in case it does not find a different preimage); we refer to an algorithm A with this property as an *F-collision finder*.

Our construction is based on an information-theoretic view of UOWHFs. The fact that F is shrinking implies that there are many preimages x' of $F(x)$ available to A. Indeed, if we consider an (inefficient) adversary $\mathsf{A}(x)$ that outputs a uniformly random preimage $x' \leftarrow F^{-1}(F(x))$ and let X be a random variable uniformly distributed on $\{0,1\}^n$, then

$$\mathrm{H}(\mathsf{A}(X) \mid X) = \mathrm{H}(X \mid F(X)) \geq n - m,$$

where $\mathrm{H}(\cdot \mid \cdot)$ denotes conditional Shannon entropy. (See Section 2 for more definitional details.) We refer to the quantity $\mathrm{H}(X \mid F(X))$ as the *real entropy of F^{-1}*.

On the other hand, *target collision resistance* means that effectively only one of the preimages is accessible to A. That is for every probabilistic polynomial-time F-collision finder A, we have $\Pr[\mathsf{A}(X) \neq X] = \mathrm{neg}(n)$, which is equivalent to requiring that:

$$\mathrm{H}(\mathsf{A}(X) \mid X) = \mathrm{neg}(n)$$

for all probabilistic polynomial-time F-collision finders A. (If A can find a collision X' with non-negligible probability, then it can achieve non-negligible conditional entropy by returning X' with probability 1/2 and returning X with probability $1/2$.) We refer to the maximum of $\mathrm{H}(\mathsf{A}(X) \mid X)$ over all efficient F-collision finders as the *accessible entropy of F^{-1}*. We emphasize that accessible entropy refers to an *upper bound* on a form of computational entropy, in contrast to the Håstad et al. notion of *pseudoentropy* [15].

Thus, a natural weakening of the UOWHF property is to simply require a noticeable gap between the real and the accessible entropies of F^{-1}. That is, for every probabilistic polynomial-time F-collision finder A, we have $\mathrm{H}(\mathsf{A}(X) \mid X) < \mathrm{H}(X \mid F(X)) - \Delta$, for some $\Delta \geq 1/\mathrm{poly}(n)$, which we refer to as the *inaccessible entropy of F*.

1.2 Our UOWHF constructions

Our constructions of UOWHFs have two parts. First, we show how to obtain a function with noticeable inaccessible entropy from any one-way function. Second, we show how to build a UOWHF from any function with inaccessible entropy.

[3]It is easy to convert any such function F into a standard UOWHF family by defining $F_z(x) = F(z+x)$.

OWFs $\implies$ inaccessible entropy. Given a one-way function $f\colon \{0,1\}^n \to \{0,1\}^m$, we show that a random truncation of f has inaccessible entropy. Specifically, we define $F(x,i)$ to be the first i bits of $f(x)$.

To see that this works, suppose for contradiction that F does not have noticeable inaccessible entropy. That is, we have an efficient adversary A that on input (x,i) can sample from the set $S(x,i) = \{x'\colon f(x')_{1\dots i} = f(x)_{1\dots i}\}$ with almost-maximal entropy, which is equivalent to sampling according to a distribution that is statistically close to the uniform distribution on $S(x,i)$. We can now use A to construct an inverter Inv for f that works as follows on input y: choose $x_0 \leftarrow \{0,1\}^n$, and then for $i = 1,\dots,n$ generate a random $x_i \leftarrow A(x_{i-1}, i-1)$ subject to the constraint that $f(x_i)_{1,\cdots,i} = y_{1,\cdots,i}$. The latter step is feasible, since we are guaranteed that $f(x_i)_{1,\dots,i-1} = y_{1,\cdots,i-1}$ by the fact that A is an F-collision finder, and the expected number of trials needed get agreement with y_i is at most $O(1)$ (since $y_i \in \{0,1\}$, and y and $f(x_i)$ are statistically close). It is not difficult to show that when run on a random output Y of f, Inv produces an almost-uniform preimage of Y. This contradicts the one-wayness of f. Indeed, we only need f to be a *distributional* one-way function [19], whereby it is infeasible to generate almost-uniform preimages under f.

Inaccessible entropy $\implies$ UOWHFs. Once we have a non-negligible amount of inaccessible entropy, we can construct a UOWHF via a series of standard transformations.

1. Repetition: By evaluating F on many inputs, we can increase the amount of inaccessible entropy from $1/\mathrm{poly}(n)$ to $\mathrm{poly}(n)$. Specifically, we take $F^t(x_1,\dots,x_t) = (F(x_1),\dots,F(x_t))$ where $t = \mathrm{poly}(n)$. This transformation also has the useful effect of converting the real entropy of F^{-1} to *real min-entropy*: with very high probability $x \leftarrow X^t$, $F^t(x)$ has large number of pre-images.

2. Hashing Inputs: By hashing the input to F (namely taking $F'(x,g) = (F(x), g(x))$ for a universal hash function g), we can reduce both the real (min-)entropy and the accessible entropy so that $(F')^{-1}$ still has a significant amount of real entropy, but has (weak) target collision resistance (on random inputs).

3. Hashing Outputs: By hashing the output to F (namely taking $F'(x,g) = g(F(x))$), we can reduce the output length of F to obtain a shrinking function that still has (weak) target collision resistance.

There are two technicalities that occur in the above steps. First, hashing the inputs only yields *weak* target collision resistance; that is, we only guarantee that an adversary fails in the target collision resistance game with probability at least $1/\mathrm{poly}(n)$ rather than $1-\mathrm{neg}(n)$. This is due to the fact that accessible Shannon entropy is an average-case measure and thus allows for the possibility that the adversary can achieve high accessible entropy most of the time. Fortunately, this weak form of target collision resistance can be amplified to full target collision resistance using another application of repetition and hashing (similar to [6]).

Second, the hashing steps require having a fairly accurate estimate of the real entropy. This can be handled similarly to [15, 26], by trying all (polynomially many) possibilities and concatenating the resulting UOWHFs, at least one of which will be target collision resistant.

A more efficient construction. We obtain a more efficient construction of UOWHFs by hashing the output of the one-way function f before truncating. That is, we define $F(x,g,i) = (g, g(f(x))_{1\cdots i})$. This function is in the spirit of the function that Rompel [26] uses as a first step, but our function uses a three-wise independent hash function instead of an n-wise independent one, and enjoys a simpler structure.[4] Our analysis of this function is simpler than Rompel's and can be viewed as providing a clean abstraction of what it achieves (namely, inaccessible entropy) that makes the subsequent transformation to a UOWHF easier.

We obtain improved UOWHF parameters over our first construction for two reasons. First, we obtain a larger amount of inaccessible entropy: $(\log n)/n$ bits instead of roughly $1/n^4$ bits. Second, we obtain a bound on a stronger form of accessible entropy, which enables us to get full target collision resistance when we hash the inputs, avoiding the second amplification step.

The resulting overall construction yields better parameters than Rompel's original construction. A one-way function of input length n yields a UOWHF with output length $\widetilde{O}(n^7)$, slightly improving Rompel's bound of $\widetilde{O}(n^8)$. Additionally, we are able to reduce the key length needed: Rompel's original construction uses a key of length $\widetilde{O}(n^{12})$, whereas our construction only needs a key of length $\widetilde{O}(n^7)$. If we allow the construction to utilize some nonuniform information (namely an estimate of the real entropy of F^{-1}), then we obtain output length $\widetilde{O}(n^5)$, improving Rompel's bound of $\widetilde{O}(n^6)$. For the key length, the improvement in this case is from $\widetilde{O}(n^7)$ to $\widetilde{O}(n^5)$. Of course, these bounds are still far from practical, but they illustrate the utility of inaccessible entropy in reasoning about UOWHFs, which may prove useful in future constructions (whether based on one-way functions or other building blocks).

1.3 Connection to average-case complexity

We use the notion of inaccessible entropy to reprove the following theorem by Impagliazzo and Levin [18], given in the realm of *average-case complexity*.

Theorem 1.2 ([18], informal)**.** *Assume there exists an* NP *language that is hard on some (efficiently) samplable distribution for* heuristics*: every efficient algorithm fails to decide the language correctly on a noticeable part of the distribution. Then there exists a language in* NP *that is hard against heuristics on the uniform distribution.*

Our proof follows to a large extent in the footsteps of [18], where the main novelty is formulating the proof in the language of inaccessible entropy, and rephrasing it to make it resembles our proof of UOWHFs from one-way functions. This draws an interesting connection between two seemingly separate lines of research: average-case complexity and universal one-way hash-functions.

As in [18], we prove Theorem 1.2 by proving its search variant: hardness of finding a witness on a samplable distribution implies hardness of finding a witness on the uniform distribution. Let $(\mathcal{R},\mathcal{D})$ be an NP-search problem (i. e., $\mathcal{R}$ is an NP relation and $\mathcal{D}$ is a samplable distribution) that is hard to solve heuristically, and let D be the algorithm sampling instances according to $\mathcal{D}$. For a family of pair-wise independent hash functions $\mathcal{G}$, consider the following NP relation:

$$\mathcal{R}' = \left\{(x' = (g,i), w' = (x,w)) \colon g \in \mathcal{G}, (D(x),w) \in \mathcal{R} \wedge g(D(x))_{1,\ldots,i} = 0^i\right\}^{5}$$

[4][26] started with the function $f'(z,g_1,g_2) := (g_2(f_0(g_1(z))), g_1, g_2)$, where g_1 and g_2 are n-wise independent hash-functions, and f_0 is defined as $f_0(x,y,i) = (f(x), y^{n-i}, 0^i)$.

[5]To keep the discussion simple, we ignore input length consideration. See Section 6 for the formal definitions.

While $\mathcal{R}'_{\mathcal{L}}$ (i. e., the language of $\mathcal{R}'$) might not be hard on the uniform distribution (interpreted as the uniform distribution over random pairs (g,i)), it is not hard to prove that the distribution is "somewhat hard." In particular, it happens noticeably often that i is the "right one for $(\mathcal{R},\mathcal{D})$"; meaning that for a fixed element in y which might be in the language, exactly one x with $D(x) = y$ satisfies $g(x) = 0^i$. Conditioned on this event, letting $\mathsf{A}(\cdot)_x$ being the x part in the witness output by A, it is not hard to show that $\mathrm{H}(\mathsf{A}(G,I)_x)$ is noticeably less than its information theoretic maximum: for any efficient algorithm A, it holds that

$$\mathrm{H}(\mathsf{A}(G,I)_x) < \mathrm{H}(\mathcal{D}|_{G(\mathcal{D})_{1,\dots,I}=0^I}),$$

where (G,I) is the parsing of a random string into a pair $(g \in \mathcal{G}, i)$, and assuming for simplicity that A never fails to provide a correct witness. Namely, the (accessible) entropy of A is smaller than the (real) entropy of $\mathcal{D}$. Using similar means to the ones used to amplify the initial UOWHFs constructions described in Section 1.2 , the above gap can be amplified to induce hardness over the uniform distribution.

1.4 Perspective

The idea of inaccessible entropy was introduced in [12] for the purpose of constructing statistically hiding commitment schemes from one-way functions and from zero-knowledge proofs. There, the nature of statistically hiding commitments necessitated more involved notions of inaccessible entropy than we present here — inaccessible entropy was defined in [12] for interactive protocols, where one considers adversaries that try to generate next-messages or next-blocks of high entropy. (See [13] for a simpler notion of inaccessible entropy that suffices for commitments based on one-way functions.)

Here, we are able to work with a much simpler form of inaccessible entropy (significantly simpler also than the notion considered in [13]). The simplicity comes from the non-interactive nature of UOWHFs (and of solving NP problems) so we only need to measure the entropy of a single string output by the adversary. Thus, the definitions here can serve as a gentler introduction to the concept of inaccessible entropy.

On the other hand, the many-round notions from [12, 13] allow for a useful "entropy equalization" transformation that avoids the need to try all possible guesses for the entropy. We do not know an analogous transformation for constructing UOWHFs. We also note that our simple construction of a function with inaccessible entropy by randomly truncating a one-way function (and its analysis) is inspired by the construction of an "inaccessible entropy generator" from a one-way function in [12].

Finally, with our constructions, the proof that one-way functions imply UOWHFs now parallels those of pseudorandom generators [15, 11] and statistically hiding commitments [10, 12], with UOWHFs and statistically hiding commitments using dual notions of entropy (high real entropy, low accessible entropy) to pseudorandom generators (low real entropy, high pseudoentropy). See the surveys [14, 30].

1.5 Related work

UOWHFs. Katz and Koo [20] gave a complete write-up of Rompel's result [26, 27], with some corrections. Prior to our paper, Rompel's result represented the state of the art for UOWHFs from arbitrary one-way functions. Since the initial publication of this work in 2010 [9], there have been several improvements in the setting of regular one-way functions. Ames et al. [1] presented an even more efficient construction of UOWHFs from (unknown) regular one-way functions. Barhum and Maurer [2] gave even a more efficient construction assuming the regularity of the one-way function is *known*, where Yu et al. [32] improved the result of [2] presenting an almost optimal construction (with respect to the known black-box impossibility results) of UOWHFs from known-regular one-way functions.

Average-case complexity. The notion of average-case complexity was first introduced by Levin [21]. We focus on the result by Impagliazzo and Levin [18] who show that if it is possible use a polynomial time sampler to pick average-case problems which are hard, then there is a different problem which is hard on average for the uniform distribution. We give a different perspective on that proof, and in particular highlight the connections to inaccessible entropy. A good overview of average-case complexity was given by Bogdanov and Trevisan [5].

Recently, Hubáček et al. [17] made a different, and very elegant, connection between constructing UOWHFs from OWFs and average-case hardness on the uniform distribution, showing that a solution to the first challenge implies a solution to the second one. Their approach is surprisingly simple: if OWFs exist, then UOWHFs also exist, which can be seen as a problem that is hard on the uniform distribution (given a UOWHF key and an input, find a colliding input). On the other hand, assuming OWFs do not exist, without loss of generality the sampler of a hard-on-the-average problem can be assumed to output its random coins (indeed, its coins can be sampled from its original output under the assumption that OWFs do not exist). So in both cases, a hard-on-the-average problem implies a hard-on-the-average problem with respect to the uniform distribution.

Organization of the paper

Formal definitions are given in Section 2, where the notion of inaccessible entropy used through the paper is defined in Section 3. In Section 4 we show how to use any one-way function to get a function with inaccessible entropy, where in Section 5 we use any function with inaccessible entropy to construct UOWHF. Finally, our result for average-case complexity is described in Section 6.

2 Preliminaries

Most of the material in this section is taken almost verbatim from [12], and missing proofs can be found in that paper.

2.1 Notation

All logarithms considered here are in base two. For $t \in \mathbb{N}$, we let $[t] = \{1, \ldots, t\}$. A function $\mu : \mathbb{N} \to [0, 1]$ is *negligible*, denoted $\mu(n) = \text{neg}(n)$, if $\mu(n) = n^{-\omega(1)}$. We let poly denote the set of all polynomials, and let PPT stand for probabilistic polynomial time. Given a distribution D, we write $d \leftarrow D$ to indicate that d is selected according to D. Similarly, given a finite set $\mathcal{S}$, we write $s \leftarrow \mathcal{S}$ to indicate that s is selected according to the uniform distribution on $\mathcal{S}$.

2.2 Random variables

Let X and Y be random variables taking values in a discrete universe $\mathcal{U}$. We adopt the convention that when the same random variable appears multiple times in an expression, all occurrences refer to the same instantiation. For example, $\Pr[X = X]$ is 1. For an event E, we write $X|_E$ to denote the random variable X conditioned on E. The *support* of a random variable X is $\mathrm{Supp}(X) := \{x\colon \Pr[X = x] > 0\}$. X is *flat* if it is uniform on its support. For an event E, we write $I(E)$ for the corresponding indicator random variable, i. e., $I(E)$ is 1 when E occurs and is 0 otherwise.

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

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

We say that X and Y are ε*-close* if $\|X - Y\| \leq \varepsilon$ and ε*-far* otherwise.

2.3 Entropy measures

In this article we shall refer to several measures of entropy. The relation and motivation of these measures is best understood by considering a notion that we will refer to as the *sample-entropy*: For a random variable X and $x \in \mathrm{Supp}(X)$, we define the sample-entropy of x with respect to X to be the quantity

$$\mathrm{H}_X(x) := \log(1/\Pr[X = x]).$$

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

$$\begin{aligned} \mathrm{H}(X) &:= \mathop{\mathrm{E}}_{x \leftarrow X}[\mathrm{H}_X(x)], \\ \mathrm{H}_\infty(X) &:= \min_{x \in \mathrm{Supp}(X)} \mathrm{H}_X(x). \end{aligned}$$

We will also discuss the *max-entropy* $\mathrm{H}_0(X) := \log(|\mathrm{Supp}(X)|)$. The term "max-entropy" and its relation to the sample-entropy will be made apparent below.

It can be shown that $\mathrm{H}_\infty(X) \leq \mathrm{H}(X) \leq \mathrm{H}_0(X)$ with equality if and only if X is flat. Thus, saying $\mathrm{H}_\infty(X) \geq k$ is a strong way of saying that X has "high entropy" and $\mathrm{H}_0(X) \leq k$ a strong way of saying that X as "low entropy".

Smoothed entropies. Shannon entropy is robust in that it is insensitive to small statistical differences. Specifically, if X and Y are ε-close then $|\mathrm{H}(X) - \mathrm{H}(Y)| \leq \varepsilon \cdot \log|\mathcal{U}|$. For example, if $\mathcal{U} = \{0,1\}^n$ and $\varepsilon = \varepsilon(n)$ is a negligible function of n (i. e., $\varepsilon = n^{-\omega(1)}$), then the difference in Shannon entropies is vanishingly small (indeed, negligible). In contrast, min-entropy and max-entropy are brittle and can change dramatically with a small statistical difference. Thus, it is common to work with "smoothed" versions of these measures, whereby we consider a random variable X to have *high entropy* if X is ε-close to some X' with $\mathrm{H}_\infty(X) \geq k$ and to have *low entropy* if X is ε-close to some X' with $\mathrm{H}_0(X) \leq k$, for some parameter k and a negligible ε.[6]

[6] The term "smoothed entropy" was coined by [25], but the notion of smoothed min-entropy has commonly been used (without a name) in the literature on randomness extractors [24].

These smoothed versions of min-entropy and max-entropy can be captured quite closely (and more concretely) by requiring that the sample-entropy be large or small, resp., with high probability:

Lemma 2.1.

1. Suppose that with probability at least $1-\varepsilon$ over $x \leftarrow X$, we have $\mathrm{H}_X(x) \geq k$. Then X is ε-close to a random variable X' such that $\mathrm{H}_\infty(X') \geq k$.

2. Suppose that X is ε-close to a random variable X' such that $\mathrm{H}_\infty(X') \geq k$. Then with probability at least $1-2\varepsilon$ over $x \leftarrow X$, we have $\mathrm{H}_X(x) \geq k - \log(1/\varepsilon)$.

Lemma 2.2.

1. Suppose that with probability at least $1-\varepsilon$ over $x \leftarrow X$, we have $\mathrm{H}_X(x) \leq k$. Then X is ε-close to a random variable X' such that $\mathrm{H}_0(X') \leq k$.

2. Suppose that X is ε-close to a random variable X' such that $\mathrm{H}_0(X') \leq k$. Then with probability at least $1-2\varepsilon$ over $x \leftarrow X$, we have $\mathrm{H}_X(x) \leq k + \log(1/\varepsilon)$.

Think of ε as inverse polynomial or a negligible function in $n = \log(|\mathcal{U}|)$. The above lemmas show that up to negligible statistical difference and a slightly super-logarithmic number of entropy bits, the min-entropy and the max-entropy are captured by a lower and an upper bound on sample-entropy, respectively.

Conditional entropies. We will also be interested in conditional versions of entropy. For jointly distributed random variables (X,Y) and $(x,y) \in \mathrm{Supp}(X,Y)$, we define the *conditional sample-entropy* to be $\mathrm{H}_{X|Y}(x \mid y) = \log(1/\Pr[X = x \mid Y = y])$. Then the standard *conditional Shannon entropy* can be written as:

$$\mathrm{H}(X \mid Y) = \mathop{\mathrm{E}}_{(x,y)\leftarrow(X,Y)}\left[\mathrm{H}_{X|Y}(x \mid y)\right] = \mathop{\mathrm{E}}_{y\leftarrow Y}\left[\mathrm{H}(X|_{Y=y})\right] = \mathrm{H}(X,Y) - \mathrm{H}(Y).$$

There is no standard definition of conditional min-entropy and max-entropy, or even their smoothed versions. For us, it will be most convenient to generalize the sample-entropy characterizations of smoothed min-entropy and max-entropy given above. Specifically we will think of X as having "high min-entropy" and "low max-entropy" given Y if with probability at least $1-\varepsilon$ over $(x,y) \leftarrow (X,Y)$, we have $\mathrm{H}_{X|Y}(x \mid y) \geq k$ and $\mathrm{H}_{X|Y}(x \mid y) \leq k$, resp.

Flattening Shannon entropy. The *asymptotic equipartition property* in information theory states that for a random variable $X^t = (X_1, \ldots, X_t)$, whose marginals X_i are independent, with high probability, the sample-entropy $\mathrm{H}_{X^t}(X_1, \ldots, X_t)$ is close to its expectation. In [15] a quantitative bound on this was shown by reducing it to the Hoeffding bound. (One cannot directly apply the Hoeffding bound, because $\mathrm{H}_X(X)$ does not have an upper bound, but one can define a related random variable which does.) We use a different bound here, which was proven in [16]. The bound has the advantage that it is somewhat easier to state, even though the proof is longer. We remark that the bound from [15] would be sufficient for our purposes.

Lemma 2.3.

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

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

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

$$\left|\mathrm{H}_{X^t|Y^t}(x \mid y) - t \cdot \mathrm{H}(X \mid Y)\right| \leq O(\sqrt{t \cdot \log(1/\varepsilon)} \cdot \log(|\mathcal{U}|))$$

The statement follows directly from [16, Thm 2].

2.4 Hashing

A family of functions $F = \{f \colon \{0,1\}^n \to \{0,1\}^m\}$ is *2-universal* if for every $x \neq x' \in \{0,1\}^n$, when we choose $f \leftarrow F$, we have $\Pr[f(x) = f(x')] \leq 1/|\{0,1\}^m|$. F is *t-wise independent* if for all distinct $x_1, \ldots, x_t \in \{0,1\}^n$, when we choose $f \leftarrow F$, the random variables $f(x_1), \ldots, f(x_t)$ are independent and each of them is uniformly distributed over $\{0,1\}^m$.

