Limits of Approximation Algorithms 15 April, 2010 (TIFR)

Lec. 11: Håstad's 3-bit PCP

Lecturer: Prahladh Harsha *Scribe: Bodhayan Roy & Prahladh Harsha*

In last lecture, we proved the hardness of label cover problem. We showed that it is NP-hard to distinguish if a given label cover instance is perfectly satisfiable or if every labelling satisfies at most a δ-fraction of edges. However, this was over some large alphabet. Today, we will be interested in constructed PCPs over smaller alphabets, in fact, the binary alphabet. We will show Håstad's PCP construction which is a 3 query PCP over the binary alphabet. We will then use this result to show tight inapproximability results for problems such as MAX3SAT, MAX3LIN2.

The references for this lecture include Lecture 7 of the DIMACS tutorial on Limits of approximation [HC09], lecture 16 from a course on PCPs by Venkatesan Guruswami and Ryan O'Donnell at the University of Washington, Seattle [GO05] and a guest column in SIGACT News on inapproximability results from long codes by Subhash Khot [Kho05].

11.1 Recap: Hardness of Label-Cover

Recall that the Label-Cover problem was defined as follows:

Definition 11.1.1 (Label-Cover)**.** *An instance I of the* Label-Cover *problem with labels L and R is specified by a pair (G, Π) where $G = (U, V, E)$ is a bipartite graph, and $\Pi = \{\pi_e : L \to R | e \in E\}$, is a set of functions (also called projections), one for each edge $(u, v) \in E$.*

A labeling $A : U \to L, B : V \to R$, is said to satisfy an edge (u, v) iff $\pi_{(u,v)}(A(u)) = B(v)$. The value *of an instance is the maximal fraction of edges satisfied by any such labeling.*

For any $\delta \in (0, 1)$, the gap problem $\text{gap}_{1,\varepsilon}$-LC$(L, R)$ is the promise problem of deciding if a given instance has value 1 or at most ε. More precisely, the YES *and* NO *of $\text{gap}_{1,\varepsilon}$-LC are given as follows.*

$$\begin{aligned} \mathsf{YES} &= \left\{ I : \exists\, (A : U \to L,\ B : V \to R) \text{ such that } \forall (u,v) \in E,\ \pi_{(u,v)}(A(u)) = B(v) \right\} \\ \mathsf{NO} &= \left\{ I : \forall\, (A : U \to L,\ B : V \to R),\ \left| \{(u,v) \in E : \pi_{(u,v)}(A(u)) = B(v)\} \right| \leq \varepsilon |E| \right\} \end{aligned}$$

In the last few lectures, we proved the following inapproximability result for Label-Cover.

Theorem 11.1.2. *$\forall \delta \in (0, 1)$, there exist alphabets L, R such that $|L|, |R| \leq \exp\left(\frac{1}{\delta}\right)$, such that the corresponding gap problem $\text{gap}_{1,\delta}$-LC(L, R) is NP-hard.*

Note, that the above hardness for Label-Cover is great, but unfortunately is over a large alphabet.

11.2 PCPs over the Boolean alphabet

Using the above NP-hardness of Label-Cover, Håstad constructed the following PCP for NP over the Boolean alphabet [Hås01].

Theorem 11.2.1 (Håstad's 3-bit PCP). *For all $\varepsilon, \delta \in (0,1)$, there exists a PCP for SAT over the Boolean alphabet with the following properties :*

1. *The verifier queries 3 bits and its predicate is of the form $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$.*
2. *Completeness: $1 - \varepsilon$.*
   I.e., if φ is satisfiable, then there exists a PCP π such that $\Pr[V^\pi(\varphi) = 1] \geq 1 - \varepsilon$.
3. *Soundness: $\frac{1}{2} + \delta$.*
   I.e., if φ is unsatisfiable, then for all proofs π, $\Pr[V^\pi(\varphi) = 1] \leq \frac{1}{2} + \delta$

Note that this result is almost optimal. If $\mathsf{NP} \neq \mathsf{P}$, we cannot reduce the number of queries to 2, since the corresponding PCP class, $\mathrm{PCP}[O(\log n), 2]$ is in fact in P (see problem set 1). Furthermore, we also know that the soundness cannot be reduced below 1/2. On the completeness side, we can show that one can actually obtain PCP (albeit adaptive PCPs) with perfect completeness (i.e., 1 as opposed to $1 - \varepsilon$) [GLST98].

