

New Bounds for Edge-Cover by Random Walk

Agelos Georgakopoulos* and Peter Winkler†

August 28, 2012

Abstract

We show that the expected time for a random walk on a (multi-)graph G to traverse all m edges of G , and return to its starting point, is at most $2m^2$; if each edge must be traversed in both directions, the bound is $3m^2$. Both bounds are tight and may be applied to graphs with arbitrary edge lengths, with implications for Brownian motion on a finite or infinite network of total edge-length m .

1 Introduction

Overview

The expected time for a random walk on a graph G to hit all the vertices of G has been extensively studied by probabilists, combinatorialists and computer scientists; applications include the construction of universal traversal sequences [2, 7], testing graph connectivity [2, 19], and protocol testing [21]. It has also been studied by physicists interested in the fractal structure of the uncovered set of a finite grid; see [11] for references and for an interesting relation between the cover time of a finite grid and Brownian motion on Riemannian manifolds.

Here we consider the time to cover all *edges* of G , and moreover we take as the fundamental parameter the number m of edges of G , rather than the number of vertices. Indeed, if the edges are taken to be of various lengths, then the number of vertices is no longer of interest and the total edge-length (also denoted by m) becomes the natural parameter by means of which to bound cover time.

Earlier Results

Let $G = \langle V, E \rangle$ be a connected, undirected graph (possibly with loops and multiple edges), where $|V| = n$ and $|E| = m$. Let $r \in V$ be a distinguished “root” vertex. A (simple) *random*

*Department of Mathematics, Technische Universität Graz, Austria; georgakopoulos@tugraz.at. The first author acknowledges the kind hospitality of the Dartmouth College and the second author.

†Department of Mathematics, Dartmouth, Hanover NH 03755-3551, USA; peter.winkler@dartmouth.edu. Research supported by NSF grant DMS-0901475.

walk on G beginning at r is a Markov chain on V , defined as follows: if the walk is currently at vertex x it chooses an edge incident to x uniformly at random, then follows that edge to determine its next state.

The expected value of the least time such that all vertices in V have been hit by the random walk will be denoted by $C_r^V(G)$ or simply C_r^V , and we call this number the (expected) *vertex cover time* of G from r . The vertex cover time $C^V(G)$ of G is defined to be the maximum of $C_r^V(G)$ over all $r \in V$.

Upper bounds on $C^V(G)$ have been widely sought by the theoretical computer science community, and many results obtained. The cover time of a simple graph (no loops or multiple edges) was shown [2] to be at most cubic in n ; more recently Feige [13] refined the bound to $\frac{4}{27}n^3$ plus lower-order terms, which is realized by the same “lollipop” graph that maximizes expected hitting time [5]—namely, a clique on $\frac{2}{3}n$ vertices attached to one end of a path on the remaining vertices.

In fact [2] shows that $C^V(G)$ is $O(mn)$, thus trivially $O(m^2)$, even if the walk is required to return to the root. The expected time to cover all vertices and return to the root is called the (vertex) *cover and return time*, here denoted by CR_r^V ; it is often easier to work with than C_r^V . Among simple graphs (no loops or multiple edges) with m edges, the path P_{m+1} on m edges—starting in the middle—provides the greatest vertex cover time, namely $m^2 + \lfloor m/2 \rfloor \lceil m/2 \rceil \sim \frac{5}{4}m^2$ [13]. The path started at one end had earlier been shown [6] to have the strictly greatest vertex cover and return time among all trees with m edges.

The expected time for a random walk beginning at r to cover all the *edges* of G will be denoted by $C_r^E(G)$; if the walk is required to cover all the *arcs*, that is, to traverse every edge of G in both directions, the expected time is denoted by $C_r^A(G)$. If the walk must return to r , we use CR_r^E and CR_r^A as above. Edge-cover is for some purposes more natural than vertex-cover; for example, it makes more sense on a graph with loops and multiple edges, where there is no bound on any kind of cover based only on the number n of vertices, but there are natural, tight bounds (as we shall see) on CR^E and CR^A in terms of the number of edges.

Bounds on expected edge-cover times were obtained by Bussian [8] and Zuckerman [28, 29]. Bussian noted that to cover the edges of a graph G , it suffices to hit each vertex somewhat more often than its degree. Recently Ding, Lee and Peres [10] proved a conjecture of Winkler and Zuckerman [27] saying roughly that the expected time to hit every vertex *often* is no more than a constant times the vertex-cover time, which for a regular graph is bounded by $6n^2$ [18]; it

follows that the edge- (or arc-) cover time for a regular graph is $O(mn)$, although the constant may be awful. For general graphs, Zuckerman [29] succeeded in showing $\text{CR}^A(G) \leq 22m^2$, using probabilistic methods.

