RANDOM WALK MEETING TIMES ON GRAPHS

ETHAN NAEGELE

Abstract. We discuss the bounds on expected meeting times for a game involving two tokens which pass from vertex to vertex randomly on an arbitrary n-vertex graph. At each unit of time, the player chooses which token to move, and the game ends when the two tokens collide. We find the coefficients of the approximate bounds for the meeting times, which are of order n^3, and the player strategies which give rise to such extremal meeting times.

Contents

1. Introduction

We investigate a game played on some connected, undirected graph G. The game to be discussed is played as follows: two tokens are placed at initial vertices (x, y) on G. At each unit of time (say, each second), the player chooses which of the two tokens is to be moved. The selected token moves according to the uniform distribution on its neighbors. When both tokens meet at the same vertex, the game ends. In this paper, we investigate the expected meeting times for each of the three main modes of play: Angel, Demon, and Random. Angel plays to minimize the meeting time; Demon plays to maximize it; Random flips a fair coin to select which token to move.

The motivation for studying such games is well summarized in [1]. These games are useful in understanding "token-management" schemes, which pass a "token," some abstract object between nodes, or "processors" in a network. In the desired state, there is generally one token in the network. Were the system to enter some undesired state in which the network contains multiple tokens, the way to return to the desired state would be to remove the tokens by colliding them together at the same processor. It becomes of our interest, then, to study the maximum amount of time we should expect to wait in the worst case scenario, described by the Demon player. We would also like to find how much shorter we should expect to wait in the worst case scenario when we are trying to end the undesired state as soon as possible, i.e., the Angel player. Naturally, we also ask what the strategies are which result in these extremal Demon and Angel games, so that we may implement them, and so that we can understand under what conditions these meeting times arise. In response to these questions, we synthesize the main results from [4], [2], and [7].

We first build motivation for understanding how this game works by finding strategies for different graphs. We find the existence of pure optimal strategies for both Angel and Demon on any graph G, as well as finding an explicit strategy for the Demon game played on a path. We also deduce a general pure optimal strategy which utilizes a *remote vertex* for Demon and a *central vertex* for Angel.

After developing an intuition for the game, a natural question arises: how long is the game expected to last? To answer this, we look for the graph which maximizes the hitting time of a random walk across all n-vertex graphs. The graph which achieves this is known as the *lollipop graph* (see Figure 1). We then calculate that the maximal hitting time on such a graph is of order $\frac{4}{27}n^3$, and we develop this calculation explicitly. We use this calculation to obtain the Demon's bound, then again in obtaining the lower bound for the Angel by employing his pure optimal strategy. Then, since the random game must have bounds which lie between these two, we obtain the bounds on each of the three modes of play.

2. Preliminaries

We assume basic knowledge of probability, random walks, and graph theory. It should be noted that we use the notation $o(1)$ to denote some function $f : \mathbb{N} \to \mathbb{R}$ such that $\lim_{n\to\infty} \frac{f(n)}{1} = 0$.

We again emphasize that G is required to be undirected and connected, so that the notion of maximum meeting time is always well-defined, and when we refer to "any graph G," we imply that G is also undirected and connected. The following definitions are required:

Definition 2.1. The *hitting time* $H_G(x, y)$ is the expected time for a random walk starting at vertex x on graph G to reach vertex y.

We may also define a special quantity involving hitting time as

$$H_G^M(x, y) := H_G(x, y) + M \cdot e(G),$$

where M is a positive real number, and $e(G)$ is the number of edges on graph G. We denote $V(G)$ as the set of vertices on G. Now let us define for fixed M the quantity $H^M(n) := \max\{H_G^M(x, y) : |V(G)| = n \text{ and } x, y \in V(G)\}$.

Definition 2.2. An n-vertex graph G is (n, M)-*extremal* if there exists a pair of vertices x, y such that $H_G^M(x, y) = H^M(n)$.

We also have some special vertices which are useful for later results. A *remote vertex* is a vertex r such that $H_G(r, x) \le H_G(x, r)$ for all x of G. Similarly, a *central vertex* is a vertex c such that $H_G(x, c) \le H_G(c, x)$ for all x. It is noted in [7] that for any G, some c and r both exist. Applying these definitions, we define the "potential function"

$$Z(z) = H_G(c, z) - H_G(z, c).$$

Intuitively, we can think of the potential of z as a measure of how far z is from being central; c has a potential of 0, whereas r has the maximum potential, which is easy to see from the aforementioned hitting time inequalities.

We also define the Demon's harmonic function as

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

where $C(x,y) = H_G(x,y) + H_G(y,x)$. We can think of $C(x,y)$ as the "round trip time," or the expected amount of time it takes for a random walk to travel from x to y, then back to x. We define the Angel's harmonic function as

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

In our investigation of which n-vertex graphs produce the maximal hitting time, it is useful to define the *lollipop graph*.

Definition 2.3. A *lollipop graph* L_n^m is an n-vertex graph containing a clique of m vertices which is connected to a path of $t := n - m$ vertices (see Figure 1).

