Simultaneous Reversible Markov Chains

Prasad Tetali[1] **and Peter Winkler**[2]

Abstract

Two tokens are placed on vertices of a connected, undirected, n-node graph G. Each token, when it moves, proceeds to any neighbor of its current location, each with equal probability. We bound the expected time T before the tokens meet.

In particular, we show that if the moves are simultaneous (that is, if both tokens move at each tick of the clock) then $T \leq 16n^3/27$; if only one token moves at a time, with the choices made at random, then $T \leq 4n^3/27$; if the choices are made so as to minimize T, then $T \leq 2n^3/27$. We give examples to show that lower bounds of order n^3 obtain in each case.

When G is provided with edge-weights, the random walks (proceeding to a neighbor with probability proportional to the connecting edge-weight) become general irreducible, reversible Markov chains. We provide a general formula for the expected time before two such chains enter the same state, under any move-strategy; further, we give a complete description of optimal strategy for minimizing T on a tree. Most of our results rely on our boundary value theorem for a class of functions on graphs which we call "biharmonic".

1 Introduction

Random walks on graphs have proved to be useful tools in several aspects of the theory of computing, as well as an elegant subject for mathematical analysis. Most notable, perhaps, of recent applications has been to polynomial-time randomized approximation algorithms e.g. [13, 8, 10], but there have also been uses in space complexity [2], on-line algorithms [6], and distributed computing [12].

Three critical parameters of random walks, closely related, are the *cover time*, the *mixing time* and the *meeting time*; all three are also tied to eigenvalues and expansion of graphs. The first measures the expected time for all vertices (states) to be covered; the second, for the state distribution to be near stationary; and the third, for two walks to reach the same state. Bounds on cover time were the subject of several articles in a single issue of the *Journal of Theoretical Probability* [14] and other works such as [4] and [5]. Mixing times have been considered in numerous recent articles such as [13, 10].

Meeting times were studied by Aldous [1] who bounds meeting time and mixing times, each in terms of the other. In [7] worst-case meeting times are tightly bounded for the case where a "schedule demon" decides which chain moves next; this work solved a distributed computing problem of Israeli and Jalfon [12] by providing a polynomial-time bound for their token management protocol to self-stabilize.

[1] AT & T Bell Labs, Murray Hill, NJ 07974
[2] Bellcore, Morristown, NJ 07962

In this paper we consider three natural alternatives to the schedule demon for move selection, each with its own meeting time. In each case we are concerned with bounding the meeting time relative to the number of vertices in G, and/or to the maximum over pairs (u, v) of vertices in G of the expected time for a random walk to proceed from u to v. We also look for ways to compute the meeting time in certain situations. Here are the three "rules":

SIMULTANEOUS MOVES: Both tokens move at the same time, at each step of the clock. Meeting occurs only when tokens reach the same vertex at time t, not, for example, if they pass each other on an edge.

RANDOM MOVES: At each tick of the clock, a fair coin is flipped to determine which token moves. This is equivalent to the case of two $continuous-time$ random walks, e.g. as studied by Aldous [1].

ANGEL MOVES: At each tick of the clock, an angel chooses which token to move, with the intent of minimizing the expected time before the tokens meet.

The maximum expected meeting time of (essentially) $4n^3/27$ for the *demon*-moves case, where move selection is made so as to maximize meeting time, was established in [7]. This quantity thus serves as an upper bound for the random and angel cases but not for the simultaneous case, which behaves quite differently and requires a non-bipartite graph.

Our main results are presented in the following order. In Section 2, we introduce the notion of biharmonic functions to prove a general formula for the expected meeting time under any move-strategy. In Section 3, we provide bounds ($\Theta(n^3)$) for each expected meeting time. In Section 4, we show that the expected meeting times can, in principle, be computed in polynomial time; however, it turns out that figuring out the expected meeting times on specific classes of graphs can be quite tricky. We work out some examples in Sections 3, 5, and 6 including a complete characterization of angel's strategy on a tree.

We conclude with several open problems, including some questions regarding analogous expected meeting times with more than a constant number of reversible Markov chains.

2 General Results

2.1 Preliminaries

We introduce some preliminary material that is handy in proving our main results.

Let G be a connected graph on n vertices. (Initially we will assume the edges of G to be unweighted, but our later results generalize.) The *hitting time* $H(x, y)$ from x to y is defined to be the expected time for a random walk on G beginning at x to reach y for the first time.

Although hitting time is in general a highly asymmetric function, it is torsion-free in the sense that for any three vertices x, y, z,

$$H(x,y) + H(y,w) + H(w,x) = H(x,w) + H(w,y) + H(y,x) .$$

This useful fact, which we call the "triangle rule," was proved in Lemma 2 of [7] and in fact turns out to be equivalent to reversibility for Markov chains [17]. The triangle rule implies that the relation $\leq$ defined by $x \leq y$ iff $H(x,y) \leq H(y,x)$ is transitive. In particular, there is always a "remote" vertex r such that $H(r,x) \leq H(x,r)$, for all $x \in G$. Analogously, there is a "central" vertex c such that $H(x,c) \leq H(c,x)$, for all $x \in G$. The intuition is that a remote (central) vertex is always harder (resp. easier) to get to than to return from.

