BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
http://dx.doi.org/10.1090/bull/1557
Article electronically published on September 8, 2016

GRAPHS, VECTORS, AND MATRICES

DANIEL A. SPIELMAN

Abstract. This survey accompanies the Josiah Williard Gibbs Lecture that I gave at the 2016 Joint Mathematics Meetings. It provides an introduction to three topics: algebraic and spectral graph theory, the sparsification of graphs, and the recent resolution of the Kadison–Singer Problem. The first and third are connected by the second.

1. Introduction

Graphs are the quintessential objects of study in discrete mathematics. They are usually described as a set of vertices, V, that are connected by a set of edges, E, each of which is a pair of vertices. Graphs encode connections and are one of the most commonly used representations of data. Mathematicians often define graphs abstractly. For example, we define a path graph to be a graph with vertex set $V = \{1, \ldots, n\}$ and $E = \{(i, i+1) : 1 \leq i < n\}$. Or a number theorist might consider a graph with $V = \{1, \ldots, n\}$ and E the set of pairs (i, j) for which i divides j. The public is more familiar with social network graphs, in which each person is a vertex and edges exist between pairs of people who are "friends". Chemists consider graphs connecting the atoms within a molecule. Physicists consider graphs describing the interactions of molecules. The graphs that arise in one discipline may have little in common with those that arise in another.

While we first learn to prove theorems about graphs through local arguments and combinatorial manipulations, much of what I want to know about a graph is revealed through the more continuous approach of algebraic graph theory. I introduce this approach in Section 2 by defining the Laplacian matrices of graphs and their associated quadratic forms, and by showing how they help us draw graphs and understand their structure. When we introduce graphs to students, we often do so through pictures. We draw the vertices as little circles and the edges as lines or curves connecting the circles representing their endpoints. While we obtain the same graph wherever we put the circles, some drawings reveal the structure of the graph much better than others. For example, consider the two drawings in Figure 1. They both represent the same graph, but the second reveals its structure much better than the first. It was drawn using two eigenvectors of the Laplacian matrix of the graph, in a way explained in Section 2.2.

In Section 3, we use Laplacian matrices to associate a vector with each edge of a graph. We then explain how every collection of vectors can be approximated by a rescaling of a small subset of those vectors, and use this association to prove that every graph can be approximated by a graph having few edges. Our notion

Received by the editors July 29, 2016.
2010 *Mathematics Subject Classification.* Primary 05C50.

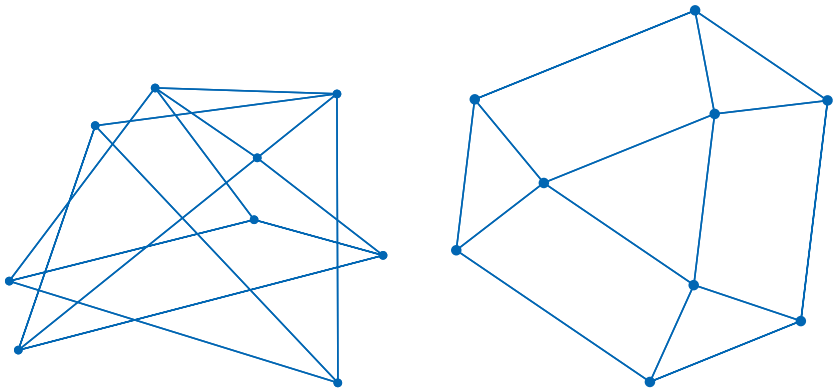

FIGURE 1. An arbitrary drawing of a graph, and a spectral drawing of that graph.

of approximation is algebraic: we say that one graph approximates another if their Laplacian matrices are similar.

The statement that every collection of vectors can be approximated by a subset of those vectors has a very close resemblance to a conjecture in discrepancy theory that Weaver [Wea04] proved would affirmatively resolve the Kadison–Singer Problem. In Section 4 we explain Weaver's conjecture and introduce some of the techniques that were recently used to prove it.

2. The graph Laplacian

We now define the Laplacian quadratic form and the Laplacian matrix associated with a weighted, undirected graph. Given an undirected graph G with vertex set V, edge set E, and positive edge weights, $w_{a,b}$ for $(a,b) \in E$, we define the Laplacian quadratic form associated with G to be the function from $x \in \mathbb{R}^V$ to real numbers given by

$$q_G(x) \stackrel{\text{def}}{=} \sum_{(a,b)\in E} w_{a,b}(x(a) - x(b))^2.$$

In this expression, the terms a and b index vertices and the unordered pair (a,b) indexes an edge.

The Laplacian matrix of G, denoted L_G, is the symmetric matrix such that

$$q_G(x) = x^T L_G x.$$

When reading these expressions, one should keep in mind that the entries of x are indexed by vertices, and so the rows and columns of L_G are also indexed by vertices. We will often think of x as being a function from V to $\mathbb{R}$.

2.1. Physical models.

The graph Laplacian has many natural physical interpretations, and these provide some of the motivations for its use. For example, we could imagine that each edge of the graph is a resistor in an electrical network. As a large weight represents a strong connection, it should correspond to a low resistance. Thus, we view an edge (a,b) of weight $w_{a,b}$ as a resistor connecting a to b of resistance $1/w_{a,b}$. If we then impose voltages $x(a)$ and $x(b)$ at vertices a and b, a current of $w_{a,b}(x(a) - x(b))$ will flow across edge (a,b) and result in an energy

dissipation $w_{a,b}(x(a)-x(b))^2$. If x provides the voltage at every vertex, then $q_G(x)$ is the sum of the energy dissipation over all resistors. In this way, the electrical properties of the resistor network are described by the graph Laplacian.

One of the most useful measurements in a resistor network is the *effective resistance* between two vertices. This is the resistance between the two vertices when we view the entire network as one complex resistor. One way to measure the effective resistance between vertices a and b is to fix the voltage of vertex a to 1, fix the voltage of vertex b to 0, and then measure how much current flows from a to b. The effective resistance is the reciprocal of the current flow. Physics tells us that the electrical current will minimize energy dissipation. Thus, the induced voltage at every vertex will be given by the vector x that minimizes $q_G(x)$ subject to $x(a) = 1$ and $x(b) = 0$. One can compute this minimum by finding the vector at which the partial derivatives of $q_G(x)$ are zero. This amounts to solving a system of linear equations in the submatrix of L_G containing the rows and columns that are not indexed by a or b. The effective resistance between vertices is known to be a distance [KR93], and it is small when there are many distinct short paths connecting them.

