The Cover Time of Cartesian Product Graphs

Mohammed Amin Abdullah Colin Cooper Tomasz Radzik

Abstract

Let $F = G\square H$ be the Cartesian product of graphs G, H. We relate the cover time $\mathbf{COV}[F]$ of F to the cover times of its factors. When one of the factors is in some sense larger than the other, its cover time dominates, and can become of the same order as the cover time of the product as a whole. Our main theorem effectively gives conditions for when this holds. The probabilistic technique which we introduce, based on the blanket time, is more general and may be of independent interest, as might some of our lemmas.

1 Introduction

For a connected graph Let G, denote by $V(G)$ and $E(G)$ the vertex and edge set respectively. The *vertex cover time* $\mathbf{COV}[G]$ of G is defined as the expected time it takes a random walk to visit all vertices of the graph, maximised over all possible starting vertices. This quantity is a fundamental area in the study of random walks has been extensively studied giving rise to a large body of theory and application. Let $n = |V(G)|$ and $m = |E(G)|$. It is a classic result of Aleliunas, Karp, Lipton, Lovász and Rackoff [1] that $\mathbf{COV}[G] \leq 2m(n-1)$. It was shown by Feige [6], [7], that for any connected graph G, the cover time satisfies $(1-o(1))n\log n \leq \mathbf{COV}[G] \leq (1+o(1))\frac{4}{27}n^3$. Between these two extremal examples, the cover time, both exact and asymptotic, has been determined for a number of different classes of graphs.

In this paper we study the cover time of a graph F that is the Cartesian product of graphs G and H. This encompasses many classes of graphs, some of which, such as grids, play a prominent role in applications such as networking.

2 Cartesian Product of Graphs: Definition, Properties, Examples

Graphs G and H are assumed to be finite, undirected, simple and connected.

2.1 Definition

Definition 1. *Let $G = (V_G, E_G)$ and $H = (V_H, E_H)$ be simple, connected, undirected graphs. The* Cartesian product*, $G\square H$ of G and H is the graph $F = (V_F, E_F)$ such that*

(i) $V_F = V_G \times V_H$

(ii) $((a,x),(b,y)) \in E_F$ *if and only if either*

1. $(a,b) \in E_G$ *and* $x = y$, *or*
2. $a = b$ *and* $(x,y) \in E_H$

We call G and H the *factors* of F, and we say that G and H are *multiplied* together.

We can think of $F = G \Box H$ in terms of the following construction: We make a copy of one of the graphs, say G, once for each vertex of the other, H. Denote the copy of G corresponding to vertex $x \in V_H$ by G_x. Let a_x denote a vertex in G_x corresponding to $a \in V_G$. If there is an edge $(x,y) \in E_H$, then add an edge (a_x, a_y) to the construction.

Notation For a graph $\Gamma = (V_\Gamma, E_\Gamma)$, let $n_\Gamma = |V_\Gamma|$ and $m_\Gamma = |E_\Gamma|$. In addition we will use use the notation N and M to stand for n_F and m_F respectively.

2.2 Properties

Commutativity and Associativity of the Cartesian Product Operation For a pair of graphs G and H, $G \Box H$ is isomorphic to $H \Box G$; that is, if vertex labels are ignored, the graphs are identical. Note, however, that by (i) of Definition 1, the two different orders on the product operation do produce different labellings.

The cartesian product is associative, and for a natural number d, we denote by G^d the d'th Cartesian power, that is, $G^d = G$ when $d = 1$ and $G^d = G^{d-1} \Box G$ when $d > 1$.

Vertices and Edges of the Product Graph The number of vertices and edges of a Cartesian product is related to the vertices and edges of its factors as follows:

(i) $N = n_G n_H$.

(ii) $M = n_G m_H + n_H m_G$.

2.3 Examples

We give examples of Cartesian product of graphs, some of which are important to the proofs of this paper. First, we remind the reader of some specific classes of graphs, and define new ones: Let P_n denote the *n-path*, the path graph of n vertices. Let $\mathbb{Z}_n$ represent the *n-cycle*, the cycle graph with n vertices. The Cartesian product of a pair of paths, $P_p \Box P_q$ is a $p \times q$ *rectangular grid*, and when $p = q = n$, is a $n \times n$ *grid*, or *lattice*. The product of a pair of cycles $\mathbb{Z}_p \Box \mathbb{Z}_q$ is a *toroid*, and when $p = q$, is a *torus*. Both grids and toroids can be generalised to higher powers in the obvious way to give $d-$dimensional grids and toroids respectively, where d is the number of paths or cycles multiplied together, respectively.

To give another - somewhat more arbitrary - example, a pictorial representation of the product of a triangle graph with a tree is given in Figure 1.

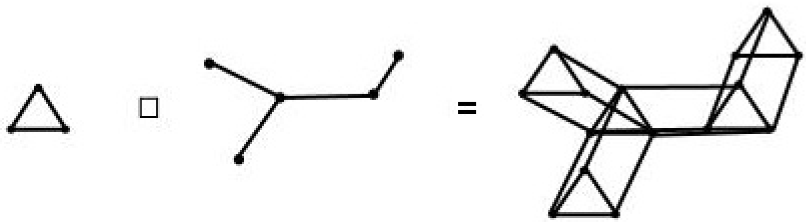

Figure 1: Cartesian product of a triangle with a tree.

3 Blanket Time

We introduce here a notion that is related to the cover time, and is an important part of the main theorem and the proof technique we use.

Definition 2 ([12]). *For a random walk $\mathcal{W}_u$ on a graph $G = (V, E)$ starting at some vertex $u \in V$, and $\delta \in [0, 1)$, define the the random variable*

$$B_{\delta,u}[G] = \min\{t : \forall v \in V, N_v(t) > \delta\pi_v t\}, \tag{1}$$

where $N_v(t)$ is the number of times $\mathcal{W}_u$ has visited v by time t and π_v is the stationary probability of vertex v. The blanket time *is*

$$\mathbf{B}_\delta[G] = \max_{u \in V} \mathbf{E}[B_{\delta,u}[G]].$$

The following was recently proved in [4]:

Theorem 1 ([4]). *For any graph G, and any $\delta \in (0, 1)$, we have*

$$\mathbf{B}_\delta[G] \le \kappa(\delta)\mathbf{COV}[G]$$

where the constant $\kappa(\delta)$ depends only on δ.

We define the following:

Definition 3 (**Blanket-Cover Time**). *For a random walk $\mathcal{W}_u$ on a graph $G = (V, E)$ starting at some vertex $u \in V$, define the the random variable*

$$\beta_u[G] = \min\{t : \forall v \in V\ N_v(t) \ge \pi_v \mathbf{COV}[G]\},$$

where $N_v(t)$ is the number of times $\mathcal{W}_u$ has visited v by time t and π_v is the stationary probability of vertex v. The blanket-cover time *is the quantity*

$$\mathbf{BCOV}[G] = \max_{u \in V} \mathbf{E}[\beta_u[G]].$$

Thus the blanket-cover time of a graph is the expected first time that each vertex v is visited at least $\pi_v\mathbf{COV}[G]$ times - which we shall refer to as the *blanket-cover criterion.*

In the paper that introduced the blanket time, [12], the following equivalence was asserted, which we conjecture to be true.

