Lecture Notes 10

23 Analyzing the Success Probability

We now complete the proof of Lemma 3 of section 21. Recall again that f is a strongly one-way and length preserving function and that

$$\begin{aligned} g(x,r) &\stackrel{\text{df}}{=} (f(x),r) \\ b(x,r) &\stackrel{\text{df}}{=} x \cdot r \bmod 2. \end{aligned}$$

Assuming that b is not a hard core for g, there is a p.p.t. algorithm G and a polynomial $p(n)$ such that G predicts b with advantage

$$\varepsilon(n) \stackrel{\text{df}}{=} \varepsilon_G(n) \geq \frac{1}{p(n)} \tag{1}$$

for all n in an infinite set N. In section 22.2 of lecture 9, we constructed an algorithm A for inverting f. We now show that A has success probability at least $\frac{1}{p(n)}$ at inverting f on length-n inputs, for all $n \in N$. This contradicts the assumption that f is strongly one-way and completes the proof of Lemma 3.

Assume for the rest of this discussion that $n \in N$. Let

$$s(x) = \Pr[G(f(x), R_n) = b(x, R_n)].$$

Here R_n is a uniformly distributed random variable over length-n strings, distinct from the identically distributed random variables U_n and X_n, which we also mention from time to time. Thus, $s(x)$ is the fine-grained success probability of G for each particular length-n string x. We know that the average of $s(x)$ taken over all length-n strings x is the overall success probability of G, so

$$\frac{\sum_x s(x)}{2^n} = \frac{1}{2} + \varepsilon(n). \tag{2}$$

Define

$$S_n = \left\{ x \in \{0,1\}^n \mid s(x) \geq \frac{1}{2} + \frac{\varepsilon(n)}{2} \right\}. \tag{3}$$

Claim 1 $|S_n| \geq \varepsilon(n) \cdot 2^n$.

Three different proofs of this claim are given in handout 3: One is algebraic, one is geometric, and one is based on Markov's inequality. We do not repeat them here but refer the reader to the handout. We only mention that all three are based on the idea that in order for the average value of $s(x)$ to exceed $\frac{1}{2} + \varepsilon(n)$, there must be a certain number of x for which $s(x) \geq \frac{1}{2} + \frac{\varepsilon(n)}{2}$. That number turns out to be $\varepsilon(n) \cdot 2^n$.

Claim 2 $\forall x \in S_n, \forall i \in \{1, \ldots, n\}$,

$$\Pr[|\{J \mid b(x, r^J) \oplus G(f(x), r^J \oplus e^i) = x_i\}| > \frac{1}{2}(2^\ell - 1)] > 1 - \frac{1}{2n}.$$

Proof: Let $x \in S_n$ and $i \in \{1, \ldots, n\}$. Let ζ^J be a random variable ranging over $\{0,1\}$ such that

$$\begin{aligned} \zeta^J = 1 \quad &\text{iff} \quad b(x, r^J) \oplus G(f(x), r^J \oplus e^i) = x_i \\ &\text{iff} \quad G(f(x), r^J \oplus e^i) = b(x, r^J \oplus e^i). \end{aligned}$$

Thus, $\zeta^J = 1$ whenever G succeeds at computing x_i. Note that r^J and $r^J \oplus e^i$ are uniformly distributed over $\{0,1\}^n$. Hence, by 1 and 3,

$$\Pr[\zeta^J = 1] = s(x) \geq \frac{1}{2} + \frac{\varepsilon(n)}{2} \geq \frac{1}{2} + \frac{1}{2p(n)}. \tag{4}$$

A key observation is that the ζ^J's are pairwise independent. This follows from the fact that the r^J's are pairwise independent. To see this, let $J \neq K$. Without loss of generality, we can choose $k \in K - J$. By definition, r^K and r^J are both sums of subsets of the independent random variables $\{s^1, \ldots, s^\ell\}$. Since $k \in K - J$, the term s^k appears in the sum for r^K but not for r^J. Therefore, s^k is independent of r^J, which implies that r^K is independent of r^J.

