Exact mixing times for random walks on trees

Andrew Beveridge[*] and Meng Wang
Department of Mathematics, Statistics and Computer Science
Macalester College
Saint Paul, MN 55105

Abstract

We characterize the extremal structures for certain random walks on trees. Let $G = (V, E)$ be a tree with stationary distribution π. For a vertex $i \in V$, let $H(\pi, i)$ and $H(i, \pi)$ denote the expected lengths of optimal stopping rules from π to i and from i to π, respectively. We show that among all trees with $|V| = n$, the quantities $\min_{i\in V} H(\pi, i)$, $\max_{i\in V} H(\pi, i)$, $\max_{i\in V} H(i, \pi)$ and $\sum_{i\in V} \pi_i H(i, \pi)$ are all minimized uniquely by the star $S_n = K_{1,n-1}$ and maximized uniquely by the path P_n.

1 Introduction

Naturally, we think of the star $S_n = K_{1,n-1}$ and the path P_n as the extremal tree structures. This is also true with respect to random walks. Herein, we consider a few different quantities associated with random walks, each of which identifies the star as the minimizing tree structure and the path as the maximizing tree structure.

Let $G = (V, E)$ be a graph. A *random walk* on G is a sequence of vertices $(w_0, w_1, \ldots, w_t, \ldots)$ such that $\Pr(w_{t+1} = j \mid w_t = i\}$ is $1/\deg(i)$ if $(i, j) \in E$ and 0 otherwise. For an introduction to random walks on graphs, see [9] or [10]. For $i \neq j$, the *hitting time* $H(i, j)$ is the expected number of steps before a walk started at i visits vertex j. We define $H(i, i) = 0$. When G is not bipartite, as t tends to infinity the distribution of w_t tends to the *stationary distribution* π, where $\pi_i = \deg(i)/2|E|$. For bipartite G, we have the same convergence if we consider a lazy walk in which we remain at the current state with probability 1/2, at the cost of doubling the expected length of any walk. For simplicity of exposition, we will consider non-lazy walks on trees. We also assume $n \geq 4$, since there is a single unique tree structure for $n \leq 3$.

---

[*] Supported in part by NSA Young Investigator Grant H98230-08-1-0064.

The study of extremal graph structures with respect to random walks has been well developed. Given a vertex i, we use i' to denote an *i-pessimal vertex* achieving

$$H(i', i) = \max_{j \in V} H(j, i).$$

Brightwell and Winkler [5] characterized the graph structure that maximizes the pessimal hitting time. The *cover time* of a graph is the expected number of steps required to visit every vertex of the graph. Feige [8] proved that the cover time of any graph is at least $(1+o(1))n \log n$. Brightwell and Winkler [6] showed that among trees, the star has the minimum cover time (starting from the center) and the path has the maximum cover time (starting from a leaf).

Proposition 24 in chapter 6 of [1] collects a number of extremal results for random walks on trees on $n \geq 3$ vertices. Among the results listed are the following.

Proposition 1.1 *For any tree G on n vertices, we have*

(a) $2(n-1) \leq \max_{i \in V} \max_{j \in V} H(j, i) = \max_{i \in V} H(i', i) \leq (n-1)^2$,

(b) $n - \frac{3}{2} \leq \sum_{i,j \in V} \pi_i \pi_j H(i, j) \leq \frac{1}{6}(2n^2 - 4n + 3)$,

(c) $1 \leq (1 - \lambda_2)^{-1} \leq (1 - \cos(\pi/(n-1)))^{-1}$.

Moreover, the minimum is attained uniquely by the star and the maximum is attained uniquely be the path.

The *relaxation time* $(1-\lambda_2)^{-1}$ from statement (c) is a well-known approximate mixing measure.

We add this to list of extremal results by considering a number of interesting quantities arising in the study of *exact stopping rules.* Given an initial distribution σ and a target distribution τ, a (σ, τ)-*stopping rule* halts a random walk whose initial state is drawn from σ so that the distribution of the final state is τ. See [11, 12] for details. The *access time* $H(\sigma, \tau)$ is the minimum expected length of all stopping rules that achieve τ when started from σ. A (σ, τ)-stopping rule is *optimal* if it achieves this minimum expected length.

For our first two results, we consider the case where our initial distribution is π and our target is a singleton. We consider rules to the worst and best target vertices. In each case, the star and the path are the extremal structures.

Theorem 1.2 *If G is a tree on n vertices then*

$$2n - \frac{7}{2} \leq \max_{i \in V} H(\pi, i) \leq \frac{4n^2 - 8n + 3}{6}.$$

The minimum value is achieved uniquely by the star S_n and the maximum value is achieved uniquely by the path P_n.

Theorem 1.3 *If G is a tree on n vertices then*

$$\frac{1}{2} \leq \min_{i \in V} H(\pi, i) \leq \begin{cases} \frac{1}{6}(n^2 - 2n + 3) & \textit{if } n \textit{ is even,} \\ \frac{1}{6}(n^2 - 2n) & \textit{if } n \textit{ is odd.} \end{cases}$$

The minimum value is achieved uniquely by the star S_n and the maximum value is achieved uniquely by the path P_n.

As per [1] (chapter 5, p. 48), a result equivalent to Theorem 1.3 appears in unpublished work of Yaron [14]. His result addressed variance of return times, rather than access times. We give a direct proof of this fact using hitting times.

Our second pair of results concern mixing times for trees. A number of parameterless mixing measures defined via stopping rules have been introduced and studied [2, 3, 4, 11, 12]. Among the most important measures are the mixing time $\mathcal{T}_{\text{mix}} = \max_{i \in V} H(i, \pi)$ and the reset time $\mathcal{T}_{\text{reset}} = \sum_{i \in V} \pi_i H(i, \pi)$. We interpret $\mathcal{T}_{\text{mix}}$ as the pessimal mixing time and $\mathcal{T}_{\text{reset}}$ as the average mixing time. Considering these mixing measures for trees, we find that once again, the star and the path are the extremal structures.

