

Chris Godsil
Gordon Royle

Algebraic Graph Theory

With 120 Illustrations



Springer

Chris Godsil
Department of Combinatorics
and Optimization
University of Waterloo
Waterloo, Ontario N2L 3G1
Canada
cgodsil@math.uwaterloo.ca

Gordon Royle
Department of Computer Science
University of Western Australia
Nedlands, Western Australia 6907
Australia
gordon@cs.uwa.edu.au

Editorial Board

S. Axler
Mathematics Department
San Francisco State
University
San Francisco, CA 94132
USA

F.W. Gehring
Mathematics Department
East Hall
University of Michigan
Ann Arbor, MI 48109
USA

K.A. Ribet
Mathematics Department
University of California
at Berkeley
Berkeley, CA 94720-3840
USA

Mathematics Subject Classification (2000): 05Cxx, 05Exx

Library of Congress Cataloging-in-Publication Data
Godsil, C.D. (Christopher David), 1949–
Algebraic graph theory / Chris Godsil, Gordon Royle.
p. cm. – (Graduate texts in mathematics; 207)
Includes bibliographical references and index.
ISBN 0-387-95241-1 (hc. : alk. paper)
ISBN 0-387-95220-9 (pbk.: alk. paper)
I. Graph theory I. Royle, Gordon. II. Title. III. Series.
QA166 .G63 2001
511'.5–dc21
00-053776

Printed on acid-free paper.

© 2001 Springer-Verlag New York, Inc.
All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag New York, Inc., 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden. The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Production managed by A. Orrantia; manufacturing supervised by Jerome Basma.
Electronically imposed from the authors' PostScript files.
Printed and bound by R.R. Donnelley and Sons, Harrisonburg, VA.
Printed in the United States of America.

9 8 7 6 5 4 3 2 1

ISBN 0-387-95241-1
ISBN 0-387-95220-9

SPIN 10793786 (hardcover)
SPIN 10791962 (softcover)

Springer-Verlag New York Berlin Heidelberg
A member of BertelsmannSpringer Science+Business Media GmbH

8

Matrix Theory

There are various matrices that are naturally associated with a graph, such as the adjacency matrix, the incidence matrix, and the Laplacian. One of the main problems of algebraic graph theory is to determine precisely how, or whether, properties of graphs are reflected in the algebraic properties of such matrices.

Here we introduce the incidence and adjacency matrices of a graph, and the tools needed to work with them. This chapter could be subtitled “Linear Algebra for Graph Theorists,” because it develops the linear algebra we need from fundamental results about symmetric matrices through to the Perron–Frobenius theorem and the spectral decomposition of symmetric matrices.

Since many of the matrices that arise in graph theory are 01-matrices, further information can often be obtained by viewing the matrix over the finite field $GF(2)$. We illustrate this with an investigation into the binary rank of the adjacency matrix of a graph.

8.1 The Adjacency Matrix

The *adjacency matrix* $A(X)$ of a directed graph X is the integer matrix with rows and columns indexed by the vertices of X , such that the uv -entry of $A(X)$ is equal to the number of arcs from u to v (which is usually 0 or 1). If X is a graph, then we view each edge as a pair of arcs in opposite directions, and $A(X)$ is a symmetric 01-matrix. Because a graph has no

loops, the diagonal entries of $A(X)$ are zero. Different directed graphs on the same vertex set have different adjacency matrices, even if they are isomorphic. This is not much of a problem, and in any case we have the following consolation, the proof of which is left as an exercise.

Lemma 8.1.1 *Let X and Y be directed graphs on the same vertex set. Then they are isomorphic if and only if there is a permutation matrix P such that $P^T A(X)P = A(Y)$. \square*

Since permutation matrices are orthogonal, $P^T = P^{-1}$, and so if X and Y are isomorphic directed graphs, then $A(X)$ and $A(Y)$ are similar matrices. The *characteristic polynomial* of a matrix A is the polynomial

$$\phi(A, x) = \det(xI - A),$$

and we let $\phi(X, x)$ denote the characteristic polynomial of $A(X)$. The *spectrum* of a matrix is the list of its eigenvalues together with their multiplicities. The spectrum of a graph X is the spectrum of $A(X)$ (and similarly we refer to the eigenvalues and eigenvectors of $A(X)$ as the eigenvalues and eigenvectors of X). Lemma 8.1.1 shows that $\phi(X, x) = \phi(Y, x)$ if X and Y are isomorphic, and so the spectrum is an invariant of the isomorphism class of a graph.

However, it is not hard to see that the spectrum of a graph does not determine its isomorphism class. Figure 8.1 shows two graphs that are not isomorphic but share the characteristic polynomial

$$(x + 2)(x + 1)^2(x - 1)^2(x^2 - 2x - 6),$$

and hence have spectrum

$$\{ -2, -1^{(2)}, 1^{(2)}, 1 \pm \sqrt{7} \}$$

(where the superscripts give the multiplicities of eigenvalues with multiplicity greater than one). Two graphs with the same spectrum are called *cospectral*.