We define the *potential* $Z(x)$ of a vertex x by

$$Z(x) = H(c,x) - H(x,c).$$

Note that $Z(x) \geq 0$. Moreover, a central vertex has potential zero, and a remote vertex has the maximum potential $Z(r)$. Some words of caution regarding our terminology are appropriate here: $Z(x)$ is not an electrical potential in the usual sense, nor does it signify attractive power for some vertices of G. It *is* the case that tokens take longer, in general, to reach points of higher potential than lower potential; and we can make this statement more precise in an electrical setting.

Regarding G as an electrical circuit with a unit resistor on each edge, we have the following hitting time formula of [16]:

$$H(x,y) = mR(x,y) + \frac{1}{2}\sum_w d(w)[R(w,y) - R(w,x)].$$

Here $R(x,y)$ denotes the effective resistance between x and y, m is the number of edges of G, and $d(w)$ is the valence of w.

(If the edges of G are instead weighted by positive reals $w(e)$, so that the random walk beomes a general reversible Markov chain, the weights are interpreted as conductances; that is, edge e is provided with resistance $1/w(e)$ instead of unit resistance.)

Letting $C(x,y) = H(x,y) + H(y,x)$ denote the "commute time" (i.e. expected round trip time) between x and y, we have the following interpretation:

$$\begin{aligned} Z(x) &= \sum_w d(w)[R(w,x) - R(w,c)] \\ \text{and } H(x,y) &= \frac{1}{2}(C(x,y) + Z(y) - Z(x)). \end{aligned}$$

In [7], we introduced the following function $\phi(x,y)$ to bound the expected meeting time under the demon's strategy. We call ϕ the "demon's harmonic", wherein the justification becomes apparent following our main theorem. Let ϕ be defined as follows:

$$\begin{aligned} \phi(x,y) &= H(x,y) + H(y,r) - H(r,y) \\ &= \frac{1}{2}(C(x,y) - Z(x) - Z(y)) + Z(r) \end{aligned}$$

Similarly, we define the "angel's harmonic" to be

$$\begin{aligned} \psi(x,y) &= H(x,y) + H(y,c) - H(c,y) \\ &= \frac{1}{2}(C(x,y) - Z(x) - Z(y)). \end{aligned}$$

At the outset, we have angel's meeting time $M_{\text{angel}}(x,y) \leq \min\{H(x,y), H(y,x)\}$, and demon's meeting time $M_{\text{demon}}(x,y) \geq \max\{H(x,y), H(y,x)\}$, since the controller can always fix one token and achieve either hitting time. Thus it is of interest to see how much better the controller can do by being willing to change horses in midstream.

In fact, we will show that the meeting time under *any* strategy is bounded between the demon's and angel's harmonics (see Corollary 2 below).

2.2 General Strategies

We prove the following very general theorem for all strategies S which move one token at a time. A "strategy" for us will be an assignment of a probability p to each pair (x,y) of vertices such that with one token at x and the other at y, S moves the token at x with probability p. (The "random moves" rule then corresponds to a strategy in which p is always $\frac{1}{2}$.) We will use $M_S(x,y)$ to denote the expected meeting time of tokens at x and y, when the controller follows S.

Let $q_S(z;x,y)$ denote the probability that under strategy S, tokens initially positioned at x and y first meet at vertex z. Then

Theorem 1

$$M_S(x,y) = \phi(x,y) - \sum_z q_S(z;x,y)[Z(r) - Z(z)]$$

Theorem 1 often helps to compute $M_S(x,y)$, by reducing the "time" question to a "space" question. As we shall see, the probabilities $q_S(z;x,y)$ are sometimes easy to determine.

As a prelude to proving Theorem 1 we introduce the notion of "biharmonic" functions. Let $g = g(x_1,\ldots,x_k)$ be a real-valued function whose arguments are vertices of G. It will be convenient to use the shorthand notation

$$g(x_1,\ldots,x_{i-1},\bar{x}_i,x_{i+1},\ldots,x_k) = \frac{1}{d(x_i)}\sum_{z\sim x_i} g(x_1,\ldots,x_{i-1},z,x_{i+1},\ldots,x_k).$$

A function $f: V \times V \to \Re$ is said to be **biharmonic** if for all x,y,

$$\min(f(\bar{x},y), f(x,\bar{y})) \leq f(x,y) \leq \max(f(\bar{x},y), f(x,\bar{y})).$$

We show that biharmonic functions satisfy the following Maximum Principle, much in the spirit of the usual (discrete or continuous) harmonic functions.

Maximum Principle. If f is biharmonic then $\max_{x,y} f(x,y) = f(z,z)$ for some z.

