

More hardness of approximation

Last time: New notion of "proof" (probabilistic proof systems)

⇒ PCP theorem

⇒ Hardness of approximation for CSPs in general,

Max-3SAT specifically:

- $\frac{15}{16}$ using first PCP thm

- $\frac{7}{8} + \epsilon$ using Hastad's 3-bit PCP thm

Today: Another notion of "proof", hardness of approximation.

Thm (last time): For Max-3SAT, it is NP-hard to distinguish instances in which all clauses satisfiable from instances in which at most $\frac{15}{16}$ of clauses are satisfiable

- YES instance: all clauses satisfiable

- NO instance: $\leq \frac{15}{16}$ fraction of clauses satisfiable

Note: $\frac{15}{16}$ worse than $\frac{7}{8}$, but completeness 1 is nice property!

Max-3SAT: Every clause has 3 literals

- Max-3SAT-S: - Every clause has 3 literals
- Every variable in 5 clauses

Standard transformation from Max-3SAT to Max-3SAT-S
- loses a constant in soundness

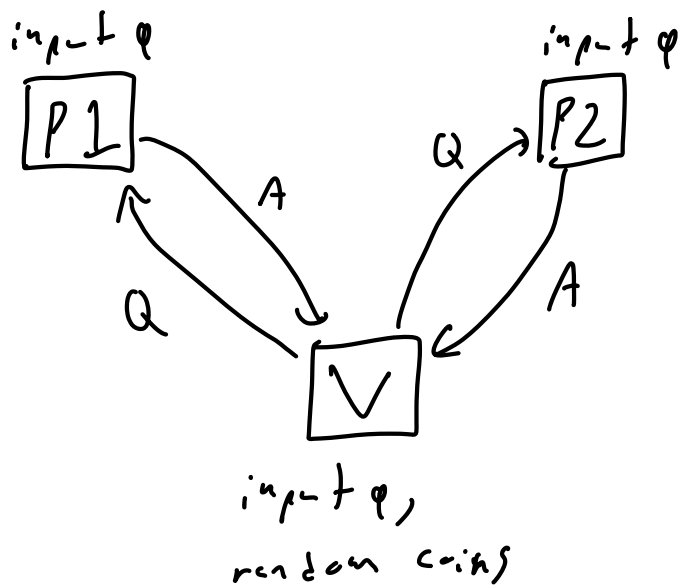
Thm: For Max-3SAT-S, it is NP-hard to distinguish between:
- instances in which all clauses satisfiable (YES instances)
- instances in which $\leq 1-\epsilon$ fraction of clauses are satisfiable (NO instances)
for some constant $\epsilon > 0$

One-round, Two-Prover proof system for language L

- Two provers, one verifier. All know input φ
- Verifier asks each prover a question (possibly different)
 - Provers answer. **Computationally Unbounded, deterministic**
- Based on responses, verifier decides whether to accept (YES) or reject (NO). **Must run in polytime**
- Provers can decide on a strategy beforehand, but can't

communicate with each other after receiving question

- Provers trying to get verifier to accept
- Verifier trying to check if $\varphi \in L$



Trivial 2-prover proof system for \exists SAT-S

- Verifier asks each prover for assignment
- Checks whether each assignment satisfies all clauses, or $\leq 1-\epsilon$ fraction of clauses


\Rightarrow if φ YES instance, provers can get verifier to accept with prob. 1 (completeness 1)

if φ NO instance, no matter what provers do, verifier accepts with prob. 0 (soundness 0).

What if we want questions, answers to be "short"?
Use randomness!

Verifier on instance φ :

- Choose clause C uniformly at random
- Choose one of the three variables in C uniformly at random. Call it x_i
- Ask prover 1 for an assignment to x_i : (T/F)
- Ask prover 2 for a **satisfying** assignment to C (7 possibilities)
- P2's answer includes assignment to x_i . Return YES if it matches P1's assignment to x_i .
Otherwise return NO.



Pick $C = x_i \vee \bar{x}_7 \vee x_{12}$
uniformly at random
Pick x_{i2} u.a.r. from $\{x_1, x_7, x_{12}\}$

Lemma: If φ a YES instance (there is an assignment satisfying all clauses), then provers can get verifier to return YES with probability 1.

Pf:

Return appropriate part of satisfying assignment!

