Penn State University
CSE 564 Complexity of Combinatorial Problems

Instructor: Pei Wu

Module

16

Unique Games and Max Cut

This lecture introduces the *Unique Games Conjecture* (UGC) and explains one of its most striking consequences: under UGC, the celebrated Goemans–Williamson semidefinite programming algorithm for MAXCUT is essentially optimal. The key technical component is a 2-query PCP with predicate "$\neq$" (equivalently, a cut constraint), obtained by composing a Unique Games instance with the Long Code and a noisy inequality test.

On the completeness side, an (almost) satisfying labeling of the Unique Games instance induces (dictator) Long Codes that pass the test with probability $(1-\rho)/2$ where ρ is the correlation parameter of the noise. On the soundness side, one shows that if the test accepts noticeably more than the Gaussian benchmark $\arccos(\rho)/\pi$, then one can decode a nontrivial labeling for the Unique Games instance, violating the assumed soundness. The analysis follows Håstad's Fourier-analytic framework and uses a "majority-is-stablest" type theorem to relate *noise stability* to *low-degree influences.*

16.1 Unique Games and the Unique Games Conjecture

A Unique Games instance is a special 2-CSP (often presented as a projection game / Label Cover) in which every constraint is a bijection.

Definition 16.1 (Unique Games instance). A *Unique Games* instance I consists of:

- a bipartite graph $G = (L, R, E)$ (often assumed k-regular for simplicity);
- a label set $\Sigma = \{1, 2, \dots, q\}$;
- for each edge $(u, v) \in E$ with $u \in L$, $v \in R$, a bijection (permutation) $\pi_{uv} : \Sigma \to \Sigma$.

A labeling $\ell : L \cup R \to \Sigma$ *satisfies* an edge (u, v) if

$$\pi_{uv}(\ell(u)) = \ell(v).$$

The value of the instance is the maximum fraction of satisfied edges:

$$\mathrm{val}(I) := \max_{\ell: L\cup R\to\Sigma} \Pr_{(u,v)\sim E}\big[\pi_{uv}(\ell(u)) = \ell(v)\big],$$

where $(u, v) \sim E$ denotes a uniformly random edge.

Conjecture 16.2 (Unique Games Conjecture (Khot)). For every $\varepsilon > 0$ and every $\gamma > 0$, it is NP-hard to distinguish between Unique Games instances I with

$$\mathrm{val}(I) \geq 1 - \varepsilon \qquad \text{and} \qquad \mathrm{val}(I) \leq \gamma.$$

Remark 16.3 (Some consequences (informal)). UGC has many striking implications for hardness of approximation. For example, UGC implies near-tight hardness for Vertex Cover (Khot–Regev) and implies optimality of the Goemans–Williamson approximation ratio for MAXCUT (Khot–Kindler–Mossel–O'Donnell). There are also results showing NP-hardness of nontrivial (c, s)-versions of Unique Games even for constant completeness c (e.g. $c = 1/2$) in various regimes.

16.2 UGC implies Max Cut inapproximability

The reduction in the notes is easiest to describe in two conceptual steps.

Step 1: UG $\Rightarrow$ a 2-query PCP with predicate "$\neq$". Fix a correlation parameter $\rho \in (-1, 0)$. From a Unique Games instance I we build a 2-query PCP whose proof alphabet is $\Sigma_{\mathrm{PCP}} = \{\pm 1\}$, and whose verifier accepts iff the two queried bits are *different*. The completeness and soundness parameters are:

$$\text{completeness} \approx \frac{1-\rho}{2}, \qquad \text{soundness} \approx \frac{\arccos(\rho)}{\pi}.$$

The gap is achieved by composing UG with the Long Code and using a noisy inequality test.

Step 2: a 2-query "$\neq$"-PCP $\Rightarrow$ a (weighted) MAXCUT **instance.** Given a 2-query PCP with a distribution over queried pairs of proof locations (i, j) and acceptance predicate "$\pi(i) \neq \pi(j)$", we create a weighted graph H:

- each proof location i becomes a vertex of H;
- for each ordered pair (i, j), put an edge of weight equal to the probability that the verifier queries (i, j).

An assignment of proof bits $\pi \in \{\pm 1\}^V$ is exactly a cut of H, and the PCP acceptance probability equals the weight of the cut (normalized to total weight 1). Hence distinguishing completeness vs. soundness of the PCP is the same as distinguishing large cut vs. small cut in H. (There is a standard gadget to convert weighted graphs to unweighted graphs while preserving the cut value up to $\pm\varepsilon$.)