Proof. Suppose not. Let $f_M = \max_{x,y} f(x,y)$ and choose vertices u,v of minimum distance (in the path metric on G) such that $f_M = f(u,v)$. Either $f(u,v) \leq f(\bar{u},v)$ or $f(u,v) \leq f(u,\bar{v})$; we may assume the former. But then

$$f(u,v) \quad \leq \quad f(\bar{u},v)$$

$$
\begin{aligned}
&= \frac{1}{d(u)} \sum_{w \sim u} f(w, v) \\
&\leq \frac{1}{d(u)} f(w', v) + \frac{d(u)-1}{d(u)} f_M \\
&< f_M
\end{aligned}
$$

where w' is a neighbor of u which is closer to v than u is; the contradiction concludes the proof. □

Of course the corresponding **Minimum Principle** also holds, by symmetry. Now let $\chi_S(x,y) = M_S(x,y) - \phi(x,y) + \sum_z q_S(z;x,y)[Z(r) - Z(z)]$.

Lemma 1 *χ_S is biharmonic.*

Proof. Let the tokens be placed at x and y. If p is the probability that the strategy S moves the token at x, then we claim that

$$\chi(x,y) = p\chi(\bar{x}, y) + (1-p)\chi(x, \bar{y})$$

For, clearly $M_S(x,y) = 1 + pM_S(\bar{x}, y) + (1-p)M_S(x,\bar{y})$. Moreover

$$
\begin{aligned}
\phi(x,y) &= 1 + \phi(\bar{x}, y) \\
&= 1 + \phi(x, \bar{y}), \quad \text{by the symmetry of } \phi \\
\text{So } \phi(x,y) &= 1 + p\phi(\bar{x}, y) + (1-p)\phi(x, \bar{y})
\end{aligned}
$$

And finally,

$$q_S(z;x,y) = p \sum_{u \sim x} \frac{1}{d(x)} q_S(z;u,y) + (1-p) \sum_{v \sim y} \frac{1}{d(y)} q_S(z;x,v) \ .$$

With slight abuse of notation we could write this as

$$q_S(z;x,y) = pq_S(z;\bar{x},y) + (1-p)q_S(z;x,\bar{y}) \ .$$

Now it is easy to verify our claim:

$$
\begin{aligned}
\chi(x,y) &= M_S(x,y) - \phi(x,y) + \sum_z q_S(z;x,y)[Z(r) - Z(z)] \\
&= p\left(M_S(\bar{x},y) - \phi(\bar{x},y) + \sum_z q_S(z;\bar{x},y)[Z(r) - Z(z)]\right) \\
&\quad +(1-p)\left(M_S(x,\bar{y}) - \phi(x,\bar{y}) + \sum_z q_S(z;x,\bar{y})[Z(r) - Z(z)]\right) \\
&= p\chi(\bar{x},y) + (1-p)\chi(x,\bar{y}) \ .
\end{aligned}
$$

□

Proof of Theorem 1. Note that

$$
\begin{aligned}
\chi_S(x,x) &= 0 - \phi(x,x) + \sum_z q_S(z;x,x)[Z(r) - Z(z)] \\
&= -(Z(r) - Z(x)) + [Z(r) - Z(x)] \quad (\text{since } q_S(x;x,x) = 1) \\
&= 0 \ .
\end{aligned}
$$

By the Maximum and Minimum principles, χ is identically zero, proving the theorem. $\square$

Corollary 1 *(a)* $M_{\text{angel}}(x,y) = \psi(x,y) + \sum_z q_{\text{angel}}(z;x,y)Z(z)$
(b) $M_{\text{demon}}(x,y) = \phi(x,y) - Z(r) + \sum_z q_{\text{demon}}(z;x,y)Z(z)$.

Proof. For any strategy S that moves only one token at a time, $M_S(x,y)$ is bounded between the angel's meeting time and the demon's meeting time. Now the corollary is obvious from Theorem 1 and by noticing that $\phi(x,y) - Z(r) = \psi(x,y)$.

Corollary 2 $\psi(x,y) \leq M_S(x,y) \leq \phi(x,y)$.

Proof. Follows from Theorem 1 and Corollary 1. However, note that ϕ and ψ are also biharmonic functions (proofs similar to the proof of Lemma 1); ϕ is non-negative on the diagonal and ψ non-positive, so that Corollary 2 follows directly from the Maximum Principle.

Corollary 3 *The angel need never move a token from a central vertex, and the demon need never move a token from a remote vertex.*

Proof. Note that $H(x,c) = \psi(x,c) \leq M_{\text{angel}}(x,y) \leq H(x,c)$. Thus the angel may as well let the token on c stay put. The argument for the demon is similar.

3 Bounds on Meeting Times

In this section we consider only the uniformly-weighted case, so that we can obtain bounds strictly in terms of the number of vertices n of G. However, the results may be generalized by limiting the ratio of largest to smallest edge-weight.

Let us note first that the demon's harmonic, and therefore the demon's, angel's and random meeting time, are all bounded by twice the maximum hitting time H_{max} of G; this (as noted in [7]) already improves Aldous' result [1] bounding only the random meeting time by a rather large constant multiple of H_{max}.