F is *explicit* if given the description of a function $f \in F$ and $x \in \{0,1\}^n$, the value $f(x)$ can be computed in time $\mathrm{poly}(n,m)$. F is *constructible* if it is explicit and there is a probabilistic polynomial-time algorithm that given $x \in \{0,1\}^n$, and $y \in \{0,1\}^m$, outputs a random $f \leftarrow F$ such that $f(x) = y$.

It is well-known that there are constructible families of t-wise independent functions in which choosing a function $f \leftarrow F$ uses only $t \cdot \max\{n,m\}$ random bits.

3 Inaccessible entropy for inversion problems

In this section we define, following the infomercial description given in the introduction, the real and accessible entropy of the inverse of a function. The inaccessible entropy of the inverse is define as the gap between the two.

3.1 Real entropy

For a function F, we define the *real entropy* of F^{-1} to be the amount of entropy left in the input after revealing the output. We measure the above entropy using Shannon entropy (average case), min-entropy and max-entropy.

Definition 3.1 (real entropy)**.** Let n be a security parameter, and $F \colon \{0,1\}^n \to \{0,1\}^m$ a function. We say that F^{-1} has real Shannon entropy k if

$$\mathrm{H}(X \mid F(X)) = k,$$

where X is uniformly distributed on $\{0,1\}^n$. We say that F^{-1} has real min-entropy at least k if there is a negligible function $\varepsilon = \varepsilon(n)$ such that

$$\Pr_{x \leftarrow X}\left[\mathrm{H}_{X|F(X)}(x \mid F(x)) \geq k\right] \geq 1 - \varepsilon(n).$$

We say that F^{-1} has real max-entropy at most k if there is a negligible function $\varepsilon = \varepsilon(n)$ such that

$$\Pr_{x \leftarrow X}\left[\mathrm{H}_{X|F(X)}(x \mid F(x)) \leq k\right] \geq 1 - \varepsilon(n).$$

It is easy to verify that, ignoring negligible terms, the min-entropy of F^{-1} is at most its Shannon entropy which in turn is at most its max-entropy, where equality holds only if F is regular. We also note that more concrete formulas for the entropies above are:

$$\begin{aligned}\mathrm{H}_{X|F(X)}(x \mid F(x)) &= \log\left|F^{-1}(F(x))\right| \\ \mathrm{H}(X \mid F(X)) &= \mathrm{E}\left[\log\left|F^{-1}(F(X))\right|\right].\end{aligned}$$

As our goal is to construct UOWHFs that are shrinking, achieving high real entropy is a natural intermediate step. Indeed, the amount by which F shrinks is a lower bound on the real entropy of F^{-1}:

Proposition 3.2. If $F\colon \{0,1\}^n \to \{0,1\}^m$, then the real Shannon entropy of F^{-1} is at least $n-m$, and the real min-entropy of F^{-1} is at least $n-m-s$ for any $s = \omega(\log n)$.

Proof. For Shannon entropy, we have

$$\mathrm{H}(X \mid F(X)) \geq \mathrm{H}(X) - \mathrm{H}(F(X)) \geq n - m.$$

For min-entropy, let $S = \{y \in \{0,1\}^m\colon \Pr[f(X) = y] < 2^{-m-s}\}$. Then $\Pr[f(X) \in S] \leq 2^m \cdot 2^{-m-s} = \mathrm{neg}(n)$, and for every x such that $f(x) \notin S$, we have

$$\mathrm{H}_{X|F(X)}(x \mid F(x)) = \log \frac{1}{\Pr[X = x \mid F(X) = f(x)]} = \log \frac{\Pr[f(X) = f(x)]}{\Pr[X = x]} \geq \log \frac{2^{-m-s}}{2^{-n}} = n - m - s.$$

□

3.2 Accessible entropy

We define accessible entropy of F^{-1} using the notion of "collision-finding" algorithm, an algorithm that aims to find a second-pre-image of $F(X)$ with "maximal entropy". The accessible entropy of F will be defined as the entropy of the best *efficient* collision-finding algorithm.

Definition 3.3 (collision finding algorithm)**.** For a function $F\colon \{0,1\}^n \to \{0,1\}^m$, an F-collision-finder is a randomized algorithm A such that for every $x \in \{0,1\}^n$ and coin tosses r for A, we have $\mathsf{A}(x;r) \in F^{-1}(F(x))$.

Note that A is required to *always* produce an input $x' \in \{0,1\}^n$ such that $F(x) = F(x')$. This is a reasonable constraint because A has the option of outputting $x' = x$ if it does not find a true collision. We consider A's goal to be maximizing the entropy of its output $x' = A(x)$, given a random input x.

It is easy to see that If we let A be *computationally unbounded*, then the optimum turns out to equal exactly the real entropy:

Proposition 3.4. Let $F\colon \{0,1\}^n \to \{0,1\}^m$. Then the real Shannon entropy of F^{-1} equals the maximum of $\mathrm{H}(\mathsf{A}(X;R) \mid X)$ over all (computationally unbounded) F-collision finders A, where the random variable X is uniformly distributed in $\{0,1\}^n$ and R is uniformly random coin tosses for A. That is,

$$\mathrm{H}(X \mid F(X)) = \max_{\mathsf{A}} \mathrm{H}(\mathsf{A}(X;R) \mid X),$$

where the maximum is taken over all F-collision finders A.

Proof. The "optimal" F-collision finder A that maximizes $\mathrm{H}(\mathsf{A}(X) \mid X)$ is the algorithm $\widetilde{\mathsf{A}}$ that, on input x, outputs a uniformly random element of $f^{-1}(f(x))$. Then

$$\mathrm{H}(\widetilde{\mathsf{A}}(X;R) \mid X) = \mathrm{E}[\log \left\| f^{-1}(f(X)) \right\|] = \mathrm{H}(X \mid F(X)).$$

□

The notion of *accessible entropy* simply restricts the above to PPT algorithms. We consider both Shanon and max-entropy variants (since we aim to upper bound the accessible entropy, we care not about the min-entropy variant).

Definition 3.5 (accessible entropy)**.** Let n be a security parameter and $F\colon \{0,1\}^n \to \{0,1\}^m$ a function. We say that F^{-1} has accessible Shannon entropy at most k if for every PPT F-collision-finder A, we have

$$\mathrm{H}(\mathsf{A}(X;R) \mid X) \le k$$

for all sufficiently large n, where the random variable X is uniformly distributed on $\{0,1\}^n$ and R is uniformly random coin tosses for A.

We say that F^{-1} has p-accessible max-entropy at most k if for every PPT F-collision-finder A, there exists a family of sets $\{\mathcal{L}(x)\}_{x \in \mathrm{Supp}(X)}$ each of size at most 2^k, such that $x \in \mathcal{L}(x)$ for all $x \in \mathrm{Supp}(X)$, and

$$\Pr[\mathsf{A}(X;R) \in \mathcal{L}(X)] \ge 1 - p$$

for all sufficiently large n, where the random variable X is uniformly distributed on $\{0,1\}^n$ and R is uniformly random coin tosses for A. In addition, if $p = \varepsilon(n)$ for some negligible function $\varepsilon(\cdot)$, then we simply say that F^{-1} has accessible max-entropy at most k.

It is easy to verify that, ignoring negligible terms, the accessible Shannon entropy of F^{-1} is at most its accessible max-entropy, i. e., if the accessible max-entropy of F^{-1} is at most k, then its accessible Shannon entropy is at most k. (We will later, Section 3.2.1, introduce an in-between variant of accessible entropy; larger than Shanon smaller than max.)

The reason that having an upper bound on accessible entropy is useful as an intermediate step towards constructing UOWHFs, is that accessible max-entropy 0 is equivalent to target collision resistance (on random inputs):

Definition 3.6 (q-collision-resistant)**.** Let $F\colon \{0,1\}^n \to \{0,1\}^m$ be a function and $q = q(n) \in [0,1]$. We say that F is q-collision-resistant on random inputs if for every PPT F-collision-finder A,

$$\Pr[\mathsf{A}(X;R) = X] \ge q,$$

for all sufficiently large n, where the random variable X is uniformly distributed on $\{0,1\}^n$ and R is uniformly random coin tosses for A. In addition, if $q = 1-\varepsilon(n)$ for some negligible function $\varepsilon(\cdot)$, we say that F is collision-resistant on random inputs.

Lemma 3.7. Let n be a security parameter and $F\colon \{0,1\}^n \to \{0,1\}^m$ be a function. Then, for any $p = p(n) \in (0,1)$, the following statements are equivalent:

(1) F^{-1} has p-accessible max-entropy 0.

(2) F is $(1-p)$-collision-resistant on random inputs.

In particular, F^{-1} has accessible max-entropy 0 iff F is collision-resistant on random inputs.

Proof. Note that (1) implies (2) follows readily from the definition. To see that (2) implies (1), simply take $\mathcal{L}(x) = \{x\}$. □

While bounding p-accessible max-entropy with negligible p is our ultimate goal, one of our constructions will work by first giving a bound on accessible Shannon entropy, and then deducing a bound on p-accessible max-entropy for a value of $p < 1$ using the following lemma:

Lemma 3.8. Let n be a security parameter and $F\colon \{0,1\}^n \to \{0,1\}^m$ be a function. If F^{-1} has accessible Shannon entropy at most k, then F^{-1} has p-accessible max-entropy at most $k/p + O(2^{-k/p})$ for any $p = p(n) \in (0,1)$.

Proof. Fix any PPT F-collision-finder A. From the bound on accessible Shannon entropy, we have that $\mathrm{H}(\mathsf{A}(X;R) \mid X) \le k$. Applying Markov's inequality, we have

$$\Pr_{x\leftarrow X, r\leftarrow R}\left[\mathrm{H}_{\mathsf{A}(X;R)|X}(\mathsf{A}(x;r) \mid x) \le k/p\right] \ge 1-p$$

Take $\mathcal{L}(x)$ to be the set:

$$\mathcal{L}(x) = \{x\} \cup \left\{x'\colon\ \mathrm{H}_{\mathsf{A}(X;R)|X}(x' \mid x) \le k/p\right\}$$

We may rewrite $\mathcal{L}(x)$ as $\{x\} \cup \left\{x'\colon\ \Pr_r[\mathsf{A}(x;r) = x'] \ge 2^{-k/p}\right\}$. It is easy to see that $|\mathcal{L}(x)| \le 2^{k/p}+1$ and thus F^{-1} has p-accessible max-entropy at most $k/p + O(2^{-k/p})$. □

Once we have a bound on p-accessible max-entropy for some $p < 1$, we need to apply several transformations to obtain a function with a good bound on $\mathrm{neg}(n)$-accessible max-entropy.

3.2.1 Accessible average max-entropy

Our second construction (which achieves better parameters), starts with a bound on a different average-case form of accessible entropy, which is stronger than bounding the accessible Shannon entropy. The benefit of this notion it that it can be converted more efficiently to $\mathrm{neg}(n)$-accessible max-entropy, by simply taking repetitions.

To motivate the definition, recall that a bound on accessible Shannon entropy means that the sample entropy $\mathrm{H}_{\mathsf{A}(X;R)|X}(x' \mid x)$ is small on average over $x \leftarrow X$ and $x' \leftarrow \mathsf{A}(x;R)$. This sample entropy may depend on both the input x and the x' output by the adversary (which in turn may depend on its coin tosses). A stronger requirement is to say that we have upper bounds $k(x)$ on the sample entropy that depend *only on* x. The following definition captures this idea, thinking of $k(x) = \log|\mathcal{L}(x)|$. (We work with sets rather than sample entropy to avoid paying the $\log(1/\varepsilon)$ loss in Lemma 2.2.)

Definition 3.9 (accessible average max-entropy)**.** Let n be a security parameter and $F\colon \{0,1\}^n \to \{0,1\}^m$ a function. We say that F^{-1} has accessible average max-entropy at most k if for every PPT F-collision-finder A, there exists a family of sets $\{\mathcal{L}(x)\}_{x\in\mathrm{Supp}(X)}$ and a negligible function $\varepsilon = \varepsilon(n)$ such that $x \in \mathcal{L}(x)$ for all $x \in \mathrm{Supp}(X)$, $\mathrm{E}[\log|\mathcal{L}(X)|] \leq k$ and

$$\Pr[\mathsf{A}(X;R) \in \mathcal{L}(X)] \geq 1-\varepsilon(n),$$

for all sufficiently large n, where the random variable X is uniformly distributed on $\{0,1\}^n$ and R is uniformly random coin tosses for A.

It is easy to verify that, ignoring negligible terms, the accessible average max-entropy of F^{-1} is at least its accessible Shannon entropy and at most its accessible max-entropy.

4 Inaccessible entropy from one-way functions

We present two constructions of inaccessible entropy functions from one-way functions. The one in Section 4.1 is extremely simple and merely trims the one-way function output. The one in Section 4.2 is somewhat more complicated (in the spirit of the first step of Rompel [26], thought still significantly simpler) that yields a more efficient overall construction.

4.1 A direct construction

The goal of this section is to prove the following theorem:

Theorem 4.1. *Let $f : \{0,1\}^n \to \{0,1\}^n$ be a one-way function and define F over $\{0,1\}^n \times [n]$ as $F(x,i) = f(x)_{1,\dots,i-1}$. Then, F^{-1} has accessible Shannon entropy at most $\mathrm{H}(Z \mid F(Z)) - \frac{1}{64n^2}$, where $Z = (X,I)$ is uniformly distributed over $\{0,1\}^n \times [n]$.*

We do not know whether the function F^{-1} has even less accessible Shannon entropy (say, with a gap of $\Omega(1/n)$). However, it seems that a significantly stronger bound would require much more effort, and even improving the bound to $\Omega(1/n)$ does not seem to yield an overall construction which is as efficient as the one resulting from Section 4.2. Therefore we aim to present a proof which is as simple as possible.

We begin with a high-level overview of our approach. Recall from Proposition 3.4 the "optimal" F-collision-finder $\widetilde{\mathsf{A}}$ that computes $F^{-1}(F(\cdot))$. The proof basically proceeds in three steps:

1. First, we show that it is easy to invert f using $\widetilde{\mathsf{A}}$ (Lemma 4.2).

2. Next, we show that if a F-collision-finder A has high accessible Shannon entropy, then it must behave very similarly to $\widetilde{\mathsf{A}}$ (Lemma 4.3).

3. Finally, we show that if A behaves very similarly to $\widetilde{\mathsf{A}}$, then it is also easy to invert f using A (Lemma 4.4).

We may then deduce that if f is one-way, any F-collision-finder A must have accessible Shannon entropy bounded away from $\mathrm{H}(Z \mid F(Z))$.

Step 1. Suppose we have an optimal collision finder $\widetilde{\mathsf{A}}(x,i;r)$ that outputs a uniform random element from $F^{-1}(F(x,i))$. In order to invert an element y, we repeat the following process: start with an arbitrary element $x^{(0)}$ and use $\widetilde{\mathsf{A}}$ to find an element $x^{(1)}$ such that $f(x^{(1)})$ has the same first bit as y. In the i'th step find $x^{(i)}$ such that the first i bits of $f(x^{(i)})$ equal $y_{1,\dots,i}$ (until $i=n$).

This is done more formally in the following algorithm for an arbitrary oracle CF which we set to $\widetilde{\mathsf{A}}$ in the first lemma we prove. The algorithm $\mathsf{ExtendOne}$ does a single step. Besides the new symbol x' which we are interested in, $\mathsf{ExtendOne}$ also returns the number of calls which it did to the oracle. This is completely uninteresting to the overall algorithm, but we use it later in the analysis when we bound the number of oracle queries made by $\mathsf{ExtendOne}$.

Algorithm $\mathsf{ExtendOne}$

Oracle: An F-collision finder CF.
Input: $x \in \{0,1\}^n$, $b \in \{0,1\}$, $i \in [n]$.

$j := 0$
repeat
    $x' := \mathsf{CF}(x,i)$
    $j := j+1$
until $f(x')_i = b$
return (x', j)

Inverter Inv

Oracle: An F-collision finder CF.
Input: $y \in \{0,1\}^n$

$x^{(0)} \leftarrow U_n$
for $i = 1$ **to** n **do**:
    $(x^{(i)}, j) := \mathsf{ExtendOne}^{\mathsf{CF}}(x^{(i-1)}, y_i, i)$
done
return $x^{(n)}$

We first show that with our optimal collision finder $\widetilde{\mathsf{A}}$, the inverter inverts with only $2n$ calls in expectation (even though it can happen that it runs forever). Towards proving that, we define $p(b \mid y_{1,\dots,i-1})$ as the probability that the i'th bit of $f(x)$ equals b, conditioned on the event that $f(x)_{1,\dots,i-1} = y_{1,\dots,i-1}$ (or 0 if $f(x)_{1,\dots,i-1} = y_{1,\dots,i-1}$ is impossible).

Lemma 4.2. The expected number of calls to $\widetilde{\mathsf{A}}$ in a random execution of $\mathsf{Inv}^{\widetilde{\mathsf{A}}}(y = f(x))$ with $x \leftarrow \{0,1\}^n$, is at most $2n$.

Proof. Fix some string $y_{1,\dots,i-1}$ in the image of F. We want to study the expected number of calls to $\widetilde{\mathsf{A}}(x^{(i-1)}, i)$ in case $F(x^{(i-1)}, i) = y_{1,\dots,i-1}$.

If we would know y_i, then this expected number of calls would be $\frac{1}{p(y_i|y_{1,\dots,i-1})}$. Since $y_i = 0$ with probability $p(0 \mid y_{1,\dots,i-1})$ we get that the expected number of calls is 1 if either of the probabilities is 0 and

$$p(0 \mid y_{1,\dots,i-1}) \cdot \frac{1}{p(0 \mid y_{1,\dots,i-1})} + p(1 \mid y_{1,\dots,i-1}) \cdot \frac{1}{1 \mid p(y_{1,\dots,i-1})} = 2$$

otherwise. Using linearity of expectation we get the result. $\square$

Step 2. Given an F-collision finder A, we define $\varepsilon(x,i)$ to be the statistical distance of the distribution of $\mathsf{A}(x,i;r)$ and the the output distribution of $\widetilde{\mathsf{A}}(x,i;r)$ (which equals the uniform distribution over $F^{-1}(F(x,i))$).

We want to show that if A has high accessible Shannon entropy, then A behaves very similarly to $\widetilde{\mathsf{A}}$. The next lemma formalizes this by stating that $\varepsilon(x,i)$ is small on average (over the uniform random choice of $x \in \{0,1\}^n$ and $i \in [n]$). The lemma follows by applying Jensen's inequality on the well known relationship between entropy gap and statistical distance.

Lemma 4.3. Assume $\mathrm{H}(\mathsf{A}(Z)) \geq \mathrm{H}(Z \mid F(Z)) - \frac{1}{64n^2}$, then $\mathrm{E}_{i \leftarrow [n], x \leftarrow \{0,1\}^n}[\varepsilon(x,i)] \leq \frac{1}{8n}$.

Proof.

$$\begin{aligned}
\left\| (Z, \widetilde{\mathsf{A}}(Z)) - (Z, \mathsf{A}(Z)) \right\| &= \mathop{\mathrm{E}}_{z \leftarrow Z}\left[\left|(z, \widetilde{\mathsf{A}}(z)) - (z, \mathsf{A}(z))\right|\right] \\
&\leq \mathop{\mathrm{E}}_{z \leftarrow Z}\left[\sqrt{\mathrm{H}(\widetilde{\mathsf{A}}(z)) - \mathrm{H}(\mathsf{A}(z))}\right] \\
&\leq \sqrt{\mathop{\mathrm{E}}_{z \leftarrow Z}[\mathrm{H}(\widetilde{\mathsf{A}}(z)) - \mathrm{H}(\mathsf{A}(z))]} \\
&\leq \frac{1}{8n}.
\end{aligned}$$

The first inequality uses the fact that if W is a random variable whose support is contained in a set S and U is the uniform distribution on S, then $\|U - W\| \leq \sqrt{\mathrm{H}(U) - \mathrm{H}(W)}$ (see [7, Lemma 11.6.1]). The second inequality follows by Jensen's inequality. The final inequality uses $\mathrm{H}(\widetilde{\mathsf{A}}(Z)) = \mathrm{H}(Z \mid F(Z))$ (Proposition 3.4). $\square$

Step 3. We have seen now that $\mathsf{Inv}^{\widetilde{\mathsf{A}}}$ inverts f with $2n$ calls in expectation and that A behaves similarly to $\widetilde{\mathsf{A}}$. We now want to show that $\mathsf{Inv}^{\mathsf{A}}$ also inverts f efficiently. The main technical difficulty is that even though $\mathsf{Inv}^{\widetilde{\mathsf{A}}}$ makes $2n$ calls to $\widetilde{\mathsf{A}}$ in expectation and A and $\widetilde{\mathsf{A}}$ are close in statistical distance, we cannot immediately deduce an upper bound on the number of calls $\mathsf{Inv}^{\mathsf{A}}$ makes to A. Indeed, our analysis below exploits the fact that Inv and $\mathsf{ExtendOne}$ have a fairly specific structure.

We will assume without loss of generality that

$$\Pr_R[\mathsf{A}(x,i;R) = \widetilde{\mathsf{A}}(x,i;R)] = 1 - \varepsilon(x,i),$$

where $\widetilde{\mathsf{A}}$ is an optimal collision finder as above. This follows from a standard coupling argument since we do not require $\widetilde{\mathsf{A}}$ to be polynomial time, and also because we can extend the number of random bits A uses (we assume it just ignores unused ones). To do this, $\widetilde{\mathsf{A}}$ first computes the statistics of A on input (x,i), and also the result of $\mathsf{A}(x,i;r)$. He checks whether $\mathsf{A}(x,i;r)$ is one of the elements which occur too often, and outputs a different, carefully chosen one, with appropriate probability if this is the case.

We now show that in most executions of $\mathsf{ExtendOne}$ it does not matter whether we use A or $\widetilde{\mathsf{A}}$ (that is, $\mathsf{ExtendOne}$ makes the same number of oracle queries to A and $\widetilde{\mathsf{A}}$, and outputs the same value).

Lemma 4.4. For any $(x,i) \in \{0,1\}^n \times [n]$, and any $y_i \in \{0,1\}$, we have

$$\Pr_R[\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R) = \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R)] \geq 1 - \frac{2\varepsilon(x,i)}{p(y_i \mid y_{1,\ldots,i-1})} \tag{4.1}$$

Note that the oracle algorithm $\mathsf{ExtendOne}$ is deterministic, and in the above expressions, R refers to the coin tosses used by the oracles that $\mathsf{ExtendOne}$ queries, namely A and $\widetilde{\mathsf{A}}$. We stress that the lemma says that both the value x' *and* the number j returned are equal with high probability.

