Assignment 2

Technion 236646

Yuval Filmus

August 10, 2022

An instance of *MAX-3XOR* consists of constraints C_i of the form $x_i \oplus x_j \oplus x_k = b$. Each constraint is accompanied by a rational positive weight w_i. The *value* of the instance is

$$\max_x \frac{w(x)}{\sum_i w_i},$$

where x goes over all truth assignments, and $w(x)$ is the total weight of constraints satisfied by x.

Unweighted MAX-3XOR is the special case in which $w_i = 1$ for all i.

Tripartite MAX-3XOR is the special case in which the variables are partitioned into three parts, and each constraint involves exactly one variable from each part.

In class we showed that for every constant $\epsilon > 0$, it is NP-hard to distinguish instances of MAX-3XOR whose value is at least $1 - \epsilon$ from those whose value is at most $1/2 + \epsilon$.

Question 1 Prove that for every constant $\epsilon > 0$, it is NP-hard to distinguish instances of *unweighted* MAX-3XOR whose value is at least $1 - \epsilon$ from those whose value is at most $1/2 + \epsilon$.

Answer For every constant $\epsilon > 0$ we will give a polynomial time reduction ϕ_ϵ from MAX-3XOR to unweighted MAX-3XOR such that

$$|\mathsf{val}(I) - \mathsf{val}(\phi_\epsilon(I))| \leq \epsilon.$$

Assuming this, given constant $\epsilon > 0$, we have shown in class that it is NP-hard to distinguish instances of MAX-3XOR whose value is at least $1 - \epsilon/2$ from those whose value is at most $1/2 + \epsilon/2$. It remains to compose the corresponding reduction with $\phi_{\epsilon/2}$.

Let us now construct ϕ_ϵ. Suppose that we are given an instance I consisting of m constraints $C_1, \ldots, C_m$ with weights $w_1, \ldots, w_m$. We can normalize the weights so that $\sum_{i=1}^m w_i = 1$. Let $N = \lceil m/\epsilon \rceil$. We let $\phi_\epsilon(I)$ consist of each constraint C_i repeated $W_i \in \{\lfloor N w_i \rfloor, \lceil N w_i \rceil\}$ times, where the floors and ceilings are chosen so that $\sum_{i=1}^m W_i = N$ (this is possible since $\sum_{i=1}^m \lfloor N w_i \rfloor \leq N \leq \sum_{i=1}^m \lceil N w_i \rceil$). Since $|W_i - N w_i| < 1$, we have $|w_{\phi_\epsilon(I)}(x) - N w_I(x)| < m$ for all x, and so $|\mathsf{val}(\phi_\epsilon(I)) - \mathsf{val}(I)| < m/N \leq \epsilon$.

Question 2 Prove that for every constant $\epsilon > 0$, it is NP-hard to distinguish instances of *tripartite* MAX-3XOR whose value is at least $1 - \epsilon$ from those whose value is at most $1/2 + \epsilon$.

Answer We closely follow the proof in Section 3.6 of the lecture notes.

Given an instance $\Pi = (U, V, E, \{\pi_e\})$ of (Σ, Δ)-Label Cover, we construct a MAX-3XOR instance with the following variables: for each $u \in U$, we have variables encoding an arbitrary function $f'_u \colon \{\pm 1\}^\Sigma \to \{\pm 1\}$ and an odd function $f''_u \colon \{\pm 1\}^\Sigma \to \{\pm 1\}$, and for each $v \in V$, we have variables encoding an odd function $f_v \colon \{\pm 1\}^\Delta \to \{\pm 1\}$. We generate the constraints by sampling u, v, x, y, z as in Section 3.6, and checking whether

$$f'_u(x) f_v(y) f''_u(x(y \circ \pi_{uv})z) = 1.$$

Note that this is indeed a tripartite constraint.

