

The Electrical Resistance of a Graph Captures its Commute and Cover Times

(Detailed Abstract)

Ashok K. Chandra*
Prabhakar Raghavan*
Walter L. Ruzzo†
Roman Smolensky‡
Prasoon Tiwari*

Abstract

View an n -vertex, m -edge undirected graph as an electrical network with unit resistors as edges. We extend known relations between random walks and electrical networks by showing that resistance in this network is intimately connected with the *lengths* of random walks on the graph. For example, the *commute time* between two vertices s and t (the expected length of a random walk from s to t and back) is precisely characterized by the effective resistance R_{st} between s and t : commute time = $2mR_{st}$. Additionally, the *cover time* (the expected length of a random walk visiting all vertices) is characterized by the maximum resistance R in the graph to within a factor of $\log n$: $mR \leq \text{cover time} \leq O(mR \log n)$. For many

graphs, the bounds on cover time obtained in this manner are better than those obtained from previous techniques such as the eigenvalues of the adjacency matrix. In particular, using this approach, we improve known bounds on cover times for various classes of graphs, including high-degree graphs, expanders, and multi-dimensional meshes. Moreover, resistance seems to provide an intuitively appealing and tractable approach to these problems.

1. Motivation and Summary

A *random walk* on a graph is the following discrete-time stochastic process: from a vertex, the walk proceeds at the next step to an adjacent vertex chosen uniformly at random. The study of random walks in graphs has many applications in the design of algorithms — in the study of space-bounded computation [2], in distributed computation [5], and in the design of approximation algorithms for some hard counting problems [11]. In this paper we explore the connection between random walks in undirected graphs and electrical network theory, building on the work reported in [7].

In [7] Doyle and Snell demonstrate many interesting relations between random walks in graphs and electrical networks. They view an undirected graph as an electrical network in which each edge of the graph is replaced by a unit resistance. For example, their work related the probability of visiting a vertex a (say) from b before visiting c to

*IBM T.J. Watson Research Center, Yorktown Heights, NY 10598.

†Dept. of Computer Science, FR-35, University of Washington, Seattle, WA 98195. Research supported in part by the National Science Foundation under Grant CCR-8703196; a portion of this work was done while the author was visiting the IBM T.J. Watson Research Center.

‡Dept. of Mathematics, University of California, Berkeley CA 94720. Research supported in part by an IBM Predoctoral Fellowship; a portion of this work was done while the author was visiting the IBM T.J. Watson Research Center.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

the effective resistances between nodes a , b and c in the electrical network. Their work deals with finite as well as infinite graphs, and highlights many tools from electrical network analysis that are useful in the study of random walks. However, they do not discuss the number of steps in a random walk, which will be our primary focus.

The main subject of our study will be the *cover time* of a graph, which is the expected number of steps for a random walk to visit all the vertices in a graph (the maximum being taken over all starting vertices). To this end we define the *electrical resistance of a graph* to be the maximum effective resistance between any pair of vertices. We show that this quantity captures the cover time to within a factor of $O(\log n)$: for n -vertex, m -edge graphs of resistance R ,

$$mR \leq \text{cover time} \leq O(mR \log n). \quad (1)$$

An important step on the way to showing this correspondence is a result we prove about the *commute time* of a random walk: for a given pair of vertices s and t , this is the expected length of a walk from s to t and back to s . We give an *equality* for commute time in terms of the effective resistance between s and t . This equality (like the equalities of Doyle and Snell) reiterates the fact that the electrical properties of the network underlying a graph are innately related to the random walk.

Prior work in the study of the cover time of graphs has used techniques from Markov-chain theory [2, 10], from combinatorics [12], from linear algebra [5] and from graph theory [11]. The electrical approach used here provides an intuitive basis for understanding a variety of phenomena about random walks which had hitherto seemed counterintuitive.

As an example, a simple and plausible conjecture is that adding more edges to a graph can only reduce its cover time since they make it “easier” to reach vertices missed so far. This is shown to be false by the following counterexample: an n -vertex chain has cover time $\Theta(n^2)$, but by adding edges it can be converted to a “lollipop graph” (an $n/2$ -vertex chain connected at one end to an $n/2$ -clique) which has cover time $\Theta(n^3)$. This can be easily explained from resistance arguments. By examining Equation 1 we see that adding edges so as to reduce the resistance R can decrease the cover time; but

adding edges in a region of the graph where R is largely unaffected will increase the cover time.

In addition to a number of new results, our methods yield alternative proofs (and often improvements) of earlier results on cover times. An added advantage of our approach is that our results are *robust*: minor perturbations in the graph (such as the deletion/addition of a few edges) usually do not change the electrical properties of the graph substantially.

The rest of this paper is organized as follows. In section 2 we relate electrical resistance to commute and cover times. Section 3 studies the electrical resistance and the cover time of dense regular graphs. Section 4 studies the relation between the maximum resistance of a graph and the eigenvalues of its adjacency matrix. We then obtain a tight upper bound on the cover time of *expanders* in section 5. We conclude with a study of the resistance and the cover time of multidimensional meshes in section 6. The remainder of this section is devoted to a technical summary of our results and a comparison to previous work.

A *commute* between two vertices s and t is a random walk from s to t and back to s ; and the *commute time* between s and t is the expected length of a commute between the two vertices. Aleliunas *et al.* [2] showed that the commute time between s and t is bounded above by $2md_{st}$, where d_{st} is the distance between s and t . We refine this, showing that the commute time is *exactly* $2mR_{st}$, where R_{st} is the effective resistance between s and t . Note $R_{st} \leq d_{st}$, with equality if and only if there is a unique simple path from s to t . On the other hand, for some graphs R_{st} may be smaller than d_{st} by as much as a factor of n . Thus, resistance not only gives exact values for commute times, these values may be much better than the estimates provided by [2]. (Section 2, Theorem 1.)

Using commute time, we are able to bound the cover time to within a factor of $O(\log n)$, as in (1) above. Letting R_{span} be the minimum resistance of a spanning tree of G , we get an alternative upper bound on cover time:

$$\text{cover time} \leq 2mR_{span}. \quad (2)$$

For many graphs this provides a better bound than (1). For example, $R_{span} = O(n)$ for the n -vertex chain and lollipop graphs, hence their cover times

are $O(n^2)$ and $O(n^3)$, respectively, which happen to be tight. Since $R_{span} \leq n - 1$ for any graph, this result refines the $2m(n - 1)$ upper bound which was one of the main results given by Aleliunas *et al.* [2]. Again, R_{span} may be much smaller than $n - 1$, as small as $O(1)$, in fact. (Section 2, Theorem 3.)

For d -regular graphs, the Aleliunas *et al.* bound for cover time is $O(dn^2)$. Kahn *et al.* [12] improved this bound for d -regular graphs to $O(n^2)$. Reexamination of their proof reveals that it supports the stronger statement that $R_{span} = O(n/d)$ for any d -regular graph, hence cover time is $O(n^2)$ by (2).

Kahn *et al.* [12] also give examples, for any $d \leq n/2 - 1$, of n -vertex, d -regular graphs with maximum resistance $\Omega(n/d)$, and hence by (1) with cover time $\Omega(n^2)$. For $d = n - 1$ (the clique), the cover time is much smaller, namely $O(n \log n)$. One might expect a gradual decline in cover time as d increases from $n/2 - 1$ to $n - 1$. Much to our surprise, this is *not* the case — there is a sharp threshold at $d = n/2$. We show that in going from $d = n/2 - 1$ to $n/2$ the maximum resistance drops from $\Omega(1)$ to $O(1/n)$, hence by (1) the cover time drops from $\Omega(n^2)$ to $O(n \log n)$ (where it remains for all $d \geq n/2$). This result has a very simple and intuitive proof. (Section 3, Theorem 6.)