Figure 8.1. Two cospectral graphs

The graphs of Figure 8.1 show that the valencies of the vertices are not determined by the spectrum, and that whether a graph is planar is not determined by the spectrum. In general, if there is a cospectral pair of graphs, only one of which has a certain property \mathcal{P} , then \mathcal{P} cannot be

determined by the spectrum. Such cospectral pairs have been found for a large number of graph-theoretical properties.

However, the next result shows that there is some useful information that can be obtained from the spectrum. A *walk* of length r in a directed graph X is a sequence of vertices

$$v_0 \sim v_1 \sim \cdots \sim v_r.$$

A walk is *closed* if its first and last vertices are the same. This definition is similar to that of a path (Section 1.2), with the important difference being that a walk is permitted to use vertices more than once.

Lemma 8.1.2 *Let X be a directed graph with adjacency matrix A . The number of walks from u to v in X with length r is $(A^r)_{uv}$.*

Proof. This is easily proved by induction on r , as you are invited to do. \square

The *trace* of a square matrix A is the sum of its diagonal entries and is denoted by $\text{tr } A$. The previous result shows that the number of closed walks of length r in X is $\text{tr } A^r$, and hence we get the following corollary:

Corollary 8.1.3 *Let X be a graph with e edges and t triangles. If A is the adjacency matrix of X , then*

- (a) $\text{tr } A = 0$,
- (b) $\text{tr } A^2 = 2e$,
- (c) $\text{tr } A^3 = 6t$. \square

Since the trace of a square matrix is also equal to the sum of its eigenvalues, and the eigenvalues of A^r are the r th powers of the eigenvalues of A , we see that $\text{tr } A^r$ is determined by the spectrum of A . Therefore, the spectrum of a graph X determines at least the number of vertices, edges, and triangles in X . The graphs $K_{1,4}$ and $K_1 \cup C_4$ are cospectral and do not have the same number of 4-cycles, so it is difficult to extend these observations.

8.2 The Incidence Matrix

The *incidence matrix* $B(X)$ of a graph X is the 01-matrix with rows and columns indexed by the vertices and edges of X , respectively, such that the uf -entry of $B(X)$ is equal to one if and only if the vertex u is in the edge f . If X has n vertices and e edges, then $B(X)$ has order $n \times e$.

The rank of the adjacency matrix of a graph can be computed in polynomial time, but we do not have a simple combinatorial expression for it. We do have one for the rank of the incidence matrix.

Theorem 8.2.1 *Let X be a graph with n vertices and c_0 bipartite connected components. If B is the incidence matrix of X , then its rank is given by $\text{rk } B = n - c_0$.*

Proof. We shall show that the null space of B has dimension c_0 , and hence that $\text{rk } B = n - c_0$. Suppose that z is a vector in \mathbb{F}^n such that $z^T B = 0$. If uv is an edge of X , then $z_u + z_v = 0$. It follows by an easy induction that if u and v are vertices of X joined by a path of length r , then $z_u = (-1)^r z_v$. Therefore, if we view z as a function on $V(X)$, it is identically zero on any component of X that is not bipartite, and takes equal and opposite values on the two colour classes of any bipartite component. The space of such vectors has dimension c_0 . \square

The inner product of two columns of $B(X)$ is nonzero if and only if the corresponding edges have a common vertex, which immediately yields the following result.

Lemma 8.2.2 *Let B be the incidence matrix of the graph X , and let L be the line graph of X . Then $B^T B = 2I + A(L)$. \square*

If X is a graph on n vertices, let $\Delta(X)$ be the diagonal $n \times n$ matrix with rows and columns indexed by $V(X)$ with uu -entry equal to the valency of vertex u . The inner product of any two distinct rows of $B(X)$ is equal to the number of edges joining the corresponding vertices. Thus it is zero or one according as these vertices are adjacent or not, and we have the following:

Lemma 8.2.3 *Let B be the incidence matrix of the graph X . Then $BB^T = \Delta(X) + A(X)$. \square*

When X is regular the last two results imply a simple relation between the eigenvalues of $L(X)$ and those of X , but to prove this we also need the following result.

Lemma 8.2.4 *If C and D are matrices such that CD and DC are both defined, then $\det(I - CD) = \det(I - DC)$.*

Proof. If

$$X = \begin{pmatrix} I & C \\ D & I \end{pmatrix}, \quad Y = \begin{pmatrix} I & 0 \\ -D & I \end{pmatrix},$$

then

$$XY = \begin{pmatrix} I - CD & C \\ 0 & I \end{pmatrix}, \quad YX = \begin{pmatrix} I & C \\ 0 & I - DC \end{pmatrix},$$

and since $\det XY = \det YX$, it follows that $\det(I - CD) = \det(I - DC)$. \square

This result implies that $\det(I - x^{-1}CD) = \det(I - x^{-1}DC)$, from which it follows that CD and DC have the same nonzero eigenvalues with the same multiplicities.