Let $m = 2^\ell - 1$. Since $\ell = \lceil \log_2(2n \cdot p(n)^2 + 1) \rceil$, we have that $2^\ell \geq 2^{\log_2(2n \cdot p(n)^2 + 1)} = 2n \cdot p(n)^2 + 1$, so $m \geq 2n \cdot p(n)^2$. We also have

$$2^\ell \leq 2^{\log_2(2n \cdot p(n)^2 + 1) + 1} = 2 \cdot (2n \cdot p(n)^2 + 1) \tag{5}$$

We use Chebyshev's inequality to bound $\Pr[\sum_J \zeta^J \leq \frac{1}{2} \cdot m]$. This is an upper bound on the probability that the majority value z_i that algorithm A computes for x_i is wrong. Recall Chebyshev's inequality

$$\Pr[|X - E(X)| \geq \delta] \leq \frac{\mathrm{Var}(X)}{\delta^2}. \tag{6}$$

Let $X = \sum_J \zeta^J$ and $\delta = \frac{m}{2p(n)}$. All of the ζ^J are identically distributed, so we drop the superscript in the following.

$$E(\zeta) = \Pr[\zeta = 1] \geq \frac{1}{2} + \frac{1}{2p(n)}$$

by equation 4. Thus,

$$E(X) = m \cdot E(\zeta) \geq \left(\frac{1}{2} + \frac{1}{2p(n)} \right) \cdot m. \tag{7}$$

We also have

$$\mathrm{Var}(\zeta) = E(\zeta^2) - E(\zeta)^2$$

Since ζ is 0-1 valued, it follows that $\zeta^2 = \zeta$. Hence,

$$\mathrm{Var}(\zeta) = E(\zeta) - E(\zeta)^2 = E(\zeta)(1 - E(\zeta)) \leq \frac{1}{4}$$

The bound of $1/4$ simply reflects the fact that the maximum value of the function $x(1-x)$, which is reached for $x = 1/2$. Since the variables ζ^J are pairwise independent, they are also uncorrelated, so

$$\mathrm{Var}(X) = \mathrm{Var}\left(\sum_J \zeta^J) \right) = \sum_J \mathrm{Var}(\zeta^J) \leq \frac{m}{4}. \tag{8}$$

Plugging the expression for δ into inequality 6 and doing some calculations using inequalities 7 and 8, we get

$$\begin{aligned}
\Pr\left[X \le \frac{1}{2}\cdot m\right] &\le \Pr\left[\left|X - \left(\frac{1}{2} + \frac{1}{2p(n)}\right)\cdot m\right| \ge \frac{1}{2p(n)}\cdot m\right] \\
&\le \frac{\mathrm{Var}(X)}{(\frac{m}{2p(n)})^2} \\
&\le \frac{m}{4}\cdot\frac{(2p(n))^2}{m^2} = \frac{p(n)^2}{m} \\
&\le \frac{p(n)^2}{2n\cdot p(n)^2} = \frac{1}{2n}.
\end{aligned}$$

This completes the proof of the claim. ■

To finish the proof of the lemma, we observe that A successfully inverts $f(x)$ if all of the following are true:

1. $x \in S_n$.
2. The guesses for the σ^i are all correct, that is, $\sigma^i = b(x, s^I)$ for all i.
3. Each z_i that A produces is correct.

The first event is true with probability at least $\varepsilon(n) \ge \frac{1}{p(n)}$ by Claim 1. The second event is true with probability

$$\frac{1}{2^\ell} \ge \frac{1}{4n\cdot p(n)^2 + 2}$$

by inequality 5. The third event is true with probability at least $(1 - \frac{1}{2n})^n > \frac{1}{2}$ for all $n \ge 2$. Multiplying these together gives us a lower bound on the success probability of f, namely,

$$\frac{1}{p(n)}\cdot\frac{1}{4n\cdot p(n)^2 + 2}\cdot\frac{1}{2} = \frac{1}{8n\cdot p(n)^3 + 4p(n)}.$$

Thus, taking $q(n) = 8n\cdot p(n)^3 + 4p(n)$, A has a success probability greater than $\frac{1}{q(n)}$ for all sufficiently large $n \in N$, contradicting the assumption that f is strongly one-way. Thus, the assumption that b is not hard-core for f and that G exists must be false. This completes the proof of Lemma 3 of section 21.

24 Hard-Core Functions

We extend the notion of a hard-core predicate of a function to a hard-core function.

