APPROXIMATE CONSTRAINT SATISFACTION REQUIRES LARGE LP RELAXATIONS

SIU ON CHAN, JAMES R. LEE, PRASAD RAGHAVENDRA, AND DAVID STEURER

Presented by: Marek Eliáš

Definition (basic definitions)**.** $f\colon \{\pm 1\}^n \to \mathbb{R}$, we write $\mathbb{E}f = 2^{-n}\sum_x \in \{\pm 1\}^n f(x)$. $\langle f, g\rangle = \mathbb{E}[fg]$.
Any such f can be written uniquely in the Fourier basis as $f = \sum_{\alpha\subseteq[n]}\langle f, \chi_\alpha\rangle\chi_\alpha$, where $\chi_\alpha = \prod_{i\in\alpha} x_i$.

Definition (d-junta)**.** $f\colon \{\pm 1\}^n \to \mathbb{R}$ is called d-junta for $d \in [n]$ if f depends only a subset $S \subseteq [n]$ of coordinates with $|S| \leq d$. In other words, f can be written as $f = \sum_{\alpha\subseteq S}\langle f, \chi_\alpha\rangle\chi_\alpha$.

Definition (density)**.** We say that f is a density if it is non-negative and satisfies $\mathbb{E}f = 1$. For such an f, we let μ_f denote the corresponding probability measure on $\{\pm 1\}^n$. Observer that for any $g\colon \{\pm 1\}^n \to \mathbb{R}$, we have $\mathbb{E}_{x\sim\mu_f}[g(x)] = \langle f, g\rangle$.

Definition ($(c, s)-$*approx.*)**.** We say that a linear programming relaxation $\mathcal{L}$ for MAX-Π_n achieves a (c, s)-approximation if $\mathcal{L}(\mathcal{I}) \leq c$ for all instances $\mathcal{I} \in$ MAX-Π_n with $\mathrm{opt}(\mathcal{I}) \leq s$.

Theorem (2.2)**.** *There exists an LP relaxation of size R that achieves a (c, s)-approximation for* MAX-Π_n *if and only if there exist non-negative functions $q_1, \ldots, q_R\colon \{\pm 1\}^n \to \mathbb{R}_{\geq 0}$ such that for every instance $\mathcal{I} \in$ MAX-Π_n with opt$(\mathcal{I}) \leq s$, the function $c - \mathcal{I}$ is a nonnegative combination of $q_1, \ldots, q_R$, i.e.*

$$c - \mathcal{I} \in \{\sum_i \lambda_i q_i \mid \lambda_i \geq 0\}.$$

Lemma (2.3)**.** *In order to show that $(c, s) -$ MAX-Π_n requires LP relaxations of size greater than R, it is sufficient to prove the following: For every collection of densities $q_1, \ldots, q_R\colon \{\pm 1\}^n \to \mathbb{R}_{\geq 0}$, there is $\epsilon > 0$, a function $H : \{\pm 1\}^n \to \mathbb{R}$ and a* MAX-Π_n *instance $\mathcal{I}$ such that*
1. $\langle H, c - \mathcal{I}\rangle < -\epsilon$
2. $\langle H, q_i\rangle \geq -\epsilon$

Lemma (2.4)**.** *Suppose that $f\colon \{\pm 1\}^n \to \mathbb{R}$ depends only on a subset of at most d coordinates $S \subseteq [n]$, then*

$$\langle H, f\rangle = \mathbb{E}_{x\sim\mu_S}[f(x)]$$

for some probability measure μ_S on $\{\pm 1\}^n$.

Theorem (Main, 3.1)**.** *Fix a positive number $d \in \mathbb{N}$. Suppose that the d-round Sherali-Adams relaxation cannot achieve a (c, s)-approximation for* MAX-Π_n *for every n. Then no sequence of LP relaxations of size at most $n^{d/2}$ can achieve a (c, s)-approximation for* MAX-Π_n *for every n.*

Theorem (3.2)**.** *Consider a function $f\colon \mathbb{N} \to \mathbb{N}$. Suppose that the $f(n)$-round Sherali-Adams relaxation cannot achieve a (c, s)-approximation for* MAX-Π_n*. Then for all sufficiently large n, no LP relaxation of size at most $n^{f(n^2)}$ can achieve a (c, s)-approximation for* MAX-Π_N *where $N \leq n^{10f(n)}$.*

Lemma (3.3)**.** *For all $1 \leq d$, $t \leq n$ and $\beta > 0$, the following holds. If μ has entropy $\geq n - t$, there exists a set $J \subseteq [n]$ of at most $\frac{td}{\beta}$ coordinates such that for all subsets $A \nsubseteq J$ with $|A| \leq d$, we have*

$$\max_{v\in A} H(X_v \mid X_{A\setminus v}) \geq 1 - \beta$$

Date: 2014-03-27.

Definition (KL-divergence)**.**

$$D(\mu||\nu) = \mathbb{E}_\mu[\log_2 \frac{\mu(x)}{\nu(x)}].$$

Lemma (3.5)**.** *Let μ be a distribution as in the statement of Lemma 3.3, and let $J \subseteq [n]$ be the corresponding set of coordinates. If $A \subseteq [n]$ satisfies $|A| \le d$ and $A \nsubseteq J$, then*

$$|\mathbb{E}_\mu[\chi_A(x)]| \le \sqrt{(\ln 4)\beta}.$$

Lemma (3.7, random restrictions)**.** *For any $d \in \mathbb{N}$, the following holds. Let Q be a collection of densities $q\colon \{\pm 1\}^n \to \mathbb{R}_{\ge 0}$ such that the corresponding measures μ_q have entropy at least $n - t$. If $|Q| \le n^{d/2}$, then for all integers m with $3 \le m \le n/4$, there exists a set $S \subseteq [n]$ such that:*
1. $|S| = m$
2. For each $q \in Q$, there is a set of at most d coordinates $J(q) \subseteq S$ such that under the distribution μ_q, all d-wise correlations in $S - J(q)$ are small. Quantitatively, we have

$$|\hat{q}(\alpha)| \le \left(\frac{32mtd}{\sqrt{n}}\right)^{1/2} \quad \forall\, \alpha \subseteq S, \alpha \nsubseteq J(q), |\alpha| \le d$$