Proof. Let $J = J(R)$ be the second coordinate of the output of $\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R)$ (i. e., the counter) and $\widetilde{J} = \widetilde{J}(R)$ the analogous output of $\mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R)$. We write

$$\begin{aligned}
&\Pr_R[\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R) = \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R)] = \\
&\sum_{j\geq 1} \Pr_R[\min(J,\widetilde{J}) = j] \cdot \Pr_R[\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R) = \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R) | \min(J,\widetilde{J}) = j] \\
&= \mathop{\mathrm{E}}_{j\leftarrow P_J}\big[\Pr_R[\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R) = \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R) | \min(J,\widetilde{J}) = j]\big] \qquad (4.2)
\end{aligned}$$

where P_J is some distribution over the integers which, as it turns out, we do not need to know.

Let now R' be the randomness used by A or $\widetilde{\mathsf{A}}$ in round j. Then,

$$\begin{aligned}
&\Pr_R[\mathsf{ExtendOne}^{\mathsf{A}}(x,y_i,i;R) = \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x,y_i,i;R) | \min(J,\widetilde{J}) = j] \\
&= \Pr_{R'}\big[\mathsf{A}(x,i;R') = \widetilde{\mathsf{A}}(x,i;R') \mid f(\mathsf{A}(x,i;R'))_i = y_i \vee f(\widetilde{\mathsf{A}}(x,i;R'))_i = y_i\big],
\end{aligned}$$

because each iteration of $\mathsf{ExtendOne}$ uses fresh independent randomness.

Let P be the distribution over $\{0,1\}^n$ produced by $\mathsf{A}(x,i;R')$, and P^* be the (uniform) distribution produced by $\widetilde{\mathsf{A}}(x,i;R')$. For $p = p(y_i \mid y_{1,\dots,i-1})$ and $\varepsilon = \varepsilon(x,i)$, it holds that

$$
\begin{aligned}
&\Pr_{R'}\left[\mathsf{A}(x,i;R') = \widetilde{\mathsf{A}}(x,i;R') \mid f(\mathsf{A}(x,i;R'))_i = y_i \vee f(\widetilde{\mathsf{A}}(x,i;R'))_i = y_i\right] \\
&\quad = \frac{\Pr\left[\mathsf{A}(x,i;R') = \widetilde{\mathsf{A}}(x,i;R') \wedge (f(\mathsf{A}(x,i;R'))_i = y_i \vee f(\widetilde{\mathsf{A}}(x,i;R'))_i = y_i)\right]}{\Pr\left[f(\mathsf{A}(x,i;R'))_i = y_i \vee f(\widetilde{\mathsf{A}}(x,i;R'))_i = y_i\right]} \\
&\quad = \frac{\sum_{x' \in F^{-1}(y_{1,\dots,i})} \Pr\left[\mathsf{A}(x,i;R') = \widetilde{\mathsf{A}}(x,i;R') = x' \wedge (\mathsf{A}(x,i;R') = x' \vee \widetilde{\mathsf{A}}(x,i;R') = x')\right]}{\sum_{x' \in F^{-1}(y_{1,\dots,i})} \Pr\left[\mathsf{A}(x,i;R') = x' \vee \widetilde{\mathsf{A}}(x,i;R') = x'\right]} \\
&\quad = \frac{\sum_{x' \in F^{-1}(y_{1,\dots,i})} \min(P(x'), P^*(x'))}{\sum_{x \in F^{-1}(y_{1,\dots,i})} \max(P(x'), P^*(x'))} = \frac{p-\varepsilon}{p+\varepsilon} = \frac{1-\frac{\varepsilon}{p}}{1+\frac{\varepsilon}{p}} \geq 1 - \frac{2\varepsilon}{p}\,.
\end{aligned}
$$

For the penultimate equality, note that

$$
\varepsilon = \mathsf{SD}(\mathsf{A}(x,i), \widetilde{\mathsf{A}}(x,i)) = \sum_{x' \in F^{-1}(y_{1,\dots,i})} P^*(x') - \min(P(x'), P^*(x')).
$$

Hence,

$$
\begin{aligned}
\sum_{x' \in F^{-1}(y_{1,\dots,i})} \min(P(x'), P^*(x')) &= \sum_{x' \in F^{-1}(y_{1,\dots,i})} \min(P(x'), P^*(x')) + P^*(x') - P^*(x') \\
&= \sum_{x' \in F^{-1}(y_{1,\dots,i})} P^*(x') - \left(\sum_{x' \in F^{-1}(y_{1,\dots,i})} P^*(x') - \min(P(x'), P^*(x'))\right) \\
&= p - \varepsilon.
\end{aligned}
$$

And similarly, $\sum_{x' \in F^{-1}(y_{1,\dots,i})} \max(P(x'), P^*(x')) = p + \varepsilon$.

Collecting the equations and inserting into (4.2) proves the lemma. □

Putting everything together. We can now finish the proof of Theorem 4.1. Consider the following random variables: let X be uniformly drawn from $\{0,1\}^n$ and let $Y = f(X)$. Run $\mathsf{Inv}^{\mathsf{A}}(Y)$ and $\mathsf{Inv}^{\widetilde{\mathsf{A}}}(Y)$ in parallel, using the same randomness in both executions. Let $\widetilde{X}^{(0)}, \dots, \widetilde{X}^{(n)}$ be the random variables which have the values assigned to $x^{(0)}, \dots, x^{(n)}$ in the run of $\mathsf{Inv}^{\widetilde{\mathsf{A}}}$. Finally, let the indicator variables Q_i be 1, iff the i'th call to $\mathsf{ExtendOne}$ in the above parallel run is the *first* call such that $\mathsf{ExtendOne}^{\mathsf{A}}(\widetilde{X}^{(i)}, Y_i, i; \cdot) \neq \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(\widetilde{X}^{(i)}, Y_i, i; \cdot)$.

We proceed to obtain an upper bound on $\Pr[Q_i = 1]$. Observe that for all $x \in \{0,1\}^n$:

$$\begin{aligned}
&\Pr[Q_i = 1 \mid \widetilde{X}^{(i-1)} = x] \\
&= \Pr[\mathsf{ExtendOne}^{\mathsf{A}}(x, Y_i, i; R) \neq \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x, Y_i, i; R) \mid \widetilde{X}^{(i-1)} = x] \\
&= \sum_{y_i \in \{0,1\}} p(y_i \mid f(x)_{1,\ldots,i-1}) \cdot \Pr[\mathsf{ExtendOne}^{\mathsf{A}}(x, y_i, i; R) \neq \mathsf{ExtendOne}^{\widetilde{\mathsf{A}}}(x, y_i, i; R)] \\
&\leq \sum_{y_i \in \{0,1\}} p(y_i \mid f(x)_{1,\ldots,i-1}) \cdot \frac{2\varepsilon(x,i)}{p(y_i \mid f(x)_{1,\ldots,i-1})} \\
&= 4\varepsilon(x,i)
\end{aligned}$$

where the inequality above follows by Lemma 4.4. Averaging over x, we have that for all $i = 1, \ldots, n$:

$$\Pr[Q_i = 1] \leq 4 \mathop{\mathrm{E}}_{x \leftarrow \{0,1\}^n}[\varepsilon(x,i)]. \tag{4.3}$$

Here, we use the fact that by induction on i, the random variable $\widetilde{X}^i$, for $i \in \{0, \ldots, n\}$, is uniformly distributed in $\{0,1\}^n$ (it is uniform preimage of a uniformly chosen output). Using Equation (4.3), we have

$$\begin{aligned}
\Pr\left[\sum_{i=1}^{n} Q_i \geq 1\right] &= \sum_{i=1}^{n} \Pr[Q_i = 1] \\
&= n \cdot \mathop{\mathrm{E}}_{i \leftarrow [n], x \leftarrow \{0,1\}^n}[Q_i] \\
&\leq 4n \cdot \mathop{\mathrm{E}}_{i \leftarrow [n], x \leftarrow \{0,1\}^n}[\varepsilon(x,i)] \\
&\leq \frac{1}{2}
\end{aligned}$$

where the last inequality follows from Lemma 4.3. Hence, with probability $1/2$, a run of $\mathsf{Inv}^{\mathsf{A}}$ and $\mathsf{Inv}^{\widetilde{\mathsf{A}}}$ produce the same output and use the same number of queries to the oracles A. Moreover, the probability that $\mathsf{Inv}^{\widetilde{\mathsf{A}}}$ uses more than $8n$ oracle queries is at most $1/4$ (by applying Markov's inequality on Lemma 4.2). Hence, with probability $1/4$, $\mathsf{Inv}^{\mathsf{A}}$ inverts f using $8n$ oracle queries in total, which contradicts the one-wayness of f. In order to make sure that $\mathsf{Inv}^{\mathsf{A}}$ runs in polynomial time, we just halt it after $8n$ calls.

4.2 A more efficient construction

The following theorem shows that a simplified variant of the first step of [26] (which is also the first step of [20]) yields inaccessible entropy with much stronger guarantees than those obtained in Section 4.1. The function we construct is $F(x,g,i) = (g(f(x))_{1,\ldots,i}, g)$, where $g : \{0,1\}^n \to \{0,1\}^n$ is a three-wise independent function. Since the composition of g and f is still a one-way function, Theorem 4.1 already implies that F^{-1} has inaccessible entropy. The benefits of the additional hashing step are that

1. we get more inaccessible entropy ($\tilde{\Theta}(1/n)$ bits rather than $\tilde{\Theta}(1/n^2)$ bits), and

2. we get a bound on accessible average max-entropy rather than accessible Shannon entropy.

These allow for a more efficient and simpler transformation of F into a UOWHF.

Theorem 4.5 (Inaccessible average max-entropy from one-way functions)**.** *Let $f\colon \{0,1\}^n \to \{0,1\}^n$ be a one-way function and let $\mathcal{G} = \{g\colon \{0,1\}^n \to \{0,1\}^n\}$ be a family of constructible, three-wise independent hash functions. Define F over* $\mathrm{Dom}(F) := \{0,1\}^n \times \mathcal{G} \times [n]$ *by*

$$F(x,g,i) = (g(f(x))_{1,\dots,i}, g, i).$$

Then, for every constant $d > 0$, F^{-1} has accessible average max-entropy at most $\mathrm{H}(Z \mid F(Z)) - (d \log n)/n$*, where Z is uniformly distributed over* $\mathrm{Dom}(F)$.

Proof. Let c be a sufficiently large constant (whose value to be determined later as a function of the constant d in the theorem statement). The sets $\{\mathcal{L}(x,g,i)\}_{x\in\{0,1\}^n, i\in[n], g\in\mathcal{G}}$ realizing the inaccessible entropy of F^{-1} are defined by

$$\mathcal{L}(x,g,i) = \{(x',g,i)\colon f(x') \in \widetilde{\mathcal{L}}(f(x),i) \land g(f(x'))_{1,\dots,i} = g(f(x))_{1,\dots,i}\} \tag{4.4}$$

where for $y \in \{0,1\}^n$ and $i \in [n]$, we let

$$\begin{aligned} \widetilde{\mathcal{L}}(y,i) &= \{y\} \cup \left\{y' \in \{0,1\}^n\colon \mathrm{H}_{f(X)}(y') \geq (i + c\cdot\log n)\right\} \\ &= \{y\} \cup \left\{y' \in \{0,1\}^n\colon |f^{-1}(y')| \leq 2^{n-i}/n^c\right\}. \end{aligned} \tag{4.5}$$

Namely, $\widetilde{\mathcal{L}}(y,i)$ consists, in addition to y itself, of "i-light" images with respect to f.[7] As a warm-up, it is helpful to write down $\widetilde{\mathcal{L}}(y,i)$ and $\mathcal{L}(x,g,i)$ for the case where f is a one-way permutation.[8]

The proof of the theorem immediately follows by the following two claims.

Claim 4.6. For every PPT F-collision-finder A and every constant $c > 0$, it holds that

$$\Pr[\mathsf{A}(Z;R) \notin \mathcal{L}(Z)] \leq \mathrm{neg}(n),$$

where Z is uniformly distributed over $\mathrm{Dom}(F)$ and R is uniformly distributed over the random coins of A.

Claim 4.7. For any constant c it holds that

$$\mathrm{E}\left[\log|\mathcal{L}(Z)|\right] \leq \mathrm{E}\left[\log\left|F^{-1}(F(Z))\right|\right] - \Omega\left(\frac{c\log n}{n}\right),$$

where Z is uniformly distributed in $\mathrm{Dom}(F)$.

□

[7]Recall that the sample entropy is defined as $\mathrm{H}_{f(X)}(y) = \log(1/\Pr[f(X)=y]) = n - \log\left|f^{-1}(y)\right|$, so the "heavy" images, where $f^{-1}(y)$ is large, have low sample entropy.

[8]If f is a permutation, then $\widetilde{\mathcal{L}}(y,i)$ is given by:

$$\widetilde{\mathcal{L}}(y,i) = \begin{cases} \{0,1\}^n & \text{if } i \leq n - c\log n \\ \{y\} & \text{otherwise.} \end{cases}$$

Then, for all $x \in \{0,1\}^n$, we have $\mathrm{E}[|\mathcal{L}(x,G,i)|] = 2^{n-i}$ for all $i \leq n - c\log n$ and $|\mathcal{L}(x,g,i)| = 1$ for all $g \in \mathcal{G}$ and all $i > n - c\log n$. This means that the entropy gap between $F^{-1}(F(Z))$ and $\mathcal{L}(X,G,I)$ is roughly $\frac{1}{n}\sum_{i>n-c\log n} n - i = \Omega(c^2\log^2 n/n)$.

4.2.1 Accessible inputs of F — Proving Claim 4.6

Proof of Claim 4.6. Recall that $Z = (X, G, I)$ is uniformly distributed over $\mathrm{Dom}(F)$, and that R is uniformly distributed over the random coins of A. Let A_1 denote the first component of A's output. It suffices to show that

$$\Pr[\mathsf{A}_1(X,G,I;R) \notin f^{-1}(\widetilde{\mathcal{L}}(f(X),I))] \leq \mathrm{neg}(n) \tag{4.6}$$

since the other two output components of A are required to equal (G,I), due to the fact that $F(X,G,I)$ determines (G,I).

We construct an inverter Inv such that for all F-collision-finders A and for c as in Equation (4.5) we have

$$\Pr[\mathsf{Inv}^{\mathsf{A}}(Y) \in f^{-1}(Y)] \geq \frac{1}{n^c} \cdot \Pr[\mathsf{A}_1(X,G,I;R) \notin f^{-1}(\widetilde{\mathcal{L}}(f(X),I))] \tag{4.7}$$

where $Y = f(X)$, and the proof of Claim 4.6 follows readily from the one-wayness of f.

Inverter $\mathsf{Inv}^{\mathsf{A}}$
Oracle: An F-collision finder A. **Input:** $y \in \{0,1\}^n$
$x \leftarrow \{0,1\}^n$ $i \leftarrow [n]$ $g' \leftarrow \mathcal{G}_{y,x,i} := \{g \in \mathcal{G} \colon g(y)_{1\ldots i} = g(f(x))_{1\ldots i}\}$ **return** $A_1(x,g',i;r)$

Observe that Inv can be implemented efficiently by sampling g' as follows: pick first $z, z^* \in \{0,1\}^n$ such that $z_{1\ldots i} = z^*_{1\ldots i}$ and use the constructibility of $\mathcal{G}$ to pick g with $g(f(x)) = z$ and $g(y) = z^*$.

We analyze the success probability of $\mathsf{Inv}^{\mathsf{A}}$. Using the short hand notation $\Pr_{g'}[\cdots]$ for $\Pr_{g' \leftarrow \mathcal{G}_{y,x,i}}[\cdots]$ we observe that

$$\begin{aligned}\Pr[\mathsf{Inv}^{\mathsf{A}}(Y) \in f^{-1}(Y)] &= \mathop{\mathrm{E}}_{x \leftarrow \{0,1\}^n, i \leftarrow [n]}\left[\sum\nolimits_{y \in \{0,1\}^n} \Pr[f(X) = y] \cdot \Pr_{g',r}[\mathsf{A}_1(x,g',i;r) \in f^{-1}(y)]\right] \\ &\geq \mathop{\mathrm{E}}_{x,i}\left[\sum\nolimits_{y \notin \widetilde{\mathcal{L}}(f(x),i)} \frac{2^{-i}}{n^c} \cdot \Pr_{g',r}[\mathsf{A}_1(x,g',i;r) \in f^{-1}(y)]\right]\end{aligned} \tag{4.8}$$

where the inequality holds since $\Pr[f(X) = y] \geq 2^{-i}/n^c$ for any $y \notin \widetilde{\mathcal{L}}(f(x),i)$.

Next, observe that for any tuple (y,x,i) such that $y \neq f(x)$, it holds that (where we distinguish $\Pr_{g'}[\cdots]$

as above from $\Pr_g[\cdots] = \Pr_{g \leftarrow \mathcal{G}}[\cdots]$)

$$\begin{aligned}\Pr_{g',r}\big[\mathsf{A}_1(x,g',i;r) \in f^{-1}(y)\big] &= \Pr_{g \leftarrow \mathcal{G},r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y) \mid g(f(x))_{1\cdots i} = g(y)_{1\cdots i}\big] \\ &= \frac{\Pr_{g,r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y) \wedge g(f(x))_{1\cdots i} = g(y)_{1\cdots i}\big]}{\Pr_{g,r}\big[g(f(x))_{1\cdots i} = g(y)_{1\cdots i}\big]} \\ &= \frac{\Pr_{g,r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y)\big]}{\Pr_{g,r}\big[g(f(x))_{1\cdots i} = g(y)_{1\cdots i}\big]} \\ &= 2^i \cdot \Pr_{g,r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y)\big].\end{aligned} \tag{4.9}$$

The second equality follows by Bayes' rule and the third uses the fact that A is a F-collision finder. The last equality follows since $\mathcal{G}$ is two-wise independent (recall we assumed that $\mathcal{G}$ is three-wise independent) and $f(x) \neq y$.

Combining the two preceding observations, and the fact that $f(x) \in \widetilde{\mathcal{L}}(f(x),i)$, we have that

$$\begin{aligned}\Pr[\mathsf{Inv}^{\mathsf{A}}(Y) \in f^{-1}(Y)] &\geq \mathop{\mathrm{E}}_{x \leftarrow \{0,1\}^n, i \leftarrow [n]}\Big[\textstyle\sum_{y \notin \widetilde{\mathcal{L}}(f(x),i)} \frac{2^{-i}}{n^c} \cdot 2^i \cdot \Pr_{g \leftarrow \mathcal{G},r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y)\big]\Big] \\ &\geq \frac{1}{n^c} \cdot \mathop{\mathrm{E}}_{x,i}\Big[\textstyle\sum_{y \notin \widetilde{\mathcal{L}}(f(x),i)} \cdot \Pr_{g,r}\big[\mathsf{A}_1(x,g,i;r) \in f^{-1}(y)\big]\Big] \\ &= \frac{1}{n^c} \cdot \Pr_{x,g,i,r}\big[\mathsf{A}_1(x,g,i;r) \notin f^{-1}(\widetilde{\mathcal{L}}(f(x),i))\big],\end{aligned}$$

and the proof of the claim follows. □

4.2.2 Upper bounding the size of $\mathcal{L}$ — Proving Claim 4.7

Recall that $Z = (X,G,I)$ is uniformly distributed over $\mathrm{Dom}(F)$. In the following we relate the size of $\mathcal{L}(Z)$ to that of $F^{-1}(F(Z))$.

We make use of the following property of *three*-wise independent hash-functions.

Claim 4.8. Let i, x and x^*, be such that $f(x) \neq f(x^*)$ and $i \leq \mathrm{H}_{f(X)}(f(x^*)) \leq \mathrm{H}_{f(X)}(f(x))$. Then,

$$\Pr_{\substack{g \leftarrow \mathcal{G} \\ z' \leftarrow F^{-1}(F(x,g,i))}}\big[z' = (x^*,g,i)\big] \geq \frac{2^{-n}}{8}.$$

Note that in the above experiment it is always the case that $(g',i') = (g,i)$, where $z' = (x',g',i')$.

Proof. Note that with probability 2^{-i} over $g \leftarrow \mathcal{G}$, it holds that $g(f(x^*))_{1\cdots i} = g(f(x))_{1\cdots i}$. Henceforth, we condition on this event that we denote by E, and let $w = g(f(x^*))_{1\cdots i} = g(f(x))_{1\cdots i}$. Observe that for a fixed g satisfying E, it holds that

$$\Pr_{z' \leftarrow F^{-1}(F(x,g,i))}\big[z' = (x^*,g,i) \mid E\big] \geq \frac{1}{|F^{-1}(F(x,g,i))|} \tag{4.10}$$

In order to obtain a lower bound on $\left|F^{-1}(F(x,g,i))\right|$, we first consider x' such that $f(x') \notin \{f(x), f(x^*)\}$. By the three-wise independence of $\mathcal{G}$,

$$\Pr_{g \leftarrow \mathcal{G}}\left[g(f(x')) = w \mid E\right] = 2^{-i} \tag{4.11}$$

This implies that the expected number of x' such that $g(f(x')) = w$ and $f(x') \notin \{f(x), f(x^*)\}$ is at most 2^{n-i}. By Markov's inequality, we have that with probability at least $1/2$ over $g \leftarrow \mathcal{G}$ (conditioned on E),

$$|F^{-1}(F(x,g,i))| \leq 2 \cdot 2^{n-i} + |f^{-1}(f(x))| + |f^{-1}(f(x^*))| \leq 4 \cdot 2^{n-i}, \tag{4.12}$$

where the second inequality uses the fact that $i \leq \mathrm{H}_{f(X)}(f(x^*)) \leq \mathrm{H}_{f(X)}(f(x))$. Putting everything together, we have that the probability we obtain x^* is at least $2^{-i} \cdot 1/2 \cdot (4 \cdot 2^{n-i})^{-1} = 2^{-n}/8$. □

We now use Claim 4.8 for proving Claim 4.7.

