Abstract

Many known optimal NP-hardness of approximation results are reductions from a problem called Label-Cover. The input is a bipartite graph $G = (L, R, E)$ and each edge $e = (x, y) \in E$ carries a projection π_e that maps labels to x to labels to y. The objective is to find a labeling of the vertices that satisfies as many of the projections as possible. It is believed that the best approximation ratio efficiently achievable for Label-Cover is of the form N^{-c} where $N = nk$, n is the number of vertices, k is the number of labels, and $0 < c < 1$ is *some* constant. Inspired by a framework originally developed for Densest k-Subgraph, we propose a "log density threshold" for the approximability of Label-Cover. Specifically, we suggest the possibility that the Label-Cover approximation problem undergoes a computational phase transition at the same threshold at which local algorithms for its random counterpart fail. This threshold is $N^{3-2\sqrt{2}} \approx N^{-0.17}$. We then design, for any $\varepsilon > 0$, a polynomial-time approximation algorithm for *semi-random* Label-Cover whose approximation ratio is $N^{3-2\sqrt{2}+\varepsilon}$. In our semi-random model, the input graph is random (or even just expanding), and the projections on the edges are arbitrary. For *worst-case* Label-Cover we show a polynomial-time algorithm whose approximation ratio is roughly $N^{-0.233}$. The previous best efficient approximation ratio was $N^{-0.25}$. We present some evidence towards an N^{-c} threshold by constructing integrality gaps for $N^{\Omega(1)}$ rounds of the Sum-of-squares/Lasserre hierarchy of the natural relaxation of Label Cover. For general 2CSP the "log density threshold" is $N^{-0.25}$, and we give a polynomial-time algorithm in the semi-random model whose approximation ratio is $N^{-0.25+\varepsilon}$ for any $\varepsilon > 0$.