Another useful physical model for a network is obtained by thinking of each edge (a, b) as being a spring with spring constant $w_{a,b}$. The potential energy in a spring with constant $w_{a,b}$ that has been stretched to length l is $w_{a,b}l^2/2$. Thus, the potential energy in the spring network corresponding to the graph G when each vertex a is fixed to location $x(a) \in \mathbb{R}$ is $q_G(x)/2$. However, we are not restricted to fixing vertices to real numbers; we can locate them in the plane! If we fix vertex a to the point $(x(a), y(a))$, then the potential energy in the spring network is

$$\frac{1}{2}\sum_{(a,b)\in E} w_{a,b}\left\|\begin{pmatrix}x(a)\\y(a)\end{pmatrix}-\begin{pmatrix}x(b)\\y(b)\end{pmatrix}\right\|^2 = \frac{1}{2}\left(q_G(x)+q_G(y)\right).$$

In an amazing paper called "How to Draw a Graph", Tutte [Tut63] proved that we can use these spring networks to obtain planar drawings of planar graphs!

A planar graph is one that can be drawn in the plane without crossing edges. That is, every vertex can be mapped to a distinct point in the plane and every edge can be drawn as a continuous curve connecting the points to which its vertices are mapped in such a way that the curves representing distinct edges do not intersect. If we remove the points representing vertices and the curves representing edges from the plane, we are left with planar regions called *faces*. Faces are typically bounded by *cycles* in the graph. Tutte suggests drawing a planar graph by fixing the locations of the vertices defining one face to the corners of a convex polygon, treating the graph as a spring network, and then placing the other vertices at the positions that minimize the potential energy of the network. This is exactly the position to which physics tells us they should settle, and it is the position in which every vertex that is not fixed is located at the center of gravity of its neighbors. Tutte proves that if the graph is planar and 3-connected,[1] then this provides a planar drawing of the graph. We provide an example of such a drawing in Figure 2.

[1]A graph is 3-connected if there is no set of two vertices whose removal disconnects the graph. The boundaries of faces in 3-connected graphs are always cycles. Graphs that are not 3-connected are the natural obstructions to this theorem: if there are vertices a and b whose removal disconnects a set of vertices S from the rest of the graph, then all of the vertices in S will lie on a line from a to b in a spring embedding of the graph formed by fixing a face that does not intersect S.

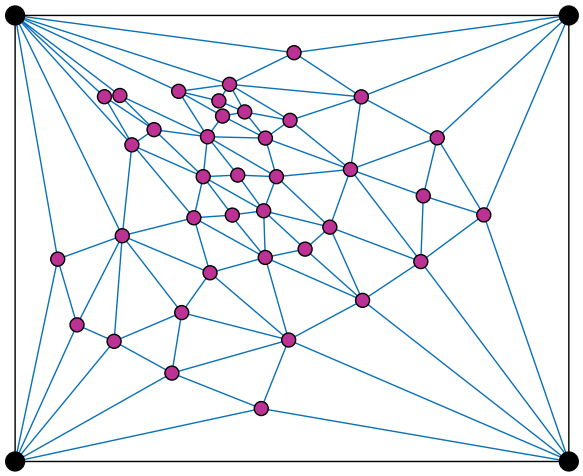

FIGURE 2. A drawing of a planar graph by a spring embedding. The locations of the four vertices at the corners of the square were fixed. The locations of every other vertex were chosen to minimize the potential energy of the spring network. The edges are drawn as straight lines.

We emphasize that this drawing was produced using only the knowledge of the connections between the vertices in the graph and does not exploit any geometric information. I find it amazing that the mathematics of spring networks realize our human desire to draw graphs without crossing edges.

2.2. **Spectral graph drawing.** A natural way to introduce spectral graph theory would be to present theorems relating combinatorial properties of a graph to the eigenvalues of its Laplacian matrix. Instead, I will try to build your intuition for why eigenvalues and eigenvectors of matrices associated with graphs should contain combinatorial information by showing how they can be used to produce nice pictures of graphs. I will also explain Cheeger's inequality, which relates the eigenvalues of the Laplacian to the sizes of boundaries of sets of vertices in the graph. For more theorems, I recommend the books of Biggs [Big74], Chung [Chu97], Godsil [God93], and Godsil and Royle [GR01].

We begin by considering the problem of drawing a graph on the line. Given a graph with n vertices, we might want to assign each vertex a number in $\{1, \ldots, n\}$. There are many reasons we might wish to do this: the first is that we need to choose some ordering on the vertices if we are going to display matrices associated with the graph. Many graph computations require an ordering of the vertices, and these often perform better if the pairs of vertices connected by edges appear close to each other. We will consider the motivation of drawing the graph. In this case, we will map the vertices of the graph to the corners of a regular n-gon. The picture we draw this way will better reveal the structure of the graph if the edges connect vertices that are close together.

For example, Figure 3 is a drawing of part of my social network from Facebook. Each vertex represents one of my friends, and I have drawn edges between people who are friends with each other.

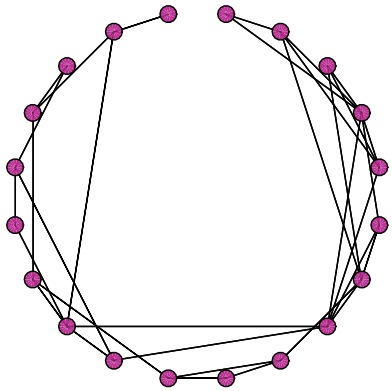

FIGURE 3. A drawing of a graph using one eigenvector wrapped around an n-gon.

An attempt to draw the best picture of this form leads us to a natural computational problem: assign distinct integers in $\{1, \ldots, n\}$ to the vertices of a graph so that the edges are as short as possible. Unfortunately, we do not know of any way of formalizing this problem that leads to something we can compute efficiently. For example, the problems of minimizing the maximum edge length or the sum of the edge lengths are NP-complete [Pap76].

Instead, we will consider a *relaxation* of this problem suggested by Hall [Hal70]. In this relaxation, we will map each vertex of the graph to a real number. I will view this mapping as a function $x : V \to \mathbb{R}$. Hall suggested minimizing the sum of the squares of the lengths of the edges: $q_G(x)$. To make the minimum unique, we need to impose two constraints. The first,

$$\sum_{a \in V} x(a) = 0, \tag{1}$$

fixes the average of the numbers. The second,

$$\sum_{a \in V} x(a)^2 = 1, \tag{2}$$

fixes the scale.

