Peter Kolman 10/1/2022

LABEL COVER

wlog $\Sigma = \{1, 2, \dots, |\Sigma|\}$

instance $(V, W, E, \{\pi_e\}, \Sigma)$ where

- (V, W, E) is bipartite graph and all vertices in V have the same degree
- for each edge $e = (u, v) \in E$ there is a function $\pi_e : \Sigma \to \Sigma$ (constraint)

A labeling $L : V \cup W \to \Sigma$ satisfies the constraint π_e, $e = (u, v)$, if $\pi_e(L(u)) = L(v)$

UNIQUE LABEL COVER

- all constraints π_e are permutations

Promise instance - either $\geq (1-\delta)|E|$ or $\leq \delta |E|$ satisfiable

GAP-ULC(δ)

for a "promisse instance" of ULC decide

- whether there is a labeling satisfying at least $(1-\delta)|E|$ constraints, or
- whether any labeling satisfies at most $\delta |E|$ constraints

UNIQUE GAMES CONJECTURE, Khot 2002

- For any $\delta > 0$, there is M s.t. GAP-ULC(δ) is NP-hard for $|\Sigma| \geq M$

REDUCTION GAP-ULC $\to$ MAX CUT

For $(V, W, E, \{\pi_e\}, \Sigma)$, regular on V, and $\underbrace{\rho}_{\text{parameter}} \in (-1, 0)$

we construct $H = (U, F)$ as follows

Vertices $U = W \times 2^{\Sigma}$ cloud

(think about 2^{Σ} as about the bits

of LONG-CODE encoding of a label

$\in \Sigma$ of a vertex from W;

LONG CODE of $\ell \in \Sigma$ is a

function $f_\ell : \{-1, 1\}^{\Sigma} \to \{-1, 1\}$

s.t. $f_\ell(x) = x_\ell$ (i.e., the ℓ-th coordinate of x))

Edges (with weights/multi) are given by the following

distribution D on $U \times U$ defined as follows:

pick uniformly at random $v \in V$,

and unif. at random $w, w' \in N(v)$

(neighbours of v)

pick u.a.r. $x \in \{-1, 1\}^{\Sigma}$

pick "noise" $\mu \in \{-1, 1\}^{\Sigma}$ s.t. $\mu_i = \begin{cases} -1 & \frac{1-\rho}{2} \\ 1 & \frac{1+\rho}{2} \end{cases}$ PROB.

2

set $y := x\mu$ (multiply coordinate-wise)

Output the pair (intuition: think about $\pi = id$)

$$\{(w, x\circ\pi_w), (w', y\circ\pi_{w'})\}$$

where $x\circ\pi$ is defined by

$$(x\circ\pi)_i = x_{\pi(i)}$$

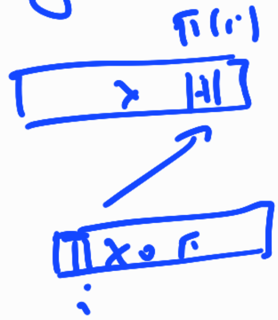

weights (or: multiplicity) of the edges in H are given by probabilities in the distribution

