ABRAHAM D. FLAXMAN

Address:
1 Microsoft Way
Redmond, WA 98052
(412) 726-0401

E-mail:
abie@microsoft.com

Webpage:
www.math.cmu.edu/~adf/

1. R. Andersen, C. Borgs, J. Chayes, U. Feige, A. D. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: an axiomatic approach. Manuscript submitted for publication.
2. A. D. Flaxman and B. Reed. The frugality of a VCG path auction in a sparse random graph. Manuscript submitted for publication.
3. A. D. Flaxman, A. M. Frieze, J. Vera. A geometric preferential attachment model of networks II. To appear in Web Algorithms Workshop (2007).
4. A. D. Flaxman, J. Vera. Bias reduction in traceroute sampling: towards a more accurate map of the Internet. To appear in Web Algorithms Workshop (2007).
5. M. Carey, A. D. Flaxman, J. Hartline, A. Karlin. Auctions for structured procurement. To appear in SODA (2008).
6. A. D. Flaxman. Algorithms for random 3-SAT. To appear in *The Encyclopedia of Algorithms.*

   Survey article about distributions over instances of 3-SAT and polynomial-time algorithms which succeed on these instances with high probability.
7. A. D. Flaxman. Expansion and lack thereof in perturbed random graphs. Microsoft Research Technical Report TR-2006-118 (2006); (to appear in *Proc. of 4th International Workshop on Algorithms and Models for the Web-Graph (WAW)* (2006)).

   Randomly perturbed graphs are not necessarily constant expanders.
8. A. D. Flaxman. A lower bound on the large deviation probability for the lower tail of the random minimum spanning tree. *Electronic Journal of Combinatorics* 14 (1) (2007) N3.

   It is known that if the edge costs of the complete graph K_n are independent random variables, uniformly distributed between 0 and 1, then the expected cost of the minimum spanning tree is asymptotically equal to $\zeta(3) = \sum_{i=1}^{\infty} i^{-3}$. Here I consider the large deviation probability of this random variable. The proof technique is simple but relies on what I consider a very slick trick based on the two-stage stochastic version of this problem.

9. A. D. Flaxman and S. Hoory. Maximum matchings in regular graphs of high girth. *Electronic Journal of Combinatorics* 14 (1) (2007) N1.

   If a regular graph has high girth then it must also have a maximum matching that is pretty close to a perfect matching. The proof combines the non-regular Moore with the linear programming formulation of the maximum matching problem.

10. H. Yu, M. Kaminsky, P. B. Gibbons, and A. D. Flaxman. Defending against Sybil attacks via social networks. *Proc. of the 2006 conference on applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, (2006) 267-278.

    The Sybil attack is a major security issue in online systems, where an attacker controls an unlimited number of agents. This paper uses assumptions about the likely structure of a social network to defend against such an attack. My contribution was in understanding the structural properties that the social network would be likely to exhibit.

11. A. D. Flaxman and G. B. Sorkin. Tools for average-case analysis of optimization heuristics. Manuscript in preparation, to appear in *Handbooks of Operations Research and Management Science: Approximation and Heuristics.*

    This is an expository chapter on techniques for analyzing random algorithms and random structures, illustrated with a variety of examples, with emphasis on optimization problems and typical approximation ratios achieved on random instances. Some of the techniques include McDiarmid's formulation of the Azuma-Hoefding Inequality, Talagrand's Inequality, the Logarithmic Sobolev Inequality, and Wormald's Differential Equations Method.

12. U. Feige, A. D. Flaxman, J. D. Hartline, R. Kleinberg. On the competitive ratio of the random sampling auction. *Proc. of the 1st International Workshop on Internet and Network Economics* (2005) 878-886.

    We give a simple analysis of the competitive ratio of the random sampling auction for digital goods (the auction consists of randomly bipartitioning the bidders, finding the best single price to offer each part, and offering that price to the other part). The random sampling auction was previously known to have a constant competitive ratio (with a bound of 7600 on the constant); our analysis improves the bound to 15. But, in my heart, I believe the true value of the competitive ratio is 4. In support of this, we show that on the equal revenue input, where any sale price gives the same revenue, random sampling is exactly a factor of 4 from optimal.

