



HAL
open science

Rendezvous of Distance-aware Mobile Agents in Unknown Graphs

Shantanu Das, Dariusz Dereniowski, Adrian Kosowski, Przemyslaw Uznanski

► **To cite this version:**

Shantanu Das, Dariusz Dereniowski, Adrian Kosowski, Przemyslaw Uznanski. Rendezvous of Distance-aware Mobile Agents in Unknown Graphs. 2014. hal-00986103v1

HAL Id: hal-00986103

<https://inria.hal.science/hal-00986103v1>

Preprint submitted on 30 Apr 2014 (v1), last revised 10 Jun 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rendezvous of Distance-aware Mobile Agents in Unknown Graphs

Shantanu Das¹, Dariusz Dereniowski², Adrian Kosowski³, and Przemysław Uznański¹

¹LIF, Aix Marseille University and CNRS, Marseille, France.

²Department of Algorithms and System Modeling, Gdańsk University of Technology,
Gdańsk, Poland.

³LIAFA, INRIA and Paris Diderot University, Paris, France.

Abstract

We study the problem of rendezvous of two mobile agents starting at distinct locations in an unknown graph. The agents have distinct labels and walk in synchronous steps. However the graph is unlabelled and the agents have no means of marking the nodes of the graph and cannot communicate with or see each other until they meet at a node. When the graph is very large we want the time to rendezvous to be independent of the graph size and to depend only on the initial distance between the agents and some local parameters such as the degree of the vertices, and the size of the agent's label. It is well known that even for simple graphs of degree Δ , the rendezvous time can be exponential in Δ in the worst case. In this paper, we introduce a new version of the rendezvous problem where the agents are equipped with a device that measures its distance to the other agent after every step. We show that these *distance-aware* agents are able to rendezvous in any unknown graph, in time polynomial in all the local parameters such the degree of the nodes, the initial distance D and the size of the smaller of the two agent labels $l = \min(l_1, l_2)$. Our algorithm has a time complexity of $O(\Delta(D + \log l))$ and we show an almost matching lower bound of $\Omega(\Delta(D + \log l / \log \Delta))$ on the time complexity of any rendezvous algorithm in our scenario. Further, this lower bound extends existing lower bounds for the general rendezvous problem without distance awareness.

Keywords: Mobile Agent, Rendezvous, Synchronous, Anonymous Networks, Distance Oracle, Lower Bounds

1 Introduction

1.1 Overview

Suppose two friends travel to a distant land and arrive at a city where all road signs are written in a language unknown to either of them. If the friends get separated and can no longer communicate with each-other, how could the two friends get together again without any help from a third person. This problem of gathering two autonomous mobile agents, called the *rendezvous* problem has been studied in many different contexts, for example for two ships lost in the sea, two astronauts that land in separate parts of a planet and so on. Initial studies on the problem were restricted to finding probabilistic strategies for movement of the two agents that minimize the expected time to rendezvous (See [1] for a survey of such results). In recent years, the deterministic version of the problem has received a lot of attention especially by the distributed computing community. The rendezvous problem for two agents moving along the edges of a graph is a typical problem of symmetry breaking and it is a primitive for distributed coordination among autonomous mobile robots. The solution to the problem depends on the structure of the graph, the capabilities of the agents and the initial knowledge available to the agents. In this paper, we consider the problem for deterministic agents with local vision moving in an initially unknown graph; the problem is solved when the two agents are simultaneously located in the same vertex. We are interested in the worst case time complexity for rendezvous. In general, rendezvous cannot always be solved deterministically when the underlying graph is highly symmetric and the agents follow identical strategies. A typical example is that of a ring network with unlabelled nodes where the two agents are placed on opposite vertices on any diameter. In this case, the distance between the two agents may never decrease if the agents use identical strategies, moving left or right at the same time.

Known solutions to the rendezvous problem are based on one of the following two approaches. The first type of solutions relies on finding a point of asymmetry in the graph and meeting at a unique point of asymmetry (e.g. such a point of asymmetry always exists in graphs where the nodes are labelled uniquely). The second type of solutions assumes that the agents are provided with distinct labels and thus, they can execute distinct strategies and ensure rendezvous. The former type of results require the agents to traverse every edge of the graph in the worst case and the time to rendezvous depends on the size of the graph. On the other hand, the latter type of solutions allow the agents to rendezvous in graphs of arbitrary size or even infinite graphs when the agents are located a finite distance apart. It has been shown that rendezvous of agents with distinct labels can be achieved in arbitrary finite graphs in time polynomial in the size of the graph and the size of the smaller of the labels assigned to the agents both in the synchronous case [17, 20] and the asynchronous case [10]. For infinite graphs, the only known results are for very specific graphs such as lines [19] or grids [3]. For a Δ -dimensional infinite grid, the optimal time to rendezvous is $\Theta(D^\Delta)$ which is already exponential in the maximum degree of the graph. In fact, in unknown graphs of degree Δ , where the agents start a distance of D apart, an agent may have to visit all vertices at a distance of D from its initial location. Since there could be D^Δ such vertices, the time cost of rendezvous would be exponential in Δ , even if the agents have complete knowledge of the graph as well as the initial distance between them. Thus the question is what additional capabilities would enable the agents to rendezvous in polynomial time.