Proof of Claim 4.7. Let $Z' = (X', G, I) \leftarrow F^{-1}(F(Z = (X, G, I)))$ (note that indeed the second and third coordinates of Z and Z' are guaranteed to match). We claim that for proving Claim 4.7 it suffices to show that

$$\Pr[X' \notin f^{-1}(\widetilde{\mathcal{L}}(f(X), I))] \in \Omega\left(\frac{c \log(n)}{n}\right). \tag{4.13}$$

Indeed, let $\overline{\mathcal{L}}(z) := F^{-1}(F(z)) \setminus \mathcal{L}(z)$, and compute

$$\begin{aligned}
\mathrm{E}\left[\log\left|F^{-1}(F(Z))\right|\right] - \mathrm{E}\left[\log|\mathcal{L}(Z)|\right] &= \mathrm{E}\left[\log\left(1 + \frac{|\overline{\mathcal{L}}(Z)|}{|\mathcal{L}(Z)|}\right)\right] \\
&\geq \mathrm{E}\left[\log\left(1 + \frac{|\overline{\mathcal{L}}(Z)|}{|F^{-1}(F(Z))|}\right)\right] \\
&\geq \frac{1}{2}\,\mathrm{E}\left[\frac{|\overline{\mathcal{L}}(Z)|}{|F^{-1}(F(Z))|}\right] \\
&= \frac{1}{2}\Pr[X' \notin f^{-1}(\widetilde{\mathcal{L}}(f(X), I))] \\
&\in \Omega\left(\frac{c \log(n)}{n}\right).
\end{aligned} \tag{4.14}$$

The first equality holds since by definition $\mathcal{L}(Z) \subseteq F^{-1}(F(z))$, the second inequality holds since $\log(1 + \alpha) \geq \frac{\alpha}{2}$ for $\alpha \in [0, 1]$ and the containment by Equation (4.13).

We prove Equation (4.13) in two steps. First, observe that for all x:

$$\begin{aligned}
&\Pr_{(g,i)\leftarrow\mathcal{G}\times[n],(x',g,i)\leftarrow F^{-1}(F(x,g,i))}\left[f(x')\notin\widetilde{\mathcal{L}}(f(x),i)\right] && (4.15)\\
&\geq \Pr_{g,i,(x',g,i)\leftarrow F^{-1}(F(x,g,i))}\left[f(x')\neq f(x)\wedge(i\leq \mathrm{H}_{f(X)}(f(x'))<i+c\log n)\right]\\
&=\frac{1}{n}\cdot\sum_{x^*:\,f(x)\neq f(x^*)}\Bigg(\sum_{i\leq \mathrm{H}_{f(X)}(f(x^*))<i+c\log n}\Pr_g\left[(x^*,g,i)\leftarrow F^{-1}(F(x,g,i))\right]\Bigg)\\
&\geq\frac{1}{n}\cdot\sum_{\substack{x^*:\,f(x)\neq f(x^*)\\ \wedge(\mathrm{H}_{f(X)}(f(x^*))\leq \mathrm{H}_{f(X)}(f(x)))}}\Bigg(\sum_{i\leq \mathrm{H}_{f(X)}(f(x^*))<i+c\log n}\frac{2^{-n}}{8}\Bigg) && \text{(by Claim 4.8)}\\
&\geq\frac{c\log n}{n}\cdot\sum_{\substack{x^*:\,f(x)\neq f(x^*)\\ \wedge(\mathrm{H}_{f(X)}(f(x^*))\leq \mathrm{H}_{f(X)}(f(x)))}}\frac{2^{-n}}{8}\\
&=\frac{c\log n}{8n}\cdot\Pr_{x^*}\left[f(x)\neq f(x^*)\wedge(\mathrm{H}_{f(X)}(f(x^*))\leq \mathrm{H}_{f(X)}(f(x)))\right].
\end{aligned}$$

It follows that

$$\begin{aligned}
\Pr[X'\notin f^{-1}(\widetilde{\mathcal{L}}(f(X),I))] &= \Pr_{(x,g,i)\leftarrow\{0,1\}^n\times\mathcal{G}\times[n]}\left[x'\notin\widetilde{\mathcal{L}}(f(x),i)\colon (x',g,i)\leftarrow F^{-1}(F(x,g,i))\right]\\
&\geq\frac{c\log n}{8n}\cdot\Pr_{(x,x^*)\leftarrow\{0,1\}^n\times\{0,1\}^n}\left[f(x)\neq f(x^*)\wedge(\mathrm{H}_{f(X)}(f(x^*))\leq \mathrm{H}_{f(X)}(f(x)))\right]\\
&\geq\frac{c\log n}{8n}\cdot\frac{1}{2}\cdot\Big(1-\Pr_{x,x^*}[f(x)\neq f(x^*)]\Big)\\
&\geq\frac{c\log n}{8n}\cdot\frac{1}{2}\cdot\frac{1}{2}.
\end{aligned}$$

The last inequality holds since the one-wayness of f yields that $\Pr_{x,x^*}[f(x)=f(x^*)]$ is negligible (otherwise inverting f is trivial). This concludes the the proof of Equation (4.13), and hence of the claim. □

5 UOWHFs from inaccessible entropy

In this section we show how to construct a UOWHF from any efficiently computable function with a noticeable gap between real Shannon entropy and either accessible average max-entropy or accessible Shannon entropy. Recall that the more efficient construction from Section 4.2 satisfies the former, and the more direct construction from Section 4.1 satisfies the latter. Combined with these constructions, we obtain two new constructions of UOWHFs from any one-way function.

In both cases, we first transform the entropy gap into a noticeable gap between real Shannon entropy and accessible *max*-entropy. We begin with the construction that starts from a gap between real Shannon entropy and accessible average max-entropy because the transformation involves fewer steps (and is also more efficient).

5.1 The more efficient UOWHF

Theorem 5.1. *Suppose there exists a polynomial-time computable function $F : \{0,1\}^{\widetilde{n}} \to \{0,1\}^m$ such that F^{-1} has a noticeable gap Δ between real Shannon entropy and accessible average max-entropy. Then, there exists a family of universal one-way hash functions with output length $O(\widetilde{n}^4 s/\Delta^3)$ and key length $O(\widetilde{n}^4 s/\Delta^3 \cdot \log n)$ for any $s = \omega(\log n)$, where n is the security parameter.*[9]

We first show how to combine this with Theorem 4.5 to get a universal one-way hash function.

Theorem 5.2. *Suppose there exists a one-way function $f : \{0,1\}^n \to \{0,1\}^n$. Then, there exists a family of universal one-way hash functions with key and output length $O(n^7)$.*

Proof. Fixing $s := \log^2(n)$, we use Theorem 4.5 to get a function $F : \{0,1\}^{\widetilde{n}} \to \{0,1\}^m$ with $\widetilde{n} = O(n)$ and gap $\Delta = \log n/n$ between real Shannon entropy and accessible average max-entropy. By Theorem 5.1 we get a family of universal one-way hash functions with output length $O(\widetilde{n}^4 s/\Delta^3) = O(n^7 \log^2(n)/\log^3(n))$ and key length $O(\widetilde{n}^4 s/\Delta^3 \cdot \log(n)) = O(n^7)$. □

Overview. The construction proceeds via a series of transformations as outlined in Section 1.2: gap amplification (via repetition), entropy reduction (by hashing inputs) and reducing output length (by hashing outputs). In each of these transformations, we use n_0 to denote the input length of the function F we start with, and n to denote the security parameter.

5.1.1 Gap amplification

Here, we show that a direct product construction increases the gap between real entropy and accessible entropy. Another useful effect of direct product (for certain settings of parameters) is turning real Shannon entropy into real min-entropy, and turning accessible average max-entropy into accessible max-entropy.

Lemma 5.3 (Gap amplification)**.** Let n be a security parameter and $F\colon \{0,1\}^{\widetilde{n}} \to \{0,1\}^m$ be a function. For $t \in \mathrm{poly}(n)$, let F^t be the t-fold direct product of F. Then, F^t satisfies the following properties:

i. If F^{-1} has real Shannon entropy at least k, then $(F^t)^{-1}$ has real min-entropy at least $t \cdot k - \widetilde{n} \cdot \sqrt{st}$ for any $s = \omega(\log n)$ and $t > s$.

ii. If F^{-1} has accessible average max-entropy at most k, then $(F^t)^{-1}$ has accessible max-entropy at most $t \cdot k + \widetilde{n} \cdot \sqrt{st}$ for any $s = \omega(\log n)$.

Proof. In the following X and $X^{(t)} = (X_1, \ldots, X_t)$ are uniformly distributed over $\{0,1\}^{\widetilde{n}}$ and $(\{0,1\}^{\widetilde{n}})^t$, respectively.

i. Follows readily from Lemma 2.3 (with $\varepsilon = 2^{-s}$).

ii. Given any PPT F^t-collision-finder A', we construct a PPT F-collision-finder A that:

> On input x, picks a random i in $[t]$ along with random $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_t$, computes $\mathsf{A}'(x_1, \ldots, x_t) \to (x'_1, \ldots, x'_t)$, and outputs x'_i.

[9] Note that Δ is not required to be efficiently computable.

By the bound on the accessible average max-entropy of F^{-1}, we know that there exists a family of sets $\{\mathcal{L}(x)\}$ such that $\mathrm{E}\big[\log|\mathcal{L}(X)|\big] \le k$, $x \in \mathcal{L}(x)$, and $\Pr[\mathsf{A}(X) \notin \mathcal{L}(X)] \le \mathrm{neg}(n)$. Consider the family of sets $\big\{\mathcal{L}'(x^{(t)})\colon x^{(t)} \in (\{0,1\}^{\widetilde{n}})^t\big\}$ given by:

$$\mathcal{L}'(x^{(t)}) = \mathcal{L}(x_1^{(t)}) \times \mathcal{L}(x_2^{(t)}) \times \cdots \times \mathcal{L}(x_t^{(t)}).$$

By linearity of expectations, we have $\mathrm{E}\big[\log|\mathcal{L}'(X_1,\ldots,X_t)|\big] \le t \cdot k$. Moreover, by the Chernoff-Hoeffding bound and using the fact that $\log|\mathcal{L}(X)|$ assumes values in $[0,\widetilde{n}]$, we have

$$\begin{aligned}&\Pr\Big[\log\Big|\mathcal{L}'(X^{(t)})\Big| \ge t\cdot k + \widetilde{n}\sqrt{st}\Big] \\ &= \Pr\Big[\log\Big|\mathcal{L}(X_1^{(t)})\Big| + \cdots + \log\Big|\mathcal{L}(X_t^{(t)})\Big| \ge t\cdot k + \widetilde{n}\sqrt{st}\Big] \le \mathrm{e}^{-2s}.\end{aligned} \tag{5.1}$$

We claim that this implies that A' has accessible max-entropy at most $t\cdot k + \widetilde{n}\sqrt{st}$. Suppose otherwise, then there exists a non-negligible function ε such that

$$\Pr[\mathsf{A}'(F^t(X^{(t)})) \notin \mathcal{L}'(X^{(t)})] \ge \varepsilon - \mathrm{e}^{-2s} \ge \varepsilon/2$$

Therefore,

$$\Pr[\mathsf{A}(F(X)) \notin \mathcal{L}(X)] = \Pr[\mathsf{A}'(F^t(X^{(t)})) \notin \mathcal{L}'(X^{(t)})]/t \ge \varepsilon/2t$$

which contradicts our assumption on A.

□

5.1.2 Entropy reduction

Next we describe a construction that given F and any parameter ℓ, reduces the accessible max-entropy of F^{-1} by roughly ℓ bits, while approximately preserving the gap between real min-entropy and accessible max-entropy.

Lemma 5.4 (Reducing entropy)**.** Let n be a security parameter and $F\colon \{0,1\}^{\widetilde{n}} \to \{0,1\}^m$ be a function. Fix a family of pairwise independent hash functions $\mathcal{G} = \big\{g\colon \{0,1\}^{\widetilde{n}} \to \{0,1\}^{\ell}\big\}$. Then, $F'\colon \{0,1\}^{\widetilde{n}} \times \mathcal{G} \to \{0,1\}^m \times \mathcal{G} \times \{0,1\}^{\ell}$ as given by $F'(x,g) = (F(x),g,g(x))$ satisfies the following properties:

i. Assuming F^{-1} has real min-entropy at least k, then $(F')^{-1}$ has real min-entropy at least $k-\ell-s$ for any $s = \omega(\log n)$.

ii. Assuming F^{-1} has accessible max-entropy at most k, then $(F')^{-1}$ has accessible max-entropy at most $\max\{k-\ell+s,0\}$ for any $s = \omega(\log n)$.

Proof. In the following X and G are uniformly distributed over $\{0,1\}^{\widetilde{n}}$ and $\mathcal{G}$, respectively.

i. Fix $g \in \mathcal{G}$ and let $S_g = \big\{z \in \{0,1\}^{\ell}\colon \Pr[g(X) = z] \le 2^{-\ell-s}\big\}$. Observe the following.

   (a) $\Pr[g(X) \in S_g] \le 2^{-s}$ (by a union bound over $z \in S_g$);

(b) Fix any $z \notin S_g$ and any $x \in \{0,1\}^n$ such that $\mathrm{H}_{X|F(X)}(x \mid F(x)) \geq k$. Then,

$$\begin{aligned}\Pr[X = x \mid F'(X,g) = (F(x),g,z)] &= \Pr[X = x \mid F(X) = F(x) \wedge g(X) = z] \\ &\leq \frac{\Pr[X = x \mid F(X) = F(x)]}{\Pr[g(X) = z]} \\ &\leq \frac{2^{-k}}{2^{-\ell-s}} = 2^{-(k-\ell-s)}.\end{aligned}$$

where the second inequality follows from our assumptions on z and x.

Combining the above two observations and the bound on the real min-entropy of F, it follows that for all $g \in \mathcal{G}$, with probability $1-2^{-s}-\mathrm{neg}(n)$ over $x \leftarrow X$, we have

$$\Pr[X = x \mid F'(X,g) = F'(x,g)] \leq 2^{-(k-\ell-s)}.$$

The bound on the real min-entropy of F' follows readily.

ii. Given a PPT F'-collision-finder A', we construct a PPT F-collision-finder A as follows:

On input x, picks a pair (g,r) uniformly at random and output $\mathsf{A}'(x,g;r)$.

By the bound on the accessible max-entropy of F^{-1}, we know that there exists a family of sets $\left\{\mathcal{L}(x) \subseteq \{0,1\}^{\widetilde{n}} \colon x \in \{0,1\}^{\widetilde{n}}\right\}$ such that $|\mathcal{L}(x)| \leq 2^k$, $x \in \mathcal{L}(x)$, and

$$\Pr\left[\mathsf{A}(X,G;R) \in \mathcal{L}(X)\right] \geq 1-\mathrm{neg}(n), \tag{5.2}$$

where R is uniformly distributed over the random coins of A.

Let $\mathcal{L}'(x,g) := \{(x',g) \colon x' \in \mathcal{L}(x) \wedge g(x') = g(x)\}$. Equation (5.2) yields that

$$\Pr\left[\mathsf{A}'(X,G;R) \in \mathcal{L}'(X,G)\right] \geq 1-\mathrm{neg}(n) \tag{5.3}$$

We next bound the size of the set $\mathcal{L}'(x,g)$. Fix any $x \in \{0,1\}^n$. For any $x' \neq x$, pairwise independence of $\mathcal{G}$ tells us that $\Pr[G(x') = G(x)] = 2^{-\ell}$. It follows from linearity of expectation that

$$\mathrm{E}\left[\left|\mathcal{L}'(x,G) \setminus \{x\}\right|\right] \leq |\mathcal{L}(x)| \cdot 2^{-\ell} \leq 2^{k-\ell}$$

Then, by Markov's inequality, we have

$$\Pr\left[\left|\mathcal{L}'(x,G)\right| \leq 2^{k-\ell+s-1}+1\}\right] \geq 1-2^{-(s-1)}, \tag{5.4}$$

Combining the last two inequalities, we obtain

$$\Pr\left[\mathsf{A}'(X,G;R) \in \mathcal{L}'(X,G) \wedge \left|\mathcal{L}'(X,G)\right| \leq \max\left\{2^{k-\ell+s},1\right\}\right] \geq 1-\mathrm{neg}(n)-2^{-(s-1)} \tag{5.5}$$

The above yields an upper bound of $\max\{k-\ell+s,0\}$ on the accessible max-entropy of $(F')^{-1}$.
□

5.1.3 Reducing output length

The next transformation gives us a way to derive a function that is both length-decreasing and collision-resistant on random inputs.

Lemma 5.5 (Reducing output length)**.** Let n be a security parameter and $F\colon \{0,1\}^{\tilde{n}} \to \{0,1\}^m$ be a function. Fix a family of pairwise independent hash functions $\mathcal{G} = \left\{g\colon \{0,1\}^m \to \{0,1\}^{\tilde{n}-\log n}\right\}$ and let $F'\colon \mathcal{G} \times \{0,1\}^{\tilde{n}} \to \{0,1\}^{\tilde{n}-\log n} \times \mathcal{G}$ be defined by $F'(x,g) = (g, g(F(x)))$. The following holds: if F^{-1} has real min-entropy at least $\omega(\log n)$ and F is collision-resistant on random inputs, then F' is collision-resistant on random inputs.

Proof. The bound on real min-entropy implies that there exists a subset $S \subseteq \{0,1\}^{\tilde{n}}$ of density at most $\operatorname{neg}(n)$, such that for all $x \notin S$ it holds that $\left|F^{-1}(F(x))\right| = n^{\omega(1)}$. Hence,

$$|\operatorname{Im}(F)| \le |F(S)| + \left|F(\bar{S})\right| \le |S| + \left|\bar{S}\right| / n^{\omega(1)} \le \operatorname{neg}(n) \cdot 2^n. \tag{5.6}$$

By the two-wise independent of $\mathcal{G}$,

$$\Pr\left[\exists y' \in \operatorname{Im}(F)\colon y' \neq F(X) \wedge G(y') = G(F(X))\right] \le \frac{|\operatorname{Im}(F)|}{2^{\tilde{n}-\log n}} \le \operatorname{neg}(n). \tag{5.7}$$

Namely, $g(F(x))$ uniquely determines $F(x)$ with high probability. In particular, a collision for $g \circ F$ is also a collision for F. Given any PPT F'-collision-finder A', we construct a PPT F-collision-finder A as follows:

> On input x, pick g and r at random and compute $x' = \mathsf{A}'(x, g; r)$. If $F(x') = F(x)$, output x', else output x.

Equation (5.7) implies that $\Pr[\mathsf{A}'(X,G;R) \neq (\mathsf{A}(X;G,R),G)] \le \operatorname{neg}(n)$. Therefore, $\Pr[\mathsf{A}'(X,G;R) = (X,G)] \ge 1 - \operatorname{neg}(n)$. Namely, F' is also collision-resistant on random inputs. □

5.1.4 Additional transformations

We present two more standard transformations that are needed to complete the construction.

Lemma 5.6 (From random inputs to targets, folklore)**.** Let n be a security parameter and $F\colon \{0,1\}^{\tilde{n}} \to \{0,1\}^m$ be a length-decreasing function. Suppose F is collision-resistant on random inputs. Then,

$$\left\{F'_y\colon \{0,1\}^{\tilde{n}} \to \{0,1\}^m\right\}_{y \in \{0,1\}^{\tilde{n}}}$$

as defined by $F'_y(x) = F(y+x)$ is a family of target collision-resistant hash functions.

Proof. Given a PPT adversary A' that breaks target collision-resistance of F'_y, we can construct a PPT adversary A that breaks F as follows:

> On input x, run $\mathsf{A}'(1^n)$ to compute (x_0, state), and then run $\mathsf{A}'(\mathsf{state}, x \oplus x_0)$ to compute x_1. Output $x \oplus x_0 \oplus x_1$.

Note that (x_0,x_1) is a collision for $F'_{x\oplus x_0}$ iff $(x,x\oplus x_0\oplus x_1)$ is a collision for F. It then follows quite readily that A breaks F with the same probability that A' breaks F'_y. □

The following result of [29] (improving on [23, 3]) shows that we can construct target collision-resistant hash functions for arbitrarily long inputs starting from one for a fixed input length.

Lemma 5.7 (Increasing the input length [29])**.** Let n be a security parameter, $t=\mathrm{poly}(n)$ be a parameter and let $\left\{F_y\colon \{0,1\}^{\tilde{n}+\log n} \to \{0,1\}^{\tilde{n}}\right\}$ be a family of target collision-resistant hash functions. Then, there exists a family of target collision-resistant hash functions

$$\left\{F'_{y'}\colon \{0,1\}^{\tilde{n}+t\log n} \to \{0,1\}^{\tilde{n}}\right\}$$

where $|y'|=O(|y|\log t)$.

5.1.5 Putting everything together

Using these transformations, we can now prove Theorem 5.1.

Proof of Theorem 5.1. Recall that we have given $F:\{0,1\}^{n_0}\to\{0,1\}^{m_0}$, $s\in\omega(\log n)$, the gap Δ and that n is the security parameter.

STEP 1 **(gap amplification):** For a parameter t, we define F_1 as $F_1(x_1,\dots,x_t)=(F(x_1),\dots,F(x_t))$, the t-fold direct product of F. We choose the parameter $t\in O(n_0^2 s/\Delta^2)$ such that

$$t\cdot k_{\mathrm{REAL}} - n_0\cdot\sqrt{st} \geq t\cdot(k_{\mathrm{REAL}}-\Delta/2)+n_0\cdot\sqrt{st}+3s.$$

Lemma 5.3 yields that this repetition increases both the real and accessible entropies of F_1 by a factor of t (comparing to F). In addition, this repetition converts real Shannon entropy to real min-entropy and accessible average max-entropy to accessible max-entropy (up to additive terms that are sub-linear in t). More precisely, we have the following properties:

- $F_1:\{0,1\}^{n_1}\to\{0,1\}^{m_1}$, where $n_1(n)=t\cdot n_0$ and $m_1(n)=t\cdot m_0$.
- F_1^{-1} has real min-entropy at least $t\cdot k_{\mathrm{REAL}}-n_0\cdot\sqrt{st}\geq t\cdot(k_{\mathrm{REAL}}-\Delta/2)+n_0\cdot\sqrt{st}+3s$.
- F_1^{-1} has accessible max-entropy at most $t\cdot(k_{\mathrm{REAL}}-\Delta)+n_0\cdot\sqrt{st}$.

In steps 2 to 4, the construction uses non-uniform advice k, which corresponds to an approximation to k_{REAL}. In step 5, we will remove this non-uniform advice via "exhaustive search". Concretely, for steps 2 to 4, we are given k satisfying

$$k\in[k_{\mathrm{REAL}},k_{\mathrm{REAL}}+\Delta/2] \tag{5.8}$$

This means that

- F_1^{-1} has real min-entropy at least $t\cdot(k-\Delta)+n_0\cdot\sqrt{st}+3s$.
- F_1^{-1} has accessible max-entropy at most $t\cdot(k-\Delta)+n_0\cdot\sqrt{st}$.

This yields a gap of $3s$ between real min-entropy and accessible max-entropy.