Conjecture 1. $\mathbf{BCOV}[G] = O(\mathbf{COV}[G])$.

In the same paper, this equivalence was proved for paths and cycles. However, we have not found a proof for the more general case. It can be shown without much difficulty that $\mathbf{BCOV}[G] = O((\mathbf{COV}[G])^2)$. Using the following lemma, we can improve upon this.

Lemma 2 ([9])**.** *Let i and j be two vertices and $k \geq 1$. Let W_k be the number of times j had been visited when i was visited the k-th time. Then for every $\varepsilon > 0$,*

$$\mathbf{Pr}\left(W_k < (1-\varepsilon)\frac{\pi_j}{\pi_i}k\right) \leq \exp\left(\frac{-\varepsilon^2 k}{4\pi_i\mathbf{COM}[i,j]}\right).$$

First we state a useful technique due to Matthews, [11], for bounding the cover time in terms of hitting times between vertices. We state a version that is slightly more general:

Theorem 3 (**Matthews' bound, [11]**)**.** *For a graph $G = (V,E)$, suppose $V' \subseteq V$, then*

$$\mathbf{H}^*_G[V'] = \max\{\mathbf{H}[u,v] : u,v \in V'\},$$

where $\mathbf{H}[u,v]$ is the hitting time from u to v in G. For a random walk on G starting at some vertex $v \in V'$, denote by $\mathbf{COV}_v[V']$ the expected time to visit all the vertices of V'. Then

$$\mathbf{COV}_v[V'] \leq \mathbf{H}^*_G[V']h(|V'|), \tag{2}$$

The notation $\mathbf{COV}[V']$ shall mean $\max_{v\in V'}\mathbf{COV}_v[V']$.

We use it thus:

Lemma 4.

$$\mathbf{BCOV}[G] = O\left((\log n)\mathbf{COM}^*[G]\right)$$

where $\mathbf{COM}^*[G] = \max_{u,v\in V(G)}\mathbf{COM}[u,v]$.

Proof. At time t some vertex i must have been visited at least $\pi_i t$ times, otherwise we would get $t = \sum_{v\in V} N_v(t) < \sum_{v\in V}\pi_v t = t$, where $N_v(t)$ is the number of times v has been visited by time t.

We let the walk run for $\tau = A(\log n)\mathbf{COM}^*[G]$ steps where A is a large constant. Some vertex i

will have been visited at least $\pi_i \tau$ times. Now we use Lemma 2 with $k = \pi_i \tau$. Then for any j,

$$\begin{aligned}
\mathbf{Pr}\left(W_k < (1-\varepsilon)\frac{\pi_j}{\pi_i}k\right) &\le \exp\left(\frac{-\varepsilon^2 k}{4\pi_i \mathbf{COM}[i,j]}\right) \\
&\le \exp\left(\frac{-\varepsilon^2 A \log n}{4}\right) \\
&\le 1/n^c
\end{aligned}$$

for some constant $c > 1$. Hence with probability at most $1/n^{c-1}$ the walk has failed to visit each vertex j at least $\pi_j \mathbf{COV}[G]$ times (by Matthews' technique, Theorem 3). We repeat the process until success. The expected number of attempts is $1 + O(n^{1-c})$. □

4 Results and Related work

Notation For a graph Γ, denote by: δ_Γ the minimum degree; θ_Γ the average degree; Δ_Γ the maximum degree; D_Γ the diameter.

The main theorem of this paper is the following. The main part of the work is the derivation of (4); the inequality (3) is relatively straightforward to derive and is included for completeness.

Theorem 5. *Let $F = (V_F, E_F) = G \Box H$ where $G = (V_G, E_G)$ and $H = (V_H, E_H)$ are simple, connected, unweighted, undirected graphs. We have*

$$\mathbf{COV}[F] \ge \max\left\{\left(1 + \frac{\delta_G}{\Delta_H}\right)\mathbf{COV}[H], \left(1 + \frac{\delta_H}{\Delta_G}\right)\mathbf{COV}[G]\right\}. \tag{3}$$

Suppose further that $n_H \ge D_G + 1$, then

$$\mathbf{COV}[F] \le K\left(\left(1 + \frac{\Delta_G}{\delta_H}\right)\mathbf{BCOV}[H] + \frac{M m_G m_H n_H \ell^2}{\mathbf{COV}[H] D_G}\right) \tag{4}$$

where $M = |E_F| = n_G m_H + n_H m_G$, $\ell = \log(D_G + 1)\log(n_G D_G)$ and K is some universal constant.

Note, by the commutativity of the Cartesian product, G and H in the may be swapped in (4), subject to the condition $n_G \ge D_H + 1$.

Theorem 5 extends much work done on the particular case of the two-dimensional toroid $\mathbb{Z}_n^2 = \mathbb{Z}_n \Box \mathbb{Z}_n$, culminating in a result of [3], which gives a tight asymptotic result for the cover time of $\mathbb{Z}_n^2$: $\mathbf{COV}[\mathbb{Z}_n^2] \sim \frac{1}{\pi} N (\log N)^2$, where $N = n^2$ is the number of vertices in the product. For $d \ge 3$, and with $N = n^d$, it is well established (see e.g., [10]) that $\mathbf{COV}[\mathbb{Z}_n^d] = \Theta(N \log N)$.

Theorem 5 also extends work done in [8] on powers G^d of general graphs G. Letting $N = n^d$ again, where n is the number of vertices in G, [8] shows $\mathbf{COV}[G^2] = O(\theta_G N \log^2 N)$ and for $d \ge 3$, $\mathbf{COV}[G^d] = O(\theta_G N \log N)$. The quantity $\theta_G = 2|E|/n$ is the average degree of G. However [8] does not address products of graphs that are different, nor does it seem that the proof techniques used could be directly extended to deal with such cases. Our proof techniques are different, but both this work and [8] make use of electrical network theory and analysis of subgraphs of the product that are isomorphic to the square grid $P_k \Box P_k$.

To prove the Theorem 5, we present a framework to bound the cover time of a random walk on a graph which works by dividing the graph up into (possibly overlapping) regions, analysing the behaviour of the walk when locally observed on those regions, and then composing the analysis of all the regions over the whole graph. Thus the analysis of the whole graph is reduced to the analysis of outcomes on local regions and subsequent compositions of those outcomes. This framework can be applied more generally than Cartesian products.

Some of the lemmas we use may be of independent interest. In particular, Lemmas 16 and 17 provide bounds on effective resistances of graph products that extend well-known and commonly used bounds for the $n \times n$ grid.

