

Almost all graphs with $2.522n$ edges are not 3-colorable

Dimitris Achlioptas ^{*}
optas@cs.toronto.edu

Michael Molloy [†]
molloy@cs.toronto.edu

Department of Computer Science
University of Toronto
Toronto, Ontario M5S 3G4, Canada.

Abstract

We prove that for $c \geq 2.522$ a random graph with n vertices and $m = cn$ edges is not 3-colorable with probability $1 - o(1)$. Similar bounds for non- k -colorability are given for $k > 3$.

1991 Mathematics Subject Classification: Primary 05C80; Secondary 05C15.

1 Introduction

Let $N(n, m, A)$ denote the number of graphs with vertices $\{1, \dots, n\}$, $m = m(n)$ edges and some property A . The term “almost all” in the title has the meaning introduced by Erdős and Rényi [5]:

$$\lim_{n \rightarrow \infty} \frac{N(n, m, A)}{\binom{\binom{n}{2}}{m}} = 1 . \quad (1)$$

Equivalently, one can consider a random graph $G = G(V, E)$ where $|V| = n$ and E is a uniformly random m -subset of all $\binom{n}{2}$ possible edges on V , i.e. the $G(n, m)$ model of random graphs. If n is an index running over probability spaces, we will say that a sequence of events \mathcal{E}_n occurs *with high probability* (w.h.p.) if $\lim_{n \rightarrow \infty} \Pr[\mathcal{E}_n] = 1$. In particular, we will say that “ $G(n, m(n))$ has property A w.h.p.” if $m(n)$ is such that (1) holds for A .

In their seminal paper introducing random graphs [5], Erdős and Rényi pointed out that a number of interesting properties exhibit a sharp threshold behavior on $G(n, m)$: for each such property A , there exists a critical number of edges $m_A(n)$ such that for m around

^{*}Research supported in part by an NSERC PGS Scholarship. Current address: Microsoft Research, One Microsoft Way, Redmond WA 98052, U.S.A. Email: optas@microsoft.com

[†]Research supported in part by an NSERC grant.

$m(n)$ the probability of $G(n, m)$ having A changes rapidly from near 0 to near 1. Such properties include having a multicyclic component, having a perfect matching, connectivity, Hamiltonicity and others.

A central property in this context is the k -colorability of $G(n, cn)$ where k is a fixed integer. For $k = 2$, this is very well-understood as bipartiteness is equivalent to containing no odd cycles. In particular, the probability of non-2-colorability is bounded away from 0 for any $c > 0$ and keeps increasing gradually with c , reaching $1 - o(1)$ during the emergence of the giant component at $c = 1/2$. For $k > 2$, though, our understanding of k -colorability is not nearly as good; moreover, the situation is conjectured to be quite different. In particular, see [5, 3], Erdős asked: for each $k > 2$, is there a constant c_k such that for any $\epsilon > 0$,

$$G(n, (c_k - \epsilon)n) \text{ is w.h.p. } k\text{-colorable} \quad \text{and} \quad G(n, (c_k + \epsilon)n) \text{ is w.h.p. not } k\text{-colorable} ? \quad (2)$$

Recently, Friedgut [6] made great progress in our understanding of threshold phenomena in random graphs by establishing necessary and sufficient conditions for a property to have a sharp threshold. Using the main theorem of [6], Friedgut and the first author [1] showed that for $k > 2$, there exists a *function* $t_k(n)$ such that (2) holds upon replacing c_k with $t_k(n)$, i.e. that indeed k -colorability has a sharp threshold. While it is widely believed that $\lim_{n \rightarrow \infty} t_k(n)$ exists, confirming this conjecture and determining the limit c_k , even for $k = 3$, seems very challenging.

