$MAXCUT$ is UGC-hard to approximate [KKMO05]

Orr Paradise*

January 2018

Abstract

We present a key result of Khot, Kindler, Mossel and O'Donnel [KKMO05] which states that if the Unique Games Conjecture [Khot02] holds then the Goemans-Williamson approximation algorithm [GW95] for $MAXCUT$ is optimal, unless $P = NP$.

Prepared for the Student Seminar in Complexity at the Weizmann Institute of Science, organized by Roei Tell.

`These notes aim to help the students (including, most importantly, me) during the presentation, so informalities and inaccuracies are to be expected.`

1 Preliminaries

1.1 Fourier analysis of Boolean functions

Recall that each $f : \{\pm 1\}^n \to \mathbb{R}$ can be uniquely written as $f = \sum_{S \subseteq [n]} \widehat{f}(S) \chi_S$ where $\chi_S(x) = \prod_{i \in S} x_i$, and the following definitions and facts

Definition 1.1. The *influence of* i *on* f *is*

$$\mathrm{Inf}_i(f) \coloneqq \sum_{S \ni i} \widehat{f}(S)^2$$

The k*-bounded influence of* i *on* f *is*

$$\mathrm{Inf}_i^{\leq k}(f) \coloneqq \sum_{S \ni i, |S| \leq k} \hat{f}(S)^2$$

where $S \ni i$ is shorthand for $S \subseteq [n]$ $s.t$ $i \in S$.

Exercise 1.2. For a boolean $f : \{\pm 1\}^n \to \{\pm 1\}$, $\sum_{i=1}^n \mathrm{Inf}_i^{\leq k}(f) \leq k$. [Guidline: $\sum_{i=1}^n \mathrm{Inf}_i^{\leq k}(f) = \sum_{i=1}^n \sum_{j=1}^k \sum_{i \ni S, |S|=j} \hat{f}(S)^2$. Show that $\sum_{i=1}^n \sum_{i \ni S, |S|=j} \hat{f}(S)^2 \leq j \sum_{|S|=j} \hat{f}(S)^2$ by counting how many time each $S \ni i, |S| = j$ is summed.

We turn to define the notion of a *stability* of a function.

Definition 1.3. For $\rho \in [-1, 1]$ and $x \in \{\pm 1\}^n$ we say that (the random variable) y is ρ*-correlated to* x and denote $y \sim N_\rho(x)$ when y is drawn by independently setting each y_i as follows:

$$y_i \coloneqq \begin{cases} x_i & \text{w.p } \frac{1}{2} + \frac{1}{2}\rho \\ -x_i & \text{w.p } \frac{1}{2} - \frac{1}{2}\rho \end{cases}$$

The process of drawing a ρ*-correlated pair* (x, y) is defined as follows:

*orr.paradise@weizmann.ac.il

1. Draw $x \sim U\left(\{\pm 1\}^n\right)$
2. Draw $y \sim N_\rho(x)$

Using this random variable we can define the *stability* of f, which measures the correlation between $f(x)$ and $f(y)$ when (x, y) is a ρ-correlated pair.

Definition 1.4. Let $f : \{\pm 1\}^n \to \mathbb{R}$ and $\rho \in [-1, 1]$. The *noise stability of f at ρ* is defined by

$$\mathbb{S}_\rho(f) = \mathbb{E}\left[f(x) f(y)\right]$$

where the expectation is taken over ρ-correlated pairs (x, y).

Remark 1.5.

1. If f is boolean then
$$\mathbb{S}_\rho(f) = 2\mathbb{P}\left[f(x) = f(y)\right] - 1$$
where the probability is taken over ρ-correlated pairs (x, y).
2. When $\rho = 1$, $y \equiv x$. When $\rho = -1$, $y \equiv -x$. When $\rho = 0$, $y \sim U\left(\{\pm 1\}^n\right)$.
3. (x, y) is a ρ-correlated pair iff for all $i \in [n]$ x_i is independent of y_i, $\mathbb{E}[x_i] = \mathbb{E}[y_i] = 0$ and $\mathbb{E}[x_i y_i] = \rho$.

We present an alternative view of the noise stability of f. Let T_ρ be the linear operator on the space $\mathbb{R}^{\{\pm 1\}^n}$ that is defined by $T_\rho(f)(x) = \mathbb{E}_{y \sim N_\rho(x)}$. Then

$$T_\rho(\chi_S)(x) = \mathbb{E}_{y \sim N_\rho(x)}\left[\prod_{i \in S} y_i\right] = \prod_{i \in S} \mathbb{E}_{y_i}[y_i] = \prod_{i \in S} \rho x_i = \rho^{|S|} \chi_S(x)$$

and since T_ρ is a linear operator this means that $T_\rho(f) = \sum_{S \subseteq [n]} \rho^{|S|} \widehat{f}(S) \chi_S$. On the other hand,

$$\begin{aligned}
\mathbb{S}_\rho(f) &= \mathbb{E}_{(x,y) \sim \rho\text{-noisy pair}}\left[f(x) f(y)\right] = \mathbb{E}_{x \sim U(\{\pm 1\}^n)}\left[f(x)\, \mathbb{E}_{y \sim N_\rho(x)}\left[f(y)\right]\right] = \langle f, T_\rho f \rangle \\
&= \sum_{S, T \subseteq [n]} \widehat{f}(S) \rho^{|T|} \widehat{f}(T) \langle \chi_S, \chi_T \rangle \\
&= \sum_{S \subseteq [n]} \rho^{|S|} \widehat{f}(S)^2
\end{aligned}$$

1.2 The Long Code

We describe a highly inefficient way of representing numbers in $[n]$.

Definition 1.6. Set $n \in \mathbb{N}$. The *Long Code* of $i \in [n]$ is $\chi_{\{i\}} : \{\pm 1\}^n \to \{\pm 1\}$. One can think of the map

$$LongCode : [n] \to \{\pm 1\}^{\{\pm 1\}^n} \qquad LongCode(i) = \chi_{\{i\}}$$

Notice that we map $\log n$ bits to 2^n bits which is a doubly-exponential blowup.

1.3 The Goemans-Williamson $MAXCUT$ approximation algorithm

Definition 1.7. To us, the *Maximal Cut problem* ($MAXCUT$) is finding a cut of maximal weight in a weighted graph. Formally, the input is a weighted graph $G = (V, E, w)$ where $w : E \to \mathbb{R}_+$ and we wish to find a set $S \subseteq V$ that maximizes $\frac{1}{|E|} \sum_{e \in \left(S \times \overline{S}\right) \cap E} w(e)$.

