

THE ELECTRICAL RESISTANCE OF A GRAPH CAPTURES ITS COMMUTE AND COVER TIMES

ASHOK K. CHANDRA, PRABHAKAR RAGHAVAN,
WALTER L. RUZZO, ROMAN SMOLENSKY,
AND PRASOON TIWARI

Dedicated to the memory of Roman Smolensky

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 = $2mR_{st}$. As a corollary, 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, we improve known bounds on cover times for high-degree graphs and expanders, and give new proofs of known results for multi-dimensional meshes. Moreover, resistance seems to provide an intuitively appealing and tractable approach to these problems.

Key words. Random walk, resistance, cover time, commute time.

Subject classifications. 60J15, 68Q99.

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 distributed computation (Broder & Karlin 1989), space-bounded computation (Aleliunas *et al.* 1979,

Borodin *et al.* 1989), time-space tradeoffs (Barnes & Feige 1996, Broder *et al.* 1994, Feige 1993), and in the design of approximation algorithms for some hard combinatorial problems (Dyer *et al.* 1991, Jerrum & Sinclair 1989).

Doyle & Snell (1984) exposed many interesting connections between random walks and electrical network theory, and traced the origins of the topic back into the nineteenth century. Building on this work, we extend the known connections to include several new properties, with an emphasis on questions about cover and commute times.

Doyle & Snell (1984) view an undirected graph as an electrical network in which each edge of the graph is replaced by a unit resistance. As an example of the interplay between electrical and probabilistic notions, their work related the effective resistance between nodes a and b in the electrical network to the probability, in a random walk starting from vertex a , of escaping to vertex b before returning to a . In particular, this probability equals $1/(d(a)R_{ab})$, where R_{ab} is the effective resistance between a and b , and $d(a)$ is the degree of a . 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 ,

$$mR \leq \text{cover time} \leq O(mR \log n). \quad (1.1)$$

The key 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 tied to the random walk.

Prior work in the study of the cover time of graphs has used techniques from Markov chain theory (Aleliunas *et al.* 1979, Göbel & Jagers 1974), from combinatorics (Kahn *et al.* 1989), from linear algebra (Broder & Karlin 1989), and from graph theory (Jerrum & Sinclair 1989). The electrical approach used here provides an intuitive basis for understanding a variety of phenomena about random walks that 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 1.1 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 or addition of a few edges) usually do not change the electrical properties of the graph substantially. Following appearance of a preliminary version of this paper (Chandra *et al.* 1989), Tetali (1991) has extended our ideas to establish a number of new relations between hitting times and effective resistance.

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.* (1979) 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 almost a factor of n . (K_n is a simple example.) Thus, resistance not only gives exact values for commute times, these values may be much better than the estimates provided by Aleliunas *et al.* (1979) (Section 2, Theorem 2.1).

These results also generalize in various ways. For example, we extend the results to characterize the expected costs incurred for walking around a di-

rected cycle, even when edge costs are arbitrary and transition probabilities are nonuniform (Section 2, Theorems 2.2 and 2.3).

Using commute time, we are able to bound the cover time to within a factor of $O(\log n)$, as in (1.1) above. Letting R_{span} be the minimum resistance of a spanning tree of G , we get an alternative upper bound on cover time:

$$\text{cover time} \leq 2mR_{span}. \quad (1.2)$$

For many graphs this provides a better bound than (1.1). 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 given by Aleliunas *et al.* (1979). Again, R_{span} may be much smaller than $n - 1$, as small as $O(1)$, in fact. The graph consisting of a pair of $n/2$ -cliques connected by a single edge provides a simple example: $R_{span} = O(1)$, hence cover time is $O(n^2)$, which is better than the bounds given by Aleliunas *et al.* (1979), or (1.1) (Section 2, Theorem 2.4).

For d -regular graphs, the Aleliunas *et al.* bound for cover time is $O(dn^2)$. Kahn *et al.* (1989) 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 (1.2).

Kahn *et al.* (1989) also give examples, for any $d \leq \lfloor n/2 \rfloor - 1$, of n -vertex, d -regular graphs with maximum resistance $\Omega(n/d)$, and hence by (1.1) 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 $\lfloor n/2 \rfloor - 1$ to $n - 1$. Much to our surprise, this is *not* the case—there is a sharp threshold at $d = \lfloor n/2 \rfloor$. We show that in going from $d = \lfloor n/2 \rfloor - 1$ to $\lfloor n/2 \rfloor$ the maximum resistance drops from $\Omega(1)$ to $O(1/n)$, hence by (1.1) the cover time drops from $\Omega(n^2)$ to $O(n \log n)$ (where it remains for all $d \geq \lfloor n/2 \rfloor$). This result has a very simple and intuitive proof (Section 3, Theorem 3.3).

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 & Karlin (1989) as corollaries. Again, we show that (1.1) gives 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 4.2).

One interesting application of our approach is to the cover time for d -regular expander graphs. Using the eigenvalue approach, Broder & Karlin (1989) showed that such graphs have cover time $O((n \log n)/(1 - \lambda_2)) = O(dn \log n)$, where $\lambda_2 = 1 - \sigma_2/d$. 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 5, Theorem 5.2).

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 $\Theta(1/d)$. Random walks on meshes have been previously considered by many authors, including some studies of cover times by Aldous (1983, 1993), Cox (1989), and Göbel & Jagers (1974) (the latter only for $d = \log_2 n$, i.e., hypercubes). Although our conclusions about cover times of meshes were previously known, our approach is novel and potentially illuminating. For example, the resistance of a graph will generally not be changed significantly by the insertion or removal of a few edges, so our results naturally suggest bounds on cover times for “imperfect” meshes, which are more difficult to treat by more classical analytical techniques (Section 6).