In this paper we are interested in designing the simplest mechanism that can help the agents to rendezvous in a large (possibly infinite) graph in time polynomial in the other parameters of the problem, e.g. the initial distance, the maximum degree of the graph and the labels assigned to the agents. We achieve this by equipping the agents with a device that can measure the distance

to the other agent in the graph after each step of the algorithm¹. In fact, our algorithm does not require knowledge of the exact distance between the agents, but instead it is sufficient if the agent can detect whether the distance to the other agent increased or decreased after each move. We assume time to be discretized into rounds; in each round an agent can either traverse one edge of the graph or stay at its current location, and at the end of each round the agent can determine whether the distance to the other agent increased, decreased or remained unchanged during this round. Note that the agents have no means of detecting the direction leading to the other agent. We call agents equipped with the above device *distance-aware* agents. We show that distance-aware agents can rendezvous in arbitrary graphs in time polynomial in the initial distance, in the degree of the graph, and in the size of the smaller of the two agent labels.

1.2 Our Contributions

We show that two distance-aware agents, starting from an initial distance D apart from each other in a connected graph with maximum degree Δ , can rendezvous in time $O(\Delta \cdot D + \Delta \cdot \log l)$ rounds, where $l = \min(l_1, l_2)$ and l_1 and l_2 are the labels of the agents. The proof is constructive and provides a deterministic algorithm for the agent that takes as input the label of the agent. The algorithm does not require any prior knowledge of the graph and works for any connected graph. We also show that our algorithm is almost optimal by providing a lower bound of $\Omega(\Delta(D + \log l / \log \Delta))$ for rendezvous of distance-aware agents. Thus, our algorithm is asymptotically optimal when the maximum degree Δ is not extremely large.

The lower bound presented in this paper holds even for agents that can compute the exact distance to each other at every step, while the algorithm requires only the knowledge of changes in distance. Moreover, this lower bound improves upon the known bounds for the general problem of rendezvous of labelled agents, for which the best known lower bound (in terms of size of agent labels) was $\Omega(\log l)$ which already holds for the simplest graph — two nodes connected by a single edge. To the best of our knowledge, no generalizations of this lower bound to arbitrary graphs have been presented before.

1.3 Related Work

This paper considers the deterministic version of rendezvous (for randomized solutions see e.g. [1]). The problem of rendezvous of autonomous mobile agents has been studied for agents moving in a discrete space i.e. a graph [18] or those moving on a continuous space (e.g. two dimensional plane [11]). In the graph setting, rendezvous of *identical* agents is possible only if the graph is asymmetric or the agents are placed in asymmetric positions on the graph. There exists a characterization of such instances (graphs and initial positions of agents) where rendezvous is solvable [21]. If the agents are asynchronous, they can take advantage of the asymmetry in their initial positions by marking their initial position by a pebble [2, 16]. On the other hand, if the agents have distinct labels then rendezvous is possible in any graph and any starting positions, without the need to mark nodes. The first deterministic synchronous rendezvous algorithm for agents with distinct labels was presented by Dessmark et al. [9] and the time complexity of the algorithm was $O(n^5 \sqrt{\tau \log l} \log n + n^{10} \log^2 n \log l)$ for a graph of n nodes, where τ is the delay in the starting time of the agents. Subsequent studies [17, 20] improved this result and removed the parameter τ from the time complexity allowing for rendezvous in time polynomial in both n and $\log l$. For the asynchronous case, De Marco et al. [12] provided an algorithm for rendezvous

¹Such a device can be implemented in practice e.g. by emitting specific signals at periodic intervals and measuring the intensity of the signal received from the other device held by the other agent.

with a cost of $O(D \cdot \log l)$ rounds, when the graph is known. For unknown arbitrary graphs, Czyzowicz et al. [6] gave the first algorithm for asynchronous rendezvous but the cost of this algorithm was exponential in the size of the graph and the size of the maximum of the two labels. Recently, Dieudonné et al. [10] provided an improvement over this result achieving asynchronous rendezvous in time polynomial in n and $\log l$. In all the above results, the solution to rendezvous depends on the size of the graph. Rendezvous of agents starting from a finite distance D in an infinite graph has been studied for the special cases when the graph is a line [12, 19] or a grid [3], assuming that the agents have a sense of orientation and they know their own location in the labelled grid.

There have been several studies on the minimum capabilities needed by the agents to solve rendezvous. For example, the minimum memory required by an agent to solve rendezvous is known to be $\Theta(\log n)$ for arbitrary graphs. Czyzowicz et al. [7] have provided a memory optimal algorithm for rendezvous, and there are studies on the tradeoff between time and space requirements for rendezvous [8]. In some papers, additional capacities are assumed for the agents to overcome other limitations, e.g. global vision is assumed to overcome memory limitations [15] or the capability to mark nodes using tokens [16] or whiteboards [5] is often used to break symmetry. The model used in this paper can be seen as a special case of the oracle model for computation [13] where the agent is allowed to query an oracle that has global knowledge of the environment. However in our case, since the only queries are distance queries, the oracle can be implemented without complete knowledge of the graph topology.

2 Model and Notations

We model the environment as an undirected connected (possibly infinite) graph $G(V, E)$. The nodes of V are unlabelled such that vertices of the same degree look identical to any agent (i.e. the nodes are anonymous). At each node of the graph, the edges incident to it are locally labelled, so that an agent arriving at a node can distinguish among them². We assume that edges incident to a node v are labelled by distinct integers (called port numbers) from the set $\{1, 2, \dots, d(v)\}$, where $d(v)$ is the degree of node v . The degree of each node is finite and bounded by the parameter Δ (which is unknown to the agent). For any two distinct vertices $u, v \in V$, the distance between them, denoted by $\text{dist}(u, v)$, is the number of edges in any shortest path from u to v in G .