$G=(V,W,E)$ $H=(V,F)$

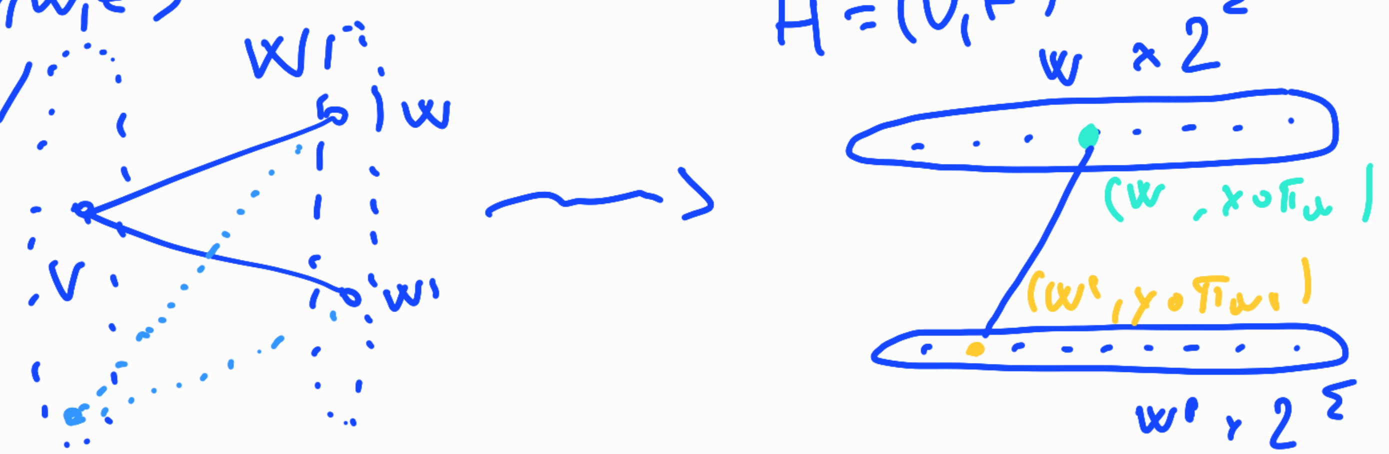

THM A: $\forall \rho \in (-1,0), \forall \varepsilon > 0\ \exists \delta > 0$ s.t.

for a GAP_ULC(Σ) instance

$I = (V, W, E, \{\pi_e\}, \Sigma)$, the above reduction satisfies:

a) $OPT(I) \geq (1-\delta)|E| \Rightarrow MAXCUT(H) \geq \frac{1-\rho}{2} - \varepsilon$

b) $OPT(I) \leq \delta|E| \Rightarrow MAXCUT(H) \leq \frac{\arccos\rho}{\pi} + \varepsilon$

Note: by the UGC, it is NP-hard to distinguish between a) and b), for Σ large enough $\Rightarrow$

Corollary: 2005

Assuming UGC, it is NP-hard to approximate MAXCUT to within

$$\min_{\rho} \frac{\frac{\arccos \rho}{\pi} + \varepsilon}{\frac{1-\rho}{2} - \varepsilon} > \min_{\rho} \frac{\frac{\arccos \rho}{\pi}}{\frac{1-\rho}{2}} = \alpha_{GW} \approx 0{,}878$$

the constant from the SDP-based alg. by Goemans, Williamson.

Proof of Thm A, part a)

a) Given $I = (V, W, E, \{\pi_e\}, \Sigma)$ we assume that $\ell : V \cup W \to \Sigma$ is a labeling satisfying $\geq (1-\delta)|E|$ constraints.

For $w \in W$, let $f_w : \{-1,1\}^{\Sigma} \to \{-1,1\}$ be the long code of $\ell(w)$, i.e.: $f_w(x) = x_{\ell(w)}$

"dictator function"

Let $(S, U \setminus S)$ be the cut of H defined by $S = \{(w,z) \in W \times \{-1,1\}^{\Sigma} \mid f_w(z) = 1\}$

(Note: $U \setminus S = \{ \quad \cdots\cdots \quad = -1\}$)

Consider an edge $e = \{(w, x \circ \pi_w), (w', y \circ \pi_{w'})\}$ selected according to the distribution D – and defined by $v, w, w', x, y = x\mu$?

$$\Pr[e \in (S, V \setminus S)] = \Pr[f_w(x \circ \pi_w) \neq f_{w'}(y \circ \pi_{w'})] = ?$$

$$f_w(x \circ \pi_w) \underset{\substack{\uparrow \\ \text{def. of } f_w}}{=} (x \circ \pi_w)_{\ell(w)} \underset{\substack{\uparrow \\ \text{def. of } x \circ \pi}}{=} x_{\pi_w(\ell(w))}$$

$$f_{w'}(y \circ \pi_{w'}) = (y \circ \pi_{w'})_{\ell(w')} = y_{\pi_{w'}(\ell(w'))} \quad = ?$$