The Courant–Fischer characterization of the eigenvalues of a symmetric matrix tells us that the minimum of $q_G(x)$ subject to (1) and (2) is achieved by an eigenvector of the second-smallest eigenvalue of the Laplacian matrix of the graph. Let $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ be the eigenvalues of L_G, and let $v_1, \ldots, v_n$ be corresponding eigenvectors, normalized to have norm 1. As L_G is symmetric, we may assume that the eigenvectors are orthogonal. The Courant–Fischer theorem tells us that

$$\lambda_i = \min_{x : v_j^T x = 0 \text{ for } 1 \leq j < i} \frac{x^T L_G x}{x^T x},$$

and that v_i is a vector on which this minimum is achieved. The smallest eigenvalue of a Laplacian matrix is 0, and the corresponding eigenvectors are the constant vectors. So, the condition (1) is equivalent to $v_1^T x = 0$. If G is connected, then $\lambda_2 > 0$. We usually ignore the first eigenvalue of 0, and call $\lambda_2, \ldots, \lambda_n$ the eigenvalues of G.

Provided that $v_2(a) \neq v_2(b)$ for all distinct vertices a and b, we can use v_2 to obtain an ordering on the vertices. I drew the picture in Figure 3 by computing v_2, using it to obtain an ordering of the vertices (from low to high) and then placing the vertices at the corners of the regular n-gon in this order. For a proof that this procedure produces the "right" ordering for certain graphs, I refer the reader to [ABH98].

To obtain more flexibility in how we draw the graph, we can assign every vertex a pair of real numbers: x and y coordinates in the plane. In this case, Hall suggested minimizing

$$\sum_{(a,b)\in E} \left\| \begin{pmatrix} x(a) \\ y(a) \end{pmatrix} - \begin{pmatrix} x(b) \\ y(b) \end{pmatrix} \right\|^2 = q_G(x) + q_G(y)$$

subject to

$$\sum_a x(a) = 0, \qquad \sum_a y(a) = 0, \tag{3}$$

$$\sum_a x(a)^2 = 1, \qquad \sum_a y(a)^2 = 1, \tag{4}$$

and

$$\sum_a x(a)y(a) = 0. \tag{5}$$

Conditions (3) and (4) are the same as those we imposed when we drew the graph in the line. Condition (5) forces x to be different from y. If we did not include it, we would obtain the solution $x = y = v_2$. With this condition, Hall proves that the solutions to this problem are given by $x = v_2$, $y = v_3$, and rotations of this solution. This is how we chose the coordinates of the vertices in the right-hand drawing in Figure 1. The edges were then drawn as straight lines. This procedure does not work with every graph, but it often provides a good start. One can often produce a nice drawing of a graph by massaging the coordinates provided by the eigenvectors. In Figure 4 we give one more intriguing example of this phenomenon. The image on the left was generated by sampling 500 points at random in the unit square, computing their Delaunay triangulation, and drawing the edges of the triangles. We form a graph by taking the union of the edges in the triangles. The image on the right is the spectral drawing of this graph. While it is not a planar drawing, the

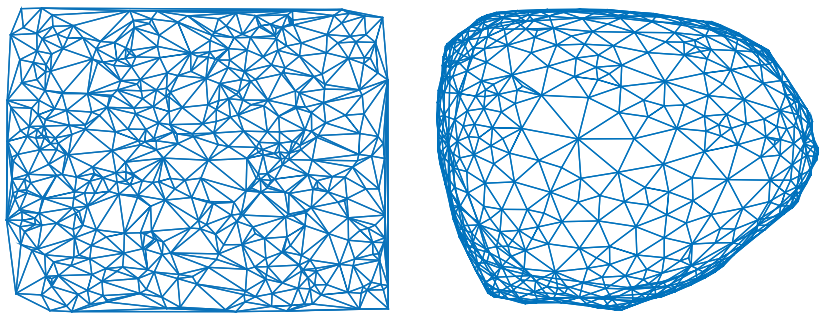

FIGURE 4. A Delaunay triangulation of random points in the square and a spectral drawing of the resulting graph.

violations of planarity only appear on the boundary. This is a robust experimental observation, and I would love to find a theorem that explains it.

Unfortunately, many of the graphs that we encounter cannot be drawn nicely. If you try to formalize the notion of a "nice" drawing, you will probably require most of the edges to be short and that the vertices not be too closely packed together. A nice drawing should reveal some structure of a graph. A random graph will not have structure, and so should not have a nice drawing. In fact, many of the graphs we encounter in real life, such as large social network graphs, cannot be drawn nicely. If λ_2 is large, say much larger than $10/\sqrt{n}$, then there is no way to draw the graph in the plane with most of the edges short and the vertices well distributed. I often ask my students to make this statement formal and prove it.

2.3. **Boundaries of sets.** Another way to understand the structure of a graph is to divide it into pieces. We usually do this by partitioning the vertices into sets so that most of the edges of the graph connect vertices in the same set. Such a partition of a social network provides a division of the people in the network into communities. Conversely, we can try to identify communities by searching for the best partitions. If we just want to identify one community, then we may just focus on that set of vertices and ignore how the partition divides the others.

Given a subset of vertices $S \subset V$, the *boundary* of S, written ∂S, is defined to be the set of edges with exactly one endpoint in S: $\{(a,b) \in E : a \in S, b \notin S\}$. The weight of the boundary of S is the sum of the weights of these edges. The Laplacian allows us to measure the weights of the boundaries of sets of vertices. If χ_S is the characteristic vector of S (1 inside S and 0 outside), then the value of the Laplacian quadratic form in χ_S measures the size of the boundary

$$\chi_S^T L_G \chi_S = \sum_{(a,b)\in E: a\in S, b\notin S} w_{a,b}.$$

For a graph in which every edge has weight 1, this is $|\partial S|$.

We think of S as providing a good community if its boundary is small relative to the size of S. We quantify this by measuring the *cut ratio* of S, written

$$\theta(S) = \frac{|\partial S|}{|S|}.$$

We should make one correction to this definition: if $|S| > |V|/2$, then we should think of the complement of S as providing the community instead of S, and then divide by $|V| - |S|$ instead.

