Meeting Times of Random Walks on Graphs

Nader H. Bshouty Lisa Higham Jolanta Warpechowska-Gruca

Computer Science
University of Calgary
Calgary, AB, T2N 1N4
Canada
(e-mail: {bshouty,higham,jolanta}@cpsc.ucalgary.ca)

ABSTRACT. We prove an upper bound on the meeting time of an arbitrary number of random walks in any connected undirected graph in terms of the meeting times of fewer random walks. We show that the bound is tight for rings, and that it is both stronger and more general than a bound suggested by Tetali and Winkler [4].

KEYWORDS: random walks, meeting time, hitting time.

1. Introduction

A *random walk* on an undirected connected graph G is a process that starts at some vertex of G, and at each time step moves to one of the neighbours of the current vertex, each of them chosen with equal probability. Since this probability does not depend on the time step, a random walk is a Markov chain that has state set the vertices of G, and with transition probabilities given by: $p(x, y) = 0$ if y is not adjacent to x, and $p(x, y) = 1/d(x)$ if y is adjacent to x, and $d(x)$ is the degree of vertex x in G.

Many tasks in distributed computing require one processor, a *leader*, to be distinguished from among several anonymous processors. This can be achieved using random walks as follows. Initially, each processor contending for leadership gets a *token*. At each step, a processor that has a token passes it randomly to one of its neighbours. When more than one token meet at one processor, they merge to one token. In a connected undirected graph, eventually all tokens will merge to one token. This idea was used by Israeli and Jalfon [3] to create a self-stabilizing mutual exclusion system. The expected time until several random walks on a graph collapse into one characterizes the performance of this protocol. In the literature, however, one random walk on a graph is usually analyzed. Aldous [1], as well as Tetali and Winkler [4], and Coppersmith, Tetali and Winkler [2] studied the expected number of steps until two random walks on a graph meet. Aldous [1] discussed the expected number of steps until random walks, starting from every vertex in the graph and coalescing upon collision, collapse into one. This leaves a gap in analysis between two random walks and random walks starting from every vertex. Tetali and Winkler [4] suggested using their technique to obtain an upper bound on the meeting time of an arbitrary number of tokens in terms of meeting times between all pairs of tokens. We derive a simple technique for proving bounds on the expected time until an arbitrary number of random walks meet (lemma 4). This technique is used to prove a bound (theorem 6) that is stronger and more general than the bound suggested by Tetali and Winkler [4]. Finally, we show that the stronger bound is tight for rings.

2. Generalized meeting times

Consider the following *meeting-tokens* game. Initially k tokens are placed at vertices $x_1, \ldots, x_k$ of graph G. At each step, a token is chosen and that token moves to a randomly chosen neighbour. When two tokens meet at a vertex, they merge to one token. The game ends when only one token remains. In general, the expected duration of the meeting-tokens game depends on the scheduler, which determines the token that moves at each step. Several definitions follow that make precise the progression of this game.

A *configuration* of k tokens on graph G is a multiset of k vertices of graph G. A configuration $\{x_1, \ldots, x_k\}$ describes a placement of k tokens on the vertices of graph G, where x_i is the *position* of token i. If $x_i = x_j$ then the tokens i and j are merged to one token. The initial placement of tokens is called an *initial configuration*. A configuration where all positions of tokens are identical is called a *terminal configuration*. A configuration that is not terminal is called *nonterminal*. For notational convenience the braces ($\{,\}$) are omitted when a configuration is the only argument of a function.

A configuration $\{y_1, \ldots, y_k\}$ is a *i-next configuration* of $\{x_1, \ldots, x_k\}$, if it satisfies the conditions:

(1) $y_i \in N(x_i)$, where $N(x_i)$ denotes the set of neighbours of x_i,
(2) for all j such that $x_j = x_i$: $y_j = y_i$, and
(3) for all j such that $x_j \neq x_i$: $y_j = x_j$.

An i-next configuration of $\{x_1, \ldots, x_k\}$ is a configuration that can result from one step of a random walk from $\{x_1, \ldots, x_k\}$, assuming that the token placed at x_i moves.

An *execution* starting at initial configuration ν_0 is a finite sequence of configurations $(\nu_0, \ldots, \nu_m)$ such that, for each $j = 0, \ldots, m-1$, if $\nu_j = \{x_1, \ldots, x_k\}$ then there exists $i \in \{1, \ldots, k\}$ for which ν_{j+1} is a i-next configuration of ν_j. Configuration ν_m is called the *current configuration* at time step m. An execution $(\nu_0, \ldots, \nu_m)$ is called a *game execution* if ν_m is terminal. In this case the *length* of the game execution is m.

