Routing permutations on graphs via matchings (extended abstract)

N. Alon
Tel Aviv University, Tel Aviv and
Bellcore, Morristown, NJ 07960

F. R. K. Chung
Bellcore, Moristown, NJ 07960 and
Harvard University, Cambridge, MA 02138

R. L. Graham
AT&T Bell Laboratories, Murray Hill, NJ 07974 and
Rutgers University, New Brunswick, NJ 08903

Abstract

We consider a class of routing problems on connected graphs G. Initially, each vertex v of G is occupied by a "pebble" which has a unique destination $\pi(v)$ in G, so that π is a permutation of the vertices of G. It is required to route all the pebbles 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.

The odd-even sorting network shows that in the very special case that G is an n-vertex path, any permutation can be routed in n steps. Here we investigate this routing problem for a variety of graphs G, including trees, complete graphs, hypercubes, Cartesian products of graphs, expander graphs and various Cayley graphs. In addition, we relate this routing problem to certain network flow problems, and to several graph invariants including diameter, eigenvalues and expansion coefficients. Three of our results are the following:

(i) Any permutation can be routed on any n-vertex connected graph in less than $3n$ steps.

(ii) For any d-regular graph on n vertices in which the absolute value of any nontrivial eigenvlaue is at most λ, any permutation can be routed in $O(d^2 \log^2 n/(d-\lambda)^2)$ steps.

(iii) For any Cayley graph of a group of n elements with respect to a polylogarithmic (in n) number of generators, any permutation can be routed in a polylogarithmic number of steps if and only if the diameter of the graph is polylogarithmic.

25th ACM STOC '93-5/93/CA,USA

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. Each pebble p has 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.

Routing problems received a considerable amount of attention (see

[10] for a survey). Our model here differs from the common models that usually allow more flexibility either by allowing queues of pebbles in the vertices or by by allowing the set of edges on which pebbles are moved in a given step to be an arbitrary set of edges and not necessarily a matching. It seems, however, that the model considered here is interesting as its basic building block is a very simple element- an edge that is capable of interchanging its endpoints in a step. Note, also, that the above model for routing has been considered in the study of routing algorithms on the hypercube. In addition, the odd-even sorting network can be viewed as a routing algorithm in this model, restricted to the special case when the graph on which we route is a path. It should be noted that our model enables only the routing of permutations, and that we consider here mainly off-line routing algorithms that are computed by using global information about the routing problem.

We will imagine that the steps during a routing process occur 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 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 $\tau = (u_1v_1)(u_2v_2)\dots(u_rv_r)$ 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.

In this work, we investigate our routing problem on a variety of graphs. These include trees, complete 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.

The rest of this extended abstract is organized as follows. After discussing some simple bounds on the routing numbers of graphs in Section 2 we show in Section 3 that any permutation on any n-vertex connected graph can be routed in less than $3n$ steps. Such a routing can be found efficiently for any given permutation. In Section 4 we consider the routing numbers of complete graphs and those of cartesian products of graphs. The final Section 5 contains upper bounds for the routing numbers of graphs in terms of their expansion properties and their eigenvalues. In particular it is shown that the routing number of any bounded-degree expander on n-vertices (see Section 5 for the precise definition) is at most $O(\log^2 n)$ (and, trivially, at least $\Omega(\log n)$.) Due to space limitations most proofs are omitted, and the only proofs which appear in details are those in Section 5.

Many additional questions regarding our routing problem can be raised; we believe that this model for routing permutations on graphs deserves further attention.

2 Some simple bounds on rt(G)

To begin with, an obvious lower bound on $rt(G)$ is the following:

$$rt(G) \geq diam(G) \qquad (1)$$

where $diam(G)$ denotes the diameter of G, i.e., the number of edges in a longest path in G.

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

Then

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

This follows by considering the permutation π which 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 $3n$, as shown in the next section.

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 3.1

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

The proof is rather lengthy and due to space limitations we only sketch it here. We note, though, that it is much easier to prove the weaker statement that $rt(T(n)) \leq O(n)$ for all n. Here is an outline of the proof of the above theorem.

It is easy and well known that for any tree T on n vertices, there always exists a vertex z of T such that each subtree T_i formed by removing z (and all incident edges) has at most $n/2$ vertices. Our proof proceeds by induction on $n = |T|$. Let us split T by removing a vertex z as above 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 $p_z(0)$ will be classified as improper.)

Our first objective will be to move all improper pebbles in (each) T' towards z', the vertex of T' adjacent to z, so that the vertices they occupy form a subtree T'' of T' containing z'. By defining and analyzing an appropriate weight function we can prove the following:

Claim 1 *The subtree T'' in T' can be formed in at most $|T'|$ steps.*

Moreover, the proof of this claim shows that the tree T'' can be constructed by a simple greedy-type algorithm (which is on-line and runs in linear time).