Our last application couples resistance-based commute bounds to the proof technique of Aleliunas *et al.* (1979) to derive new upper bounds for universal traversal sequences, namely $O(mR \log(ng))$, where g is the number of labeled graphs in the family under consideration. This gives improved upper bounds for universal traversal sequences for many classes of graphs, including dense graphs and expanders (Section 2, Theorem 2.6). 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 connected 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 $\mathcal{N}(G)$, i.e., the voltage induced between u and v by passing a current flow of one ampere between them.

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 2.1. *For any two vertices u and v in G , the commute time C_{uv} equals $2mR_{uv}$.*

PROOF. For any x in V , let $d(x)$ denote the degree of x in G . 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

$$H_{uv} = \phi_{uv} \quad \forall u \in V. \quad (2.3)$$

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

$$d(u) = \sum_{w \in N(u)} (\phi_{uw} - \phi_{vw}) \quad \forall u \in V - \{v\}. \quad (2.4)$$

By elementary probability theory,

$$H_{uv} = \sum_{w \in N(u)} \frac{1}{d(u)} (1 + H_{vw}) \quad \forall u \in V - \{v\}. \quad (2.5)$$

Equations (2.4) and (2.5) are both linear systems with unique solutions; furthermore, they are identical if we identify ϕ_{uv} in (2.4) with H_{uv} in (2.5). This proves (2.3). 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 $2m$ units of current are injected at u , and $d(v)$ units are removed from all nodes $v \in V$. 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 injected into and removed from all nodes except u and v cancel, resulting in C_{uv} being the voltage between u and v when $2m$ 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, P. G. Doyle, A. R. Karlin, and J. L. Snell all have derived alternative proofs of Theorem 2.1 using similar methods from renewal theory. We sketch this alternate proof below. For any vertex u , let T be a random time at which the walk from u returns to u . By standard techniques from renewal theory (e.g., Ross 1989, Prop. 7.4.1), for any vertex x (or directed edge (x, y)) the expected number of visits to x (traversals of (x, y)) before T is exactly $\mathcal{E}[T]$ times the steady-state probability of visiting x (traversing (x, y) , respectively). Choosing T to be the time of first return to u after visiting v , and $x = u$, and noting that the steady-state probability of visiting u is $d(u)/(2m)$, we conclude that the expected number of returns to u during a commute to v is $\mathcal{E}[T]d(u)/(2m) = C_{uv}d(u)/(2m)$. The result mentioned in our introduction from Doyle & Snell (1984, Section 3.3) implies that during a random commute from u to v , the expected number of returns to u is $d(u) \cdot R_{uv}$. Combining these expressions yields the result.

Although Theorem 2.1 suffices for most of our applications, it is interesting to note that it easily generalizes to walks on graphs with self-loops, and with non-uniform transition probabilities and costs. With each ordered pair of vertices $(u, v) \in V \times V$ we associate a positive real *resistance* $r_{uv} = r_{vu}$. Non-edges are represented by infinite resistances. Additionally, let each directed edge have a real *cost* f_{uv} . (We do not require that $f_{uv} = f_{vu}$, unless, of course, $u = v$.) We now consider a random walk on G defined by the following discrete-time process: when at a vertex $u \in V$, step to vertex v with probability inversely proportional to r_{uv} (the resistance of edge $\{u, v\}$), i.e., with probability

$$p_{uv} = \frac{1/r_{uv}}{\sum_{w \in V} 1/r_{uw}}.$$

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 the standard random walk on a graph is the special case where all costs are 1, all edges have resistance 1, and all non-edges have infinite resistance.

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 pair of vertices $\{u, v\}$ in V ,

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 2.2. *Let $F = \sum_{(x,y) \in V \times V} f_{xy}/r_{xy}$. For any two vertices u, v in G , the commute cost $C_{uv}^f = F \cdot R_{uv}$.*

PROOF. The proof is identical to that of Theorem 2.1, except that the current injected into node x is $f_x = \sum_{y \in V} f_{xy}/r_{xy}$ for all $x \in V$. □

Theorem 2.1 is obviously a corollary when there are no self-loops and all edges have resistance and cost 1 ($F = 2m$).

Aleliunas *et al.* (1979) 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 2.2 by setting all resistances to one, and all costs to zero, except for an arbitrary directed edge, which is given cost one. Further, we find that $\tau = R_{uv}$.

For non-unit resistances, Doyle & Snell (1984) 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 2.2 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 & Karlin (1989).

We can generalize Theorem 2.2 to the expected cost of a trip around a directed cycle (a commute being the special case where the cycle has length two). Let $D = (v_1, v_2, \dots, v_{k+1} = v_1)$ be a sequence of vertices in G . Let C_D^f denote the expected cost of a random walk starting at v_1 and stopping upon returning to v_1 after visiting v_2, \dots, v_k in order.

