Limits of Approximation Algorithms 22 April, 2010 (TIFR)

Lec. 12: Unique-Games Hardness of MAXCUT

Lecturer: Prahladh Harsha *Scribe: Girish Varma*

The main topic of today's lecture is the inapproximability of MAXCUT. Recall that MAXCUT has a $\alpha_{GW}\left(=\min_{-1<\rho<1}\frac{2\cos^{-1}(\rho)}{\pi(1-\rho)}=0.87856\right)$-approximation algorithm due to Goemans and Williamson. We will show that if we assume the hardness of a an easier problem than Label-Cover (in fact, a special case called Unique-Label-Cover, we can show that $\text{GAP}_{\frac{1-\rho}{2}-\varepsilon,\frac{\cos^{-1}(\rho)}{\pi}+\varepsilon}$-MAXCUT is hard for every $\rho\in(-1,0)$ and $\varepsilon>0$ [KKMO07]. In particular, choosing the appropriate ρ, we can show that the Goemans-Williamson algorithm is tight.

The references for this lecture include Lecture 9 of the DIMACS tutorial on Limits of approximation [HC09], lectures 17,18 and 19 from a course on PCPs by Venkatesan Guruswami and Ryan O'Donnell at the University of Washington, Seattle [GO05], the original MAXCUT inapproximability paper [KKMO07], and a guest column in SIGACT News on inapproximability results from long codes by Subhash Khot [Kho05].

12.1 Dictator Tests

As in the case MAX3LIN2, we will first design a tailor made dictator test which has the same predicate as MAXCUT (i.e., it is a 2-query test with the $\neq$ predicate) and then compose this test with the Label-Cover problem. In this section, we will be concerned with constructing a suitable dictator test which distinguishes between dictators and functions "far-from-dictators" and defer the actual inapproximability result to later.

Given a boolean function $f:\{0,1\}^m\to\{0,1\}$, we will check whether it is a dictator function(ie. $f(x)=x_i$ for some $i\in[n]$) or not using only values of f at two random points in its domain. We will translate f to the ±1 world by considering 0 as 1 and 1 as -1. So from now on will assume it to be of the form $f:\{\pm1\}^m\to\{\pm1\}$. Furthermore the test must be constrained by the fact that it can query only 2 locations and accept if the value of the function is not the same at these locations.

Below is a first stab at such a dictator test.

$\neq$-test (first attempt)

1. Pick x uniformly at random from $\{\pm1\}^m$
2. Set $y\leftarrow -x$
3. Accept iff $f(x)\neq f(y)$

Though dictators pass this test, any χ_S with $|S|$ odd will also pass the test with probability 1. Another function that passes this test is the majority function($f(x)=\text{sign}(\sum_i x_i)$) assuming n is odd. In fact, any odd function (i.e, $f(-x)=-f(x)$) will pass this test with probability 1.

12.1.1 2-query $\neq$-test for Dictators

Khot, Kindler Mossel and O'Donnell [KKMO07] suggested the following test which picks a pair of points (x, y), which instead of being antipodal as above, satisfies that $\mathbb{E}[\text{dist}(x, y)] = (1-\rho)/2$ for some $\rho \in (-1, 0)$.

$\neq_\rho$-test

1. Pick $x \leftarrow_R \{\pm 1\}^m$
2. Set μ as follows:
   $$\forall i \in [m], \text{ indepedently set } \mu_i \leftarrow \begin{cases} -1 & \text{with probability } \frac{(1-\rho)}{2} \\ 1 & \text{with probability } \frac{(1+\rho)}{2}. \end{cases}$$
3. Set $y \leftarrow x\mu$
4. Accept iff $f(x) \neq f(y)$

Since $\rho \in (-1, 0)$, with probability more that half, each μ_i is -1 and $\mathbb{E}[\mu_i] = \rho$.

Completeness: If $f(x) = x_i$ for some $i \in [m]$ (i.e., f is a dictator) then $\neq_\rho$-test passes when

$$f(x) \neq f(y) \iff x_i \neq x_i\mu_i \iff \mu_i \neq 1$$

which happens with probability $(1-\rho)/2$, which is greater than $1/2$ since $\rho \in (-1, 0)$.