Our Results

The main result of this paper is that if G is a multigraph with m edges then the expected time it takes for a simple random walk started at any vertex to traverse all the edges of G (and return to its starting point) is at most $2m^2$; in fact, we can even fix an orientation of each edge beforehand, and $2m^2$ steps will be enough on average to traverse each edge in its pre-specified direction. If instead we insist that each edge be traversed both ways, then we have to wait at most $3m^2$ steps on average.

As an intermediate step in our proof of the above bounds we obtain a refinement of a well-known formula of Chandra et. al. [9] concerning the commute time of random walk.

These results apply more generally if the edges of G are assigned positive real lengths and the transition probabilities of the random walk, as well as the time it takes to make a step, are appropriately adapted. In that case the above bounds apply when m is the sum of the lengths of the edges of G .

We call a graph with edge-lengths a *network*, and interpret it as an *electrical network* by equating edge-length with resistance along an edge. The usual connections between random walk and electrical networks (as in [12, 23] or see below) apply provided that in the definition of random walk, the probability that an edge adjacent to the current position is traversed is proportional to its conductance—that is, inversely proportional to its resistance (length). Such random walk has the same distribution as the sequence of vertices visited by standard Brownian motion.

To effect the correct scaling, and to make our results work for standard Brownian motion, we say that traversing an edge of length ℓ takes time ℓ^2 . In fact the expected time for a standard Brownian particle to proceed from one vertex to a neighboring vertex is *not* generally the square of the length of the edge; that the “ ℓ^2 ” model for discrete random walk permits extending our results to Brownian motion is a consequence of an averaging argument to be elucidated in Section 5.1.

For example, the expected time for a random walk to proceed from one end to the other of a path P_{m+1} of length m (that is, a path consisting of $m+1$ vertices and m unit-length edges)

is exactly m^2 ; this matches the time taken by standard Brownian motion to travel from 0 to m on the non-negative X -axis, and the time for our generalized random walk to travel from one end to the other of any linear network, regardless of the number and placement of vertices on it. Thus it works in all senses to say that the expected edge-cover time $C_a^E(P_{m+1})$ of a path of length m from an endpoint a is m^2 .

Our bounds have interesting implications for Brownian motion on infinite networks. Motivated by earlier work of the first author [15], Georgakopoulos and Kolesko [17] study Brownian motion on infinite networks in which the total edge length m is finite. Applied in this context, our results imply that Brownian motion will cover all edges in expected time at most $2m^2$, thus almost surely in finite time. See Section 5.2 for more details.

2 Definitions

2.1 Random Walk

We will denote by $N = \langle G, \ell \rangle$ a network with underlying multigraph G and edge-lengths $\ell : E(G) \rightarrow \mathbb{R}^+$, also interpreted as resistances. An *arc* of N is an oriented edge; in our context, a loop comprises two arcs. A *random walk* on N begins at some vertex r and when at vertex x , traverses the arc (x, y) to y with probability

$$\frac{1/\ell(x, y)}{\sum_{z \sim x} 1/\ell(x, z)}. \quad (1)$$

Our random walks take place in continuous time. The time it takes our random walker to perform a step, i.e. to move from a vertex x to one of its neighbors, depends on the lengths of the edges incident with x . There are various ways to define this dependence. In this paper we will consider two natural models for this dependence. Our main results will apply to both models with identical proofs, except that it takes a different argument to prove (2) below.

The Brownian model

Brownian motion on a line extends naturally to Brownian motion on a network (see, e.g., [3, 4, 14, 24]); when at a vertex a Brownian particle moves with equal likelihood onto any incident edge (with the understanding that an incident loop counts for this purpose as two incident edges). The edges incident to a vertex constitute a “Walsh spider” (see, e.g., [26, 4]) with equiprobable legs, and it is easily verified that in such a setting the probability of traversing

a particular incident edge (or oriented loop) first is proportional to the reciprocal of the length of that edge.

In our model, we care only about where the particle is after an edge-traversal, and how long it took to get there. The latter should thus be the time taken by a Brownian particle to traverse a given edge incident to its starting vertex, given that it traversed that edge first; this time is a random variable whose distribution and expectation (the latter to be computed later) depend not only on the length of the specified edge, but on the lengths of the other incident edges as well.

In the case of *simple random walk*, i.e. when all edges have length 1, the expected time for this random walk to take a step is 1.

The ℓ^2 model

In this model, the time it takes for the random walk to take a step is less random: traversing an edge of length ℓ always takes time ℓ^2 . Thus time is governed by (1) alone. The ℓ^2 model and the Brownian model differ only in timing; probabilities of walks are identical, as in both cases the next edge to be traversed is chosen according to the distribution in (1) for incident edges. If all edges of the network are of the same length, then in expectation, timing is identical as well. Readers interested only in simple random walk will lose nothing by assuming throughout that all edges are of length 1.