Coming back to the non-adaptive PCP in Theorem 11.2.1, because of the nature of thePCP verifier's queries and predicate ($x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$), we have the following hardness result for approximating the number of satisfied equations of a given MAX3LIN2 system.

Corollary 11.2.2 (Hardness of MAX3LIN2). *For all $\varepsilon, \delta \in (0,1)$, it is NP hard to distinguish if a given system of linear equations over $\mathbb{GF}(2)$ (with 3 variables per equation) is at least $(1 - \varepsilon)$-satisfiable or at most $(\frac{1}{2} + \delta)$-satisfiable.*

The above corollary states that it is NP-hard to approximate the number of satisfied linear equations over $\mathbb{GF}(2)$ to a factor better than $1/2 + \varepsilon$ for any $\varepsilon > 0$. Note this is optimal since it is trivial to give a 1/2-approximation, a random assignment will do.

Let us now show how this corollary implies a hardness for MAX3SAT. We will do this by encoding every 3LIN2 condition by a constant number of 3SAT conditions. More precisely, we will do the following. Replace the the 3LIN2 condition $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = 1$ by the following 4 3SAT conditions $x_{i_1} \vee x_{i_2} \vee x_{i_3}$, $x_{i_1} \vee \overline{x_{i_2}} \vee \overline{x_{i_3}}$, $\overline{x_{i_1}} \vee \overline{x_{i_2}} \vee x_{i_3}$ and $\overline{x_{i_1}} \vee x_{i_2} \vee \overline{x_{i_3}}$ and the 3LIN2 condition $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = 0$ by the 4 3SAT conditions $\overline{x_{i_1}} \vee \overline{x_{i_2}} \vee \overline{x_{i_3}}$, $x_{i_1} \vee x_{i_2} \vee \overline{x_{i_3}}$, $x_{i_1} \vee \overline{x_{i_2}} \vee x_{i_3}$ and $\overline{x_{i_1}} \vee x_{i_2} \vee x_{i_3}$. It is easy to check that if the source 3LIN2 condition is satisfied, all the 4 target 3SAT conditions are satisfied, while if source 3LIN2 condition is not satisfied, exactly 3 of the 4 3SAT conditions are satisfied. Using this transformation, if the MAX3LIN2 is α-satisfiable, then the target MAX3SAT instance is exactly $(\frac{\alpha+3}{4})$-satisfiable (since $\alpha \cdot 1 + (1 - \alpha) \cdot \frac{3}{4} = \frac{\alpha+3}{4}$). We thus have the following corollary.

Corollary 11.2.3 (Hardness of MAX3SAT). *For all $\varepsilon, \delta \in (0,1)$, it is NP hard to distinguish if a MAX3SAT instance is at least $(1 - \varepsilon)$-satisfiable or at most $(\frac{7}{8} + \delta)$-satisfiable.*

Once again, this is optimal, since it is trivial to get a 7/8-approximation for MAX3SAT, a random assignment will do.

11.3 Proof Overview of Theorem 11.2.1

How does one convert the hardness of Label-Cover into a 3-query PCP with the following property? If the instance is an YES instance, then there exists a proof such that the verifier

must accept the PCP with probability at least $1-\varepsilon$, and if it is a NO instance, then for any proof,the verifier must accept with probability at most $\frac{1}{2}+\delta$.

In the LABEL-COVER world, the verifier could expect as proof the label for every vertex in $U \cup V$. However, this is not over the Boolean alphabet. So, instead the verifier now expects the prover to provide the encoding of every label over a binary alphabet (say given by f_u). Given such an encoding the verifier needs to check two things for a random $(u,v) \in E$.

- *Codeword Test:*
  There exist $l \in L$ and $r \in R$ such that f_u is a valid encoding of l and f_v is a valid encoding of r.
- *Consistency test:*
  The above l and r obey $\pi_{(u,v)}(l) = r$.