In the setup that follows, we will have the starting point x of our random walk on L_n^m be some vertex which is in the clique, z is the vertex in the clique which is also an endpoint of the path (it is the point where the path and the clique intersect), and y is the vertex which is the endpoint of the path that is farthest from the clique.

An L_9^6 lollipop graph is shown in Figure 1 below, with the points x, y, and z as defined above. Note that this particular graph, as we will see, is part of a class of lollipop graphs which maximize $H_G(x,y)$, because the number of vertices in the clique of the graph is $\frac{2}{3}$ the total number of vertices:

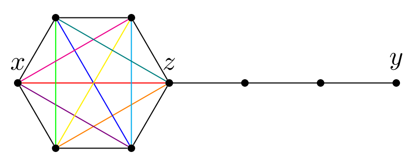

FIGURE 1. The lollipop graph L_9^6.

From Figure 1, we develop an intuition for why $H_G(u,v) \neq H_G(v,u)$ for any pair u, v of vertices. In particular, if we consider the graph from Figure 1, we see that $H(x,z) = 6-1 = 5$, since $\deg u = 5$ for any vertex $u \neq z$ in the clique. However, we have $H(z,x) > 5$, since the random walk starting on z can not only reach any other vertex in the clique but may also retreat further down the path with probability $p > 0$, thereby extending the expected time to reach x.

If w is a neighboring vertex of x, then we denote this as $w \sim x$. As a convention used for a general function f on vertices of graphs, we write $f(\bar{x}, y)$ to denote the average value of $f(w,y)$ over all the neighbors w of x and with some second, fixed input y.

A "strategy" in the context of the token game is a set of criteria used by the player to assign some probability p such that with tokens on vertices x and y, the player moves the token on x with probability p. For example, the Random player implements a strategy which states that at every move, $p = \frac{1}{2}$ is assigned to either token.

The goal of the token game depends upon the player. Demon wishes to avoid token collision as long as possible; Angel wishes to collide the tokens as soon as possible; Random is a "neutral" player which has no ultimate goal.

Functions with a subscript or superscript of S, such as χ_S, denote the value of the function when the player follows strategy S. The function $M_G(x, y)$ denotes the maximum expected meeting time over all strategies S with starting vertices x, y.

Lastly, some proofs require a notion of a "distance" metric on a graph G. The "distance" between two vertices x and y on G is the number of edges on the shortest path from x to y. Note that such a path must exist since G is assumed to be connected.

3. Strategies

We begin with a general result about strategies, and from the proof we also learn of a common proof technique in studying this game. Note that a *pure* strategy is a strategy which only depends on the current positions of the tokens. An *optimal* strategy for the Demon is a strategy which maximizes the expected meeting time, while an optimal strategy for the Angel minimizes the expected meeting time.

Proposition 3.1. *On any graph G, Demon has a pure optimal strategy.*

We imitate the proof found in [4].

Proof. Let $M_G(x, y)$ be the maximum expected meeting time with tokens initially on vertices x and y on G across all strategies for Demon. Let $S(x, y)$ be a strategy maximizing $M^{S(x,y)}$, the expected meeting time of the tokens for strategy $S(x, y)$. We define a tournament T such that x wins over y when $S(x, y)$ moves x, and T is similarly defined for y. Let us define a strategy S which corresponds to T. Now, S is indeed pure; a pure strategy is equivalent to a tournament on the vertices of G, as noted in [4]. (The intuition is that, given some tournament T on $V(G)$, each pair of vertices x, y will have a winner over the other. The winner depends solely on the pair x, y we choose. Then we may create a strategy S which assigns a probability p to moving the token on x depending on the winner of the tournament. Here, we assign $p = 1$ to the winner. Therefore, S only assigns p based on the current location of the tokens (x, y); hence S is pure.) We now wish to show that S is optimal.

Assume to the contrary that $M_G(x, y) > M^S(x, y)$. Then define

$$\alpha := \max\{M_G(x, y) - M^S(x, y) : x, y \in V(G)\},$$

where $V(G)$ is the set of vertices of G. There may be several pairs x, y which achieve α; let us choose the one with the minimal distance between x and y. Without loss of generality, assume x is the winner over y in T, and $S(x, y)$ moves x with probability $p > 0$. Then we have

$$M^{S(x,y)} = 1 + pM_G(\bar{x}, y) + (1 - p)M_G(\bar{x}, y)$$

Furthermore, we must have

$$M^{S(x,y)} = 1 + M_G(\bar{x}, y)$$

because if not, then moving y would be superior and hence x would not be the winner over y in T.

