## Lecture Notes 24

## 114 Bit-prediction

One important property of the uniform distribution $U$ on bit-strings $b_{1}, \ldots, b_{\ell}$ is that the individual bits are statistically independent from each other. This means that the probability that a particular bit $b_{i}=1$ is unaffected by the values of the other bits in the sequence. Thus, any algorithm that attempts to predict $b_{i}$, even knowing other bits of the sequence, will be correct only $1 / 2$ of the time. We now translate this property of unpredictability to pseudorandom sequences.

### 114.1 Next-bit prediction

One property we would like a pseudorandom sequence to have is that it be difficult to predict the next bit given the bits that came before.

We say that an algorithm $A$ is an $\epsilon$-next-bit predictor for bit $i$ of a PRSG $G$ if

$$
\operatorname{prob}\left[A\left(b_{1}, \ldots, b_{i-1}\right)=b_{i}\right] \geq \frac{1}{2}+\epsilon
$$

where $\left(b_{1}, \ldots, b_{i}\right)=G_{i}(S)$. To explain this notation, $S$ is a uniformly distributed random variable ranging over the possible seeds for $G . G(S)$ is a random variable (i.e., probability distribution) over the output strings of $G$, and $G_{i}(S)$ is the corresponding probability distribution on the length- $i$ prefixes of $G(S)$.

Next-bit prediction is closely related to indistinguishability, introduced in section 110. We will show later that if if $G(S)$ has a next-bit predictor for some bit $i$, then $G(S)$ is distinguishable from the uniform distribution $U$ on the same length strings, and conversely, if $G(S)$ is distinguishable from $U$, then there is a next-bit predictor for some bit $i$ of $G(S)$. The precise definitions under which this theorem is true are subtle, for one must quantify both the amount of time the judge and next-bit predictor algorithms are permitted to run as well as how much better than chance the judgments or predictions must be in order to be considered a successful judge or next-bit predictor. We defer the mathematics for now and focus instead on the intuitive concepts that underlie this theorem.

### 114.2 Building a judge from a next-bit predictor

Suppose a PRSG $G$ has an $\epsilon$-next-bit predictor $A$ for some bit $i$. Here's how to build a judge $J$ that distinguishes $G(S)$ from $U$. The judge $J$, given a sample string drawn from either $G(S)$ or from $U$, runs algorithm $A$ to guess bit $b_{i}$ from bits $b_{1}, \ldots, b_{i-1}$. If the guess agrees with the real $b_{i}$, then $J$ outputs 1 (meaning that he guesses the sequence came from $G(S)$ ). Otherwise, $J$ outputs 0 . For sequences drawn from $G(S), J$ will output 1 with the same probability that $A$ successfully predicts bit $b_{i}$, which is at least $1 / 2+\epsilon$. For sequences drawn from $U$, the judge will output 1 with probability exactly $1 / 2$. Hence, the judge distinguishes $G(S)$ from $U$ with advantage $\epsilon$.

It follows that no cryptographically strong PRSG can have an $\epsilon$-next-bit predictor. In other words, no algorithm that attempts to predict the next bit can have more than a "small" advantage $\epsilon$ over chance.

### 114.3 Previous-bit prediction

Previous-bit prediction, while perhaps less natural, is analogous to next-bit prediction. An $\epsilon$ -previous-bit predictor for bit $i$ is an algorithm that, given bits $b_{i+1}, \ldots, b_{\ell}$, correctly predicts $b_{i}$ with probability at least $1 / 2+\epsilon$.

As with next-bit predictors, it is the case that if $G(S)$ has a previous-bit predictor for some bit $b_{j}$, then some judge distinguishes $G(S)$ from $U$. Again, I am being vague with the exact conditions under which this is true. The somewhat surprising fact follows that $G(S)$ has an $\epsilon$-next-bit predictor for some bit $i$ if and only if it has an $\epsilon^{\prime}$-previous-bit predictor for some bit $j$ (where $\epsilon$ and $\epsilon^{\prime}$ are related but not necessarily equal).