    The proof that the constant is at most 15 relies on a probability-of-ruin calculation for a particular asymmetric random walk, which can be calculated exactly, but requires solving a quartic equation. The proof that random sampling is 4-competitive on the equal revenue distribution is computer-aided.

13. A. D. Flaxman, D. Gamarnik, and G. B. Sorkin. First-passage percolation and VCG-overpayment on a width-2 strip. *Proc. of the 2nd International Workshop on Internet and Network Economics* (2006) 99-111.

We consider first-passage percolation on the width-2 strip (which is alternatively known as the length of the shortest path), for edge weights with Bernoulli or uniform distribution. Our main result is to determine exact rate constants. The same technique also allows us to determine the overpayment incurred by the VCG mechanism when each edge is controlled by a selfish agent.

The proofs in this paper are similar in spirit to Aldous' Objective Method (also called Local Weak Convergence), which was used to prove that the cost of a minimum assignment under random edge weights is $\pi^2/6$. However, since the width-2 strip is such a simple graph, in this context the technique is more easily accessible.

14. D. Coppersmith, A. D. Flaxman, and C. Smyth. On sets of permutations summing to permutations. Unpublished manuscript, Carnegie Mellon University, Pittsburgh (2004).

A permutation sum (resp. difference) set on a group G is a set $\mathcal{F}$ of bijections from G to G such that fg (resp. $f^{-1}g$) is again a bijection for all $f, g \in \mathcal{F}$ (resp. $f, g \in \mathcal{F}$ with $f \neq g \in S$), where $(fg)(x) := f(x)g(x)$ for all $x \in G$. The maximum size $d(G)$ of a permutation difference set has been well studied, with many connections drawn between such sets and combinatorial objects such as families of pairwise orthogonal Latin squares. Here we primarily study its natural counterpart, $s(G)$, the maximum size of a permutation sum set.

The two quantities often differ widely. If p is a prime, we have $d(\mathbb{Z}_{p-1}) = p - 1$ while $\max(p \times 2^{(p-1)/k}, \binom{p}{2}) \leq s(\mathbb{Z}_p) \leq p((p-1)/2)^{(p-3)/2}$ where k is the multiplicative order of -2 mod p. For example $d(\mathbb{Z}_{1613}) = 1612$ while $s(\mathbb{Z}_{1613}) \geq 1613 \times 2^{31} > 3 \times 10^{12}$.

15. M. Dyer, A. D. Flaxman, A. M. Frieze, and E. Vigoda. Randomly coloring sparse random graphs with fewer colors than the maximum degree. *Random Structures and Algorithms* 29, (2006) 450-465.

We analyze Markov chains for generating a random k-coloring of a random graph $G_{n,d/n}$. When the average degree d is constant, a random graph has maximum degree $\log n/ \log\log n$, with high probability. We efficiently generate a random k-coloring when $k = \Omega(\log\log n/ \log\log\log n)$, i.e., with many fewer colors than the maximum degree. Previous results hold for a more general class of graphs, but always require more colors than the maximum degree. For $k \geq (\ln n)^\alpha$ we can also show that the Glauber Dynamics mix in polynomial time.

16. A. D. Flaxman, A. M. Frieze, and J. C. Vera. On the average case performance of some greedy approximation algorithms for the uncapacitated facility location problem. *Proc. of the 37th ACM Symposium on the Theory of Computing (STOC)* (2005) 441-449.

We analyze the performance of three greedy approximation algorithms for the uncapacitated facility location problem on random instances. The random instances are formed by placing n points uniformly at random in the unit square. We find that these three algorithms do not find asymptotically optimal solutions with high probability, while a simple plane-partitioning heuristic does find an asymptotically optimal solution with high probability.

It turned out to be hard work to do an average-case analysis of these algorithms, because of the dependent nature of the random variables. My hope is that this sort of analysis will give an alternative to the worst-case competitive ratio for measuring the effectiveness of approximation algorithms. But for that, we need to analyze alternative algorithms under this model of randomness, such as heuristics which use a relax-and-round approach.

17. A. D. Flaxman and B. Pryzdatek. Solving medium-density subset sum problems in expected polynomial time. *Proc. of the 22nd Symposium on Theoretical Aspects of Computer Science (STACS)* (2005) 305-314.

