Spectral Graph Theory Lecture 6

High-Frequency Eigenvalues

Daniel A. Spielman September 17, 2018

6.1 Overview

In this lecture we will see how high-frequency eigenvalues of the Laplacian matrix can be related to independent sets and graph coloring. Some of the bounds we obtained will be more easilys stated in terms of the adjacency matrix, $\boldsymbol{M}$. Recall the we number the Laplacian matrix eigenvalues in increasing order:

$$0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n.$$

We call the adjacency matrix eigenvalues $\mu_1, \ldots, \mu_n$, and number them in the reverse order:

$$\mu_1 \geq \cdots \geq \mu_n.$$

The reason is that for d-regular graphs, $\mu_i = d - \lambda_i$.

6.2 Graph Coloring and Independent Sets

A coloring of a graph is an assignment of one color to every vertex in a graph so that each edge connects vertices of different colors. We are interested in coloring graphs while using as few colors as possible. Formally, a k-coloring of a graph is a function $c : V \rightarrow \{1, \ldots, k\}$ so that for all $(u, v) \in V$, $c(u) \neq c(v)$. A graph is k-colorable if it has a k-coloring. The chromatic number of a graph, written χ_G, is the least k for which G is k-colorable. A graph G is 2-colorable if and only if it is bipartite. Determining whether or not a graph is 3-colorable is an NP-complete problem. The famous 4-Color Theorem [AH77a, AH77b] says that every planar graph is 4-colorable.

A set of vertices S is *independent* if there are no edges between vertices in S. In particular, each color class in a coloring is an independent set. The problem of finding large independent sets in a graph is NP-Complete, and it is very difficult to even approximate the size of the largest independent set in a graph.

However, for some carefully chosen graphs one can obtain very good bounds on the sizes of independent sets by using spectral graph theory. We may later see some uses of this theory in the analysis of error-correcting codes and sphere packings.

6.3 Hoffman's Bound

One of the first results in spectral graph theory was Hoffman's proof the following upper bound on the size of an independent set in a graph G.

Theorem 6.3.1. *Let $G = (V, E)$ be a d-regular graph, and let μ_n be its smallest adjacency matrix eigenvalue. Then*

$$\alpha(G) \leq n \frac{-\mu_n}{d - \mu_n}.$$

Recall that $\mu_n < 0$. Otherwise this theorem would not make sense. We will prove a generalization of Hoffman's theorem due to Godsil and Newman [GN08]:

Theorem 6.3.2. *Let S be an independent set in G, and let $d_{ave}(S)$ be the average degree of a vertex in S. Then,*

$$|S| \leq n \left(1 - \frac{d_{ave}(S)}{\lambda_n}\right).$$

This is a generalization because in the d-regular case $d_{ave} = d$ and $\lambda_n = d - \mu_n$. So, these bounds are the same for regular graphs:

$$1 - \frac{d_{ave}(S)}{\lambda_n} = \frac{\lambda_n - d}{\lambda_n} = \frac{-\mu_n}{d - \mu_n}.$$

Proof. Let S be an independent set of vertices and let $d(S)$ be the sum of the degrees of vertices in S.

Recall that

$$\lambda_n = \max_{\boldsymbol{x}} \frac{\boldsymbol{x}^T \boldsymbol{L} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}}.$$

We also know that the vector $\boldsymbol{x}$ that maximizes this quantity is $\boldsymbol{\psi}_n$, and that $\boldsymbol{\psi}_n$ is orthogonal to $\boldsymbol{\psi}_1$. So, we can refine this expression to

$$\lambda_n = \max_{\boldsymbol{x} \perp \boldsymbol{1}} \frac{\boldsymbol{x}^T \boldsymbol{L} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}}.$$

Consider the vector

$$\boldsymbol{x} = \boldsymbol{\chi}_S - s\boldsymbol{1},$$

where $s = |S|/n$. As S is independent, we have

$$\boldsymbol{x}^T \boldsymbol{L} \boldsymbol{x} = |\partial(S)| = d(S) = d_{ave}(S)\,|S|\,.$$

