

ROUTING PERMUTATIONS ON GRAPHS VIA MATCHINGS*

NOGA ALON[†], F. R. K. CHUNG[†], AND R. L. GRAHAM[§]

Abstract. A class of routing problems on connected graphs G is considered. Initially, each vertex v of G is occupied by a “pebble” that has a unique destination $\pi(v)$ in G (so that π is a permutation of the vertices of G). It is required that all the pebbles be routed to their respective destinations by performing a sequence of moves of the following type: A disjoint set of edges is selected, and the pebbles at each edge’s endpoints are interchanged. The problem of interest is to minimize the number of steps required for any possible permutation π .

This paper investigates this routing problem for a variety of graphs G , including trees, complete graphs, hypercubes, Cartesian products of graphs, expander graphs, and Cayley graphs. In addition, this routing problem is related to certain network flow problems, and to several graph invariants including diameter, eigenvalues, and expansion coefficients.

Key words. eigenvalues, diameters, expanders

AMS subject classifications. 05C, 68

1. Introduction. Routing problems on graphs 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. A simple (though fundamental) problem of this type is the following. Suppose we are given a connected graph $G = (V, E)$, where V and E represent the vertex and edge sets, respectively, of G . We denote the cardinality $|V|$ of V by n . Initially, each vertex v of G is occupied by a unique marker or “pebble” p . To each pebble p 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 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 $t = 0, 1, 2, \dots$. Thus, for any t , the set $\{p_v(t) : v \in V\}$ is just a permutation of V . We 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 required 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 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 .

* Received by the editors August 31, 1992; accepted for publication (in revised form) June 3, 1993.

[†] Bell Communications Research, Morristown, New Jersey 07960, and Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel.

[‡] Bell Communications Research, Morristown, New Jersey 07960.

[§] AT&T Bell Laboratories, Murray Hill, New Jersey 07974, and Department of Mathematics, Rutgers University, New Brunswick, New Jersey 08903.

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 a destination that 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.

In this paper, we investigate routing on a variety of graphs. These include trees, complete graphs, complete bipartite graphs, hypercubes, Cartesian products of graphs, Cayley graphs, and expander graphs. We also consider a related continuous version of the routing problem, the so-called flow problem, which is of independent interest. Furthermore, we relate the routing problem on a given graph to several invariants of it, including its diameter, its resistance, and its expansion coefficients and eigenvalues.

2. General bounds on $rt(G)$. To begin, an obvious lower bound on $rt(G)$ is the following:

$$(1) \quad rt(G) \geq \text{diam}(G),$$

where $\text{diam}(G)$ denotes the diameter of G , i.e., the number of edges in a longest path in G . It would be interesting (but probably difficult) to characterize graphs for which this equality holds.

Suppose C is a cutset of vertices, and let A and B be subsets of V separated by the removal of C . Then

$$(2) \quad rt(G) \geq \frac{2}{|C|} \min(|A|, |B|).$$

This follows by considering the permutation π that maps all pebbles starting in $|A|$ into $|B|$ (where we assume, without loss of generality, that $|A| \leq |B|$). All pebbles in A (i.e., those p_i with $p_i(0) \in A$) must pass through some vertex v of C , and it takes two steps for p_i to pass through v : one to move it from A onto v , and one to move it from v into B (which exchanges it with some pebble from B).

Almost the same argument applies if C is a cutset of edges of G , giving the following similar bound:

$$rt(G) \geq \frac{2}{|C|} \min(|A|, |B|) - 1.$$

This is tight for paths of even length.

Let $\mu(G)$ denote the size of a maximum matching in G . For a routing assignment π , define $D(G, \pi)$ by

$$D(G, \pi) := \sum_v d_G(v, \pi(v)),$$

where d_G is the usual (path-) metric on G . Then, setting

$$D(G) := \max_{\pi} D(G, \pi),$$

we have the bound

$$rt(G) \geq \frac{D(G)}{2\mu(G)}.$$

This can be seen by noting that $D(G)$ can only be decreased by at most $2\mu(G)$ at each step.

Since for any spanning subgraph H of G we have

$$rt(G) \leq rt(H),$$

then $rt(G)$ is bounded above by $rt(T)$ for any spanning subtree of G . For any graph G on n vertices, this last quantity is less than $3n$, by Theorem 1, below. We next consider the routing number of trees.

3. Trees. Let $T(n)$ denote some arbitrary fixed tree on n vertices. The following result gives a reasonably good upper bound on $rt(T(n))$.

THEOREM 1.

$$(3) \quad rt(T(n)) < 3n.$$

Proof. We need the following simple and known fact, which can be easily proved (by induction, for example).

Fact. For any tree T on n vertices, there always exists a vertex z of T (see Fig. 1) such that each subtree T_i formed by removing z (and all incident edges) satisfies

$$(4) \quad |T_i| \leq n/2.$$




FIG. 1. Decomposing a tree T .

The proof of Theorem 1 is by induction on $n = |T|$. Let us apply (4) and let T' denote any one of the subtrees T_i . Consider a pebble $p = p_v(0)$ initially placed on a vertex v of T' . Let us call p *proper* if the destination of p under the routing assignment π belongs to T' ; otherwise, call p *improper*. For the special vertex z (the “root”), the pebble $p_z(0)$ is classified as improper.

Our first objective is to move all improper pebbles in (each) T' toward z' , the vertex of T' adjacent to z , so that the vertices they occupy form a subtree T'' of T' containing z' .

