

Fast Rendezvous on a Cycle by Agents with Different Speeds

Ofer Feinerman^{1,*}, Amos Korman^{2,**}, Shay Kutten^{3,***}, and Yoav Rodeh⁴

¹ The Shlomo and Michla Tomarin Career Development Chair,
Weizmann Institute of Science, Rehovot, Israel

² CNRS and University Paris Diderot, Paris, France

³ Faculty of IR&M, Technion, Haifa 32000, Israel

⁴ Jerusalem College of Engineering

Abstract. The difference between the speed of the actions of different processes is typically considered as an obstacle that makes the achievement of cooperative goals more difficult. In this work, we aim to highlight potential *benefits* of such asynchrony phenomena to tasks involving symmetry breaking. Specifically, in this paper, identical (except for their speeds) mobile agents are placed at arbitrary locations on a (continuous) cycle of length n and use their speed difference in order to rendezvous fast. We normalize the speed of the slower agent to be 1, and fix the speed of the faster agent to be some $c > 1$. (An agent does not know whether it is the slower agent or the faster one.) The straightforward *distributed-race (DR)* algorithm is the one in which both agents simply start walking until rendezvous is achieved. It is easy to show that, in the worst case, the rendezvous time of DR is $n/(c-1)$. Note that in the interesting case, where c is very close to 1 (e.g., $c = 1 + 1/n^k$), this bound becomes huge. Our first result is a lower bound showing that, up to a multiplicative factor of 2, this bound is unavoidable, even in a model that allows agents to leave arbitrary marks (the *white board* model), even assuming sense of direction, and even assuming n and c are known to agents. That is, we show that under such assumptions, the rendezvous time of any algorithm is at least $\frac{n}{2(c-1)}$ if $c \leq 3$ and slightly larger (specifically, $\frac{n}{c+1}$) if $c > 3$. We then manage to construct an algorithm that precisely matches the lower bound for the case $c \leq 2$, and almost matches it when $c > 2$. Moreover, our algorithm performs under weaker assumptions than those stated above, as it does not assume sense of direction, and it allows agents to leave only a single mark (a pebble) and only at the place where they start the execution. Finally, we investigate the setting in which no marks can be used at all, and show tight bounds for $c \leq 2$, and almost tight bounds for $c > 2$.

Keywords: rendezvous, asynchrony, heterogeneity, speed, cycle, pebble, white board, mobile agents.

* Supported by the Clore Foundation, the Israel Science Foundation (FIRST grant no. 1694/10) and the Minerva Foundation.

** Supported by the ANR project DISPLEXITY, and by the INRIA project GANG.

*** Supported in part by the ISF and by the Technion TASP center.

1 Introduction

1.1 Background and Motivation

The difference between the speed of the actions of different entities is typically considered disruptive in real computing systems. In this paper, we illustrate some advantages of such phenomena in cases where the difference remains *fixed* throughout the execution¹. We demonstrate the *usefulness* of this manifestation of asynchrony to tasks involving symmetry breaking. More specifically, we show how two mobile agents, identical in every aspect save their speed, can lever their speed difference in order to achieve fast rendezvous.

Symmetry breaking is a major issue in distributed computing that is completely absent from traditional sequential computing. Symmetry can often prevent different processes from reaching a common goal. Well known examples include leader election [3], mutual exclusion [14], agreement [4,25] and renaming [6]. To address this issue, various differences between processes are exploited. For example, solutions for leader election often rely on unique identifiers assumed to be associated with each entity (e.g., a process) [3]. Another example of a difference is the location of the entities in a network graph. Entities located in different parts of a non-symmetric graph can use this knowledge in order to behave differently; in such a case, a leader can be elected even without using unique identifiers [26]. If no differences exist, breaking symmetry deterministically is often impossible (see, e.g., [3,27]) and one must resort to randomized algorithms, assuming that different entities can draw different random bits [19].

We consider mobile agents aiming to rendezvous. See, e.g., [7,13,22,23,24,27]. As is the case with other symmetry breaking problems, it is well known that if the agents are completely identical then rendezvous is, in some cases, impossible. In fact, a large portion of the research about rendezvous dealt with identifying the conditions under which rendezvous was possible, as a result of some asymmetries. Here, the fact that agents have different speeds implies that the mere feasibility of rendezvous is trivial, and our main concern is therefore the time complexity, that is, the time to reach a rendezvous. More specifically, we study the case where the agents are identical except for the fact that they have different speeds of motion. Moreover, to isolate the issue of the speed difference, we remove other possible differences between the agents. For example, the agents, as well as the graph over which the agents walk, are assumed to be anonymous. To avoid solutions of the kind of [26], we consider a symmetric graph, that is, specifically, a cycle topology. Further symmetry is obtained by hiding the graph features. That is, an agent views the graph as a continuous cycle of length n , and cannot even distinguish between the state it is at a node and the state it is walking over an edge.

