CS 880: Advanced Complexity Theory	3/12/2008

Lecture 20: Inapproximability of Minimum Vertex Cover

Instructor: Dieter van Melkebeek | Scribe: Mark Liu

Last time we examined a generic approach for inapproximability results based on the Unique Games Conjecture. Before, we had already shown that approximating MAX-3-LIN to within a constant factor larger than $\frac{1}{2}$ is NP-hard. To do this we used a tweaked version of our dictatorship test that we came up with earlier in the semester. Last time we (re)proved that approximating MAX-3-LIN to within a constant larger than $\frac{1}{2}$ is UG-hard. The latter is a weaker statement than the earlier NP-hardness result, but the argument used the dictatorship test as a blackbox. In this lecture we show that when we replace the dictatorship test by a noise sensitivity test, then we obtain that MAX-CUT is UG-hard to approximate to within any factor larger than ρ_{GW}, where ρ_{GW} refers to the approximation ratio achieved by the Goemans-Williamson algorithm. The numerical value of ρ_{GW} is approximately .878.

We end the lecture wiht a proof sketch that approximating MIN-VC to within any constant factor smaller than 2 is UG-hard.

1 The Generic Approach of Last Lecture

We briefly go over the framework we saw last lecture. We start by restating the necessary properties of the family of dictatorship tests T_α, $\alpha \in A$.

- Dictatorship test T_α^g ($\alpha \in A$):
  - Completeness Condition: If g is a dictator then $\Pr[T_\alpha^g \text{ accepts}] \geq c_\alpha$.
  - Soundness Condition: For every α and $\delta > 0$ there exists $\tau > 0$ and $d > 0$ such that the following holds for every distribution $g = \mathrm{E}_u[f_u]$ over Boolean functions f_u in l variables: If $\Pr[T_\alpha^g \text{ accepts}] \geq s_\alpha + \delta$ then there exists $i \in [l]$ with $I_j^{\leq d}(g) \geq \tau$.

Given such a test T_α^g, we developed the following test S_α^f, where f represents the purported long encoding of the labels of the left-hand side of a constraint graph game $G = (L, R, E, [l], [r], C)$ of permutation type with underlying permutations π_e for $e \in E$.

- Test S_α^f:
  - Pick $v \in R$ at random.
  - Run $T_\alpha^{g_v}$, where on $g_v(y) = \mathrm{E}_{e=(u,v)\in \mathrm{E}}[f_u(y \circ \pi_e)]$ and accept iff $T_\alpha^{g_v}$ accepts.

Note that S_α uses T_α as a blackbox. We established the following key properties in the case where G is right-regular.

- Completeness: There exists an f such that $\Pr[S_\alpha^f \text{ accepts}] \geq c_\alpha \cdot \nu^\star(G)$.
- Soundness: If there exists an f such that $\Pr[S_\alpha^f \text{ accepts}] \geq s_\alpha + 2\delta$, then $\nu(G) \geq \frac{\delta\tau^2}{4d}$.

We can view $\Pr[S_\alpha^f \text{ accepts}]$ as the fraction of test conditions that are satisfied by f. The completeness result for S_α^f implies that if $\nu^\star(G) \geq 1-\gamma$ then MAX-SAT$(S_\alpha) \geq c_\alpha(1-\gamma)$, were MAX-SAT$(S_\alpha)$ is the maximum probability of acceptance of the test S_α over all possible functions f. By the soundness result, we have that if $\nu(G) \leq \gamma < \frac{\delta\tau^2}{4d}$ then MAX-SAT$(S_\alpha) < s_\alpha + 2\delta$. From this we can conclude that it is UG-hard to approximate any problem that contains all problems of the form MAX-SAT(S_α) for $\alpha \in A$, to within a constant factor $\rho > \inf_{\alpha\in A}(\frac{s_\alpha}{c_\alpha})$.

2 Applications

We first recall teh application to MAX-3-LIN from last class and then develop the new one for MAX-CUT.