There are exactly two agents a_1 and a_2 , and each agent a_i has a distinct label $\ell(A_i) \in \{0, 1, \dots, L - 1\}$ for some integer $L \geq 1$. An agent knows its own label but not that of the other agent. The agents have no prior knowledge of the graph. Each agent starts from a distinct node of the graph and moves along the edges of the graph in synchronous steps following a deterministic algorithm. In other words, time is discretized into regular intervals called rounds; in each round, an agent at a node v can either move to an adjacent node of G or remain stationary at v . If the agent moves to an adjacent node w , the agent becomes aware of the port number of the edge through which it entered w . The agent has no means of marking a node that it visits and the agents cannot communicate with each other. An agent can see the other agent only when both agents are on the same node (in particular the agents do not see each other if they cross on the same edge from opposite directions). The two agents start the algorithm in the same round (called round 0) and rendezvous is achieved in the earliest round T when the two agents are at the same node. We denote by D the distance between the starting locations of the agents.

Contrary to previous studies on rendezvous, we assume that the agent is equipped with a device that measures the distance to the other agent. An agent at a node v in round t can make

²Such an assumption is necessary to allow the agent to navigate in the graph.

a query to this device (modelled as a function call `distance()`) which returns the value $\text{dist}(v, u)$, where node u is the location of the other agent in this round. In each round the agent can make one call to `distance()` and depending on the value returned, the value of the agent's label, the current state of the agent, and the degree of the current node, the agent chooses a number between 0 and $d(v)$ and leaves the current node v through this port. We assume that the port number 0 corresponds to a self-loop at node v and if the agent chooses 0 it remains at the same node v . In this paper we do not restrict the memory of an agent in any way. Thus, the agent can memorize its complete history of moves up to the current round and store this as its internal state.

3 Lower Bound for Distance-aware Rendezvous

In this section we provide lower bounds on the rendezvous time for distance-aware agents. Observe that a trivial lower bound is $\Omega(D)$ since at least one of the two agents must traverse $D/2$ edges to achieve rendezvous. We could easily obtain a better lower bound for graphs of degree Δ . The result below is folklore and we include it only for completeness.

Lemma 3.1 *Rendezvous of two agents that are initially at a distance of D apart in an unknown graph of maximum degree Δ requires $\Omega(D \cdot \Delta)$ rounds in the worst case.*

Proof: Consider the caterpillar graph shown in Figure 1 obtained by taking a line of $D + 1$ nodes and replacing each node with a star of size $\Delta - 1$. Suppose that the two agents start at the endpoints of the path of length D in this graph, as shown. If an agent traverses an edge leading to one of the leaves, it has no other option than to return and try another port. For any deterministic algorithm, an adversary can assign the port numbers in such a way that the agent needs to traverse all the Δ incident edges before it gets any closer to the other agent. Thus after $2\Delta - 1$ rounds the agents can reduce the distance between them by 2. Using a repetition of this argument, the result follows. ■




Figure 1: The two agents are located at the specially marked nodes at the two ends of a Caterpillar graph.

The above lower bound is independent of the agent's labels. We now present a bound which includes the parameter L which bounds the size of the set of possible agent labels. It is known that rendezvous of two agents with distinct labels requires $\Omega(\log L)$ rounds [9]. However this lower bound is for the simplest graph consisting of two nodes and a single edge connecting them. We provide below a better lower bound for arbitrary graphs of maximum degree Δ .

Our method is constructive, i.e., given a deterministic rendezvous algorithm \mathcal{A} , we describe a graph and a procedure for the adversary to label the ports in a graph and to choose starting positions that will keep the agents executing \mathcal{A} from meeting for the desired time period.

Definition 3.1 *A k -clique- p -butterfly is a $(k + 3)$ -regular graph on $(k \times p)$ vertices (denoted by $v_{0,0}, \dots, v_{k-1,p-1}$), constructed as follows:*

- for each $0 \leq j < p$, we connect all pairs of vertices from $v_{0,j}, v_{1,j}, \dots, v_{k-1,j}$ (so they form a k -vertex clique, named j -th clique),
- for any given $0 \leq j, j' < p$, $0 \leq i, i' < k$, we connect $v_{i,j}$ and $v_{i',j'}$ if $j' = (j + 1) \bmod p$ and $i' \in \{2i \bmod k, (2i + 1) \bmod k\}$.




Figure 2: A k -clique- p -butterfly graph where the two agents are located at the marked nodes at a distance $D > \log k$.

For two vertices, v from j -th clique and v' from j' -th clique, we say that their horizontal distance is $\min((j - j') \bmod p, (j' - j) \bmod p)$. By the properties of the butterfly-type interconnections between the cliques, we observe that, if for two vertices the horizontal distance is at least $\log k$, then their actual distance in the graph is equal to the horizontal distance. See Figure 2 for an example of a k -clique- p -butterfly graph where the agents are at a distance $D > \log k$.

Theorem 3.1 *For any odd $k \geq 3$ and $p \geq 2 \cdot \lceil \log k \rceil$, given a k -clique- p -butterfly (denoted further as G), for any deterministic algorithm \mathcal{A} , and integers $L > 0$ and $D \geq \log k = \log(\Delta - 3)$, there exist labels $0 \leq \ell(1), \ell(2) < L$ and a port numbering of G such that if two agents with these labels start at two vertices at distance D in this graph, they will remain at distance D for at least $\Theta(\Delta \log L / \log \Delta)$ steps.*

Proof: The main idea of the proof is as follows. We will choose four special port numbers that will be assigned in every vertex to the “bridge” edges connecting it to an adjacent clique (two ports to advance to the next clique and two ports to go backward). Note that moving inside the clique doesn’t change the distance to the other agent. By carefully picking the agent labels, we will ensure that both agents will choose the forward ports at the same rounds and the backward ports also at the same rounds, thus they will maintain a constant distance between themselves, for a sufficient amount of time and during this time period, any queries to the distance oracle would be useless.