CLAIM. *The subtree T'' in T' can be formed in at most $|T'|$ steps.*

Proof of Claim. Let $z' = v_1, v_2, \dots, v_m$ be the vertices on some path M in T' . After the i th step in the process (i.e., after “time i ”), there is a certain distribution of pebbles on M . We let $p(v, i)$ denote the pebble occupying vertex v at time i . More generally, we use the index i to denote the value of a parameter at time i . In

particular, let $I(i)$ denote the set of *improper* pebbles on M at time i that are further from z' than *some* proper pebble on M (where distance on T is measured by the usual path metric, i.e., the number of edges in the unique path connecting two vertices). Let $x(i)$ denote the set of all improper pebbles in T' that are not in the path M . Also, let $P(i)$ denote the set of *proper* pebbles on M that are closer to z' than some improper pebbles on M . Further, let $C(i)$ denote the set of proper pebbles p on M that are adjacent to an improper pebble on M further from z . Finally, define the function $\phi(i)$, called the *potential*, by

$$(5) \quad \phi(i) := |I(i)| + |P(i)| + \min\{|P(i)|, |x(i)|\} - |C(i)|.$$

For example, for the distribution (on the path M) shown in Fig. 2 (where \bullet denotes a proper pebble, and \circ denotes an improper pebble) we have: $|I| = 6$, $|P| = 5$, $|C| = 3$, and (assuming $|x| \geq 5$) $\phi = 13$.




FIG. 2. Pebbles in a path.

The algorithm we employ for reaching the desired state is simply a greedy algorithm: Whenever we can interchange an improper and proper pebble so as to bring the improper pebble closer to z' , we do it. More specifically, at each step we choose a maximal set of disjoint pairs of this type, and perform the interchanges. We now argue that if we have not yet reached the desired state (i.e., the set of all improper pebbles in T' does not yet span a subtree of T' containing z'), then the potential $\phi(i)$ (computed for some specific path M to be chosen later) must decrease at the next step.

To see this, observe that since our greedy algorithm must eventually terminate, we can find some improper pebble \hat{p} that is moved during the last step. Consider the path $M = (z' = v_1, \dots, v_m)$, where v_m is the location of \hat{p} at time 0, i.e., $p(v_m, 0) = \hat{p}$. By the definition of our algorithm, no improper pebble is ever moved off of M . On the other hand, it is quite possible that new improper pebbles are moved onto M . Let us denote the pebble distribution on M in terms of alternating blocks of improper and proper pebbles (see Fig. 3). P_j denotes the j th block of proper pebbles (with size




FIG. 3. Pebbles on M .

$|P_j|$) and I_j denotes the j th block of improper pebbles. Each P_j and I_j , $1 \leq j \leq r$, is nonempty (although I_0 may be empty). By definition, $\phi(i)$ depends only on $x(i)$, P_j and I_j , $1 \leq j \leq r$, and is given by

$$\phi(i) = \sum_{j=1}^r |I_j| + \sum_{j=1}^r |P_j| - r + \min \left(\sum_{j=1}^r |P_j|, |x(i)| \right).$$

Now, when we go to time $i + 1$, various changes in M can occur. To begin, the last (i.e., right-most) proper pebble in each P_j is replaced by some improper pebble, either the *first* pebble in I_j or some other improper pebble from outside of M . Observe that

if $|I(i)|$ increases during a step, then both $|P(i)|$ and $|x(i)|$ must decrease by at least the same amount. By keeping track of all the possible changes that can occur at the next step, it is not difficult (though it is somewhat tedious) to verify that in all cases, $\phi(i + 1) \leq \phi(i) - 1$. We omit the somewhat lengthy details. Since the potential can never exceed, by definition, the number of vertices in T' , this completes the proof of the claim. \square

The next step in the proof of Theorem 1 is to move each component's improper pebbles to their correct components. With T_1, \dots, T_r denoting the subtrees formed by the removal of the root z , let $\bar{I}(T_j)$ denote the set of improper pebbles in T_j , and let $\bar{P}(T_j)$ denote the set of proper pebbles in T_j . It can easily be shown that using at most three steps, two improper pebbles can be moved to their correct destination components. In fact, if t denotes the largest $|\bar{I}(T_j)|$, then we can move at least $2t$ improper pebbles to their correct destination components in at most $2t + 1$ steps.

Following this procedure, we can guarantee that all pebbles are in their correct components (and z is occupied by its proper pebble) in at most

$$\frac{3}{2} \left(\sum_j |\bar{I}(T_j)| - 2t \right) + 2t + 1$$

steps.

Note that by the claim, T'' can be formed (in T_j) in at most $|T_j|$ steps.

Now, since by induction each T_j can now be routed in fewer than $3|T_j|$ steps, then they can all be routed (in parallel) in fewer than $3 \max |T_j|$ steps. Thus, T can be routed in less than

$$\max_j |T_j| + \frac{3}{2} \left(\sum_j |\bar{I}(T_j)| \right) - t + 1 + 3 \max |T_j|$$

steps. However, $\max_j |T_j| - t$ does not exceed the number of proper pebbles in the largest subtree. Let y denote this number of proper pebbles. Then, clearly

$$\sum_j |\bar{I}(T_j)| \leq n - 1 - y,$$

and hence the previous quantity is at most

$$y + \frac{3}{2}(n - 1 - y) + 1 + 3 \max |T_j| < \frac{3}{2}n + 3(n/2) = 3n.$$

This completes the induction step; since (3) holds for $n = 2$, Theorem 1 is proved. \square

The bound in Theorem 1 can perhaps be improved. For example, it seems clear that one should *not* wait to start moving pebbles across z and routing within T_i 's until all improper pebbles in each T_i have been moved close to z_i (i.e., all of these steps can be made in parallel). In fact, the correct value of the constant may be half as large, as suggested by the following conjecture.