The lower bound in Theorem 5 implies that $\mathbf{COV}[G \Box H] \geq \mathbf{COV}[H]$ (and $\mathbf{COV}[G \Box H] \geq \mathbf{COV}[G]$), and the upper bound can be viewed as providing conditions sufficient for $\mathbf{COV}[G \Box H] = O(\mathbf{BCOV}[H])$ (or $\mathbf{COV}[G \Box H] = O(\mathbf{BCOV}[G])$). For example, since paths and cycles have $\mathbf{BCOV}[G] = \Theta(\mathbf{COV}[G])$, then $\mathbf{COV}[\mathbb{Z}_p \Box \mathbb{Z}_q] = \Theta(\mathbf{COV}[\mathbb{Z}_q]) = \Theta(q^2)$ subject to the condition $p \log^4 p = O(q)$. Thus for this example, the lower and upper bounds in Theorem 5 are within a constant factor.

Before we discuss the proof of Theorem 5 and the framework use to produce it, we discuss related work, and give examples of the application of the theorem to demonstrate how it extends that work.

5 Cover Time: Examples and Comparisons

In this section, we shall apply Theorem 5 to some examples and make comparisons to established results.

5.1 Two-dimensional Toroid with a Dominating Factor

We consider $\mathbf{COV}[\mathbb{Z}_p \Box \mathbb{Z}_q]$:

(i) $G \equiv \mathbb{Z}_p$;

(ii) $H \equiv \mathbb{Z}_q$

(iii) $\Delta_G = \Delta_H = \delta_G = \delta_H = 2$;

(iv) $\mathbf{BCOV}[H] = \Theta(\mathbf{COV}[H])$.

(v) $m_G = n_G = p$;

(vi) $m_H = n_H = q$;

(vii) $D_G = \lfloor \frac{p}{2} \rfloor$.

(viii) Thus $M = 2pq$, and

(ix) $\ell = \log(\lfloor \frac{p}{2} \rfloor + 1) \log(p \lfloor \frac{p}{2} \rfloor)$.

(x) $\mathbf{COV}[\mathbb{Z}_p] = \frac{p(p-1)}{2}$ and $\mathbf{COV}[\mathbb{Z}_q] = \frac{q(q-1)}{2}$.

Thus,

$$\begin{aligned}
\mathbf{COV}[F] &\leq K\left(\left(1+\frac{\Delta_G}{\delta_H}\right)\mathbf{BCOV}[H] + \frac{Mm_G m_H n_H \ell^2}{\mathbf{COV}[H]D_G}\right) \\
&= O\left(q^2 + \frac{2pqpqq\ell^2}{q^2 p}\right) \\
&= O\left(q^2 + pq\log^4 p\right) \\
&= O(q^2)
\end{aligned}$$

if $p\log^4 p = O(q)$.

Comparing this to the lower bound of Theorem 5,

$$\mathbf{COV}[F] \geq \max\left\{\left(\frac{\delta_G}{\Delta_H}+1\right)\mathbf{COV}[H], \left(\frac{\delta_H}{\Delta_G}+1\right)\mathbf{COV}[G]\right\},$$

which implies

$$\mathbf{COV}[F] = \Omega\left(\left(\frac{\delta_G}{\Delta_H}+1\right)\mathbf{COV}[H]\right) = \Omega(q^2).$$

Thus, Theorem 5 gives upper and lower bounds within a constant a multiple for this example. That is, it tells us $\mathbf{COV}[\mathbb{Z}_p \square \mathbb{Z}_q] = \Theta(\mathbf{COV}[\mathbb{Z}_q]) = \Theta(q^2)$ subject to the condition $p\log^4 p = O(q)$. Looking at it another way, it gives conditions for when the cover time of the product $F = G\square H$ is within a constant multiple of the cover time of one of it's factors. We describe that factor as the *dominating factor*.

6 Electrical Networks and Random Walks

In this section we give an introduction to the electrical network metaphor of random walks on graphs and present some of the concepts and results from the literature that are used in subsequent parts of this paper. Although a purely mathematical construction, the metaphor of electrical networks facilitates the expression of certain properties and behaviours of random walks on networks, and provides a language for which to describe these properties and behaviours. The classical treatment of the topic is [5]. The recent book [10] presents material within the more general context of Markov chains.

We first present some definitions: An *electrical network* is a connected, undirected, finite, graph $G = (V, E, c)$ where each edge $e \in E$ is has a strictly positive weight $c(e)$. The weight is called the *conductance*. Define the *resistance* $r(e) = r(u, v)$ of an edge $e = (u, v)$ as the inverse of the conductance: $r(e) = 1/c(e)$.

A random walk on an electrical network is a random walk on a weighted graph $G = (V, E, c)$, and is an example of a reversible Markov chain. Suppose the walk $\mathcal{W}$ is on a vertex u. Then if v is a neighbour, the probability of transitioning a particular edge $e = (u, v)$ (there may be more than one) is given by $c(e)/c(u)$ where $c(u) = \sum_{e:e=(u,x)} c(e)$. We also define $c(G) = \sum_{u \in V} c(u) = 2 \sum_{e \in E} c(e)$. In terms of random walks, an unweighted graph is the same as a uniformly weighted graph, that is, one where c is a constant function.

6.1 Effective Resistance

The theory of electrical networks defines a number of other notions that we describe informally, and we refer the reader to the book [5] for precise definitions and a comprehensive treatment of the subject. Our interest is in some key results, which we will state precisely.

Given two vertices a and z (not necessarily neighbours), we can set z to have a "ground", i.e., zero *potential* and a to have a positive potential (a value in $\mathbb{R}^+$). There is then a *potential difference* or *voltage* which develops "across" a and z, and causes a *current flow* from a to z via other other vertices. The magnitude of the current flow is a function of the voltage, the graph structure and the edge resistances. The consequence of this is that each vertex $u \in V \setminus \{a, z\}$ develops a potential that lies between those of a and z (mathematically, the potential is a harmonic function with boundary points a, z). Equivalently, we may specify a current flow in the graph from a to z, and the potential at of a vertex u is relative to z, which by convention is assumed to have a potential zero. The *effective resistance* $R(a, z)$ between a and z is the potential that develops at a when a current flow of unit magnitude is set from a to z. If a and z are connected then $R(a, z) \leq r(a, z)$, since generally speaking, the current has many paths to flow through from a to z besides the edge between them.

It is important to note that the *resistance* $r(u, v)$ of an edge (u, v) is different to the *effective resistance* $R(u, v)$ between the vertices u, v. Resistance $r(u, v)$ is $1/c(u, v)$, the inverse of conductance, which is part of the definition of the network $G = (V, E, c)$, and is the weighting function c defined on an *edge*. Effective resistance, on the other hand, is a property of the network, but not explicitly given in the tuple (V, E, c), and it is defined between *a pair of vertices*.

Theorem 6 (**Thomson's Principle**). *For any network $G = (V, E, c)$ and any pair of vertices $u, v \in V$,*

$$R(u, v) = \min\{\mathcal{E}(\varphi) : \varphi \text{ is a unit flow from } u \text{ to } v\}. \tag{5}$$

The unit current flow is the unique φ that gives the minimum element of the above set.

6.1.1 Rayleigh's Monotonicity Law, Cutting & Shorting

Rayleigh's Monotonicity Law, as well as the related Cutting and Shorting Laws, are intuitive principles that play important roles in our work. They are very useful means of making statements about bounds on effective resistance in a network when the network is somehow altered. With minor alterations of notation, we quote [10] Theorem 9.12, including proof.