Definition: Let $h : \{0,1\}^* \to \{0,1\}^*$ be a polynomial-time computable length-regular function.[1] Let $\ell = |h(1^n)|$. h is a *hard core* of f if for all p.p.t. algorithms D', all positive polynomials $p(\cdot)$, and all sufficiently large n,

$$|\Pr[D'(f(X_n), h(X_n)) = 1] - \Pr[D'(f(X_n), R_{\ell(n)}) = 1] < \frac{1}{p(n)},$$

where X_n and $R_{\ell(n)}$ are independent random variables, uniformly distributed over $\{0,1\}^n$ and $\{0,1\}^{\ell(n)}$, respectively.

[1] Recall that h is length-regular if $|h(x)| = |h(y)|$ whenever $|x| = |y|$.

Intuitively, h is a hard core for f if the value of $h(x)$ is indistinguishable from a random string, even knowing the value of $f(x)$. On the surface, this looks to be an even stronger condition than unpredictability. Obviously, if one could predict $h(x)$ from $f(x)$, then one could distinguish $h(x)$ from random. Namely, if the given string equals the prediction, output 1, otherwise output 0. On the other hand, it's not a priori obvious how being able to distinguish $h(x)$ from random would be useful at prediction.

Theorem 1 *Let f be strongly one-way. Let $c > 0$ and let $\ell(n) = \min\{n, \lceil c \log_2 n \rceil\}$. Let x be a string of length n and s a string of length $2n$. Define*

$$\begin{aligned} g_2(x,s) &= (f(x), s), \\ b_i(x,s) &= x \cdot (s_{i+1}, \ldots, s_{i+n}), \text{ for } i = 1, \ldots, \ell(n), \\ h(x,s) &= b_1(x,s) \ldots b_{\ell(|x|)}(x,s). \end{aligned}$$

Then h is a hard core of g_2.

We omit the non-trivial proof of this theorem and remark only that hard core functions with logarithmic lengths are known for RSA and other cryptographic collections, assuming the corresponding collections are one-way. Details are in the textbook.

25 Probability Ensembles

To begin our formal development of pseudorandom sequence generation, we define a probability ensemble, analogous to the previous definition of a collection of one-way functions.

Let I be a countable set. An *ensemble* indexed by I is a sequence of random variables $X = \{X_i\}_{i \in I}$ indexed by I.

Typical index sets are the natural numbers $\mathbb{N}$ or binary strings $\{0,1\}^*$. Typically, $X = \{X_n\}_{n \in \mathbb{N}}$ has X_n ranging over strings of length $\mathrm{poly}(n)$, and $X = \{X_w\}_{w \in \{0,1\}^*}$ has X_w ranging over strings of length $\mathrm{poly}(|w|)$.

26 Polynomial Time Indistinguishability

We give two definitions of polynomial time indistinguishable ensembles, depending on the index set.

Variant 1: Ensembles $X = \{X_n\}_{n \in \mathbb{N}}$ and $Y = \{Y_n\}_{n \in \mathbb{N}}$ are *indistinguishable in polynomial time* if for all probabilistic polynomial time algorithms D, all positive polynomials $p(\cdot)$, and all sufficiently large n

$$|\Pr[D(X_n, 1^n) = 1] - \Pr[D(Y_n, 1^n) = 1]| < \frac{1}{p(n)}.$$

Variant 2: Ensembles $X = \{X_w\}_{w \in \{0,1\}^*}$ and $Y = \{Y_w\}_{w \in \{0,1\}^*}$ are *indistinguishable in polynomial time* if for all probabilistic polynomial time algorithms D, all positive polynomials $p(\cdot)$, and all sufficiently large n

$$|\Pr[D(X_w, w) = 1] - \Pr[D(Y_w, w) = 1]| < \frac{1}{p(|w|)}.$$

Easy consequences

Let D be a p.p.t. algorithm. Let $d(\alpha)$ be the probability that $D(\alpha) = 1$. Let $d_X(n) = E[d(X_n)]$, $d_Y(n) = E[d(Y_n)]$, be the expected value of D's output when given a string from X or from Y, respectively. Let

$$\delta(n) = |d_X(n) - d_Y(n)|.$$

Then X and Y are indistinguishable by D iff δ is negligible in n.

Let $\mathrm{wt}(\alpha) =$ # 1's in α. If X, Y are polynomial-time indistinguishable, then

$$\left|\Pr\left[\mathrm{wt}(X_n) < \frac{n}{2}\right] - \Pr\left[\mathrm{wt}(Y_n) < \frac{n}{2}\right]\right|$$

regarded as a function of n is negligible. (Why?)

Note: The same applies to any polynomial-time computable string statistic.