Conjecture. For any tree T_n on n vertices,

$$(6) \quad rt(T_n) \leq \left\lfloor \frac{3(n-1)}{2} \right\rfloor.$$

Furthermore, we suspect that the equality can only be achieved when the tree is the star S_n on n vertices.

The fact that equality in (6) holds for S_n was pointed out to us by Goddard [7]. This can be seen as follows.

Denote the center of the star S_n by x . Each cycle $(a_1 \dots a_r)$ of $r \geq 2$ vertices not containing x requires exactly $r + 1$ steps to achieve the desired routing. If $(a_1 \dots a_r)$ contains x , then exactly r steps are required. Therefore

$$rt(S_n, \pi) = n - 1 + t,$$

where t denotes the number of cycles of π of size at least 2 that do not contain x . Since $t \leq \lfloor \frac{n-1}{2} \rfloor$, we obtain the desired bound.

For the case where T_n is a path P_n on n vertices, our routing problem reduces to a well-studied problem in parallel sorting networks (see [9] for a comprehensive survey). In this case, it can be shown that $rt(P_n) = n$. 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, \dots, e_{n-1} and only making interchanges with *even* edges e_{2k} on *even* steps and *odd* edges e_{2k+1} on *odd* steps.

4. Complete graphs. Let K_n denote the complete graph on n vertices. In this case, because K_n is so highly connected, the routing number K_n is as small as we could hope.

THEOREM 2. *For the complete graph K_n on $n \geq 3$ vertices,*

$$(7) \quad rt(K_n) = 2.$$

Proof. To see that $rt(K_n) \geq 2$, it is enough to consider the permutation $\pi = (abc)$ consisting of a 3-cycle on K_n . It is clear that such a π cannot be achieved in a single step.

To show that $rt(K_n) \leq 2$, it suffices to show that any *cyclic* permutation can be achieved in two steps, since any permutation π can be factored into disjoint cycles, which can then all be routed in parallel. Thus, let π_m denote the cyclic permutation on $\{1, 2, \dots, m\}$ given by

$$\begin{aligned} \pi(i) &= i - 1, \quad 1 < i \leq m, \\ \pi(1) &= m. \end{aligned}$$

Consider the following two routing steps:

$$S_1 : (1, m + 1 - 1)(2, m + 1 - 2) \cdots (i, m + 1 - i) \cdots,$$

and

$$S_2 : (1, m - 1)(2, m - 2) \cdots (j, m - j) \cdots.$$

We confirm that the composition $S_1 \circ S_2$ sends the following:

$$\begin{aligned} i \rightarrow m + 1 - i \rightarrow m - (m + 1 - i) &= i - 1, \quad i \neq 1, \\ 1 &\rightarrow m. \end{aligned}$$

This map achieves the desired permutation in two steps. Consequently, $rt(K_n) \leq 2$, and the theorem is proved. \square

The following result is due to Goddard [7].

THEOREM 3. *For the complete bipartite graph $K_{n,n}$ with $n \geq 3$,*

$$(8) \quad rt(K_{n,n}) = 4.$$

Proof. Suppose $K_{n,n}$ has vertex sets A and B , where the edges are all between A and B . To see that $rt(K_{n,n}) \geq 4$, we consider the permutation $\pi = (a_1 a_2 a_3)$, where a_i 's are in A . It is not difficult to show that π cannot be achieved in three steps.

A pebble is said to be an A -pebble if its destination is in A . Otherwise, it is called a B -pebble. In at most one step, we can move all A -pebbles to B and all B -pebbles to A . To prove that $rt(K_{n,n}) \leq 4$, it suffices to show that any cyclic permutation $\pi = (1, 2, \dots, 2m)$ can be achieved in the following three routing steps:

$$\begin{aligned}
 S_1 &: (1, 2)(3, 4) \cdots \left(2 \left\lfloor \frac{m}{2} \right\rfloor - 1, 2 \left\lfloor \frac{m}{2} \right\rfloor\right) \left(2 \left\lfloor \frac{m}{2} \right\rfloor + 2, 2 \left\lfloor \frac{m}{2} \right\rfloor + 3\right) \cdots (2m - 2, 2m - 1), \\
 S_2 &: (1, 2m)(3, 2m - 2) \cdots \left(2 \left\lfloor \frac{m}{2} \right\rfloor - 1, 2 \left\lfloor \frac{m}{2} \right\rfloor + 2\right), \\
 S_3 &: (3, 2m) \cdots \left(2 \left\lfloor \frac{m}{2} \right\rfloor + 1, 2 \left\lfloor \frac{m}{2} \right\rfloor + 2\right).
 \end{aligned}$$

This proves the theorem. \square

More generally, it is not difficult to show that for a general complete bipartite graph $K_{m,n}$, $m \leq n$, we have

$$(9) \quad rt(K_{m,n}) \leq 2 \left\lfloor \frac{n}{m} \right\rfloor + 2,$$

since in at most two steps, m pebbles (in fact, B -pebbles, as defined above,) can be routed to their destinations.

5. Cartesian products. For graphs $G = (V, E)$, $G' = (V', E')$, we define the Cartesian product graph $G \times G'$ to be the graph with vertex set $V \times V' = \{(v, v') \mid v \in V, v' \in V'\}$ and with $(u, u')(v, v')$ an edge of $G \times G'$ if and only if either $u = v$, $u'v' \in E'$, or $u' = v'$, $uv \in E$. Thus, the n -cube Q^n is just the Cartesian product of K_2 with itself n times.

The following theorem can be traced back to the early work of Beněš [5]. It was also proved by Baumslag and Annexstein [4].