Note that the verifier is allowed only one test of the form $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$ to do both the above checks. What is a suitable encoding that allows the verifier to do all of this?

11.3.1 The Long Code

Bellare, Goldreich and Sudan [BGS98] proposed the *long code* as a natural encoding to do the above tests. The long code, as the name suggests, is the longest possible binary encoding without any redundancy. The long code is defined as follows. A label $l \in L$ is mapped to a string of length $2^{|L|}$. Given any $f \in [2^{|L|}]$, we can interpret this f as a function $f : L \to \{1,-1\}$. The f-th bit of the long code of l is defined to be $f(l)$. Thus, the long code of l is the concatenation of the evaluations of every possible function from L to $\{1,-1\}$ at the the point l, i.e., $\text{long}_l = (f(l)|f : L \to \{1,-1\})$. Notice that the long code maps a string of length $\log |L|$ to a string of length $2^{|L|}$ (a doubly exponential blowup!) The long_l is a mapping from the set of functions from L to $\{1,-1\}$ to $\{1,-1\}$.

Long code can be alternately viewed as dictator functions. A function $f : \{1,-1\}^m \to \{1,-1\}$ is said to be a dictator if there exists $i \in [m]$ such that $f(x_1,\ldots,x_m) = x_i$ for all $(x_1,\ldots,x_m) \in \{1,-1\}^m$. Recall that a long code is a function $\text{long}_l : \{1,-1\}^L \to \{1,-1\}$ where $\text{long}_l(x_1,\ldots,x_L) = x_l$ Here we are viewing the functions from L to $\{1,-1\}$ in its truth table form – a string of length $2^{|L|}$. Depending on the context, we will refer to long-codes as dictators or vice-versa.

We now need to design a suitable test that will both check if the given words are actually the long codes of some labels and furthermore that these labels satisfy the projection constraint. We will first see how we can achieve just the former goal (check if a given function is dictator) and later see how we can adapt this test to also perform the consistency test.

11.4 Dictator Tests/Long-code Tests

We need to check if a given function $f : \{1,-1\}^m \to \{1,-1\}$ is a dictator test. Furthermore, we are constrained by the fact that the test can query the function at only 3 locations and the acceptance predicate must be of the form $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$. We now observe that dictator functions are linear functions, in fact, $\text{long}_l = \chi_{\{l\}}$. That is, the long codes are precisely

the linear functions corresponding to the singleton sets. Recall that the BLR-test checks if a given function is linear and furthermore, it is precisely of the form "$x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$".

$\mathsf{BLR\text{-}Test}^f$: 1. Choose $x, y \in_R \{1,-1\}^m$
2. Accept if $f(x)f(y) = f(xy)$.

Clearly, dictators, being linear functions, pass the above test with probability 1 and functions that are "far from linear" are rejected. On the other hand, other linear functions such as χ_S for non-singleton sets S also pass the above test with probability 1. How does one modify the above test such that the acceptance probability corresponding to χ_S for non-singleton S is reduced. Håstad suggested the following modification by adding a slight noise to the BLR-Test which reduces the acceptance probability for χ_S for large $|S|$ by compromising on perfect completeness.

ε-$\mathsf{perturbed\text{-}BLR\text{-}Test}^f$: 1. Choose $x, y \in_R \{1,-1\}^m$
2. Pick $\mu \in \{1,-1\}^L$ (noise vector) as follows:
$$\mu_i \longleftarrow \begin{cases} 1 & \text{with probability } 1-\varepsilon \\ -1 & \text{with probability } \varepsilon \end{cases}$$
3. Set $z \leftarrow xy\mu$, (i.e., $z_i = x_i y_i \mu_i, \forall i \in [m]$)
4. Accept if $f(z) = f(x)f(y)$

Clearly, if f is a dictator (i.e, $f = \chi_{\{i\}}$) then the probability that f passes the test is precisely the probability that $\mu_i = 1$, which is $1-\varepsilon$. We thus have the following completeness claim.

Claim 11.4.1 (Completeness)**.** *If $f = \chi_{\{i\}}$ for some $i \in [m]$ (i.e., f is a dictator), then*

$$\Pr_{x,y,\mu}[f(z) = f(x)f(y)] = 1 - \varepsilon.$$

Soundness of the ε-perturbed-BLR-Test is given by the following claim.