Claim 2.3.2 tells us that the square of the norm of $\boldsymbol{x}$ is

$$\boldsymbol{x}^T \boldsymbol{x} = n(s - s^2).$$

So,

$$\lambda_n \geq \frac{d_{ave}(S)\,|S|}{n(s-s^2)} = \frac{d_{ave}(S)sn}{n(s-s^2)} = \frac{d_{ave}(S)}{1-s}.$$

Re-arranging terms, this gives

$$1 - \frac{d_{ave}(S)}{\lambda_n} \geq s,$$

which is equivalent to the claim of the theorem. □

6.4 Application to Paley graphs

Let's examine what Hoffman's bound on the size of the largest independent set tells us about Paley graphs.

If G is a Paley graph and S is an independent set, we have $n = p$, $d = (p-1)/2$, and $\lambda_n = (p+\sqrt{p})/2$, so Hoffman's bound tells us that

$$\begin{aligned} |S| &\leq n\left(1 - \frac{d_{ave}(S)}{\lambda_n}\right) \\ &= p\left(1 - \frac{p-1}{p+\sqrt{p}}\right) \\ &= p\left(\frac{\sqrt{p}+1}{p+\sqrt{p}}\right) \\ &= \sqrt{p}. \end{aligned}$$

One can also show that every clique in a Paley graph has size at most $\sqrt{p}$.

A graph is called a k-Ramsey graph if it contains no clique or independent set of size k. It is a challenge to find large k-Ramsey graphs. Equivalently, it is challenging to find k-Ramsey graphs on n vertices for which k is small. In one of the first papers on the Probabilistic Method in Combinatorics, Erdös proved that a random graph on n vertices in which each edge is included with probability $1/2$ is probably $2\log_2 n$ Ramsey [Erd47].

However, constructing explicit Ramsey graphs has proved much more challening. Until recently, Paley graphs were among the best known. A recent construction of Barak, Rao, Shatltiel and Wigderson [BRSW12] constructs explicit graphs that are $2^{(\log n)^{o(1)}}$ Ramsey.

6.5 Lower Bound on the chromatic number

As a k-colorable graph must have an independent set of size at least n/k, an upper bound on the sizes of independent sets gives a lower bound on its chromatic number. However, this bound is not always a good one.

For example, consider a graph on $2n$ vertices consisting of a clique on n vertices and n vertices of degree 1, each of which is connected to a different vertex in the clique. The chromatic number of this graph is n, because each of the vertices in the clique must have a different color. However, the graph also has an independent set of size n, which would only give a lower bound of 2 on the chromatic number.

Hoffman proved the following lower bound on the chromatic number of a graph that does not require the graph to be regular. Numerically, it is obtained by dividing n by the bound in Theorem 6.3.1. But, the proof is very different because that theorem only applies to regular graphs.

Theorem 6.5.1.

$$\chi(G) \geq \frac{\mu_1 - \mu_n}{-\mu_n} = 1 + \frac{\mu_1}{-\mu_n}.$$

The proof of this theorem relies on one inequality that I will not have time to cover in class. So, I will put its proof in Section B.

Lemma 6.5.2. *Let*

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} & \cdots & \boldsymbol{A}_{1,k} \\ \boldsymbol{A}_{1,2}^T & \boldsymbol{A}_{2,2} & \cdots & \boldsymbol{A}_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{A}_{1,k}^T & \boldsymbol{A}_{2,k}^T & \cdots & \boldsymbol{A}_{k,k} \end{bmatrix}$$

be a block-partitioned symmetric matrix with $k \geq 2$. Then

$$(k-1)\lambda_{min}(\boldsymbol{A}) + \lambda_{max}(\boldsymbol{A}) \leq \sum_i \lambda_{max}(\boldsymbol{A}_{i,i}).$$

Proof of Theorem 6.5.1. Let G be a k-colorable graph. After possibly re-ordering the vertices, the adjacency matrix of G can be written

$$\begin{bmatrix} \boldsymbol{0} & \boldsymbol{A}_{1,2} & \cdots & \boldsymbol{A}_{1,k} \\ \boldsymbol{A}_{1,2}^T & \boldsymbol{0} & \cdots & \boldsymbol{A}_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{A}_{1,k}^T & \boldsymbol{A}_{2,k}^T & \cdots & \boldsymbol{0} \end{bmatrix}.$$

Each block corresponds to a color.