To give some intuition into why such a fact might be true, we look at the special case of $\ell=2$, that is, of 2-bit sequences. The probability distribution $G(S)$ can be described by four probabilities

$$
p_{u, v}=\operatorname{prob}\left[b_{1}=u \wedge b_{2}=v\right], \text { where } u, v \in\{0,1\} .
$$

Written in tabular form, we have


We describe an algorithm $A(v)$ for predicting $b_{1}$ given $b_{2}=v . A(v)$ predicts $b_{1}=0$ if $p_{0, v}>$ $p_{1, v}$, and it predicts $b_{1}=1$ if $p_{0, v} \leq p_{1, v}$. In other words, the algorithm chooses the value for $b_{1}$ that is most likely given that $b_{2}=v$. Let $a(v)$ be the value predicted by $A(v)$.

Theorem 1 If $A$ is an $\epsilon$-previous-bit predictor for $b_{1}$, then $A$ is an $\epsilon$-next-bit predictor for either $b_{1}$ or $b_{2}$.

Proof: Assume $A$ is an $\epsilon$-previous-bit predictor for $b_{1}$. Then $A$ correctly predicts $b_{1}$ given $b_{2}$ with probability at least $1 / 2+\epsilon$. We show that $A$ is an $\epsilon$-next-bit predictor for either $b_{1}$ or $b_{2}$.

We have two cases:
Case 1: $a(0)=a(1)$. Then algorithm $A$ does not depend on $v$, so $A$ itself is also an $\epsilon$-next-bit predictor for $b_{1}$.
Case 2: $a(0) \neq a(1)$. The probability that $A(v)$ correctly predicts $b_{1}$ when $b_{2}=v$ is given by the conditional probability

$$
\operatorname{prob}\left[b_{1}=a(v) \mid b_{2}=v\right]=\frac{\operatorname{prob}\left[b_{1}=a(v) \wedge b_{2}=v\right]}{\operatorname{prob}\left[b_{2}=v\right]}=\frac{p_{a(v), v}}{\operatorname{prob}\left[b_{2}=v\right]}
$$

The overall probability that $A\left(b_{2}\right)$ is correct for $b_{1}$ is the weighted average of the conditional probabilities for $v=0$ and $v=1$, weighted by the probability that $b_{2}=v$. Thus,

$$
\begin{aligned}
\operatorname{prob}\left[A\left(b_{2}\right) \text { is correct for } b_{1}\right] & =\sum_{v \in\{0,1\}} \operatorname{prob}\left[b_{1}=a(v) \mid b_{2}=v\right] \cdot \operatorname{prob}\left[b_{2}=v\right] \\
& =\sum_{v \in\{0,1\}} p_{a(v), v} \\
& =p_{a(0), 0}+p_{a(1), 1}
\end{aligned}
$$

Now, since $a(0) \neq a(1)$, the function $a$ is one-to-one and onto. A simple case analysis shows that either $a(v)=v$ for $v \in\{0,1\}$, or $a(v)=\neg v$ for $v \in\{0,1\}$. That is, $a$ is either the identity
or the complement function. In either case, $a$ is its own inverse. Hence, we may also use algorithm $A(u)$ as a predictor for $b_{2}$ given $b_{1}=u$. By a similar analysis to that used above, we get

$$
\begin{aligned}
\operatorname{prob}\left[A\left(b_{1}\right) \text { is correct for } b_{2}\right] & =\sum_{v \in\{0,1\}} \operatorname{prob}\left[b_{2}=a(u) \mid b_{1}=u\right] \cdot \operatorname{prob}\left[b_{1}=u\right] \\
& =\sum_{v \in\{0,1\}} p_{u, a(u)} \\
& =p_{0, a(0)}+p_{1, a(1)}
\end{aligned}
$$

But

$$
p_{a(0), 0}+p_{a(1), 1}=p_{0, a(0)}+p_{1, a(1)}
$$

since either $a$ is the identity function or the complement function. Hence, $A\left(b_{1}\right)$ is correct for $b_{2}$ with the same probability that $A\left(b_{2}\right)$ is correct for $b_{1}$. Therefore, $A$ is an $\epsilon$-next-bit predictor for $b_{2}$.

In both cases, we conclude that $A$ is an $\epsilon$-next-bit predictor for either $b_{1}$ or $b_{2}$.