¹ Advantages can also be exploited in cases where the difference in speed follows some stochastic distribution, however, in this initial study, we focus on the simpler fully deterministic case. That is, we assume a speed heterogeneity that is arbitrary but fixed throughout the execution.

1.2 The Model and the Problem

The problem of rendezvous on a cycle: Consider two identical deterministic agents placed on a cycle of length n (in some distance units). To ease the description, we name these agents A and B but these names are not known to the agents. Each agent is initially placed in some location on the cycle by an adversary and both agents start the execution of the algorithm simultaneously. An agent can move on the cycle at any direction. Specifically, at any given point in time, an agent can decide to either start moving, continue in the same direction, stop, or change its direction. The agents' goal is to *rendezvous*, namely, to get to be co-located somewhere on the cycle². We consider continuous movement, so this rendezvous can occur at any location along the cycle. An agent can detect the presence of another agent at its location and hence detect a rendezvous. When agents detect a rendezvous, the rendezvous task is considered completed.

Orientation issues: We distinguish between two models regarding orientation. The first assumes that agents have the *sense of direction* [8], that is, we assume that the agents can distinguish clockwise from the anti-clockwise. In the second model, we do not assume this orientation assumption. Instead, each agent has its own perception of which direction is clockwise and which is anti-clockwise, but there is no guarantee that these perceptions are globally consistent. (Hence, e.g., in this model, if both agents start walking in their own clockwise direction, they may happen to walk in opposite directions, i.e., towards each other).

The pebble and the white board models: Although the agents do not hold any direct means of communication, in some cases we do assume that an agent can leave marks in its current location on the cycle, to be read later by itself and by the other agent. In the *pebble* model, an agent can mark its location by dropping a pebble [10,11]. Both dropping and detecting a pebble are local acts taking place only on the location occupied by the agent. We note that in the case where pebbles can be dropped, our upper bound employs agents that drop a pebble only once and only at their initial location [1,9,24]. On the other hand, our corresponding lower bound holds for any mechanism of (local) pebble dropping. Moreover, this lower bound holds also for the seemingly stronger *'white board* model, in which an agent can change a memory associated with its current location such that it could later be read and further manipulated by itself or the other agent [20,16,17].

Speed: Each agent moves at the same fixed speed at all times; the *speed* of an agent A , denoted $s(A)$, is the inverse of the time t_α it takes agent A to

² In some sense, this rendezvous problem reminds also the *cow-path* problem, see, e.g., [5]. Here, the agents (the cow and the treasure she seeks to find) are both mobile (in the cow-path problem only one agent, namely, the cow, is mobile). It was shown in [5] that if the cow is initially located at distance D from the treasure on the infinite line then the time to find the treasure can be $9D$, and that 9 is the best multiplicative constant (up to lower order terms in D).

traverse one unit of length. For agent B , the time t_β and speed $s(B)$ are defined analogously. Without loss of generality, we assume that agent A is faster, i.e., $s(A) > s(B)$ but emphasize that this is unknown to the agents themselves. Furthermore, for simplicity of presentation, we normalize the speed of the slower agent B to one, that is, $s(B) = 1$ and denote $s(A) = c$ where $c > 1$. We stress that the more interesting cases are when c is a function of n and arbitrarily close to 1 (e.g., $c = 1 + 1/n^k$, for some constant k). We assume that each agent has a pedometer that enables it to measure the distance it travels. Specifically, a (local) *step* of an agent is a movement of one unit of distance (not necessarily all in the same direction, e.g., in one step, an agent can move half a step in one direction and the other half in the other direction). Using the pedometer, agents can count the number of steps they took (which is a real number at any given time). In some cases, agents are assumed to possess some knowledge regarding n and c ; whenever used, this assumption will be mentioned explicitly.

Time complexity: The *rendezvous time* of an algorithm is defined as the worst case time bound until rendezvous, taken over all pairs of initial placements of the two agents on the cycle. Note, a lower bound for the rendezvous time that is established assuming sense of direction holds trivially for the case where no sense of direction is assumed. All our lower bounds hold assuming sense of direction.

The Distributed Race (DR) algorithm: Let us consider a trivial algorithm, called *Distributed Race* (DR), in which an agent starts moving in an arbitrary direction, and continues to walk in that direction until reaching rendezvous. Note that this algorithm does not assume knowledge of n and c , does not assume sense of direction and does not leave marks on the cycle. The worst case for this algorithm is that both agents happen to walk on the same direction. Without loss of generality, assume this direction is clockwise. Let d denote the clockwise distance from the initial location of A to that of B . The rendezvous time t thus satisfies $t \cdot s(A) = t \cdot s(B) + d$. Hence, we obtain the following.

Observation 1. *The rendezvous time of DR is $d/(c-1) < n/(c-1)$.*

Note that in the cases where c is very close to 1, e.g., $c = 1 + 1/n^k$, for some constant k , the bound on the rendezvous time of DR is very large.