Soundness: We will now try to understand the probability of passing the $\neq$-test by functions that are not dictators. Let us first understand the acceptance probability of $\neq_\rho$-test.

$$\begin{aligned} \Pr[\neq_\rho\text{-test accepts } f] &= \Pr_{x,\mu}[f(x) \neq f(\mu x)] \\ &= \mathbb{E}_{x,\mu}\left[\frac{1 - f(x)f(\mu x)}{2}\right] \\ &= \frac{1}{2} - \frac{1}{2}\mathbb{E}_{x,\mu}[f(x)f(\mu x)] \end{aligned}$$

It will be convenient to define $\mathbb{E}_{x,\mu}[f(x)f(\mu x)]$ as the stability $\text{Stab}_\rho(f)$ of a function f.

Definition 12.1.1 (Noise Stability)**.** *For a (noise) parameter $\rho \in (-1, 1)$, the stability of a function $f : \{\pm 1\}^m \to \{\pm 1\}$ denoted by $Stab_\rho(f)$ is defined as*

$$Stab_\rho(f) = \mathbb{E}_{x,\mu}[f(x)f(\mu x)]$$

where x is picked uniformly at random from $\{\pm 1\}^m$ and $\mu \in \{0, 1\}^m$ such that each μ_i is -1 with probability $(1-\rho)/2$ and 1 otherwise.

It can be easily shown that

$$\text{Stab}_\rho = \sum_S \hat{f}^2(S)\rho^{|S|}$$

where $\hat{f}(S)$ are the Fourier coefficients of f. Thus,

$$\Pr[\neq_\rho\text{-}\mathsf{test} \text{ accepts } f] = \frac{1}{2} - \frac{1}{2}\mathrm{Stab}_\rho(f).$$

It is easy to see that dictators have stability ρ. Furthermore, for any linear functions χ_S, stability goes to 0 as $|S|$ becomes large, which means that dictators pass the $\neq_\rho$-test with probability $(1-\rho)/2$ as noted above and χ_S with probability $(1-\rho^{|S|})/2$ which goes to $1/2$ as $|S|$ becomes large (i.e., they essentially behave like random functions with respect to this test). Thus, at the very least, the test distinguishes dictators from other linear functions which are functions of a large number of variables.

Another function to keep in mind is the majority function (i.e, $\mathrm{majority}_m(x_1,\ldots,x_m) = \mathrm{sign}(\sum x_i)$) which passes the test with probability approaching $\cos^{-1}(\rho)/\pi$ as $m \to \infty$. Let us see how various other functions behave with respect to this test[1].

Remark 12.1.2.

- *Dictators and their negations pass the test with probability* $(1-\rho)/2$.
- *The constant functions passes the test with probability 0.*
- *The majority function passes the test with probability approaching* $\frac{\cos^{-1}\rho}{\pi}$ *as* $m \to \infty$*..*
- *Functions of the form* $f(x) = sign(\sum_{i=1}^{m} a_i x_i)$ *pass the test with probability* $\approx \frac{\cos^{-1}\rho}{\pi}$*, for most choices of* a_i*'s.*
- *The majority function on a small subset of bits, say* $maj_3(x) = maj(x_1, x_2, x_3)$ *passes the test with probability* $p = \frac{1}{2} + \frac{3}{8}\rho - \frac{1}{8}\rho^3$ *which is closer to* $(1-\rho)/2$*. So these numbers are in the order* $\frac{\cos^{-1}\rho}{\pi} < p < (1-\rho)/2$*.*
- *The behaviour of the function* $f(x) = sign(Ax_1 + x_2 \cdots x_n)$ *varies according to parameter* A*. For* $A = 1$*, it is just the majority, but as* A *increases, it becomes more dependent on* x_1*. So the probability of* f *passing the test also increases, and above a critical regime of* $A = \Theta(\sqrt{n})$*, the probability is close to the one for dictators. Below this regime, it is still around* $\frac{\cos^{-1}\rho}{\pi}$*.*

All the above seem to suggest that there are functions like majority which also pass the test with probability significantly more than 1/2. But can there be worse examples? In fact the above seem to suggest, that only functions in which some variable is considerably more "influential" then the others (eg: dictator, maj_3 have acceptance probability in the range $(\cos^{-1}(\rho)/\pi, (1-\rho)/2)$. Informally, we would like to show the following.