The Goemans-Williamson algorithm [GW95] approximates $MAXCUT$ with a ratio[1] of α_{GW}, which is given by minimizing some trigonometric expression

$$\alpha_{GW} \coloneqq \min_{\theta\in(0,\pi)} \frac{\theta/\pi}{\frac{1}{2}-\frac{1}{2}\cos\theta} \approx 0.878567 \tag{1.1}$$

The above ratio is derived from the geometric nature of the GW algorithm and seems somewhat arbitrary for a combinatoric problem such as $MAXCUT$. Surprisingly, we show that $MAXCUT$ cannot be approximated with a better ratio, assuming certain conjectures – one is the famous $P \neq NP$, while the other will be described in more details in 3.1.

For this talk no familiarity with the algorithm is needed, however a high-level understanding of it could provide deeper insight. Those may be obtained by glossing over [Cai03, Section 2].

2 Hardness of approximation

2.1 Approximation

To show that a problem Π is tractable, one needs to prove the existence of a polynomial-time algorithm that solves Π (e.g by constructing said algorithm), but showing that is is hard to solve Π efficiently requires more sophisticated tools, namely using NP-hardness. NP-hardness aids us in proving that it is hard to solve all instances of a problem *exactly*. Specifically when dealing with an optimization problem, if $P \neq NP$ then we cannot always find the optimal solution. The study of *approximation algorithms* offers to trade the optimality of the output witht the runtime of the computation.

Definition 2.1. Let $\Pi \in NP$ be some maximization problem with value function[2] $v : \{0,1\}^* \to [0,1]$. Algorithm A is a *ρ-approximation algorithm for* Π if for every $I \in \{0,1\}^*$ instance to Π it holds that

$$\rho v^*(I) \leq v(A(I))$$

where v^* denotes the value of the maximal solution (w.r.t v) of I.

As in the opening paragraph, showing that we can approximate a problem Π to a factor ρ is as simple as constructing an ρ-approximation algorithm (which is sometimes very hard!). But how can we show that it is *hard* to approximate Π to a ratio ρ?

2.2 Gap reductions

Just as in the precise case, a framework of *hardness* will be the solution to our problem. The building block of such a framework is the *reduction*, which we need to adapt to our impercise relaxation.

Definition 2.2. Let $s_1 \leq c_1, s_2 \leq c_2 \in [0,1]$. A *$(s_1, c_1, s_2, c_2)$-gap reduction* from Π_1 to Π_2 (viewed as maximization problems with value functions v_1, v_2 resp.) is a poly-time computable function $g : \{0,1\}^* \to \{0,1\}^*$ that maps instances of Π_1 to instances of Π_2 such that the following holds:

$$\begin{aligned} c_1 \leq v_1^*(I) \quad &\Longrightarrow \quad c_2 \leq v_2^*(g(I)) \\ v_1^*(I) < s_1 \quad &\Longrightarrow \quad v_2^*(g(I)) < s_2 \end{aligned}$$

Notice that there is no constraint on the behavior of the reduction on Is for which $v_1^*(I) \in (s_1|I|, c_1|I|)$

We say that Π is *NP-hard to (s,c)-distinguish* if every (decision) problem $\Delta \in NP$ is $(1,1,s,c)$-gap reducible to Π, where Δ is endowed with the value function χ_Δ[3] of Δ. In words, the reduction(s) should map $x \notin \Delta$ to instance I (of Π) with value at most s, and $x \in \Delta$ to instance I with value at least c.

[1] The reader should be familiar with the notion of an approximation algorithm, although a formal definition will be given.

[2] We assume all value functions are efficiently computable w.r.t the input and the solution

[3] Let S be a set. The characteristic function of S is defined by $\chi_S(x) = \begin{cases} 1 & x \in S \\ 0 & x \notin S \end{cases}$

We justify the existence of these definitions with the following theorem:

Theorem 2.3. *(Gap reduction)*

1. *Assume* Π_1 *is* (s_1, c_1, s_2, c_2)*-gap reducible to* Π_2*. If* Π_1 *is* NP*-hard to* (s_1, c_1)*-distinguish then* Π_2 *is* NP*-hard to* (s_2, c_2)*-distinguish.*
2. *Assume* Π *is* NP*-hard to* (s, c)*-distinguish. If there exists a polynomial time* $\frac{s}{c}$*-approximation algorithm for* Π*, then* $P = NP$*.*

Proof.

1. Set $\Delta \in NP$. Take g, h to be the gap reductions $\Pi_0 \xrightarrow{g} \Pi_1 \xrightarrow{h} \Pi_2$ with gaps as in the theorem statement. Then $h \circ g$ is poly-time computable, and indeed

$$\begin{aligned} x \in \Delta &\implies c_1 \leq v_1^* \left(g\left(x\right)\right) \implies c_2 \leq v_2^* \left(\left(h \circ g\right)\left(x\right)\right) \\ x \notin \Delta &\implies v_1^* \left(g\left(x\right)\right) < s_1 \implies v_2^* \left(\left(h \circ g\right)\left(x\right)\right) < s_2 \end{aligned}$$

2. Assume there is a polynomial time $\frac{s}{c}$-approximation algorithm for Π. Let $\Delta \in NP$ and take the corresponding gap reduction g. The polynomial time decider for Δ, upon receiving input x, will output "Yes" iff $s \leq v\left(A\left(g\left(x\right)\right)\right)$. It is indeed polynomial, and it is correct:
   - If $x \in \Delta$ then $c \leq v^*\left(g\left(x\right)\right)$, so

   $$s \leq \frac{s}{c} v^*\left(g\left(x\right)\right) \leq v\left(A\left(x\right)\right)$$

   - If $x \notin \Delta$ then $v^*\left(g\left(x\right)\right) < s$, by definition of v^* we have $v\left(A\left(g\left(x\right)\right)\right) < s$

So gap reductions provide us with an NP-hardness theory for approximation problems. Since we know of problems that are NP-hard to approximate ([SG76]), we can develop this theory. Since we are particularly interested in $MAXCUT$, the following observation will prove useful: □

Remark 2.4. Crescenzi, Silvestri and Trevisan proved that the approximation thresholds of $MAXCUT$ and $unweighted-MAXCUT$ are equal, which for our interest means that there exists an α such that $MAXCUT$ and $unweighted - MAXCUT$ are α-approximable, but for any $\varepsilon > 0$ if there is an $(\alpha + \varepsilon)$-approximation algorithm for $MAXCUT$ or $unweighted - MAXCUT$ then $P = NP$ [CST01].

