CHAPTER 4

Paths, flows, and routing

4.1. Paths and sets of paths

One of the main themes in graph theory concerns paths joining pairs of vertices. For example, the Hamiltonian path problem is to decide if a graph has a simple path containing every vertex of the graph. Some diameter and distance problems involve finding shortest paths. There are many basic problems depending on sets of paths that are either vertex-disjoint or edge-disjoint. These path problems arise naturally in a variety of guises, such as the study of communicating processes on networks, data flow on parallel computers, and the analysis of routing algorithms on VLSI chips. Some path problems appear to be quite difficult computationally. For example, the Hamiltonian path problem is well known to be NP-complete. The problem of finding disjoint paths between given pairs of vertices even in very special graphs [**139**] is also NP-complete. Nevertheless, we will see that eigenvalue techniques are amazingly effective in providing good solutions for a range of path problems.

Before we proceed, we first define several types of disjoint paths that we call *flow*, *route set*, and *routing*. Consider a graph G with vertex set V and edge set E. Suppose X and Y are two equinumerous subsets of vertices of G. In general, X and Y can be multisets and it is not necessary to require $X \cap Y = \emptyset$.

For $|X| = |Y| = m$, a *flow* F from X to Y consists of m paths in G joining the vertices in X to the vertices in Y. We call X the *input* of the flow F and Y the *output* of F. Paths in F join vertices of X to vertices in Y in a one-to-one fashion, but we do not care about "who is talking to whom." We *do* care that the paths be chosen so that no edge is overused. For example, the paths might be required to be edge-disjoint or vertex-disjoint or with small "congestion" in the sense that every edge (or vertex) of G is used in relatively few paths of F. We will define "congestion" precisely later.

A *route set* is a flow with input-output assignments. Namely, for a specified assignment $A = \{(x_i, y_i) : x_i \in X, y_i \in Y\}$, a route set consists of paths P_i joining x_i to y_i for each i. In other words, an assignment specifies "who is talking to whom."

Roughly speaking, a *routing* R is a dynamic version of a route set. It can be defined as a pebble game. Initially, there is a pebble p_i placed at each input vertex x_i with destination y_i for each of the assignments (x_i, y_i) in A. At each time unit, a pebble can be moved to some adjacent vertex. The routing R is then a route set together with a strategy for moving pebbles to their destinations. Additional requirements can be imposed. For example, at each time unit, the edges used for moving pebbles should be (vertex- or edge-) disjoint or all edges must have small congestion.

Flow and routes are very useful in establishing lower bounds for Cheeger constants as well as providing lower bounds for eigenvalues (see Sections 4.2 and 4.5). Conversely, for graphs with good eigenvalue lower bounds, short routes and effective routing schemes exist with small congestion which will be described in Sections 4.3 and 4.4.

4.2. Flows and Cheeger constants

Flows are closely related to cuts as evidenced by the max flow-min cut theorem which was used in the previous chapter. In fact, there is a direct connection between the Cheeger constants and flow problems on graphs. Although these observations are quite easy, we will state them here since they are useful for bounding eigenvalues. We follow the definition for Cheeger constants h_G and h'_G as given in Sections 2.2 and 2.5.

LEMMA 4.1. *For a graph G on n vertices, suppose there is a set of $\binom{n}{2}$ paths joining all pairs of vertices such that each edge of G is contained in at most m paths. Then*

$$h'_G = \inf_S \frac{|E(S,\bar{S})|}{\min(|S|,|\bar{S}|)} \geq \frac{n}{2m}.$$

PROOF. The proof follows from the simple fact that for any set $S \subseteq V$ with $|S| \leq |\bar{S}|$, we have

$$\begin{aligned} |E(S,\bar{S})| \cdot m &\geq |S| \cdot |\bar{S}| \\ &\geq |S| \cdot \frac{n}{2}. \end{aligned}$$

□

As an immediate consequence, we have the following:

COROLLARY 4.2. *For a k-regular graph G on n vertices, suppose there is a set P of $\binom{n}{2}$ paths joining all pairs of vertices such that each edge of G is contained in at most m paths in P. Then the Cheeger constant h_G satisfies*

$$h_G = \inf_S \frac{|E(S,\bar{S})|}{k\min(|S|,|\bar{S}|)} \geq \frac{n}{2mk}.$$

By using Cheeger's inequality in Chapter 2 and the above lower bound for the Cheeger constant derived from a flow, we can establish eigenvalue lower bounds for a regular graph. In fact, we can derive a better lower bound for λ_1 directly from a flow in a general graph. We first prove a simple version for a regular graph.