As each diagonal block is all-zero, Lemma 6.5.2 implies

$$(k-1)\lambda_{min}(\boldsymbol{A}) + \lambda_{max}(\boldsymbol{A}) \leq 0.$$

Recalling that $\lambda_{min}(\boldsymbol{A}) = \mu_n < 0$, and $\lambda_{max}(\boldsymbol{A}) = \mu_1$, a little algebra yields

$$1 + \frac{\mu_1}{-\mu_n} \leq k.$$

□

To return to our example of the n clique with n degree-1 vertices attached, I examined an example with $n = 6$. We find $\mu_1 = 5.19$ and $\mu_{12} = -1.62$. This gives a lower bound on the chromatic number of 4.2, which implies a lower bound of 5. We can improve the lower bound by re-weighting the edges of the graph. For example, if we give weight 2 to all the edges in the clique and weight 1 to all the others, we obtain a bound of 5.18, which agrees with the chromatic number of this graph which is 6.

6.6 Coloring and The Adjacency Matrix

I would also like to show how to use spectral graph theory to color a graph. I do not know how to do this using the Laplacian matrix, so we will work with the Adjacency matrix. This will provide me with a good opportunity to cover some material about adjacency matrices that I have neglected.

6.7 The Largest Eigenvalue, μ_1

We now examine μ_1 for graphs which are not necessarily regular. Let G be a graph, let d_{max} be the maximum degree of a vertex in G, and let d_{ave} be the average degree of a vertex in G.

Lemma 6.7.1.

$$d_{ave} \leq \mu_1 \leq d_{max}.$$

Proof. The lower bound follows by considering the Rayleigh quotient with the all-1s vector:

$$\mu_1 = \max_{\boldsymbol{x}} \frac{\boldsymbol{x}^T \boldsymbol{M} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}} \geq \frac{\boldsymbol{1}^T \boldsymbol{M} \boldsymbol{1}}{\boldsymbol{1}^T \boldsymbol{1}} = \frac{\sum_{i,j} \boldsymbol{M}(i,j)}{n} = \frac{\sum_i \boldsymbol{d}(i)}{n}.$$

To prove the upper bound, Let $\boldsymbol{\psi}_1$ be an eigenvector of eigenvalue μ_1. Let v be the vertex on which it takes its maximum value, so $\boldsymbol{\psi}_1(v) \geq \boldsymbol{\psi}_1(u)$ for all u, and assume without loss of generality that $\boldsymbol{\psi}_1(v) \neq 0$. We have

$$\mu_1 = \frac{(\boldsymbol{M}\boldsymbol{\psi}_1)(v)}{\boldsymbol{\psi}_1(v)} = \frac{\sum_{u \sim v} \boldsymbol{\psi}_1(u)}{\boldsymbol{\psi}_1(v)} = \sum_{u \sim v} \frac{\boldsymbol{\psi}_1(u)}{\boldsymbol{\psi}_1(v)} \leq \sum_{u \sim v} 1 \leq d(v) \leq d_{max}. \tag{6.1}$$

□

Lemma 6.7.2. *If G is connected and $\mu_1 = d_{max}$, then G is d_{max}-regular.*

Proof. If we have equality in (6.1), then it must be the case that $d(v) = d_{max}$ and $\boldsymbol{\phi}_1(u) = \boldsymbol{\phi}_1(v)$ for all $(u, v) \in E$. Thus, we may apply the same argument to every neighbor of v. As the graph is connected, we may keep applying this argument to neighbors of vertices to which it has already been applied to show that $\boldsymbol{\phi}_1(z) = \boldsymbol{\phi}_1(v)$ and $d(z) = d_{max}$ for all $z \in V$. □

6.8 Wilf's Theorem

While we may think of μ_1 as being a related to the average degree, it does behave differently. In particular, if we remove the vertex of smallest degree from a graph, the average degree can increase. On the other hand, μ_1 can only decrease when we remove a vertex. Let's prove that now.