3 Two tools

3.1 The Unique Games Conjecture

Though there is a plethora of hardness of approximation results , some much desired results seem to be too elusive without making further assumptions. In particular, we prove a hardness of approximation result for $MAXCUT$ assuming an important conjecture regarding the Unique Games problem which improves on the previously best-known ratio, $16/17 \approx 0.941176$ due to [Håstad01].

Definition 3.1. An instance to the *Unique M-Label Cover* problem ($ULC\left(M\right)$) is composed of a bipartite graph $G = (V, W, E)$ (we assume all edges are oriented from V to W) and a set of constraints $\left\{\pi_e\right\}_{e \in E}$ such that for every $e \in E$, π_e is a permutation of $[M]$. A solution is an assignment $\sigma : V \sqcup W \to [M]$ whose value is defined

$$v\left(\sigma\right) := \Pr_{(v,w) \sim U(E)} \left[\pi_{v,w}\left(\sigma\left(w\right)\right) = \sigma\left(v\right)\right]$$

that is, $v\left(\sigma\right)$ is the fraction of constraints it satisfies.

Notice that for every M, deciding whether an instance to $ULC(M)$ is completely satisfiable (that is, whether $v^*(I) = 1$) is tractable: Assuming G is connected, choose an arbitrary $v \in V$. Then for each $i \in [M]$ set $\sigma(v) = i$ – by uniqueness of the constraints, this completely decides the rest of the labels so we may iteratively label all vertices in the graph with BFS. If a contradiction is reached (that is, we attempt to relabel a vertex with a different label) then continue to the next $i \in [M]$, otherwise we found a satisfying σ.

The Unique Games Conjecture (originally formulated in [Khot02]) states that it is hard to distinguish between highly and slightly satisfiable instances, and has a fascinating history ([Klarreich11]).

Conjecture 3.2. *(Unique Games)For every $s \leq c \in (0,1)$ there exists an $M \in \mathbb{N}$ such that $ULC(M)$ is NP-hard to (s,c)-distinguish.*

Assuming the Unique Games Conjecture provides us with a family of problems that are NP-hard to (s,c)-distinguish, which will be one of two key tools used in our main result. A problem that is NP-hard to approximate under this assumption is sometimes called UG-hard to approximate – but note that UGC is not a class of problems so there is some abuse in this term.

We end the introduction of this key player with the following observation.

Remark 3.3. Without loss of generality, we assume that all instances to ULC are regular on the V side, that is that for all $v, v' \in V$ $\deg v = \deg v'$. This is a step in the proof that the weighted and unweighted versions of UGC are equivalent, specifically [KR08, Lemma 3.4].

3.2 Majority is Stablest (revisited)

Another key player in our proof will be the Majority Is Stablest theorem (which was only a conjecture when [KKMO05] was originally published!). Although the below formulation is somewhat different from the one we saw a few weeks ago, we will use it as a black-box for our proof:

Theorem 3.4. *(Majority Is Stablest) Let $\rho \in [0,1)$. For any $\varepsilon > 0$ there is $\delta > 0$ such that if $f : \{\pm 1\}^n \to [-1,1]$ satisfies*

$$\mathbb{E}[f] = 0, \quad \forall i \in [n] \ \mathrm{Inf}_i(f) \leq \delta$$

then

$$\mathbb{S}_\rho(f) \leq 1 - \frac{2}{\pi}\arccos\rho + \varepsilon$$

Defining $\mathrm{Dict}_n^i, \mathrm{Maj}_n : \{\pm 1\}^n \to \{\pm 1\}$ by

$$\mathrm{Dict}_n^i(x) \coloneqq x_i, \quad \mathrm{Maj}_n(x) \coloneqq \mathrm{sign}\left(\frac{1}{2} + \sum_{i=1}^n x_i\right) = \begin{cases} 1 & |\{i | x_i = -1\}| \leq |\{i | x_i = 1\}| \\ -1 & \text{else} \end{cases}$$

it can be shown easily and less easily (resp.) be shown that

$$\forall i, n \ \mathbb{S}_\rho\left(\mathrm{Dict}_n^i\right) = \rho \quad \lim_{n\to\infty} \mathbb{S}_\rho(\mathrm{Maj}_n) = 1 - \frac{2}{\pi}\arccos\rho$$

so theorem 3.4 means that of all (zero-meaned) functions that aren't close to being dictatorships (meaning no single coordinate has significant influence), the majority function is stablest.

In fact, we will expand our black-box slightly to include a slightly different formulation of the Majority Is Stablest theorem, tailor-made for our main proof.

Proposition 3.5. *Let $\rho \in (-1,0]$. For any $\varepsilon > 0$ there is $\delta > 0$ and $k \in \mathbb{N}$ such that if $f : \{\pm 1\}^n \to [-1,1]$ satisfies*

$$\forall i \in [n] \ \mathrm{Inf}_i^{\leq k}(f) \leq \delta$$

then

$$\mathbb{S}_\rho(f) > 1 - \frac{2}{\pi}\arccos\rho - \varepsilon$$

The differences are that we take negative ρ (and switch the inequality accordingly), and generalize to all (not necessarily zero-meaned) fs that have small k-bounded influence. The proposition follows from the two following rather technical claims.

Claim 3.6. For all $\rho \in [0,1)$,$\varepsilon > 0$ there are $\delta > 0$,$k \in \mathbb{N}$ such that if $f : \{\pm 1\}^n \to [-1,1]$ satisfies

$$\mathbb{E}[f] = 0, \quad \forall i \in [n] \ \mathrm{Inf}_i(f) \leq \delta$$

then

$$\mathbb{S}_\rho(f) \leq 1 - \frac{2}{\pi}\arccos\rho + \varepsilon$$

Proof. Let $\rho \in [0,1)$,$\varepsilon > 0$. Take[4] $\gamma \in (0,1)$ small enough such that such that for all k $\rho^k \left(1 - (1-\gamma)^{2k}\right) < \frac{\varepsilon}{4}$. Take $\delta' > 0$ from 3.4 applied to ρ and $\frac{\varepsilon}{4}$. Choose $\delta = \frac{\delta'}{2}$, and k such that $(1-\gamma)^{2k} < \delta$.

