Assignment 2

Technion 236646

Yuval Filmus

May 15, 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 $1/2 + \epsilon$.

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 $1/2 + \epsilon$.

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 $1/2 + \epsilon$.

---

[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 $n^{1+o(1)}(1/\gamma)^{O(1)}$.