Lemma 8.2.5 *Let X be a regular graph of valency k with n vertices and e edges and let L be the line graph of X . Then*

$$\phi(L, x) = (x + 2)^{e-n} \phi(X, x - k + 2).$$

Proof. Substituting $C = x^{-1}B^T$ and $D = B$ into the previous lemma we get

$$\det(I_e - x^{-1}B^T B) = \det(I_n - x^{-1}BB^T),$$

whence

$$\det(xI_e - B^T B) = x^{e-n} \det(xI_n - BB^T).$$

Noting that $\Delta(X) = kI$ and using Lemma 8.2.2 and Lemma 8.2.3, we get

$$\det((x - 2)I_e - A(L)) = x^{e-n} \det((x - k)I_n - A(X)),$$

and so

$$\phi(L, x - 2) = x^{e-n} \phi(X, x - k),$$

whence our claim follows. \square

8.3 The Incidence Matrix of an Oriented Graph

An *orientation* of a graph X is the assignment of a direction to each edge; this means that we declare one end of the edge to be the *head* of the edge and the other to be the *tail*, and view the edge as oriented from its tail to its head. Although this definition should be clear, we occasionally need a more formal version. Recall that an arc of a graph is an ordered pair of adjacent vertices. An orientation of X can then be defined as a function σ from the arcs of X to $\{-1, 1\}$ such that if (u, v) is an arc, then

$$\sigma(u, v) = -\sigma(v, u).$$

If $\sigma(u, v) = 1$, then we will regard the edge uv as oriented from tail u to head v .

An *oriented graph* is a graph together with a particular orientation. We will sometimes use X^σ to denote the oriented graph determined by the specific orientation σ . (You may, if you choose, view oriented graphs as a special class of directed graphs. We tend to view them as graphs with extra structure.) Figure 8.2 shows an example of an oriented graph, using arrows to indicate the orientation.

The *incidence matrix* $D(X^\sigma)$ of an oriented graph X^σ is the $\{0, \pm 1\}$ -matrix with rows and columns indexed by the vertices and edges of X , respectively, such that the uf -entry of $D(X^\sigma)$ is equal to 1 if the vertex u is the head of the edge f , -1 if u is the tail of f , and 0 otherwise. If X




Figure 8.2. An oriented graph

has n vertices and e edges, then $D(X^\sigma)$ has order $n \times e$. For example, the incidence matrix of the graph of Figure 8.2 is

$$\begin{pmatrix} -1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 & 0 & 1 \end{pmatrix}.$$

Although there are many different ways to orient a given graph, many of the results about oriented graphs are independent of the choice of orientation. For example, the next result shows that the rank of the incidence matrix of an oriented graph depends only on X , rather than on the particular orientation given to X .

Theorem 8.3.1 *Let X be a graph with n vertices and c connected components. If σ is an orientation of X and D is the incidence matrix of X^σ , then $\text{rk } D = n - c$.*

Proof. We shall show that the null space of D has dimension c , and hence that $\text{rk } D = n - c$. Suppose that z is a vector in \mathbb{R}^n such that $z^T D = 0$. If uv is an edge of X , then $z_u - z_v = 0$. Therefore, if we view z as a function on $V(X)$, it is constant on any connected component of X . The space of such vectors has dimension c . \square

We note the following analogue to Lemma 8.2.3.

Lemma 8.3.2 *If σ is an orientation of X and D is the incidence matrix of X^σ , then $DD^T = \Delta(X) - A(X)$.* \square

If X is a plane graph, then each orientation of X determines an orientation of its dual. This orientation is obtained by viewing each edge of X^* as arising from rotating the corresponding edge of X through 90° clockwise (as in Figure 8.3). We will use σ to denote the orientation of both X and X^* .




Figure 8.3. Orienting the edges of the dual

Lemma 8.3.3 *Let X and Y be dual plane graphs, and let σ be an orientation of X . If D and E are the incidence matrices of X^σ and Y^σ , then $DE^T = 0$.*

Proof. If u is an edge of X and F is a face, there are exactly two edges on u and in F . Denote them by g and h and assume, for convenience, that g precedes h as we go clockwise around F . Then the uF -entry of DE^T is equal to

$$D_{ug}E_{gF}^T + D_{uh}E_{hF}^T.$$

If the orientation of the edge g is reversed, then the value of the product $D_{ug}E_{gF}^T$ does not change. Hence the value of the sum is independent of the orientation σ , and so we may assume that g has head u and that h has tail u . This implies that the edges in Y corresponding to g and h both have head F , and a simple computation now yields that the sum is zero. \square

8.4 Symmetric Matrices

In this section we review the main results of the linear algebra of symmetric matrices over the real numbers, which form the basis for the remainder of this chapter.

Lemma 8.4.1 *Let A be a real symmetric matrix. If u and v are eigenvectors of A with different eigenvalues, then u and v are orthogonal.*

Proof. Suppose that $Au = \lambda u$ and $Av = \tau v$. As A is symmetric, $u^T Av = (v^T Au)^T$. However, the left-hand side of this equation is $\tau u^T v$ and the right-hand side is $\lambda u^T v$, and so if $\tau \neq \lambda$, it must be the case that $u^T v = 0$. \square

Lemma 8.4.2 *The eigenvalues of a real symmetric matrix A are real numbers.*

Proof. Let u be an eigenvector of A with eigenvalue λ . Then by taking the complex conjugate of the equation $Au = \lambda u$ we get $A\bar{u} = \bar{\lambda}\bar{u}$, and so \bar{u} is also an eigenvector of A . Now, by definition an eigenvector is not zero, so $u^T \bar{u} > 0$. By the previous lemma, u and \bar{u} cannot have different eigenvalues, so $\lambda = \bar{\lambda}$, and the claim is proved. \square

We shall now prove that a real symmetric matrix is diagonalizable. For this we need a simple lemma that expresses one of the most important properties of symmetric matrices. A subspace U is said to be *A-invariant* if $Au \in U$ for all $u \in U$.

