13

The Laplacian of a Graph

The Laplacian is another important matrix associated with a graph, and the Laplacian spectrum is the spectrum of this matrix. We will consider the relationship between structural properties of a graph and the Laplacian spectrum, in a similar fashion to the spectral graph theory of previous chapters. We will meet Kirchhoff's expression for the number of spanning trees of a graph as the determinant of the matrix we get by deleting a row and column from the Laplacian. This is one of the oldest results in algebraic graph theory. We will also see how the Laplacian can be used in a number of ways to provide interesting geometric representations of a graph. This is related to work on the Colin de Verdière number of a graph, which is one of the most important recent developments in graph theory.

13.1 The Laplacian Matrix

Let σ be an arbitrary orientation of a graph X, and let D be the incidence matrix of X^σ. Then the *Laplacian* of X is the matrix $Q(X) = DD^T$. It is a consequence of Lemma 8.3.2 that the Laplacian does not depend on the orientation σ, and hence is well-defined.

Lemma 13.1.1 *Let X be a graph with n vertices and c connected components. If Q is the Laplacian of X, then* $\operatorname{rk} Q = n - c$.

Proof. Let D be the incidence matrix of an arbitrary orientation of X. We shall show that $\operatorname{rk} D = \operatorname{rk} D^T = \operatorname{rk} DD^T$, and the result then follows from Theorem 8.3.1. If $z \in \mathbb{R}^n$ is a vector such that $DD^T z = 0$, then

$z^T DD^T z = 0$. But this is the squared length of the vector $D^T z$, and hence we must have $D^T z = 0$. Thus any vector in the null space of DD^T is in the null space of D^T, which implies that $\operatorname{rk} DD^T = \operatorname{rk} D$. □

Let X be a graph on n vertices with Laplacian Q. Since Q is symmetric, its eigenvalues are real, and by Theorem 8.4.5, $\mathbb{R}^n$ has an orthogonal basis consisting of eigenvectors of Q. Since $Q = DD^T$, it is positive semidefinite, and therefore its eigenvalues are all nonnegative. We denote them by $\lambda_1(Q)$, $\dots$, $\lambda_n(Q)$ with the assumption that

$$\lambda_1(Q) \leq \lambda_2(Q) \leq \cdots \leq \lambda_n(Q).$$

We use $\lambda_i(X)$ as shorthand for $\lambda_i(Q(X))$, or simply λ_i when Q is clear from the context or unimportant. We will also use λ_∞ to denote λ_n. For any graph, $\lambda_1 = 0$, because $Q\mathbf{1} = 0$. By Lemma 13.1.1, the multiplicity of zero as an eigenvalue of Q is equal to the number of components of X, and so for connected graphs, λ_2 is the smallest nonzero eigenvalue. Much of what follows will concentrate on the information determined by this particular eigenvalue.

If X is a regular graph, then the eigenvalues of the Laplacian are determined by the eigenvalues of the adjacency matrix.

Lemma 13.1.2 *Let X be a regular graph with valency k. If the adjacency matrix A has eigenvalues $\theta_1, \dots, \theta_n$, then the Laplacian Q has eigenvalues $k - \theta_1, \dots, k - \theta_n$.*

Proof. If X is k-regular, then $Q = \Delta(X) - A = kI - A$. Thus every eigenvector of A with eigenvalue θ is an eigenvector of Q with eigenvalue $k - \theta$. □

This shows that if two regular graphs are cospectral, then they also have the same Laplacian spectrum. However, this is not true in general; the two graphs of Figure 8.1 have different Laplacian spectra.

The next result describes the relation between the Laplacian spectrum of X and the Laplacian spectrum of its complement $\overline{X}$.

Lemma 13.1.3 *If X is a graph on n vertices and $2 \leq i \leq n$, then $\lambda_i(\overline{X}) = n - \lambda_{n-i+2}(X)$.*

Proof. We start by observing that

$$Q(X) + Q(\overline{X}) = nI - J. \tag{13.1}$$

The vector $\mathbf{1}$ is an eigenvector of $Q(X)$ and $Q(\overline{X})$ with eigenvalue 0. Let x be another eigenvector of $Q(X)$ with eigenvalue λ; we may assume that x is orthogonal to $\mathbf{1}$. Then $Jx = 0$, so

$$nx = (nI - J)x = Q(X)x + Q(\overline{X})x = \lambda x + Q(\overline{X})x.$$

Therefore, $Q(\overline{X})x = (n - \lambda)x$, and the lemma follows. □

Note that $nI - J = Q(K_n)$; thus (13.1) can be rewritten as

$$Q(X) + Q(\overline{X}) = Q(K_n).$$

From the proof of Lemma 13.1.3 it follows that the eigenvalues of $Q(K_n)$ are n, with multiplicity $n-1$, and 0, with multiplicity 1. Since $K_{m,n}$ is the complement of $K_m \cup K_n$, we can use this fact, along with Lemma 13.1.3, to determine the eigenvalues of the complete bipartite graph. We leave the pleasure of this computation to the reader, noting only the result that the characteristic polynomial of $Q(K_{m,n})$ is

$$t(t-m)^{n-1}(t-n)^{m-1}(t-m-n).$$

We note another useful consequence of Lemma 13.1.3.

Corollary 13.1.4 *If X is a graph on n vertices, then $\lambda_n(X) \leq n$. If $\overline{X}$ has $\bar{c}$ connected components, then the multiplicity of n as an eigenvalue of $Q(X)$ is $\bar{c}-1$.* □

Our last result in this section is a property of the Laplacian that will provide us with a lot of information about its eigenvalues.

Lemma 13.1.5 *Let X be a graph on n vertices with Laplacian Q. Then for any vector x,*

$$x^T Q x = \sum_{uv \in E(X)} (x_u - x_v)^2.$$

Proof. This follows from the observations that

$$x^T Q x = x^T D D^T x = (D^T x)^T (D^T x)$$

and that if $uv \in E(X)$, then the entry of $D^T x$ corresponding to uv is $\pm(x_u - x_v)$. □

13.2 Trees

In this section we consider a classical result of algebraic graph theory, which shows that the number of spanning trees in a graph is determined by the Laplacian.

First we need some preparatory definitions. Let X be a graph, and let $e = uv$ be an edge of X. The graph $X \setminus e$ with vertex set $V(X)$ and edge set $E(X) \setminus e$ is said to be obtained by *deleting* the edge e. The graph X/e constructed by identifying the vertices u and v and then deleting e is said to be obtained by *contracting* e. Deletion and contraction are illustrated in Figure 13.1. If a vertex x is adjacent to both u and v, then there will be multiple edges between x and the newly identified vertex in X/e. Furthermore, if X itself has multiple edges, then any edges between

u and v other than e itself become loops on the newly identified vertex in X/e. Depending on the situation, it is sometimes possible to ignore loops, multiple edges, or both.

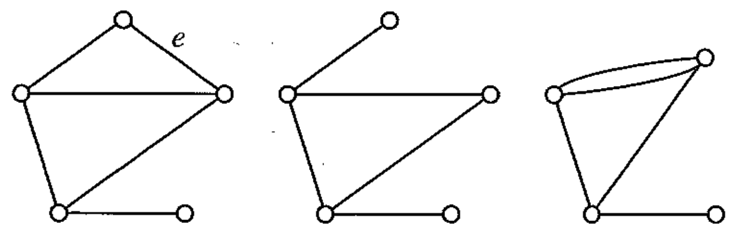

Figure 13.1. Graph Y, deletion $Y \setminus e$, and contraction Y/e

If M is a symmetric matrix with rows and columns indexed by the set V and if $S \subseteq V$, then let $M[S]$ denote the submatrix of M obtained by deleting the rows and columns indexed by elements of S.

Theorem 13.2.1 *Let X be a graph with Laplacian matrix Q. If u is an arbitrary vertex of X, then* $\det Q[u]$ *is equal to the number of spanning trees of X.*

Proof. We prove the theorem by induction on the number of edges of X.

Let $\tau(X)$ denote the number of spanning trees of X. If e is an edge of X, then every spanning tree either contains e or does not contain e, so we can count them according to this distinction. There is a one-to-one correspondence between spanning trees of X that contain e and spanning trees of X/e, so there are $\tau(X/e)$ such trees. Any spanning tree of X that does not contain e is a spanning tree of $X \setminus e$, and so there are $\tau(X \setminus e)$ of these. Therefore,

$$\tau(X) = \tau(X/e) + \tau(X \setminus e). \tag{13.2}$$

In this situation, multiple edges are retained during contraction, but we may ignore loops, because they cannot occur in a spanning tree.

Now, assume that $e = uv$, and let E be the $n \times n$ diagonal matrix with E_{vv} equal to 1, and all other entries equal to 0. Then

$$Q[u] = Q(X \setminus e)[u] + E,$$

from which we deduce that

$$\det Q[u] = \det Q(X \setminus e)[u] + \det Q(X \setminus e)[u, v]. \tag{13.3}$$

Note that $Q(X \setminus e)[u, v] = Q[u, v]$.