We relate the resistance of a graph to the second smallest eigenvalue σ_2 of a matrix closely related to its adjacency matrix, thus obtaining some of the results of Broder and Karlin [5] as corollaries. Again, we show that (1) gives much tighter bounds on cover time than are possible in terms of σ_2 alone. Specifically, we show that

$$\frac{1}{n\sigma_2} \leq R \leq \frac{2}{\sigma_2},$$

and exhibit graphs where each inequality is tight. Thus, σ_2 only weakly captures resistance, hence is also weak in estimating cover time (whereas resistance captures cover time to within an $O(\log n)$ factor). (Section 4, Theorem 7.)

One interesting application of our approach is to the cover time for d -regular expander graphs. Using the eigenvalue approach, Broder and Karlin [5] showed that such graphs have cover time $O((n \log n)/(1 - \lambda_2)) = O(dn \log n)$, where $\sigma_2 = d(1 - \lambda_2)$. No better bound is possible using their approach, since there are d -regular expanders having second eigenvalue $\lambda_2 = 1 - \Theta(1/d)$. We are

able to show that the resistance of an expander is $\Theta(1/d)$, and hence the cover time is $O(n \log n)$.

Expanders have potential practical application in the design of efficient, fault-tolerant communication networks, where the expansion properties of the graph make it likely that many communication paths will remain open even in the face of congestion and/or failure of certain links. Larger degree translates to greater robustness to failure and/or congestion. The cover time of the graph is an appropriate metric for the performance of certain kinds of randomized broadcast or routing algorithms. Thus, it is pleasant that increased robustness can be had without significantly increasing the cost of these algorithms — cover time is essentially independent of degree. (Section 4, Theorem 12.)

Using resistance, we also derive upper bounds for covering d -dimensional meshes. We show that a 2-dimensional mesh of size $\sqrt{n} \times \sqrt{n}$ has resistance $\Theta(\log n)$, whereas d -dimensional meshes for $3 \leq d \leq \log_2 n$ have resistance $\Theta(1/d)$. These results generalize recent results of Aldous [1] and Cox [6] who obtained the same results for fixed d , and of Göbel and Jagers [10], who considered $d = \log_2 n$ (hypercubes). (Section 6.)

Our last application of resistance is to derive new upper bounds for universal traversal sequences, namely $O(mR \log(n\log n))$, where g is the number of labeled graphs in the family under consideration. This gives improved upper bounds for universal traversal sequences for many of the classes of graphs considered in this paper, including dense graphs, meshes and expanders. (Section 2, Theorem 4.) We also find the first known family of labeled graphs with a tight bound on UTS length. (Section 6.)

2. Basic Relations

Let $G = (V, E)$ be an undirected graph on $|V| = n$ vertices with $|E| = m$ edges. Let $\mathcal{N}(G)$ be the electrical network having a node for each vertex in V , and, for every edge E , having a one Ohm resistor between the corresponding nodes in $\mathcal{N}(G)$. For two vertices $u, v \in V$, R_{uv} denotes the effective resistance between the corresponding nodes in $\mathcal{N}(G)$.

Let H_{uv} (the *hitting-*, or *first passage time*) denote the expected number of steps in a random walk that starts at u and ends upon first reaching

v . We define C_{uv} , the *commute time* between u and v , by $C_{uv} = H_{uv} + H_{vu}$.

Theorem 1: For any two vertices u and v in G , the commute time $C_{uv} = 2 \cdot m \cdot R_{uv}$.

Proof: Let $d(x)$ denote the degree of x in G , $\forall x \in V$. Let ϕ_{uv} denote the voltage at u in $\mathcal{N}(G)$ with respect to v , if $d(x)$ units of current are injected into each node $x \in V$, and $2m$ are removed from v . Let $N(x)$ denote the set of vertices in V that are adjacent to x in G . We will first prove

$$H_{uv} = \phi_{uv} \quad \forall u \in V. \quad (3)$$

By Kirchoff's current law, Ohm's law, and the fact that all edges have unit resistance, the ϕ_{uv} satisfy

$$d(u) = \sum_{w \in N(u)} (\phi_{uw} - \phi_{wu}) \quad \forall u \in V - \{v\}. \quad (4)$$

By elementary probability theory,

$$H_{uv} = \sum_{w \in N(u)} \frac{1}{d(u)} (1 + H_{vw}) \quad \forall u \in V - \{v\}. \quad (5)$$

Equations (4) and (5) are both linear systems with unique solutions; furthermore, they are identical if we identify ϕ_{uv} in (4) with H_{uv} in (5). This proves (3). To complete the proof of the theorem, we note that H_{vu} is the voltage ϕ_{vu} at v in $\mathcal{N}(G)$ measured with respect to u , when currents are injected into all nodes and removed from u . Changing signs, ϕ_{vu} is also the voltage at u relative to v when current is injected at u , and removed from all other nodes. Since resistive networks are linear, we can derive an expression for $C_{uv} = H_{uv} + H_{vu}$ by superposing the two networks on which ϕ_{uv} and ϕ_{vu} are measured. Currents at all nodes except u and v cancel, resulting in C_{uv} being the voltage between u and v when $\sum_{w \in V} d(w) = 2 \cdot m$ units of current are injected into u and removed from v , which yields the theorem by Ohm's law. \square

D. Aldous, A. Z. Broder and A. R. Karlin, and P. G. Doyle and J. L. Snell all have derived alternative proofs of Theorem 1 using similar methods from renewal theory. For the benefit of readers familiar with renewal methods, we sketch this alternate proof below. A proposition in Section 3.3 of [7] implies that on a walk starting from u , the expected number of returns to u before hitting v is $d(u) \cdot R_{uv}$. Define the stopping time T to be the

time of the first return to u , starting from u , after visiting v at least once. Clearly $C_{uv} = \mathcal{E}[T]$. By a theorem from renewal theory [13, Prop. 9–58], the expected number of returns to u up to and including stopping (at T) is exactly equal to the steady state probability of state u , namely $d(u)/(2m)$, times the expected walk length $\mathcal{E}[T] = C_{uv}$; combining this with the Doyle-Snell expression for the expected number of returns yields the result.

Although Theorem 1 suffices for most of our applications, it is interesting to note that it easily generalizes to walks with non-uniform transition probabilities and costs. Let each edge $\{u, v\} \in E$ have a positive real *resistance* r_{uv} , and let each directed edge have a real *cost* f_{uv} . We now consider a random walk on G defined by the following discrete-time process: when at a vertex $u \in V$, take the edge to neighboring vertex v with probability inversely proportional to the resistance of edge $\{u, v\}$, i.e., with probability

$$p_{uv} = \frac{1/r_{uv}}{\sum_{w \in N(u)} 1/r_{uw}}.$$

For a T -step walk traversing the sequence of (not necessarily distinct) directed edges $(u_0, u_1), (u_1, u_2), \dots, (u_{T-1}, u_T)$, the *cost* of the walk is defined to be $\sum_{j=1}^T f_{u_{j-1}u_j}$. Note that when all resistances and costs are 1, the process we are considering is the standard random walk, and the cost measures the number of steps in the walk.

Let $\mathcal{N}(G)$ be the electrical network derived from G as follows: there is a node in $\mathcal{N}(G)$ for each vertex in V , and for every edge $\{u, v\} \in E$, there is a resistor between the corresponding nodes in $\mathcal{N}(G)$ whose value is r_{uv} . Again, for two vertices $u, v \in V$, R_{uv} denotes the effective resistance between the corresponding nodes in $\mathcal{N}(G)$.