THEOREM 4. *We have*

$$(10) \quad rt(G \times G') \leq 2rt(G) + rt(G').$$

Note that since $G \times G'$ and $G' \times G$ are isomorphic graphs, then (10) can be written in the symmetric form

$$(10') \quad rt(G \times G') \leq \min\{2rt(G) + rt(G'), 2rt(G') + rt(G)\}.$$

We briefly describe the proof of Theorem 4 here. We can picture $G \times G'$ as an array $V \times V'$, with each row spanning a copy of G' and each column spanning a copy of G as illustrated in Fig. 4. To route in $G \times G'$, we:

- (i) route in columns (copies of G); then
- (ii) route in rows (copies of G'); then
- (iii) route in columns (copies of G).

Let π be the desired routing permutation we are trying to achieve. Each pebble p has some destination $(\sigma(p), \sigma'(p))$, where $\sigma(p) \in V$, $\sigma'(p) \in V'$. Let us first classify the pebbles according to their second coordinates. Since π is a permutation on $V \times V'$, for each $v' \in V'$, there are exactly $|V|$ pebbles with $\sigma'(p) = v'$. Hence, by the well-known marriage theorem of Hall (see, for example, [11]), we can select a set of distinct representatives from the columns, i.e., one pebble from each column so that their second coordinates are all distinct. Furthermore, we can now repeat this procedure




FIG. 4. $G \times G'$.

(again by Hall's theorem) to obtain another set of distinct representatives, and so on. At the end, we see that we can in fact arrange the pebbles in each column so that the pebbles in each row of the rearranged columns all have distinct values of σ' . By hypothesis, this rearrangement can be accomplished in at most $rt(G)$ (parallel) steps.

Next, we rearrange the pebbles in each row (i.e., copy of G') so that pebbles p in the column indexed by $v' \in V'$ have $\sigma'(p) = v'$. This can be done (by hypothesis) in $rt(G')$ more steps, and guarantees, when completed, that the pebbles in each column have distinct values of σ .

The final step, permuting each column (copy of G), can be done in $rt(G)$ more steps. Thus, the whole process requires at most $2rt(G) + rt(G')$ steps.

COROLLARY 1. For the n -cube Q^n ,

$$rt(Q^n) \leq 2n - 1.$$

COROLLARY 2. For the m by n grid graph $P_m \times P_n$, $m \leq n$,

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

Remarks. Routing on the n -cube Q^n is a very natural question in view of the popular use the n -cube structure for models of parallel computation and communication. Indeed, it was this context (through the work of Ramras [10]) that first motivated our considerations of these questions.

Corollary 1 is well known in the literature. The exact value of $rt(Q^n)$ is still unknown. It is easy to see that $rt(Q^n) \geq n$, since $\text{diam}(Q^n) = n$. The permutations shown in Fig. 5 can be checked to show that $rt(Q^n) \geq n + 1$ for $n = 2, 3$. It is reasonable to conjecture that we always have $rt(Q^n) \geq n + 1$ for $n \geq 2$. Certainly $rt(Q^n) \sim \alpha n$ for some $\alpha \in [1, 2]$. Again, we suspect that the correct value of α is closer to 1 than to 2, but this seems difficult to prove.

6. Flow problems on graphs. Ordinarily, one might expect that $rt(G \times G)$ is substantially larger than $rt(G)$, e.g., as large as $2rt(G)$. However, this is not always the case, as the following result shows.

Let G_n denote the graph consisting of two copies of K_n joined by an edge e (see Fig. 6). It is easy to see that

$$rt(G_n) = 2n + O(1).$$




FIG. 5. Bad permutations on Q^n , $n \leq 3$.




FIG. 6.

It turns out that $rt(G_n \times G_n)$ is not much larger.

THEOREM 5. We have

$$(11) \quad rt(G_n \times G_n) = (1 + o(1))2n.$$

Proof. We can view $G_n \times G_n$ as consisting of 4 copies of $K_n \times K_n$ joined to each other by n parallel edges to form a 4-cycle (see Fig. 7). Within each $V_i = K_n \times K_n$, the two sets of n vertices incident to “crossing” edges have exactly one common vertex. Initially, each V_i has pebbles with destinations lying in the other various V_j ’s. We group each of V_i ’s pebbles into sets S_{ij} of size n , according to their destinations (so that all pebbles in a set S_{ij} have the same V_j destination; there may be mixed groups of n left over). We move the S_{ij} ’s as a unit, in that all pebbles in each S_{ij} cross from one V_k to another V_l at the same time. Further, we restrict our routing algorithm so that crossing moves are only made at even times. Of course, permutations within a V_i can occur at all times. Since $rt(K_n \times K_n) \leq 6$ (by Theorems 2 and 4), it is not difficult to see that if we have m S_{ij} ’s in V_k that should cross over to V_l (which is adjacent to it), then this can be done in $2m + O(1)$ steps. Thus, our problem can be reduced




FIG. 7. $G_n \times G_n$.

to the following continuous flow problem on the 4-cycle C_4 . We are given one unit of “mass” on each vertex v of C_4 (where we have rescaled our sets of n^2 pebbles at each vertex of C_4 to have total mass 1). Thus, mass is required to “flow” along the edges of C_4 in order to satisfy a 4×4 doubly stochastic *circulation matrix* $C = (C(u, v))$ where, for vertices u, v of C_4 , $C(u, v)$ denotes the amount of mass initially at u that must end up at v . Since C is assumed to be doubly stochastic, then $C(u, v) \geq 0$ and

$$\sum_v C(u, v) = 1 = \sum_u C(u, v).$$

Therefore, each vertex of C_4 also ends up with a total of one unit of mass (hence, our use of the terminology “circulation”).