Theorem 1.4 *If G is a tree on n vertices then*

$$\frac{3}{2} \leq \mathcal{T}_{\text{mix}} \leq \frac{2n^2 - 4n + 3}{6}.$$

The minimum value is achieved uniquely by the star S_n and the maximum value is achieved uniquely by the path P_n.

Theorem 1.5 *If G is a tree on n vertices then*

$$1 \leq \mathcal{T}_{\text{reset}} \leq \begin{cases} \frac{1}{4}(n^2 - 2n + 2) & \textit{if } n \textit{ is even,} \\ \frac{1}{4}(n - 1)^2 & \textit{if } n \textit{ is odd.} \end{cases}$$

The minimum value is achieved uniquely by the star S_n and the maximum value is achieved uniquely by the path P_n.

In summary, by using the theory of exact stopping rules, we can extend the list of extremal characterizations for random walks on trees. In particular, the exact mixing results of Theorems 1.4 and 1.5 complement the approximate mixing result of Proposition 1.1(c).

The organization of this paper is as follows. In Section 2, we recall previous results about stopping rules and random walks on trees. In Section 3, we prove a lemma concerning pessimal hitting times. In Section 4, we prove our theorems concerning $H(\pi, i)$ and in Section 5, we prove our results on exact mixing measures.

2 Preliminaries

2.1 Random walks and stopping rules

We summarize some essential results from [11]. When σ and τ are concentrated on single states i and j respectively (we write $\sigma = i$, $\tau = j$), then the access time $H(i, j)$ is the hitting time from i to j. In this instance, the only optimal stopping rule is "walk until you hit j." Whenever our target is concentrated on a vertex, the analogous rule is optimal for any starting distribution:

$$H(\sigma, j) = \sum_{i \in V} \sigma_i H(i, j).$$

Typically $H(\sigma, \tau) < \sum_{i \in V} \sigma_i H(i, \tau)$ when τ is not a singleton.

Let Γ be a (σ, τ)-stopping rule. For each $i \in V$ the *exit frequency* $x_i(\Gamma)$ is the expected number of times the walk leaves state i before halting. Exit frequencies partition the expected length of the stopping rule: $E[\Gamma] = \sum_{i \in V} x_i(\Gamma)$. A key observation due to Pitman [13] is that exit frequencies satisfy the *conservation equation*

$$\sum_{i \in V} p_{ij} x_i(\Gamma) - x_j(\Gamma) = \tau_j - \sigma_j, \ \ \forall j \in V.$$

An important consequence of this equation is that the exit frequencies for two rules from σ to τ differ by $K\pi_i$ where K is the difference between the expected lengths of these rules. Hence the distributions σ, τ uniquely determine the exit frequencies for an optimal stopping rule between them. We denote these *optimal exit frequencies* by $x_i(\sigma, \tau)$. Furthermore,

$$\Gamma \text{ is an optimal stopping rule } \iff \exists k \in V, x_k(\Gamma) = 0. \tag{1}$$

We call such a vertex k a (σ, τ)-*halting state*, or simply a *halting state* when the initial and target distributions are clear. Note that an optimal rule may have multiple halting states, but we need only identify one such state to ensure that a rule is optimal.

Finally, we state the very useful *random target identity* (c.f. [10], equation (3.3)):

$$\sum_{k \in V} \pi_k H(i,k) = \sum_{k \in V} \pi_k H(j,k) \text{ for all } i,j \in V. \tag{2}$$

2.2 Exact mixing measures

Stopping rules give rise to a number of exact mixing measures (as opposed to approximate mixing measures that incorporate an approximation parameter $\epsilon > 0$). We collect some results that were originally proven in [12], and subsequently developed into a broader framework in [4]. The *mixing time* $\mathcal{T}_{\text{mix}}$ is the expected length of an optimal mixing rule starting from the worst initial vertex:

$$\mathcal{T}_{\text{mix}} = \max_{i \in V} H(i,\pi).$$

A vertex achieving this maximum is called *mixing pessimal.* The *forget time* $\mathcal{T}_{\text{forget}}$ is the smallest t for which there exists a distribution μ such that for every starting vertex, the expected time to attain μ via an optimal rule is at most t:

$$\mathcal{T}_{\text{forget}} = \max_{i \in V} H(i,\mu) = \min_{\tau} \max_{i \in V} H(i,\tau).$$

The forget time is attained by a unique distribution given by

$$\mu_i = \pi_i \left(1 + \sum_{j \in V} p_{ij} H(j,\pi) - H(i,\pi) \right). \tag{3}$$

Furthermore, if a vertex is mixing pessimal then it is also pessimal for μ.

The *reset time*

$$\mathcal{T}_{\text{reset}} = \sum_{i \in V} \pi_i H(i,\pi)$$

can be viewed as an average mixing time. We have the remarkable equality

$$\mathcal{T}_{\text{forget}} = \mathcal{T}_{\text{reset}} \tag{4}$$

for a random walk on an undirected graph. Moreover, for an undirected graph we have $\mathcal{T}_{\text{reset}} \leq \mathcal{T}_{\text{mix}} \leq 4\mathcal{T}_{\text{reset}}$. For a tree, we have the tighter upper bound $\mathcal{T}_{\text{mix}} \leq 2\mathcal{T}_{\text{reset}}$.