Let H_{uv}^f denote the expected cost (relative to cost function f) of a random walk that starts at u and ends upon first reaching v , and let $C_{uv}^f = H_{uv}^f + H_{vu}^f$.

The surprising fact is that even in this general setting, commute costs are still determined by effective resistances, although the constant of proportionality is no longer simply $2m$.

Theorem 2: Let $F = \sum_{\{x,y\} \in E} (f_{xy} + f_{yx})/r_{xy}$. For any two vertices uv in G , the commute cost $C_{uv}^f = F \cdot R_{uv}$.

Proof: The proof is identical to that of Theorem 1, except that the current injected into node x is $f_x = \sum_{y \in N(x)} f_{xy}/r_{xy}$ for all $x \in V$. \square

Theorem 1 is obviously a corollary when all resistances and costs are 1 ($F = 2m$).

Aleliunas *et al.* [2] showed that during a commute between u and v , every directed edge is traversed the same expected number, τ , of times. This follows easily from Theorem 2 by setting all resistances to one, and all costs are zero, except for an arbitrary directed edge, given cost one. Further, we find that $\tau = R_{uv}$.

For non-unit resistances, Doyle and Snell [7] have shown that the class of random processes considered here is exactly the class of “reversible ergodic Markov chains.” Thus, with general resistances, but unit costs, Theorem 2 determines the number of steps in commutes in such chains. Our results below can then be used to bound the cover time for reversible ergodic Markov chains, a problem also considered by Broder and Karlin [5].

Throughout the remainder of the paper, unless otherwise stated, graphs are assumed to be unweighted, i.e. we will consider only the basic unit-resistance version of the random walk problems. We now turn to cover times.

Let $R = \max_{u,v \in V} R_{uv}$. Let $\mathcal{N}'(G)$ be a network in which there is a node u' for every vertex $u \in V$, and an edge $(u', v') \forall u', v' \in V$ whose length equals R_{uv} . Let R_{span} be the length of the minimum spanning tree in $\mathcal{N}'(G)$. Let C_u denote the expected length of a walk that starts at u and ends upon visiting every vertex in G at least once. Let C_G be the cover time of G , i.e., $C_G = \max_u C_u$.

Theorem 3:

$$m \cdot R \leq C_G \leq \min((2 + o(1)) \cdot m \cdot R \cdot \ln n, 2 \cdot m \cdot R_{span})$$

Proof: The proof of the lower bound follows from the fact that there exist vertices u, v such that $R = R_{uv}$ and $\max(H_{uv}, H_{vu}) \geq C_{uv}/2$; the result then follows from Theorem 1. Matthews [15] has shown that the cover time is at most $(1 + o(1))H \ln n$, where $H = \max_{uv} H_{uv}$. The first upper bound follows from the observation that $H \leq \max_{uv} C_{uv} = 2mR$. (A similar upper bound with a somewhat larger constant can be obtained from a simple argument like that used in Theorem 4 below.) The proof of the second upper bound

follows directly from the spanning-tree argument of [2]. \square

Note that the bounds in Theorem 3 cannot in general be improved; the upper bounds are tight (within constant factors) for the complete graph and the chain, resp., and the lower bound is also tight for the chain. There are also graphs for which none of the bounds above is tight.

Let \mathcal{G} be a family of labeled d -regular graphs on n vertices. Let $U(\mathcal{G})$ denote the length of the shortest universal traversal sequence (see [2] or [4] for definitions) for all the labeled graphs in \mathcal{G} . Let $R(\mathcal{G})$ denote the maximum resistance between any pair of vertices in any graph in \mathcal{G} .

Theorem 4: $U(\mathcal{G}) \leq 5 \cdot m \cdot R(\mathcal{G}) \cdot \log_2(n \cdot |\mathcal{G}|)$.

Proof: The proof is by a probabilistic argument similar to that in [2]. Given a labeled graph $G \in \mathcal{G}$, let v be a vertex of G . Consider a random walk of length $5 \cdot m \cdot R(\mathcal{G}) \cdot \log_2(n \cdot |\mathcal{G}|)$, divided into $\log_2(n \cdot |\mathcal{G}|)$ “epochs” each of length $5 \cdot m \cdot R(\mathcal{G})$. The probability that the walk fails to visit v in one of these epochs is at most $2/5$ by Theorem 1 and Markov’s inequality, regardless of the vertex of G at which the epoch began. The probability that v was not visited during any of the epochs is thus at most $(n \cdot |\mathcal{G}|)^{-c}$ for a value of $c > 1$. Summing this probability that v is not visited over all n choices of the vertex v and all $|\mathcal{G}|$ choices of the graph G , the probability that the random walk (sequence) is not universal is less than one. Thus there is a sequence of this length that is universal for the class. \square

The constant 5 in Theorem 4 can be improved.

3. Dense Graphs

In this section we demonstrate for d -regular graphs the threshold in resistance, and hence cover time, at $d = \lfloor n/2 \rfloor$.

A simple fact we will use several times to help bound resistances is the following.

Rayleigh’s “Short/Cut” Principle [7, 16]: Resistance is never raised by lowering the resistance on an edge, e.g., by “shorting” two nodes together, and is never lowered by raising the resistance on an edge, e.g., by “cutting” it. Similarly, resistance is never lowered by “cutting” a node, leaving each incident edge attached to only one of the two “halves” of the node.

As one very simple application, notice that in a graph with minimum degree d , $R \geq 1/d$: short all nodes except the one of minimum degree. This lower bound will prove useful later.

Another simple application is the following lemma.

Lemma 5: If G contains p edge-disjoint paths of length less than or equal to l from s to t , then $R_{st} \leq l/p$.

Proof: Extract from G a network H as follows. Cut all edges not on one of the p paths. Split nodes if necessary to make the paths vertex-disjoint. Note that the paths are edge-disjoint, so it is possible to do this without duplicating edges. Raise the resistance of each edge in a path of length $l' < l$ to l/l' Ohms. Clearly R_{st} is exactly l/p in H . Hence, by the “short/cut” principle, $R_{st} \leq l/p$ in G . \square

When n is even and $d = \lfloor n/2 \rfloor - 1$, there are d -regular graphs having maximum resistance $\Theta(1)$. To see this, take two $n/2$ -vertex cliques, remove one edge (a_i, b_i) from clique i , $0 \leq i \leq 1$, and join the two cliques with edges (a_i, b_{1-i}) , $0 \leq i \leq 1$. By the “short/cut” principle above, the resistance between any two vertices not in the same $n/2$ -clique must be at least $1/2$ Ohm — shorting all the nodes in each clique leaves a two-node network with two 1 Ohm resistors in parallel. Thus, by (1), the cover time for this graph is $\Omega(n^2)$; this bound is tight by the results of Kahn *et al.* [12]. A similar construction works for odd n and even $d \leq \lfloor n/2 \rfloor - 1$.

When $d = \lfloor n/2 \rfloor$, the situation changes radically. Intuitively, one can't add another $\lfloor n/2 \rfloor$ edges to the graph above without making it so highly connected that the resistance drops sharply. This is proved below.

Theorem 6: For any n -vertex, d -regular graph G with $d \geq \lfloor n/2 \rfloor$, $R \leq 4/d = O(1/n)$. Hence $C_G = O(n \log n)$.