Claim 11.4.2 (Soundness)**.** *Suppose $f : \{1,-1\}^m \to \{1,-1\}$ satisfies*

$$\Pr[\ f \text{ passes } \varepsilon\text{-perturbed-BLR-Test }] \geq \frac{1}{2} + \eta$$

then f "resembles" a dictator in the following sense:

$$\exists S \subseteq [L], \text{ such that } |S| \leq O\left(\frac{1}{\varepsilon}\log\frac{1}{\eta}\right) \text{ and } |\hat{f}(S)| \geq 2\eta.$$

In other words, if f passes the εperturbed-BLR-Test with probability bounded away from $1/2$, then it must be the case that f is co-related with a linear function χ_S where $|S|$ is small (i.e, even though f is not a dictator it is very much like a function of a few variables)!

Proof. We can write the acceptance probability of εperturbed-BLR-Test as follows.

$$\Pr[\ f \text{ passes } \varepsilon\text{-perturbed-BLR-Test }] = \mathop{\mathbb{E}}_{x,y,\mu}\left[\frac{1+f(x)f(y)f(z)}{2}\right].$$

Thus, if $\Pr[\text{ acc }] \geq 1/2+\eta$, it must be the case that $\mathbb{E}_{x,y,\mu}[f(x)f(y)f(x,y\mu)] \geq 2\eta$. Using the Fourier expansion $f(x) = \sum \hat{f}_S \chi_S(x)$, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{x,y,\mu}[f(x)f(y)f(xy\mu)] &= \sum_{S,T,U} \hat{f}_S\hat{f}_T\hat{f}_U \cdot \mathop{\mathbb{E}}_{x,y,\mu}[\chi_S(x)\chi_T(y)\chi_U(xy\mu)] \\
&= \sum_{S,T,U} \hat{f}_S\hat{f}_T\hat{f}_U \cdot \mathop{\mathbb{E}}_{x}[\chi_S(x)\chi_U(x)] \cdot \mathop{\mathbb{E}}_{y}[\chi_T(y)\chi_U(y)] \cdot \mathop{\mathbb{E}}_{\mu}[\chi_U(\mu)] \\
&= \sum_{S} \hat{f}_S^3 \cdot \mathop{\mathbb{E}}_{\mu}[\chi_S(\mu)]
\end{aligned}$$

We can now compute $\mathbb{E}_{\mu}[\chi_S(\mu)]$ as follows:

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mu}[\chi_S(\mu)] &= \mathop{\mathbb{E}}_{\mu}\left[\prod_{i\in S}\mu_i\right] = \prod_{i\in S}\mathop{\mathbb{E}}_{\mu_i}[\mu_i] \\
&= \prod_{i\in S}(1-2\varepsilon) = (1-2\varepsilon)^{|S|}.
\end{aligned}$$

We thus, have

$$\begin{aligned}
\mathop{\mathbb{E}}_{x,y,\mu}[f(x)f(y)f(xy\mu)] &= \sum_{S} \hat{f}_S^3 \cdot (1-2\varepsilon)^{|S|} \\
&= \sum_{S} \hat{f}_S^2 \cdot \left(\hat{f}_S \cdot (1-2\varepsilon)^{|S|}\right) \\
&= \mathop{\mathbb{E}}_{S\sim\hat{f}_S^2}\left[\hat{f}_S \cdot (1-2\varepsilon)^{|S|}\right]
\end{aligned}$$

Since $\mathbb{E}\left[\hat{f}_S \cdot (1-2\varepsilon)^{|S|}\right] \geq 2\eta$, there exists an S such that $\hat{f}_S \cdot (1-2\varepsilon)^{|S|} \geq 2\eta$. Clearly, this S satisfies that $\hat{f}_S \geq 2\eta$ and $(1-2\varepsilon)^{|S|} \geq 2\eta$ which implies that $|S| \leq O\left(\frac{1}{\varepsilon}\log\frac{1}{\eta}\right)$. □

11.5 Håstad's 3-bit PCP

We now need to compose the above 3-bit dictator test with the Label-Cover problem to obtain a 3-bit PCP. We proceed as follows. Let $I = (G = (U, V, E), \Pi)$ be the Label-Cover. The verifier needs to check if the instance I is an YES or a NO instance of $\text{gap}_{1,\delta}$-LC. For this purpose, the verifier expects as proof the long code of all the labels of the vertices in $U \cup V$ (i.e., $f_u : \{1,-1\}^L \to \{1,-1\}, \forall u \in U$ and $f_v : \{1,-1\}^R \to \{1,-1\}, \forall v \in V$). The verifier checks if all these functions are valid long codes and that the long codes are encodings of label that satisfy the projections Π as follows.