Then we have

$$\begin{aligned} M^{S(x,y)}(x,y) &= 1 + M_G(\bar{x},y) \\ &\leq 1 + M^S(\bar{x},y) + \alpha \\ &= M^S(x,y) + \alpha \\ &= M_G(x,y) \\ &= M^{S(x,y)}(x,y). \end{aligned}$$

However, there is at least one vertex z adjacent to x which is closer to y than x is. If $z \neq y$, then α is no longer attained at z, implying $\alpha > M_G(z,y) - M^S(z,y)$. Therefore, $\alpha > M_G(\bar{x},y) - M^S(\bar{x},y)$. Hence, the above inequality is strict. If $z = y$, then $M_G(z,y) = 0 = M^S(x,y)$. In any case, we reach a contradiction. $\square$

As noted in [1], Proposition 3.1 also holds for Angel.

Corollary 3.2. *On any graph G, Angel has a pure optimal strategy.*

Proof. The proof is analogous to that of Proposition 3.1. We instead let $M_G(x,y)$ denote the minimum expected meeting time with tokens initially on vertices x and y on G across all strategies, and we let $S(x,y)$ be a strategy which minimizes $M^{S(x,y)}$. The tournament T is defined as in the proof of Proposition 3.1. Assume to the contrary that $M_G(x,y) < M^S(x,y)$, and define α as the maximum difference between $M^S(x,y)$ and $M_G(x,y)$ across all vertex pairs x, y. The remainder of the proof is similar to the proof of Proposition 3.1. $\square$

We may search for optimal strategies on certain graphs, and it turns out some are not hard to find. If we consider a path, for example, we reason intuitively that if Demon wishes to keep the tokens apart, he should keep one token on an endpoint of the path, and only move it if necessary. This intuition is correct and is formalized in Proposition 3.3. Note that a player following a *degree strategy* with tokens on x and y moves the token on x whenever $\deg(y) < \deg(x)$.

Proposition 3.3. *Demon's strategy on a path is optimal if and only if it is a degree strategy.*

The proof here from [1] mimics that from [4]. The following proof technique is an example of *random walk coupling* (see [5]).

Proof. Assume G is a path with $V(G) = \{v_0, ..., v_{n-1}\}$. Consider a copy of G, denoted G', with $V(G') = \{v'_0, ..., v'_{n-1}\}$. Furthermore, consider a "distance token" on G' such that with tokens on G at v_i and v_j, the location of the distance token is v'_k, where $k = |i - j|$.

If Demon follows a degree strategy, then as the pair of tokens moves on G, the distance token takes a uniform random walk on G' with absorbing state v'_0 and reflecting state v'_{n-1}. The expected length of the game under a degree strategy is $H_{G'}(v'_k, v'_0)$ where k is the initial position of the distance token. However, if at any pair v_0, v_i or v_i, v_{n-1}, with $0 < i < n-1$, Demon has positive probability of moving the token at the endpoint, then the probabilities along the edges leading from the corresponding vertex on G' are skewed toward v'_0, thus decreasing the expected duration of the game. $\square$

We may also look to see if there is a general rule which Demon and Angel can follow in order to achieve their respective goals. It turns out that seeking this general rule leads us to great progress on bounding the meeting times. First, however, we must do some analysis to get there. Define $q_S(z;x,y)$ as the probability that if a player is following strategy S, tokens initially at x and y will meet for the first time at z. Let us begin our first main goal with a theorem.

Theorem 3.4.
$$M^S(x,y) = \phi(x,y) - \sum_z q_S(z;x,y)[Z(r) - Z(z)]$$

In order to develop the tools to prove this theorem, the following definition is helpful.

Definition 3.5. A function $f : V(G) \times V(G) \to \mathbb{R}$ is *biharmonic* if
$$\min\{f(\bar{x},y), f(x,\bar{y})\} \le f(x,y) \le \max\{f(\bar{x},y), f(x,\bar{y})\}.$$

Lemma 3.6 (Maximum Principle)**.** *If f is biharmonic, then* $\max_{x,y} f(x,y) = f(z,z)$ *for some z.*

Proof. The proof here from [7] is similar to that of Proposition 3.1. Assume the contrary. Define $f_M = \max_{x,y} f(x,y)$. There may be several pairs of x and y such that $f_M = f(x,y)$; choose the pair of minimal distance between the two. Because f is biharmonic, we must have that either $f(x,y) \le f(\bar{x},y)$ or $f(x,y) \le f(x,\bar{y})$. Assume the first case without loss of generality. We have
$$\begin{aligned} f_M = f(x,y) &\le f(\bar{x},y) \\ &= \frac{1}{\deg(x)} \sum_{w \sim x} f(w,y), \end{aligned}$$
by definition. But since there must be at least one neighbor of x, denoted w', such that w' is closer to y than x is, f_M must not be achieved at all neighbors w, so we have that
$$\begin{aligned} \frac{1}{\deg(x)} \sum_{w \sim x} f(w,y) &\le \frac{1}{\deg(x)} f(w',y) + \frac{\deg(x)-1}{\deg(x)} f_M \\ &< f_M, \end{aligned}$$
in contradiction. □

The analogous Minimum Principle holds with a similar proof.

