14.2.2008

PCPP (Assignment Testers)

Let $P \subseteq \{-1,1\}^m$ be some property of m-bit strings (e.g.. the graph represented by the string is 3-colorable). Assume someone tries to convince us that a string w has the property P. However, we can only look at a constant # of bits. To help us the prover has a "proof" that $w \in P$.

Def: A property P of m-bit strings has PCPPs of length $\ell(m)$ if there exists a tester T making $O(1)$ non-adaptive queries to a string "w, π" where $\pi \in \{-1,1\}^{\ell(m)}$ s.t.

(1). If $w \in P$ then $\exists \pi_w$ s.t. $\Pr[T \text{ accepts } w, \pi_w] = 1$.

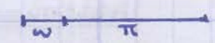

(2). If w is ϵ-far from P, then $\forall \pi$, $\Pr[T \text{ accepts } w, \pi] < 1 - \Omega(\epsilon)$.

Thm: [BHSV + DR] Every property $P \subseteq \{-1,1\}^m$ has PCPPs of length 2^{2^m} with only 3 queries.

Remark: This is about 25% of Dinur's proof of the PCP thm.

Proof: We will ~~expect~~ identify $\{-1,1\}^m$ with $[n] = \{1, .., n\}$ where $n = 2^m$. So $w \in [n]$ and $P \subseteq [n]$. We set π_w to be the truth table of $\chi_{\{w\}} : \{0,1\}^n \to \{-1,1\}$.

The tester should check: (1). π is a dictator on a coordinate in P. (2). This coordinate is w. We achieve this by running one of the following tests with equal probability: (1). The "subset of dictators" test using P on π.

(2). Choose $i \in [m]$ uniformly. Use local decoding on π to compute its value on the string $x \in \{0,1\}^n$ given by $x_j = 1$ iff $j_i = -1$ (where we think of j as both $j \in [n]$ and $j \in \{-1,1\}^m$). Also query w_i and accept if both are equal.

Hardness of Approximation

Influence: For a Boolean function $f: \{0,1\}^n \to \{-1,1\}$ we define the influence of the ith coordinate to be $Inf_i(f) = \Pr_x[f(x) \neq f(x \oplus e_i)]$.

Examples: 1. Dictatorship $\chi_{\{i\}}$: $Inf_i = 1$; $Inf_j = 0$ (for $j \neq i$).

2. Parity : $\forall i.\ Inf_i = 1$. 3. Majority: $\forall i.\ Inf_i \approx \frac{1}{\sqrt{n}}$.

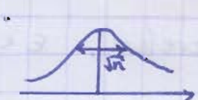

4. Constant function: $\forall i.\ Inf_i = 0$.

Let's try to express Inf in terms of the Fourier coeff. Define $g(x) = f(x \oplus e_i)$. Then, as we saw, $\hat{g}(S) = \hat{f}(S)$ if $i \notin S$, and $\hat{g}(S) = -\hat{f}(S)$ if $i \in S$. Now, define $h(x) = f(x) - g(x)$. Then h is 0 if $f(x) = f(x \oplus e_i)$ and ± 2 otherwise. Hence, $Inf_i(f) = \frac{1}{4} E[h(x)^2] =$

$$= \frac{1}{4}\sum_S \hat{h}(S)^2 = \frac{1}{4}\sum_S (\hat{f}(S) - \hat{g}(S))^2 = \sum_{i \in S} \hat{f}(S)^2.$$

Def: For $f:\{0,1\}^n \to \mathbb{R}$ define $Inf_i(f) = \sum_{S \ni i} \hat{f}(S)^2$.

Total influence: For $f:\{0,1\}^n \to \mathbb{R}$ define its total influence to be $I(f) = \sum_i Inf_i(f)$.

Claim: $I(f) = \sum_S |S| \cdot \hat{f}(S)^2$

Examples: Dictator - 1, Parity - n, Maj - $\sim\sqrt{n}$, Const - 0.

Claim: For any $f:\{0,1\}^n \to \{-1,1\}$ with $E[f]=0$, $I(f) \geq 1$.