Theorem 7 (**Rayleigh's Monotonicity Law**). *If $G = (V, E)$ is a network and c, c' are two different weightings of the network such that $r(e) \leq r'(e)$ for all $e \in E$, (recall, $r(e) = 1/c(e)$), then for any $u, v \in V$,*

$$R(u, v) \leq R'(u, v)$$

where $R(u, v)$ is the effective resistance between u and v under the weighting c (or r), and $R'(u, v)$ under weighting c' (or r').

Lemma 8 (**Cutting Law**). *Removing an edge e from a network cannot decrease the effective resistance between any vertices in the network.*

Lemma 9 (**Shorting Law**). *To short a pair of vertices u, v in a network G, replace u and v with a single vertex w and do the following with the edges: Replace each edge (u, x) or (v, x) where $x \notin \{u, v\}$ with an edge (w, x). Replace each edge (u, v) with a loop (w, w). Replace each loop (u, u) or (v, v) with a loop (w, w). A new edge has the same conductance as the edge it replaced. Let G' denote the network after this operation, and let R and R' represent effective resistance in G and G' respectively. Then, for a pair of vertices $a, z, \notin \{u, v\}$, $R'(a, z) \leq R(a, z)$, $R'(a, w) \leq R(a, u)$ and $R'(a, w) \leq R(a, v)$.*

Sometimes the Shorting Law is defined as putting a zero-resistance edge between u, v, but since zero-resistance (infinite-conductance) edges are not defined in our presentation, we refer to the act of "putting a zero-resistance edge" between a pair of vertices as a metaphor for shorting as defined above.

6.1.2 Commute Time Identity

The following theorem, first given in [2] is a fundamental tool in our analysis of random walks on graphs. It provides a link between random walks and electrical network theory.

Theorem 10 ([2]). *Let $G = (V, E, c)$ be a network. Then for a pair of vertices $u, v \in V$*

$$\mathbf{COM}[u, v] = c(G)R(u, v).$$

(The reader is reminded that $c(G) = \sum_{v \in V} c(v) = 2 \sum_{e \in E} c(e)$).

6.2 Parallel and Series Laws

The parallel and series laws are rules that establish equivalences between certain structures in a network. They are useful for reducing a network G to a different form G', where the latter may be more convenient to analyse. We quote from [10], with minor modifications for consistency in notation.

Lemma 11 (**Parallel Law**). *Conductances in parallel add.*

Suppose edges e_1 and e_2, with conductances $c(e_1)$ and $c(e_2)$ respectively, share vertices u and v as endpoints. Then e_1 and e_2 can be replaced with a single edge e with $c(e) = c(e_1) + c(e_2)$, without affecting the rest of the network. All voltages and currents in $G \setminus \{e_1, e_2\}$ are unchanged and the current $I(\overrightarrow{e}) = I(\overrightarrow{e_1}) + I(\overrightarrow{e_2})$. For a proof, check Ohm's and Kirchhoff's laws with $I(\overrightarrow{e}) = I(\overrightarrow{e_1}) + I(\overrightarrow{e_2})$.

Lemma 12 (**Series Law**). *Resistances in series add.*

If $v \in V \setminus \{a, z\}$, where a and z are source and sink, is a node of degree 2 with neighbours v_1 and v_2, the edges (v_1, v) and (v, v_2) can be replaced with a single edge (v_1, v_2) with resistance $r(v_1, v_2) = r(v_1, v) + r(v, v_2)$. All potentials and currents in $G \setminus \{v\}$ remain the same and the current that flows from v_1 to v_2 is $I(\overrightarrow{v_1, v_2}) = I(\overrightarrow{v_1, v}) = I(\overrightarrow{v, v_2})$. For a proof, check Ohm's Law and Kirchhoff's Law with $I(\overrightarrow{v_1, v_2}) = I(\overrightarrow{v_1, v}) = I(\overrightarrow{v, v_2})$.

7 Preliminaries

7.1 Some Notation

For clarity, and because a vertex u may be considered in two different graphs, we may use $d_G(u)$ to explicitly denote the degree of u in graph G.

$h(n)$ denotes the n'th harmonic number, that is, $h(n) = \sum_{i=1}^{n} 1/i$. Note $h(n) = \log n + \gamma + O(1/n)$ where $\gamma \approx 0.577$. All logarithms are base-e.

In the notation $(., y)$, the '.' is a place holder for some unspecified element, which may be different from one tuple to another. For example, if we refer to two vertices $(., a), (., b) \in G\square H[S]$, the first elements of the tuples may or may not be the same, but $(., a)$, for example, refers to a *particular* vertex, not a set of vertices $\{(x, a) : a \in V(G)\}$.

7.2 The Square Grid

The $k \times k$ grid graph P_k^2, where P_k is the k-path, plays an important role in our work. We shall analyse random walks on subgraphs isomorphic to this structure. It is well known in the literature (see, e.g. [5], [10]) that for any pair of vertices $u, v \in V(P_k^2)$, we have $R(u, v) \leq C \log k$ where C is some universal constant. We shall quote part of [8] Lemma 3.1 in our notation and refer the reader to the proof there.

Lemma 13 ([8], Lemma 3.1(a))**.** *Let u and v be any two vertices of P_k^2. Then $R(u, v) < 8h(k)$, where $h(k)$ is the k'th harmonic number.*

8 Locally Observed Random Walk

Let $G = (V, E)$ be a connected, unweighted (equiv., uniformly weighted) graph. Let $S \subset V$ and let $G[S]$ be the subgraph of G induced by S. Let $B = \{v \in S : \exists x \notin S, (v, x) \in E\}$. Call B the *boundary* of S, and the vertices of $V \setminus S$ *exterior vertices*. If $v \in S$ then $d_G(v)$ (the degree of v in G) is partitioned into $d(v, in) = |N(v, in)| = |N(v) \cap S|$ and $d(v, out) = |N(v, out)| = |N(v) \cap (V \setminus S)|$, (*inside* and *outside* degree). Here $N(v)$ denotes the neighbour set of v.

Let $u, v \in B$. Say that u, v are *exterior-connected* if there is a (u, v)-path $u, x_1, ...x_k, v$ where $x_i \in V \setminus S, k \geq 1$. Thus all vertices of the path except u, v are exterior, and the path contains at least one exterior vertex. Let $A(B) = \{(u, v) : u, v \text{ are exterior-connected } \}$. Note $A(B)$ may include self-loops.

Call edges of $G[S]$ *interior*, edges of $A(B)$ *exterior*. We say that a walk $\omega = (u, x_1, ...x_k, v)$ on G is an exterior walk if $u, v \in S$ and $x_i \notin S$, $1 \le i \le k$.

We derive a weighted multi-graph H from G and S as follows: $V(H) = S$, $E(H) = E(G[S]) \cup A(B)$. Note if $u, v \in B$ and $(u, v) \in E$ then $(u, v) \in E(G[S])$, and if, furthermore, u, v are exterior connected, then $(u, v) \in A(B)$ *and these edges are distinct*, hence, H may not only have self-loops but also parallel edges, i.e., $E(H)$ is a multiset.

