Amplifying Hardness: XOR and Hardcore Lemmas

March 30, 2006

From last time: Assumption 1: $\exists f \in \mathbf{E}$ such that $CC(f) \geq 2^{\epsilon n}$.

Assumption 2: $\exists f \in \mathbf{E}$ such that $CC_{1/2+2^{-\epsilon n}}(f) \geq 2^{\epsilon n}$. That is, for every large enough n and $2^{\epsilon n}$ sized circuit C,

$$\Pr_{x \leftarrow_{\mathrm{R}} \{0,1\}^n}[f(x) = C(x)] \leq \tfrac{1}{2} + 2^{-\epsilon n}$$

Define $R_{C,f}(x) = +1$ if $C(x) = f(x)$ and -1 if $C(x) \neq f(x)$. Then, an equivalent form[1] is that

$$\mathbb{E}_{x \leftarrow_{\mathrm{R}} \{0,1\}^n}[R_{C,f}(x)] \leq 2^{-\epsilon n}$$

Theorem 1 (NW94)**.** *If Assumption 2 holds then* $\mathbf{BPP} = \mathbf{P}$.

Proof was by a pseudorandom generator from $c \log m$-long strings to m-long strings for some constant $c > 1$ or equivalently from ℓ-long strings to $2^{\epsilon \ell}$-long strings for some constant $1 > \epsilon > 0$.

Different range of parameters Define a weaker assumption:

Assumption 3: $\exists f \in \mathbf{E}$ such that $CC_{1/2+2^{-n^\epsilon}}(f) \geq 2^{n^\epsilon}$. That is, for every large enough n and 2^{n^ϵ} sized circuit C,

$$\mathbb{E}_{x \leftarrow_{\mathrm{R}} \{0,1\}^n}[R_{C,f}(x)] \leq 2^{n^\epsilon}$$

Theorem 2 (NW94)**.** *If Assumption 3 holds then* $\mathbf{BPP} = \mathbf{QuasiP} = \mathbf{DTIME}(2^{\mathrm{polylog}(n)})$.

Proof will show this time a pseudorandom generator from ℓ-long strings to 2^{ℓ^ϵ}-long strings or equivalently from $\log^c m$-long strings to m-long strings.

Our goal today: Assumption 3 is still pretty strong in the sense that it says that no circuit can guess $f(x)$ much better than the trivial $1/2$. We will show that it is implied by the seemingly much weaker assumption that there's some function f such that no circuit can compute $f(x)$ with probability $1 - 1/n^c$ for some constant $c > 0$.

Assumption 4: $\exists f \in \mathbf{E}$ such that $CC_{1-n^{-c}}(f) \geq 2^{n^\epsilon}$. That is, for every large enough n and 2^{n^ϵ} sized circuit C,

$$\mathbb{E}_{x \leftarrow_{\mathrm{R}} \{0,1\}^n}[R_{C,f}(x)] \leq 1 - n^{-c}$$

Theorem 3 (Yao)**.** *Assumption 4 implies Assumption 3.*

Yao's XOR Lemma The proof of the theorem follows from the following lemma:

[1] Up to a factor of two which we can ignore.

Lemma 4. *For any* $f : \{0,1\}^n \to \{0,1\}$ *let* $\overline{f} : \{0,1\}^{nk} \to \{0,1\}$ *be defined as follows:* $\overline{f}(x_1,\ldots,x_k) = f(x_1) \oplus f(x_2) \oplus \cdots \oplus f(x_k)$. *If* $CC_{1-\delta}(f) \geq S$ *then for every* $\epsilon > 2(1-\delta)^k$

$$CC_{1/2+\epsilon}(\overline{f}) \geq \tfrac{\epsilon^2}{100\log(1/\delta\epsilon)} S$$

From this lemma, plugging in $\delta = \frac{1}{n^c}$, $\epsilon = 2^{-n^\epsilon/20}$ and $S = 2^{n^\epsilon}$, and $k = n^{c+1}$ we get the theorem.

Proving the XOR Lemma The proof will go through an intersting characterization of functions f that have $CC_{1-\delta}(f) \geq S$.