By plugging the characteristic vector of S into the Laplacian quadratic form and by treating eigenvalues as the solution of optimization problems as we did in the previous section, one can easily show that

$$\theta(S) \stackrel{\text{def}}{=} \frac{|\partial S|}{\min\left(|S|, |V| - |S|\right)} \geq \lambda_2/2, \tag{6}$$

where λ_2 is the second-smallest eigenvalue of L_G. To see this, consider the vector $x \stackrel{\text{def}}{=} \chi_S - |S|/|V|$. The vector x is orthogonal to the constant vectors, and so the Courant–Fischer theorem tells us that

$$\lambda_2 \leq \frac{x^T L_G x}{x^T x}.$$

Inequality (6) then follows from the observations that

$$x^T L_G x = \chi_S^T L_G \chi_S$$

and

$$x^T x \geq \min\left(|S|, |V| - |S|\right)/2.$$

This provides a way of proving that a graph for which λ_2 is large does not have a nice drawing: every graph that has a nice drawing has a subset of vertices S for which $\theta(S)$ is small.

This result is the motivation for *spectral partitioning heuristics* that use eigenvectors of matrices associated with graphs to find sets of vertices of low cut ratio. The most basic of these considers sets of vertices of the form

$$S_t \stackrel{\text{def}}{=} \{a : v_2(a) \leq t\},$$

where v_2 is the eigenvector of λ_2 and t is a threshold that determines the size of the set. Discrete analogs of Cheeger's inequality [Che70, Dod84, AM85, LS88, SJ89, Var85] imply that if a graph has a set of small cut ratio, then one can find it this way. If we choose t to minimize $\theta(S_t)$, then we obtain a set that satisfies

$$\theta(S_t) \leq \sqrt{2 d_{\max} \lambda_2},$$

where $d_{\max}$ is the maximum number of edges attached to a vertex in G. Cleaner bounds that do not involve terms like $d_{\max}$ can be obtained if we count edges by their weights, count vertices by the sum of the weights of edges attached to them, and instead consider the eigenvalues of the *normalized Laplacian*, $D^{-1/2} L_G D^{-1/2}$, where D is the diagonal matrix of weighted degrees of vertices.

I study and teach spectral graph theory for three reasons. The first is that it contains many beautiful theorems. The second is that we can compute good approximations of the eigenvalues and eigenvectors of a Laplacian matrix very quickly [ST14]. Thus, they provide computationally tractable analyses and can serve as the basis of many useful heuristics. The third is that they have many generalizations: semidefinite programming generalizes optimization with eigenvalues, and eigenvectors are just one of many objects that may be computed as fixed points of natural iterative procedures.

3. Sparsification

We will say that a graph G is an ϵ-approximation of a graph H with the same vertex set if for all $x \in \mathbb{R}^V$,

$$(1+\epsilon) q_H(x) \geq q_G(x) \geq (1+\epsilon)^{-1} q_H(x). \tag{7}$$

We may express this condition in a linear algebraic manner by introducing the notation $A \succcurlyeq B$ for symmetric matrices A and B to indicate that $x^T A x \geq x^T B x$ for all vectors x. This is equivalent to saying that $A - B$ had no negative eigenvalues. That is, it is a positive semidefinite matrix. With this notation, (7) becomes

$$(1+\epsilon) L_H \succcurlyeq L_G \succcurlyeq (1+\epsilon)^{-1} L_H.$$

For small ϵ this is a very strong condition. Among other things, it implies that L_G and L_H have approximately the same eigenvalues, that the weight of the boundary of every set of vertices is approximately the same in G and H, that the effective resistance between every pair of vertices is approximately the same in G and H,

and that linear equations in L_G have approximately the same solutions as linear equations in L_H. The last statement can be expressed algebraically as

$$(1+\epsilon)L_H^+ \succcurlyeq L_G^+ \succcurlyeq (1+\epsilon)^{-1}L_H^+,$$

where L_G^+ denotes the pseudo-inverse of L_G: an operator that acts as the inverse of L_G on its span. In terms of numerical linear algebra, this says that L_H is a good preconditioner for L_G. For more implications of this approximation, I refer the reader to [ST11, BSST13].

A sparse graph is a graph in which the number of edges is a small multiple of the number of vertices. Every graph may be approximated by a sparse graph, where the number of edges in the sparse graph depends on the quality of the approximation. The strongest result of this form that we presently know comes from the following theorem of [BSS12].

Theorem 3.1. *For every weighted graph G on n vertices and every $\epsilon > 0$, there is a weighted graph H having at most $\lceil n/\epsilon^2 \rceil$ edges so that*

$$(1+\epsilon)^2 L_H \succcurlyeq L_G \succcurlyeq (1-\epsilon)^2 L_H.$$

The proof of this theorem is purely linear algebraic and relies on the association of vectors with the edges of a graph. We define the *vector associated with* edge (a,b) to be the vector

$$u_{a,b} \stackrel{\text{def}}{=} e_a - e_b,$$

where e is the elementary unit vector in direction a. That is, $u_{a,b}$ has 1 in position a, -1 in position b, and is zero everywhere else. For a vector $x \in \mathbb{R}^V$,

$$x(a) - x(b) = u_{a,b}^T x,$$

and thus

$$(x(a) - x(b))^2 = (u_{a,b}^T x)^2 = x^T (u_{a,b} u_{a,b}^T) x.$$

So, we can write L_G as

$$\sum_{(a,b)\in E} w_{a,b} u_{a,b} u_{a,b}^T.$$

In [BSS12], we derive Theorem 3.1 as a consequence of the following theorem about collections of vectors.

Theorem 3.2. *Let $u_1, \ldots, u_m$ be vectors in $\mathbb{R}^n$, and let $\epsilon > 0$. Then, there exists a subset $S \subseteq \{1, \ldots, m\}$ of size at most $\lceil n/\epsilon^2 \rceil$ and real numbers $s_i > 0$ so that for*

$$A = \sum_{i=1}^m u_i u_i^T \qquad \text{and} \qquad B = \sum_{i\in S} s_i u_i u_i^T,$$

$$(1+\epsilon)^2 B \succcurlyeq A \succcurlyeq (1-\epsilon)^2 B.$$

When we use Theorem 3.2 to prove Theorem 3.1, we obtain a graph H whose edges are a subset of the edges of G, although with different weights. Theorem 3.1 would be interesting even if the edges used in H were not a subset of the edges of G. However, Theorem 3.2 would be uninteresting if we did not require the vectors from which we form B to be a subset of those used to form A, as one can always use the eigenvector decomposition to write a positive semidefinite matrix as a sum of outer products of n vectors.