    The subset sum problem (SSP) (given n numbers and a target bound B, find a subset of the numbers summing to B), is a classic NP-hard problem. The hardness of SSP varies greatly with the density of the problem. In particular, when m, the logarithm of the largest input number, is at least $c \cdot n$ for some constant c, the problem can be solved by a reduction to finding a short vector in a lattice. On the other hand, when $m = O(\log n)$ the problem can be solved in polynomial time using dynamic programming or some other algorithms especially designed for dense instances. However, as far as we are aware, all known algorithms for dense SSP take at least $\Omega(2^m)$ time, and no polynomial time algorithm is known which solves SSP when $m = \omega(\log n)$ (and $m = o(n)$).

    We present an expected polynomial time algorithm for solving uniformly random instances of the subset sum problem over the domain $\mathbb{Z}[M]$, with $m = O((\log n)^2)$. To the best of our knowledge, this is the first algorithm working efficiently beyond the magnitude bound of $O(\log n)$, thus narrowing the interval of hard-to-solve SSP instances.

    It turns out that this algorithm has an application to solving the shortest vector problem with a quantum computer. If the density could be increased to anything significantly larger than $(\log n)^2$ this would yield a faster quantum algorithm (but still super-polynomial time) for finding the shortest vector in a lattice.

18. A. D. Flaxman, A. T. Kalai, and H. B. McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *Proc. of the 16th Symposium on Discrete Algorithms (SODA)* (2005) 385-394.

    We study a general online convex optimization problem. We have a convex set S and an unknown sequence of cost functions $(c_1, c_2, \ldots)$, and, in each round, we choose a feasible point x_t in S, and learn the cost $c_t(x_t)$. If the function c_t is also revealed after each round then, as Zinkevich has shown, gradient descent can be used on these functions to get regret bounds of $O(\sqrt{n})$. That is, after n rounds, the total cost incurred will be $O(\sqrt{n})$ more than the cost of the best single feasible decision chosen with the benefit of hindsight, $\min_{x \in S} \sum c_t(x)$.

    We extend this to the "bandit" setting where, during each period, only the cost $c_t(x_t)$ is revealed, and we obtain an upper-bound on the expected regret of $O(n^{3/4})$.

    Our approach uses a simple approximation of the gradient that is computed from evaluating c_t at a single (random) point. We show that this biased estimate is sufficient to approximate gradient descent on the sequence of functions. In other words, it is possible to use gradient descent in the online setting without seeing anything more than the value of the functions at a single point.

For the online linear optimization problem, algorithms with low regrets in the bandit setting have recently been given against oblivious and adaptive adversaries. In contrast to these algorithms, which divide time into explicit *explore* and *exploit* phases, our algorithm can be interpreted as doing a small amount of exploration in each round.

19. A. D. Flaxman, A. M. Frieze, and J. C. Vera. Adversarial deletions in a scale-free random graph process. *Combinatorics, Probability, and Computing* 16 (2007) 261-270; (extended abstract appeared in *Proc. of the 16th Symposium on Discrete Algorithms (SODA)* (2005) 287-292).

    We study a dynamically evolving random graph which adds vertices and edges using preferential attachment and is "attacked by an adversary." At time t, we add a new vertex x_t and m random edges incident with x_t, where m is constant. The neighbors of x_t are chosen with probability proportional to degree. After adding the edges, the adversary is allowed to delete vertices. The only constraint on the adversarial deletions is that the total number of vertices deleted by time n must be no larger than δn, where δ is a constant. We show that if δ is sufficiently small, then at time n the generated graph has a giant component, with high probability.

    As with many theoretical results about preferential attachment graphs, the proof here is rather involved. However, it is made much easier through an elegant (relatively speaking) coupling argument which shows that the graph we are studying can be approximated in a certain sense by Erdős-Renyi graphs with appropriate parameters. This is an extension of a very nice idea of Bollobás and Riordan.

20. A. D. Flaxman, A. M. Frieze, and M. Krivelevich. On the random 2-stage minimum spanning tree. *Random Structures and Algorithms* 28 (1) (2006) 24-36. (Extended abstract appeared as *Proc. of the 16th Symposium on Discrete Algorithms (SODA)* (2005) 919-928).

