

Computing Without Communicating: Ring Exploration by Asynchronous Oblivious Robots*

Paola Flocchini [†] David Ilcinkas [‡] Andrzej Pelc [§] Nicola Santoro [¶]

Abstract

We consider the problem of exploring an anonymous unoriented ring by a team of k identical, oblivious, asynchronous mobile robots that can view the environment but cannot communicate. This weak scenario is standard when the spatial universe in which the robots operate is the two-dimensional plane, but (with one exception) has not been investigated before for networks. Our results imply that, although these weak capabilities of robots render the problem considerably more difficult, ring exploration by a small team of robots is still possible.

We first show that, when k and n are not co-prime, the problem is not solvable in general, e.g., if k divides n there are initial placements of the robots for which gathering is impossible. We then prove that the problem is always solvable provided that n and k are co-prime, for $k \geq 17$, by giving an exploration algorithm that always terminates, starting from arbitrary initial configurations. Finally, we consider the minimum number $\rho(n)$ of robots that can explore a ring of size n . As a consequence of our positive result we show that $\rho(n)$ is $O(\log n)$. We additionally prove that $\Omega(\log n)$ robots are necessary for infinitely many n .

Keywords: mobile robots, asynchronous, oblivious, exploration, ring.

Corresponding Author: David Ilcinkas

LaBRI, bât A30, Université Bordeaux I

351 cours de la Libération

F-33405 TALENCE Cedex

FRANCE

Phone: +33 540 006 912

Fax: +33 540 006 669

*A preliminary version of this paper appeared in [27].

[†]School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, Ontario K1N 6N5, Canada.
E-mail: flocchin@eecs.uottawa.ca

[‡]LaBRI, CNRS & Université de Bordeaux, France. E-mail: david.ilcinkas@labri.fr.

[§]Département d'informatique, Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada. E-mail: pelc@uqo.ca.

[¶]School of Computer Science, Carleton University, Ottawa, Ontario, K1S 5B6, Canada.
E-mail: santoro@scs.carleton.ca

1 Introduction

1.1 Framework

Recently a lot of attention has been devoted to the computational and complexity issues arising in systems of autonomous mobile entities located in a spatial universe \mathcal{U} . The entities have storage and processing capabilities, exhibit the same behavior (i.e., execute the same protocol), can move in \mathcal{U} (their movement is constrained by the nature of \mathcal{U}), and are asynchronous in their actions. Depending on the context, the entities are sometimes called *agents*, other times *robots*; in the following, we use the latter. The research concern is on determining what tasks can be performed by such entities, under what conditions, and at what cost. In particular, a central question is to determine what minimal hypotheses allow a given problem to be solved.

Depending on the nature of \mathcal{U} , there are two basic settings in which autonomous mobile entities are being investigated. The first setting, called sometimes the *continuous universe*, is when \mathcal{U} is the two-dimensional plane (e.g., [1, 11, 12, 13, 30, 39, 29, 41, 42]). The second setting, sometimes called the *graph world* or the *discrete universe*, is when \mathcal{U} is a simple graph (e.g., [4, 5, 8, 17, 31, 32]). In both settings, each robot operates in a *Look - Compute - Move* cycle. The robot observes and interacts with the environment (Look), then, based on this observation and interaction, it decides to stay idle or to move (Compute), and in the latter case it moves towards its destination (Move).

Interestingly, in spite of the common features of the two settings, the researchers investigating them usually operate under two radically different assumptions on the robots' capabilities.

(1) *Communication vs Vision* - In the investigations in the graph world, the robots are assumed to communicate with each other directly; e.g., by means of tokens [7, 8, 28], or whiteboards [17, 31], or when they meet [14, 16]. Instead, in the studies on the continuous universe, the robots do not communicate in any explicit way; however, they see the position of the other robots and can acquire knowledge from this information (e.g., see [1, 11, 12, 13, 30, 39, 29, 41, 42]).