Proof: The key point is to show that there are $\Omega(d)$ edge-disjoint paths of length at most 4 between any pair of vertices. The result then follows by application of Lemma 5. Consider any two vertices s and t . Let k be 1 if $\{s, t\} \in E$, else $k = 0$. Let k' be the number of vertices ($\neq s, t$) mutually adjacent to s and t . Then there are exactly $j = d - k - k'$ vertices which are adjacent to s but not to t , and vice versa. Call these vertices

s_1, \dots, s_j and t_1, \dots, t_j , resp. Let k'' be the size of a maximum matching between the s_i 's and the t_i 's, and WLOG assume that $\{\{s_i, t_i\} \mid 1 \leq i \leq k''\}$ are the matching edges. Because $d \geq \lfloor n/2 \rfloor$, every pair of vertices in G either are neighbors or have a common neighbor. In particular, s_i and t_i have a common neighbor m_i , $k'' < i \leq j$. Thus, we have d paths of length at most 4 from s to t , namely k of length 1, k' of length 2, k'' of length 3 ($\langle s, s_i, t_i, t \rangle$, $1 \leq i \leq k''$), and $d - k - k' - k''$ of length 4 ($\langle s, s_i, m_i, t_i, t \rangle$, $k'' < i \leq j$). Note that the m_i 's are not necessarily distinct from each other or from the other vertices mentioned. Despite this, it's not hard to see that the d paths are edge-disjoint. Thus, there are d edge-disjoint paths of length at most 4 from s to t , hence $R_{st} \leq 4/d = O(1/n)$ by Lemma 5. \square

Theorem 6 shows a sharp threshold in cover time at $d = \lfloor n/2 \rfloor$. Applying Theorem 4 we see that the length of universal traversal sequences for d -regular graphs, for any $d \geq \lfloor n/2 \rfloor$, is $O(n^3 \log n)$. This bound was previously known to hold only for cliques ($d = n - 1$). Interestingly, recent lower bounds for universal traversal sequences [4] are $\Omega(n^4)$ for linear $d \leq n/3 - 2$. Thus, length of universal traversal sequences also declines somewhere between $d = n/3 - 2$ and $d = \lfloor n/2 \rfloor$; whether there is a sharp threshold at $d = \lfloor n/2 \rfloor$ as in the case of cover time is unknown.

4. Resistance and Eigenvalues

Consider a connected graph G with vertices numbered $1, 2, \dots, n$. Let \mathbf{D} be the diagonal matrix whose i^{th} diagonal entry is $d(i)$, the degree of the vertex i . Let \mathbf{A} be the adjacency matrix of G , and define $\mathbf{K} = \mathbf{D} - \mathbf{A}$. Since \mathbf{K} is a real symmetric matrix, all its eigenvalues are real and it has a set of n orthonormal eigenvectors (see, for example, [8]). It is easy to verify that zero is an eigenvalue of \mathbf{K} , and that the vector of all ones is a corresponding eigenvector. By Gershgorin's theorem ([8]) zero is also the smallest eigenvalue, and has multiplicity one since G is connected. Define $\sigma(G)$ to be the second smallest eigenvalue of \mathbf{K} .

We will use the following inner product in this section.

Definition: Let $\mathbf{x} = [x_1, x_2, \dots, x_n]$ and $\mathbf{y} = [y_1, y_2, \dots, y_n]$ be vectors of n components. The

the inner product of \mathbf{x} and \mathbf{y} , denoted by (\mathbf{x}, \mathbf{y}) , is given by $\sum_{i=1}^n (x_i y_i)$. The length of \mathbf{x} , denoted by $\|\mathbf{x}\|$, is given by $\sqrt{(\mathbf{x}, \mathbf{x})}$.

Let $\sigma_1 < \sigma_2 \leq \sigma_3 \leq \dots \leq \sigma_n$ be the eigenvalues of \mathbf{K} , and let $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ be the corresponding orthonormal eigenvectors, i.e.,

$$(\mathbf{u}_i, \mathbf{u}_j) = \begin{cases} 1 & \text{if } i = j, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

It is well known that all components of \mathbf{u}_j can be chosen to be real. By the discussion above $\sigma_1 = 0$, and all components of \mathbf{u}_1 are equal to $1/\sqrt{n}$. Also, note that $\sigma(G) = \sigma_2$.

Let \mathbf{U} be the $n \times n$ unitary matrix whose j^{th} column is \mathbf{u}_j , and let $\mathbf{\Sigma}$ be the $n \times n$ diagonal matrix whose i^{th} diagonal entry is σ_i . Then $\mathbf{U}^T \mathbf{U} = \mathbf{I}$, and $\mathbf{K} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^T$.

Let u_{ij} be the i^{th} component of \mathbf{u}_j . Since \mathbf{U} is the inverse of \mathbf{U}^T , i.e., $\mathbf{U} \mathbf{U}^T = \mathbf{I}$, we get that $\sum_{k=1}^n u_{ik} u_{jk} = 0$, unless $i = j$, in which case the sum is 1.

Theorem 7: If G is a graph on n vertices, then

$$\frac{1}{n\sigma(G)} \leq R \leq \frac{2}{\sigma(G)}.$$

Proof: Let $\mathbf{v} = [v_1, v_2, \dots, v_n]^T$ be the vector of voltages in $\mathcal{N}(G)$, relative to node t , when a unit current is injected into node s and removed from node t . Clearly, $v_t = 0$, and $0 \leq v_k \leq R$ for all k . Let $\mathbf{c} = \mathbf{e}_s - \mathbf{e}_t$, where \mathbf{e}_k is an n component vector whose k^{th} component is one and all other components are zero. Then $\mathbf{K} \mathbf{v} = \mathbf{c}$, and therefore

$$\mathbf{v} = \delta \mathbf{u}_1 + \sum_{k=2}^n \frac{\alpha_k}{\sigma_k} \mathbf{u}_k, \quad (6)$$

where δ is \sqrt{n} times the average voltage in the network, and $\alpha_k = (\mathbf{c}, \mathbf{u}_k)$. Notice that $\alpha_1 = 0$ and $\sum_{k=1}^n \alpha_k^2 = \|\mathbf{U}^T \mathbf{c}\|^2 = \|\mathbf{c}\|^2 = 2$.

For the upper bound, choose s and t above so that $R = R_{st}$. Note that $R = v_s = v_s - v_t$, so by Equation 6, we get

$$\begin{aligned} R &= v_s - v_t \\ &= \sum_{k=2}^n \frac{\alpha_k (u_{sk} - u_{tk})}{\sigma_k} \\ &= \sum_{k=2}^n \frac{\alpha_k^2}{\sigma_k} \end{aligned}$$

$$\leq \frac{2}{\sigma(G)}.$$

For the lower bound, proceed as above, this time choosing s and t so that $(\mathbf{c}, \mathbf{u}_2) \geq 1/\sqrt{n}$. Such a pair exist since some component of \mathbf{u}_2 must have magnitude at least $1/\sqrt{n}$, and not all are of the same sign, since $(\mathbf{u}_1, \mathbf{u}_2) = 0$. Note that $0 \leq v_i \leq R$, so $\|\mathbf{v}\| \leq R\sqrt{n}$. But

$$\|\mathbf{v}\| = \sqrt{\delta^2 + \sum_{k=2}^n \left(\frac{\alpha_k}{\sigma_k}\right)^2} \geq \frac{\alpha_2}{\sigma_2} \geq \frac{1}{\sigma_2 \sqrt{n}}.$$

This implies that $R \geq 1/(n\sigma(G))$. \square

Theorem 3 immediately implies the following corollary.

Corollary 8: $C_G \leq (4 + o(1))m \ln n / \sigma(G)$.

