CS 127/CSCI E-127: Introduction to Cryptography

Prof. Salil Vadhan Fall 2013

Lecture Notes 10:

Hardcore Bits

Reading.

- Katz–Lindell 2nd edition §7.1.3–7.4.1 (except 7.3.3) OR 1st edition §6.1.3–6.4.1 (except 6.3.3).

1 Hardcore Bits

Motivation: Why can't we directly use a OWF for encryption? That is, encrypt a message m as ciphertext $c = f(m)$. What are the difficulties with this idea?

If f is a OWF, it is hard to determine x from $f(x)$, but is it also hard to compute a particular bit of x from $f(x)$, say the first bit of x? Random guessing gives a probability of success of $\frac{1}{2}$ but some bits might be even easier to guess. A few examples:

A one-way function can reveal a large part of its input: is there a fraction of the bits of the input which is always "well-hidden"? (i.e. any polynomial-time algorithm cannot have a nonnegligible advantage over random guessing when computing those bits from the output of the function) The answer is no, because we can construct one-way functions such that each bit of x can be obtained from $f(x)$ with high probability. Thus, we instead look for some "bit of information" which is hard to compute.

Definition 1 *$b : \{0,1\}^* \to \{0,1\}$ is a* hardcore predicate *(or* hardcore bit*) for one-way function f if*

- *b is polynomial-time computable.*
- *For every PPT A, there is a negligible function ε such that*

$$\Pr[A(1^n, f(X)) = b(X)] \leq \frac{1}{2} + \varepsilon(n) \qquad \forall n,$$

where the probability is over $X \stackrel{R}{\leftarrow} \{0,1\}^n$ and the coin tosses of A.

Definition 2 *$\{b_I : D_I \to \{0,1\}\}_{I \in \mathcal{I}}$ is a* hardcore predicate family *for the one-way function family $\mathcal{F} = \{f_I : D_I \to R_I\}$ if*

- *Given $I \in \mathcal{I}$ and $x \in D_I$, $b_I(x)$ can be computed in polynomial time.*
- *For every PPT A, there is a negligible function ε such that*

$$\Pr\left[A(1^n, I, f_I(X)) = b_I(X)\right] \leq \frac{1}{2} + \varepsilon(n) \qquad \forall n,$$

where the probability is taken over $I \xleftarrow{R} \mathsf{Gen}(1^n)$, $X \xleftarrow{R} D_I$, and the coin tosses of A.

2 Examples

RSA functions

- Under the RSA Assumption, the least significant bit is a hardcore bit for RSA:

$$\mathrm{lsb}_{N,e} : \mathbb{Z}_N^* \mapsto \{0,1\}$$

Given $N, e, x^e \bmod N$, we cannot compute $\mathrm{lsb}_{N,e}(x)$ with a nonnegligible advantage over random guessing. (It has been shown that if we could, then we could also invert RSA with nonnegligible probability.)

- Define $\mathrm{half}_N(x)$ by $\mathrm{half}_N(x) = 0$ if $0 \leq x < N/2$ and 1 otherwise ($\mathrm{half}_N(x)$ is like the most significant bit of x). $\mathrm{half}_N(x)$ is a hardcore bit for RSA (again under the RSA Assumption).

Rabin's functions

- The least significant bit is a hardcore bit for Rabin's functions (under the Factoring Assumption):

$$\mathrm{lsb}_N : \mathbb{Z}_N^* \mapsto \{0,1\}$$

Given $N, x^2 \bmod N$, we cannot compute $\mathrm{lsb}_N(x)$ with a nonnegligible advantage over random guessing.

- $\mathrm{half}_N(x)$ is a hardcore bit for Rabin's functions (under the Factoring Assumption).

Modular Exponentiation/Discrete Log $\mathrm{half}_{p-1}(x)$ is a hardcore bit for Modular Exponentiation (under the Discrete Log Assumption).

3 Goldreich–Levin hardcore bit

Does every one-way function have a hardcore bit? The following theorem proves that from any arbitrary OWF, we can construct a OWF with a hardcore bit by taking the XOR of a random subset of bits. For $x, r \in \{0,1\}^n$, define $\langle x, r\rangle = \sum_i x_i r_i \bmod 2 = \oplus_{i|r_i=1} x_i$.

Theorem 3 (Goldreich–Levin hardcore bit) *Let f be any one-way function, and define $f'(x,r) = (f(x), r)$ for $|x| = |r|$. Then $\langle x, r\rangle$ is a hardcore predicate for f'.*

This theorem is most interesting when f is one-to-one. Note that if f is one-to-one, then so is f'.

Proof ideas:

Reducibility argument: Suppose that there exists a PPT A that predicts $\langle x, r\rangle$ from $(f(x), r)$ with nonnegligible advantage over random guessing. We construct a PPT B that uses A to invert f with nonnegligible probability.

"Easy" case: Assume that $A(f(x), r)$ computes the hardcore bit $\langle x, r\rangle$ with probability 1.

"Medium" case We assume that $A(f(x), r)$ computes the hardcore bit $\langle x, r\rangle$ with probability $\geq \frac{3}{4} + \varepsilon$, where $\varepsilon = \varepsilon(n)$ is a nonnegligible function and the probability is taken over the random input x and the coin tosses of A. This imples that for at least $\varepsilon/2$ fraction of x, $\Pr[A(f(x), R) = \langle x, R\rangle] \geq 3/4 + \varepsilon/2$ (probability just over R and the coin tosses of A). We will give a PPT that inverts f with high probability on these *good* x's, which will contradict the one-wayness of f. So assume for the rest of the proof that x is good.

We have a problem generalizing the argument used in the easy case because A is only guaranteed to succeed on *random* r: we do not know how A behaves if r is not random (such as for $r = e^{(i)}$).

Idea ("random self-reducibility"): reduce the task of computing $\langle x, r\rangle$ for a particular value of r (namely $r = e^{(i)}$) to computing it on several uniformly random values of r:

Computing $\langle x, e^{(i)}\rangle$ with probability $1/2 + \varepsilon$:

Computing $\langle x, e^{(i)}\rangle$ with high probability:

Computing x with high probability:

General case (A computes hardcore bit with probability $1/2 + \varepsilon$) requires additional ideas.

Theorem 4 (Goldreich-Levin hardcore bit for collections) *Let $\mathcal{F} = \{f_i : \mathcal{D}_i \to \mathcal{R}_i\}$ be any one-way function family, and define $g_{i,r}(x) = f_i(x)$, $b_{i,r}(x) = \langle x, r\rangle$. Then $\{b_{i,r} : \mathcal{D}_i \to \mathcal{R}_i\}$ is a family of hardcore predicates for the family of one-way functions $\{g_{i,r} : \mathcal{D}_i \to \mathcal{R}_i\}$.*

4 Hardcore Bits $\Rightarrow$ PRGs with 1-bit Stretch

Theorem 5 *If f is a one-way permutation with hardcore predicate b, then $G(s) = f(s)b(s)$ is a pseudorandom generator.*

Proof:

1. Suppose there is a PPT D that distinguishes between $G(S) = f(S)b(S)$ and $U_{n+1} = f(S)R$ with nonnegligible advantage ε (where $S \xleftarrow{R} \{0,1\}^n$ and $R \xleftarrow{R} \{0,1\}$).
2. Then D distinguishes between $Y_0 = f(S)b(S)$ and $Y_1 = f(S)\overline{b(S)}$ with advantage 2ε:
3. The following PPT A predicts $b(S)$ from $f(S)$ with probability at least $1/2 + \varepsilon$:

■