

Randomized Rendez-Vous with Limited Memory

EVANGELOS KRANAKIS

Carleton University

and

DANNY KRIZANC

Wesleyan University

and

PAT MORIN

Carleton University

We present a tradeoff between the expected time for two identical agents to rendez-vous on a synchronous, anonymous, oriented ring and the memory requirements of the agents. In particular, we show there exists a $2t$ state agent, which can achieve rendez-vous on an n node ring in expected time $O(n^2/2^t + 2^t)$ and that any $t/2$ state agent requires expected time $\Omega(n^2/2^t)$. As a corollary we observe that $\Theta(\log \log n)$ bits of memory are necessary and sufficient to achieve rendez-vous in linear time.

Categories and Subject Descriptors: H.3.4 [Systems and Software]: Distributed Systems; G.3 [Probability and Statistics]: —*Probabilistic Algorithms*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Rendez-Vous, Distributed Computing

1. INTRODUCTION

The problem of *rendez-vous* (the gathering of agents widely dispersed in some domain at a common place and time) has been studied under many guises and in many settings [Alpern and Gal 2003; Marco et al. 2005; Das et al. 2008; Dessmark et al. 2003; Dobrev et al. 2004; Flocchini et al. 2004; Flocchini et al. 2004; Gasieniec et al. 2006; Kowalski and Pelc 2004; Kowalski and Malinowski 2006; Kranakis et al. 2003; Marco et al. 2006; Kranakis et al. 2006; Roy and Dudek 2001; Sawchuk 2004; Suzuki and Yamashita 1999; Yu and Yung 1996]. (See Kranakis et al. 2006 for a survey of recent results.) In this paper we consider the problem of autonomous mobile software agents gathering in a distributed network. This is a fundamental operation useful in such applications as web-crawling, distributed search, meet-

Authors' addresses: E. Kranakis and P. Morin, School of Computer Science, Carleton University, 1125 Colonel By Drive, Ottawa, ON, CANADA K1S 5B6.

D. Krizanc, Computer Science Group, Mathematics Department, Wesleyan University, Middletown, CT 06459 USA.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0000-0000/20YY/0000-0001 \$5.00

ing scheduling, etc. In particular, we study the problem of two identical agents attempting to rendez-vous on a synchronous anonymous ring.

1.1 The Model

Network rendez-vous problems have many different parameters, including the topology of the underlying network, the agents' knowledge (if any) of the network, labellings (if any) of the edges and vertices of the network, the capabilities of the agents (unlimited memory, access to random bits, the ability to clone, the ability to leave messages at vertices), the types of agents (identical or not), the synchrony of the agents, and the reliability of the agents.

We consider the standard model of a synchronous anonymous oriented n -node ring [Santoro 2006]. The nodes are assumed to have no identities, the computation proceeds in synchronous steps and the edges of the ring are labelled `clockwise` and `counterclockwise` in a consistent fashion. We model the agents as identical probabilistic finite automata $A = \langle S, \delta, s_0 \rangle$ where S is the set of states of the automata including s_0 the initial state and the special state `halt`, and $\delta : S \times C \times P \rightarrow S \times M$ where $C = \{H, T\}$ represents a random coin flip, $P = \{\text{present}, \text{notpresent}\}$ represents a predicate indicating the presence of the other agent at a node, and $M = \{-1, 0, +1\}$ represents the potential moves the agent may make, $+1$ representing clockwise, -1 counterclockwise and 0 stay at the current node. During each synchronous step, depending upon its current state, the answer to a query for the presence of the other agent, and the value of an independent random coin flip with probability of `heads` equal to $1/2$, the agent uses δ in order to change its state and either move across the edge labelled `clockwise`, move across the edge labelled `counterclockwise` or stay at the current node. We assume that the agent halts once it detects the presence of the other agent at a node. Rendez-vous occurs when both agents halt on the same node. The complexity measures we are interested in are the expected time (the number of synchronous steps) to rendez-vous (where the expectation is taken over all sequences of coin flips of the two agents) and the size ($|S|$) or memory requirement ($\log_2 |S|$) of the agents.

We assume for simplicity that n , the size of the ring is an even number and that the two agents start an even distance apart. This avoids the possibility that the two agents simultaneously cross the same edge in opposite directions without achieving a rendez-vous. This assumption does not significantly affect our (asymptotic) results. We can achieve a similar rendez-vous time without this assumption by having agents use a coin periodically to determine whether they should pause, for one unit of time, at their currently location.

1.2 Related Work and New Results

A number of researchers have observed that using random walks one can design $O(1)$ state agents that will rendez-vous in polynomial $O(n^3)$ number steps on any network [Coppersmith et al. 1993]. For the ring the expected time for two random walks to meet is easily shown to be $O(n^2)$. (See Reference [Kranakis and Krizanc 2007] for an example proof of this fact.)