1.3 Our Results

Our first result is a lower bound showing that, up to a multiplicative approximation factor of 2, the bound of DR mentioned in Observation 1 is unavoidable, even in the white board model, even assuming sense of direction, and even assuming n and c are known to agents. That is, we show that under such assumptions, the rendezvous time of any algorithm is at least $\frac{n}{2(c-1)}$ if $c \leq 3$ and slightly larger (specifically, $\frac{n}{c+1}$) if $c > 3$. We then manage to construct an algorithm that matches the lower bound precisely for the case $c \leq 2$, and almost matches it when $c > 2$. Specifically, when $c \leq 2$, our algorithm runs in time $\frac{n}{2(c-1)}$ and when $c > 2$, the rendezvous time is n/c (yielding a $(2 - \frac{2}{c})$ -approximation when

$2 < c \leq 3$, and a $(\frac{c+1}{c})$ -approximation when $c > 3$). Moreover, our algorithm performs under weaker assumptions than those stated above, as it does not assume sense of direction, and allows agents to leave only a single mark (a pebble) and only at the place where they start the execution.

Finally, we investigate the setting in which no marks can be used at all, and show tight bounds for $c \leq 2$, and almost tight bounds for $c > 2$. Specifically, for this case, we establish a tight bound of $\frac{cn}{c^2-1}$ for the rendezvous time, in case agents have sense of direction. With the absence of sense of direction, the same lower bound of $\frac{cn}{c^2-1}$ holds, and we obtain an algorithm matching this bound for the case $c \leq 2$, and rather efficient algorithms for the case $c > 2$. Specifically, the rendezvous time for the case $c \leq 2$ is $\frac{cn}{c^2-1}$, for the case $2 < c \leq 3$, the rendezvous time is $\frac{2n}{c+1}$, and for the case $c > 3$, the rendezvous time is $\frac{n}{c-1}$.

2 Lower Bound for the White Board Model

The following lower bound implies that DR is a 2-approximation algorithm, and it becomes close to optimal when c goes to infinity.

Theorem 1. *Any rendezvous algorithm in the white board model requires at least $\max\{\frac{n}{2(c-1)}, \frac{n}{c+1}\}$ time, even assuming sense of direction and even assuming n and c are known to the agents.*

Proof. We assume that agents have sense of direction; hence, both agents start walking at the same direction. We first show that any algorithm in the white board model requires $\frac{n}{2(c-1)}$ time. Consider the case that the adversary locates the agents at *symmetric locations* of the cycle, i.e., they are at distance $n/2$ apart. Now consider any algorithm used by the agents. Let us fix any c' such that $1 < c' < c$, and define the (continuous) interval

$$I := \left[0, \frac{nc'}{2(c-1)}\right].$$

For every (real) $i \in I$, let us define the following (imaginary) scenario S_i . In scenario S_i , each agent executes the algorithm for i steps and terminates³. We claim that for every $i \in I$, the situation at the end of scenario S_i is completely symmetric: that is, the white board at symmetric locations contain the same information and the two agents are at symmetric locations. We prove this claim by induction. The basis of the induction, the case $i = 0$, is immediate. Let us assume that the claim holds for scenario S_i , for (real) $i \in I$, and consider scenario $S_{i+\epsilon}$, for any positive ϵ such that $\epsilon < \frac{n}{4}(1 - \frac{c'}{c})$. Our goal is to show that the claim holds for scenario $S_{i+\epsilon}$ ⁴.

³ We can think of this scenario as if each agent executes another algorithm B , in which it simulates precisely i steps of the original algorithm and then terminates.

⁴ Note that for some $i \in I$ and some $\epsilon < \frac{n}{4}(1 - \frac{c'}{c})$, we may have that $i + \epsilon \notin I$. Our proof will show that the claim for $S_{i+\epsilon}$ holds also in such cases. However, since we wish to show that the claim holds for S_j , where $j \in I$, we are not really interested in those cases, and are concerned only with the cases where $i + \epsilon \in I$ and $i \in I$.

Consider scenario $S_{i+\epsilon}$. During the time interval $[0, \frac{i}{c}]$, both agents perform the same actions as they do in the corresponding time interval in scenario S_i . Let a denote the location of agent A at time i/c . Now, during the time period $[\frac{i}{c}, \frac{i+\epsilon}{c}]$, agent A performs some movement all of which is done at distance at most ϵ from a (during this movement it may write information at various locations it visits); then, at time $\frac{i+\epsilon}{c}$, agent A terminates.