Putting the two steps together gives the following informal consequence: for every $\varepsilon > 0$ and every $\rho \in (-1, 0)$, under UGC it is NP-hard to distinguish graphs with

$$\mathrm{MAXCUT}(H) \geq \frac{1-\rho}{2} - \varepsilon \qquad \text{from graphs with} \qquad \mathrm{MAXCUT}(H) \leq \frac{\arccos(\rho)}{\pi} + \varepsilon.$$

Optimizing over ρ yields the Goemans–Williamson threshold and explains why UGC implies the optimality of their approximation ratio.

16.3 The PCP from a Unique Games instance

Fix a Unique Games instance $I = (L, R, E, \Sigma, \{\pi_{uv}\})$ with $\Sigma = \{1, \ldots, q\}$.

The proof (Long Codes). For each right vertex $v \in R$, the prover supplies a Boolean function

$$f_v : \{-1, 1\}^\Sigma \to \{-1, 1\}.$$

An *honest* proof corresponds to an underlying labeling $\ell : L \cup R \to \Sigma$ and sets each f_v to be the dictator on coordinate $\ell(v)$:

$$f_v(x) = x_{\ell(v)}.$$

Permutations acting on strings. Given a permutation $\pi : \Sigma \to \Sigma$, we let it act on strings $x \in \{-1, 1\}^\Sigma$ by

$$(\pi(x))_i := x_{\pi^{-1}(i)}.$$

With this convention, if $\ell(v) = \pi(\ell(u))$ then

$$f_v(\pi(x)) = (\pi(x))_{\ell(v)} = x_{\pi^{-1}(\ell(v))} = x_{\ell(u)}.$$

The verifier: a 2-query inequality test. The verifier performs:

1. Sample $u \in L$ uniformly. Independently sample $v, v' \in N(u)$ uniformly from the neighborhood of u.
2. Sample $x \in \{-1, 1\}^\Sigma$ uniformly. Sample $y \sim_\rho x$ by correlating each coordinate: for each $i \in \Sigma$,
$$y_i = x_i \text{ with probability } \frac{1+\rho}{2}, \qquad y_i = -x_i \text{ with probability } \frac{1-\rho}{2}.$$
3. Query the proof at two locations and accept iff the answers differ:
$$\text{accept iff} \quad f_v(\pi_{uv}(x)) \neq f_{v'}(\pi_{uv'}(y)).$$

16.4 Completeness

Claim 16.4 (Completeness calculation). *If* $\mathrm{val}(I) \geq 1 - \varepsilon$*, then there is a proof (dictators coming from a nearly-satisfying labeling) for which the verifier accepts with probability at least* $(1-\rho)/2 - \varepsilon$*.*

Proof. Let ℓ be a labeling satisfying at least a $(1 - \varepsilon)$ fraction of edges, and use the honest proof $f_v(x) = x_{\ell(v)}$.

For a random choice of $u \in L$ and random neighbors $v, v' \in N(u)$, with probability at least $1 - \varepsilon$ both edges (u, v) and (u, v') are satisfied, i.e.

$$\pi_{uv}(\ell(u)) = \ell(v) \qquad \text{and} \qquad \pi_{uv'}(\ell(u)) = \ell(v').$$

Conditioned on this event, using the permutation action convention,

$$f_v(\pi_{uv}(x)) = x_{\ell(u)}, \qquad f_{v'}(\pi_{uv'}(y)) = y_{\ell(u)}.$$

Therefore the verifier accepts iff $x_{\ell(u)} \neq y_{\ell(u)}$.

By construction, each coordinate pair (x_i, y_i) has correlation ρ, hence

$$\mathbb{P}[x_{\ell(u)} \neq y_{\ell(u)}] = \frac{1-\rho}{2}.$$

Accounting for the at most ε probability that one of the sampled constraints is not satisfied gives acceptance probability at least $(1-\rho)/2 - \varepsilon$. □

16.5 Fourier preliminaries: influences and truncated influences

Soundness is the main difficulty. The analysis uses Fourier expansion on $\{-1,1\}^n$ and the notion of (in particular, low-degree) influence.

Fourier expansion. For $f : \{-1,1\}^n \to \mathbb{R}$, write

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S)\, \chi_S(x), \qquad \chi_S(x) := \prod_{i \in S} x_i.$$

Definition 16.5 (Influence and truncated influence). For Boolean $f : \{-1,1\}^n \to \{-1,1\}$, the (Boolean) influence of coordinate i is

$$\mathrm{Inf}_i[f] \;:=\; \mathbb{P}_x\big[f(x) \neq f(x^{\oplus i})\big] \;=\; \sum_{S \ni i} \widehat{f}(S)^2,$$

where $x^{\oplus i}$ is x with the ith bit flipped.

