Chapter 3

Graph Spectrum

The linear algebraic approach to algorithmic graph theory is to view graphs as matrices, and use concepts and tools in linear algebra to design and analyze algorithms for graph problems. Spectral graph theory uses eigenvalues and eigenvectors of matrices associated with the graph to study its combinatorial properties. In this chapter, we consider the adjacency matrix and the Laplacian matrix of a graph, and see some basic results in spectral graph theory. A general reference for this chapter is the upcoming book by Spielman [Spi19].

3.1 Adjacency Matrix

Definition 3.1 (Adjacency Matrix)**.** *Given an undirected graph $G = (V, E)$ with $V(G) = [n]$, the adjacency matrix $A(G)$ is an $n \times n$ matrix with $A_{ij} = A_{ji} = 1$ if $ij \in E(G)$ and $A_{ij} = A_{ji} = 0$ otherwise.*

The adjacency matrix of an undirected graph is symmetric. So, by the spectral theorem for real symmetric matrices in Theorem 2.5, the adjacency matrix has an orthonormal basis of eigenvectors with real eigenvalues. We denote the eigenvalues of the adjacency matrix by

$$\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n.$$

It is not clear that these eigenvalues should provide any useful information about the combinatorial properties of the graph, but they do, and surprisingly much information can be obtained from them. Let's start with some examples and compute their spectrum.

Example 3.2 (Complete Graphs)**.** *If G is a complete graph, then $A(G) = J - I$ where J denotes the all-one matrix. Any vector is an eigenvector of I with eigenvalue 1. Hence the eigenvalues of A are one less than that of J. Since J is of rank 1, there are $n-1$ eigenvalues of 0. The all-one vector is an eigenvector of J with eigenvalue n. So, $n-1$ is an eigenvalue of A with multiplicity one, and -1 is an eigenvalue of A with multiplicity $n-1$. The is the example with the largest gap between the largest eigenvalue and the second largest eigenvalue.*

Example 3.3 (Complete Bipartite Graphs)**.** *Let $K_{p,q}$ be the complete bipartite graph with p vertices on one side and q vertices on the other side. Its adjacency matrix $A(K_{p,q})$ is of rank 2, so 0 is an eigenvalue with multiplicity $p+q-2$, and there are two non-zero eigenvalues α and β. By Fact 2.32, the sum of the eigenvalues is equal to the trace of A, which is equal to zero as there* *are no self-loops, and so* $\alpha = -\beta$. *To determine* α*, we consider the characteristic polynomial* $\det(xI - A) = (x-\alpha)(x+\alpha)x^{p+q-2} = x^{p+q} - \alpha^2 x^{p+q-2}$. *Using the Leibniz formula for determinant in Fact 2.23, any term that contributes to* x^{p+q-2} *must have* $p+q-2$ *diagonal entries, and the remaining two entries must be* $-A_{ij}$ *and* $-A_{ji}$ *for some* i, j. *There are totally* pq *such terms, one for each edge, where the sign of the corresponding permutation is* -1 *as it only has one inversion pair. So,* $\alpha^2 = pq$*, and thus* $|\alpha| = \sqrt{pq}$. *To conclude, the spectrum is* $(\sqrt{pq}, 0, \ldots, 0, -\sqrt{pq})$*, where* 0 *is an eigenvalue with multiplicity* $p+q-2$.

Bipartiteness

It turns out that bipartite graphs can be characterized by the spectrum of their adjacency matrix. The following lemma says that the spectrum of a bipartite graph must be symmetric around the origin on the real line.

Lemma 3.4 (Spectrum of Bipartite Graph is Symmetric)**.** *If* G *is a bipartite graph and* α *is an eigenvalue of* $A(G)$ *with multiplicity* k*, then* $-\alpha$ *is an eigenvalue of* $A(G)$ *with multiplicity* k.