Lemma 6.8.1. *Let A be a symmetric matrix with largest eigenvalue α_1. Let B be the matrix obtained by removing the last row and column from A, and let β_1 be the largest eigenvalue of B. Then,*

$$\alpha_1 \geq \beta_1.$$

Proof. For any vector $\boldsymbol{y} \in \mathbb{R}^{n-1}$, we have

$$\boldsymbol{y}^T B \boldsymbol{y} = \begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}^T A \begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}.$$

So, for $\boldsymbol{y}$ an eigenvector of B of eigenvalue β_1,

$$\beta_1 = \frac{\boldsymbol{y}^T B \boldsymbol{y}}{\boldsymbol{y}^T \boldsymbol{y}} = \frac{\begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}^T A \begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}}{\begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}^T \begin{pmatrix} \boldsymbol{y} \\ 0 \end{pmatrix}} \leq \max_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\boldsymbol{x}^T A \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}}.$$

□

Of course, this holds regardless of which row and column we remove, as long as they are the same row and column.

It is easy to show that every graph is $(d_{max} + 1)$-colorable. Assign colors to the vertices one-by-one. As each vertex has at most d_{max} neighbors, there is always some color one can assign that vertex that is different that those assigned to its neighbors. The following theorem of Wilf improves upon this bound.

Theorem 6.8.2.

$$\chi(G) \leq \lfloor \mu_1 \rfloor + 1.$$

Proof. We prove this by induction on the number of vertices in the graph. To ground the induction, consider the graph with one vertex and no edges. It has chromatic number 1 and largest eigenvalue zero[1]. Now, assume the theorem is true for all graphs on $n-1$ vertices, and let G be a graph on n vertices. By Lemma 6.7.1, G has a vertex of degree at most $\lfloor \mu_1 \rfloor$. Let v be such a vertex and let $G - \{v\}$ be the graph obtained by removing this vertex. By Lemma 6.8.1 and our induction hypothesis, $G - \{v\}$ has a coloring with at most $\lfloor \mu_1 \rfloor + 1$ colors. Let c be any such coloring. We just need to show that we can extend c to v. As v has at most $\lfloor \mu_1 \rfloor$ neighbors, there is some color in $\{1, \ldots, \lfloor \mu_1 \rfloor + 1\}$ that does not appear among its neighbors, and which it may be assigned. Thus, G has a coloring with $\lfloor \mu_1 \rfloor + 1$ colors. □

For an example, consider a path graph with at least 3 vertices. We have $d_{max} = 2$, but $\alpha_1 < 2$.

6.9 Perron-Frobenius Theorey

The eigenvector corresponding to the largest eigenvalue of the adjacency matrix of a graph is usually not a constant vector. However, it is always a positive vector if the graph is connected.

This follows from the Perron-Frobenius theory. In fact, the Perron-Frobenius theory says much more, and it can be applied to adjacency matrices of strongly connected directed graphs. Note that

[1] If this makes you uncomfortable, you could use both graphs on two vertices

these need not even be diagonalizable! We will defer a discussion of the general theory until we discuss directed graphs, which will happen towards the end of the semester. If you want to see it now, look at the third lecture from my notes from 2009.

In the symmetric case, the theory is made much easier by both the spectral theory and the characterization of eigenvalues as extreme values of Rayleigh quotients.

Theorem 6.9.1. *[Perron-Frobenius, Symmetric Case] Let G be a connected weighted graph, let $\boldsymbol{M}$ be its adjacency matrix, and let $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n$ be its eigenvalues. Then*

a. $\mu_1 \geq -\mu_n$, *and*

b. $\mu_1 > \mu_2$,

c. The eigenvalue μ_1 has a strictly positive eigenvector.

Before proving Theorem 6.9.1, we will prove a lemma that will be useful in the proof and a few other places today. It says that non-negative eigenvectors of non-negative adjacency matrices of connected graphs must be strictly positive.