We need the following lemma to compare the preceding theorem to some previously known results. Let \mathbf{P} be the transition matrix of the Markov chain corresponding to the random walk on a graph G . Since $\mathbf{P} = \mathbf{A} \mathbf{D}^{-1}$ and $\mathbf{Q} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-1} \mathbf{D}^{\frac{1}{2}}$ are similar matrices, they have the same set of eigenvalues. Moreover, all these eigenvalues are real because \mathbf{Q} is a real symmetric matrix. Let $\lambda_1 > \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of \mathbf{P} (and \mathbf{Q}). For an ergodic Markov chain, it is well known that $1 = \lambda_1 > \lambda_2$. Observe that the λ_i 's are arranged in the descending order whereas the σ_i 's are arranged in the ascending order. Define $\lambda(G) = \lambda_2$. Since \mathbf{Q} is symmetric, it has a set of orthonormal eigenvectors $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ where $\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{w}_i = \lambda_i \mathbf{w}_i$.

Lemma 9: Let G be a connected graph with minimum and maximum degrees given by d_{min} and d_{max} , respectively. Then

$$(1 - \lambda(G))d_{min} \leq \sigma(G) \leq (1 - \lambda(G))d_{max}.$$

Don Coppersmith pointed out to us the following elegant proof of this lemma.

Proof: If \mathbf{B} is an $n \times n$ symmetric real matrix with real eigenvalues $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$, and corresponding orthonormal eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$, then Rayleigh's principle [8] gives the following expressions for the eigenvalues:

$$\alpha_i = \min_{\mathbf{x} \perp \{\mathbf{v}_{i+1}, \mathbf{v}_{i+2}, \dots, \mathbf{v}_n\}} \frac{\mathbf{x}^T \mathbf{B} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}, \quad (7)$$

$$= \max_{\mathbf{x} \perp \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{i-1}\}} \frac{\mathbf{x}^T \mathbf{B} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}. \quad (8)$$

Observe that $\mathbf{x} \perp \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{i-1}\}$ if and only if that \mathbf{x} is in the span of $\{\mathbf{v}_i, \mathbf{v}_{i+1}, \dots, \mathbf{v}_n\}$. This observation will be used later.

With the \mathbf{u}_i 's as before, consider the set of $n+1$ vectors $\{\mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_n, \mathbf{D}^{-\frac{1}{2}}\mathbf{w}_1, \mathbf{D}^{-\frac{1}{2}}\mathbf{w}_2\}$. Since there are more than n vectors in this set, they are linearly dependent, i.e., there exist constants $a_2, a_3, \dots, a_n, b_1, b_2$, not all zero, such that

$$\sum_{i=2}^n a_i \mathbf{u}_i = \mathbf{D}^{-\frac{1}{2}}(b_1 \mathbf{w}_1 + b_2 \mathbf{w}_2). \quad (9)$$

Let us denote the left hand side of this equation by \mathbf{z} . If $\mathbf{z} = \mathbf{0}$, then $a_2 = a_3 = \dots = a_n = 0$, and $b_1 = b_2 = 0$ because each of the two sets of vectors $\{\mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_n\}$ and $\{\mathbf{w}_1, \mathbf{w}_2\}$ are independent. Therefore, $\mathbf{z} \neq \mathbf{0}$, and without loss of generality, we may assume that \mathbf{z} is a unit vector. Equation 7 implies that

$$\begin{aligned} \sigma(G) &\leq \mathbf{z}^T(\mathbf{D} - \mathbf{A})\mathbf{z}; \text{ and} \\ \lambda(G) &\leq \frac{(\mathbf{D}^{\frac{1}{2}}\mathbf{z})^T(\mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}})(\mathbf{D}^{\frac{1}{2}}\mathbf{z})}{\mathbf{z}^T\mathbf{D}\mathbf{z}}. \end{aligned}$$

The first of these two inequalities yields an upper bound on $\mathbf{z}^T\mathbf{A}\mathbf{z}$. Substituting this upper bound in the second inequality, we arrive at $\sigma(G) \leq \mathbf{z}^T\mathbf{D}\mathbf{z}(1 - \lambda(G))$. Finally, observe that $\mathbf{z}^T\mathbf{D}\mathbf{z} \leq d_{max}$, which establishes proves the upper bound on $\sigma(G)$ asserted in the statement of the theorem.

The lower bound can be proved in a similar manner by starting with the set of $n+1$ vectors $\{\mathbf{w}_2, \mathbf{w}_3, \dots, \mathbf{w}_n, \mathbf{D}^{\frac{1}{2}}\mathbf{u}_1, \mathbf{D}^{\frac{1}{2}}\mathbf{u}_2\}$ and using Equation 8 instead of Equation 7. \square

The following example will be useful in showing where the inequalities in Theorem 7 are tight.

Definition: Let $Z_n = \{0, 1, \dots, n-1\}$. For $n_1, n_2, \dots, n_d \geq 2$, the $n_1 \times n_2 \times \dots \times n_d$ d -dimensional (toroidal) mesh is an undirected graph $G(V, E)$ where $V = Z_{n_1} \times Z_{n_2} \times \dots \times Z_{n_d}$, and any vertex (k_1, k_2, \dots, k_d) is connected to vertices $(k_1, \dots, k_{i-1}, k_i \pm 1 \bmod n_i, k_{i+1}, \dots, k_d)$, for each $i = 1, 2, \dots, d$.

A $k \times k \times \dots \times k$ d -dimensional mesh will be called a (k, d) mesh for short.

Theorem 10: The multiset

$$\left\{ 2 \sum_{i=1}^d \cos\left(\frac{2\pi k_i}{n_i}\right) : (k_1, k_2, \dots, k_d) \in Z_{n_1} \times Z_{n_2} \times \dots \times Z_{n_d} \right\}$$

contains all the eigenvalues (with correct multiplicity) of the adjacency matrix of the $n_1 \times n_2 \times \dots \times n_d$ d -dimensional mesh.

Proof: Let ω_i be the n_i^{th} root of unity and let $n = \prod_{i=1}^d n_i$. Choose any $(k_1, k_2, \dots, k_d) \in Z_{n_1} \times Z_{n_2} \times \dots \times Z_{n_d}$. Let \mathbf{u} be a vector of n components whose component corresponding to vertex (j_1, j_2, \dots, j_d) is given by $\prod_{i=1}^d \omega_i^{k_i j_i}$. Check that \mathbf{u} is an eigenvector of the adjacency matrix of the $n_1 \times n_2 \times \dots \times n_d$ mesh, with eigenvalue $\sum_{i=1}^d (\omega_i^{k_i} + \omega_i^{-k_i})$. \square

Corollary 11: If G is the $n_1 \times n_2 \times \dots \times n_d$ d -dimensional mesh, then $\sigma(G) = 2(1 - \cos \frac{2\pi}{n_i}) \approx (\frac{2\pi}{n_i})^2$, where n_i is the largest of the n_j 's.

We now discuss some consequences of Theorem 7. The lower bound on resistance given by Theorem 7 is tight to within a constant factor for the n -node cycle (the $(n, 1)$ -mesh). Observe that for this graph $R = \Theta(n)$, and from Corollary 11 $\sigma(n\text{-cycle}) \approx (\frac{2\pi}{n})^2$. The upper bound on resistance given by Theorem 7 is exactly tight for the n -node complete graph. Observe that for this graph $R = 2/n$, and $\sigma(K_n) = n$. In view of the last two remarks, it is not possible to improve the inequalities in Theorem 7, except perhaps the constant factor in the lower bound, for all graphs. On the other hand, both the inequalities in Theorem 7 are weak for $(n^{1/d}, d)$ -meshes, for any $d \geq 2$. The maximum resistance in multidimensional meshes can be determined by other techniques. This is the subject of Section 6.