Let us focus on agent B (in scenario $S_{i+\epsilon}$) during the time period $[\frac{i}{c}, i]$. We claim that during this time period, agent B is always at distance at least ϵ from a . Indeed, as long as it is true that agent B is at distance at least ϵ from a , it performs the same actions as it does in scenario S_i (because it is unaware of any action made by agent A in scenario $S_{i+\epsilon}$, during the time period $[\frac{i}{c}, \frac{i+\epsilon}{c}]$). Therefore, if at some time $t' \in [\frac{i}{c}, i]$, agent B is (in scenario $S_{i+\epsilon}$) at distance less than ϵ from a then the same is true also for scenario S_i . However, in scenario S_i , by the induction hypothesis, agent B was at time i at \bar{a} , the symmetric location of a , that is, at distance $n/2$ from a . Thus, to get from a distance less than ϵ from a to \bar{a} , agent B needs to travel a distance of $n/2 - \epsilon$, which takes $n/2 - \epsilon$ time. This is impossible since

$$i - i/c \leq \frac{nc'}{2c} < \frac{n}{2} - \epsilon,$$

where the first inequality follows from the definition of I and the second follows from the definition of ϵ . It follows that during the time period from $[\frac{i}{c}, i]$, agent B behaves (in scenario $S_{i+\epsilon}$) the same as it does in the corresponding time period in scenario S_i . Therefore, according to the induction hypothesis, at the corresponding i 'th steps in scenario $S_{i+\epsilon}$, agent B is at distance $n/2$ from where agent A is at its i 'th step (recall, agent A is at a at its i 'th step), and the cycle configuration (including the white boards) is completely symmetric. Now, since $\epsilon < n/4$, during the time period $[i, i + \epsilon]$, agent B is still at a distance more than ϵ from a and remains unaware of any action made by agent A , during the time period $[\frac{i}{c}, \frac{i+\epsilon}{c}]$. (Similarly, agent A , during the time period $[\frac{i}{c}, \frac{i+\epsilon}{c}]$, is unaware of any action made by agent B during this time period.) Hence, at each time $i' \in [i, i + \epsilon]$, agent B takes the same action as agent A in the corresponding time i'/c . This establishes the induction proof. To sum up, we have just shown that for any $i \in I$, the cycle configuration at the end of scenario S_i is completely symmetric.

Now assume by contradiction that the rendezvous time t is less than the claimed one, that is, $t < \frac{n}{2(c-1)}$. At time t , both agents meet at some location u . Since $t \in I$, the above claim holds for S_t . Hence, at time t/c , agent A is at \bar{u} , the symmetric location of u . Since rendezvous happened at time t , this means that agent A traveled from \bar{u} to u (i.e., a distance of $n/2$) during the time period $[\frac{t}{c}, t]$. Therefore $t(1 - \frac{1}{c})c \geq n/2$, contradicting the assumption that $t < \frac{n}{2(c-1)}$. This establishes that any algorithm requires $\frac{n}{2(c-1)}$ time.

We now show the simpler part of the theorem, namely, that the rendezvous time of any algorithm in the white board model is at least $\frac{n}{c+1}$. Let us represent the cycle as the reals modulo n , that is, we view the cycle as the continuous

collection of reals $[0, n]$, where n coincides with 0. Assume that the starting point of agent A is 0. Consider the time period $T = [0, \frac{n}{c+1} - \epsilon]$, for some small positive ϵ . In this time period, agent A moves a total length of less than $\frac{nc}{c+1}$. Let r (and ℓ , correspondingly) be the furthest point from 0 on the cycle that A reached while going clockwise (or anti-clockwise, correspondingly), during that time period. Note that there is a gap of length larger than $n - \frac{nc}{c+1} = \frac{n}{c+1}$ between ℓ and r . This gap corresponds to an arc not visited by agent A during this time period. On the other hand, agent B walks a total distance of less than $\frac{n}{c+1}$ during the time period T . Hence, the adversary can locate agent B initially at some point in the gap between r and ℓ , such that during the whole time period T , agent B remains in this gap. This establishes the $\frac{n}{c+1}$ time lower bound, and concludes the proof of the theorem.

3 Upper Bound for the Pebble Model

Note that the assumptions of the pebble model are weaker than the white board model. Hence, in view of Theorem 1, the following theorem establishes a tight bound for the case where $c \leq 2$, a $(2 - \frac{2}{c})$ -approximation for the case $2 < c \leq 3$, and a $(c + 1)/c$ -approximation for the case $c > 3$.

Theorem 2. *There exists an algorithm that in the pebble model whose rendezvous time is $\max\{\frac{n}{2(c-1)}, \frac{n}{c}\}$. Moreover, this algorithm does not assume sense of direction, uses only one pebble and drops it only once: at the initial position of the agent. The algorithm assumes that agents know n and whether or not $c > 2$.*

Proof. Consider the following algorithm. Each agent (1) leaves a pebble at its initial position and then starts walking in an arbitrary direction while counting the distance travelled. If (2) an agent reaches a location with a pebble for the first time and (3) the distance it walked is strictly less than $\tau := \min\{n/2, n/c\}$, then (4) the agent turns around and walks continuously in the other direction.