Notation: Given $x \in \{1,-1\}^R$ and $\pi : L \to R$, let $x \circ \pi \in \{1,-1\}^L$ be defined as follows: $(x \circ \pi)_i = x_{\pi(i)}, \forall i \in L$.

Håstad's 3-bit PCP_ε:

Input: Label-Cover instance $I = (G = (U, V, E), \Pi)$.

Proof: $f_u : \{1,-1\}^L \to \{1,-1\}, \forall u \in U$ and $f_v : \{1,-1\}^R \to \{1,-1\}, \forall v \in V$.

1. Pick $(u, v) \in_R E$.
2. Perform the following consistency check for $(f_u, f_v, \pi_{uv}, \varepsilon)$
   (for ease of notation, denote f_u by f, f_v by g, and π_{uv} by π)
   (a) Pick $x \in_R \{1,-1\}^R$, $y \in_R \{1,-1\}^L$.
   (b) Construct $\mu \in \{1,-1\}^L$ as follows
   $$\mu_i = \begin{cases} 1 & \text{with probability } 1-\varepsilon, \\ -1 & \text{with probability } \varepsilon. \end{cases}$$
   (c) Set $z \leftarrow (x \circ \pi) y \mu$.
   (d) Accept if $f(z) = g(x) f(y)$.

Clearly, the above PCP has the desired acceptance predicate. Completeness is given by the proposition below.

Proposition 11.5.1 (Completeness)**.** *If the instance* (G, Π) *is an* YES *instance of* $\mathrm{gap}_{1,\delta}$-$\mathrm{LC}(L, R)$*, then there exists* $f_u : \{1,-1\}^L \to \{1,-1\}, \forall u \in U$ *and* $f_v : \{1,-1\}^R \to \{1,-1\}, \forall v \in V$ *such that* $\Pr[$ Håstad's 3-bit PCP_ε *accepts* $] = 1 - \varepsilon$.

Proof. Let $A : U \to L$ and $B : V \to R$ be the labeling that satisfies all the edges of the Label-Cover instance. Define functions f_u and f_v to be long codes of the labels as follows.

$$f_u(x_1, \ldots, x_L) = x_{A(u)}, \quad f_v(x_1, \ldots, x_R) = x_{B(v)}.$$

It is easy to see that

$$\begin{aligned} \Pr[\text{ Håstad's 3-bit } \mathsf{PCP}_\varepsilon \text{ accepts }] &= \Pr_{(u,v),x,y,\mu} [f_u((x \circ \pi_{u,v}) y \mu) = f_v(x) f_u(y)] \\ &= \Pr\left[(x \circ \pi_{u,v})_{A(u)} \, y_{A(u)} \mu_{A(u)} = x_{B(v)} y_{A(u)} \right] \\ &= \Pr\left[x_{\pi_{u,v}(A(u))} \mu_{A(u)} = x_{B(v)} \right] \\ &= \Pr[\mu_{A(u)} = 1] \qquad [\text{ Since } \pi_{u,v}(A(u)) = B(v) \text{ }] \\ &= 1 - \varepsilon \end{aligned}$$

□

11.5.1 Soundness analysis

We would like to show that if $Pr[acc] > \frac{1}{2} + \eta$ then there exists a labelling $A : U \to L, B : V \to R$ that satisfies at least δ fraction of the edge constraints. But is this true?

11.5.1.1 Folding: Derailed by all the 1's proof

Consider the all 1's proof. In other words the set of functions $f_u(x) = 1, f_v(y) = 1, \forall u \in U, v \in V, x, y$. Clearly, the PCP verifier accepts this proof with probability 1. This in the $\{0,1\}$ world corresponds to the case that when all the checks $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b$ had $b = 0$, then clearly there exists a solution that satisfies all the equations, the all $0's$ solution.