Before we proceed to prove the theorem, we need to fix some notation and provide some basic lemmas. We will consider the k -clique- p -butterfly graphs with a special type of port numbering where the ports on the two ends of each edge always form one of the pairs of values: $\{1, \Delta\}, \{2, \Delta - 1\}, \dots, \{\Delta/2, \Delta/2 + 1\}$. In other words, whenever an agent chooses to leave a node by port j , it arrives at the adjacent node by port $\Delta + 1 - j$. In a Δ -regular graph with such a port numbering, an agent can never distinguish between any two vertices and it can learn nothing new about the graph by traversing it. Thus, the algorithm \mathcal{A} must choose a predefined sequence of ports to follow during the first $t > 0$ rounds when the distance between the agents remains D . For every $0 \leq \ell < L$ we denote the sequence of ports chosen by algorithm \mathcal{A} (with input label ℓ) by $(\mathcal{P}_\ell(i))_{i=0}^{i < t}$.

Lemma 3.2 *There exist p_1, p_2 , $0 < p_1, p_2 \leq \Delta/2$, and a set $\mathcal{L} \subset \{0, 1, \dots, L - 1\}$, $|\mathcal{L}| \geq L/2$, such that for each $\ell \in \mathcal{L}$, at most $8t/\Delta$ elements of $(\mathcal{P}_\ell(i))_{i=0}^{i < t}$ are equal to one of the four values $\{p_1, p_2, \Delta + 1 - p_1, \Delta + 1 - p_2\}$ (in total).*

Proof: In G , the port numbers on the edges are paired together as follows: $\{1, \Delta\}, \dots, \{\Delta/2, \Delta/2+1\}$. For each pair $\{p, \Delta+1-p\}$, let S_p be the number of pairs (i, ℓ) , $0 \leq i < t, 0 \leq \ell < L$, such that $\mathcal{P}_\ell(i) \in \{p, \Delta+1-p\}$. Since $S_1 + S_2 + \dots + S_{\Delta/2} \leq t \cdot L$, there must exist p_1, p_2 such that $S_{p_1} + S_{p_2} \leq t \cdot L \cdot 2/(\Delta/2) = t \cdot L \cdot 4/\Delta$. Thus, there must exist a subset \mathcal{L} containing at least half of the possible labels, such that for each $\ell \in \mathcal{L}$, the number of i 's for which $\mathcal{P}_\ell(i) \in \{p_1, p_2, \Delta+1-p_1, \Delta+1-p_2\}$ is not greater than $2 \cdot \frac{t \cdot L \cdot 4/\Delta}{L} = 8t/\Delta$. ■

Given the ports p_1 and p_2 as in the lemma above, the adversary can assign the pairs of port numbers $(p_1, \Delta+1-p_1)$ and $(p_2, \Delta+1-p_2)$ to those edges of G that connect two adjacent cliques, such that ports p_1 and p_2 will take an agent forward to the next clique, ports $\Delta+1-p_1$ and $\Delta+1-p_2$ would take an agent backward to the previous clique and any other port will keep an agent within the same clique. We partition the set of port numbers into the subsets:

$$\begin{aligned} A &= \{p_1, p_2\}, \\ B &= \{\Delta+1-p_1, \Delta+1-p_2\}, \\ C &= \{0, 1, \dots, \Delta\} \setminus (A \cup B). \end{aligned}$$

Informally, the ports in A take an agent one step ahead, the ports in B take it one step back, and the C ports keep it in the same clique.

Lemma 3.3 *There exist $0 \leq \ell(1) < \ell(2) < L$ such that $\forall i \leq \lfloor \frac{\log L}{2 \log \Delta} \rfloor \cdot \frac{\Delta}{8}$,*

$$\mathcal{P}_{\ell(1)}(i) \in A \text{ iff } \mathcal{P}_{\ell(2)}(i) \in A, \text{ and} \quad (1)$$

$$\mathcal{P}_{\ell(1)}(i) \in B \text{ iff } \mathcal{P}_{\ell(2)}(i) \in B. \quad (2)$$

Proof: Let $t > 0$ be the first round such that for $i = t$, either (1) or (2) does not hold for any possible pair of labels $\ell(1)$ and $\ell(2)$ from the set $\{0, 1, \dots, L-1\}$. We will bound the value of t . Note that the sequence of ports $(\mathcal{P}_\ell(i))_{i=0}^{i < t}$ can be written as a sequence of A, B and C 's. We choose the set of labels \mathcal{L} from Lemma 3.2 and for any $\ell \in \mathcal{L}$, let us count the number of possible sequences $(\mathcal{P}_\ell(i))_{i=0}^{i < t}$ which correspond to distinct words from the alphabet $\{A, B, C\}$; let us denote this number as X . Recall that any such sequence can have at most $8t/\Delta$ A -ports and B -ports in total (cf. Lemma 3.2). Assuming for the sake of notation that $T = 8t/\Delta$ is an integer and using Stirling's approximation, we get:

$$\begin{aligned} X &= \sum_{i=0}^{\lfloor T \rfloor} \binom{t}{i} \cdot 2^i \leq 1 + T \cdot \binom{t}{T} \cdot 2^T \leq 1 + T \cdot \frac{(2t)^T}{(8t/\Delta)^T} \\ &\leq 1 + T \cdot \left(\frac{2t}{8t/\Delta \cdot \frac{1}{e}} \right)^T = 1 + T \left(\frac{e}{4} \Delta \right)^T. \end{aligned}$$