THEOREM 4.3. *For a k-regular graph G on n vertices, suppose there is a set P of $\binom{n}{2}$ paths joining all pairs of vertices such that each path in P has length at most l and each edge of G is contained in at most m paths in P. Then the eigenvalue λ_1 satisfies*

$$\lambda_1 \geq \frac{n}{kml}.$$

PROOF. Using the definition (1.5) of the eigenvalues, we consider the harmonic eigenfunction $f : V(G) \to \mathbb{R}$ achieving λ_1.

$$\lambda_1 = \frac{n \displaystyle\sum_{\{x,y\} \in E(G)} (f(x) - f(y))^2}{k \displaystyle\sum_{x,y} (f(x) - f(y))^2}$$

We note that for $x, y \in V(G)$ and the path $P(x,y)$ joining x and y in G, we have

$$(f(x) - f(y))^2 \leq |P(x,y)| \sum_{e \in P(x,y)} f^2(e) \leq l \sum_{e \in P(x,y)} f^2(e),$$

where $f^2(e) = (f(x) - f(y))^2$ for $e = \{x, y\}$, and $|P(x,y)|$ denotes the number of edges of G in $P(x,y)$. Hence

$$\begin{aligned} m \sum_{e \in E(G)} f^2(e) &\geq \sum_{x,y} \sum_{e \in P(x,y)} f^2(e) \\ &\geq \frac{1}{l} \sum_{x,y} (f(x) - f(y))^2. \end{aligned}$$

Therefore we have

$$\lambda_1 \geq \frac{n}{kml}.$$

This completes the proof of Theorem 4.3. □

For a general graph, the above theorems can be generalized as follows:

THEOREM 4.4. *For an undirected graph G, replace each edge $\{u, v\}$ by two directed edges (u, v) and (v, u). Suppose there is a set P of $4e^2$ paths such that for each (ordered) pair of directed edges there is a directed path joining them. In addition, assume that each directed edge of G is contained in at most m directed paths in P. Then the Cheeger constant h_G satisfies*

$$h_G = \min_S \frac{|E(S, \bar{S})|}{\min(\text{vol } S, \text{vol } \bar{S})} \geq \frac{\text{vol } G}{2m}.$$

PROOF. For any $S \subseteq V(G)$, we have

$$m|E(S, \bar{S})| \geq \text{vol } S \text{ vol } \bar{S} \geq \frac{\text{vol } S \text{ vol } G}{2}.$$

□

THEOREM 4.5. *For an undirected graph G, replace each edge $\{u, v\}$ by two directed edges (u, v) and (v, u). Suppose there is a set P of $4e^2$ directed paths such that for each (ordered) pair of directed edges there is a directed path joining them, each of length at most l. In addition, assume that each directed edge of G is contained in at most m directed paths in P. Then the eigenvalue λ_1 satisfies*

$$\lambda_1 \geq \frac{\text{vol } G}{ml}.$$

The proof of Theorem 4.5 is very similar to that of Theorem 4.3 and will be omitted.

We remark that Theorems 4.3 and 4.5 can be generalized in a number of ways. For example, instead of having one path joining two vertices, we can ask for a number of paths or weighted paths with fixed total capacities (in the spirit of the max flow-min cut theorem). Another direction is to derive the comparison theorems which will be discussed in Section 4.5.

4.3. Eigenvalues and routes with small congestion

In a graph G, a random walk of length l starting at a vertex v of G is a randomly chosen sequence $v = v_0, v_1, \ldots, v_l$, where each v_{i+1} is chosen, uniformly at random and independently, among the neighbors of v_i, for $i = 0, \ldots, l-1$. We say that the walk visits v_i at time i.