A *strategy* for k tokens on graph G is a probabilistic function S that, given time step m and an execution $\sigma = (\nu_0, \ldots, \nu_m)$ up to time step m, where the current configuration is $\nu_m = \{x_1, \ldots, x_k\}$, returns a probability space over $\{x_1, \ldots, x_k\}$. A strategy defines a scheduler for the token movement.

The *meeting time under strategy S on graph G with initial configuration* $\{x_1, \ldots, x_k\}$, denoted $M_G^S(x_1, \ldots, x_k)$, is the expected length of the game execution assuming that token movement is scheduled by strategy S. That is, each step t with current configuration ν_t satisfies:

(1) $\nu_0 = \{x_1, \ldots, x_k\}$,

(2) given execution $(\nu_0, \ldots, \nu_t)$, where ν_t is nonterminal, the configuration ν_{t+1} is chosen uniformly and randomly from all i-next configurations of ν_t, where i is determined by x_i chosen according to the probability space $S(\nu_0, \ldots, \nu_t)$.

The meeting time under S captures the expected time until all tokens collapse to one assuming each takes a random walk, merging when they collide at a vertex, and the random walks interleave according to strategy S.

The *meeting time* on graph G for initial configuration $\{x_1, \ldots, x_k\}$, denoted $M_G(x_1, \ldots, x_k)$, is the maximum $M_G^S(x_1, \ldots, x_k)$, over all strategies S. A strategy that achieves $M_G(x_1, \ldots, x_k)$ is called *optimal.*

We first prove that, as far as meeting time is concerned, the set of strategies can be restricted without loss of generality to strategies that depend only on the current configuration and are deterministic. Such strategies are called *pure deterministic* strategies (introduced by Tetali and Winkler [4]). A pure deterministic strategy is a function specifying, for each configuration ν, a token position contained in ν. We first eliminate dependence of an optimal strategy on the whole game execution; next we eliminate the need for randomization.

Lemma 1. *For any graph G and every initial configuration, there exists an optimal strategy that depends at each step only on the current configuration of the tokens.*

Proof. Let S be an optimal strategy. Suppose that σ_1 and σ_2 are two executions that have the same current configuration $\nu = \{x_1, \ldots, x_k\}$, and that strategy S returns two different probability spaces over the multiset $\{x_1, \ldots, x_k\}$. Let the probabilities be $\{p_1, \ldots, p_k\}$ in the case of execution σ_1, and $\{q_1, \ldots, q_k\}$ in the case of execution σ_2. Let L_1 and L_2 be the expected length of the continuation of the game after the configuration ν in the case of execution σ_1 and execution σ_2, respectively. Since S is optimal, $L_1 = L_2$. Otherwise, if for example $L_1 > L_2$, then the strategy that returns x_i with probability p_i, $i = 1, \ldots, n$, for execution σ_2 yields longer expected duration of the game than does S. But then the strategy that returns x_i with probability p_i regardless of whether execution σ_1 or σ_2 occurred, yields an expected game execution at least as long as S. $\square$

Theorem 2. *For any graph G and every initial configuration, there exists a pure deterministic optimal strategy.*

Proof. Let S be an optimal strategy. By lemma 1 we can assume that S depends only on the current configuration. Suppose that, for the current configuration $\nu = \{x_1, \ldots, x_k\}$, S returns x_i with probability p_i, $i = 1, \ldots, n$. For each i, let L_i be the expected length of the continuation after configuration ν, given that the token in the position x_i is chosen. Let j be any index satisfying $L_j = \max L_i$. Then the strategy that chooses token j with probability 1 would yield an expected duration of the game that is at least as long as that of strategy S. This new strategy is deterministic. $\square$

3. Upper bounds on the meeting time

Let G be a simple connected graph. To bound the meeting time of k tokens on G under scheduler S, we first convert the meeting times problem on G to a simpler problem, called hitting time, on a larger graph. For k tokens on graph G and a pure deterministic strategy S, the *configuration graph* $G_{k,S} = (V(G_{k,S}), E(G_{k,S}))$ is the directed graph defined as follows:

$V(G_{k,S}) = \{\nu : \nu \text{ is a nonterminal configuration of } k \text{ tokens on } G\} \cup \{TERM\}$
$E(G_{k,S}) = A_1 \cup A_2 \cup A_3$
where

$A_1 = \{(\nu_1, \nu_2) : \nu_1 = \{x_1, \dots, x_k\} \neq TERM \text{ and } \nu_2 \neq TERM \text{ and } S(\nu_1) = x_i \text{ and } \nu_2 \text{ is an } i\text{-next configuration of } \nu_1\}$

$A_2 = \{(\nu, TERM) : \nu = \{x_1, \dots, x_k\} \neq TERM \text{ and } S(\nu) = x_i \text{ and there exists a terminal } i\text{-next configuration of } \nu\}$

$A_3 = \{(TERM, TERM)\}$

Since graph G is connected and undirected, the configuration graph $G_{k,S}$ is connected. The meeting-tokens game with k tokens on graph G under a pure deterministic strategy S corresponds to a random walk of one token on the directed configuration graph $G_{k,S}$. This random walk is a Markov chain with one absorbing state $TERM$; the other states (corresponding to nonterminal configurations) are transient. Order the nodes in the configuration graph $G_{k,S}$ so that all transient nodes precede the absorbing node: $\nu_1, \nu_2, \dots, \nu_t, \nu_{t+1} = TERM$. Let $\mathrm{P}_{k,S}$ be the transition matrix of this Markov chain. Specifically, $\mathrm{P}_{k,S}$ is a $((t+1) \times (t+1))$ matrix satisfying:

$$\mathrm{P}_{k,S}(\nu_l, \nu_m) = \begin{cases} \frac{1}{d(x_i)} & \text{if } (\ \nu_l = \{x_1, \dots, x_k\} \neq TERM \text{ and } \nu_m \neq TERM \text{ and } S(\nu_l) = x_i \text{ and} \\ & \nu_m \text{ is an } i\text{-next configuration of } \nu_l\ ) \text{ or} \\ & (\ \nu_l = \{x_1, \dots, x_k\} \neq TERM \text{ and } \nu_m = TERM \text{ and } S(\nu_l) = x_i \\ & \text{ and there exists a terminal } i\text{-next configuration of } v_l\ ) \\ 1 & \text{if } \nu_l = \nu_m = TERM \\ 0 & \text{otherwise} \end{cases}$$

The *hitting time* of a Markov chain from state ν_i to ν_j, $H(\nu_i, \nu_j)$, is the expected number of steps until a Markov chain starting at ν_i enters ν_j for the first time. Of course, $H(\nu_i, \nu_i) = 0$.

Let $\mathrm{H}_{G_{k,S}}$ denote the vector of hitting times to the absorbing node $TERM$:

$$\mathrm{H}_{G_{k,S}} = (H_{G_{k,S}}(\nu_1, TERM), \dots, H_{G_{k,S}}(\nu_t, TERM), 0)^T$$

Let ξ be a $(t+1)$-element column vector containing all ones except zero in the last position. Then $\mathrm{H}_{G_{k,S}}$ is the unique solution, that has zero on the last position, of the following matrix equation:

$$\mathrm{H}_{G_{k,S}} = \mathrm{P}_{k,S} \cdot \mathrm{H}_{G_{k,S}} + \xi \tag{1}$$

The expected number of steps to reach node $TERM$ in $G_{k,S}$ starting from a node ν is equal to the expected number of steps to reach a terminal configuration in G starting from configuration ν, assuming the tokens move according to the strategy S.

Observation 3. *If S is a pure deterministic strategy then for any configuration $\{x_1, \ldots, x_k\}$*

$$M_G^S(x_1, \ldots, x_k) = H_{G_{k,S}}(\{x_1, \ldots, x_k\}, TERM)$$

For any vectors $\mathrm{v} = (v_1, \ldots, v_n)$ and $\mathrm{w} = (w_1, \ldots, w_n)$, define $\mathrm{v} \geq \mathrm{w}$ if and only if $v_i \geq w_i$, for all $i \in \{1, \ldots, n\}$. Observation 3 leads to a technique for proving upper and lower bounds on the meeting time.

