Chapter 6
Cover Times

David Aldous
Department of Statistics
University of California
Berkeley, CA 94720

James Allen Fill
Department of Mathematical Sciences
The Johns Hopkins University
Baltimore, MD 21218-2692

October 31, 1994

The maximal mean hitting time $\max_{i,j} E_i T_j$ arises in many contexts. In Chapter 5 we saw how to compute this in various simple examples, and the discussion of τ^* in Chapter 4 indicated general methods (in particular, the electrical resistance story) for upper bounding this quantity. But what we've done so far doesn't answer questions like "how large can τ^* be, for random walk on a n-vertex graph". Such questions are dealt with in this Chapter, in parallel with a slightly different topic. The *cover time* for a n-state Markov chain is the random time C taken for the entire state-space I to be visited. Formally,

$$C \equiv \max_j T_j.$$

It is sometimes mathematically nicer to work with the "cover-and-return" time

$$C^+ \equiv \min\{t \geq C : X_t = X_0\}.$$

There are several reasons why cover times are interesting.

- Several applications involve cover times directly: graph connectivity algorithms (section 8.2), universal traversal sequences (section 8.1), the "white screen problem" (Chapter 1 yyy)
- There remains an interesting "computability" open question (section 8.3)
- In certain "critical" graphs, the uncovered subset at the time when the graph is almost covered is believed to be "fractal" (see the Notes on Chapter 7).

We are ultimately interested in random walks on unweighted graphs, but some of the arguments have as their natural setting either reversible Markov chains or general Markov chains, so we sometimes switch to those settings. Results are almost all stated for discrete-time walks, but we occasionally work with continuized chains in the proofs, or to avoid distracting complications in statements of results. Results often can be simplified or sharpened under extra symmetry conditions, but such results and examples are deferred until Chapter 7.

xxx contents of chapter

1 The spanning tree argument

Except for Theorem 1, we consider in this section random walk on an n-vertex unweighted graph. Results can be stated in terms of the number of edges $|\mathcal{E}|$ of the graph, but to aid comparison with results involving minimal or maximal degree it is helpful to state results in terms of average degree $\bar{d}$:

$$\bar{d} = \frac{2|\mathcal{E}|}{n}; \qquad |\mathcal{E}| = n\bar{d}/2.$$

The argument for Theorem 1 goes back to Aleliunas et al [3]. Though elementary, it can be considered the first (both historically and logically) result which combines Markov chain theory with graph theory in a non-trivial way.

Consider random walk on a weighted graph. Recall from Chapter 3 yyy the *edge-commute inequality*: for an edge (v, x)

$$\begin{aligned} E_v T_x + E_x T_v &\le w/w_{vx} \quad \text{(weighted)} && (1) \\ &\le \bar{d}n \quad \text{(unweighted)}. && (2) \end{aligned}$$

One can alternatively derive these inequalities from the commute interpretation of resistance (Chapter 3 yyy), since the resistance between x and v is at most $1/w_{vx}$.

Theorem 1 *For random walk on a weighted graph,*

$$\max_v E_v C^+ \le w \min_{\mathcal{T}} \sum_{e \in \mathcal{T}} 1/w_e$$

where the min *is over spanning trees $\mathcal{T}$. In the unweighted case*

$$\max_v E_v C^+ \le \bar{d}n(n-1).$$

Proof. Given a spanning tree $\mathcal{T}$ and a vertex v, there is a path $v = v_0, v_1, \ldots, v_{2n-2} = v$ which traverses each edge of the tree once in each direction, and in particular visits every vertex. So

$$\begin{aligned} E_v C^+ &\leq \sum_{j=0}^{2n-3} E_{v_j} T_{v_{j+1}} \\ &= \sum_{e=(v,x)\in\mathcal{T}} (E_v T_x + E_x T_v) \\ &\leq \sum_{e\in\mathcal{T}} w/w_e \text{ by (1)} \end{aligned}$$

This gives the weighted case, and in the unweighted case $w = \bar{d}n$ and each spanning tree has $\sum_{e\in\mathcal{T}} 1/w_e = n-1$. $\square$

Note that in the unweighted case, the bound is at most $n(n-1)^2$. On the barbell (Chapter 5 Example yyy) it is easy to see that $\min_i E_i C = \Omega(n^3)$, so the maximal values of any formalization of "mean cover time", over n-vertex graphs, is $\Theta(n^3)$. Results and conjectures on the optimal numerical constants in the $\Theta(n^3)$ upper bounds are given in section 3.

Corollary 2 *On an unweighted n-vertex tree, $E_v C_v^+ \leq 2(n-1)^2$, with equality iff the tree is the n-path and v is a leaf.*

Proof. The inequality follows from Theorem 1. On the n-path with leaves v, z we have $E_v C_v^+ = E_v T_z + E_z T_v = 2(n-1)^2$. $\square$

It is worth dissecting the proof of Theorem 1. Two different inequalities are used in the proof. Inequality (2) is an equality iff the edge is essential, so the second inequality in the proof is an equality iff the graph is a tree. But the first inequality in the proof bounds C^+ by the time to traverse a spanning tree *in a particular order*, and is certainly not sharp on a general tree, but only on a path. This explains Corollary 2. More importantly, these remarks suggest that the bound $\bar{d}n(n-1)$ in Theorem 1 will be good iff there is some fixed "essential path" in the graph, and the dominant contribution to C is from the time taken to traverse that path (as happens on the barbell).

There are a number of variations on the theme of Theorem 1, and we will give two. The first (due to Zuckerman [30], whose proof we follow) provides a nice illustration of probabilistic technique.

Proposition 3 *Write C_e for the time to cover all edges of an unweighted graph, i.e. until each edge (v, w) has been traversed in each direction. Then*

$$\max_v E_v C_e \leq 11\bar{d}n^2.$$

Proof. Fix a vertex v and a time t_0. Define "excursions", starting and ending at v, as follows. In each excursion, wait until all vertices have been visited, then wait t_0 longer, then end the excursion at the next visit to v. Writing S_i for the time at which the i'th excursion ends, and N for the (random) number of excursions required to cover each edge in each direction, we have

$$S_N = \min\{S_i : S_i \geq C_e\}$$

and so by Wald's identity (yyy refs)

$$E_v C_e \leq E_v S_N = E_v N \times E_v S_1. \tag{3}$$

Clearly

$$E_v S_1 \leq E_v C + t_0 + \max_w E_w T_v \leq t_0 + 2 \max_i E_i C.$$

To estimate the other factor, we shall first show

$$P_v(N > 2) \leq m^3/t_0^2 \tag{4}$$

where $m \equiv \bar{d}n$ is the number of directed edges. Fix a directed edge (w, x), say. By Chapter 3 Lemma yyy the mean time, starting at x, until (w, x) is traversed equals m. So the chance, starting at x, that (w, x) is not traversed before time t_0 is at most m/t_0. So using the definition of excursion, the chance that (v, w) is not traversed during the first excursion is at most m/t_0, so the chance it is not traversed during the first *two* excursions is at most $(m/t_0)^2$. Since there are m directed edges, (4) follows.

Repeating the argument for (4) gives

$$P_v(N > 2j) \leq \left(\frac{m^3}{t_0^2}\right)^j ; \; j \geq 0$$

and hence, assuming $m^3 < t_0^2$,

$$E_v N \leq \frac{2}{1 - m^3/t_0^2}.$$

Putting $t_0 = \lceil 2m^{3/2} \rceil$ gives $E_v N \leq 8/3$. Substituting into (3),

$$\max_v E_v C_e \leq \frac{8}{3}(\lceil 2m^{3/2} \rceil + 2 \max_v E_v C).$$

Now Theorem 1 says $\max_v E_v C \leq m(n-1) \leq mn - 1$, so

$$\begin{aligned}
\max_v E_v C_e &\leq \frac{8}{3}(2m^{3/2} + 2mn) \\
&= \frac{16}{3} m(m^{1/2} + n) \\
&\leq \frac{32}{3} mn
\end{aligned}$$

establishing the Proposition. $\square$

Another variant of Theorem 1, due to Kahn et al [24] (whose proof we follow), uses a graph-theoretical lemma to produce a "good" spanning tree in graphs of high degree.

Theorem 4 *Writing* $d_* = \min_v d_v$,

$$\max_v E_v C^+ \leq \frac{6\bar{d}n^2}{d_*} \tag{5}$$

and so on a regular graph

$$\max_v E_v C^+ \leq 6n^2. \tag{6}$$

To appreciate (6), consider

Example 5 Take an even number $j \geq 2$ cliques of size $d \geq 3$, distinguish two vertices v_i, v_i' in the i'th clique (for each $0 \leq i < j$), remove the edges (v_i, v_i') and add the edges $(v_i', v_{(i+1) \bmod j})$. This creates a $(d-1)$-regular graph with $n = jd$ vertices.

