Lecture 12: Goldrecih-Levin Hardcore Predicate

Recall: Overall Construction of PRG from OWP

- Let $f\colon \{0,1\}^n \to \{0,1\}^n$ be a OWP
- Given f construct a new OWP that has a hardcore predicate. Let $g\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n \times \{0,1\}^n$ be a OWP defined by $g(x,r) = (f(x), r)$ and $h(x,r) = \langle x, r\rangle$ be the corresponding hardcore predicate
- Given a OWP with a hardcore predicate, construct a one-bit extension PRG. Let $G\colon \{0,1\}^{2n} \to \{0,1\}^{2n+1}$ be the one-bit extension PRG defined by
  $G(x,r) = (g(x,r), h(x,r)) \equiv (f(x), r, \langle x, r\rangle)$
- Given the one-bit extension PRG G, construct an arbitrary polynomial-stretch PRG. Let $H\colon \{0,1\}^{2n} \to \{0,1\}^{\ell}$ be the arbitrary stretch PRG, where $\ell > 2n$ and ℓ is a polynomial in n. We define
  $H(x,r) = \left(\langle x, r\rangle, \langle f(x), r\rangle, \left\langle f^2(x), r\right\rangle, \dots, \left\langle f^{\ell-1}(x), r\right\rangle\right)$

Proofs

- We have seen the proofs of all the steps except the following: $h(x, r)$ is a hardcore predicate of $g(x, r)$.
- To show this result, we need to show the following equivalent result: f is a OWP $\implies$ Given $(f(x), r)$ for random x, r, it only possible to predict $\langle x, r \rangle$ with negligible advantage
- We consider the contrapositive of this statement
- We are given: There exists an efficient adversary $\mathcal{A}^*$ that takes as input $(f(x), r)$ and correctly guesses $\langle x, r \rangle$ with $1/n^c$ advantage
- We need to show: There exists an efficient adversary $\widetilde{\mathcal{A}}$ that can invert f at $1/n^d$ fraction of inputs
- This is Goldreich-Levin Hardcore Predicate Theorem
- We will only see a restricted proof of this result

Restricted Proof: Version 1

- So, we are given:

$$\Pr[x \sim U_{\{0,1\}^n}, r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r \rangle] \geqslant \frac{1}{2} + \frac{1}{n^c}$$

- In this restriction we consider:

$$\Pr[x \sim U_{\{0,1\}^n}, r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r \rangle] = 1$$

- Consider the following algorithm for $\widetilde{\mathcal{A}}(y)$
  - For $i \in \{1, \dots, n\}$: Let $\widetilde{x}_i = \mathcal{A}^*(y, e_i)$, where $e_i = (\overbrace{0, \dots, 0}^{(i-1)}, 1, \overbrace{0, \dots, 0}^{(n-i)})$
  - Return $(\widetilde{x}_1, \dots, \widetilde{x}_n)$
- Note that $\widetilde{x}_i = x_i$ for all i and hence the algorithm completely recovers x with probability 1

Restricted Proof: Version 2

- In this restriction we consider: For $\varepsilon = 1/n^c$

$$\Pr[x \sim U_{\{0,1\}^n}, r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r \rangle] \geqslant \frac{3}{4} + \varepsilon$$

- Define the following subset

$$G = \left\{ x \colon \Pr_{r \sim U_{\{0,1\}^n}} [\mathcal{A}^*(f(x), r) = \langle x, r \rangle] \geqslant \frac{3}{4} + \frac{\varepsilon}{2} \right\}$$

- Intuition: G is the set of all those "good" x where the adversary successfully finds the hardcore predicate with "good probability." We will invert the function f for $x \in G$

Claim

$|G| \geqslant (\varepsilon/2) \cdot 2^n$

Proof of the Claim

Overview:

- This argument is a general argument referred to as: Averaging Argument, Pigeon-hole Principle, or Markov Inequality
- English Version of this Inequality: If for random (x, r) an algorithm is "successful" with "overwhelming probability." Then the fraction of inputs that are "good values of x" where the algorithm succeeds with "good enough probability" is "noticeable"
- In our setting "successful" is the even that $\mathcal{A}^*$ correctly outputs $\langle x, r\rangle$, "overwhelming probability" is $3/4 + \varepsilon$, "good enough probability" is $3/4 + \varepsilon/2$, "good values of x" are those xs where for random r the algorithm finds the bit $\langle x, r\rangle$ with good enough probability, and "noticeable" is $\varepsilon/2$

Proof of the Claim

Perspective:

- Note that

$$\Pr[x \sim U_{\{0,1\}^n}, r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r\rangle] \geqslant \frac{3}{4} + \varepsilon$$

  implies that there exists one x such that:

$$\Pr_{r \sim U_{\{0,1\}^n}}[\mathcal{A}^*(f(x), r) = \langle x, r\rangle] \geqslant \frac{3}{4} + \varepsilon$$

- The claim weakens the threshold from $\frac{3}{4} + \varepsilon$ to $\frac{3}{4} + \varepsilon/2$ and expects to find a lot of xs