The expected time to cover all edges (respectively, arcs) of N by a random walk (in either model) will be denoted by $C_r^E(N)$ (resp., $C_r^A(N)$); to cover all edges or arcs and return, by $CR_r^E(N)$ or $CR_r^A(N)$. Maximizing over r gives $C^E(N)$, $C^A(N)$, $CR^E(N)$, and $CR^A(N)$.

The *effective resistance* \mathcal{R}^{xy} between vertices x and y is defined in electrical terms as the reciprocal of the amount of current that flows from x to y in N when a unit voltage difference is applied to them, assuming that each edge offers resistance equal to its length. See [12] or [20] for the basic definitions concerning electrical networks. Effective resistances sum in series: if all paths from x to y go through z , then $\mathcal{R}^{xy} = \mathcal{R}^{xz} + \mathcal{R}^{zy}$. The reciprocals of effective resistances, i.e., effective conductances, sum in parallel: if A and B are otherwise disjoint networks containing x and y , then in the union of the two networks,

$$\frac{1}{\mathcal{R}^{xy}} = \frac{1}{\mathcal{R}_A^{xy}} + \frac{1}{\mathcal{R}_B^{xy}}.$$

3 Commute Times

The *commute time* $T^{x \leftrightarrow y}$ from vertex x to y in N is the (random) time it takes for a random walk to travel from x to y and back to x . The expected commute time $\mathbb{E}T^{x \leftrightarrow y}$ between two vertices x, y of a network has an elegant expression proved in [9], and well known for the case of unit edge lengths:

$$\mathbb{E}T^{x \leftrightarrow y} = 2m\mathcal{R}^{xy}, \tag{2}$$

where $m := \sum_{e \in E(G)} \ell(e)$ is the total length of the network and \mathcal{R}^{xy} the effective resistance as defined above.

In fact a more general identity is proved in [9] (Theorem 2.2 there): suppose each traversal of an arc \vec{e} comes with a cost $f(\vec{e})$, where f is a possibly asymmetric cost function, and transition probabilities are still given by (1). Then the expected cost of an x - y commute is $F\mathcal{R}^{xy}$, where F is the sum of $f(\vec{e})/\ell(e)$ over all arcs \vec{e} . Substituting $\ell(e)^2$ for $f(\vec{e})$ gives the familiar $2m\mathcal{R}^{xy}$ for commute time in the ℓ^2 model, with m now understood as the total length of the network.

That the same formula applies to the Brownian model follows by approximating edge-lengths with rational numbers, then subdividing so that every edge has the same length; this has no effect on commute time in the Brownian model but brings traversal times in line with the ℓ^2 model without changing the formula.

In the rest of this section we refine (2) by differentiating between various commute tours according to their behavior with respect to subnetworks. We will need this refinement in order to prove our main results in the next section.

Suppose that N is the union of two subnetworks A, B such that $A \cap B = \{x, y\}$ (as in Fig. 1 below). For example, A could be an x - y edge and B could be the rest of N ; this is in fact the case that is needed later. We define the following events for random walks on N starting at x :

- (i) An A -commute from x to y is a closed walk starting at x and containing either an x - y subwalk via A or a y - x subwalk via A ;
- (ii) An \overrightarrow{A} -commute from x to y is a closed walk starting at x and containing an x - y subwalk via A ;
- (iii) An \overleftarrow{A} -commute from x to y is a closed walk starting at x and containing a y - x subwalk via A ;

Figure 1: Networks A and B meeting only at vertices x and y .

- (iv) An \overleftrightarrow{A} -commute from x to y is a closed walk starting at x and containing both an x - y subwalk via A and a y - x subwalk via A .

Define $T_A^{xy}, T_{A \rightarrow}^{xy}, T_{A \leftarrow}^{xy}, T_{A \leftrightarrow}^{xy}$ to be the time for random walk on N starting at x to perform an A -commute, a \overrightarrow{A} -commute, a \overleftarrow{A} -commute, and a \overleftrightarrow{A} -commute from x to y , respectively. Let $\mathbb{E}T_A^{xy}, \mathbb{E}T_{A \rightarrow}^{xy}, \mathbb{E}T_{A \leftarrow}^{xy}, \mathbb{E}T_{A \leftrightarrow}^{xy}$ denote the corresponding expected times.