Lemma 6.9.2. *Let G be a connected weighted graph (with non-negative edge weights), let $\boldsymbol{M}$ be its adjacency matrix, and assume that some non-negative vector $\boldsymbol{\phi}$ is an eigenvector of $\boldsymbol{M}$. Then, $\boldsymbol{\phi}$ is strictly positive.*

Proof. Assume by way of contradiction that $\boldsymbol{\phi}$ is not strictly positive. So, there is some vertex u for which $\boldsymbol{\phi}(u) = 0$. Thus, there must be some edge (u, v) for which $\boldsymbol{\phi}(u) = 0$ but $\boldsymbol{\phi}(v) > 0$. We would then

$$(\boldsymbol{M}\boldsymbol{\phi})(u) = \sum_{(u,z)\in E} w(u,z)\boldsymbol{\phi}(z) \geq w(u,v)\boldsymbol{\phi}(v) > 0,$$

as all the terms $w(u,z)$ and $\boldsymbol{\phi}(z)$ are non-negative. But, this must also equal $\mu\boldsymbol{\phi}(u) = 0$, where μ is the eigenvalue corresponding to $\boldsymbol{\phi}$. This is a contradiction.

So, we conclude that $\boldsymbol{\phi}$ must be strictly positive. □

We probably won't have time to say any more about Perron-Frobenius theory, so I defer the proof of the theorem to the appendix.

A Perron-Frobenius, continued

Proof of Theorem 6.9.1. Let $\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_n$ be the eigenvectors corresponding to $\mu_1, \ldots, \mu_n$.

We start with part *c*. Recall that

$$\mu_1 = \max_{\boldsymbol{x}} \frac{\boldsymbol{x}^T \boldsymbol{M} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}}.$$

Let $\boldsymbol{\phi}_1$ be an eigenvector of μ_1, and construct the vector $\boldsymbol{x}$ such that

$$\boldsymbol{x}(u) = |\boldsymbol{\phi}_1(u)|, \text{for all } u.$$

We will show that $\boldsymbol{x}$ is an eigenvector of eigenvalue μ_1.

We have $\boldsymbol{x}^T\boldsymbol{x} = \boldsymbol{\phi}_1^T\boldsymbol{\phi}_1$. Moreover,

$$\boldsymbol{\phi}_1^T\boldsymbol{M}\boldsymbol{\phi}_1 = \sum_{u,v}\boldsymbol{M}(u,v)\boldsymbol{\phi}_1(u)\boldsymbol{\phi}_1(v) \leq \sum_{u,v}\boldsymbol{M}(u,v)\left|\boldsymbol{\phi}_1(u)\right|\left|\boldsymbol{\phi}_1(v)\right| = \boldsymbol{x}^T\boldsymbol{M}\boldsymbol{x}.$$

So, the Rayleigh quotient of $\boldsymbol{x}$ is at least μ_1. As μ_1 is the maximum possible Rayleigh quotient, the Rayleigh quotient of $\boldsymbol{x}$ must be μ_1 and $\boldsymbol{x}$ must be an eigenvector of μ_1.

So, we now know that $\boldsymbol{M}$ has an eigenvector $\boldsymbol{x}$ that is non-negative. We can then apply Lemma 6.9.2 to show that $\boldsymbol{x}$ is strictly positive.

To prove part *b*, let $\boldsymbol{\phi}_n$ be the eigenvector of μ_n and let $\boldsymbol{y}$ be the vector for which $\boldsymbol{y}(u) = |\boldsymbol{\phi}_n(u)|$. In the spirit of the previous argument, we can again show that

$$|\mu_n| = |\boldsymbol{\phi}_n\boldsymbol{M}\boldsymbol{\phi}_n| \leq \sum_{u,v}\boldsymbol{M}(u,v)\boldsymbol{y}(u)\boldsymbol{y}(v) \leq \mu_1\boldsymbol{y}^T\boldsymbol{y} = \mu_1.$$