Theorem 7 also improves a bound due to Landau and Odlyzko [14] (and Corollary 17 of [5]). In [14] it is proved that $(1 - \lambda(G)) \geq 1/((d_{max} + 1)\Delta n)$ where d_{max} and Δ are the maximum degree and the diameter of G , respectively. Using the resistance bound from Theorem 7, and Lemma 9, we get $(1 - \lambda(G)) \geq 1/(d_{max} R n)$. This is an improvement because $\Delta \geq R$.

Some upper bounds on cover times due to Broder and Karlin [5] are implied as a consequence of Theorem 7. For example, Corollary 8 and Lemma 9 imply that $C_G \leq ((4 + o(1))m \ln n)/(d_{min}(1 - \lambda(G)))$. For most graphs, this is stronger than Corollary 8 of [5], which states that $C_G \leq (1 + o(1))n^2 \ln n/(1 - \lambda(G))$.

Finally, Theorem 7 also implies that the resistance between any pair of vertices in any family

of bounded degree expander graphs (see the next section, or [3]) is bounded by $O(1)$.

In the rest of this paper we study resistance in two graph families: (i) families of expanders whose maximum degree may be a function of n ; and (ii) multidimensional meshes. Neither the results in [5] nor Theorem 7 yield good bounds on the cover time of these graphs.

5. Expanders

We will use the following definition of expanders, also used by Broder and Karlin [5]

Definition: An (n, d, α) -expander is a graph $G = (V, E)$ on n vertices, of maximal degree d , such that every subset $X \subseteq V$ satisfying $|X| \leq n/2$ has $|N(X) - X| \geq \alpha \cdot |X|$. Recall $N(X) = \{v \mid \{u, v\} \in E \text{ for some } u \in X\}$.

Note that $\alpha \leq 1$, and $\alpha > 0$ if G is connected.

There is some inconsistency in the literature concerning the definition of “expanders”. For instance, Alon [3] calls graphs with the above property “magnifiers”, reserving the term “expander” for bipartite graphs with a similar property. He shows very close connections between the two notions, so there seems to be no essential loss of generality in choosing the above definition, which is more convenient for our purposes. Further, Ronitt Rubinfeld [18] has shown a result analogous to our Theorem 12 for graphs which are “expanders” according to the definition of Peleg and Upfal [17], giving further evidence that the basic result of this section is reasonably insensitive to variations in the definition.

Alon [3] has shown that if G is an (n, d, α) -expander, then $\sigma(G) \geq \alpha^2/(4 + 2\alpha^2)$, hence by Theorem 7, $R \leq (2 + \alpha^2)/\alpha^2$. The main result of this section sharpens this estimate, reducing it by a factor of order d . When d is a function of n , this considerably improves the bounds of Broder and Karlin [5] on the cover time of these graphs.

Theorem 12: A connected (n, d', α) -expander G , with minimum degree d , has resistance at most $24/(\alpha^2(d + 1))$.

Proof: Let s, t be two vertices in G such that $R_{s,t} = R$. In the electrical network $\mathcal{N}(G)$, connect a unit voltage source between s and t , with t grounded. We will show by contradiction that the current flow from s to t in $\mathcal{N}(G)$ is at

least $\alpha^2(d + 1)/(8(1 + \alpha/2)(1 + \alpha))$, hence $R \leq (8(1 + \alpha/2)(1 + \alpha))/(\alpha^2(d + 1))$, which is at most $24/(\alpha^2(d + 1))$, since $\alpha \leq 1$.

The basic idea is that any set T of “low voltage” nodes has a relatively large set U of neighbors, since G is an expander. Further, the bulk of the nodes in U must be at voltages “near” those in T , for otherwise there would be a “large” current flow from U to T . Repeating this argument inductively, we show that, unless the current is “large”, more than half the nodes have voltage less than $1/2$; a similar argument for sets S of “high voltage” nodes shows that more than half have voltage greater than $1/2$, a contradiction. Thus the current must be “large”. These ideas are quantified and made precise below.

Let

$$c = \frac{1}{2} \left(\sum_{i=0}^{\infty} (1 + \alpha/2)^{-i} \right)^{-1} = \frac{\alpha}{4(1 + \alpha/2)}$$

$$v_k = c \sum_{i=0}^k (1 + \alpha/2)^{-i}, \text{ for all } k \geq 0,$$

and define

$$\begin{aligned} T_k &= \{a \mid \text{node } a \text{ of } \mathcal{N}(G) \text{ has voltage} < v_k\} \\ S_k &= \{a \mid \text{node } a \text{ of } \mathcal{N}(G) \text{ has voltage} > 1 - v_k\} \\ t_k &= |T_k| \\ s_k &= |S_k|. \end{aligned}$$

Note that $0 < v_0 < v_1 < \dots < 1/2$.

First we make the following claim.

Claim: $t_0 \geq (d + 1)(1 + \alpha/2)/(1 + \alpha)$, and for all $k \geq 1$, if $t_{k-1} \leq n/2$ then $t_k \geq (1 + \alpha/2)t_{k-1}$, and so $t_k \geq (1 + \alpha/2)^{k+1}(d + 1)/(1 + \alpha)$.

The claim is proved by induction on k .

BASIS ($k = 0$):

Suppose $t_0 < (d + 1)(1 + \alpha/2)/(1 + \alpha)$. Then at least $(d - (t_0 - 1))$ of t 's neighbors are at voltage at least v_0 , hence the current flow into t is at least

$$\begin{aligned} &(d - (t_0 - 1))v_0 \\ &> \left(d - \left((d + 1) \frac{(1 + \alpha/2)}{(1 + \alpha)} - 1 \right) \right) \cdot \\ &\quad \frac{\alpha}{4(1 + \alpha/2)} \\ &= (d + 1) \left(\frac{\alpha}{2(1 + \alpha)} \right) \frac{\alpha}{4(1 + \alpha/2)} \end{aligned}$$

$$= \frac{\alpha^2(d+1)}{8(1+\alpha/2)(1+\alpha)},$$

contradicting the assumption that the current is less than the later quantity.

INDUCTION ($k \geq 1$, and $t_{k-1} \leq n/2$):

If $t_{k-1} \leq n/2$, then by the fact that G is an expander, $U = N(T_{k-1}) - T_{k-1}$ has size at least αt_{k-1} . If T_k is small, then more than half of the nodes of U are not in T_k , hence at voltage at least v_k . In this case, the current flow from U to T_k would be too large. More precisely, if $t_k < (1 + \alpha/2)t_{k-1}$, then the current will be greater than

$$\begin{aligned} & \frac{\alpha}{2} t_{k-1} (v_k - v_{k-1}) \\ & \geq \frac{\alpha}{2} \left((d+1) \frac{(1+\alpha/2)^k}{(1+\alpha)} \right) \cdot (c(1+\alpha/2)^{-k}) \\ & = \frac{\alpha^2(d+1)}{8(1+\alpha/2)(1+\alpha)} \end{aligned}$$

again contradicting the assumption that the current is less than the later quantity. Thus, $t_k \geq (1 + \alpha/2)t_{k-1}$. This completes the proof of the claim.

As a consequence of the claim, there is a $k \geq 0$ such that $t_k > n/2$, i.e. more than half the nodes have voltage strictly less than $1/2$ volt. By a similar argument about the high-voltage sets S_i , there is a k' such that $s_{k'} > n/2$, i.e., more than half the nodes *also* have voltage strictly *greater* than $1/2$, an impossibility. Thus, the current from s to t must be at least $\alpha^2(d+1)/(8(1+\alpha/2)(1+\alpha))$. \square

It is unknown whether the quadratic dependence on $1/\alpha$ is necessary.

