Columbia University - Crypto Reading Group Apr 27, 2011

Inaccessible Entropy and its Applications

Igor Carboni Oliveira

We summarize the constructions of PRGs from OWFs discussed so far and introduce the notion of inaccessible entropy [HILL99, HRVW09].

1 Review: Psedorandom Generators from One-Way Functions

Remember that we are trying to construct objects that look random (PRGs) from an assumption about hardness of computation (OWFs). So far we have seen that it is possible to construct PRGs from OWFs if the OWF has some nice structural property.

One-way Permutations. Goldreich-Levin hardcore predicate [GL89] combined with the equivalence between next-bit unpredictability and computational indistinguishability [Yao82].

1-1 One-way Functions. Pairwise-independent hash families, leftover hash lemma. See description in [Gol01]. Note: we have not covered this construction.

Regular One-Way Functions. One construction is based on a reduction (again using pairwise independence) to the 1-1 case [GKL93]. However, we have only discussed the relatively recent construction using the randomized iterate method [HHR05].

Exponentially Hard One-Way Functions. We discussed the recent proof based on the randomized iterate method.

Arbitrary One-Way Functions. The construction of PRGs is much more difficult in this case. Known constructions rely on computational notions of entropy: psudoentropy [HILL99] and next-block pseudoentropy [HRV10] (which is inspired by the notion of inaccessible entropy [HRVW09]).

2 The Shannon Entropy Function [Sh48]

Let Ω be a probability distribution and suppose that X is a random variable defined over Ω. Intuitively, the entropy of the random variable X measures the number of random bits (or the amount of uncertainty) that is possible to extract (on average) when we sample an element r according to Ω an have access to the value $X(r)$.

Definition 1. [Entropy of a Random Variable/Distribution]
Let X be a random variable taking values in some range $B \subseteq \mathbb{R}$. The entropy of X, denoted by $H[X]$, is defined by

$$H[X] = \sum_{b \in B} - \Pr[X = b] \log_2 \Pr[X = b],$$

where $0 \log_2 0$ *is interpreted as* 0*.*

Definition 2. [Conditional Entropy]
If E *is some event over* Ω*, then we can define the conditional entropy of* X *given* E *by*

$$H[X|E] = \sum_{b \in B} - \Pr[X = b \mid E\ ] \log_2 \Pr[X = b \mid E\ ].$$

In other words, this is just the entropy of the random variable X over the new probability space obtained from Ω under the assumption that event E happened.

Definition 3. *If* Y *is a random variable taking values in some range* A*, we define the conditional entropy of* X *given* Y *by*

$$H[X|Y] = \sum_{a \in A} H[X \mid Y = a\ ] \Pr[Y = a].$$

Useful Inequalities. The following results follow from the definition of entropy and the convexity of the function $x \log_2 x$. It is interesting to think about each one of them using the intuitive interpretation of the entropy function as the amount of randomness that we can extract on average from the random variable.

(a) [*uniform distribution has maximum entropy*]. If $|B| = 2^n$, then $0 \leq H[X] \leq n$.

(b) [*range-size bound*]. In general, if $S \subseteq B$, then $\sum_{b \in S} - \Pr[X = b] \log_2 \Pr[X = b] \leq \log_2 |S|$.

(c) [*no randomness from determinism*]. $H[X|Y] = 0$ if and only if $X = f(Y)$ for some function f.

(d) [*independence rule*]. $H[X|Y] = H[X]$ if and only if X and Y are independent random variables.

(e) [*subadditivity*]. $H[X, Y] \leq H[X] + H[Y]$, where X,Y denotes the random variable $\langle X, Y \rangle$.

(f) [*chain rule*]. $H[X, Y] = H[Y] + H[X|Y]$.

(g) [*elimination inequality*]. $H[X|Y, Z] \leq H[X|Y]$.

(h) [*conditional entropy bounds*]. $H[X] - H[Y] \leq H[X|Y] \leq H[X]$.

(i) [*inverting formula*]. $H[X|Y] = H[Y|X] + H[X] - H[Y]$.

(j) [*entropy of a function of a random variable*]. $H[f(X)] \leq H[X]$ for any function f.

We will refer to $H(X)$ as the real entropy of X.