Perhaps the main tool in attacking the question of k -colorability for small values of $k > 2$ has been the elementary fact that if a graph has no subgraph with minimum degree at least k , then it is k -colorable. In particular, first Łuczak [11] proved that w.h.p. $G(n, cn)$ remains 3-colorable after the emergence of the giant component by showing that for $c \leq 0.50005$, w.h.p. $G(n, cn)$ has no subgraph of minimum degree 3. Shortly afterwards, Chvátal [4] improved this greatly by showing that $G(n, cn)$ w.h.p. has no subgraph with minimum degree 3 for $c \leq 1.44$ and Reed and the second author [13] improved the bound even further to $c \leq 1.67$. Finally, Pittel, Spencer and Wormald [16], proved that, in fact, for all $k > 2$ there exists γ_k such that for $c < \gamma_k$, $G(n, cn)$ w.h.p. has no subgraph with minimum degree at least k , while for $c > \gamma_k$ it has such a subgraph w.h.p. Moreover, they determined γ_k exactly for all k , in particular yielding $c_3 \geq \gamma_3 = 1.675\dots$ Following that, and in answering a question of Bollobás [3], the second author [14] proved $c_k > \gamma_k$ for all $k \geq 4$ and conjectured $c_3 \neq \gamma_3$ as well. This conjecture was verified recently by the authors [2] after analyzing the performance of a greedy “list-coloring” heuristic on $G(n, cn)$. That argument yielded $c_3 > 1.923$, which is the best known lower bound for c_3 .

In this paper, after briefly reviewing the known upper bounds for c_k , we show how a technique of Kirousis et al. [8], developed for random k -SAT, can be used to yield an improved upper bound for c_k for small values of k . For example, we obtain

Theorem 1

$$c_3 < 2.522 \ .$$

2 The first moment method

Grimmett and McDiarmid [7] gave the first lower bound on the chromatic number of random graphs by determining α_k such that $G(n, m = \alpha_k n^2)$ w.h.p. has no independent set of size n/k , and thus $\chi(G) > k$ (here $k \rightarrow \infty$). Moreover, they conjectured that the lower bound derived by this argument is tight, and as evidence for this they showed that the expected number of k -colorings of $G(n, \alpha n^2)$ tends to infinity for $\alpha < \alpha_k$. Devroye (see [4]) later observed that when k is fixed, letting the expected number of k -colorings go to 0 as $n \rightarrow \infty$ yields much better lower bounds for the chromatic number than letting the number of (suitably large) independent sets go to 0 as $n \rightarrow \infty$.

Our proof can be viewed as a refinement of Devroye's argument which we will reproduce below to introduce some ideas and notation. Before doing so, let us recall that in the $G(n, m)$ model the edge set is a random m -subset of the set of all $\binom{n}{2}$ possible edges. Equivalently, we can say that the edges of the graph are selected from the set of all possible edges one-by-one, uniformly, independently and *without* replacement. For the calculations in this paper it will be convenient to consider a slight modification of the $G(n, m)$ model in that the selection is done *with* replacement, i.e. multiple edges are allowed. We will denote this model by $G^r(n, m)$. Intuitively, it is clear that for any monotone increasing property A and any value of m , the probability of A holding in $G(n, m)$ is no smaller than it is in $G^r(n, m)$ since "additional occurrences of an edge do not help". Formally, this is Theorem 5 in [9] and, for our purposes, it will imply that if for a given $m(n)$, $G^r(n, m(n))$ is w.h.p. non- k -colorable then so is $G(n, m(n))^*$.

We will distinguish between a proper k -coloring of a graph and one in which some adjacent vertices might have the same color by referring to them as a " k -coloring" and a " k -partition", respectively. In fact, it will be helpful to think of a k -coloring of a graph $G(V, E)$ as a k -partition of V such that every $e \in E$ has its endpoints in distinct blocks of the partition, so that each block is an independent set.

Let $P = V_1, \dots, V_k$ be an arbitrary (ordered) k -partition of V and let C_P denote the event that P is a k -coloring of G . For C_P to hold, every edge of the random graph has to connect vertices from two different blocks. Introducing

$$T(P) = \sum_{i < j} |V_i| \cdot |V_j| , \quad (3)$$

the total number of pairs of vertices belonging to different blocks, we have

$$\Pr[C_P] = \left(\frac{T(P)}{\binom{n}{2}} \right)^m . \quad (4)$$

*In fact, it turns out that since the expected number of multiple edges in $G^r(n, cn)$ is $O(1)$ the converse holds as well, i.e. if $G(n, cn)$ is w.h.p. non- k -colorable then so is $G^r(n, cn)$. Thus, by switching to the $G^r(n, m)$ model we are not giving anything away with respect to bounding c_k .