To assist us in proving Theorem 3.4, define
$$\chi_S(x,y) := M_S(x,y) - \phi(x,y) + \sum_z q_S(z;x,y)[Z(r) - Z(z)].$$

Lemma 3.7. *χ_S is biharmonic.*

Proof. The proof is from [7]. Let the tokens have starting positions x and y. Let p be the probability that S moves the token at x. Our goal is to show the following:
$$\chi(x,y) = p\chi(\bar{x},y) + (1-p)\chi(x,\bar{y}).$$

We know that $M^S(x,y) = 1 + pM_S(\bar{x},y) + (1-p)M^S(x,\bar{y})$. We also have
$$\begin{aligned} \phi(x,y) &= 1 + \phi(\bar{x},y) \\ &= 1 + \phi(x,\bar{y}). \end{aligned}$$

Furthermore,

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

We may rewrite this as

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

Now,

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

Now, assume without loss of generality that $\chi(\bar{x},y) \geq \chi(x,\bar{y})$. Then

$$\begin{aligned}
0\cdot(\chi(\bar{x},y) - \chi(x,\bar{y})) + \chi(x,\bar{y}) &\leq \chi(x,y) \\
&= p(\chi(\bar{x},y) - \chi(x,\bar{y})) + \chi(x,\bar{y}) \\
&\leq 1(\chi(\bar{x},y) - \chi(x,\bar{y})) + \chi(x,\bar{y}).
\end{aligned}$$

Therefore, χ is biharmonic. □

Now, with the proper tools at hand, it will be merely an application of definitions and the now-established lemmas to show that χ is equal to zero. We prove Theorem 3.4 as follows.

Proof. By our definitions, as shown in [7], we get

$$\begin{aligned}
\chi_S(x,x) &= M^S(x,x) - \phi(x,x) + \sum_z q_S(z;x,x)[Z(r)-Z(z)] \\
&= 0 - (Z(r) - Z(x)) + [Z(r) - Z(x)],
\end{aligned}$$

since $q_S(z;x,x) = 0$ for $z \neq x$ and equals 1 for $z = x$. Then $\chi_S(x,x) = 0$. Then by the Maximum and Minimum principles, at an arbitrary vertex x, $\chi_S(x,x)$ is equal to its maximum and minimum. Therefore, χ is simply zero. □

Much of the difficult work of bounds has already been done. We now easily bound the meeting times by the harmonic functions which will lead to an easy constant bound of H_{max}, the maximum hitting time over all starting positions x, y on all n-vertex graphs G.

Corollary 3.8.

$$\psi(x,y) \leq M^S(x,y) \leq \phi(x,y)$$

Proof. The proof combines several results from [7]. By Theorem 3.4, we get

$$\begin{aligned} M^{Angel} &= \phi(x,y) - \sum_z q_{Angel}(z;x,y)[Z(r)] + \sum_z q_{Angel}(z;x,y)[Z(z)] \\ &= \phi(x,y) - Z(r) + \sum_z q_{Angel}(z;x,y)[Z(z)] \\ &= \psi(x,y) + \sum_z q_{Angel}(z;x,y)[Z(z)] \text{ (see preliminaries).} \end{aligned}$$

Similarly, for the Demon,

$$M^{Demon}(x,y) = \phi(x,y) - \sum_z q_{Demon}(z;x,y)[Z(r) - Z(z)],$$

with the subtracted term being nonnegative. $M^S(x,y)$ is bounded from above by Demon's meeting time, and $M^S(x,y)$ is bounded from below by Angel's meeting time, giving the inequality in the corollary statement. □

Corollary 3.9. $M_G(x,y) \leq 2H_{max}$

Proof. We immediately get

$$\begin{aligned} M_G(x,y) &\leq \phi(x,y) \\ &= H_G(y,x) + H_G(x,r) - H_G(r,x) \\ &\leq H_{max} + H_{max} - 0 \\ &= 2H_{max}. \end{aligned}$$

□

We see that Demon, even using an optimal strategy, cannot prevent token collision forever.

Corollary 3.10. *For any remote vertex r of graph G, Demon has a pure optimal strategy that never moves a token off r. Similarly, Angel has a pure optimal strategy that never moves a token off a central vertex c.*

Proof. The Demon case is proven in [4], with the Angel case being analogous, but also proven in [7]. If $y = r$, then we have

$$\begin{aligned} M_{Demon}(x,r) \leq \phi(x,r) &= H_G(x,r) + H_G(r,r) - H_G(r,r) \\ &= H_G(x,r). \end{aligned}$$

Therefore, the Demon can never achieve a larger meeting time than if he were to place a token on r and wait for the token on x to hit it. Similarly for Angel, if $y = c$, then we have

$$\psi(x,c) = H_G(x,c) \leq M_{Angel}(x,c).$$ □

Now we have developed strategies for the two players, and we have found an upper bound for the maximum expected meeting time on a graph G in terms of H_{max}. We now seek the value of H_{max} to achieve an explicit bound.