In a graph G with $\lambda_1 > 0$, a random walk starting from any vertex converges roughly in $\frac{\log n}{\lambda_1}$ steps to the stationary distribution (if G is bipartite, we use a lazy random walk; see Section 1.5). We will use this property to derive the following fact.

THEOREM 4.6. *Let G be a graph on n vertices and suppose $l \geq \log n/\lambda_1$. Suppose for any $v \in V(G)$ there are d_v random walks of length l starting at v. For any edge q, let $I(q)$ denote the total number of walks containing q. Then, almost surely (i.e., with probability tending to 1 as n tends to infinity), there is no edge q so that*

$$I(q) > 10l.$$

PROOF. Let P denote the transition matrix defined by

$$P(u, v) = \begin{cases} 1/d_u & \text{if } u \text{ and } v \text{ are adjacent,} \\ 0 & \text{otherwise.} \end{cases}$$

The probability that a random walk $W(u)$ starting at u visits a vertex x at time i is precisely $\psi_u P^i (\psi_x)^*$ where ψ_y is the unit vector having 1 in coordinate y and 0 in every other coordinate. For a directed edge (u, v), the probability that a random walk $W(x)$ visits u at time i and v at time $i+1$ is

$$\psi_x P^i \psi_u^* / d_u.$$

With d_v walks starting at v, the sum of the probabilities that there exists a walk $W(x)$ that visits $q = \{u, v\}$ is

$$\begin{aligned} I(q) &\leq \sum_{i=0}^{l-1} \sum_x d_x \psi_x P^i(\psi_u^*/d_u + \psi_v^*/d_v) \\ &= \sum_{i=0}^{l-1} (\mathbf{1}T) P^i(\psi_u^*/d_u + \psi_v^*/d_v) \end{aligned}$$

where $\mathbf{1}$ denotes the all 1's vector, T is the diagonal matrix with entries $T(v,v) = d_v$, and

$$(\mathbf{1}T)P^j = \mathbf{1}T.$$

Therefore

$$\begin{aligned} I(q) &\leq \sum_{i=0}^{l-1} \mathbf{1}T(\psi_u^*/d_u + \psi_v^*/d_v) \\ &= 2l. \end{aligned}$$

Therefore, for each fixed v, the expectation of the random variable $I(e)$ is no more than l. We observe that this random variable is a sum of $|E(G)|$ independent indicator random variables (see, e.g., [**12**] Theorem A.12, page 237), and that for each fixed edge q, the probability that $I(q)$ exceeds, say $10l$, is at most

$$(\frac{(2e)^9}{10^{10}})^l \ll \frac{1}{n^3}$$

where e here denotes the natural logarithm. Since there are at most n^2 edges, it follow that the probability that there is an edge with $I(q) > 10l$ is much smaller than $\frac{1}{n}$. This completes the proof of Theorem 4.6. □

The above estimate can in fact be proved directly. Suppose there is a set of m independent events such that the probability of the i-th event is p_i. Furthermore, suppose that $\sum p_i \leq \gamma$. Then the probability that at least s events occur is bounded by

$$\begin{aligned} \sum_{\substack{S\subset\{1,\ldots,m\} \\ |S|=s}} \prod_{i\in S} p_i &\leq \frac{1}{s!}(\sum p_i)^s \\ &\leq (\frac{\gamma e}{s})^s. \end{aligned}$$

The proof of Theorem 4.6 follows by choosing $\gamma = l$ and $s = 10\gamma$.