3 Pseudoentropy [HILL99]

The entropy function defined in the previous section is an information-theoretic notion. To prove the equivalence between PRGs and OWFs, [HILL99] introduced a computational analogue of entropy known as *pseudoentropy*. It tries to capture the idea that a distribution may behave like one of higher entropy.

Definition 4. [Pseudoentropy]
A random variable X *has pseudoentropy (at least)* k *if there exists a random variable* Y *of entropy (at least)* k *such that* X *and* Y *are computationally indistinguishable.*

The output of a pseudorandom generator $G : \{0,1\}^l \to \{0,1\}^n$ on a uniformly random seed has entropy at most l by the range-size bound. However, its pseudoentropy is n by the definition of pseudorandom generator.

Hastad et al. used this notion to prove the equivalence between PRGs and OWFs.

Theorem 5. [HILL99] *It is possible to construct from any one-way function an efficiently samplable distribution X whose pseudoentropy is noticeably larger than its real entropy.*

Theorem 6. [HILL99] *From any efficiently samplable distribution X whose pseudoentropy is noticeably larger than its real entropy, it is possible to construct a pseudorandom generator.*

4 Inaccessible Entropy [HRVW09, HHRVW10]

In [HRVW09] Haitner et al. put forth a new computational notion of entropy that tries to capture instead the situation where a distribution can behave like one of much lower entropy.

4.1 Inaccessible Entropy of Protocols

Consider the following 3-message protocol between parties (A, B):

1. B selects a random function $h : \{0,1\}^n \to \{0,1\}^m$ from a family of collision-resistant hash functions (where $m \ll n$) and sends h to A.

2. A selects a random x from $\{0,1\}^n$, sets $y = h(x)$, and sends y to B.

3. A sends x to B.

Now let's consider the entropy of A's messages during the different rounds of the protocol (i.e, we condition according to the history of communication). Initially, X (random bits of A) is a random variable uniformly distributed over the set of n-bit strings. Hence $H[X] = n$ before A receives the first message from B. After B reveals the random function f to A, we still have $H[X|H] = H[X] = n$, since X and H are independent random variables (remember the independence property of the entropy function). In the next message B receives $Y = h(X)$ from A. The entropy of A is now:

$$H[X|Y] = H[X|h(X)] \geq H[X] - H[h(X)] = n - H[h(X)] \geq n - m,$$

where the first inequality follows from the conditional entropy bound and the second inequality follows from the range-size bound ($h(X)$ is a string of size 2^m). Since by assumption $m \ll n$, we see that the entropy (amount of uncertainty) associated to the random bits of A is still close to n.

Now consider an adversary (cheating strategy) $\hat{A}$ that, using the current view/history of the protocol, tries to deviate from the behaviour of the honest party A. Perhaps due to some new external information, $\hat{A}$ may want to come up with a different x' such that $h(x) = h(x')$ to send to B as the final message. However, by the collision-resistant property of h, $\hat{A}$ will only succeed with negligible probability. Hence, although there is real entropy in X (random bits of A) conditioned on the previous messages exchanged between the parties, this entropy is computationally inaccessible to $\hat{A}$, to whom X has entropy 0 for all practical purposes.

[HRVW09]: “Thus, in constrast to pseudoentropy, accessible entropy is useful for expressing the idea that the computational entropy in a distribution is *smaller* than its real entropy. We refer to the difference (real entropy) − (accessible entropy) as the *inaccessible entropy* of the protocol.”

See the paper [HRVW09] for a formal definition of the real and accessible entropy of a protocol.

4.1.1 Application: Construction of Statistically Hiding Commitment Schemes

Commitment schemes are important for the construction of zero-knowledge proof systems for NP (see Goldreich’s book for details).

[HRVW09]: “In this section we present the main technical contribution of this paper, showing how any protocol with a noticeable gap between its real and accessible entropies can be converted into a *statistically* hiding and computationally binding commitment scheme.”

We note that it is not very difficult to construct *computationally* hiding commitment schemes assuming the existence of one-way functions.

Theorem 7. [HRVW09] One-Way Function to Entropy Gap, informal.
Given any one-way function $f : \{0,1\}^n \to \{0,1\}^n$, we can construct an $O(n/\log n)$-round protocol (A, B) in which the real entropy of A’s messages is noticeably larger than their accessible (max-)entropy.