Theorem 12.1.3. *(Informal) If* $\neq_\rho$-test *passes with probability* $\geq \frac{\cos^{-1}\rho}{\pi} + \varepsilon$ *then* f *is in some way "similar" to a dictator (i.e., there is an influential variable).*

Let us first define what we mean by an "influential" variable

[1] These examples are from Ryan O'Donnell's lecture notes (lecture 18 in the course on PCPs by Venkatesan Guruswami and Ryan O'Donnell at the University of Washington, Seattle [GO05]).

Definition 12.1.4 (Influence)**.** *The influence of the ith variable on* $f : \{\pm 1\}^m \to \{\pm 1\}$ *denoted by* $Inf_i(f)$ *is defined as*

$$Inf_i(f) = \Pr_x[f(x) \neq f(xe_i)] = \mathbb{E}_x\left[\frac{1 - f(x)f(xe_i)}{2}\right] = \sum_{S: i \in S} \hat{f}(S)^2$$

where $e_i \in \{\pm 1\}^m$ *with* -1 *only in the ith position.*

Let us see the influences of variables for some functions

$$\mathrm{Inf}_i(x_j) = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}$$

$$\mathrm{Inf}_i(\chi_S(x)) = \begin{cases} 0 & \text{if } i \notin S \\ 1 & \text{if } i \in S \end{cases}$$

$$\mathrm{Inf}_i(\mathrm{majority}_m) = \binom{m-1}{(m-1)/2}\frac{1}{2^{m-1}} = \Theta(1/\sqrt{m})$$

$$\mathrm{Inf}_i(\mathrm{parity}_m) = 1$$

Given the above behaviour of various functions with respect to the $\neq_\rho$-test, Khot, Kindler, Mossel and O'Donnell [KKMO07] conjectured the following theorem, which was later proved by Mossel, O'Donnell and Oleszkiewicz [MOO05]

Theorem 12.1.5 ("Majority is the Stablest"(MIS) [MOO05, KKMO07])**.** *For* $-1 < \rho < 0$ *and* $\varepsilon > 0$ *there exists* $\tau > 0$*, such that if* $f : \{\pm\}^m \to \{\pm 1\}$ *has* $Inf_i(f) < \tau$, $\forall i \in [m]$*, then*

$$\Pr[f \text{ passes } \neq_\rho\text{-test }] < \frac{\cos^{-1}\rho}{\pi} + \varepsilon.$$

The above theorem states that if f passes the test with probability more that $\frac{\cos^{-1}\rho}{\pi} + \varepsilon$ then there exists $i \in [m]$ such that the ith variable has high influence ($\mathrm{Inf}_i(f) \geq \tau$). It is in this sense that we said in that f is similar to a dictator function.

To prove hardness of MAXCUT, we will need a strengthening of the above "Majority is the Stablest" theorem that applies to more general $f : \{\pm 1\}^m \to [-1, 1]$. These are functions which can be viewed as randomized functions (i.e., given $f : \{\pm 1\}^m \to [-1, 1]$, it defines a function $f'(x)$ that is 1 with probability $(1 + f(x))/2$ and -1 with probability $(1 - f(x))/2$.)

Furthermore, we will need to work with low-degree influence instead of influence (i.e., sum of only the low order Fourier coefficients). For $1 \leq d \leq m$¡ define the d-degree influence as follows:

$$\mathrm{Inf}_i^{\leq d}(f) = \sum_{S: i \in S, |S| \leq d} \hat{f}(S)^2.$$

The following genalization of the MIS theorem can be proven from the above more standard version.

Theorem 12.1.6 (Generalized MIS Theorem)**.** *For any* $-1 < \rho < 0$ *and* $\varepsilon > 0$*, there exists* $\tau > 0$ *and* $d < \infty$ *such that if* $f : \{\pm 1\}^m \to [-1, 1]$ *has* $Inf_i^{\leq d}(f) < \tau$ $\forall i \in [m]$*, then*

$$\frac{1}{2} - \frac{1}{2}\sum_S \rho^{|S|}\hat{f}^2(S) < \frac{\cos^{-1}\rho}{\pi} + \varepsilon.$$

12.2 Unique Games

In the previous lectures, we proved that for any δ, it is NP-Hard to distuinguish between instances of LABEL-COVER for which all the constrants are satisfied and only δ fraction are satisfied. This was then used to show that MAX3LIN2 is NP-hard to approximate beyond a factor of $1/2+\varepsilon$ for any $\varepsilon > 0$. For the case of MAXCUT, this can be used to show that it is NP-hard to approximate to factors better than $16/17+\varepsilon$ for all $\varepsilon > 0$. It was found that assuming a stronger version of the above result called the Unique Games Conjecture, tighter inapproximability of many problems could be proved. In particular one could prove that the Goemans-Williamson algorithm for MAXCUT is optimal. For stating this conjecture we will define a restriction of the LABEL-COVER problem where the labels for each side of the bipartite graph are from the same set and the constraints are permutations.