As $\mathrm{OPT}(L) \geq (1-\delta)|E|$, the probability that $\underbrace{\pi_w(\ell(w)) = \ell(v) = \pi_{w'}(\ell(w'))}_{*}$ is $\geq 1 - 2\delta$

If $*$ occurs, then $\Pr[e \in (S, V \setminus S)] = \frac{1-\rho}{2}$

Overall: $\Pr[e \in (S, U \setminus S)] \geq (1-2\delta)\frac{1-\rho}{2} =$

$$= \frac{1-\rho}{2} - \delta + \delta\rho = \frac{1-\rho}{2} - \varepsilon \quad \text{if } \delta = \frac{\varepsilon}{1-\rho}$$

$\Rightarrow$ the total weight of the cut is $\geq \frac{1-\rho}{2} - \varepsilon$ fraction of the weight of all edges. $\blacksquare$

Parameters of PCPS: # random bits
queries

PCP

DISGRESSION

Probabilistically checkable proof system with completeness c and soundness s, $0 \le s \le c \le 1$, for problem P, is a system of a prover and verifier:

- given a claimed solution x for P, the prover constructs a "proof" that x solves P
  ("string $\in \{-1,1\}^*$")
- the verifier checks the proof and either decides to accept it or to reject it;
  if $x \in$ ULC, then the accept. prob. $\ge c$
  if $x \notin$ ULC $\le s$.

For $P =$ GAP-ULC,

the proof is just the labeling of $V \cup W$ represented by long codes (prover).

View it as a cut of H ... $(S, V \setminus S)$

$\approx$ coloring $f : V \to \{1, -1\}$

The verifier selects $(u,v) \in F$ according to D and accepts the proof if $f(u) \neq f(v)$,
rejects if $f(u) = f(v)$.

THM A rephrased: under the same assumptions, there is a PCPS reading just two bits s.t.

a) $OPT(I) \ge (1-\delta)|E| \Rightarrow Prob[Accept] \ge \frac{1-\rho}{2} - \varepsilon$

b) $OPT(I) \le \delta |E| \Rightarrow Prob[Accept] \le \frac{\arccos \rho}{\pi} + \varepsilon$.

Proof of THM A, part b) - HIGH LEVEL IDEA

GOAL - prove $OPT(I) \le \delta|E| \Rightarrow \mathrm{Prob}[Acc] \le \frac{\arccos \rho}{\pi} + \varepsilon$

By contrapositive, assume $\mathrm{Prob}[Acc] > \frac{\arccos \rho}{\pi} + \varepsilon$.

We show that then

interpretation of the col in cloud as a function

$\exists j$ s.t. "influence" of the j-th coordinate of f_w and $f_{w'}$ is large $\approx$ the functions resemble dictators

$\approx\!\!>$ use j to color w, w' s.t. $> \delta|E|$ constraints satisfied - - a contradiction.

machinery - Fourier analysis of Boolean functions

Proof of THM A, part b) - more details

$$\Pr[\text{Accept}] = \mathop{\mathbb{E}}_{v,w,w',x,y,\mu}\left[\frac{1}{2} - \frac{1}{2} f_w(x\circ\pi_w) f_{w'}(y\circ\pi_{w'})\right]$$

$$= \frac{1}{2} - \frac{1}{2}\mathop{\mathbb{E}}_{v,x,\mu}\left[\mathop{\mathbb{E}}_{w\in N(v)}[f_w(x\circ\pi_w)]\mathop{\mathbb{E}}_{w'\in N(v)}[f_{w'}(y\circ\pi_{w'})]\right]$$

define $g_v(x) = \mathop{\mathbb{E}}_{w\in N(v)} f_w(x\circ\pi_w)$, then

$$= \frac{1}{2} - \frac{1}{2}\mathop{\mathbb{E}}_{v,x,\mu}\left[g_v(x)\cdot g_v(y)\right]$$

define $S_\rho(g_v) = \mathop{\mathbb{E}}_{x,y}\left[g_v(x)\cdot g_v(y)\right]$ ... $y_i := \begin{cases} -x_i & \text{w.p. } \frac{1-\rho}{2} \\ x_i & \frac{1+\rho}{2}\end{cases}$