THEOREM 2.3. *Let F be defined as in Theorem 2.2, and let $R_D = \sum_{i=1}^k R_{v_i, v_{i+1}}$. Then $C_D^f = F \cdot R_D/2$.*

PROOF. Let $D' = (v_1 = v_{k+1}, v_k, v_{k-1}, \dots, v_1)$. Now

$$C_D^f + C_{D'}^f = \sum_{i=1}^k H_{v_i, v_{i+1}}^f + \sum_{i=k}^1 H_{v_{i+1}, v_i}^f$$

The right-hand side can be re-written as

$$\sum_{i=1}^k C_{v_i, v_{i+1}}^f = F \cdot \sum_{i=1}^k R_{v_i, v_{i+1}} = F \cdot R_D.$$

Since the Markov chain corresponding to our walk is reversible, the probability of walking any cycle is equal to the probability of its reversal (Doyle & Snell 1984, Exercise 3.1.3), and so the expected lengths of random walks traversing D and D' are equal. Applying the renewal theory result cited below the proof of Theorem 2.1, we see that edge costs accrue at the equilibrium rate during either walk, hence $C_D^f = C_{D'}^f$, which yields the result. \square

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. Known relationships between cover time and hitting time allow us to frame a nearly tight relationship between cover time and resistance.

Let $R = \max_{u,v \in V} R_{uv}$. Let $\mathcal{N}'(G)$ be an edge-weighted complete graph having a vertex u' for every vertex u in V , and having an edge $\{u', v'\}$ of weight R_{uv} for each pair of (not necessarily adjacent) vertices u, v in V . Let R^* be the weight of a 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 2.4.

$$mR \leq C_G \leq 2m \cdot \min(R(1 + \ln n), R^*)$$

PROOF. The lower bound follows from Theorem 2.1, the fact that there exist vertices u, v such that $R = R_{uv}$, and the fact that $\max(H_{uv}, H_{vu}) \geq C_{uv}/2$. Matthews (1988) has shown that the cover time is at most Hh_n , where $H = \max_{uv} H_{uv}$, and h_n is the n^{th} harmonic number, $h_n = \sum_{i=1}^n 1/i \leq 1 + \ln n$. 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 2.6 below.) The proof of the second upper bound follows directly from the spanning-tree argument of Aleliunas *et al.* (1979), which is a special case of Theorem 2.3. \square

Let R_{span} be the minimum, over all spanning trees T of G , of the sum of the effective resistances of T 's edges. R_{span} is often easier to determine than R^* . However, since $R^* \leq R_{span}$, we also have the following corollary.

COROLLARY 2.5. $C_G \leq 2mR_{span}$.

Note that the bounds in Theorem 2.4 cannot in general be improved by more than constant factors; the upper bounds are tight for the complete graph and the chain, respectively, 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 for all the labeled graphs in \mathcal{G} . (See Aleliunas *et al.* 1979 or Borodin *et al.* 1992 for definitions.) Let $R(\mathcal{G})$ denote the maximum resistance between any pair of vertices in any graph in \mathcal{G} .

THEOREM 2.6. $U(\mathcal{G}) \leq (4 + o(1))mR(\mathcal{G})\log_2(n|\mathcal{G}|)$.

PROOF. The proof is by a probabilistic argument similar to that in Aleliunas *et al.* (1979). Given a labeled graph $G \in \mathcal{G}$, let v be a vertex of G , and let $t = \lfloor 4mR(\mathcal{G}) \rfloor + 1$. Consider a random walk of length kt , divided into k “epochs” of length t . Let A_i be the event that the walk fails to visit v during the i^{th} epoch, $1 \leq i \leq k$. Then, for $1 \leq i \leq k$, the probability of event A_i is less than $1/2$ by Theorem 2.1 and Markov’s inequality, regardless of the vertex of G at which the epoch began. If these events were independent, the probability that the full walk would fail to visit v would be less than 2^{-k} . Unfortunately, the events are not independent, since obviously epoch i starts at the vertex where epoch $i - 1$ ends. However, we will show by another method that the 2^{-k} upper bound is still valid.

Let $p(a, b)$ be the probability that a walk of length t starting at vertex a ends at vertex b and fails to visit v . Note that for each fixed $b \neq v$, given that epoch $i - 1$ ends at vertex b (and consequently epoch i starts at b) the events “fails to visit v during epoch $i - 1$ ” and “fails to visit v during epoch i ” are independent, by the Markov property of random walks. So, assuming epoch $k - 1$ starts on vertex $a \neq v$, we see that the walk fails to visit v during the last *two* epochs with probability

$$\begin{aligned} Pr(A_{k-1} \wedge A_k \mid \text{epoch } k-1 \text{ starts at vertex } a \neq v) &= \sum_{b, c \neq v} p(a, b)p(b, c) \\ &= \sum_{b \neq v} p(a, b) \sum_{c \neq v} p(b, c) \\ &< \sum_{b \neq v} p(a, b)(1/2) \\ &< (1/2)(1/2) = 1/4. \end{aligned}$$

Proceeding similarly, we can show that the probability that v was not visited during any of the k epochs is less than 2^{-k} . Choosing $k = \lceil \log_2(n|\mathcal{G}|) \rceil$, we see that the probability of avoiding v is less than $(n|\mathcal{G}|)^{-1}$. 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 techniques of Matthews (1988) can be used to give a different proof of Theorem 2.6.

We study one final random variable associated with a random walk on G : let C_E denote the expected number of steps to traverse every *edge* of G at least once, taking the maximum of the expectations over all starting vertices. By a traversal we mean that each edge can be traversed in either direction; a simple extension of the method applies to the case when we require each edge to be traversed in both directions.

Let G' be the graph derived from G by inserting a new vertex in the middle of each edge of G . More precisely, G' is defined as follows: there is a vertex v in G' corresponding to each vertex v in G ; call these *real vertices*. In addition, there is a vertex uv in G' corresponding to each edge $\{u, v\}$ in G . Thus G' has $m + n$ vertices in all. The edges of G' are as follows: for each edge $\{u, v\}$ in G there is a pair of edges $\{u, uv\}$ and $\{v, uv\}$ in G' .