2.3 Random walks on trees

We summarize some results from [3]. For adjacent vertices u, v we define $V_{u:v}$ to be the set of vertices in the subtree rooted at u after the removal of edge (u, v). For any two adjacent vertices i and j on tree $G = (V, E)$,

$$H(i,j) = \sum_{k \in V_{i:j}} \deg(k) = 2|E| \sum_{k \in V_{i:j}} \pi_k = 2|E| \pi(V_{i:j}).$$

Let $d(i, j)$ denote the path distance between vertices i and j. Define $\ell(i, k; j) = \frac{1}{2}(d(i, j) + d(k, j) - d(i, k))$ where i, j, and k are vertices on G. The quantity $\ell(i, k; j)$ measures the length of the intersection of the (i, j)-path and the (k, j)-path. For any pair i, j,

$$H(i,j) = \sum_{k \in V} \ell(i,k;j) \deg(k). \tag{5}$$

Indeed, we can confirm the well-known formula for the hitting times on a path P_n on vertices $\{1, 2, \ldots, n\}$:

$$H(i,j) = \begin{cases} (j-1)^2 - (i-1)^2 & \text{for } i \leq j, \\ (n-j+1)^2 - (n-i+1)^2 & \text{for } i > j. \end{cases} \tag{6}$$

There are two distinct centers for random walks on trees: an "average center" and an "extremal" center. The average center is the *barycenter*, which is the vertex (or two adjacent vertices) that achieves $\min_{i \in V} \sum_{j \in V} d(i, j)$. In other words, the barycenter minimizes the total distance to all other vertices. The following statements for a vertex c are equivalent:

$$\begin{array}{ll} (a) & \text{The vertex } c \text{ is a barycenter of the tree.} \\ (b) & \text{The vertex } c \text{ satisfies } H(i,c) \leq H(c,i) \text{ for all vertex } i. \\ (c) & \sum_{k \in V} \pi_k H(k,c) = \min_{i \in V} \sum_{k \in V} \pi_k H(k,i). \\ (d) & \text{For every vertex } i \text{ adjacent to } c,\ \pi(V_{i:c}) = \sum_{k \in V_{i:c}} \pi_k \leq 1/2. \end{array} \tag{7}$$

The extremal center of the tree is the *focus* (or two adjacent *foci*), introduced in [3]. The focus (or foci) is the "extremal" center of the tree. There are two types: primary foci and secondary foci. A *primary focus* is a vertex $a \in V$ satisfying $H(a', a) = \min_{j \in V} \max_{i \in V} H(i, j)$. For example, the path P_{2k+1} has vertex $k + 1$ as its unique primary focus, while the path P_{2k} has two primary foci: vertices k and $k + 1$.

A tree with one focus is called a *focal* tree. A tree is focal if and only if there are a-pessimal vertices in at least two components of $G - a$. For a focal tree, we use a' and a'' to denote a-pessimal points in distinct components of $G - a$. Asymmetric trees typically have two foci, and such trees are called *bifocal.* Once we have identified the primary focus a on a bifocal tree, the other focus b will be the neighbor of a on the unique (a, a')-path. Vertex b is a primary focus when $H(b', b) = H(a', a)$ (for example: P_{2k}). Vertex b is a *secondary focus* when $H(b', b) > H(a', a)$. In summary, a tree must either have a single primary focus or two adjacent foci. In the latter case, a tree may have two primary foci, or one primary focus and one secondary focus.

We reserve the symbols a, b to denote the foci of the tree, and symbols a', a'' and b' for their pessimal vertices, as appropriate. The foci of the tree play a special role for mixing and forgetting walks on the tree. In particular, if u is the focus that is closest to vertex i then

$$H(i, \pi) = H(i, u) + H(u, \pi). \tag{8}$$

Finally, the forget distribution μ is concentrated on the foci of the tree. For a focal tree, μ is the singleton distribution on a and we have $\mathcal{T}_{\text{forget}} = H(a', a) = H(a'', a)$. For a bifocal tree, we have

$$\mu_a = \frac{H(b', b) - H(a', b)}{2|E|} \quad \text{and} \quad \mu_b = \frac{H(a', a) - H(b', a)}{2|E|},$$

with

$$\mathcal{T}_{\text{forget}} = H(a', \mu) \;=\; H(b', \mu) = \frac{H(a, b)\, H(a', a) + H(b, a)\, H(b', a)}{2|E|}. \tag{9}$$

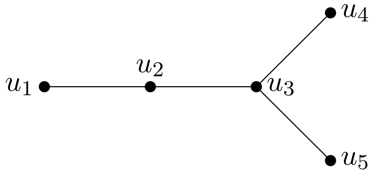

Figure 2.1: A tree on 5 vertices

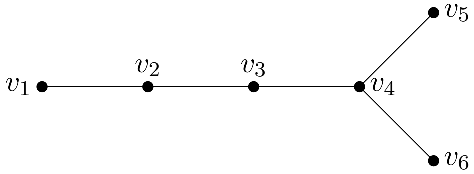

Figure 2.2: A tree on 6 vertices

We use equation (5) to calculate the foci and the barycenters for two examples. The tree in Figure 2.1 has primary focus $a = u_3$ with $H(a', a) = H(u_1, u_3) = 4$ and secondary focus $b = u_2$ with $H(b', b) = H(u_4, u_2) = 6$. The vertex u_3 is the unique barycenter of the tree. The tree in Figure 2.2 has primary focus $a = v_3$ with $H(a', a) = H(v_5, v_3) = 6$ and secondary focus $b = v_4$ with $H(b', b) = H(v_1, v_4) = 9$. Both v_3 and v_4 are barycenters, as each satisfy condition (d) of equation (7).

In general, foci are distinct from barycenters. Indeed, let $B_{r,s}$ denote the broom graph consisting of a star with s leaves $u_1, u_2, \ldots, u_s$ centered at vertex c, along with a path of length s on vertices $c = v_0, v_1, \ldots, v_r$. For the broom graph $B_{4k,4k}$, the vertex c is the unique barycenter. Simple calculations using formula (5) show that vertex $a = v_k$ is the primary focus with $H(v_k', v_k) = H(u_1, v_k) = 9k^2 + 1$ and that $b = v_{k-1}$ is the secondary focus with $H(v_{k-1}', v_{k-1}) = H(v_{4k}, v_{k-1}) = (3k+1)^2$.

