COM S 6830 – Cryptography Sep 27, 2011

Lecture 9: Computational Indistinguishability and Pseudorandomness

Instructor: Rafael Pass *Scribe: Anthony Chang*

1 Recap

1.1 Ensemble of distributions

An **ensemble of distributions** $\{X_n\}_{n\in\mathbb{N}}$ ($\{X_n\}$ for short) is a sequence of probability distributions $X_1, X_2, \ldots$

1.2 Computational indistinguishability

Two ensembles of distributions $\{X_n\}$ and $\{Y_n\}$ are said to be **computationally indistinguishable** ($\{X_n\} \approx \{Y_n\}$) if:

$$\forall\, \text{nuPPT}\, \mathcal{D}\ \exists\, \text{neg}\, \epsilon\ \forall\, n \in \mathbb{N}\ |\Pr[t \leftarrow X_n : \mathcal{D}(1^m, t) = 1] - \Pr[t \leftarrow X_n : \mathcal{D}(1^m, t) = 1]| \leq \epsilon(n)$$

1.3 Two properties of (in)distinguishability

1.3.1 Closure under efficient operations

$\{X_n\} \approx \{Y_n\}$ implies $\{M(X_n)\} \approx \{M(Y_n)\}$ for any nuPPT M.

1.3.2 Hybrid lemma

For a sequence of probability distributions $X_1, X_2, \ldots, X_m$, if there exists a machine $\mathcal{D}$ that distinguishes X_1 from X_m with probability ϵ, then there exists some i such that $\mathcal{D}$ distinguishes X_i from X_{i+1} with probability at least $\frac{\epsilon}{m}$.

2 Prediction lemma

A third property of distinguishability, the **prediction lemma**, intuitively states that if you can distinguish two distributions, then you should be able to guess which distribution an arbitrary sample came from as well.

Lemma 1 *For ensembles $\{X_n^0\}$ and $\{X_n^1\}$ where each X_n^0 and X_n^1 is a distribution over $\{0,1\}^{m(n)}$ for some polynomial m, let $\mathcal{D}$ be a nuPPT that distinguishes $\{X_n^0\}$ and $\{X_n^1\}$ with probability $\mu(n)$ for infinitely many n. Then there exists an nuPPT $\mathcal{A}$ such that for infinitely many n,*

$$\Pr[b \leftarrow \{0,1\}, t \leftarrow X_n^b : \mathcal{A}(t) = b] \geq \frac{1}{2} + \frac{\mu(n)}{2}$$

Proof. Assume without loss of generality that $\mathcal{D}$ outputs 1 with higher probability when getting a sample from X_n^1. This assumption is safe because either $\mathcal{D}(t)$ or $1 - \mathcal{D}(t)$ will work for infinitely many n, or alternatively, because $\mathcal{D}$ can accept additional information about whether to invert its output as a nuPPT.

We'll show that $\mathcal{D}$ actually satisfies the above conditions for $\mathcal{A}$, so $\mathcal{D}$ is a predictor:

$\Pr[b \leftarrow \{0,1\} : t \leftarrow X_n^b : \mathcal{D}(t) = b]$

$= \frac{1}{2}\Pr[t \leftarrow X_n^1 : \mathcal{D}(t) = 1] + \frac{1}{2}\Pr[t \leftarrow X_n^0 : \mathcal{D}(t) \neq 1]$

$= \frac{1}{2}\Pr[t \leftarrow X_n^1 : \mathcal{D}(t) = 1] + \frac{1}{2}(1 - \Pr[t \leftarrow X_n^0 : \mathcal{D}(t) = 1])$

$= \frac{1}{2} + \frac{1}{2}(\Pr[t \leftarrow X_n^1 : \mathcal{D}(t) = 1] - \Pr[t \leftarrow X_n^0 : \mathcal{D}(t) = 1])$

$= \frac{1}{2} + \frac{\mu(n)}{2}$ ■

Note that there are no restrictions on μ, but the predictor $\mathcal{A}$ will have some special properties if μ is polynomial. This will be discussed in a later lecture.

3 Pseudorandomness

With these three lemmas, we can define pseudorandomness as indistinguishability from the uniform distribution: $\{X_n\}$ is pseudorandom if $\{X_n\} \approx \{U_{m(n)}\}$, where X_n is over $\{0,1\}^{m(n)}$, m is polynomial, and U is the uniform distribution.

We can show that this definition of pseudorandomness is equivalent to passing the next-bit test using Yao's theorem.

3.1 Next-bit test

Definition 1 *An ensemble $\{X_n\}$ over $\{0,1\}^{m(n)}$ passes the* ***next-bit test*** *if*

$$\forall\, nuPPT\, \mathcal{A}\; \exists\, neg\, \epsilon\; \forall\, n \in \mathbb{N}, i \in [0, 1, \ldots, m(n)-1] Pr[t \leftarrow X_n : \mathcal{A}(1^n, t_{0\to i}) = t_{i+1}] \leq \frac{1}{2} + \epsilon(n)$$

where $t_{0\to i}$ denotes the first $i+1$ bits of t.