    It was previously known that if the edge costs of the complete graph K_n are independent random variables, uniformly distributed between 0 and 1, then the expected cost of the minimum spanning tree is asymptotically equal to $\zeta(3) = \sum_{i=1}^{\infty} i^{-3}$. Here we consider the following stochastic two-stage version of this optimization problem. There are two sets of edge costs $c_M \colon E \to \mathbb{R}$ and $c_T \colon E \to \mathbb{R}$, called Monday's prices and Tuesday's prices, respectively. For each edge e, both costs $c_M(e)$ and $c_T(e)$ are independent random variables, uniformly distributed in $[0, 1]$. The Monday costs are revealed first. The algorithm has to decide on Monday for each edge e whether to buy it at Monday's price $c_M(e)$, or to wait until its Tuesday price $c_T(e)$ appears. The set of edges X_M bought on Monday is then completed by the set of edges X_T bought on Tuesday to form a spanning tree. If both Monday's and Tuesday's prices were revealed simultaneously, then the optimal solution would have expected cost $\zeta(3)/2 + o(1)$. We show that in the case of two-stage optimization, the expected value of the optimal cost exceeds $\zeta(3)/2$ by an absolute constant $\epsilon > 0$. We also consider a threshold heuristic, where, on

Monday, the algorithm buys edges of cost less than α, and, on Tuesday, the algorithm completes the tree in an optimal way. We show that the optimal choice for α is $\alpha = 1/n$ for which the expected cost is $\zeta(3) - 1/2 + o(1)$. The threshold heuristic is shown to be sub-optimal. Finally we discuss the directed version of the problem, where the task is to construct a spanning out-arborescence rooted at a fixed vertex r. We show, somewhat surprisingly, that in this case a simple variant of the threshold heuristic gives the asymptotically optimal value $1 - 1/e + o(1)$.

I find it quite intriguing that the limiting value of so many natural combinational optimization problems turn out to be miraculous constants related to Riemann's zeta function. The calculations which show that the solution found by the threshold heuristic has value asymptotic to $\zeta(3) - 1/2$ is particularly miraculous (and unfortunately offers no "reason" why this value appears.)

21. A. D. Flaxman, A. M. Frieze, and J. C. Vera. A geometric preferential attachment model of networks. *Proc. of 3rd International Workshop on Algorithms and Models for the Web-Graph* (2004) 44-55.

    We study a random graph G_n that combines certain aspects of geometric random graphs and preferential attachment graphs. The vertices of G_n are n sequentially generated points $x_1, x_2, \ldots, x_n$ chosen uniformly at random from the unit sphere in $\mathbb{R}^3$. After generating x_t, we randomly connect it to m points from those points in $x_1, x_2, \ldots, x_{t-1}$ which are within distance r. Neighbors are chosen with probability proportional to their current degree. We show that if m is sufficiently large and if $r \geq \log n / n^{1/2-\beta}$ for some constant β then with high probability at time n the number of vertices of degree k follows a power law with exponent 3. Unlike the preferential attachment graph, this geometric preferential attachment graph has small separators, similar to some experimental observations. We also show that if $m \geq K \log n$, K sufficiently large, then G_n is connected and has diameter $O(m/r)$ with high probability.

    The main results of this paper rely on a lemma with an erroneous proof. We are currently working to fix it.

22. A. D. Flaxman and A. M. Frieze. The diameter of randomly perturbed digraphs and some applications. *Proc. of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems and 8th International Workshop on Randomization and Computation* (2004) 345-356.

    The central observation of this paper is that if ϵn random edges are added to any n-node connected graph or digraph then the resulting graph has diameter $\mathcal{O}(\log n)$ with high probability. We apply this to smoothed analysis of algorithms and property testing.

    Smoothed Analysis: Recognizing strongly connected digraphs is a basic computational task in graph theory. Even for graphs with bounded out-degree, it is NL-complete. By XORing an arbitrary bounded out-degree digraph with a sparse random digraph $R \sim \mathbb{G}_{n,\epsilon/n}$ we obtain a "smoothed" instance. We show that, with high probability, a log-space algorithm will correctly determine if a smoothed instance is strongly connected.

We also show that if NL $\not\subseteq$ almost-L then no heuristic can recognize similarly perturbed instances of (s,t)-connectivity.