3.1. **Complete graphs and expanders.** Even the problem of sparsifying the complete graph is interesting. Recall that the complete graph on n vertices is the graph with every possible edge. If K is the complete graph on n vertices and x is orthogonal to the constant vectors, then

$$q_K(x) = n \|x\|^2 .$$

So, a sparse approximation of the complete graph is a graph whose eigenvalues are all close to n. Sparse graphs whose eigenvalues are concentrated in this way are called *expander graphs* (see [HLW06]) and they have proved incredibly useful in computer science and combinatorics. They are called *expanders* because all sets of vertices in these graphs have unusually large boundaries. Most constructions of expander graphs produce regular, unweighted graphs. The Ramanujan graphs, first constructed by Margulis [Mar88] and Lubotzky, Phillips and Sarnak [LPS88], are the best possible unweighted, regular approximations of complete graphs: they are infinite families of d-regular graphs with all eigenvalues differing from d by at most $2\sqrt{d-1}$. If we assign the weight n/d to every edge, then all their eigenvalues are close to n. Alon and Boppana (see [Nil91]) prove that no regular unweighted graph approximates a complete graph significantly better than the Ramanujan graphs: they prove that for every $\epsilon > 0$, every sufficiently large d-regular graph has $\lambda_2 \leq d - 2\sqrt{d-1} + \epsilon$. We do not know if there are better approximations of the complete graph by weighted or irregular graphs of the same number of edges. We conjecture that there are not. See [BSS12, Proposition 4.2] for some bounds.

Theorem 3.1 says that every graph can be approximated almost as well as Ramanujan graphs approximate complete graphs: less than twice as many edges are required to obtain the same quality of approximation. We do not know if better approximations of arbitrary graphs exist.

A natural way to try to form a sparse approximation of the complete graph is to choose a random subset of its edges. A good approximation is obtained if one chooses $O(n \log n)$ edges at random. However, the "Coupon Collector" phenomenon tells us that choosing fewer edges does not work: if one chooses fewer than $n \ln n$ edges, then the resulting random graph is likely to have a vertex of degree zero. This will result in it having $\lambda_2 = 0$, and so it will be a poor approximation of the complete graph. Vertices of the wrong degree are essentially the only obstacle to approximating the complete graph by a random graph. Friedman [Fri08] proved that if we choose a random graph in which every vertex has degree d, then it is probably very close to being a Ramanujan graph.

In [SS11] we show that, by carefully choosing the probabilities with which we sample edges, we can use a similar random sampling procedure to sparsify any graph. We include edge (a, b) in H with probability proportional to

$$q_{a,b} \stackrel{\text{def}}{=} w_{a,b} u_{a,b}^T L_G^+ u_{a,b},$$

where L_G^+ is the pseudo-inverse of L_G. $L_G^+ u_{a,b}$ is the vector x of sum zero so that $L_G x = u_{a,b}$. The quantity $u_{a,b}^T L_G^+ u_{a,b}$ is the *effective resistance* of edge (a, b), which we described in Section 2.1. The sum of the quantities $q_{a,b}$ may be shown to equal $n - 1$. We decide whether or not to include each edge (a, b) in the graph independently, including it with probability

$$p_{a,b} \stackrel{\text{def}}{=} \min(C(\log n) q_{a,b}/\epsilon^2, 1),$$

where C is some absolute constant. If we do include edge (a,b), then we increase its weight by dividing it by $p_{a,b}$. Thus, if we view not including an edge as setting its weight to zero, then the expected weight of every edge in H is its original weight in G. This implies that the expectation of the Laplacian of H is the Laplacian of G:

$$\mathbb{E}\, L_H = \sum_{(a,b)\in E} p_{a,b} \left(\frac{w_{a,b}}{p_{a,b}} u_{a,b} u_{a,b}^T \right) = \sum_{(a,b)\in E} w_{a,b} u_{a,b} u_{a,b}^T = L_G.$$

This procedure will result in a graph with an expected number of edges equal to $O(n \log n/\epsilon^2)$. One can use recent concentration bounds for random matrices, such as [Tro12, Theorem 1.1], to prove that the resulting Laplacian matrix is an ϵ approximation of the original. It is worth pointing out that Tropp's theorem has nothing to do with graphs; it is just a statement about a sum of random matrices. The original analysis of this procedure in [SS11] instead used Rudelson's concentration bounds [Rud99] to analyze a variant of this procedure that sampled edges with replacement.

Both analyses begin by considering a convenient rescaling of the problem. One can show that if M is nonsingular, then

$$A \preccurlyeq B \quad \Longleftrightarrow \quad MAM^T \preccurlyeq MBM^T.$$

So, we may prove that L_H is an ϵ-approximation of L_G by proving that

$$(1+\epsilon) L_G^{+/2} L_H L_G^{+/2} \succcurlyeq \Pi \succcurlyeq (1+\epsilon)^{-1} L_G^{+/2} L_H L_G^{+/2},$$

where

$$\Pi = L_G^{+/2} L_G L_G^{+/2}$$

is the projection onto the span of L_G, and should be thought of as the identity, and $L^{+/2}$ is the square root of the pseudo-inverse of L_G. Under this rescaling we are led to consider the vectors

$$v_{a,b} \stackrel{\text{def}}{=} w_{a,b}^{1/2} L_G^{+/2} u_{a,b},$$

and to seek a subset S and scaling factors $s_{a,b}$ so that

$$\sum_{(a,b)\in S} s_{a,b} v_{a,b} v_{a,b}^T \approx_\epsilon \Pi.$$

These vectors $v_{a,b}$ are special in that, within their span, the sum of their outer products is the identity

$$\sum v_{a,b} v_{a,b}^T = \sum w_{a,b} L_G^{+/2} u_{a,b} u_{a,b}^T L_G^{+/2} = \sum w_{a,b} L_G^{+/2} L_G L_G^{+/2} = \Pi.$$

The proof of Theorem 3.2 operates on the same normalization but has a very different flavor: it analyzes an algorithmic procedure that adds vectors to S one by one. The proof of Theorem 3.2 does not provide any control of the scaling factors s_i. They are proved to exist, and a procedure is given for computing them. This is rather different from the scaling given by the random sampling procedure, in which the scaling factors are the reciprocals of the probabilities with which edges are included. As we will explain in the next section, proving a variant of Theorem 3.2 with the right scaling factors is almost equivalent to proving Weaver's conjecture KS_2.