The next step in the proof is to move each component's improper pebbles to their correct components. It is not too difficult to show that if t denotes the maximum number of improper pebbles in a subtree T_i and if x is the total number of improper pebbles this task can be accomplished in at most $\frac{3}{2}x - t + 1$ steps.

Since by induction each subtree T_j can now be routed in fewer than $3|T_j|$ steps, then they can all be routed (in parallel) in fewer than $3max|T_j|$ steps. Combining these estimates the assertion of the theorem can be deduced. □

We note that the proof easily yields an efficient, nearly linear time algorithm for finding (off line) a routing as above for any given permutation π.

The bound in the last theorem can perhaps be improved, although it is easy to see that the best possible estimate in its statement is at least $\lfloor\frac{3(n-1)}{2}\rfloor$ which may very well be the truth.

For the case that 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 the edges in P_n according to their order 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.

4 Complete graphs and cartesian products

4.1 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 of K_n is as small as one could hope for and it is not difficult to prove the following statement, by exhibiting explicitly the appropriate routing for any given permutation.

Proposition 4.1 *For the complete graph K_n on $n \geq 3$ vertices,*

$$rt(K_n) = 2 \ . \qquad (4)$$

As observed by Wayne Goddard [7], a similar statement can be proved for complete bipartite graphs with equal color classes.

Proposition 4.2 *For the complete bipartite graph $K_{n,n}$ with $n \geq 3$,*

$$rt(K_{n,n}) = 4 \ . \qquad (5)$$

The proof can be given again by a direct construction. We omit the details.

4.2 Cartesian products

For graphs $G = (V, E)$, $G' = (V', E')$, we define the Cartesian product graph $G \times G'$ to be (as usual) 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, for example, 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 Beněs [5]. It was also proved by Baumslag and Annexstein[4]. The proofs are algorithmic.

Theorem 4.3

$$rt(G \times G') \leq 2rt(G) + rt(G') \ . \qquad (6)$$

Note that since $G \times G'$ and $G' \times G$ are isomorphic graphs then (6) can be written in the following symmetric form

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

Corollary 4.4 *For the n-cube Q^n,*

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

Corollary 4.5 *For the m by n grid graph $P_m \times P_n$, $m \leq n$,*

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

Note that in both cases the routing number is less than twice the diameter. Routing on the n-cube Q^n is a very natural question in view of the popular use of the n-cube structure for models of parallel computation and communication. Indeed, it was this context (through the work of Ramras [11]) which first motivated our considerations of these questions.

Corollary 4.4 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 $diam(Q^n) = n$ and that $rt(Q^n) \geq n + 1$ for $n = 2, 3$. It is reasonable to suspect that $rt(Q^n)$ is closer to n than to $2n$ but this seems difficult to prove.

4.3 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 vertex disjoint copies of K_n joined by an edge e.

It is easy to see that

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

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

Theorem 4.6

$$rt(G_n \times G_n) = (1 + o(1))2n . \qquad (7)$$

The proof of this theorem requires several preparations and will not be given here. The main idea is to reduce the problem 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 and, 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 which 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 and more formally, a C-circulation ϕ on a graph $G = (V, E)$ is a set of assignments $\phi_{uv} : E \rightarrow R^+$, $u, v \in V$, such that for all u, v,

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

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

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

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 $||\phi||$, is defined to be the maximum amount of mass

$$\phi(e) = \sum_{u,v} \phi_{uv}(e)$$

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

It turns out that the main step in the proof of Theorem 4.6 is the following lemma, whose proof, which applies certain cannonical balanced flows corresponding to the spanning trees of C_4, will appear in the full version of the paper.

Lemma 4.7 *For all circulation matrices C on C_4, there always exists some balanced C-circulation ϕ with $||\phi|| \leq 1$.*

This lemma can be used to establish Theorem 4.6. Moreover, the whole argument can be extended 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 may be of interest in its own right. This can indeed be done, by proving the following result whose detailed proof is omitted.

Proposition 4.8 *Let F be a doubly stochastic circulation matrix on Q^k. Then there always exists a balanced C-circulation ϕ on Q^k with $||\phi|| \leq 1$.*

Corollary 4.9 *For fixed k, if G^k denotes the k-fold Cartesian product $G \times \cdots \times G$ of the graph $G = G_n$ considered in Theorem 4.6, then*

$$rt(G^k) = (1 + o(1))2n .$$

Let us define *circ*(G), the *circulation index* of a (connected) graph G by

$$circ(G) := \ sup_C \ inf_\phi ||\phi|| \qquad (8)$$

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 *circ*(G) is the *resistance* of G, defined by

$$res(G) := \max_C \frac{1}{|\widehat{C}|} \min(|A(\widehat{C})|, |B(\widehat{C})|) \qquad (9)$$

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

$$circ(G) \geq res(G) \tag{10}$$

follows by considering the circulation matrix which sends all $|A(\widehat{C})|$ units of mass into $B(\widehat{C})$, for an extremal cutset $\widehat{C}$, where we assume $|A(\widehat{C})| \leq |B(\widehat{C})|$. It is interesting to note that (10) holds with *equality* for Q^k. This is not true in general as can be shown by considering any bounded degree expander graph G on n vertices. In this case, we can have

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

It will be interesting to know other classes of graphs G for which equality holds in 10. Our technique can be easily used to show that the set of even cycles is such a class.

