Spectral Graph Theory Lecture 3

Fundamental Graphs

Daniel A. Spielman September 5, 2018

3.1 Overview

We will bound and derive the eigenvalues of the Laplacian matrices of some fundamental graphs, including complete graphs, star graphs, ring graphs, path graphs, and products of these that yield grids and hypercubes. As all these graphs are connected, they all have eigenvalue zero with multiplicity one. We will have to do some work to compute the other eigenvalues.

We derive some meaning from the eigenvalues by using them to bound isoperimetric numbers of graphs, which I recall are defined by

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

We bound this using the following theorem from last lecture.

Theorem 3.1.1. *For every $S \subset V$*

$$\theta(S) \geq \lambda_2(1-s),$$

where $s = |S| / |V|$. In particular,

$$\theta_G \geq \lambda_2/2.$$

3.2 The Laplacian Matrix

We beging this lecture by establishing the equivalence of multiple expressions for the Laplacian. These will be necessary to derive its eigenvalues.

The Laplacian Matrix of a weighted graph $G = (V, E, w)$, $w : E \to \mathbb{R}^+$, is designed to capture the Laplacian quadratic form:

$$\boldsymbol{x}^T \boldsymbol{L}_G \boldsymbol{x} = \sum_{(a,b)\in E} w_{a,b}(\boldsymbol{x}(a) - \boldsymbol{x}(b))^2. \tag{3.1}$$

We will now use this quadratic form to derive the structure of the matrix. To begin, consider a graph with just two vertices and one edge. Let's call it $G_{1,2}$. We have

$$\boldsymbol{x}^T \boldsymbol{L}_{G_{1,2}} \boldsymbol{x} = (\boldsymbol{x}(1) - \boldsymbol{x}(2))^2. \tag{3.2}$$

Consider the vector $\boldsymbol{\delta}_1 - \boldsymbol{\delta}_2$, where by $\boldsymbol{\delta}_i$ I mean the elementary unit vector with a 1 in coordinate i. We have

$$\boldsymbol{x}(1) - \boldsymbol{x}(2) = \boldsymbol{\delta}_1^T \boldsymbol{x} - \boldsymbol{\delta}_2^T \boldsymbol{x} = (\boldsymbol{\delta}_1 - \boldsymbol{\delta}_2)^T \boldsymbol{x},$$

so

$$(\boldsymbol{x}(1) - \boldsymbol{x}(2))^2 = \left((\boldsymbol{\delta}_1 - \boldsymbol{\delta}_2)^T \boldsymbol{x}\right)^2 = \boldsymbol{x}^T (\boldsymbol{\delta}_1 - \boldsymbol{\delta}_2)(\boldsymbol{\delta}_1 - \boldsymbol{\delta}_2)^T \boldsymbol{x} = \boldsymbol{x}^T \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \boldsymbol{x}.$$

Thus,

$$\boldsymbol{L}_{G_{1,2}} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Now, let $G_{a,b}$ be the graph with just one edge between u and v. It can have as many other vertices as you like. The Laplacian of $G_{a,b}$ can be written in the same way: $\boldsymbol{L}_{G_{a,b}} = (\boldsymbol{\delta}_a - \boldsymbol{\delta}_b)(\boldsymbol{\delta}_a - \boldsymbol{\delta}_b)^T$. This is the matrix that is zero except at the intersection of rows and columns indexed by u and v, where it looks looks like

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Summing the matrices for every edge, we obtain

$$\boldsymbol{L}_G = \sum_{(a,b)\in E} w_{a,b}(\boldsymbol{\delta}_a - \boldsymbol{\delta}_b)(\boldsymbol{\delta}_a - \boldsymbol{\delta}_b)^T = \sum_{(a,b)\in E} w_{a,b} \boldsymbol{L}_{G_{a,b}}.$$

You should check that this agrees with the definition of the Laplacian from the first class:

$$\boldsymbol{L}_G = \boldsymbol{D}_G - \boldsymbol{A}_G,$$

where

$$\boldsymbol{D}_G(a,a) = \sum_b w_{a,b}.$$

This formula turns out to be useful when we view the Laplacian as an operator. For every vector $\boldsymbol{x}$ we have

$$(\boldsymbol{L}_G \boldsymbol{x})(a) = d(a)\boldsymbol{x}(a) - \sum_{(a,b)\in E} w_{a,b}\boldsymbol{x}(b) = \sum_{(a,b)\in E} w_{a,b}(\boldsymbol{x}(a) - \boldsymbol{x}(b)). \tag{3.3}$$

3.3 The complete graph

The complete graph on n vertices, K_n, has edge set $\{(a,b) : a \neq b\}$.

