Approximation Algorithms for Label Cover and The Log-Density Threshold

Paper by Eden Chlamtáč, Pasin Manurangsi, Dana Moshkovitz, Aravindan Vijayaraghavan
Presenter: Jimmy Wu

Label Cover/Projection Game

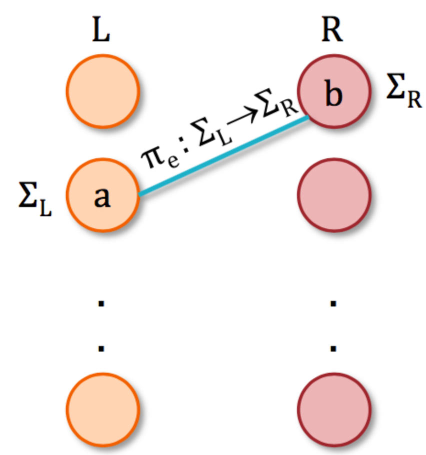

Input:
- Bipartite graph $G = (L, R, E)$
- Alphabets Σ_L, Σ_R
- $\forall e \in E$, **projection/constraint function** $\pi_e : \Sigma_L \to \Sigma_R$

Goal:
- Give a **labeling/assignment** $\phi_L : L \to \Sigma_L$ and $\phi_R : R \to \Sigma_R$
- Edge $e = (a, b)$ **satisfied/consistent** if $\pi_e(\phi_L(a)) = \phi_r(b)$
- Maximize the fraction of satisfied edges.

δ-Gap Label Cover

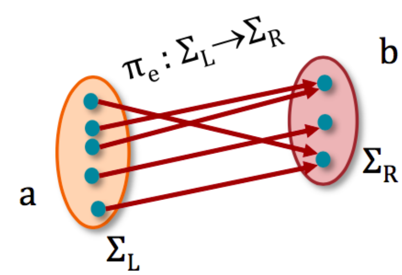

Distinguish between:
- (YES) There is a labeling that satisfies every edge
- (NO) No labeling satisfies more than δ fraction of edges

Fundamental problem in hardness of approximation.

Notation

$n := |L| + |R|$: number of vertices in G
$k := |\Sigma_L| \geq |\Sigma_R|$: size of left alphabet
$N := nk$: "size of the instance"
δ: gap

Previous Bounds for Label Cover

Algorithms
- [Charikar, Hajiaghayi, Karloff] $O(N^{1/3})$-approximation.
- [Manurangsi, Moshkovitz] $O(N^{1/4})$-approximation for fully satisfiable instances.

Hardness
- [Dinur, Steurer '13]: NP-hard: $\delta = 1/\log^c N$ for every $c > 0$.
- Assuming NP not in quasipoly-time: hard for $\delta = 2^{-\Omega(\sqrt{\log N})}$.
- Projection Games Conjecture: hard for $\delta = 1/N^c$ for some $c > 0$.
  What is the correct c?
  This paper suggests: $c = 3 - 2\sqrt{2} \approx 0.17$ given by the **"log density threshold"**.

Main Results

- $N^{3-2\sqrt{2}+\varepsilon} \approx N^{0.17}$-approximation algorithm for **semi-random** Label Cover in time $N^{O(1/\varepsilon^2)}$
- $N^{0.23}$-approximation algorithm for **worst-case** Label Cover
- $N^{1/8-\varepsilon}$ integrality gap for N^{ε}-level Lasserre/SoS relaxation

The Log-Density Method

1. Study "random vs. planted"
2. Identify and count "witnesses"
3. Identify threshold at which witness algorithms start to work (log density)
4. Use insights to devise algorithm for worst-case

Random Label Cover

Graphs:
- Erdős-Rényi: $G(n/2, n/2, p = \Delta/n)$
- $n/2$ vertices on each side, left- and right-regular
- Size of right label set is k/d

Distinguish:
- Random/NO instance: each π_e is a random d-to-1 function
- Planted/YES instance:
  - Plant a total labeling ϕ
  - Each $\pi_{(a,b)}$ is random d-to-1 function s.t. $\pi_{(a,b)}(\phi(a)) = \phi(b)$

Distinguishing via Witnesses

Witness: constant-size subgraph W such that:
- W appears in G w.h.p.
- In NO case, there is no satisfying assignment for W w.h.p.

