Lecture 7: Coloring 3-colorable graphs

Uriel Feige

February 8

Let G be a 3-colorable graph on n vertices. In this lecture we design algorithms for approximate coloring, in the sense that they do legally color G, but use more than 3-colors. We remark that it is known that coloring 3-colorable graphs with 4 colors is NP-hard.

Every graph of maximum degree Δ can be colored by $\Delta + 1$ colors, using inductive coloring. An early approximate coloring algorithm of Wigderson is based on the observation that the neighborhood of a vertex in a 3-colorable graph is bipartite, and hence can efficiently be colored by at most two colors. Hence as long as the graph has a vertex of degree at least $\sqrt{n}$, color and remove its neighborhood, and when no such vertex remains, apply inductive coloring. This uses $O(\sqrt{n})$ vertices.

Karger, Motwani and Sudan [1] show that if Δ is large and the graph happens to be colorable with much fewer colors (say, 3), then there is a polynomial time algorithm that colors it with $o(\Delta)$ colors. Their paper is very readable, and here we sketch only some of the details as specialized to 3-colorable graphs.

Recall the SDP for the theta function as presented in an earlier lecture. We show how it can be used to obtain approximation algorithms for coloring.

For a 3-colorable graph with n vertices and m edges, using the SDP we obtain a unit vector solution in which for every edge (i, j), $\langle v_i, v_j \rangle = -1/2$. (In fact, as SDPs are not guaranteed to be solved exactly, one would get a solution with $\langle v_i, v_j \rangle \le -\frac{1}{2} + \epsilon$, where ϵ can be made arbitrarily small. For sufficiently small ϵ, say $\epsilon = \frac{1}{n^2}$, the analysis goes through with only negligible effect on the number of colors used.)

Consider a representation of R^n in the standard orthonormal basis $(e_1, \ldots, e_n)$. To round the vector solution of the SDP, pick a random vector r in which the value of each coordinate of r is chosen independently at random, distributed according to the standard normal distribution $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ with mean 0 and variance 1. This is known to give a spherically symmetric distribution where each direction is also distributed according to the normal distribution. (Moreover, orthogonal directions have independent distributions. This last fact is not needed for the analysis of the approximation ratio.)

Here is an intuitive argument showing that the distribution is spherically symmetric. By the central limit theorem, a standard normal random variable can be thought of as the sum on N^2 independent auxiliary random variables of values $\pm\frac{1}{N}$, for some very large N. Multiplying the random variable by α is essentially equivalent to multiplying the number of auxiliary random variables by α^2. Adding two normal random variables gives a new normal random variable based on the sum of the two sets of auxiliary random variables. Representing an arbitrary unit vector $x = \sum_{i=1}^{n} \alpha_i e_i$ in the standard basis, the coefficients satisfy $\sum(\alpha_i)^2 = 1$, hence the component of r in direction x is also distributed like the sum of N^2 auxiliary random variables.

In algebraic form, spherical symmetry follows by observing that the density function at a point $x = \sum \alpha_i e_i$ is $\phi(x) = \frac{1}{\sqrt{2\pi}} \prod_{i=1}^{n} e^{-\alpha_i^2/2} = \frac{1}{\sqrt{2\pi}} e^{-\sum_{i=1}^{n} \alpha_i^2/2} = \frac{1}{\sqrt{2\pi}} e^{-|x|^2/2}$, where $|X|$ is the Euclidean norm of x. Hence $\phi(x)$ depends only on the norm of x and not on the direction of x, implying spherical symmetry.

Let $N(x) = \int_x^\infty \phi(y) d(y)$ denote the tail of the normal distribution. It is known that

$$\phi(x)(\frac{1}{x} - \frac{1}{x^3}) \leq N(x) \leq \phi(x)\frac{1}{x}$$

(We do not prove the above inequality, but remark that for $x \geq 2$ it is easy to see the $N(x) = \Theta(\frac{\phi(x)}{x})$. The lower bound follows from $N(x) \geq \frac{1}{x}\phi(x + \frac{1}{x})$. The upper follows because starting at $y = x$ and advancing in steps of size $1/x$, the value of $\phi(y)$ drops at least in an exponential rate, and hence the integral can be upper bounded by a geometric sum.)

Let c be a value to be optimized later. The rounding technique first selects a set S of vertices that contains all vertices i with $\langle v_i, r\rangle \geq c$, and then removes one endpoint of every edge induced by S so as to remain with an independent set. The analysis of this rounding technique is based on linearity of expectation (a principle that applies even if random variables are not independent).

The expected number of vertices in S is $nN(c)$. To compute the expected number of edges in S, let (i, j) be an edge. We need $\langle v_i, r\rangle \geq c$ and $\langle v_j, r\rangle \geq c$ implying $\langle v_i + v_j, r\rangle \geq 2c$. Note that $v_i + v_j$ is also a unit vector because $(v_i + v_i)^2 = 1 + 2(-\frac{1}{2}) + 1 = 1$. Hence the expected number of edges is $mN(2c)$. We shall pick the smallest c such that $nN(c) - mN(2c) \geq nN(c)/2$. (Then, in expectation removing one endpoint of every edge is S, at least half the expected vertices remain and form an independent set.) Let $d = 2m/n$ denote the average degree. Hence we need $n\phi(c)(\frac{1}{c} - \frac{1}{c^3}) \geq dn\phi(2c)\frac{1}{2c}$. When d is a sufficiently large constant (and then we will have $\frac{1}{c^3} \leq \frac{1}{2c}$), this simplifies to a sufficient condition of $e^{3c^2/2} \geq d$. Hence we shall choose $c = \sqrt{\frac{2\ln d}{3}}$. At this point the expected number of vertices in S is $nN(c) \geq \frac{n}{2c\sqrt{2\pi}} e^{-\frac{\ln d}{3}} \geq \Omega(\frac{n}{d^{1/3}\sqrt{\log d}})$. By repeatedly removing independent sets in this manner, we obtain a coloring with $O(\Delta^{1/3}\sqrt{\log \Delta} \log n)$ colors. (The maximum degree Δ serves as a convenient upper bound on all average degrees d encountered in the iterations of the algorithm. A tighter parameter than Δ is the average degree of the vertex induced subgraph of maximum average degree. This parameter can easily be approximated within a ratio of 2 by a greedy algorithm that repeatedly removes vertices of lowest degree, and can also be computed exactly in polynomial time using flow techniques.)

The above rounding technique is randomized. Our analysis provides a lower bound on the expected size of the independent set that is found in a single iteration. As always in such situations, one can transform this to a high probability of success, by repeating each iteration many times with independent choices of random normal vectors r, and taking the largest independent set that is found. Likewise, often randomized algorithm whose analysis is based on expectations can be derandomized (e.g., through the method of conditional expectations), but this topic is beyond the scope of this course.

There is no approximation known for 3-coloring that uses fewer than $d^{1/3}$ colors (where $d = \frac{2m}{n}$ is the average degree of the graph). As a function of n, there have been various improvements to the approximation ratio. For example, by using the observation by Wigderson referred to above (for vertices of degree above $n^{3/4}$), one can color 3-colorable graph by $n^{1/4}$ colors. Further improvements are known, to slightly below $n^{1/5}$.

Adapting the KMS analysis to k-colorable graphs gives a coloring using $\tilde{O}(\Delta^{1-\frac{2}{k}})$ colors.

References

[1] David R. Karger, Rajeev Motwani, Madhu Sudan: Approximate Graph Coloring by Semidefinite Programming. J. ACM 45(2): 246-265 (1998).