This expected time bound can be improved by considering the following strategy. Repeat the following until rendez-vous is achieved: flip a fair coin and walk $n/2$ steps clockwise if the result is heads, $n/2$ steps counterclockwise if the result is

tails. If the two agents choose different directions (which they do with probability $1/2$) then they will rendez-vous (at least in the case where they start at an even distance apart). The expected time to rendez-vous in this case satisfies

$$T \leq (1/2)(n/2) + (1/2)(n/2 + T)$$

and is therefore at most $3n/2$. Alpern refers to this strategy as *Coin Half Tour* and studies it in detail [Alpern 1995]. A variant of Coin Half Tour, in which each agent either travels $n - 1$ steps in the same direction or remains stationary for $n - 1$ time units, was studied by Alpern *et al.* [Alpern et al. 1999]. When the agents have uniformly distributed starting positions this strategy achieves an expected meeting time, for odd n , of $n - \Theta(1)$. The case for even values of n is complicated by the fact that the agents are more likely to pass each other without meeting. In this case, the agents can still rendez-vous in $1.254122768n + O(1)$ expected time. A generalization of the above strategy, in which each agent either searches exhaustively for $2n$ steps or waits for $2n$ steps, allows two agents to rendez-vous in any n vertex graph in expected time at most $4n$ [Alpern et al. 1999, Section 4].

Note that these strategies require the agent to count up to at least $n/2$ and thus require $\Omega(n)$ states or $\Omega(\log n)$ bits of memory. The main result of this paper is that this memory requirement can be reduced to $O(\log \log n)$ bits while still achieving rendez-vous in $O(n)$ expected time, and this is optimal.

Below we show a tradeoff between the (memory) size of the agents and the time required for them to rendez-vous. We prove there exists a $2t$ state algorithm, which can achieve rendez-vous on an n node ring in expected time $O(n^2/2^t + 2^t)$ and that any $t/2$ state algorithm requires expected time $\Omega(n^2/2^t)$. As a corollary we observe that $\Theta(\log \log n)$ bits of memory are necessary and sufficient to achieve rendez-vous in linear time.

A preliminary version of these results was presented at the 8th Latin American Theoretical Informatics Conference (LATIN 2008) [Kranakis et al. 2008]. Section 2 contains some preliminary results, Section 3 our upper bound and Section 4 the lower bound.

2. PRELIMINARIES

2.1 Martingales, Stopping Times, and Wald's Equations

In this section, we review some results on stochastic processes that are used several times in our proofs. The material in this section is based on the presentation in Ross' textbook [Ross 2002, Chapter 6]. Let $X = X_1, X_2, X_3, \dots$ be a sequence of random variables and let $Q = Q_1, Q_2, Q_3, \dots$ be a sequence of random variables in which Q_i is a function of X_1, \dots, X_i . Then we say that Q is a *martingale with respect to X* if, for all i , $E[|Q_i|] < \infty$ and $E[Q_{i+1} | X_1, \dots, X_i] = Q_i$.

A positive integer-valued random variable T is a *stopping time* for the sequence X_1, X_2, X_3, \dots if the event $T = i$ is determined by the values X_1, \dots, X_i . In particular, the event $T = i$ is independent of the values X_{i+1}, X_{i+2}, \dots . Some of our results rely on the Martingale Stopping Theorem:

THEOREM 1 MARTINGALE STOPPING THEOREM. *If Q_1, Q_2, Q_3, \dots is a martingale with respect to X_1, X_2, X_3, \dots and T is a stopping time for X_1, X_2, X_3, \dots then*

$$E[Q_T] = E[Q_1]$$

provided that at least one of the following holds:

- (1) Q_i is uniformly bounded for all $i \leq T$,
- (2) T is bounded, or
- (3) $E[T] < \infty$ and there exists an $M < \infty$ such that

$$E[|Q_{i+1} - Q_i| \mid X_1, \dots, X_i] < M .$$

If X_1, X_2, X_3, \dots is a sequence of i.i.d. random variables with expected value $E[X] < \infty$ and variance $\text{var}(X) < \infty$ then the sequence $Q_i = \sum_{j=1}^i (X_j - E[X])$ is a Martingale and the assumption that $\text{var}(X) < \infty$ implies that this sequence satisfies Condition 3 of Theorem 1, so we obtain *Wald's Equation for expectation*:

$$E\left[\sum_{i=1}^T X_i\right] = E[T] \cdot E[X] \quad (1)$$

whenever T is a stopping time for X_1, X_2, X_3, \dots . Similarly, we can derive a version of *Wald's Equation for variance* by considering the martingale

$$Q_i = \left(\sum_{j=1}^i (X_j - E[X])\right)^2 - i \cdot \text{var}(X)$$

to obtain

$$\text{var}\left(\sum_{i=1}^T X_i\right) = E\left[\left(\sum_{i=1}^T (X_i - E[X_i])\right)^2\right] = E[T] \cdot \text{var}(X) . \quad (2)$$