STEP 2 **(entropy reduction):** We next apply entropy reduction to F_1 to obtain $F_2^{(k)}$. That is, $F_2^{(k)}(x,g) = (F_1(x), g, g(x))$, where $g\colon \{0,1\}^{n_1} \to \{0,1\}^{\ell}$ is selected from a family of pairwise independent hash functions with $\ell = t\cdot(k-\Delta) + n_0 \cdot \sqrt{st} + s = O(tn_0)$. Lemma 5.4 yields that this additional hashing reduces the real min-entropy and accessible max-entropy by ℓ (up to an additive term of s). More exactly, we have the following properties:

- $F_2^{(k)}\colon \{0,1\}^{n_2} \to \{0,1\}^{m_2}$ where $n_2(n,k) = O(tn_0)$ and $m_2(n,k) = O(tn_0)$. Note that in particular n_2 and m_2 also depend on k (unlike n_1 and m_1).
- If (5.8) holds, then $(F_2^{(k)})^{-1}$ has real min-entropy at least s.
- If (5.8) holds, then $(F_2^{(k)})^{-1}$ has accessible max-entropy at most 0. Hence, $F_2^{(k)}$ is collision-resistant on random inputs (by Lemma 3.7).

STEP 3 **(reducing the output length):** We next reduce the output length of $F_2^{(k)}$ by hashing the output to $n_2 - \log n$ bits. That is, $F_3^{(k)}(x,g) = (g, g(F_2^{(k)}(x)))$ where $g\colon \{0,1\}^{m_2} \to \{0,1\}^{n_2 - \log n}$ is selected from a family of pairwise-independent hash functions.

- $F_3^{(k)}\colon \{0,1\}^{n_3} \to \{0,1\}^{m_3}$ where $n_3(n,k) = O(tn_0)$ and $m_3(n,k) = n_3 - \log n$.
- By Lemma 5.5, $F_3^{(k)}$ is collision-resistant on random inputs, assuming that (5.8) holds.

STEP 4 **(adding random shifts)** We then transform $F_3^{(k)}$ into a family $\left\{G_y^{(k)}\right\}$ of target collision-resistant hash functions via a random shift, following Lemma 5.6. That is, $G_y^{(k)}(x) = F_3^{(k)}(y+x)$. We then have that

- $G_y^{(k)}(x) : \{0,1\}^{n_3} \to \{0,1\}^{m_3}$ and $G_y^{(k)}$ uses a key y of length $n_3(n,k)$.
- If (5.8) holds, then $\left\{G_y^{(k)}\right\}$ is target collision-resistant.

STEP 5 **(removing non-uniformity):** To remove the non-uniform advice k, we "try all possibilities" from 0 to n_0 in steps of size $\Delta/2$, similar to the approach used in [26] (see also [20, Section 3.6])

i. First, we construct $\kappa = n_0 \cdot 2/\Delta$ families of functions $\left\{G_y^{(k)}\right\}$, where we instantiate $\left\{G_y^{(k)}\right\}$ for all $k \in \left\{\frac{\Delta}{2}, 2\cdot\frac{\Delta}{2}, 3\cdot\frac{\Delta}{2}, \ldots, n_0\right\}$. These κ families of functions satisfy the following properties:
   - Each of $G_y^{(k)}$ is length-decreasing; in particular, $G_y^{(k)}$ has input length $n_3(n,k)$ and output length $n_3(n,k) - \log n$. Note that $G_y^{(n_0)}$ has the longest input length, i. e., $n_3(n, i\Delta/2) \le n_3(n,n_0)$ for all i because $\ell(n,k)$ increases as a function of k. We may then assume that all κ functions $G_y^1, \ldots, G_y^{\kappa}$ have the same input length $n_3(n,n_0)$ and the same output length $n_3(n,n_0) - \log n$ by padding "extra part" of the input to the output.
   - At least one of the $\left\{G_y^{(k)}\right\}$ is target collision-resistant; this is because $k_{\text{REAL}} \in [0,n_0]$, and so (5.8) holds for some k which we picked.

ii. Next, for each k, we construct a family of functions $\left\{\tilde{G}_{\tilde{y}}^{(k)}\right\}$ from $\left\{G_{y}^{(k)}\right\}$ with input length $\kappa \cdot n_3(n,n_0)$, key length $O(n_3(n,n_0)\cdot\log n)$ and output length $n_3(n,n_0)-\log n$, by following the construction given by Lemma 5.7. Again, at least one of the $\left\{\tilde{G}_{\tilde{y}}^{(k)}\right\}$ for k as above is target collision-resistant.

iii. Finally, we define a family of functions $\{G_{\tilde{y}_1,\dots,\tilde{y}_\kappa}\}$ to be the concatenation of all $\tilde{G}_{\tilde{y}}^{(k)}$ on the same input. That is, $G_{\tilde{y}_1,\dots,\tilde{y}_\kappa}(x) = \tilde{G}_{\tilde{y}_1}^{(\Delta/2)}(x)\circ\cdots\circ\tilde{G}_{\tilde{y}_\kappa}^{(n_0)}(x)$.

- Note that G has input length $\kappa \cdot n_3(n,n_0)$ and output length $\kappa\cdot(n_3(n,n_0)-\log n)$, so G is length-decreasing.
- Moreover, since at least one of $\left\{\tilde{G}_{\tilde{y}_1}^{(\Delta/2)}(x)\right\},\dots,\left\{\tilde{G}_{\tilde{y}_\kappa}^{(n_0)}\right\}$ is target collision-resistant, $\{G_{\tilde{y}_1,\dots,\tilde{y}_\kappa}\}$ must also be target collision-resistant. This is because a collision for $G_{\tilde{y}_1,\dots,\tilde{y}_\kappa}$ is a collision for each of $\tilde{G}_{\tilde{y}_1}^{(\Delta/2)},\dots,\tilde{G}_{\tilde{y}_\kappa}^{(n_0)}$.

The family $\{G_{\tilde{y}_1,\dots,\tilde{y}_\kappa}\}$ is the universal one-way hash function we wanted to construct, and so this finishes the proof of Theorem 5.1. □

5.2 UOWHF via a direct construction

Theorem 5.8. *Suppose there exists a polynomial-time computable function $F:\{0,1\}^{\widetilde{n}}\to\{0,1\}^m$ such that F^{-1} has a noticeable gap Δ between real Shannon entropy and accessible Shannon entropy. Then, there exists a family of universal one-way hash functions with output length $O(\widetilde{n}^8 s^2/\Delta^7)$ and key length $O(\widetilde{n}^8 s^2/\Delta^7\cdot\log n)$ for any $s=\omega(\log n)$.*

As before, we can use Theorem 5.8 together with results from the previous sections to get a universal one-way hash function.

Theorem 5.9. *Suppose there exists a one-way function $f:\{0,1\}^n\to\{0,1\}^n$. Then, there exists a family of universal one-way hash functions with key and output length $\widetilde{O}(n^{22})$.*

Proof. We set $s:=\log^2(n)$, and use Theorem 4.1 to get a F with $\widetilde{n}=O(n)$, and $\Delta:=O(\frac{1}{n^2})$. Using F in Theorem 5.8 this gives key and output length $\widetilde{O}(n^{22})$. □

In order to prove Theorem 5.8, we show how to transform a noticeable gap between real Shannon entropy and accessible Shannon entropy to one between real Shannon entropy and accessible max-entropy, and then follow the construction from the previous section. This step is fairly involved as we are unable to show that parallel repetition directly transforms an upper bound on accessible Shannon entropy into one for accessible max-entropy. We proceed by first establishing some additional properties achieved by gap amplification and entropy reduction.

Lemma 5.10 (Gap amplification, continued)**.** Let n be a security parameter and $F:\{0,1\}^{\widetilde{n}}\to\{0,1\}^m$ be a function. For $t\in\mathrm{poly}(n)$, let F^t be the t-fold direct product of F. Then, the following holds:

i. If F^{-1} has real Shannon entropy at most k, then $(F^t)^{-1}$ has real max-entropy at most $t\cdot k+\widetilde{n}\cdot\sqrt{st}$ for any $s=\omega(\log n)$ and $t>s$.

ii. If F^{-1} has real min-entropy at least k, then $(F^t)^{-1}$ has real min-entropy at least $t \cdot k$.

iii. If F^{-1} has real max-entropy at most k, then $(F^t)^{-1}$ has real max-entropy at most $t \cdot k$.

iv. If F^{-1} has accessible Shannon entropy at most k, then $(F^t)^{-1}$ has accessible Shannon entropy at most $t \cdot k$.

v. If F^{-1} has accessible max-entropy at most k, then $(F^t)^{-1}$ has accessible max-entropy at most $t \cdot k$.

vi. If F is q-collision-resistant on random inputs and F^{-1} has real max-entropy at most k, then $(F^t)^{-1}$ has accessible max-entropy at most $(1-q/8) \cdot tk + t$, provided that $t = \omega((1/q) \cdot \log n)$.

Proof. Again, X and $X^{(t)} = (X_1, \ldots, X_t)$ are uniformly distributed over $\{0,1\}^{\widetilde{n}}$ and $(\{0,1\}^{\widetilde{n}})^t$, respectively.

i. Follows readily from Lemma 2.3.

ii. This follows from a union bound and that fact that for all $x_1, \ldots, x_t$:

$$\mathrm{H}_{X^{(t)}}(x_1, \ldots, x_t \mid F^t(x_1, \ldots, x_t)) = \sum_{i=1}^{t} \mathrm{H}_{X|F(X)}(x_i \mid F(x_i))$$

iii. Same as previous part.

iv. Given any PPT F^t-collision-finder A', we construct the following PPT F-collision-finder A:

> On input x, pick a random i in $[t]$ along with random $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_t$, compute $\mathsf{A}'(x_1, \ldots, x_t) \to (x_1', \ldots, x_t')$, and output x_i'.

Define the random variables $(X_1', \ldots, X_t') = \mathsf{A}'(X_1, \ldots, X_t)$. Then,

$$\begin{array}{lll} \multicolumn{2}{l}{\mathrm{H}(X_1', \ldots, X_t' \mid X_1, \ldots, X_t)} & \\ \leq & \mathrm{H}(X_1' \mid X_1) + \cdots + \mathrm{H}(X_t' \mid X_t) & \text{subadditivity of conditional Shannon entropy} \\ = & t \cdot \mathrm{H}(X_I' \mid X_I) & \text{where } I \text{ has the uniform distribution over } [t] \\ = & t \cdot \mathrm{H}(\mathsf{A}(X) \mid X) & \text{by definition of } \mathsf{A} \\ \leq & t \cdot k & \text{by the bound on accessible Shannon entropy of } F^{-1} \end{array}$$

v. Analogous to Lemma 5.3 part ii, but simpler, since we do not have to use the Chernoff-Hoeffding bound.

vi. Suppose on the contrary that there exists a PPT F^t-collision-finder A' that violates the guarantee on accessible max-entropy. For $x^{(t)} \in (\{0,1\}^{\widetilde{n}})^t$,

let

$$B(x^{(t)}) := \left\{ x'^{(t)} \in (\{0,1\}^{\widetilde{n}})^t : F^t(x^{(t)}) = F^t(x'^{(t)}) \wedge \left| \left\{ i \in [t] : x'^{(t)}_i = x^{(t)}_i \right\} \right| \geq qt/8 \right\}.$$

By the bound on real max-entropy, we have that

$$\Pr\left[\exists i \in [t] : \left| F^{-1}(F(X_i^{(t)})) \right| > 2^k \right] \leq t \cdot \mathrm{neg}(n) = \mathrm{neg}(n).$$

Hence,

$$\Pr\left[\left|B(X^{(t)})\right| > \binom{t}{qt/8} 2^{(1-q/8)tk}\right] \leq \mathrm{neg}(n) \tag{5.9}$$

Since A' achieves accessible max-entropy greater than $(1-q/8)tk+t$, there must exists a non-negligible function ε such that $\Pr[\mathsf{A}'(X^{(t)};R') \notin B(X^{(t)})] \geq \varepsilon - t \cdot \mathrm{neg}(n) \geq \varepsilon/2$, where R' is uniformly distributed over the random coins of A'. Namely, A' finds collisions on at least a $1-q/8$ fraction of the coordinates with non-negligible probability.

Since F is q-collision resistant, this violates a standard Chernoff-type direct product theorem. We provide a proof sketch, following a similar analysis done for standard collision resistance in [6]. Consider the following PPT F-collision-finder A:

> On input $x \in \{0,1\}^{\widetilde{n}}$, pick a random $i \in [t]$ along with random $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_t$, compute $\mathsf{A}'(x_1, \ldots, x_t) \to (x'_1, \ldots, x'_t)$, and output x'_i.

To analyze the success probability of A', fix any subset S of $\{0,1\}^{\widetilde{n}}$ of density $q/2$. If $t = \omega(\log n/q)$, then a Chernoff bound yields that

$$\Pr[\mathsf{A}'(X^{(t)}) \notin B(X^{(t)}) \wedge \left|\left\{i \in [t] \colon X_i^{(t)} \in S\right\}\right| \geq q/4] \geq \varepsilon/4.$$

This means that

$$\Pr_{i \leftarrow [t]}[\mathsf{A}'(X^{(t)}) \to (X'_1, \ldots, X'_t) \wedge X_i \in S \wedge X'_i \neq X_i] \geq \varepsilon/4 \cdot q/8.$$

We may then deduce (following the same calculations in [6, Prop 2]) that

$$\Pr_{x \leftarrow X}\left[\Pr[\mathsf{A}(x;R) \neq x] \geq \varepsilon/4 \cdot q/8 \cdot 2/q\right] \geq 1 - q/2.$$

where R is uniformly distributed over the random coins of A. By repeating A a sufficient number of times, we may find collisions on random inputs of F with probability $1-q$, contradicting our assumption that F is q-collision-resistant on random inputs.

□

Lemma 5.11 (Reducing entropy, continued)**.** Let n be a security parameter and $F : \{0,1\}^{\widetilde{n}} \to \{0,1\}^m$ be a function. Fix a family of 2-universal hash functions $\mathcal{G} = \left\{g \colon \{0,1\}^{\widetilde{n}} \to \{0,1\}^{\ell}\right\}$. Then, $F' \colon \{0,1\}^{\widetilde{n}} \times \mathcal{G} \to \{0,1\}^m \times \mathcal{G} \times \{0,1\}^{\ell}$ as given by $F'(x,g) = (F(x), g, g(x))$ satisfies the following properties:

i. If F^{-1} has real max-entropy at most k, then $(F')^{-1}$ has real max-entropy at most $\max\{k-\ell+s, 0\}$ for any $s = \omega(\log n)$.

ii. If F^{-1} has p-accessible max-entropy at most k, then $(F')^{-1}$ has $p + 2^{-\Omega(s)}$-accessible max-entropy at most $\max\{k-\ell+s, 0\}$ for any s.

Proof. In the following X and G are uniformly distributed over $\{0,1\}^{\widetilde{n}}$ and $\mathcal{G}$, respectively.

i. Fix an x such that $\left|F^{-1}(F(x))\right| \leq 2^k$. By 2-universal hashing,

$$\mathrm{E}\left[\left|G^{-1}(G(x)) \cap (F^{-1}(F(x)) \setminus \{x\})\right|\right] \leq (2^k - 1) \cdot 2^{-\ell} \leq 2^{k-\ell}.$$

The bound on the real max-entropy of F^{-1} and Markov's inequality yield that

$$\Pr\left[\left|G^{-1}(G(X)) \cap (F^{-1}(F(X)) \setminus \{x\})\right| \geq 2^{(s-1)} \cdot 2^{k-\ell}\right] \leq 2^{-(s-1)} + \mathrm{neg}(n).$$

The bound on the real max-entropy of $(F')^{-1}$ follows.

ii. Readily follows from the proof of Lemma 5.4 part ii.

□

5.2.1 Putting everything together

Proof of Theorem 5.8. Recall that we start out with a function $F : \{0,1\}^{n_0} \to \{0,1\}^{m_0}$ with a gap Δ between real Shannon entropy and accessible Shannon entropy. Let k_{REAL} denote the real Shannon entropy of F^{-1}.

STEP 1 **(gap amplification):** Let F_1 be the t-fold direct product of F for a sufficiently large t to be determined later. That is, $F_1(x_1, \ldots, x_t) = (F(x_1), \ldots, F(x_t))$.

Lemma 5.3 yields that this repetition increases both the real and accessible entropies of F_1 by a factor of t. In addition, the repetition converts real Shannon entropy to real min-entropy and real max-entropy (up to an additive $o(t)$ term). More precisely:

- $F_1 \colon \{0,1\}^{n_1} \to \{0,1\}^{m_1}$ where $n_1(n) = t \cdot \widetilde{n}$ and $m_1(n) = t \cdot m$.
- F_1^{-1} has real min-entropy at least $t \cdot k_{\mathrm{REAL}} - n_0\sqrt{st}$ and real max-entropy at most $t \cdot k_{\mathrm{REAL}} + n_0\sqrt{st}$.
- F_1^{-1} has accessible Shannon entropy at most $t \cdot k_{\mathrm{REAL}} - t\Delta$.

From the next step on, the construction again uses an additional parameter k. We will be especially interested in the case

$$k \in [k_{\mathrm{REAL}}, k_{\mathrm{REAL}} + \Delta^2/128n_0]. \tag{5.10}$$

In case this holds,

- F_1^{-1} has accessible Shannon entropy at most $tk - t\Delta$. Lemma 3.8 yields that F_1^{-1} has $(1 - \Delta/4k)$-accessible max-entropy at most $tk - t\Delta/2$.

STEP 2 **(entropy reduction):** Apply entropy reduction to F_1 with $\ell = tk - t\Delta/2 + s$ to obtain $F_2^{(k)}$. That is, $F_2^{(k)}(x, g) = (F_1(x), g, g(x))$, where $g \colon \{0,1\}^{n_1} \to \{0,1\}^{\ell}$ is selected from a family of 2-universal hash functions.

By Lemma 5.4 and Lemma 5.11, this reduces the accessible max-entropy to 0, which allows us to deduce that $F_2^{(k)}$ is weakly collision-resistant on random inputs. Assuming Equation (5.10) we have

- $F_2^{(k)}\colon \{0,1\}^{n_2} \to \{0,1\}^{m_2}$ where $n_2(n,k) = O(tn_0 + \ell(n,k)) = O(tn_0)$ and $m_2(n,k) = O(tm_0 + \ell(n,k)) = O(tn_0)$.
- $(F_2^{(k)})^{-1}$ has real min-entropy at least $t \cdot (k_{\text{REAL}} - k + \Delta/2) - n_0\sqrt{st} - 2s$, which is at least

$$t \cdot (\Delta/2 - \Delta^2/128n_0) - n_0\sqrt{st} - 2s$$

and real max-entropy at most $t \cdot (k_{\text{REAL}} - k + \Delta/2) + n_0\sqrt{st} \le t \cdot \Delta/2 + n_0\sqrt{st}$.
- $(F_2^{(k)})^{-1}$ has $(1 - \Delta/4k + 2^{-\Omega(s)})$-accessible max-entropy at most 0. Thus, $F_2^{(k)}$ is q-collision-resistant on random inputs (by Lemma 3.7), for $q = \Delta/4k - 2^{-\Omega(s)}$.

STEP 3 **(gap amplification):** $F_3^{(k)}$ is t'-fold direct product of $F_2^{(k)}$, where $t' = s/q = O(ks/\Delta)$. That is, $F_3^{(k)}(x_1, \ldots, x_{t'}) = (F_2^{(k)}(x_1), \ldots, F_2^{(k)}(x_{t'}))$.

By Lemma 5.10, this allows us to amplify the weak collision-resistance property of $F_2^{(k)}$ to obtain a gap between real min-entropy and accessible max-entropy in $F_3^{(k)}$, again assuming Equation (5.10).

- $(F_3^{(k)})^{-1}$ has real min-entropy at least

$$t' \cdot \big(t \cdot (\Delta/2 - \Delta^2/128n_0) - n_0\sqrt{st} - 2s\big).$$

- $(F_3^{(k)})^{-1}$ has accessible max-entropy at most $t' \cdot \big((1 - q/8) \cdot (t\Delta/2 + n_0\sqrt{st}) + 1\big)$, which is at most:

$$t' \cdot \big(t \cdot (\Delta/2 - \Delta q/16) + n_0\sqrt{st}) + 1\big).$$

Now, $k \le n_0$, so $q = \Delta/4k - 2^{-\Omega(s)} \ge \Delta/4n_0 - 2^{-\Omega(s)}$. This means $(F_3^{(k)})^{-1}$ has accessible max-entropy at most:

$$t' \cdot \big(t \cdot (\Delta/2 - \Delta^2/64n_0 + 2^{-\Omega(s)}) + n_0\sqrt{st}) + 1\big).$$

Note that the gap is at least $t' \cdot \big(t \cdot \Delta^2/128n_0 - 2^{-\Omega(s)} - (2n_0\sqrt{st} + 2s + 1)\big)$, which is at least $3s$ as long as:

$$t \cdot \Delta^2/128n_0 \ge 2^{-\Omega(s)} + 2n_0\sqrt{st} + 2s + 1 + 3s/t'$$

Since $3s/t' = 3q \le 3\Delta$, we can set $t = O(n_0/\Delta + n_0 s/\Delta^2 + n_0^4 s/\Delta^4) = O(n_0^4 s/\Delta^4)$ so that $(F_3^{(k)})^{-1}$ has a gap of $3s$ between real min-entropy and accessible max-entropy, and moreover, we know where this gap is (given k).

STEP 4: We follow steps 2, 3, 4, and 5 in the previous construction, with the following modifications in the parameters:

- We apply entropy reduction first, with

$$\ell = t' \cdot \big(t \cdot (\Delta/2 - \Delta q/16) + n_0\sqrt{st}) + 1\big) + s.$$

- To remove the non-uniform advice k, we "try all possibilities" from 0 to n_0 in steps of size $\Delta^2/128n_0$.

We then obtain a non-uniform construction of UOWHFs with output and key length $O(n_0 \cdot t \cdot t') = O(n_0^6 s^2/\Delta^5)$, since $t = O(n_0^4 s/\Delta^4)$ and $t' = O(n_0 s/\Delta)$. We also obtain a uniform construction with output length $O(n_0/(\Delta^2/n_0) \cdot n_0 \cdot t \cdot t' \cdot \log n) = O(n_0^8 s^2/\Delta^7)$ and key length $O(n_0^8 s^2/\Delta^7 \cdot \log n)$.

This finishes the proof of Theorem 5.8. □

6 Connection to average-case complexity

In this section we use the notion of inaccessible entropy to reprove a result by Impagliazzo and Levin [18], given in the realm of *average-case complexity*. Our proof follows to a large extent the footstep of [18], where the main novelty is formulating the proof in the language of inaccessible entropy, and rephrasing it to make it resembles our proof of UOWHFs from one-way functions.