Theorem 3.1. *For every network $N = (G, \ell)$ and any two subnetworks A, B having precisely two vertices x, y in common,*

$$(i) \quad \mathbb{E}T_A^{xy} = 2m\mathcal{R}_A \frac{\mathcal{R}_A + \mathcal{R}_B}{2\mathcal{R}_A + \mathcal{R}_B};$$

$$(ii) \quad \mathbb{E}T_{A \rightarrow}^{xy} = 2m\mathcal{R}_A \frac{2\mathcal{R}_A + \mathcal{R}_B}{2\mathcal{R}_A + \mathcal{R}_B} = 2m\mathcal{R}_A = \mathbb{E}T_{A \leftarrow}^{xy};$$

$$(iii) \quad \mathbb{E}T_{A \leftrightarrow}^{xy} = 2m\mathcal{R}_A \frac{3\mathcal{R}_A + \mathcal{R}_B}{2\mathcal{R}_A + \mathcal{R}_B},$$

where $m := \sum_{e \in E(G)} \ell(e)$ is the total edge-length, and \mathcal{R}_X is the effective resistance between x and y in the network X .

Remark: Notice that the three expressions differ only in the coefficient of \mathcal{R}_A in the numerator.

Proof. In all cases, let the particle start at x and walk on N ‘forever’, and consider the times when the particle completes a commute from x to y . We are interested in the expectation of

the time T when the first A -commute, \overrightarrow{A} -commute, or \overleftarrow{A} -commute is completed according to which of the cases (i) - (iii) we are considering.

Note that any A -commute, \overrightarrow{A} -commute, or \overleftarrow{A} -commute from x to y can be decomposed into a sequence of disjoint commutes. Thus we can define a random variable Y to be the number of x - y commutes until the first A -commute, \overrightarrow{A} -commute, or \overleftarrow{A} -commute from x to y is completed, according to which of the cases (i) - (iii) we are considering. Let $X_i, i = 1, 2, \dots$ be the duration of the i th x - y commute. Then our stopping time T satisfies

$$T = \sum_{1 \leq i \leq Y} X_i.$$

Note that the expectation of each X_i is $\mathbb{E}T^{x \leftrightarrow y} = 2m\mathcal{R}^{xy}$ by (2). By Wald's identity [25] the expectation of T equals the expectation of Y times the expectation of X_i , and so we have

$$\mathbb{E}T = 2m\mathcal{R}^{xy}\mathbb{E}Y. \quad (3)$$

Thus it only remains to determine $\mathbb{E}Y$ in each of the cases (i) - (iii).

In order to compute $\mathbb{E}Y$ it is useful to first calculate the probability p_A that the first visit to y of a random walk starting at x will be via A . For this it is convenient to use the electrical network technique. Disconnect A and B at their common vertex y to obtain a network consisting of A, B connected 'in series' at x . Denote by y_A the vertex of A corresponding to y , and by $\mathcal{C}_A = 1/\mathcal{R}_A$ the effective conductance between x and y in A ; similarly for B . From [12] we have the probability that a random walk started at x hits y_A before y_B is $\mathcal{C}_A/\mathcal{C}_B$, thus $p_A = \frac{\mathcal{C}_A}{\mathcal{C}_A + \mathcal{C}_B} = \frac{\mathcal{R}^{xy}}{\mathcal{R}_A}$.

By the same argument, a random walk from y to x will go via A with the same probability p_A . We can now determine $\mathbb{E}Y$ in each of the three cases. For case (ii), note that a commute from x to y is a \overrightarrow{A} -commute if the forward trip is via A , which occurs with probability p_A . Thus we have $\mathbb{E}Y =: \mathbb{E}Y_{ii} = 1/p_A = \frac{\mathcal{R}_A}{\mathcal{R}^{xy}}$, since the expected number of Bernoulli trials until the first success is the reciprocal of the success probability of one trial. Plugging this into (3) yields $\mathbb{E}T = \mathbb{E}T_{A \rightarrow}^{xy} = 2m\mathcal{R}_A$ as claimed. By similar arguments this expression also equals $\mathbb{E}T_{A \leftarrow}^{xy}$.

For case (i), note that a commute from x to y fails to be an A -commute if and only if both trips fail to be via A . Since the two trips are independent, the probability that a random x - y commute is an A -commute is $1 - (1 - p_A)^2 = p_A(2 - p_A)$. By a similar argument as above we obtain $\mathbb{E}Y =: \mathbb{E}Y_i = \frac{1}{p_A(2 - p_A)}$ and so $\mathbb{E}T =: \mathbb{E}T_A^{xy} = 2m\mathcal{R}^{xy} \frac{\mathcal{R}_A}{\mathcal{R}^{xy}} \frac{1}{2 - p_A} = 2m\mathcal{R}_A \frac{1}{2 - p_A} = 2m\mathcal{R}_A \frac{1}{2 - \mathcal{R}_B/(\mathcal{R}_A + \mathcal{R}_B)} = 2m\mathcal{R}_A \frac{\mathcal{R}_A + \mathcal{R}_B}{2\mathcal{R}_A + \mathcal{R}_B}$.