2.2 A Lemma on Random Walks

Let $X_1, X_2, X_3, \dots \in \{-1, +1\}$ be independent random variables with

$$\Pr\{X_i = -1\} = \Pr\{X_i = +1\} = 1/2$$

and let $S_i = \sum_{j=1}^i X_j$. The sequence S_1, S_2, S_3, \dots is a *simple random walk* on the line, where each X_i represents a step to the left ($X_i = -1$) or a step to the right ($X_i = +1$). Define the *hitting time* h_m as

$$h_m = \min\{i : |S_i| = m\} ,$$

which is the number of steps in a simple random walk before it travels a distance of m from its starting location. The following result is well-known (see, e.g., Reference [Mitzenmacher and Upfal 2005]):

LEMMA 1. $E[h_m] = m^2$.

Applying Markov's Inequality with Lemma 1 yields the following useful corollary

COROLLARY 1. $\Pr\{\max\{|S_i| : i \in \{1, \dots, 2m^2\}\} \geq m\} \geq 1/2$.

In other words, Corollary 1 says that, at least half the time, at some point during the first $2m^2$ steps of a simple random walk, the walk is at distance m from its starting location.

Let Y_1, \dots, Y_m be i.i.d. non-negative random variables with finite expectation $r = E[Y_i]$, independent of X_1, \dots, X_m , and with the property that

$$\Pr\{Y_i \geq \alpha r\} \geq 1/2$$

for some constant $\alpha > 0$. The following lemma considers a modified random walk in which the i th step is of length Y_i :

LEMMA 2. *Let X_1, \dots, X_m and Y_1, \dots, Y_m be defined as above. Then there exists constants $\beta, \kappa > 0$ such that*

$$\Pr \left\{ \max \left\{ \left| \sum_{i=1}^{m'} X_i Y_i \right| : m' \in \{1, \dots, m\} \right\} \geq \beta r \sqrt{m} \right\} \geq \kappa .$$

PROOF. We will define 3 events E_1, E_2, E_3 such that $\Pr\{E_1 \cap E_2 \cap E_3\} \geq 1/8$ and, if E_1, E_2 , and E_3 all occur, then there exists a value $m' \in \{1, \dots, m\}$ such that $\left| \sum_{i=1}^{m'} X_i Y_i \right| \geq \alpha r \sqrt{m}/2^{3/2}$. This will prove the lemma for $\kappa = 1/8$ and $\beta = \alpha/2^{3/2}$.

Let E_1 be the event that there exists a value $m' \in \{1, \dots, m\}$ such that

$$\left| \sum_{i=1}^{m'} X_i \right| \geq \sqrt{m/2} .$$

By Corollary 1, $\Pr\{E_1\} \geq 1/2$. Assume E_1 occurs and, without loss of generality, assume $\sum_{i=1}^{m'} X_i > 0$.

Let $I^+ = \{i \in \{1, \dots, m'\} : X_i = +1\}$ and $I^- = \{1, \dots, m'\} \setminus I^+$. We further partition I^+ into two sets I_1^+ and I_2^+ where I_1^+ contains the smallest $|I^-|$ elements of I^+ and I_2^+ contains the remaining elements. Note that, with these definitions, $|I_1^+| = |I^-|$ and that $|I_2^+| = \sum_{i=1}^{m'} X_i$. Let E_2 be the event that

$$\sum_{i \in I_1^+} X_i Y_i + \sum_{i \in I^-} X_i Y_i \geq 0$$

which is equivalent to $\sum_{i \in I_1^+} Y_i \geq \sum_{i \in I^-} Y_i$ and observe that, by symmetry, $\Pr\{E_2|E_1\} \geq 1/2$.

Finally, let E_3 be the event

$$\sum_{i \in I_2^+} X_i Y_i \geq \alpha r |I_2^+|/2$$

To bound $\Pr\{E_3|E_1 \cap E_2\}$, let $T = |\{i \in I_2^+ : Y_i \geq \alpha r\}|$ and observe that $T \geq |I_2^+|/2$ implies E_3 . Now, T is a binomial($|I_2^+|, p$) random variable for $p \geq 1/2$ so its median value is at least $p|I_2^+| \geq |I_2^+|/2$ and therefore $\Pr\{E_3|E_1 \cap E_2\} \geq \Pr\{T \geq |I_2^+|/2\} \geq 1/2$.