(2) *Persistency vs Obliviousness* - In addition to its program, each robot has a local memory (sometimes called notebook or workspace), used for computations and to store information obtained during the cycles. In all the investigations in the graph world, the local memory is possibly limited (e.g., each robot is a finite-state automaton) but almost always *persistent*: unless explicitly erased by the robot, all the information contained in the workspace will persist throughout the robot's cycles. Instead, in the majority of the studies on the continuous universe, the robots are *oblivious*: all the information contained in the workspace is *cleared* at the end of each cycle. In other words, the robots have *no* memory of past actions and computations, and the computation is based solely on what has been determined in the current cycle. The importance of obliviousness comes from its link to *self-stabilization* and *fault-tolerance*.

Let us point out that there is nothing inherent in the nature of \mathcal{U} that forces these differences in the assumptions. In other words, there is no reason why robots in a graph should not be oblivious; on the contrary, an oblivious solution would be highly desirable ensuring fault-tolerance and self-stabilization. Similarly, there is nothing in the continuous domain that forbids robots from communicating explicitly; indeed, in the recent investigations on mobile sensor networks, the robots

do communicate wirelessly [37].

Surprisingly, nobody has investigated how to solve problems in the discrete universe if the robots have the capabilities and limitations standard in the continuous one. In fact, with one exception, there were no previous studies on how a collection of asynchronous oblivious robots endowed with vision can perform a non-trivial task without any communication. The only exception is the recent investigation of the *gathering* problem in the ring [34].

In this paper, we continue this investigation and focus on a basic problem in the graph world, *exploration*, that is the process by which every node of the graph is visited by at least one robot, and we study this problem in a *ring*.

1.2 Our results

We consider the problem of exploring an anonymous ring of size n by k oblivious anonymous asynchronous robots scattered in the ring. The robots are endowed with vision but they are unable to communicate. Within finite time and regardless of the initial placement of the robots, each node must be visited by at least one robot and the robots must be in a configuration in which they all remain idle.

We first show that this problem is in general unsolvable if $k|n$ (i.e., there are initial configurations for which the exploration cannot be done). We then prove that, whenever $\gcd(n, k) = 1$, for $k \geq 17$, the robots can explore the ring terminating within finite time. The proof is constructive: we present a terminating protocol that explores the ring starting from an arbitrary initial configuration, and prove its correctness.

Finally, we consider the minimum number $\rho(n)$ of robots that can explore a ring of size n . As a consequence of our positive result we show that $\rho(n)$ is $O(\log n)$. We also prove that $\rho(n) = \Omega(\log n)$ for infinitely many n . More precisely, there exists a constant c such that, for infinitely many n , we have $\rho(n) \geq c \log n$. Interestingly, for some specific values of n , the number $\rho(n)$ is constant. This is the case, for example, of any prime n greater than 17, in which case 17 robots can always explore the ring.

1.3 Related Work

Algorithms for graph exploration by mobile entities (robots) have been intensely studied in recent literature. Several scenarios have been considered. Most of the research is concerned with the case of a single robot exploring the graph. In [2, 7, 8, 19, 25] the robot explores strongly connected directed graphs and it can move only in the direction from head to tail of an edge, not vice-versa. In particular, [19] investigates the minimum time of exploration of directed graphs, and [2, 25] give improved algorithms for this problem in terms of the deficiency of the graph (i.e., the minimum number of edges to be added to make the graph Eulerian). Many papers, e.g., [20, 23, 24, 3, 38] study the scenario where the explored graph is undirected and the robot can traverse edges in both directions. In [20] the authors investigate the problem of how the availability of a map influences

the efficiency of exploration. In [38] it is shown that a graph with n nodes and e edges can be explored in time $e + O(n)$. In some papers, additional restrictions on the moves of the robot are imposed. It is assumed that the robot has either a restricted tank [6, 9], forcing it to periodically return to the base for refueling, or that it is tethered, i.e., attached to the base by a rope or cable of restricted length [24].