4. Bounds

The following theorem shows that the lollipop graph L_n^m is the graph which maximizes the hitting time across n-vertex graphs for values close to $m = \frac{2}{3}n$, and the expected hitting time as a result is approximately $\frac{4}{27}n^3$. Let us develop why, on a lollipop graph, these results for m and n occur, basing our preliminary discussion off the one from [2].

We start with x on a clique, with z and y in their aforementioned positions (see Figure 1). All vertices in a clique are connected, and the degree of any vertex (besides z) in the m-vertex clique is $m-1$. So on any given step, the probability of reaching any other vertex in the clique, in particular z, is $\frac{1}{m-1}$. Once we are at z, the probability of continuing on the path (instead of going back to the clique) is $\frac{1}{m}$. Moreover, the probability of reaching the end of the path of length t before reaching z again is $\frac{1}{t}$ (this is exactly the Gambler's Ruin problem, about which one may consult any number of resources. See [6]). Therefore, the expected number of steps to reach y from x is about m^2t.

Then if we define for fixed n the function $H(m) = m^2t = m^2(n-m)$, we may find the maximum using elementary calculus. Solving $H'(m) = 2mn - 3m^2 = 0$ gives $m = \frac{2}{3}n$, from which $H(\frac{2}{3}n) = \frac{4}{27}n^3$.

Theorem 4.1 below, from [2], proves that the lollipop graph is the appropriate graph to maximize the hitting time, and that a value of m as near to $\frac{2}{3}n$ as possible is always what is required.

Theorem 4.1.

(1) If $M > n-1$, the complete graph K_n is the only (n, M) extremal graph.

(2) If $0 \leq M \leq n-1$, then the lollipop graph L_n^m with $m = \left\lfloor \dfrac{2n+M+1}{3} \right\rfloor$ is (n, M) extremal. It is the only (n, M) extremal graph, unless $2n+M+1$ is divisible by 3. Then $m = \dfrac{2n+M-2}{3}$ is also an (n, M) extremal lollipop graph.

The full proof of this amazing result takes the entirety of [2]. What this tells us immediately is the following: we know precisely which graphs maximize $H_G(x, y)$, we know the positions of x and y on the graph, and we obtain a bound on what this hitting time should be. Using a brief calculation from [2], we are able to find the exact value of H_{max} in terms of n.

Set $M = 0$ in $H_G^M(x, y)$. Let $G = L_n^m$, and let the path vertices be $v_0 = z, v_1, \ldots, v_t = y$ and $t = n - m$ (see Figure 1). We have that

$$\begin{aligned} H_G(x, y) &= H_G(x, z) + H_G(z, y) \\ &= m - 1 + H_G(z, y). \end{aligned}$$

Calculating the value of $H_G(z, y)$, we get

$$\begin{aligned} H_G(z, y) &= \sum_{i=1}^{t} H_G(v_{i-1}, v_i) \\ &= \sum_{i=1}^{t} H_{G\setminus\{v_{i+1}, \ldots, v_t\}}(v_{i-1}, v_i), \end{aligned}$$

since we need not consider any part of the graph which lies beyond v_i for any i such that $1 \leq i \leq t$. Proceeding, we get

$$\begin{aligned} H_G(z,y) &= \sum_{i=1}^{t}(H_{G\setminus\{v_{i+1},...,v_t\}}(v_i,v_i)-1) \\ &= \sum_{i=1}^{t}(2e(G\setminus\{v_{i+1},...,v_t\})-1) \\ &= 2t\binom{m}{2}+\sum_{i=1}^{t}(2i-1) \\ &= tm(m-1)+t^2. \end{aligned}$$

Thus

$$H_G(x,y) = (n-m)(m)(m-1)+(n-m)^2+m-1.$$

Define $H(n)$ to be the function above when $m = \lfloor\frac{2n+1}{3}\rfloor$. The floor function behavior splits $H(n)$ into a piecewise function which depends on the value $n \bmod 3$, so we have the three cases (all cubic polynomials) which we may denote H_0, H_1, H_2. One may take any four values of n such that $n \equiv 0 \bmod 3$, and record the outputs of them, so we get the four points $(n_1, H_0(n_1)), (n_2, H_0(n_2)), (n_3, H_0(n_3)), (n_4, H_0(n_4))$. One can then determine the polynomial $H_0(n)$ completely, and we may repeat the process for the other cases. Then we find, in agreement with [4], that

$$H(n) = \begin{cases} \frac{4}{27}n^3-\frac{1}{9}n^2+\frac{2}{3}n-1, & n \equiv 0 \bmod 3 \\ \frac{4}{27}n^3-\frac{1}{9}n^2+\frac{4}{9}n-\frac{13}{27}, & n \equiv 1 \bmod 3 \\ \frac{4}{27}n^3-\frac{1}{9}n^2+\frac{2}{3}n-\frac{29}{27}, & n \equiv 2 \bmod 3 \end{cases}.$$

As we have seen in Corollary 3.9, $M_G(x,y)$ is bounded by twice the above value. As we now progress, we will see that the maximum hitting time here shown provides a lower bound for $M_G(x,y)$, because $M_G(x,y)$ may exceed $H_G(x,y)$ (see [4] for examples), and in the aforementioned lollipop, we will see that they are equal. The following results will provide a stricter upper bound, squeezing our target variable within a rather narrow margin.