Proof. If G is a bipartite graph, then we can permute the rows and columns of G to obtain the form

$$A(G) = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}.$$

Suppose $u = \begin{pmatrix} x \\ y \end{pmatrix}$ is an eigenvector of $A(G)$ with eigenvalue α. Then

$$\begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \alpha \begin{pmatrix} x \\ y \end{pmatrix} \qquad \Longrightarrow \qquad B^T x = \alpha y \ \text{ and } \ By = \alpha x.$$

It follows that

$$\begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ -y \end{pmatrix} = \begin{pmatrix} -By \\ B^T x \end{pmatrix} = \begin{pmatrix} -\alpha x \\ \alpha y \end{pmatrix} = -\alpha \begin{pmatrix} x \\ -y \end{pmatrix},$$

and thus $\begin{pmatrix} x \\ -y \end{pmatrix}$ is an eigenvector of $A(G)$ with eigenvalue $-\alpha$. By this construction, note that k linearly independent eigenvectors with eigenvalue α would give k linearly independent eigenvectors with eigenvalue $-\alpha$, and so their multiplicity is the same. □

The next lemma shows that the converse is also true.

Lemma 3.5 (Symmetric Spectrum Implies Bipartiteness)**.** *Let* G *be an undirected graph and let* $\alpha_1 \geq \ldots \geq \alpha_n$ *be the eigenvalues of its adjacency matrix. If* $\alpha_i = -\alpha_{n-i+1}$ *for each* $1 \leq i \leq n$*, then* G *is a bipartite graph.*

Proof. Let k be any positive odd number. Then $\sum_{i=1}^n \alpha_i^k = 0$, by the symmetry of the spectrum. Note that $\alpha_1^k \geq \alpha_2^k \geq \ldots \geq \alpha_n^k$ are the eigenvalues of A^k, because if $Av = \alpha v$ then $A^k v = \alpha^k v$. By Fact 2.32, it follows that $\mathrm{Tr}(A^k) = \sum_{i=1}^n \alpha_i^k = 0$. Observe that $A^k_{i,j}$ is the number of length k walks from i to j in G, which can be proved by a simple induction. So, if G has an odd cycle of length k, then $A^k_{i,i} > 0$ for each vertex i in the odd cycle, and this would imply that $\mathrm{Tr}(A^k) > 0$ as each diagonal entry $A^k_{i,i}$ is non-negative. Therefore, since $\mathrm{Tr}(A^k) = 0$, G must have no odd cycles and is thus a bipartite graph. □

Combining Lemma 3.4 and Lemma 3.5, a graph is bipartite if and only if the spectrum of its adjacency matrix is symmetric around the origin.

Proposition 3.6 (Spectral Characterization of Bipartite Graphs)**.** *Let G be an undirected graph and let $\alpha_1 \geq \ldots \geq \alpha_n$ be the eigenvalues of its adjacency matrix. Then G is a bipartite graph if and only if $\alpha_i = -\alpha_{n-i+1}$ for each $1 \leq i \leq n$.*

When the graph is connected, the characterization is even simpler.

Problem 3.7 (Spectral Characterization of Connected Bipartite Graphs)**.** *Let G be a connected undirected graph and let $\alpha_1 \geq \ldots \geq \alpha_n$ be the eigenvalues of its adjacency matrix. Prove that G is bipartite if and only if $\alpha_1 = -\alpha_n$.*

You may need to use the result of Perron-Frobenius in Theorem 2.16 and also the optimization formulation of eigenvalues in Definition 2.9 to solve this problem.

Largest Eigenvalue

We see some upper and lower bounds on the largest eigenvalue of the adjacnecy matrix in this subsection.

Lemma 3.8 (Max Degree Upper Bound)**.** *Let $G = (V, E)$ be an undirected graph with maximum degree d and let $\alpha_1 \geq \ldots \geq \alpha_n$ be the eigenvalues of its adjacency matrix. Then $\alpha_1 \leq d$.*

Proof. Let v be an eigenvector with eigenvalue α_1. Let j be a vertex with $v(j) \geq v(i)$ for all $i \in V(G)$. Then

$$\alpha_1 \cdot v(j) = (Av)(j) = \sum_{i:ij \in E(G)} v(i) \leq \sum_{i:ij \in E(G)} v(j) = \deg(j) \cdot v(j) \leq d \cdot v(j),$$

which implies that $\alpha_1 \leq d$. □

Look at the proof more closely, we can characterize the connected graphs with α_1 equal to the maximum degree.