Arguing as in the barbell example (Chapter 5 yyy), as $d \to \infty$ with j varying arbitrarily,

$$\max_{v,w} E_v T_w \sim \frac{d}{2} \times d \times \frac{j^2}{4} \sim \frac{n^2}{8}.$$

Thus the $O(n^2)$ bound in (6) can't be improved, even as a bound for the smaller quantity $\max_{v,w} E_v T_w$. (Note that in the example, $d/n \leq 1/2$. From the results in Chapter 5 Example yyy and Matthews' method one gets $EC = O(n \log n)$ for regular graphs with d/n bounded above $1/2$.)

Here is the graph-theory lemma needed for the proof of Theorem 4.

Lemma 6 *Let* G *be an* n*-vertex graph with minimal degree* d_*. *There exists a family of* $\lceil d_*/2 \rceil$ *spanning forests* F_i *such that*
(i) Each edge of G *appears in at most* 2 *forests*
(ii) Each component of each forest has size at least $\lceil d_*/2 \rceil$.

Proof. Replace each edge (i,j) of the graph by two directed edges $(i \to j)$, $(j \to i)$. Pick an arbitrary v_1 and construct a path $v_1 \to v_2 \to \ldots v_q$ on distinct vertices, stopping when the path cannot be extended. That is the first stage of the construction of F_1. For the second stage, pick a vertex v_{q+1} not used in the first stage and construct a path $v_{q+1} \to v_{q+2} \to \ldots v_r$ in which no second-stage vertex is revisited, stopping when a first-stage vertex is hit or when the path cannot be extended. Continue stages until all vertices have been touched. This creates a directed spanning forest F_1. Note that all the neighbors of v_q must be amongst $\{v_1, \ldots, v_{q-1}\}$, and so the size of the component of F_1 containing v_1 is at least $d_* + 1$, and similarly for the other components of F_1.

Now delete from the graph all the directed edges used in F_1. Inductively construct forests $F_2, F_3, \ldots, F_{\lceil d_*/2 \rceil}$ in the same way. The same argument shows that each component of F_i has size at least $d_* + 2 - i$, because at a "stopping" vertex v at most $i-1$ of the directed edges out of v were used in previous forests.

Proof of Theorem 4. Write m for the number of (undirected) edges. For an edge $e = (v,x)$ write $b_e = E_v T_x + E_x T_v$. Chapter 3 Lemma yyy says $\sum_e b_e = 2m(n-1)$. Now consider the $\lceil d_*/2 \rceil$ forests F_i given by Lemma 6. Since each edge appears in at most two forests,

$$\sum_i \sum_{e \in F_i} b_e \leq 2 \sum_e b_e \leq 4mn,$$

and so there exists a forest F with $\sum_{e \in F} b_e \leq 4mn/\lceil d_*/2 \rceil \leq 8mn/d_*$. But each component of F has size at least $\lceil d_*/2 \rceil$, so F has at most $2n/d_*$ components. So to extend F to a tree $\mathcal{T}$ requires adding at most $2n/d_* - 1$ edges (e_j), and for each edge e we have $b_e \leq 2m$ by (2). This creates a spanning tree $\mathcal{T}$ with $\sum_{e \in \mathcal{T}} b_e \leq 12mn/d_*$. As in the proof of Theorem 1, this is an upper bound for $E_v C^+$.

2 Simple examples of cover times

There are a few (and only a few) examples where one can study EC by bare-hands exact calculations. Write h_n for the *harmonic sum*

$$h_n = \sum_{i=1}^{n} i^{-1} \sim \log n. \qquad (7)$$

(a) *The coupon collector's problem.* Many textbooks discuss this classical problem, which involves C for the chain $(X_t; t \geq 0)$ whose values are independent and uniform on an n-element set, i.e. random walk on the complete graph with self-loops. Write (cf. the proof of Matthews' method, Chapter 2 yyy) C^m for the first time at which m distinct vertices have been visited. Then each step following time C^m has chance $(n-m)/n$ to hit a new vertex, so $E(C^{m+1}-C^m) = n/(n-m)$, and so

$$EC = \sum_{m=1}^{n-1} E(C^{m+1}-C^m) = nh_{n-1}. \tag{8}$$

(By symmetry, E_vC is the same for each initial vertex, so we just write EC) It is also a textbook exercise (e.g. [17] p. 124) to obtain the limit distribution

$$n^{-1}(C - n\log n) \ \stackrel{d}{\to}\ \xi \tag{9}$$

where ξ has the extreme value distribution

$$P(\xi \le x) = \exp(-e^{-x}), \quad -\infty < x < \infty. \tag{10}$$

We won't go into the elementary derivations of results like (9) here, because in Chapter 7 yyy we give more general results.

(b) *The complete graph.* The analysis of C for random walk on the complete graph (i.e. without self-loops) is just a trivial variation of the analysis above. Each step following time C^m has chance $(n-m)/(n-1)$ to hit a new vertex, so

$$EC = (n-1)h_n \sim n\log n. \tag{11}$$

And the distribution limit (9) still holds. Because $E_vT_w = n-1$ for $w \ne v$, we also have

$$EC^+ = EC + (n-1) = (n-1)(1+h_{n-1}) \sim n\log n. \tag{12}$$

(c) *The n-star (Chapter 5 Example yyy).* Here the visits to the leaves (every second step) are exactly i.i.d., so we can directly apply the coupon collector's problem. For instance, writing v for the central vertex and l for a leaf,

$$\begin{aligned} E_lC &= 2(n-1)h_{n-2} \sim 2n\log n \\ E_vC^+ &= 1 + E_lC + 1 = 2(n-1)h_{n-1} \sim 2n\log n \end{aligned}$$

and $C/2$ satisfies (9). Though we won't give the details, it turns out that a clever inductive argument shows these are the minima over all trees.

Proposition 7 (Brightwell - Winkler [7]) *On an n-vertex tree,*

$$\min_v E_v C \geq 2(n-1)h_{n-2}$$

$$\min_v E_v C^+ \geq 2(n-1)h_{n-1}.$$

(d) *The n-cycle.* Random walk on the n-cycle is also easy to study. At time C^m the walk has visited m distinct vertices, and the set of visited vertices must form an interval $[j, j+m-1]$, say, where we add modulo n. At time C^m the walk is at one of the endpoints of that interval, and $C^{m+1} - C^m$ is the time until the first of $\{j-1, j+m\}$ is visited, which by Chapter 5 yyy has expectation $1 \times m$. So

$$EC = \sum_{m=1}^{n-1} E(C^{m+1} - C^m) = \sum_{i=1}^{n-1} i = \frac{1}{2}n(n-1).$$

There is also an expression for the limit distribution (see Notes).

The n-cycle also has an unexpected property. Let V denote the last vertex to be hit. Then

$$\begin{aligned} P_0(V = v) &= P_0(T_{v-1} < T_{v+1})P_{v-1}(T_{v+1} < T_v) \\ &\quad + P_0(T_{v+1} < T_{v-1})P_{v+1}(T_{v-1} < T_v) \\ &= \frac{n-(v+1)}{n-2}\frac{1}{n-1} + \frac{v-1}{n-2}\frac{1}{n-1} \\ &= \frac{1}{n-1}. \end{aligned}$$

In other words, the n-cycle has the property

> For any initial vertex v_0, the last-visited vertex V is uniform on the states excluding v_0.

Obviously the complete graph has the same property, by symmetry. Lovasz and Winkler [25] gave a short but ingenious proof that these are the *only* graphs with that property, a result rediscovered in [22].

3 More upper bounds

We remain in the setting of random walk on an unweighted graph. Theorems 1 and 4 show that the mean cover times, and hence mean hitting times, are $O(n^3)$ on irregular graphs and $O(n^2)$ on regular graphs, and examples such as the barbell and the n-cycle show these bounds are the right order of magnitude. Quite a lot of attention has been paid to sharpening the constants in such bounds. We will not go into details, but will merely record a very simple argument in section 3.1 and the best known results in section 3.2.