Lemma 8.4.3 *Let A be a real symmetric $n \times n$ matrix. If U is an A -invariant subspace of \mathbb{R}^n , then U^\perp is also A -invariant.*

Proof. For any two vectors u and v , we have

$$v^T(Au) = (Av)^T u.$$

If $u \in U$, then $Au \in U$; hence if $v \in U^\perp$, then $v^T Au = 0$. Consequently, $(Av)^T u = 0$ whenever $u \in U$ and $v \in U^\perp$. This implies that $Av \in U^\perp$ whenever $v \in U^\perp$, and therefore U^\perp is A -invariant. \square

Any square matrix has at least one eigenvalue, because there must be at least one solution to the polynomial equation $\det(xI - A) = 0$. Hence a real symmetric matrix A has at least one real eigenvalue, θ say, and hence at least one real eigenvector (any vector in the kernel of $A - \theta I$, to be precise). Our next result is a crucial strengthening of this fact.

Lemma 8.4.4 *Let A be an $n \times n$ real symmetric matrix. If U is a nonzero A -invariant subspace of \mathbb{R}^n , then U contains a real eigenvector of A .*

Proof. Let R be a matrix whose columns form an orthonormal basis for U . Then, because U is A -invariant, $AR = RB$ for some square matrix B . Since $R^T R = I$, we have

$$R^T AR = R^T RB = B,$$

which implies that B is symmetric, as well as real. Since every symmetric matrix has at least one eigenvalue, we may choose a real eigenvector u of B with eigenvalue λ . Then $ARu = RBu = \lambda Ru$, and since $u \neq 0$ and the columns of R are linearly independent, $Ru \neq 0$. Therefore, Ru is an eigenvector of A contained in U . \square

Theorem 8.4.5 *Let A be a real symmetric $n \times n$ matrix. Then \mathbb{R}^n has an orthonormal basis consisting of eigenvectors of A .*

Proof. Let $\{u_1, \dots, u_m\}$ be an orthonormal (and hence linearly independent) set of $m < n$ eigenvectors of A , and let M be the subspace that they span. Since A has at least one eigenvector, $m \geq 1$. The subspace M is A -invariant, and hence M^\perp is A -invariant, and so M^\perp contains a (normalized) eigenvector u_{m+1} . Then $\{u_1, \dots, u_m, u_{m+1}\}$ is an orthonormal set of $m + 1$ eigenvectors of A . Therefore, a simple induction argument shows that a set consisting of one normalized eigenvector can be extended to an orthonormal basis consisting of eigenvectors of A . \square

Corollary 8.4.6 *If A is an $n \times n$ real symmetric matrix, then there are matrices L and D such that $L^T L = LL^T = I$ and $LAL^T = D$, where D is the diagonal matrix of eigenvalues of A .*

Proof. Let L be the matrix whose rows are an orthonormal basis of eigenvectors of A . We leave it as an exercise to show that L has the stated properties. \square

8.5 Eigenvectors

Most introductory linear algebra courses impart the belief that the way to compute the eigenvalues of a matrix is to find the zeros of its characteristic polynomial. For matrices with order greater than two, this is false. Generally, the best way to obtain eigenvalues is to find eigenvectors: If $Ax = \theta x$, then θ is an eigenvalue of A .

When we work with graphs there is an additional refinement. First, we stated in Section 8.1 that the rows and columns of $A(X)$ are indexed by the vertices of X . Formally, this means we are viewing $A(X)$ as a linear mapping on $\mathbb{F}^{V(X)}$, the space of real functions on $V(X)$ (rather than on the isomorphic vector space \mathbb{F}^n , where $n = |V(X)|$). If $f \in \mathbb{F}^{V(X)}$ and $A = A(X)$, then the image Af of f under A is given by

$$(Af)(u) = \sum A_{uv} f(v);$$

since A is a 01-matrix, it follows that


$$(Af)(u) = \sum_{v \sim u} f(v).$$

In words, the value of Af at u is the sum of the values of f on the neighbours of u . If we suppose that f is an eigenvector of A with eigenvalue θ , then $Af = \theta f$, and so

$$\theta f(u) = \sum_{v \sim u} f(v).$$

In words, the sum of the values of f on the neighbours of u is equal to θ times the value of f at u . Conversely, any function f that satisfies this condition is an eigenvector of X . Figure 8.4 shows an eigenvector of the Petersen graph. It can readily be checked that the sum of the values on the neighbours of any vertex is equal to the value on that vertex; hence we have an eigenvector with eigenvalue one. (The viewpoint expressed in this paragraph is very fruitful, and we will make extensive use of it.)