First note that if both agents happen to start walking in opposite directions (due to lack of sense of direction), then they walk until they meet. In this simple case, their relative speed is $c + 1$, hence rendezvous happens in time at most $\frac{n}{c+1} < \max\{\frac{n}{2(c-1)}, \frac{n}{c}\}$. For the remaining of the proof, we consider the case that both agents start walking at the same direction, which is without loss of generality, the clockwise direction. Let d be the initial clockwise distance from A to B . Consider three cases.

1. $d = \tau$.

Here, no agent turns around. In other words, they behave exactly as in DR. If $d = n/2$, Observation 1 implies that the rendezvous time is $\frac{n}{2(c-1)}$. Otherwise, $c > 2$ and $d = n/c$. By Observation 1, the rendezvous is reached earlier, specifically, by time $\frac{d}{c-1} = \frac{n}{c(c-1)}$.

2. $d < \tau$.

In this case, Agent A will reach B 's starting point v_B , at time d/c , before B reaches A 's starting point v_A . Moreover, agent B does not turn, since its

initial distance to A 's starting point is at least τ . At time d/c , agent B is at distance d/c clockwise from v_B . By the algorithm, Agent A then turns and walks anti-clockwise. The anti-clockwise distance from A to B is then $n - d/c$. Their relative speed is $c + 1$. Hence, they will rendezvous in an additional time of $\frac{n-d/c}{1+c}$, since no agent may turn around after time d/c . Hence, the total time for reaching rendezvous is at most

$$d/c + \frac{n - d/c}{1 + c} = \frac{d + n}{1 + c}.$$

This function is maximized when $d = \tau$ where it is $\frac{\tau+n}{1+c}$. Now, if $c \leq 2$, we have $\tau = n/2$ and the rendezvous time is therefore $\frac{3n}{2(1+c)}$. Since $c \leq 2$, the later bound is at most $\frac{n}{2(c-1)}$. On the other hand, if $c > 2$, we have $\tau = n/c$ and the rendezvous time is n/c .

3. $d > \tau$.

In this case, A doesn't turn when it hits B 's initial position. Consider the following sub-cases.

(a) The agents meet before B reaches A 's initial position.

In this case, the rendezvous time (as in DR) is $d/(c-1)$. On the other hand, the rendezvous time is at most $n-d$ since B did not reach A 's initial position. So $d/(c-1) \leq n-d$. A simple calculation now implies that the rendezvous time $d/(c-1)$ is at most n/c .

(b) Agent B reaches A 's initial position before rendezvous.

In this case, Agent B walks for $d' = n - d$ time to first reach A 's initial position. We first claim that $d' < \tau$. One case is that $c \leq 2$, and thus, $\tau = n/2$. Since $d > \tau$, we have $d' < n/2 = \tau$. The other case is that $c > 2$, so $\tau = n/c$. We claim that also in this case, we have $d' < \tau$. Otherwise, we would have had $d' \geq n/c$, which would have meant that the faster agent A would have had, at least, n/c time before B reached the initial position of A . So much time would have allowed it to cover the whole cycle. This contradicts the assumption that B reached A 's initial position before rendezvous. This establishes the fact that, regardless of the value of c , we have $d' < \tau$. This fact implies that when agent B reaches A 's initial position, it turns around and both agents go towards each other. By the time B turns around, A has walked a distance of cd' . Hence, at that point in time, they are $n - cd'$ apart. This implies to the following rendezvous time:

$$d' + \frac{n - cd'}{1 + c} = \frac{2n - d}{1 + c}.$$

Now recall that we are in the case that agent B reaches A 's initial position before they rendezvous. This implies that $n - d < n/c$. Hence, the running time is at most

$$\frac{2n - d}{1 + c} < \frac{n + \frac{n}{c}}{1 + c} = \frac{n}{c}.$$

4 Rendezvous without Communication

In this section, we consider the case that agents cannot use marks (e.g., pebbles) to mark their location. More generally, the agents cannot communicate in any way (before rendezvous).

Theorem 3. *Consider the setting in which no communication is allowed, and assume both n and c are known to the agents.*

1. *The rendezvous time of any algorithm is, at least, $\frac{cn}{c^2-1}$, even assuming sense of direction.*
2. *Assuming sense of direction, there exists an algorithm whose rendezvous time is $\frac{cn}{c^2-1}$.*
3. *Without assuming sense of direction, there exists an algorithm whose rendezvous time is:*
 - $\frac{cn}{c^2-1}$, for $c \leq 2$,
 - $\frac{2n}{c+1}$, for $2 < c \leq 3$,
 - $\frac{n}{c-1}$, for $c > 3$.

Proof. Let us first show the first part of the theorem, namely, the lower bound.