3.1 Simple upper bounds for mean hitting times

Obviously $\max_j(E_iT_j + E_jT_i) \leq E_iC^+$, so maximizing over i gives

$$\tau^* \leq \max_i E_iC^+ \tag{13}$$

and the results of section 1 imply upper bounds on τ^*. But implicit in earlier results is a direct bound on τ^*. The edge-commute inequality implies that, for arbitrary v, x at distance $\Delta(v, x)$,

$$E_vT_x + E_xT_v \leq \bar{d}n\Delta(v, x) \tag{14}$$

and hence

Corollary 8 *$\tau^* \leq \bar{d}n\Delta$, where Δ is the diameter of the graph.*

It is interesting to compare the implications of Corollary 8 with what can be deduced from (13) and the results of section 1. To bound Δ in terms of n alone, we have $\Delta \leq n - 1$, and then Corollary 8 gives the same bound $\tau^* \leq \bar{d}n(n-1)$ as follows from Theorem 1. On the other hand, the very simple graph-theoretic Lemma 10 gives (with Corollary 8) the following bound, which removes a factor of 2 from the bound implied by Theorem 4.

Corollary 9 *$\tau^* \leq \frac{3\bar{d}n^2}{d_*}$ and so on a regular graph $\tau^* \leq 3n^2$.*

Lemma 10 *$\Delta \leq 3n/d_*$.*

Proof. Consider a path $v_0, v_1, \ldots, v_\Delta$, where vertices v_0 and v_Δ are distance Δ apart. Write A_i for the set of neighbors of v_i. Then A_i and A_j must be disjoint when $|j - i| \geq 3$. So a given vertex can be in at most 3 of the A's, giving the final inequality of

$$(\Delta + 1)d_* \leq \sum_{i=0}^{\Delta} d_{v_i} = \sum_{i=0}^{\Delta} |A_i| \leq 3n. \qquad \square$$

3.2 Known and conjectured upper bounds

Here we record results without giving proofs. Write **max** for the maximum over n-vertex graphs. The next result is the only case where the exact extremal graph is known.

Theorem 11 (Brightwell-Winkler [8]) **max** $\max_{v,x} E_v T_x$ *is attained by the lollipop (Chapter 5 Example yyy) with* $m_1 = \lfloor (2n+1)/3 \rfloor$, *taking* x *to be the leaf.*

Note that the implied asymptotic behavior is

$$\mathbf{max}\ \max_{v,w} E_v T_w \sim \frac{4}{27} n^3. \tag{15}$$

Further asymptotic results are given by

Theorem 12 (Feige [20, 18])

$$\mathbf{max} \max_v E_v C^+ \sim \frac{4}{27} n^3 \tag{16}$$

$$\mathbf{max} \min_v E_v C^+ \sim \frac{3}{27} n^3 \tag{17}$$

$$\mathbf{max} \min_v E_v C \sim \frac{2}{27} n^3 \tag{18}$$

The value in (16) is asymptotically attained on the lollipop, as in Theorem 11. Note that (15) and (16) imply the same $4n^3/27$ behavior for intermediate quantities such as τ^* and $\max_v E_v C$. The values in (17) and (18) are asymptotically attained by the graph consisting of a $n/3$-path with a $2n/3$-clique attached at the middle of the path.

The corresponding results for τ_0 and τ_2 are not known. We have $\tau_2 \leq \tau_0 \leq \min_v E_v C$, the latter inequality from the random target lemma, and so (18) implies

$$\mathbf{max}\tau_0 \text{ and } \mathbf{max}\tau_2 \ \leq (\frac{2}{27} + o(1)) n^3. \tag{19}$$

But a natural guess is that the asymptotic behavior is that of the barbell, giving the values below.

Open Problem 13 Prove the conjectures

$$\mathbf{max}\ \tau_0 \sim \frac{1}{54} n^3, \quad \mathbf{max}\ \tau_2 \sim \frac{1}{54} n^3.$$

For regular graphs, none of the asymptotic values are known exactly. A natural candidate for extremality is the necklace graph (Chapter 5 yyy), where the time parameters are asymptotically 3/4 times the parameters for the n-path. So the next conjecture uses the numerical values from the necklace graph.

Open Problem 14 Prove the conjectures that, over the class of regular n-vertex graphs

$$\mathbf{max} \max_{i,j} E_i T_j \sim \frac{3}{4} n^2$$

$$\mathbf{max}\ \tau^* \sim \frac{3}{2} n^2$$

$$\mathbf{max} \max_{v} E_v C^+ \sim \frac{3}{2} n^2$$

$$\mathbf{max} \max_{v} E_v C \sim \frac{15}{16} n^2$$

$$\mathbf{max} \min_{v} E_v C \sim \frac{3}{4} n^2$$

$$\mathbf{max}\ \tau_0 \sim \frac{1}{4} n^2$$

$$\mathbf{max}\ \tau_2 \sim \frac{3}{2\pi^2} n^2$$

The best bounds known are those implied by the following result.

Theorem 15 (Feige [18]) *On a d-regular graph,*

$$\begin{aligned} \max_v E_v C &\le 2n^2 \\ \max_v E_v C_v^+ &\le 2n^2 \left(1 + \tfrac{d-2}{(d+1)^2}\right) \le 13n^2/6. \end{aligned}$$

4 Short-time bounds

It turns out that the bound "$\tau^* \le 3n^2$ on a regular graph" given by Corollary 9 can be used to obtain bounds concerning the short-time behavior of random walks. Such bounds, and their applications, are the focus of this section. We haven't attempted to optimize numerical constants (e.g. Theorem 15 implies that $\tau^* \le 13n^2/6$ on regular graphs). More elaborate arguments (see Notes) can be used to improve constants and to deal with the irregular case, but we'll restrict attention to the regular case for simplicity.

Write $N_i(t)$ for the number of visits to i before time t, i.e. during $[0, t-1]$.

Proposition 16 *Consider random walk on an n-vertex regular graph G. Let A be a proper subset of vertices and let $i \in A$.*
(i) $E_i T_{A^c} \leq 4|A|^2$.
(ii) $E_i N_i(T_{A^c}) \leq 5|A|$.
(iii) $E_i N_i(t) \leq 5t^{1/2},\ 0 \leq t < 5n^2$.
(iv) $P_\pi(T_i < t) \geq \frac{1}{5n} \min(t^{1/2}, n)$.

Remarks. For part (i) we give a slightly fussy argument repeating ingredients of the proof of Corollary 9, since these are needed for (ii). The point of (iv) is to get a bound for $t \ll E_\pi T_i$. On the n-cycle, it can be shown that the probability in question really is $\Theta(\min(t^{1/2}/n, 1))$, uniformly in n and t.

Proof of Proposition 16. Choose a vertex $b \in A^c$ at minimum distance from i, and let $i = i_0, i_1, \ldots, i_j, i_{j+1} = b$ be a minimum-length path. Let G^* be the subgraph on vertex-set A, and let G^{**} be the subgraph on vertex-set A together with all the neighbors of i_j. Write superscripts * and ** for the random walks on G^* and G^{**}. Then

$$E_i T_{A^c} \leq E_i T^{**}_{A^c} = E_i T^*_{i_j} + E_{i_j} T^{**}_{A^c}$$

The inequality holds because we can specify the walk on G in terms of the walk on G^{**} with possibly extra chances of jumping to A^c at each step (this is a routine *stochastic comparison* argument, written out as an example in Chapter 14 yyy). The equality holds because the only routes in G^{**} from i to A^c are via i_j, by the minimum-length assumption. Now write $\mathcal{E}, \mathcal{E}^*, \mathcal{E}^{**}$ for the edge-sets. Using the commute interpretation of resistance,

$$E_i T^*_{i_j} \leq 2|\mathcal{E}^*| j. \tag{20}$$

Writing $q \geq 1$ for the number of neighbors of i_j in A^c, the effective resistance in G^{**} between i_j and A^c is $1/q$, so the commute interpretation of resistance give the first equality in

$$E_{i_j} T^{**}_{A^c} = 2|\mathcal{E}^{**}|\frac{1}{q} - 1 = 2\frac{|\mathcal{E}^*|}{q} + 1 \leq 2|\mathcal{E}^*| + 1 \leq |A|^2.$$

The neighbors of $i_0, i_1, \ldots, i_{j-1}$ are all in A, so the proof of Lemma 10 implies

$$j \leq 3|A|/d \tag{21}$$

where d is the degree of G. Since $2|\mathcal{E}^*| \leq d|A|$, the bound in (20) is at most $3|A|^2$, and part (i) follows.