Property Testing: A digraph is called k-linked if for every choice of $2k$ distinct vertices $s_1, t_1, \ldots, s_k, t_k$, the graph contains k vertex disjoint paths joining s_r to t_r for $r = 1, \ldots, k$. Recognizing k-linked digraphs is NP-complete for $k \geq 2$. We describe a polynomial time algorithm for bounded degree digraphs which accepts k-linked graphs with high probability, and rejects all graphs which are at least ϵn edges away from being k-linked.

I think both applications of the central observation are great, and I think the theorem is quite interesting in its own right. It provides some theoretical explanation of the "six degrees of separation" phenomenon observed experimentally by Stanley Milgram, by showing that any large-enough network which contains a small amount of uniform randomness will have low diameter.

23. A. D. Flaxman. A sharp threshold for a random constraint satisfaction problem. *Discrete Math.* 285/1-3 (2004) 301-305.

    Considers random instances I of a constraint satisfaction problem generalizing k-SAT: given a set of ordered k-tuples over n literals, and a set of q "bad" clause assignments, find an assignment which does not set any of the k-tuples to a bad clause assignment. We consider the case where $k = \Omega(\log n)$, and study the probability of satisfiability for a random instance I formed by including every k-tuple of literals independently with probability p. Appropriate choice of the bad clause assignments results in random instances of k-SAT and not-all-equal k-SAT. A second moment method calculation yields the sharp threshold

$$\lim_{n\to\infty} \Pr[I \text{ is satisfiable}] = \begin{cases} 1, & \text{if } p \leq (1-\epsilon)\frac{\ln 2}{qn^{k-1}}; \\ 0, & \text{if } p \geq (1+\epsilon)\frac{\ln 2}{qn^{k-1}}. \end{cases}$$

    This calculation uses a nice trick to turn a messy sum into a messy random variable, with which it turns out to be easier to deal.

24. A. D. Flaxman, A. M. Frieze, and E. Upfal. Efficient communication in an ad-hoc network. *J. Algorithms* 52 (1) (2004) 1-7.

    Considers a communication protocol for a large number of low power sensors distributed uniformly at random in a large geographic area. The sensors can communicate directly only with neighbors within a relatively small radius so communication to a distant unit must propagate through a large number of intermediate units. We show that an algorithm similar to local flooding succeeds with high probability in sending messages between any 2 points within the giant component.

    These random geometric graphs are a joy to work with, and required combining ideas from plane geometry, discrete probability, and combinatorial optimization.

25. A. D. Flaxman, D. Gamarnik, and G. B. Sorkin. Embracing the giant component. *Proc. of the 6th Conference of Latin American Theoretical Informatics* (2004) 69-79. (To appear in *Random Structures and Algorithms.*)

Consider a game in which edges of a graph are provided a pair at a time, and the player selects one edge from each pair, attempting to construct a graph with a component as large as possible. This is in the spirit of recent papers on *avoiding* a giant component, and the general investigation of the power of 2 choices.

If the instance has $\frac{1}{4}(1+\epsilon)n$ random edge pairs, with $0 < \epsilon \leq 0.003$, then any online algorithm generates a component of size $O((\log n)^{3/2})$ with high probability, while the optimal offline solution contains a component of size $\Omega(n)$ with high probability. For "dense" random instances, the average-case competitive ratio is much smaller. If the instance has $\frac{1}{2}(1-\epsilon)n$ random edge pairs, with $0 < \epsilon \leq 0.015$, we give an online algorithm which finds a component of size $\Omega(n)$ with high probability.

The constants are a mess, but the take-home message is that there is a range where an online algorithm can achieve a constant competitive ration and another range where no online algorithm can be $n/\operatorname{polylog}(n)$ competitive, even on a simple average-case input.

26. A. D. Flaxman, A. W. Harrow, and G. B. Sorkin. Strings with maximum numbers of distinct subsequences and substrings. *Elec. J. Combinatorics* 11 (1) R8 (2004).

    A natural problem in extremal combinatorics is to maximize the number of distinct subsequences for any length-n string over a finite alphabet Σ; this value grows exponentially, but slower than 2^n. We use the probabilistic method to determine the maximizing string, which is a cyclically repeating string. The number of distinct subsequences is exactly enumerated by a generating function, from which we also derive asymptotic estimates. For the alphabet $\Sigma = \{1, 2\}$, $(1, 2, 1, 2, \dots)$ has the maximum number of distinct subsequences, namely $\mathrm{Fib}(n+3) - 1 \sim \left((1+\sqrt{5})/2\right)^{n+3}/\sqrt{5}$.