Witness exists when:
log-density of constraint graph > log-density of the projections

$$2\log\Delta/\log n > \log d/\log k$$

Algorithm for (distinguishing) Random Label Cover:
- Fix a small set of vertices $U = (u_1, \ldots, u_r)$ and labels $(\sigma_1, \ldots, \sigma_r)$ for them.
- For each W containing U, try to assign labels consistently
- Repeat for all small sets U and possible labelings

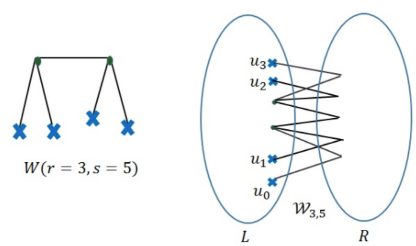

Semi-Random Label Cover

- Constraint graph G is (still) random
- Projections π_e are arbitrary functions satisfied by planted labeling

Algorithm:
Case 1: $2\log\Delta/\log n \leq \log d/\log k$. Take best of:

- d/k-approx: random assignment
- $1/\Delta$-approx: satisfy edges of a perfect matching

Case 2: $2 \log \Delta / \log n > \log d / \log k$

- Reduce/sparsify the alphabet of each $v \in V$.

Result: $N^{3-2\sqrt{2}} \approx N^{0.17}$-approximation.

Worst-Case Label Cover

Issue: no expansion properties needed for alphabet reduction
Workaround: Partition into dense subgraphs and solve separately
Result: $N^{0.23}$-approximation (lose objective from recombining subproblems)

Integrality Gap for Label Cover Lasserre SDP

Lasserre SDP for Projection Games

$$\text{maximize} \sum_{(u,v)\in E} \sum_{\sigma\in[k]} ||U_{(\{u,v\},\{u\to\sigma, v\to\pi_{(u,v)}(\sigma)\})}||^2$$

subjects to

$$||U_{(\emptyset,\emptyset)}|| = 1$$

$$\langle U_{(S_1,\alpha_1)}, U_{(S_1,\alpha_1)} \rangle = 0 \qquad \forall S_1, S_2, \alpha_1, \alpha_2 \text{ s.t. } \alpha_1(S_1) \neq \alpha_2(S_2)$$

$$\langle U_{(S_1,\alpha_1)}, U_{(S_2,\alpha_2)} \rangle = \langle U_{(S_3,\alpha_3)}, U_{(S_4,\alpha_4)} \rangle \qquad \forall S_1, S_2, S_3, S_4, \alpha_1, \alpha_2, \alpha_3, \alpha_4 \text{ s.t. } \alpha_1 \circ \alpha_2 = \alpha_3 \circ \alpha_4$$

$$\langle U_{(S_1,\alpha_1)}, U_{(S_1,\alpha_1)} \rangle \geq 0 \qquad \forall S_1, S_2, \alpha_1, \alpha_2$$

$$\sum_{\sigma\in[k]} ||U_{(\{v\},\sigma)}||^2 = 1 \qquad \forall v \in L \cup R$$

Integrality gap via reduction: from random Max k-CSP (gap given by [Tulsiani '09], [BCVGZ '12])
Results:

- $N^{1/8-\varepsilon}$ integrality gap for $N^{\Omega(\varepsilon)}$-level Lasserre/SoS relaxation
- $N^{\Omega(\varepsilon)}$ integrality gap for $N^{1-\varepsilon}$-level Lasserre/SoS relaxation
  Nearly matches a trivial algorithm.

Remarks:

- SoS cannot refute Projection Games Conjecture (evidence that it's true)
- Gap instances are semi-random, so $N^{0.17}$-approximation algorithms apply