For part (ii), by the electrical network analogy (Chapter 3 yyy) the quantity in question equals

$$\frac{1}{P_i(T_{A^c} < T_i^+)} = w_i r(i, A^c) = d r(i, A^c) \tag{22}$$

where $r(i, A^c)$ is the effective resistance in G between i and A^c. Clearly this effective resistance is at most the distance ($j+1$, in the argument above) from i to A^c, which by (21) is at most $3|A|/d\ +1$. Thus the quantity (22) is at most $3|A| + d$, establishing the desired result in the case $d \leq 2|A|$. If $d > 2|A|$ then there are at least $d - |A|$ edges from i to A^c, so $r(i, A^c) \leq \frac{1}{d-|A|}$ and the quantity (22) is at most $\frac{d}{d-|A|} \leq 2 \leq 5|A|$.

For part (iii), fix a state i and an integer time t. Write $N_i(t)$ for the number of visits to i before time t, i.e. during times $\{0, 1, \ldots, t-1\}$. Then

$$\frac{t}{n} = E_\pi N_i(t) \leq P_\pi(T_i < t) E_i N_i(t) \tag{23}$$

the inequality by conditioning on T_i. Now choose real s such that $ns \geq t$. Since $\sum_j E_i N_j(t) = t$, the set

$$A \equiv \{j : E_i N_j(t) > s\}$$

has $|A| < t/s \leq n$, so part (ii) implies

$$E_i N_i(T_{A^c}) \leq 5t/s. \tag{24}$$

Now by regularity we can rewrite A as $\{j : E_j N_i(t) > s\}$, and so by conditioning on T_{A^c}

$$E_i N_i(t) \leq E_i N_i(T_{A^c}) + s.$$

Setting $s = \sqrt{5t}$ and combining with (24) gives (iii). The bound in (iv) now follows from (iii) and (23).

4.1 Covering by multiple walks

The first application is a variant of work of Broder et al [10] discussed further in section 8.2.

Proposition 17 *On a regular n-vertex graph, consider K independent random walks, each started at a uniform random vertex. Let $C^{[K]}$ be the time until every vertex has been hit by some walk. Then*

$$EC^{[K]} \leq \frac{(25 + o(1)) n^2 \log^2 n}{K^2} \quad \textit{as } n \to \infty \textit{ with } K \geq 6 \log n.$$

Remarks. The point is the $\frac{1}{K^2}$ dependence on K. On the n-cycle, for $K \sim \varepsilon n$ it can be shown that initially the largest gap between adjacent walkers is $\Theta(\log n)$ and that $EC^{[K]} = \Theta(\log^2 n)$, so in this respect the bound is sharp. Of course, for $K \leq \log n$ the bound would be no improvement over Theorem 4.

Proof. As usual write T_i for the hitting time on i for a single walk, and write $T_i^{[K]}$ for the first time i is visited by some walk. Then

$$\begin{aligned} P_\pi(T_i^{[K]} \geq t) &= (P_\pi(T_i \geq t))^K \\ &= (1 - P_\pi(T_i < t))^K \\ &\leq \exp(-K P_\pi(T_i < t)) \\ &\leq \exp\left(-\frac{Kt^{1/2}}{5n}\right) \end{aligned}$$

by Proposition 16 (iii), provided $t \leq n^2$. So

$$P(C^{[K]} \geq t) \leq \sum_i P(T_i^{[K]} \geq t) \leq n \exp\left(-\frac{Kt^{1/2}}{5n}\right), \ t \leq n^2. \tag{25}$$

The bound becomes 1 for $t_0 = \frac{25n^2}{K^2} \log^2 n$. So

$$\begin{aligned} EC^{[K]} &= \sum_{t=1}^{\infty} P(C^{[K]} \geq t) \\ &\leq \lceil t_0 \rceil + \sum_{t=\lceil t_0 \rceil + 1}^{n^2-1} n \exp\left(-\frac{Kt^{1/2}}{5n}\right) + \sum_{t=n^2}^{\infty} P(C^{[K]} \geq t) \\ &= \lceil t_0 \rceil + S_1 + S_2, \text{ say,} \end{aligned}$$

and the issue is to show that S_1 and S_2 are $o(t_0)$. To handle S_2, split the set of K walks into subsets of sizes $K-1$ and 1. By independence, for $t \geq n^2$ we have $P(C^{[K]} \geq t) \leq P(C^{[K-1]} \geq n^2) P(C^{[1]} \geq t)$. Then

$$\begin{aligned} S_2 &\leq P(C^{[K-1]} \geq n^2) EC^{[1]} \text{ by summing over } t \\ &\leq n \exp(-(K-1)/5) \cdot 6n^2 \text{ by (25) and Theorem 4} \\ &= o(t_0) \text{ using the hypothesis } K \geq 6 \log n. \end{aligned}$$

To bound S_1 we start with a calculus exercise: for $u > 1$

$$\int_{u^2}^{\infty} \exp(-x^{1/2}) \, dx = \int_u^{\infty} 2y \ \exp(-y) \, dy \text{ by putting } x = y^2$$

$$
\begin{aligned}
&\leq 2e^{-1}u \int_u^\infty \exp(-\left(\frac{u-1}{u}\right) y)\ dy \text{ , using } \frac{y}{u} \leq \exp(\frac{y}{u} - 1) \\
&= \frac{2u^2 \exp(-u)}{u-1}.
\end{aligned}
$$

The sum S_1 is bounded by the corresponding integral over $[t_0, \infty)$ and the obvious calculation, whose details we omit, bounds this integral by $2t_0/(\log n - 1)$.

4.2 Bounding point probabilities

Our second application is to universal bounds on point probabilities. A quite different universal bound will be given in Chapter yyy.

Proposition 18 *For continuous-time random walk on a regular n-vertex graph,*

$$
\begin{aligned}
P_i(X_t = j) &\leq 5t^{-1/2},\ t \leq n^2 \\
&\leq \frac{1}{n} + \frac{K_1}{n} \exp\left(\frac{-t}{K_2 n^2}\right),\ t \geq n^2
\end{aligned}
$$

where K_1 and K_2 are absolute constants.

In discrete time one can get essentially the same result, but with the bounds multiplied by 2, though we shall not give details (see Notes).

Proof. $P_i(X_t = i)$ is decreasing in t, so

$$
P_i(X_t = i) \leq t^{-1} \int_0^t P_i(X_s = i) ds = t^{-1} E_i N_i(t) \leq 5t^{-1/2}
$$

where the last inequality is Proposition 16 (iii), whose proof is unchanged in continuous time, and which holds for $t \leq n^2$. This gives the first inequality when $i = j$, and the general case follows from Chapter 3 yyy.

For the second inequality, recall the definition of *separation* $s(t)$ from Chapter 4 yyy. Given a vertex i and a time t, there exists a probability distribution θ such that

$$
P_i(X_t \in \cdot) = (1 - s(t))\pi + s(t)\theta.
$$

Then for $u \geq 0$,

$$
P_i(X_{t+u} = j) - \frac{1}{n} = s(t) \left(P_\theta(X_u = j) - \frac{1}{n}\right).
$$

Thus, defining $q(t) = \max_{i,j}\left(P_i(X_t = j) - \frac{1}{n}\right)$, we have proved

$$q(t+u) \leq s(t)q(u); \ t, u \geq 0. \tag{26}$$

Now $q(n^2) \leq 4/n$ by the first inequality of the Proposition, and $s(\tau_1^{(1)}) = e^{-1}$ by definition of $\tau_1^{(1)}$ in Chapter 4 yyy, so by iterating (26) we have

$$q(n^2 + m\tau_1^{(1)}) \leq \frac{4}{n}\, e^{-m}, \ m \geq 1. \tag{27}$$

But by Chapter 4 yyy we have $\tau_1^{(1)} \leq K\tau^*$ for an absolute constant K, and then by Corollary 9 we have $\tau_1^{(1)} \leq 3Kn^2$. The desired inequality now follows from (27).

4.3 A cat and mouse game

Here we reconsider the cat and mouse game discussed in Chapter 4 section yyy. Recall that the cat performs continuous-time random walk on a n-vertex graph, and the mouse moves according to some arbitrary deterministic strategy. Let M be the first meeting time, and let m^* be the maximum of EM over all pairs of initial vertices and all strategies for the mouse.