In general, a C -circulation φ on $G = (V, E)$ is a set of assignments $\varphi_{uv} : E \rightarrow \mathbb{R}^+$, $u, v \in V$, such that, for all u, v ,

$$\sum_{ux \in E} \varphi_{uv}(ux) = C(u, v) = \sum_{yv \in E} \varphi_{uv}(yv),$$

while for any $w \neq u, v$,

$$\sum_{sw \in E} \varphi_{uv}(sw) = \sum_{wt \in E} \varphi_{uv}(wt).$$

Intuitively, these equations specify that for each pair $u, v \in V$, $C(u, v)$ units of mass flow from u to v . The norm of φ , denoted by $\|\varphi\|$, is defined to be the maximum amount of mass

$$\varphi(e) = \sum_{u,v} \varphi_{uv}(e)$$

assigned to any edge e of E , where we distinguish between $e = ij \in E$ and the edge $-e = ji$ with the reverse orientation. We say that φ is *balanced* if $\varphi(e) = \varphi(-e)$ for all edges of G .

A little reflection shows that we have proved Theorem 5 if we establish the following lemma.

LEMMA 1. For all circulation matrices C on C_4 , there always exists some balanced C -circulation φ with $\|\varphi\| \leq 1$.

Proof of Lemma 1. Consider a spanning tree T on C_4 . There are four such trees; these are all paths of length 3. (See Fig. 8.) Note that there is a unique balanced F -flow φ_T on T . The amount $\varphi_T(e)$ that φ_T assigns to e is just

$$\varphi_T(e) = \sum_{\substack{u \in A \\ v \in B}} C(u, v),$$

where A and B are the components of T formed by the removal of $e = ij$, and $i \in A, j \in B$. We form our desired F -flow φ as a convex combination of φ_T 's, as





FIG. 8.

T ranges over the spanning subtrees of G . This guarantees that φ is a C -circulation and is balanced (since each φ_T is). In fact, we take the simplest possible convex combination, namely,

$$\varphi := \frac{1}{4} \sum_T \varphi_T,$$

where T ranges over all four spanning trees of C_4 . To compute $\|\varphi\|$, we need to bound the value of $\varphi(e)$ for each edge e . For any edge e (by symmetry), the mass $\varphi(e)$ assigned to e is equal to that assigned to the edge in the figure, which is the following:

$$\begin{aligned} \varphi(e) &= 0 && \text{from (a),} \\ &+ \frac{1}{4}(f(1, 2) + f(4, 2) + f(3, 2)) && \text{from (b),} \\ &+ \frac{1}{4}(f(4, 2) + f(4, 3) + f(1, 2) + f(1, 3)) && \text{from (c),} \\ &+ \frac{1}{4}(f(1, 2) + f(1, 3) + f(1, 4)) && \text{from (d).} \end{aligned}$$




FIG. 9. *Spanning trees of C_4 .*

However, observe that in Fig. 9, the maximum value that can flow from A to B is just $\min(|A|, |B|)$. Consequently,

$$\begin{aligned} f(1, 2) + f(4, 2) + f(3, 2) &\leq 1, \\ f(4, 2) + f(4, 3) + f(1, 2) + f(1, 3) &\leq 2, \\ f(1, 2) + f(1, 3) + f(1, 4) &\leq 1, \end{aligned}$$

and so,

$$\varphi(e) \leq 1.$$

Since e was arbitrary, $\|\varphi\| \leq 1$ and the lemma is proved. \square

Now, because φ is balanced, we can reinterpret it as pebble movements, where small time delays (due to the nonuniform S_{ij} , or bounded movement within the V_i 's) are negligible as $n \rightarrow \infty$. We can then conclude that

$$rt(G_n \times G_n) = (1 + o(1))2n,$$

as claimed. \square

The same argument applies, with the same conclusion, for the k -fold product

$$G_n^k = \overbrace{G_n \times \cdots \times G_n}^k,$$

provided we prove the corresponding flow result on Q^k , the k -cube, which actually is of interest in its own right. We do this with the following theorem.

THEOREM 6. *Let F be a doubly stochastic circulation matrix on Q^k . Then there always exists a balanced C -circulation φ on Q^k with $\|\varphi\| \leq 1$.*

Proof. We follow the same strategy as in the case of $C_4 = Q^2$, and build φ as a (uniform) convex combination of tree flows. Define the spanning tree T_k on Q^k recursively as follows:

- (i) T_1 is just an edge, which is all of Q^1 ;
- (ii) T_k is formed by adding the edge $e_k = \{(00 \dots 0), (10 \dots 0)\}$ to join the two copies of T_{k-1} in the corresponding two copies of Q^{k-1} that make up Q^k , namely, $\{\bar{x} = (x_1, \dots, x_k) \mid x_1 = 0\}$ and $\{\bar{x} = (x_1, \dots, x_k) \in Q^k \mid x_1 = 1\}$.

We observe that for any edge e of T_k , we can always find a minimum-sized component $A(e)$ of $T_k - \{e\}$ that *does not contain the origin* $(00 \dots 0)$. Hence we conclude (by

induction) that

$$\begin{aligned}
 g(k) &:= \sum_{e \in T} |A(e)| \\
 &= \sum_{e=e_k} |A(e)| + \sum_{e \neq e_k} |A(e)| \\
 &= 2^{k-1} + 2 \cdot g(k-1) \\
 &= 2^{k-1} + 2 \cdot (k-1) \cdot 2^{k-2} = k \cdot 2^{k-1},
 \end{aligned}
 \tag{12}$$

where we check that $g(1) = 1$ satisfies (12) to start the induction.