Associate with an orientation $(\vec{u,v})$ of an edge $(u, v) \in A(B)$ the set of all exterior walks $\omega = (u, x_1, ...x_k, v)$, $k \ge 1$ that start at u and end at v, and associate with each such walk the value $p(\omega) = 1/(d_G(u)d_G(x_1)...d_G(x_k))$ (note, the $d(x_i)$ is not ambiguous, since $x_i \notin V(H)$, but we leave the 'G' subscript in for clarity). This is precisely the probability that the walk ω is taken by a simple random walk on G starting at u. Let

$$p_H(\vec{u,v}) = \sum_{k \ge 1} \sum_{\omega=(u,x_1...x_k,v)} p(\omega),$$

where the sum is over all exterior walks ω.

We set the edge conductances (weights) of H as follows: If e is an interior edge, $c(e) = 1$. If it is an exterior edge $e = (u, v)$ define $c(e)$ as

$$c(e) = d_G(u)p_H(\vec{u,v}) = \sum_{k \ge 1} \sum_{\omega=(u,x_1...x_k,v)} \frac{1}{d_G(x_1)...d_G(x_k)} = d_G(v)p_H(\vec{v,u}).$$

Thus the edge weight is consistent. A weighted random walk on H is thus a finite reversible Markov chain with all the associated properties that this entails.

Definition 4. *The weighted graph H derived from (G, S) is termed* the local observation of G at S*, or* G locally observed at S*. We shall denote it as $H = Loc(G, S)$.*

The intuition in the above is that we wish to observe a random walk $\mathcal{W}(G)$ on a subset S of the vertices. When $\mathcal{W}(G)$ makes an external transition at the border B, we cease observing and resume observing if/when it returns to the border. It will thus appear to have transitioned a virtual edge between the vertex it left off and the one it returned on. It will therefore appear to be a weighted random walk on H. This equivalence is formalised thus

Definition 5. *Let G be a graph and $S \subset V(G)$. For an (unweighted) random walk $\mathcal{W}(G)$ on G starting at $x_0 \in S$, derive the Markov chain $\mathcal{M}(G, S)$ on the states of S as follows: (i) $\mathcal{M}(G, S)$ starts on x_0 (ii) If $\mathcal{W}(G)$ makes a transition through an internal edge (u, v) then so does $\mathcal{M}(G, S)$ (iii)If $\mathcal{W}(G)$ takes an exterior walk $\omega = (u, x_1...x_k, v)$ then $\mathcal{M}(G, S)$ remains at u until the walk is complete and subsequently transitions to v. We call $\mathcal{M}(G, S)$* the local observation of $\mathcal{W}(G)$ at S*, or* $\mathcal{W}(G)$ locally observed at S.

Lemma 14. *For a walk $\mathcal{W}(G)$ and a set $S \subset V(G)$, the local observation of $\mathcal{W}(G)$ at S, $\mathcal{M}(G, S)$ is equivalent to the weighted random walk $\mathcal{W}(H)$ where $H = Loc(G, S)$.*

Proof. The states are clearly the same so it remains to show that the transition probability $P_{\mathcal{M}}(u,v)$ from u to v in $\mathcal{M}(G,S)$ is the same as $P_{\mathcal{W}(H)}(u,v)$ in $\mathcal{W}(H)$. Recall that B is the border of the induced subgraph $G[S]$. If $u \notin B$ then an edge $(u,v) \in E(H)$ is internal and so has unit conductance in H, as it does in G. Furthermore, for an internal edge e, $e \in E(H)$ if and only if $e \in E(G)$, thus $d_H(u) = d_G(u)$ when $u \notin B$. Therefore $P_{\mathcal{W}(H)}(u,v) = 1/d_H(u) = 1/d_G(u) = P_{\mathcal{M}}(u,v)$.

Now suppose $u \in B$. Let $E(u)$ denote the set of all edges incident with u in H and recall $A(B)$ above is the set of exterior edges. The total conductance (weight) of the exterior edges at u is

$$\begin{aligned}\sum_{e\in E(u)\cap A(B)} c_H(e) &= \sum_{x\in N(u,out)} \sum_{v\in B} \mathbf{Pr}(\text{walk from } x \text{ returns to } B \text{ at } v)\\ &= \sum_{x\in N(u,out)} 1\\ &= d(u,out).\end{aligned}$$

(Note the H subscript in $c_H(e)$ above is redundant since exterior edges are only defined for H, but we leave it for clarity).

Thus for $u \in B$

$$\begin{aligned}c_H(u) = \sum_{e\in E(u)} c_H(e) &= \sum_{e\in E(u)\cap G[S]} 1 + \sum_{e\in E(u)\cap A(B)} c_H(e)\\ &= d(u,in) + d(u,out)\\ &= d_G(u).\end{aligned}$$

Now

$$P_{\mathcal{M}}(u,v) = \mathbf{1}_{\{(u,v)\in G[S]\}}\frac{1}{d_G(u)} + \sum_{k\geq 1} \sum_{\omega=(u,x_1\ldots x_k,v)} \frac{1}{d_G(u)d_G(x_1)\ldots d_G(x_k)}$$

where the sum is over all exterior walks ω. Thus

$$P_{\mathcal{M}}(u,v) = \mathbf{1}_{\{(u,v)\in G[S]\}}\frac{1}{d_G(u)} + p_H(\vec{u,v})$$

$$\begin{aligned}P_{\mathcal{W}(H)}(u,v) &= \frac{1}{c_H(u)}\left[\mathbf{1}_{\{(u,v)\in G[S]\}} + \mathbf{1}_{\{(u,v)\in A(B)\}}c_H(u,v)\right]\\ &= \frac{1}{d_G(u)}\left[\mathbf{1}_{\{(u,v)\in G[S]\}} + \mathbf{1}_{\{(u,v)\in A(B)\}}d_G(u)p_H(\vec{u,v})\right]\\ &= \mathbf{1}_{\{(u,v)\in G[S]\}}\frac{1}{d_G(u)} + \mathbf{1}_{\{(u,v)\in A(B)\}}p_H(\vec{u,v})\\ &= P_{\mathcal{M}}(u,v).\end{aligned}$$

$\square$

9 Effective Resistance Lemmas

For the upper bound of Theorem 5, we require the following lemmas.

Lemma 15. *Let G be an undirected graph. Let $G' \subseteq G$, be any subgraph such that such that $V(G') = V(G)$. For any $u, v \in V(G)$,*

$$R(u, v) \leq R'(u, v)$$

where $R(u, v)$ is the effective resistance between u and v in G and $R'(u, v)$ similarly in G'.

Proof. Since $V(G') = V(G)$, G' can be obtained from G by only removing edges. The lemma follows by the Cutting Law (Lemma 8). □

Denote by $R_{max}(G)$ the maximum effective resistance between any pair of vertices in a graph G.