THEOREM 4.7. *Let G denote a graph on n vertices. Let $A = \{(x_i, y_i) : x_i \in X, y_i \in Y\}$ denote any assignment such that each vertex v is in X with multiplicity d_v and in Y with multiplicity d_v. Then there are paths P_i joining x_i to y_i of length at most $\frac{2}{\lambda_1}\log n$ such that each edge of G is contained in at most $\frac{20}{\lambda_1}\log n$ paths P_j.*

PROOF. Let P_i denote a random walk of length $2l$ between x_i and y_i where $l \approx \frac{\log n}{\lambda_1}$. Using an argument of Valiant in his work on parallel routing [**245**] (also see [**40**]), we may assume that each walk consists of two random walks of length

l, one starting from x_i and the other from y_i. The reason for this is that by our eigenvalue condition, the distribution of the random walk of length l is close to its stationary distribution and hence one may view the walk P_i as being chosen by first choosing its middle point (according to the stationary distribution) and then choosing its two halves. The proof of Theorem 4.7 then follows from Theorem 4.6. □

The above proofs for Theorem 4.6 and 4.7 were adapted from the following simpler version for regular graphs in [**7**]:

THEOREM 4.8. *Let G denote a k-regular graph on n vertices. Let π denote a permutation of the vertices of G. Then there are paths P_x joining x to $\pi(x)$ of length at most $\frac{2}{\lambda_1}\log n$ such that each edge of G is contained in at most $\frac{20}{k\lambda_1}\log n$ paths P_y.*

4.4. Routing in graphs

In this section, we consider a simple (though fundamental) problem of the following type: Suppose we are given a connected graph G with vertex set V and edge set E. Initially, each vertex v of G is occupied by a unique marker or "pebble" p_v. To each pebble p_v is associated a destination vertex $\pi(v) \in V$, so that distinct pebbles have distinct destinations. Pebbles can be moved to different vertices of G according to the following basic procedure: At each step a disjoint collection of edges of G is selected and the pebbles at each edge's two endpoints are interchanged. Our goal is to move or "route" the pebbles to their respective destinations in a minimum number of steps.

We will imagine the steps occurring at discrete times, and we let $p_v(t) \in V$ denote the location of the pebble with initial position v at time $t = 0, 1, 2, \ldots$. Thus, for any t, the set $\{p_v(t) : v \in V\}$ is a permutation of V. We will denote our target permutation that takes v to $\pi(v)$, $v \in V$, by π. Define $rt(G, \pi)$ to be the minimum possible number of steps to achieve π. Finally, define $rt(G)$, the *routing number* of G, by

$$rt(G) = \max_{\pi} \ rt(G, \pi)$$

where π ranges over all destination permutations on G. (Sometimes we will also call π a routing assignment.)

In more algebraic terms, the problem is simply to determine for G the largest number of terms $\tau = (u_1v_1)(u_2v_2)\cdots(u_rv_r)$ ever required to represent any permutation in the symmetric group on $n = |V|$ symbols, where each permutation τ consists of a product of disjoint transpositions (u_kv_k) with all pairs $\{u_k, v_k\}$ required to be edges of G.

To see that $rt(G)$ always exists, let us restrict our attention to some spanning subtree T of G. It is clear that if p has destination which is a leaf of T, then we can first route p to its destination u, and then complete the routing on $T \setminus \{u\}$ by induction.

An obvious lower bound on $rt(G)$ is the following:

$$rt(G) \geq D(G)$$

where $D(G)$ denotes the diameter of G.

For P_n a path on n vertices, our routing problem reduces to a well studied problem in parallel sorting networks, the so-called *odd-even transposition sort* (see [**174**] for a comprehensive survey). In this case, it can be shown that $rt(P_n) = n$ for $n \geq 3$. In fact, any permutation π on P_n can be sorted in n steps by labelling consecutive edges in P_n as $e_1, e_2, \ldots, e_{n-1}$ and only making interchanges with *even* edges e_{2k} on *even* steps and *odd* edges e_{2k+1} on *odd* steps.

Let K_n denote the complete graph on n vertices. In this case, because K_n is so highly connected, the routing number of K_n is as small as one could hope for (see [**7**]):

$$rt(K_n) = 2.$$

For the complete bipartite graph $K_{n,n}$ with $n \geq 3$, the following result is due to Wayne Goddard [**142**].

$$rt(K_{n,n}) = 4.$$

For any tree T_n on n vertices, it was proved in [**7**] that

$$rt(T_n) < 3n.$$

However, the correct value of the constant may be half as large, as suggested by the following:

Conjecture. For any tree T_n on n vertices,

$$rt(T_n) \leq \lfloor \frac{3(n-1)}{2} \rfloor.$$

Furthermore, equality holds only when T_n is a star S_n on n vertices.