Finally, to determine $\mathbb{E}Y$ in case (iii) we argue as follows. To begin with, we have to make an expected Y_i tries until we go via A in at least one of the trips of some x - y commute, i.e. until we achieve our first A -commute. Then, unless our first A -commute was an \overleftarrow{A} -commute, we will have to make another Y_{ii} tries in expectation to go via A in the other direction. By elementary calculations, an A -commute fails to be an \overleftarrow{A} -commute with probability $q = \frac{2(1-p_A)}{(2-p_A)} = 2\frac{R_A}{2R_A+R_B}$. Summing up, we have $\mathbb{E}Y =: \mathbb{E}Y_{ii} = \mathbb{E}Y_i + q\mathbb{E}Y_{ii}$. Using our earlier calculations and (3) this yields $\mathbb{E}T =: \mathbb{E}T_{A\leftrightarrow}^{xy} = 2mR_A \left(\frac{R_A+R_B}{2R_A+R_B} + 2\frac{R_A}{2R_A+R_B} \right) = 2mR_A \frac{3R_A+R_B}{2R_A+R_B}$. \square

4 Main Results

For a random walk on a network N starting at a vertex x , define the random variable CR_x^E , called the *edge cover-and-return time*, to be the first time when each edge of N has been traversed and the particle is back at x . Similarly, if N is a digraph, we define CR_x^A , called the *arc cover-and-return time*, to be the first time when each arc of N has been traversed and the particle is back at x . Here, the directions of the edges do not affect the behavior of the random walk; the particle is allowed to traverse arcs backwards and its transition probabilities are always given by (1).

If N is undirected we interpret it as a digraph by replacing each edge (loops included) by an arc in each direction; thus, in that case, CR_x^A is the first time when each edge of N has been traversed in both directions and the particle is back at x .

Theorem 4.1. *Let $N = (G, \ell)$ be an undirected network and let $m := \sum_{e \in E(G)} \ell(e)$. Then*

$$(i) \quad \mathbb{E}\text{CR}^E(N) \leq 2m^2 \text{ and}$$

$$(ii) \quad \mathbb{E}\text{CR}^A(N) \leq 3m^2.$$

Moreover, if \vec{N} is the result of orienting the edges of N in any way, then

$$(iii) \quad \mathbb{E}\text{CR}^A(\vec{N}) \leq 2m^2.$$

Proof. Our proof follows the lines of the ‘spanning tree argument’ used in [1] for the vertex cover time, the main difference being that we apply Lemma 3.1 instead of (2).

Let σ be a closed walk in G starting at an arbitrary fixed root r and traversing each edge of G precisely once in each direction. Such a walk always exists: if G is a tree then a depth-first

search will do, and if not then one can consider a spanning tree T of G , construct such a walk on T , and then extend it to capture the chords.

We are going to use σ in order to define epochs for the random walk in N starting at x , and then apply Theorem 3.1 to bound the time spent between pairs of such epochs. These epochs will differ depending on which of the two cases we are considering.

Cases (i) and (iii): We prove the stronger result that $\mathbb{E}\text{CR}_r^A(\vec{N}) \leq 2m^2$. let $e_i, i = 1, 2, \dots, 2|E(G)|$, be the i th edge traversed by σ , with endvertices x_i, y_i appearing in σ in that order. For $i = 1, 2, \dots, 2|E(G)|$ define the i th epoch τ_i as follows. If e_i is directed from x_i to y_i , then τ_i is the first time after τ_{i-1} when σ is at y_i and has gone there from x_i using the edge e_i in that step, where we set $\tau_0 = 0$. If e_i is directed the other way, then we just let τ_i be the first time after τ_{i-1} when σ is at y_i .

Note that at time $\tau_{2|E(G)|}$ our random walk is back to r and has performed an arc cover-and-return tour. Thus $\mathbb{E}\text{CR}_r^E(\vec{N})$ is at most the expectation of $\tau_{2|E(G)|}$. Now the latter can be bounded using Theorem 3.1 (ii) as follows. For every edge $e = xy$ of G , there are precisely two time intervals bounded by the above epochs corresponding to e : If j, k are the two indices for which $e_j = e_k = e$, then these are the time intervals $I_e^1 := [\tau_{j-1}, \tau_j]$ and $I_e^2 := [\tau_{k-1}, \tau_k]$. Note that, by the definition of τ_i , the motion of the random walker in the union of these two intervals is in fact an \vec{e} -commute or an \overleftarrow{e} -commute (as defined in Section 3) from x to y , according to whether e is directed from x to y or the other way round. Thus, applying Theorem 3.1 (ii) with $A = \{x, y; e\}$ and $B = G - \{e\}$ yields that the expected value \mathcal{E}_e of $|I_e^1| + |I_e^2|$ is $2m\ell(e)$. Since the time interval $[0, \tau_{2|E(G)|}]$ is the union of all such pairs of intervals, one pair for each $e \in E(G)$, we have

$$\mathbb{E}\tau_{2|E(G)|} = \sum_{e \in E(G)} \mathcal{E}_e = \sum_{e \in E(G)} 2m\ell(e) = 2m^2.$$

This yields $\mathbb{E}\text{CR}_r^A(\vec{N}) \leq 2m^2$, thus also $\mathbb{E}\text{CR}_r^E(N) \leq 2m^2$, as claimed.