Now, we will find the eigenvalues of the cycle C_n . Take the vertex set of C_n to be $\{0, 1, \dots, n-1\}$. Let τ be an n th root of unity (so τ is probably

Figure 8.4. An eigenvector of P with eigenvalue 1

not a real number) and define $f(u) := \tau^u$. Then for all vertices u ,

$$\sum_{v \sim u} f(v) = (\tau^{-1} + \tau)\tau^u,$$

and therefore $\tau^{-1} + \tau$ is an eigenvalue of C_n . Note that this is real, even if τ is not. By varying our choice of τ we find the n eigenvalues of C_n . This argument is easily extended to any circulant graph.

By taking $\tau = 1$ we see that the vector with all entries equal to one is an eigenvector of C_n with eigenvalue two. We shall denote this eigenvector by $\mathbf{1}$. It is clear that $\mathbf{1}$ is an eigenvector of a graph X with eigenvalue k if and only if X is regular with valency k . We can say more about regular graphs.

Lemma 8.5.1 *Let X be a k -regular graph on n vertices with eigenvalues $k, \theta_2, \dots, \theta_n$. Then X and its complement \bar{X} have the same eigenvectors, and the eigenvalues of \bar{X} are $n - k - 1, -1 - \theta_2, \dots, -1 - \theta_n$.*

Proof. The adjacency matrix of the complement \bar{X} is given by

$$A(\bar{X}) = J - I - A(X),$$

where J is the all-ones matrix. Let $\{\mathbf{1}, u_2, \dots, u_n\}$ be an orthonormal basis of eigenvectors of $A(X)$. Then $\mathbf{1}$ is an eigenvector of \bar{X} with eigenvalue $n - k - 1$. For $2 \leq i \leq n$, the eigenvector u_i is orthogonal to $\mathbf{1}$, and so

$$A(\bar{X})u_i = (J - I - A(X))u_i = (-1 - \theta_i)u_i.$$

Therefore, u_i is an eigenvector of $A(\bar{X})$ with eigenvalue $-1 - \theta_i$. \square

Finally, suppose that X is a semiregular bipartite graph with bipartition $V(X) = V_1 \cup V_2$, and let k and ℓ be the valencies of the vertices in V_1 and V_2 , respectively. Assume that u_1 is a vertex with valency k , and u_2 is a vertex with valency ℓ . We look for an eigenvector f that is constant on the two parts of the bipartition. If f is such an eigenvector and has eigenvalue

θ , then

$$\theta f(u_1) = kf(u_2), \quad \theta f(u_2) = \ell f(u_1).$$

Because an eigenvector is a nonzero vector, we can multiply the two equations just given to obtain

$$\theta^2 = k\ell.$$

Thus, if $\theta = \pm\sqrt{k\ell}$, then defining f by

$$f(u) = \begin{cases} 1, & \text{if } u \in V_1, \\ \theta/k, & \text{if } u \in V_2, \end{cases}$$

yields two eigenvectors of X .

We comment on a feature of the last example. If A is the adjacency matrix of a graph X , and f is a function on $V(X)$, then so is Af . If X is a semiregular bipartite graph, then the space of functions on $V(X)$ that are constant on the two parts of the bipartition is A -invariant. (Indeed, this is equivalent to the fact that X is bipartite and semiregular.) By Lemma 8.4.4, an A -invariant subspace must contain an eigenvector of A ; in the above example this subspace has dimension two, and the eigenvector is easy to find. In Section 9.3 we introduce and study equitable partitions, which provide many further examples of A -invariant subspaces.

8.6 Positive Semidefinite Matrices

A real symmetric matrix A is *positive semidefinite* if $u^T Au \geq 0$ for all vectors u . It is *positive definite* if it is positive semidefinite and $u^T Au = 0$ if and only if $u = 0$. (These terms are used only for symmetric matrices.) Observe that a positive semidefinite matrix is positive definite if and only if it is invertible.

There are a number of characterizations of positive semidefinite matrices. The first we offer involves eigenvalues. If u is an eigenvector of A with eigenvalue θ , then

$$u^T Au = \theta u^T u,$$

and so we see that a real symmetric matrix is positive semidefinite if and only if its eigenvalues are nonnegative.

Our second characterization involves a factorization. If $A = B^T B$ for some matrix B , then

$$u^T Au = u^T B^T B u = (Bu)^T B u \geq 0,$$

and therefore A is positive semidefinite. The *Gram matrix* of vectors u_1, \dots, u_n from \mathbb{F}^m is the $n \times n$ matrix G such that $G_{ij} = u_i^T u_j$. Note that $B^T B$ is the Gram matrix of the columns of B , and that any Gram

matrix is positive semidefinite. The next result shows that the converse is true.