Lemma 3.3.1. *The Laplacian of K_n has eigenvalue 0 with multiplicity 1 and n with multiplicity $n-1$.*

Proof. To compute the non-zero eigenvalues, let $\boldsymbol{\psi}$ be any non-zero vector orthogonal to the all-1s vector, so

$$\sum_a \boldsymbol{\psi}(a) = 0. \tag{3.4}$$

We now compute the first coordinate of $\boldsymbol{L}_{K_n}\boldsymbol{\psi}$. Using (3.3), we find

$$(\boldsymbol{L}_{K_n}\boldsymbol{\psi})(1) = \sum_{v \geq 2} (\boldsymbol{\psi}(1) - \boldsymbol{\psi}(b)) = (n-1)\boldsymbol{\psi}(1) - \sum_{v=2}^{n} \boldsymbol{\psi}(b) = n\boldsymbol{\psi}(1), \quad \text{by (3.4).}$$

As the choice of coordinate was arbitrary, we have $\boldsymbol{L}\boldsymbol{\psi} = n\boldsymbol{\psi}$. So, every vector orthogonal to the all-1s vector is an eigenvector of eigenvalue n. □

Alternative approach. Observe that $\boldsymbol{L}_{K_n} = n\boldsymbol{I} - \mathbf{1}\mathbf{1}^T$. □

We often think of the Laplacian of the complete graph as being a scaling of the identity. For every $\boldsymbol{x}$ orthogonal to the all-1s vector, $\boldsymbol{L}\boldsymbol{x} = n\boldsymbol{x}$.

Now, let's see how our bound on the isoperimetric number works out. Let $S \subset [n]$. Every vertex in S has $n - |S|$ edges connecting it to vertices not in S. So,

$$\theta(S) = \frac{|S|\,(n - |S|}{|S|} = n - |S| = \lambda_2(\boldsymbol{L}_{K_n})(1 - s),$$

where $s = |S|/n$. Thus, Theorem 3.1.1 is sharp for the complete graph.

3.4 The star graphs

The star graph on n vertices S_n has edge set $\{(1, a) : 2 \leq a \leq n\}$.

To determine the eigenvalues of S_n, we first observe that each vertex $a \geq 2$ has degree 1, and that each of these degree-one vertices has the same neighbor. Whenever two degree-one vertices share the same neighbor, they provide an eigenvector of eigenvalue 1.

Lemma 3.4.1. *Let $G = (V, E)$ be a graph, and let a and b be vertices of degree one that are both connected to another vertex c. Then, the vector $\boldsymbol{\psi} = \boldsymbol{\delta}_a - \boldsymbol{\delta}_b$ is an eigenvector of $\boldsymbol{L}_G$ of eigenvalue 1.*

Proof. Just multiply $\boldsymbol{L}_G$ by $\boldsymbol{\psi}$, and check (using (3.3)) vertex-by-vertex that it equals $\boldsymbol{\psi}$. □

As eigenvectors of different eigenvalues are orthogonal, this implies that $\boldsymbol{\psi}(a) = \boldsymbol{\psi}(b)$ for every eigenvector with eigenvalue different from 1.

Lemma 3.4.2. *The graph S_n has eigenvalue* 0 *with multiplicity* 1*, eigenvalue* 1 *with multiplicity* $n - 2$*, and eigenvalue* n *with multiplicity* 1*.*

Proof. Applying Lemma 3.4.1 to vertices i and $i+1$ for $2 \leq i < n$, we find $n-2$ linearly independent eigenvectors of the form $\boldsymbol{\delta}_i - \boldsymbol{\delta}_{i+1}$, all with eigenvalue 1. As 0 is also an eigenvalue, only one eigenvalue remains to be determined.