How do we get around this hurdle? Observe that a valid proof which comprises of long codes cannot be all $1's$, since a long code comprises an equal number of 1's and -1's. A long code is balanced since $\chi_{\{i\}}(-x) = -\chi_{\{i\}}(x)$. If we ensure that all the functions f_u and f_v satisfy $f_u(-x) = -f_u(x)$ and $f_v(-y) = -f_v(y)$, then the all 1's proof is avoided. But how can we do this without increasing the number of queries. We say that a proof is folded, if $f(-x) = -f(x), \forall x$. We can ensure the proof is folded by observing the following probing convention. For each of the functions f_u (similarly f_v) and for every pair of inputs $(x, -x)$ only one of $f_u(x)$ or $f_u(-x)$ is given in the proof. If the verifier needs to probe $f_u(x)$, if it is present in the proof it reads of the value, else it reads $f_u(-x)$ and sets $f_u(x) = -f_u(-x)$. By the above probing convention, we can assume that the proof is folded and the all 1's spurious proof is avoided. It can be checked that a folded proof satisfies the following lemma.

Claim 11.5.2. *If $f : \{1,-1\}^L \to \{1,-1\}$ is folded (i.e. $f(x) = -f(-x), \forall x$), then for even sized sets S, we have $\hat{f}_S = 0$ (in particular, $\hat{f}_\emptyset = 0$).*

Proof. We will only prove the case $S = \emptyset$, which is what we require in soundness analysis.

$$\hat{f}_\emptyset = \mathop{\mathbb{E}}_x[f(x)] = \mathop{\mathbb{E}}_x\left[\frac{f(x) + f(-x)}{2}\right] = 0.$$

The general case for $|S| > 0$ is similar. □

11.5.1.2 Soundness analysis continued

Let us resume the soundness analysis. First for some notation. Given any $S \subseteq L$, and $\pi : L \to R$, define the following sets.

$$\begin{aligned} \pi(S) &= \{j \in R | \text{ there exist an } l \text{ in } S \text{ such that } \pi(l) = r\} \\ \pi_2(S) &= \{j \in R | \text{ there exist an odd number of } l\text{'s in } S \text{ such that } \pi(l) = r\} \end{aligned}$$

Claim 11.5.3. *For all $x \in \{1,-1\}^R$, we have $\chi_S(x \circ \pi) = \chi_{\pi_2(S)}(x)$.*

Proof.

$$\chi_S(x \circ \pi) = \prod_{i \in S}(x \circ \pi)_i = \prod_{i \in S} x_{\pi(i)} = \prod_{j \in \pi(S)} x_j = \prod_{j \in \pi_2(S)} x_j = \chi_{\pi_2(S)}(x).$$

□

Coming back to the soundness analysis, the acceptance probability of the PCP verifier can be written as follows.

$$\Pr[\text{ Håstad's 3-bit PCP}_\varepsilon \text{ accepts }] = \mathop{\mathbb{E}}_{(u,v),x,y,\mu}\left[\frac{1 + f_u\left((x \circ \pi_{u,v})y\mu\right) f_v(x) f_u(y)}{2}\right].$$

Thus, if the PCP verifier accepts with probability at least $\frac{1}{2} + \delta$, we have that

$$\mathop{\mathbb{E}}_{(u,v),x,y,\mu} [f_u ((x \circ \pi_{u,v})y\mu)\, f_v(x) f_u(y)] \geq 2\delta.$$