Since H_{max} itself is bounded by $\frac{4}{27}n^3$ plus lower-order terms [3] we obtain a cubic bound on the three meeting times, but the resulting constant $8/27$ can be improved. Indeed, from [7] we have that

$$C(x,y) + C(y,z) + C(z,x) \leq (\frac{8}{27} + o(1))n^3$$

for any three points x, y and z; since this expression is twice the demon's harmonic $\phi(x,y)$ when $z = r$, the upper bound for meeting time drops to $\frac{4}{27}n^3$. Moreover, putting $y = z$ bounds the *commute* time of G by $\frac{4}{27}n^3$, thus the angel's meeting time is bounded by $\frac{2}{27}n^3$ since she can fix one token and achieve the lower of the two hitting times.

The simultaneous moves case is a wholly different matter. Here we must assume to begin with that G is not bipartite, else tokens which start on opposite sides will never meet. Even then we cannot achieve a bound which is a constant multiple of $H_{\max}$ on account of the following example: let G^- be the complete bipartite graph with parts $\{u_1, \ldots, u_{n/2}\}$ and $\{v_1, \ldots, v_{n/2}\}$. Let G be G^- with the single additional edge $\{u_2, u_3\}$, and put tokens at u_1 and v_1. We have $H_{\max}(G) \sim H_{\max}(G^-) \sim n$, but for the tokens to meet in G, one of them must reach u_2 or u_3 and then traverse the new edge. It is not difficult to see that even with two tokens "on the move" this takes expected time of order n^2.

It is nonetheless possible to give an order n^3 bound for the simultaneous meeting time, assuming G is not bipartite. Let $X(t)$ and $Y(t)$ denote the positions of the two tokens at time t. Let $M_{\rm sim}(x,y)$ denote the expected time before the tokens meet, beginning from positions x and y; that is,

$$M_{\rm sim}(x,y) = \mathbf{E}[\ \min t : X(t) = Y(t) \mid X(0) = x; Y(0) = y]\ .$$

We also define the "near meeting time" $N(x,y)$ as follows:

$$N(x,y) = \mathbf{E}[\ \min t :\ (X(t) = Y(t) \text{ or } X(t) = Y(t-1)) \mid X(0) = x; Y(0) = y]\ .$$

We first show that the near meeting time is bounded by $(2/27 + \mathrm{o}(1))n^3$. To see this, suppose that the tokens move not simultaneously but alternately, first the token at x, then the one at y, etc. Their meeting time will then be no more than the demon's meeting time, which is bounded by $(\frac{4}{27} + \mathrm{o}(1))n^3$. Fixing the random sequence of vertices occupied by each token, we have that $X(t) = Y(t)$ iff the alternating tokens meet at an "even" move $2t$ and $X(t) = Y(t-1)$ if they meet at move $2t-1$.

We now reduce the simultaneous meeting time on G to the near meeting time on a different graph B. For the vertices of B we take two copies of $V(G)$; let us say x_0 and x_1 in B correspond to the vertex x of G. For each edge $\{x,y\}$ in G, we create two edges $\{x_0, y_1\}$ and $\{x_1, y_0\}$ in B. It is immediate that there is a canonical, one-to-one, probability-preserving correspondence between random walks beginning at x in G, random walks beginning at x_0 in B, and random walks beginning at x_1 in B.

Note that because G is not bipartite, B is connected. If we begin simultaneous random walks in B at x_0 and y_0, then their near meeting time is bounded by $\frac{2}{27}(2n)^3 = \frac{16}{27}n^3$. But we cannot have $X(t) = Y(t-1)$ because in the bipartite graph B, $X(t)$ and $Y(t-1)$ will never have the same subscript. Hence the simultaneous meeting time in B, and thus also in G, is equal to the near meeting time in B.

What about lower bounds for worst-case meeting times? For the demon, the lollipop of [3] shows that $\frac{4}{27}n^3$ is essentially correct (this graph consists of a clique on $2n/3$ vertices, including x, attached to a path on $n/3$ vertices ending in y). However, the random, angel and simultaneous meeting times for this graph are all of order only n^2.

To foil the angel we instead employ a "barbell" graph consisting of two cliques on $n/3$ vertices each, connected by a path on the remaining $n/3$ vertices. The midpoint of the path will be a central vertex of this graph, hence the angle's best strategy, confronted with one token in each clique, is to move each until it reaches the central vertex; this takes time $2\frac{4}{27}(n/2)^3 = \frac{1}{27}n^3$. Interestingly, random moves is only marginally worse for this graph.