Recall that the trace of a matrix equals both the sum of its diagonal entries and the sum of its eigenvalues. We know that the trace of $\boldsymbol{L}_{S_n}$ is $2n-2$, and we have identified $n-1$ eigenvalues that sum to $n-2$. So, the remaining eigenvalue must be n.

To determine the corresponding eigenvector, recall that it must be orthogonal to the other eigenvectors we have identified. This tells us that it must have the same value at each of the points of the star. Let this value be 1, and let x be the value at vertex 1. As the eigenvector is orthogonal to the constant vectors, it must be that

$$(n-1)+x=0,$$

so $x=-(n-1)$. $\square$

3.5 Products of graphs

We now define a product on graphs. If we apply this product to two paths, we obtain a grid. If we apply it repeatedly to one edge, we obtain a hypercube.

Definition 3.5.1. *Let $G=(V,E)$ and $H=(W,F)$ be graphs. Then $G\times H$ is the graph with vertex set $V\times W$ and edge set*

$$\begin{aligned}&\Big((a,b),(\hat{a},b)\Big) \text{ where } (a,\hat{a})\in E \text{ and}\\ &\Big((a,b),(a,\hat{b})\Big) \text{ where } (b,\hat{b})\in F.\end{aligned}$$

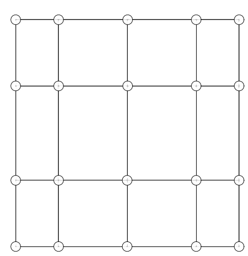

Figure 3.1: An m-by-n grid graph is the product of a path on m vertices with a path on n vertices. This is a drawing of a 5-by-4 grid made using Hall's algorithm.

Theorem 3.5.2. *Let $G=(V,E)$ and $H=(W,F)$ be graphs with Laplacian eigenvalues $\lambda_1,\ldots,\lambda_n$ and $\mu_1,\ldots,\mu_m$, and eigenvectors $\boldsymbol{\alpha}_1,\ldots,\boldsymbol{\alpha}_n$ and $\boldsymbol{\beta}_1,\ldots,\boldsymbol{\beta}_m$, respectively. Then, for each $1\le i\le n$ and $1\le j\le m$, $G\times H$ has an eigenvector $\boldsymbol{\gamma}_{i,j}$ of eigenvalue $\lambda_i+\mu_j$ such that*

$$\boldsymbol{\gamma}_{i,j}(a,b)=\boldsymbol{\alpha}_i(a)\boldsymbol{\beta}_j(b).$$

Proof. Let $\boldsymbol{\alpha}$ be an eigenvector of L_G of eigenvalue λ, let $\boldsymbol{\beta}$ be an eigenvector of L_H of eigenvalue μ, and let $\boldsymbol{\gamma}$ be defined as above.

To see that $\boldsymbol{\gamma}$ is an eigenvector of eigenvalue $\lambda + \mu$, we compute

$$
\begin{aligned}
(L\boldsymbol{\gamma})(a,b) &= \sum_{(a,\hat{a})\in E} (\boldsymbol{\gamma}(a,b) - \boldsymbol{\gamma}(\hat{a},b)) + \sum_{(b,\hat{b})\in F} \left(\boldsymbol{\gamma}(a,b) - \boldsymbol{\gamma}(a,\hat{b})\right) \\
&= \sum_{(a,\hat{a})\in E} (\boldsymbol{\alpha}(a)\boldsymbol{\beta}(b) - \boldsymbol{\alpha}(\hat{a})\boldsymbol{\beta}(b)) + \sum_{(b,\hat{b})\in F} \left(\boldsymbol{\alpha}(a)\boldsymbol{\beta}(b) - \boldsymbol{\alpha}(a)\boldsymbol{\beta}(\hat{b})\right) \\
&= \sum_{(a,\hat{a})\in E} \boldsymbol{\beta}(b)\left(\boldsymbol{\alpha}(a) - \boldsymbol{\alpha}(\hat{a})\right) + \sum_{(b,\hat{b})\in F} \boldsymbol{\alpha}(a)\left(\boldsymbol{\beta}(b) - \boldsymbol{\beta}(\hat{b})\right) \\
&= \sum_{(a,\hat{a})\in E} \boldsymbol{\beta}(b)\lambda\boldsymbol{\alpha}(a) + \sum_{(b,\hat{b})\in F} \boldsymbol{\alpha}(a)\mu\boldsymbol{\beta}(b) \\
&= (\lambda + \mu)(\boldsymbol{\alpha}(a)\boldsymbol{\beta}(b)).
\end{aligned}
$$