We have just shown that $\Pr\{E_1 \cap E_2 \cap E_3\} \geq 1/8$. To complete the proof we observe that, if E_1, E_2 and E_3 occur then

$$\sum_{i=1}^{m'} X_i Y_i = \sum_{i \in I_1^+} X_i Y_i + \sum_{i \in I^-} X_i Y_i + \sum_{i \in I_2^+} X_i Y_i$$

$$\begin{aligned}
&\geq \sum_{i \in I_2^+} X_i Y_i \\
&\geq \alpha r |I_2^+| / 2 \\
&\geq \alpha r \sqrt{m} / 2^{3/2} .
\end{aligned}$$

□

2.3 An Approximate Counter

In the previous section we have shown that, if we can generate random variables Y_i that are frequently large, then we can speed up the rate at which a random walk moves away from its starting location. In this section we consider how to generate these frequently-large random variables. Consider a random variable Y generated by the following algorithm:

BIGRAND(t)

```

1:  $Y \leftarrow C \leftarrow 0$ 
2: while  $C < t$  do
3:    $Y \leftarrow Y + 1$ 
4:   if a coin toss comes up heads then
5:      $C \leftarrow C + 1$ 
6:   else
7:      $C \leftarrow 0$ 
8:   end if
9: end while
10: return  $Y$ 

```

LEMMA 3. *Let Y be the output of Algorithm BIGRAND(t). Then*

- (1) $E[Y] = 2^t(2 - 1/2^{t-1})$ and
- (2) $\text{var}(Y) \leq 2^{t+1} < \infty$
- (3) $\Pr\{Y \geq E[Y]/2\} \geq 1/2$.

PROOF. To compute the expected value of Y we observe that the algorithm begins by tossing a sequence of $i - 1$ heads and then either (a) returning to the initial state if the i th coin toss is a tail or (b) terminating if $i = 2^t$. The first case occurs with probability $1/2^i$ and the second case occurs with probability $1/2^t$. Call the interval between consecutive visits to the initial state a *round*. The number of rounds, T , is a geometric($1/2^t$) random variable and therefore $E[T] = 2^t$. The length X_i of the i th round is dominated¹ by a geometric($1/2$) random variable and its expectation and variance are easily shown to satisfy $E[X_i] = 2 - 1/2^{t-1}$ and $\text{var}(X_i) \leq 2$. Therefore, Parts 1 and 2 of the lemma follow from Wald's Equations for expectation and variance of $Y = \sum_{i=1}^T X_i$, respectively.

To prove the second part of the lemma, consider only the random variable T that counts the number of rounds. Since T is a geometric($1/2^t$) random variable, its median is $\lceil -\log 2 / \log(1 - 1/2^t) \rceil$ and is therefore at least 2^t , for $t \geq 1$. Since the value of T is a lower bound on the value of Y , this completes the proof. □

¹A random variable X dominates a random variable Y if $\Pr\{X > x\} \geq \Pr\{Y > x\}$ for all $x \in \mathbb{R}$.

3. THE RENDEZ-VOUS ALGORITHM

Consider the following algorithm used by an agent to make a random walk on a ring. The agent repeatedly performs the following steps: (1) toss a coin to determine a direction $d \in \{\text{clockwise}, \text{counterclockwise}\}$ then (2) run algorithm `BIGRAND`(t) replacing each increment of the variable Y with a step in direction d . By using t states for a clockwise counter and t states for a counterclockwise counter this algorithm can be implemented by a $2t$ state finite automata. (Or using one bit to remember the direction d and $\log t$ bits to keep track of the counter C in the `BIGRAND` algorithm, it can be implemented by an agent having only $1 + \log_2 t$ bits of memory.)

We call m iterations of the above algorithm a *round*. Together, Lemma 2 and Lemma 3 imply that, during a round, with probability at least κ , an agent will travel a distance of at least $\beta 2^t \sqrt{m}$ from its starting location. Set

$$m = \left\lceil \frac{n^2}{\beta^2 2^{2t}} \right\rceil$$

and consider what happens when two agents A and B both execute this rendez-vous algorithm. During the first round of A 's execution, with probability at least κ , agent A will have visited agent B 's starting location. Furthermore, with probability at least $1/2$ agent B will not have moved away from A when this occurs, so the paths of agents A and B will cross, and a rendez-vous will occur, with probability at least $\kappa/2$. If we define T as the round in which agents A and B rendez-vous, we therefore have $E[T] \leq 2/\kappa$.

By Lemma 3, the expected number of steps taken for A to execute the i th round is at most

$$E[M_i] \leq m 2^t (2 - 1/2^{t-1})$$

and

$$\text{var}(M_i) \leq m 2^{t+1}$$

The variables M_1, M_2, \dots are independent and the algorithm terminates when A and B rendez-vous. The time for two agents to rendez-vous is bounded by

$$\sum_{i=1}^T M_i .$$

Note that the event $T = j$ is independent of M_{j+1}, M_{j+2}, \dots so T is a stopping time for the sequence M_1, M_2, \dots . Furthermore, $\text{var}(M_i) < \infty$, so by Wald's Equation for the expectation

$$E \left[\sum_{i=1}^T M_i \right] = E[T] \cdot E[M_1] \leq \frac{2}{\kappa} \cdot m 2^t (2 - 1/2^{t-1}) .$$

This completes the proof of our first theorem.