Assume that in forming X/e we contract u onto v, so that $V(X/e) = V(X) \setminus u$. Then $Q(X/e)[v]$ has rows and columns indexed by $V(X) \setminus \{u, v\}$ with the xy-entry being equal to Q_{xy}, and so we also have that $Q(X/e)[v] = Q[u, v]$.

Thus we can rewrite (13.3) as

$$\det Q[u] = \det Q(X \setminus e)[u] + \det Q(X/e)[v].$$

By induction, $\det Q(X\setminus e)[u] = \tau(X\setminus e)$ and $\det Q(X/e)[v] = \tau(X/e)$; hence (13.2) implies the theorem. □

It follows from Theorem 13.2.1 that $\det Q[u]$ is independent of the choice of the vertex u.

Corollary 13.2.2 *The number of spanning trees of K_n is n^{n-2}.*

Proof. This follows directly from the fact that $Q[u] = nI_{n-1} - J$ for any vertex u. □

If M is a square matrix, then denote by $M(i,j)$ the matrix obtained by deleting row i and column j from M. The *ij-cofactor* of M is the value

$$(-1)^{i+j} \det M(i,j).$$

The transposed matrix of cofactors of M is called the *adjugate* of M and denoted by $\operatorname{adj} M$. The ij-entry of $\operatorname{adj} M$ is the ji-cofactor of M. The most important property of the adjugate is that

$$M \operatorname{adj}(M) = (\det M)I.$$

If M is invertible, it implies that $M^{-1} = (\det M)^{-1} \operatorname{adj}(M)$. Theorem 13.2.1 implies that if Q is the Laplacian of a graph, then the diagonal entries of $\operatorname{adj}(Q)$ are all equal. The full truth is somewhat surprising: All of the entries of $\operatorname{adj}(Q)$ are equal.

Lemma 13.2.3 *Let $\tau(X)$ denote the number of spanning trees in the graph X and let Q be its Laplacian. Then* $\operatorname{adj}(Q) = \tau(X)J$.

Proof. Suppose that X has n vertices. Assume first that X is not connected, so that $\tau(X) = 0$. Then Q has rank at most $n-2$, so any submatrix of Q of order $(n-1) \times (n-1)$ is singular and $\operatorname{adj}(Q) = 0$.

Thus we may assume that X is connected. Then $\operatorname{adj}(Q) \neq 0$, but nonetheless $Q \operatorname{adj}(Q) = 0$. Because X is connected, $\ker Q$ is spanned by $\mathbf{1}$, and therefore each column of $\operatorname{adj}(Q)$ must be a constant vector. Since $\operatorname{adj}(Q)$ is symmetric, it follows that it is a nonzero multiple of J; now the result follows at once from Theorem 13.2.1. □

To prove the next result we need some information about the characteristic polynomial of a matrix. If A and B are square $n \times n$ matrices, then $\det(A+B)$ may be computed as follows. For each subset S of $\{1, \ldots, n\}$, let A_S be the matrix obtained by replacing the rows of A indexed by elements of S with the corresponding rows of B. Then

$$\det(A + B) = \sum_S \det A_S.$$

Applying this to $tI + (-A)$, we deduce that the coefficient of t^{n-k} in $\det(tI - A)$ is $(-1)^k$ times the sum of the determinants of the principal $k \times k$ submatrices of A. (This is a classical result, due to Laplace.)

Lemma 13.2.4 *Let X be a graph on n vertices, and let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of the Laplacian of X. Then the number of spanning trees in X is $\frac{1}{n}\prod_{i=2}^{n} \lambda_i$.*

Proof. The result clearly holds if X is not connected, so we may assume without loss that X is connected. Let $\phi(t)$ denote the characteristic polynomial $\det(tI - Q)$ of the Laplacian Q of X. The zeros of $\phi(t)$ are the eigenvalues of Q. Since $\lambda_1 = 0$, its constant term is zero and the coefficient of t is

$$(-1)^{n-1} \prod_{i=2}^{n} \lambda_i.$$

On the other hand, by our remarks just above, the coefficient of the linear term in $\phi(t)$ is

$$(-1)^{n-1} \sum_{u \in V(X)} \det Q[u].$$

This yields the lemma immediately. □

13.3 Representations

Define a *representation* ρ of a graph X in $\mathbb{R}^m$ to be a map ρ from $V(X)$ into $\mathbb{R}^m$. Informally, we think of a representation as the positions of the vertices in an m-dimensional drawing of a graph. Figure 13.2 shows a representation of the cube in $\mathbb{R}^3$.

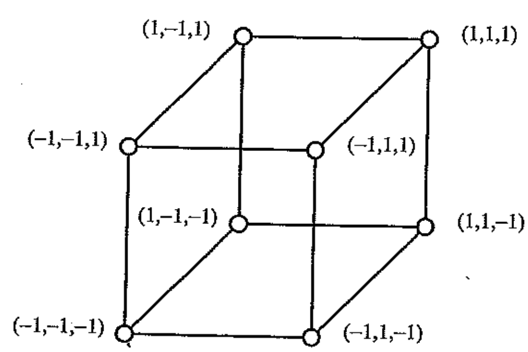

Figure 13.2. The cube in $\mathbb{R}^3$

We regard the vectors $\rho(u)$ as row vectors, and thus we may represent ρ by the $|V(X)| \times m$ matrix R with the images of the vertices of X as its rows.

Suppose then that ρ maps $V(X)$ into $\mathbb{R}^m$. We say ρ is *balanced* if

$$\sum_{u \in V(X)} \rho(u) = 0.$$

Thus ρ is balanced if and only if $\mathbf{1}^T R = 0$. The representation of Figure 13.2 is balanced. A balanced representation has its "centre of gravity" at the origin, and clearly we can translate any representation so that it is balanced without losing any information. Henceforth we shall assume that a representation is balanced.

If the columns of the matrix R are not linearly independent, then the image of X is contained in a proper subspace of $\mathbb{R}^m$, and ρ is just a lower-dimensional representation embedded in $\mathbb{R}^m$. Any maximal linearly independent subset of the columns of R would suffice to determine all the properties of the representation. Therefore, we will furthermore assume that the columns of R are linearly independent.

We can imagine building a physical model of X by placing the vertices in the positions specified by ρ and connecting adjacent vertices by identical springs. It is natural to consider a representation to be better if it requires the springs to be less extended. Letting $\|x\|$ denote the Euclidean length of a vector x, we define the *energy* of a representation ρ to be the value

$$\mathcal{E}(\rho) = \sum_{uv \in E(X)} \|\rho(u) - \rho(v)\|^2,$$

and hope that natural or good drawings of graphs correspond to representations with low energy. (Of course, the representation with least energy is the one where each vertex is mapped to the zero vector. Thus we need to add further constraints, to exclude this.)

We can go further by dropping the assumption that the springs are identical. To model this, let ω be a function from the edges of X to the positive real numbers, and define the energy $\mathcal{E}(\rho)$ of a representation ρ of X by

$$\mathcal{E}(\rho) = \sum_{uv \in E(X)} \omega_{uv} \|\rho(u) - \rho(v)\|^2,$$

where ω_{uv} denotes the value of ω on the edge uv. Let W be the diagonal matrix with rows and columns indexed by the edges of X and with the diagonal entry corresponding to the edge uv equal to ω_{uv}.

The next result can be viewed as a considerable generalization of Lemma 13.1.5.

Lemma 13.3.1 *Let ρ be a representation of the edge-weighted graph X, given by the $|V(X)| \times m$ matrix R. If D is an oriented incidence matrix* *for X, then*

$$\mathcal{E}(\rho) = \operatorname{tr} R^T D W D^T R.$$

Proof. The rows of $D^T R$ are indexed by the edges of X, and if $uv \in E(X)$, then the uv-row of $D^T R$ is $\pm(\rho(u) - \rho(v))$. Consequently, the diagonal entries of $D^T R R^T D$ have the form $\|\rho(u) - \rho(v)\|^2$, where uv ranges over the edges of X. Hence

$$\mathcal{E}(\rho) = \operatorname{tr} W D^T R R^T D = \operatorname{tr} R^T D W D^T R$$

as required. □

We may view $Q = DWD^T$ as a *weighted Laplacian*. If $uv \in E(X)$, then $Q_{uv} = -\omega_{uv}$, and for each vertex u in X,

$$Q_{uu} = \sum_{v \sim u} \omega_{uv}.$$

Thus $Q\mathbf{1} = 0$. Conversely, any symmetric matrix Q with nonpositive off-diagonal entries such that $Q\mathbf{1} = 0$ is a weighted Laplacian.

Note that $R^T DWD^T R$ is an $m \times m$ symmetric matrix; hence its eigenvalues are real. The sum of the eigenvalues is the trace of the matrix, and hence the energy of the representation is given by the sum of the eigenvalues of $R^T DWD^T R$.