□

3.5.1 The Hypercube

The d-dimensional hypercube graph, H_d, is the graph with vertex set $\{0,1\}^d$, with edges between vertices whose names differ in exactly one bit. The hypercube may also be expressed as the product of the one-edge graph with itself $d-1$ times, with the proper definition of graph product.

Let H_1 be the graph with vertex set $\{0,1\}$ and one edge between those vertices. It's Laplacian matrix has eigenvalues 0 and 2. As $H_d = H_{d-1} \times H_1$, we may use this to calculate the eigenvalues and eigenvectors of H_d for every d.

Using Theorem 3.1.1 and the fact that $\lambda_2(H_d) = 2$, we can immediately prove the following isoperimetric theorem for the hypercube.

Corollary 3.5.3.

$$\theta_{H_d} \geq 1.$$

In particular, for every set of at most half the vertices of the hypercube, the number of edges on the boundary of that set is at least the number of vertices in that set.

This result is tight, as you can see by considering one face of the hypercube, such as all the vertices whose labels begin with 0. It is possible to prove this by more concrete combinatorial means. In fact, very precise analyses of the isoperimetry of sets of vertices in the hypercube can be obtained. See [Har76] or [Bol86].

3.6 Bounds on λ_2 by test vectors

We can reverse our thinking and use Theorem 3.1.1 to prove an upper bound on λ_2. If you recall the proof of that theorem, you will see a special case of proving an upper bound by a test vector.

By Theorem 2.1.3 we know that every vector $\boldsymbol{v}$ orthogonal to $\mathbf{1}$ provides an upper bound on λ_2:

$$\lambda_2 \leq \frac{\boldsymbol{v}^T L \boldsymbol{v}}{\boldsymbol{v}^T \boldsymbol{v}}.$$

When we use a vector $\boldsymbol{v}$ in this way, we call it a *test vector.*

Let's see what a test vector can tell us about λ_2 of a path graph on n vertices. I would like to use the vector that assigns i to vertex a as a test vector, but it is not orthogonal to $\mathbf{1}$. So, we will use the next best thing. Let $\boldsymbol{x}$ be the vector such that $\boldsymbol{x}(a) = (n+1) - 2a$, for $1 \leq a \leq n$. This vector satisfies $\boldsymbol{x} \perp \mathbf{1}$, so

$$\begin{aligned}
\lambda_2(P_n) &\leq \frac{\sum_{1 \leq a < n} (x(a) - x(a+1))^2}{\sum_a x(a)^2} \\
&= \frac{\sum_{1 \leq a < n} 2^2}{\sum_a (n+1-2a)^2} \\
&= \frac{4(n-1)}{(n+1)n(n-1)/3} \qquad \text{(clearly, the denominator is } n^3/c \text{ for some } c) \\
&= \frac{12}{n(n+1)}. \qquad (3.5)
\end{aligned}$$

We will soon see that this bound is of the right order of magnitude. Thus, Theorem 3.1.1 does not provide a good bound on the isoperimetric number of the path graph. The isoperimetric number is minimized by the set $S = \{1, \ldots, n/2\}$, which has $\theta(S) = 2/n$. However, the upper bound provided by Theorem 3.1.1 is of the form c/n^2. Cheeger's inequality, which we will prove later in the semester, will tell us that the error of this approximation can not be worse than quadratic.

The Courant-Fischer theorem is not as helpful when we want to prove lower bounds on λ_2. To prove lower bounds, we need the form with a maximum on the outside, which gives

$$\lambda_2 \geq \max_{S : \dim(S) = n-1} \min_{v \in S} \frac{\boldsymbol{v}^T \boldsymbol{L} \boldsymbol{v}}{\boldsymbol{v}^T \boldsymbol{v}}.$$

This is not too helpful, as it is difficult to prove lower bounds on

$$\min_{v \in S} \frac{\boldsymbol{v}^T \boldsymbol{L} \boldsymbol{v}}{\boldsymbol{v}^T \boldsymbol{v}}$$

over a space S of large dimension. We will see a technique that lets us prove such lower bounds next lecture.