Section 6.1 introduces the basic notion and definitions used through the section, and in particular what "success on the average" means. It also formally describes the result of Impagliazzo and Levin [18], a result that we reprove in Section 6.2.

6.1 Preliminaries and the Impagliazzo and Levin result

We start by introducing some basic notions from average-case complexity.[10]

6.1.1 Algorithms that err

Let L be some language, and suppose $\mathsf{A}(y;r)$ is a randomized algorithm with input $y \in \{0,1\}^s$, randomness r, and output domain $\{0,1,\perp\}$.[11] It is useful to think that $\mathsf{A}(y,\cdot)$ is trying to decide L, where 0 and 1 are guesses whether $y \in \mathsf{L}$, and $\perp$ signals that A refuses to guess.

Definition 6.1. A randomized algorithm A is *α-correct* on input $y \in \{0,1\}^s$ with respect to a language $\mathsf{L} \subseteq \{0,1\}^s$, if $\Pr[\mathsf{A}(y;R) = \mathsf{L}(y)] \geq \alpha$, where we identify languages with their characteristic functions, and R is uniformly distributed over the possible random coins of A.

Recall that when defining a worst case complexity class (e. g., BPP), one requires $\Pr[\mathsf{A}(y;R) = \mathsf{L}(y)] \geq 2/3$ for *any* $y \in \{0,1\}^s$ (the choice of the constant $2/3$ is somewhat arbitrary). In other words, we require that an algorithm is $2/3-$correct for every input.

In contrast, in average-case complexity an algorithm is allowed to be wrong on some inputs. Specifically, the success probability of a decider A is measured not with respect to a *single* input, but with respect to a *distribution* over the elements of $\{0,1\}^s$.

[10] We limit the following discussion only to notions that we actually use, so this section should not be regarded as an comprehensive introduction to the field of average-case complexity (see Bogdanov and Trevisan [5] for such an introduction). While the definitions given here are equivalent to the ones given in [5], some of them are formulated somewhat differently (the interested reader is welcome to check their equivalence).

[11] We make the less common choice of using y as the input variable (and not x), since in our applications the element y is sampled as the output of a one-way function.

To make this formal, we first a problem as such a pair of language and distribution family.

Definition 6.2 (Problem)**.** A problem is a pair $(\mathsf{L},\mathcal{D})$ of language and distribution family, where $\mathsf{L} \subseteq \{0,1\}^s$ and $\mathcal{D} = \{D_i\}_{i\in\mathbb{N}}$ is a family of distributions over $\{0,1\}^s$.

The problem class $\mathsf{HeurBPP}$ (see, e. g., [5, Definition 15]) contains those problems for which there exists an efficient algorithm that is correct on all but a "small" fraction of the inputs.

Definition 6.3 ($\mathsf{HeurBPP}$)**.** A problem $(\mathsf{L},\mathcal{D})$ is in $\mathsf{HeurBPP}$, if there exists a four-input algorithm A such that the following holds for every $(n,\delta) \in \mathbb{N}\times(0,1]$: $\mathsf{A}(\cdot,1^n,\delta;\cdot)$ runs in time $p(n,1/\delta)$ for some $p \in \mathrm{poly}$, and

$$\Pr_{Y\leftarrow D_n}[\mathsf{A}(Y,1^n,\delta;\cdot) \text{ is } \tfrac{2}{3}\text{-correct on } Y] \geq 1-\delta.$$

In one second we will restrict ourselves to the case that the sampler runs in polynomial time. Then, this means that A is allowed to run in time polynomial in the "sampling complexity" of the instance, and inverse polynomial in the probability with which it is allowed to err.

6.1.2 Samplable distributions

We next study the families of distributions $\mathcal{D}$ to consider. While it is most natural to focus on efficiently samplable distributions, the definition of $\mathsf{HeurBPP}$ does not pose such limits on the distributions considered; a pair $(\mathsf{L},\mathcal{D})$ can be decided efficiently on average even if sampling the distribution $\mathcal{D}$ is a computationally hard problem. For the reduction, however, we restrict ourselves to polynomial-time samplable distributions. This limitation is crucial, since there exist (not efficiently samplable) distributions $\mathcal{D}$ with the property that $(L,\mathcal{D}) \in \mathsf{HeurBPP}$ if and only if $L \in \mathsf{BPP}$ (see [22] or [5, Section 2.5]).

Definition 6.4 (Polynomial-time samplable distributions, $(\mathsf{NP},\mathsf{PSamp})$)**.** A distribution family $\mathcal{D} = \{D_n\}_{n\in\mathbb{N}}$ is polynomial-time samplable, denoted $\mathcal{D} \in \mathsf{PSamp}$, if there exists a polynomial-time computable function D and a polynomial $p(n)$, such that $\mathsf{D}(1^n,U_{p(n)})$ is distributed according to D_n for every $n \in \mathbb{N}$, where U_m is the uniform distribution over m-bit strings.

The product set $(\mathsf{NP},\mathsf{PSamp})$ denotes the set of all pairs $(\mathsf{L},\mathcal{D})$ with $\mathsf{L} \in \mathsf{NP}$ and $\mathcal{D} \in \mathsf{PSamp}$.

Note that the input $U_{p(n)}$ to D above is the *only* source of randomness used to sample the elements of $\mathcal{D}$. In the following we make use the distribution family $\mathcal{U} = \{U_n\}_{n\in\mathbb{N}}$ (clearly, $\mathcal{U} \in \mathsf{PSamp}$).

6.1.3 Impagliazzo and Levin result

In the above terminology the result of Impagliazzo and Levin [18] can be stated as follows (cf., [5, Thm. 29]):

Theorem 6.5 ([18])**.** *If* $(\mathsf{NP},\mathsf{PSamp}) \not\subseteq \mathsf{HeurBPP}$*, then* $\exists \mathsf{L} \in \mathsf{NP}$ *with* $(\mathsf{L},\mathcal{U}) \notin \mathsf{HeurBPP}$.

In other words, suppose there exists an average-case hard problem whose distribution is polynomial-time samplable, then there exists an average-case hard problem whose distribution is uniform.

$$\begin{tikzcd}
\boxed{(\mathsf{L},\mathcal{D}) \notin \mathsf{HeurBPP}} \arrow[d, "\text{Trivial}"] & & \boxed{(\mathsf{L}',\mathcal{U}) \notin \mathsf{HeurBPP}} \\
\boxed{(\mathcal{R},\mathcal{D}) \notin \mathsf{SearchHeurBPP}} \arrow[rr, "\text{Lemma 6.9}"] & & \boxed{(\mathcal{V},\mathcal{U}) \notin \mathsf{SearchHeurBPP}} \arrow[u, "\text{Theorem 6.8}"']
\end{tikzcd}$$

Figure 1: Schematic proof of Theorem 6.5.

6.1.4 Search problems

The proof of Theorem 6.5 uses the notion of "NP-search problems":

Definition 6.6 (Search problems)**.** A search problem is a pair $(\mathcal{R},\mathcal{D})$, where $\mathcal{R} \subseteq \{0,1\}^s \times \{0,1\}^s$ is a binary relation, and $\mathcal{D} = \{D_i\}_{i\in\mathbb{N}}$ is a family of distributions over $\{0,1\}^s$. In case $\mathcal{R}$ is an NP-relation, then $(\mathcal{R},\mathcal{D})$ is an NP-search problem.

For a relation $\mathcal{R}$, let $\mathcal{R}_{\mathsf{L}}$ be the corresponding language, i. e., $\mathcal{R}_{\mathsf{L}} = \{y\colon \exists w\colon (y,w) \in \mathcal{R}\}$.

The notion of heuristics is naturally generalized to NP-search problems. The only change is that in case the algorithm claims $y \in \mathcal{R}_{\mathsf{L}}$, it additionally has to provide a witness to prove that. A search algorithm A always outputs a pair, and we let A_1 be the first component of this pair, and A_2 be the second component.

Definition 6.7 (SearchHeurBPP)**.** An NP-search problem $(\mathcal{R},\mathcal{D})$ is in SearchHeurBPP, if there exists an algorithm A outputting pairs in $\{0,1,\perp\} \times \{0,1\}^s$ such that the following holds: (1) A_1 is a heuristic for $(\mathcal{R}_{\mathsf{L}},\mathcal{D})$ and (2) $\mathsf{A}_1(y,1^n,\delta;r) = 1 \implies (y,\mathsf{A}_2(y,1^n,\delta;r)) \in \mathcal{R}$.

Algorithm A is called a (randomized) heuristic search algorithm for $(\mathcal{R},\mathcal{D})$.

6.1.5 Search problems vs. decision problems

Suppose $(\mathsf{L},\mathcal{D}) \in (\mathsf{NP},\mathsf{PSamp})$ is a "difficult decision problem". Then any NP-relation associated with L gives a "difficult NP-search problem", because finding a witness also solves the decision problem.

The converse direction is less obvious (recall that even in worst-case complexity, one invokes self-reducibility). Nevertheless, [4] prove the following (see also [5, Thm. 4.5]):

Theorem 6.8 ([4])**.** *Suppose there is an* NP*-relation* $\mathcal{R}$ *with* $(\mathcal{R},\mathcal{U}) \notin \mathsf{SearchHeurBPP}$*, then* $\exists \mathsf{L} \in \mathsf{NP}$ *with* $(\mathsf{L},\mathcal{U}) \notin \mathsf{HeurBPP}$.

Using Theorem 6.8 the proof of Theorem 6.5 proceeds as follows (see Figure 1): suppose there is a pair $(\mathsf{L},\mathcal{D}) \in (\mathsf{NP},\mathsf{PSamp}) \setminus \mathsf{HeurBPP}$ and let $\mathcal{R}$ be an NP-relation for L. Then $(\mathcal{R},\mathcal{D}) \notin \mathsf{SearchHeurBPP}$ (if $(\mathcal{R},\mathcal{D})$ would have a search heuristic algorithm, the first component of this algorithm, that outputs its left hand side output, would place $(\mathsf{L},\mathcal{D}) \in \mathsf{HeurBPP}$.) The following lemma states that in this case there is a pair $(\mathcal{V},\mathcal{U}) \notin \mathsf{SearchHeurBPP}$, and therefore Theorem 6.8 yields the conclusion.

Lemma 6.9 (Impagliazzo and Levin [18] main lemma, reproved here)**.** Assume that there exists an NP-search problem $(\mathcal{R},\mathcal{D})$ with $\mathcal{D} \in \mathsf{PSamp}$ such that $(\mathcal{R},\mathcal{D}) \notin \mathsf{SearchHeurBPP}$, then there is an NP-relation $\mathcal{V}$ such that $(\mathcal{V},\mathcal{U}) \notin \mathsf{SearchHeurBPP}$.

Intuitively Lemma 6.9 states the following: suppose some sampling algorithm $\mathsf{D}(1^n,\cdot)$ generates hard search problems, then there exist an NP-search problem that is hard over the uniform distribution.

Consider the following application of Lemma 6.9; suppose that one-way functions exist and let $f\colon \{0,1\}^n \to \{0,1\}^n$ be a length-preserving one-way function. Let $\mathsf{D}(1^n,r)$ be the algorithm that applies $f\colon \{0,1\}^n \to \{0,1\}^n$ on the input randomness $x=r$ and outputs $f(x)$, and set $\mathcal{D}=\{D_n\}_{n\in\mathbb{N}}$ to the corresponding distribution family. Furthermore, consider the NP-search problem given by the relation $\mathcal{R}=\{(f(x),x)\colon x\in\{0,1\}^s\}$. It is easy to verify that the NP-search problem $(\mathcal{R},\mathcal{D})$ is not in $\mathsf{SearchHeurBPP}$.

Lemma 6.9 implies that hard problems exist for some uniform distribution, under the assumption that one-way functions exist. However, we knew this before: if one-way functions exist, then UOWHFs also exist. Let $\mathcal{F}_k$ be such a family as in Definition 1.1, and consider the relation $\mathcal{V}=\{((z,x),x')\colon F_z(x)=F_z(x'),x\neq x'\}$, which asks us to find a non-trivial collision in F_z with a given x. By the security property of UOWHF, if we pick (z,x) uniformly at random, then this is a hard problem, and it is possible to show that $(\mathcal{V},\mathcal{U})\notin \mathsf{SearchHeurBPP}$. Thus, Rompel's result gives the conclusion of Lemma 6.9 in case we have the stronger assumption that one-way functions exist.

Given the above, it seems natural to ask whether the strategy used for constructing UOWHF from one-way functions, can be used for proving the general case stated in Lemma 6.9. In the following section we answer the above question in the affirmative. Specifically, we present an alternative proof for Lemma 6.9 following a similar approach to that taken in the first part of this paper for constructing UOWHF from one-way functions.

6.1.6 The Valiant-Vazirani lemma

We make use of the Valiant-Vazirani Lemma, originated in [31].

Lemma 6.10 ([31])**.** Let $\mathcal{S}\subseteq\{0,1\}^s$ be a set such that $2^k\le|\mathcal{S}|\le 2^{k+1}$, and let

$$\mathcal{G}=\left\{g\colon \{0,1\}^s\to\{0,1\}^{k+2}\right\}$$

be a family of pairwise independent hash-functions. Then for any $x\in\mathcal{S}$, it holds that

$$\Pr_{g\leftarrow\mathcal{G}}[g^{-1}(0^{k+2})=\{x\}]\ge 1/2^{k+3}.$$

A more usual formulation of the Valiant-Vazirani lemma states that with probability at least $\frac{1}{8}$ there is exactly one element $x\in\mathcal{S}$ with $g(x)=0^{k+2}$. This follows immediately from the above form.

Proof. The probability that $g(x)=0^{k+2}$ for a fixed $x\in\mathcal{S}$ is $1/2^{k+2}$. Conditioned on this event, due to the pairwise independence of $\mathcal{G}$, the probability that any other element of $\mathcal{S}$ is mapped to 0^{k+2} is at most $|\mathcal{S}|/2^{k+2}\le 1/2$. $\square$

6.2 Proving Lemma 6.9 via inaccessible entropy

Recall the basic idea underlying the two constructions of UOWHF from one-way functions presented in the first part of this paper. In a first step, we use the one-way function f to construct a function F that induces a gap between its real and accessible entropy (i. e., F has "inaccessible entropy"). Roughly, the distribution induced by the output of any efficient "collision-finder" algorithm getting a random x and returning a random $x' \in F^{-1}(F(x))$, has a *smaller* entropy than that induced by random preimage of $F(x)$. Afterwards, we use F to build the UOWHF.

We want to redo this first step in the current setting. Now, however, it is not anymore important to talk about collisions.[12] Thus, we can instead define F such that $F^{-1}(y)$ has some inaccessible entropy for a uniform random y. This is in fact compatible with the construction given in Section 4.2: it is possible to show that the image of F is close to uniform in case $i \approx \mathrm{H}_{f(X)}(f(x))$ (recall that i is the number of bits hashed out from $f(x)$ in the definition of F).

Let now $(\mathcal{R}, \mathcal{D})$ be an NP search problem with $\mathcal{D} \in \mathsf{PSamp}$ which is not in $\mathsf{SearchHeurBPP}$. We would like to use a similar approach as above to define a relation with limited accessible max-entropy. One might suggest that the following search problem has inaccessible entropy: given a four tuple (n, i, g, z), where g is a pairwise independent hash-function, and z has i bits, find as solution an input x such that $g(\mathsf{D}(1^n, x))_{1,\dots,i} = z$. However, it turns out that one does not in fact need the randomness inherent in the choice of z (note that a typical pairwise independent hash-function XORs the output with a random string anyhow). Instead, it makes no difference to fix $z = 0^i$, and so we adopt this to simplify the notation, so that the suggested search problem becomes to find x with $g(\mathsf{D}(1^n, x))_{1,\dots,i} = 0^i$ for a given triple (n, i, g).

Problems with the above intuition and postprocessing the witness. A moment of thought reveals that there can be cases where this suggested search problem is easy. For example if the sampler $\mathsf{D}(1^n, x)$ simply outputs $y = x$ itself, which is possible if finding w with $(y, w) \in \mathcal{R}$ is difficult for a uniform random y. The solution is easy: ask the solving algorithm to output also a matching witness w with $(\mathsf{D}(1^n, x), w) \in \mathcal{R}$ (ignore invalid outputs).

Thus, the suggested search problem becomes: "given (n, i, g), find (x, w) such that $g(\mathsf{D}(1^n, x))_{1\dots i} = 0^i$ and $(\mathsf{D}(1^n, x), w) \in \mathcal{R}$". The hope is then that this search problem has limited accessible entropy in the coordinate corresponding to x (we do not want to talk about the entropy in w because it arise from the number of witnesses which $\mathcal{R}$ has, and at this point we have no control over this number).

There is a last little problem to take care of: it is not obvious how to encode n into the search problem, as (n, i, g) does not look like a uniform bitstring of a certain length, even if i and g look random. However, it is possible to ensure that the length of (i, g) uniquely define n, and we assume that this is done in such a way that n can be easily computed from the length of (i, g).

6.2.1 A Relation with bounded accessible average max-entropy

Using the above discussion, we now finally have enough intuition to define the relation $\mathcal{Q}$. For $\mathcal{D} \in \mathsf{PSamp}$, we let $\mathrm{Canon}(\mathcal{D})$ be an arbitrary polynomial-time sampler for $\mathcal{D}$.

Construction 6.11. Let $\mathcal{R}$ be an NP relation, let $\mathcal{D} \in \mathsf{PSamp}$, let $\mathsf{D} = \mathrm{Canon}(\mathcal{D})$ and let $d \in \mathbb{N}$ be such that D's running time on input $(1^n, \cdot)$ is bounded by n^d. Let $\mathcal{G}$ be an explicit and constructible family of pairwise independent hash functions, where the family $\mathcal{G}_m$ maps all strings of length at most m to strings of length m.

[12] One advantage of using an "inversion problem" instead of a "collision problem" is that it becomes possible to use two-wise independent hash-functions (instead of three-wise independent).

For $n \in \mathbb{N}$, define

$$\mathcal{Q}^{(n)} := \left\{ \big((i,g),(x,w)\big) : x \in \{0,1\}^{n^d}, i \in [n^d], g \in \mathcal{G}_{n^d}, g(\mathsf{D}(1^n,x))_{1\ldots i} = 0^i, (\mathsf{D}(1^n,x),w) \in \mathcal{R} \right\},$$

and let $\mathcal{Q} := \bigcup_{n\in\mathbb{N}} \mathcal{Q}^{(n)}$.

Note that the elements of $\mathcal{G}_m$ in Construction 6.11 have domain $\bigcup_{i=0}^{m}\{0,1\}^i$. This somewhat unusual requirement is needed since the sampler might output strings of arbitrary lengths (up to n^d).

From now on, we will only consider the case where we have some fixed sampler D in mind. In this case, whenever n is given, we will assume that (i,g) are elements satisfying the conditions in Construction 6.11. Furthermore, we assume without loss of generality that (the encoding of) a uniform random bitstring, of the right length, induces the uniform distribution on $\mathcal{G}_{n^d} \times [n^d]$.

6.2.2 Accessible average max-entropy

We next define what it means for an NP-search problem to have limited accessible max-entropy, with respect to a part of its witness. This notion is modeled by introducing a function f that outputs the "interesting part" of the witness.

Definition 6.12. Let $\mathcal{Q}$ be an NP relation, and $f\colon \{0,1\}^s \to \{0,1\}^s$ a function. For $y \in \{0,1\}^s$ let

$$\mathcal{S}_{\mathcal{Q},f}(y) := \{f(w)\colon (y,w) \in \mathcal{Q}\}.$$

The real average max-entropy of $(\mathcal{Q},\mathcal{D})$ with respect to f, is the function

$$\mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{D},f}(m) = \mathop{\mathrm{E}}_{Y\leftarrow D_m}[\log(|\mathcal{S}_{\mathcal{Q},f}(Y)|)]$$

letting $\log(0) := -1$.[13]

In case the relation $\mathcal{R}$ and f are clear from the context, we sometimes write $\mathcal{S}(y)$ instead of $\mathcal{S}_{\mathcal{Q},f}(y)$.

We next define a useful notion of limited accessible max-entropy in this setting. Here, one should think of algorithm A as an algorithm which, on input y produces a witness w with $(y,w) \in \mathcal{Q}$. It furthermore "aims" to produce witnesses w for which $f(w)$ has as much entropy as possible.

Definition 6.13. Let $f\colon \{0,1\}^s \to \{0,1\}^s$, $p, k\colon \mathbb{N}\times(0,1] \to \mathbb{R}$, $\mathcal{Q}$ an NP-relation, and $\mathsf{D} = \{D_n\}_{n\in\mathbb{N}}$ a family of distributions. The pair $(\mathcal{Q},\mathcal{D})$ has i.o. (infinitely often) ρ-accessible average max-entropy at most k with respect to f, if for every four-input algorithm $\mathsf{A}(y,1^m,\varepsilon;r)$ running in time $\ell = \ell(m,1/\varepsilon)$ for some $\ell \in \mathrm{poly}$, there exists infinitely many m's in $\mathbb{N}$, a function $\varepsilon(m) = \varepsilon \in (0,1]$ and an ensemble of set families

$$\left\{ \{\mathcal{L}_m(y) \subseteq \{0,1\}^s\}_{y\in\mathrm{Supp}(D_m)} \right\}_{m\in\mathbb{N}},$$

such that

$$\Pr_{Y\leftarrow D_m, R\leftarrow\{0,1\}^\ell}\left[\Gamma(Y,\mathsf{A}(Y,1^m,\varepsilon;R)) \in \big(\mathcal{L}_m(Y)\cup\{\perp\}\big)\right] \geq 1-\rho(m,\varepsilon) \tag{6.1}$$

[13] The convention $\log(0) = -1$ helps us to simplify notation in a few places. (We need that $\log(0) < \log(1)$ since an algorithm which produces no valid witness should produce less entropy than an algorithm which produces some valid witness).

where $\Gamma = \Gamma_{\mathcal{Q},f}(y,w)$ equals $f(w)$ in case $(y,w) \in \mathcal{Q}$ and equals $\perp$ otherwise, and

$$\mathop{\mathrm{E}}_{Y \leftarrow D_m}\left[\log(|\mathcal{L}_m(Y)|)\right] \leq k(m,\varepsilon) \tag{6.2}$$

The following lemma, proven in Section 6.2.3, states that the relation $\mathcal{Q}$ defined in Construction 6.11 has limited accessible max-entropy with respect to the function $(x,w) \to x$.