Exploration of anonymous graphs presents different difficulties. In this case it is impossible to explore arbitrary graphs by a single robot if no marking of nodes is allowed. Hence the scenario adopted in [7, 8] allows the use of *pebbles* which the robot can drop on nodes to recognize already visited ones, and then remove them and drop in other places. The authors concentrate attention on the minimum number of pebbles allowing efficient exploration and mapping of arbitrary directed n -node graphs. (In the case of undirected graphs, one pebble suffices for efficient exploration.) In [8] the authors compare exploration power of one robot with a constant number of pebbles to that of two cooperating robots, and give an efficient exploration algorithm for the latter scenario. In [7] it is shown that one pebble is enough if the robot knows an upper bound on the size of the graph, and $\Theta(\log \log n)$ pebbles are necessary and sufficient otherwise.

In all the above papers, except [8], exploration is performed by a single robot. Exploration by many robots has been investigated mostly in the context when moves of the robots are centrally coordinated. In [32], approximation algorithms are given for the collective exploration problem in arbitrary graphs. In [4, 5] the authors construct approximation algorithms for the collective exploration problem in weighted trees. On the other hand, in [31] the authors study the problem of distributed collective exploration of trees of unknown topology. However, the robots performing exploration have memory and can directly communicate with each other. Exploration of arbitrary anonymous graphs by a team of robots communicating through *whiteboards* has been studied in [17].

The very weak assumption of asynchronous identical robots that cannot send any messages and that communicate with the environment only by observing it, has not been used in the context of graph exploration prior to the announcement of the preliminary results of this paper [27].

It has been used quite extensively, however, in the case of robots moving freely in the plane (e.g., see [1, 11, 12, 13, 15, 18, 30, 39, 42]), where the robots were oblivious, i.e., it was assumed that they do not have any memory of past observations. Oblivious robots operate in Look-Compute-Move cycles, similar to those described in our scenario. The differences are in the amount of synchrony assumed in the execution of the cycles. In [1, 18, 42] cycles were executed synchronously in rounds by all active robots, and the adversary could only decide which robots are active in a given cycle. In [11, 12, 13, 15, 30, 39] they were executed asynchronously: the adversary could interleave operations arbitrarily, stop robots during the move, and schedule Look operations of some robots while others were moving. Our scenario is very similar to the asynchronous model used in [30, 39]. The only difference with respect to [30, 39] is in the execution of Move operations. All possibilities of the adversary concerning interleaving operations performed by various robots as well as the characteristics of the robots are the same as in the model from [30, 39].

This scenario has been introduced in the discrete context of graphs in [34] and used in [33, 36] to

study the gathering problem in the ring. After the conference version of this paper [27], the exploration problem has been investigated also in this model. In particular, in [26] we have considered the tree network showing that there are n -node trees where $\Omega(n)$ robots are necessary and that, if the maximum degree is three, it is possible to explore the tree with an optimal team of $O(\frac{\log n}{\log \log n})$ robots. Exploration in the ring has been object of investigation in [21, 22, 35]. In particular, in [22] the authors have shown that in a probabilistic setting four robots are necessary and sufficient to perform exploration without any constraint on the relationship between number of nodes and number of robots and they have given an optimal algorithm for any ring size greater than eight. In [21] an optimal algorithm for the probabilistic exploration of rings of size smaller than eight was given. In [35], on the other hand, it has been shown that no deterministic exploration is feasible with less than five robots whenever the size of the ring is even, and that five robots are sufficient for any n that is co-prime with five. Finally, in [10] the authors investigated a similar but stronger model with labeled edges and provided necessary and sufficient conditions for explorability of arbitrary graphs for $k = 3, 4$ robots and $k > 4$ odd robots.