Theorem 8. [HRVW09] Entropy Gap to Commitment, informal.
If there is an efficient protocol (A, B) in which the real entropy of A’s messages is noticeably larger than their accessible entropy, then statistically hiding commitment schemes exist.

4.2 The Inaccessible Entropy of Hash Functions [HHRVW10]

Universal one-way hash functions (UOWHF) are a weaker form of collision-resistant hash functions. In particular, any collision-resistant hash function is a UOWHF. For simplicity, we will use the following equivalent variant of UOWHF.

Definition 9. [Universal One-Way Hash Functions]
Consider a sequence of functions $\mathcal{F}_k : \{0,1\}^{n(k)} \to \{0,1\}^{m(k)}$, where $k \in \mathbb{N}$ (security parameter). For simplicity, we will often omit the index k and write $F : \{0,1\}^n \to \{0,1\}^m$. We say that F is a universal one-way hash function if:

i) F can be efficiently computed.

ii) F is shrinking, i.e., $n(k) = O(poly(k))$ and $m(k) < n(k)$.

iii) F satisfies a weaker notion of collision resistance. Consider the following game:

An adversary A_F is given a random $x \leftarrow \{0,1\}^n$, and outputs x' such that $F(x) = F(x')$. The adversary A_F succeeds if and only if $x \neq x'$.

We require that for all PPT *adversaries A_F, the success probability of A_F is negligible.*

Note that we can assume w.l.o.g. that the adversary always output an x' such that $F(x) = F(x')$. If A_F is a (not necessarily efficient) probabilistic algorithm with this property, we say that A_F is an F-collision finder.

It is possible to convert any such F into a standard family of universal one-way hash functions by setting $F_z(x) = F(z + x)$. An important application of UOWHFs is the construction of secure digital signatures schemes. We note that it is possible to construct OWFs from any UOWHF. On the other hand, the original construction of UOWHF from OWF is very complicated [Rom90]. The paper [HHRVW10] presents a new construction using inaccessible entropy.

Consider the (rather inefficient) probabilistic algorithm A_F^* that, on input x, outputs a uniform $x' \leftarrow F^{-1}(F(x))$. Let X be a random variable uniformly distributed on $\{0,1\}^n$.

Claim 10. *$H[A_F^*(X,R)|X] = H[X|F(X)] \geq n - m$.*

Proof. For convenience we denote the random variable $A_F^*(X,R)$ by $A_F^*(X)$, i.e., we omit the random bits R used by algorithm A^*. Note that, conditioning on the random variable on the right-hand side of each expression $H[\ .\ |\ .\ ]$, we have the same distribution on the left-hand side. Hence, to prove the equality, it is enough to collect conditional probabilities. The inequality follows from basic properties of the entropy function. Formally, we have:

$$\begin{aligned}
H[A_F^*(X)|X] &= H[Z \leftarrow F^{-1}(F(X))|X] \\
&= \sum_{x \in \{0,1\}^n} H[Z \leftarrow F^{-1}(F(X))|X = x] \Pr_{X \leftarrow U_n}[X = x] \\
&= \sum_{x \in \{0,1\}^n} H[Z \leftarrow F^{-1}(F(x))] \Pr_{X \leftarrow U_n}[X = x] \\
&= \sum_{y \in Im(F)} H[Z \leftarrow F^{-1}(y)] \Pr_{X \leftarrow U_n}[F(X) = y] \\
&= \sum_{y \in Im(F)} H[X|F(X) = y] \Pr_{X \leftarrow U_n}[F(X) = y] \\
&= H[X|F(X)] \\
&\geq H[X] - H[F(X)] \\
&\geq n - m.
\end{aligned}$$

□

Definition 11. [Real Entropy of a Function]
The real entropy of F is defined as $H[X|F(X)]$.

In other words, the real entropy of F is defined as the expected entropy left in the input after revealing the output. It follows from the preceding claim that the amount by which F shrinks is a lower bound on its real entropy.

Now remember that, given a probabilistic polynomial-time F-collision finder A_F, by the collision resistance property of F, we have $\Pr[A_F(X) \neq X] = negl(n)$. This is actually equivalent to require that $H[A_F(X)|X] = negl(n)$. In what follows we prove one direction of this equivalence.