"stability of g"

then

$$= \frac{1}{2} - \frac{1}{2}\mathop{\mathbb{E}}_{v}\left[S_\rho(g_v)\right]$$

As $\Pr[\text{Acc}] > \frac{\arccos\rho}{\pi} + \varepsilon$, there must be at least $\frac{\varepsilon}{2}$-fraction of vertices (called V-good)

s.t. $\Pr[\text{Acc} \mid v \text{ is good}] \geq \frac{\arccos\rho}{\pi} + \frac{\varepsilon}{2}$

(otherwise, the total success probability is at most

$$\left(1-\frac{\varepsilon}{2}\right)\left(\frac{\arccos\rho}{\pi} + \frac{\varepsilon}{2}\right) + \frac{\varepsilon}{2}\cdot 1 < \frac{\arccos\rho}{\pi} + \varepsilon \;\lightning)$$

$\Rightarrow$ for a V-good value v:

$$\frac{1}{2} - \frac{1}{2} S_\rho(g_v) \geq \frac{\arccos \rho}{\pi} + \frac{\varepsilon}{2}$$

i.e. $$S_\rho(g_v) \leq 1 - \frac{2 \arccos \rho}{\pi} - \varepsilon$$

Thm (Majority is the stablest) (as in Newman)

Let $\rho \in (-1, 0)$. For any $\varepsilon > 0$ there is $\delta > 0$ and $k > 0$ s.t. if $f : \{-1,1\}^m \to \{-1,1\}$ satisfies
- $\mathbb{E}[f] = 0$
- and $\forall i \in \{1, \dots, m\}$: $\mathrm{Inf}_i^{\leq k}(f) < \delta$.

("f is far from a dictator")

then $$S_\rho(f) > 1 - \frac{2 \arccos \rho}{\pi} - \varepsilon .$$

Thus, for any good value v there exists j s.t. $\mathrm{Inf}_j^{\leq k}(g_v) \geq \delta$ - label v by j

Then (Fourier analysis) one can show

$$\delta \leq \mathbb{E}_{w \in N(v)} \left[\mathrm{Inf}^{\leq k}_{\pi^{-1}_w(j)} (f_w) \right]$$

Thus, for at least $\frac{\delta}{2}$ fraction of neighbours w of v it holds

$$\mathrm{Inf}^{\le k}_{\pi_w^{-1}(j)}(f_w) \ge \frac{\delta}{2}$$

$\Uparrow$ (averaging argument, as at the bottom of p. 8)

we call these W-good vertices

For a W-good w, let

$$\mathrm{cand}(w) = \left\{ i \in \Sigma \;\middle|\; \mathrm{Inf}^{\le k}_{\pi_w^{-1}(j)}(f_w) \ge \frac{\delta}{2} \right\}$$

label candidates for w

Since $\sum_{i \in \Sigma} \mathrm{Inf}_i^{\le k}(f) \le k$

$$\left(= \sum_{|S| \le k} |S| \hat{f}(S)^2 \le k \cdot \sum_S \hat{f}(S)^2 = k \right),$$

Fourier coefficients

$$|\mathrm{cand}(w)| \le \frac{2k}{\delta} \qquad (\text{i.e., small}).$$

label w with a label chosen uniformly at random from cand(w)

$\Rightarrow$ The probability that edge (v, w) satisfied is

$$\ge \frac{\varepsilon}{2} \cdot \frac{\delta}{2} \cdot \frac{1}{|\mathrm{cand}(w)|} \ge \frac{\varepsilon}{2} \cdot \frac{\delta}{2} \cdot \frac{\delta}{2k} =: \varepsilon'(\varepsilon, \delta)$$

fraction of V-good

W-good

$\Rightarrow$ our labeling satisfies $\ge \varepsilon'$-fraction of all edges – a contradiction ☒