2 Preliminaries

2.1 Terminology and definitions

The network we consider is a ring of n nodes, u_0, u_1, \dots, u_{n-1} ; i.e., u_i is connected to both u_{i-1} and u_{i+1} . Here and in the following, all operations on the indices are modulo n . The indices are used for notation purposes; in fact, the nodes are anonymous (i.e., identical) and the ring is unoriented. Operating in the ring are k identical robots; initially, at time $t = 0$, there is at most one robot in each node. During the exploration, robots move, and at any time they occupy some nodes of the ring.

Notice that the assumption that in the initial configuration there is at most one robot in each node cannot be removed. In the absence of this assumption any configuration could be initial, but then no terminating exploration algorithm for any $k < n$ would be possible. Indeed, consider any hypothetical terminating exploration algorithm that works correctly for all initial configurations for some parameters $k < n$. Suppose that starting from some initial configuration I it finishes exploration in configuration I' . Now suppose that the initial configuration is I' . Since in the first scenario this is a final configuration (in which no robots move), in the second scenario no robots move either (i.e., they never move) and thus some nodes are never visited.

We shall indicate by $d_i(t)$ the multiplicity of robots present at node u_i at time t ; more precisely $d_i(t) = 0$ indicates that there are no robots, $d_i(t) = 1$ indicates that there is exactly one robot, and $d_i(t) = 2$ indicates that there is more than one robot at u_i at time t . If $d_i(t) = 2$, we will say that there is a *tower* in u_i at time t .

Let $\delta^{+j}(t)$ denote the sequence $\delta^{+j}(t) = \langle d_j(t) \ d_{j+1}(t) \ \dots \ d_{j+n-1}(t) \rangle$, and let $\delta^{-j}(t)$ denote the sequence $\delta^{-j}(t) = \langle d_j(t) \ d_{j-1}(t) \ \dots \ d_{j-(n-1)}(t) \rangle$. The unordered* pair of sequences $\delta^{+j}(t)$ and

*Since the ring is unoriented, agreement on only one sequence might be impossible, and the pair cannot be ordered.

$\delta^{-j}(t)$ describes the configuration of the system at time t viewed from node u_j . Let $\Delta^+(t) = \{\delta^{+j}(t) : 0 \leq j < n\}$ and $\Delta^-(t) = \{\delta^{-j}(t) : 0 \leq j < n\}$.

We will denote by $\delta_{max}(t)$ the ascending lexicographically maximum sequence in $\Delta^+(t) \cup \Delta^-(t)$. Since the lexicographic order is total, there is at most one maximal sequence in each of $\Delta^+(t)$ and $\Delta^-(t)$. A configuration is said to be *symmetric* if the maximal sequences in $\Delta^+(t)$ and $\Delta^-(t)$ are equal, and *asymmetric* otherwise.

As an example of the above definitions, consider a ring of 7 nodes, a time t and a node u_j , for which $\delta^{+j}(t) = \langle 1201101 \rangle$ and $\delta^{-j}(t) = \langle 1101102 \rangle$. Thus $\Delta^+(t)$ is composed of sequences $\langle 1201101 \rangle$, $\langle 2011011 \rangle$, $\langle 0110112 \rangle$, $\langle 1101120 \rangle$, $\langle 1011201 \rangle$, $\langle 0112011 \rangle$, $\langle 1120110 \rangle$, and $\Delta^-(t)$ is composed of sequences $\langle 1101102 \rangle$, $\langle 1011021 \rangle$, $\langle 0110211 \rangle$, $\langle 1102110 \rangle$, $\langle 1021101 \rangle$, $\langle 0211011 \rangle$, $\langle 2110110 \rangle$. The maximal sequence in $\Delta^+(t)$ is $\langle 2011011 \rangle$ and in $\Delta^-(t)$ it is $\langle 2110110 \rangle$, hence the configuration is asymmetric. We have $\delta_{max}(t) = \langle 2110110 \rangle$.

