Phase Transitions in k-Colorability of Random Graphs

Cris Moore
Santa Fe Institute
University of New Mexico

Anna Olson
Carnegie Mellon University

August 15, 2008

Abstract

We used the first- and second-moment probabilistic and the Potts spin-glass model to find improved upper and lower bounds on the critical average degree d_k governing $k-$colorability of random large graphs, and found numerical results regarding upper and lower bounds, d_k^+ and d_k^- respectively, such that for large k, $d_k^- \approx d_k^+ - 1$. Moreover, it seems that d_k^+ and d_k^- differ from the naive upper bound, $\frac{2 \ln k}{\ln k - \ln k - 1}$ by some small constant.

1 Introduction

The question of whether a given graph is $k-$colorable is known to be in NP, and as such, we know of no polynomial-time algorithm to answer this question, and if $P \neq NP$, there is no such algorithm. However, one can use probabilistic methods to shed light on the likelihood that a very large random graph is $k-$colorable. It has been suggested that there is a real number d_k such that if the average degree d of a graph G is less than d_k, then as $n \to \infty$, the probability of G being $k-$colorable approaches 1, and if $d > d_k$, then this probability approaches 0. Similar behavior has been observed in $k-$SAT. Our goal was to find bounds d_k^- and d_k^+ on d_k, such that $d_k^- \leq d_k \leq d_k^+$ for general k.

In this research, we used the Potts spin-glass model. Conventionally, when studying $k-$colorability of graphs, we check to see if there is a coloring of a graph such that every edge connects vertices of differing colors (given that each vertex is colored with one of k colors) - in other terms, if we have vertices u, v are connected, and $c(u), c(v)$ are the colors of u, v respectively, then $c(u) \neq c(v)$ for any connected pair u, v in a proper coloring. In the spin-glass model, each edge (u, v) is associated with a random permutation $\pi_{(u,v)} \in S_k$, such that $\pi_{(u,v)}(c(u)) \neq c(v)$, and $\pi^{-1}_{(u,v)}(c(v)) \neq c(u)$.

The Potts spin-glass model has the advantage of symmetry and easy calculations, and it is conjectured that a critical degree d_k is valid both in the spin-glass model and in the strict chromatic model. The ease of calculation with the spin-glass model becomes apparent when one considers a small cycle - the probability of each edge being valid in a $k-$coloring is not symmetric in the strict chromatic model, but is symmetric in the spin-glass model. As small cycles are very rare and thus negligible, we assumed that the spin-glass model was valid for our purposes.

2 Lower Bound

We first take a moment to calculate the naive upper bound, using the random variable X =number of colorings. Each edge has probability $(1 - \frac{1}{k})$ of being compliant, so $E[X] = (k(1 - \frac{1}{k})^{\frac{d}{2}})^n$. When $E[X] = 0$, we have a value of d, noted d_k^{++}, such that $d_k^{++} \geq d_k$. $E[X] = 0$ for $k(1 - \frac{1}{k})^{\frac{d}{2}} < 1$, as $n \to \infty$. d_k^{++} occurs when $k(1 - \frac{1}{k})^{\frac{d}{2}} = 1$, and, after some algebra, is found:

$$d_k^{++} = \frac{2 \ln k}{\ln k - \ln (k-1)} \tag{2.1}$$

This is the naive upper bound, and may be approximated as $2k \ln k - \ln k$.

To determine the lower bound on d_k, referred to as d_k^-, we used the second moment method. We know that for any random variable W, $Pr[W > 0] \geq \frac{E[W]^2}{E[W^2]}$, so if we concern ourselves with X as defined for the naive upper bound, remember that X is a function of d, k, then d_k^- is the the highest value of d such that for $C = \frac{E[X]^2}{E[X^2]}$, $C > 0$ in the limit $n \to \infty$.