Since for any two labels in \mathcal{L} the algorithm chooses distinct sequences, we have $L/2 \leq |\mathcal{L}| \leq X$, which gives us $T \geq \lfloor \frac{\log L}{2 \log \Delta} \rfloor$ and thus $t \geq \lfloor \frac{\log L}{2 \log \Delta} \rfloor \cdot \frac{\Delta}{8}$. ■

We now return to the proof of the Theorem 3.1. From Lemma 3.3, we know that there exist two agents with labels $\ell(1), \ell(2)$ such that these agents, when executing algorithm \mathcal{A} , will use the ports going forward or backward simultaneously during the first $t = \lfloor \frac{\log L}{2 \log \Delta} \rfloor \cdot \frac{\Delta}{8}$ rounds. Since the distance in the graph is independent of relative positions inside cliques (for $D \geq \log k$), the distance between agents is maintained to be D for at least $t = \Theta(\Delta \log L / \log \Delta)$ steps. ■

The results of this section provide a lower bound of $\Omega(\Delta(D + \log L / \log \Delta))$ rounds for rendezvous of distance-aware agents.

4 Upper Bound and Algorithm for Rendezvous

In this section we provide an algorithm that guarantees rendezvous of two distance-aware agents in $O(\Delta(D + \log l))$ rounds. The algorithm is divided into four procedures. For each procedure, first we provide some intuitions on its behavior, then we give its formal description in the form of pseudo-code. After the descriptions of all procedures, the proof of correctness and analysis of rendezvous time of the algorithm are given.

We start with a procedure **TestPorts** (see Algorithm 1), which attempts to decrease the distance between the two agents. The input consists of a positive integer δ and a bit $b \in \{0, 1\}$. Recall that the command **distance()** performs the oracle query and provides the current distance from the other agent. We also introduce a command **move(x)**, where x is an integer, which behaves as follows. If $x \in \{1, \dots, d(v)\}$, where v is the node currently occupied by the executing agent, then **move(x)** forces the executing agent to move from v by taking port x , and the value returned by the **move(x)** command is the entry port at the arrival node. Otherwise, that is when $x \notin \{1, \dots, d(v)\}$, the executing agent stays idle in the given round, and **move(x)** returns 0. (We pass 0 as an argument to deliberately make an agent idle.)

Algorithm 1 Procedure **TestPorts**(δ, b)

Input: Two integers $\delta \geq 1$ and $b \in \{0, 1\}$.

Output: **success** if the distance between agents decreases in some round; **failure** otherwise.

```

for  $i \leftarrow 1$  to  $\delta$  do
   $x \leftarrow$  distance()
   $p \leftarrow$  move( $i \cdot b$ )
   $y \leftarrow$  distance()
  if  $y < x$  then
    return success {When the distance to the other agent decreased.}
  end if
  move( $p \cdot b$ ) {Going back along the same edge.}
end for
return failure {When the distance to the other agent never decreased.}

```

The interpretation of the input variable b is that whenever $b = 0$, then the agent does not perform any movements during the execution of **TestPorts**. Note that if $b = 1$, then **move($i \cdot b$)** forces an agent to move only if there exists an edge with port i at v .

Suppose that, in the same round, both agents perform calls to procedure **TestPorts** with input parameters δ_1, b_1 and δ_2, b_2 , respectively. We will always ensure that $\delta = \delta_1 = \delta_2$. Informally, $b_1 = b_2 = 1$ implies that both agents iteratively take ports $1, \dots, \delta$ (skipping the ones not present at the current node), ending the process if the distance between them decreases. Clearly, if $b_1 = b_2 = 0$, then both agents just stay idle during the 2δ rounds and the procedure returns **failure**. If $b_1 \neq b_2$, then one agent stays idle while the other ‘tests ports’. If we ensure that δ exceeds the degree of the node occupied by the agent that performs the movements, then the procedure will return **success** whenever $b_1 \neq b_2$.

We now describe a procedure **BoundDegrees** with input variable $b \in \{0, 1\}$ (see Algorithm 2). Informally speaking, for each 2^l such that $2^l < d(v)$, where v is the node occupied at the beginning of the execution of the procedure, the executing agent stays idle for 2^{l+1} rounds (this is achieved by the call to **TestPorts**($2^l, 0$)). This part is independent of b . Then, **TestPorts**($2^{\lceil \log_2 d(v) \rceil}, b$) is called. If $b = 0$, then the agent stays idle for another $2^{\lceil \log_2 d(v) \rceil + 1}$ rounds. If $b = 1$, then the

agent sequentially explores all ports at v (Note that the value of $2^{\lceil \log_2 d(v) \rceil}$ exceeds the degree of v). The above process is interrupted whenever the agents observe that the distance between them decreased, and the procedure returns **success** in that case.

Algorithm 2 Procedure **BoundDegrees**(b)

Input: An integer $b \in \{0, 1\}$.

Output: **success** or **failure**.

Let v be the currently occupied node.

for $l \leftarrow 0$ **to** $\lceil \log_2 d(v) \rceil - 1$ **do**

$s \leftarrow \text{TestPorts}(2^l, 0)$

if $s = \text{success}$ **then**

return success

end if

end for

$s \leftarrow \text{TestPorts}(2^{\lceil \log_2 d(v) \rceil}, b)$

return s

The observation given below follows directly from the formulation of procedure **BoundDegrees**.