For the normalized representation of the cube we have (with $W = I$)

$$R^T = \frac{1}{\sqrt{8}} \begin{pmatrix} 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \end{pmatrix},$$

$$Q = \begin{pmatrix} 3 & -1 & 0 & -1 & -1 & 0 & 0 & 0 \\ -1 & 3 & -1 & 0 & 0 & -1 & 0 & 0 \\ 0 & -1 & 3 & -1 & 0 & 0 & -1 & 0 \\ -1 & 0 & -1 & 3 & 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 & 3 & -1 & 0 & -1 \\ 0 & -1 & 0 & 0 & -1 & 3 & -1 & 0 \\ 0 & 0 & -1 & 0 & 0 & -1 & 3 & -1 \\ 0 & 0 & 0 & -1 & -1 & 0 & -1 & 3 \end{pmatrix},$$

which implies that

$$R^T Q R = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

and $\mathcal{E}(\rho) = 6$. This can be confirmed directly by noting that each of the 12 edges of the cube has length $1/\sqrt{2}$.

13.4 Energy and Eigenvalues

We now show that the energy of certain representations of a graph X are determined by the eigenvalues of the Laplacian of X. If M is an invertible $m \times m$ matrix, then the map that sends u to $\rho(u)M$ is another representation of X. This representation is given by the matrix RM and provides as much information about X as does ρ. From this point of view the representation is determined by its column space. Therefore, we may assume that the columns of R are orthogonal to each other, and as above that each column has norm 1. In this situation the matrix R satisfies $R^T R = I_m$, and the representation is called an *orthogonal representation.*

Theorem 13.4.1 *Let X be a graph on n vertices with weighted Laplacian Q. Assume that the eigenvalues of Q are $\lambda_1 \leq \cdots \leq \lambda_n$ and that $\lambda_2 > 0$. The minimum energy of a balanced orthogonal representation of X in $\mathbb{R}^m$ equals $\sum_{i=2}^{m+1} \lambda_i$.*

Proof. By Lemma 13.3.1 the energy of a representation is $\operatorname{tr} R^T QR$. From Corollary 9.5.2, the energy of an orthogonal representation in $\mathbb{R}^\ell$ is bounded below by the sum of the ℓ smallest eigenvalues of Q. We can realize this lower bound by taking the columns of R to be vectors $x_1, \ldots, x_\ell$ such that $Qx_i = \lambda_i x_i$.

Since $\lambda_2 > 0$, we must have $x_1 = \mathbf{1}$, and therefore by deleting x_1 we obtain a balanced orthogonal representation in $\mathbb{R}^{\ell-1}$, with the same energy. Conversely, we can reverse this process to obtain an orthogonal representation in $\mathbb{R}^\ell$ from a balanced orthogonal representation in $\mathbb{R}^{\ell-1}$ such that these two representations have the same energy. Therefore, the minimum energy of a balanced orthogonal representation of X in $\mathbb{R}^m$ equals the minimum energy of an orthogonal representation in $\mathbb{R}^{m+1}$, and this minimum equals $\lambda_2 + \cdots + \lambda_{m+1}$. $\square$

This result provides an intriguing automatic method for drawing a graph in any number of dimensions. Compute an orthonormal basis of eigenvectors $x_1, \ldots, x_n$ for the Laplacian Q and let the columns of R be $x_2, \ldots, x_{m+1}$. Theorem 13.4.1 implies that this yields an orthogonal balanced representation of minimum energy. The representation is not necessarily unique, because it may be the case that $\lambda_{m+1} = \lambda_{m+2}$, in which case there is no reason to choose between x_{m+1} and x_{m+2}.

Figure 13.3 shows that such a representation (in $\mathbb{R}^2$) can look quite appealing, while Figure 13.4 shows that it may be less appealing.

Both of these graphs are planar graphs on 10 vertices, but in both cases the drawing is not planar. Worse still, in general there is no guarantee that the images of the vertices are even distinct. The representation of the cube in $\mathbb{R}^3$ given above can be obtained by this method.

More generally, any pairwise orthogonal triple of eigenvectors of Q provides an orthogonal representation in $\mathbb{R}^3$, and this representation may have

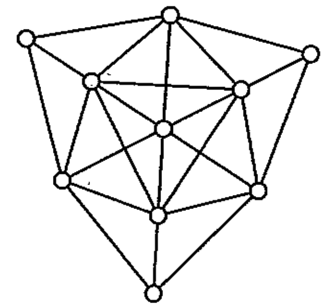

Figure 13.3. A planar triangulation represented in $\mathbb{R}^2$

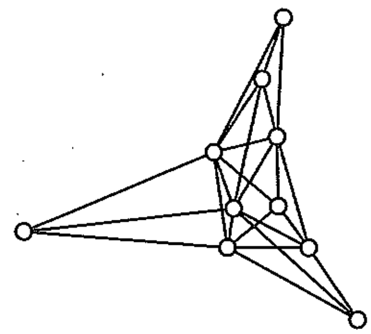

Figure 13.4. A planar triangulation represented in $\mathbb{R}^2$

pleasing properties, even if we do not choose the eigenvectors that minimize the energy.

We finish this section with a corollary to Theorem 13.4.1.

Corollary 13.4.2 *Let X be a graph on n vertices. Then the minimum value of*

$$\frac{\sum_{uv \in E(X)} (x_u - x_v)^2}{\sum_u x_u^2},$$

as x ranges over all nonzero vectors orthogonal to $\mathbf{1}$, is $\lambda_2(X)$. The maximum value is $\lambda_\infty(X)$. □

13.5 Connectivity

Our main result in this section is a consequence of the following bound.

Theorem 13.5.1 *Suppose that S is a subset of the vertices of the graph X. Then $\lambda_2(X) \le \lambda_2(X \setminus S) + |S|$.*

Proof. Let z be a unit vector of length n such that (when viewed as a function on $V(X)$) its restriction to S is zero, and its restriction to $V(X)\backslash S$ is an eigenvector of $Q(X \setminus S)$ orthogonal to $\mathbf{1}$ and with eigenvalue θ. Then by Corollary 13.4.2

$$\lambda_2(X) \leq \sum_{uv \in E(X)} (z_u - z_v)^2.$$

Hence by dividing the edges into those with none, one, or two endpoints in $X \setminus S$ we get

$$\lambda_2(X) \leq \sum_{u \in S} \sum_{v \sim u} z_v^2 + \sum_{uv \in E(X \setminus S)} (z_u - z_v)^2 \leq |S| + \theta.$$

We may take $\theta = \lambda_2(X \setminus S)$, and hence the result follows. □

If S is a vertex-cutset, then $X \setminus S$ is disconnected, so $\lambda_2(X \setminus S) = 0$, and we have the following bound on the vertex connectivity of a graph.

Corollary 13.5.2 *For any graph X we have $\lambda_2(X) \leq \kappa_0(X)$.* □

It follows from our observation in Section 13.1 or from Exercise 4 that the characteristic polynomial of $Q(K_{1,n})$ is $t(t-1)^{n-1}(t-n-1)$. This provides one family of examples where λ_2 equals the vertex connectivity.

Provided that X is not complete, the vertex connectivity of X is bounded above by the edge connectivity, which, in turn, is bounded above by the minimum valency $\delta(X)$ of a vertex in X. We thus have the following useful inequalities for noncomplete graphs:

$$\lambda_2(X) \leq \kappa_0(X) \leq \kappa_1(X) \leq \delta(X).$$

Note that deleting a vertex may increase λ_2. For example, suppose $X = K_n$, where $n \geq 3$, and Y is constructed by adding a new vertex adjacent to two distinct vertices in X. Then $\lambda_2(Y) \leq 2$, since $\delta(Y) = 2$, but $\lambda_2(X) = n$.

Recall that a bridge is an edge whose removal disconnects a graph, and thus a graph has edge-connectivity one if and only if it has a bridge. In this case, the above result shows that $\lambda_2(X) \leq 1$ unless $X = K_2$. It has been noted empirically that λ_2 seems to give a fairly natural measure of the "shape" of a graph. Graphs with small values of λ_2 tend to be elongated graphs of large diameter with bridges, whereas graphs with larger values of λ_2 tend to be rounder with smaller diameter, and larger girth and connectivity.

For cubic graphs, this observation can be made precise, at least as regards the minimum values for λ_2. If $n \geq 10$ and $n \equiv 2 \bmod 4$, the graphs shown in Figure 13.5 have the smallest value of λ_2 among all cubic graphs on n vertices. If $n \geq 12$ and $n \equiv 0 \bmod 4$, the graphs shown in Figure 13.6 have the smallest value of λ_2 among all cubic graphs on n vertices. In both cases these graphs have the maximum diameter among all cubic graphs on n vertices.