Lemma 8.6.1 *If A is a positive semidefinite matrix, then there is a matrix B such that $A = B^T B$.*

Proof. Since A is symmetric, there is a matrix L such that

$$A = L^T \Lambda L,$$

where Λ is the diagonal matrix with i th entry equal to the i th eigenvalue of A . Since A is positive semidefinite, the entries of Λ are nonnegative, and so there is a diagonal matrix D such that $D^2 = \Lambda$. If $B = L^T D L$, then $B = B^T$ and $A = B^2 = B^T B$, as required. \square

We can now establish some interesting results about the eigenvalues of graphs, the first being about line graphs.

Let $\theta_{\max}(X)$ and $\theta_{\min}(X)$ respectively denote the largest and smallest eigenvalues of $A(X)$.

Lemma 8.6.2 *If L is a line graph, then $\theta_{\min}(L) \geq -2$.*

Proof. If L is the line graph of X and B is the incidence matrix of X , we have

$$A(L) + 2I = B^T B.$$

Since $B^T B$ is positive semidefinite, its eigenvalues are nonnegative and all eigenvalues of $B^T B - 2I$ are at least -2 . \square

What is surprising about this lemma is how close it comes to characterizing line graphs. We will study this question in detail in Chapter 12.

Lemma 8.6.3 *Let Y be an induced subgraph of X . Then*

$$\theta_{\min}(X) \leq \theta_{\min}(Y) \leq \theta_{\max}(Y) \leq \theta_{\max}(X).$$

Proof. Let A be the adjacency matrix of X and abbreviate $\theta_{\max}(X)$ to θ . The matrix $\theta I - A$ has only nonnegative eigenvalues, and is therefore positive semidefinite. Let f be any vector that is zero on the vertices of X not in Y , and let f_Y be its restriction to $V(Y)$. Then

$$0 \leq f^T(\theta I - A)f = f_Y^T(\theta I - A(Y))f_Y,$$

from which we deduce that $\theta I - A(Y)$ is positive semidefinite. Hence $\theta_{\max}(Y) \leq \theta$. A similar argument applied to $A - \theta_{\min}(X)I$ yields the second claim of the lemma. \square

It is actually true that if Y is any subgraph of X , and not just an induced subgraph, then $\theta_{\max}(Y) \leq \theta_{\max}(X)$. Furthermore, when Y is a proper subgraph, equality can hold only when X is not connected. We return

to this when we discuss the Perron–Frobenius theorem in the next two sections.

Finally, we clear a debt incurred in Section 5.10. There we claimed that the matrix

$$(r - \lambda)I + \lambda J$$

is invertible when $r > \lambda \geq 0$. Note that $(r - \lambda)I$ is positive definite: All its eigenvalues are positive and $\lambda J = \lambda \mathbf{1}\mathbf{1}^T$ is positive semidefinite. But the sum of a positive definite and a positive semidefinite matrix is positive definite, and therefore invertible.

8.7 Subharmonic Functions

In this section we introduce subharmonic functions, and use them to develop some properties of nonnegative matrices. We will use similar techniques again in Section 13.9, when we show how linear algebra can be used to construct drawings of planar graphs.

If A is a square matrix, then we say that a nonnegative vector x is λ -subharmonic for A if $x \neq 0$ and $Ax \geq \lambda x$. When the value of λ is irrelevant, we simply say that x is subharmonic. We note one way that subharmonic vectors arise. Let $|A|$ denote the matrix obtained by replacing each entry of A with its absolute value. If x is an eigenvector for A with eigenvalue θ , then

$$|\theta| |x_i| = |\theta x_i| = |(Ax)_i| = \left| \sum_j A_{ij} x_j \right| \leq \sum_j |A_{ij}| |x_j|,$$

from which we see that $|x|$ is $|\theta|$ -subharmonic for $|A|$.

Let A be an $n \times n$ real matrix. The *underlying directed graph* of A has vertex set $\{1, \dots, n\}$, with an arc from vertex i to vertex j if and only if $A_{ij} \neq 0$. (Note that this directed graph may have loops.) A square matrix is irreducible if its underlying graph is strongly connected.

Lemma 8.7.1 *Let A be an $n \times n$ nonnegative irreducible matrix. Then there is a maximum real number ρ such that there is a ρ -subharmonic vector for A . Moreover, any ρ -subharmonic vector x is an eigenvector for A with eigenvalue ρ , and all entries of x are positive.*

Proof. Let

$$F(x) = \min_{i: x_i \neq 0} \frac{(Ax)_i}{x_i}$$

be a function defined on the set of nonnegative vectors, and consider the values of F on the vectors in the set

$$S = \{x : x \geq 0, \mathbf{1}^T x = 1\}.$$

It is clear that any nonnegative vector x is $F(x)$ -subharmonic, and so we wish to show that there is some vector $y \in S$ such that F attains its maximum on y . Since S is compact, this would be immediate if F were continuous on S , but this is not the case at the boundary of S . As A is irreducible, Lemma 8.1.2 shows that the matrix $(I + A)^{n-1}$ is positive. Therefore, the set

$$T = (I + A)^{n-1}S$$

contains only positive vectors, and F is continuous on T . Since T is also compact, it follows that F attains its maximum value ρ at a point $z \in T$. If we set

$$y = \frac{z}{\mathbf{1}^T z},$$

then $y \in S$ and $F(y) = F(z) = \rho$. Moreover, for any vector x we have

$$F((I + A)^{n-1}(x)) \geq F(x),$$

and therefore by the choice of z , there is no vector $x \in S$ with $F(x) > \rho$.