Case (ii): To bound $\mathbb{E}\text{CR}_r^A(N)$, we follow the same arguments as before, except that we define the epochs slightly differently: the edges of G are not directed now, and we always let τ_i be the first time after τ_{i-1} when our random walk is at y_i and has gone there from x_i using the edge e_i in that step. We define the time intervals I_e^1 and I_e^2 as above, but this time we note that the motion of the random walker in the union of these two intervals is an \overleftarrow{e} -commute. Applying Theorem 3.1 (iii) with $A = \{x, y; e\}$ and $B = G - \{e\}$, we obtain $\mathcal{E}_e \leq 2m\ell(e) \frac{3\ell(e) + R_B}{2\ell(e) + R_B}$. This expression attains its maximum value when $R_B = 0$, and so we

obtain $\mathcal{E}_e \leq 2m\ell(e)\frac{3\ell(e)}{2\ell(e)} = 3m\ell(e)$. Adding up all $\mathcal{E}_e, e \in E(G)$ as above we conclude that $\mathbb{E}\text{CR}_r^A(N) \leq 3m^2$, as claimed. \square

The bounds of Theorem 4.1 are tight in the following situations. For *Case (i)* and thus also *(iii)*, we have already noted that a path of length m takes time $2m^2$ to cover all edges and return. For *Case (ii)*, a network consisting of a single vertex and a loop (of any length m) takes one step to cover that loop in one direction, then on average two more to catch the other direction, so for this network $\mathbb{E}\text{CR}^A = 3m^2$.

5 Application to Brownian Motion and Infinite Networks

5.1 Finite networks

We begin this section by showing directly that (expected) edge-cover-and-return time—and, indeed, any kind of return time—is the same in the Brownian model of random walk as it is in the ℓ^2 model as defined in Section 2. Fix a network N and suppose that vertex x of N has incident edges e_1, \dots, e_k of lengths ℓ_1, \dots, ℓ_k respectively. (As usual loops must be represented twice in this list, once in each direction.) The mean time taken by a walk in the ℓ^2 model to traverse one of these edges, starting from x , is easily calculated using (1):

$$\sum_{i=1}^k \frac{1}{\ell_i} \frac{1}{C_x} \ell_i^2 = \frac{1}{C_x} \sum_{i=1}^k \ell_i$$

where $C_x := \sum_{j=1}^k (1/\ell_j)$. The same is true in the Brownian model, because we may identify the endpoints of the edges, calling the unified vertex y , and compute the commute time between x and y using the formula from Section 3, which holds in both models. By symmetry, the desired quantity is half the commute time.

Since the expected time taken by a cover-and-return tour in either model is the sum over the vertices of N of the expected time spent exiting those vertices, this quantity is the same for both models. Note, however, that this does not mean that the expected time taken in the Brownian model for a *particular* cover-and-return tour is the sum of the squares of the lengths of its edges, and indeed that is not generally the case.

To compute the latter we need to know what the expected edge-traversal times are in the Brownian model. This is not needed to apply our bounds, but the computation is easy and as far as we know, has not appeared elsewhere. We make use of a couple of simple (and known, but proved here as well) facts about Brownian motion.

Lemma 5.1. *Consider standard Brownian motion on the positive real half axis, and let t_b be the time of the first visit to $\ell \in \mathbb{R}^+$, and t_a the time of the last visit to 0 before t_b . Then the expected value T_ℓ of $t_b - t_a$ is $\ell^2/3$.*

Proof. From scaling properties of Brownian motion (see, e.g., [22]) we know that T_ℓ must be a multiple of ℓ^2 , say $\alpha\ell^2$. Then $T_{\ell/2}$ is $\alpha\ell^2/4$. Consider the first time t_c that the particle reaches the point ℓ^2 . From there, it takes expected time $\ell^2/4$ to reach either 0 or ℓ for the first time again, reaching each of them first with equal probability $1/2$. Now conditioning on the event that ℓ was reached before 0 after t_c , we expect $t_b - t_a$ to be $T_{\ell/2} + \ell^2/4$, while if 0 was reached first then the experiment is effectively restarted, and we expect $t_b - t_a$ to be its overall expectation T_ℓ . Combining we get $T_\ell = 1/2(T_{\ell/2} + \ell^2/4) + 1/2T_\ell$. Plugging in the above formulas for T_ℓ and $T_{\ell/2}$ we can now solve for α , and we get $\alpha = 1/3$. \square