Exercise 3.9 (Tight Max Degree Upper Bound)**.** *Let G be a connected undirected graph with maximum degree d and the largest eigenvalue $\alpha_1 = d$. Then G is d-regular and the eigenvalue α_1 is of multiplicity one.*

The maximum degree upper bound can be far from tight. The following problem provides such an example, whose bound is also important in the study of Ramanujan graphs in the second part of the course.

Problem 3.10 (Largest Eigenvalue of a Tree)**.** *Prove that the maximum eigenvalue of the adjacency matrix of a tree of maximum degree d is at most $2\sqrt{d-1}$.*

On the other hand, the average degree is a lower bound on the largest eigenvalue. More generally, the largest eigenvalue is at least the average degree of the densest induced subgraph. One corollary of the following exercise is that the largest eigenvalue is at least the size of a maximum clique minus one.

Exercise 3.11 (Average Degree Lower Bound). *Let $G = (V, E)$ be an undirected graph with largest eigenvalue α_1. For a subset $S \subseteq V$ and a vertex $v \in S$, let $\deg_S(v) := \big|\{u \mid uv \in E \text{ and } u \in S\}\big|$ be the degree of v induced in S. Then*

$$\alpha_1 \geq \max_{S:S\subseteq V} \frac{1}{|S|} \sum_{v \in S} \deg_S(v).$$

We remark that the largest eigenvalue of the adjacency matrix of a connected graph is always of multiplicity one by the Perron-Frobenius Theorem 2.16, and the spectrum of the adjacency matrix satisfies

$$d \geq \alpha_1 > \alpha_2 \geq \ldots \geq \alpha_n \geq -d.$$

In general, I do not know of a nice combinatorial characterization of the largest eigenvalue of the adjacency matrix.

Question 3.12 (Largest Eigenvalue of the Adjacency Matrix). *Is there a better "combinatorial" characterization of the largest eigenvalue of the adjacency matrix of an undirected graph?*

3.2 Laplacian Matrix

The Laplacian matrix plays a more important role in spectral graph theory than the adjacency matrix, as we will see some reasons soon.

Definition 3.13 (Diagonal Degree Matrix). *Let $G = (V, E)$ be an undirected graph with $V(G) = [n]$. The diagonal degree matrix $D(G)$ of G is the $n \times n$ diagonal matrix with $D_{i,i} = \deg(i)$ for each $1 \leq i \leq n$.*

Definition 3.14 (Laplacian Matrix). *Let G be an undirected graph. The Laplacian matrix $L(G)$ of G is defined as $L(G) := D(G) - A(G)$, where $D(G)$ is the diagonal degree matrix in Definition 3.13 and $A(G)$ is the adjacency matrix in Definition 3.1.*

When G is a d-regular graph, the diagonal degree matrix $D(G)$ is simply $d \cdot I_n$, and so the spectrums of the adjacency matrix and the Laplacian matrix are basically the same. That is, let $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$ be the eigenvalues of the adjacency matrix, and $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$ be the eigenvalues of the Laplacian matrix. When G is d-regular, it holds that $\lambda_i = d - \alpha_i$ for $1 \leq i \leq n$, and so the i-th largest eigenvalue of A corresponds to the i-th smallest eigenvalue of L. We will use this convention throughout that the eigenvalues of A are denoted by α_i and the eigenvalues of L are denoted by λ_i, and also the eigenvalues of A are ordered in non-increasing order while the eigenvalues of L are ordered in non-decreasing order. So, later on, when we say the k-th eigenvalue of a graph, we either mean the k-th largest eigenvalue of the adjacency matrix or the k-th smallest eigenvalue of the Laplacian matrix.

When the graph is not a regular graph, it may not be easy to relate the eigenvalues of the adjacency matrix and the Laplacian matrix. On one hand, as we mentioned in the previous subsection, it is not so easy to give a characterization of α_1 when the graph is not regular. On the other hand, λ_1 is equal to zero for every graph as we will soon see. We define a matrix for the proof which will also be useful later.