Lemma 16. *For a graph G and tree T, $R_{max}(G \square T) < 4R_{max}(G \square P_r)$ where $|V(T)| \leq r \leq 2|V(T)|$ and P_r is the path on r vertices.*

Proof. Note first the following:

(i) By the parallel law, an edge (a, b) of unit resistance can be replaced with two parallel edges between a, b, each of resistance 2.

(ii) By the shorting law, a vertex a can be replaced with two vertices a_1, a_2 with a zero-resistance edge between them and the ends of edges incident on a distributed arbitrarily between a_1 and a_2.

(iii) By the same principle of the cutting law, this edge can be broken without decreasing effective resistance between any pair of vertices.

Transformations **(i)** and **(ii)** do not alter the effective resistance $R(u, v)$ between a pair of vertices u, v in the network. For any vertex $u \notin \{a_1, a_2, a\}$, $R(u, a_1) = R(u, a_2)$ and these are equal to $R(u, a)$ before the operation.

Points **(ii)** and **(iii)** require elaboration. In this paper, we do not define zero-resistance (infinite conductance) edges. As stated in section 6.1.1, to say that a zero-resistance edge is placed between a_1 and a_2, is another way of referring to shorting as defined in Lemma 9. It would seem then, that **(ii)**, in fact says nothing. However, it serves as a useful short hand for talking about operations on the graph when used in conjunction with **(iii)**. If **(ii)** and **(iii)** are always used together, that is, if a zero-resistance edge created from **(ii)** is always cut by **(iii)**, then this is equivalent to the reverse of process of shorting two vertices a_1 and a_2 into a_3, as per Lemma 9. Hence, these two operations together are sound.

We continue thus:

1. Let $F = G \square T$. Let each edge of F have unit resistance. In what follows, we shall modify F, but shall continue to refer to the modified graphs as F.

2. Starting from some vertex v in T, perform a depth-first search (DFS) of T stopping at the first return to v after all vertices in T have been visited. Each edge of T is traversed twice; once in each orientation. Each vertex x will be visited $d(x)$ times.

3. Let (e_i) be the sequence of oriented edges generated by the search. The idea is to use (e_i) to construct a transformation from $F = G\square T$ to $G\square P_r$. From (e_i), we derive another sequence (a_i), which is generated by following (e_i) and if we have edges e_i, e_{i+1} with $e_i = (a, b)$, $e_{i+1} = (b, c)$ such that it is neither the first time nor the last time b is visited in the DFS, then we replace e_i, e_{i+1} with (a, c). We term such an operation an *aggregation*. Observe that in the sequence (a_i), all leaf vertices of T appear only once (just as in (e_i)), and a non-leaf vertex appears twice.

4. By **(i)** above, we can replace each (unit resistance) edge in F by a pair of parallel edges each of resistance 2.

5. For a pair of parallel edges in the T dimension, arbitrarily label one of them with an orientation, and label the other with the opposite orientation. Note, orientations are only an aid to the proof, and are not a flow restriction. We therefore see that (e_i) can be interpreted as a sequence of these parallel oriented edges.

6. We further modify F using (a_i): If $(a, b), (b, c)$ was aggregated to (a, c), then replace each pair of **oriented** edges $((x, a), (x, b))$ and $((x, b), (x, c))$ in F with an oriented edge $((x, a), (x, c))$. The resistances of $((x, a), (x, b))$ and $((x, b), (x, c))$ were $r((x, a), (x, b)) = 2$ and $r((x, b), (x, c)) = 2$. Set the resistance $r((x, a), (x, c)) = r((x, a), (x, b)) + r((x, b), (x, c)) = 4$.

7. The above operation is the same as restricting flow through $((x, a), (x, b))$ and $((x, b), (x, c))$ to only going from one to the other at vertex (x, b), without the possibility of going through other edges. The infimum of the energies of this subset of flows is at least the infimum of the energies of the previous set and so by Thomson's principle, the effective resistance cannot be decreased by this operation.

8. For each copy G_i of G in F excluding those that correspond to a leaf of T, we can create a "twin" copy G'_i. Associate with each vertex $x \in V(F)$ (except those excluded) a newly-created twin vertex x' with no incident edges. Thus, $V(G_i)$ has a twin set $V(G'_i)$, though the latter has no edges yet.

9. Recall the parallel edges created initially from all the edges of F; we did not manipulate those in the G dimension, but we do so now: redistribute half of the parallel edges of G_i in the G dimension to the set of twin vertices $V(G'_i)$ so as to make G'_i a copy of G (isomorphic to it). Now put a zero-resistance edge between x and x'. By **(ii)**, effective resistance is unchanged by this operation.

10. We now redistribute the oriented parallel edges in the T dimension so as to respect the sequence (a_i). We do this as follows: follow the sequence (a_i) by traversing edges in their orientation. Consider the following event: In the sequence (a_i) there is an element $a_j = (a, b)$ and b has appeared in some element a_i such that $i < j$. Then a_j is the second time that b has occurred in the sequence. Now change each edges $((x, a), (x, b)) \in F$ to $((x, a), (x, b)')$. If $b = v$, then stop; otherwise, a_j is followed by $a_{j+1} = (b, c)$, for some $c \in V(T)$. In this case, also change all $((x, b), (x, c)) \in F$ to $((x, b)', (x, c))$. Continue in the same manner to the end

of the sequence (a_i).

11. We then remove the zero-resistance edges between each pair of twin vertices, and by **(iii)**, this cannot decrease the effective resistance.

Using the sequence (a_i) to trace a path of copies of G, we see that the resulting structure is isomorphic to $G \Box P_r$. Since the aggregation process only aggregates edges that pass through a previously seen vertex, r is at least $|V(T)|$. Also, because each edge is traversed at most once in each direction, r is at most $2|V(T)|$. Each edge has resistance at most 4, and so the lemma follows. $\square$

Lemma 17. *For graphs G, H suppose $D_G+1 \leq n_H \leq \alpha(D_G+1)$, for some α. Then $R_{max}(G \Box H) < \zeta\alpha \log(D_G + 1)$, where ζ is some universal constant.*

Proof. Let $(a, x), (b, y)$ be any two vertices in $G \Box H$. Let D be some diametric path of G. Let $\langle a, D\rangle$ represent the shortest path from a to D in G (which may trivially be a if it is on D). Similarly with $\langle b, D\rangle$. Let $T_D = D \cup \langle a, D\rangle \cup \langle b, D\rangle$. Let $k = D_G + 1$. Note $k \leq |V(T_D)| \leq 3k$. Now let T_H be any spanning tree of H. Applying Lemma 16 twice we have

$$R_{max}(T_D \Box T_H) < 4R_{max}(T_D \Box P_s) < 16R_{max}(P_r \Box P_s)$$

where $k \leq r \leq 6k$ and $k \leq s \leq 2\alpha k$. Considering a series of connected P_k^2 subgraphs and using Lemma 13 and the triangle inequality for effective resistance, we have $R_{max}(P_r \Box P_s) \leq 16(6 + 2\alpha)8h(k)$, where $h(k)$ is the k'th harmonic number. Since $T_D \Box T_H \subseteq G \Box H$, the lemma follows by Lemma 15. $\square$

A diametric path D is involved in the proof of Lemma 17 because the use of D means that the dimension of P_r is effectively maximised, and we can break up the grid $P_r \Box P_s$ roughly into $k \times k$ square grids, each with maximum effective resistance $O(\log k) = O(\log D_G)$. If, for example, the shortest path between a and b is used, the product $P_r \Box P_s$ may have r much smaller than s, looking like a long thin grid, which may have a high effective resistance.