Let a particular graph G be defined using the Erdős-Rényi $G(n, m)$ model. We can then define $E[X^2]$ as the sum over all pairs of colorings c_1, c_2 of the probability raised to the m^{th} power that if an edge (u, v) is chosen randomly, then it is compliant in both c_1 and c_2. We can define such a probability in terms of a, where a is the number of nodes on which c_1, c_2 differ. Assuming every coloring is possible (though not necessarily compliant!), there are $k^n \binom{n}{a} (k-1)^a$ pairs of colorings c_1, c_2 that differ on a nodes. This gives:

$$E[X^2] = \sum_{a=0}^{n} k^n \binom{n}{a} (k-1)^a f(a)^m \tag{2.2}$$

We now require an expression for $f(a)$. Let c refer to the set of $n-a$ nodes on which c_1, c_2 are the same, and let $\bar{c}$ refer to the set of a nodes on which c_1, c_2 differ. Let (u, v) be an edge in G, and let $P_{comp}[(u, v)]$ be defined:

$$P_{comp}[(u,v)] = \begin{cases} 1 - \frac{1}{k} & \text{if both } u, v \in c \\ 1 - \frac{2}{k} + \frac{1}{k(k-1)} & \text{if both } u, v \in \bar{c} \\ 1 - \frac{2}{k} & \text{if either } u \in c \text{ and } v \in \bar{c}, \text{ or } u \in \bar{c} \text{ and } v \in c \end{cases} \tag{2.3}$$

We get $P_{comp}[(u, v)]$ using principles of probability. It is important to note that this calculation is far easier in the Potts spin-glass model, as the strict chromatic model requires a matrix of overlaps, instead of a simple Hamming distance. For the first case, as the vertices have the same color in both colorings, we simply have the probability that a random edge with randomly colored vertices is $k-$compliant. For the second, we recognize the probability that (u, v) is valid in both colorings is $(1 - Pr[(u, v)$ isn't valid in both colorings$])$. This probability can be written as $Pr[(u, v)$ invalid in $c_1]+Pr[(u, v)$ invalid in $c_2]-Pr[(u, v)$invalid in both$]$, or $\frac{2}{k} - \frac{1}{k(k-1)}$. For the third, assume w.l.o.g. that $u \in c$ and $v \in \bar{c}$, and that u is colored with color i. For (u, v) to be compliant in both c_1, c_2, we must have $\pi(i) \neq j_1$ and $\pi(i) \neq j_2$ (c_1, c_2 respectively), and there are $(k-1)(k-2)$ ways of doing this, out of $k(k-1)$ ways of coloring u, v such that $u \in c$ and $v \in \bar{c}$. This gives the probability $\frac{k-2}{k}$, or $1 - \frac{2}{k}$.

Let edges be chosen with replacement, and directed, such that there are n^2 possibilities for edges. The probability that both $u, v \in c$ is $(n-a)^2$, the probability that both $u, v \in \bar{c}$ is a^2, and the probability that only one of u, v is in c is $2a(n-a)$. This gives us an expression for $f(a)$, after some algebra:

$$f(a) = \frac{a^2k - 2a(k-1)n + (k-1)^2n^2}{k(k-1)n^2} \tag{2.4}$$

Of course, as we are examining the behavior of $Pr[G$ is $k-$colorable$]$ as $n \to \infty$, we can define $\alpha = \frac{a}{n}$ and integrate instead of summing. This requires the use of a binomial approximation, where $\binom{n}{\alpha n} \approx \frac{e^{h(\alpha)}}{\sqrt{2\pi\alpha(1-\alpha)n}}$, and $h(\alpha) = -\alpha \ln \alpha - (1-\alpha) \ln (1-\alpha)$. We know also that $m = \frac{dn}{2}$. Using (2) and (4), and performing some algebra, we get:

$$E[X^2] = n \int_{\alpha=0}^{1} \frac{1}{\sqrt{2\pi\alpha(1-\alpha)n}} \left(e^{h(\alpha)} k(k-1)^{\alpha} \left(\frac{\alpha^2 k - 2\alpha(k-1) + (k-1)^2}{k(k-1)} \right)^{\frac{d}{2}} \right)^{n} d\alpha \tag{2.5}$$

(Note that the term $d\alpha$ at the end refers to the fact that we are integrating over α, and that it is not an extra $d \cdot \alpha$ term.)