Proposition 19 *On a regular graph, $m^* \leq KN^2$ for some absolute constant K.*

Proof. The proof relies on Proposition 18, whose conclusion implies there exists a constant K such that

$$p^*(t) \equiv \max_{x,v} p_{vx}(t) \leq \frac{1}{n} + Kt^{-1/2}; \ 0 \leq t < \infty.$$

Consider running the process forever. The point is that, regardless of the initial positions, the chance that the cat and mouse are "together" (i.e. at the same vertex) at time u is at most $p^*(u)$. So in the case where the cat starts with the (uniform) stationary distribution,

$$\begin{aligned}
P(\text{ together at time } s\ ) &= \int_0^s f(u) P(\text{ together at time } s | M = u)\ du \\
&\quad \text{where } f \text{ is the density function of } M \\
&\leq \int_0^s f(u) p^*(s-u) du \\
&\leq \frac{1}{n} P(M \leq s)\ + K \int_0^s f(u)(s-u)^{-1/2} du.
\end{aligned}$$

So

$$
\begin{aligned}
\frac{t}{n} &= \int_0^t P(\text{ together at time } s \text{ })\, ds \text{ by stationarity} \\
&\leq \frac{1}{n}\int_0^t P(M \leq s)ds \ + K\int_0^t f(u)du \int_u^t (s-u)^{-1/2}ds \\
&= \frac{t}{n} - \frac{1}{n}\int_0^t P(M > s)ds + 2K\int_0^t f(u)(t-u)^{1/2}du \\
&\leq \frac{t}{n} - \frac{1}{n}E\min(M,t) + 2Kt^{1/2}.
\end{aligned}
$$

Rearranging, $E\min(M,t) \leq 2Knt^{1/2}$. Writing $t_0 = (4Kn)^2$, Markov's inequality gives $P(M \leq t_0) \geq 1/2$. This inequality assumes the cat starts with the stationary distribution. When it starts at some arbitrary vertex, we may use the definition of separation $s(u)$ (recall Chapter 4 yyy) to see $P(M \leq u + t_0) \geq (1 - s(u))/2$. Then by iteration, $EM \leq \frac{2(u+t_0)}{1-s(u)}$. So appealing to the definition of $\tau_1^{(1)}$,

$$m^* \leq \frac{2}{1-e^{-1}}(t_0 + \tau_1^{(1)}).$$

But results from Chapter 4 and this chapter show $\tau_1^{(1)} = O(\tau^*) = O(n^2)$, establishing the Proposition.

5 Hitting time bounds and connectivity

The results so far in this chapter may be misleading in that upper bounds accommodating extremal graphs are rather uninformative for "typical" graphs. For a family of n-vertex graphs with $n \to \infty$, consider the property

$$\tau^* = O(n). \tag{28}$$

(in this order-of-magnitude discussion, τ^* is equivalent to $\max_{v,x} E_v T_x$). Recalling from Chapter 3 yyy that $\tau^* \geq 2(n-1)$, we see that (28) is equivalent to $\tau^* = \Theta(n)$. By Matthews' method (repeated as Theorem 26 below), (28) implies $EC = O(n \log n)$, and then by Theorem 31 we have $EC = \Theta(n \log n)$. Thus understanding when (28) holds is fundamental to understanding order-of-magnitude questions about cover times. But surprisingly, this question has not been studied very carefully. An instructive example in the d-dimensional torus (Chapter 5 Example yyy), where (28) holds iff $d \geq 3$. This example, and other examples of vertex-transitive graphs satisfying (28) discussed in Chapter 8, suggest that (28) is frequently true. More concretely, the torus example suggests that the following condition ("the isoperimetric property in $2+\varepsilon$ dimensions") may be sufficient.

Open Problem 20 *Show that for real $1/2 < \gamma < 1$ and $\delta > 0$, there exists a constant $K_{\gamma,\delta}$ with the following property. Let G be a regular n-vertex graph such that, for any subset A of vertices with $|A| \leq n/2$, there exist at least $\delta|A|^\gamma$ edges between A and A^c. Then $\tau^* \leq K_{\gamma,\delta}\ n$.*

The $\gamma = 1$ case is implicit in results from previous chapters. Chapter 3 yyy gave the bound $\max_{i,j} E_i T_j \leq 2 \max_j E_\pi T_j$, and Chapter 3 yyy gave the bound $E_\pi T_j \leq \tau_2/\pi_j$. This gives the first assertion below, and the second follows from Cheeger's inequality.

Corollary 21 *On a regular graph,*

$$\max_{v,x} E_v T_x \leq 2n\tau_2 \leq 16 n \tau_c^2.$$

Thus the "expander" property that $\tau_2 = O(1)$, or equivalently that $\tau_c = O(1)$, is sufficient for (28), and the latter is the $\gamma = 1$ case of Open Problem 20.

5.1 Edge-connectivity

At the other end of the spectrum from expanders, we can consider graphs satisfying only a little more than connectivity.

xxx more details in proofs – see Fill's comments.

Recall that a graph is *r-edge-connected* if for each proper subset A of vertices there are at least r edges linking A with A^c. By a variant of Menger's theorem (e.g. [13] Theorem 5.11), for each pair (a,b) of vertices in such a graph, there exist r paths $(a = v_0^i, v_1^i, v_2^i, \ldots, v_{m_i}^i = b)$, $i = 1, \ldots, r$ for which the edges (v_j^i, v_{j+1}^i) are all distinct.

Proposition 22 *For a r-edge-connected graph,*

$$\tau^* \leq \frac{\bar{d} n^2 \psi(r)}{r^2}$$

where ψ is defined by

$$\psi\left(\frac{i(i+1)}{2}\right) = i$$

$$\psi(\cdot) \textit{ is linear on } \left[\frac{i(i+1)}{2}, \frac{(i+1)(i+2)}{2}\right].$$

Note $\psi(r) \sim \sqrt{2r}$. So for a d-regular, d-edge-connected graph, the bound becomes $\sim 2^{1/2}d^{-1/2}n^2$ for large d, improving on the bound from Corollary 9. Also, the Proposition improves on the bound implied by Chapter 4 yyy in this setting.

Proof. Given vertices a, b, construct a unit flow from a to b by putting flow $1/r$ along each of the r paths $(a = v_0^i, v_1^i, v_2^i, \ldots, v_{m_i}^i = b)$. By Chapter 3 Theorem yyy

$$E_aT_b + E_bT_a \le \bar{d}n(1/r)^2M$$

where $M = \sum_i m_i$ is the total number of edges in the r paths. So the issue is bounding M. Consider the digraph of all edges (v_j^i, v_{j+1}^i). If this digraph contained a directed cycle, we could eliminate the edges on that cycle, and still create r paths from a to b using the remaining edges. So we may assume the digraph is acyclic, which implies we can label the vertices as $a = 1, 2, 3, \ldots, n = b$ in such a way that each edge (j, k) has $k > j$. So the desired result follows from

Lemma 23 *In a digraph on vertices $\{1, 2, \ldots, n\}$ consisting of r paths $1 = v_0^i < v_1^i < v_2^i < \ldots v_{m_i}^i = n$ and where all edges are distinct, the total number of edges is at most $n\psi(r)$.*

Proof.

xxx give proof and picture.

Example 24 Take vertices $\{0, 1, \ldots, n-1\}$ and edges $(i, i+u \bmod n)$ for all i and all $1 \le u \le \kappa$.

This example highlights the "slack" in Proposition 22. Regard κ as large and fixed, and $n \to \infty$. Random walk on this graph is classical random walk (i.e. sums of independent steps) on the n-cycle, where the steps have variance $\sigma^2 = \frac{1}{\kappa}\sum_{u=1}^{\kappa} u^2$, and it is easy to see

$$\tau^* = 2E_0T_{\lfloor n/2 \rfloor} \sim \frac{(n/2)^2}{\sigma^2} = \Theta(n^2/\kappa^2).$$

This is the bound Proposition 22 would give if the graph were $\Theta(\kappa^2)$-edge-connected. And for a "typical" subset A such as an interval of length greater than κ there are indeed $\Omega(\kappa^2)$ edges crossing the boundary of A. But by considering a singleton A we see that the graph is really only 2κ-edge-connected, and Proposition 22 gives only the weaker $O(n^2/\kappa^{1/2})$ bound.

xxx tie up with similar discussion of τ_2 and connectivity being affected by small sets; better than bound using τ_c only.

5.2 Equivalence of mean cover time parameters

Returning to the order-of-magnitude discussion at the start of section 5, let us record the simple equivalence result. Recall (cf. Chapter 4 yyy) we call parameters *equivalent* if their ratios are bounded by absolute constants.