Each robot operates in Look-Compute-Move cycles described in section 1.1. Cycles are performed asynchronously for each robot: the time between Look, Compute, and Move operations is finite but unbounded, and is decided by the adversary for each action of each robot. The only constraint is that moves are instantaneous, as in [34], and hence any robot performing a Look operation sees all other robots at nodes of the ring and not on edges. However, a robot \mathcal{R} may perform a Look operation at some time t , perceiving robots at some nodes, then Compute a target neighbor at some time $t' > t$, and Move to this neighbor at some later time $t'' > t'$ in which some robots are in different nodes from those previously perceived by \mathcal{R} because in the meantime they performed their Move operations. Hence robots may move based on significantly outdated perceptions. We assume that the robots can perceive, during the Look operation, if there is one or more robots in a given location; this ability, called *multiplicity detection* is a standard assumption in the continuous model [11, 34, 39]. We now describe formally what a robot perceives when performing a Look operation. Consider a robot \mathcal{R} that, at time t is at node u_j and performs a Look; the result of this operation, called the *view* of \mathcal{R} at time t , is precisely the unordered pair of sequences $\{\delta^{+j}(t), \delta^{-j}(t)\}$, that is, the configuration of the system at time t viewed from node u_j . We order all views as follows: we use the ascending lexicographic order on the sequences and order the views according to the order of the maximum of the two sequences forming them. From its view, the robot can determine $\delta_{max}(t)$, decide whether or not it is unique, and compute views of all other robots.

Let robot \mathcal{R} perform in the same cycle a Look operation at time t' and a Move operation at time $t'' > t'$. We will say that \mathcal{R} is *engaged to move* (or, simply *engaged*) in the open interval (t', t'') ; that is, \mathcal{R} is engaged at any time t , where $t' < t < t''$.

One final precision has to be added, concerning the decisions of robots made during the Compute action. Every such decision is based on the snapshot obtained during the last Look action. However, it may happen that both edges incident to a node v currently occupied by the deciding robot look identical in this snapshot, i.e., v lies on a symmetry axis of the configuration. In this case, if the robot decides to take one of these edges, it may take any of the two. We assume the worst-case decision in such cases, i.e., that the actual edge among the identically looking ones is chosen by an adversary.

We say that exploration of an n -node ring is possible with k robots, if there exists an algorithm which, starting from any initial configuration of the k robots without towers, allows the robots to explore the entire ring and brings all robots to a configuration in which they all remain idle. Obviously, if $n = k$, the exploration is already accomplished, hence we always assume that $k < n$.

2.2 Basic restriction

Lemma 2.1 *Let $k < n$. If $k \nmid n$ then the exploration of an n -node ring with k robots is not possible.*

Proof: By contradiction, let P be a solution protocol. Choose as the initial configuration an equidistant placement of the k robots in the ring (it exists since $k \mid n$). Thus, initially the states of all robots are identical, say $\sigma(0)$. Clearly this state is not a terminal state. Otherwise, since $k < n$, P would terminate without exploring the ring, thus contradicting the correctness of P . Consider now an adversary that uses a synchronous scheduler and a consistent orientation of the ring. Then, at each time step t , the states of all robots continue to be identical, say $\sigma(t)$, and furthermore they are the same as those of previous steps; i.e., $\sigma(t) = \sigma(0)$ for all t . Hence the robots will never enter a terminal state, contradicting the fact that P leads within finite time to a configuration in which all robots remain idle. \square

In the following we will consider the case when $\gcd(n, k) = 1$, and design an algorithm that allows $k \geq 17$ robots to explore a n -node ring whenever $\gcd(n, k) = 1$. Observe that if $\gcd(n, k) = 1$, the configuration is aperiodic and either asymmetric or it is symmetric with respect to a single axis of symmetry. Therefore at most two robots can have a given view. In the symmetric case, the adjective symmetric will be used with respect to this unique axis of symmetry. Note that symmetric robots have the same view.

3 Exploration of a ring

3.1 Overview of the algorithm

The overall structure of the algorithm can be seen as a sequence of three distinct phases: *Set-Up*, *Tower-Creation*, and *Exploration*.