We remark that Louxin Zhang [**255**] has proved an asymptotical version of the above conjecture by showing $rt(T_n) = 3n/2 + O(\log n)$.

The following result on routing on the hypercube can be traced back to the early work of switching networks (see Beněs [**20**]) and has appeared frequently in the literature on parallel computing:

$$rt(Q_n) \leq 2n - 1.$$

The exact value of $rt(Q_n)$ is still unknown. It is easy to see that $rt(Q_n) \geq n$ since the diameter of Q_n is n. For small cases, it can be checked that $rt(Q_n) \geq n+1$ for $n = 2, 3$.

Problem: Is it true that for the n-cube Q_n,

$$rt(Q_n) = n + o(n)\,?$$

Perhaps, $rt(Q_n) = n + o(n)$ for all sufficiently large n.

For the m by n grid graph $P_m \times P_n$, $m \leq n$,

$$rt(P_m \times P_n) \leq 2m + n.$$

In general, for the cartesian product of two graphs, we have [**7, 179**]

$$rt(G \times G') \leq 2rt(G) + rt(G').$$

Note that since $G \times G'$ and $G' \times G$ are isomorphic graphs, this can be written in the symmetric form

$$rt(G \times G') \leq \min\{2rt(G) + rt(G'),\ 2rt(G') + rt(G).\}$$

Problem: Is it true that for every connected graph G,

$$rt(G \times G) \ \geq \ rt(G)\,?$$

From the above results and partial results, we can see that the problem of determining the routing number is quite difficult even for very special graphs. It is indeed surprising in a way that by using eigenvalues we can get very good approximations for the routing number problem. The following arguments are basically adapted from [**7**]. In the remaining part of this section, we assume G is a regular graph.

THEOREM 4.9. *Let G be a regular graph on n vertices and suppose $l \geq \log n/\lambda_1$. For each $v \in V$, independently, let $W(v)$ denote a random walk of length l starting at v. Let $I(v)$ denote the total number of other walks $W(u)$ such that there exists a vertex x and two indices $0 \leq i, j \leq l,\ |i - j| < 5$, so that $W(v)$ visits x at time i and $W(u)$ visits x at time j. Then, almost surely (i.e., with probability tending to 1 as n tends to infinity), there is no vertex v so that $I(v) > 100l$.*

The proof is very similar to that of Theorem 4.6 and is omitted.

THEOREM 4.10. *Let G denote a regular graph on n vertices and let σ be a permutation of order two on V (i.e., a product of pairwise disjoint transpositions). Put $l = \frac{10}{\lambda_1}\log n$. Then there is a set of n walks $W(v)$, $v \in V$, each of length $2l$, where both $W(v)$ and $W(\sigma(v))$ connect v and $\sigma(v)$ and traverse the same set of edges (in different directions) satisfying the following: If $I(v)$ denotes the total number of other walks $W(u)$ such that there exists a vertex x and two indices $0 \leq i, j \leq l,\ |i - j| < 5$, so that $W(v)$ visits x at time i and $W(u)$ visits x at time j or at time $2l - j$, then $I(v) \leq 400l$ for all v.*