Claim 12. *If for every probabilistic polynomial-time F-collision finder B_F we have $H[B_F(X)|X] = negl(n)$, then for every probabilistic polynomial-time F-collision finder A_F we have $\Pr[A_F(X) \neq X] = negl(n)$.*

Proof. We prove the contrapositive. Suppose there exists an A_F such that $\Pr[A_F(X) \neq X] \geq 1/p(n)$ for infinitely many values of n. It is not necessarily the case that $H[A_F(X)|X]$ is non-negligible, since A_F may map several different inputs to the same pre-image x', resulting in non-negligible probability success but only negligible conditional entropy. Using A_F, we construct an efficient F-collision finder B_F with non-negligible entropy as follows:

Construction of B_F. Given an input string X, let $X' = A_F(X)$. Algorithm B_F simply toss a fair coin and output X or X' with probability exactly $1/2$.

Clearly, B_F is an efficient probabilistic algorithm. To finish the proof, we need to prove that $H[B_F(X)|X]$ is non-negligible. Consider the following set of good input strings x:

$$S = \{x \in \{0,1\}^n \ : \ A_F(x) \neq x \text{ with probability at least } \tfrac{1}{2p(n)}\},$$

where the probability is taken over the random bits of A_F. By a standard Markov argument, it follows that $|S| \geq 2^n/(2p(n))$. For any given input string x, let R_x be an indicator random variable such that $R_x(r) = 1$ if and only if $A_F(x,r) \neq x$. The result now follows from basic properties of the entropy function:

$$\begin{aligned}
H[B_F(X)|X] &= \sum_{x\in\{0,1\}^n} H[B_F(X)|X = x] \Pr_{X\leftarrow U_n}[X = x] \\
&= 2^{-n} \sum_{x\in\{0,1\}^n} H[B_F(x)] \\
&= 2^{-n}(\sum_{x\notin S} H[B_F(x)] + \sum_{x\in S} H[B_F(x)]) \\
&\geq 2^{-n}(0 + \sum_{x\in S} H[B_F(x)|R_x]) \\
&= 2^{-n} \sum_{x\in S}(H[B_F(x)|R_x = 0]) \Pr[R_x = 0] + H[B_F(x)|R_x = 1]) \Pr[R_x = 1]) \\
&\geq 2^{-n} \sum_{x\in S} H[B_F(x)|R_x = 1]) \frac{1}{2p(n)} \\
&\geq \frac{2^{-n}}{2p(n)} \sum_{x\in S} 1 \geq \frac{2^{-n}}{2p(n)} \frac{2^n}{2p(n)} = \frac{1}{4p(n)^2}.
\end{aligned}$$

□

Definition 13. [Accessible Entropy of a Function]
Let $F : \{0,1\}^n \to \{0,1\}^m$ be a function. We refer to the maximum of $H[A_F(X)|X]$ over all efficient F-collision finders A_F as the accessible entropy of F. Formally, we say that F has accessible entropy at most k if, for all efficient F-collision finders A_F, we have $H[A_F(X)|X] \leq k$ for all sufficiently large n.

[HHRVW10]: "Thus, a natural weakening of the UOWHF property is to simple require a noticeable gap between the real and accessible entropies of F. That is, for every probabilistic polynomial-time

F-collision finder A_F, we have $H[A_F(X)|X] < H(X|F(X)) - \triangle$, for some noticeable $\triangle$, which we refer to as the *inaccessible entropy* of F." (note: the notation used in the original quote is slightly different).

4.2.1 Constructing UOWHFs via Inaccessible Entropy

Conceptually, the proof that OWF implies UOWHF is similar to the construction of statistically hiding commitment schemes from OWF. First, given any one-way function f, we construct a function F that is efficiently computed and have noticeable inaccessible entropy. After that, using several different tools, we obtain a new function that in addition has the shrinking property and a large entropy gap (enough for us to apply claim 12 and obtain the weaker collision resistance property).

We describe the first part of this construction (from any one-way function to noticeable entropy gap). This will set the stage for the discussion of the last paper of our original list of papers: how to construct PRGs from arbitrary OWFs [HRV10].