Lemma 25 *The parameters* $\max_i E_iC^+$, $E_\pi C^+$, $\min_i E_iC^+$, $\max_i E_iC$ *and* $E_\pi C$ *are equivalent for reversible chains, but* $\min_i E_iC$ *is not equivalent to these.*

Proof. Of the five parameters asserted to be equivalent, it is clear that $\max_i E_iC$ is the largest, and that either $\min_i E_iC^+$ or $E_\pi C$ is the smallest, so it suffices to prove

$$\max_i E_iC^+ \leq 4E_\pi C \tag{29}$$

$$\max_j E_jC^+ \leq 3 \min_i E_iC^+. \tag{30}$$

Inequality (30) holds by concatenating three "cover-and-return" cycles starting at i and considering the first hitting time on j in the first and third cycles. In more detail, write

$$\Gamma(s) = \min\{u > s : (X_t : s \leq t \leq u) \text{ covers all states}\}.$$

For the chain started at i write $C^{++} = \Gamma(C^+)$ and $C^{+++} = \Gamma(C^{++})$. Since $T_j < C^+$ we have $\Gamma(T_j) \leq C^{++}$. So the chain started at time T_j has covered all states and returned to j by time C^{+++}, implying $E_jC^+ \leq EC^{+++} = 3E_iC^+$. For inequality (29), recall the *random target lemma*: the mean time to hit a π-random state V equals τ_0, regardless of the initial distribution. The inequality

$$E_iC^+ \leq \tau_0 + E_\pi C + \tau_0 + E_\pi T_i$$

follows from the four-step construction:
(i) Start the chain at i and run until hitting a π-random vertex V at time T_V;
(ii) continue until time $\Gamma(T_V)$;
(iii) continue until hitting an independent π-random vertex V';
(iv) continue until hitting i.
But $E_\pi T_i \leq E_\pi C$, and then by the random target lemma $\tau_0 \leq E_\pi C$, so (29) follows.

For the final assertion, on the lollipop graph (Chapter 5 Example yyy) one has $\min_i E_i C = \Theta(n^2)$ while the other quantities are $\Theta(n^3)$. One can also give examples on regular graphs (see Notes).

6 Lower bounds

6.1 Matthews' method

We restate Matthews' method (Chapter 2 yyy) as follows. The upper bound is widely useful: we have already used it several times in this chapter, and will use it several more times in the sequel.

Theorem 26 *For a general Markov chain,*

$$\max_v E_v C \le h_{n-1} \max_{i,j} E_i T_j.$$

And for any subset A of states,

$$\min_v E_v C \ge h_{|A|-1} \min_{i \neq j:\ i,j \in A} E_i T_j.$$

In Chapter 2 we proved the lower bound in the case where A was the entire state space, but the result for general A follows by the same proof, taking the J's to be a uniform random ordering of the states in A. One obvious motivation for the more general formulation comes from the case of trees, where for a leaf l we have $\min_j E_l T_j = 1$, so the lower bound with A being the entire state space would be just h_{n-1}. We now illustrate use of the more general formulation.

6.2 Balanced trees

We are accustomed to finding that problems on trees are simpler than problems on general graphs, so it is a little surprising to discover that one of the graphs where studying the mean cover time is difficult is the balanced r-tree of height H (Chapter 5 Example yyy). Recall this tree has $n = (r^{H+1} - 1)/(r - 1)$ vertices, and that (by the commute interpretation of resistance)

$$E_i T_j = 2m(n-1) \text{ for leaves } (i,j) \text{ distance } 2m \text{ apart.}$$

Now clearly E_iT_j is maximized by some pair of leaves, so $\max_{i,j} E_iT_j = 2H(n-1)$. Theorem 26 gives

$$\max_v E_vC \leq 2H(n-1)h_{n-1} \sim 2Hn\log n.$$

To get a lower bound, consider the set S_m of r^{H+1-m} vertices at depth $H+1-m$, and let A_m be a set of leaves consisting of one descendant of each element of S_m. The elements of A_m are at least $2m$ apart, so applying the lower bound in Theorem 26

$$\begin{aligned}\min_v E_vC &\geq \max_m 2m(n-1)\ h_{r^{H+1-m}}\\ &\sim 2n\log r\ \max_m m(H-m)\\ &\sim \frac{1}{2}H^2 n\log r.\end{aligned}$$

It turns out that this lower bound is asymptotically off by a factor of 4, while the upper bound is asymptotically correct.

Theorem 27 ([1]) *On the balanced r-tree, as $H \to \infty$ for arbitrary starting vertex,*

$$EC \sim 2Hn\log n \sim \frac{2H^2r^{H+1}\log r}{r-1}$$

Improving the lower bound to obtain this result is not easy. The natural approach (used in [1]) is to seek a recursion for the cover time distribution $C^{(H+1)}$ in terms of $C^{(H)}$. But the appropriate recursion is rather subtle (we invite the reader to try to find it!) so we won't give the statement or analysis of the recursion here.

6.3 A resistance lower bound

Our use of the commute interpretation of resistance has so far been only to obtain *upper* bounds on commute times. One can also use "shorting" ideas to obtain lower bounds, and here is a very simple implementation of that idea.

Lemma 28 *The effective resistance between $r(v,x)$ between vertices v and x in a weighted graph satisfies*

$$\frac{1}{r(v,x)} \leq w_{v,x} + \frac{1}{\frac{1}{w_v - w_{v,x}} + \frac{1}{w_x - w_{v,x}}}.$$

In particular, on an unweighted graph

$$\begin{aligned} r(v,x) &\geq \frac{d_v + d_x - 2}{d_v d_x \; - 1} \textit{ if } (v,x) \textit{ is an edge} \\ &\geq \frac{1}{d_v} + \frac{1}{d_w} \textit{ if not} \end{aligned}$$

and on an unweighted d-regular graph

$$\begin{aligned} r(v,x) &\geq \frac{2}{d+1} \textit{ if } (v,x) \textit{ is an edge} \\ &\geq \frac{2}{d} \textit{ if not .} \end{aligned}$$

So on an unweighted d-regular n-vertex graph,

$$\begin{aligned} E_v T_x + E_x T_v &\geq \frac{2dn}{d+1} \textit{ if } (v,x) \textit{ is an edge} \\ &\geq 2n \textit{ if not .} \end{aligned}$$

Proof. We need only prove the first assertion, since the others follow by specialization and by the commute interpretation of resistance. Let A be the set of vertices which are neighbors of either v or x, but exclude v and x themselves from A. Short the vertices of A together, to form a single vertex a. In the shorted graph, the only way current can flow from v to x is directly $v \to x$ or indirectly as $v \to a \to x$. So, using $'$ to denote the shorted graph, the effective resistance $r'(v,x)$ in the shorted graph satisfies

$$\frac{1}{r'(v,x)} = w'_{v,x} + \frac{1}{\frac{1}{w'_{v,a}} + \frac{1}{w'_{x,a}}}.$$

Now $w'_{x,v} = w_{x,v}$, $w'_{v,a} = w_v - w_{v,x}$ and $w'_{x,a} = w_x - w_{v,x}$. Since shorting decreases resistance, $r'(v,x) \leq r(v,x)$, establishing the first inequality.

6.4 General lower bounds

Chapter 3 yyy shows that, over the class of random walks on n-vertex graphs or the larger class of reversible chains on n states, various mean hitting time parameters are minimized on the complete graph. So it is natural to anticipate a similar result for cover time parameters. But the next example shows that some care is required in formulating conjectures.

Example 29 Take the complete graph on n vertices, and add an edge (v, l) to a new leaf l.

Since random walk on the complete graph has mean cover time $(n-1)h_{n-1}$, random walk on the enlarged graph has

$$E_l C = 1 + (n-1)h_{n-1} + 2\mu$$

where μ is the mean number of returns to l before covering. Now after each visit to v, the walk has chance $1/n$ to visit l on the next step, and so the mean number of visits to l before visiting some other vertex of the complete graph equals $1/(n-1)$. We may therefore write μ in terms of expectations for random walk on the complete graph as

$$\begin{aligned}
\mu &= \frac{1}{n-1} E_v(\text{ number of visits to } v \text{ before } C) \\
&= \frac{1}{n-1} E_v(\text{ number of visits to } v \text{ before } C^+) \\
&= \frac{1}{n-1}\ \frac{1}{n} E_v C^+ \text{ by Chapter 2 Proposition yyy} \\
&= \frac{1+h_{n-1}}{n} \text{ by (12).}
\end{aligned}$$

This establishes an expression for $E_l C$, which (after a brief calculation) can be rewritten as

$$E_l C = n h_n - \left(1 - \frac{2}{n}\right)\left(h_n - \frac{1}{n}\right).$$

Now random walk on the complete $(n+1)$-graph has mean cover time nh_n, so $E_l C$ is smaller in our example than in the complete graph.