Definition 3.15 (Edge Incidence Matrix)**.** *Let $G = (V, E)$ be an undirected graph and with $V(G) = [n]$ and $m = |E|$. For each edge $e = ij \in E$, let b_e be an n-dimensional vector with the i-th position equal to $+1$ and the j-th position equal to -1 and all other positions equal to 0. Let $B(G)$ be the $n \times m$ edge incidence matrix whose columns are $\{b_e \mid e \in E\}$.*

For an edge $e \in E$, let L_e be its Laplacian matrix with $(L_e)_{i,i} = (L_e)_{j,j} = 1$ and $(L_e)_{i,j} = (L_e)_{j,i} = -1$. Note that the Laplacian L_e of an edge e can be written as $b_e b_e^T$, and the Laplacian of the graph G can be written as

$$L(G) = \sum_{e \in E} L_e = \sum_{e \in E} b_e b_e^T = B(G) \cdot B(G)^T.$$

With this definition in place, the proof that zero is the smallest eigenvalue of Laplacian matrix is straightforward.

Lemma 3.16 (Smallest Eigenvalue of Laplacian Matrix)**.** *The Laplacian matrix $L(G)$ of an undirected graph G is positive semidefinite, and its smallest eigenvalue is zero with the all-one vector being a corresponding eigenvector.*

Proof. As L can be written as BB^T as shown in Definition 3.15, it follows that L is a positive semidefinite matrix by Fact 2.7, and so all eigenvalues of L are non-negative. It is easy to check that $L\vec{1} = 0$, and so 0 is the smallest eigenvalue and $\vec{1}$ is a corresponding eigenvector. □

Having a trivial smallest eigenvalue and a simple corresponding eigenvector is one reason that Laplacian matrix is easier to work with. Another reason is that the Laplacian matrix has a nice quadratic form.

Lemma 3.17 (Quadratic Form for Laplacian Matrix)**.** *Let L be the Laplacian matrix of an undirected graph $G = (V, E)$ with $V(G) = [n]$. For any vector $x \in \mathbb{R}^n$,*

$$x^T L x = \sum_{ij \in E} \big(x(i) - x(j)\big)^2.$$

Proof. Using the decomposition of L in Definition 3.15,

$$x^T L x = x^T \Big(\sum_{ij \in E} L_{ij}\Big) x = x^T \Big(\sum_{ij \in E} b_{ij} b_{ij}^T\Big) x = \sum_{ij \in E} x^T b_{ij} b_{ij}^T x = \sum_{ij \in E} \big(x(i) - x(j)\big)^2.$$

□

We will use Lemma 3.16 and Lemma 3.17 to write a nice formulation for the second smallest eigenvalue when we study Cheeger's inequality.

Connectedness

The second smallest eigenvalue of the Laplacian matrix can be used to determine whether the graph is connected.

Proposition 3.18 (Spectral Characterization of Connected Graphs)**.** *Let G be an undirected graph and let $\lambda_1 \leq \ldots \leq \lambda_n$ be the eigenvalues of its Laplacian matrix L. Then G is a connected graph if and only if $\lambda_2 > 0$.*

Proof. Suppose G is disconnected. Then the vertex set can be partitioned into two sets S_1 and S_2 such that there are no edges between them. For a subset $S \subseteq V$, let $\chi_S \in \mathbb{R}^n$ be the characteristic vector of S. Check that both χ_{S_1} and χ_{S_2} are eigenvectors of L with eigenvalue 0. Since χ_{S_1} and χ_{S_2} are linearly independent, it follows that 0 is an eigenvalue with multiplicity at least 2, and thus $\lambda_2 = 0$.

Suppose G is connected. Let x be an eigenvector with eigenvalue 0. Then its quadratic form $x^T Lx = 0$, and so $\sum_{ij \in E}(x(i)-x(j))^2 = 0$ by Lemma 3.17, which implies that $x(i) = x(j)$ for every edge $ij \in E$. Since G is connected, it follows that $x = c \cdot \vec{1}$ for some c, and thus the eigenspace of eigenvalue 0 is of one dimension. Therefore, the eigenvalue 0 has multiplicity one and thus $\lambda_2 > 0$. □

The proof of Proposition 3.18 can be extended to prove the following generalization.