Now, using the fact $\sum_i |V_i| = n$ and the Cauchy-Schwartz inequality, respectively, we bound

$$\begin{aligned} T(P) &= \frac{n^2}{2} - \frac{1}{2} \sum_i |V_i|^2 \\ &\leq \frac{n^2}{2} - \frac{1}{2} \cdot \frac{n^2}{k} \\ &= \frac{k-1}{2k} n^2 . \end{aligned}$$

Thus, (4) yields

$$\Pr[C_P] \leq \left(\frac{k-1}{k}\right)^m \left(\frac{n}{n-1}\right)^m . \quad (5)$$

Since the number of k -partitions of V is k^n , (5) implies that the expected number of k -colorings of G , for $m = cn$, is of order

$$\left[k \left(\frac{k-1}{k}\right)^c \right]^n .$$

Hence, if $c > \frac{\ln k}{\ln k - \ln(k-1)}$ the expected number of k -colorings of $G^r(n, cn)$ tends to 0 as $n \rightarrow \infty$ implying that $G^r(n, cn)$ is w.h.p. non- k -colorable. For $k = 3$, this argument yields $c_3 < 2.71$ and in general $c_k < k \ln k$.

It is worth noting that this simple argument is asymptotically tight: the upper bound on $\chi(G(n, m))$ given by Łuczak [10] implies that for any $\epsilon > 0$ and all $k \geq k_0(\epsilon)$, $c_k > (1-\epsilon)k \ln k$. On the other hand, the following two observations can be used to show that for $k > 2$ this argument is not exact: (a) if a k -colorable graph has s_i vertices of degree i then it has at least $\prod_{i=0}^{k-1} (k-i)^{s_i}$ distinct k -colorings and (b) with extremely high probability, for every fixed i , $G(n, cn)$ has $\Omega(n)$ vertices of degree i . If X is the number of k -colorings of $G(n, cn)$, using (a),(b), one can show that there are values of c such that for some $a > b > 1$: (i) $\mathbf{E}[X] \approx b^n$ and (ii) w.h.p. if $X > 0$ then $X > a^n$. Hence, for such c , $\Pr[X > 0] \leq (a/b)^n + o(1) = o(1)$, while $\mathbf{E}[X]$ is exponentially large. Thus, it is not the case that $G(n, cn)$ is w.h.p. k -colorable for exactly those values of c for which its expected number of k -colorings is large.

Indeed, Reed and the second author [13] proved that this “naive” first moment argument is quite a bit off the mark for $k = 3$. To that end, they first extended the argument to uniformly random pseudographs on a given degree sequence (for a definition see also [15]). In particular, they proved that such a pseudograph with ρn edges is w.h.p. non- k -colorable if $\rho > \frac{\ln k}{\ln k - \ln(k-1)}$. Then, in order to improve over the naive bound, they considered the random pseudograph resulting by repeatedly (20 times) removing all vertices of degree less than 3 from $G(n, m = cn)$. They proved that this pseudograph (i) is uniformly random with respect to its degree sequence and (ii) if $c \geq c_0 = 2.571\dots$, then w.h.p. it has at least ρn edges where $\rho > \frac{\ln 3}{\ln 3 - \ln 2}$. Hence, w.h.p. $G(n, m = c_0 n)$ contains a non-3-colorable subgraph, implying $c_3 < 2.572$.

Inspired by the work of Kirousis et al. [8], we will take a less direct but more fruitful approach towards accounting for the wastefulness of the first moment method. Instead of

focusing on the low degree vertices explicitly, we will prove the following: if P is a k -coloring of $G \in G(n, cn)$ and we randomly pick a vertex v , then with probability $\phi = \phi(k, c) > 0$ we can assign a different color to v and still have a k -coloring of G . This suggests that when k -colorings exist, they tend to appear in large “clusters” of similar colorings. The approach of Kirousis et al. [8], when translated to coloring, suggests that instead of counting all the k -colorings of a random graph (as the first moment does) we should only count a few “representative” ones. Following this idea we will consider as representatives those k -colorings satisfying a certain “local maximality” condition and determine their expected number in $G^r(n, cn)$. Letting that expectation go to 0 as $n \rightarrow \infty$ will yield $c_3 < 2.522$.