Lemma 4. *Let $G_{k,S}$ be the configuration graph for k tokens on an undirected connected graph G under a pure deterministic strategy S, and let $P_{k,S}$ be the corresponding transition matrix. Let $W : V(G_{k,S}) \to \mathbb{R}$ such that $W(TERM) \geq 0$. If the vector $W = (W(\nu_1), \ldots, W(\nu_t), W(TERM))^T$, where $\{\nu_1, \ldots, \nu_t\}$ are nonterminal configurations, satisfies*

$$W \geq P_{k,S} \cdot W + \xi \tag{2}$$

then, for any configuration $\nu = \{x_1, \ldots, x_k\}$, $W(\nu) \geq M_G^S(x_1, \ldots, x_k)$.

Proof. By observation 3, it suffices to prove that for any configuration ν, $W(\nu) \geq H_{G_{k,S}}(\nu, TERM)$. Suppose vector W satisfies inequality 2. Subtracting equation 1 yields $\mathrm{W} - \mathrm{H}_{G_{k,S}} \geq \mathrm{P}_{k,S} \cdot (\mathrm{W} - \mathrm{H}_{G_{k,S}})$. Matrix $\mathrm{P}_{k,S}$ is nonnegative, so for any vectors u and w, $\mathrm{u} \geq \mathrm{w}$ implies that $\mathrm{P}_{k,S} \cdot \mathrm{u} \geq \mathrm{P}_{k,S} \cdot \mathrm{w}$. Therefore $\mathrm{P}_{k,S} \cdot (\mathrm{W} - \mathrm{H}_{G_{k,S}}) \geq \mathrm{P}_{k,S}^2 \cdot (\mathrm{W} - \mathrm{H}_{G_{k,S}})$. By induction, $\forall n \in \mathbb{N} : \mathrm{W} - \mathrm{H}_{G_{k,S}} \geq \mathrm{P}_{k,S}^n \cdot (\mathrm{W} - \mathrm{H}_{G_{k,S}})$. Let $\mathrm{w} = (w_1, \ldots, w_t, w_{t+1})^T$ denote the vector $(\mathrm{W} - \mathrm{H}_{G_{k,S}})$. The last inequality states that $\forall n \in \mathbb{N} : \mathrm{w} \geq \mathrm{P}_{k,S}^n \cdot \mathrm{w}$. Matrix $\mathrm{P}_{k,S}^n$ contains probabilities of transition from one node of $G_{k,S}$ to another in exactly n steps. Since $\nu_{t+1} = TERM$ is the only absorbing state, the other states being transient, $\mathrm{P}_{k,S}^n$ tends, as n tends to infinity, to the matrix having the last column one and all other entries zero. The inequality holds in the limit as well. Therefore

$$\begin{pmatrix} w_1 \\ \vdots \\ w_{t+1} \end{pmatrix} \geq \begin{pmatrix} 0 & \cdots & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} w_1 \\ \vdots \\ w_{t+1} \end{pmatrix}$$

The right side of the above inequality evaluates to the constant vector $(w_{t+1}, \cdots, w_{t+1})^T$ so the inequality reduces to: $\forall i = 1, \ldots, t+1 : w_i \geq w_{t+1}$. But $w(t+1) = W(TERM) \geq 0$. Therefore for each i, $w(i) = W(\nu_i) - H_{G_{k,S}}(\nu_i, TERM) \geq 0$, and the lemma follows. $\square$

A similar proof to that of lemma 4, can be used to establish a lower bound on meeting times.

Lemma 5. *Let $G_{k,S}$ be the configuration graph for k tokens on an undirected connected graph G and a pure deterministic strategy S. Let $W : V(G_{k,S}) \to \mathbb{R}$ such that $W(TERM) \leq 0$. If the* *vector* $W = (W(\nu_1), \ldots, W(\nu_t), W(TERM))^T$, *where* $\{\nu_1, \ldots, \nu_t\}$ *are nonterminal configurations, satisfies*

$$W \leq P_{k,S} \cdot W + \xi \tag{3}$$

then, for any configuration $\nu = \{x_1, \ldots, x_k\}$, $W(\nu) \leq M_G^S(x_1, \ldots, x_k)$.

Let $(x_1, \ldots, x_k)$ be a sequence of k vertices of G (that is, a configuration with an imposed order on the positions of tokens). Let $2 \leq l \leq k$. An *l-segment* of $(x_1, \ldots, x_k)$ is a contiguous subsequence $(x_i, \ldots, x_{i+l-1})$ of some cyclic permutation of $(x_1, \ldots, x_k)$. (Indices are assumed to be reduced modulo k to yield indices in the range from 1 to k.) The following theorem bounds the meeting time of k tokens $\{x_1, \ldots, x_k\}$ in terms of the meeting times of l tokens over all l-segments of $(x_1, \ldots, x_k)$.