Now, take f that satisfies $\mathbb{E}[f] = 0$ and $\mathrm{Inf}_i^{\leq k}(f) \leq \delta$ for all i, and let $T_{1-\gamma}f$ where $T_{1-\gamma}f$ is the noise operator with noise $1-\gamma$. Then

$$\begin{aligned}
\mathrm{Inf}_i(T_{1-\gamma}f) &= \sum_{i \ni S} \widehat{T_{1-\gamma}f}(S)^2 = \sum_{i \ni S} (1-\gamma)^{2|S|}\hat{f}(S)^2 \\
&\leq \sum_{i \ni S, |S| \leq k} (1-\gamma)^{2|S|}\hat{f}(S)^2 + (1-\gamma)^{2k} \sum_{i \ni S, |S| > k} \hat{f}(S)^2 \\
&\leq \sum_{i \ni S, |S| \leq k} \hat{f}(S)^2 + (1-\gamma)^{2k} \\
&= \mathrm{Inf}_i^{\leq k}(f) + (1-\gamma)^k \\
&\leq 2\delta = \delta'
\end{aligned}$$

since $\mathbb{E}[T_{1-\gamma}f] = 0$, from 3.4 it must holds that $\mathbb{S}_\rho(T_{1-\gamma}f) \leq 1 - \frac{2}{\pi}\arccos\rho + \frac{\varepsilon}{2}$. Finally, notice that

$$\begin{aligned}
\mathbb{S}_\rho(T_{1-\gamma}f) &= \sum_{S \subseteq [n]} \rho^{|S|}\widehat{T_{1-\gamma}f}(S)^2 = \sum_{S \subseteq [n]} \rho^{|S|}(1-\gamma)^{2|S|}\hat{f}(S)^2 \\
&= \sum_{S \subseteq [n]} \left(\rho^{|S|}(1-\gamma)^{2|S|}\hat{f}(S)^2 - \rho^{|S|}\hat{f}(S)^2\right) + \mathbb{S}_i(f) \\
&= \mathbb{S}_i(f) + \underbrace{\sum_{S \subseteq [n]} \left((1-\gamma)^{2|S|} - 1\right)\rho^{|S|}\hat{f}(S)^2}_{A}
\end{aligned}$$

where

$$-A = \sum_{S \subseteq [n]} \left(1 - (1-\gamma)^{2|S|}\right)\rho^{|S|}\hat{f}(S)^2 \leq \sum_{S \subseteq [n]} \frac{\varepsilon}{4}\hat{f}(S)^2 \leq \frac{\varepsilon}{4}$$

Which gives us

$$\mathbb{S}_i(f) = \mathbb{S}_\rho(T_{1-\gamma}f) - A \leq \mathbb{S}_\rho(T_{1-\gamma}f) \leq 1 - \frac{2}{\pi}\arccos\rho + \frac{1}{2}\varepsilon + \frac{1}{4}\varepsilon$$

□

Claim 3.7. For all $\rho \in (-1,0]$, $\varepsilon > 0$ there is $\delta > 0$ such that if $f : \{\pm 1\}^n \to [-1,1]$ satisfies for all $i \in [n]$ $\mathrm{Inf}_i(f) \leq \delta$ then $\mathbb{S}_\rho(f) \geq 1 - \frac{2}{\pi}\arccos\rho - \varepsilon$.

[4] For all $\gamma \in (0,1)$ the LHS tends to 0 as $k \to \infty$. Fix an arbitrary γ, then there is a K s.t for every $k \geq K$ the inequality holds. Now shrink γ so that the inequality holds also for $k < K$.

Proof. Let ρ, ε and take δ corresponding to $-\rho, \varepsilon$ in 3.4. Let $g(x) := \frac{f(x)-f(-x)}{2} = \sum_{|S| \text{ odd}} \widehat{f}(S) \chi_S(x)$, where the latter inequality can be seen by noticing that $\langle f(-x), \chi_S(x) \rangle = (-1)^{|S|} \widehat{f}(S)$. Then, $\mathbb{E}[g] = 0$ and $\text{Inf}_i(g) \le \text{Inf}_i(f) \le \delta$ so

$$\mathbb{S}_{-\rho}(g) \le 1 - \frac{2}{\pi} \arccos -\rho + \varepsilon = 1 - \frac{2}{\pi}(\pi - \arccos \rho) + \varepsilon = -\left(1 + \frac{2}{\pi} \arccos \rho - \varepsilon\right)$$

Therefore

$$\mathbb{S}_\rho(f) = \sum_{S \subseteq [n]} \rho^{|S|} \widehat{f}(S) \ge \sum_{S \subseteq [n], |S| \text{ odd}} \rho^{|S|} \widehat{f}(S) \, \mathbb{S}_\rho(g) = -\mathbb{S}_{-\rho}(g) \ge 1 + \frac{2}{\pi} \arccos \rho - \varepsilon$$

□

4 Main result

Without further ado, we present and prove the main result

Theorem 4.1. *Assume the Unique Games Conjecture 3.2. For any $\rho \in (-1, 0)$ and $\varepsilon > 0$, $MAXCUT$ is NP-hard to $\left(\frac{\arccos \rho}{\pi} + \varepsilon, \frac{1}{2} - \frac{1}{2}\rho\right)$-distinguish.*

Corollary 4.2. *Assuming the Unique Games Conjecture, if $MAXCUT$ can be approximated with ratio greater than α_{GW} then $P = NP$.*

First, let's see how 4.2 follows from theorem 4.1.

Proof. Assume UGC and theorem 4.1, and let $\varepsilon' > 0$. Let

$$r : [-1, 0] \to \mathbb{R} \quad r(\rho) := {}^{\frac{\arccos \rho}{\pi}} / {}_{\frac{1}{2} - \frac{1}{2}\rho}$$

Notice that r obtains a unique minimum at $\rho^* := \text{argmin}(r) \approx -0.689$ – either by some analysis (Weierstrass tells us it gets a minimum, find roots of the derivative) or by picture (figure 4.1). Furthermore, $r(\rho^*) = \alpha_{GW}$ – again, either because minimizing r over $\rho \in [-1, 0]$ is equivalent to minimizing it the the fraction in equation (1.1) over $\theta \in \left[\frac{\pi}{2}, \pi\right]$ which in turn is equivalent to minimizing that fraction over $\theta \in [0, \pi]$, or because of a figure 4.1.

theorem 4.1 with $\rho = \rho^* \in (-1, 0)$ and $\varepsilon = \left(\frac{1}{2} - \frac{1}{2}\rho^*\right) \varepsilon' > 0$ tells us that $MAXCUT$ is NP-hard to (s, c)-distinguish, where $s = \frac{\arccos \rho^*}{\pi} + \left(\frac{1}{2} - \frac{1}{2}\rho^*\right) \varepsilon'$ and $c = \frac{1}{2} - \frac{1}{2}\rho^*$. Notice that

$$\frac{s}{c} = \frac{\frac{\arccos \rho^*}{\pi} + \left(\frac{1}{2} - \frac{1}{2}\rho^*\right) \varepsilon'}{\frac{1}{2} - \frac{1}{2}\rho^*} = \frac{\frac{\arccos \rho^*}{\pi}}{\frac{1}{2} - \frac{1}{2}\rho^*} + \varepsilon' = \alpha_{GW} + \varepsilon'$$

so by theorem 2.3, if there exists a poly-time algorithm that achieves approximation ratio $\alpha_{GW} + \varepsilon'$ then $P = NP$.