Lemma 5.2. *Suppose a Brownian particle begins at vertex x in a network N and proceeds until it traverses one of the incident edges e_1, \dots, e_k , and let \mathcal{T} be the (random) time spent before the particle departs x for the last time. Then \mathcal{T} is independent of the index of the edge traversed.*

Proof. Since the past is irrelevant to the particle, there is a fixed distribution σ on $\{1, 2, \dots, k\}$ for which edge is traversed after the particle departs x for the last time (namely, the distribution whose probabilities are proportional to the reciprocals of the edge-lengths). Thus, the index of the traversed edge is independent of \mathcal{T} , and necessarily, vice-versa. \square

We are now ready to derive our formula. Lemma 5.2 implies that in particular $\mathbb{E}\mathcal{T}$ is independent of which edge is traversed. Combining with Lemma 5.1, the expected time taken by the particle to traverse an edge from x , given that it traversed edge e_i first, is $\mathbb{E}\mathcal{T} + \ell_i^2/3$. It follows that the expected time to traverse *some* edge from x is

$$\sum_{i=1}^k (\mathbb{E}\mathcal{T} + \ell_i^2/3) \frac{1}{\ell_i C_x}$$

which we know must be equal to $\frac{1}{C_x} \sum_{i=1}^k \ell_i$. Solving gives $\mathbb{E}\mathcal{T} = \frac{2}{3} \frac{1}{C_x} \sum_{i=1}^k \ell_i$, and we have proved:

Theorem 5.3. *Suppose a standard Brownian particle on a network N begins at vertex x with incident edges e_1, \dots, e_k of lengths ℓ_1, \dots, ℓ_k respectively. Then the expected time taken by the particle to traverse edges e_i , given that it traversed e_i first, is $\frac{1}{3}\ell_i^2 + \frac{2}{3} \sum_{j=1}^k \ell_j$.*

A simple example of an edge-cover-and return tour with different expected times for the two models takes place on the network N consisting of two vertices x and y , connected by an edge e of length 1 and an edge f of length 2. Then the cover tour from x consisting of e then f has expected time $5/3 + 8/3 = 13/3$ in the Brownian model, but constant time $1^2 + 2^2 = 9$ in the ℓ^2 model. The expected time for an edge-cover-and-return tour on N is 7.7 in either model, although when conditioned on (say) having started with edge e , the results are quite different.

5.2 Infinite networks of finite total length

We conclude with some implications for Brownian motion on infinite networks.

As an illustrative example, consider the infinite binary tree T , with edge-lengths 4^{-k} at level k (the edges incident to the root counting as level 1). To this network of total length 1 it is natural to append a *boundary* ∂T : considered as a metric space, T has a metric completion $|T|$ and ∂T is the set of completion points. An equivalent way to define ∂T is as the set of infinite paths starting at the root of T , which admits a natural bijection to the set of infinite binary sequences. Note that ∂T is homeomorphic to the Cantor set.

It is possible to prove that starting at the root of T , Brownian motion or random walk in the ℓ^2 model will almost surely reach ‘infinity’, i.e. ∂T , after finite time. Georgakopoulos and Kolesko [17] show that it is possible to let the particle continue its random motion afterwards: they construct a random process on $|T|$ whose sample paths are continuous with respect to the topology of $|T|$ and behave like standard Brownian motion in the neighborhood of each vertex of T . This construction was motivated by the results of [15] and an attempt to extend the theory of [12], relating electrical networks and random processes, to the infinite case.

Applied in this context, our results have a somewhat surprising implication: Brownian motion on $|T|$ will cover all edges of T —and all of the continuum-many boundary points ∂T —in expected time at most 2, thus almost surely in finite time. (We proved our results here for finite networks only, and so they cannot be directly applied to infinite ones. However, the Brownian motion of [17] is constructed as a limit of the Brownian motions on an increasing sequence of finite subgraphs of T , and as our results apply to each member of this sequence, they can be extended to the limit; see [17] for details.)

In fact, the Brownian motion of [17] is defined not only for trees, but also for arbitrary networks of finite total length. It was shown in [16] that every such network admits a boundary

as above. Moreover, this boundary, with its corresponding topology, is well known and has had many applications. To these results we have here added a (tight) bound on expected cover time, perhaps adding further impetus to the study of networks of finite total length.

6 Acknowledgments

The portions of this research concerning Brownian motion benefited greatly from communications with Steve Evans (U.C. Berkeley) and with Yuval Peres and David B. Wilson (Microsoft Research).