Proof of Part 1: Given an algorithm, let \hat{t} denote the rendezvous time of the algorithm, that is, the maximum time (over all initial placements and all cycles of length n) for the agents (executing this algorithm) to reach rendezvous. Recall, in this part of the theorem, we assume that agents have sense of direction. Without loss of generality, we assume that the direction an agent starts walking is clockwise.

Consider, first, two identical cycles C_A and C_B of length n each. Let us mark a location $v \in C_A$ and a location $u \in C_B$. Let us examine the (imaginary) scenario in which agent A (respectively, B) is placed on the cycle C_A (respectively, C_B) alone, that is, the other agent is not on the cycle. Furthermore, assume that agent A is placed at v and agent B is placed at u . In this imaginary scenario, agents A and B start executing the protocol at the same time, separately, on each of the corresponding cycles. Viewing u as homologous to v , a homologous location $h(x) \in C_B$ can be defined for each location in $x \in C_A$ in a natural way (in particular, $h(v) = u$).

For each time $t \in [0, \hat{t}]$, let $d(t)$ denote the clockwise distance between the location $b_t \in C_B$ of the slower agent B at time t and the homologous location $h(a_t) \in C_B$ of the location $a_t \in C_A$ of the faster agent A at time t . Note that $d(t)$ is a real value in $[0, n)$. Initially, we assume that all reals in $[0, n)$ are colored *white*. As time passes, we *color* the corresponding distances by black, that is, at every time t , we color the distance $d(t) \in [0, n)$ by black. Note that the size of the set of black distances is monotonously non-decreasing with time.

We first claim that, by time \hat{t} , the whole range $[0, n)$ is colored black. To prove by contradiction, assume that there is a real $d \in [0, n)$ that remains white. Now, consider the execution on a single cycle of length n , where agent A is initially

placed at anti-clockwise distance d from agent B . In such a case, rendezvous is not reached by time \hat{t} , contrary to the assumption that it is. This implies that the time T it takes until all reals in $[0, n)$ are colored black in the imaginary scenario, is a lower bound on the rendezvous time, that is, $T \leq \hat{t}$. It is left to analyze the time T .

With time, the change in the distance $d(t)$ has to follow two rules:

- **R1.** After one time unit, the distance can change by at most $1 + c$ (the sum of the agents' speeds).
- **R2.** At time x , the distance, in absolute value is, at most, $x(c - 1)$.

To see why Rule R2 holds, recall that the programs of the agents are identical. Hence, if agent A is at some point $a \in C_A$, at A 's k th step, then agent B is at point $b = h(a)$ at its own k th step. This happens for agent B at time k and for agent A at time k/c . Since A 's speed is c , the maximum it can get away from point a during the time period from k/c until k is $c(k - k/c) = k(c - 1)$.

Consider the path $P := d(t)$ covering the range $[0, n)$ in T time. First, note that this path P may go several times through the zero point (i.e., when $d(t) = 0$). At a given time s , we say that the path is *on the right* side, if the last time it left the zero point before time s was while going towards 1. Similarly, the path is *on the left*, if the last time it left the zero point before time s was while going towards $n - 1$. Let x denote the largest point on the range $[0, n)$ reached by the path while the path was on the right side. Let $y = n - x$. By time T , path P had to go left to distance y from the zero point. Assume, w.l.o.g. that $x < y$. (In particular, x may be zero.) The fastest way to color these two points (and all the rest of the points, since those lie between them), would be to go from zero to the right till reaching x , then return to zero and go to distance y on the left. Hence, T will be at least: $T \geq \frac{x}{c-1} + \frac{y}{c+1}$. Indeed, Rule R2, applied to the time of reaching distance x , implies the first term above. The second term uses Rule R1 to capture the time that starts with the distance reaching x , proceeds with the distance reaching the zero point and ends when the distance reaches y when going left. Since $c > 1$, we obtain

$$x \leq \left(T - \frac{n}{c+1} \right) (c-1). \quad (1)$$

On the other hand, applying Rule R2 to the final destination y , we have $T(c - 1) \geq y = n - x$. This implies that:

$$x \geq n - T(c-1). \quad (2)$$

Combining Equations 1 and 2, we get $T \geq \frac{cn}{c^2-1}$. This establishes the first part of the theorem. We now prove the second and third parts of the theorem, namely, the upper bounds.

Proof of Parts 2 and 3: Note that the DR algorithm does not assume sense of direction, and its complexity is the one required by the third part of the theorem for the case $c > 3$. We now revise DR and consider an algorithm, called

Turn(k), which consists of two stages. At the first stage, the agent walks over its own clockwise direction for k steps. Subsequently (if rendezvous hasn't occurred yet), the agent executes the second stage of the algorithm: it turns around and goes in the other direction until rendezvous.