The purpose of the *Set-Up* phase is to transform the (arbitrary) initial configuration into one from a predetermined set of configurations (called *no-towers-final*) with special properties. More precisely, in the *Set-Up* phase, the robots create a configuration where there is a single set of consecutive nodes occupied by robots, or two such sets of the same size (called blocks). When the configuration is *no-towers-final*, the next phase begins.

The purpose of the *Tower-Creation* phase is to transform the *no-towers-final* configuration created in the previous phase, into one from a predetermined set of configurations (called *towers-completed*) in which everything is prepared for exploration to begin. More precisely, in the *Tower-Creation*

phase, one or two towers are created inside each block (the number depending on the parity of the size of the block); furthermore a number of robots become uniquely identified as explorers. As soon as the configuration is *towers-completed*, the next phase begins.

During the *Exploration* phase, the ring is actually being explored. The configuration reached upon exploration depends solely on the configuration at the beginning of this phase. The set of these special *terminal* configurations is uniquely identified, and once in a configuration of this type, no robots will make any further move.

Since the robots are oblivious (i.e., they have no recollection of actions and computations made in previous cycles), there is no explicit way for them to record which phase is the current one. This information is derived by a robot solely based on the configuration currently observed (i.e., the one obtained as a result of the Look operation). Since the determination of the phase should be non-ambiguous, each reachable configuration should be assigned to exactly one phase.

For any possible configuration we will identify a set of *players*, which are the robots deciding to move, if they perform a Look operation in this configuration, and corresponding *destinations*, i.e., target neighbors. The exploration algorithm (which contains the rules describing the Compute actions in the robot's cycle) can be succinctly formulated as follows.

<p>Algorithm RING EXPLORATION</p> <p>If I am a <i>player</i> move to my <i>destination</i></p>
--

3.2 Set-Up Phase

The first phase of the protocol is the *Set-Up*. The fact of being in this phase is easily recognizable by the robots since, unlike those of the other phases, the configurations of this phase contain no towers. Precisely because they contain no towers, any configuration of this phase can be an initial configuration.

We define the *interdistance* of a configuration as the minimum distance taken over all pairs of distinct robots in the configuration. Given an arbitrary configuration of interdistance d , a *block* is a maximal set of robots, of size at least 2, forming a line with a robot every d nodes. The *size* of a block is the number of robots it contains. The *border* of a block are the two nodes occupied by the two extremal robots of the block. A robot not in a block is said to be *isolated*. A robot is said to be a *neighbor* of a block/robot if in at least one direction there is no robot between itself and the block/robot. A *leader* of a configuration is a robot from which the view is the maximal in the configuration, with respect to the order defined in Section 2.1. A block containing a leader is called a *leading block*. Otherwise it is called a *non-leading block*.

The *Set-Up* phase is described by identifying four types of configurations that form a partition of all possible configurations without towers. For each type we indicate the players and their destinations.

Type A. A configuration of *type A* is a configuration of interdistance $d \geq 1$ with at least one isolated robot. Consider an arbitrary configuration of *type A* and let S be the maximum among the sizes of the blocks that are neighbors of at least one isolated robot. Let \mathcal{I} be the set of isolated robots that are neighbors of a block of size S such that no other isolated robot is closer to a block of size S . The players in a configuration of *type A* are all the robots in \mathcal{I} . The destination of a player is its adjacent node in the direction of the closest neighboring block of size S . If there are two such blocks, the choice is arbitrary.

Type B. A configuration of *type B* is a configuration of interdistance $d \geq 1$, without isolated robots, and containing at least one non-leading block. More precisely, if all blocks have the same size then the configuration is of *type B1*. Otherwise, it is of *type B2*.