We draw a correspondence between a walk on G traversing all its edges and a certain walk on G' visiting all its vertices. Consider a random walk on G' . Each time we take a step out of a real vertex, say v , we proceed to a vertex vw that is not real; from there, we proceed to another real vertex w with probability $1/2$, or return to v with probability $1/2$. Call a pair of such steps *useful* if we proceed to a new real vertex such as w , and *wasted* if we return to v . Thus, each pair of steps is useful with probability $1/2$.

We consider a modified random walk on G , one in which at each step the walk may choose to *idle* in its present vertex with probability $1/2$; clearly in this modified walk the expected number of steps (including idle steps) to traverse all the edges is $2C_E$. We now draw the obvious correspondence between this modified walk and walks on G' starting at the same (real) vertex: each wasted pair of steps in the walk on G' corresponds to an idle step in the modified walk on G . A useful pair of steps $v - vw - w$ in the walk on G' corresponds to the traversal of edge vw in G . Conditioned by the probability that a pair of steps is useful, the probability distribution on the real vertex w we reach is the uniform distribution on the neighbors of v in G . In a walk on G' , if every non-real vertex is visited during a useful step pair, then every edge of

G has been traversed in the corresponding modified walk on G . Thus every sequence of idle and edge-traversal steps in the modified walk on G has the same probability measure and number of steps as the corresponding sequence of wasted and useful step pairs in the walk on G' . From this we can prove the following result.

THEOREM 2.7. C_E is $O((m+n)(R+1)\log n)$.

PROOF. The proof is very similar to that of Theorem 2.6. The maximum resistance in G' is $O(R+1)$, and the maximum commute time in it is $O((m+n)(R+1))$. For a fixed non-real vertex vw , the probability that it is not visited in an epoch of length $O((m+n)(R+1))$ is a constant, and so is the probability that it is not visited during a useful pair of steps. Thus, by calculations similar to those in the proof of Theorem 2.6, the probability that a walk of length $k((m+n)(R+1)\log n)$ does not cover all edges declines exponentially with k . \square

Note that C_E is within a factor of $\log n$ of C_G , except perhaps when $R = o(1)$. Zuckerman (1991) shows that C_E , like C_G , is $O(mn)$ for all graphs.

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.

PROPOSITION 3.1. (Rayleigh's "Short/Cut" Principle, Doyle & Snell 1984, Maxwell 1918.) *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 3.2. *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 Theorem 2.4, the cover time for this graph is $\Omega(n^2)$; this bound is tight by the results of Kahn *et al.* (1989). A similar construction works for odd n and $d \leq \lfloor n/2 \rfloor - 1$.

When $d = \lfloor n/2 \rfloor$, the situation changes radically. Intuitively, one cannot 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 3.3. *For any n -vertex graph G with minimum degree $d \geq \lfloor n/2 \rfloor$, $R \leq 4/d = O(1/n)$. Hence $C_G = O(n \log n)$.*

PROOF. The key point is to show that there are d edge-disjoint paths of length at most 4 between any pair of vertices. The result then follows by application of Lemma 3.2. 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 both s and t . Then there are at least $j = d - k - k'$ vertices which are adjacent to s but not to t , and vice versa. Choose any j of each, and call them s_1, \dots, s_j and t_1, \dots, t_j , respectively. Let k'' be the size of a maximum matching between the s_i 's and the t_i 's, and without loss of generality 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 is 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 3.2. \square

Theorem 3.3, when combined with Theorem 2.4, shows a sharp threshold in cover time at minimum degree $d = \lfloor n/2 \rfloor$. Specifically, the cover time may be $\Omega(n^2)$ when $d = \lfloor n/2 \rfloor - 1$, but drops to $O(n \log n)$ when $d = \lfloor n/2 \rfloor$. Applying Theorem 2.6 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, lower bounds for universal traversal sequences (Borodin *et al.* 1992) are $\Omega(n^4)$ for linear $d \leq n/3 - 2$. Thus, the 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

Let G be a connected graph with vertices numbered $1, 2, \dots, n$, and a positive real resistance r_{ij} associated with each edge. The *conductance* of edge $\{i, j\}$ is defined to be the reciprocal of its resistance: $1/r_{ij}$. Let $\mathcal{N}(G)$ be the corresponding electrical network, as defined immediately above Theorem 2.2. Let $d(i)$ be the sum of the conductances connected to node i , and let \mathbf{D} be the diagonal matrix whose i -th diagonal entry is $d(i)$. Let \mathbf{A} be the matrix whose ij -th entry is the conductance on the edge from i to j . 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, Franklin 1968.) 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 (Franklin 1968) zero is also the smallest eigenvalue. Using the same theorem and the fact that G is connected, it can be shown that zero is an eigenvalue of multiplicity one. Define $\sigma(G)$ to be the second smallest eigenvalue of \mathbf{K} .

It is worth pointing out that if G is a graph with unit resistances, then $d(i)$ is the degree of vertex i in G , \mathbf{A} is the adjacency matrix of G , and \mathbf{K} is the Laplacian.

We will use the following inner product in this section.

DEFINITION 4.1. 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 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 .