To show that the multiplicity of μ_1 is 1 (that is, $\mu_2 < \mu_1$), consider an eigenvector $\boldsymbol{\phi}_2$. As $\boldsymbol{\phi}_2$ is orthogonal to $\boldsymbol{\phi}_1$, it must contain both positive and negative values. We now construct the vector $\boldsymbol{y}$ such that $\boldsymbol{y}(u) = |\boldsymbol{\phi}_2(u)|$ and repeat the argument that we used for $\boldsymbol{x}$. We find that

$$\mu_2 = \frac{\boldsymbol{\phi}_2^T\boldsymbol{M}\boldsymbol{\phi}_2}{\boldsymbol{\phi}_2\boldsymbol{\phi}_2} \leq \frac{\boldsymbol{y}^T\boldsymbol{M}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \leq \mu_1.$$

From here, we divide the proof into two cases. First, consider the case in which $\boldsymbol{y}$ is never zero. In this case, there must be some edge (u,v) for which $\boldsymbol{\phi}_2(u) < 0 < \boldsymbol{\phi}_2(v)$. Then the above inequality must be strict because the edge (u,v) will make a negative contribution to $\boldsymbol{\phi}_2^T\boldsymbol{M}\boldsymbol{\phi}_2$ and a positive contribution to $\boldsymbol{y}^T\boldsymbol{M}\boldsymbol{y}$.

We will argue by contradiction in the case that $\boldsymbol{y}$ has a zero value. In this case, if $\mu_2 = \mu_1$ then $\boldsymbol{y}$ will be an eigenvector of eigenvalue μ_1. This is a contradiction, as Lemma 6.9.2 says that a non-negative eigenvector cannot have a zero value. So, if $\boldsymbol{y}$ has a zero value then $\boldsymbol{y}^T\boldsymbol{M}\boldsymbol{y} < \mu_1$ and $\mu_2 < \mu_1$ as well. $\square$

The following characterization of bipartite graphs follows from similar ideas.

Proposition A.1. *If G is a connected graph, then $\mu_n = -\mu_1$ if and only if G is bipartite.*

Proof. First, assume that G is bipartite. That is, we have a decomposition of V into sets U and W such that all edges go between U and W. Let $\boldsymbol{\phi}_1$ be the eigenvector of μ_1. Define

$$\boldsymbol{x}(u) = \begin{cases} \boldsymbol{\phi}_1(u) & \text{if } u \in U, and \\ -\boldsymbol{\phi}_1(u) & \text{if } u \in W. \end{cases}$$

For $u \in U$, we have

$$(\boldsymbol{M}\boldsymbol{x})(u) = \sum_{(u,v)\in E}\boldsymbol{x}(v) = -\sum_{(u,v)\in E}\boldsymbol{\phi}(v) = -\mu_1\boldsymbol{\phi}(u) = -\mu_1\boldsymbol{x}(u).$$

Using a similar argument for $u \notin U$, we can show that $\boldsymbol{x}$ is an eigenvector of eigenvalue $-\mu_1$.

To go the other direction, assume that $\mu_n = -\mu_1$. We then construct $\boldsymbol{y}$ as in the previous proof, and again observe

$$|\mu_n| = |\boldsymbol{\phi}_n \boldsymbol{M} \boldsymbol{\phi}_n| = \left| \sum_{u,v} \boldsymbol{M}(u,v)\boldsymbol{\phi}_n(u)\boldsymbol{\phi}_n(v) \right| \leq \sum_{u,v} \boldsymbol{M}(u,v)\boldsymbol{y}(u)\boldsymbol{y}(v) \leq \mu_1 \boldsymbol{y}^T \boldsymbol{y} = \mu_1.$$

For this to be an equality, it must be the case that $\boldsymbol{y}$ is an eigenvalue of μ_1, and so $\boldsymbol{y} = `\boldsymbol{\phi}_1$. For the first inequality above to be an equality, it must also be the case that all the terms $\boldsymbol{\phi}_n(u)\boldsymbol{\phi}_n(v)$ have the same sign. In this case that sign must be negative. So, we every edge goes between a vertex for which $\boldsymbol{\phi}_n(u)$ is positive and a vertex for which $\boldsymbol{\phi}_n(v)$ is negative. Thus, the signs of $\boldsymbol{\phi}_n$ give the bi-partition. □

B Proofs for Hoffman's lower bound on chromatic number

To prove Lemma 6.5.2, we begin with the case of $k = 2$. The general case follows from this one by induction.