Theorem 6. *For any undirected connected graph* G, *any* $(x_1, \ldots, x_k) \in (V(G))^k$, *and for any* $2 \leq l \leq k$,

$$M_G(x_1, \ldots, x_k) \leq \frac{1}{l} \sum_{j=1}^{k} M_G(x_j, x_{j+1}, \ldots, x_{j+l-1})$$

Proof. Let G be any undirected connected graph. Let S be an optimal pure deterministic strategy for k tokens on vertices of graph G. For a fixed ordering of tokens, define $W : V(G_{k,S}) \to \mathbb{R}$ by:

$$W(\nu) = \begin{cases} \frac{1}{l}\sum_{i=1}^{k} M_G(x_i, x_{i+1}, \ldots, x_{i+l-1}) & \text{if } \nu = \{x_1, \ldots, x_k\} \\ 0 & \text{if } \nu = TERM \end{cases}$$

For notational convenience extend the definition of W to terminal configurations: $W(x_1, \ldots, x_k) = \frac{1}{l}\sum_{i=1}^{k} M_G(x_i, x_{i+1}, \ldots, x_{i+l-1}) = 0$ if $x_1 = \ldots = x_k$. Notice that $W(\nu) = W(TERM)$ when ν is terminal.

By lemma 4 it is enough to prove that W satisfies inequality 2. Let $W(\nu)[y/x_i]$ denote function W evaluated at the configuration obtained by replacing each occurrence of x_i by y in configuration ν. By the definition of matrix $P_{k,S}$, inequality 2 is reduced to:

$$W(x_1, \ldots, x_k) \geq \frac{1}{d(x_i)} \sum_{y \in N(x_i)} W(x_1, \ldots, x_k)[y/x_i] + 1 \tag{4}$$

when $\{x_1, \cdots, x_k\}$ is a nonterminal configuration, and $S(x_1, \ldots, x_k) = x_i$.

Let $\nu = \{x_1, \ldots, x_k\}$ be a nonterminal configuration and suppose $S(x_1, \ldots, x_k) = x_i$. Then, by definition of W, establishing inequality 4 is equivalent to showing the following inequality.

$$\sum_{j=1}^{k} M_G(x_j, \ldots, x_{j+l-1}) \geq l + \sum_{j=1}^{k} \frac{1}{d(x_i)} \sum_{y \in N(x_i)} M_G(x_j, \ldots, x_{j+l-1})[y/x_i]$$

Let A be the set of indices m such that the sequence $(x_m, \ldots, x_{m+l-1})$ contains at least one and at most $(l-1)$ occurrences of x_i. The theorem now follows by observing:

(1) for $j \notin A$:

$$M_G(x_j,\ldots,x_{j+l-1}) = \frac{1}{d(x_i)} \sum_{y\in N(x_i)} M_G(x_j,\ldots,x_{j+l-1})[y/x_i]$$

(2) for $j \in A$:

$$M_G(x_j,\ldots,x_{j+l-1}) \geq \frac{1}{d(x_i)} \sum_{y\in N(x_i)} M_G(x_j,\ldots,x_{j+l-1})[y/x_i] + 1$$

with equality achieved when moving the token at x_i is optimal for configuration $\{x_j,\ldots,x_{j+l-1}\}$;

(3) $|A| \geq l$ since $(x_1,\ldots,x_k)$ is nonterminal, and it contains at least one occurrence of x_i.

□

Observation 7. *Since theorem 6 holds for any $2 \leq l \leq k$, and for any order of the positions in configuration $\{x_1,\ldots,x_k\}$, the bound can be optimized by taking the value of l and the order on $\{x_1,\ldots,x_k\}$ that minimizes the value of expression $\frac{1}{l}\sum_{j=1}^{k} M_G(x_j,x_{j+1},\ldots,x_{j+l-1})$.*

In their additional remarks Tetali and Winkler [4] sketch a proof that

$$M_G(x_1,\ldots,x_k) \leq \frac{1}{k-1}\sum_{i<j} M_G(x_i,x_j)$$

The next theorem 8 generalizes this result. Theorem 8 can be proved directly using lemma 4 (see Warpechowska-Gruca [5]). Here we prove it by showing that the bound provided by theorem 6, together with observation 7 is in general even stronger.