THEOREM 2. *There exists a rendez-vous algorithm in which each agent has at most $2t$ states and whose expected rendez-vous time is $O(n^2/2^t + 2^t)$.*

4. THE LOWER BOUND

Next we show that the algorithm in Section 3 is optimal.

The model of computation for the lower bound represents a rendez-vous algorithm \mathcal{A} as a probabilistic finite automata having t states. Each vertex of the automata has two outgoing edges representing the two possible results of a coin toss and each edge e is labelled with a real number $\ell(e) \in [-1, +1]$. The edge label of e represents a step of length $|\ell(e)|$ with this step being counterclockwise if $\ell(e) < 0$ and clockwise if $\ell(e) > 0$. As before, both agents use identical automata and start in the same state. The rendez-vous process is complete once the distance between the two agents is at most 1. This model is stronger than the model used for upper bound of Theorem 2 since the edge labels are no longer restricted to be in the discrete set $\{-1, 0, +1\}$ and the definition of a rendez-vous has been slightly relaxed.

4.1 Well-Behaved Algorithms and Reset Times

We say that an algorithm is *well-behaved* if the directed graph of its state machine has only one strongly connected component that contains all nodes. We are particularly interested in intervals between consecutive visits to the start state, which we will call *rounds*.

LEMMA 4. *Let R be the number of steps during a round. Then $E[R] \leq 2^t$ and $E[R^2] \leq 2 \cdot 2^{2t}$.*

PROOF. For each state v of \mathcal{A} 's automata fix a shortest path (a sequence of edges) leading from v to the start state. For an automata that is currently at v we say that the next step is a *success* if it traverses the first edge of this path, otherwise we say that the next step is a *failure*.

Each round can be further refined into *phases*, where every phase consists of 0 or more successes followed by either a failure or by reaching the start vertex. Let X_i denote the length of the i th phase and note that X_i is dominated by a geometric(1/2) random variable X'_i , so $E[X_i] \leq E[X'_i] = 2$. On the other hand, if a phase lasts $t - 1$ steps then the start vertex is reached. Therefore, the probability of reaching the start vertex during any particular phase is at least $1/2^{t-1}$ and the number T of phases is dominated by a geometric($1/2^{t-1}$) random variable T' , so $E[T] \leq E[T'] \leq 2^{t-1}$. Therefore, by Wald's Equation

$$E[R] = E \left[\sum_{i=1}^T X_i \right] \leq E \left[\sum_{i=1}^{T'} X'_i \right] = E[T'] \cdot E[X'_1] = 2^t .$$

For the second part of the lemma, we can apply Wald's Equation for the variance (2) to obtain

$$\begin{aligned} E[R^2] &= E \left[\left(\sum_{i=1}^T X_i \right)^2 \right] \\ &\leq E \left[\left(\sum_{i=1}^{T'} X'_i \right)^2 \right] \end{aligned}$$

$$\begin{aligned}
 &= \text{var} \left(\sum_{i=1}^{T'} X'_i \right) + (\mathbb{E}[T'] \cdot \mathbb{E}[X'_1])^2 \\
 &= \mathbb{E}[T'] \cdot \text{var}(X_1) + (\mathbb{E}[T'] \cdot \mathbb{E}[X'_1])^2 \\
 &\leq 2^{t-1} \cdot 2 + 2^{2t} \\
 &\leq 2 \cdot 2^{2t}
 \end{aligned}$$

as required. \square

4.2 Unbiasing Algorithms

Note that $\mathbb{E}[R]$ can be expressed another way: For an edge e of the state machine, let $f(e)$ be the expected number of times the edge e is traversed during a round. The *reset time* of algorithm \mathcal{A} is then defined as

$$\text{reset}(\mathcal{A}) = \sum_e f(e) = \mathbb{E}[R] .$$

The *bias* of a well-behaved algorithm \mathcal{A} is defined as

$$\text{bias}(\mathcal{A}) = \sum_e f(e) \cdot \ell(e) ,$$

which is the expected sum of the edge labels encountered during a round. We say that \mathcal{A} is *unbiased* if $\text{bias}(\mathcal{A}) = 0$, otherwise we say that \mathcal{A} is *biased*.

Biased algorithms are somewhat more difficult to study. However, observe that, for any algorithm \mathcal{A} we can replace every edge label $\ell(e)$ with the value $\ell(e) - x$ for any real number x and obtain an equivalent algorithm in the sense that, if two agents A and B execute the modified algorithm following the same sequence of state transitions then A and B will rendez-vous after exactly the same number of steps. In particular, if we replace each edge label $\ell(e)$ with the value

$$\ell'(e) = \ell(e) - \frac{\text{bias}(\mathcal{A})}{\text{reset}(\mathcal{A})}$$

then we obtain an algorithm \mathcal{A}' with $\text{bias}(\mathcal{A}') = 0$. Furthermore, since $|\text{bias}(\mathcal{A})| \leq \text{reset}(\mathcal{A})$, every edge label $\ell'(e)$ has $-2 \leq \ell'(e) \leq 2$. This gives the following relation between biased and unbiased algorithms:

LEMMA 5. *Let \mathcal{A} be a well-behaved t -state algorithm with expected rendez-vous time R . Then there exists a well-behaved unbiased t -state algorithm \mathcal{A}' with expected rendez-vous time at most $2R$.*

4.3 The Lower Bound for Well-Behaved Algorithms

We now have all the tools in place to prove the lower bound for the case of well-behaved algorithms.