Proof of the Claim

- Consider a $2^n \times 2^n$ matrix where the rows are indexed by x and the columns are indexed by r. The (x, r)-th entry is 1 or depending on whether $\mathcal{A}^*(f(x), r) = \langle x, r \rangle$ or not. The entry that is 1 will be referred to as "shaded"
- The statement

$$\Pr[x \sim U_{\{0,1\}^n}, r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r \rangle] \geqslant \frac{3}{4} + \varepsilon$$

  is equivalent to saying that at least $3/4 + \varepsilon$ fraction of the entries of the matrix are shaded
- We say that "x is below threshold" if the following is true

$$\Pr[r \sim U_{\{0,1\}^n} \colon \mathcal{A}^*(f(x), r) = \langle x, r \rangle] < \frac{3}{4} + \frac{\varepsilon}{2}$$

- This is same as saying that the row corresponding to x is shaded at $< \frac{3}{4} + \frac{\varepsilon}{2}$ fraction of entries

Proof of the Claim

- Suppose all x are below threshold.
  - Then every row is shaded $< \frac{3}{4} + \frac{\varepsilon}{2}$ fraction of entries
  - Therefore, the whole matrix is shaded $< \frac{3}{4} + \frac{\varepsilon}{2}$ fraction of entries
- Suppose all x are below threshold; except one x
  - Then $(2^n - 1)$ rows are shaded $< \frac{3}{4} + \frac{\varepsilon}{2}$ fraction of entries, and one row is shaded $\leqslant 1$ fraction of entries
  - Therefore, the whole matrix is shaded $< \frac{2^n-1}{2^n}\left(\frac{3}{4} + \frac{\varepsilon}{2}\right) + \frac{1}{2^n} \cdot 1 = \left(\frac{3}{4} + \frac{\varepsilon}{2}\right) + \frac{1}{2^n} \cdot \left(\frac{1}{4} - \frac{\varepsilon}{2}\right)$ fraction of entries

Proof of the Claim

- Suppose all (except $\alpha 2^n$) x are below threshold
  - Then $(2^n - \alpha 2^n)$ rows are shaded $< \frac{3}{4} + \frac{\varepsilon}{2}$ fraction of entries, and $\alpha 2^n$ rows are shaded $\leqslant 1$ fraction of entries
  - Therefore, the whole matrix is shaded $\left(\frac{3}{4} + \frac{\varepsilon}{2}\right) + \alpha \cdot \left(\frac{1}{4} - \frac{\varepsilon}{2}\right)$ fraction of entries
- Note that if $\alpha < \varepsilon/2$ then the matrix is shaded at $< \left(\frac{3}{4} + \frac{\varepsilon}{2}\right) + \alpha \cdot \left(\frac{1}{4} - \frac{\varepsilon}{2}\right) < \left(\frac{3}{4} + \frac{\varepsilon}{2}\right) + (\varepsilon/2) \cdot 1 = \frac{3}{4} + \varepsilon$
- This contradicts the fact that the matrix is shaded at $\geqslant \frac{3}{4} + \varepsilon$ fraction of entries
- So, it must be the case that $\alpha \geqslant (\varepsilon/2)$

Using G to Invert

For any $x \in G$, we have the following properties:

- $\Pr_{r \sim U_{\{0,1\}^n}}[\mathcal{A}^*(f(x), r) = \langle x, r \rangle] \geqslant \frac{3}{4} + \frac{\varepsilon}{2}$
- $\Pr_{r \sim U_{\{0,1\}^n}}[\mathcal{A}^*(f(x), r + e_i) = \langle x, r + e_i \rangle] \geqslant \frac{3}{4} + \frac{\varepsilon}{2}$, for all e_i
- Therefore, by union bound, we have

$$\Pr_{r \sim U_{\{0,1\}^n}}[\mathcal{A}^*(f(x), r) + \mathcal{A}^*(f(x), r + e_i) = \langle x, e_i \rangle] \geqslant \frac{1}{2} + \varepsilon$$

Consider the following algorithm $\mathcal{B}(y, i)$

- Let $m = \mathrm{poly}(n/\varepsilon)$
- For $r^{(1)}, \dots, r^{(m)} \sim U_{\{0,1\}^n}$ compute
  $b^{(k)} = \mathcal{A}^*(f(x), r^{(k)}) + \mathcal{A}^*(f(x), r^{(k)} + e_i)$
- Output the majority of $\{b^{(1)}, \dots, b^{(m)}\}$

For a suitable polynomial m, the probability that $\mathcal{B}(y, i)$ outputs x_i (when $x \in G$), is at least $(1 - 2^n)$ [This part uses Chernoff Bound]

Using G to Invert

Consider the following algorithm $\widetilde{\mathcal{A}}(y)$

- Output $(\mathcal{B}(y,1),\ldots,\mathcal{B}(y,n))$

For $x \in G$, the probability that $\widetilde{\mathcal{A}}(y)$ outputs x is at least $1 - n \cdot 2^{-n} \geqslant 1/2$ (using union bound) So, $\widetilde{\mathcal{A}}$ inverts all y with probability $1/2$, if $x \in G$. Therefore, $\widetilde{\mathcal{A}}$ successfully inverts y with probability at least $\frac{|G|}{2^n} \cdot \frac{1}{2} \geqslant \varepsilon/4$