Observation 1 *Suppose that agent A_i occupies node v_i , $i \in \{1, 2\}$, and executes procedure **BoundDegrees** at the beginning of round r . If procedure **TestPorts** does not return **success** in the first j iterations of **BoundDegrees**, then:*

- (i) A_i occupies v_i at the end of the j -th iteration of **BoundDegrees**,
- (ii) both agents end the execution of the j -th iteration in round $r + 2^j - 1$.

Denote $I_0 = (0, 1]$ and $I_j = (2^{j-1}, 2^j]$ for $j \geq 1$. We say that two nodes u and v are *similar* if there exists $j \geq 0$ such that $d(u) \in I_j$ and $d(v) \in I_j$.

Lemma 4.1 *Let r be some integer. Suppose that agent A_i is present at v_i , $i \in \{1, 2\}$, and calls in round r procedure **BoundDegrees** with input value b_i . Then:*

- (i) for $b_1 = b_2 = 1$, if both calls to **BoundDegrees** return **failure**, then the nodes v_1 and v_2 are similar, and
- (ii) if $b_1 \neq b_2$ and v_1 and v_2 are similar, then both calls to **BoundDegrees** return **success**.

Proof: We first prove (ii). Suppose without loss of generality that $b_1 = 1$ and $b_2 = 0$. Since the nodes v_1 and v_2 are similar, $x = \lceil \log_2 d(v_1) \rceil = \lceil \log_2 d(v_2) \rceil$ and therefore the executions of procedure **BoundDegrees** have the same number of iterations of the ‘for’ loop. If, in one of those iterations, the execution of procedure **BoundDegrees** ends, then (ii) holds and hence suppose that this is not the case. Thus, the calls to **TestPorts**($2^x, b_1$) and **TestPorts**($2^x, b_2$) are made by the agents. Moreover, by Observation 1, both calls are made in the same round and when agent A_i is at v_i , $i \in \{1, 2\}$. During these calls, A_2 stays idle during 2^{x+1} rounds (because $b_2 = 0$) while A_1 explores all ports at v_1 during the same 2^{x+1} rounds (because $b_1 = 1$ and $2^{\lceil \log_2 d(v_1) \rceil} \geq d(v_1)$). This guarantees that the latter calls to **TestPorts** return **success**, which completes the proof of (ii).

We now prove (i). Let $b_1 = b_2 = 1$. We argue that if the nodes v_1 and v_2 are not similar, then the call to **BoundDegrees** results in returning **success**. Suppose without loss of generality that $d(v_i) \in I_{j_i}$, $i \in \{1, 2\}$, where $j_1 < j_2$. The number of iterations of the ‘for’ loop of procedure

BoundDegrees executed by A_i is $\lceil \log_2 d(v_i) \rceil = j_i$, $i \in \{1, 2\}$. Thus, by Observation 1(i), A_i occupies v_i at the end of j_1 -th iteration of **BoundDegrees** for each $i \in \{1, 2\}$. Moreover, by Observation 1(ii), after finishing the execution of the ‘for’ loop, A_1 calls **TestPorts**($2^{j_1}, b_1$) while A_2 calls **TestPorts**($2^{j_1}, 0$) in the $(j_1 + 1)$ -st iteration of the ‘for’ loop of **BoundDegrees**. Also, both of the above-mentioned calls to **TestPorts** are made in the same round r' . By similar arguments as when proving (ii), we obtain that condition (i) holds. ■

Note that if procedure **BoundDegrees** returns **success**, then the agents get closer, and we can repeat the same process for at their current locations. However, for some nodes procedure **BoundDegrees** may return **failure** and then procedure **CompareLabels** described below (see Algorithm 3) helps to break the symmetry. Procedure **CompareLabels** uses the notion of extended

Algorithm 3 Procedure **CompareLabels**

Input: None.

Output: $\xi(\ell)_i$ such that i is distinguishing for the extended labels of the agents.

Let $\xi(\ell)$ be the extended label of the executing agent.

```

for  $i \leftarrow 1$  to  $\lceil \log_2 \xi(\ell) \rceil$  do
   $s \leftarrow$  BoundDegrees( $\xi(\ell)_i$ )
  if  $s =$  success then
    return  $\xi(\ell)_i$ 
  end if
end for

```

labels. The *extended label* $\xi(\ell)$ of a label ℓ is an integer whose j -th bit is defined as follows:

$$\xi(\ell)_j = \begin{cases} \ell_{\lceil j/2 \rceil}, & \text{for } j \in \{1, 3, 5, \dots, 2\lceil \log_2 \ell \rceil - 1\}, \\ 1, & \text{for } j = 2\lceil \log_2 \ell \rceil, \\ 0, & \text{otherwise.} \end{cases}$$

The index $j = 2\lceil \log_2 \ell \rceil$ is called the *terminating bit* of $\xi(\ell)$ (this is the last bit set to 1). Informally, the odd positions of the extended label are the bits of ℓ while the even positions are all zeros, except for the terminating bit. We say that an index i is *distinguishing* for two extended labels $\xi(\ell)$ and $\xi(\ell')$ if $\xi(\ell)_i \neq \xi(\ell')_i$.

Informally, procedure **CompareLabels** iterates over the bits of the extended label of the executing agent in order to find an index i that is distinguishing for the extended labels of the two agents. The construction of extended labels guarantees that there exists a distinguishing index i not greater than the smaller length of the two extended labels and hence **BoundDegrees** returns **success** at the latest in the i -th iteration of the ‘for’ loop of procedure **CompareLabels** (the formal proof is given later; we remark here that if we used the label instead of the extended label, then the number of iterations of the ‘for loop’ would have to be equal to the length of the greater label to ensure rendezvous).