3 Extremal Hitting Times

We begin by setting some terminology. In the remainder of this paper, we will often compare quantities for two trees $G = (V, E)$ and $G^* = (V^*, E^*)$. The tree G^* will be one of S_n, P_n or a slight modification of G. We use a starred notation for quantities related to G^*: for example its stationary distribution, degrees and hitting times are denoted by π^*, $\deg^*(i)$ and $H^*(i,j)$, respectively.

Let G be a tree on n vertices. Clearly, if $H(a,b) = \max_i \max_j H(j,i) = \max_i H(i', i)$ then a, b must be leaves. We start by proving a hitting time result that is complementary to statement (a) of Proposition 1.1.

Lemma 3.1 *For any tree G on n vertices,*

$$1 \leq \min_i H(i', i) \leq \begin{cases} \frac{1}{4}n^2 & \text{if } n \text{ is even,} \\ \frac{1}{4}(n-1)^2 & \text{if } n \text{ is odd.} \end{cases}$$

The lower bound is achieved uniquely by the star. For n is even, the path is the unique maximizing structure. For n odd, the maximum is achieved by the path, as well as any bifocal tree such that the induced subgraph on $V_{b:a}$ is $P_{(n-1)/2}$, and $\frac{1}{4}(n-1)^2 - n \leq H(b', a) < \frac{1}{4}(n-1)^2$.

Proof. Clearly, a star achieves the lower bound. If G is not a star, then for every vertex i, there exists a vertex j such that $d(i,j) \geq 2$, so $H(i', i) > 1$. We now consider the upper bounds. Using equation (6), we can verify that the central vertex (or vertices) of the path P_n is the best target and achieves the stated upper bound.

Case 1: n is even. Suppose that G is a focal tree with $H(a', a) = H(a'', a) \geq n^2/4$. Let A' and A'' be the components of $G - a$ containing a' and a'', respectively. At least one of these components (say, A') has size at most $\lfloor \frac{1}{2}(n-1) \rfloor = \frac{1}{2}(n-2)$. This forces $H(a', a) \leq \frac{1}{4}(n-2)^2 < \frac{1}{4}n^2$ by equation (6), a contradiction. So G cannot be a focal tree.

Consider a bifocal G with foci a, b where $H(a', a) \leq H(b', b)$. If one of $V_{a:b}$, $V_{b:a}$ has size at most $\frac{1}{2}n - 1$, then $\min_i H(i', i) \leq (n/2 - 1)^2$, a contradiction. So $|V_{a:b}| = |V_{b:a}| = \frac{1}{2}n$. If either of these components is not a path, then $\min_i H(i', i) < \frac{1}{4}n^2$ by statement (a) of Proposition 1.1. In conclusion, G must be a path.

CASE 2: n is odd. Suppose that G is a focal tree with unique focus a, such that $H(a', a) = H(a'', a) \geq \frac{1}{4}(n-1)^2$. By Proposition 1.1(a), the components of $G - a$ containing a' and a'' must both be of size at least $\frac{1}{2}(n-1)$. All vertices are accounted for, so these are the only two components of $G - a$. Furthermore, to achieve this hitting time, both components must be paths.

Finally, suppose that G is bifocal with foci a, b with $H(a', a) = \min_{i \in V} H(\pi, i)$. At least one of the components $V_{a:b}$, $V_{b:a}$ has at most $\frac{1}{2}(n-1)$ vertices. Similar to our argument above, this component must be a path of length $\frac{1}{2}(n-1)$. Therefore $H(a', a) \leq \frac{1}{4}(n-1)^2$.

It remains to characterize the structure of odd bifocal trees attaining the upper bound. We know that $V_{b:a}$ must be a path on $\frac{1}{2} * (n-1)$ vertices and $|V_{a:b}| = \frac{1}{2}(n+1)$. In order for both a and b to be foci, we must have $H(b', a) < H(a', a) = \frac{1}{4}(n-1)^2$. Furthermore, this tree attains the upper bound, so $\frac{1}{4}(n-1)^2 \leq H(b', b) = H(b', a) + H(a, b) = H(b', a) + n$. Putting the bounds together gives $\frac{1}{4}(n-1)^2 - n \leq H(b', a) < \frac{1}{4}(n-1)^2$. ∎

For large odd n, there is a wide variety of extremal bifocal tree structures achieving $\min_{i \in V} H(i', i) = \frac{1}{4}(n-1)^2$. We limit ourselves to describing the tree structures for the two extreme cases for $d(a, b')$. Suppose that G is a bifocal graph on $2k + 1$ vertices with $H(a', a) = \min_{i \in V} H(i', i) = k^2$. The hitting time restriction for $H(b', a)$ in Lemma 3.1 translates to

$$k^2 - 2k - 1 \leq H(b', a) < k^2, \tag{10}$$

with $k^2 - 2k - 1 = H(b', a)$ if and only if b is also a primary focus. Let F be the subgraph on $k + 1$ vertices induced by $V_{a:b}$, and let $a = v_1, v_2, \ldots, v_r = b'$ be the unique (a, b') path, where $r \leq k$ (since G is bifocal).

First, consider the case $r = k$, so that $d(a, b') = k - 1$. Such a graph F is a k-path with a pendant edge (v_i, u) where $1 \leq i \leq k - 1$. Setting $F^* = P_k$ and using equation (5), we have $H(b', a) = H(v_k, v_1) = H^*(v_k, v_1) + \ell(v_k, u; v_1) + \ell(v_k, v_i; v_1) = (k-1)^2 + 2(i-1)$. Equation (10) is satisfied with strict inequality for $1 \leq i \leq k - 1$. So every such graph structure is valid, and b is always a secondary focus.