4. Weaver's conjecture and the Kadison–Singer Problem

Weaver's conjecture KS_2 concerns a collection of complex vectors $u_1, \ldots, u_m$ such that

$$\sum_{i=1}^{m} u_i u_i^* = I. \tag{8}$$

As these vectors are complex, we multiply each by its conjugate transpose. But, for most purposes, it suffices to consider the outer products of real vectors with their transposes. Collections of vectors that satisfy (8) are said to be in *isotropic position*, and are also called a *Parseval frame*. The sum in this expression is also known as a *decomposition of the identity*.

We would like to know conditions under which a set of vectors in isotropic position is guaranteed to contain a subset whose sum of outer products approximates half the identity. If

$$\sum_{i \in S} u_i u_i^*$$

is an ϵ-approximation of $I/2$, then all of the eigenvalues of the matrix in this sum must lie between $1/2(1+\epsilon)$ and $(1+\epsilon)/2$. The most obvious obstacle to this happening is if one of the vectors u_i has large norm. For example, if u_i has norm 1, then the sum will have an eigenvalue of 1 if $i \in S$ and an eigenvalue of 0 if $i \notin S$. Weaver conjectured [Wea04] that vectors of large norm are the only obstacle.

Conjecture 4.1 ([Wea04])**.** *There are positive constants α and ϵ so that for every collection of vectors $u_1, \ldots, u_m$ in isotropic position such that $\|u_i\| \leq \alpha$ for all i, there exists a subset $S \subseteq \{1, \ldots, m\}$ so that all eigenvalues of*

$$\sum_{i \in S} u_i u_i^*$$

lie between ϵ and $1 - \epsilon$.

Using results of Akemann and Anderson [AA91], Weaver proved that the truth of this conjecture would imply a positive resolution to the Kadison–Singer Problem [KS59].

A natural approach to proving this conjecture would be to select the subset S uniformly at random. One can use Tropp's matrix concentration bound [Tro12, Theorem 1.1] to prove that this would work if every vector u_i had norm at most $c/\sqrt{\log n}$, where n is the dimension of the space in which the vectors lie and c is some constant. However, Conjecture 4.1 requires a constant bound on the norms of the vectors that is independent of the dimension. With only a constant bound on the norms of the vectors, we can form collections of vectors for which random subsets almost always fail. For example, consider the collection of vectors that contains k vectors of form $(1/k)e_i$, for each elementary unit vector e_i. If k is held constant while the dimension grows and one chooses a subset S of the vectors at random, then there will probably be an i so that the set S contains all of the vectors in direction e_i and there will probably be a j so that S contains none of the vectors in direction e_j. Either of these conditions would prevent the sum of the outer products of the vectors in S from approximating half the identity.

Of course, it would suffice to show that a random set S works with some nonzero probability. Examples of such arguments appear in [BFU94] and [FM99], both of

which show that it is possible to partition the edges of an expander graph into two disjoint expander graphs. To see how this problem is related, observe that the vectors associated with the edges of an expander are very close to being in isotropic position. One can bring them into isotropic position by projecting orthogonally to the constant vectors, and then applying a linear transformation that only changes the lengths of the vectors slightly. This produces a set of vectors in isotropic position that all have approximately the same norm. Frieze and Malloy use the Lovász Local Lemma to prove that there is a nonzero probability that the vectors in a random subset approximate half the identity.

Theorem 3.2 looks like it comes close to proving Conjecture 4.1. It seems that the most difficult cases of the conjecture should involve vectors that are as large as possible, and so they should all have the same norm. If we set ϵ so that $\lceil n/\epsilon^2 \rceil$ equals $m/2$, we would hope that the s_i would all be equal to or close to 2. This is what the random sampling procedure suggests, it is what we expected would be true, and it would imply Conjecture 4.1. Unfortunately, the proof of Theorem 3.2 does not provide any such guarantees about the scaling factors s_i.

In [MSS15b] we use a very different argument to prove the following theorem, which implies a strong version of Conjecture 4.1. This in turn provides a positive resolution of the Kadison–Singer Problem.

Theorem 4.2. *For every constant $\alpha > 0$ and every collection of vectors $u_1, \ldots, u_m$ in isotropic position such that $\|u_i\| \leq \alpha$ for all i, there exists a subset $S \subseteq \{1, \ldots, m\}$ so that all eigenvalues of*

$$\sum_{i \in S} u_i u_i^*$$

lie between $1/2 - \beta$ and $1/2 + \beta$ for

$$\beta = \sqrt{2}\alpha^{1/4} + \alpha^{1/2}.$$

The next section explains the main technique used in the proof of this theorem.

4.1. **Interlacing families of polynomials.** Given a subset $S \subseteq \{1, \ldots, m\}$, let $p_S(z)$ be the characteristic polynomial of the matrix

$$\sum_{i \in S} u_i u_i^*.$$

Theorem 4.2 requires a set S such that all the roots of $p_S(z)$ are between $1/2 - \beta$ and $1/2 + \beta$. Let $\bar{S}$ denote the complement of S. As the vectors are in isotropic position, the roots of $p_{\bar{S}}(z)$ are one minus the roots of $p_S(z)$. So, it suffices to find a set S so that all of the roots of both $p_S(z)$ and $p_{\bar{S}}(z)$ are at most $1/2 + \beta$. We define

$$q_S(z) \stackrel{\text{def}}{=} p_S(z) p_{\bar{S}}(z),$$

and seek a set S so that all of the roots of $q_S(z)$ are at most $1/2 + \beta$.

We do this by considering the expected polynomial obtained by choosing S uniformly at random:

$$Q_0(z) \stackrel{\text{def}}{=} \mathop{\mathbb{E}}_S q_S(z).$$

We prove

1. all of the roots of $Q_0(z)$ are real,
2. all of the roots of $Q_0(z)$ are less than $1/2 + \beta$, and
3. there is a set S so that the largest root of $q_S(z)$ is at most the largest root of $Q_0(z)$.

The first of these statements is somewhat surprising. Even though each polynomial $q_S(z)$ is real rooted, the average of real rooted polynomials does not need to be real rooted. It becomes less surprising if you are aware of the large family of combinatorially defined polynomials that turn out to be real rooted, such as the matching polynomials of graphs [HL72, God81].