But, first we compute the eigenvalues and eigenvectors of the path graph exactly.

3.7 The Ring Graph

The ring graph on n vertices, R_n, may be viewed as having a vertex set corresponding to the integers modulo n. In this case, we view the vertices as the numbers 0 through $n-1$, with edges $(a, a+1)$, computed modulo n.

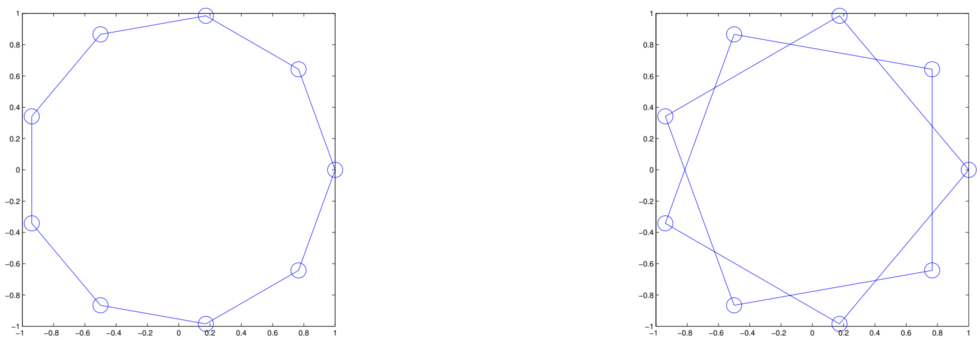

(a) The ring graph on 9 vertices. (b) The eigenvectors for $k = 2$.

Figure 3.2:

Lemma 3.7.1. *The Laplacian of R_n has eigenvectors*

$$\begin{aligned} \boldsymbol{x}_k(a) &= \cos(2\pi k a/n), \text{ and} \\ \boldsymbol{y}_k(a) &= \sin(2\pi k a/n), \end{aligned}$$

for $0 \leq k \leq n/2$, ignoring $\boldsymbol{y}_0$ which is the all-zero vector, and for even n ignoring $\boldsymbol{y}_{n/2}$ for the same reason. Eigenvectors $\boldsymbol{x}_k$ and $\boldsymbol{y}_k$ have eigenvalue $2 - 2\cos(2\pi k/n)$.

Note that $\boldsymbol{x}_0$ is the all-ones vector. When n is even, we only have $\boldsymbol{x}_{n/2}$, which alternates ± 1.

Proof. We will first see that $\boldsymbol{x}_1$ and $\boldsymbol{y}_1$ are eigenvectors by drawing the ring graph on the unit circle in the natural way: plot vertex u at point $(\cos(2\pi a/n), \sin(2\pi a/n))$.

You can see that the average of the neighbors of a vertex is a vector pointing in the same direction as the vector associated with that vertex. This should make it obvious that both the x and y coordinates in this figure are eigenvectors of the same eigenvalue. The same holds for all k.

Alternatively, we can verify that these are eigenvectors by a simple computation.

$$\begin{aligned} \left(L_{R_n} \boldsymbol{x}_k\right)(a) &= 2\boldsymbol{x}_k(a) - \boldsymbol{x}_k(a+1) - \boldsymbol{x}_k(a-1) \\ &= 2\cos(2\pi k a/n) - \cos(2\pi k(a+1)/n) - \cos(2\pi k(a-1)/n) \\ &= 2\cos(2\pi k a/n) - \cos(2\pi k a/n)\cos(2\pi k/n) + \sin(2\pi k a/n)\sin(2\pi k/n) \\ &\qquad - \cos(2\pi k a/n)\cos(2\pi k/n) - \sin(2\pi k a/n)\sin(2\pi k/n) \\ &= 2\cos(2\pi k a/n) - \cos(2\pi k a/n)\cos(2\pi k/n) - \cos(2\pi k a/n)\cos(2\pi k/n) \\ &= (2 - 2\cos(2\pi k/n))\cos(2\pi k a/n) \\ &= (2 - \cos(2\pi k/n))\boldsymbol{x}_k(a). \end{aligned}$$

The computation for $\boldsymbol{y}_k$ follows similarly. $\square$

3.8 The Path Graph

We will derive the eigenvalues and eigenvectors of the path graph from those of the ring graph. To begin, I will number the vertices of the ring a little differently, as in Figure 3.3.