We postpone the analysis of procedure **CompareLabels** (given in Lemma 4.2) as it depends on the context at which it is called by the main procedure.

We finally describe the main procedure **Rendezvous** (see Algorithm 4). We start with its intuitive description. The first ‘while’ loop iteratively calls procedure **BoundDegrees**(1) as long as its execution gets the agents closer to each other. If a call to **BoundDegrees**(1) does not achieve that, then (as we formally prove later) the agents observed the same distance between each other while both explored all ports at their respective locations. This is significant as both agents

Algorithm 4 Procedure `Rendezvous`

```
s ← success
while rendezvous not achieved and s = success do
  s ← BoundDegrees(1)
end while
b ← CompareLabels
while rendezvous not achieved do
  BoundDegrees(b)
end while
```

learn that they occupy nodes whose degrees are in the same interval I_j for some $j \geq 0$. In other words, the agents learn an asymptotically tight upper bound on both degrees. Then, procedure `CompareLabels` is called and uses the above fact as well as the labels of the agents to break the symmetry that occurs at the current agents' nodes. Note that `CompareLabels` returns either 0 or 1 and in this case different values are returned for both agents (see Lemma 4.2 below). Thus, the agent whose execution of `CompareLabels` returned 0 stays idle from now on. The other agents continues making calls to `BoundDegrees(1)` and since each execution results in exploring all ports at the currently occupied node, each execution gets the agent one step closer to the one that is idle. We also remark that the respective calls to procedure `BoundDegrees` in the second 'while' loop of `CompareLabels` are not necessarily 'synced', that is, the j -th of those calls can be made in different rounds by the agents. This, however, is not important as one of the agents stays idle and the other one performs appropriate movements.

Lemma 4.2 analyzes the only call to procedure `CompareLabels` made by procedure `Rendezvous`. Then, Theorem 4.1 provides the upper bound on the rendezvous time for distance-aware agents in arbitrary networks.

Lemma 4.2 *Whenever procedure `CompareLabels` is called by both agents during the execution of procedure `Rendezvous`, both agents finish the execution of `CompareLabels` in the same round and the values returned by `CompareLabels` are different for the two agents.*

Proof: Since both agents call `CompareLabels`, none of the preceding calls to procedure `BoundDegrees` returns `success` and the agents do not rendezvous prior to the call to `CompareLabels`. By Observation 1 and a simple inductive argument, the calls to procedure `CompareLabels` made by both agents end in the same round r .

It remains to prove that the calls to `CompareLabels` return different values. By Lemma 4.1(i), at the beginning of round r the agents A_1 and A_2 are, respectively, at two similar nodes v_1 and v_2 . Let $j \geq 0$ be the minimum distinguishing index for agents' labels, i.e., $\xi(\ell(A_1))_j \neq \xi(\ell(A_2))_j$ and $\xi(\ell(A_1))_{j'} = \xi(\ell(A_2))_{j'}$ for each $1 \leq j' < j$. Such an index j exists because the labels of the agents are different and hence the extended labels have a distinguishing index. Moreover,

$$j \leq \min \{ \lceil \log_2 \xi(\ell(A_1)) \rceil, \lceil \log_2 \xi(\ell(A_2)) \rceil \}.$$

Indeed, if the two labels are of the same length, then the extended labels are of the same length. If, on the other hand, the labels have different lengths, then the terminating bits are at different positions, which in particular implies that the terminating bit of the smaller label is at position that is distinguishing for the two extended labels.

By assumption, Observation 1 and an inductive argument, A_i is at v_i at the beginning of the j' -th call to procedure `BoundDegrees` during the execution of procedure `CompareLabels`, where

$j' \leq j$. The last execution of procedure `BoundDegrees` preceding the call to `CompareLabels` returns `failure`. Hence, if $\xi(\ell(A_1))_{j'} = \xi(\ell(A_2))_{j'} = 1$, then the j' -th call to `BoundDegrees` returns `failure`. Clearly, if $\xi(\ell(A_1))_{j'} = \xi(\ell(A_2))_{j'} = 0$, then the agents stay idle during the j' -th call to procedure `BoundDegrees` which also implies that it returns `failure`.

Thus, the above proves that the j -th call to procedure `BoundDegrees` takes place during the execution of procedure `CompareLabels` and, by Lemma 4.1(ii), it returns `success` for both agents. Thus, procedure `CompareLabels` returns the respective bits of the extended label at position that is distinguishing, which completes the proof. ■

Theorem 4.1 *Suppose that agent A_i with label $\ell(A_i)$ initially occupies node v_i , $i \in \{1, 2\}$. Procedure `Rendezvous` guarantees that A_1 and A_2 rendezvous within $O(\Delta \cdot (D + \min_i \{\log \ell(A_i)\}))$ rounds where $D = \text{dist}(v_1, v_2)$.*

Proof: We first prove that the execution of procedure `Rendezvous` guarantees rendezvous. If the agents rendezvous during the execution of the first ‘for’ loop, then the claim follows and hence suppose that this is not the case. Denote by b_i the value of the variable b returned by the call to procedure `CompareLabels` by agent A_i , $i \in \{1, 2\}$. By Lemma 4.2, $b_1 \neq b_2$. Let without loss of generality, $b_1 = 0$ and $b_2 = 1$. This implies that the agent A_1 stays idle indefinitely. The agent A_2 , during execution of procedure `BoundDegrees`(b_2) called in the second ‘for’ loop of procedure `Rendezvous`, explores all ports of the currently occupied node. Thus, the distance between agents decreases in some round which implies that each such call to procedure `BoundDegrees`(b_2) returns `success`. Thus, the agents rendezvous eventually.