Proof: $I(f) = \sum_S |S| \cdot \hat{f}(S)^2 \geq \sum_S \hat{f}(S)^2 = 1$. ∎

Ex: Moreover, if $I(f)=1$ and $E[f]=0$ then f is a dictator or antidictator.

Attenuated influence: Consider a situation where with some prob. $\delta > 0$ each vote is miscounted. In such a situation the expected outcome is $T_{1-2\delta} f(x)$ (and not $f(x)$ as before). So here we should look at $\frac{1}{4} \mathop{E}_x \left[\left(T_{1-2\delta} f(x) - T_{1-2\delta} f(x \oplus e_i) \right)^2 \right] = \sum_{S \ni i} (1-2\delta)^{2|S|} \cdot \hat{f}(S)^2$.

Def: For $-1 \leq \rho \leq 1$ define the ρ-attenuated influence as $Inf_i^{(\rho)}(f) = \sum_{S \ni i} \rho^{|S|} \cdot \hat{f}(S)^2$.

Examples: $\chi_{\{i\}}$: $Inf_j^{(\rho)} = \begin{cases} \rho & j=i \\ 0 & \text{o.w.} \end{cases}$; $\chi_{[n]}$: $\forall i.\ Inf_i^{(\rho)} = \rho^n$ (very small).

Maj : $\forall i$, ρ constant, $Inf_i^{(\rho)} \approx \frac{1}{n}$; $Inf_i^{(\rho)}(f)$ is increasing as ρ goes from 0 to 1.

$Inf_i^{(1)}(f) = Inf_i(f)$; $Inf_i^{(0)}(f) = \hat{f}(\phi)^2$.

Proposition: For $f:\{0,1\}^n \to [-1,1]$, $(0 \leq \delta \leq 1)$ $\sum_{i=1}^n Inf_i^{(\rho)}(f) \leq 1/\delta$ [with $\rho = 1-\delta$], and hence $\#\{i : Inf_i^{(1-\delta)}(f) \geq \epsilon\} \leq 1/\epsilon\delta$.

Proof: First notice that $\forall k \geq 0$, $k \cdot (1-\delta)^k \leq \sum_{i=0}^{\infty} (1-\delta)^i = 1/\delta$.

So, $\sum_{i=1}^n Inf_i^{(1-\delta)}(f) = \sum_{i=1}^n \sum_{S \ni i} (1-\delta)^{|S|} \cdot \hat{f}(S)^2 = \sum_S |S| \cdot (1-\delta)^{|S|} \cdot \hat{f}(S)^2 \leq \frac{1}{\delta} \cdot \sum_S \hat{f}(S)^2 = \frac{1}{\delta} \cdot E[f(x)^2] \leq \frac{1}{\delta}$. ∎

Quasirandomness: We say that a function $f:\{0,1\}^n \to [-1,1]$ is (ϵ,δ)-quasirandom if $\forall i \in [n]$, $Inf_i^{(1-\delta)} < \epsilon$.

Examples: Maj, Parity, Random are all quasirandom. Dictator, 5-junta are not.

Thm: If the Hast-Odd$_\delta$ accepts with prob. $> \frac{1}{2} + \frac{1}{2}\epsilon$ the function $h:\{0,1\}^n \to [-1,1]$, then h is not (ϵ^2, δ)-quasirandom.

Proof: As we saw the acceptance prob of Hast-Odd$_\delta$ is $\frac{1}{2} + \frac{1}{2} \sum_{\substack{S \\ |S| \text{ odd}}} (1-2\delta)^{|S|} \cdot \hat{h}(S)^3$.

Therefore, $\epsilon < \sum_{\substack{S \\ |S| \text{ odd}}} (1-2\delta)^{|S|} \hat{h}(S)^3 \leq \max_{|S| \text{ odd}} \{ (1-2\delta)^{|S|} \hat{h}(S) \} \cdot \sum_{\substack{S \\ |S| \text{ odd}}} \hat{h}(S)^2 \leq \max_{|S| \text{ odd}} \{ (1-2\delta)^{|S|} \cdot \hat{h}(S) \}$.