Hence for atleast δ-fraction of edges (u, v), we have

$$\mathop{\mathbb{E}}_{x,y,\mu} [f((x \circ \pi)y\mu)g(x)f(y)] \geq \delta,$$

where we have followed the convention that $f_u = f, f_v = g$ and $\pi_{u,v} = \pi$. We can now rewrite this expression using Fourier expansion

$$\begin{aligned}
\mathop{\mathbb{E}}_{x,y,\mu}[f((x \circ \pi)y\mu)g(x)f(y)] &= \sum_{S,T,U} \hat{f_S}\hat{g_T}\hat{f_U} \cdot \mathbb{E}[\chi_S((x \circ \pi)y\mu) \cdot \chi_T(x) \cdot \chi_U(y)] \\
&= \sum_{S,T,U} \hat{f_S}\hat{g_T}\hat{f_U} \cdot \mathop{\mathbb{E}}_{x}[\chi_S(x \circ \pi) \cdot \chi_T(x)] \cdot \mathop{\mathbb{E}}_{y}[\chi_S(y)\chi_U(y)] \cdot \mathop{\mathbb{E}}_{\mu}[\chi_S(\mu)] \\
&= \sum_{S,T} \hat{f_S^2} \cdot \hat{g_T} \cdot \mathop{\mathbb{E}}_{x}[\chi_{\pi_2(S)}(x) \cdot \chi_T(x)] \cdot \mathop{\mathbb{E}}_{\mu}[\chi_S(\mu)] \\
&= \sum_{S} \hat{f_S^2} \cdot \hat{g}_{\pi_2(S)} \cdot (1 - 2\varepsilon)^{|S|}
\end{aligned}$$

We thus have $\sum_S \hat{f_u^2}(S) \cdot \hat{f_v}(\pi_2(S)) \cdot (1 - 2\varepsilon)^{|S|} > \delta$ for at least δ-fraction of edges.

11.5.1.3 Decoding a labeling from the f_u's and f_v's

Since the $f_u : \{1, -1\}^L \longrightarrow \{1, -1\}, \forall u \in U$ are Boolean functions, we have that $\sum_S \hat{f_u}^2(S) = 1$. This lets us define the following (randomized) labeling $A : U \to L$ as follows.

For each $u \in U$ do

1. Pick set S with probability $\hat{f_u^2}(S)$.
2. Set $A(u) \longleftarrow_R S$ (i.e., a random element from the set S)

Note that since $\hat{f_u}(\emptyset)$, we never pick the empty set is step 1. Similarly, we can define a labeling $B : V \to R$ as follows.

For each $v \in V$ do

1. Pick set T with probability $\hat{f_v^2}(T)$.
2. Set $B(v) \longleftarrow_R T$ (i.e., a random element from the set T)

What is the fraction of edges satisfied by this labeling?

$$\Pr_{(u,v)\in_R E} [(u, v) \text{ is satisfied }] = \Pr_{(u,v)\in_R E} [\pi_{u,v}(A(u)) = B(v)] \geq \sum_S \sum_{T \subseteq \pi(S)} \hat{f_u^2}(S) \cdot \hat{f_v^2}(T) \cdot \frac{1}{|S|}$$

This follows from the below argument: first pick an S with probability $\hat{f_u^2}(S)$, then pick T with probability $\hat{f_v^2}(T)$ where $T \subset \pi(S)$, now set the label of v to be a random element of

T. Since T is picked such that $T \subseteq \pi(S)$, clearly this random element has a pre-image in S. The probability that this pre-image is picked to be the label of u is at least $1/|S|$. We now relate this expression to $\sum_S \hat{f}_u^2(S) \cdot \hat{f}_v(\pi_2(S)) \cdot (1-2\varepsilon)^{|S|}$ which we know is at least δ for at least δ-fraction of the edges.

$$\begin{aligned}
\Pr_{(u,v)\in_R E}[(u,v) \text{ is satisfied }] &= \sum_{S,T\subseteq\pi(S)} \hat{f}_u^2(S)\hat{f}_v^2(T)\frac{1}{|S|} \\
&\geq \sum_{S} \hat{f}_u^2(S)\hat{f}_v^2(\pi_2(S))\frac{1}{|S|} \qquad \text{(dropping all the remaining positive terms)} \\
&= \sum_{S} \left(\hat{f}_u(S)\hat{f}_v(\pi_2(S))\frac{1}{\sqrt{|S|}}\right)^2 \cdot \sum_S \hat{f}_u^2(S) \qquad \text{(by Parsevals)} \\
&\geq \left(\sum_{S} \hat{f}_u^2(S)\hat{f}_v(\pi_2(S)) \cdot \frac{1}{\sqrt{|S|}}\right)^2 \qquad \text{(Cauchy-Schwarz)} \\
&\geq 4\varepsilon \left(\sum_{S} \hat{f}_u^2(S)\hat{f}_v^2(\pi_2(S))(1-2\varepsilon)^{|S|}\right)^2
\end{aligned}$$