Exercise 3.19 (Spectral Characterization of Number of Components)**.** *Prove that the Laplacian matrix $L(G)$ of an undirected graph G has 0 as its eigenvalue with multiplicity k if and only if the graph G has k connected components.*

3.3 Normalized Adjacency and Laplacian Matrix

Recall that the spectrum of the adjacency matrix satisfies

$$d \geq \alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n \geq -d,$$

where the upper bound and the lower bound depend on the maximum degree d of the graph. So, when we relate the eigenvalues of the adjacency matrix to some combinatorial parameters, there is usually a dependency on the maximum degree of the graph.

To remove this dependency and state the Cheeger's inequality nicely, we will use the following normalized version of the adjacency matrix and the Laplacian matrix.

Definition 3.20 (Normalized Adjacency and Laplacian Matrix)**.** *Let G be an undirected graph with no isolated vertices. The normalized adjacency matrix $\mathcal{A}(G)$ of G is defined as $\mathcal{A}(G) := D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$, where D is the diagonal degree matrix in Definition 3.13 and A is the adjacency matrix in Definition 3.1.*

The normalized Laplacian matrix $\mathcal{L}(G)$ of G is defined as $\mathcal{L}(G) := D^{-\frac{1}{2}}LD^{-\frac{1}{2}}$, where L is the Laplacian matrix in Definition 3.14. Note that $\mathcal{L}(G) = I - \mathcal{A}(G)$.

We will overload notations and still use $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$ to denote the eigenvalues of $\mathcal{A}(G)$ and $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$ to denote the eigenvalues of $\mathcal{L}(G)$. Since $\mathcal{L}(G) = I - \mathcal{A}(G)$ as stated in Definition 3.20, the spectrums of $\mathcal{L}(G)$ and $\mathcal{A}$ are basically equivalent such that $\lambda_i = 1 - \alpha_i$ for $1 \leq i \leq n$. After the normalization, the eigenvalues are bounded as follows.

Lemma 3.21 (Normalized Eigenvalues)**.** *Let G be an undirected graph with no isolated vertices. Let $\alpha_1 \geq \ldots \geq \alpha_n$ be the eigenvalues of its normalized adjacency matrix and $\lambda_1 \leq \ldots \leq \lambda_n$ be the eigenvalues of its normalized Laplacian matrix. Then $1 = \alpha_1 \geq \alpha_n \geq -1$ and $0 = \lambda_1 \leq \lambda_n \leq 2$.*

Proof. First we prove that $\lambda_1 = 0$. Note that 0 is an eigenvalue of $\mathcal{L}$, as

$$\mathcal{L}\big(D^{\frac{1}{2}}\vec{1}\big) = \big(D^{-\frac{1}{2}}LD^{-\frac{1}{2}}\big)\big(D^{\frac{1}{2}}\vec{1}\big) = \big(D^{-\frac{1}{2}}L\vec{1}\big) = 0.$$

Note also that

$$\mathcal{L} = D^{-\frac{1}{2}} L D^{-\frac{1}{2}} = D^{-\frac{1}{2}} B B^T D^{-\frac{1}{2}} = \left(D^{-\frac{1}{2}} B\right)\left(D^{-\frac{1}{2}} B\right)^T$$

where B is the edge incidence matrix in Definition 3.15. It follows that $\mathcal{L} = I - \mathcal{A}$ is a positive semidefinite matrix by Fact 2.7, and thus 0 is the smallest eigenvalue of $\mathcal{L}$ and hence $\lambda_1 = 0$. This implies that $\alpha_1 = 1$ as $\lambda_1 = 1 - \alpha_1$.