Next, we determine the lower bound for $r = d(a, b') + 1$. For a fixed r, the hitting time $H(b', a)$ is maximized by appending the remaining $k+1-r$ vertices as leaves adjacent to v_{r-1}. In other words, F is a broom graph with a at the handle tip. Setting $F^* = P_r$, we use equation (5) to yield $H(b', a) = H(v_r, v_1) = H^*(v_r, v_1) + 2(k-2) = (r-1)^2 + 2(k+1-r)(r-2)$. The first inequality of equation (10) is then equivalent to the condition $r \geq k + 2 - \sqrt{2k+2}$.

Considering a more specific example, if $k = 7$ then we can satisfy this condition with equality. This results in an asymmetric extremal tree with two primary foci! Indeed, the resulting tree G is the broom graph $B_{11,4}$ on 15 vertices, for which $H(a', a) = H(b', b) = 49$. More generally, we can construct extremal trees for $k + 2 - \sqrt{2k+2} \leq r \leq k$ by strategically placing the remaining $k+1-r$ vertices to guarantee that equation (10) is satisfied. For example, Figure 3.1, enumerates all extremal structures for the induced subgraph on $V_{a:b}$ when $k = 6$.

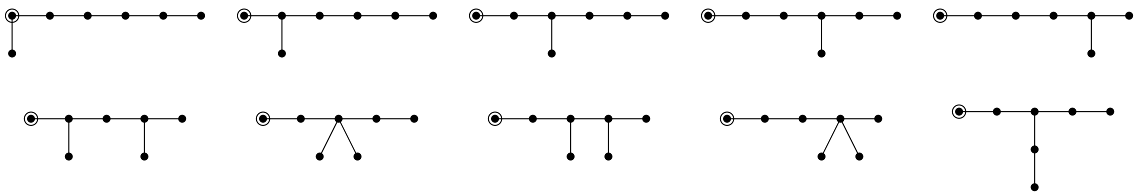

Figure 3.1: The 10 possible structures for the induced 7-vertex subgraph on $V_{a:b}$ for bifocal trees on 13 vertices that achieve $\min_{i \in V} H(i', i) = 36$. The primary focus a is the circled upper left hand vertex of each tree. The remaining (nondisplayed) 6 vertices in $V_{b:a}$ are a 6-path connected to vertex a.

4 Access times from the stationary distribution

In this section, we prove Theorems 1.2 and 1.3.

Proof of Theorem 1.2. The target vertex achieving $\max_{i \in V} H(\pi, i)$ must be a leaf. Indeed, let u be a non-leaf with neighbors $v_1, \ldots v_d$ where $\pi(V_{v_1:u}) \leq \pi(V_{v_2:u}) \leq \cdots \leq \pi(V_{v_d:u})$. This ordering ensures that $\pi(V_{v_1:u}) < \pi(V_{u:v_1})$, so that

$$H(\pi, v_1) - H(\pi, u) = \pi(V_{u:v_1}) H(u, v_1) - \pi(V_{v_1:u}) H(v_1, u) = 2|E| \left(\pi(V_{u:v_1})^2 - \pi(V_{v_1:u})^2 \right) > 0,$$

where the second equality follows from equation (5). We may repeat this process until reaching leaf, strictly increasing the access time with each iteration.

Consider the lower bound. Let $G^* = S_n$ denote the star of n vertices with center c. If y, z are distinct leaves then $H^*(c,z) = 2n-3$ and $H^*(y,z) = 2n-2$. We have $H^*(\pi^*, z) = (1-\pi^*_z - \pi^*_c)H^*(y,z) + \pi^*_c H^*(c,z) = 2n - 7/2$.

Let G be a tree on n vertices with leaf u attaining $H(\pi, u) = \max_i H(\pi, i)$. Let the unique neighbor of u be v, so that $H(v,u) = 2n-3$. Since G is not a star, $\pi_v < \pi^*_c$. We have

$$\begin{aligned} H(\pi,u) - H^*(\pi^*, z) &= \pi_v H(v,u) - \pi^*_c H^*(c,z) + \sum_{i \in V - v} \pi_i H(i,u) - \sum_{j \in V^* - c} \pi^*_j H(j,u) \\ &> (\pi_v - \pi^*_c)(2n-3) + \big((1-\pi_v - \pi_u) - (1 - \pi^*_c - \pi^*_z)\big)(2n-2) \\ &= \pi^*_c - \pi_v \, > \, 0. \end{aligned}$$

The first inequality follows from the existence of a nontrivial component of $G - v$, giving at least two vertices w such that $H(w,u) > 2n-2$. This strict inequality shows that the star is the unique minimizing structure.

Next, we consider the upper bound. Taking $G = P_n$, we confirm that

$$\begin{aligned} H(\pi, 1) &= \frac{1}{n-1}\sum_{k=1}^{n-1}((n-1)^2 - (k-1)^2) - \frac{(n-1)^2}{2(n-1)} \\ &= (n-1)^2 - \frac{1}{n-1}\sum_{j=1}^{n-2} j^2 - \frac{n-1}{2} \;=\; \frac{4n^2 - 8n + 3}{6}. \end{aligned} \tag{11}$$

Given any tree $G \neq P_n$ on n vertices, we construct a related tree G^* by moving one leaf, and then show that $\max_{i \in V} H(\pi, i) < \max_{j \in V^*} H^*(\pi^*, j)$. Let the leaf $u \in V$ achieve $H(\pi, u) = \max_i H(\pi, i)$. The u-pessimal vertex u' must also be a leaf. Let z be nearest vertex to u' on G with degree at least 3. Let x be a leaf in $V \backslash \{u, u'\}$ such that the (x, u')-path contains z. Let the unique neighbor of x be y (possibly $y = z$). Let G^* be the graph obtained by removing the leaf x and then adding a leaf x^* adjacent to u'. We claim that

$$\max_{i \in V} H(\pi, i) = H(\pi, u) < H^*(\pi^*, u) \leq \max_{j \in V^*} H^*(\pi^*, j).$$

We partition the vertices of G into 4 sets. The forest $G - z$ has $\deg(z)$ components. Let $U_{u'}, U_x$ be the components containing u' and x, respectively. Our partition of V is $V_x = \{x\}$, $V_{u'} = U_{u'}$, $V_z = U_x - x + z$ and $W = V \backslash (V_x \cup V_{u'} \cup V_z)$. We use an analogous partition for V^*, replacing V_x with $V_{x^*} = \{x^*\}$.