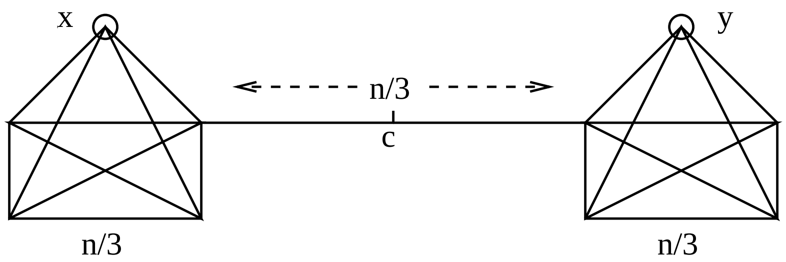

Figure 1:

For the simultaneous case we modify the lollipop, replacing the clique by a balanced complete bipartite graph on $2n/3$ vertices, and ending the attached path with a triangle. Now one of the tokens must reach the triangle and traverse it an odd number of times before returning to the fat end (or hitting the other token en route). The near end of the path is hit by one token or the other about one time in $n/3$, from which the token will proceed to the next point in the path with probability $3/n$ and then accomplish its odd traversal with probability $3/2n$. Altogether this produces a meeting time of about $\frac{n}{3} \cdot \frac{n}{3} \cdot \frac{2n}{3} = \frac{2}{27}n^3$.

We put all this together as follows.

Theorem 2 *Let M_{demon} be the maximum demon's meeting time over all n-vertex graphs G and all starting positions x, y; similarly for M_{demon}, M_{rand} and M_{sim}. Then*

$$(\frac{4}{27} + o(1))n^3 \le M_{\mathrm{demon}} \le (\frac{4}{27} + o(1))n^3 \ ;$$

$$(\frac{1}{27} + o(1))n^3 \le M_{\mathrm{angel}} \le (\frac{4}{27} + o(1))n^3 \ ;$$

$$(\frac{1}{27} + o(1))n^3 \le M_{\mathrm{rand}} \le (\frac{4}{27} + o(1))n^3 \ ;$$

$$(\frac{2}{27} + o(1))n^3 \le M_{\mathrm{sim}} \le (\frac{16}{27} + o(1))n^3 \ .$$

4 Algorithmic Issues

It is easy to see that the meeting times for the random moves game may be computed in time polynomial in n; one can simply solve the linear system

$$M(x, x) = 0 \ ;$$

$$M(u, v) = 1 + \frac{1}{2}\sum_{u' \sim u} \frac{1}{d(u)} M(u', v) + \frac{1}{2}\sum_{v' \sim v} \frac{1}{d(v)} M(u, v') \ .$$

The simultaneous moves case is equally easy, but for the angel and demon there appear to be $2^{\binom{n}{2}}$ strategies to consider. Nonetheless both cases reduce to linear programming problems, and are thus also computable in polynomial time. In the angel's case we *maximize* $M(u,v)$ subject to the following constraints:

$$M(x,x) = 0\ ;$$
$$M(u,v) \le 1 + \sum_{u' \sim u} \frac{1}{d(u)} M(u',v)\ ;$$
$$M(u,v) \le 1 + \sum_{v' \sim v} \frac{1}{d(v)} M(u,v')\ .$$

The formulation in the case of unequal edge-weights is a straightforward generalization.

To see that this optimization problem (and its mate in the demon's case) actually does produce the angel's meeting time, note that maximization forces one of the two inequalities to become an equality, assuring that $M(u,v)$ is a meeting time for *some* strategy; the remaining inequality guarantees that the strategy moves the token which produces the smaller of the two meeting times. (This result may also be derived from techniques in the theory of *Markovian decision processes* as in e.g. [9].)

Before proceeding, let us remark that the linear programming approach does not necessarily help us determine the angel's and demon's strategies, and resulting meeting times, for a parameterized *class* of graphs. We look at some examples below.

5 Examples and Counterexamples

5.1 Angel vs Demon

Corollary 3 suggests that the angel's and demon's strategies, as well as their goals, are diametrically opposite. Nonetheless situations can arise in which they must both move the same token.

Consider the following graph (see Fig. 2) with tokens on c and d.

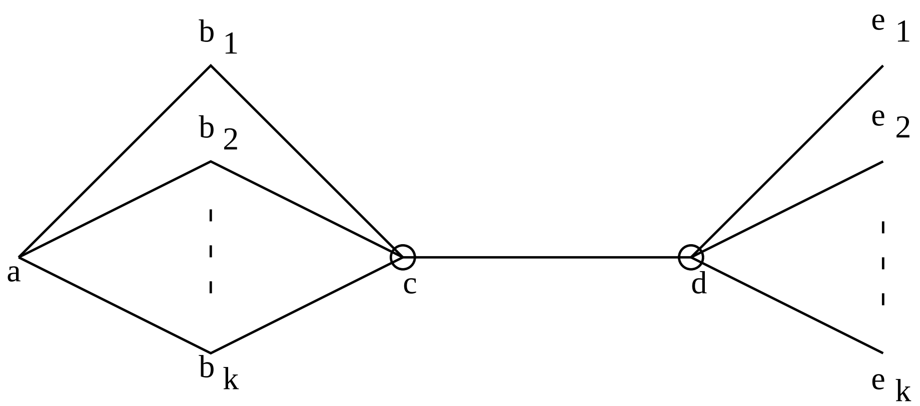

Figure 2:

Let $k \ge 4$. It is easy to see that c is the unique central vertex, and $e_1, \ldots, e_k$ are the remote vertices of the graph. From Corollary 3 we know that the angel stays forever on c, and in particular moves the token on d at this time. We claim that the demon's best strategy is also to move from d—and never to move that token again.