The second statement is surprising because we have seen that there are sets of vectors in isotropic position for which almost every set S will produce a polynomial $q_S(z)$ whose largest root is 1. Substantial cancellation is required for the sum of these polynomials to have all roots close to 1/2.

The most surprising step in this argument is probably the last one. There is no particular reason to expect an upper bound on the roots of a sum of polynomials to provide an upper bound on the roots of a particular polynomial in the sum. It happens in this case because the polynomials that we are averaging form an *interlacing family*. To define this, we recall that two degree d polynomials $p_1(z)$ and $p_2(z)$ are said to *have a common interlacing* if there exist real numbers $r_0 \leq \cdots \leq r_d$ so that for every $1 \leq i \leq d$ both polynomials $p_1(z)$ and $p_2(z)$ have exactly one root in the interval $[r_{i-1}, r_i]$ (carefully treating roots of multiplicity greater than one). The terminology comes from the fact that the roots of both $p_1(z)$ and $p_2(z)$ interlace those of the polynomial $\prod(z - r_i)$. The following fact about polynomials that have a common interlacing is easy to prove.

Proposition 4.3 ([MSS15a, Lemma 4.2])**.** *If $q_1(z)$ and $q_2(z)$ are polynomials of the same degree with positive leading coefficient that have a common interlacing, then the largest root of $q_1(z)$ or $q_2(z)$ is at most the largest root of $q_1(z) + q_2(z)$.*

We cannot exploit this proposition directly, as the polynomials $q_S(z)$ do not all have common interlacings. Instead, for each $0 \leq k \leq n$ and each $R \subseteq \{1, \ldots, k\}$, we define the polynomials

$$Q_{k,R}(z) \stackrel{\text{def}}{=} \mathop{\mathbb{E}}_{T \subseteq \{k+1,\ldots,m\}} q_{R \cup T}(z).$$

That is, $Q_{k,R}(z)$ is the expectation of $q_S(z)$ when we choose S at uniformly at random subject to $S \cap \{1, \ldots, k\} = R$. We also prove that for every k and $R \subseteq \{1, \ldots, k\}$, the polynomials

$$Q_{k+1,R}(z) \qquad \text{and} \qquad Q_{k+1,R \cup \{k+1\}}(z)$$

have a common interlacing. We define sets of polynomials $q_S(z)$ that satisfy a condition like this to be an *interlacing family*. This condition implies that for every k and $R \subseteq \{1, \ldots, k\}$, there is a set $\widehat{R} \subseteq \{1, \ldots, k+1\}$ so that the largest root of $Q_{k+1,\widehat{R}}(z)$ is at most the largest root of $Q_{k,R}(z)$. An inductive argument then tells us that there is a set S so that the largest root of $Q_{n,S}(z)$ is at most the largest root of $Q_{0,\emptyset}(z)$. As $Q_{n,S}(z) = q_S(z)$ and $Q_{0,\emptyset}(z) = Q_0(z)$, the set S is the one that we require.

5. Conclusions

There are many interesting connections between these topics that I have not been able to cover, such as the relation between the Kadison–Singer Problem and the construction of Ramanujan graphs [MSS15a, MSS15c]. To learn more about related topics, I suggest the survey [MSS14] or the blog posts of Srivastava [Sri13] and Tao [Tao13]. To learn more about sparsification, I recommend the surveys [BSST13, Nao12]. For another perspective on the solution of the Kadison–Singer Problem, I suggest [Brä14]. I maintain a web page [Spi] at which I keep pointers to the latest developments in sparsification.

About the author

Daniel Spielman is professor of computer science, mathematics, and applied mathematics at Yale University. While his first love was linear algebra, he also studies algorithms, machine learning, network science, combinatorics, and digital communication. His many awards include a MacArthur Fellowship.

References

[AA91] C. A. Akemann and J. Anderson, *Lyapunov theorems for operator algebras*, Mem. Amer. Math. Soc. **94** (1991), no. 458, iv+88, DOI 10.1090/memo/0458. MR1086563

[ABH98] J. E. Atkins, E. G. Boman, and B. Hendrickson, *A spectral algorithm for seriation and the consecutive ones problem*, SIAM J. Comput. **28** (1999), no. 1, 297–310, DOI 10.1137/S0097539795285771. MR1630473

[AM85] N. Alon and V. D. Milman, λ_1, *isoperimetric inequalities for graphs, and superconcentrators*, J. Combin. Theory Ser. B **38** (1985), no. 1, 73–88, DOI 10.1016/0095-8956(85)90092-9. MR782626

[BFU94] A. Z. Broder, A. M. Frieze, and E. Upfal, *Existence and construction of edge-disjoint paths on expander graphs*, SIAM J. Comput. **23** (1994), no. 5, 976–989, DOI 10.1137/S0097539792232021. MR1293269

[Big74] N. Biggs, *Algebraic graph theory*, Cambridge University Press, London, 1974. Cambridge Tracts in Mathematics, No. 67. MR0347649

[Brä14] Petter Brändén. *Hyperbolic polynomials and the Marcus-Spielman-Srivastava theorem.* arXiv preprint arXiv:1412.0245, 2014.

[BSS12] J. Batson, D. A. Spielman, and N. Srivastava, *Twice-Ramanujan sparsifiers*, SIAM J. Comput. **41** (2012), no. 6, 1704–1721, DOI 10.1137/090772873. MR3029269

[BSST13] Joshua Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. *Spectral sparsification of graphs: Theory and algorithms.* Commun. ACM, **56** (2013), no. 8. 87–94.