THEOREM 4.2. *If G is a connected graph on n vertices, then*

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

PROOF. For any distinct pair of vertices s and t , 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 1 and all other components are 0. Then, as in the proof of Theorem 2.1, $\mathbf{K}\mathbf{v} = \mathbf{c}$, and therefore

$$\mathbf{v} = \delta \mathbf{u}_1 + \sum_{k=2}^n \frac{\alpha_k}{\sigma_k} \mathbf{u}_k, \tag{4.6}$$

where δ is \sqrt{n} times the average voltage in the network, and $\alpha_k = (\mathbf{c}, \mathbf{u}_k) = u_{sk} - u_{tk}$. 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 4.6, we get

$$\begin{aligned} R &= v_s - v_t = \sum_{k=2}^n \frac{\alpha_k(u_{sk} - u_{tk})}{\sigma_k} \\ &= \sum_{k=2}^n \frac{\alpha_k^2}{\sigma_k} \leq \frac{1}{\sigma(G)} \sum_{k=2}^n \alpha_k^2 = \frac{2}{\sigma(G)}. \end{aligned}$$

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 exists 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 $R \geq 1/(n\sigma(G))$. □

Theorem 2.4 immediately implies the following corollary.

COROLLARY 4.3. *If G is a graph with unit resistances and unit costs, then $C_G \leq (4 + o(1))m \ln n / \sigma(G)$.*

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 descending order whereas the σ_i 's are arranged in ascending order. 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 4.4. *Let G be a connected graph with minimum and maximum degrees given by d_{min} and d_{max} , respectively. Then, for all $1 \leq k \leq n$,*

$$(1 - \lambda_k)d_{min} \leq \sigma_k \leq (1 - \lambda_k)d_{max}.$$

PROOF. Coppersmith devised the following elegant proof of this lemma.

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 (Franklin 1968) gives the following expressions for the eigenvalues:

$$\alpha_i = \min_{\mathbf{x} \perp \{\mathbf{v}_{i+1}, \mathbf{v}_{i+2}, \dots, \mathbf{v}_n\}} \frac{\mathbf{x}^T \mathbf{B} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}, \tag{4.7}$$

$$= \max_{\mathbf{x} \perp \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{i-1}\}} \frac{\mathbf{x}^T \mathbf{B} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}. \tag{4.8}$$

Note for later use that $\mathbf{x} \perp \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{i-1}\}$ if and only if \mathbf{x} is in the span of $\{\mathbf{v}_i, \mathbf{v}_{i+1}, \dots, \mathbf{v}_n\}$.

With the \mathbf{u}_i 's as before, consider the set of $n + 1$ vectors

$$\{\mathbf{u}_k, \mathbf{u}_{k+1}, \dots, \mathbf{u}_n, \mathbf{D}^{-\frac{1}{2}}\mathbf{w}_1, \mathbf{D}^{-\frac{1}{2}}\mathbf{w}_2, \dots, \mathbf{D}^{-\frac{1}{2}}\mathbf{w}_k\}.$$

Since there are more than n vectors in this set, they are linearly dependent, i.e., there exist constants $a_k, a_{k+1}, \dots, a_n, b_1, b_2, \dots, b_k$, not all zero, such that

$$\sum_{i=k}^n a_i \mathbf{u}_i = \mathbf{D}^{-\frac{1}{2}} \sum_{i=1}^k b_i \mathbf{w}_i. \tag{4.9}$$

Let us denote the left hand side of this equation by \mathbf{z} . If $\mathbf{z} = 0$, then $a_k = a_{k+1} = \dots = a_n = 0$, and $b_1 = b_2 = \dots = b_k = 0$ because each of the two sets of vectors $\{\mathbf{u}_k, \mathbf{u}_{k+1}, \dots, \mathbf{u}_n\}$ and $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_k\}$ are independent. Therefore, $\mathbf{z} \neq 0$, and, without loss of generality, we may assume that \mathbf{z} is a unit vector. Equation 4.7 implies that

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

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 $\sigma_k \leq \mathbf{z}^T\mathbf{D}\mathbf{z}(1 - \lambda_k)$. Finally, observe that $\mathbf{z}^T\mathbf{D}\mathbf{z} \leq d_{max}$, which establishes the upper bound on σ_k 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 $\{\mathbf{w}_k, \mathbf{w}_{k+1}, \dots, \mathbf{w}_n, \mathbf{D}^{\frac{1}{2}}\mathbf{u}_1, \mathbf{D}^{\frac{1}{2}}\mathbf{u}_2, \dots, \mathbf{D}^{\frac{1}{2}}\mathbf{u}_k\}$ and using Equation 4.8 instead of Equation 4.7. □

The following example will be useful in showing where the inequalities in Theorem 4.2 are tight.

DEFINITION 4.5. 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 $(k_1, \dots, k_{i-1}, k_i \pm 1 \text{ mod } n_i, k_{i+1}, \dots, k_d)$, 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 4.6. *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 $\sum_{i=1}^d (\omega_i^{k_i} + \omega_i^{-k_i})$. □

COROLLARY 4.7. *If G is the $n_1 \times n_2 \times \cdots \times n_d$ 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 4.2. The *lower* bound on resistance given by Theorem 4.2 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 4.7 $\sigma(n\text{-cycle}) \approx (\frac{2\pi}{n})^2$. The *upper* bound on resistance given by Theorem 4.2 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 4.2, except perhaps the constant factor in the lower bound, for *all* graphs. On the other hand, both the inequalities in Theorem 4.2 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 4.2 also improves a bound due to Landau & Odlyzko (1981) (and Corollary 17 of Broder & Karlin 1989). Landau & Odlyzko proved that $(1 - \lambda_2) \geq 1/((d_{max} + 1)\Delta n)$ where d_{max} and Δ are the maximum degree and the diameter of G , respectively. Using the resistance bound from Theorem 4.2, and Lemma 4.4, we get $(1 - \lambda_2) \geq 1/(d_{max} R n)$. This is an improvement because $\Delta \geq R$, and may be a large improvement. For example, for K_n , $\Delta = 1$ and $R = \frac{2}{n}$.