To see this, let W be the probability space each member of which is a sequence of positions beginning with the tokens at c and d, and ending with collision. Obviously W depends on the scheduling strategy; let W_c be the result of moving c whenever the tokens are on c and d, and otherwise maximizing meeting time; and similarly for W_d. We demonstrate a one-to-one probability-preserving correspondence between W_c and W_d for which the meeting time is always equal or greater in W_d; it follows that moving the token at D must have been the demon's correct choice. We argue case by case.

The token on c moving to d (in W_c) corresponds to the token on d moving to c, in which case the meeting time is just 1. Movement from c to b_1 corresponds to movement from d to e_1; in the former case the demon next moves from d, and in the latter from c. Thus the two situations – c to b_1 followed by d to e_2 (say) and d to e_1 followed by c to b_2 – correspond and have the same meeting time.

This leaves us with: (1) c to b_1 followed by d to c versus (2) d to e_1 followed by c to d. We either meet, go back to the previous position, or proceed to: (3) tokens at b_1 and b_2 versus (4) tokens at e_1 and e_2. Now we either go back to (1) versus (2) or proceed to (5) tokens at a and b_1 versus (2). So it suffices to show that the demon prefers (2) to (5). This is obvious if in situation (5) the demon wants to move off a. The proof is complete by noticing that for k large enough ($k \geq 4$ suffices), the demon always moves from a when in situation (5). □

5.2 The Path

Strategies and meeting time for the demon and angel on a path are easy to compute. Since the angel never moves from a central vertex she may move either token when they are on opposite sides of the central vertex; otherwise she moves the token closest to the end. As established in [7] the demon merely makes sure that he never moves a token from an endpoint unless he has to. The random mover has no strategy decision to make, but computing the meeting time in this case turns out to be relatively tricky.

Suppose, for example, that G is the path of length n on vertices 0,1,...,n, with token X on vertex 0 and token Y on vertex 1. Obviously the angel moves X to achieve meeting time 1, while the demon moves Y throughout to achieve meeting time $H(1,0) = 2n-1$.

It follows that $M_{\text{rand}}(0,1)$ is somewhere between constant and linear, but where? Consider the following problem on the same path: U and V are two tokens, both at vertex 1, which are performing random walks according to the *simultaneous* moves rule. We ignore collisions of U and V; they are welcome to occupy the same vertex at any time. Let T be the expected minimum number of steps before U or V reaches vertex 0; we claim that $T = M_{\text{rand}}$. In fact, the numbers of steps in each case have the same distribution.

To see this, extend G to a ring of length $2n$ as shown below, with vertices $0, 1, \ldots, n-1, n = n^*, n^*, (n-1)^*, \ldots, 2^*, 1^*$ and back to $0 = 0^*$, reading clockwise. When token X or Y takes a random step we may imagine that it steps clockwise or counterclockwise with equal probability, but we have to consider that the tokens have collided if one is at i and the other at i^* for some i (as well as if they land on the same vertex).

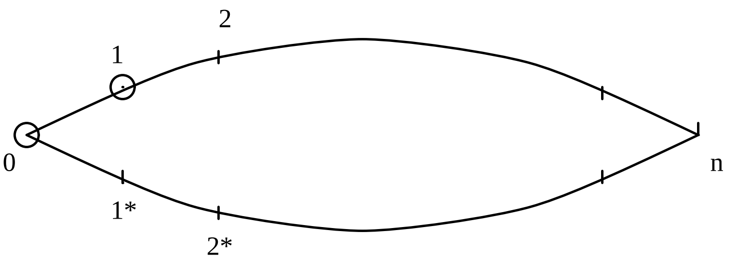

Figure 3:

At any time t let $U(t) = d_{\text{ring}}(X(t), Y(t))$ be the length of the shortest path from X's position to Y's *in the ring*, and $V(t) = d_{\text{ring}}(X(t), (Y(t))^*)$. Then it is easy to check that U and V behave as described above, with collision occurring when either hits 0.

Since the precise distribution of the number of steps for (say) U to reach 0 is known (see e.g. Feller [11]), T can be expressed as a sum whose value turns out to be of order $\log n$.

□

The above example illustrates the fact that computing the expected meeting time under random moves tends to be, in general, quite difficult. Thus it would be interesting to find better upper bounds for M_{rand}. It might seem reasonable to guess that random meeting time is bounded above by commute time (this is true, for example, for Cayley graphs). However, the following counterexample [15] shows that in general $C(x, y)$ can exceed $M_{\text{rand}}(x, y)$ by an order of magnitude.