3 A refinement of the first moment method

Recall that for a k -partition $P = V_1, V_2, \dots, V_k$ of V , C_P denotes the event that P is a k -coloring of G . Let us say that a vertex $v \in V_i$ is *unmovable* in P if for every $j > i$ the partition resulting by moving v to V_j is not a k -coloring of G . We will say that P is a *rigid* k -coloring of G if C_P holds and every vertex is unmovable in P . We will denote this event by R_P . Note now that if we consider the k -partitions of V as strings of length n over $\{1, \dots, k\}$ then, clearly, the lexicographically last k -coloring of G (if any k -coloring exists) is rigid by definition. Hence, G has a rigid k -coloring iff it is k -colorable, implying that the probability that $G^r(n, cn)$ is k -colorable is bounded by the expected number of rigid k -colorings of $G^r(n, cn)$. With this in mind, we take $m = cn$ and seek $c = c(k)$ for which this last expectation tends to 0 as $n \rightarrow \infty$.

Remark: Note that requiring k -colorings to be rigid, immediately eliminates all the redundant counting caused by vertices of degree $k - 1$ or less; only the k -colorings which assign every such vertex the greatest possible color get counted.

3.1 Probability Calculations

For every k -partition $P = V_1, V_2, \dots, V_k$ of V we let

$$\alpha_i = \alpha_i(P) = \frac{|V_i|}{n} .$$

Also, recalling (3), we let

$$\tau = \tau(P) = \frac{T(P)}{n^2} . \tag{6}$$

It is well-known that for any $c > 0$, the largest independent set of $G(n, cn)$ w.h.p. contains only a constant fraction of all vertices. Thus, the probability that $G(n, cn)$ has a k -coloring where only one color class contains $\Omega(n)$ vertices is $o(1)$. Hence, in the following we only consider partitions P in which at least two blocks have $\Omega(n)$ vertices (and bound the expected number of rigid k -colorings among such partitions).

We will first bound $\Pr[R_P]$. For this, using (5), it suffices to bound $\Pr[R_P | C_P]$.

For a given k -coloring P , any i , any vertex $v \in V_i$, and any $j > i$ we let $\mathcal{E}(v, j)$ denote the event “ v cannot be moved to V_j ”. Thus,

$$\Pr[R_P | C_P] = \Pr \left[\bigcap_{\substack{i < j \\ v \in V_i}} \mathcal{E}(v, j) \mid C_P \right] . \tag{7}$$

Letting $E(v, j) = \{\{v, w\} : w \in V_j\}$, we see that $\mathcal{E}(v, j)$ occurs iff at least one member of $E(v, j)$ is an edge of G . Note that since we have conditioned on C_P , only two-element sets $\{v, w\}$ enumerated by $T(P)$ can appear in the graph. Thus, since the edges of G were chosen uniformly, independently and with replacement,

$$\Pr[\mathcal{E}(v, j) | C_P] = 1 - \left(1 - \frac{|V_j|}{T(P)} \right)^m \tag{8}$$

$$= 1 - e^{-\alpha_j c/\tau} + O(n^{-1}) , \tag{9}$$

where the passage from (8) to (9) relies on the fact that P has more than one blocks with $\Omega(n)$ vertices and, thus, $T(P) = \Omega(n^2)$. (This is our only use of the fact that there are more than one blocks with $\Omega(n)$ vertices.)

To bound $\Pr[R_P | C_P]$ using (7),(9) we first observe that the sets $E(v, j)$ induce a partition of the set of two-element sets $\{v, w\}$ enumerated by $T(P)$, since each $\{v, w\}$ where $v \in V_i$, $w \in V_j$ and $i < j$ belongs to exactly one such set, namely $E(v, j)$. Since the total number of edges is fixed and each event $\mathcal{E}(v, j)$ “consumes” at least one edge of E , it is intuitively clear that the events $\mathcal{E}(v, j)$ should be negatively correlated. To prove this assertion, we view the formation of E (conditional on C_P) as an allocation scheme with m distinguishable balls, $T(P)$ boxes, and a partition of the set of boxes into disjoint subsets $E(v, j)$, ($v \in V_i$, $i < j$). Thus, the occurrence of $\mathcal{E}(v, j)$ simply means that the total occupancy of boxes from $E(v, j)$ is at least one. Now, the negative correlation of the events $\mathcal{E}(v, j)$ follows from a classical result of McDiarmid [12]. As a result we get

$$\Pr[R_P | C_P] \leq \prod_{\substack{i < j \\ v \in V_i}} \Pr[\mathcal{E}(v, j)] , \tag{10}$$

and, thus, using (7),(9) and (10) we get

$$\begin{aligned} \Pr[R_P | C_P] &\leq \prod_{1 \leq i < j \leq k} (1 - e^{-\alpha_j c/\tau} + O(n^{-1}))^{\alpha_i n} \\ &= \left(\prod_{2 \leq j \leq k} (1 - e^{-\alpha_j c/\tau})^{\sum_{i < j} \alpha_i} \right)^n \times O(1) . \end{aligned} \tag{11}$$