Let us now concentrate on proving the second part of the theorem, and assume that agents have sense of direction. Here we consider Algorithm **Turn**(k), with parameter $k = \frac{cn}{c^2-1}$. Since we assume sense of direction, both agents walk initially in the same direction (clockwise). Assume by contradiction that rendezvous hasn't occurred by time k . By that time, agent B took k steps. Agent A took those k steps by time k/c . At that time, agent A turns. (However, agent B will turn only at time k). Hence, between those two turning times, there is a time duration $k(1 - \frac{1}{c})$ where the two agents walk towards each other. Hence, at each time unit they shorten the distance between them by $1 + c$. Let d' denote the maximum distance between the agents at time k/c . It follows that $d' > k(1 - \frac{1}{c})(1 + c) = n$. A contradiction. This establishes the second part of the theorem.

We now prove the third part of the theorem, which focuses on the case where no sense of direction is assumed. Let us first consider the case $c \leq 2$. Here we apply the same algorithm above, namely Algorithm **Turn**(k), with parameter $k = \frac{cn}{c^2-1}$. As proved before, if the two agents happen to start at the same direction then rendezvous occurs by time $\frac{cn}{c^2-1}$. Hence, let us consider the case that both agents walk initially at opposite directions. Assume by contradiction, that rendezvous hasn't occurred by time k/c . In this case, by time k/c the faster agent A walked k steps toward B , and the slower agent B walked k/c steps towards A . Hence, the initial distance between them must be greater than $k(1 + 1/c) = n/(c - 1) > n$, a contradiction. This proves the first item in Part 3 of the Theorem.

Note that Algorithm DR establishes the third item in Part 3 of the Theorem. Hence, it is left to prove the second item in Part 3, namely, the case $2 < c \leq 3$. For this case, we apply Algorithm **Turn**(k), with parameter $k = \frac{cn}{c+1}$. First note, if the two agents happen to start walking towards each other they continue to do that until time $k/c = \frac{n}{c+1}$, hence they would meet by this time. Therefore, we may assume that initially, both agents start in the same direction. In this case, if rendezvous haven't occurred by time k/c , then at this time agent A turns around, and both agents walk towards each other for at least $k(1 - 1/c)$ more time (the time when agent B is supposed to turn). Since $c > 2$, we have $k(1 - 1/c) > k/c$, and hence, from the time agent A turns, both agents walk towards each other for at least $\frac{k}{c} = \frac{n}{c+1}$ time, which means they must meet by this time. Altogether, the time to rendezvous is $\frac{2k}{c} = \frac{2n}{c+1}$, as desired.

5 Discussion and Future Work

We show how some form of asynchrony could be useful for solving a symmetry breaking problem efficiently. Our study could be considered as a first attempt to harness the (unknown) heterogeneity between individuals in a cooperative population towards more efficient functionality.

There are many natural ways of further exploring this idea in future work. First, we have studied the exploitation of asynchrony for a specific kind of problems. It seems that it can be useful for other symmetry breaking problems as well. Another generalization: the “level” of asynchrony considered in this paper is very limited: the ratio c between the speeds of the agents is the same throughout the execution, and is known to the agents. Exploiting a higher level of asynchrony should also be studied, for example, the case that the speed difference is stochastic and changes through time.

Our main interest was the exploitation of asymmetry, rather than the specific problem of rendezvous. Still, even for the rendezvous problem, this current study leaves many open questions. First, even for the limited case we study, not all our bounds are completely tight, and it would be interesting to close the remaining gaps between our lower and upper bounds (these gaps hold for some cases when $c > 2$). In addition, it would be interesting to study further the uniform case [21], in which n and c are not known to agents. Another direction is to generalize the study to multiple agents (more than two, see, e.g., [15,18]) and to other graph classes. This would also allow one to study interactions between various means of breaking symmetry (such as different speeds together with different locations on a graph).

References

1. Alpern, S.: Rendezvous Search: A Personal Perspective. LSE Research Report, CDAM-2000-05, London School of Economics (2000)
2. Alpern, S., Gal, S.: The Theory of Search Games and Rendezvous. Kluwer Academic Publishers (2003)
3. Angluin, D.: Local and global properties in networks of processors. In: ACM STOC 1980, pp. 82–93 (1980)
4. Attiya, H., Gorbach, A., Moran, S.: Computing in Totally Anonymous Asynchronous Shared Memory Systems. *Information and Computation* 173(2), 162–183 (2002)
5. Baeza-Yates, R.A., Culberson, J.C., Rawlins, G.J.E.: Searching in the plane. *Inform. and Comput.* 106(2), 234–252 (1993)
6. Attiya, H., Bar-Noy, A., Dolev, D., Peleg, D., Reischuk, R.: Renaming in an asynchronous environment. *Journal of the ACM* 37(3), 524–548 (1990)
7. Bampas, E., Czyzowicz, J., Gašieniec, L., Ilcinkas, D., Labourel, A.: Almost Optimal Asynchronous Rendezvous in Infinite Multidimensional Grids. In: Lynch, N.A., Shvartsman, A.A. (eds.) DISC 2010. LNCS, vol. 6343, pp. 297–311. Springer, Heidelberg (2010)
8. Barrière, L., Flocchini, P., Fraigniaud, P., Santoro, N.: Rendezvous and Election of Mobile Agents: Impact of Sense of Direction. *Theory Comput. Syst.* 40(2), 143–162 (2007)
9. Baston, V., Gal, S.: Rendezvous search when marks are left at the starting points. *Naval Research Logistics* 48(8), 722–731 (2001)
10. Bender, M.A., Fernandez, A., Ron, D., Sahai, A., Vadhan, S.: The power of a pebble: Exploring and mapping directed graphs. In: Proc. 30th ACM Symp. on Theory of Computing (STOC), pp. 269–287 (1998)