□

Before moving on to the proof, we introduce one final piece of notation: For a vector $x \in \{\pm 1\}^n$ and a permutation $\pi \in Sym_n$, we obtain the vector $x^\pi \in \{\pm 1\}^n$ is defined by $x^\pi_i := x_{\pi(i)}$ for all $i \in [n]$.

4.1 The reduction

Let $\rho \in (-1, 0), \varepsilon > 0$. We construct a $\left(\gamma, 1 - \eta, \frac{\arccos \rho}{\pi} + \varepsilon, \frac{1}{2} - \frac{1}{2}\rho\right)$-gap reduction from $ULC(M)$ to $MAXCUT$, where $M \in \mathbb{N}$ is the one corresponding to $(\gamma, 1 - \eta)$ in 3.2, and we will choose γ, η to be sufficiently small later.

Figure 4.1: $(r(\rho) - \alpha_{GW})$ and its minimum

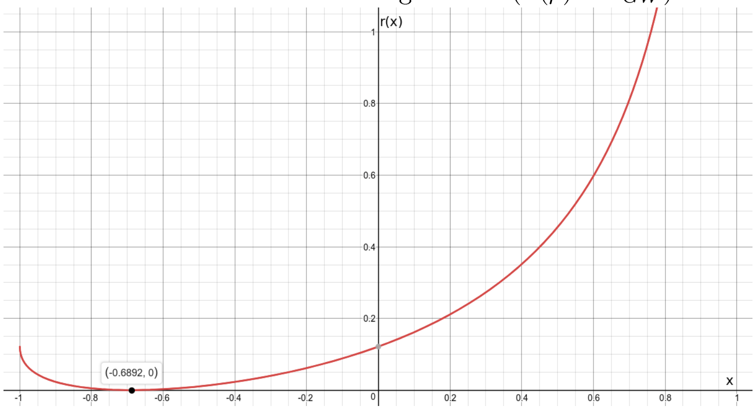

4.1.1 Consistency Test

First, we define the (probabilstic) Consistency Test that takes a weighted instance to $ULC(M)$ denoted by graph $G = (V, W, E, m)$ and constraints $\{\pi_e\}_{e \in E}$, and a set $\{f_w\}_{w \in W}$ where for all $w \in W$, $f_w : \{\pm 1\}^M \to \{\pm 1\}$.

1. Pick $v \sim U(V)$, and then pick two of its neighbours w, w' indepently and uniformly from $\Gamma(v)$.
   (a) Let $\pi := \pi_{(v,w)}$ and $\pi' := \pi_{(v,w')}$ be the respective constraints on these edges.
2. Choose $x \sim U\left(\{\pm 1\}^M\right)$ and $\mu \sim N_\rho(\mathbf{1})$ independently, where $\mathbf{1} \in \{\pm 1\}^M$ is the all-ones vector.
3. Accept iff $f(x^\pi) \neq f'\left(x^{\pi'}\mu\right)$.

Notice that $(x, x\mu)$ is a ρ-correlated pair so $(x\mu, x)$ is the same, therefore the test is symmetric in the inputs f, f' .

Exercise 4.3. If $f = f' = LongCode(i)$ for some $i \in [M]$ then the above test accepts with probability $\left(\frac{1}{2} - \frac{1}{2}\rho\right)$.

4.1.2 The actual reduction

The reduction itself takes as input an instance to $ULC(M)$ as in 4.1.1 and outputs a weighted complete graph $G' = (W', W' \times W', m')$, where W' and m' are defined as follows:

- For each $w \in W$ we construct 2^M vertices in W' that correspond to the truth table of a function $f_w : \{\pm 1\}^M \to \{\pm 1\}$. Formally,

$$W' = \left\{ f_w(x) \,|\, w \in W, x \in \{\pm 1\}^M \right\}$$

- The weight of the edge $\{f_w(x), f_{w'}(y)\}$ is the probability that the test $f_w(x) \overset{?}{\neq} f_{w'}(y)$ is performed in the run of the Consistency Test on input $G = (V, W, E)$.

4.2 Runtime

The reduction seems immensely inefficient, but the trick is that M is constant. G' is a graph on $|W| \cdot 2^M$ vertices, and to compute the weights the reduction needs to simulate all possible "coins" (randomness) of the Consistency Test on the given input. Notice that there are at most $|V| \cdot |W|^2 \cdot 2^{M+1}$ possible outcomes for the random choices, and the run of the test on each choice is polynomial in $|G|$. So, the reduction is polynomial in its input.

4.3 Correctness

We prove that for all $\rho \in (-1,0)$,$\varepsilon, \eta > 0$ there is a γ such that a reduction from $(\gamma, 1-\eta)$-gap-$ULC\,(M)$ to $\left(\frac{\arccos\rho}{\pi} + \varepsilon, \left(\frac{1}{2} - \frac{1}{2}\rho\right)(1-2\eta)\right)$-gap-$MAXCUT$ exists. We can get rid of the $(1-2\eta)$ factor in the completeness by "trading off soundness for completion". More (but not entirely) formally, for a given ε,ρ, take ρ', ε' such that $\left(\frac{1}{2} - \frac{1}{2}\rho'\right)(1-2\eta) = \left(\frac{1}{2} - \frac{1}{2}\rho\right)$ and $\frac{\arccos\rho'}{\pi} + \varepsilon' = \frac{\arccos\rho}{\pi} + \varepsilon$, and take $\eta > 0$ sufficiently small such that $\rho' \in (-1,0)$ and $\varepsilon > 0$. Applying the reduction with$\rho', \varepsilon', \eta$ gives us the gap

$$\left(\frac{\arccos\rho'}{\pi} + \varepsilon', \left(\frac{1}{2} - \frac{1}{2}\rho'\right)(1-2\eta)\right) = \left(\frac{\arccos\rho}{\pi} + \varepsilon, \left(\frac{1}{2} - \frac{1}{2}\rho\right)\right)$$