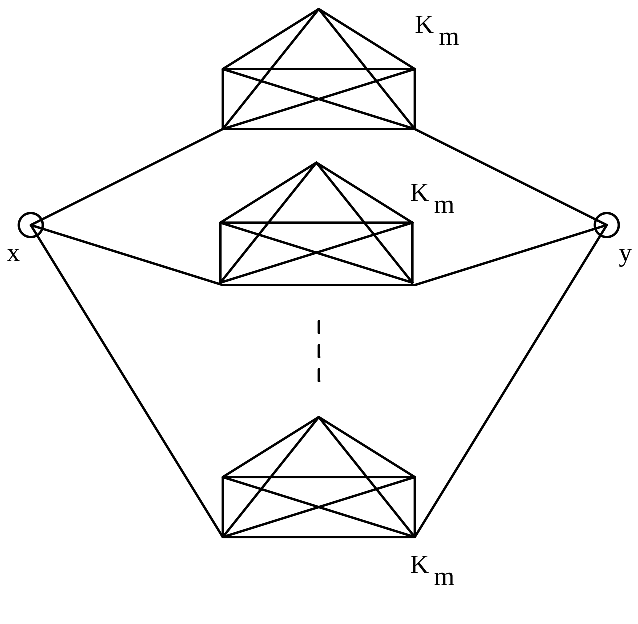

The vertices x and y are connected to l copies of a clique of size m. For ease of calculation imagine l tending to ∞ . As $l \to \infty$, the commute time between x and y is $C(x, y) \sim 2m^2$, while the random meeting time $M_{rand}(x,y) \sim \frac{1}{2}m^4$. Thus, for l large enough we get a counterexample.

6 Trees

Let c be a central vertex of the tree T, and let x and y be vertices in the same component of $T - \{c\}$. We define the "join" of x and y with respect to c to be the (unique) vertex which lies in the intersection of the paths from x to c, y to c and x to y. Intuition suggests that the angel's strategy should be to make the tokens meet at their join, and for a change, intuition is correct. As a result we obtain a complete characterization of the angel's strategies on a tree.

Theorem 3 *The "angel time" $M_{\text{angel}}(x,y) = H(x,w) + H(y,w)$, where w is the join of x and y with respect to any central vertex c.*

Proof. Suppose not. Then let

$$m = \max_{x,y}(H(x,w) + H(y,w) - M_{\text{angel}}(x,y)).$$

Further, let x and y be the closest pair that achieves this. i.e. among all pairs that realize m, $\{x,y\}$ is the (not necessarily unique) pair with minimum distance $d(x,y)$ between them. Note that we can assume that x and y are distinct, and different from c (otherwise, the theorem is already known to be true by Corollary 3.)

Suppose first that the angel moves a token which is not on w, so that the join w does not change. We may assume the token at x is the one which is moved, in which case

$$\begin{aligned} M_{\text{angel}}(x,y) &= 1 + M_{\text{angel}}(\bar{x},y) \\ &> 1 + H(\bar{x},w) + H(y,w) - m \\ &= H(x,w) + H(y,w) - m = M_{\text{angel}}(x,y)\ , \end{aligned}$$

a contradiction.

Thus we are left with the case where $y = w$ and the angel chooses to move the token at y. We make use of the following trivial observation: if a, b, and a' are three vertices of a tree with a and a' both adjacent to b, then $H(a,b) < H(b,a')$. To see this let T_a be the subtree containing a when the edge $\{b,a'\}$ is removed, and T_b be $T_a \cup \{b,a'\}$. Then $H(a,b)$ is one less than the commute time between a and b in T_a, and $H(b,a')$ is one less than the commute time between b and a' in T_b. Since T_b is bigger than T_a the observation follows.

Now let $y_1, y_2, ..., y_d$ be the neighbors of y with y_1 closest to x and y_2 closest to the central vertex c. Then

$$\begin{aligned} M_{\text{angel}}(x,y) &= 1 + M_{\text{angel}}(x,\bar{y}) \\ &> 1 + (1/d)[(H(x,y_1) - m) + (H(x,y_2) - m) \\ &\quad +(H(x,y) + H(y_3,y) - m) + ... + (H(x,y) + H(y_d,y) - m)] \end{aligned}$$

$$
\begin{aligned}
&> \quad 1 + (1/d)[H(x, y_1) + H(x, y_2) + (d-2)H(x, y)] - m \\
&= \quad 1 + (1/d)[(H(x, y) - H(y_1, y)) + (H(x, y) + H(y, y_2)) + (d-2)H(x, y)] - m \\
&> \quad 1 + H(x, y) - m = 1 + M_{\text{angel}}(x, y)
\end{aligned}
$$

once again, a contradiction. In the last inequality we made use of the observation above, with y_1, y and y_2 playing the role of a, b and a'. □

It is clear from the theorem that the angel has many equivalent, optimal strategies (including probabilistic ones) but that a strategy is optimal for her if and only if every move is consistent with the theorem's "join rule" and the central point rule of Corollary 3, for some choice of central point. It is perhaps worth commenting on one particular intuitive strategy, however.