Theorem 14. [From one-way functions to noticeable entropy gap]
Let $f : \{0,1\}^n \to \{0,1\}^n$ be a one-way function. Let $F : \{0,1\}^n \times [n] \to \{0,1\}^{i \leq n}$ be the function defined by $F(x,i) = f(x)_{1 \ldots i}$, i.e., $F(x,i)$ truncates the output of $f(x)$ to the first i bits. Then F has accessible entropy at most $H(Z|F(Z)) - \frac{1}{2^{12}n^4 \log^2 n}$, where $Z = (X,I)$ is uniformly distributed over $\{0,1\}^n \times [n]$.

Proof Idea. Remember that the entropy of F is the entropy of the perfect F-collision finder A^* that, given an input $z = (x,i)$, always samples uniformly from the set of pre-images of $F(z)$. Now suppose towards a contradiction that there is an *efficient* F-collision finder A such that the entropy $H[A(Z)|Z]$ of A is very close to $H(A^*(Z)|Z)$, implying an entropy gap smaller than $\triangle = \frac{1}{2^{12}n^4 \log^2 n}$. The following lemma can be used to argue that A is *almost* as good as A^*.

Lemma 15. *If W is a random variable whose support is contained in a set S and U is the uniform distribution on S, then $||U - W|| \leq \sqrt{H[U] - H[W]}$, where $||U - W||$ denotes the statistical distance between U and W.*

So from now on we make the simplifying assumption ($\star$) that we have an *efficient* F-collision finder A that is *as good* as A^* (based on the fact that a statistical distance bounded by $\sqrt{\triangle}$ is sufficiently small for our purposes). Remember that we want to invert the one-way function f under the assumption that we have an A that is able to sample uniformly from the pre-images of $F(x,i)$ when it has access to the *input* (x,i). It is not immediately clear whether this is enough to invert f if we are only given access to the *output* $y = f(x) = F(x,n)$. We will use A to construct an inverter B that does substantially more than inverting f.

Consider the following game: instead of selecting $x \in \{0,1\}^n$ at random and presenting $y = f(x)$ to B, algorithm B is given a string $y = y_1 \ldots y_n$ bit by bit, and we require that for every $0 \leq i \leq n$, B is able to output an $x^{(i)} \in \{0,1\}^n$ such that $f(x^{(i)})_{1 \ldots i} = y_{1 \ldots i}$. In other words, B succeeds at stage i of the game if it is able to "invert" the first i bits of y. We say that B succeeds on input $y \in \{0,1\}^n$ if B succeeds at every stage of the game when it is given the bits of y as input. Clearly, if B is able to win this game with non-negligible probability, then it is possible to invert f (in the original sense) with non-negligible probability.

Note that if $i = 0$ then any random input string $x^{(0)}$ is a good string for B. Thus let's assume that B has succeeded in the previous i rounds of the game. In other words, B knows a string $x^{(i)}$ such that

$f(x^{(i)})_{1\dots i} = y_{1\dots i}$ and is now given the next bit y_{i+1}. By our assumption about A, we can run this algorithm on input $(x^{(i)}, i)$ to sample uniformly at random from the following set of strings:

$$F^{-1}(F(x^{(i)}, i)) = F^{-1}(y_{1\dots i}) = \{(x, i) \mid x \in \{0,1\}^n \text{ and } f(x)_{1\dots i} = y_{1\dots i}\}.$$

In other words, there is an efficient subroutine that allows us to sample uniformly at random from the set:

$$S_{y_{1\dots i}} = \{x \in \{0,1\}^n \mid f(x)_{1\dots i} = y_{1\dots i}\}.$$

To succeed in the $i+1$ round, B has to find an $x \in S_{y_{1\dots i}}$ that in addition satisfies $f(x)_{i+1} = y_{i+1}$. We sample enough times from S until we get a good x. Formally, let:

$$\begin{aligned} p_{\langle y_{1\dots i},0\rangle} &= \Pr_{x \leftarrow S_{y_{1\dots i}}}[f(x)_{i+1} = 0] \\ p_{\langle y_{1\dots i},1\rangle} &= \Pr_{x \leftarrow S_{y_{1\dots i}}}[f(x)_{i+1} = 1] \end{aligned}$$