Theorem 8. *For any undirected connected graph G, any configuration $\{x_1,\ldots,x_k\}$ and for any $2 \leq l \leq k$,*

$$M_G(x_1,\ldots,x_k) \leq \frac{1}{\binom{k-1}{l-1}} \sum_{\{y_1,\ldots,y_l\}\subseteq\{x_1,\ldots,x_k\}} M_G(y_1,\ldots,y_l)$$

Proof. Without loss of generality assume that $(x_1,\ldots,x_k)$ is a permutation of $\{x_1,\ldots,x_k\}$ that minimizes the value of the expression $\sum_{i=1}^{k} M_G(x_i,x_{i+1},\ldots,x_{i+l-1})$. We prove that

$$\frac{1}{l}\sum_{i=1}^{k} M_G(x_i,x_{i+1},\ldots,x_{i+l-1}) \leq \frac{1}{\binom{k-1}{l-1}} \sum_{\{y_1,\ldots,y_l\}\subseteq\{x_1,\ldots,x_k\}} M_G(y_1,\ldots,y_l)$$

The theorem then follows from theorem 6 and observation 7.

Let $\sigma_1,\ldots,\sigma_{k!}$ be the permutations of $\{x_1,\ldots,x_k\}$, where $\sigma_j = (\sigma_j(1),\ldots,\sigma_j(k))$ and $\sigma_1 = (x_1,\ldots,x_k)$. Construct a matrix, A, of meeting times of l-subsets, with $k!$ columns and k rows, where the columns of A correspond to the permutations $\sigma_1,\ldots,\sigma_{k!}$, and the rows of A in column j correspond to the k cyclic permutations of σ_j. Specifically, $(\mathrm{A})_{ij} = M_G(\sigma_j(i),\sigma_j(i+1),\ldots,\sigma_j(i+l-1))$. Then the sum of elements in the first column of matrix A is less than or equal to the sum of elements in any column, and thus it is less than or equal to the average column sum. Therefore

$$\sum_{i=1}^{k} M_G(x_i,x_{i+1},\ldots,x_{i+l-1}) \leq \frac{\sum_{j=1}^{k!}\sum_{i=1}^{k}(\mathrm{A})_{ij}}{k!}$$

However, for any l-subset $\{y_1, \ldots, y_l\} \subseteq \{x_1, \ldots, x_k\}$, $M_G(y_1, \ldots, y_l)$ appears in matrix A exactly $k \cdot l! \cdot (k-l)!$ times. Thus

$$\begin{aligned}\sum_{i=1}^{k} M_G(x_i, x_{i+1}, \ldots, x_{i+l-1}) &\leq \frac{k \cdot l! \cdot (k-l)! \cdot \sum_{\{y_1,\ldots,y_l\}\subseteq\{x_1,\ldots,x_k\}} M_G(y_1, \ldots, y_l)}{k!} \\ &= \frac{l}{\binom{k-1}{l-1}} \sum_{\{y_1,\ldots,y_l\}\subseteq\{x_1,\ldots,x_k\}} M_G(y_1, \ldots, y_l)\end{aligned}$$

□

Tetali and Winkler[4] show that the expected time before k tokens collapse to one is bounded by k times the maximum meeting time for two tokens in the graph. The following corollary generalizes and strengthens their result.

Corollary 9. *Let M_G^k denote the maximum, over all initial configurations $\{x_1, \ldots, x_k\}$, of the meeting time $M_G(x_1, \ldots, x_k)$. Then for any natural number $2 \leq l \leq k$*

$$M_G^k \leq \frac{k}{l} \cdot M_G^l$$

Proof. By theorem 8,

$$M_G(x_1, \ldots, x_k) \leq \frac{\sum_{\{y_1,\ldots,y_l\}\subseteq\{x_1,\ldots,x_k\}} M_G(y_1, \ldots, y_l)}{\binom{k-1}{l-1}}$$

Therefore

$$M_G(x_1, \ldots, x_k) \leq \frac{\binom{k}{l} M_G^l}{\binom{k-1}{l-1}} = \frac{k}{l} \cdot M_G^l$$

Since this holds for any initial configuration of k tokens, the corollary follows. □

4. Meeting time on rings

Theorem 6 establishes an upper bound on the meeting time for k tokens in terms of the meeting times of l tokens, $2 \leq l \leq k$. We show that this upper bound is tight when $l = 2$ in the case of rings.