11. Blum, M., Kozen, D.: On the power of the compass (or, why mazes are easier to search than graphs). In: 19th Annual Symposium on Foundations of Computer Science (FOCS 1978), October 16-18 (1978)
12. Czyzowicz, J., Dobrev, S., Krizanc, D., Kranakis, E.: The Power of Tokens: Rendezvous and Symmetry Detection for Two Mobile Agents in a Ring. In: Geffert, V., Karhumäki, J., Bertoni, A., Preneel, B., Návrat, P., Bieliková, M. (eds.) SOFSEM 2008. LNCS, vol. 4910, pp. 234–246. Springer, Heidelberg (2008)
13. Czyzowicz, J., Ilcinkas, D., Labourel, A., Pelc, A.: Asynchronous deterministic rendezvous in bounded terrains. *Theor. Comput. Sci.* 412(50), 6926–6937 (2011)
14. Dijkstra, E.W.: Self-stabilizing systems in spite of distributed control. *Communications of the ACM* 17(11), 643–644 (1974)
15. Dobrev, S., Flocchini, P., Prencipe, G., Santoro, N.: Multiple Agents Rendezvous in a Ring in Spite of a Black Hole. In: Papatriantafyllou, M., Huneil, P. (eds.) OPODIS 2003. LNCS, vol. 3144, pp. 34–46. Springer, Heidelberg (2004)
16. Flocchini, P., Nayak, A., Schulz, A.: Cleaning an Arbitrary Regular Network with Mobile Agents. In: Chakraborty, G. (ed.) ICDCIT 2005. LNCS, vol. 3816, pp. 132–142. Springer, Heidelberg (2005)
17. Flocchini, P., Ilcinkas, D., Santoro, N.: Ping Pong in Dangerous Graphs: Optimal Black Hole Search with Pebbles. *Algorithmica* 62(3-4), 1006–1033 (2012)
18. Flocchini, P., Kranakis, E., Krizanc, D., Santoro, N., Sawchuk, C.: Multiple Mobile Agent Rendezvous in a Ring. In: Farach-Colton, M. (ed.) LATIN 2004. LNCS, vol. 2976, pp. 599–608. Springer, Heidelberg (2004)
19. Itai, A., Rodeh, M.: Symmetry Breaking in Distributive Networks. In: FOCS, pp. 150–158 (1981)
20. Korach, E., Kutten, S., Moran, S.: A Modular Technique for the Design of Efficient Distributed Leader Finding Algorithms. In: ACM PODC 1985, pp. 163–174 (1985)
21. Korman, A., Sereni, J.-S., Viennot, L.: Toward More Localized Local Algorithms: Removing Assumptions Concerning Global Knowledge. In: ACM PODC 2011, pp. 49–58 (2011)
22. Kranakis, E., Krizanc, D., Markou, E.: The Mobile Agent Rendezvous Problem in the Ring. Morgan & Claypool Publishers (2010)
23. Kranakis, E., Krizanc, D., Rajsbaum, S.: Mobile Agent Rendezvous: A Survey. In: Flocchini, P., Gašieniec, L. (eds.) SIROCCO 2006. LNCS, vol. 4056, pp. 1–9. Springer, Heidelberg (2006)
24. Kranakis, E., Santoro, N., Sawchuk, C.: Mobile Agent Rendezvous in a Ring. In: ICDCS 2003, pp. 592–599 (2003)
25. Suzuki, I., Yamashita, M.: Distributed Anonymous Mobile Robots: Formation of Geometric Patterns. *SIAM J. Computing*. 28(4), 1347–1363 (1999)
26. Yamashita, M., Kameda, T.: Computing on an Anonymous Network. In: ACM PODC 1988, pp. 117–130 (1988)
27. Yu, X., Yung, M.: Agent Rendezvous: A Dynamic Symmetry-Breaking Problem. In: Meyer auf der Heide, F., Monien, B. (eds.) ICALP 1996. LNCS, vol. 1099, pp. 610–621. Springer, Heidelberg (1996)