10 A General bound

In this section, we prove Theorem 5, starting with the lower bound.

10.1 Lower Bound

Proof. In order for the walk $\mathcal{W}$ to cover F, it needs to have covered the H dimension of F. That is, each copy of G in F needs to have been visited at least once. The probability of a transition in the H dimension is distributed as a geometric random variable with success probability at most $\frac{\Delta_H}{\Delta_H+\delta_G}$. Thus, the expectation of the number of steps of $\mathcal{W}$ per transition in the H dimension is at least $\frac{\Delta_H+\delta_G}{\Delta_H}$. Transitions of $\mathcal{W}$ in the H dimension are independent of the location of $\mathcal{W}$ in the

G dimension, and have the same distribution (in the H dimension) as a walk on H. This proves

$$\mathbf{COV}[F] \geq \left(1 + \frac{\delta_G}{\Delta_H}\right) \mathbf{COV}[H].$$

By commutativity,

$$\mathbf{COV}[F] \geq \left(1 + \frac{\delta_H}{\Delta_G}\right) \mathbf{COV}[G].$$

□

10.2 Upper Bound

The following proves the upper bound in Theorem 5. It is envisaged that theorem is used with the idea in mind that G is small relative to H, and so the cover time of the product is essentially dominated by the cover time of H.

Proof. Let $k = D_G + 1$. We group the vertices of H into sets such that for any set S and the subgraph of H induced by S, $H[S]$: (i)$|S| \geq k$, (ii)$H[S]$ is connected, (iii) The diameter of $H[S]$ is at most $4k$. We do this through the following decomposition algorithm on H: Choose some arbitrary vertex $v \in V(H)$ as the root, and using a breadth-first search (BFS) on H, descend from v at most distance k. The resulting tree $T(v) \subseteq H$ will have diameter at most $2k$. For each leaf l of $T(v)$, continue the BFS using l as a root. If $T(l)$ has fewer than k vertices, append it to $T(v)$. If not, recurse on the leaves of $T(l)$. The set of vertices of each tree thus formed satisfies the three conditions above. The root is part of a new set, unless it has been appended to another tree.

In the product F we refer to copies of G as *columns*. In F we have a natural association of each column with the set $S \subseteq V(H)$ defined above. We define $Block[S] = (G \square H[S])$.

[Refer to section 7.1 for a reminder of the notation $(., y)$]. For any two vertices $(., a), (., b) \in G \square H[S]$ there exists a tree $T\langle a, b\rangle$ subgraph of the tree T in H that generated S such that a and b are connected in $T\langle a, b\rangle$ and $k \leq |V(T\langle a, b\rangle)| \leq 4k$. Then using Lemmas 17 and 15, we can upper bound the effective resistance $R((., a), (., b))$ in $B = Block[S]$,

$$R_{max}(B) \leq 4\zeta \log(D_G + 1). \tag{6}$$

Furthermore, if $B' = Loc(F, V(B))$ (*Loc* is defined in Definition 4), then $B \subseteq B'$ so by Lemma 15,

$$R_{max}(B') \leq 4\zeta \log(D_G + 1). \tag{7}$$

We use the following two-phase approach to bound the cover time of $F = G \square H$.

Phase 1 Perform a random walk $\mathcal{W}(F)$ on F until the blanket-cover criterion is satisfied for the H dimension.

Phase 2 Starting from the end of phase 1, perform a random walk on F until all vertices of F not visited in phase 1 are visited.

Phase 1 can be thought of in the following way: We couple $\mathcal{W}(F)$ with a walk $\mathcal{W}(H)$ such that (i) if $\mathcal{W}(F)$ starts at $(., x)$, then $\mathcal{W}(H)$ starts at x, and (ii) $\mathcal{W}(H)$ moves to a new vertex y from a vertex x when and only when $\mathcal{W}(F)$ moves from $(., x)$ to $(., y)$. This coupled process runs until $\mathcal{W}(H)$ satisfies the blanket-cover criteria for H, i.e., when each vertex $v \in V(H)$ has been visited at least $\pi(v)\mathbf{COV}[H]$ times. An implication is that the corresponding column G_v in F will have been visited at least that many times.

Having grouped F into blocks, we analyse the outcome of phase 1 by relating $\mathcal{W}(F)$ to the local observation on each block. A particular block B will have some vertices unvisited by $\mathcal{W}(F)$ if and only if $\mathcal{W}(F)$ locally observed on B fails to visit all vertices. We refer to such a block as *failed*. Consider the weighted random walk $\mathcal{W}(B')$ on $B' = Loc(F, V(B))$. This has the same distribution as $\mathcal{W}(F)$ locally observed on B. Hence, we bound the probability of $\mathcal{W}(F)$ failing to cover B by bounding the probability that $\mathcal{W}(B')$ fails to cover B'. Done for all blocks, we can bound the expected time it takes phase 2 to cover the failed blocks. We think of phase 1 as doing most of the "work", and phase 2 as a "mopping up" phase. Mopping up a block in phase 2 is costly, but if there are few of them, the overall cost is within a small factor of phase 1.

We bound $\mathbf{Pr}(\mathcal{W}(B')$ fails) by exploiting the fact that $\mathcal{W}(B')$ will have made some minimal number of transitions t. This is guaranteed because phase 1 terminates only when $\mathcal{W}(H)$ has satisfied the blanket-cover criterion on H. If κ counts the number of steps of a walk $\mathcal{W}(B')$ until B' is covered, then

$$\mathbf{Pr}(\mathcal{W}(B') \text{ fails to cover} B') \leq \mathbf{Pr}(\kappa > t) \leq \frac{\mathbf{E}[\kappa]}{t} \tag{8}$$

by Markov's inequality.

Definition 6. *For graphs $I = J \square K$, and $S \subseteq V(I)$, denote by $S.K$ the* projection of S on to K*, that is, $S.K = \{v \in K : (., v) \in S\}$.*

For a weighted graph G, recall that $c(G)$ is the twice the sum of the conductances (weights) of all edges of G. By the definition of $G \square H[S]$ and section 8,

$$c(B') \leq m_G |V(B).H| + n_G \sum_{u \in V(B).H} d(u). \tag{9}$$

Using (7) and Theorem 10 we therefore have for any $u, v \in V(B')$, $\mathbf{COM}[u, v] \leq Kc(B') \log(D_G + 1)$ for some universal constant K. (In what follows K will change, but we shall keep the same symbol, with an understanding that what we finish with is a universal constant). Hence, by 3,

$$\mathbf{COV}[B'] \leq Kc(B') \log(D_G + 1) \log(|V(B')|).$$

For a block B, the number of transitions on the H dimension - and therefore the number of transitions on B - as demanded by the blanket-cover criterion is at least

$$\tau = \sum_{u \in V(B).H} \pi_H(u)\mathbf{COV}[H] = \frac{\mathbf{COV}[H]}{2m_H} \sum_{u \in V(B).H} d_H(u), \tag{10}$$

where $\pi_H(u)$ and $d_H(u)$ denote the stationary probability and degree of u in H.