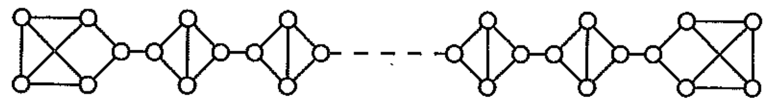

Figure 13.5. Cubic graph with minimum λ_2 on $n \equiv 2 \bmod 4$ vertices

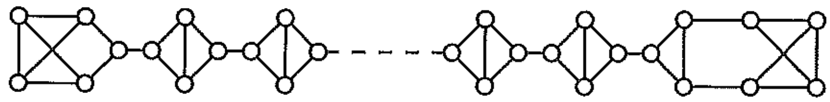

Figure 13.6. Cubic graph with minimum λ_2 on $n \equiv 0 \bmod 4$ vertices

13.6 Interlacing

We now consider what happens to the eigenvalues of $Q(X)$ when we add an edge to X.

Lemma 13.6.1 *Let X be a graph and let Y be obtained from X by adding an edge joining two distinct vertices of X. Then*

$$\lambda_2(X) \leq \lambda_2(Y) \leq \lambda_2(X) + 2.$$

Proof. Suppose we get Y by joining vertices r and s of X. For any vector z we have

$$z^T Q(Y) z = \sum_{uv \in E(Y)} (z_u - z_v)^2 = (z_r - z_s)^2 + \sum_{uv \in E(X)} (z_u - z_v)^2.$$

If we choose z to be a unit eigenvector of $Q(Y)$, orthogonal to $\mathbf{1}$, and with eigenvalue $\lambda_2(Y)$, then by Corollary 13.4.2 we get

$$\lambda_2(Y) \geq \lambda_2(X) + (z_r - z_s)^2. \tag{13.4}$$

On the other hand, if we take z to be a unit eigenvector of $Q(X)$, orthogonal to $\mathbf{1}$, with eigenvalue $\lambda_2(X)$, then by Corollary 13.4.2 we get

$$\lambda_2(Y) \leq \lambda_2(X) + (z_r - z_s)^2. \tag{13.5}$$

It follows from (13.4) that $\lambda_2(X) \leq \lambda_2(Y)$. We can complete the proof by appealing to (13.5). Since $z_r^2 + z_s^2 \leq 1$, it is straightforward to see that $(z_r - z_s)^2 \leq 2$, and the result is proved. □

A few comments on the above proof. If we add an edge joining the two vertices in $2K_1$ (to get K_2), then λ_2 increases from 0 to 2. Although this example might not be impressive, it does show that the upper bound can be tight. The full story is indicated in Exercise 8.

Next, the reader may well have thought that we forgot to insist that the edge added to X in the lemma has to join two distinct and nonadjacent vertices. In fact, the proof works without alteration even if the two vertices chosen are adjacent. We say no more, because here we are not really interested in graphs with multiple edges.

Theorem 13.6.2 *Let X be a graph with n vertices and let Y be obtained from X by adding an edge joining two distinct vertices of X. Then $\lambda_i(X) \leq \lambda_i(Y)$, for all i, and $\lambda_i(Y) \leq \lambda_{i+1}(X)$ if $i < n$.*

Proof. Suppose we add the edge uv to X to get Y. Let z be the vector of length n with u-entry and v-entry 1 and -1, respectively, and all other entries equal to 0. Then $Q(Y) = Q(X) + zz^T$, and if we use Q to denote $Q(X)$, we have

$$tI - Q(Y) = tI - Q - zz^T = (tI - Q)(I - (tI - Q)^{-1}zz^T).$$

By Lemma 8.2.4,

$$\det(I - (tI - Q)^{-1}zz^T) = 1 - z^T(tI - Q)^{-1}z,$$

and therefore

$$\frac{\det(tI - Q(Y))}{\det(tI - Q(X))} = 1 - z^T(tI - Q)^{-1}z.$$

The result now follows from Theorem 8.13.3, applied to the rational function $\psi(t) = 1 - z^T(tI - Q)^{-1}z$, and the proof of Theorem 9.1.1. □

One corollary of this and Theorem 13.4.1 is that if X is a spanning subgraph of Y, then the energy of any balanced orthogonal representation of Y can never be less than the energy of the induced representation of X.

As another corollary of the theorem, we prove again that the Petersen graph does not have a Hamilton cycle. The eigenvalues of the adjacency matrix of the Petersen graph are 3, 1, and -2, with multiplicities 1, 5, and 4, respectively. Therefore, the eigenvalues of the Laplacian matrix for the Petersen graph are 0, 2, and 5, with multiplicities 1, 5, and 4, respectively. The eigenvalues of the adjacency matrix of C_{10} are $2\cos(\pi r/5)$, for $r = 0, 1, \ldots, 9$. It follows that

$$\lambda_6(C_{10}) = (3 + \sqrt{5})/2 > \lambda_6(P) = 2.$$

Consequently, the eigenvalues of the Laplacian matrix of C_{10} do not interlace the eigenvalues of the Laplacian matrix of the Petersen graph, and therefore the Petersen graph does not have a Hamilton cycle.

We present two further examples in Figure 13.7; we can prove that these graphs are not hamiltonian by considering their Laplacians in this fashion. These two graphs are of some independent interest. They are cubic hypohamiltonian graphs, which are somewhat rare. The first graph, on 18 vertices, is one of the two smallest cubic hypohamiltonian graphs after the Petersen graph. Like the Petersen graph it cannot be 3-edge coloured (it is one of the *Blanuša snarks*). The second graph, on 22 vertices, belongs to an infinite family of hypohamiltonian graphs.

It is interesting to note that the technique described in Section 9.2 using the adjacency matrix is not strong enough to prove that these two graphs are not hamiltonian. However, there are cases where the adjacency matrix technique works, but the Laplacian technique does not.

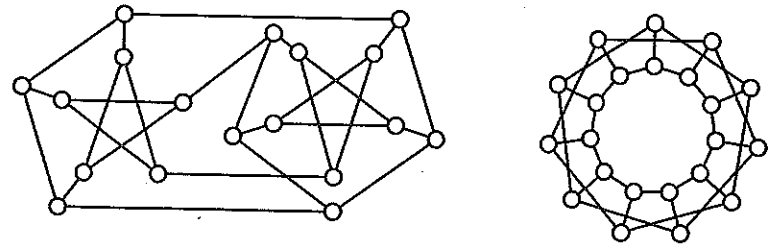

Figure 13.7. Two nonhamiltonian graphs

13.7 Conductance and Cutsets

We now come to some of the most important applications of λ_2. If X is a graph and $S \subseteq V(X)$, let ∂S denote the set of edges with one end in S and the other in $V(X) \setminus S$.

Lemma 13.7.1 *Let X be a graph on n vertices and let S be a subset of $V(X)$. Then*

$$\lambda_2(X) \leq \frac{n|\partial S|}{|S|(n-|S|)}.$$

Proof. Suppose $|S| = a$. Let z be the vector (viewed as a function on $V(X)$) whose value is $n-a$ on the vertices in S and $-a$ on the vertices not in S. Then z is orthogonal to $\mathbf{1}$, so by Corollary 13.4.2

$$\lambda_2(X) \leq \frac{\sum_{uv \in E(X)} (z_u - z_v)^2}{\sum_u z_u^2} = \frac{|\partial S| n^2}{a(n-a)^2 + (n-a)a^2}.$$

The lemma follows immediately from this. □

By way of a simple example, if S is a single vertex with valency k, then the lemma implies that $\lambda_2(X) \leq kn/(n-1)$. This is weaker than Fiedler's result that λ_2 is no greater than the minimum valency of X (Theorem 13.5.1), although not by much.

Our next application is much more important. Define the *conductance* $\Phi(X)$ of a graph X to be the minimum value of

$$\frac{|\partial S|}{|S|},$$

where S ranges over all subsets of $V(X)$ of size at most $|V(X)|/2$. (Many authors refer to this quantity as the *isoperimetric number* of a graph. We follow Lovász, which seems safe.) From Lemma 13.7.1 we have at once the following:

Corollary 13.7.2 *For any graph X we have $\Phi(X) \geq \lambda_2(X)/2$.* □

The real significance of this bound is that λ_2 can be computed to a given number of digits in polynomial time, whereas determining the conductance of a graph is an NP-hard problem. A family of graphs with constant valency and conductance bounded from below by a positive constant is called a family of *expanders*. These are important in theoretical computer science, if not in practice.

The *bisection width* of a graph on n vertices is the minimum value of $|\partial S|$, for any subset S of size $\lfloor n/2 \rfloor$. Again, this is NP-hard to compute, but we do have the following:

Corollary 13.7.3 *The bisection width of a graph X on $2m$ vertices is at least $m\lambda_2(X)/2$.* □

We apply this to the k-cube Q_k. In Exercise 13 it is established that $\lambda_2(Q_k) = 2$, from which it follows that the bisection width of the k-cube is at least 2^{k-1}. Since this value is easily realized, we have thus found the exact value.

Let $\text{bip}(X)$ denote the maximum number of edges in a spanning bipartite subgraph of X. This equals the maximum value of $|\partial S|$, where S ranges over all subsets of $V(X)$ with size at most $|V(X)|/2$.