4.3.1 Completeness

Assume the $ULC\,(M)$ instance has a labeling σ of value at least $(1-\eta)$, that is it satisfies a $(1-\eta)$-fraction of constraints. The cut in G' is obtained by assigning f_w the truth table of the long code of $\sigma(w)$. Formally, the cut in G' is $W' = W'_1 \sqcup W'_{-1}$ where

$$W'_1 = \left\{ f_w(x) \,|\, \chi_{\{\sigma(w)\}}(x) = 1 \right\}, \quad W'_{-1} = \left\{ f_w(x) \,|\, \chi_{\{\sigma(w)\}}(x) = -1 \right\}$$

We argue that this cut has value $(1-2\eta)\left(\frac{1}{2} - \frac{1}{2}\rho\right)$. We say that σ satisfies $(v,w) \in E$ if $\pi_{v,w}(\sigma(w)) = \sigma(v)$. Well, taking probability over $v \sim U(V)$ and $w, w' \sim U(\Gamma(v))$,

$$\begin{aligned} \mathbb{P}\left[\sigma \text{ satisfies } (v,w) \text{ and } (v,w')\right] &= 1 - \mathbb{P}\left[\sigma \text{ doesn't satisfy } (v,w) \text{ or } (v,w')\right] \\ &\geq 1 - 2\mathbb{P}\left[\sigma \text{ doesn't satisfy } (v,w)\right] \\ &= 1 - 2\eta \end{aligned}$$

where the last equality uses the fact that G is regular on the V side (3.3), so that choosing v uniformly and then choosing w uniformly from $\Gamma(v)$ is akin to choosing uniform edge from E.

If σ satsifies (v,w) and (v,w') then

$$\begin{aligned} f_w\left(x^{\pi_{v,w}}\right) &= x_{\pi_{v,w}(\sigma(w))} = x_{\sigma(v)} \\ f_{w'}\left(x^{\pi_{v,w'}}\mu\right) &= x_{\pi_{v,w'}(\sigma(w'))}\mu_{\sigma(w')} = x_{\sigma(v)}\mu_{\sigma(w')} \end{aligned}$$

Combining the above, the Consistency Test chooses satisfied v, w, w' with probability at least $1-2\eta$, and for these vertices the probability that $f_w\left(x^{\pi_{v,w}}\right) \neq f_{w'}\left(x^{\pi_{v,w'}}\mu\right)$ is precisely the probability that $\mu_{\sigma(v)} = -1$, which is $\left(\frac{1}{2} - \frac{1}{2}\rho\right)$. Since μ is chosen independently of v, w, w' we have that the probability acceptance of the Consistency Test is at least $(1-2\eta)\left(\frac{1}{2} - \frac{1}{2}\rho\right)$. Since an edge $\{f_w(x), f_{w'}(x')\}$ is in the cut iff $f_w(x) \neq f_{w'}(x')$ and the weight of such edge is the probabilty that the test $f_w(x) \stackrel{?}{\neq} f_{w'}(x')$ is performed, the weight of the cut is the probability that the Consistency Test accepts, giving us the required lower bound on the value.□

4.3.2 Soundness

We prove that contrapositive. Assume that we have a graph G' with a cut of weight at least $\frac{\arccos\rho}{\pi}+\varepsilon$, and we show that it was obtained from a $ULC(M)$ instance that has an assignment satisfying at least a $\gamma'(\varepsilon,\rho)=\gamma'$-fraction of constraints. Then, since γ' does not depend on M we can take $\gamma<\gamma'$ (enlarging M) to obtain the required result.

From a cut $W'=W_1'\sqcup W_{-1}'$ of weight at least $\frac{\arccos\rho}{\pi}+\varepsilon$ obtain functions $\left\{f_w:\{\pm1\}^M\to\{\pm1\}\right\}_{w\in W}$ by letting $f_w(x)=1$ iff $f_w(x)\in W_1'$. Say that $v\in V$ is good if $\mathbb{P}[\mathrm{acc}|v]\geq\frac{\arccos\rho}{\pi}+\frac{\varepsilon}{2}$, that is if the test accepts with probability higher than the r.h.s when v is drawn. Since[5] $\mathbb{P}[\mathrm{acc}|v\text{ is not good}]<\frac{\arccos\rho}{\pi}+\frac{\varepsilon}{2}$, using the law of total probability we have

$$\frac{\arccos\rho}{\pi}+\varepsilon\leq\mathbb{P}[\mathrm{acc}]\leq\mathbb{P}[v\text{ is good}]+\mathbb{P}[\mathrm{acc}|v\text{ is not good}]<\mathbb{P}[v\text{ is good}]+\frac{\arccos\rho}{\pi}+\frac{\varepsilon}{2}$$

and so $\mathbb{P}[v\text{ is good}]\geq\frac{\varepsilon}{2}$. Second, for fixed $v\in V$ we have the below arithmetization

$$\begin{aligned}\mathbb{E}\left[f_w\left(x^{\pi_{v,w}}\right)f_{w'}\left(x^{\pi_{v,w'}}\mu\right)\right] &= 1\cdot\mathbb{P}\left[f_w\left(x^{\pi_{v,w}}\right)=f_{w'}\left(x^{\pi_{v,w'}}\mu\right)\right]+(-1)\,\mathbb{P}\left[f_w\left(x^{\pi_{v,w}}\right)\neq f_{w'}\left(x^{\pi_{v,w'}}\mu\right)\right]\\ &= \mathbb{P}[\mathrm{rej}|v]-\mathbb{P}[\mathrm{acc}|v]\\ &= 1-2\mathbb{P}[\mathrm{acc}|v]\end{aligned}$$

Where $\mathbb{E}$ and $\mathbb{P}$ are taken over $w,w'\sim U(\Gamma(v))$ and x,μ are drawn as in the Consistency Test. This implies that

$$\begin{aligned}\mathbb{P}[\mathrm{acc}|v] &= \frac{1}{2}-\frac{1}{2}\mathbb{E}_{w,w',x,\mu}\left[f_w\left(x^{\pi_{v,w}}\right)f_{w'}\left(x^{\pi_{v,w'}}\mu\right)\right]\\ \mu\text{'s elements are independent} &= \frac{1}{2}-\frac{1}{2}\mathbb{E}_{w,w',x,\mu}\left[f_w\left(x^{\pi_{v,w}}\right)f_{w'}\left((x\mu)^{\pi_{v,w'}}\right)\right]\\ \text{Law of Total Exp.} &= \frac{1}{2}-\frac{1}{2}\mathbb{E}_{x,\mu}\left[\mathbb{E}_{w,w'}\left[f_w\left(x^{\pi_{v,w}}\right)f_{w'}\left((x\mu)^{\pi_{v,w'}}\right)\right]\right]\\ w,w'\text{ are independent} &= \frac{1}{2}-\frac{1}{2}\mathbb{E}_{x,\mu}\left[\mathbb{E}_{w}\left[f_w\left(x^{\pi_{v,w}}\right)\right]\mathbb{E}_{w'}\left[f_{w'}\left((x\mu)^{\pi_{v,w'}}\right)\right]\right]\\ &= \frac{1}{2}-\frac{1}{2}\mathbb{E}_{x,\mu}\left[g_v(x)\,g_v(x\mu)\right]\\ \text{Law of Total Exp.} &= \frac{1}{2}-\frac{1}{2}\mathbb{S}_\rho(g_v)\end{aligned}$$