LEMMA 6. *Let \mathcal{A} be a well-behaved t -state algorithm. Then the expected rendez-vous time of \mathcal{A} is $\Omega(n^2/2^t)$.*

PROOF. Suppose the agents are placed at antipodal locations on a ring of size n , so that the distance between them is $n/2$. We will show that there exists constants $c > 0$ and $p > 0$ such that, after $cn^2/2^t$ steps, with probability at least p neither

agent will have travelled a distance greater than $n/4$ from their starting location. Thus, the expected rendez-vous time is at least $pcn^2/2^t = \Omega(n^2/2^t)$.

By Lemma 5 we can assume that \mathcal{A} is unbiased. Consider the actions of a single agent starting at location 0. The actions of the agent proceed in rounds where, during the i th round, the agent takes R_i steps and the sum of edge labels encountered during these steps is X_i . Note that the random variables X_1, X_2, \dots are i.i.d. with expectation $E[X] = 0$ and variance $E[X^2]$. Since the absolute value of X_i is bounded from above by R_i , we have the inequalities $E[|X_i|] \leq E[R_i]$ and $E[X_i^2] \leq E[R_i^2]$.

Let $S_i = \left| \sum_{j=1}^i X_j \right|$, for $i = 0, 1, 2, \dots$ be the agent's distance from their starting location at the end of the i th round. Let $Q_i = S_i^2 - iE[X^2]$ and observe that the sequence Q_1, Q_2, \dots is a martingale with respect to the sequence X_1, X_2, \dots [Ross 2002, Example 6.1d]. Define

$$T = \min\{i : S_i \geq m\} ,$$

and observe that this is equivalent to

$$T = \min\{i : Q_i \geq m^2 - iE[X^2]\} .$$

The random variable T is a *stopping time* for the martingale Q_1, Q_2, \dots . Furthermore,

$$\begin{aligned} E[|Q_{i+1} - Q_i| \mid X_1, \dots, X_i] &= E[|S_{i+1}^2 - (i+1)E[X^2] - S_i^2 + iE[X^2]| \mid X_1, \dots, X_i] \\ &= E[|S_{i+1}^2 - S_i^2 - E[X^2]| \mid X_1, \dots, X_i] \\ &= E\left[\left|\left(\sum_{j=1}^{i+1} X_j\right)^2 - \left(\sum_{j=1}^i X_j\right)^2 - E[X^2]\right| \mid X_1, \dots, X_i\right] \\ &= E\left[\left|\sum_{j=1}^{i+1} X_{i+1}X_j - E[X^2]\right| \mid X_1, \dots, X_i\right] \\ &= E[|X_{i+1}^2 - E[X^2]| \mid X_1, \dots, X_i] \\ &= E[|X_{i+1}^2 - E[X^2]|] \\ &\leq E[X^2] \leq E[R^2] < \infty . \end{aligned}$$

Therefore, by the Theorem 1

$$E[Q_T] = E[Q_1] = E[(X_1)^2 - E[X^2]] = 0 . \quad (3)$$

However, by definition $Q_T \geq m^2 - T \cdot E[X^2]$, so

$$E[Q_T] \geq E[m^2 - T \cdot E[X^2]] = m^2 - E[T] \cdot E[X^2] . \quad (4)$$

Equating the right hand sides of (3) and (4) gives

$$E[T] \geq \frac{m^2}{E[X^2]} .$$

Furthermore, the expected number of steps taken by the agent during these T

rounds is, by Wald's Equation,

$$\mathbb{E} \left[\sum_{i=1}^T R_i \right] = \mathbb{E}[T] \cdot \mathbb{E}[R_1] \geq \frac{m^2 \mathbb{E}[R]}{\mathbb{E}[X^2]} \geq \frac{m^2 \mathbb{E}[R]}{\mathbb{E}[R^2]} \geq \frac{m^2 \mathbb{E}[R]}{c2^{2t}} \geq \frac{m^2}{c2^{2t}}, \quad (5)$$

where the last two inequalities follow from Lemma 4 and the fact that $R \geq 1$. \square