Lemma 4.2. *If x,y,z are arbitrary vertices of a connected, undirected graph G, then*

$$H_G(x,y)+H_G(y,z)+H_G(z,x) = H_G(x,z)+H_G(z,y)+H_G(y,x)$$

As noted in [7], this is equivalent to the reversibility of Markov chains, and the full proof is included in [4].

Lemma 4.3. *In an n-vertex graph G, where $n \geq 13$, for any three distinct vertices x,y,z of G, we have*

$$C_G(x,y)+C_G(y,z)+C_G(z,x) \leq \frac{8}{27}n^3+\frac{8}{3}n^2+\frac{4}{9}n-\frac{592}{27}$$

Proof Sketch. We outline the full proof from [4]. A result from [3] shows us that

$$C_G(x,y) = 2e(G)R_G(x,y), \tag{4.4}$$

where $R_G(x,y)$ denotes the effective resistance between x and y when treating the graph as a circuit.

Define $d_G(x,y)$ to be the distance between x and y. We have that

$$R_G(x,y) \leq d_G(x,y)$$

from which

$$C_G(x,y) + C_G(y,z) + C_G(z,x) \leq 2e(G)(d_G(x,y) + d_G(y,z) + d(z,x)).$$

In the steps that follow, we consider the case when the sum of the distances in the above inequality are even and when they are odd. In the first case, we have

$$d_G(x,y) + d_G(y,z) + d(z,x) = 2k$$

for some integer k. Then there are some nonnegative integers a, b, and c for which $d_G(x,y) = a+b$, $d_G(y,z) = b+c$, $d_G(z,x) = c+a$, and $k = a+b+c$. The n vertices may be partitioned into $a+b+c+1$ nonempty disjoint subsets: the set of all vertices t such that $d_G(x,t) = 0$, and for $d_G(x,t) = 1$, and so on, up until $d_G(x,t) = a-1$, and similarly for $d_G(y,t)$ and $d_G(z,t)$, and the final set is Residue, which contains everything else.

Let F contain one vertex from each of the $k+1$ subsets, and let $H = G - F$ contain the remaining $n-k-1$ vertices. F contains at most k edges; H contains at most $\frac{(n-k-1)(n-k-2)}{2}$ edges; the maximum amount of edges between H and F is $4(n-k-1)$. Therefore, the total number of edges is bounded by $k + \frac{(n-k-1)(n-k-2)}{2} + 4(n-k-1)$.

It follows, then, from (4.4) and from the previous calculations, that the sum of the commute times is bounded by

$$2 \cdot 2k \left(k + \frac{(n-k-1)(n-k-2)}{2} + 4(n-k-1) \right),$$

which is maximized at $k = \left\lceil \frac{n+3}{3} \right\rceil$ for $n \geq 13$. Depending on the value of n modulo 3, the optimal value of k gives different polynomial bounds.

We repeat a similar process for the case when the sum of the distances is equal to $2k+1$. Similar calculations ensue, and one obtains the expression

$$2(2k+1) \left(k + 2 + \frac{(n-k-2)(n-k-3)}{2} + 4(n-k-2) \right).$$

Again, the polynomial bounds depend on n modulo 3; each such value results in different polynomials. The largest such polynomial bound between the two cases is found in the first case; the bound is $\frac{8}{27}n^3 + \frac{8}{3}n^2 + \frac{4}{9}n - \frac{592}{27}$. $\square$

With an application of Lemma 4.3, we can now greatly improve upon our upper bound estimate compared to our earlier bound given by $2H_{max}$. The following Theorem nears us to our main result.

Theorem 4.5. *For $n \geq 13$,*

$$M_G(x,y) \leq \frac{4}{27}n^3 + \frac{4}{3}n^2 + \frac{2}{9}n - \frac{296}{27}.$$

Proof. The proof is again from [4]. By definition,

$$C_G(x,y) = H_G(x,y) + H_G(y,x).$$

Combining the results of Corollary 3.8 and Lemma 4.2 for $z = r$ gives us

$$M_G(x,y) \leq H_G(x,y) + H_G(y,z) - H_G(z,y) = H_G(y,x) + H_G(x,z) - H_G(z,x). \tag{4.6}$$

Since the above inequalities hold, we certainly have

$$M_G(x,y) \leq H_G(x,y) + H_G(y,z) + H_G(z,y) \tag{4.7}$$

$$M_G(x,y) \leq H_G(y,x) + H_G(x,z) + H_G(z,x). \tag{4.8}$$

Adding (4.7) and (4.8) gives

$$\begin{aligned} 2M_G(x,y) &\leq (H_G(x,y)+H_G(y,x))+(H_G(y,z)+H_G(z,y))+(H_G(x,z)+H_G(z,x)) \\ &= C_G(x,y) + C_G(y,z) + C_G(z,x). \end{aligned}$$

If z is distinct from x and y, we may use Lemma 4.3 to divide both sides by 2 and obtain the theorem statement immediately.