Definition 12.2.1 (UNIQUELC(m)). *An instance of the* UNIQUELC(m) *problem is of the form*

$$I = (G = (U, V, E), \Pi = \{\pi_{(u,v)} : [m] \to [m] \text{ for } (u, v) \in E\})$$

where G is a bipartite graph with n vertices on each side, $[m]$ is a set of labels and π_e's are permutations on the set $[m]$ which are to be viewed as constraints on the labellings of vertices of G. The objective is to find a labelling $\sigma : U \cup V \to [m]$ that satisfies (ie. $\pi_{(u,v)}(\sigma(u)) = \sigma(v)$) the maximum number of constraints ie. π_e's.

Given an instance of UNIQUELC(m), one can check if it has a labelling that satisfies *all* the constraints in polynomial time. The algorithm is as follows: For each connected component in G, pick a vertex and try all possible labels to the vertex and for each, try to label all other connected vertices satisfying the constraints. But this doesnt say anything about the following problem.

12.2.1 Unique Games Conjecture

Definition 12.2.2 ($\text{GAP}_{1-\delta,\delta}$-UNIQUELC($m$)). $\text{GAP}_{1-\delta,\delta}$-UNIQUELC($m$) *is a promise problem with set of yes and no instances given by*

$$YES = \{I : \exists \text{ labelling } \sigma \text{ such that it satisfies } \geq 1-\delta \text{ fraction of } \pi_e's\}$$

$$NO = \{I : \forall \text{ labellings } \sigma, \text{ only } \leq \delta \text{ fraction of constraints are satisfied}\}$$

Khot showed that if the above problem is hard then it implies the hardness of several other problems. So, it is natural to consider the following conjecture.

Theorem 12.2.3 (Unique Games Conjecture [Kho02]). *For any $\delta > 0$, there exists a sufficiently large m such that* $\text{GAP}_{1-\delta,\delta}$-UNIQUELC($m$) *is not in P.*

We will say that a problem Γ is UG-hard if for some δ and all m there exists a polynomial time reduction from $\text{GAP}_{1-\delta,\delta}$-UNIQUELC($m$) to the problem Γ. Note that if a problem Γ is UG-hard, then $\Gamma \notin P$ assuming the UGC.

Remark 12.2.4.
- *At present, there is no compelling reason to believe or disbelieve the UGC conjecture. On the other hand, the UGC implies the tight inapproximability results for several problems. On the algorithmic front, there have been several efforts to refute the UGC including a recent subexponential algorithm due to Arora, Barak and Steurer [ABS10]. We will discuss some of these algorithms in the next lecture.*

- *Assuming UGC, one can prove that is hard to distinguish between instances of MAX2LIN(q) which are at least $1-\delta$-satisfiable and those that are at most $1/q^{\delta/2+\Omega(\delta^2)}$-satisfiable. MAX2LIN($q$) refers to the problem of finding the assignment that satisfies the most number of linear equations $\pmod q$ where each equation is over two variables. Observe that MAX2LIN(q) is itself a specific type of UG, where the constraints are linear constraints of the $ax + by = c \pmod q$.*

- *Refuting/proving UGC is a major open problem and the best known polynomial time algorithm for approximating* UNIQUELC *are due to Charikar, Makarychev and Makarychev [CMM06]. They show that there is an efficient algorithm A such that $\forall \delta$, given a* UNIQUELC *instance I which satisfies $1-\delta$ of the constraints, A outputs a labelling that satisfies atleast $1/m^{\delta/2+O(\delta^2)}$ of the constraints. Any improvement on this result will disprove UGC as can be noted from the previous remark.*

12.3 UG-hardness of MAXCUT

In this section we will prove the following theorem which clearly gives the optimality of the Goemans-Williamson algorithm under the unique games conjecture.