Having bounded $\Pr[R_P | C_P]$, we bound the expected number of rigid k -colorings, $\mathbf{E}[R(G)]$, as follows. For k -partitions $P_1 = V_1^1, \dots, V_k^1$ and $P_2 = V_1^2, \dots, V_k^2$, we say that

P_1 is isomorphic to P_2 if $|V_i^1| = |V_i^2|$, for all i . Clearly, if P_1, P_2 are isomorphic then $\Pr[R_{P_1}] = \Pr[R_{P_2}]$. Let \mathcal{P} be any maximal set of non-isomorphic k -partitions of V . Then

$$\begin{aligned} \mathbf{E}[R(G)] &= \sum_P \Pr[R_P] \\ &= \sum_{P \in \mathcal{P}} \binom{n}{\alpha_1 n, \dots, \alpha_k n} \Pr[R_P] \\ &\leq \max_{P \in \mathcal{P}} \left[\binom{n}{\alpha_1 n, \dots, \alpha_k n} \Pr[R_P] \right] n^{k-1}, \end{aligned} \tag{12}$$

as there are at most n^{k-1} (ordered) partitions of n into k integers. Moreover, if $n > 0$ and all $\alpha_i n$ are integers it is well-known that

$$\binom{n}{\alpha_1 n, \dots, \alpha_k n} < \left(\frac{1}{\alpha_1^{\alpha_1} \dots \alpha_k^{\alpha_k}} \right)^n, \text{ where } 0^0 \equiv 1. \tag{13}$$

Thus, combining (4),(6) and (11)–(13) we have

$$\mathbf{E}[R(G)] \leq \left(\max_{P \in \mathcal{P}} f(P) \right)^n \times O(n^{k-1}) \tag{14}$$

where

$$f(P) = \frac{\left(2 \sum_{i < j} \alpha_i \alpha_j \right)^c}{\alpha_1^{\alpha_1} \dots \alpha_k^{\alpha_k}} \prod_{2 \leq j \leq k} \left(1 - e^{-\alpha_j c / \tau} \right)^{\sum_{i < j} \alpha_i}. \tag{15}$$

Letting $Q = \{q/n : q \in \{0, \dots, n\}\}$, it is clear that maximizing f over $P \in \mathcal{P}$ amounts to maximizing the right-hand side of (15) over Q^k subject to $\sum_i \alpha_i = 1$. Naturally, we still get an upper bound on $\mathbf{E}[R(G)]$ if we relax each such α_i to an arbitrary real number in $[0, 1]$ and maximize the extended function, g , over $D = [0, 1]^k$ subject to $\sum_i \alpha_i = 1$. If for some $c^* = c^*(k)$ the resulting maximum of g is strictly less than 1, then (14) implies that $\mathbf{E}[R(G)] \rightarrow 0$ as $n \rightarrow \infty$ and, thus, that $G^r(n, c^*n)$ is w.h.p. non- k -colorable.

It is straightforward to verify that g is continuous, differentiable and its gradient is bounded on D . As a result, g can be maximized numerically with arbitrarily good, guaranteed precision (we used Maple [18] and the code in [17]). For example, for $k = 3$ we have

$$g(\alpha_1, \alpha_2, \alpha_3) = \frac{(2\tau_3)^c (1 - e^{-\alpha_2 c / \tau_3})^{\alpha_1} (1 - e^{-\alpha_3 c / \tau_3})^{\alpha_1 + \alpha_2}}{\alpha_1^{\alpha_1} \alpha_2^{\alpha_2} \alpha_3^{\alpha_3}}$$

where $\tau_3 = \alpha_1 \alpha_2 + \alpha_1 \alpha_3 + \alpha_2 \alpha_3$. For $c^* = 2.5217$, g is maximized around $\alpha_1 = 0.30746$, $\alpha_2 = 0.33527$, $\alpha_3 = 0.35727$ and at that vicinity it is strictly less than 0.9999744. Thus, $G^r(n, m = c^*n)$ is w.h.p. non- k -colorable, implying $c_3 < 2.522$.