Therefore $\exists S^* \neq \phi$, for which $(1-2\delta)^{|S^*|} \cdot \hat{h}(S^*) > \epsilon$.

Let $i \in S^*$ be any index in S^*. Then, $Inf_i^{(1-\delta)}(h) = \sum_{S \ni i} (1-\delta)^{|S|} \hat{h}(S)^2 \geq (1-\delta)^{|S^*|} \cdot \hat{h}(S^*)^2 > \epsilon^2$. ∎

We are given a list of equations mod 2 with 3 variables in each. Assume we are promised that there is a solution satisfying 99% of equations. Can we find it?

$$\begin{cases} x_3 + x_9 + x_{10} = 0 \pmod 2 \\ x_2 + x_6 + x_9 = 1 \pmod 2 \\ \vdots \\ x_5 + x_{100} + x_4 = 0 \pmod 2 \end{cases}$$

We can satisfy 50% by assigning it randomly. Can we do 51%? No, it is NP-hard. [Hastad 97].

<u>Label Cover Problem</u>: The label cover problem - we are given a graph and a constraint $\phi_e: [k]\times[k] \to \{F, T\}$ associated to each $e \in E$. Our goal is to assign a value in $[k]$ to each vertex so as to maximize the number of satisfied constraints.

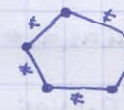

$k=2$
val $=4/5$.

<u>Remark</u>: This is an instance of 2-CSP (Constraint Satisfaction Problem).

<u>Thm</u> (PCP Theorem [AS92, ALMSS92]) + (Parallel Repetition [Raz95])

$\forall \varepsilon > 0, \exists k$, s.t. it is NP-hard given a label cover with alphabet $[k]$ to tell if val $=1$ or val $\le \varepsilon$. Moreover, the graph is bipartite and all constraints are "projections".

<u>Proof outline</u>: The PCP theorem says that it is NP-hard to tell if a given 3CNF is satisfiable or $<99\%$ satisfied. We convert a 3CNF formula to label cover:

$(x_2 \vee \bar{x}_5 \vee x_4) \wedge (x_1 \vee \bar{x}_2 \vee \bar{x}_4) \wedge \dots$

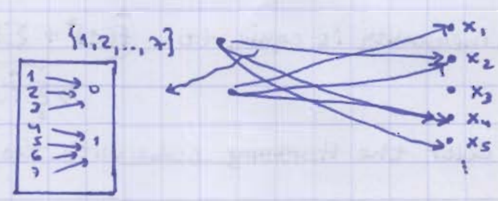

If the 3CNF is satis. then val $=1$.
If the 3CNF is $1-\delta$ satis. (for $\delta = \frac{1}{100}$) then val $\le 1-\delta/3$. We now apply the parallel repetition theorem, which makes $(1, 1-\delta/3)$ into $(1, \varepsilon)$. It works by replacing the label cover instance by a power of itself.

<u>Def</u>: A constraint $\phi: [k]\times[k] \to \{T, F\}$ is called unique if for $\forall a \in [k]$ there exists exactly one $b \in [k]$ s.t. $\phi(a,b) = T$ and vice versa. In other words, ϕ is a permutation on $[k]$.

<u>Claim</u>: There is a poly-time algorithm to check if val $=1$ or not for any given unique-label-cover. (Note that without "uniqueness" this is NP-hard).

<u>Unique Games Conjecture</u>: [Khot 02] $\forall \varepsilon, \delta > 0, \exists k$ s.t. it is NP-hard given a unique label cover instance, to tell if val $\ge 1-\varepsilon$ or val $\le \delta$.