Now if we have $z = x$, we get from (4.6) that

$$\begin{aligned} M_G(x,y) &\leq H_G(x,y) + H_G(y,x) - H_G(x,y) \\ &= H_G(y,x) \\ &< C_G(x,y). \end{aligned}$$

Now we select some z' which is distinct from x and y. Since

$$C_G(y,z') + C_G(z',x) \geq C_G(x,y),$$

we apply the same procedure as before, obtaining

$$2M_G(x,y) \leq C_G(x,y) + C_G(y,z') + C_G(z',x)$$

and apply Lemma 4.3 as before to get the theorem statement as desired. □

Theorem 4.5 has provided us an upper bound on all expected meeting times. If we are to find a lower bound for the Angel's game in his worst case scenario, we will have achieved our final goal. It turns out that we now have all the results needed to determine such a bound, providing both bounds for the expected meeting times for the three main modes of play. The final result is given in the theorem below.

Theorem 4.9. *Define Demon's maximum expected meeting time over all n-vertex graphs G and all initial vertices x, y to be M_D. Define M_A and M_R similarly for Angel and Random meeting times, respectively. Then*

$$\begin{aligned} \left(\frac{4}{27} + o(1)\right) n^3 \leq M_D \leq \left(\frac{4}{27} + o(1)\right) n^3 \\ \left(\frac{1}{27} + o(1)\right) n^3 \leq M_A \leq \left(\frac{4}{27} + o(1)\right) n^3 \\ \left(\frac{1}{27} + o(1)\right) n^3 \leq M_R \leq \left(\frac{4}{27} + o(1)\right) n^3. \end{aligned}$$

Proof. The proof is based off the discussion in [7]. First, Theorem 4.5 immediately bounds our three values of interest from above.

Now, we obtain the lower bound of M_D through the use of Theorem 4.1. The maximum expected meeting time on a graph G may, in general, be larger than the maximum hitting time on G, and in the case of $G = L_n^m$ with m as near as possible to $\frac{2}{3}n$, the maximum hitting time is the same as the maximum expected meeting time. This follows from the fact that Demon, faced with a token on x in the clique and on y at the end of the path, may leave the token on y and only move the token initially at x as his best strategy; this is due to Corollary 3.10. (The fact that y is indeed a remote vertex is noted in [4]). So we have $H_{max} = \left(\frac{4}{27} + o(1)\right) n^3 \leq M_D$.

To show the lower bound for M_A, we consider a "barbell" graph which joins two lollipop graphs, each of about $\frac{n}{2}$ vertices, with each half being (approximately) symmetrical lollipop graphs which have about $\frac{2}{3}n$ total vertices contained in a clique, and $\frac{1}{3}n$ vertices on the path connecting the cliques (see Remark 4.10). We will prove that the barbell graph creates the worst case scenario for the Angel.

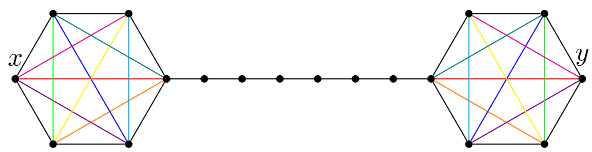

FIGURE 2. A barbell graph with 18 vertices.

Assume that Angel is playing on a barbell graph, with a token on each clique. Recall that Angel's best strategy is to move the token on (say) x until it reaches the center c, and then move the token on y until reaching c where the tokens collide. Assume Angel is using this strategy. However, since each half of the graph is a lollipop graph with m as close as possible to $\frac{2}{3}$ the number of the vertices of the lollipop, if Angel uses this strategy, we know that the expected time for the token initially on x to reach c is maximized, and the same is true for the token initially on y. Therefore, the expected meeting time of the tokens, which is in this strategy the sum $H_G(x, c) + H_G(y, c)$, is as large as possible. Then Angel is simultaneuously playing his best strategy, while faced with the worst case scenario for this strategy. So the meeting time in this instance is about $2 \cdot \frac{4}{27}(n/2)^3 = \frac{1}{27}n^3$, giving the lower bound to the maximum expected meeting time for Angel as stated in the theorem.

The random game must be bounded between the upper bound of M_D and the lower bound of M_A, since the random game, in the best case scenario, plays the optimal Angel strategy by chance, and in the worst case scenario, plays the optimal Demon strategy by chance. Thus we have the bounds on the random game as specified in the theorem statement. $\square$

Remark 4.10. Constructing these graphs in practice gives rise to issues of parity, since of course, we are not always able to divide the graph's vertices as desired. If n is odd, then we may construct two identical lollipop graphs of $\lfloor \frac{n}{2} \rfloor + 1$ vertices sharing the midpoint of the path. Then we also may choose the clique of each half to be of $\left\lfloor \frac{2\left(\lfloor \frac{n}{2} \rfloor + 1\right) + 1}{3} \right\rfloor$ vertices as desired. If n is even, if we wish to construct two halves sharing the midpoint, the cliques are asymmetrical, since joining two identical $L^m_{\frac{n}{2}}$ graphs will not intersect at a vertex (Figure 2).