We now prove that any ρ -subharmonic vector is an eigenvector for A , necessarily with eigenvalue ρ . If x is ρ -subharmonic, define $\sigma(x)$ by

$$\sigma(x) = \{i : (Ax)_i > \rho x_i\}.$$

Clearly, x is an eigenvector if and only if $\sigma(x) = \emptyset$. Assume by way of contradiction that $\sigma(x) \neq \emptyset$. The support of a vector v is the set of nonzero coordinates of v and is denoted by $\text{supp}(v)$. Let h be a nonnegative vector with support equal to $\sigma(x)$ and consider the vector $y = x + \epsilon h$.

We have

$$(Ay)_i - \rho y_i = (Ax)_i - \rho x_i + \epsilon(Ah)_i - \epsilon \rho h_i.$$

If $i \in \sigma(x)$, then $(Ax)_i > \rho x_i$, and so for all sufficiently small values of ϵ , the right side of (8.7) is positive. Hence

$$(Ay)_i > \rho y_i.$$

If $i \notin \sigma(x)$, then $(Ax)_i = \rho x_i$ and $h_i = 0$, so (8.7) yields that

$$(Ay)_i - \rho y_i = \epsilon(Ah)_i.$$

Provided that $\epsilon > 0$, the right side here is nonnegative. Since A is irreducible, there is at least one value of i not in $\sigma(x)$ such that $(Ah)_i > 0$, and hence $\sigma(y)$ properly contains $\sigma(x)$.

If $|\sigma(y)| = n$, it follows that y is ρ' -subharmonic, where $\rho' > \rho$, and this is a contradiction to our choice of ρ . Otherwise, y is ρ -subharmonic but $|\sigma(y)| > |\sigma(x)|$, and we may repeat the above argument with y in place of x . After a finite number of iterations we will arrive at a ρ' -subharmonic vector, with $\rho' > \rho$, again a contradiction.

Finally, we prove that if x is ρ -subharmonic, then $x > 0$. Suppose instead that $x_i = 0$ for some i . Because $\sigma(x) = \emptyset$, it follows that $(Ax)_i = 0$, but

$$(Ax)_i = \sum_j A_{ij}x_j,$$

and since $A \geq 0$, this implies that $x_j = 0$ if $A_{ij} \neq 0$. Since A is irreducible, a simple induction argument yields that all entries of x must be zero, which is the required contradiction. Therefore, x must be positive. \square

The spectral radius $\rho(A)$ of a matrix A is the maximum of the moduli of its eigenvalues. (If A is not symmetric, these eigenvalues need not be real numbers.) The spectral radius of a matrix need not be an eigenvalue of it, e.g., if $A = -I$, then $\rho(A) = 1$. One consequence of our next result is that the real number ρ from the previous lemma is the spectral radius of A .

Lemma 8.7.2 *Let A be an $n \times n$ nonnegative irreducible matrix and let ρ be the greatest real number such that A has a ρ -subharmonic vector. If B is an $n \times n$ matrix such that $|B| \leq A$ and $Bx = \theta x$, then $|\theta| \leq \rho$. If $|\theta| = \rho$, then $|B| = A$ and $|x|$ is an eigenvector of A with eigenvalue ρ .*

Proof. If $Bx = \theta x$, then

$$|\theta||x| = |\theta x| = |Bx| \leq |B||x| \leq A|x|.$$

Hence $|x|$ is $|\theta|$ -subharmonic for A , and so $|\theta| \leq \rho$. If $|\theta| = \rho$, then $A|x| = |B||x| = \rho|x|$, and by the previous lemma, $|x|$ is positive. Since $A - |B| \geq 0$ and $(A - |B|)|x| = 0$, it follows that $A = |B|$. \square

Lemma 8.7.3 *Let A be a nonnegative irreducible $n \times n$ matrix with spectral radius ρ . Then ρ is a simple eigenvalue of A , and if x is an eigenvector with eigenvalue ρ , then all entries of x are nonzero and have the same sign.*

Proof. The ρ -eigenspace of A is 1-dimensional, for otherwise we could find a ρ -subharmonic vector with some entry equal to zero, contradicting Lemma 8.7.1. If x is an eigenvector with eigenvalue ρ , then by the previous lemma, $|x|$ is a positive eigenvector with the same eigenvalue. Thus $|x|$ is a multiple of x , which implies that all the entries of x have the same sign.

Since the geometric multiplicity of ρ is 1, we see that $K = \ker(A - \rho I)$ has dimension 1 and the column space C of $A - \rho I$ has dimension $n - 1$. If C contains x , then we can find a vector y such that $x = (A - \rho I)y$. For any k , we have $(A - \rho I)(y + kx) = x$, and so by taking k sufficiently large, we may assume that y is positive. But then y is ρ -subharmonic and hence is a multiple of x , which is impossible. Therefore, we conclude that $K \cap C = 0$, and that \mathbb{F}^n is the direct sum of K and C . Since K and C are A -invariant, this implies that the characteristic polynomial $\varphi(A, t)$ of A is the product of $t - \rho$ and the characteristic polynomial of A restricted to C . As x is not in C , all eigenvectors of A contained in C have eigenvalue different from ρ , and so we conclude that ρ is a simple root of $\varphi(A, t)$, and hence has algebraic multiplicity one. \square

8.8 The Perron–Frobenius Theorem

The Perron–Frobenius theorem is the most important result on the eigenvalues and eigenvectors of nonnegative matrices.