2.1 Application to MAX-3-LIN

When we use the 3-query dictatorship test T_α from the beginning of the course for $\alpha \in (-1,1)$, then MAX-SAT(S_α) becomes an instance of MAX-3-LIN. We argued that $c_\alpha = 1-\epsilon$ for $\epsilon = \frac{1-\alpha}{2}$, and $s_\alpha = \frac{1}{2}$. We conclude that approximating MAX-3-LIN to within any constant factor $rho >$ $\inf_{\alpha\in(-1,1)}(\frac{s_\alpha}{c_\alpha}) = frac12$ is UG-hard.

2.2 Application to MAX-CUT

To achieve our MAX-CUT inapproximability result, we will replace our 3-query dictatorship test with a 2-query noise sensitivity test defined as follows.

- Test T_α^g:
  - Pick $x \in \{-1,1\}^l$ at random
  - Pick $y \sim_\alpha x$
  - Accept iff $g(x) \neq g(y)$

Note that the conditions induced by the resulting test S_α^f are all of the form $f_u(x) \neq f_v(y)$ for $u, v \in L$ and $x, y \in \{-1,1\}^l$. Thus, when we consider a graph H with a vertex for each combo $f_u(x)$ and an edge $(f_u(x), f_v(y))$ for each condition $f_u(x) \neq f_v(y)$, then MAX-SAT(S_α) is equivalent to finding a two-coloring of the vertices of H that maximizes the number of mixed edges. In other words, MAX-SAT(S_α) is equivalent to finding a maximum cut in H.

The above test is an effectivization of the notion of noise sensitivity. As a result, we have that $\Pr[T_\alpha^g \text{ accepts}] = NS_\alpha(g)$. This allows us to establish the parameters c_α and s_α for which this test meets the conditions listed at the end of Section 1. As for the completeness condition, we can set $c_\alpha = \alpha$ since $NS_\epsilon(\text{dictator}) = \epsilon$. As for the soundness condition, recall the Reverse Majority Stablest Theorem: For any $g : \{-1,1\}^l \to [-1,1]$ and all α such that $\frac{1}{2} < \alpha < 1$, if for all $i \in [l], I_i^{\leq \log(1/\tau)}(g) \leq \tau$, then $NS_\alpha(g) \leq \frac{1}{\pi}\arccos(1-2\alpha) + O(\frac{1}{1-\alpha}\frac{\log\log\frac{1}{\tau}}{\log\frac{1}{\tau}})$. The contrapositive essentially gives us the soundness condition we need to hold. Note that for any $\alpha \in (\frac{1}{2}, 1)$ we can make the error term $\delta = O(\frac{1}{1-\alpha}\frac{\log\log\frac{1}{\tau}}{\log\frac{1}{\tau}})$ arbitrarily small by choosing τ sufficiently small. Setting

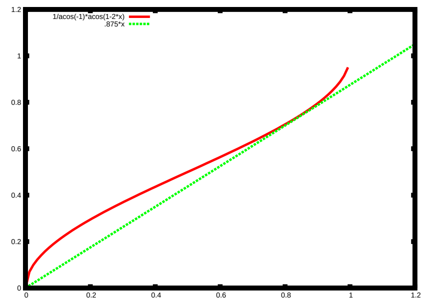

Figure 1:

$d = \log \frac{1}{\tau}$ then gives the soundness property with $s_\alpha = \frac{1}{\pi} \arccos(1 - 2\alpha)$. We conclude that it is UG-hard to approximate MAX-CUT to within any constant

$$\rho > \inf_{\frac{1}{2} < \alpha < 1} \left(\frac{\frac{1}{\pi} \arccos(1 - 2\alpha)}{\alpha} \right). \tag{1}$$

To see the infimum visually, we plot $h(x) = \frac{1}{\pi} \arccos(1 - 2\alpha)$ against α in Figure 1. The infimum is reached where a line drawn through the origin is tangent to the curve.

We claim that the right-hand side of (1) is exactly the approximation factor ρ_{GW} guaranteed by the Goemans-Williamson algorithm. To see why, we quickly review the algorithm and its analysis.

For a given graph $H = (V, E)$, we have

$$\text{MAX-CUT}(H) \leq \max_{||x_v||=1, v \in V} \sum_{e=(u,v) \in E} \frac{1 - \langle x_u, x_v \rangle}{2}, \tag{2}$$

where the x_v's range over all vectors in $\mathbb{R}^d$ for unrestricted d. The problem on the right-hand side of (2) is a semidefinite program, which can be solved in polynomial time. It is a relaxation of the MAX-CUT problem, which is equivalent to the problem on the right-hand side with teh additional restriction that $d = 1$.