The definition of a good edge stays the same. If an edge (u, v) is good then the same argument that derives (1) in Section 3.6 implies the following inequality:

$$\sum_{S \subseteq [\Sigma]} (1 - 2\delta)^{|S|} \hat{f}'_u(S) \hat{f}''_u(S) \hat{f}_v(\pi^2_{uv}(S)) \geq \epsilon.$$

Applying the Cauchy–Schwarz inequality, this implies that

$$\begin{aligned} \epsilon &\leq \sqrt{\sum_S \hat{f}'_u(S)^2} \sqrt{\sum_S |S|(1 - 2\delta)^{2|S|} \hat{f}''_u(S)^2 \hat{f}_v(\pi^2_{uv}(S))} \\ &\leq \sqrt{\sum_S |S|(1 - 2\delta)^{2|S|} \hat{f}''_u(S)^2 \hat{f}_v(\pi^2_{uv}(S))}. \end{aligned}$$

For every $u \in U$, we choose $c(u)$ as follows: choose S from the spectral sample of f''_u, and then choose a random element $c(u) \in S$ (note that S is non-empty since f''_u is odd). We color elements of V analogously using f_v. As in Section 3.6, the probability that $\pi_{uv}(c(u)) = c(v)$ is at least

$$\sum_S \frac{1}{|S|} \hat{f}''_u(S)^2 \hat{f}_v(\phi^2_{uv}(S))^2.$$

The rest of the argument in Section 3.6 goes through, replacing f_u with f''_u.

Question 3 Moshkovitz and Raz[1] proved that for every $\gamma \geq \frac{1}{n^{O(1)}}$,[2] it is NP-hard to distinguish satisfiable instances of (Σ, Δ)-Label Cover from instances whose value is at most γ, where $|\Sigma| \leq 2^{(1/\gamma)^{O(1)}}$ and $|\Delta| \leq (1/\gamma)^{O(1)}$.[3]

Prove that for some function $\epsilon(n) = o(1)$, it is NP-hard to distinguish instances of MAX-3XOR on n variables whose value is at least $1 - \epsilon$ from those whose value is at most $1/2 + \epsilon$.

Answer We closely follow the argument in Section 3.6 of the lecture notes.

Let us summarize the main points of the construction:

- Given a parameter δ and an instance Π of (Σ, Δ)-Label Cover on N vertices, where $|\Sigma|, |\Delta| \leq M$, we construct an instance Ψ of MAX-3LIN with $n \leq 2^M N$ variables. The construction is polynomial as long as 2^M is polynomial in N.
- If Π is satisfiable then Ψ has value at least $1 - \delta$, which needs to be at least $1 - \epsilon$.
- If Ψ has value more than $1/2 + \epsilon$ then Π has value at least $2e\delta\epsilon^3$, which needs to exceed γ.

It suffices to require $\delta \leq \epsilon$ and $\delta\epsilon^3 \geq \gamma$, that is,

$$\frac{\gamma}{\epsilon^3} \leq \delta \leq \epsilon.$$

Such a choice is possible as long as $\epsilon \geq \gamma^{1/4}$.

Since $M = 2^{(1/\gamma)^{O(1)}}$ and 2^M needs to be polynomial in N, we need to choose $\gamma = 1/(\log\log N)^{\Omega(1)} = 1/(\log\log n)^{\Omega(1)}$. This means that we can choose

$$\epsilon = \frac{1}{(\log\log n)^{\Omega(1)}}.$$

Moshkovitz and Raz conjecture that the size of the alphabet Σ can be reduced to $(1/\gamma)^{O(1)}$, in which case we would be able to choose $\epsilon = \frac{1}{(\log n)^{\Omega(1)}}$. The linear Projection Games Conjecture would improve this to $\epsilon = \frac{1}{n^{\Omega(1)}}$.

[1] Dana Moshkovitz and Raz Raz, Two-query PCP with subconstant error, Journal of the ACM, vol. 57(5), 2010, pp. 1–29 (with a 100 page appendix).

[2] That is, γ needs to satisfy $\gamma \geq 1/n^C$ for some constant $C > 0$ of our choice.

[3] In contrast, applying parallel repetition to the PCP theorem, we obtain that for every *constant* $\gamma > 0$, it is NP-hard to distinguish satisfiable instances of (Σ, Δ)-Label Cover from instances whose value is at most γ, where $|\Sigma|, |\Delta| \leq (1/\gamma)^{O(1)}$. We can only take γ to be constant since parallel repetition blows up an instance of size n to one of size $n^{\Theta(\log(1/\gamma))}$. In contrast, the Moshkovitz–Raz construction blows up the instance to one of size $n^{1+o(1)}(1/\gamma)^{O(1)}$.