Theorem 12.3.1 (UG-hardness of MAXCUT [KKMO07])**.** *For all $\rho \in (-1,0), \varepsilon \in (0,1)$,*

$$\text{GAP}_{\frac{1-\rho}{2}-\varepsilon, \frac{\cos^{-1}\rho}{\pi}+\varepsilon}\text{-MAXCUT is UG-hard.}$$

Setting $\rho = -06934..$, we get that Goemans-Williamson is the best approximation algorithm for MAXCUT assuming UGC. We will first see a different way of looking at the UG-Hardness of $\text{GAP}_{c,s}$-MAXCUT.

Claim 12.3.2. *$\text{GAP}_{c,s}$-MAXCUT is UG-Hard iff there exist δ and for all m there is a 2 query PCP over $\{0,1\}$ for $\text{GAP}_{1-\delta,\delta}$-UNIQUELC($m$) with a $\neq$ predicate with completeness c and soundness s.*

Proof. $(\Rightarrow)$: First reduce the instance I of $\text{GAP}_{1-\delta,\delta}$-UNIQUELC to an instance $G = (V, E)$ of MAXCUT. The binary PCP is a string indexed by elements of V, obtained from the max cut of G by setting the entries of all vertices on one side of the cut to 1 and -1 for the other side. The verifier randomly picks an edge, queries the value of the ends and accepts iff the values are not equal.
$(\Leftarrow)$: Given a PCP π of the said form, one can construct a graph $G(V, E)$ and weights W as follows. V is $[|\pi|]$, ie the set of locations in the proof. An edge (u, v) is give a weight equal to the probability that the verifier generates (u, v) as query locations. □

Thus, to prove Theorem 12.3.1 it suffices to construct a PCP of the said form with the required completeness and soundness from an instance of $\text{GAP}_{1-\delta,\delta}$-UniqueLC.

Given any $x \in \{\pm 1\}^m$ and $\pi : [m] \to [m]$, let $x \circ \pi \in \{\pm 1\}^m$ denote the string $(x_{\pi(1)}, x_{\pi(2)}, \cdots, x_{\pi(m)})$ (i.e., $(x \circ \pi)_i = x_{\pi(i)}$).

First Attempt: Given an instance of $\text{GAP}_{,}$-UniqueLC, $I = (G(U, V, E), \{\pi_e\})$, the PCP Verifier expects as proofs $f_u : \{\pm 1\}^m \to \{\pm 1\}$ for all $u \in U \cup V$, which in the right proof is supposed to be the long codes of the label of each vertex in $U \cup V$. The PCP verifier then checks the long codes as follows.

1. Pick an edge (u, v) at random from E.
2. Pick x uniformly at random from $\{\pm 1\}^m$.
3. Set $y \leftarrow \mu x$ where $\mu \in \{\pm 1\}^m$ such that μ_i is -1 with probability $(1-\rho)/2$ and 1 otherwise.
4. Accept iff $f_u(x) \neq f_v(y \circ \pi_{(v,u)})$

Clearly, a proof constructed honestly, when the $\text{GAP}_{,}$-UniqueLC is $(1-\delta)$-satisfiable passes the PCP test with probability roughly $(1-\rho)/2$. On the other hand, consider the proof $f_u(x) = 1, \forall u \in U$ and $f_v(x) = -1, \forall v \in V$. Clearly, this proof passes the PCP with probability 1. In fact, we notice that any PCP which makes one query each on either side of the Label-Cover problem is doomed to fail for precisely the above reason. We will instead work with a different PCP verifier, which makes queries into only one side of the Label-Cover problem. In other words, this new PCP verifier expects as proof, $f_v : \{\pm 1\}^m \to \{\pm 1\}$ only for all $v \in V$, which in the good case is supposed to be the long code of the labels of v in V.

MAXCUT-PCP Verifier

Input: UniqueLC(m) instance (G, Π).