The example motivates the following as the natural "exact extremal conjecture".

Open Problem 30 *Prove that, for any reversible chain on n states,*

$$E_\pi C \geq (n-1)h_{n-1}$$

(the value for random walk on the complete graph).

The related asymptotic question was open for many years, and was finally proved by Feige [19].

Theorem 31 *For random walk on an unweighted n-vertex graph,*

$$\min_v E_v C \geq c_n,$$

where $c_n \sim n \log n$ as $n \to \infty$.

The proof is an intricate mixture of many of the techniques we have already described.

7 Distributional aspects

In many examples one can apply the following result to show that hitting time distributions become exponential as the size of state space increases.

Corollary 32 *Let i, j be arbitrary states in a sequence of reversible Markov chains.*
(i) If $E_\pi T_j / \tau_2 \to \infty$ then

$$P_\pi \left(\frac{T_j}{E_\pi T_j} > x \right) \to e^{-x}, \ 0 < x < \infty.$$

(ii) If $E_i T_j / \tau_1 \to \infty$ and $E_i T_j \geq (1 - o(1)) E_\pi T_j$ then $E_i T_j / E_\pi T_j \to 1$ and

$$P_i \left(\frac{T_j}{E_i T_j} > x \right) \to e^{-x}, \ 0 < x < \infty.$$

Proof. In continuous time, assertion (i) is immediate from Chapter 3 Proposition yyy. The result in discrete time now holds by continuization: if T_j is the hitting time in discrete time and T_j' in continuous time, then $E_\pi T_j' = E_\pi T_j$ and $T_j' - T_j$ is order $\sqrt{E_\pi T_j}$. For (ii) we have (cf. Chapter 4 section yyy) $T_j \leq U_i + T_j^*$ where T_j is the hitting time started at i, T_j^* is the hitting time started from stationarity, and $E_i U_i \leq \tau_1^{(2)}$. So $ET_j \leq ET^* + O(\tau_1)$, and the hypotheses of (ii) force $ET_j / ET_j^* \to 1$ and force the limit distribution of T_j / ET_j to be the same as the limit distribution of T_j^* / ET_j^*, which is the exponential distribution by (i) and the relation $\tau_2 \leq \tau_1$. □

In the complete graph example, C has mean $\sim n \log n$ and s.d. $\Theta(n)$, so that $C/EC \to 1$ in distribution, although the convergence is slow. The next result shows this "concentration" result holds whenever the mean cover time is essentially larger than the maximal mean hitting time.

Theorem 33 ([2]) *For states i in a sequence of (not necessarily reversible) Markov chains,*

$$\text{if } E_iC/\tau^* \to \infty \text{ then } P_i\left(\left|\frac{C}{E_iC} - 1\right| > \varepsilon\right) \to 0,\ \varepsilon > 0.$$

The proof is too long to reproduce.

8 Algorithmic aspects

Many of the mathematical results in this chapter arose originally from algorithmic questions, so let me briefly describe the questions and their relation to the mathematical results.

8.1 Universal traversal sequences

This was one motivation for the seminal paper [3]. Consider an n-vertex d-regular graph G, with a distinguished vertex v_0, and where for each vertex v the edges at v are labeled as $1, 2, \ldots, d$ in some way – it is not required that the labels be the same at both ends of an edge. Now consider a sequence $\mathbf{i} = (i_1, i_2, \ldots, i_L) \in \{1, \ldots, d\}^L$. The sequence defines a deterministic walk (x_i) on the vertices of G via

$x_0 = v_0$

(x_{j-1}, x_j) is the edge at x_{j-1} labeled i_j.

Say $\mathbf{i}$ is a *traversal sequence* for G if the walk $(x_i : 0 \le i \le L)$ visits every vertex of G. Say $\mathbf{i}$ is a *universal traversal sequence* if it is a traversal sequence for every graph G in the set $\mathcal{G}_{n,d}$ of edge-labeled graphs with distinguished vertices.

Proposition 34 (Aleliunas et al [3]) *There exists a universal traversal sequence of length $(6e + o(1))dn^3 \log(nd)$ as $n \to \infty$ with d varying arbitrarily.*

Proof. It is enough to show that a uniform random sequence of that length has non-zero chance to be a universal traversal sequence. But for such a random sequence, the induced walk on a fixed G is just simple random walk on the vertices of G. Writing $t_0 = \lceil 6en^2 \rceil$, Theorem 4 implies

$$P_v(C > t_0) \le \frac{E_vC}{t_0} \le \frac{6n^2}{t_0} \le e^{-1} \text{ for all initial } v$$

and so inductively (cf. Chapter 2 section yyy)

$$P_{v_0}(C > Kt_0) \le e^{-K}, \quad K \ge 1 \text{ integer} .$$

Thus by taking K sufficiently large that

$$e^{-K}|\mathcal{G}_{n,d}| < 1$$

there is non-zero chance that the induced walk on *every* G covers before time Kt_0. The crude bound $|\mathcal{G}_{n,d}| \le (nd)^{nd}$ means we may take $K = \lceil nd \log(nd) \rceil$.

8.2 Graph connectivity algorithms

Another motivation for the seminal paper [3] was the time-space tradeoff in algorithms for determining connectivity in graphs. Here is a highly informal presentation, illustrated by the two *Mathematician graphs.* The vertices are all mathematicians (living or dead). In the first graph, there is an edge between two mathematicians if they have written a joint paper; in the second, there is an edge if they have written two or more joint papers. A well known Folk Theorem asserts that the first graph has a giant component containing most famous mathematicians; a lesser known and more cynical Folk Theorem asserts that the second graph doesn't. Suppose we actually want to answer a question of that type – specifically, take two mathematicians (say, the reader and Paul Erdos) and ask if they are in the same component of the first graph. Suppose we have a database which, given a mathematician's name, will tell us information about their papers and in particular will list all their co-authors.

xxx continue story

Broder et al [10]

8.3 A computational question

Consider the question of getting a numerical value for E_iC (up to error factor $1 \pm \varepsilon$, for fixed ε) for random walk on a n-vertex graph. Using Theorem 1 it's clear we can do this by Monte Carlo simulation in $O(n^3)$ steps.

xxx technically, using s.d./mean bounded by submultiplicitivity.

Open Problem 35 Can E_iC be deterministically calculated in a polynomial (in n) number of steps?

It's clear one can compute mean *hitting* times on a n-step chain in polynomial time, but to set up the computation of E_iC as a hitting-time problem one has to incorporate the subset of already-visited states into the "current state", and thus work with hitting times for a $n \times 2^{n-1}$-state chain.

9 Notes on Chapter 6

Attributions for what I regard as the main ideas were given in the text. The literature contains a number of corollaries or variations of these ideas, some of which I've used without attribution, and many of which I haven't mentioned at all. A number of these ideas can be found in Zuckerman [29, 31], Palacios [27, 26] and the Ph.D. thesis of Sbihi [28], as well as papers cited elsewhere.

Section 1. The conference proceedings paper [3] proving Theorem 1 was not widely known, or at least its implications not realized, for some years. Several papers subsequently appeared proving results which are consequences (either obvious, or via the general relations of Chapter 4) of Theorem 1. I will spare their authors embarrassment by not listing them all here!

The spanning tree argument shows, writing b_e for the mean commute time across an edge e, that

$$\max_v E_v C^+ \leq \min_{\mathcal{T}} \sum_{e \in \mathcal{T}} b_e.$$

Coppersmith et al [15] give a deeper study and show that the right side is bounded between γ and $10\gamma/3$, where

$$\gamma = \left(\sum_v d_v \right) \left(\sum_v \frac{1}{d_v + 1} \right).$$

The upper bound is obtained by considering a random spanning tree, cf. Chapter yyy.

Section 2. The calculations in these examples, and the uniformity property of V on the n-cycle, are essentially classical. For the cover time C_n on the n-cycle there is a non-degenerate limit distribution $n^{-2}C_n \xrightarrow{d} C$. From the viewpoint of weak convergence (Chapter yyy), C is just the cover time for Brownian motion on the circle of unit circumference, and its distribution is known as part of a large family of known distributions for maximal-like statistics of Brownian motion: Imhof [23] eq. (2.4) gives the density as

$$f_C(t) = 2^{3/2}\pi^{-1/2}t^{-3/2}\sum_{m=1}^{\infty}(-1)^{m-1}m^2\exp(-\frac{m^2}{2t}).$$

Sbihi [28] gives a direct derivation of a different representation of f_C.