The "hitting time strategy" for an angel calls for moving the token at x whenever $H(x, y) < H(y, x)$ and the other token is at y; equivalently, the vertex of higher potential is vacated. The opposite strategy is *not* necessarily optimal for the demon on a tree—an example is given in [7]. Here, we have

Corollary 4 *The hitting time strategy is optimal on a tree for the angel.*

Proof. In view of the above theorem, it suffices to prove that potential is always increasing towards the leaves, and is unimodal.

Let x be any vertex on a tree, and let y be a neighbor of x. Then it is easy to verify that

$$Z(y) = Z(x) + 2n - 4t_y$$

where t_y is the order of the subtree containing y but not x. With this it is easy to show that the potential of the neighborhood of x always increases as long as x is not a leaf. Let d be the degree of x. Then

$$
\begin{aligned}
Z(\bar{x}) &= \frac{1}{d} \sum_{y \sim x} Z(y) \\
&= \frac{1}{d} \sum_{y \sim x} [Z(x) + 2n - 4t_y] \\
&= Z(x) + 2n - \frac{4}{d} \sum_{y \sim x} [t_y] \\
&= Z(x) + 2n - \frac{4}{d}(n-1)
\end{aligned}
$$

Thus for $d \geq 2$ (i.e. x not a leaf), $Z(\bar{x}) \geq Z(x) + 2$.

The proof is completed by noticing that moving the higher potential point corresponds to meeting on the join with respect to a central vertex. □

7 Open Problems

Of the many questions that have arisen in this work, we list a few of the more hopeful.

1. Since the (reversed) hitting time strategy for the demon fails on a tree, it remains an interesting question to find a simple strategy for the demon on a tree. We do suspect that it is never right for the demon to move from a leaf if he is not forced to do so.

2. Is the hitting time strategy optimal for the angel, on a general graph? Perhaps not, but we do not know of a counterexample.

3. Not having any better examples, we expect that all the lower bounds given in Theorem 2 are also upper bounds.

4. Many questions remain in the case where there are more than two tokens. In particular, with a token on *every* vertex, and tokens coalescing as they collide, our previous conjecture [7] that $M_{\text{demon}} = \mathrm{O}(n^3)$ remains open. This question is unresolved even with random moves as well (see Conjecture 3 of [1]); but note that the angel can always move one token and scoop up all the others in cover time $= \mathrm{O}(n^3)$. In case of simultaneous moves, it is easy to see that $\mathrm{O}(n^3 \log n)$ is possible, by considering meeting of pairs of tokens simultaneously; however, the correct bound may be $\mathrm{O}(n^3)$.

5. The k-token case also strains our complexity observations since the number of variables in our optimization problem is order n^k. So, under the various move-rules we consider, can meeting time of k tokens—with k part of the input—be solved in polynomial time?

References

[1] D.J. Aldous, Meeting times for independent Markov chains, *Stochastic Processes and Applications* (1991), to appear.

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

[3] G. Brightwell and P. Winkler, Maximum hitting time for random walks on graphs, *J. Random Structures and Algorithms* **1** No. 3 (1990), 263-276.

[4] G. Brightwell and P. Winkler, Extremal cover time for random walks on trees, *J. Graph Theory* **14** No. 5 (1990), 547-554.

[5] A. Condon and D. Hernek, Random Walks in Colored Graphs, preprint.

[6] D. Coppersmith, P. Doyle, P. Raghavan and M. Snir, Random walks on weighted graphs, and applications to on-line algorithms, *Proc. 22nd Annual ACM Symposium on Theory of Computing* (1990), 369-378.

[7] D. Coppersmith, P. Tetali, and P. Winkler, Collisions among Random Walks on a Graph, *SIAM J. on Discrete Mathematics*, (to appear).

[8] P. Dagum, M. Luby, M. Mihail, and U. Vazirani, Polytopes, permanents and graphs with large factors, *29th Annual IEEE Symposium on Foundations of Computer Science* (1989), 412-421.

[9] C. Derman, *Finite State Markovian Decision Processes*, Academic Press, New York NY 1970.

[10] M. Dyer, A. Frieze and R. Kannan, A random polynomial time algorithm for estimating volumes of convex bodies, *Proc. 21st Annual ACM Symposium on the Theory of Computing* (1989), 375-381.

[11] W. Feller, *An Introduction to Probability Theory and its Applications* Volume I, John Wiley and Sons, New York 1950.

[12] A. Israeli and M. Jalfon, Token management schemes and random walks yield self stabilizing mutual exclusion, *Proc. 9th Annual ACM Symposium on Principles of Distributed Computing*, Quebec City, Canada (1990), 119-131.

[13] M. Jerrum and A. Sinclair, Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved, *Proc. 20nd Annual ACM Symposium on Theory of Computing* (1988), 235-243.

[14] *Journal of Theoretical Probability*, Vol. **2** (1989).

[15] L.A. Shepp, private communication.

[16] P. Tetali, Random walks and effective resistance of networks, *J. Theoretical Prob.* No. 1 (1991), 101-109.

[17] P. Tetali, An extension of Foster's network theorem, preprint.