Next we prove that $\alpha_n \geq -1$. We will show that $D + A$ is also a positive semidefinite matrix. Then the same argument as in the above paragraph can be used to show that $I + \mathcal{A} = D^{-\frac{1}{2}}(D + A)D^{-\frac{1}{2}}$ is also a positive semidefinite matrix, and this would imply that $1 + \alpha_n \geq 0$ and thus $\alpha_n \geq -1$. There are at least two ways to see that $D + A$ is positive semidefinite. One way is to define $\bar{B}$ to be the "unsigned" matrix of B where $\bar{B}_{ij} = |B_{ij}|$ for all $i, j \in V$, and go through the same argument in Definition 3.15 and check that $D + A = \bar{B}\bar{B}^T$. Another way is to use a similar decomposition as in Definition 3.15 and see that the quadratic form of $D + A$ can be written as

$$x^T(D + A)x = \sum_{ij \in E} \left(x_i + x_j\right)^2,$$

which is a sum of squares and thus non-negative. This implies that $\lambda_n \leq 2$ as $\lambda_n = 1 - \alpha_n$. □

3.4 Robust Generalizations

So far we have used the graph spectrum to deduce some simple combinatorial properties of the graph, such as bipartiteness and connectedness, which are easy to deduce directly by simple combinatorial methods such as breadth first search and depth first search. So one may wonder why these spectral characterizations are useful. The key feature of the spectral characterizations is that they can be generalized quantitatively to prove the following robust generalizations of the basic results:

- λ_2 is close to zero if and only if the graph is close to being disconnected. This is the content of Cheeger's inequality.
- λ_n is close to 2 if and only if the graph has a structure close to a bipartite component. This is an analog of Cheeger's inequality for λ_n.
- λ_k is close to zero if and only if the graph is close to having k connected components. This is a generalization called the higher-order Cheeger's inequality.

We will make these statements precise in the next two chapters.

3.5 Problems

The following are some additional problems that are relevant and interesting.

Problem 3.22 (Cycles). *Compute the Laplacian spectrum of C_n, the cycle with n vertices.*

Problem 3.23 (Hypercubes). *A hypercube of n-dimension is an undirected graph with 2^n vertices. Each vertex corresponds to a string of n bits. Two vertices have an edge if and only if their corresponding strings differ by exactly one bit.*

1. *Given two undirected graphs* $G = (V, E)$ *and* $H = (U, F)$*, we define* $G \times H$ *as the undirected graph with vertex set* $V \times U$ *and two vertices* (v_1, u_1)*,* (v_2, u_2) *have an edge if and only if either (1)* $v_1 = v_2$ *and* $u_1 u_2 \in F$ *or (2)* $u_1 = u_2$ *and* $v_1 v_2 \in E$*. Let* x *be an eigenvector of the Laplacian of* G *with eigenvalue* α*, and let* y *be an eigenvector of the Laplacian of* H *with eigenvalue* β*. Show that we can use* x *and* y *to construct an eigenvector of the Laplacian of* $G \times H$ *with eigenvalue* $\alpha + \beta$*.*

2. *Use (1), or otherwise, to compute the Laplacian spectrum of the hypercube of* n *dimension.*

Problem 3.24 (Number of Spanning Trees)**.** *Let* $G = (V, E)$ *be an undirected graph with* $V = [n]$*.*

1. *Let* B *be the edge incidence matrix of* G *in Definition 3.15. Prove that the determinant of any* $(n-1) \times (n-1)$ *submatrix of* B *is* ± 1 *if and only if the* $n-1$ *edges corresponding to the columns form a spanning tree of* G*.*

2. *Let* L *be the Laplacian matrix of* G *and let* L' *be the matrix obtained from* L *by deleting the last row and last column. Use (1), or otherwise, to prove that* $\det(L')$ *is equal to the number of spanning trees in* G*. You can use the Cauchy-Binet formula in Fact 2.27 to solve this problem.*

Problem 3.25 (Wilf's Theorem)**.** *Let* G *be an undirected graph and* α_1 *be the largest eigenvalue of its adjacency matrix. Prove that* $\chi(G) \le \lfloor \alpha_1 \rfloor + 1$*, where* $\chi(G)$ *is the chromatic number of* G*. You may find the Cauchy interlacing Theorem 2.13 useful.*

3.6 References

[Spi19] Daniel A. Spielman. *Spectral and Algebraic Graph Theory*. 2019. 17