Lemma 6.14. Let $(\mathcal{R},\mathcal{D})$ be an NP relation with $\mathcal{D} \in \mathsf{PSamp}$ and $(\mathcal{R},\mathcal{D}) \notin \mathsf{SearchHeurBPP}$. Define $\mathcal{Q}$ from $(\mathcal{R},\mathcal{D})$ as in Construction 6.11, and let $f\colon \{0,1\}^s \times \{0,1\}^s \to \{0,1\}^s$ be given by $f(x,w) = x$. Then, for any $c \in \mathbb{N}$, $(\mathcal{Q},\mathcal{U})$ has i.o. $(\frac{\varepsilon}{m})^c$-accessible average max-entropy at most

$$k(m,\varepsilon) = \mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{U},f}(m) - \sqrt[c]{\varepsilon} \cdot m^c$$

with respect to f.

Lemma 6.14 is proven below, and in Section 6.2.4 we use Lemma 6.14 for proving Lemma 6.9. The latter is done by additionally fixing the value of $h(x)_{1\ldots j}$, where h is an additional random hash function, and j is a random integer (in a certain range). The ratio is that an algorithm producing (x,w) with $h(x)_{1\ldots j} = 0^j$, can be used to access max-entropy roughly 2^j.

6.2.3 Proving Lemma 6.14

The proof of Lemma 6.14 follows similar lines to the proof of Theorem 4.5.

Proof (of Lemma 6.14). Let A be an algorithm that "aims to produce max-entropy for $\mathcal{Q}$". Without loss of generality, we assume that A either outputs a valid witness (x,w) for a given input (i,g) or $\perp$. We show how to find infinitely many $m \in \mathbb{N}$, $\varepsilon = \varepsilon(m) \in (0,1]$, and ensemble of set families

$$\left\{\mathcal{L}_m = \{\mathcal{L}_m(i,g)\}_{(i,g)\in\mathcal{Q}_{\mathsf{L}}}\right\}_{m\in\mathbb{N}}$$

with the properties as required in the lemma (we write $\mathcal{L}_m(i,g)$ instead of $\mathcal{L}_m(y)$ because the elements of $\mathcal{Q}_{\mathsf{L}}$ are pairs (i,g)). Towards achieving the above, consider the following candidate algorithm B for a search heuristics for $(\mathcal{R},\mathcal{D})$.

Let $\beta \in \mathbb{N}$ be a constant to be determined by the analysis, let $d \in \mathbb{N}$ be such that n^d is an upper bound on the runtime of the sampler D (recall that we have fixed D above) on input $(1^n,\cdot)$, and let $\ell = \ell(m,\varepsilon)$ be an upper bound on the running time of A on parameters m and ε. Let $m(n)$ be the description length of a pair in $[n^d] \times \mathcal{G}_{n^d}$ and let

$$\varepsilon(n,\delta) = \left(\delta(n)/n^{\beta}\right)^{\beta}.$$

Search Heuristics B^{A}

Oracle: A // Entropy generator for $(\mathcal{Q},\mathcal{U})$.
Input: $y \in \{0,1\}^s$, 1^n and $\delta \in (0,1]$.

$\varepsilon = \varepsilon(n,\delta)$; $m = m(n)$; $\ell = \ell(m,\varepsilon)$
repeat $n \cdot (n/\delta)^{\beta}$ **times:**
  $i \leftarrow [n^d]$
  $g \leftarrow \left\{g' \in \mathcal{G}_{n^d} : g'(y)_{1\ldots i} = 0^i\right\}$
  $r \leftarrow \{0,1\}^{\ell}$
  $(x,w) = \mathsf{A}((i,g),1^m,\varepsilon;r)$
  if $(y,w) \in \mathcal{R}$ **and** $\mathsf{D}(1^n,x) = y$
    return $(1,w)$
return 0

It is clear that, for any $\beta \in \mathbb{N}$, the running time of B is $\operatorname{poly}(n,1/\delta)$. Recall that, by assumption, $(\mathcal{R},\mathcal{D})$ has no search heuristics. Algorithm B satisfies item (2) of Definition 6.7, and since $(\mathcal{R},\mathcal{D})$ has no search heuristics, it must fail to satisfy item (1) (i. e., B_1 cannot be a randomized heuristic). Note now that algorithm B_1 is always correct if $y \notin \mathsf{L}$ (always outputs 0). Thus, B must fail to be correct on some $y \in \mathsf{L}$. In fact, there exist infinitely many n's in $\mathbb{N}$ and a function $\delta = \delta(n) \in (0,1]$. such that $\mathsf{B}_1(y,1^n,\delta;\cdot)$ is not $\frac{2}{3}$-correct for more than a fraction δ of the inputs $y \in \mathcal{R}_{\mathsf{L}}$ produced by $\mathsf{D}(1^n,\cdot)$.

The following discussion is with respect to any fixed pair $(n,\delta = \delta(n))$ from the above infinite set.

We present a family of sets

$$\{\mathcal{L}_m(i,g)\}_{(i,g)\in\mathcal{G}_{n^d}\times[n^d]}$$

for which Equations (6.1) and (6.2) holds with respect to algorithm A and f, for the parameters $m = m(n)$, $\varepsilon = \varepsilon(n,\delta) = \varepsilon(m)$, and ρ and k as stated in the lemma. Since this holds for *any* such pair (n,δ) and since $m(n) \in \Omega(n)$ (and thus, there are infinitely many different m's) the proof of the lemma would follow.

Consider the following set

$$\mathcal{Y} = \left\{ y \in \mathsf{L} \colon \Pr_{\substack{I\leftarrow[n^d]\\ G\leftarrow\mathcal{G}_{n^d}, R\leftarrow\{0,1\}^{\ell}}}\left[\mathsf{A}_1((I,G),1^m,\varepsilon;R) \in \mathsf{D}^{-1}(1^n,y) \colon G(y)_{1\ldots I} = 0^I\right] < \left(\frac{\varepsilon}{n}\right)^{\beta} \right\}, \tag{6.3}$$

letting

$$\mathsf{D}^{-1}(1^n,y) := \left\{x \in \{0,1\}^{n^d} \colon \mathsf{D}(1^n,x) = y\right\}$$

and let $\ell = \ell(m)$. Note that $\mathcal{Y}$ contains all the y's in L for which B_1 is not $\frac{2}{3}$-correct, and the above discussion implies

$$\Pr_{Y\leftarrow D_n}[Y \in \mathcal{Y}] > \delta \tag{6.4}$$

Towards defining the sets $\{\mathcal{L}_m(i,g)\}$, we partition the preimages of the elements in $\mathcal{Y}$ into buckets; for $i \in \left\{0,\ldots,n^d-1\right\}$ let

$$\overline{\mathcal{L}}(i) := \left\{x \in \{0,1\}^{n^d} \colon \mathsf{D}(1^n,x) \in \left(\mathcal{Y} \cap \left\{y \colon \mathrm{H}_{D_n(y)}(y) \in [i,i+1)\right\}\right)\right\}, \tag{6.5}$$

where $\mathrm{H}_{D_n(y)}(y) = -\log(1/D_n(y))$ is the sample entropy of y with respect to the distribution D_n. In words: $\overline{\mathcal{L}}(i)$ are those x for which $y = \mathsf{D}(1^n, x) \in \mathsf{L}$ is an element for which $\mathsf{B}(y)$ is unlikely to produce a witness, and for which y has roughly $2^{(n^d)-i}$ preimages.

For $(i,g) \in \mathcal{G}_{n^d}$, the set $\mathcal{L}_m(i,g)$ is defined as

$$\mathcal{L}_m(i,g) := \mathcal{S}(i,g) \setminus \overline{\mathcal{L}}(i), \tag{6.6}$$

where $\mathcal{S}(i,g)$ is taken from Definition 6.12. In the remaining of the proof we show that, for the right choice of β, Equations (6.1) and (6.2) holds with respect to the family $\{\mathcal{L}_m(i,g)\}$ for the functions $\rho(m,\varepsilon) = (\frac{\varepsilon}{m})^c$ and $k(m,\varepsilon) = \mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{U},f}(m) - \sqrt[c]{\varepsilon} \cdot m^c$. The proof easily follow by the next two claims.

Claim 6.15. We have

$$\Pr_{\substack{I \leftarrow [n^d], G \leftarrow \mathcal{G}_{n^d} \\ R \leftarrow \{0,1\}^\ell}} [\mathsf{A}_1((I,G), 1^m, \varepsilon; R) \in \overline{\mathcal{L}}(I)] \leq 2n^d \cdot \left(\frac{\varepsilon}{n}\right)^\beta$$

Claim 6.16. For $i \in [n^d]$ and $g \in \mathcal{G}_{n^d}$, let $\tilde{\mathcal{S}}(i,g) = \mathcal{S}(i,g)$ in case this set in non-empty and $\tilde{\mathcal{S}}(i,g) = \{\perp\}$ otherwise, then

$$\Pr_{\substack{I \leftarrow [n^d], G \leftarrow \mathcal{G}_{n^d} \\ X \leftarrow \tilde{\mathcal{S}}(I,G)}} [X \in \overline{\mathcal{L}}(I)] \geq \frac{\delta}{10n^d}$$

Before proving the above claims, we first use them to conclude the proof of the lemma.

Claim 6.15 implies that for large enough β

$$\Pr_{\substack{I \leftarrow [n^d], G \leftarrow \mathcal{G}_{n^d} \\ R \leftarrow \{0,1\}^\ell}} [\Gamma\big((I,G), \mathsf{A}((I,G), 1^m, \varepsilon; R)\big) \in (\mathcal{L}_m(I,G) \cup \{\perp\})] \geq 1 - 2n^d \cdot \left(\frac{\varepsilon}{n}\right)^\beta \geq 1 - \left(\frac{\varepsilon}{m}\right)^c \tag{6.7}$$

yielding that Equation (6.1) holds for $\{\mathcal{L}_m(i,g)\}$ and ρ.

Applying Markov's inequality on Claim 6.16, yields that

$$\Pr_{X \leftarrow \mathcal{S}(i,g)} [\mathsf{D}(1^n, X) \in \overline{\mathcal{L}}(i)] \geq \frac{\delta}{20n^d} \tag{6.8}$$

for at least $\frac{\delta}{20n^d}$ fraction of the pairs $(i,g) \in [n^d] \times \mathcal{G}_{n^d}$. Where for any such pair it holds that

$$\begin{aligned} \log(|\mathcal{L}_m(i,g)|) &\leq \log((1 - \frac{\delta}{20n^d}) \cdot |\mathcal{S}(i,g)|) \\ &\leq \log(|\mathcal{S}(i,g)|) - \frac{\delta}{20n^d}. \end{aligned} \tag{6.9}$$

It follows that for large enough β

$$\begin{aligned} \mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{U},f}(m) - \mathrm{E}[\log(|\mathcal{L}_m(I,G)|)] &= \mathop{\mathrm{E}}_{(I,G) \leftarrow [n^d] \times \mathcal{G}_{n^d}} [\log(|\mathcal{S}(I,G)|) - \log(|\mathcal{L}(m,I,G)|)] \\ &\geq \frac{\delta^2}{400n^{2d}} = \frac{(\sqrt[\beta]{\varepsilon} \cdot m^\beta)^2}{400n^{2d}} \\ &\geq \sqrt[c]{\varepsilon} \cdot m^c \end{aligned}$$

Hence, Equation (6.2) holds for $\{\mathcal{L}_m(i,g)\}$ and k, and the proof of the lemma follows. $\square$

Proof of Claim 6.15. Compute

$$n^d \cdot \left(\frac{\varepsilon}{n}\right)^\beta \geq \sum_{y \in \mathcal{Y}} D_n(y) \cdot n^d \cdot \left(\frac{\varepsilon}{n}\right)^\beta \tag{6.10}$$
$$\geq \sum_{i=1}^{n^d} \sum_{y \in \mathsf{D}(1^n, \overline{\mathcal{L}}(i))} D_n(y) \cdot \Pr_{G \leftarrow \mathcal{G}_{n^d}, R \leftarrow \{0,1\}^\ell}[\mathsf{A}_1((i,G),1^m,\varepsilon;R) \in \mathsf{D}^{-1}(1^n,y) \mid G(y)_{1\dots i} = 0^i],$$

letting $\mathsf{D}(1^n, \overline{\mathcal{L}}(i)) := \left\{\mathsf{D}(1^n,x) \colon x \in \overline{\mathcal{L}}(i)\right\}$. In addition, for any (i,r) it holds that

$$\begin{aligned}
&\sum_{y \in \mathsf{D}(1^n, \overline{\mathcal{L}}(i))} D_n(y) \cdot \Pr_{G \leftarrow \mathcal{G}_{n^d}}[\mathsf{A}_1((i,G),1^m,\varepsilon;r) \in \mathsf{D}^{-1}(1^n,y) \mid G(y)_{1\dots i} = 0^i] \\
&= \sum_{y \in \mathsf{D}(1^n, \overline{\mathcal{L}}(i))} D_n(y) \cdot 2^i \cdot \Pr_{G \leftarrow \mathcal{G}_{n^d}}[\mathsf{A}_1((i,G),1^m,\varepsilon;r) \in \mathsf{D}^{-1}(1^n,y) \wedge G(y)_{1\dots i} = 0^i] \\
&= \sum_{y \in \mathsf{D}(1^n, \overline{\mathcal{L}}(i))} D_n(y) \cdot 2^i \cdot \Pr_{G \leftarrow \mathcal{G}_{n^d}}[\mathsf{A}_1((i,G),1^m,\varepsilon;r)) \in \mathsf{D}^{-1}(1^n,y)] \\
&\geq \frac{1}{2} \cdot \sum_{y \in \mathsf{D}(1^n, \overline{\mathcal{L}}(i))} \Pr_{G \leftarrow \mathcal{G}_{n^d}}[\mathsf{A}_1((i,G),1^m,\varepsilon;r)) \in \mathsf{D}^{-1}(1^n,y)] \\
&= \frac{1}{2} \cdot \Pr_{G \leftarrow \mathcal{G}_{n^d}}[\mathsf{A}_1((i,G),1^m,\varepsilon;r)) \in \overline{\mathcal{L}}(i)].
\end{aligned}$$

Collecting the equations yields the claim. □

For the proof of Claim 6.16, we need first a pairwise independence analogue of Claim 4.8. The proof is exactly the same, except a bit simpler as we fix the output w instead of fixing another preimage. We provide it for completeness.

Claim 6.17. Let $i \in [n^d]$, $w \in \{0,1\}^i$ and $x^* \in \{0,1\}^{n^d}$ be such that $\mathsf{H}_{f(X)}(f(x^*)) \geq i$. Then,

$$\Pr_{\substack{g \leftarrow \mathcal{G}_{n^d} \\ x \leftarrow (g \circ f)^{-1}(w)}}[x = x^*] \geq \frac{2^{-n^d}}{10},$$

letting $(g \circ f)^{-1}(w)$ equals the set $\{x \colon g(f(x))_{1\dots i} = w\}$ in case this set is not empty, and $\{\perp\}$ otherwise.

Proof. Let G be uniformly distributed over $\mathcal{G}_{n^d}$, and let E be the event that $G(f(x^*)) = w$. Note that $\Pr[E] = 2^{-i}$ and that

$$\Pr_{x \leftarrow (G \circ f)^{-1}(w)}[x = x^* \mid E] = \frac{1}{|(G \circ f)^{-1}(w)|} \tag{6.11}$$

The pairwise independence of $\mathcal{G}_{n^d}$ yields that $\Pr[G(f(x)) = w \mid E] = 2^{-i}$ for any $x \in \{0,1\}^{\leq n^d} \setminus f^{-1}(f(x^*))$. Hence, $\mathrm{E}\left[\left|(G \circ f)^{-1}(w) \setminus f^{-1}(f(x^*))\right| \mid E\right] \leq 2 \cdot 2^{-i+n^d}$ and by Markov's inequality

$$\Pr\left[\left|(G \circ f)^{-1}(w) \setminus f^{-1}(f(x^*))\right| \leq 4 \cdot 2^{-i+n^d} \mid E\right] \geq \frac{1}{2} \tag{6.12}$$

Combining the above inequality and the assumption $\mathrm{H}_{f(X)}(f(x^*)) \geq i$, we get

$$\Pr\left[\left|(G\circ f)^{-1}(w)\right| + |f^{-1}(f(x^*)| \leq 5\cdot 2^{-i+n^d} \mid E\right] \geq \frac{1}{2}, \tag{6.13}$$

and conclude that

$$\begin{aligned}\Pr_{x\leftarrow(G\circ f)^{-1}(w)}[x=x^*] &= \Pr[E]\cdot \Pr_{x\leftarrow(G\circ f)^{-1}(w)}[x=x^* \mid E]\\ &\geq 2^{-i}\cdot\frac{1}{2}\cdot\frac{1}{5\cdot 2^{-i+n^d}}\\ &= \frac{2^{-n^d}}{10}.\end{aligned}$$

□

Proof of Claim 6.16. For $i \in \left\{0\ldots,n^d\right\}$ and $x \in \overline{\mathcal{L}}(i)$, Claim 6.17 yields that

$$\Pr_{G\leftarrow\mathcal{G}_{n^d},X\leftarrow\tilde{\mathcal{S}}(i,G)}[X=x] \geq \frac{2^{-n^d}}{10} \tag{6.14}$$

By Equation (6.4) it holds that $\Pr_{Y\leftarrow D_n}[Y\in\mathcal{Y}] > \delta$, and therefore

$$\Pr_{X\leftarrow\{0,1\}^{n^d}}\left[X\in\bigcup_{i=1}^{n^d}\overline{\mathcal{L}}(i)\right] = \Pr_{Y\leftarrow D_n}[Y\in\mathcal{Y}] > \delta.$$

We conclude that

$$\begin{aligned}\Pr_{\substack{I\leftarrow[n^d],G\leftarrow\mathcal{G}_{n^d}\\ X\leftarrow\tilde{\mathcal{S}}(I,G)}}[X\in\overline{\mathcal{L}}(I)] &\geq \mathrm{E}\left[\left|\overline{\mathcal{L}}(I)\right|\right]\cdot\frac{2^{-n^d}}{10}\\ &\geq \frac{\delta\cdot 2^{n^d}}{n^d}\cdot\frac{2^{-n^d}}{10} = \frac{\delta}{10n^d}.\end{aligned}$$

□

6.2.4 A difficult problem for the uniform distribution

In this section we show how to transform a uniform search problem with a gap between its real and accessible entropy, into a uniform search problem for which no heuristic search algorithm exists (i. e., the problem is not in SearchHeurBPP). Combining it with Lemma 6.14 concludes the proof of Lemma 6.9.

The transformation is achieved by adding additional restriction on the witness of the given search problem. Specifically, requiring its "hash value" with respect to a randomly chosen pairwise independent hash function to be the all zero string.

We use the following construction:

Construction 6.18. Let $\mathcal{Q}$ be an NP-relation, let $f\colon \{0,1\}^s \to \{0,1\}^s$ be a function, let $\mathcal{G} = \{\mathcal{G}_m\}$ be family of pairwise independent hash functions, where the functions of $\mathcal{G}_k$ map strings of length at most k to string of length m (as in Construction 6.11), and let $d \in \mathbb{N}$ be such that $(y,w) \in \mathcal{Q} \implies |f(w)| \le |y|^d$. For $n \in \mathbb{N}$ let

$$\mathcal{V}^{(n)} := \left\{ \left((y,j,g),w\right)\colon y \in \{0,1\}^n, j \in [n^d+2], g \in \mathcal{G}_{n^d+2}, (y,w) \in \mathcal{Q}, g(f(w))_{1\dots j} = 0^j \right\}$$

and let $\mathcal{V} := \bigcup_{n\in\mathbb{N}} \mathcal{V}^{(n)}$.

As in Construction 6.18, we assume that the tuples (y,j,g)'s above can be encoded such that a uniformly random string, of the right length, decodes to a uniformly random tuple in $\{0,1\}^n \times [n^d+2] \times \mathcal{G}_{n^d+2}$.

Lemma 6.19. Let $\mathcal{Q}$, f, d and $\mathcal{V}$ be as in Construction 6.18. Suppose that $(\mathcal{Q},\mathcal{U})$ has i.o. $\frac{\varepsilon^2}{50m^d}$-accessible average max-entropy at most $\mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{U},f}(m) - 5\varepsilon m^d$ with respect to f, then $(\mathcal{V},\mathcal{U}) \notin \mathsf{SearchHeurBPP}$.

Proof. We assume towards a contradiction that $(\mathcal{V},\mathcal{U}) \in \mathsf{SearchHeurBPP}$, and show that $(\mathcal{Q},\mathcal{U})$ has too high accessible average max-entropy.

Let A be a randomized search heuristics for $(\mathcal{V},\mathcal{U})$. The following algorithm B contradicts the assumption that $(\mathcal{Q},\mathcal{U})$ has i.o. $\frac{\varepsilon^2}{50m^d}$-accessible average max-entropy at most $\mathrm{H}^{\mathrm{Real}}_{\mathcal{Q},\mathcal{U},f}(m) - 5\varepsilon m^d$ with respect to f.

Let $\ell = \ell(n,\delta)$ be an upper bound on the running time of A on parameters n and δ. Let $n(m)$ be the description length of a triplet in $\{0,1\}^m \times [m^d+2] \times \mathcal{G}_{m^d+2}$ and let $\delta(m,\varepsilon) = \frac{\varepsilon^2}{100m^d}$.

Entropy generator B^{A} for $(\mathcal{Q},\mathcal{U})$

Oracle: A // Search heuristics for $(\mathcal{V},\mathcal{U})$.
Input: $y \in \{0,1\}^m$, 1^m and $\varepsilon \in (0,1]$

$\delta = \delta(m,\varepsilon)$; $n = n(m)$; $\ell = \ell(n,\delta)$
$j \leftarrow \left\{2,\dots,m^d+2\right\}$
$g \leftarrow \mathcal{G}_{m^d+2}$
$r \leftarrow \{0,1\}^\ell$
$(b,w) = \mathsf{A}(y,j,g,1^n,\delta;r)$
if $b=1$ **return** w, **else return** $\perp$

It is clear that the running time of B is $\mathrm{poly}(m,1/\varepsilon)$. We show that B achieves high max-entropy for all (except maybe finitely many) values of m and any ε. Specifically, that for all (except maybe finitely many) (m,ε), there exists no family $\{\mathcal{L}_m(y)\}_{y\in\{0,1\}^m}$ as described in Definition 6.13.