Lemma 13.7.4 *If X is a graph with n vertices, then* $\text{bip}(X) \le n\lambda_\infty(X)/4$.

Proof. By applying Lemma 13.7.1 to the complement of X we get

$$|\partial S| \le |S|(n - |S|)\lambda_\infty(X)/n \le n\lambda_\infty(X)/4,$$

which is the desired inequality. □

13.8 How to Draw a Graph

We will describe a remarkable method, due to Tutte, for determining whether a 3-connected graph is planar.

Lemma 13.8.1 *Let S be a set of points in $\mathbb{R}^m$. Then the vector x in $\mathbb{R}^m$ minimizes $\sum_{y\in S} \|x - y\|^2$ if and only if*

$$x = \frac{1}{|S|}\sum_{y\in S} y.$$

Proof. Let $\hat{y}$ be the centroid of the set S, i.e.,

$$\hat{y} = \frac{1}{|S|}\sum_{y\in S} y.$$

Then

$$\sum_{y\in S} \|x - y\|^2 = \sum_{y\in S} \|(x - \hat{y}) + (\hat{y} - y)\|^2$$

$$= |S|\,\|x-\hat{y}\|^2 + \sum_{y\in S}\|\hat{y}-y\|^2 + 2\sum_{y\in S}\langle x-\hat{y}, \hat{y}-y\rangle$$
$$= |S|\,\|x-\hat{y}\|^2 + \sum_{y\in S}\|\hat{y}-y\|^2.$$

Therefore, this is a minimum if and only if $x = \hat{y}$. □

We say that a representation ρ of X is *barycentric* relative to a subset F of $V(X)$ if for each vertex u not in F, the vector $\rho(u)$ is the centroid of the images of the neighbours of u. A barycentric representation can easily be made balanced, but will normally not be orthogonal. If the images of the vertices in F are specified, then a barycentric embedding has minimum energy. Our next result formalizes the connection with the Laplacian.

Lemma 13.8.2 *Let F be a subset of the vertices of X, let ρ be a representation of X, and let R be the matrix whose rows are the images of the vertices of X. Let Q be the Laplacian of X. Then ρ is barycentric relative to F if and only if the rows of QR corresponding to the vertices in $X \setminus F$ are all zero.*

Proof. The vector x is the centroid of the vectors in S if and only if

$$\sum_{y\in S}(x-y) = 0.$$

If u has valency d, the u-row of QR is equal to

$$d\rho(u) - \sum_{v\sim u}\rho(v) = \sum_{v\sim u}\rho(u)-\rho(v).$$

The lemma follows. □

Lemma 13.8.3 *Let X be a connected graph, let F be a subset of the vertices of X, and let σ be a map from F into $\mathbb{R}^m$. If $X \setminus F$ is connected, there is a unique m-dimensional representation ρ of X that extends σ and is barycentric relative to F.*

Proof. Let Q be the Laplacian of X. Assume that we have

$$Q = \begin{pmatrix} Q_1 & B^T \\ B & Q_2 \end{pmatrix},$$

where the rows and columns of Q_1 are indexed by the vertices of F. Let R be the matrix describing the representation ρ. We may assume

$$R = \begin{pmatrix} R_1 \\ R_2 \end{pmatrix},$$

where R_1 gives the values of σ on F. Then ρ extends σ and is barycentric (relative to F) if and only if

$$\begin{pmatrix} Q_1 & B^T \\ B & Q_2 \end{pmatrix}\begin{pmatrix} R_1 \\ R_2 \end{pmatrix} = \begin{pmatrix} Y_1 \\ 0 \end{pmatrix}.$$

Then $BR_1 + Q_2R_2 = 0$, and so if Q_2 is invertible, this yields that

$$R_2 = -Q_2^{-1}BR_1, \qquad Y_1 = (Q_1 - B^TQ_2B)R_1.$$

We complete the proof by showing that since $X \setminus F$ is connected, Q_2 is invertible. Let $Y = X\setminus F$. Then there is a nonnegative diagonal matrix Δ_2 such that

$$Q_2 = Q(Y) + \Delta_2.$$

Since X is connected, $\Delta_2 \neq 0$. We prove that Q_2 is positive definite. We have

$$x^TQ_2x = x^TQ(Y)x + x^T\Delta_2x.$$

Because $x^TQ(Y)x = \sum_{ij\in E(Y)}(x_i - x_j)^2$, we see that $x^TQ(Y)x \geq 0$ and that $x^TQ(Y)x = 0$ if and only if $x = c\mathbf{1}$ for some c. But now $x^T\Delta_2x = c^2\mathbf{1}^T\Delta_2\mathbf{1}$, and this is positive unless $c = 0$. Therefore, $x^TQ_2x > 0$ unless $x = 0$; in other words, Q_2 is positive definite, and consequently it is invertible. □

Tutte showed that each edge in a 3-connected graph lies in a cycle C such that no edge not in C joins two vertices of C and $X\setminus C$ is connected. He called these *peripheral cycles*. For example, any face of a 3-connected planar graph can be shown to be a peripheral cycle.

Suppose that C is a peripheral cycle of size r in a 3-connected graph X and suppose that we are given a mapping σ from $V(C)$ to the vertices to a convex r-gon in $\mathbb{R}^2$, such that adjacent vertices in C are adjacent in the polygon. It follows from Lemma 13.8.3 that there is a unique barycentric representation ρ of X relative to F. This determines a drawing of X in the plane, with all vertices of $X \setminus C$ inside the image of C. Tutte proved the truly remarkable result that this drawing has no crossings if and only if X is planar.

Peripheral cycles can be found in polynomial time, and given this, Lemma 13.8.3 provides an automatic method for drawing 3-connected planar graphs. Unfortunately, from an aesthetic viewpoint, the quality of the output is variable. Sometimes there is a good choice of outside face, maybe a large face as in Figure 13.8 or one that is preserved by an automorphism as in Figure 13.9.

However, particularly if there are a lot of triangular faces, the algorithm tends to produce a large number of faces-within-faces, many of which are minuscule.

13.9 The Generalized Laplacian

The rest of this chapter is devoted to a generalization of the Laplacian matrix of a graph. There are many generalized Laplacians associated with

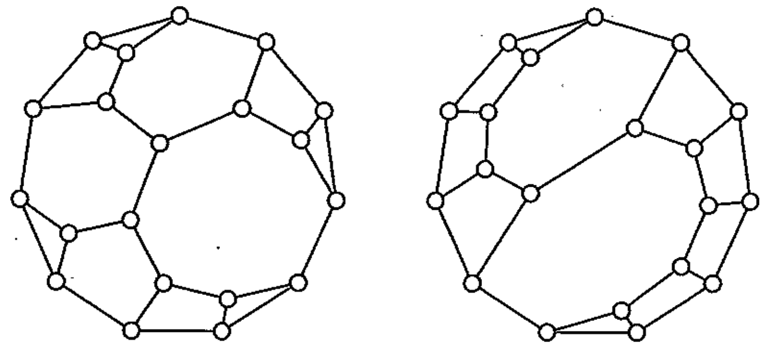

Figure 13.8. Tutte embeddings of cubic planar graphs

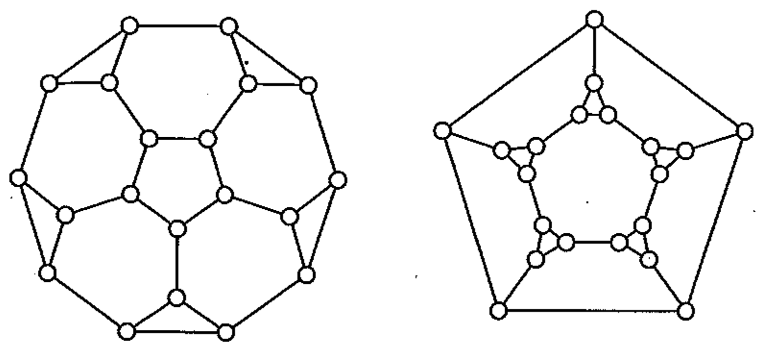

Figure 13.9. Different Tutte embeddings of the same graph

each graph, which at first sight seem only tenuously related. Nevertheless, graph-theoretical properties of a graph constrain the algebraic properties of the entire class of generalized Laplacians associated with it. The next few sections provide an introduction to this important and recent development.

Let X be a graph with n vertices. We call a symmetric $n \times n$ matrix Q a *generalized Laplacian* of X if $Q_{uv} < 0$ when u and v are adjacent vertices of X and $Q_{uv} = 0$ when u and v are distinct and not adjacent. There are no constraints on the diagonal entries of Q; in particular, we do not require that $Q\mathbf{1} = 0$. The ordinary Laplacian is a generalized Laplacian, and if A is the adjacency matrix of X, then $-A$ is a generalized Laplacian.