<u>Homework 1 - Solutions</u>

2h. $g(x_1, \dots, x_n) = f(x_1, x_1, x_2, x_2, \dots, x_{\frac{n}{2}}, x_{\frac{n}{2}})$. The transformation $f \mapsto g$ is linear (from $\{0,1\}^n \to \mathbb{R}$ to $\{0,1\}^{n/2} \to \mathbb{R}$). Therefore it suffices to analyze it on $\{\chi_S\}$. $\chi_\phi \mapsto \chi_\phi$; $\chi_{\{1\}} \mapsto \chi_{\{1\}}$; $\chi_{\{2\}} \mapsto \chi_{\{1\}}$; $\chi_{\{1,2\}} \mapsto \chi_\phi$. In general, $\chi_S \mapsto \chi_{\pi_1(S)}$ where $i \in \pi_1(S)$ iff exactly one of $\{2i-1, 2i\}$ is in S. So, $g = \sum_S \hat{f}(S)\chi_{\pi_1(S)}$. $\hat{g}(T) = \sum_{S:\pi_1(S)=T} \hat{f}(S)$.

2d. $\hat{g}(S) = \hat{f}(S+y)$ $\quad (-1)^{\langle x,y\rangle} \cdot \chi_S(x) = \chi_{S+y}(x)$.

2j. $S=\phi$. In this case, $g(x) = \mathbb{E}_{y\in\{0,1\}^n}[f(x,y)]$. $\chi_T \mapsto \begin{cases} \chi_T & T \subseteq I \\ 0 & \text{o.w.} \end{cases}$ $T \subseteq [n]$. Therefore, $g(x) = \sum_{T \subseteq I} \hat{f}(T)\chi_T$.

In general case, $g(x) = \sum_{T \in \bar{I}} \hat{f}(T \cup S) \chi_T$.

5. For $a_1, a_2, a_3 \in \{-1, 1\}$, $NAE(a_1,a_2,a_3) = \frac{3}{4} - \frac{1}{4}a_1a_2 - \frac{1}{4}a_2a_3 - \frac{1}{4}a_1a_3$.

Hence, $\Pr[NAE \text{ accepts } f] = \frac{3}{4} - \frac{1}{4}E[f(x)f(y)] - \frac{1}{4}E[f(y)f(z)] - \frac{1}{4}E[f(x)f(z)]$ where for each i, (x_i, y_i, z_i) is chosen uniformly from $\{001, 010, 100, 011, 101, 110\}$. By symmetry,

$= \frac{3}{4} - \frac{3}{4}E[f(x)f(y)]$ where for each i (x_i, y_i) is chosen according to $\begin{cases} 00 & 1/6 \\ 01 & 1/3 \\ 10 & 1/3 \\ 11 & 1/6 \end{cases}$.

Equivalently, we choose x uniformly and choose $y = x + w$ where $w \sim \mu_{2/3}$. $\Rightarrow = \frac{3}{4} - \frac{3}{4}\langle f, T_{-\frac{1}{3}} f\rangle$.

$= \frac{3}{4} - \frac{3}{4}\sum_S (-\frac{1}{3})^{|S|} \hat{f}(S)^2$.

3a. Estimate $\hat{f}(S)$ to within $\pm\varepsilon$ with conf. of $1-\delta$. $\hat{f}(S) = E[f(x)\chi_S(x)]$. By Chernoff, number of samples is $O(\log 1/\delta / \varepsilon^2)$.

6. $\sum_{S \in [n]} \hat{f}(S)^2 = E_x[f(x)^2]$, so again Chernoff.

7. $I(f) = \sum_S |S| \hat{f}(S)^2$, $4 \cdot \Pr[f=1] \cdot \Pr[f=-1] = 4 \cdot \frac{1}{2}(1+\hat{f}(\phi)) \cdot \frac{1}{2}(1-\hat{f}(\phi)) = 1 - \hat{f}(\phi)^2$.

So the inequality is equiv. to $\underbrace{\hat{f}(\phi)^2 + \sum |S| \cdot \hat{f}(S)^2}_{\geq \sum_S \hat{f}(S)^2 = 1} \geq 1$.

If we color the Hamming cube with two colors in a balanced way, then # of non-monochromatic edges is $\geq 2^{n-1}$! In general $E(V', V-V') \geq \frac{4|V'| \cdot |V-V'|}{2 \cdot 2^n}$.