## 115 Foundations of Cryptography

[This section is based on the second textbook, John Talbot and Dominic Welsh, Complexity and Cryptograpy: An Introduction, §6.1-6.3.]

We have been rather vague throughout the course about what it means for a problem to be "hard" to solve. A public key system such as RSA can always be broken given enough time. Eve could simply factor $n$ and find the decryption key, or she could try all possible messages $m$ to see which one encrypts to the known ciphertext $c$ given the public encryption key. So we can't hope to always prevent Eve from decrypting our messages. The best we can hope for is that she rarely succeeds. We can state this as follows:

Goal: If Eve uses a probabilistic polynomial time algorithm, the probability of her successfully decrypting $c=E_{k}(m)$ for a random message $m$ is negligible.

We have to say what "negligible" means. A function $r: \mathbb{N} \rightarrow \mathbb{N}$ is negligible if for all polynomial functions $p: \mathbb{N} \rightarrow \mathbb{N}$, there exists $k_{0}$ such that for all $k \geq k_{0}$,

$$
r(k)<\frac{1}{p(k)} .
$$

(We assume $p(k) \geq 1$ for all integers $k \geq 0$.)
This definition says that $r(k) \rightarrow 0$ as $k \rightarrow \infty$. But it says more. It also says that $r(k)$ goes to 0 "faster" than one over any polynomial. For all large enough $k, r(k)$ is smaller than $1 / k^{10}$, for example. It's tempting to equate negligible with inverse exponential. While $1 / 2^{k}$ is indeed a negligible function, there are larger functions that are also negligible by our definitions, for example, $1 / 2^{\sqrt{k}}$ and $1 / k^{\log k}$.

Definition: A function $f$ is strong one-way if

1. $f$ is polynomial time computable.
2. Any probabilistic polynomial time algorithm for inverting $f(x)$ when given a random instance $y=f(x)$ has negligible chance of finding a preimage of $y$.

By "random instance", this means a $y$ of the form $y=f(x)$ for a uniformly distributed random $x$.
Even this definition hides some technical difficulties that are done carefully in the text book. In particular, the difficulty of inverting $f$ should remain hard even when the inverter is told the length of $x$.

The following function related to the discrete $\log$ problem is often assumed to be strong one:

$$
\operatorname{dexp}(p, g, x)=\left(p, g, g^{x}(\bmod p)\right)
$$

It is easy to compute, but to invert it requires solving the discrete log problem.
Although functions like dexp and factoring are believed to be strong one-way functions, no function has been proved to be strong one-way. The following theorem explains why it is so hard to prove that functions are strong one-way - doing so would also settle the famous $P=N P$ ? problem by showing that $P \neq N P$.

Theorem 2 If strong one-way functions exist, then $P \neq N P$.
Proof: Assume $f$ is strong one-way. Define the language

$$
L_{f}=\left\{\left(x, y, 1^{k}\right) \mid \exists u \in\{0,1\}^{k}: f(x u)=y\right\}
$$

$L_{f}$ is easily seen to be in $N P$. To test if a given $\left(x, y, 1^{k}\right)$ triple is in $L_{f}$, simply guess a string $u$ of length $k$, compute $f(x u)$, and check that $f(x u)=y$.

If $P=N P$, then there is a polynomial time algorithm $A$ for testing membership in $L_{f} . A$ can be used to determine the bits of a preimage of $y$ one by one. At stage $i$, assume that the first $i-1$ bits $x_{i-1}$ of a preimage have already been determined. To determine bit $i$, test if $\left(x_{i-1} 0, y, 1^{k-i}\right) \in L_{f}$. If so, then $x_{i}=0 x_{i-1}$ are the first $i$ bits of some preimage. If not, then $\left(x_{i-1} 1, y, 1(k-i)\right)$ must be in $L_{f}$ since $x_{i 01}$ is known to be the prefix of a preimage. Hence, one can take $x_{i}=x_{i-1} 1$ in this case. After stage $k$, the number $x_{k}$ is a preimage of $y$, as desired. The entire procedure requires only $k$ evaluations of $A$, so it is also a polynomial time algorithm for inverting $f$, contradicting the assumption that $f$ is one-way.

We conclude that $P \neq N P$ if one-way functions exist.