As with the usual Laplacian, we will denote the eigenvalues of a generalized Laplacian Q by

$$\lambda_1(Q) \leq \lambda_2(Q) \leq \cdots \leq \lambda_n(Q).$$

We will be concerned with the eigenvectors in the λ_2-eigenspace of Q. If Q is a generalized Laplacian of X, then for any c, the matrix $Q - cI$ is a generalized Laplacian with the same eigenvectors as Q. Therefore, we can freely assume that $\lambda_2(Q) = 0$, whenever it is convenient to do so.

Lemma 13.9.1 *Let X be a graph with a generalized Laplacian Q. If X is connected, then $\lambda_1(Q)$ is simple and the corresponding eigenvector can be taken to have all its entries positive.*

Proof. Choose a constant c such that all diagonal entries of $Q - cI$ are nonpositive. By the Perron–Frobenius theorem (Theorem 8.8.1), the largest eigenvalue of $-Q+cI$ is simple and the associated eigenvector may be taken to have only positive entries. □

If x is a vector with entries indexed by the vertices of X, then the *positive support* $\mathrm{supp}_+(x)$ consists of the vertices u such that $x_u > 0$, and the *negative support* $\mathrm{supp}_-(x)$ of the vertices u such that $x_u < 0$. A *nodal domain* of x is a component of one of the subgraphs induced by $\mathrm{supp}_+(x)$ or $\mathrm{supp}_-(x)$. A nodal domain is *positive* if it is a component of $\mathrm{supp}_+(x)$; otherwise, it is *negative*.

If Y is a nodal domain of x, then x_Y is the vector given by

$$(x_Y)_u = \begin{cases} |x_u|, & u \in Y; \\ 0, & \text{otherwise.} \end{cases}$$

If Y and Z are distinct nodal domains with the same sign, then since no edges of X join vertices in Y to vertices in Z,

$$x_Y^T Q x_Z = 0. \tag{13.6}$$

Lemma 13.9.2 *Let x be an eigenvector of Q with eigenvalue λ and let Y be a positive nodal domain of x. Then $(Q - \lambda I)x_Y \le 0$.*

Proof. Let y denote the restriction of x to $V(Y)$ and let z be the restriction of x to $V(X) \setminus \mathrm{supp}_+(x)$. Let Q_Y be the submatrix of Q with rows and columns indexed by $V(Y)$, and let B_Y be the submatrix of Q with rows indexed by $V(Y)$ and with columns indexed by $V(X) \setminus \mathrm{supp}_+(x)$. Since $Qx = \lambda x$, we have

$$Q_Y y + B_Y z = \lambda y. \tag{13.7}$$

Since B_Y and z are nonpositive, $B_Y z$ is nonnegative, and therefore

$$Q_Y y \le \lambda y.$$

□

It is not necessary for x to be an eigenvector for the conclusion of this lemma to hold; it is sufficient that $(Q - \lambda I)x \le 0$. Given our discussion in Section 8.7, we might say that it suffices that x be λ-*superharmonic*.

Corollary 13.9.3 *Let x be an eigenvector of Q with eigenvalue λ, and let U be the subspace spanned by the vectors x_Y, where Y ranges over the positive nodal domains of x. If $u \in U$, then $u^T(Q - \lambda I)u \le 0$.*

Proof. If $u = \sum_Y a_Y x_Y$, then using (13.6), we find that

$$u^T(Q - \lambda I)u = \sum_Y a_Y^2\, x_Y^T(Q - \lambda I)x_Y,$$

and so the claim follows from the previous lemma. □

Theorem 13.9.4 *Let X be a connected graph, let Q be a generalized Laplacian of X, and let x be an eigenvector for Q with eigenvalue $\lambda_2(Q)$. If x has minimal support, then* $\mathrm{supp}_+(x)$ *and* $\mathrm{supp}_-(x)$ *induce connected subgraphs of X.*

Proof. Suppose that v is a λ_2-eigenvector with distinct positive nodal domains Y and Z. Because X is connected, λ_1 is simple and the span of v_Y and v_Z contains a vector, u say, orthogonal to the λ_1-eigenspace.

Now, u can be expressed as a linear combination of eigenvectors of Q with eigenvalues at least λ_2; consequently, $u^T(Q - \lambda_2 I)u \geq 0$ with equality if and only if u is a linear combination of eigenvectors with eigenvalue λ_2.

On the other hand, by Corollary 13.9.3, we have $u^T(Q - \lambda_2 I)u \leq 0$, and so $u^T(Q - \lambda_2 I)u = 0$. Therefore, u is an eigenvector of Q with eigenvalue λ_2 and support equal to $V(Y) \cup V(Z)$.

Any λ_2-eigenvector has both positive and negative nodal domains, because it is orthogonal to the λ_1-eigenspace. Therefore, the preceding argument shows that an eigenvector with distinct nodal domains of the same sign does not have minimal support. Therefore, since x has minimal support, it must have precisely one positive and one negative nodal domain. □

Lemma 13.9.5 *Let Q be a generalized Laplacian of a graph X and let x be an eigenvector of Q. Then any vertex not in* $\mathrm{supp}(x)$ *either has no neighbours in* $\mathrm{supp}(x)$*, or has neighbours in both* $\mathrm{supp}_+(x)$ *and* $\mathrm{supp}_-(x)$.

Proof. Suppose that $u \notin \mathrm{supp}(x)$, so $x_u = 0$. Then

$$0 = (Qx)_u = Q_{uu}x_u + \sum_{v \sim u} Q_{uv}x_v = \sum_{v \sim u} Q_{uv}x_v.$$

Since $Q_{uv} < 0$ when v is adjacent to u, either $x_v = 0$ for all vertices adjacent to u, or the sum has both positive and negative terms. In the former case u is not adjacent to any vertex in $\mathrm{supp}(x)$; in the latter it is adjacent to vertices in both $\mathrm{supp}_+(x)$ and $\mathrm{supp}_-(x)$. □

13.10 Multiplicities

In this section we show that if X is 2-connected and outerplanar, then λ_2 has multiplicity at most two, and that if X is 3-connected and planar, then λ_2 has multiplicity at most three. In the next section we show that if equality holds in the latter case, then the representation provided by the λ_2-eigenspace yields a planar embedding of X.

Lemma 13.10.1 *Let Q be a generalized Laplacian for the graph X. If X is 3-connected and planar, then no eigenvector of Q with eigenvalue $\lambda_2(Q)$ vanishes on three vertices in the same face of any embedding of X.*

Proof. Let x be an eigenvector of Q with eigenvalue λ_2, and suppose that u, v, and w are three vertices not in $\text{supp}(x)$ lying in the same face. We may assume that x has minimal support, and hence $\text{supp}_+(x)$ and $\text{supp}_-(x)$ induce connected subgraphs of X. Let p be a vertex in $\text{supp}_+(x)$. Since X is 3-connected, Menger's theorem implies that there are three paths in X joining p to u, v, and w such that any two of these paths have only the vertex p in common. It follows that there are three vertex-disjoint paths P_u, P_v, and P_w joining u, v, and w, respectively, to some triple of vertices in $N(\text{supp}_+(x))$. Each of these three vertices is also adjacent to a vertex in $\text{supp}_-(x)$. Since both the positive and negative support induce connected graphs, we may now contract all vertices in $\text{supp}_+(x)$ to a single vertex, all vertices in $\text{supp}_-(x)$ to another vertex, and each of the paths P_u, P_v, and P_w to u, v, and w, respectively. The result is a planar graph which contains a copy of $K_{2,3}$ with its three vertices of valency two all lying on the same face. This is impossible. □

Corollary 13.10.2 *Let Q be a generalized Laplacian for the graph X. If X is 3-connected and planar, then $\lambda_2(Q)$ has multiplicity at most three.*

Proof. If λ_2 has multiplicity at least four, then there is an eigenvector in the associated eigenspace whose support is disjoint from any three given vertices. Thus we conclude that λ_2 has multiplicity at most three. □

The graph $K_{2,n}$ is 2-connected and planar. Its adjacency matrix A has eigenvalues $\pm\sqrt{2n}$, both simple, and 0 with multiplicity $n-2$. Taking $Q = -A$, we see that we cannot drop the assumption that X is 3-connected in the last result.

Lemma 13.10.3 *Let X be a 2-connected plane graph with a generalized Laplacian Q, and let x be an eigenvector of Q with eigenvalue $\lambda_2(Q)$ and with minimal support. If u and v are adjacent vertices of a face F such that $x_u = x_v = 0$, then F does not contain vertices from both the positive and negative support of x.*

Proof. Since X is 2-connected, the face F is a cycle. Suppose that F contains vertices p and q such that $x_p > 0$ and $x_q < 0$. Without loss of generality we can assume that they occur in the order u, v, q, and p clockwise around the face F, and that the portion of F from q to p contains only vertices not in $\text{supp}(x)$. Let v' be the first vertex not in $\text{supp}(x)$ encountered moving anticlockwise around F from q, and let u' be the first vertex not in $\text{supp}(x)$ encountered moving clockwise around F

from p. Then u', v', q, and p are distinct vertices of F and occur in that order around F. Let P be a path from v' to p all of whose vertices other than v' are in $\mathrm{supp}_+(x)$, and let N be a path from u' to q all of whose vertices other than u' are in $\mathrm{supp}_-(x)$. The existence of the paths P and N is a consequence of Corollary 13.9.4 and Lemma 13.9.5. Because F is a face, the paths P and N must both lie outside F, and since their endpoints are interleaved around F, they must cross. This is impossible, since P and N are vertex-disjoint, and so we have the necessary contradiction. □

We call a graph *outerplanar* if it has a planar embedding with a face that contains all the vertices. Neither $K_{2,3}$ nor K_4 is outerplanar, and it is known that a graph is outerplanar if and only if it has no minor isomorphic to one of these two graphs. (A *minor* of a graph X is a graph obtained by contracting edges in a subgraph of X.)