References

- [1] D. Aldous and J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, book in preparation, <http://www.stat.berkeley.edu/~aldous/RGW/book.html> (as of 2011).
- [2] R. Aleliunas, R. Karp, R. Lipton, L. Lovász and C. Rackoff, Random walks, universal traversal sequences, and the complexity of maze traversal, *Proc. 20th IEEE Symp. Found. Computer Science* (1979), 218–233.
- [3] J. R. Baxter and R. V. Chacon, The equivalence of diffusions on networks to Brownian motion, *Contemp. Math.* **26** (1984), 33–47.
- [4] M. T. Barlow, J. W. Pitman and M. Yor, On Walsh’s Brownian motions, *Séminaire de Probabilités (Strasbourg)* **23** (1989), 275–293.
- [5] G. R. Brightwell and P. Winkler, Maximum hitting time for random walks on graphs, *J. Random Structures and Algorithms* **1** #3 (1990), 263–276.
- [6] G. R. Brightwell and P. Winkler, Extremal cover time for random walks on trees, *J. Graph Theory* **14** #5 (1990), 547–554.
- [7] A. Broder, Universal sequences and graph cover times; a short survey, *Sequences* (Naples/Positano, 1988) 109–122, Springer, New York, 1990.
- [8] E. R. Bussian, *Bounding the Edge Cover Time of Random Walks on Graphs*, PhD Thesis, Dept. of Mathematics, Georgia Institute of Technology (1996), <http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA311706&Location=U2&doc=GetTRDoc.pdf>.

- [9] 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, *Proc. 21st ACM Symp. Theory of Computing* (1989), 574–586.
- [10] G. Ding, J. Lee and Y. Peres, Cover times, blanket times, and majorizing measures, <http://arxiv.org/abs/1004.4371>.
- [11] A. Dembo, Y. Peres, J. Rosen and O. Zeitouni, Cover times for Brownian motion and random walks in two dimensions, *Acta. Math.* **186** #2 (2001), 239–270.
- [12] P. Doyle and J. L. Snell, *Random Walks and Electrical Networks*, Mathematical Association of America, 1984.
- [13] U. Feige, A tight upper bound on the cover time for random walks on graphs, *Random Structures & Algorithms* **6** #1 (1995).
- [14] H. F. Frank and S. Durham, Random motion on binary trees, *J. Appl. Prob.* **21** (1984), 58–69.
- [15] A. Georgakopoulos, Uniqueness of electrical currents in a network of finite total resistance, *J. London Math. Soc.* **82** #1 (2010), 256–272.
- [16] A. Georgakopoulos, Graph topologies induced by edge lengths. In Infinite Graphs: Introductions, Connections, Surveys. Special issue of *Discrete Math* **311** (2011), 1523–1542.
- [17] A. Georgakopoulos and K. Kolesko, In preparation.
- [18] J. D. Kahn, N. Linial, N. Nisan and M. E. Saks, On the cover time of random walks on graphs, *J. Theoretical Probability* **2** (1989), 121–128.
- [19] A. Karlin and P. Raghavan, Random walks and undirected graph connectivity; a survey, *Discrete Probability and Algorithms* (Minneapolis 1993), 95–101, IMA Vol. Math. Appl. **72**, Springer, New York 1995.
- [20] R. Lyons with Y. Peres, *Probability on Trees and Networks*, Cambridge University Press, to appear; 2011 draft at <http://mypage.iu.edu/~rdlyons/prbtree/book.pdf>.
- [21] M. Mihail and C. Papadimitriou, On the random walk method for protocol testing, *Computer Aided Verification* (Stanford, CA), 132–141, Lecture Notes on Comp. Sci. **818**, Springer, Berlin 1994.

- [22] P. Mörters and Y. Peres, *Brownian Motion*, Cambridge University Press, Cambridge 2010.
- [23] P. Tetali, Random walks and the effective resistance of networks, *J. Theoretical Prob.* **4** (1991) 101–109.
- [24] N. Th. Varopoulos, Long range estimations for Markov chains, *Bull. Sci. Math.* **109** (1985) 225–252.
- [25] A. Wald, On cumulative sums of random variables, *Ann. Math. Stat.* **15** #3 (Sept. 1944), 283–296, doi:10.1214/aoms/1177731235. JSTOR 2236250. MR10927. Zbl 0063.08122.
- [26] J. Walsh, A diffusion with discontinuous local time, in *Temps Locaux, Astérisque* **52-53** (1978), 37–45.
- [27] P. Winkler and D. Zuckerman, Multiple cover time, *Random Structures & Algorithms* **9** (1996), 403–411.
- [28] D. Zuckerman, Covering times of random walks on bounded-degree trees and other graphs, *J. Theoretical Prob.* **2** #1 (1989) 147–157.
- [29] D. Zuckerman, On the time to traverse all edges of a graph, *Information Processing Letters* **38** (1991) 335–337.