Aside on distributions and convex combinations If X and Y are distributions over $\{0,1\}^n$ and $\alpha \in [0,1]$ is a number then by $Z = \alpha X + (1-\alpha)Y$ we denote the distribution obtained by taking with probability α a random element of X and with probability $1-\alpha$ a random element of Y. This is called a *convex combination* of X and Y. We have that for every function $f : \{0,1\}^n \to \mathbb{R}$ $\mathbb{E}[f(Z)] = \alpha\mathbb{E}[f(X)] + (1-\alpha)\mathbb{E}[f(Y)]$.

Note that this is not the standard linearity of expectation since αX does not denote multiplying the output of X by α (which makes no sense for a string) but rather if we think of X as a vector of probabilities of 2^n numbers between 0 and 1 then we multiply this vector by α (hence making it sum up to α instead of to 1).

We can generalize this to more than two distribution and we say that Z is a convex combination of $X_1,\ldots,X_k$ if there are non-negative $\alpha_1,\ldots,\alpha_k$ that sum up to one such that Z can be thought of as choosing with probability α_i to output an element of X_i. Again, thinking of Z and $X_1,\ldots,X_k$ as vectors of probabilities in $\mathbb{R}^{2^n}$ we write $Z = \alpha_1 X_1 + \cdots \alpha_k X_k$. Note also that by the standard averaging argument if $\mathbb{E}[Z] \geq \mu$ then there exists some i such that $\mathbb{E}[X_i] \geq \mu$.

For any distribution X, we define max-pr(X) to be the largest probability that a particular element is attained by X. Note that

- if max-pr$(X_1),\ldots,$max-pr$(X_k) \leq \epsilon$ and Z is a convex combination of $X_1,\ldots,X_k$ then max-pr$(Z) \leq \epsilon$. Indeed, for every $x \in \{0,1\}^n$ we have that $\Pr[Z = x] = \alpha_1 \Pr[X_1 = x] + \cdots + \alpha_k \Pr[X_k = x]$ and a weighted average of things smaller than ϵ is smaller than ϵ.
- For every X, max-pr$(X) \geq 2^{-n}$. Indeed, if all elements are attained with probability less than 2^{-n} then the probabilities will sum up to less than one. Note that the only distribution with max-pr $= 2^{-n}$ is the uniform distribution over $\{0,1\}^n$, denoted by U_n.
- If max-pr$(X) \leq \frac{1}{\delta}2^{-n}$ then we can write $U_n = \delta X + (1-\delta)Y$ for some distribution Y. Indeed, δX is a vector summing up to δ in which all numbers are between 0 and 2^{-n}. Thus, we can add some positive vector Y' to δX to form the uniform distribution. The sum of this vector Y' will necessarily be $1-\delta$ and hence Y' is of the form $(1-\delta)Y$ for some probability distribution Y.

  In this case we say that X has *density* δ in $\{0,1\}^n$. One example for such a distribution is the uniform distribution over some subset S of size $\delta 2^n$. In fact (as you'll also see in the exercise) this is a good example to think about as often we can restrict ourselves to such distributions without loss of generality.

Impagliazzo's hard core lemma Suppose that f is a function that is "moderately hard" for S-sized circuits in the sense that $CC_{1-\delta}(f) \geq S$. Intuitively, one can think that the functions could be hard in two forms: **(a)** the hardness is sort of "spread" all over the inputs, and it is roughly $1-\delta$-hard on every significant set of inputs or **(b)** there's a set H of inputs of density roughly δ such that on H the function is *extremely hard* (cannot be computed better than $\frac{1}{2}+\epsilon$ for some tiny ϵ) and on the rest of the inputs the functions may be even very easy.

Surprisingly, it turns out that we can always assume we are in the case (b):

Lemma 5. *Suppose that $CC_{1-\delta}(f) \geq S$ and let $1 > \epsilon > 0$ be any number. Then there exists a distribution H with density $\geq \delta$ such that $CC^H_{1/2+\epsilon}(f) \geq \frac{\epsilon^2 S}{100 \log(\delta\epsilon)}$. That is, for every S' sized circuit where $S' \leq \frac{\epsilon^2 S}{100\log(\delta\epsilon)}$ we have that*

$$\Pr_{x \leftarrow_R H}[C(x) = f(x)] \leq \tfrac{1}{2} + \epsilon$$

We note that it's possible to get the same result for a distribution H that is uniform over some set S of size at least $\delta 2^n$ (just choose x to be in S with probability $\delta 2^n \Pr[H = x]$, you can show that it will be both be of the right size and will be hard for all circuits using Chernoff bounds and a union bound over all circuits).