where $g_v(z):=\mathbb{E}_{w\sim U(\Gamma(v))}\left[f_w\left(z^{\pi_{v,w}}\right)\right]$ and the expectation is taken over $w,w'\sim U(\Gamma(v))$ and x,μ as in the Consistency Test. So for good v's we have

$$\frac{\arccos\rho}{\pi}+\frac{\varepsilon}{2}\leq\mathbb{P}[\mathrm{acc}|v]=\frac{1}{2}-\frac{1}{2}\mathbb{S}_\rho(g_v)$$

therefore if v is good then

$$\mathbb{S}_\rho(g_v)\leq1-2\frac{\arccos\rho}{\pi}-\varepsilon$$

Finally, from the counterpositive to 3.5, there is a large enough k such that for each good v there exists $\sigma(v)\in[n]$ with $\mathrm{Inf}_{\sigma(v)}^{\leq k}(g_v)>\delta$ (if there is more than one $\sigma(v)$, we fix one arbitrarily). This completes the task of labeling good v's, and we label the rest of V arbitrarily.

What's left is to find labels $\sigma(w)$. Let the candidate set of $w\in W$ be

$$\mathrm{Cand}(w):=\left\{j\in[M]\,|\mathrm{Inf}_j^{\leq k}(f_w)\geq\frac{\delta}{2}\right\}$$

[5] If $c\geq0$ and A,$\{B_i\}_i$ are events s.t $\{B_i\}_i$ are pairwise disjoint and $\mathbb{P}[B_i]>0$ and $\mathbb{P}[A|B_i]<c$ for all i, then $\mathbb{P}\left[A|\bigsqcup_i B_i\right]<c$

From 1.2 we have

$$|\mathrm{Cand}(w)| \cdot \frac{\delta}{2} \leq \sum_{i\in \mathrm{cand}(W)} \mathrm{Inf}_i^{\leq k}(f_w) \leq k \implies |\mathrm{Cand}(w)| \leq \frac{2\delta}{k}$$

and on the other hand, notice that for all good vs

$$\begin{aligned}
\delta &\leq \mathrm{Inf}_\rho^{\leq k}(g_v) \\
&= \sum_{S\ni\sigma(v),|S|\leq k} \widehat{g_v}(S)^2 \\
(*) &= \sum_{S\ni\sigma(v),|S|\leq k} \mathbb{E}_w\left[\widehat{f_w}\left(\pi_{v,w}^{-1}(S)\right)\right]^2 \\
\text{Jensen's ineq.} &\leq \sum_{S\ni\sigma(v),|S|\leq k} \mathbb{E}_w\left[\widehat{f_w}\left(\pi_{v,w}^{-1}(S)\right)^2\right] \\
&= \mathbb{E}_w\left[\sum_{S\ni\sigma(v),|S|\leq k} \widehat{f_w}\left(\pi_{v,w}^{-1}(S)\right)^2\right] \\
&= \mathbb{E}_w\left[\mathrm{Inf}_{\pi_{v,w}^{-1}(\sigma(v))}^{\leq k}(f_w)\right]
\end{aligned}$$

where $(*)$ is obtained by noticing that $\widehat{g_v}(S) = \mathbb{E}_w\left[\widehat{f_w}\left(\pi_{v,w}^{-1}(S)\right)\right]$, by definition of the fourier coefficient and perhaps Fubini's theorem[6]. It follows[7] that for any good v, at least a $\frac{\delta}{2}$-fraction of $w \in \Gamma(v)$ satisfy $\mathrm{Inf}_{\pi_{v,w}^{-1}(\sigma(v))}^{\leq k}(f_w) \geq \frac{\delta}{2}$ so for those ws it holds that $\pi_{v,w}^{-1}(\sigma(v)) \in \mathrm{Cand}(w)$.

Consider the random process that labels $\sigma(w) \sim U(\mathrm{cand}(w))$ if $\mathrm{Cand}(w) \neq \emptyset$, else labels arbitrarily, and keep in mind the desired outcome $\pi_{v,w}(\sigma(w)) = \sigma(v)$ which is iff $\sigma(w) = \pi_{v,w}^{-1}(\sigma(v))$. Then we have an assignment satisfying a $\gamma' = \frac{\varepsilon}{2}\cdot\frac{\delta}{2}\cdot\frac{\delta}{2k}$ fraction of constraints of G – details follow.

From 3.3, drawing $\{v,w\} \sim U(E)$ is equivalent to first choosing $v \sim U(V)$ and then choosing $w \in \Gamma(v)$. Taking $\mathbb{P}$ over $\{v,w\} \sim U(E)$, from the foregoing analysis we have

$$\begin{aligned}
&\mathbb{P}[\{v,w\} \text{ satisfied by } \sigma] \\
\geq& \frac{\varepsilon}{2}\mathbb{P}[\{v,w\} \text{ satisfied by } \sigma | v \text{ is good}] \\
\geq& \frac{\varepsilon}{2}\mathbb{P}\left[\pi_{v,w}^{-1}(\sigma(v)) \in \mathrm{Cand}(w) | v \text{ is good}\right]\mathbb{P}\left[\{v,w\} \text{ satisfied by } \sigma | v \text{ is good}, \pi_{v,w}^{-1}(\sigma(v)) \in \mathrm{Cand}(w)\right] \\
\geq& \frac{\varepsilon}{2}\cdot\frac{\delta}{2}\mathbb{P}\left[\{v,w\} \text{ satisfied by } \sigma | v \text{ is good}, \pi_{v,w}^{-1}(\sigma(v)) \in \mathrm{Cand}(w)\right]
\end{aligned}$$