4.4 Badly-Behaved Algorithms

Finally, we consider the case where the algorithm \mathcal{A} is not well-behaved. In this case, \mathcal{A} 's automata contains a set of *terminal components*. These are disjoint sets of vertices of the automata that are strongly connected and that have no outgoing edges (edges with source in the component and target outside the component). From each terminal component, select an arbitrary vertex and call it the *terminal start state* for that terminal component. An argument similar to that given in Lemma 4 proves:

LEMMA 7. *The expected time to reach some terminal start state is at most 2^t .*

Observe that each terminal component defines a well-behaved algorithm. Let c be the number of terminal components and let t_1, \dots, t_c be the sizes of these terminal components. When two agents execute the same algorithm \mathcal{A} , Lemma 7 and Markov's Inequality imply that the probability that both agents reach the same terminal component after at most 2^{t+2} steps is at least $1/2c$. By applying Lemma 6 to each component, we can therefore lower bound the expected rendez-vous time by

$$\frac{1}{2c} \Omega(n^2/2^{t-c}) \geq \Omega(n^2/2^{2t}).$$

Substituting $t' = t/2$ into the above completes the proof of our second theorem:

THEOREM 3. *Any $t/2$ -state rendez-vous algorithm has expected rendez-vous time $\Omega(n^2/2^t)$.*

4.5 Linear Time Rendez-vous

We observe that Theorems 2 and 3 immediately imply:

THEOREM 4. *$\Theta(\log \log n)$ bits of memory are necessary and sufficient to achieve rendez-vous in linear time on an n node ring.*

5. CONCLUSIONS

We have given upper and lower bounds on the expected rendez-vous time for two identical agents to rendez-vous on a ring as a function of the ring size n and the the memory available to the agents. In particular, we have shown that $O(\log \log n)$ bits of memory are necessary and sufficient for two agents to rendez-vous in $O(n)$ expected time.

A gap remains in our upper and lower bounds. When expressed in terms of the number of states t available to the agents, our upper and lower bounds differ by a factor of 4. We believe that the upper bound is tight and this gap is an artifact of the lower bound proof. Closing the gap remains an open problem.

The current paper studies symmetric rendez-vous with limited memory when the underlying graph is a ring. Another possibility is to consider rendez-vous with limited memory in other graphs. Possibilities include rendez-vous on an n -vertex torus or, more generally, on any n -vertex vertex-transitive graph.² With complete knowledge of the underlying graph and unlimited memory, rendezvous can be achieved in $O(n)$ expected time for any graph [Alpern and Gal 2003, Section 15.2]. On the other hand, if both agents take a random walk (which requires no memory and no knowledge of the underlying graph) then their expected meeting time is $O(n^3)$ and this is tight for some graphs [Coppersmith et al. 1993].

Another possibility is to consider the effects of memory limitations on randomized algorithms for the rendezvous of multiple (greater than two) agents on an n node ring. In particular, what is the expected time required for k identical agents, each having t states, to achieve rendez-vous on a synchronous, anonymous, oriented n node ring?

REFERENCES

- ALPERN, S. 1995. The rendezvous search problem. *SIAM Journal of Control and Optimization* 33, 673–683.
- ALPERN, S., BASTON, V. J., AND ESSEGAIER, S. 1999. Rendezvous search on a graph. *Journal of Applied Probability* 36, 1, 223–231.
- ALPERN, S. AND GAL, S. 2003. *The Theory of Search Games and Rendezvous*. Kluwer Academic Publishers, Norwell, Massachusetts.
- COPPERSMITH, D., TETALI, P., AND WINKLER, P. 1993. Collisions among random walks on a graph. *SIAM Journal of Discrete Mathematics* 6, 363–374.
- DAS, S., MIHALÁK, R. S., VICARI, E., AND WIDMAYER, P. 2008. Rendezvous of mobile agents when tokens fail anytime. In *Proceedings of the 12th International Conference on Principles of Distributed Systems (OPODIS2008)*, T. P. Baker, A. Bui, and S. Tixeuil, Eds. LNCS, vol. 5401. Springer, Berlin, 463–480.
- DESSMARK, A., FRAIGNAUD, P., AND PELC, A. 2003. Deterministic rendezvous in graphs. In *Proceedings of the 11th European Symposium on Algorithms (ESA 2003)*, G. D. Battista and U. Zwick, Eds. LNCS, vol. 2832. Springer, Berlin, 184–195.
- DOBREV, S., FLOCCHINI, P., PRENCIPE, G., AND SANTORO, N. 2004. Multiple agents rendezvous in a ring in spite of a black hole. In *Proceedings 8th International Conference on Principles of Distributed Systems (OPODIS 2004)*, T. Higashino, Ed. LNCS, vol. 3144. Springer, Berlin, 34–46.
- FLOCCHINI, P., KRANAKIS, E., KRIZANC, D., LUCCIO, F., SANTORO, N., AND SAWCHUK, C. 2004. Mobile agent rendezvous when tokens fail. In *Proceedings of the 11th Colloquium on Structural Information and Communication Complexity (SIROCCO 2004)*, R. Královic and O. Sýkora, Eds. LNCS, vol. 3104. Springer, Berlin, 161–172.
- FLOCCHINI, P., KRANAKIS, E., KRIZANC, D., SANTORO, N., AND SAWCHUK, C. 2004. Multiple mobile agent rendezvous in the ring. In *Proceedings of the 6th Latin American Symposium on Theoretical Informatics (LATIN 2004)*, M. Farach-Colton, Ed. LNCS, vol. 2976. Springer, Berlin, 599–608.
- GASINIENEC, L., KRANAKIS, E., KRIZANC, D., AND ZHANG, X. 2006. Optimal memory rendezvous of anonymous mobile agents in a uni-directional ring. In *Proceedings of the 32nd International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM 06)*, J. Widemann, G. Tel, J. Pokorný, M. Bieliková, and J. Štuller, Eds. LNCS, vol. 3831. Springer, Berlin, 282–292.