Lemma: If φ a NO instance (every assignment satisfies at most $1-\epsilon$ fraction of clauses), then no matter what provers do,

$$P[\text{Verifier returns YES}] \leq 1 - \frac{\epsilon}{3}$$

Pf: Note: provers are deterministic

P1: Has some assignment, returns $x_i = T/F$ depending on assignment

\Rightarrow satisfies $\leq 1-\epsilon$ of clauses

\Rightarrow with prob. $\geq \epsilon$, we choose C not satisfied

P2: returns **satisfying** assignment for C

\Rightarrow disagrees with P1 on at least one of the three vars

⇒ We choose which var to ask 11 u.a.v from the 3

⇒ find disagreement with prob. $\geq \frac{\epsilon}{3}$

New computational problem: Find best strategy for prover.

Label cover:

Input: - Bipartite graph $G = (L, R, E)$

- Alphabet Σ_L


- Alphabet Σ_R

- Relation $\Pi_e \subseteq \Sigma_L \times \Sigma_R$ for each $e \in E$

Feasible: Assignment $f: L \rightarrow \Sigma_L$ and $g: R \rightarrow \Sigma_R$

Objective: max fraction of edges (u, v) s.t. $(f(u), g(v)) \in \Pi_{u,v}$

Ex: $\Sigma_L = \Sigma_R = \{ \bullet, \circ, \cdot \}$



$$\Pi_{u,v} = \{ (\bullet, \circ), (\bullet, \cdot) \}$$


$$\Pi_{w,v} = \{ (\cdot, \bullet), (\cdot, \circ), (\cdot, \cdot), (\circ, \cdot) \}$$

Informal claim: this is the problem of finding the best strategy for provers!

On input φ with n variables and $m = \frac{5}{3}n$ clauses;

L = variables (vertex for each variable)

R = clauses (vertex for each clause)



E : add edge b/w every vertex and clause it appears in

\Rightarrow left nodes have degree 5

right nodes have degree 3

$\Sigma_L = \{T, F\}$

$\Sigma_R = \{7 \text{ satisfying assignments}\}$

$\Pi_{(x,y)} = 7$ pairs out of 14 that are consistent

Regularity \Rightarrow choosing random C , random $x_i \in C$ same as choosing random edge

\Rightarrow LC solution f is a strategy for provers where

$\Pr[\text{verifier accepts}] =$ fraction of edges whose relation is satisfied by f

= LC objective

Thm: There is some constant $\epsilon > 0$ s.t. it is NP-hard to distinguish between instances of Label Cover where

- All edges can be satisfied
- At most $1-\epsilon$ fraction of edges can be satisfied

\Rightarrow NP-hard to approximate LC better than $1-\epsilon$

Turns out LC **much** harder.

Back to 2-prover proof system: how to boost soundness from $1-\epsilon/3$ to something smaller?

(How to boost probability of catching provers in inconsistency?)

Obvious approach: **repetition**

Repeat K times \Rightarrow

$$P[\text{never detect inconsistency}] \leq (1-\epsilon/3)^K$$

Works great! But to maintain connection to LC, need to maintain 1 round

Idea: repeat in parallel

Verifier:

- Choose k random clauses C_1, C_2, \dots, C_k
- From each clause C_i , choose random variable x_i from C_i
- Ask prover 1 for assignment for every x_i
- Ask prover 2 for satisfying assignment for every C_i
- Return YES if consistent on all k ,
NO otherwise

Gives LC instance:

$$L = [n]^k \quad R = [m]^k \quad \Sigma_L = [2]^k \quad \Sigma_R = [7]^k$$

$$\Pi_{(x_1, \dots, x_k), (C_1, \dots, C_k)} = \left\{ ((\alpha_1, \dots, \alpha_k), (\beta_1, \dots, \beta_k)) \in [2]^k \times [7]^k : \right. \\ \left. (\alpha_i, \beta_i) \in \Pi_{(x_i, C_i)} \text{ for all } i \in [k] \right\}$$

Q: Is asking questions in parallel same as repetitively?

Intuition: yes. How can provers cheat by knowing questions in parallel?