THEOREM 4.11. *Let σ denote a permutation on the vertex set of G. Then*

$$rt(G, \sigma) = O(\frac{1}{\lambda_1}\log^2 n).$$

PROOF. Let G denote a regular graph on n vertices. It suffices to consider a permutation σ of order two on V since any permutation is a product of at most two such permutations (as proved in the proof of $rt(K_n) = 2$; also see [**7**]). We set $l = \frac{10}{\lambda_1}\log n$. We want to show that $rt(G, \sigma) = O(l^2)$. Let $W(v)$ be a system of walks of length $2l$ satisfying the assumptions of the previous theorem. Let H be the graph whose vertices are the walks $W(v)$ in which $W(u)$ and $W(v)$ are adjacent if there exists a vertex x and two indices $0 \leq i, j \leq l,\ |i - j| < 5$ so that $W(v)$ visits x at time i and $W(u)$ visits x at time j or at time $2l - j$. Then the maximum degree of H is $O(l)$ and hence it is $O(l)$-colorable. It follows that one can split all our paths $W(v)$ into $O(l)$ classes of paths such that the paths in each class are not adjacent in H. Consider now the following routing algorithm. For each set of paths as above, perform $2l$ steps, where the steps numbered i and $2l+1-i$ correspond to flipping the pebbles along edges numbered i and $2l+1-i$ in each of the paths in the set for all $1 \leq i \leq l$. One can check that by the end of these $2l$ steps, the ends of each path exchange pebbles, and all the other pebbles stay in their original places. (Note that some pebbles that are not at the ends of any of the paths may move several times during these steps, but the symmetric way these are performed guarantees that such pebbles will return to their original places at the completion of the $2l$ steps.) By repeating the above procedure for all the path-classes, the result follows. □

We mention here several problems closely related to the routing number of a graph. One such problem is the following:

Suppose $G = (V, E)$ is a connected graph on n vertices. For a permutation π, we consider a *route set* P, which is just some set of paths P_i joining each vertex v_i to its destination vertex $\pi(v_i)$, for $i = 1, \ldots, n$. For each edge e of G, we consider the number $rc(e, G, \pi, P)$ of paths P_i in P which contain e. The *route covering number* $rc(G)$ of G is defined to be

$$rc(G) = \max_{\pi} \min_{P} \max_{e \in E} rc(e, G, \pi, P).$$

In other words, for each permutation we want to choose the route set so that the maximum number of occurrences of any edge in the paths of the route set is minimized.

For example, for the n-cube Q_n, the method [**7**] used to establish the route set gives

$$rc(Q_n) \leq 4.$$

In the other direction, by choosing π to be the permutation of vertices in Q_n so that the distance between v and $\pi(v)$ is n for every vertex v, it can be easily seen that

$$rc(Q_n) \geq \frac{\sum_v d(v, \pi(v))}{|E(Q_n)|} = 2.$$

Conjecture:

$$rc(Q_n) = 2.$$

Also of interest is a "symmetric" version of the route covering problem, especially for Q_n :

A *pairing* for Q_n is a partition of the vertex set of Q_n into subsets of size 2. Is it possible to find edge-disjoint paths joining vertices of each pair for any pairing of Q_n ?

The answer is negative when n is even. However, for odd n this problem remains open.

4.5. Comparison theorems

We can often bound the eigenvalues of one graph by the eigenvalues of another provided pairs of adjacent vertices in the first graph can be joined by "short" paths in the second graph. Although the proofs for these comparison theorems are quite easy, the applications are abundant. Interesting examples along this line are given in Diaconis and Stroock [**103**] and numerous other papers [**98, 99, 124**] for comparing various different card shuffling schemes. We remark that the comparison theorems in this section can be viewed as generalized versions of the so-called "Poincaré" inequalities [**103**].

THEOREM 4.12. *Let G and G' be two connected regular graphs, with eigenvalues λ_1 and λ'_1 and degrees k and k', respectively. Suppose that the vertex set of G is the same as the vertex set of G'. We assume that for each edge $\{x, y\}$ in G, there is a path $P(x, y)$ in G' joining x and y of length at most l. Furthermore, suppose that every edge in G' is contained in at most m paths $P(x, y)$. Then we have*

$$\lambda'_1 \geq \frac{k\lambda_1}{k'lm}.$$

PROOF. Using the definition of the eigenvalues, we consider the harmonic eigenfunction f achieving λ_1 in G'.

$$\begin{aligned}
\lambda'_1 &= \frac{\displaystyle\sum_{\{x,y\}\in E(G')} (f(x)-f(y))^2}{k'\sum f^2(x)} \\
&= \frac{k\displaystyle\sum_{\{x,y\}\in E(G')} (f(x)-f(y))^2}{k'\displaystyle\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2} \cdot \frac{\displaystyle\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2}{k\sum f^2(x)}
\end{aligned}$$