    We also consider the same problem with sub*strings* in lieu of sub*sequences*. Here, we show that an appropriately truncated de Bruijn word attains the maximum. For both problems, we compare the performance of random strings with that of the optimal ones.

    The charm of this paper is the proof that the cyclically repeating string maximizes the number of distinct subsequences. We manage to use the probabilistic method not to show that a random string has the property under consideration, but that a very particular string has the property under consideration.

27. A. D. Flaxman. Note on the hardness of finding the minimum cut in space. Unpublished manuscript, Carnegie Mellon University, Pittsburgh (2003).

    We are given a collection of points (examples) in n-dimensional space, some are labeled "+", some are labeled "-", and some are not labeled at all. In addition, some of the points are connected by edges. The goal is to find a linear separator with the positive labeled examples on one side and the negative labeled examples on the other, that cuts as few of the edges as possible. Alternatively, we can think of the input as just a collection of labeled points, together with some unlabeled convex bodies. Then the goal is to separate the points, while cutting as few of the convex bodies as possible.

We show that it is NP-hard to find the linear separator that minimizes the number of edges cut. The proof is a reduction from 3SAT.

Proving that new optimization problems are NP-hard is a fun diversion. And someone needs to do it, especially when there is a problem that seems similar to something polynomially solvable. This short gadget-style reduction is a nice exercise in high-dimensional geometry.

28. A. D. Flaxman, A. M. Frieze, and T. Fenner. High degree vertices and eigenvalues in the preferential attachment graph, *Proc. 7th International Workshop on Randomization and Approximation Techniques in Computer Science* (2003) 264-274. (To appear in *Internet Mathematics.*)

    The preferential attachment graph is a random graph formed by adding a new vertex at each time step, with a single edge which points to a vertex selected at random with probability proportional to its degree. Every m steps the most recently added m vertices are contracted into a single vertex, so at time t there are roughly t/m vertices and exactly t edges. This process yields a graph which has been proposed as a simple model of the world wide web. For any constant k, let $\Delta_1 \geq \Delta_2 \geq \cdots \geq \Delta_k$ be the degrees of the k highest degree vertices. We show that at time t, for any function f with $f(t) \to \infty$ as $t \to \infty$, $\frac{t^{1/2}}{f(t)} \leq \Delta_1 \leq t^{1/2} f(t)$, and for $i = 2, \ldots, k$, $\frac{t^{1/2}}{f(t)} \leq \Delta_i \leq \Delta_{i-1} - \frac{t^{1/2}}{f(t)}$, with high probability. We use this to show that at time t the largest k eigenvalues of the adjacency matrix of this graph have $\lambda_k = (1 \pm o(1))\Delta_k^{1/2}$ with high probability.

    This paper contains some of the most difficult calculations I have ever had to slog through. The most complicated proof was greatly simplified after the conference version by looking at the expected value of the k-th rising factorial of a random variable instead of the k-th power.

29. A. D. Flaxman. A spectral technique for random satisfiable 3CNF formulas, *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms* (2003) 357-363.

    Let I be a random 3CNF formula generated by choosing a truth assignment ϕ for variables $x_1, \ldots, x_n$ uniformly at random and including every clause with i literals set true by ϕ with probability p_i, independently. For such instances, we develop a spectral algorithm similar to the graph coloring algorithm of Alon and Kahale. We show that for any constants $0 \leq \eta_2, \eta_3 \leq 1$ there is a constant d_{min} so that for all $d \geq d_{min}$ the algorithm finds a satisfying assignment with high probability for $p_1 = d/n^2$, $p_2 = \eta_2 d/n^2$, and $p_3 = \eta_3 d/n^2$. Appropriately setting the η_i's yields natural distributions on satisfiable 3CNFs, not-all-equal-sat 3CNFs, and exactly-one-sat 3CNFs.

    The spectral part of the algorithm is very magical at first sight. However, the hard work is the non-spectral part, which required some careful analysis (which is follows this trail blazed by the prior works of Alon and Kahale and of Chen and Frieze). This algorithm resolved a conjecture of Koutsoupias and Papadimotriou from 1992.