Let R be a ring consisting of n vertices. Let $(x_1, \ldots, x_k)$ be a sequence of *cyclically ordered positions* of k tokens on R. That is, for each $i = 1, \ldots, k$: the token x_i is the first token encountered travelling clockwise from x_{i-1} (indices are reduced modulo k to yield indices in the range from 1 to k). Let d_i, where $i = 1, \ldots, k$, be the number of edges of R between tokens x_i and x_{i+1}.

Theorem 10. *On ring R, for any cyclically ordered positions $(x_1, \ldots, x_k)$ with distances between successive tokens $d_1, \ldots, d_k$:*

$$M_R(x_1, \ldots, x_k) = \sum_{i<j} d_i \cdot d_j$$

Proof. Let S be an optimal pure deterministic strategy for the k tokens $(x_1,\ldots,x_k)$ on ring R. Let $R_{k,S}$ be a configuration graph for k tokens on ring R with strategy S. Define $W : V(R_{k,S}) \to \mathbb{R}$ by:

$$W(\nu) = \begin{cases} \sum_{i<j} d_i \cdot d_j & \text{if } \nu = \{x_1,\ldots,x_k\} \text{ is a nonterminal configuration,} \\ 0 & \text{if } \nu = TERM \end{cases}$$

For notational convenience extend the definition of W to terminal configurations by $W(x_1,\ldots,x_k) = \sum_{i<j} d_i \cdot d_j = 0$ if $x_1 = \ldots = x_k$. Notice that $W(\nu) = W(TERM)$ when ν is terminal.

Since the hitting time equation 1 has a unique solution with zero in the last position, by observation 3, it is enough to prove that vector W $= (W(\nu_1),\ldots,W(\nu_t),0)$ satisfies equation 1. We need to prove that if configuration $\{x_1,\ldots,x_k\}$ is nonterminal and $S(x_1,\ldots,x_k) = x_i$ then

$$W(x_1,\ldots,x_k) = \frac{1}{d(x_i)} \sum_{y \in N(x_i)} W(x_1,\ldots,x_k)[y/x_i] + 1$$

Suppose configuration $\{x_1,\ldots,x_k\}$ is nonterminal. Since W does not depend on the cyclic shift of the token numbers, we can assume without loss of generality that $S(x_1,\ldots,x_k) = x_1$. Suppose t tokens are placed at x_1. Therefore $x_1 = \ldots = x_t$ and $d_1 = \ldots = d_{t-1} = 0$. To simplify notation, let x denote the position x_1. On the ring, x has two neighbours y_1 and y_2, with the pair of distances to the tokens x_{t+1} and x_k equal (d_t+1, d_k-1) and (d_t-1, d_k+1), respectively. We need to show:

$$W(x,\ldots,x,x_{t+1},\ldots,x_k) = \frac{1}{2}W(y_1,\ldots,y_1,x_{t+1},\ldots,x_k) + \frac{1}{2}W(y_2,\ldots,y_2,x_{t+1},\ldots,x_k) + 1$$

Since $d_1 = \ldots = d_{t-1} = 0$, $W(x,\ldots,x,x_{t+1},\ldots,x_k) = \sum_{t \le i < j \le k} d_i d_j + 1$. Thus

$$\begin{aligned} & W(x,\ldots,x,x_{t+1},\ldots,x_k) \\ = \; & \sum_{t \le i < j \le k} d_i d_j = \sum_{t<j<k} d_t d_j + \sum_{t<i<j<k} d_i d_j + \sum_{t<i<k} d_i d_k + d_t d_k \\ = \; & \frac{1}{2}[\sum_{t<j<k} (d_t+1) d_j + \sum_{t<i<j<k} d_i d_j + \sum_{t<i<k} d_i (d_k - 1) + (d_t+1)(d_k-1)] \\ & + \frac{1}{2}[\sum_{t<j<k} (d_t-1) d_j + \sum_{t<i<j<k} d_i d_j + \sum_{t<i<k} d_i (d_k + 1) + (d_t-1)(d_k+1)] + 1 \\ = \; & \frac{1}{2}W(y_1,\ldots,y_1,x_{t+1},\ldots,x_k) + \frac{1}{2}W(y_2,\ldots,y_2,x_{t+1},\ldots,x_k) + 1 \end{aligned}$$