We see that $E[X]$ has the form f_1^n, and $E[X^2]$ has form f_2^n (by the Laplace method - really, $E[X^2] = \int_0^1 g(\alpha)^n d\alpha$), we want $f_2 \sim f_1^2$, and thus $g_{max} \sim f_2 \sim f_1^2$. g_{max} occurs at $\alpha = 1 - \frac{1}{k}$, and the lowest d for which $g_{max} \sim f_2 \sim f_1^2$ is the d_k^- we want. After some algebra, we find that:

$$d_k^- \leq \frac{2\ln\left((1-\alpha)^{1-\alpha}\alpha^{\alpha}k(k-1)^{-\alpha}\right)}{\ln\left(k(k-1)^{-3}(\alpha^2 k - 2\alpha(k-1) + (k-1)^2\right)} \tag{2.6}$$

We found d_k^- numerically. It appears that $\lim_{k\to\infty} d_k^- = d_k^{++} - 2$, although for $k \geq 9$ it also appears that $d_k^- \leq d_k^{++} - 2$. As seen in Figures 1 and 2, the difference between d_k^- and d_k^{++} is increasingly small, as to be seen as negligible past a certain k, depending on the application.

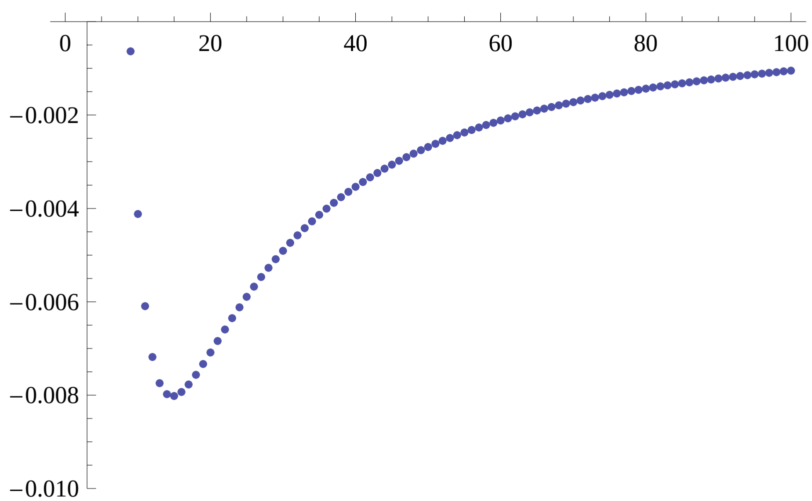

Figure 1: A plot of $d_k^- - (d_k^{++} - 2)$ as a function of k, showing that the greatest difference in values, -0.008..., occurs at $k = 15$

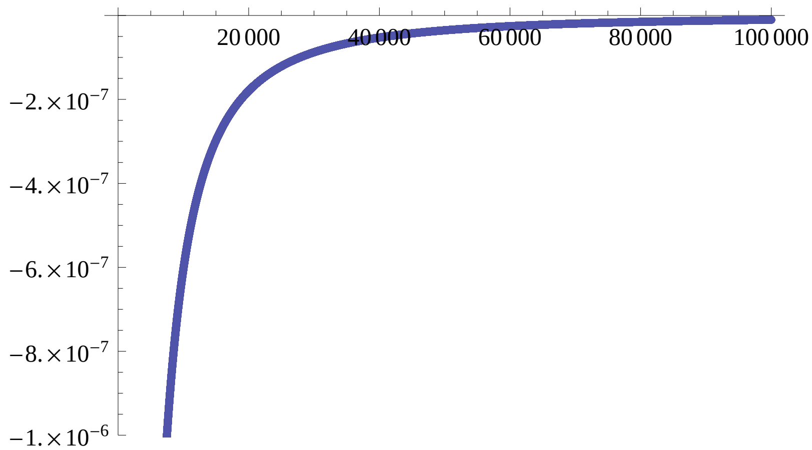

Figure 2: A plot of $d_k^- - (d_k^{++} - 2)$ as a function of k for large values of k

3 Upper Bound

It is known that bounds may be improved by calculating the expected value of some other random variable that doesn't necessarily count the number of solutions, but has some characteristics that show whether a solution exists. Such random variables might have other desirable characteristics such as symmetry, which makes a problem easier to solve.