When do we fail? We fail only if, say, we need to have $f(x)_{i+1} = 0$ but the probability p_0 that this event happens is very small. Fortunately, if p_0 is small, then it must be the case that we get $y_{i+1} = 0$ with very low probability. It follows from the definition of the set $S_{y_{1\dots i}}$ that:

$$\begin{aligned} p_{\langle y_{1\dots i},0\rangle} &= \Pr_{X \leftarrow U_n,\ Y = f(X)}[Y_{i+1} = 0 \mid Y_{1\dots i} = y_{1\dots i}] \\ p_{\langle y_{1\dots i},1\rangle} &= \Pr_{X \leftarrow U_n,\ Y = f(X)}[Y_{i+1} = 1 \mid Y_{1\dots i} = y_{1\dots i}] \end{aligned}$$

Therefore, conditioning on $Y_{1\dots i} = y_{1\dots i}$, what is the probability that we fail in round $i+1$ if we sample independently at random k strings from $S_{y_{1\dots i}}$? For convenience, let's write $p_0 = p_{\langle y_{1\dots i},0\rangle}$ and $p_1 = p_{\langle y_{1\dots i},1\rangle}$. Then

$$\begin{aligned} \Pr[\mathit{Fail}(i+1)] &= \Pr[\mathit{Fail}(i+1) \mid Y_{i+1} = 0]. \Pr[Y_{i+1} = 0] + \Pr[\mathit{Fail}(i+1) \mid Y_{i+1} = 1]. \Pr[Y_{i+1} = 1] \\ &= (p_1)^k p_0 + (p_0)^k p_1 \\ &= (1-p_0)^k p_0 + (1-p_1)^k p_1. \end{aligned}$$

Fortunately, regardless of the values p_0 and p_1, we can make this probability sufficiently small by setting $k = 16n \log n$.

$$\begin{aligned} (1-p_0)^k p_0 + (1-p_1)^k p_1 &\le \frac{p_0}{e^{p_0 k}} + \frac{p_1}{e^{p_1 k}} \\ &= \frac{p_0}{n^{16np_0}} + \frac{p_1}{n^{16np_1}} \\ &\le \frac{1}{4n} + \frac{1}{4n} = \frac{1}{2n}, \end{aligned}$$

where in the last inequality it is enough to consider the cases $q \le \frac{1}{4n}$ and $q > \frac{1}{4n}$ for each $q \in \{p_0, p_1\}$. To sum up, we fail in a specific round i (for any i and for any previous bits $y_1, \dots, y_i$) with probability at most $\frac{1}{2n}$. By an union bound, B as described fails for any given input y with probability at most $\frac{1}{2}$.

If given that $Y_{1\dots i} = y_{1\dots i}$ the next bit Y_{i+1} is 0 with probability p_0 and 1 with probability p_1, then it follows that in this game $Y = Y_1 \dots Y_n$ is distributed exactly as $Y = f(U_n)$. Thus B can be used to invert f with constant probability, contradicting the assumption that f is a one-way function. Hence there is no efficient F-collision finder algorithm A with entropy close to $H[Z|F(Z)]$. It follows that F has entropy gap of at least $\triangle$.

Finally, we note that in order to prove that this construction works without relying on the $(\star)$ assumption, we need to proceed more carefully because A might not perform so well when conditioned on previous biased events. The analysis is completed using hybrid distributions.

□

Theorem 16. [From noticeable entropy gap to UOWHF]
Let $F : \{0,1\}^\lambda \to \{0,1\}^m$ be an efficiently computable function with a noticeable gap $\triangle$ between its real and accessible entropy. Then there exists a family of universal one-way hash functions with output length $O(\lambda^8 s^2/\triangle^7)$ and input length $O(\log n \ \lambda^8 s^2/\triangle^7)$ for any $s = \omega(\log n)$.

Proof Idea. The entropy gap is amplified by the function F^t, where $t \in poly(n)$. We apply standard hashing techniques along the way to reduce the output length and obtain the shrinking property.

□

Combining the previous theorems, we can construct from any one-way function $f : \{0,1\}^n \to \{0,1\}^n$ (where n is the security parameter) a family of universal one-way hash functions of output length $O(s^2 n^{36} \log^{14} n)$ and key length $O(s^2 n^{36} \log^{15} n)$.