We will briefly sketch an alternative proof of Theorem 12. It is in some ways more complex than the foregoing, but still intuitive, and also seems considerably more general. In fact, we originally proved both the dense graph result and a somewhat weaker version of the expander result (Theorems 6 and 12) using the approach outlined below, before finding the more direct proofs given above. The technique is also similar to the one we use in the mesh proofs in Section 6. Peter Doyle contributed an important refinement to the technique.

Let $G = (V, E)$, s, t be as above. Build an auxiliary layered graph H , with $2l + 1$ layers (l defined below), each layer consisting of a copy of V , and with an edge between vertices u and v in adjacent

layers if and only if $\{u, v\}$ is an edge in G . Delete all vertices not on a shortest path (length $2l$) from s' , the copy of s in the topmost layer, to t' , the copy of t in the bottommost layer. We will first estimate the resistance between s' and t' in an electrical network derived from H .

Intuitively, we hope that when a voltage is applied between s' and t' the layers of H will be good approximations to the equipotential surfaces, and in fact we can adjust resistances, using the ‘‘cut’’ principle, so that this becomes true.

Edges are given capacities, exponentially decreasing towards the middle layer. Specifically, all edges between layers k and $k + 1$, (counting from the nearer of s' and t'), are given capacity $c_k = (1 + \alpha)^{-k}$. The expansion property of G prevents H from having a small s' - t' cut, since edge capacities are decreased at the same rate as expansion increases the number of relevant edges. More precisely, let S (T) be the set of vertices connected to s' (t') after the cut is made. If the cut is small, then not enough edges have been cut to prevent some expansion within S from one layer to the next. Choose l large enough so that S contains more than half of the middle layer. By the same argument T contains over half of the middle layer, too, a contradiction. Hence by the max-flow/min-cut theorem there is a large ($O(d)$) s' - t' flow D .

Next, convert the flow to an electrical current flow by constructing an electrical network $\mathcal{N}(H)$ from H by assigning each edge of capacity c_k carrying flow $f \leq c_k$ a resistance $(c_k/f) \cdot (c_k)^{-1/2}$. Then the flow in H is exactly the electrical current flow in $\mathcal{N}(H)$, and there is a voltage drop of exactly $c_k^{1/2}$ between layers k and $k + 1$. Thus, the resistance between s' and t' in $\mathcal{N}(H)$ is exactly $2(\sum_{k=0}^l c_k^{1/2})/D = O(1/(\alpha d))$.

Finally, short together all copies of each vertex in G . The result is essentially a subgraph of G , except with up to $2l$ parallel edges for each edge of G . Since $c_k/f \geq 1$ above, it is easily verified that the effective resistance of any such set of parallel edges is at least $1/(2 \sum_{k=0}^l c_k^{1/2}) = \Omega(\alpha)$. Thus, by the ‘‘short’’ principle, R_{st} in G is bounded above by $R_{s't'}/\alpha$ in $\mathcal{N}(H)$, which gives the result. \square

Rubinfeld’s proof [18] uses yet a third technique: she applies a result of Friedman and Pippenger [9] to find large trees in G rooted at s and t , uses the max flow/min cut theorem to find many short

paths joining the leaves of the two trees, and finally uses the short/cut principle to bound the resistance.

6. Meshes

In this section we consider the resistance of regular meshes. Recall (from section 4) that a (k, d) mesh is a d -dimensional mesh of side k .

Resistance of infinite meshes has been previously considered. In particular, it is the focus of a portion of Doyle and Snell's monograph [7]. They show that the resistance from the origin to infinity in an infinite two-dimensional mesh is infinite, but in a three (or higher) dimensional mesh resistance is bounded. Their motivation for this question was to obtain an elementary proof of Pólya's beautiful theorem that random walks in two dimensional meshes are recurrent while those in three or higher dimensions are transient. Resistance of the infinite mesh settles this question, since, as Doyle and Snell also show, the resistance to infinity determines the probability of escape to infinity.

Resistance of finite meshes seems not to have been considered before. Our approach is similar to [7, Section 8.7].

It is easy to see that a $(k, 1)$ mesh has resistance $n/4 - o(1)$. For higher dimensions we have:

Theorem 13: The (k, d) mesh with $n = k^d$ nodes has resistance

$$R_G = \begin{cases} \Theta(\log n) & \text{for } d = 2, \\ \Theta(\frac{1}{d}) & \text{for } d \geq 3. \end{cases}$$

Before outlining the proof of this theorem, we need to develop some machinery from circuit theory. The following triangle inequality for resistances proves useful.

Lemma 14: For any three vertices u, v, w in G ,

$$R_{uv} \leq R_{uw} + R_{wv}.$$

Definition: Given an electrical network $\mathcal{G}(V, E, r)$, with resistance $r(e)$ for each edge e , a flow c is a function from $V \times V$ to the reals, having the property that $c(u, v) = 0$ unless $\{u, v\} \in E$, and c is antisymmetric, i.e., $c(u, v) = -c(v, u)$. The net flow out of a node will be denoted $c(u) = \sum_{v \in V} c(u, v)$, and the flow along an edge $e = \{u, v\}$

is $c(e) = |c(u, v)|$. A source (respectively, sink) is a node u with $c(u) > 0$ (respectively, $c(u) < 0$). Given two flows c_1, c_2 , we can obtain a new flow $c = c_1 + c_2$ given by $c(u, v) = c_1(u, v) + c_2(u, v)$. The power $P(c)$ in a flow is $P(c) = \sum_{e \in E} r(e)c^2(e)$. A flow is a current flow if it satisfies Kirchoff's law, i.e. for any directed cycle $u_0, u_1, \dots, u_{k-1}, u_0$, $\sum_{i=0}^{k-1} c(u_i, u_{i+1 \bmod k}) \cdot r(u_i, u_{i+1 \bmod k}) = 0$.

Lemma 15: (The Minimum Power Principle [19]; also known as Thomson's Principle [20, 7].) For any electrical network (V, E, r) and flow c with only one source u , one sink v , and $c(u) = -c(v) = 1$, we have $R_{u,v} \leq P(c)$.

Lemma 16: For any two flows c_1, c_2 in an electrical network,

$$P(c_1 + c_2) \leq 2(P(c_1) + P(c_2)).$$

Proof: Straightforward. \square

Proof of Theorem 13: To prove the upper bound, construct a flow c_0 in a $(k+1, d)$ mesh as follows. For any node $u = (k_1, \dots, k_d)$, $k_i < k+1$, let its length from the origin be defined as $l(u) = \sum k_i$. For any node $v = (k_1, k_2, \dots, k_d)$, and $u = (k_1, k_2, \dots, k_i - 1, \dots, k_d)$, with $k_i \geq 1$, $l = l(v) \leq k$, we let $c_0(u, v) = -c_0(v, u) = k_i / (l \binom{l+d-1}{d-1})$. The flow in all other edges is zero. The flow c_0 has the following properties: (a) the only source is the origin $u_0 = (0, 0, \dots, 0)$ with $c_0(u_0) = 1$; (b) the sinks are nodes u at length k from the origin, each with $c_0(u) = -1 / \binom{k+d-1}{d-1}$; and (c) $P(c_0) = O(\log n)$, if $d = 2$, and $P(c_0) = O(1/d)$, if $d \geq 3$. To verify the conditions (a), (b), note that for a node $u = (k_1, \dots, k_d)$ with $l = l(u) < k$, the sum of the flows from u to all adjacent nodes at length $l+1$ is $\sum_i (k_i + 1) / ((l+1) \binom{l+d}{d-1})$, which is $(l+d) / ((l+1) \binom{l+d}{d-1}) = 1 / \binom{l+d-1}{d-1}$. Likewise, if $0 < l = l(u) \leq k$, the sum of flows to u from all adjacent nodes at length $l-1$ is $\sum_i k_i / (l \binom{l+d-1}{d-1}) = 1 / \binom{l+d-1}{d-1}$. To verify (c), consider first the case $d = 2$. There are $O(l)$ edges between nodes at length l and $l+1$, each carrying flow $O(1/l)$, for a cumulative contribution of $O(1/l)$ to the power, and hence $P(c_0) = O(\log n)$. For the case $d \geq 3$, the $d \binom{l+d-1}{d-1}$ edges between nodes at length l and $l+1$ carry flow no more than $1 / \binom{l+d}{d-1}$ each, for a total power of $O(1/d)$, the dominant contribution being the edges where $l = 0$.