Consider an arbitrary configuration of *type B1*. If there is only one leader, then the player is the leader and its destination is its adjacent node outside the block it belongs to. From now on, we assume that there are two leaders. This implies that the configuration is symmetric. There are two cases. The first case is when the blocks are of size 2. Since $k \geq 17$, there are at least 9 blocks and hence there exist two symmetric blocks separated by at least three blocks on each side. (Observe that this property does not hold for $k = 16$.) The players in such a configuration are the robots of such two blocks, having the smallest view. The destination of a player is its adjacent node outside the block it belongs to. We consider now the second case, that is when the blocks are of size larger than 2. The players in such a configuration are the pair of symmetric robots that are the closest to each other among the robots at the border of a block and such that these two robots are not neighbors. Note that such a pair is unique, otherwise it would imply that both n and k are even, which is impossible since $\gcd(n, k) = 1$. The destination of a player is the adjacent node outside the block it belongs to.

Consider an arbitrary configuration of *type B2* and let s be the minimum size of a block in the configuration. Let S be the maximum among the sizes of the blocks that are neighbors of a block of size s and let d be the minimal distance between a block of size s and a block of size S . We define \mathcal{T} as the set of robots belonging to a block of size s , neighbors of a block of size S , and at distance d from it. The players in a configuration of *type B2* are the robots in \mathcal{T} with the largest view. The destination of a player is its adjacent node in the direction of its neighboring block of size S .

Type C. A configuration of *type C* is a configuration of interdistance $d \geq 2$, without isolated robots, and such that each of its blocks is a leading block. Note that this implies that either all robots are in the same block or the robots are divided into two blocks of the same size. Moreover, there are exactly two leaders because the configuration is symmetric. The players in a configuration of *type C* are the two leaders. The destination of a player is its adjacent node in the direction of the block it belongs to. (This is not ambiguous because leaders are always located at the border of a block.)

Type D. A configuration of *type D* is a configuration of interdistance $d = 1$, without isolated robots, and such that each of its blocks is a leading block. Type *D* is the set of configurations *no-towers-final*. When such a configuration is reached, the *Set-Up* phase ends and the *Tower-Creation* phase begins.

Note that types *A*, *B*, *C* and *D* form a partition of all possible initial configurations (when $\gcd(n, k) = 1$).

The general idea of the *Set-Up* phase is to create few compact blocks (interdistance 1). Each decrease of interdistance is accomplished by first decreasing the number of blocks. The following lemmas show how this progress is achieved. Theorem 3.1 shows that a *no-towers-final* configuration is always reached at the end.

Lemma 3.1 *Assume that at some time t the configuration is of type A and that the only engaged robots are isolated robots engaged to move toward a neighboring block. Then after finite time, the configuration is of type B , C or D , of the same interdistance as in time t , and no robots are engaged.*

Proof: Assume that at some time t the configuration is of type A and that the only engaged robots are isolated robots engaged to move toward a block. Starting at time t and until the first time when the configuration changes to a type different from A , it remains true that the only engaged robots are isolated robots that are neighbors of a block and are engaged to move toward this block. Thus no blocks are created and the interdistance remains the same. Moreover, to prove the lemma, it is sufficient to prove that there exists a time $t' > t$ where the number of isolated robots has decreased by at least one. Let t'' be the first time when all robots engaged at time t have moved. If at that time, the number of isolated robots has decreased, we are done. Thus we may assume that no isolated robots joined a block between time t and t'' . Starting at time t'' and until some robot joins a block, any isolated robot will always move toward the same block because the neighboring blocks will keep their size unchanged and will not move. Hence an isolated robot will keep moving toward the same neighboring block: the larger, or the closer if its neighboring blocks are of the same size. Consequently, at least one isolated robot will join a block after finite time, which concludes the proof of the lemma. \square

Lemma 3.2 *Assume that at some time t the configuration is asymmetric, of type $B1$, and that no robots are engaged. Then after finite time, the configuration is of type $B2$, of the same interdistance as in time t , no robots are engaged, and there is one block fewer than at time t .*