Section 4. Use of Lemma 10 in the random walk context goes back at least to Flatto et al [21].

Barnes and Feige [5] give a more extensive treatment of short-time bounds in the irregular setting, and their applications to covering with multiple walks (cf. Proposition 17 and section 8.2). They also give bounds on the mean time taken to cover μ different edges or ν different vertices – their bound for the latter becomes $O(\nu^2\log\nu)$ on regular graphs.

Proposition 18 implies that on an infinite regular graph $P_i(X_t = j) \leq Kt^{-1/2}$. Carlen et al [11] Theorem 5.14 prove this as a corollary of results using more sophisticated machinery. Our argument shows the result is fairly elementary. In discrete time the analog of the first inequality can be proved using the "CM proxy" property than $P_i(X_{2t} = i) + P_i(X_{2t+1} = i)$ is decreasing, but the analog of the second inequality requires different arguments because we cannot exploit the $\tau_1^{(1)}$ inequalities.

Section 5. Variations on Corollary 21 are given in Broder and Karlin [9] and Chandra et al. [12].

Upper bounds on mean hitting times imply upper bounds on the relaxation time τ_2 via the general inequalities $\tau_2 \leq \tau_0 \leq \frac{1}{2}\tau^*$. In most concrete examples these bounds are too crude to be useful, but in "extremal" settings these bounds are essentially as good as results seeking to bound τ_2 directly. For instance, in the setting of a d-regular r-edge-connected graph, a direct bound (Chapter 4 Proposition yyy) gives

$$\tau_2 \leq \frac{d}{4r\sin^2\frac{\pi}{2n}} \sim \frac{dn^2}{\pi^2 r}.$$

Up to the numerical constant, the same bound is obtained from Proposition 22 and the general inequality $\tau_2 \leq \tau^*/2$.

xxx contrast with potential and Cheeger-like arguments ?

To sketch an example of a regular graph where $\min_i E_iC$ has a different order than $\max_i E_iC$, make a regular $m_1 + m_2$-vertex graph from a m_1-vertex graph with mean cover time $\Theta(m_1\log m_1)$ and a m_2-vertex graph (such as the necklace) with mean cover time $\Theta(m_2^2)$, for suitable values of the m's. Starting from a typical vertex of the former, the mean cover time is $\Theta(m_1 \log m_1 + m_1 m_2 + m_2^2)$ whereas starting from the unattached end of the necklace the mean cover time is $\Theta(m_1 \log m_1 + m_2^2)$. Taking $m_1 \log m_1 + m_2^2) = o(m_1 m_2)$ gives the desired example.

Section 6. The "subset" version of Matthews' lower bound (Theorem 26) and its application to trees were noted by Zuckerman [31], Sbihi [28] and others. As well as giving a lower bound for balanced trees, these authors give several lower bounds for more general trees satisfying various constraints (cf. the unconstrained result, Proposition 7). As an illustration, Devroye - Sbihi [16] show that on a tree

$$\min_v E_v C \geq \frac{(1+o(1)) n \log^2 n}{2 \log(d^* - 1)} \text{ if } d^* \equiv \max_v d_v = n^{o(1)}.$$

I believe that the recursion set-up in [1] can be used to prove Open Problem 35 on trees, but I haven't thought carefully about it.

The "shorting" lower bound, Lemma 28, was apparently first exploited by Coppersmith et al [15].

Section 7. Corollary 32 encompasses a number of exponential limit results proved in the literature by *ad hoc* calculations in particular examples.

Section 8.1. Proposition 34 is one of the neatest instances of "Erdos's Probabilistic Method in Combinatorics", though surprisingly it isn't in the recent book [4] on that subject. Constructing explicit universal traversal sequences is a hard open problem: see Borodin et al [6] for a survey.

Section 8.2. See [10] for a more careful discussion of the issues. The alert reader of our example will have noticed the subtle implication that the reader has written fewer papers than Paul Erdos, otherwise (why?) it would be preferable to do the random walk in the other direction.

Miscellaneous. Condon and Hernek [14] study cover times in the following setting. The edges of a graph are colored, a sequence (c_t) of colors is prespecified and the "random walk" at step t picks an edge uniformly at random from the color-c_t edges at the current vertex.

References

[1] D.J. Aldous. Random walk covering of some special trees. *J. Math. Analysis Appl.*, 157:271–283, 1991.

[2] D.J. Aldous. Threshold limits for cover times. *J. Theoretical Probab.*, 4:197–211, 1991.

[3] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze traversal. In *Proc. 20th IEEE Symp. Found. Comp. Sci.*, pages 218–233, 1979.

[4] N. Alon and J. H. Spencer. *The Probabilistic Method.* Wiley, 1992.

[5] G. Barnes and U. Feige. Short random walks on graphs. 1992.

[6] A. Borodin, W.L. Ruzzo, and M. Tompa. Lower bounds on the length of universal traversal sequences. *J. Computer Systems Sci.*, 45:180–203, 1992.

[7] G. Brightwell and P. Winkler. Extremal cover times for random walks on trees. *J. Graph Theory*, 14:547–554, 1990.

[8] G. Brightwell and P. Winkler. Maximum hitting time for random walks on graphs. *Random Struct. Alg.*, 1:263–276, 1990.

[9] A. Broder and A.R. Karlin. Bounds on the cover time. *J. Theoretical Probab.*, 2:101–120, 1989.

[10] A. Broder, A.R. Karlin, P. Raghavan, and E. Upfal. Trading space for time in undirected $s-t$ connectivity. In *Proc. 21st ACM Symp. Theory of Computing*, pages 543–549, 1989.

[11] E.A. Carlen, S. Kusuoka, and D.W. Stroock. Upper bounds for symmetric Markov transition functions. *Ann. Inst. H. Poincare Sect. B*, Suppl. 2:245–287, 1987.

[12] 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. 21st ACM Symp. Theory of Computing*, pages 574–586, 1989.

[13] G. Chartrand and L. Lesniak. *Graphs and Digraphs.* Wadsworth, 1986.

[14] A. Condon and D. Hernek. Random walks on colored graphs. *Random Struct. Alg.*, 5:285–303, 1994.

[15] D. Coppersmith, U. Feige, and J. Shearer. Random walks on regular and irregular graphs. Technical Report CS93-15, Weizmann Institute, 1993.

[16] L. Devroye and A Sbihi. Inequalities for random walks on trees. In A. Frieze and T. Luczak, editors, *Random Graphs*, volume 2, pages 35–45. Wiley, 1992.

[17] R. Durrett. *Probability: Theory and Examples*. Wadsworth, Pacific Grove CA, 1991.

[18] U. Feige. Collecting coupons on trees, and the analysis of random walks. Technical Report CS93-20, Weizmann Institute, 1993.

[19] U. Feige. A tight lower bound on the cover time for random walks on graphs. Unpublished, 1993.

[20] U. Feige. A tight upper bound on the cover time for random walks on graphs. *Random Struct. Alg.*, 6, 1995.

[21] L. Flatto, A.M. Odlyzko, and D.B. Wales. Random shuffles and group representations. *Ann. Probab.*, 13:154–178, 1985.

[22] K.J. Harrison and M.W. Short. The last vertex visited in a random walk on a graph. Technical report, Murdoch University, Australia, 1992.

[23] J. P. Imhof. On the range of Brownian motion and its inverse process. *Ann. Probab.*, 13:1011–1017, 1985.

[24] J.D. Kahn, N. Linial, N. Nisan, and M.E. Saks. On the cover time of random walks on graphs. *J. Theoretical Probab.*, 2:121–128, 1989.

[25] L. Lovász and P. Winkler. A note on the last new vertex visited by a random walk. *J. Graph Theory*, 17:593–596, 1993.

[26] J. L. Palacios. A bound for the covering time of random walks on graphs. *Statistics and Probab. Lett.*, 14:9–11, 1992.

[27] J.L. Palacios. On a result of Aleliunas et al. concerning random walks on graphs. *Prob. Engineering Inform. Sci.*, 4:489–492, 1990.

[28] A.M. Sbihi. *Covering Times for Random Walks on Graphs*. PhD thesis, McGill University, 1990.

[29] D. Zuckerman. Covering times of random walks on bounded degree trees and other graphs. *J. Theoretical Probab.*, 2:147–157, 1989.

[30] D. Zuckerman. On the time to traverse all edges in a graph. *Information Proc. Letters.*, 38:335–337, 1991.

[31] D. Zuckerman. A technique for lower bounding the cover time. *SIAM J. Disc.Math.*, 5:81–87, 1992.