We note that for $\{x, y\} \in E(G)$ and path $P(x, y)$ joining x and y in G', we have

$$(f(x) - f(y))^2 \leq |P(x,y)| \sum_{e\in P(x,y)} f^2(e) \leq l \sum_{e\in P(x,y)} f^2(e)$$

where $f^2(e) = (f(x) - f(y))^2$ for $e = \{x, y\}$, and $|P(x, y)|$ denotes the number of edges of G' in $P(x, y)$. Hence

$$\begin{aligned}
m \sum_{e\in E(G')} f^2(e) &\geq \sum_{\{x,y\}\in E(G)} \sum_{e\in P(x,y)} f^2(e) \\
&\geq \frac{1}{l} \sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2.
\end{aligned}$$

Therefore we have

$$\begin{aligned}\lambda_1' &\geq \frac{k}{k'lm} \cdot \frac{\displaystyle\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2}{\displaystyle\sum f^2(x)k} \\ &\geq \frac{k}{k'lm}\lambda_1.\end{aligned}$$

This completes the proof of Theorem 4.12. □

It is not surprising that the above proof is quite similar to some of those in Section 4.2. There are several generalizations of Theorem 4.12.

THEOREM 4.13. *Let G and G' be two connected graphs, with eigenvalues λ_1 and λ_1', respectively. Suppose that the vertex set of G is the same as the vertex set of G'. Assume that for each edge $\{x,y\}$ in G, there is a path $P(x,y)$ in G' of length at most l, and for each vertex v, the degree d_v of v in G is at least ad_v', where d_v' is the degree of v in G'. Furthermore, suppose every edge in G' is contained in at most m paths $P(x,y)$. Then we have*

$$\lambda_1' \geq \frac{a\lambda_1}{lm}.$$

Instead of proving Theorem 4.13, we will prove the following generalization:

THEOREM 4.14. *Let G and G' be two connected graphs, with eigenvalues λ_1 and λ_1', respectively. Suppose that the vertex set of G can be embedded into the vertex set of G' under the mapping $\varphi : V(G) \to V(G')$. Suppose φ satisfies the following conditions for fixed positive values a, l, m:*

- **(a):** *Each edge $\{x,y\}$ in $E(G)$ is associated with a path, denoted by $P(x,y)$, joining $\varphi(x)$ to $\varphi(y)$ in G' of length at most l.*
- **(b):** *Let d_v, d_v' denote the degrees of v in G and in G', respectively. For any v in $V(G')$, we have*

$$\sum_{x\in\varphi^{-1}(v)} d_x \geq ad_v'.$$

- **(c):** *Each edge in G' is contained in at most m paths $P(x,y)$.*

Then we have

$$\lambda_1' \geq \frac{a\lambda_1}{lm}.$$

PROOF. The proof is very similar to that of Theorem 4.12. For a harmonic eigenfunction g of G', we define $f : V(G) \to \mathbb{R}$ as follows: For a vertex x in $V(G)$,

$$f(x) = g(\varphi(x)) - c$$

where the constant c is chosen to satisfy

$$\sum_x f(x)d_x = 0.$$

We note that

$$\begin{aligned}\sum_{x\in V(G)} f^2(x)d_x &= \sum_{x\in V(G)} (g(\varphi(x))-c))^2 d_x\\ &= \sum_{v\in V(G')} (g(v)-c)^2 \sum_{\varphi^{-1}(v)=x} d_x\\ &\geq a\sum_{v\in V(G')} (g(v)-c)^2 d'_v. \end{aligned}\tag{4.1}$$