[Che70] J. Cheeger, *A lower bound for the smallest eigenvalue of the Laplacian*, in Problems in analysis (Papers dedicated to Salomon Bochner, 1969), Princeton Univ. Press, Princeton, N. J., 1970, pp. 195–199. MR0402831

[Chu97] F. R. K. Chung, *Spectral graph theory*, CBMS Regional Conference Series in Mathematics, vol. 92, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 1997. MR1421568

[Dod84] J. Dodziuk, *Difference equations, isoperimetric inequality and transience of certain random walks*, Trans. Amer. Math. Soc. **284** (1984), no. 2, 787–794, DOI 10.2307/1999107. MR743744

[FM99] A. M. Frieze and M. Molloy, *Splitting an expander graph*, J. Algorithms **33** (1999), no. 1, 166–172, DOI 10.1006/jagm.1999.1023. MR1712697

[Fri08] J. Friedman, *A proof of Alon's second eigenvalue conjecture and related problems*, Mem. Amer. Math. Soc. **195** (2008), no. 910, viii+100, DOI 10.1090/memo/0910. MR2437174

[God81] C. D. Godsil, *Matchings and walks in graphs*, J. Graph Theory **5** (1981), no. 3, 285–297, DOI 10.1002/jgt.3190050310. MR625070

[God93] C. D. Godsil, *Algebraic combinatorics*, Chapman and Hall Mathematics Series, Chapman & Hall, New York, 1993. MR1220704

[GR01] C. Godsil and G. Royle, *Algebraic graph theory*, Graduate Texts in Mathematics, vol. 207, Springer-Verlag, New York, 2001. MR1829620

[Hal70] K. M. Hall. *An r-dimensional quadratic placement algorithm*. Management Science, **17** (1970), 219–229.

[HL72] O. J. Heilmann and E. H. Lieb, *Theory of monomer-dimer systems*, Comm. Math. Phys. **25** (1972), 190–232. MR0297280

[HLW06] S. Hoory, N. Linial, and A. Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.) **43** (2006), no. 4, 439–561 (electronic), DOI 10.1090/S0273-0979-06-01126-8. MR2247919

[KR93] D. J. Klein and M. Randić, *Resistance distance*, J. Math. Chem. **12** (1993), no. 1-4, 81–95, DOI 10.1007/BF01164627. Applied graph theory and discrete mathematics in chemistry (Saskatoon, SK, 1991). MR1219566

[KS59] R. V. Kadison and I. M. Singer, *Extensions of pure states*, Amer. J. Math. **81** (1959), 383–400. MR0123922

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277, DOI 10.1007/BF02126799. MR963118

[LS88] G. F. Lawler and A. D. Sokal, *Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality*, Trans. Amer. Math. Soc. **309** (1988), no. 2, 557–580, DOI 10.2307/2000925. MR930082

[Mar88] G. A. Margulis, *Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators* (Russian), Problemy Peredachi Informatsii **24** (1988), no. 1, 51–60; English transl., Problems Inform. Transmission **24** (1988), no. 1, 39–46. MR939574

[MSS14] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava, *Ramanujan graphs and the solution of the Kadison-Singer problem*. In Proceedings of the International Congress of Mathematicians, volume III, pages 363–386. Kyung Moon SA, 2014.

[MSS15a] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families I: bipartite Ramanujan graphs of all degrees*, Ann. of Math. (2) **182** (2015), no. 1, 307–325, DOI 10.4007/annals.2015.182.1.7. MR3374962

[MSS15b] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families II: Mixed characteristic polynomials and the Kadison–Singer problem*, Ann. of Math. (2) **182** (2015), no. 1, 327–350, DOI 10.4007/annals.2015.182.1.8. MR3374963

[MSS15c] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families IV: bipartite Ramanujan graphs of all sizes*, 2015 IEEE 56th Annual Symposium on Foundations of Computer Science—FOCS 2015, IEEE Computer Soc., Los Alamitos, CA, 2015, pp. 1358–1377. MR3473375

[Nao12] A. Naor, *Sparse quadratic forms and their geometric applications [following Batson, Spielman, and Srivastava]*, Astérisque **348** (2012), Exp. No. 1033, viii, 189–217. Séminaire Bourbaki: Vol. 2010/2011. Exposés 1027–1042. MR3050716

[Nil91] A. Nilli, *On the second eigenvalue of a graph*, Discrete Math. **91** (1991), no. 2, 207–210, DOI 10.1016/0012-365X(91)90112-F. MR1124768

[Pap76] Ch. H. Papadimitriou, *The NP-completeness of the bandwidth minimization problem* (English, with German summary), Computing **16** (1976), no. 3, 263–270. MR0411243

[Rud99] M. Rudelson, *Random vectors in the isotropic position*, J. Funct. Anal. **164** (1999), no. 1, 60–72, DOI 10.1006/jfan.1998.3384. MR1694526

[SJ89] A. Sinclair and M. Jerrum, *Approximate counting, uniform generation and rapidly mixing Markov chains*, Inform. and Comput. **82** (1989), no. 1, 93–133, DOI 10.1016/0890-5401(89)90067-9. MR1003059

[Spi] Daniel Spielman. Laplacian linear equations, sparsification, local graph clustering, low-stretch spanning trees, etc. from `http://cs-www.cs.yale.edu/homes/spielman/precon/precon.html`.

[Sri13] N. Srivastava, *Windows on theory: Discrepancy, graphs, and the Kadison-Singer problem*, Asia Pac. Math. Newsl. **3** (2013), no. 4, 15–20. MR3156130

[SS11] D. A. Spielman and N. Srivastava, *Graph sparsification by effective resistances*, SIAM J. Comput. **40** (2011), no. 6, 1913–1926, DOI 10.1137/080734029. MR2863199

[ST11] D. A. Spielman and S.-H. Teng, *Spectral sparsification of graphs*, SIAM J. Comput. **40** (2011), no. 4, 981–1025, DOI 10.1137/08074489X. MR2825307

[ST14] D. A. Spielman and S.-H. Teng, *Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems*, SIAM J. Matrix Anal. Appl. **35** (2014), no. 3, 835–885, DOI 10.1137/090771430. MR3228466

[Tao13] Terence Tao. *What's new: Real stable polynomials and the Kadison-Singer problem*, 2013. from https://terrytao.wordpress.com/2013/11/04/real-stable-polynomials-and-the-kadison-singer-problem/.

[Tro12] J. A. Tropp, *User-friendly tail bounds for sums of random matrices*, Found. Comput. Math. **12** (2012), no. 4, 389–434, DOI 10.1007/s10208-011-9099-z. MR2946459

[Tut63] W. T. Tutte, *How to draw a graph*, Proc. London Math. Soc. (3) **13** (1963), 743–767. MR0158387

[Var85] N. Th. Varopoulos, *Isoperimetric inequalities and Markov chains*, J. Funct. Anal. **63** (1985), no. 2, 215–239, DOI 10.1016/0022-1236(85)90086-2. MR803093

[Wea04] N. Weaver, *The Kadison-Singer problem in discrepancy theory*, Discrete Math. **278** (2004), no. 1-3, 227–239, DOI 10.1016/S0012-365X(03)00253-X. MR2035401

Department of Computer Science, Yale University