where in the final step we have used the inequality $1/\sqrt{x} \geq \sqrt{4\varepsilon}(1-2\varepsilon)^x$[1] Now, since $\sum_S \hat{f}_u^2(S) \cdot \hat{f}_v(\pi_2(S)) \cdot (1-2\varepsilon)^{|S|} > \delta$ for at least δ-fraction of edges, we have

$$\Pr_{(u,v)\in_R E}[(u,v) \text{ is satisfied }] \geq \delta \cdot 4\varepsilon\delta^2 = 4\varepsilon\delta^3$$

More precisely, we have shown the following

Proposition 11.5.4 (Soundness)**.** *For every $\delta, \varepsilon > 0$, there exists a $\delta' = 4\varepsilon\delta^3$ such that the following holds. Let $I = (G, \Pi)$ be a* Label-Cover *instance such that there exists f_u, f_v's for which*

$$\Pr[\ \text{Håstad's 3-bit PCP}_\varepsilon \text{ accepts }] \geq \frac{1}{2} + \delta,$$

then there exist labelings $A : U \to L$ and $B : V \to R$ that satisfy at least δ'-fraction of the edges in Label-Cover *instance (in other words (G, Π) is not a* NO *instance of* $\text{gap}_{1,\delta'}\text{-LC}(L, R)$*).*

This completes the construction and analysis of Håstad's 3-bit PCP.

11.6 Other inapproximability results

Recall the broad outline for proving inapproximability of MAX3LIN2. In the first step, we designed a dictator test whose predicate is exactly the constraint of MAX3LIN2. In the second step, we composed this dictator test with the hardness result for Label-Cover to obtain a PCP with the appropriate acceptance predicate. This is a general recipe for

[1]Proof of inequality: $1/\sqrt{4\varepsilon x} \geq e^{-2\varepsilon x}$ (since $1/z \geq e^{-z}$ for $z > 0$ and hence $1/\sqrt{z} \geq e^{-z/2}$. Now, $e^{-2\varepsilon x} \geq (1-2\varepsilon)^x$ since $e^{-z} \geq 1 - z$ for $z > 0$).

proving other inapproximability results. We first design a tailor-made dictator test for the problem and then compose it with the labelcover hardness result.

For instance, for the problem MAX3SAT, we can design the following dictator test.

MAX3SAT-Dictator Test:

1. Pick $x, y \in_R \{1, -1\}^L$
2. Set string $z \in \{1, -1\}^L$ as follows:

If $x_i = 1, z_i = -y_i$

If $x_i = -1, z_i = y_i$ with probability $1 - \varepsilon$, and $-y_i$ with probability ε

3. Accept unless $f(x) = f(y) = f(z) = 1$.

Composing this with the labelcover hardness result we obtain the following result.

Theorem 11.6.1 (Hardness of MAX3SAT [Hås01])**.** *For all $\delta \in (0, 1)$, it is NP hard to distinguish if a MAX3SAT instance is perfectly satisfiable or at most $(\frac{7}{8} + \delta)$-satisfiable.*

Note that this theorem is an improvement over Corollary 11.2.3 as it achieves perfect completeness.

References

[BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. *Free bits, PCPs, and nonapproximability—towards tight results.* SIAM J. Computing, 27(3):804–915, June 1998. (Preliminary Version in *36th FOCS*, 1995). eccc:TR95-024, doi:10.1137/S0097539796302531.

[GLST98] Venkatesan Guruswami, Daniel Lewin, Madhu Sudan, and Luca Trevisan. *A tight characterization of NP with 3-query PCPs.* In *Proc. 39th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 18–27. 1998. doi:10.1109/SFCS.1998.743424.

[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).

[Hås01] Johan Håstad. *Some optimal inapproximability results.* J. ACM, 48(4):798–859, July 2001. (Preliminary Version in *29th STOC*, 1997). doi:10.1145/502090.502098.

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

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