[6]

$$\begin{aligned}
\widehat{g_v}(S) &= \mathbb{E}_x[g_v(x)\chi_S(x)] = \mathbb{E}_x\left[\mathbb{E}_w\left[\sum_{T\subseteq[n]} \widehat{f_w}(T)\chi_T(x^{\pi_{v,w}})\right]\chi_S(x)\right] \\
&= \sum_{T\subseteq[n]} \mathbb{E}_w\left[\widehat{f_w}(T)\mathbb{E}_x[\chi_T(x^{\pi_{v,w}})\chi_S(x)]\right] = \sum_{T\subseteq[n]} \mathbb{E}_w\left[\widehat{f_w}(T)\mathbb{E}_x\left[\chi_{\pi_{v,w}(T)}(x)\chi_S(x)\right]\right] \\
&= \sum_{T\subseteq[n]} \mathbb{E}_w\left[\widehat{f_w}(T)\mathbb{E}_x\left[\chi_T(x)\chi_{\pi_{v,w}^{-1}(S)}(x)\right]\right] = \mathbb{E}_w\left[\widehat{f_w}\left(\pi_{v,w}^{-1}(S)\right)\right]
\end{aligned}$$

[7] Let $X \leq 1$ be some RV of expectation at least δ, and denote $p = \mathbb{P}\left[X \geq \frac{\delta}{2}\right]$. Then

$$\delta \leq \mathbb{E}[X] \leq 1\cdot p + \frac{\delta}{2}(1-p) \leq p + \frac{\delta}{2} \implies \frac{\delta}{2} \leq p$$

And now taking $\mathbb{E}$ over σ as in the above random process and $\mathbb{P}$ over $\{v,w\} \sim U(E)$ we have

$$\begin{aligned}&\mathbb{E}\left[\mathbb{P}\left[\{v,w\} \text{ satisfied by } \sigma\right]\right]\\ \geq&\frac{\varepsilon}{2}\cdot\frac{\delta}{2}\mathbb{E}\left[\mathbb{P}\left[\{v,w\} \text{ satisfied by } \sigma | v \text{ is good}, \pi_{v,w}^{-1}(\sigma(v)) \in \text{Cand}(w)\right]\right]\end{aligned}$$

so all that's left to show is that $\mathbb{E}_{\sigma}\left[\mathbb{P}_{v,w}\left[\{v,w\} \text{ satisfied by } \sigma | v \text{ is good}, \pi_{v,w}^{-1}(\sigma(v)) \in \text{Cand}(w)\right]\right] \geq \frac{\delta}{2k}$, since then we are guaranteed the existence of a σ for which $\mathbb{P}_{v,w}\left[\{v,w\} \text{ satisfied by } \sigma\right] \geq \frac{\varepsilon}{2}\cdot\frac{\delta}{2}\cdot\frac{\delta}{2k} = \gamma'$. Say that (v,w) are great if v is good and $\pi_{v,w}^{-1}(\sigma(v)) \in \text{Cand}(w)$, we have

$$\begin{aligned}
& & &\mathbb{E}_{\sigma}\left[\mathbb{P}_{v,w}\left[\{v,w\} \text{ satisfied by } \sigma | (v,w) \text{ are great}\right]\right]\\
\text{Law of Total Prob.} &= & &\mathbb{E}_{\sigma}\left[\sum_{v_0,w_0 \text{ great}} \mathbb{P}_{v,w}\left[v=v_0, w=w_0 | v,w \text{ are great}\right]\mathbb{P}_{v,w}\left[\{v,w\} \text{ satisfied by } \sigma | v=v_0, w=w_0\right]\right]\\
\text{Linearity, def of indicator RV} &= & &\sum_{v_0,w_0 \text{ great}} \mathbb{P}_{v,w}\left[v=v_0, w=w_0 | v,w \text{ are great}\right]\mathbb{E}_{\sigma}\left[1_{v_0,w_0 \text{ sat by } \sigma}\right]\\
&= & &\sum_{v_0,w_0 \text{ great}} \mathbb{P}_{v,w}\left[v=v_0, w=w_0 | v,w \text{ are great}\right]\mathbb{P}_{\sigma}\left[\{v_0,w_0\} \text{ sat by } \sigma\right]\\
&\geq & &\sum_{v_0,w_0 \text{ great}} \mathbb{P}_{v,w}\left[v=v_0, w=w_0 | v,w \text{ are great}\right]\frac{\delta}{2k}\\
&= & &\frac{\delta}{2k}\sum_{v_0,w_0 \text{ great}} \mathbb{P}_{v,w}\left[v=v_0, w=w_0 | v,w \text{ are great}\right]\frac{\delta}{2k}\\
&= & &\frac{\delta}{2k}\cdot 1
\end{aligned}$$

where the inequality is because $\sigma \sim U(\text{Cand}(w_0))$, $|\text{Cand}(w_0)| \geq \frac{\delta}{2}$ and $\pi_{v_0,w_0}^{-1}(\sigma(v)) \in \text{Cand}(w_0)$ for great v_0, w_0.

References

[KKMO05] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'donnell. 2005. Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs? (2005). Retrieved from https://www.cs.cmu.edu/~odonnell/papers/maxcut.pdf

[Chekuri06] Chandra Chekuri. 2006. Hardness of Approximation. Retrieved from http://chekuri.cs.illinois.edu/teaching/fall2006/lect21and22.pdf

[SG76] Sartaj Sahni and Teofilo Gonzalez. 1976. P-complete approximation problems. Journal of the ACM (JACM) 23, 3 (1976), 555–565.

[Khot02] Subhash Khot. 2002. On the Power of Unique 2-Prover 1-Round Games. (2002). Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.133.5651&rep=rep1&type=ps

[Klarreich11] Erica Klarreich. 2011. Approximately Hard: The Unique Games Conjecture. Simons Foundation. Retrieved from https://www.simonsfoundation.org/2011/10/06/approximately-hard-the-unique-games-conjecture/

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. 2005. Noise stability of functions with low influences: invariance and optimality. In Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on, 21–30.

[CST01] Pierluigi Crescenzi, Riccardo Silvestri, and Luca Trevisan. 2001. On Weighted vs Unweighted Versions of Combinatorial Optimization Problems. Information and Computation 167, 1 (May 2001), 10–26. DOI:https://doi.org/10.1006/inco.2000.3011

[KR08] Subhash Khot and Oded Regev. 2008. Vertex cover might be hard to approximate to within $2-\varepsilon$. Journal of Computer and System Sciences 74, 3 (May 2008), 335–349. DOI:https://doi.org/10.1016/j.jcss.2007.06.019

[GW95] Michel X. Goemans and David P. Williamson. 1995. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM (JACM) 42, 6 (1995), 1115–1145.

[Cai03] Jin-Yi Cai. 2003. Goemans-Williamson MAXCUT Approximation Algorithm. Retrieved from http://pages.cs.wisc.edu/~jyc/02-810notes/lecture20.pdf

[Håstad01] Johan Håstad. 2001. Some optimal inapproximability results. Journal of the ACM (JACM) 48, 4 (2001), 798–859.