Truth: No! Provers can convince verifier with prob. $> (1 - \epsilon/2)^k$

But parallel almost as good:

Raz's Parallel Repetition Lemma:

If every assignment satisfies $\leq 1-\epsilon$ fraction of clauses, then there is some constant $c > 0$ s.t. $\forall k$, no matter what provers do in k -parallel repetition,

$$P[\text{Verifier returns YES}] \leq (1-\epsilon)^{ck} \quad \text{--- } (1-\epsilon)^c)^k$$

Implication to Label Cover:

Thm: There is some $\epsilon > 0$ and $c > 0$ s.t. $\forall k \geq 1$,

unless $NP \subseteq \text{DTIME}(n^{O(k)})$, there is no polytime algorithm which can distinguish between instances of Label Cover where:

- all edges can be satisfied
- $\leq (1-\epsilon)^{ck}$ fraction of edges can be satisfied

Note: Instead of assuming $P \neq NP$, assuming $NP \not\subseteq \text{DTIME}(n^{O(k)})$

b/c size of LC instance $\approx n^k$

For any constant k , $\text{DTIME}(n^{O(k)}) \subseteq P$

Corollary: For any constant $0 < \alpha \leq 1$, unless $P = NP$ there is no polynomial time α -approximation algorithm for Label Cover

$$\text{SpS set } k = \Theta\left(\log^{\frac{1-\epsilon}{\epsilon}} n\right)$$

$$\Rightarrow \text{LC graph has size } \approx N = n^k = n^{\Theta\left(\log^{\frac{1-\epsilon}{\epsilon}} n\right)}$$

$$\Rightarrow \log N \approx \Theta\left(\log^{\frac{1-\epsilon}{\epsilon}} n \cdot \log n\right) = \Theta\left(\log^{\frac{1}{\epsilon}} n\right)$$

$$\Rightarrow \log n = \Theta\left(\log^{\epsilon} N\right)$$

$$\Rightarrow \text{inapproximability} \approx (1-\epsilon)^{ck}$$

$$= (1-\epsilon)^{c \cdot \log^{\frac{1-\epsilon}{\epsilon}} n}$$

$$= 2^{-c' \log^{\frac{1-\epsilon}{\epsilon}} n}$$

$$= 2^{-c' \frac{\log^{\frac{1}{\epsilon}} n}{\log n}}$$

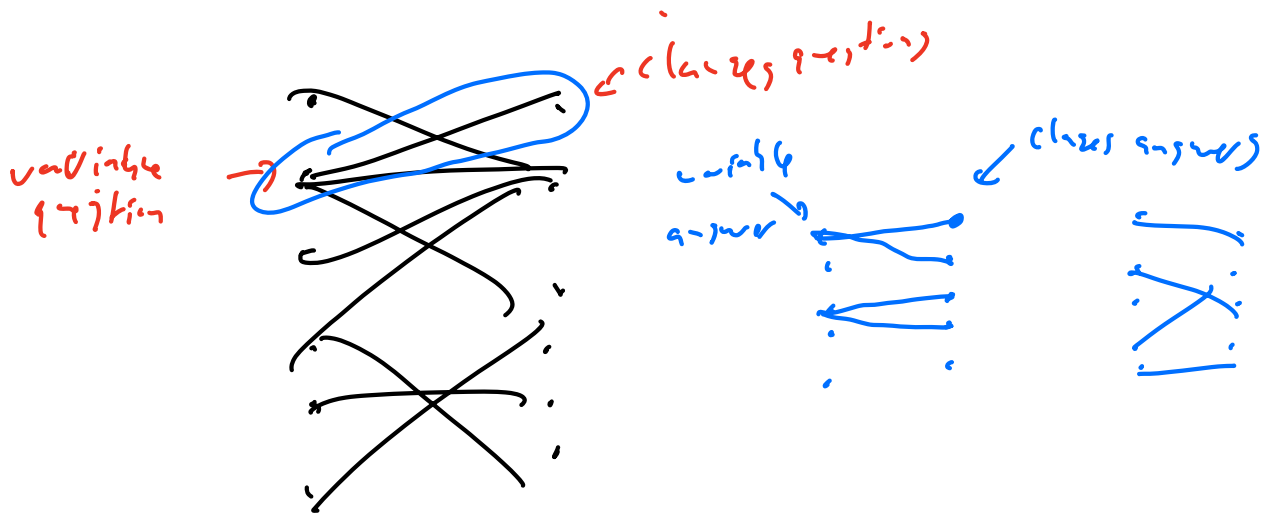
$$= 2^{-c' \cdot \frac{\log N}{\log^{\epsilon} N}}$$

$$\leq 2^{-\log^{1-\epsilon} N}$$

Quasipolytime: time $O(n^{\text{poly}(\log n)})$

Thm: For any $\epsilon > 0$, unless NP has quasipolytime algorithms, there is no polytime algorithm for Label Cover with approximation better than $2^{-\log^{1-\epsilon} n}$

Unique Games:



u.g.c.: If \exists NP-hard to distinguish $1-\epsilon$ vs. ϵ in Unique Games constant ϵ