Lemma B.1. *Let*

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{B} & \boldsymbol{C} \\ \boldsymbol{C}^T & \boldsymbol{D} \end{bmatrix}$$

be a symmetric matrix. Then

$$\lambda_{min}(\boldsymbol{A}) + \lambda_{max}(\boldsymbol{A}) \leq \lambda_{max}(\boldsymbol{B}) + \lambda_{max}(\boldsymbol{D}).$$

Proof. Let $\boldsymbol{x}$ be an eigenvector of $\boldsymbol{A}$ of eigenvalue $\lambda_{max}(\boldsymbol{A})$. To simplify formulae, let's also assume that $\boldsymbol{x}$ is a unit vector. Write $\boldsymbol{x} = \begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \end{pmatrix}$, using the same partition as we did for $\boldsymbol{A}$.

We first consider the case in which neither $\boldsymbol{x}_1$ nor $\boldsymbol{x}_2$ is an all-zero vector. In this case, we set

$$\boldsymbol{y} = \begin{pmatrix} \frac{\|\boldsymbol{x}_2\|}{\|\boldsymbol{x}_1\|}\boldsymbol{x}_1 \\ -\frac{\|\boldsymbol{x}_1\|}{\|\boldsymbol{x}_2\|}\boldsymbol{x}_2 \end{pmatrix}.$$

The reader may verify that $\boldsymbol{y}$ is also a unit vector, so

$$\boldsymbol{y}^T A \boldsymbol{y} \geq \lambda_{min}(\boldsymbol{A}).$$

We have

$$
\begin{aligned}
\lambda_{max}(\boldsymbol{A}) + \lambda_{min}(\boldsymbol{A}) &\leq \boldsymbol{x}^T A \boldsymbol{x} + \boldsymbol{y}^T A \boldsymbol{y} \\
&= \boldsymbol{x}_1^T \boldsymbol{B} \boldsymbol{x}_1 + \boldsymbol{x}_1^T \boldsymbol{C} \boldsymbol{x}_2 + \boldsymbol{x}_2^T \boldsymbol{C}^T \boldsymbol{x}_1 + \boldsymbol{x}_2^T \boldsymbol{D} \boldsymbol{x}_2 + \\
&\quad + \frac{\|\boldsymbol{x}_2\|^2}{\|\boldsymbol{x}_1\|^2} \boldsymbol{x}_1^T \boldsymbol{B} \boldsymbol{x}_1 - \boldsymbol{x}_1^T \boldsymbol{C} \boldsymbol{x}_2 - \boldsymbol{x}_2^T \boldsymbol{C}^T \boldsymbol{x}_1 + \frac{\|\boldsymbol{x}_1\|^2}{\|\boldsymbol{x}_2\|^2} \boldsymbol{x}_2^T \boldsymbol{D} \boldsymbol{x}_2 \\
&= \boldsymbol{x}_1^T \boldsymbol{B} \boldsymbol{x}_1 + \boldsymbol{x}_2^T \boldsymbol{D} \boldsymbol{x}_2 + \frac{\|\boldsymbol{x}_2\|^2}{\|\boldsymbol{x}_1\|^2} \boldsymbol{x}_1^T \boldsymbol{B} \boldsymbol{x}_1 + \frac{\|\boldsymbol{x}_1\|^2}{\|\boldsymbol{x}_2\|^2} \boldsymbol{x}_2^T \boldsymbol{D} \boldsymbol{x}_2 \\
&\leq \left(1 + \frac{\|\boldsymbol{x}_2\|^2}{\|\boldsymbol{x}_1\|^2}\right) \boldsymbol{x}_1^T \boldsymbol{B} \boldsymbol{x}_1 + \left(1 + \frac{\|\boldsymbol{x}_1\|^2}{\|\boldsymbol{x}_2\|^2}\right) \boldsymbol{x}_2^T \boldsymbol{D} \boldsymbol{x}_2 \\
&\leq \lambda_{max}(\boldsymbol{B}) \left(\|\boldsymbol{x}_1\|^2 + \|\boldsymbol{x}_2\|^2\right) + \lambda_{max}(\boldsymbol{D}) \left(\|\boldsymbol{x}_1\|^2 + \|\boldsymbol{x}_2\|^2\right) \\
&= \lambda_{max}(\boldsymbol{B}) + \lambda_{max}(\boldsymbol{D}),
\end{aligned}
$$

as $\boldsymbol{x}$ is a unit vector.