An example of this is in the problem of SAT. Suppose we have a set of solutions S, such that $S \subseteq \{0,1\}^n$, where there are n variables, and 0 corresponds to false and 1 to true. S may be viewed as a subset of points on an $n-$dimensional hypercube (H), where each edge has two nodes (0 and 1). For two points σ, τ in H, $d(\sigma, \tau)$ is defined as the distance between σ and τ. We have a random variable Y:

$$Y = \sum_{\sigma \in S} 2^{-|\{\tau \in S \,:\, d(\sigma,\tau)=1\}|} \tag{3.1}$$

It is known that $Y \geq 1$ iff $S \neq \emptyset$.

To represent the set of all possible colorings (valid or not) of a graph $G(n,m)$, we use a similar hypercube with n dimensions, and k nodes on each edge. Each dimension corresponds to a single vertex in G. The set of valid colorings is therefore $S \subseteq \{1, 2, ..., k\}^n = [k]^n$. Instead of simply expanding Y to allow for k possible assignments to each variable/vertex, we define a "weight" for each σ, and define Z:

$$Z = \sum_{\sigma \in S} \frac{1}{w_\sigma} \tag{3.2}$$

where we define w_σ:

$$w_\sigma = \prod_{i=1}^{n} c(v_i) \tag{3.3}$$

Here, $c(v_i)$ is the number of colorings that exist for vertex v_i, such that G still has a valid coloring. If v_i has no neighbors, then $c(v_i) = k$, by definition. If v_i has $\geq k-1$ neighbors, such that each of the neighbors blocks a different color, then $c(v_i) = 1$.

Like Y, $Z \geq 1$ iff $S \neq \emptyset$. We prove this using induction on n. Assume that $Z \geq 1$ iff $S \neq \emptyset$ is true for n, and we examine a hypercube with $n+1$ dimensions, with $S \subseteq \{0,1\}^{n+1}$. We wish to show that by taking all the members σ of S with $\sigma_i = 1$ and copying them across all planes in the i^{th} dimension, unioning each copy with extant members of S in the given plane, we necessarily decrease the value of Z until $S = \{0,1\}^{n+1}$, and $Z = 1$. We start by assuming this "twinning" process has been performed for all values of i up to n. Let v be the set of σ for which $\sigma_{n+1} \neq 1$, let v' be the set of σ for which $\sigma_{n+1} = 1$ and at least one neighbor τ exists with $\tau_{n+1} \neq 1$, and let $\overline{v'}$ be the set of σ for which $\sigma_{n+1} = 1$ and no neighbors τ exist with $\tau_{n+1} \neq 1$. We get a formula for Z:

$$\begin{aligned} Z &= weight[v] + weight[v'] + weight[\overline{v'}] \\ &= weight[v] + weight[v'] + \frac{1}{k^n}|\overline{v'}| \\ &= weight[v] + weight[v'] + \frac{k^n - |v'|}{k^n} \end{aligned}$$

For each element in v', there are between 1 and $k-1$ elements in v that share the same coordinates save for the $(n+1)^t h$. It is therefore clear that the only way for the weight of v to not make up for the weight "loss" caused by v' is for $|v| > (k-1)(k^n)$, which is impossible in this circumstance. Thus, the process of twinning only reduces the value of Z until all of the hypercube is filled, and $Z \geq 1$.

But what is the probability that exactly j colors are available for a given vertex? Because we are working in the Potts spin-glass model, each edge can block any color with equal probability. Suppose we associate with each vertex a set of $k-1$ bins, into which balls may be thrown - if a bin is empty, then the color associated with that bin is allowed. Assuming that the all reds (or all 1's) assignment is valid (which is true with probability $(1-\frac{1}{k})^{\frac{dn}{2}}$), then we think of some number of balls, b, being thrown at $k-1$ bins, and wondering if $j-1$ bins are unoccupied. Let $P_k(b,j)$ be the probability that given b edges, j out of k colors will be available (assuming all reds valid!), or that given b balls thrown into $k-1$ bins, exactly $j-1$ bins will be empty. Of course, it is equally valid to assume that any particular coloring works - the point is to assume that there is a valid coloring, and that at least one color is available to each vertex.