Proving Yao's XOR lemma from the hard-core lemma

Lemma 6 (Yao's XOR lemma, restated)**.** *For any $f : \{0,1\}^n \to \{0,1\}$ let $\overline{f} : \{0,1\}^{nk} \to \{0,1\}$ be defined as follows: $\overline{f}(x_1, \ldots, x_k) = f(x_1) \oplus f(x_2) \oplus \cdots \oplus f(x_k)$. If $CC_{1-\delta}(f) \geq S$ then for every $\epsilon > 2(1-\delta)^k$*

$$CC_{1/2+\epsilon}(\overline{f}) \geq \tfrac{\epsilon^2}{100\log(1/\delta\epsilon)} S$$

Proof. Let H be the δ-density distribution we get from the hard-core lemma running it with $\epsilon/4$. Thus, we know that

$$\mathbb{E}[R_{C,f}(H)] < \epsilon/2 \tag{1}$$

Write $U_n = \delta H + (1-\delta)Y$, then we have that U_n^k equal to $\alpha_0 Z_0 + \alpha_1 Z_1 + \cdots + \alpha_m Z_m$ where the Z_i's are distributions of k independent copies of either H or Y and the α's sum up to one. We let Z_0 be the distribution Y^k and so $\alpha_0 = (1-\delta)^k$. We have that

$$\sum_{i=0}^{m} \alpha_i \mathbb{E}[R_{C,f}(Z_i)] \geq \epsilon$$

and so

$$\sum_{i=1}^{m} \alpha_i \mathbb{E}[R_{C,f}(Z_i)] \geq \epsilon - (1-\delta)^k \geq \epsilon/2$$

which implies that there exists some $i > 0$ such that

$$\mathbb{E}[R_{C,f}(Z_i)] \geq \epsilon/2$$

Z_i is a distribution of k independent copies $Z_i^1 \cdots Z_i^k$ where each of them is either H or Y and at least one of them, say Z_i^1 is H. By an averaging argument there exists a string z such that

$$\mathbb{E}[R_{C,f}(Z_i^1, z)] = \mathbb{E}[R_{C,f}(H, z)] \geq \epsilon/2$$

but, hardwiring the value z to the circuit C, this implies a contradiction to (1)

$\square$

Proving the hard-core lemma

Lemma 7 (Impagliazzo's hardcore lemma, restated)**.** *Suppose that* $CC_{1-\delta}(f) \geq S$ *and let* $1 > \epsilon > 0$ *be any number. Then there exists a distribution* H *with density* $\geq \delta$ *such that* $CC^{H}_{1/2+\epsilon}(f) \geq \frac{\epsilon^2 S}{100\log(\delta\epsilon)}$*. That is, for every* S' *sized circuit where* $S' \leq \frac{\epsilon^2 S}{100\log(\delta\epsilon)}$ *we have that*

$$\Pr_{x \leftarrow_R H}[C(x) = f(x)] \leq \tfrac{1}{2} + \epsilon$$

Proof. For every circuit C and distribution H define $adv(C,H)$ to be $\mathbb{E}[R_{C,f}(H)]$. Fix S' as above and think of the following game between two parties, which we'll call Russell and Noam.

Noam plays by presenting a circuit C of size S'. Russell plays by presenting a distribution H of density at least δ. At the end Russell pays to Noam $\$adv(C,H)$.

Clearly, if Russell plays second then he can ensure that he never has to pay to Noam any positive amount, since for every circuit C of size S (and in particular S') he can find $\delta 2^n$ inputs on which that circuit is wrong. However, we want to ensure that Russell can ensure that he does not pay more than $\$\epsilon$ *even if he plays first*.

Since what Russell wins in this game Noam loses and vice versa, this game is a a *zero sum game*, for such games we have von-Neumman's min-max theorem that says it does not matter who plays first *as long as we allow randomized moves*. That is, consider the following variant: Noam produces a *distribution* $\mathcal{C}$ of size-S' circuits and Russell produces a distribution $\mathcal{H}$ of distributions of δ-density, and Russell pays Noam

$$\mathbb{E}_{C \leftarrow_{\mathrm{R}} \mathcal{C}, H \leftarrow_{\mathrm{R}} \mathcal{H}}[adv(C,H)]$$

(in fact, since a convex combination of δ-density distributions is a δ-density distribution, we can think of Russell as choosing a single distribution.)