Corollary 13.10.4 *Let X be a graph on n vertices with a generalized Laplacian Q. If X is 2-connected and outerplanar, then $\lambda_2(Q)$ has multiplicity at most two.*

Proof. If λ_2 had multiplicity greater than two, then we could find an eigenvector x with eigenvalue λ_2 such that x vanished on two adjacent vertices in the sole face of X. However, since x must be orthogonal to the eigenvector with eigenvalue λ_1, both $\mathrm{supp}_+(x)$ and $\mathrm{supp}_-(x)$ must be nonempty. □

The tree $K_{1,n}$ is outerplanar, but if A is its adjacency matrix, then $-A$ is a generalized Laplacian for it with λ_2 having multiplicity greater than two. Hence we cannot drop the assumption in the corollary that X be 2-connected.

13.11 Embeddings

We have seen that if X is a 3-connected planar graph and Q is a generalized Laplacian for X, then $\lambda_2(Q)$ has multiplicity at most three. The main result of this section is that if $\lambda_2(Q)$ has multiplicity exactly three, then the representation ρ provided by the λ_2-eigenspace of Q provides a planar embedding of X on the unit sphere.

As a first step we need to verify that in the case just described, no vertex is mapped to zero by ρ. This, and more, follows from the next result.

Lemma 13.11.1 *Let X be a 3-connected planar graph with a generalized Laplacian Q such that $\lambda_2(Q)$ has multiplicity three. Let ρ be a representation given by a matrix U whose columns form a basis for the λ_2-eigenspace of Q. If F is a face in some planar embedding of X, then the images under ρ of any two vertices in F are linearly independent.*

Proof. Assume by way of contradiction that u and v are two vertices in a face of X such that $\rho(u) = \alpha\rho(v)$ for some real number α, and let w be a third vertex in the same face. Then we can find a linear combination of the columns of U that vanishes on the vertices u, v, and w, thus contradicting Lemma 13.10.1. □

If ρ is a representation of X that maps no vertex to zero, then we define the *normalized representation* $\hat{\rho}$ by

$$\hat{\rho} = \|\rho(u)\|^{-1}\rho(u).$$

Suppose that X is a 3-connected planar graph with a generalized Laplacian Q such that $\lambda_2(Q)$ has multiplicity three, and let ρ be the representation given by the λ_2-eigenspace. By the previous lemma, the corresponding normalized representation $\hat{\rho}$ is well-defined and maps every vertex to a point of the unit sphere. If u and v are adjacent in X, then $\hat{\rho}(u) \neq \pm\hat{\rho}(v)$, so there is a unique geodesic on the sphere joining the images of u and v. Thus we have a well-defined embedding of the graph X on the unit sphere, and our task is to show that this embedding is planar, i.e., distinct edges can meet only at a vertex.

If $C \subseteq \mathbb{R}^n$, then the *convex cone* generated by C is the set of all nonnegative linear combinations of the elements of C. A subset of the unit sphere is *spherically convex* if whenever it contains points u and v, it contains all points on any geodesic joining u to v. The intersection of the unit sphere with a convex cone is spherically convex. Suppose that F is a face in some planar drawing of X, and consider the convex cone C generated by the images under $\hat{\rho}$ of the vertices of F. This meets the unit sphere in a convex spherical polygon, and by Lemma 13.10.3, each edge of F determines an edge of this polygon.

This does not yet imply that our embedding of X on the sphere has no crossings; for example, the images of distinct faces of X could overlap. Our next result removes some of the difficulty.

Lemma 13.11.2 *Let X be a 2-connected planar graph. Suppose it has a planar embedding where the neighbours of the vertex u are, in cyclic order, $v_1, \dots, v_k$. Let Q be a generalized Laplacian for X such that $\lambda_2(Q)$ has multiplicity three. Then the planes spanned by the pairs $\{\rho(u), \rho(v_i)\}$ are arranged in the same cyclic order around the line spanned by $\rho(u)$ as the vertices v_i are arranged around u.*

Proof. Let x be an eigenvector with eigenvalue λ_2 with minimal support such that $x(u) = x(v_1) = 0$. (Here we are viewing x as a function on $V(X)$.) By Lemma 13.10.1, we see that neither $x(v_2)$ nor $x(v_k)$ can be zero, and replacing x by $-x$ if needed, we may suppose that $x(v_2) > 0$. Given this, we prove that $x(v_k) < 0$.

Suppose that there are some values h, i, and j such that $2 \leq h < i < j \leq k$ and $x(v_h) > 0$, $x(v_j) > 0$, and $x(v_i) \leq 0$. Since $\text{supp}_+(x)$ is connected, the vertices v_h and v_j are joined in X by a path with all vertices in $\text{supp}_+(x)$. Taken with u, this path forms a cycle in X that separates v_1 from v_i. Since X is 2-connected, there are two vertex-disjoint paths P_1 and P_i joining v_1 and v_i respectively to vertices in $N(\text{supp}_+(x))$. The end-vertices of these paths other than v_1 and v_i are adjacent to vertices in $\text{supp}_-(x)$, and thus we have found two vertices in $\text{supp}_-(x)$ that are separated by vertices in $\text{supp}_+(x)$. This contradicts the fact that $\text{supp}_-(x)$ is connected.

It follows that there is exactly one index i such that $x(v_i) > 0$ and $x(v_{i+1}) \le 0$. Since $x(u) = 0$ and $x(v_2) > 0$, it follows from Lemma 13.9.5 that u has a neighbour in $\text{supp}_-(x)$, and therefore $x(v_k)$ must be negative.

From this we see that if we choose x such that $x(u) = x(v_i) = 0$ and $x(v_{i+1}) > 0$, then $x(v_{i-1}) < 0$ (where the subscripts are computed modulo k). The lemma follows at once from this. □

We now prove that the embedding provided by $\hat{\rho}$ has no crossings. The argument is topological.

Suppose that X is drawn on a sphere S_a without crossings. Let S_b be a unit sphere, with X embedded on it using $\hat{\rho}$, as described above. The normalized representation $\hat{\rho}$ provides an injective map from the vertices of X in S_a to the vertices of X in S_b. By our remarks preceding the previous lemma, this map extends to a continuous map ψ from S_a to S_b, which injectively maps each face on S_a to a spherically convex region on S_b. From Lemma 13.11.2, it even follows that ψ is injective on the union of the faces of X that contain a given vertex. Hence ψ is a continuous locally injective map from S_a to S_b.

It is a standard result that such a map must be injective; we outline a proof. First, since ψ is continuous and locally injective, there is an integer k such that $|\psi^{-1}(x)| = k$ for each point x on S_b. Let Y be any graph embedded on S_b, with v vertices, e edges, and f faces. Then $\psi^{-1}(Y)$ is a plane graph on S_a with kv vertices, ke edges, and kf faces. By Euler's formula,

$$2 = kv - ke + kf = k(v - e + f) = 2k,$$

and therefore $k = 1$.

Thus we have shown that ψ is injective, and therefore it is a homeomorphism. We conclude that $\hat{\rho}$ embeds X without crossings.

Exercises

1. If D is the incidence matrix of an oriented graph, then show that any square submatrix of D has determinant 0, 1, or -1.

2. Show that the determinant of a square submatrix of $B(X)$ is equal to 0 or $\pm 2^r$, for some integer r.

3. If M is a matrix, let $M(i|j)$ denote the submatrix we get by deleting row i and column j. Define a 2-forest in a graph to be a spanning forest with exactly two components. Let Q be the Laplacian of X. If u, p, and q are vertices of X and $p \neq q$, show that $\det Q[u](p|q)$ is equal to the number of 2-forests with u in one component and p and q in the other.

4. Determine the characteristic polynomial of $Q(K_{m,n})$.

5. An *arborescence* is an acyclic directed graph with a root vertex u such that u has in-valency 0 and each vertex other than u has in-valency 1 and is joined to u by a directed path. (In other words, it is a tree oriented so that all arcs point away from the root.) Let Y be a directed graph with adjacency matrix A and let D be the diagonal matrix with ith diagonal entry equal to the in-valency of the ith vertex of Y. Show that the number of spanning arborescences in Y rooted at a given vertex u is equal to $\det((D-A)[u])$.