Now

$$\begin{aligned}
\mathbf{Pr}(\mathcal{W}(F) \text{ fails on } B) &= \mathbf{Pr}(\mathcal{W}(B') \text{ fails on } B') \\
&\le Kc(B')\log(D_G+1)\log(|V(B')|)/\tau, \qquad (11)
\end{aligned}$$

as per (8). For convenience, we left $l_B = \log(D_G+1)\log(|V(B)|)$ (recall $V(B) = V(B')$). Hence, using (10) with (9) and (11),

$$\begin{aligned}
\mathbf{Pr}(\mathcal{W}(F) \text{ fails on } B) &\le \frac{Kl_B m_H}{\mathbf{COV}[H]} \frac{m_G|V(B).H| + n_G \sum_{u\in V(B).H} d(u)}{\sum_{u\in V(B).H} d_H(u)} \\
&= \frac{Kl_B m_H}{\mathbf{COV}[H]} \left(n_G + \frac{m_G|V(B).H|}{\sum_{u\in V(B).H} d_H(u)} \right). \qquad (12)
\end{aligned}$$

Phase 2 consists of movement between failed blocks, and covering a failed block it has arrived at. The total block-to-block movement is upper bounded by the time is takes to cover the H dimension of F (in other words, for each column to have been visited at least once). We denote this by $\mathbf{COV}_F[H]$. Let $\mathbf{COV}_F[B]$ denote the cover time of the set of vertices of a block B by the walk $\mathcal{W}(F)$. Let the random variables ϕ_1 and ϕ_2 represent the time it takes to complete phase 1 and phase 2 respectively.

$$\mathbf{E}[\phi_2] \le \mathbf{COV}_F[H] + \sum_{B\in F} \mathbf{Pr}(\mathcal{W}(F) \text{ fails on } B)\mathbf{COV}_F[B].$$

For $\mathcal{W}(H)$, the random variable $\beta_H = \min\{t : (\forall v) N_v(t) \ge \pi(v)\mathbf{COV}[H]\}$ counts the time it takes to satisfy the blanket-cover criterion on H.

The expected number of movements on F per movement on the H dimension is at most $(\Delta_G + \delta_H)/\delta_H$. Therefore,

$$\mathbf{E}[\phi_1] \le \frac{\Delta_G + \delta_H}{\delta_H}\mathbf{E}[\beta_H] = \frac{\Delta_G + \delta_H}{\delta_H}\mathbf{BCOV}[H].$$

Similarly,

$$\mathbf{COV}_F[H] \le \frac{\Delta_G + \delta_H}{\delta_H}\mathbf{COV}[H].$$

Using (6), Lemma 15, and Theorems 10 and 3 on B, we have

$$\mathbf{COV}_F[B] \le K'c(F)l_B \qquad (13)$$

where $c(F) = 2|E(F)| = 2M$.

Hence,

$$\begin{aligned}\mathbf{COV}[F] &\leq \mathbf{E}[\phi_1]+\mathbf{E}[\phi_2]\\ &\leq K\frac{\Delta_G+\delta_H}{\delta_H}\mathbf{BCOV}[H]+\sum_{B\in F}\mathbf{Pr}(\mathcal{W}(F)\text{ fails on }B)\mathbf{COV}_F[B].\end{aligned}$$

We have, using (12) and (13),

$$\sum_{B\in F}\mathbf{Pr}(\mathcal{W}(F)\text{ fails on }B)\mathbf{COV}_F[B]\leq K\frac{Mm_H}{\mathbf{COV}[H]}\sum_{B\in F}\left(n_G+\frac{m_G|V(B).H|}{\sum_{u\in V(B).H}d_H(u)}\right)l_B^2. \tag{14}$$

Since $\sum_{u\in V(B).H}d(u)\geq |V(B).H|$, the outer summation in (14) can be bounded thus

$$\sum_{B\in F}\left(n_G+\frac{m_G|V(B).H|}{\sum_{u\in V(B).H}d_H(u)}\right)l_B^2\leq m_G\log^2(D_G+1)\sum_{B\in F}\log^2(|V(B)|). \tag{15}$$

Since each block $B\in F$ has at least D_G+1 columns, we can upper bound the sum in the RHS of (15) by assuming all blocks have this minimum. The number of such blocks in F will be $|V(H)|/(D_G+1)$, each block having $(D_G+1)n_G$ vertices. Hence

$$\sum_{B\in F}\log(|V(B)|)^2\leq\frac{n_H}{D_G}\log^2(n_G(D_G+1)). \tag{16}$$

Putting together (14), (15) and (16), we get

$$\sum_{B\in F}\mathbf{Pr}(\mathcal{W}(F)\text{ fails on }B)\mathbf{COV}_F[B]\leq K\frac{Mm_Gm_Hn_H\ell^2}{\mathbf{COV}[H]D_G}$$

where $\ell=\log(D_G+1)\log(n_GD_G)$. □

References

[1] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász and C. Rackoff, Random walks, universal traversal sequences, and the complexity of maze problems. *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science*, 218-223 (1979).

[2] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, P. Tiwari, The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6 312-340 (1997)

[3] A. Dembo, Y. Peres, J. Rosen, O. Zeitouni, Cover times for Brownian motion and random walks in two dimensions. *Ann. Math.*, 160 433-464 (2004).

[4] J. Ding, J. Lee, Y. Peres, Cover times, blanket times, and majorizing measures. *Annals of Mathematics*, (to appear).

[5] P.G. Doyle and J.L. Snell, *Random Walks and Electrical Networks*, (2006).

[6] U. Feige, A tight upper bound for the cover time of random walks on graphs. *Random Structures and Algorithms*, 6 51-54 (1995).

[7] U. Feige, A tight lower bound for the cover time of random walks on graphs. *Random Structures and Algorithms*, 6 433-438 (1995).

[8] J. Jonasson, An upper bound on the cover time for powers of graphs. *Discrete Mathematics*, 222 181-190 (2000).

[9] J. Kahn, J. H. Kim, L. Lovász, V. H. Vu. The cover time, the blanket time, and the Matthews bound. *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, 4675. IEEE Comput. Soc. Press, Los Alamitos, CA (2000).

[10] D. A. Levin, Y. Peres, E. L. Wilmer. *Markov Chains and Mixing Times*, AMS Press (2008).

[11] P. Matthews, Covering problems for Brownian motion on spheres. *Ann. Prob.*, 16:189-199 (1988).

[12] P. Winkler, D. Zuckerman, Multiple cover time. *Random Structures and Algorithms*, 9 403-411 (1996).