For an integer $k \geq 1$, the *degree-k truncated influence* is

$$\mathrm{Inf}_i^{\leq k}[f] \;:=\; \sum_{\substack{S \ni i \\ |S| \leq k}} \widehat{f}(S)^2.$$

Fact 16.6 (Only few coordinates can have large truncated influence). *Fix $k \geq 1$ and $\delta > 0$. For any $f : \{-1,1\}^n \to [-1,1]$, the set*

$$S(f) := \{\, i \in [n] : \mathrm{Inf}_i^{\leq k}[f] > \delta \,\}$$

has size at most k/δ.

Proof. We sum truncated influences:

$$\sum_{i=1}^{n} \mathrm{Inf}_i^{\leq k}[f] = \sum_{i=1}^{n} \sum_{\substack{S \ni i \\ |S| \leq k}} \widehat{f}(S)^2 = \sum_{|S| \leq k} |S|\; \widehat{f}(S)^2 \leq \sum_{|S| \leq k} k\, \widehat{f}(S)^2 \leq k \sum_{S} \widehat{f}(S)^2.$$

By Parseval, $\sum_S \widehat{f}(S)^2 = \mathbb{E}[f(x)^2] \leq 1$, hence $\sum_i \mathrm{Inf}_i^{\leq k}[f] \leq k$. If more than k/δ coordinates had truncated influence $> \delta$, this sum would exceed k. □

16.6 Soundness: noise stability and decoding a labeling

Noise stability. For $\rho \in [-1,1]$ and a function $g : \{-1,1\}^q \to \mathbb{R}$, define its ρ-noise stability as

$$\mathrm{Stab}_\rho(g) := \mathbb{E}_{(x,y)\ \rho\text{-correlated}}[g(x)g(y)].$$

Acceptance probability in terms of stability. Work in $\{\pm 1\}$ notation. The predicate "$\neq$" can be written as

$$\mathbf{1}[a \neq b] = \frac{1-ab}{2} \qquad \text{for } a,b \in \{-1,1\}.$$

Thus the verifier acceptance probability is

$$\mathbb{P}[\mathrm{acc}] = \mathbb{E}_{u,v,v'}\mathbb{E}_{x\sim_\rho y}\left[\frac{1 - f_v(\pi_{uv}(x))\, f_{v'}(\pi_{uv'}(y))}{2}\right]. \tag{16.1}$$

Define the averaged function

$$g_u(x) := \mathbb{E}_{v\in N(u)}\big[f_v(\pi_{uv}(x))\big] \in [-1,1].$$

Using independence of v and v' and linearity of expectation,

$$\mathbb{P}[\mathrm{acc}] = \mathbb{E}_u\mathbb{E}_{x\sim_\rho y}\left[\frac{1-g_u(x)\, g_u(y)}{2}\right] = \mathbb{E}_u\left[\frac{1}{2} - \frac{1}{2}\mathrm{Stab}_\rho(g_u)\right]. \tag{16.2}$$

Hence if the verifier accepts noticeably more than $\arccos(\rho)/\pi$, then the average stability of the g_u's is noticeably smaller than the Gaussian benchmark $1 - 2\arccos(\rho)/\pi$.

Proposition 16.7 (Majority-is-stablest (as used in the notes)). *Fix $\rho \in (-1,0)$ and $\varepsilon > 0$. There exist parameters $k = k(\rho,\varepsilon)$ (large enough) and $\delta = \delta(\rho,\varepsilon) > 0$ (small enough) such that for every function $f : \{-1,1\}^n \to [-1,1]$,*

$$\Big(\forall i,\ \mathrm{Inf}_i^{\le k}[f] \le \delta\Big) \quad \Longrightarrow \quad \mathrm{Stab}_\rho(f)\ \ge\ 1 - \frac{2\arccos(\rho)}{\pi} - \varepsilon.$$

From high acceptance to an influential coordinate. Combining (16.2) with Proposition 16.7 gives the key contrapositive statement used in the decoding argument: if for some u we have

$$\mathrm{Stab}_\rho(g_u) < 1 - \frac{2\arccos(\rho)}{\pi} - \varepsilon,$$

then necessarily $\max_i \mathrm{Inf}_i^{\le k}[g_u] > \delta$.

Decoding a labeling from influences (proof sketch). Assume $\mathbb{P}[\mathrm{acc}] > \arccos(\rho)/\pi + \varepsilon$ for the given proof. Then by (16.2),

$$\mathbb{E}_u[\mathrm{Stab}_\rho(g_u)] < 1 - \frac{2\arccos(\rho)}{\pi} - 2\varepsilon.$$

By averaging, for at least an $\varepsilon/2$ fraction of vertices $u \in L$,

$$\mathrm{Stab}_\rho(g_u) < 1 - \frac{2\arccos(\rho)}{\pi} - \varepsilon,$$

and hence (by Proposition 16.7) these u satisfy $\max_i \mathrm{Inf}_i^{\le k}[g_u] > \delta$.