Fix m and ε, and let the random variables Y, J, G and R be uniformly chosen from $\{0,1\}^m \times [m^d + 2] \times \mathcal{G}_{m^d+2} \times \{0,1\}^\ell$. For $y \in \{0,1\}^m$, let $\eta(y)$ be the probability that $\mathsf{A}(y,J,G,1^n,\delta;\cdot)$ is not $\frac{2}{3}$-correct. Since $\mathrm{E}[\eta(Y)] \le \delta = \frac{\varepsilon^2}{100m^d}$, it holds that

$$\Pr\left[\eta(Y) \le \frac{\varepsilon}{8m^d}\right] \ge 1-\varepsilon \tag{6.15}$$

Fix $y \in \{0,1\}^m$ and let $\mathcal{S}(y) = \mathcal{S}_{\mathcal{Q},f}(y) = \{f(w)\colon (y,w) \in \mathcal{Q}\}$. Intuitively, since $f(w)$ is the "interesting part" of a $\mathcal{Q}$-witness w for y, $\mathcal{S}(y)$ is the set of interesting parts of witnesses of y.

For $x \in \mathcal{S}(y)$ let $\mathcal{E}(y,x) \subseteq [m^d+2] \times \mathcal{G}_{m^d+2}$ be the set of pairs (j,g) for which x is the only element in $\mathcal{S}(y)$ with $g(x)_{1\dots j} = 0^j$. Going back to our intuition again, if x is the "interesting part" of a $\mathcal{Q}$-witness for y, then $\mathcal{E}(y,x)$ is the set of pairs (j,g), such that on input (y,j,g), algorithm A *must* output w with $f(w) = x$ to be successful.

By Lemma 6.10, for every fixed pair (y,x) with $x \in \mathcal{S}(y)$:

$$\Pr[(J,G) \in \mathcal{E}(y,x)] \geq \frac{1}{8|\mathcal{S}|m^d} \tag{6.16}$$

The following argument would become simpler if Equation (6.16) would hold with equality, or if at least $\Pr[(J,G) \in \mathcal{E}(y,x)]$ was the same for every x with $x \in \mathcal{S}(y)$. This may of course not be the case, but since all pairs (j,g) have the same probability we can simply discard elements from the sets $\mathcal{E}(y,x)$ until they all have the same size. We call the resulting sets $\mathcal{E}'(y,x)$. Hence we can assume that $\mathcal{E}'(y,x) \subseteq \mathcal{E}(y,x)$ and

$$\Pr[(J,G) \in \mathcal{E}'(y,x_1)] = \Pr[(J,G) \in \mathcal{E}'(y,x_2)] \geq \frac{1}{8|\mathcal{S}|m^d}$$

for all $x_1, x_2 \in \mathcal{S}(y)$.

We next let $\mathcal{E}'(y) = \bigcup_{x \in \mathcal{S}} \mathcal{E}'(y,x)$. Intuitively, $\mathcal{E}'(y)$ are those pairs (j,g) such that A on input (y,j,g) is forced to answer with some unique x. Let $x_y(g,j)$ to be the unique element of $\mathcal{S}(y)$ with $g(x)_{1\dots j} = 0^j$ in case it exists, and $\perp$ otherwise, and let $X_y = x_y(G,J)$. Because of the above trickery with $\mathcal{E}'$, conditioned on $(J,G) \in \mathcal{E}'(y)$, the random variable X_y is uniformly distributed over $\mathcal{S}(y)$. Furthermore, since $\Pr[(J,G) \in \mathcal{E}'(y)] \geq \frac{1}{8m^d}$, it holds that

$$\begin{aligned}
&\Pr[\mathsf{A}(y,J,G,1^n,\delta;\cdot) \text{ is } \textit{not } \tfrac{2}{3}\text{-correct} \mid (J,G) \in \mathcal{E}'(y)] \\
&\leq 8m^d \cdot \Pr[\mathsf{A}(y,J,G,1^n,\delta;\cdot) \text{ is } \textit{not } \tfrac{2}{3}\text{-correct} \wedge (J,G) \in \mathcal{E}'(y)] \\
&\leq 8m^d \cdot \eta(y).
\end{aligned} \tag{6.17}$$

By definition,

$$\Pr[f(\mathsf{A}_2(y,j,g,1^n,\delta;R)) = x_y(j,g)] \geq \frac{2}{3} \tag{6.18}$$

for every $(j,g) \in \mathcal{E}'(y)$ such that $\mathsf{A}(y,j,g,1^n,\delta;\cdot)$ is $\frac{2}{3}$-correct.

For each $(j,g) \in \mathcal{E}'(y)$, we now construct a set $\mathcal{R}(j,g)$ of random strings as follows: first pick all random strings r for which $f(\mathsf{A}_2(y,j,g,1^n,\delta;r)) = x_y(j,g)$ is satisfied. Afterwards, if necessary, add other random strings or discard some of the picked random strings such that $|\mathcal{R}(j,g)| = \frac{2}{3}2^\ell$.

Note that

$$\begin{aligned}
&\Pr[f(\mathsf{A}_2(y,J,G,1^n,\delta;R)) = X_y \mid (J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)] \\
&\geq \Pr[\mathsf{A}(y,J,G,1^n,\delta;\cdot) \text{ is } \tfrac{2}{3}\text{-correct.} \mid (J,G) \in \mathcal{E}'(y)] \\
&\geq 1 - 8m^d\eta(y).
\end{aligned} \tag{6.19}$$

Hence, Equation (6.15) yields that

$$\Pr[f(\mathsf{A}_2(y,J,G,1^n,\delta;R)) = X_y \mid (J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)] \geq 1-\varepsilon \tag{6.20}$$

for a $(1-\varepsilon)$ fraction of the y's in $\{0,1\}^m$.

It remains to show that no family of sets $\{\mathcal{L}(y)\}_{y\in\{0,1\}^m}$ can be used to show that A has $\frac{\varepsilon^2}{50m^d}$-accessible max-entropy at most $\mathrm{E}[\log(|\mathcal{S}(Y)|)] - 5\varepsilon m^d$. Fix a family $\{\mathcal{L}(y)\}$ with $\mathrm{E}[\log(|\mathcal{L}(Y)|)] \leq \mathrm{E}[\log(|\mathcal{S}(Y)|)] - 5\varepsilon m^d$. The following claim concludes the proof of the lemma.

Claim 6.20. It holds that $\Pr[\Gamma(Y,\mathsf{A}(Y,1^n,\varepsilon;R)) \notin \mathcal{L}(Y) \cup \{\perp\}] \geq \frac{\varepsilon^2}{50m^d}$.

Proof. Note that $\log(|\mathcal{S}(y)|) - \log(|\mathcal{L}(y)|) > 2\varepsilon m^d$ holds for at least 2ε fraction of the y's in $\{0,1\}^m$ (otherwise, $\mathrm{E}\big[\log(|\mathcal{S}(Y)|) - \log(|\mathcal{L}(Y)|)\big] \leq (1-2\varepsilon)2\varepsilon m^d + 2\varepsilon\cdot(m^d+1) < 5\varepsilon n^d$). Since Equation (6.20) holds for a $(1-\varepsilon)$ fraction of the y's in $\{0,1\}^m$, there exists a set $\mathrm{Good} \subseteq \{0,1\}^m$ of density ε such that for every $y \in \mathrm{Good}$, both

$$\log(|\mathcal{L}(y)|) < \log(|\mathcal{S}(y)|) - 2\varepsilon m^d, \text{ and} \tag{6.21}$$

$$\Pr[f(\mathsf{A}_2(y,J,G,1^n,\varepsilon;R)) = X_y \mid (J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)] \geq 1-\varepsilon \tag{6.22}$$

hold. It follows that for every $y \in \mathrm{Good}$

$$|\mathcal{L}(y)| < |\mathcal{S}(y)|\,(1-1.5\varepsilon). \tag{6.23}$$

Equation (6.23) trivially holds in case $\mathcal{L}(y) = 0$, where in case $|\mathcal{L}(y)| > 0$, it holds that

i. $|\mathcal{L}(y)| < |\mathcal{S}(y)| \cdot \mathrm{e}^{-2\varepsilon\cdot\log|\mathcal{S}(y)|}$ (since $|\mathcal{S}(y)| > |\mathcal{L}(y)| \geq 1$) and

ii. $\mathrm{e}^{-2\varepsilon\cdot\log|\mathcal{S}(y)|} \leq |\mathcal{S}(y)|\cdot(1-1.5\varepsilon)$ (since $\mathrm{e}^{-\varepsilon\kappa} \leq \mathrm{e}^{-\varepsilon} < 1-0.75\varepsilon$ for $\kappa \geq 1$ and $\varepsilon < \frac{1}{2}$).

Equation (6.23) yields that

$$\begin{aligned}\Pr\big[X_y \in \mathcal{S}(y)\setminus\mathcal{L}(y) | (J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)\big] &= \Pr\big[X_y \in \mathcal{S}(y)\setminus\mathcal{L}(y) | (J,G) \in \mathcal{E}'(y)\big] \\ &\geq 1.5\varepsilon\end{aligned} \tag{6.24}$$

for every $y \in \mathrm{Good}$, and therefore Equation (6.22) yields that

$$\Pr\Big[f(\mathsf{A}_2(y,J,G,1^n,\delta;R)) \in \mathcal{S}(y)\setminus\mathcal{L}(y) \,\Big|\, (J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)\Big] \geq \frac{\varepsilon}{2} \tag{6.25}$$

for every $y \in \mathrm{Good}$.

Since $\Pr[Y \in \mathrm{Good}] \geq \varepsilon$ and since $\Pr[(J,G) \in \mathcal{E}'(y) \wedge R \in \mathcal{R}'(J,G)] \geq \frac{2}{3}(1-\varepsilon)\frac{1}{8m^d}$ for every $y \in \mathrm{Good}$, it follows that

$$\Pr[f(\mathsf{A}_2(Y,J,G,1^n,\delta;R)) \notin \mathcal{L}(y) \cup \{\perp\}] \geq \varepsilon \cdot \tfrac{2}{3} \cdot (1-\varepsilon) \cdot \frac{1}{8m^d} \cdot \frac{\varepsilon}{2} \geq \frac{\varepsilon^2}{50m^d},$$

proving the claim and thus the lemma. □

□

6.2.5 Putting it together

We now use Lemmas 6.14 and 6.19 to prove Lemma 6.9 (and as we have seen, this implies Theorem 6.5).

Proof of Lemma 6.9. Lemma 6.14 yields that $(\mathcal{Q},\mathcal{U})$ has i.o. $(\varepsilon/m)^{2d}$-accessible average max-entropy at most $v - \sqrt[2d]{\varepsilon} \cdot m^{2d}$ with respect to f, for $v = \mathrm{E}_{I,G}[\log(|\mathcal{S}_{\mathcal{Q},f}(I,G)|)]$. It follows that $(\mathcal{Q},\mathcal{U})$ has i.o. $\frac{\varepsilon}{50m^d}$-accessible average max-entropy at most $v - 5\varepsilon m^d$ with respect to f, and the proof follows by Lemma 6.19. $\square$

Acknowledgments

We are thankful to Ran Raz and Chiu-Yuen Koo for useful conversations.

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 *Proc. 18th Internat. Conf. Theory Appl. of Cryptology and Inform. Security (ASIACRYPT'12)*, pp. 154–171. Springer, 2012. [doi:10.1007/978-3-642-34961-4_11] 7

[2] KFIR BARHUM AND UELI MAURER: UOWHFs from OWFs: Trading regularity for efficiency. In *Progress in Cryptology (LATINCRYPT'12)*, pp. 234–253. Springer, 2012. [doi:10.1007/978-3-642-33481-8_13] 8

[3] MIHIR BELLARE AND PHILLIP ROGAWAY: Collision-resistant hashing: Towards making UOWHFs practical. In *Proc. 17th Ann. Internat. Cryptology Conf. (CRYPTO'97)*, pp. 470–484. Springer, 1997. [doi:10.1007/BFb0052256] 30

[4] SHAI BEN-DAVID, BENNY CHOR, ODED GOLDREICH, AND MICHAEL LUBY: On the theory of average case complexity. *J. Comput. System Sci.*, 44(2):193–219, 1992. Preliminary version in SCT'89. [doi:10.1016/0022-0000(92)90019-F] 39

[5] ANDREJ BOGDANOV AND LUCA TREVISAN: Average-case complexity. *Found. and Trends in Theoret. Comp. Sci.*, 2(1):1–106, 2006. [doi:10.1561/0400000004, arXiv:cs/0606037] 8, 37, 38, 39

[6] RAN CANETTI, RONALD L. RIVEST, MADHU SUDAN, LUCA TREVISAN, SALIL P. VADHAN, AND HOETECK WEE: Amplifying collision resistance: A complexity-theoretic treatment. In *Proc. 27th Ann. Internat. Cryptology Conf. (CRYPTO'07)*, pp. 264–283. Springer, 2007. [doi:10.1007/978-3-540-74143-5_15] 5, 34

[7] THOMAS M. COVER AND JOY A. THOMAS: *Elements of Information Theory*. Wiley-Interscience, 2006. [doi:10.1002/047174882X] 17

[8] RONALD CRAMER AND VICTOR SHOUP: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Comput.*, 33(1):167–226, 2003. Preliminary version in CRYPTO'98. [doi:10.1137/S0097539702403773] 3

[9] IFTACH HAITNER, THOMAS HOLENSTEIN, OMER REINGOLD, SALIL VADHAN, AND HOETECK WEE: Universal one-way hash functions via inaccessible entropy. In *Proc. 39th Ann. Internat. Conf. on Theory Appl. of Cryptographic Techniques (EUROCRYPT'10)*, pp. 616–637. Springer, 2010. [doi:10.1007/978-3-642-13190-5_31] 1, 7

[10] IFTACH HAITNER, MINH NGUYEN, SHIEN JIN ONG, OMER REINGOLD, AND SALIL VADHAN: Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009. [doi:10.1137/080725404] 2, 3, 7

[11] IFTACH HAITNER, OMER REINGOLD, AND SALIL VADHAN: Efficiency improvements in constructing pseudorandom generators from one-way functions. *SIAM J. Comput.*, 42(3):1405–1430, 2013. Preliminary version in STOC'10. [doi:10.1137/100814421] 7

[12] IFTACH HAITNER, OMER REINGOLD, SALIL VADHAN, AND HOETECK WEE: Inaccessible entropy. In *Proc. 41st STOC*, pp. 611–620. ACM Press, 2009. [doi:10.1145/1536414.1536497] 2, 3, 4, 7, 8

[13] IFTACH HAITNER, OMER REINGOLD, SALIL VADHAN, AND HOETECK WEE: Inaccessible entropy I: Inaccessible entropy generators and statistically hiding commitments from one-way functions, 2017. Preliminary version in STOC'09. [arXiv:2010.05586] 7

[14] IFTACH HAITNER AND SALIL P. VADHAN: The many entropies in one-way functions. In YEHUDA LINDELL, editor, *Tutorials on the Foundations of Cryptography — Dedicated to Oded Goldreich*, pp. 159–217. Springer, 2017. [doi:10.1007/978-3-319-57048-8_4, ECCC:TR17-084] 7

[15] JOHAN HÅSTAD, RUSSELL IMPAGLIAZZO, LEONID A. LEVIN, AND MICHAEL LUBY: A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999. Preliminary versions in STOC'89 and STOC'90. [doi:10.1137/S0097539793244708] 2, 3, 4, 5, 7, 10

[16] THOMAS HOLENSTEIN AND RENATO RENNER: On the randomness of independent experiments. *IEEE Trans. Inform. Theory*, 57(4):1865–1871, 2011. [doi:10.1109/TIT.2011.2110230, arXiv:cs/0608007] 10, 11

[17] PAVEL HUBÁČEK, MONI NAOR, AND EYLON YOGEV: The journey from NP to TFNP hardness. In *Proc. 8th Innovations in Theoret. Comp. Sci. (ITCS'17)*, pp. 60:1–60:21. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. [doi:10.4230/LIPIcs.ITCS.2017.60] 8

[18] RUSSELL IMPAGLIAZZO AND LEONID LEVIN: No better ways to generate hard NP instances than picking uniformly at random. In *Proc. 31st FOCS*, pp. 812–821. IEEE Comp. Soc., 1990. [doi:10.1109/FSCS.1990.89604] 2, 6, 8, 37, 38, 39

[19] RUSSELL IMPAGLIAZZO AND MICHAEL LUBY: One-way functions are essential for complexity based cryptography. In *Proc. 30th FOCS*, pp. 230–235. IEEE Comp. Soc., 1989. [doi:10.1109/SFCS.1989.63483] 5

[20] JONATHAN KATZ AND CHIU-YUEN KOO: On constructing universal one-way hash functions from arbitrary one-way functions. Cryptology ePrint Archive, Report 2005/328, 2005. 3, 7, 20, 31

[21] LEONID A. LEVIN: Average case complete problems. *SIAM J. Comput.*, 15(1):285–286, 1986. Preliminary version in STOC'84. [doi:10.1137/0215020] 8

[22] MING LI AND PAUL M. B. VITÁNYI: Average case complexity under the universal distribution equals worst-case complexity. *Inform. Process. Lett.*, 42(3):145–149, 1992. [doi:10.1016/0020-0190(92)90138-L] 38

[23] MONI NAOR AND MOTI YUNG: Universal one-way hash functions and their cryptographic applications. In *Proc. 21st STOC*, pp. 33–43. ACM Press, 1989. [doi:10.1145/73007.73011] 2, 3, 30

[24] NOAM NISAN AND DAVID ZUCKERMAN: Randomness is linear in space. *J. Comput. System Sci.*, 52(1):43–52, 1996. [doi:10.1006/jcss.1996.0004] 9

[25] RENATO RENNER AND STEFAN WOLF: Smooth Rényi entropy and applications. In *Internat. Symp. Inform. Theory (ISIT'04)*, p. 233. IEEE Comp. Soc., 2004. [doi:10.1109/ISIT.2004.1365269] 9

[26] JOHN T. ROMPEL: One-way functions are necessary and sufficient for secure signatures. In *Proc. 22nd STOC*, pp. 387–394. ACM Press, 1990. [doi:10.1145/100216.100269] 2, 3, 5, 6, 7, 15, 20, 31

[27] JOHN T. ROMPEL: *Techniques for Computing With Low-Independence Randomness*. Ph. D. thesis, Massachusetts Institute of Technology, 1990. ACM DL. 3, 7

[28] ALFREDO DE SANTIS AND MOTI YUNG: On the design of provably secure cryptographic hash functions. In *Proc. 19th Ann. Internat. Conf. on Theory Appl. of Cryptographic Techniques (EUROCRYPT'90)*, pp. 412–431. Springer, 1990. [doi:10.1007/3-540-46877-3_37] 3

[29] VICTOR SHOUP: A composition theorem for universal one-way hash functions. In *Proc. 29th Ann. Internat. Conf. on Theory Appl. of Cryptographic Techniques (EUROCRYPT'00)*, pp. 445–452. Springer, 2000. [doi:10.1007/3-540-45539-6_32] 30

[30] SALIL VADHAN: Computational entropy. In ODED GOLDREICH, editor, *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pp. 693–726. ACM Press, 2019. [doi:10.1145/3335741.3335767] 7

[31] LESLIE G. VALIANT AND VIJAY V. VAZIRANI: NP is as easy as detecting unique solutions. *Theoret. Comput. Sci.*, 47(1):85–93, 1986. Preliminary version in STOC'85. [doi:10.1016/0304-3975(86)90135-0] 40

[32] YU YU, DAWU GU, XIANGXUE LI, AND JIAN WENG: (Almost) optimal constructions of UOWHFs from 1-to-1, regular one-way functions and beyond. In *Proc. 35th Ann. Internat. Cryptology Conf. (CRYPTO'15)*, pp. 209–229. Springer, 2015. [doi:10.1007/978-3-662-48000-7_11] 8

AUTHORS

Iftach Haitner
Tel Aviv University
Tel Aviv Israel
iftachh@tau.ac.il
www.cs.tau.ac.il/~iftachh

Thomas Holenstein
Google
Zurich, Switzerland
thomas.holenstein@gmail.com

Omer Reingold
Stanford University
Stanford, California
reingold@stanford.edu
engineering.stanford.edu/people/omer-reingold

Salil Vadhan
Harvard University
Cambridge, Massachusetts
salil_vadhan@harvard.edu
people.seas.harvard.edu/~salil

Hoeteck Wee
ENS
Paris, France
wee@di.ens.fr
www.di.ens.fr/~wee

ABOUT THE AUTHORS

IFTACH HAITNER is an associate professor at the School of Computer Science at Tel Aviv University and the director of the Check Point Institute for Information Security. He received his Ph. D. under the supervision of Omer Reingold at the Weizmann Institute in 2008; the title of his dissertation was "New Implications and Improved Efficiency of Constructions Based on One-way Functions." He enjoys doing playback theater and swimming.

THOMAS HOLENSTEIN obtained his Ph. D. from ETH Zürich in 2006, advised by Ueli Maurer. He then spent two years as a postdoc at Microsoft Research, and one year as a postdoc at Princeton University. After this, he was a professor at ETH Zürich for 6 years, and then joined Google, where he currently works on the cryptographic library *Tink*.

OMER REINGOLD is the Rajeev Motwani Professor of Computer Science at Stanford University. Omer received his Ph. D. in 1999 from the Weizmann Institute under the direction of Moni Naor. Omer's past positions include Samsung Research America, the Weizmann Institute of Science, Microsoft Research, the Institute for Advanced Study in Princeton, NJ and AT&T Labs. His research is in the Foundations of Computer Science and most notably in Computational Complexity and the Foundations of Cryptography with emphasis on randomness, derandomization and explicit combinatorial constructions. He has a keen interest in the societal impact of computation, as reflected in his work on privacy and algorithmic fairness.

Omer is an ACM Fellow and among his distinctions are the 2005 Grace Murray Hopper Award and the 2009 Gödel Prize.

SALIL VADHAN is the Vicky Joseph Professor of Computer Science and Applied Mathematics at the Harvard John A. Paulson School of Engineering & Applied Sciences. He received his Ph. D. under the supervision of Shafi Goldwasser at MIT in 1999; the title of his dissertation was "A Study of Statistical Zero-Knowledge Proofs." His other research interests include the theory of pseudorandomness and the theory and practice of data privacy. He enjoys spending leisure time with his wife and two daughters, as well as learning to surf in the cold waters of New England.

HOETECK WEE joined ENS as a CNRS researcher in 2013, after teaching in the US for several years. He obtained his Ph. D. from UC Berkeley in 2007 under the supervision of Luca Trevisan. He is the recipient of a NSF CAREER Award, a Humboldt Research Fellowship, a Google Faculty Research Award, an ERC Starting Grant, and the Best paper award at Eurocrypt 2016; he is also a contributor to the TLS 1.3 standard.