For $i \in V_{u'} \cup V_z \cup W$, we have

$$H^*(i,z) - H(i,z) = \begin{cases} 0 & \text{if } i \in W, \\ -\ell(i,y;z) - \ell(i,x;z) \;=\; -2\ell(i,y;z) & \text{if } i \in V_z, \\ \ell(i,u';z) + \ell(i,x^*;z) \;=\; 2\ell(i,u';z) & \text{if } i \in V_{u'}. \end{cases} \tag{12}$$

Furthermore, for $i \in V \backslash \{x, y, u'\}$ we have $\deg^*(i) = \deg(i)$, while $\deg^*(y) = \deg(y) - 1$, $\deg^*(u') = \deg(u') + 1$ and $\deg^*(x^*) = \deg(x) = 1$.

We decompose the quantity

$$H^*(\pi^*, u) - H(\pi, u) = \sum_{i \in V^*} \pi_i^* H(i,u) - \sum_{j \in V} \pi_j H(j,u) \tag{13}$$

according to our partition. For $i \in W$, we have $\pi_i H(i,u) = \pi_i^* H^*(i,u)$, so

$$\sum_{i \in W} \big(\pi_i^* H^*(i,u) - \pi_i H(i,u)\big) = 0.$$

We calculate

$$\begin{aligned} \sum_{i \in V_z} \big(\pi_i^* H^*(i,u) - \pi_i H(i,u)\big) &= -\frac{1}{2|E|} H(y,u) + \sum_{i \in V_z} \pi_i^* \big(H^*(i,u) - H(i,u)\big) \\ &= -\frac{1}{2|E|} H(y,u) + \sum_{i \in V_z} \pi_i^* \big(H^*(i,z) - H(i,z)\big) \\ &= \frac{1}{2|E|} \left(-H(y,u) - 2 \sum_{i \in V_z} \deg^*(i) \ell(y,i;z) \right) \\ &= \frac{1}{2|E|} \left(-H(y,u) - 2H^*(y,z) \right) \end{aligned}$$

where the third and fourth equalities come from equations (12) and (5), respectively. Next, we have

$$\begin{aligned} \sum_{i \in V_{u'}} \big(\pi_i^* H^*(i,u) - \pi_i H(i,u)\big) &= \frac{1}{2|E|} H^*(u',u) + \sum_{i \in V_{u'}} \pi_i \big(H^*(i,u) - H(i,u)\big) \\ &= \frac{1}{2|E|} \left(H^*(u',u) + \sum_{i \in V_{u'}} \deg(i) \big(H^*(i,z) - H(i,z)\big) \right) \\ &= \frac{1}{2|E|} \left(H^*(u',u) + 2 \sum_{i \in V_{u'}} \deg(i) \ell(u', i, z) \right) \\ &= \frac{1}{2|E|} \big(H^*(u',u) + 2H(u',z) \big), \end{aligned}$$

where as above, the last two equalities are due to equations (12) and (5). Finally, we can expand equation (13) to find that

$$\begin{aligned} & H^*(\pi^*, u) - H(\pi, u) \\ & \quad = \frac{-H(y,u) - 2H^*(y,z) + H^*(u',u) + 2H(u',z) + H^*(x^*,u) - H(x,u)}{2|E|} \\ & \quad = \frac{-H(y,u) - H^*(y,z) + H^*(u',u) + H(u',z)}{|E|}. \end{aligned}$$

By equation (5), $H(z,u) = H^*(z,u)$, so that we may use $-H(y,u) + H^*(u',u) = -H(y,z) + H^*(u',z)$ to yield

$$\begin{aligned} H^*(\pi^*, u) - H(\pi, u) &= \frac{\big(H(u',z) - H(y,z)\big) + \big(H^*(u',z) - H^*(y,z)\big)}{|E|} \\ &> \frac{H(u',z) - H(x,z)}{|E|} \geq 0, \end{aligned} \tag{14}$$

where the final inequality follows from $H(u',z) + H(z,u) = H(u',u) \geq H(x,u) = H(x,z) + H(z,u)$. ∎

Proof of Theorem 1.3. By Proposition 1.1, $\min_{i \in V} H(\pi, i)$ is achieved by the barycenter of the tree. Verifying the lower bound is easy: if c is the central vertex in S_n, then $\min_{i \in V(S_n)} H(\pi, i) = H(\pi, c) = 1/2$. For $G \neq S_n$, let u achieve $H(\pi, u) = \min_{iinV} H(\pi, i)$. We must have $\pi_u < 1/2$ and therefore $H(\pi, u) \geq (1 - \pi_u) > 1/2$.