Now we analyze the rendezvous time. Each call to `BoundDegrees` takes $O(\Delta)$ rounds. Moreover, each such call, except for at most one, made directly by procedure `Rendezvous` ensures that the distance between the agents decreases. This follows immediately for calls to `BoundDegrees` preceding the call to procedure `CompareLabels` since those calls, possibly except for the last one, return `success`. As for the remaining calls to `BoundDegrees`, the above claim is due to the fact that the input values are different for both agents due to Lemma 4.2. Thus, the total number of rounds due to all calls to `BoundDegrees` made directly by procedure `Rendezvous` is $O(\Delta \cdot \text{dist}(v_1, v_2))$. The number of iterations of the ‘for’ loop of procedure `CompareLabels` is $O(\min\{\log \ell(A_1), \log \ell(A_2)\})$, each resulting in $O(\Delta)$ rounds (the call to `BoundDegrees`). ■

5 Conclusions

This paper presented a new model for mobile agent computation by providing the agents with the capability of measuring distances to each other (or detecting changes in distances) at each step. We show that this simple mechanism allows us to reduce the time to rendezvous from exponential to polynomial in the degree of the graph. Since such a distance measuring device can be easily implemented in practice, one could ask what other problems can be solved more easily using this additional capability. For example, the agents could use this mechanism for communication at distance by moving back and forth, when there are no other means of communication. This opens up a new area of research which is worth investigating.

References

- [1] S. Alpern and S. Gal. *The Theory of Search Games and Rendezvous*. Kluwer, 2003.
- [2] V. Baston and S. Gal. Rendezvous search when marks are left at the starting points. *Naval Research Logistics*, 48(8):722–731, 2001.
- [3] E. Bampas, J. Czyzowicz, L. Gasieniec, D. Ilcinkas, A. Labourel. Almost optimal asynchronous rendezvous in infinite multidimensional grids, In Proc. 24th International Symposium on Distributed Computing (DISC), LNCS 6343, 297-311, 2010.
- [4] J. Chalopin, S. Das, and N. Santoro. Rendezvous of Mobile Agents in Unknown Graphs with Faulty Links. In Proc. 21st International Symposium on Distributed Computing (DISC), pp. 108–122, 2007.
- [5] J. Chalopin, S. Das, and P. Widmayer. Rendezvous of Mobile Agents in Directed Graphs. In Proc. 24th International Symposium on Distributed Computing (DISC), pp. 282-296, 2010.
- [6] J. Czyzowicz, A. Labourel, and A. Pelc, How to meet asynchronously (almost) everywhere, In Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 22-30, 2010.
- [7] J. Czyzowicz, A. Kosowski, A. Pelc. How to meet when you forget: Log-space rendezvous in arbitrary graphs, *Distributed Computing*, 25: 165-178, 2012.
- [8] J. Czyzowicz, A. Kosowski, A. Pelc, Time vs. space trade-offs for rendezvous in trees, *Distributed Computing*, 27(2): 95-109, 2014.
- [9] A. Dessmark, P. Fraigniaud, D. Kowalski, A. Pelc, Deterministic rendezvous in graphs, *Algorithmica*, 46: 69-96, 2006.
- [10] Y. Dieudonné, A. Pelc, Vincent Villain. How to meet asynchronously at polynomial cost. In Proc 32nd Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 92-99, 2013.
- [11] P. Flocchini, G. Prencipe, N. Santoro, P. Widmayer: Gathering of asynchronous robots with limited visibility. *Theoretical Computer Science*, 337(1-3):147-168, 2005.
- [12] G. De Marco, L. Gargano, E. Kranakis, D. Krizanc, A. Pelc, and U. Vaccaro. Asynchronous deterministic rendezvous in graphs. *Theoretical Computer Science*, 355(3):315–326, 2006.
- [13] P. Fraigniaud, D. Ilcinkas, and A. Pelc. Oracle size: A new measure of difficulty for communication problems. In Proc. 25th Ann ACM Symposium on Principles of Distributed Computing (PODC), pp. 179-187, 2006.
- [14] P. Fraigniaud, and A. Pelc, Deterministic Rendezvous in Trees with Little Memory. In Proc. of 22nd Int. Conference on Distributed Computing (DISC), pp. 242-256, 2008.
- [15] R. Klasing, E. Markou, and A. Pelc, Gathering asynchronous oblivious mobile robots in a ring, *Theoretical Computer Science*, 390(1):27-39, 2008.
- [16] E. Kranakis, D. Krizanc, and E. Markou. *The mobile agent rendezvous problem in the ring*, Morgan and Claypool Publishers, 2010.

- [17] D. R. Kowalski and A. Malinowski, How to meet in anonymous network, *Theoretical Computer Science*, 399(1-2):141–156, 2008.
- [18] A. Pelc, Deterministic rendezvous in networks: A comprehensive survey, *Networks*, 59: 331–347, 2012.
- [19] G. Stachowiak. Asynchronous Deterministic Rendezvous on the Line, In Proc 35th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM), pp. 497-508, 2009.
- [20] A. Ta-Shma and U. Zwick. Deterministic rendezvous, treasure hunts and strongly universal exploration sequences. In Proc 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 599–608, 2007.
- [21] M. Yamashita and T. Kameda. Computing on anonymous networks: Part I—Characterizing the solvable cases. *IEEE Transactions on Parallel and Distributed Systems*, 7(1):69–89, 1996.
- [22] X. Yu and M. Yung. Agent rendezvous: A dynamic symmetry-breaking problem. In Proc. International Colloquium on Automata, Languages, and Programming (ICALP), pp. 610–621, 1996.