Some upper bounds on cover times due to Broder & Karlin (1989) are implied as a consequence of Theorem 4.2. For example, Corollary 4.3 and Lemma 4.4 imply that $C_G \leq ((4 + o(1))m \ln n)/(d_{min}(1 - \lambda_2))$. For most graphs, this is stronger than Corollary 8 of Broder & Karlin (1989), which states that $C_G \leq (1 + o(1))n^2 \ln n/(1 - \lambda_2)$. For example, note that $4m/d_{min} < n^2$ whenever $d_{min} \geq 2$ or $d_{max} < n/2$.

Finally, Theorem 4.2 also implies that the resistance between any pair of vertices in any family of *bounded* degree expander graphs (see the next section, or Alon 1986) 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 Broder & Karlin (1989) nor Theorem 4.2 yield good bounds on the cover time of these graphs.

5. Expanders

We will use the following definition of expanders, also used by Broder & Karlin (1989).

DEFINITION 5.1. *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 (1986) 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, Rubinfeld (1990) has shown a result analogous to our Theorem 5.2 for graphs which are “expanders” according to the definition of Peleg & Upfal (1989), giving further evidence that the basic result of this section is reasonably insensitive to variations in the definition.

Alon (1986) has shown that if G is an (n, d, α) -expander, then $\sigma(G) \geq \alpha^2/(4+2\alpha^2)$, hence by Theorem 4.2, $R \leq 4(2+\alpha^2)/\alpha^2$. The main result of this section sharpens this estimate, reducing it by a factor of order d . For large d , this considerably improves the bounds of Broder & Karlin (1989) on the cover time of these graphs.

THEOREM 5.2. *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 $\alpha^2(d+1)/(8(1+\alpha/2)(1+\alpha))$, implying $R \leq (8(1+\alpha/2)(1+\alpha))/(\alpha^2(d+1))$, 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

$$T_k = \{a \mid \text{node } a \text{ of } \mathcal{N}(G) \text{ has voltage} < v_k\},$$

$$S_k = \{a \mid \text{node } a \text{ of } \mathcal{N}(G) \text{ has voltage} > 1 - v_k\},$$

$$t_k = |T_k|,$$

$$s_k = |S_k|.$$

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

First we make the following claim.

CLAIM 5.3. $t_0 \geq (d + 1)(1 + \alpha/2)/(1 + \alpha)$, 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

$$\begin{aligned} (d - (t_0 - 1))v_0 &> \left(d - \left((d + 1) \frac{(1 + \alpha/2)}{(1 + \alpha)} - 1 \right) \right) \frac{\alpha}{4(1 + \alpha/2)} \\ &= (d + 1) \left(\frac{\alpha}{2(1 + \alpha)} \right) \frac{\alpha}{4(1 + \alpha/2)} \\ &= \frac{\alpha^2(d + 1)}{8(1 + \alpha/2)(1 + \alpha)}, \end{aligned}$$

contradicting the assumption that the current is less than the latter 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-1} would be too large. More precisely, if $t_k < (1 + \alpha/2)t_{k-1}$, then the current will be greater than

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

again contradicting the assumption that the current is less than the latter 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 $\alpha^2(d+1)/(8(1 + \alpha/2)(1 + \alpha))$. □

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

We will briefly sketch an alternative proof of Theorem 5.2. 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 3.3 and 5.2) 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 (1988) 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 $(\Theta(d))$ 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.

Rubinfeld's proof (1990) uses yet a third technique: she applies a result of Friedman & Pippenger (1987) 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 (toroidal) mesh of side k , hence $n = k^d$ vertices.

Resistance of infinite meshes has been previously considered. In particular, it is the focus of a portion of Doyle and Snell's monograph (1984). 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. We give a direct proof, although the results below largely follow from Theorems 2.1

and 2.4 and known results on hitting and cover times for meshes (c.f. Aldous 1983, 1993, Cox 1989, and Göbel & Jagers 1974, for example). Our proof technique, which is similar to (Doyle & Snell 1984, Section 8.7), is of interest, and, unlike the more “classical” analysis of hitting time for meshes, can likely be extended to less symmetrical graphs which remain “mesh-like.”

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

THEOREM 6.1. *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.

PROPOSITION 6.2. *For any three vertices u, v, w in G ,*

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

DEFINITION 6.3. *Given an electrical network (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 voltage law, i.e., for any directed cycle $u_0, u_1, \dots, u_{k-1}, u_0$, $\sum_{i=0}^{k-1} c(u_i, u_{i+1 \bmod k}) \cdot r(u_i, u_{i+1 \bmod k}) = 0$.*

PROPOSITION 6.4. *(The Minimum Power Principle, Synge 1951; also known as Thomson’s Principle, Thomson & Tait 1879, Doyle & Snell 1984, Section 3.5.) 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 $R_{u,v} \leq P(c)$, with equality when the flow is a current flow.*

Also, for a flow c with one source u and more than one sink, $R_{u,v} \leq P(c)$ holds for all sinks v , provided that $c(v) = -1$. This is easy to see by noting that the flow to all sinks except v can be returned to the source without increasing the flow through any edge.

LEMMA 6.5. 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. □

PROOF (OF THEOREM 6.1). 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 $u = (l_1, \dots, l_d)$. 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 6.4 and 6.5.