Next, we construct the family of trees that is used to form φ . Let $\text{Aut}(Q^k)$ denote the automorphism group of Q^k . It is easy to see that $|\text{Aut}(Q^k)| = 2^k \cdot k!$ and that for any $h \in \text{Aut}(Q^k)$, $h(T_k)$ is also a spanning tree of Q^k . Furthermore, for each edge e of Q^k and each edge e' of T_k , there are exactly $2(k-1)!$ choices of $h \in \text{Aut}(Q^k)$ that map e onto e' (as undirected edges; this accounts for the factor of 2. Of course, this does not depend on e' being in T_k .) Define

$$\varphi = \frac{1}{2^k k!} \sum_T \varphi_T,$$

where T ranges over all $2^k \cdot k!$ trees $h(T_k)$, $h \in \text{Aut}(Q^k)$. Note that this is just what we did for C_4 . To bound $\|\varphi\|$, we compute for any e ,

$$\varphi(e) \leq \frac{1}{2^k \cdot k!} \sum_{h \in \text{Aut}(Q^k)} \min(|A_h(e)|, |B_h(e)|) = \frac{1}{2^k k!} \cdot 2(k-1)!g(k) = 1.$$

This completes the proof of Theorem 6. \square

COROLLARY 3. *For fixed k , if G^k denotes the k -fold Cartesian product $G \times \dots \times G$ of the graph G shown in Fig. 6, then*

$$\text{rt}(G^k) = (1 + o(1))2n.$$

Let us define $\text{circ}(G)$, the *circulation index* of a (connected) graph G , by

$$\text{circ}(G) := \sup_C \inf_{\varphi} \|\varphi\|,
 \tag{13}$$

where φ ranges over all balanced C -circulations on G , and C ranges over all doubly stochastic circulation matrices C for G . A trivial lower bound for $\text{circ}(G)$ is the *resistance* of G , defined by

$$\text{res}(G) := \max_X \frac{1}{|X|} \min(|A(X)|, |B(X)|),
 \tag{14}$$

where X ranges over all *cutsets* of G (i.e., minimal sets of edges whose removal disconnects G), and $A(X)$ and $B(X)$ are the connected components formed by removing X . The inequality

$$\text{circ}(G) \geq \text{res}(G)
 \tag{15}$$

follows by considering the circulation matrix that sends all $|A(X)|$ units of mass into $B(X)$, for an extremal cutset X , where we assume $|A(X)| \leq |B(X)|$. It is interesting

to note that (15) holds with *equality* for Q^k . This is not true in general, as can be shown by considering bounded degree expander graphs G^* on n vertices. In that case, we can have

$$\text{res}(G) = O(1) \quad \text{and} \quad \text{circ}(G) > c \log n.$$

Other classes of graphs G for which equality holds in (15) prove interesting. The technique above can be easily used to show that the set of even cycles is such a class.

It seems likely that the space of balanced C -circulations on any graph is spanned by (convex combinations of) the *tree* circulations on G , i.e., the C -circulations in the spanning trees of G .

7. Eigenvalues, random walks and routing. Here we consider d -regular graphs for which all the eigenvalues of the adjacency matrix besides the trivial one have a small absolute value. Let us call a graph G an (n, d, λ) -graph, if it is a d -regular graph on n vertices and the absolute value of every eigenvalue of its adjacency matrix besides the trivial one is at most λ . If λ is small with respect to d , then a random walk on such a graph starting from any vertex converges quickly to the uniform distribution on its vertices. We use this property to derive the following theorem, the proof of which is given later.

THEOREM 7. *Let $G = (V, E)$ be an (n, d, λ) -graph and let σ denote a permutation. Then*

$$rt(G, \sigma) \leq O\left(\frac{d^2}{(d - \lambda)^2} \log^2 n\right).$$

Note that in §2, it is shown that $rt(G)$ is lower bounded by the diameter of G and therefore the routing number of a d -regular graph is at least $\log n / \log(d - 1)$ and at most

$$O\left(\frac{d^2}{(d - \lambda)^2} \log^2 n\right).$$

Now define the *expansion coefficient* α of G to be the minimum, over all subsets X of at most half the vertices of G , of the ratio $|N(X) - X|/|X|$, where $N(X)$ is the set of all neighbors of X in G . From §2, we know that the routing number is bounded below by $2/\alpha$. As an immediate corollary of Theorem 7, up to a polylogarithmic factor, $rt(G)$ is bounded above by a polynomial in $1/\alpha$ for any regular graph with polylogarithmic degrees.

COROLLARY 4. *If $G = (V, E)$ is a d -regular graph on n vertices with expansion coefficient α , then*

$$rt(G) \leq O\left(\frac{d^2}{\alpha^4} \log^2 n\right).$$

Proof. The main result of [1] states that if α is the expansion coefficient of a d -regular graph, then the second largest eigenvalue of its adjacency matrix is at most $d - \frac{\alpha^2}{4 + 2\alpha^2}$. Suppose, first, that this is an upper bound for the absolute value of every negative eigenvalue as well. Then, by Theorem 7, $rt(G) = O(l^2)$ for $l = O(\frac{d}{\alpha^2} \log n)$. If there are negative eigenvalues of large absolute value, we first add d loops in every vertex and apply the result to the new graph. This completes the proof. \square

In a similar way, we define the *edge expansion coefficient* β of G to be the minimum, over all subsets X of at most half the vertices of G , of the ratio $|\Gamma(X)|/|X|$,

where $\Gamma(X)$ is the set of edges of G leaving X (i.e., with exactly one endpoint in X). We remark that the inverse of the edge expansion coefficient is exactly the resistance of G . From §2, we know that the routing number is bounded below by $2/\beta - 1$. As a corollary of Theorem 7, up to a polylogarithmic factor, $rt(G)$ is bounded above by a polynomial in $1/\beta$ for any regular graph with polylogarithmic degrees.