Consider our approximate construction of the ideal graph in the following way. Given even n, fix a midpoint of the path c. To one side of c (say, to the left), construct an $L^m_{\frac{n}{2}+1}$ graph with $m_1 = \left\lfloor \dfrac{2(\frac{n}{2} + 1) + 1}{3} \right\rfloor$ and c a vertex of this graph. Define k to be the number of vertices in this graph excluding c and excluding those in the clique, and w will be similarly defined for the right-hand graph. To the right of c, construct an $L^m_{\frac{n}{2}}$ graph. If $m_2 = \left\lfloor \dfrac{2(\frac{n}{2})+1}{3} \right\rfloor$, then

$$k = \frac{n}{2} - \left\lfloor \frac{2(\frac{n}{2}+1)+1}{3} \right\rfloor$$
$$w = \frac{n}{2} - \left(1 + \left\lfloor \frac{2(\frac{n}{2})+1}{3} \right\rfloor\right).$$

If we demand symmetry of the path with respect to c, then we want that $k = w$. Then we must have $\lfloor \frac{n+3}{3} \rfloor = \lfloor \frac{n+1}{3} \rfloor + 1$. This fails if and only if $n+1$ is divisible by 3. If so, then $2(n/2)+1$ is divisible by 3. Recall from Theorem 4.1 that this implies $\frac{2(n/2)-2}{3}$ is also an $(n/2, 0)$-extremal lollipop graph. If this new extremal graph requires that one less vertex resides in the clique, then the difference of 1 between k and w is eliminated. One may check that this is true:

$$\left\lfloor \frac{2(\frac{n}{2})+1}{3} \right\rfloor - \frac{2(\frac{n}{2})-2}{3} = 1 \text{ if and only if } n+1 \equiv 0 \bmod 3.$$

In summary, if $n+1$ is not divisible by 3 and n is even, then construct a barbell graph out of an $L^{m_1}_{\frac{n}{2}+1}$ graph and an $L^{m_2}_{\frac{n}{2}}$ graph, with $m_1 = \lfloor \frac{2(\frac{n}{2}+1)+1}{3} \rfloor$ and $m_2 = \lfloor \frac{2(\frac{n}{2})+1}{3} \rfloor$. If $n+1$ is divisible by 3 and n is even, then on one side, construct an $L^{m_1}_{\frac{n}{2}+1}$ graph and an $L^{m_2}_{\frac{n}{2}}$ graph, with $m_1 = \lfloor \frac{2(\frac{n}{2}+1)+1}{3} \rfloor$ and $m_2 = \frac{2n-2}{3}$, with the two graphs sharing a center. If n is odd, one may construct two identical $L^m_{\frac{n}{2}+1}$ graphs sharing a center, where $m = \lfloor \frac{2(\frac{n}{2})+1}{3} \rfloor$. The result is a graph which well approximates the ideal case used in the proof.

Acknowledgments. It is a pleasure to thank my mentor, Victor Gardner, for his tremendous help and guidance in finding intruiging reading material, and for the effort he put forth in teaching the basics of measure-theoretic probability, without which such a paper would not have been possible.

References

[1] Bañuelos, J. (2011). *A Simultaneous Random Walk Game* (Master's thesis, Macalaster College, 2011) (pp. 8-27). DigitalCommons@MacalasterCollege.

[2] Brightwell, G., & Winkler, P. (1990). Maximum hitting time for random walks on graphs. *Random Structures & Algorithms,* 1(3), 263-276. doi:10.1002/rsa.3240010303

[3] Chandra, A. K., Raghavan, P., Ruzzo, W. L., & Smolensky, R. (1989). The electrical resistance of a graph captures its commute and cover times. *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing - STOC '89.* doi:10.1145/73007.73062

[4] Coppersmith, D., Tetali, P., & Winkler, P. (1993). Collisions Among Random Walks on a Graph. *SIAM Journal on Discrete Mathematics,* 6(3), 363-374. doi:10.1137/0406029

[5] den Hollander, F. (n.d.). Probability Theory: The Coupling Method. Retrieved August 11, 2020, from http://websites.math.leidenuniv.nl/probability/lecturenotes/CouplingLectures.pdf

[6] Lalley, S. (n.d.). ONE-DIMENSIONAL RANDOM WALKS. Retrieved June 26, 2020, from http://galton.uchicago.edu/ lalley/Courses/312/RW.pdf

[7] Tetali, P., & Winkler, P. (1993). Simultaneous Reversible Markov Chains. *Combinatorics, Paul Erdös Is Eighty, 1,* 433-452.