Similarly, we get the following new bounds for c_k for $3 \leq k \leq 7$. (The choice of 7 is rather arbitrary, as the numerical computations remain manageable for substantially larger k .)

k	3	4	5	6	7
First moment bound	2.710	4.819	7.213	9.828	12.714
New bound	2.522	4.587	6.948	9.539	12.316

The above table gives an idea of how our improvement over the first moment bound scales with k . Recalling that the first moment bound is asymptotically tight, we see that already for $k = 7$ the improvement has dropped to less than 3% from 7% for $k = 3$.

It seems clear that one could improve the upper bound on c_k somewhat further by imposing a stricter local maximality condition. For example, one could consider conditions that involve “moving” two vertices at a time. Unfortunately, the lack of “independence” between the outcomes of different moves in that setting seems to complicate matters greatly.

Acknowledgements

We would like to thank the authors of [8] for providing us with an early draft of their paper and an anonymous referee for many valuable comments.

References

- [1] D. Achlioptas, E. Friedgut, A sharp threshold for k -colorability, *Random Structures & Algorithms*, **14** (1999), 63–70.
- [2] D. Achlioptas, M. Molloy, The analysis of a list-coloring algorithm on a random graph, *38th Annual Symposium on Foundations of Computer Science*, Miami, FL (1997), 204–212.
- [3] N. Alon, J.H. Spencer, **The Probabilistic Method**, with an Appendix of open problems by P. Erdős, J. Wiley & Sons, New York, 1992.
- [4] V. Chvátal, Almost all graphs with $1.44n$ edges are 3-colorable, *Random Structures & Algorithms*, **2** (1991), 11–28.
- [5] P. Erdős, A. Rényi, On the evolution of random graphs, Publication of the Mathematical Institute of the Hungarian Academy of Sciences, **5** (1960), 17–61.
- [6] E. Friedgut, Necessary and sufficient conditions for sharp thresholds of graph properties, and the k -SAT problem, *Journal of the American Mathematical Society*, **12** (1999), 1017–1054.
- [7] G.R. Grimmett, C.J.H. McDiarmid, On colouring random graphs, *Mathematical Proceedings of the Cambridge Philosophical Society*, **77** (1975), 313–324.
- [8] L. M. Kirousis, E. Kranakis, D. Krizanc, and Y. Stamatiou, Approximating the unsatisfiability threshold of random formulas, *Random Structures & Algorithms*, **12** (1998), 253–269.

- [9] L. M. Kirousis, Y. Stamatiou, An inequality for reducible increasing properties of random words, Technical Report, Computer Technology Institute, Greece, (1997), 1–3.
- [10] T. Łuczak, The chromatic number of random graphs, *Combinatorica*, **11** (1991), 45–54.
- [11] T. Łuczak, Size and connectivity of the k -core of a random graph, *Discrete Mathematics*, **91** (1991), 61–68.
- [12] C.J.H. McDiarmid, On a correlation inequality of Farr, *Combinatorics, Probability and Computing*, **1** (1992), 157–160.
- [13] M. Molloy, The chromatic number of sparse random graphs, M. Math Thesis, University of Waterloo, (1992).
- [14] M. Molloy, A gap between the appearances of a k -core and a $(k + 1)$ -chromatic graph, *Random Structures & Algorithms*, **8** (1996), 159–160.
- [15] M. Molloy, B. Reed, A critical point for random graphs with a given degree sequence, *Random Structures & Algorithms*, **6** (1995), 161–179.
- [16] B. Pittel, J.H. Spencer, and N.C. Wormald, Sudden emergence of a giant k -core in a random graph, *Journal of Combinatorial Theory Series B*, **67** (1996), 111–151.
- [17] W.H. Press, S. Teukolsky, W.T. Vetterling, and B.P. Flannery, **Numerical recipes in C**, Cambridge University Press, Cambridge, 1992.
- [18] D. Redfern, **The Maple Handbook: Maple V Release 3**, Springer, New York, 1994.