Once we have an optimal solution x_v for $v \in V$ to the semidefinite program, we need to extract a cut in H whose value is not too much worse than the right-hand side of (2). We do so by randomized rounding: We pick a random hyperplane through the origin of $\mathbb{R}^d$ and use it to partition V in those vertices v for which x_v is on one side of the hyperplane, and the others. The resulting cut in H has the property that for any fixed edge $e = (u, v) \in E$,

$$\Pr[e = (u, v) \text{ is cut}] = \frac{\theta}{\pi} = \frac{\arccos(\langle x_u, x_v \rangle)}{\pi}. \tag{3}$$

This follows from an analysis in the plane spanned by the vectors x_u and x_v (see Figure 2). The intersection of the random hyperplane with this plane is a random line through the origin. The probability that this line separates x_u and x_v equals $\frac{\theta}{\pi}$, where θ is the angle between the two vectors, i.e., $\cos(\theta) = \langle x_u, x_v \rangle$.

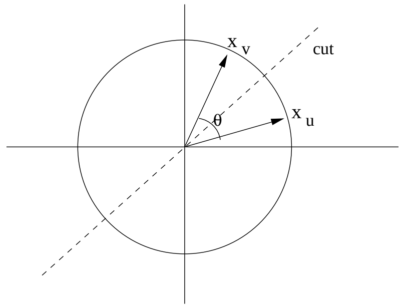

Figure 2: Probability of a Random Cut

By linearity of expectation and comparing the right-hand side of (3) with the corresponding term on the right-hand side of (2), we have

$$\begin{aligned} \mathrm{E}[\#\text{ edges cut}] &= \sum_{e=(u,v)\in E} \frac{\arccos(\langle x_u, x_v\rangle)}{\pi} \\ &\geq \min_{e=(u,v)\in E} \left(\frac{\frac{\arccos(\langle x_u,x_v\rangle)}{\pi}}{\frac{1-\langle x_u,x_v\rangle}{2}} \right) \cdot \sum_{e=(u,v)\in E} \frac{1-\langle x_u, x_v\rangle}{2} \\ &\geq \min_{-1\leq z\leq 1} \left(\frac{\frac{\arccos(z)}{\pi}}{\frac{1-z}{2}} \right) \cdot \sum_{e=(u,v)\in E} \frac{1-\langle x_u, x_v\rangle}{2}, \end{aligned}$$

where the last line follows by substituting $z = \langle x_u, x_v\rangle$ and further relaxation. Since x_v is the optimal solution for the right-hand side of (2), we conclude that

$$\mathrm{E}[\#\text{ edges cut}] \geq \rho_{GW} \cdot \text{MAX-CUT}(H),$$

where

$$\rho_{GW} = \min_{-1\leq z\leq 1} \left(\frac{\frac{\arccos(z)}{\pi}}{\frac{1-z}{2}} \right). \tag{4}$$

By substituting $z = 1 - 2\alpha$ and observing that the minimum is reached for $z \geq 0$, the right-hand side of (4) transforms into the right-hand side of (1), which is what we wanted to argue.

Note that the Goemans-Williamson algorithm can be derandomized to yield a deterministic approximation ρ_{GW}-approximation algorithm for MAX-CUT.

3 Minimum Vertex Cover

We now switch to another standard optimization problem – minimum vertex cover or MIN-VC for short. The following theorem shows that approximating MIN-VC to within any constant factor less than 2 is UG-hard.