To prove the upper bound in the theorem, it suffices to prove the resistance bound in a (k, d) mesh from the origin u_0 to an arbitrary vertex $u = (k_1, \dots, k_d)$. We construct three flows c_1, c_2, c_3 , each with power $O(1/d)$ ($O(\log n)$, if $d = 2$), such that the sum of the three flows has a single source $u_0, c(u_0) = 1$, and a single sink u . The result then follows from Lemmas 15 and 16. Flow c_1 is obtained from c_0 by identifying vertices of the form $(0, 0, \dots, 0, k, 0, \dots, 0)$ in the $(k + 1, d)$ mesh with u_0 in the (k, d) mesh; c_3 is the reverse of c_1 except the origin is translated to vertex u ; and c_2 connects the two flows. It can be shown that the claim regarding the power consumptions in c_1, c_2, c_3 holds.

For the lower bound it is immediate that the resistance between the origin and any other vertex is at least $1/2d$ (by shorting all other vertices to one another). For $d = 2$, the resistance between the origin and $(k/2, k/2)$ is seen to be $\Omega(\log n)$, by shorting, for each $l \geq 0$, vertices at length l from the origin. \square

Theorem 13 implies the following upper bounds on the cover times of d -dimensional meshes: $O(n \log^2 n)$ for $d = 2$, and $O(n \log n)$ for $d > 2$. These upper bounds are tight due to recent matching lower bounds of Zuckerman [21]. The upper bounds on cover time were known previously for some cases: e.g., for fixed d [1, 6], and for the hypercube ($d = \log_2 n$) [10]. An advantage of our proofs here is that they are fairly robust under the insertion or deletion of edges since the resistance of a mesh is also robust under these operations.

From Theorems 13, 3 and 4, we have:

Corollary 17: Minimal length universal traversal sequences for an n vertex mesh are given as follows:

1. If G is a two dimensional mesh, then $U(G) = O(n^2 \log n)$.
2. If G is a d -dimensional mesh, $3 \leq d \leq \log_2 n$, then $U(G) = O(n^2 d \log d)$.
3. If G is a hypercube, then $U(G) = O(n^2 \log n \log \log n)$.

We close with a class of graphs for which a tight bound can be obtained on the length of universal traversal sequences covering all members of the class under all labelings. This is the first known

class with this property. In fact, we show several such classes.

Definition: A (k, d, r) mesh sequence is a sequence G_0, G_1, \dots, G_{r-1} of (k, d) meshes connected as follows. Each vertex in G_{2p} , $p \geq 0$, is connected to the corresponding vertex in G_{2p+1} . G_{2p+1} and G_{2p+2} are connected by any nonempty set of criss-crossed edges: any edge (u, v) in G_{2p+1} and the corresponding edge (\hat{u}, \hat{v}) in G_{2p+2} can be deleted, and replaced with the criss-crossed edges (u, \hat{v}) and (\hat{u}, v) .

Theorem 18: For fixed $d \geq 3$ and $r \geq 4$, $r = 0 \pmod 4$, length $\Theta(n^2)$, where $n = rk^d$, is necessary and sufficient for a sequence to be universal for the family of graphs containing all labelings all (k, d, r) mesh sequences.

Proof: (Sketch): Using Theorem 13 and Lemma 14, it can be shown that (k, d, r) mesh sequences have resistance $\Theta(1)$ (for fixed $d \geq 3$, $r \geq 1$), and from a theorem of [4], if $r \geq 4$, $r = 0 \pmod 4$, (k, d, r) meshes can be shown to have UTS length $\Omega(n^2)$. \square

Acknowledgements

We appreciate the helpful comments of David Aldous, Paul Beame, Allan Borodin, Andrei Broder, Don Coppersmith, Se June Hong, Anna Karlin, Ronitt Rubinfeld, J. Laurie Snell, Jim Shearer, Martin Tompa, and David Zuckerman. Peter Doyle deserves special mention for providing an alternative proof for the expander resistance result, plus many valuable pointers and much general encouragement.

References

- [1] David J. Aldous. Random walks on large graphs, 1988. Unpublished Monograph, Berkeley.
- [2] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual Symposium on Foundations of Computer Science*, pages 218–223, San Juan, Puerto Rico, October 1979.

- [3] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [4] A. Borodin, W. L. Ruzzo, and M. Tompa. Lower bounds on the length of universal traversal sequences. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, Seattle, WA, May 1989. ACM.
- [5] A. Z. Broder and A. R. Karlin. Bounds on covering times. In *29th Annual Symposium on Foundations of Computer Science*, pages 479–487, White Plains, NY, October 1988.
- [6] J.T. Cox. Coalescing random walks and voter model consensus times on the torus, 1987. unpublished manuscript, Cornell.
- [7] Peter G. Doyle and J. Laurie Snell. *Random Walks and Electrical Networks*. The Mathematical Association of America, 1984.
- [8] J. N. Franklin. *Matrix Theory*. Prentice-Hall, 1968.
- [9] J. Friedman and N. Pippenger. Expanding graphs contain all small trees. *Combinatorica*, pages 71–76, 1987.
- [10] F. Göbel and A. A. Jagers. Random walks on graphs. *Stochastic Processes and their Applications*, 2:311–336, 1974.
- [11] Mark Jerrum and Alistair Sinclair. Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 235–244, May 1988.
- [12] Jeff D. Kahn, Nathan Linial, Noam Nisan, and Michael E. Saks. On the cover time of random walks in graphs. *Journal of Theoretical Probability*, 1, 1989. To Appear.
- [13] J.G. Kemeny, J. L. Snell, and A.W. Knapp. *Denumerable Markov Chains*. The University Series in Higher Mathematics. Van Nostrand, Princeton, NJ, 1966.
- [14] H. J. Landau and A. M. Odlyzko. Bounds for eigenvalues of certain stochastic matrices. *Linear Algebra and Its Applications*, 38:5–15, 1981.
- [15] Peter Mathews. Covering problems for random walks on spheres and finite groups. Technical Report TR 234, Stanford, Department of Statistics, 1985.
- [16] J. C. Maxwell. *A Treatise on Electricity and Magnetism*. Clarendon, 1918.
- [17] David Peleg and Eli Upfal. Constructing disjoint paths on expander graphs. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 264–273, May 1987.
- [18] Ronitt Rubinfeld. Personal communication. 1989.
- [19] J.L. Synge. The fundamental theorem of electrical networks. *Quarterly of Applied Math.*, 9:113–127, 1951.
- [20] W. Thomson and P.G. Tait. *Treatise on Natural Philosophy*. Cambridge, 1879.
- [21] David I. Zuckerman. A technique for lower bounding the cover time, 1989. Unpublished manuscript, Berkeley.