COROLLARY 5. *If $G = (V, E)$ is a d -regular graph on n vertices with edge expansion coefficient β , then*

$$rt(G) \leq O\left(\frac{d^4}{\beta^4} \log^2 n\right).$$

Proof. The proof follows from the well-known fact (see [8]) that if β is the edge expansion coefficient of a d -regular graph, then the second largest eigenvalue of its adjacency matrix is at most $d - \frac{\beta^2}{2d}$. \square

Note that by the above two corollaries, $rt(G) \leq O(\log^2 n)$ for any bounded degree expander on n vertices (i.e., any regular bounded degree graph on n vertices with expansion coefficient or edge expansion coefficient bounded away from 0).

Many interconnection networks studied in the literature are, in fact, Cayley graphs. A simple corollary of the above theorem implies that the routing number of a Cayley graph is intimately related to its diameter.

COROLLARY 6. *For any Cayley graph G of a group of n elements with a polylogarithmic (in n) number of generators, the diameter of G is polylogarithmic, if and only if the routing number $rt(G)$ is polylogarithmic.*

Proof. As shown in [3], a Cayley graph of polylogarithmic diameter has an inverse polylogarithmic expansion coefficient; hence the result follows from Corollary 4. \square

The proof of Theorem 7 follows from the following lemmas. Our first lemma holds for any d -regular graph G . A *random walk* of length l starting at a vertex v of G is a randomly chosen sequence $v = v_0, v_1, \dots, v_l$, where each v_{i+1} is chosen, randomly and independently, among the neighbors of v_i , ($0 \leq i < l$). We say that the walk *visits* v_i at time i . We make no attempt to optimize the constants here and in what follows.

LEMMA 2. *Let $G = (V, E)$ be a d -regular graph on n vertices, and suppose that $l \geq \log n$. For any $v \in V$ independently, let $P(v)$ denote a random walk of length l starting at v . Let $I(v)$ denote the total number of other walks $P(u)$, such that there exists a vertex x and two indices $0 \leq i, j \leq l$, $|i - j| < 5$ so that $P(v)$ visits x at time i and $P(u)$ visits x at time j . Then, almost surely (i.e., with probability that tends to 1 as n tends to infinity), there is no vertex v such that $I(v) > 100(l + 1)$.*

Proof. Let A be the normalized adjacency matrix of G , i.e., the matrix $A = (a_{uv})_{u,v \in V}$ defined by $a_{uv} = l(u, v)/d$, where $l(u, v)$ is the number of edges between u and v . The probability that the random walk $P(u)$ visits x at time i is precisely $e(x)^t A^i e(u)$, where $e(y)$ is the unit vector having 1 in coordinate y and 0 in any other coordinate. Given the random walk $P(v)$ and a value of i , $0 \leq i \leq l$, there is a unique vertex $x = x(v, i)$ in which $P(v)$ visits at time i . For any given $u \neq v$, the conditional probability that for some j satisfying $|i - j| < 5$, the walk $P(u)$ visits x at time j is thus at most $e(x)^t \sum_{j:|j-i|<5} A^j e(u)$. The probability $p(v, u)$ that there exists some vertex x and two indices $0 \leq i, j \leq l$, $|i - j| < 5$ follows, so that $P(v)$ visits x at time i and $P(u)$ visits x at time j can be bounded by

$$p(v, u) \leq \sum_{i=0}^l e(x(v, i))^t \sum_{j:|j-i|<5} A^j e(u).$$

By summing over all possible starting points u (including v itself, where this last summand corresponds to adding another independent random walk starting at v , an addition which may only increase the expectation of $I(v)$), we conclude that the expectation of $I(v)$ is at most

$$\sum_{u \in V} p(v, u) \leq \sum_{i=0}^l e(x(v, i))^t \sum_{j:|j-i|<5} A^j e,$$

where e is the all 1 vector. Since e is an eigenvector of A with eigenvalue 1, the last expression can be computed precisely by showing that it is strictly less than $10(l + 1)$. We have thus shown that for each fixed v , the expectation of the random variable $I(v)$ is strictly less than $10(l + 1)$. Observe that this random variable is a sum of $n - 1$ independent indicator random variables whose expectations are the quantities $p(v, u)$. It thus follows easily from the known estimates for large deviations of sums of independent indicator random variables (see, e.g., [2, Thm. A.12, p. 237]), that for each fixed v , the probability that $I(v)$ exceeds, say, $100(l + 1)$ is at most

$$(e^9/10^{10})^{10(l+1)} \ll 1/n^2.$$

(A similar estimate can in fact be proved directly. Given a set of m independent events, with the probability of the i th event being p_i , suppose that $\sum p_i \leq r$. Then, the probability that at least s events occur can be bounded by

$$\sum_{S \subset \{1, \dots, m\}, |S|=s} \prod_{i \in S} p_i \leq \frac{1}{s!} \left(\sum p_i \right)^s \leq (re/s)^s.$$

In our case, we have $r = 10(l + 1)$ and $s = 10r$.)