As for the upper bound, we first calculate $\min_i H(\pi, i)$ for the path. By condition (c) of equation (7), $\min_i H(\pi, i)$ is achieved by the barycenter. For odd n, if we set $G^* = P_{(n+1)/2}$ then by symmetry

$$\min_i H(\pi, i) = H(\pi, (n+1)/2) = 2 \cdot \frac{1}{2} H^*(\pi^*, 1) = \frac{n^2 - 2n}{6}.$$

For an even path, setting $G^* = P_{n/2}$ and $G^{**} = P_{n/2+1}$, we have

$$\min_i H(\pi, i) = H(\pi, n/2) = \frac{n-2}{2n-2} H^*(\pi^*, 1) + \frac{n}{2n-2} H^{**}(\pi^{**}, 1) = \frac{n^2 - 2n + 3}{6}.$$

Let $G \neq P_n$ be a tree on n vertices with barycenter c. We construct a tree G^* such that $H(\pi, c) = \min_{i \in V} H(\pi, i) < \min_{j \in V^*} H^*(\pi^*, j)$. Let $G_1, G_2, ..., G_d$ be the components of $G - c$, where $\pi(G_1) \geq \pi(G_2) \geq \cdots \geq \pi(G_d)$, and $u_k \in V(G_k)$ is the vertex adjacent to c for $1 \leq k \leq d := \deg(c)$. We claim that each component of $G - c$ must be a path. Indeed, let G_k

be the tree with $V_k := V(G_k) \cup \{c\}$ and $E_k := E(G_k) \cup \{(b_k, c)\}$. Using $H^{(k)}(u,v)$ and $\pi^{(k)}$ to denote hitting times and the stationary distribution for G_k, we have

$$H(\pi, c) = \sum_{k=1}^{d} \sum_{v \in V(G_k)} \pi_v H(v,c) = \sum_{k=1}^{d} \sum_{v \in V(G_k)} \frac{|E_k|}{|E|} H^{(k)}(\pi^{(k)}, c).$$

Fixing the size of the components, maximizing each of the $H^{(k)}(\pi^{(k)}, c)$ will result in the maximum $H(\pi, c)$. By Theorem 1.2, each of these components must be a path.

Since G is not a path, we must have $d > 2$. Since $\pi(G_1) \geq \pi(G_2) \geq \pi(G_3)$, and c is the barycenter, $\pi(G_1) \leq \frac{1}{2}$, so $\pi(G_2) \leq 1 - \pi(G_1 \cup \{c\}) \leq \frac{1}{2} - \frac{3}{2(n-1)}$. Let $x \neq c$ be the end leaf in the path G_3 and let y be its unique neighbor (it is possible that $x = u_3$ and $y = c$). Let $z \neq c$ be the end leaf in the path G_2.

We construct G^* by removing the leaf x and then adding a leaf x^* adjacent to z. For $1 \leq k \leq d$, let $G_k^* = (V_k^*, E_k^*)$ be the path subgraphs of G^* defined and indexed analogously to the G_k subgraphs of G. The vertex c will still be the barycenter of G^* since $\pi^*(G_2^*) = \pi^*(G_2) + \frac{2}{2(n-1)} < \frac{1}{2}$. We have

$$\begin{aligned} & H^*(\pi^*, c) - H(\pi, c) \\ = \ & \sum_{k=1}^{d} \frac{|E_k^*|}{|E|} H^{*(k)}(\pi^{*(k)}, c) - \sum_{k=1}^{d} \frac{|E_k|}{|E|} H^{(k)}(\pi^{(k)}, c) \\ = \ & \left(\frac{|E_2^*|}{|E|} H^{*(2)}(\pi^{*(2)}, c) - \frac{|E_2|}{|E|} H^{(2)}(\pi^{(2)}, c) \right) + \left(\frac{|E_3^*|}{|E|} H^{*(3)}(\pi^{*(3)}, c) - \frac{|E_3|}{|E|} H^{(3)}(\pi^{(3)}, c) \right). \end{aligned}$$

Let $r = |V(G_2)|$ and $s = |V(G_3)|$ so that $r \geq s \geq 2$. We have $G_2 = P_r$, $G_2^* = P_{r+1}$, $G_3 = P_s$ and $G_3^* = P_{s-1}$. Using formula (11) (which is actually valid for $n \geq 2$) and simplifying, this becomes

$$H^*(\pi^*, c) - H(\pi, c) = \frac{1}{6(n-1)} \cdot \begin{cases} (2r-1)(6r+1) - (2s-3)(6s-5) & \text{for } s > 2, \\ (2r-1)(6r+1) - 3 & \text{for } s = 2. \end{cases}$$

This quantity is always positive, so if G is not a path then this tree is not the maximal structure for $\max_i H(\pi, i)$. ▮

5 Mixing measures

In this section, we prove Theorems 1.4 and 1.5.

Proof of Theorem 1.4. Let $z \in V$ achieve $H(z,\pi) = \max_i H(i,\pi)$. The vertex z must be a leaf. Indeed, suppose that z is not a leaf and let u be focus nearest to z. Let j be a leaf such that z is on the unique (j,u)-path. By equation (8),

$$H(j,\pi) = H(j,u) + H(u,\pi) = H(j,z) + H(z,u) + H(u,\pi) > H(z,u) + H(u,\pi) = H(z,\pi),$$

a contradiction.

We now prove the lower bound. If $G = S_n$ then $H(z,\pi) = 3/2$. Consider any tree $G \neq S_n$. Since z is a leaf, z is not a focus. Suppose the nearest focus to z is u. By equation (8), $H(z,\pi) = H(z,u) + H(u,\pi) \geq 1 + H(u,\pi)$, with equality holding if and only if z is a leaf adjacent to u. Since $G \neq S_n$, we have $\pi_u < 1/2$, so that $H(u,\pi) > 1/2$ since we must exit u with probability at least $1 - \pi_u$. Therefore $H(z,\pi) > 3/2$, proving that the star is the unique minimizing tree structure.