6. Show that if X is connected and has n vertices, then
$$\lambda_2(X) = \min_x \frac{n \sum_{ij \in E(X)} (x_i - x_j)^2}{\sum_{i<j} (x_i - x_j)^2},$$
where the minimum is taken over all nonconstant vectors x.

7. Show that if T is a tree with at least three vertices, then $\lambda_2(T) \leq 1$, with equality if and only if T is a star (i.e., is isomorphic to $K_{1,n}$).

8. Let r and s be distinct nonadjacent vertices in the graph X. If $e \in E(X)$, show that $\lambda_2(X \setminus e) = \lambda_2(X) - 2$ if and only if X is complete.

9. Let D be an oriented incidence matrix for the graph X. Let d_i denote the valency of the vertex i in X. Show that the largest eigenvalue of $D^T D$ is bounded above by the maximum value of $d_i + d_j$, for any two adjacent vertices i and j in X. Deduce that this is also an upper bound on λ_∞. (And for even more work, show that this bound is tight if and only if X is bipartite and semiregular.)

10. Let X be a connected graph on n vertices. Show that there is a subset of $V(X)$ such that $2|S| \leq n$,
$$\frac{|\partial S|}{|S|} = \Phi(X);$$
and the subgraphs induced by S and $V \setminus S$ are both connected.

11. Let X be a graph on n vertices with diameter d. Show that $\lambda_2 \geq 1/nd$.

12. If X is the Cartesian product of two graphs Y and Z, show that $\lambda_2(X)$ is the minimum of $\lambda_2(Y)$ and $\lambda_2(Z)$. (Hint: Find eigenvectors for X, and hence determine all eigenvalues of X in terms of those of its factors.)

13. Use Exercise 12 to show that $\lambda_2(Q_k) = 2$.

14. If X is an arc-transitive graph with diameter d and valency r, show that $\Phi(X) \geq r/2d$.

15. Show that a cycle in a 3-connected planar graph is a peripheral cycle if and only if it is a face in every planar embedding of the graph.

16. Let X be a connected graph and let z be an eigenvector of $Q(X)$ with eigenvalue λ_2. Call a path $u_1, \ldots, u_r$ strictly decreasing if the values of z on the vertices of the path form a strictly decreasing sequence. Show that if $u \in V(X)$ and $z_u > 0$, then u is joined by a strictly decreasing path to some vertex v such that $z_v \leq 0$.

17. Let X be a connected graph. Show that if $Q(X)$ has exactly three distinct eigenvalues, then there is a constant μ such that any pair of distinct nonadjacent vertices in X have exactly μ common neighbours. Show further that there is a constant $\bar{\mu}$ such that any pair of distinct nonadjacent vertices in $\overline{X}$ have exactly $\bar{\mu}$ common neighbours. Find a graph X with this property that is not regular. (A regular graph would be strongly regular.)

18. Let Q be a generalized Laplacian for a connected graph X. If x is an eigenvector for Q with eigenvalue λ_2 and u is a vertex in X such that x_u is maximal, prove that
$$Q_{uu} + \sum_{v \sim u} Q_{uv} \leq \lambda_2.$$

19. Let Q be a generalized Laplacian for a connected graph X and consider the representation ρ provided by the λ_2-eigenspace. Show that if $\rho(u)$ does not lie in the convex hull of the set
$$N := \{\rho(v) : v \sim u\} \cup \{0\},$$
then there is a vector a such that $a^T\rho(u) > a^T\rho(v)$, for any neighbour v of u. (Do not struggle with this; quote a result from optimization.) Deduce that if $\rho(u)$ does not lie in the convex hull of N, then
$$Q_{uu} + \sum_{v \sim u} Q_{uv} < \lambda_2.$$

20. Let Q be a generalized Laplacian for a path. Show that all the eigenvalues of Q are simple.

21. Let Q be a generalized Laplacian for a connected graph X and let x be an eigenvector for Q with eigenvalue λ_2. Show that if no entries of x are zero, then both $\operatorname{supp}_+(x)$ and $\operatorname{supp}_-(x)$ are connected.

22. Let Q be a generalized Laplacian for a connected graph X, let x be an eigenvector for Q with eigenvalue λ_2 and let C be the vertex set of some component of $\operatorname{supp}(x)$. Show that $N(C) = N(\operatorname{supp}(x))$.

Notes

Theorem 13.4.1 comes from Pisanski and Shawe-Taylor [8], and our discussion in Section 13.3 and Section 13.4 follows their treatment. Fiedler [2] introduced the study of λ_2. He called it the *algebraic connectivity* of a graph. In [3], Fiedler proves that if z is an eigenvector for the connected graph X with eigenvalue λ_2 and $c \leq 0$, then the graph induced by the set

$$\{u \in V(X) : z_u \geq c\}$$

is connected. Exercise 16 shows that it suffices to prove this when $c = 0$.

Our work in Section 13.6 is a modest extension of an idea due to Mohar, which we treated in Section 9.2. Van den Heuvel [11] offers further applications of this type.

Alon uses Lemma 13.7.4 to show that there is a positive constant c such that for every e there is a triangle-free graph with e edges and

$$\operatorname{bip}(X) \leq \frac{e}{2} + ce^{4/5}.$$

Lovász devotes a number in exercises in Chapter 11 of [4] to conductance.

Section 13.8 is, of course, based on [9], one of Tutte's many masterpieces.

The final sections are based on work of van der Holst, Schrijver, and Lovász [12], [13], [5]. These papers are motivated by the study of the *Colin de Verdière number* of a graph. This is defined to be the maximum corank of a generalized Laplacian Q that also satisfies the additional technical condition that there is no nonzero matrix B such that $QB = 0$ and $B_{uv} = 0$ when u is equal or adjacent to v. For an introduction to this important subject we recommend [13].

For a solution to Exercise 5, see [4]. The result in Exercise 9 comes from [1]. Exercise 10 comes from Mohar [6]. B. D. McKay proved that if X has n vertices and diameter d, then $d \geq 4/n\lambda_2$. This is stronger than the result of Exercise 11, and is close to optimal for trees. For a proof of the stronger result, see Mohar [7]; for the weaker bound, try [4]. It might appear that we do not need lower bounds on the diameter, as after all, it can be computed in polynomial time. The problem is that this is polynomial time in the number of vertices of a graph. However, we may wish to bound the diameter of a Cayley graph given by its connection set; in this case we need to compute the diameter in time polynomial in the size of the connection set, i.e., the valency of the graph. Exercise 17 is based on van Dam and Haemers [10].

The Colin de Verdière number of a graph is less than or equal to three if and only if the graph is planar, and in this case, we can find a generalized Laplacian of maximum corank. The null space of this generalized Laplacian then yields a planar representation of the graph, using the method described in Section 13.11. However, for a general graph X, we do not know how to find a suitable generalized Laplacian with corank equal to the Colin de

Verdière number of X, nor do we know any indirect method to determine it.

References

[1] W. N. ANDERSON JR. AND T. D. MORLEY, *Eigenvalues of the Laplacian of a graph*, Linear and Multilinear Algebra, 18 (1985), 141–145.

[2] M. FIEDLER, *Algebraic connectivity of graphs*, Czechoslovak Math. J., 23(98) (1973), 298–305.

[3] ——, *A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory*, Czechoslovak Math. J., 25(100) (1975), 619–633.

[4] L. LOVÁSZ, *Combinatorial Problems and Exercises*, North-Holland Publishing Co., Amsterdam, second edition, 1993.

[5] L. LOVÁSZ AND A. SCHRIJVER, *On the null space of a Colin de Verdière matrix*, Ann. Inst. Fourier (Grenoble), 49 (1999), 1017–1026.

[6] B. MOHAR, *Isoperimetric numbers of graphs*, J. Combin. Theory Ser. B, 47 (1989), 274–291.

[7] ——, *Eigenvalues, diameter, and mean distance in graphs*, Graphs Combin., 7 (1991), 53–64.

[8] T. PISANSKI AND J. SHAWE-TAYLOR, *Characterising graph drawing with eigenvectors*, J. Chem. Inf. Comput. Sci, 40 (2000), 567–571.

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

[10] E. R. VAN DAM AND W. H. HAEMERS, *Graphs with constant μ and $\overline{\mu}$*, Discrete Math., 182 (1998), 293–307.

[11] J. VAN DEN HEUVEL, *Hamilton cycles and eigenvalues of graphs*, Linear Algebra Appl., 226/228 (1995), 723–730.

[12] H. VAN DER HOLST, *A short proof of the planarity characterization of Colin de Verdière*, J. Combin. Theory Ser. B, 65 (1995), 269–272.

[13] H. VAN DER HOLST, L. LOVÁSZ, AND A. SCHRIJVER, *The Colin de Verdière graph parameter*, in Graph theory and combinatorial biology (Balatonlelle, 1996), János Bolyai Math. Soc., Budapest, 1999, 29–85.