Proof: $f_v : \{\pm 1\}^m \to \{\pm 1\}, \forall v \in V$

1. Pick u at random from U with probability proportional to the degree of u.
2. Pick v, v' at random from the neigbours of u. Let $\pi_{(v,u)}, \pi_{(v',u)}$ be the inverses of the corresponding constraints.
3. Pick x uniformly at random from $\{\pm 1\}^m$.
4. Set $y \leftarrow \mu x$ where $\mu \in \{\pm 1\}^m$ such that μ_i is -1 with probability $(1-\rho)/2$ and 1 otherwise.
5. Accept iff $f_v(x \circ \pi_{(v,u)}) \neq f_{v'}(y \circ \pi_{(v',u)})$

Completeness: Suppose the instance of LABEL-COVER that we started of with had a labelling σ satisfying $(1-\delta)$ fraction of the constraints. Now the probability that σ satisfy both (u,v) and $(u,v^{'})$ is $\geq 1-2\delta$. Suppose this is the case then the test accepts when

$$f_v(x \circ \pi_{(v,u)}) \neq f_{v'}(y \circ \pi_{(v',u)}) \iff (x \circ \pi_{(v,u)})_{\sigma(v)} \neq (y \circ \pi_{(v',u)})_{\sigma(v')} \iff x_{\pi_{(v,u)}(\sigma(v))} \neq (\mu x)_{\pi_{(v',u)}(\sigma(v'))}$$

Since $\pi_{(v,u)}(\sigma(v)) = \sigma(u) = \pi_{(v',u)}(\sigma(v^{'}))$, this happen when $\mu_{\sigma(u)} = -1$. Hence the completeness is

$$(1-2\delta)(1-\rho)/2 \geq \frac{1-\rho}{2} - \varepsilon$$

since $\varepsilon > 2\delta$.

Soundness:

Claim 12.3.3. *Given ρ, ε, let τ and d be as in Generalized MIS Theorem 12.1.6 for parameters $\varepsilon/2$ and ρ. Let $\delta = \varepsilon\tau^2/8d$. If*

$$\Pr[\textsf{MAXCUT-PCP} \textit{ accepts}] \geq \frac{\cos^{-1}\rho}{\pi} + \varepsilon$$

then there exists a labelling σ that satisfies more that δ fraction of the constraints.

Proof. Suppose

$$\Pr[\textsf{MAXCUT-PCP} \text{ accepts}] = \mathbb{E}_{u,v,v',x,\mu}\left[\frac{1 - f_v(x \circ \pi_{(v,u)}) f_{v'}(x\mu \circ \pi_{(v',u)})}{2}\right] \geq \frac{\cos^{-1}\rho}{\pi} + \varepsilon$$

By an averaging argument there exists at least an $\varepsilon/2$ fraction of u's for which

$$\mathbb{E}_{v,v'}\mathbb{E}_{x,\mu}\left[\frac{1 - f_v(x \circ \pi_{(v,u)}) f_{v'}(x\mu \circ \pi_{(v',u)})}{2}\right] \geq \frac{\cos^{-1}\rho}{\pi} + \varepsilon/2$$

Lets call these the good u's. Let $g_u(x) = \mathbb{E}_v[f_v(x \circ \pi_{(v,u)})]$, then we rewrite the above as

$$\mathbb{E}_{x,\mu}\left[\frac{1 - \mathbb{E}_v[f_v(x \circ \pi_{(v,u)})]\mathbb{E}_{v'}[f_{v'}(x\mu \circ \pi_{(v',u)})]}{2}\right] = \mathbb{E}_{x,\mu}\left[\frac{1 - g_u(x) g_u(x\mu)}{2}\right] \geq \frac{\cos^{-1}\rho}{\pi} + \varepsilon/2$$

Then from the Generalized MIS Theorem 12.1.6, we know that g_u must have atleast one variable j_u with large d degree influence. In other words, for each of the good u's there exists $j_u \in [m]$ such that

$$\mathrm{Inf}_{j_u}^{\leq d}(g_u) \geq \tau$$

So lets assign the label of u, $\sigma(u) = j_u$. Now we need to label the neighbours v of u. For this we relate the Fourier coefficients of g_u and f_v.

$$\begin{aligned} g_u(x) &= \mathbb{E}_v[f_v(x \circ \pi_{(v,u)})] & &= \mathbb{E}_v\left[\sum \hat{f}_v(S)\chi_S(x \circ \pi_{(v,u)})\right] \\ &= \mathbb{E}_v\left[\sum \hat{f}_v(S)\chi_{\pi_{(v,u)}(S)}(x)\right] & &= \sum \mathbb{E}_v\left[\hat{f}_v(\pi_{(v,u)}^{-1}(T))\right]\chi_T(x) \end{aligned}$$