It is known that the number of edges that each vertex has follows a Poisson distribution (such that given average degree d, the number of vertices with degree b is $\frac{e^{-d}d^b}{b!}n + O(n^{\frac{2}{3}})$), so the probability that there are b edges is $q(b,d) = \frac{e^{-d}d^b}{b!}$. For a single vertex, we can define $x(b)$ as $E[|\{available\ colors\}|^{-1}]$, or $\sum_{j=1}^{k} \frac{P_k(b,j)}{j}$. This gives us an expression for $E[Z]$:

$$E[Z] = k^n(1-\frac{1}{k})^{\frac{dn}{2}} \left(\prod_{b=0}^{\infty} \left(\sum_{j=1}^{k} \frac{P_k(b,j)}{j} \right)^{q(b,d)} \right)^n \tag{3.4}$$

We replaced $P_k(b,j)$ with an exact expression $\frac{\binom{k}{k-j} S_{b,k-j}(k-j)!}{k^b}$, where $S_{b,k-j}$ is the Stirling number of the second kind. We also replaced the ∞ in the product term with $10d$, as the probability of $b \geq 10d$ is negligible.

We found numerical results for d_k^+ by graphing $E[Z]'$ ($E[Z]$ with $n=1$ for ease of computation) as a function of k and r, where the value of d in $E[Z]'$ is set to be $d_k^{++} - r$. The value of d at which $E[Z]' = 1$ represented d_k^+; when $n \to \infty$, if $E[Z]' < 1$, then $E[Z] \to 0$. The numerical results shown in Figure 3 suggest that $\lim_{k\to\infty} d_k^+ = d_k^{++} - 1$, although it appears that $d_k^+ \geq d_k^{++} - 1$. Because the $S_{b,k-j}$ term is so computationally expensive, we used values of k up to 50, enough to illustrate the trend.

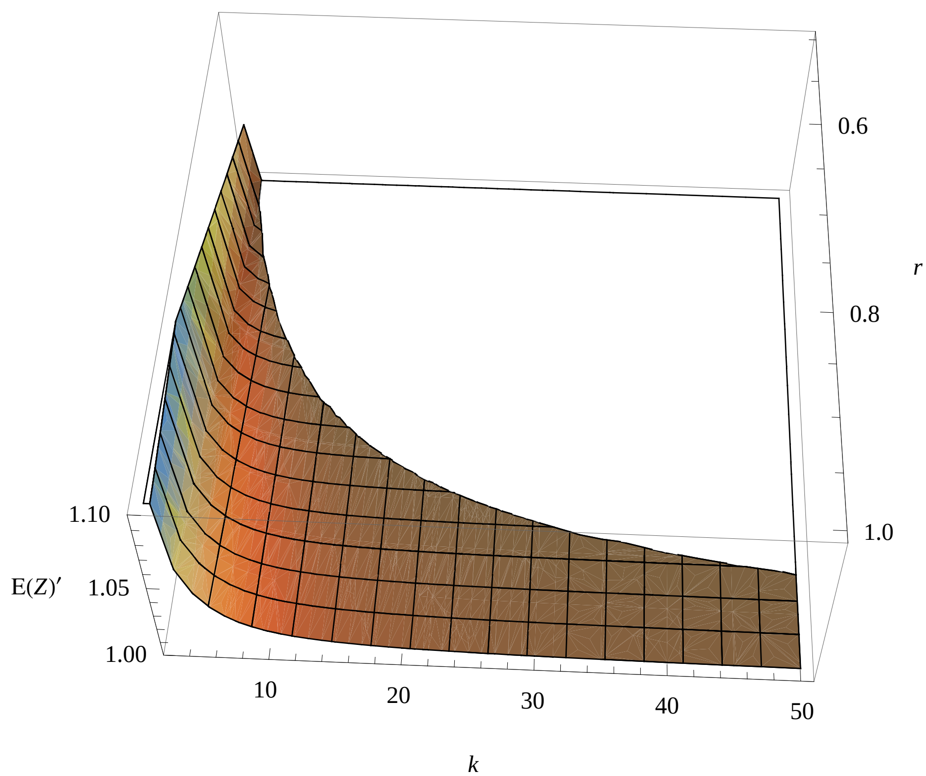

Figure 3: A 3D plot of points in the form $(k, r, E[Z]')$, showing only those points for which $E[Z]' \geq 1$. It is easily seen that $r \to 1$ as k grows, indicating $d_k^+ \to d_k^{++} - 1$.