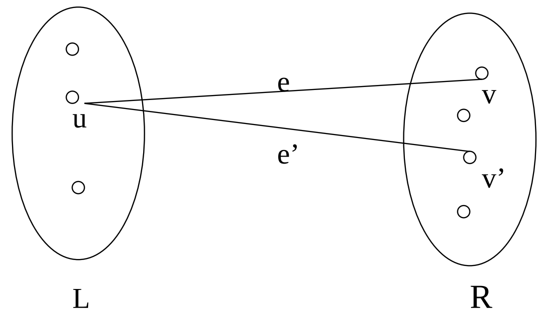

Figure 3: Unique Constraint Graph Games

Theorem 1. *For all $\delta > 0$, there exists a $\gamma > 0$ such that there is a reduction from unique constraint graph games over $[l]$ to graphs H such that the following hold:*

$$\begin{aligned} \nu^\star(G) \geq 1 - \gamma \quad &\Rightarrow \quad \textit{MIN-VC}(H) \leq \frac{1}{2} + \delta \\ \nu(G) \leq \gamma \quad &\Rightarrow \quad \textit{MIN-VC}(H) \geq 1 - \delta. \end{aligned}$$

We will not prove Theorem 1 in full. We only describe the underlying construction and argue the completeness property, but not the soundness property.

Given a unique constraint graph game $G = (L, R, E, [l], [r], C)$ we construct H as follows. The vertices of H represent the bits of the purported long codes g_v of the labels of vertices $v \in R$. More precisely, for each $v \in R$ and $y \in \{-1, 1\}^r$, we include a vertex representing $g_v(y)$ in H. In order to define the edge relationship $\sim$, we view strings in $\{-1, 1\}^r$ as the characteristic vector of subsets of $[r]$. For $v, v' \in R$ with $v \neq v'$ and $y, y' \in \{-1, 1\}^r$, we stipulate the following.

- $(v, y) \sim (v, y') \Leftrightarrow y \cap y' = \emptyset$.
- $(v, y) \sim (v', y') \Leftrightarrow$ there exists a $u \in L$ such that $e = (u, v) \in E$, $e' = (u, v') \in E$ and $y \circ \pi_e \cap y' \circ \pi_{e'} = \emptyset$.

Note that the first condition can be viewed as a special case of the second one for $v = v'$. For $v \neq v'$, the second condition involves some kind of consistency check. See Figure 4 for a visual aid.

To argue the completeness property in Theorem 1, we will think in terms of independent sets rather than vertex covers. Recall that a set B is a vertex cover iff $A = \overline{B}$ is an independent set. Thus, to establish the completeness property we need to exhibit an independent set of relative size at least $\frac{1}{2} - \delta$ whenever $\nu^*(G) \geq 1 - \gamma$.

Fix a labeling and a set $R^* \subseteq R$ realizing $\nu^*(G)$. The set R^* denotes the vertices in R all of whose incident edge constraints are met by the underlying labeling. Also, let $A = \{(v, y) \,|\, v \in R^* \text{ and } g_v(y) = -1\}$, where g_v denotes the long code of the label given to v, i.e., $g_v(y) = y_j$ where $j = \text{value}(v)$.

Claim 1. *A is an independent set of relative size at least $\frac{1-\gamma}{2}$.*

The claim about the relative size follows because the relative size of R^* is at least $1 - \gamma$ and for each $v \in R$ exactly half of the $y \in \{-1, 1\}^r$ map to -1 under the dictator g_v.

For the claim of independence, consider any $v \in R^*$ with label j. We have that $g_v(y) = -1$ iff $y_j = -1$ iff $j \in y$. Thus, if $(v, y) \in A$ and $(v, y') \in A$, then $j \in y \cap y'$, which means that there is no edge between (v, y) and (v, y'). In addition, consider any $v' \in R^*$ with value j' and such that $v' \neq v$. Let u be any common neighbor of v and v'. Since the edge constraints on $e = (u, v)$ and $e' = (u, v')$ are satisfied by the underlying labeling, we have that $\pi_e^{-1}(j) = \pi_{e'}^{-1}(j')$. Now, if $(v, y) \in A$ then $j \in y$ so $\pi_e^{-1}(j) \in y \circ \pi_e$. So, if both $(v, y) \in A$ and $(v', y') \in A$ then $\pi_e^{-1}(j) = \pi_{e'}^{-1}(j') \in (y \circ \pi_e) \cap (y' \cap \pi_{e'})$. Since this holds for all common neighbors u of v and v', there is no edge between (v, y) and (v', y').

This finishes the proof of the completeness property in Theorem 1. The proof of the soundness property is significantly more involved and we will not cover it due to lack of time.