²A vertex-transitive graph is a graph such that, for any pair of vertices u and w , there is an isometry that maps u onto w [Alpern and Gal 2003, Chapter 15].

- KOWALSKI, D. AND MALINOWSKI, A. 2006. How to meet in an anonymous network. In *Proceedings of the 13th Colloquium on Structural Information and Communication Complexity (SIROCCO 2006)*, P. Flocchini and L. Gasieniec, Eds. LNCS, vol. 4056. Springer, Berlin, 44–58.
- KOWALSKI, D. AND PELC, A. 2004. Polynomial deterministic rendezvous in arbitrary graphs. In *Proceedings of the 15th International Symposium on Algorithms and Computation (ISAAC 2004)*, R. Fleischer and G. Trippen, Eds. LNCS, vol. 3341. Springer, Berlin, 644–656.
- KRANAKIS, E. AND KRIZANC, D. 2007. An algorithmic theory of mobile agents. In *2nd Symposium on Trustworthy Global Computing*, R. Bruni and U. Montanari, Eds. Number 4661 in LNCS. Springer, Berlin, 89–97.
- KRANAKIS, E., KRIZANC, D., AND MARKOU, E. 2006. Mobile agent rendezvous in a synchronous torus. In *Proceedings of the 5th Latin American Symposium on Theoretical Informatics (LATIN 2006)*, J. Correa, A. Hevia, and M. Kiwi, Eds. LNCS, vol. 3887. Springer, Berlin, 653–664.
- KRANAKIS, E., KRIZANC, D., AND MORIN, P. 2008. Randomized rendez-vous with limited memory. In *Proceedings of the 8th Latin American Theoretical Informatics Symposium (LATIN2008)*, E. S. Laber, C. Bornstein, L. T. Nogueira, and L. Faria, Eds. LNCS, vol. 4957. Springer, Berlin, 605–616.
- KRANAKIS, E., KRIZANC, D., SANTORO, N., AND SAWCHUK, C. 2003. Mobile agent rendezvous search problem in the ring. In *Proceedings of the International Conference on Distributed Computing Systems (ICDCS 2003)*. IEEE Computer Society, Los Alamitos, CA, USA, 592–599.
- MARCO, G. D., GARGANO, L., KRANAKIS, E., KRIZANC, D., PELC, A., AND VACARO, U. 2005. Asynchronous deterministic rendezvous in graphs. In *Proceedings of the 30th International Symposium on Mathematical Foundations of Computer Science (MFCS 2005)*, B. Rovan and P. Vojtáš, Eds. LNCS, vol. 2747. Springer, Berlin, 271–282.
- MARCO, G. D., GARGANO, L., KRANAKIS, E., KRIZANC, D., PELC, A., AND VACCARO, U. 2006. Asynchronous deterministic rendezvous in graphs. *Theoretical Computer Science* 355, 315–326.
- MITZENMACHER, M. AND UPFAL, E. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA.
- ROSS, S. M. 2002. *Probability Models for Computer Science*. Harcourt Academic Press, Berkeley, CA.
- ROY, N. AND DUDEK, G. 2001. Collaborative robot exploration and rendezvous: Algorithms, performance bounds and observations. *Autonomous Robots* 11, 117–136.
- SANTORO, N. 2006. *Design and Analysis of Distributed Algorithms*. Wiley, Hoboken, NJ.
- SAWCHUK, C. 2004. Mobile agent rendezvous in the ring. Ph.D. thesis, Carleton University, School of Computer Science, Ottawa, Canada.
- SUZUKI, I. AND YAMASHITA, M. 1999. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM Journal on Computing* 28, 1347–1363.
- YU, X. AND YUNG, M. 1996. Agent rendezvous: A dynamic symmetry-breaking problem. In *Proceedings of the 23rd International Colloquium on Automata, Languages and Programming (ICALP 96)*, F. M. auf der Heide and B. Monien, Eds. LNCS, vol. 1099. Springer, Berlin, 610–621.