We turn to the upper bound. Given any tree G, suppose that the vertex u achieves $H(u,\pi) = \max_{i \in V} H(i,\pi)$. Clearly, $H(u,\pi) \leq \sum_k \pi_k H(u,k)$. By the random target identity (2) and Proposition 1.1(b), the latter quantity is at most $\frac{1}{6}(2n^2 - 4n + 3)$ and this upper bound is achieved uniquely by the path P_n. Furthermore, if z is an endpoint of the path P_n then $H(z,\pi) = \sum_k \pi_k H(z,k)$. Indeed, the other endpoint of P_n is a halting state for the naive rule "choose a target vertex k according to π and then walk until you reach k." The existence of a halting state guarantees that this rule is optimal, by equation (1). ∎

Proof of Theorem 1.5. By equation (4), we have $\mathcal{T}_{\text{reset}} = \mathcal{T}_{\text{forget}}$, so our proof will consider the forget time in place of the reset time.

As usual, the lower bound is easy to prove. Clearly the star has forget time equal to 1. Consider a tree $G \neq S_n$. If a is a primary focus of G then $d(a',a) \geq 2$ and therefore $H(a',a) \geq 4$. If G is focal then $\mathcal{T}_{\text{forget}} = H(a',a) \geq 4$. If G is bifocal with foci a, b, then $\mathcal{T}_{\text{forget}} = \mu_b H(a',b) + \mu_a H(a',a) \geq \mu_b + 4(1 - \mu_b) = 4 - 3\mu_b > 1$.

Considering the upper bound, suppose that G is focal with unique focus a, so that $\mathcal{T}_{\text{forget}} = H(a',a)$. Let the components of $G - a$ be $G_1, G_2, \ldots, G_d$ where $d = \deg(a)$ such that for $i = 1, 2$ we have $z_i \in V(G_i)$ and $H(a',a) = \max_{i \in V} H(i,a) = H(z_i,a)$. By Lemma 3.1,

$$H(a',a) \leq \min\{|V(G_1)|^2, |V(G_2)|^2\} \leq \left\lfloor \frac{1}{2}(n-1) \right\rfloor^2. \tag{15}$$

If n is odd, then equality holds if and only if G is a path. If n is even, then equality holds if and only if G is a path of length $n - 1$ with a single leaf attached to the central vertex.

Consider a bifocal tree G with foci a, b. We have $\mathcal{T}_{\text{forget}} = H(a', \mu) = H(b', \mu)$ by equation (9). Let $r = |V_{a:b}|$ and $n - r = |V_{b:a}|$, so that $H(a, b) = 2r - 1$, $H(b, a) = 2(n - r) - 1$. Suppose that n is odd. We have $\min\{r, n-r\} \leq (n-1)/2$. Therefore, $\mathcal{T}_{\text{forget}} < \min\{H(a', a), H(b', b)\} \leq \frac{1}{4}(n-1)^2$, so the odd path has a larger forget time than G. Therefore the odd path is the unique maximizing structure for odd trees.

Suppose that n is even. By Proposition 1.1(a), $H(b', a) \leq (r - 1)^2$ and $H(a', a) \leq (n - r)^2$. Using the formula from equation (9), we find that

$$\begin{aligned} \mathcal{T}_{\text{forget}} &= \frac{H(a,b)H(a',a) + H(b,a)H(b',a)}{2|E|} \leq \frac{(2r-1)(n-r)^2 + (2(n-r)-1)(r-1)^2}{2|E|} \\ &= \frac{1}{2|E|}\left(2(n-1)r(n-r) - (n-1)^2\right) = \frac{1}{2}\left(2r(n-r) - (n-1)\right) \end{aligned}$$

which is maximized for $r = n/2$. This value is attained uniquely by the even path, giving a value of $\frac{1}{4}(n^2 - 2n + 2)$. This value is strictly larger than the forget time of equation (15), so the even path is the unique maximizing structure for even trees. ▮

References

[1] D. Aldous and J. Fill. Reversible Markov Chains and Random Walks on Graphs (unpublished manuscript). `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[2] D. Aldous, L. Lovász and P. Winkler. Mixing times for uniformly ergodic Markov chains, *Stochastic Processes and their Applications*, 71 (1997), 165–185.

[3] A. Beveridge. Centers for random walks on trees. *SIAM J. Discrete Math* 23:1 (2009), 300-318.

[4] A. Beveridge and L. Lovász. Exit frequency matrices for finite Markov chains. Combinatorics, Probability and Computing, 19 (2010), 541–560.

[5] G. Brightwell and P. Winkler. Maximum hitting time for random walks on graphs. *J. Random Structures and Algorithms*, 1:3 (1990), 263–276.

[6] G. Brightwell and P. Winkler. Extremal cover time for random walks on trees. *J. Graph Theory*, 14:5 (1990), 547–554.

[7] D. Coppersmith, P. Tetali, and P. Winkler. Collisions among random walks on a graph. *SIAM J. on Discrete Mathematics*, 6 No. 3 (1993), 363–374.

[8] U. Feige. A tight lower bound on the cover time for random walks on graphs. *Random Structures and Algorithms*, 6:4, (1995), 433–438.

[9] O. Häggström. Finite Markov Chains and Algorithmic Applications. *Cambridge University Press*, 2002.

[10] L. Lovász. Random walks on graphs: a survey. *Combinatorics, Paul Erdős is Eighty.* Vol II (eds. D. Miklós, V.T. Sós and T. Szőnyi) J. Bolyai Math. Soc. (1996), 355–397.

[11] L. Lovász and P. Winkler. Efficient stopping rules for Markov chains. *Proceedings of the 27th ACM Symposium on the Theory of Computing* (1995), 76–82.

[12] L. Lovász and P. Winkler. Reversal of Markov chains and the forget time. *Combin., Probab. Comput.* 7 (1998), 189–204.

[13] J.W. Pitman. Occupation measures for Markov chains. *Adv. Appl. Prob.* 9 (1977), 69–86.

[14] O. Yaron. Random walks on trees. Technical report, Hebrew University, Jerusalem, (1988).