Now, for $\{x, y\} \in E(G)$ with $\varphi(x) = u, \varphi(y) = v$, let $P(x, y)$ denote the path corresponding to $\{x, y\}$ joining u and v in G'. We have

$$(g(u)-g(v))^2 \leq |P(x,y)| \sum_{e\in P(x,y)} g^2(e) \leq l \sum_{e\in P(x,y)} g^2(e),$$

where $g^2(e) = (g(a) - g(b))^2$ for $e = \{a, b\}$. Hence we have

$$\begin{aligned} m\sum_{e\in E(G')} g^2(e) &= \sum_{e\in E(G')} mg^2(e)\\ &\geq \sum_{\{x,y\}\in E(G)} \sum_{e\in P(x,y)} g^2(e)\\ &\geq \sum_{\{x,y\}\in E(G)} \sum_{\substack{u=\varphi(x)\\ v=\varphi(y)}} \frac{1}{l}(g(u)-g(v))^2\\ &\geq \frac{1}{l}\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2. \end{aligned}\tag{4.2}$$

Combining inequalities (4.1),(4.2), we have

$$\begin{aligned} \lambda'_1 &= \sup_t \frac{\sum_{\{u,v\}\in E(G')} (g(u)-g(v))^2}{\sum_{v\in V(G')} (g(v)-t)^2 d'_v}\\ &\geq \frac{\sum_{\{u,v\}\in E(G')} (g(u)-g(v))^2}{\sum_{v\in V(G')} (g(v)-c)^2 d'_v}\\ &= \frac{\sum_{\{u,v\}\in E(G')} (g(u)-g(v))^2}{\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2} \cdot \frac{\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2}{\sum_{x\in V(G)} f^2(x)d_x} \cdot \frac{\sum_{x\in V(G)} f^2(x)d_x}{\sum_{v\in V(G')} (g(v)-c)^2 d'_v}\\ &\geq \frac{1}{ml} \cdot \frac{\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2}{\sum_{x\in V(G)} f^2(x)d_x} \cdot a. \end{aligned}$$

Since

$$\sum_x f(x)d_x = 0,$$

we have

$$\frac{\displaystyle\sum_{\{x,y\}\in E(G)} (f(x)-f(y))^2}{\displaystyle\sum_{x\in V(G)} f^2(x)d_x} \geq \lambda_1.$$

Hence

$$\lambda_1' \geq \frac{a}{ml}\lambda_1$$

and the proof of Theorem 4.14 is complete. □

We remark that Theorems 4.3 and 4.5 in Section 4.2 are just special cases of Theorem 4.14 in which G is taken to be a complete graph and G' is chosen arbitrarily.

We also remark that the generalized version in Theorem 4.14 can often give stronger results for certain problems. For example, we consider the following simple and natural random walk problem on generating sets of groups which arises in computational group theory.

Example 4.15. Let H denote a graph on n vertices each of which is labelled by an element of a group Γ. At each unit of time, one of the vertices, say v with label g, can be changed to gf where f or f^{-1} is a label of a neighbor u of v. Suppose we start with the case that the set of all vertex labels generates the group Γ. The problem of interest is to determine how rapidly this processes mixes, i.e., how many steps it requires to be close to a "random" generating set.

By using Theorem 4.14, we can obtain an upper bound of the form cDn^2, where c depends only on the size of Γ, and D denotes the diameter of H (see [**70**]). Similar bounds have also been obtained by Diaconis and Saloff-Coste in [**102**] using more complex comparison techniques. However, all of these bounds are rather far from what is believed to be the truth, namely, that order $n \log n$ steps (under total variation distance). In [**71**], it is proved that in fact this bound is achieved for the case that $H = \mathbb{Z}_2$. Interestingly, for relative pointwise distance, $O(n^2)$ is proved to be the correct bound [**71**].

Notes

Path arguments were used early on to compute isoperimetric constants of a graph. For example, in the early work of Bhatt and Leighton [**24**] on VLSI design and parallel computation, path arguments were extensively utilized. Jerrum and Sinclair [**165**] used path arguments to bound the Cheeger constant in order to bound the eigenvalues in their seminal work of estimating permanents. Diaconis and Stroock [**103**] used path arguments to directly bound eigenvalues.

Sections 4.3 and 4.4 on paths and routing are mainly based on [**7**]. Some variations of the comparison theorems in Sections 4.2 and 4.5 can be found in [**70**].