We now return to the case in which $\|\boldsymbol{x}_2\| = 0$ (or $\|\boldsymbol{x}_1\| = 0$, which is really the same case). Lemma 6.8.1 tells us that $\lambda_{max}(\boldsymbol{B}) \leq \lambda_{max}(\boldsymbol{A})$. So, it must be the case that $\boldsymbol{x}_1$ is an eigenvector of eigenvalue $\lambda_{max}(\boldsymbol{A})$ of $\boldsymbol{B}$, and thus $\lambda_{max}(\boldsymbol{B}) = \lambda_{max}(\boldsymbol{A})$. To finish the proof, also observe that Lemma 6.8.1 implies

$$
\lambda_{max}(\boldsymbol{D}) \geq \lambda_{min}(\boldsymbol{D}) \geq \lambda_{min}(\boldsymbol{A}).
$$

□

Proof of Lemma 6.5.2. For $k = 2$, this is exactly Lemma B.1. For $k > 2$, we apply induction. Let

$$
B = \begin{bmatrix}
\boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} & \cdots & \boldsymbol{A}_{1,k-1} \\
\boldsymbol{A}_{1,2}^T & \boldsymbol{A}_{2,2} & \cdots & \boldsymbol{A}_{2,k-1} \\
\vdots & \vdots & \ddots & \vdots \\
\boldsymbol{A}_{1,k-1}^T & \boldsymbol{A}_{2,k-1}^T & \cdots & \boldsymbol{A}_{k-1,k-1}
\end{bmatrix}.
$$

Lemma 6.8.1 now implies.

$$
\lambda_{min}(\boldsymbol{B}) \geq \lambda_{min}(\boldsymbol{A}).
$$

Applying Lemma B.1 to B and the kth row and column of $\boldsymbol{A}$, we find

$$
\begin{aligned}
\lambda_{min}(\boldsymbol{A}) + \lambda_{max}(\boldsymbol{A}) &\leq \lambda_{max}(\boldsymbol{B}) + \lambda_{max}(\boldsymbol{A}_{k,k}) \\
&\leq -(k-2)\lambda_{min}(\boldsymbol{B}) + \sum_{i=1}^{k-1} \lambda_{max}(\boldsymbol{A}_{i,i}) + \lambda_{max}(\boldsymbol{A}_{k,k}) \qquad \text{(by induction)} \\
&\leq -(k-1)\lambda_{min}(\boldsymbol{A}) + \sum_{i=1}^{k} \lambda_{max}(\boldsymbol{A}_{i,i}),
\end{aligned}
$$

which proves the lemma. □

References

[AH77a] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable part i. discharging. *Illinois Journal of Mathematics*, 21:429–490, 1977.

[AH77b] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable part ii. reducibility. *Illinois Journal of Mathematics*, 21:491–567, 1977.

[AK97] Alon and Kahale. A spectral technique for coloring random 3-colorable graphs. *SICOMP: SIAM Journal on Computing*, 26, 1997.

[BRSW12] Boaz Barak, Anup Rao, Ronen Shaltiel, and Avi Wigderson. 2-source dispersers for $n^{o(1)}$ entropy, and ramsey graphs beating the frankl-wilson construction. *Annals of Mathematics*, pages 1483–1543, 2012.

[Erd47] Paul Erdös. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947.

[GN08] C.D. Godsil and M.W. Newman. Eigenvalue bounds for independent sets. *Journal of Combinatorial Theory, Series B*, 98(4):721 – 734, 2008.