In this game it does not matter who plays first. This can be viewed as follows: let A be a matrix with columns for every possible circuit of Noam and rows for every possible distribution of Russell.[2] We let $A_{C,H} = adv(C,H)$. In the deterministic game Noam chose a column and Russell chose a row. In the probabilistic game Noam and Russell each choose probability vectors, denoted $\vec{p}$ and $\vec{q}$ respectively with non-negative entries summing up to one and the value of the game is $\vec{q}A\vec{p}$. What we need to prove is that if (*) for every probability vector $\vec{p}$ there exists a probability vector $\vec{q}$ such that $\vec{q}A\vec{p} > 0$ then there exists a probability vector $\vec{q^*}$ such that $\vec{q^*}A\vec{p} > 0$ for every $\vec{p}$. (By moving from A to $aA + bI$ for some $a \neq 0, b$, this implies the general theorem for any game). However this follows because $\{A\vec{p}\}$ is a convex set: if $A\vec{p}$ is in this set and $A\vec{q}$ is in this set then so is $\alpha A\vec{p} + (1-\alpha)A\vec{q} = A(\alpha\vec{p} + (1-\alpha)\vec{q})$. Also note by (*), all members of the set have all coordinates non-negative. Let $\vec{x}$ be the vector with smallest two-norm in that set and $\vec{q^*}$ be a normalization of $\vec{x}$ so that it sums up to one. We claim that for every $\vec{y} = A\vec{p}$ it holds that $\langle \vec{q^*}, \vec{y} \rangle > 0$. Indeed, it's enough to

[2] We ignore the fact that there are infinitely many of them, as we can round them. In fact, we can work with finite matrix by using the fact that the δ-density distributions are all convex combinations of uniform distributions over sets of size $\delta 2^n$.

prove that $\langle \vec{x}, \vec{y} \rangle > 0$ but if $\langle \vec{x}, \vec{y} \rangle \leq 0$ for some $\vec{y}$ in the set then for *every* $\alpha > 0$, the vector $\vec{z} = \alpha\vec{x} + (1-\alpha)\vec{y}$ is in the set and so should satisfy that the norm of $\vec{z}$ is at least as large as the norm of $\vec{x}$. However, by taking the definition of the norm squared as the inner product and taking α small enough one can derive a contradiction.

By the reasoning above we see that all we need to prove is that for any *distribution* $\mathcal{C}$ on S'-sized circuits, Russell can come up with a distribution H on inputs such that $\mathbb{E}_{C \leftarrow_{\mathrm{R}} \mathcal{C}}[adv(C, H)] \leq \epsilon$. However, for any such distribution $\mathcal{C}$ construct the following circuit C: choose $C_1, ..., C_t$ for $t = O(\frac{\log(\delta\epsilon)}{\epsilon^2})$ at random from the distribution and take their majority (on any input x, $C(x)$ will return the majority of $C_1(x), \ldots, C_t(x)$. This is a circuit of size $\leq S$ and so we have $\delta 2^n$ inputs on which it makes a mistake. We let H be the distribution over these inputs.

Suppose that $\mathbb{E}_{C \leftarrow_{\mathrm{R}} \mathcal{C}}[adv(C, H)] \geq \epsilon$. This means that for at least an ϵ fraction of the inputs $x \in H$ (i.e., a total of at least $\epsilon\delta 2^n$ inputs) $\mathbb{E}_{C \leftarrow_{\mathrm{R}} \mathcal{C}}[C(x) = f(x)] \geq \epsilon$ let's call such an x a "surprisingly good" x (since the majority of $C_1, \ldots, C_t$ made a mistake on x but a random $C \leftarrow_{\mathrm{R}} \mathcal{C}$ actually has ϵ advantage on x). However, if we choose $C_1, \ldots, C_t$ at random then by the Chernoff bound for every x such that $\mathbb{E}_{C \leftarrow_{\mathrm{R}} \mathcal{C}}[C(x) = f(x)] \geq \epsilon$, the probability that $Maj(C_1, \ldots, C_t)(x) \neq f(x)$ is, say, less than $\epsilon\delta/10$. Thus the expected number of surprisingly good x's is at most $(\epsilon\delta)2^n/10$ and so with probability at least 0.9 there do not exist $\epsilon\delta 2^n$ of them.

$\square$