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: for $0 \leq k_i < k$,

$$\begin{aligned} c_1((k_1, \dots, k_d), (k_1, k_2, \dots, (k_i + 1) \bmod k, \dots, k_d)) = \\ c_0((k_1, \dots, k_d), (k_1, k_2, \dots, (k_i + 1), \dots, k_d)). \end{aligned}$$

Since c_0 has nonzero flow only incident on nodes at length no more than k , $P(c_1) = P(c_0)$. Flow c_1 has a single source at the origin u_0 with $c_1(u_0) =$

$1 - d/\binom{k+d-1}{d-1}$. Flow c_2 is defined to be the reverse of c_1 , except the origin is translated to vertex u : $c_2(v, w) = c_1(w + u, v + u)$ where $+$ is understood to be component-wise addition mod k . Clearly $P(c_2) = P(c_1)$. Finally, flow c_3 connects up c_1 and c_2 as follows. Let $P = (u_0, u_1, \dots, u_r = u)$ be the path from the origin u_0 to u where u_i is at length i , and if (u_i, u_{i+1}) is an edge along dimension d' then (u_{i+1}, u_{i+2}) is an edge along dimension $\geq d'$. The only positive flows in c_3 are as follows. For any node v at length k there is a flow of $1/\binom{k+d-1}{d-1}$ along path $P + v = (u_0 + v, \dots, u_r + v)$, i.e., $c_3(v + u_i, v + u_{i+1}) = 1/\binom{k+d-1}{d-1}$, and a flow of $d/\binom{k+d-1}{d-1}$ along $P + u_0 = P$. Note that all the above paths $P + v, P + u_0$ are edge disjoint. Since $r < kd$ and the number of nodes at length k is no more than $\binom{k+d-1}{d-1}$, we have $P(c_3) \leq kd/\binom{k+d-1}{d-1} + kd^3/\binom{k+d-1}{d-1}^2 = O(1/d)$ for $k, d \geq 3$ (and is $O(1)$ for $d = 2$). Finally, it can be checked that flow $c = c_1 + c_2 + c_3$ indeed has a single source u_0 with $c(u_0) = 1$ and a single sink at u . This completes the proof of the upper bound.

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 6.1 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 matching lower bounds of Zuckerman (1992). The upper bounds on cover time were known previously for some cases: e.g., Göbel & Jagers (1974) for $d = \log_2 n$ (the hypercube), and Aldous (1983, 1993), and Cox (1989) for general d . An advantage of our proofs 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. For example, from the proof of Theorem 6.1, it can be seen that the resistance $R_{u_0, u}$ between two arbitrary nodes is not much affected by the deletion of several edges, provided they are not too close to u_0 or to u —specifically, an edge whose endpoints are at distance at least l from both u_0 and u , carries a flow of no more than $2/\binom{l+d}{d-1} + d/\binom{k+d-1}{d-1}$, and edges carrying total flow of $(1 - c)$ may be deleted without increasing the upper bound by more than a factor of $1/c$.

From Theorems 6.1, 2.4 and 2.6, we have:

COROLLARY 6.6. *Minimal length universal traversal sequences for the family of labeled graphs defined by an n vertex mesh under all labelings 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 $U(G) = O(n^2 \log n \log \log n)$.

Note that we do not have matching lower bounds on universal traversal sequence lengths for meshes, so it remains possible that these lengths are the same for two- and three-dimensional meshes, even though the resistances differ.

We close with a class of graphs for which we do have a tight bound 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.

Fix $d \geq 3$. Let \mathcal{H} be any family of d -regular labeled graphs such that (1) \mathcal{H} contains $2^{O(n)}$ members with n vertices, and (2) each member of \mathcal{H} has a set of $\Omega(dn)$ edges, called *switchable* edges, whose removal leaves a (connected) graph with $R = O(1)$. Let \mathcal{G} be the set of all graphs formed in the following way from some $H \in \mathcal{H}$ and some non-empty subset X of H 's switchable edges: join two copies H_1 and H_2 of H via the criss-crossed edges dictated by X , i.e., for each edge (u, v) in X , delete the corresponding edges (u_1, v_1) and (u_2, v_2) from H_1 and H_2 , respectively, and replace them with the criss-crossed edges (u_1, v_2) and (u_2, v_1) .

Families \mathcal{H} with the properties required for this construction exist. For example, let A be a (k, d') -mesh, $d' \geq 3$, and let B be an arbitrary $n' = k^{d'}$ vertex d' -regular graph. Connect A to B by an arbitrary matching. The resulting graph H is an n -vertex, d -regular graph, ($n = 2n'$, $d = d' + 1$) with the desired properties (Theorem 6.1), when the switchable edges are taken to be the edges of B . Alternatively, A could be a d' -regular expander (Theorem 5.2).

THEOREM 6.7. *Let \mathcal{G} be any family of labeled, d -regular graphs constructed as described above. The shortest universal sequences for \mathcal{G} have length $\Theta(n^2)$.*

PROOF. By the Short/Cut Principle, each of the $2^{O(n)}$ n -vertex members of \mathcal{G} has resistance $O(1)$, hence by Theorem 2.6, \mathcal{G} has universal traversal sequences of length $O(n^2)$. From techniques of Borodin *et al.* (1992), \mathcal{G} 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, Jim Shearer, J. Laurie Snell, 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. Thanks also to the referees for careful reading and useful suggestions.

Our research was supported in part by the National Science Foundation under grant CCR-8703196; a portion of this work was done while the third author was visiting the IBM T.J. Watson Research Center.