Intuitively, a prefix of a sample of $\{X_n\}$ cannot be used to predict the next bit in the sample with high probability.

3.2 Half of Yao's theorem

Theorem 2 *Any ensemble $\{X_n\}$ over $\{0,1\}^{m(n)}$ that passes the next-bit test is pseudorandom.*

Proof. The proof will proceed by contradiction, as usual. Assume for the sake of contradiction that there exists a $\mathcal{D}$ distinguishing $\{X_n\}$ from $\{U_{m(n)}\}$ with probability $\frac{1}{p(n)}$ for polynomial p, so $\{X_n\}$ is not pseudorandom. We will use this $\mathcal{D}$ to predict the next bit of any sample from $\{X_n\}$.

Consider the hybrid distributions $H_n^i = \{l \leftarrow X_n, r \leftarrow U_{m(n)} : l_{0\to i}||r_{i+1\to m(n)}\}$. The first i bits of H_n^i come from X_n, while the rest are uniformly random (so we can generate them ourselves).

$H_n^0 = U_{m(n)}$ and $H_n^{m(n)} = X_n$, with each H_n^i in between injecting i bits from X_n. By our assumption, $\{X_n\}$ is distinguishable from $\{U_{m(n)}\}$, so for infinitely many n, $H_n^{m(n)} = X_n$ is distinguishable by $\mathcal{D}$ from $H_n^0 = U_{m(n)}$ with probability at least $\frac{1}{p(n)}$.

Applying the hybrid lemma, there exists some i such that $\mathcal{D}$ distinguishes H_n^i from H_n^{i+1} for each of these n. The only difference between these distributions is that bit $i+1$ of H_n^i is uniformly random, whereas bit $i+1$ of H_n^{i+1} is drawn from X_n.

Define another hybrid $\tilde{H}_n^{i+1} = \{l \leftarrow X_n, r \leftarrow U_{m(n)} : l_{0\to i}||1 - l_{i+1}||r_{i+2\to m}\}$. $\tilde{H}_n^{i+1}$ is exactly H_n^{i+1} with bit $i+1$ flipped; if $\mathcal{D}$ can distinguish H_n^i from H_n^{i+1}, then it can certainly distinguish H_n^{i+1} from $\tilde{H}_n^{i+1}$. We can show this with some algebra:

$$|\Pr[t \leftarrow H_n^{i+1} : \mathcal{D}(t) = 1] - \Pr[t \leftarrow H_n^i : \mathcal{D}(t) = 1]|$$

$$= |\Pr[t \leftarrow H_n^{i+1} : \mathcal{D}(t) = 1] - \tfrac{1}{2}\Pr[t \leftarrow H_n^{i+1} : \mathcal{D}(t) = 1] - \tfrac{1}{2}\Pr[t \leftarrow \tilde{H}_n^{i+1} : \mathcal{D}(t) = 1]|$$

$= |\frac{1}{2}\Pr[t \leftarrow H_n^{i+1} : D(t) = 1] - \frac{1}{2}\Pr[t \leftarrow \tilde{H}_n^{i+1} : D(t) = 1]|$

$\geq \frac{1}{p(n)m(n)}$ by assumption on $\mathcal{D}$ and hybrid lemma

where the second line follows because H_n^i can be expressed as $\frac{1}{2}H_n^{i+1} + \frac{1}{2}\tilde{H}_n^{i+1}$ (making bit $i+1$ uniformly random).

Now that we can distinguish H_n^{i+1} and $H_n^{\tilde{i}+1}$, we can apply the prediction lemma to show that there exists a nuPPT $\mathcal{A}$ for infinitely many $n \in \mathbb{N}$ such that

$$\Pr[b \leftarrow \{0,1\}, t \leftarrow H_n^{b,i+1} : \mathcal{A}(t) = b] \geq \frac{1}{2} + \frac{1}{p(n)m(n)}$$

where $H_n^{0,i+1} = \tilde{H}_n^{i+1}$ and $H_n^{1,i+1} = H_n^{i+1}$.

$\mathcal{A}$ can be used to construct a predictor for bit $i+1$ of infinitely many X_n: construct $\mathcal{A}'(1^n, y)$ (where y is an $i+1$ bit prefix of a sample of some such X_n) to pick a random $r \leftarrow \{0,1\}^{m(n)-i}$. If $\mathcal{A}(y||r) = 1$, then $\mathcal{A}'$ outputs r_i, otherwise if $\mathcal{A}(y||r) = 0$, then $\mathcal{A}'$ outputs $1 - r_i$. $\mathcal{A}'$ effectively predicts bit $i+1$ of $\{X_n\}$ because it satisfies

$$\Pr[t \leftarrow X_n : \mathcal{A}'(1^n, t_{0\rightarrow i}) = t_{i+n}] \geq \frac{1}{2} + \frac{1}{p(n)m(n)}$$

for infinitely many n. ■