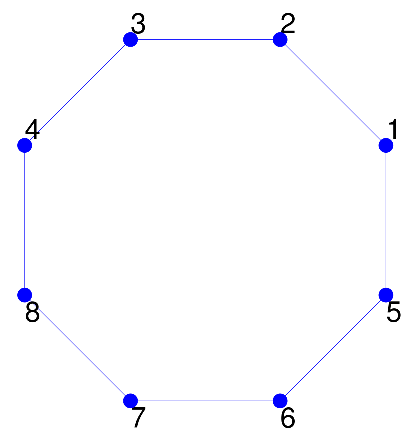

Figure 3.3: The ring on 8 vertices, numbered differently

Lemma 3.8.1. *Let $P_n = (V, E)$ where $V = \{1, \dots, n\}$ and $E = \{(a, a+1) : 1 \leq a < n\}$. The Laplacian of P_n has the same eigenvalues as R_{2n}, excluding 2. That is, P_n has eigenvalues namely $2(1 - \cos(\pi k/n))$, and eigenvectors*

$$\boldsymbol{v}_k(a) = \cos(\pi k u/n - \pi k/2n).$$

for $0 \leq k < n$

Proof. We derive the eigenvectors and eigenvalues by treating P_n as a quotient of R_{2n}: we will identify vertex u of P_n with vertices u and $u+n$ of R_{2n} (under the new numbering of R_{2n}). These are pairs of vertices that are above each other in the figure that I drew.

Let $\boldsymbol{I}_n$ be the n-dimensional identity matrix. You should check that

$$\begin{pmatrix} \boldsymbol{I}_n & \boldsymbol{I}_n \end{pmatrix} \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} = 2\boldsymbol{L}_{P_n}.$$

If there is an eigenvector $\boldsymbol{\psi}$ of R_{2n} with eigenvalue λ for which $\boldsymbol{\psi}(a) = \boldsymbol{\psi}(a+n)$ for $1 \leq a \leq n$, then the above equation gives us a way to turn this into an eigenvector of P_n: Let $\boldsymbol{\phi} \in \mathbb{R}^n$ be the vector for which

$$\boldsymbol{\phi}(a) = \boldsymbol{\psi}(a), \text{for } 1 \leq a \leq n.$$

Then,

$$\begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\phi} = \boldsymbol{\psi}, \quad \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\phi} = \lambda \boldsymbol{\psi}, \quad \text{and} \quad \begin{pmatrix} \boldsymbol{I}_n & \boldsymbol{I}_n \end{pmatrix} \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\psi} = 2\lambda \boldsymbol{\phi}.$$

So, if we can find such a vector $\boldsymbol{\psi}$, then the corresponding $\boldsymbol{\phi}$ is an eigenvector of P_n of eigenvalue λ.

As you've probably guessed, we can find such vectors $\boldsymbol{\psi}$. I've drawn one in Figure 3.3. For each of the two-dimensional eigenspaces of R_{2n}, we get one such a vector. These provide eigenvectors of eigenvalue

$$2(1 - \cos(\pi k/n)),$$

for $1 \leq k < n$. Thus, we now know $n - 1$ distinct eigenvalues. The last, of course, is zero. $\square$

The type of quotient used in the above argument is known as an *equitable partition*. You can find a extensive exposition of these in Godsil's book [God93].

References

[Bol86] Béla Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors, and combinatorial probability*. Cambridge University Press, 1986.

[God93] Chris Godsil. *Algebraic Combinatorics*. Chapman & Hall, 1993.

[Har76] Sergiu Hart. A note on the edges of the n-cube. *Discrete Mathematics*, 14(2):157–163, 1976.