□

Corollary 11. *On ring R, for any cyclically ordered positions $(x_1,\ldots,x_k)$ of tokens:*

$$M_R(x_1,\ldots,x_k) = \frac{1}{2}\sum_{i=1}^{k} M_R(x_i, x_{i+1}) \tag{5}$$

Proof. Let the distances between successive tokens be $d_1,\ldots,d_k$. By straightforward calculations, the meeting time of two tokens on a ring with the arcs connecting them equal to a and b is equal to the product $a \cdot b$. Evaluating the righthand side of the equation 5:

$$\frac{1}{2}\sum_{i=1}^{k} M_R(x_i, x_{i+1}) = \frac{1}{2}\sum_{i=1}^{k}[d_i(d_1 + \ldots + d_{i-1} + d_{i+1} + \ldots + d_k)] = \frac{1}{2}\sum_{i<j} 2d_i d_j = M_R(x_1, \ldots, x_k)$$

□

Notice that theorem 6 with observation 7 for $l = 2$ gives a bound for $M_G(x_1, \ldots, x_k)$ equal to that of corollary 11.

5. Concluding remarks and open problems

Theorem 6 bounds the meeting time of several random walks on an undirected connected graph by a function of the meeting times of fewer random walks. Corollary 11 proves that this bound is tight in the case of rings. Corollary 9 bounds the maximum, over initial configurations, of the meeting times of k random walks in terms of maximum meeting time of l random walks, for any $2 \leq l \leq k$.

The idea of expressing several random walks as one random walk on the configuration graph can be applied to other games or special classes of graphs. Warpechowska-Gruca [5] studies two variants of the meeting-tokens game that employ different collision and termination rules. One game ends at the first collision of tokens. In the other game, the tokens that have collided are not allowed to move further, and the game ends when there is no move possible. An extensive case study for these games on rings is presented with formulas and bounds for the meeting times in several cases. Warpechowska-Gruca [5] discusses the meeting times in the class of *hitting-time-symmetric graphs*, that is undirected connected graphs G, such that for any $x, y \in G : H_G(x, y) = H_G(y, x)$. The bound in theorem 6 is proven tight in the case of two or three random walks on a hitting-time-symmetric graph.

The upper bound on the expected duration of the meeting-tokens game is shown to be tight only in some special cases. It still remains to determine better upper bounds for most graphs. Furthermore, for other variants of the meeting-tokens game on general graphs, our only bounds are those implied by the bound presented in theorem 6. Also, we did not obtain any non-trivial lower bound on the minimum, over interleaving strategies, of the meeting time of several random walks.

The question of the initial configuration that maximizes the meeting time is surprisingly difficult even if the number of tokens is restricted to two, and the class of graphs is the class of trees. There are examples (see [5]) that show that the initial placement of two tokens on most distant vertices, or *remote* vertices (defined in Tetali and Winkler [4]), or vertices maximizing the hitting time, fails to maximize the meeting time.

Similarly, not much is known about strategies that maximize meeting time. Consider the game that ends with the first collision. The strategy that never moves a token adjacent to another token unless there is no other choice seemed to be a good heuristic to maximize the expected duration of the game. Warpechowska-Gruca [5] shows by way of a counterexample that this strategy is not necessarily optimal.

Finally, the exact characterization of the class of hitting-time-symmetric graphs is still open. Our conjecture is that it includes the class of regular graphs, and is included in the class of distance-regular graphs. However, our partial results [5] only establish that there is no equality in either case.

References

1. D. Aldous. Meeting times for independent Markov chains. *Stochastic Processes and their Applications*, 38:185–193, 1991.
2. D. Coppersmith, P. Tetali, and P. Winkler. Collisions among random walks. Manuscript, 1992.
3. A. Israeli and M. Jalfon. Token management schemes and random walks yield self stabilizing mutual exclusion. In *Proceedings of the Nineth Annual ACM Symposium on Principles of Distributed Computing*, pages 119–131, 1990.
4. P. Tetali and P. Winkler. On a random walk problem arising in self-stabilizing token management. In *Proceedings of the Tenth Annual ACM Symposium on Principles of Distributed Computing*, pages 273–280, 1991.
5. J. Warpechowska-Gruca. Meeting times of random walks on graphs. Technical Report 94/535/04, University of Calgary, 1994.