So $\hat{g_u}(T) = \mathbb{E}_v\left[\hat{f}_v(\pi^{-1}_{(v,u)}(T))\right]$. Also

$$\begin{aligned}
\tau \leq \mathrm{Inf}^{\leq d}_{j_u}(g_u) &= \sum_{S: j_u \in S \text{ and } |S| \leq d} \hat{g_u}^2(S) \\
&= \sum_{S: j_u \in S \text{ and } |S| \leq d} \left(\mathbb{E}_v[\hat{f}_v(\pi^{-1}_{(v,u)}(S))]\right)^2 \\
&\leq \sum_{S: j_u \in S \text{ and } |S| \leq d} \mathbb{E}_v\left[\hat{f}_v^{\,2}(\pi^{-1}_{(v,u)}(S))\right] \qquad \text{(by Cauchy Schwarz)} \\
&= \mathbb{E}_v\left[\sum_{S: j_u \in S \text{ and } |S| \leq d} \hat{f}_v^{\,2}(\pi^{-1}_{(v,u)}(S))\right] \\
&= \mathbb{E}_v\left[\mathrm{Inf}^{\leq d}_{\pi^{-1}_{(v,u)}(j_u)}(f_v)\right]
\end{aligned}$$

Thus again by an averaging argument there is a $\tau/2$ fraction of u's neighbours v such that the set $S_v = \{j : \mathrm{Inf}^{\leq d}_{\pi^{-1}_{(v,u)}(j_u)} \geq \tau/2\}$ is non empty. So we will label v, with a random element of S_v. Hence the fraction of edges that are satisfied by the labelling is

$$\Pr_{e \in E}[(u,v) \text{ is satisfied }] \geq \frac{\varepsilon}{2}\frac{\tau}{2}\frac{1}{|S_v|}.$$

The above expression is accounted as follows: with probability at least $\varepsilon/2$ u is good, with probability $\tau/2$, v is a good neighbour of u and with probability $1/|S_v|$, the label $\pi^{-1}(j_u)$ is chosen for v. Now we will show that $|S_u| \leq 2d/\tau$.

$$\sum_{i=1}^{m} \mathrm{Inf}^{\leq d}_i(f_v) = \sum_{|S| \leq d} |S| \hat{f}_v(S)^2 \leq d \sum_S \hat{f}_v(S)^2 = d$$

So $|S_v|\tau/2 \leq d$ and

$$\Pr_{e \in E}[(u,v) \text{ is satisfied }] \geq \frac{\varepsilon}{2}\frac{\tau}{2}\frac{1}{|S_v|} \geq \delta \qquad [\text{since } \delta = \varepsilon\tau^2/8d]$$

□

This completes the proof of Theorem 12.3.1.

References

[ABS10] Sanjeev Arora, Boaz Barak, and David Steurer. *Subexponential algorithms for unique games and related problems*, 2010. (manuscript).

[CMM06] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. *Near-optimal algorithms for unique games*. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, pages 205–214. 2006. doi:10.1145/1132516.1132547.

[GO05] Venkatesan Guruswami and Ryan O'Donnell. *CSE 533: The PCP Theorem and hardness of approximation*, 2005. A course on PCPs at the University of Washington, Seattle (Autumn 2005).

[HC09] Prahladh Harsha and Moses Charikar. *Limits of approximation algorithms: PCPs and unique games*, 2009. (DIMACS Tutorial, July 20-21, 2009). arXiv:1002.3864.

[Kho02] Subhash Khot. *On the power of unique 2-prover 1-round games.* In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 767–775. 2002. doi:10.1145/509907.510017.

[Kho05] ———. *Guest column: inapproximability results via long code based PCPs.* SIGACT News, 36(2):25–42, 2005. doi:10.1145/1067309.1067318.

[KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?* SIAM J. Computing, 37(1):319–357, 2007. (Preliminary version in *45th FOCS*, 2004). eccc:TR05-101, doi:10.1137/S0097539705447372.

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. *Noise stability of functions with low influences invariance and optimality.* In *Proc. 46th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 21–30. 2005. arXiv:math/0503503, doi:10.1109/SFCS.2005.53.