For each such $u \in L$, pick a label

$$\ell(u) \in \arg\max_{i\in\Sigma} \mathrm{Inf}_i^{\le k}[g_u].$$

For each $v \in R$, define the set of *candidate labels* with large truncated influence:

$$S_v := \{\, j \in \Sigma : \ \mathrm{Inf}_j^{\le k}[f_v] \ge \delta/2 \,\}.$$

By the "few large influences" fact, $|S_v| \le 2k/\delta$ for all v. Now assign $\ell(v)$ by choosing a uniformly random element of S_v (if $S_v = \emptyset$, pick any label).

It remains to show that this decoding satisfies a nontrivial fraction of edges. The key inequality is that large influence of g_u on coordinate i forces many neighbors v to have large influence of f_v on the corresponding permuted coordinate $\pi_{uv}(i)$:

$$\mathrm{Inf}_i^{\le k}[g_u] \ \le \ \mathbb{E}_{v\in N(u)}\big[\mathrm{Inf}_{\pi_{uv}(i)}^{\le k}[f_v]\big].$$

(We justify this inequality below using Fourier coefficients and Cauchy–Schwarz.) Therefore, if $\mathrm{Inf}_{\ell(u)}^{\le k}[g_u] > \delta$, then for at least a $\delta/2$ fraction of neighbors $v \in N(u)$ we have $\mathrm{Inf}_{\pi_{uv}(\ell(u))}^{\le k}[f_v] \ge \delta/2$, i.e. $\pi_{uv}(\ell(u)) \in S_v$. Conditioned on such an edge (u, v), our random choice of $\ell(v) \in S_v$ satisfies

$$\mathbb{P}[\ell(v) = \pi_{uv}(\ell(u))] \ \ge \ \frac{1}{|S_v|} \ \ge \ \frac{\delta}{2k}.$$

Hence, for a "good" u,

$$\mathbb{P}_{v\sim N(u)}\big[\pi_{uv}(\ell(u)) = \ell(v)\big] \ \ge \ \frac{\delta}{2} \cdot \frac{\delta}{2k} = \frac{\delta^2}{4k}.$$

Since at least an $\varepsilon/2$ fraction of $u \in L$ are good, the expected fraction of satisfied edges is at least

$$\frac{\varepsilon}{2} \cdot \frac{\delta^2}{4k} = \frac{\varepsilon\delta^2}{8k}.$$

In particular, if the verifier acceptance is $> \arccos(\rho)/\pi + \varepsilon$, then we can decode a labeling with $\mathrm{val}(I) \ge \varepsilon\delta^2/(8k)$, contradicting the Unique Games soundness assumption for sufficiently small $\mathrm{val}(I)$. This proves the desired soundness bound.

Relating Fourier coefficients under permutations. Finally, we record the Fourier calculation that underlies the influence inequality above. Let

$$g_u(x) = \mathbb{E}_{v\in N(u)}[f_v(\pi_{uv}(x))].$$

For every $S \subseteq \Sigma$,

$$\begin{aligned} \widehat{g}_u(S) &= \mathbb{E}_x\big[g_u(x)\chi_S(x)\big] = \mathbb{E}_x\mathbb{E}_v\big[f_v(\pi_{uv}(x))\chi_S(x)\big] \\ &= \mathbb{E}_v\mathbb{E}_x\big[f_v(\pi_{uv}(x))\chi_S(x)\big] = \mathbb{E}_v\mathbb{E}_x\big[f_v(x)\chi_S(\pi_{uv}^{-1}(x))\big] \\ &= \mathbb{E}_v\mathbb{E}_x\big[f_v(x)\chi_{\pi_{uv}(S)}(x)\big] = \mathbb{E}_v\widehat{f}_v(\pi_{uv}(S)). \end{aligned} \tag{16.3}$$

Using (16.3) and Cauchy–Schwarz,

$$\widehat{g}_u(S)^2 = \Big(\mathbb{E}_v \widehat{f}_v(\pi_{uv}(S))\Big)^2 \leq \mathbb{E}_v \widehat{f}_v(\pi_{uv}(S))^2.$$

Summing over sets S with $i \in S$ and $|S| \leq k$ gives

$$\mathrm{Inf}_i^{\leq k}[g_u] = \sum_{\substack{S \ni i \\ |S| \leq k}} \widehat{g}_u(S)^2 \leq \mathbb{E}_v \sum_{\substack{S \ni i \\ |S| \leq k}} \widehat{f}_v(\pi_{uv}(S))^2 = \mathbb{E}_v \, \mathrm{Inf}_{\pi_{uv}(i)}^{\leq k}[f_v],$$

which is the desired influence inequality.