References

D. J. ALDOUS, On the time taken by random walks on finite groups to visit every state. *Z. Wahrsch. Verw. Gebiete* **62**(3) (1983), 361–374.

D. J. ALDOUS, *Reversible Markov chains and random walks on graphs*. 1993. Draft, University of California, Berkeley.

R. ALELIUNAS, R. M. KARP, R. J. LIPTON, L. LOVÁSZ, AND C. W. RACKOFF, Random walks, universal traversal sequences, and the complexity of maze problems. In *Proc. 20th Ann. IEEE Symp. Found. Comput. Sci.*, San Juan, Puerto Rico, 1979, IEEE, 218–223.

N. ALON, Eigenvalues and expanders. *Combinatorica* **6**(2) (1986), 83–96.

G. BARNES AND U. FEIGE, Short random walks on graphs. *SIAM J. Disc. Math.* **9**(1) (1996), 19–28.

A. BORODIN, S. A. COOK, P. W. DYMOND, W. L. RUZZO, AND M. TOMPA, Two applications of inductive counting for complementation problems. *SIAM J. Comput.* **18**(3) (1989), 559–578. See also 18(6): 1283, December 1989.

A. BORODIN, W. L. RUZZO, AND M. TOMPA, Lower bounds on the length of universal traversal sequences. *J. Comput. System Sci.* **45**(2) (1992), 180–203.

A. Z. BRODER AND A. R. KARLIN, Bounds on the cover time. *J. Theoret. Probability* **2**(1) (1989), 101–120.

A. Z. BRODER, A. R. KARLIN, P. RAGHAVAN, AND E. UPFAL, Trading space for time in undirected s - t connectivity. *SIAM J. Comput.* **23**(2) (1994), 324–334.

A. K. CHANDRA, P. RAGHAVAN, W. L. RUZZO, R. SMOLENSKY, AND P. TIWARI, The electrical resistance of a graph captures its commute and cover times. In *Proc. Twenty-first Ann. ACM Symp. Theor. Comput.*, Seattle, WA, 1989, 574–586.

J. T. COX, Coalescing random walks and voter model consensus times on the torus in Z^d . *The Annals of Probability* **17**(4) (1989), 1333–1366.

P. G. DOYLE, Personal communication, 1988.

P. G. DOYLE AND J. L. SNELL, *Random Walks and Electrical Networks*. The Mathematical Association of America, 1984.

M. DYER, A. M. FRIEZE, AND R. KANNAN, A random polynomial-time algorithm for approximating the volume of convex bodies. *J. Assoc. Comput. Mach.* **38**(1) (1991), 1–17.

U. FEIGE, A randomized time-space tradeoff of $\tilde{O}(m\hat{R})$ for USTCON. In *Proc. 34th Ann. Symp. Found. Comput. Sci.*, Palo Alto, CA, 1993, IEEE, 238–246.

J. N. FRANKLIN, *Matrix Theory*. Prentice-Hall, 1968.

J. FRIEDMAN AND N. J. PIPPENGER, Expanding graphs contain all small trees. *Combinatorica* (1987), 71–76.

F. GÖBEL AND A. A. JAGERS, Random walks on graphs. *Stochastic Processes and their Applications* **2** (1974), 311–336.

M. JERRUM AND A. SINCLAIR, Approximating the permanent. *SIAM J. Comput.* **18**(6) (1989), 1149–1178.

J. D. KAHN, N. LINIAL, N. NISAN, AND M. E. SAKS, On the cover time of random walks on graphs. *J. Theoret. Probability* **2**(1) (1989), 121–128.

H. J. LANDAU AND A. M. ODLYZKO, Bounds for eigenvalues of certain stochastic matrices. *Linear Algebra and Its Applications* **38** (1981), 5–15.

P. MATTHEWS, Covering problems for Brownian motion on spheres. *The Annals of Probability* **16**(1) (1988), 189–199.

J. C. MAXWELL, *A Treatise on Electricity and Magnetism*. Clarendon, 1918.

D. PELEG AND E. UPFAL, Constructing disjoint paths on expander graphs. *Combinatorica* **9**(3) (1989), 289–313.

S. M. ROSS, *Introduction to Probability Models*. Academic Press, fourth edition, 1989.

R. RUBINFELD, The cover time of a regular expander is $O(n \log n)$. *Inform. Process. Lett.* **35**(1) (1990), 49–51.

J. L. SYNGE, The fundamental theorem of electrical networks. *Quarterly of Applied Math.* **9** (1951), 113–127.

P. TETALI, Random walks and effective resistance of networks. *J. Theoret. Probability* **4**(1) (1991), 101–109.

W. THOMSON AND P. G. TAIT, *Treatise on Natural Philosophy*. Cambridge, 1879.

D. I. ZUCKERMAN, On the time to traverse all edges of a graph. *Inform. Process. Lett.* **38**(6) (1991), 335–337.

D. I. ZUCKERMAN, A technique for lower bounding the cover time. *SIAM J. Disc. Math.* **5**(1) (1992), 81–87.

Manuscript received 27 April, 1996

ASHOK K. CHANDRA
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120
ashok@almaden.ibm.com

WALTER L. RUZZO
Dept. of Computer Science and Engineering,
University of Washington,
Box 352350,
Seattle, WA 98195-2350
ruzzo@cs.washington.edu

PRABHAKAR RAGHAVAN
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120
pragh@almaden.ibm.com

PRASOON TIWARI
IBM T.J. Watson Research Center
Box 704
Yorktown Heights, NY 10598
ptiwari@watson.ibm.com