5 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. This property can be used to derive the following theorem.

Theorem 5.1 *Let $G = (V, E)$ be an (n, d, λ)-graph and let σ denote a permutation. Then*

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

Note that in Section 2 it is observed that $rt(G)$ is lower bounded by the diameter of G and therefore the routing number of a d-regular graph as above is at least $\frac{\log n}{\log(d-1)}$ and at most $O(\frac{d^2}{(d-\lambda)^2} \log^2 n)$.

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 Section 2, we know that the routing number is bounded below by $2/\alpha$. As an immediate corollary of Theorem 5.1, up to a polylogarithmic factor, $rt(G)$ is bounded above by a polynomial in $1/\alpha$ for any regular graph with polylogarithmic degrees.

Corollary 5.2 *If $G = (V, E)$ is a d-regular graph on n vertices with expansion coefficient α, then*

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

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 5.1, $rt(G) = O(l^2)$ 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. □

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 resistence of G. From Section 2, we know that the routing number is bounded below by $2/\beta - 1$. As a corollary of Theorem 5.1, up to a polylogarithmic factor, $rt(G)$ is bounded above by a polynomial in $1/\beta$ for any regular graph with polylogarithmic degrees.

Corollary 5.3 *If $G = (V, E)$ is a d-regular graph on n vertices with edge expansion coefficient β, then*

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

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 $d - \frac{\beta^2}{2d}$. □

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 5.4 *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, and hence the result follows from Corollary 5.2. □

Here is an outline of the proof of Theorem 5.1. The first observation in the proof is that by Proposition 4.1 we can restrict our attention to permutations of order 2. For these, one can apply an argument similar to the main idea in [6] and [12] by choosing a random collection of paths between any vertex and its destination and showing that with high probability the overlap between the paths will not be "too large". Next we partition these paths by a simple coloring procedure to classes so that paths in a class are, in a sense, nearly disjoint, and route in each class in its turn according to a certain order which guarantees a successful completion of the task. The whole process is algorithmic and supplies, for a given routing assignment, a polynomial time randomized algorithm for finding the required routing. The detailed proof that requires a few lemmas follows.

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 $v = v_0, v_1, \ldots, v_l$, 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 5.5 *Let $G = (V, E)$ be a d-regular graph on n vertices and suppose $l \geq \log n$. 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 so 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 $e(x)^t \sum_{j:|j-i|<5} A^j e(u)$. It follows that the probability $p(v, u)$ that there exists some 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 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 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], Theorem A.12, page 237), that for each fixed v, the probability that $I(v)$ exceeds, say, $100(l+1)$ is at most

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

(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

$$\sum_{S\subset\{1,\cdot,m\},|S|=s} \Pi_{i\in S}\ p_i \leq \frac{1}{s!}(\sum p_i)^s$$

$$\leq (re/s)^s.$$

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

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

Lemma 5.6 *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 5.5. □

Proof of Theorem 5.1

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 stated in Proposition 4.1). 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 one can split all the 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 steps number i and $2l+2-i$ correspond to flipping the pebbles along edges number 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 $l+1$. One 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. □.

Acknowledgement We would like to thanks W. Goddard for many helpful comments.

References

[1] N. Alon, Eigenvalues and expanders, *Combinatorica* 6 (1986), 83-96.

[2] N. Alon and J. H. Spencer, *The Probabilistic Method*, Wiley, 1991.

[3] L. Babai and M. Szegedy, Local expansion of symmetrical graphs, *Combinatorics, Probability and Computing* 1 (1991), 1-12.

[4] M. Baumslag and F. Annexstein, A unified Framework for off-line permutation routing in parallel networks, Mathematical Systems Theory, 24 (1991), 233-251.

[5] V. E. Beneš, Mathematical Theory of Connecting Networks, Academic Press, New York 1965.

[6] A. Broder, A. Frieze and E. Upfal, Existence and construction of edge disjoint paths on expander graphs, *Proc.* 24^{th} *Annual ACM STOC*, ACM Press (1992), 140-149.

[7] W. Goddard, private communication.

[8] M. Jerrum and A. Sinclair, Approximating the permanent, *SIAM J. Computing* 18 (1989), 1149-1178.

[9] D. E. Knuth, The Art of Computer Programming, vol. 3, p. 241, Addison-Wesley, Reading, MA, 1973.

[10] F. T. Leighton, Methods for message routing in parallel machines, *Proc.* 24^{th} *Annual ACM STOC*, ACM Press (1992), 77-96.

[11] M. Ramras, Routing permutations on a graph, preprint.

[12] L. Valiant, A scheme for fast parallel communication, *SIAM J. on Computing*, 11 (1982), 350-361.