Since there are only n vertices v , the probability that there is a vertex v with $I(v) > 100(l + 1)$ is (much) smaller than $1/n$, completing the proof. \square

LEMMA 3. *Let $G = (V, E)$ be an $(n, d, (1 - \epsilon)d)$ -graph and let σ be a permutation of order two of V (i.e., a product of pairwise disjoint transpositions). Put $l = \frac{10}{\epsilon} \log n$. Then there is a set of $n/2$ walks $P(v) = P(\sigma(v))$, $v \in V$, of length $2l$ each, where $P(v)$ connects v and $\sigma(v)$ such that the following holds. Let $I(v)$ denote the total number of other walks $P(u)$ such that there exists a vertex x and two indices $0 \leq i, j \leq l$, $|i - j| < 5$, so that $P(v)$ visits x at time i , and $P(u)$ visits x at time j or at time $2l - j$. Then $I(v) \leq 400(l + 1)$ for all v .*

Proof. Let $P(v)$ be a random walk of length $2l$ between v and $\sigma(v)$. As shown in [6] (using an argument similar to the one used previously in [12]), we may assume that each walk $P(v)$ consists of two random walks of length l each, one starting from v and one from $\sigma(v)$. The reason for this is that by our eigenvalue condition, a random walk of length l is almost uniformly distributed on the vertices of G , and hence one may view the walk $P(v)$ as being chosen by first choosing its middle point (according to a uniform distribution) and then by choosing its two halves. For more details, see [6]. The result thus follows from Lemma 2. \square

Proof of Theorem 7. Let $G = (V, E)$ be an (n, d, λ) -graph. It suffices to consider a permutation σ of order two of V (i.e., a product of pairwise disjoint transpositions), since any permutation is a product of at most two such permutations (as proved in Theorem 2). We set $\epsilon = 1 - \frac{\lambda}{d}$ and $l = \frac{10}{\epsilon} \log n$. We want to show that $rt(G, \sigma) < O(l^2)$. Let $P(v)$ be a system of walks of length $2l$ satisfying the assumption of the previous corollary. Let H be the graph whose vertices are the walks $P(v)$ in which $P(u)$ and $P(v)$ are adjacent, if there exists a vertex x and two indices $0 \leq i, j \leq l$, $|i - j| < 5$,

so that $P(v)$ visits x at time i , and $P(u)$ 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 we can split all our paths $P(v)$ 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 $2l + 1$ steps, where the step numbers i and $2l + 2 - i$ correspond to flipping the pebbles along edge numbers i and $2l + 1 - i$ in each of the paths in the set for all $i < l$. Step number l flips edge l , and step $l + 1$ flips edge $l + 1$. We can check that by the end of these $2l + 1$ 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 $2l + 1$ steps.) By repeating the above for all the path-classes, the result follows. \square

8. The route covering number of a graph. We next discuss 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 $i = 1, \dots, n$. 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. It is easy to see that the route covering problem is a special case of C -circulation obtained by choosing C to satisfy $C(u, v) = 1$ if $v = \pi(u)$, and 0 otherwise, for each permutation π , and by insisting on integer valued circulation.

For example, for the n -cube Q^n , the method described in Theorem 4 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 easily be seen that

$$rc(Q^n) \geq \frac{\sum_v \text{dist}(v, \pi(v))}{|E(Q^n)|} = 2.$$

The problem of determining the exact value of $rc(Q^n)$ for general n remains unresolved. Also of interest is a “symmetric” version of the route covering problem especially for Q^n .

An *assignment* for Q^n is a partition of the vertex set of Q^n into subsets of size 2 or less. Is it possible to find edge-disjoint paths joining vertices in the same subset for any assignment of Q^n ?

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

9. Concluding remarks. Numerous unanswered questions remain, some of which we now mention.

(1) Is it true that, for any tree T_n on n vertices,

$$rt(T_n) \leq \frac{3}{2}n + o(n)? \quad \frac{3}{2}n + O(1)?$$

- (2) Is it true that, for the n -cube Q^n ,

$$rt(Q^n) = n + o(n)? \quad n + O(1)?$$

- (3) Is it true that, for every graph G ,

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

- (4) Is it true that, for an expander graph G of bounded degree,

$$rt(G) = O(\log n)?$$

- (5) Characterize graphs G with $\text{circ}(G) = \text{res}(G)$.

- (6) Are the balanced C -circulations on a graph always spanned by the spanning tree C -circulations on the graph?

- (7) What is the computational complexity of determining $rt(G)$?

REFERENCES

- [1] N. ALON, *Eigenvalues and expanders*, *Combinatorica*, 6 (1986), pp. 83–96.
- [2] N. ALON AND J. H. SPENCER, *The Probabilistic Method*, John Wiley, New York, 1991.
- [3] L. BABAI AND M. SZEGEDY, *Local expansion of symmetrical graphs*, *Combin., Probab. and Comput.*, 1 (1991), pp. 1–12.
- [4] M. BAUMSLAG AND F. ANNEXSTEIN, *A unified framework for off-line permutation routing in parallel networks*, *Math. Systems Theory*, 24 (1991), pp. 233–251.
- [5] V. E. BENÉŠ, *Mathematical Theory of Connecting Networks*, Academic Press, New York, 1965.
- [6] A. BRODER, A. FRIEZE, AND E. UPPAL, *Existence and construction of edge disjoint paths on expander graphs*, in *Proc. 24th ACM Sympos. on Theory of Comput.*, Victoria, British Columbia, Canada, May 1992, pp. 140–149.
- [7] W. GODDARD, private communication.
- [8] M. JERRUM AND A. SINCLAIR, *Approximating the permanent*, *SIAM J. Comput.*, 18 (1989), pp. 1149–1178.
- [9] D. E. KNUTH, *The Art of Computer Programming*, Vol. 3, Addison-Wesley, Reading, MA, 1973, p. 241.
- [10] M. RAMRAS, *Routing permutations on a graph*, *Networks*, 23 (1993), pp. 391–398.
- [11] H. J. RYSER, *Combinatorial Mathematics*, Carus Monograph No. 14, Mathematical Association of America, John Wiley, New York, 1963.
- [12] L. VALIANT, *A scheme for fast parallel communication*, *SIAM J. Comput.*, 11 (1982), pp. 350–361.