Theorem 8.8.1 *Suppose A is a real nonnegative $n \times n$ matrix whose underlying directed graph X is strongly connected. Then:*

- (a) $\rho(A)$ is a simple eigenvalue of A . If x is an eigenvector for ρ , then no entries of x are zero, and all have the same sign.
- (b) Suppose A_1 is a real nonnegative $n \times n$ matrix such that $A - A_1$ is nonnegative. Then $\rho(A_1) \leq \rho(A)$, with equality if and only if $A_1 = A$.
- (c) If θ is an eigenvalue of A and $|\theta| = \rho(A)$, then $\theta/\rho(A)$ is an m th root of unity and $e^{2\pi ir/m}\rho(A)$ is an eigenvalue of A for all r . Further, all cycles in X have length divisible by m . \square

The first two parts of this theorem follow from the results of the previous section. We discuss part (c), but do not give a complete proof of it, since we will not need its full strength.

Suppose A is the adjacency matrix of a connected graph X , with spectral radius ρ , and assume that θ is an eigenvalue of A such that $|\theta| = \rho$. If $\theta \neq \rho$, then $\theta = -\rho$, and so θ/ρ is a root of unity. If z_0 and z_1 are eigenvectors with eigenvalues θ and ρ , respectively, then they are linearly independent, and therefore the eigenspace of A^2 with eigenvalue ρ^2 has dimension at least two. However, it is easy to see that ρ^2 is the spectral radius of A^2 . As A^2 is nonnegative, it follows from part (a) of the theorem that the underlying graph of A^2 cannot be connected, and given this, it is easy to prove that X must be bipartite.

It is not hard to see that if X is bipartite, then there is a graph isomorphic to X with adjacency matrix of the form

$$A = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix},$$

for a suitable 01-matrix B . If the partitioned vector (x, y) is an eigenvector of A with eigenvalue θ , then it is easy to verify that $(x, -y)$ is an eigenvector of A with eigenvalue $-\theta$. It follows that θ and $-\theta$ are eigenvalues with the same multiplicity. Thus we have the following:

Theorem 8.8.2 *Let A be the adjacency matrix of the graph X , and let ρ be its spectral radius. Then the following are equivalent:*

- (a) X is bipartite.
- (b) The spectrum of A is symmetric about the origin, i.e., for any θ , the multiplicities of θ and $-\theta$ as eigenvalues of A are the same.
- (c) $-\rho$ is an eigenvalue of A . \square

There are two common applications of the Perron–Frobenius theorem to connected regular graphs. Let X be a connected k -regular graph with adjacency matrix A . Then the spectral radius of A is the valency k with corresponding eigenvector $\mathbf{1}$, which implies that every other eigenspace of A is orthogonal to $\mathbf{1}$. Secondly, the graph X is bipartite if and only if $-k$ is an eigenvalue of A .

8.9 The Rank of a Symmetric Matrix

The rank of a matrix is a fundamental algebraic concept, and so it is natural to ask what information about a graph can be deduced from the rank of its adjacency matrix. In contrast to what we obtained for the incidence matrix, there is no simple combinatorial expression for the rank of the adjacency matrix of a graph. This section develops a number of preliminary results about the rank of a symmetric matrix that will be used later.

Theorem 8.9.1 *Let A be a symmetric matrix of rank r . Then there is a permutation matrix P and a principal $r \times r$ submatrix M of A such that*

$$P^T A P = \begin{pmatrix} I \\ R \end{pmatrix} M (I \quad R^T).$$

Proof. Since A has rank r , there is a linearly independent set of r rows of A . By symmetry, the corresponding set of columns is also linearly independent. The entries of A in these rows and columns determine an $r \times r$ principal submatrix M . Therefore, there is a permutation matrix P such that

$$P^T A P = \begin{pmatrix} M & N^T \\ N & H \end{pmatrix}.$$

Since the first r rows of this matrix generate the row space of $P^T A P$, we have that $N = RM$ for some matrix R , and hence $H = RN^T = RMR^T$. Therefore,

$$P^T A P = \begin{pmatrix} M & MR \\ RM & RMR^T \end{pmatrix} = \begin{pmatrix} I \\ R \end{pmatrix} M (I \quad R^T)$$

as claimed. \square

We note an important corollary of this result.

Corollary 8.9.2 *If A is a symmetric matrix of rank r , then it has a principal $r \times r$ submatrix of full rank.* \square

If a matrix A has rank one, then it is necessarily of the form $A = xy^T$ for some nonzero vectors x and y . It is not too hard to see that if a matrix can be written as the sum of r rank-one matrices, then it has rank at most