Proof: Assume that at some time t the configuration is asymmetric, of type $B1$, with blocks of size s , and that no robots are engaged. In such a configuration the unique player \mathcal{R} is the leader. Let b be the block of \mathcal{R} and let b' be the block neighbor of \mathcal{R} . If $s = 2$, then b is dismantled when \mathcal{R} moves. After finite time, the two robots of the former block b join some other blocks. The configuration is then of type $B2$ and no robots are engaged, which proves the lemma in this case. If $s > 2$, then b is of size $s - 1$ when \mathcal{R} moves. Therefore \mathcal{R} joins block b' , now of size $s + 1$. This

implies that the configuration is still asymmetric, but now of type $B2$. The new player is again in block b , neighbor of b' . After finite time, all robots in b have moved to b' , and no robots are engaged. \square

Lemma 3.3 *Assume that at some time t the configuration is symmetric, of type $B1$, with blocks of size 2, and that no robots are engaged. Then after finite time, the configuration is of type C or D , of the same interdistance as in time t , and no robots are engaged.*

Proof: Assume that at some time t the configuration is symmetric (two leaders), of type $B1$, with blocks of size 2, and that no robots are engaged. Let \mathcal{R}_1 and \mathcal{R}_2 be the two players, \mathcal{R}_1 being the first of the two to move. Let \mathcal{R}'_1 , resp. \mathcal{R}'_2 , be the other robot in \mathcal{R}_1 's block, resp. \mathcal{R}_2 's block. Finally let b_1 , resp. b_2 , b'_1 , b'_2 , be the neighboring block of \mathcal{R}_1 , resp. \mathcal{R}_2 , \mathcal{R}'_1 , \mathcal{R}'_2 . Note that b_1 and b_2 , resp. b'_1 and b'_2 , are different blocks since they are separated by at least one other block b_3 , resp. b'_3 .

Let \mathcal{R}_1 move, at time t_1 . Hence \mathcal{R}'_1 becomes isolated. Therefore the configuration changes to type A . Before the configuration changes again to a type different from A , the only other possible isolated robots are \mathcal{R}_2 and \mathcal{R}'_2 , which are separated from \mathcal{R}_1 and \mathcal{R}'_1 by several blocks and thus will not interfere with how these two robots will merge with the neighboring blocks. By definition of the players in a configuration of type $B1$, we know that \mathcal{R}_1 is now closer to b_1 than \mathcal{R}'_1 is to b'_1 . Thus by definition of the players in a configuration of type A , \mathcal{R}_1 goes toward block b_1 and \mathcal{R}'_1 will not be a player before \mathcal{R}_1 merges with block b_1 . When this happens, b_1 is larger than b'_1 and thus \mathcal{R}'_1 also merges with b_1 , forming a block of size 4.

If \mathcal{R}_2 is not engaged at time t_1 , then, just after \mathcal{R}_1 and \mathcal{R}'_1 have merged with b_1 , there is one block of size 4, many blocks of size 2, and no isolated or engaged robots. Thus, after finite time, all blocks of size 2 join the large block and no robots are engaged. This proves the lemma in this case. Otherwise, let t_2 be the time where \mathcal{R}_2 moves. Between time t_1 and t_2 , engaged robots other than \mathcal{R}_2 are necessarily neighbors of b_1 and they are engaged to move toward b_1 . Indeed, this is the case before \mathcal{R}'_1 joins b_1 and this is also the case after, because of the definition of the players in configurations of types $B2$ and A .

Consider the first moment $t_3 \geq t_2$ where the configuration is not of type A . At this time, b_1 is of size at least 4. Assume that there exists some other block than b_1 of size larger than 2. Note that this block b is formed by \mathcal{R}_2 and its neighbor other than \mathcal{R}'_2 (that is the robot of the former block b_2 closer to \mathcal{R}_2), and by one or two robots among \mathcal{R}'_2 and the other robot of the former block b_2 . Assume further that one robot of b is engaged. Then from the discussion of the previous paragraph, it is necessarily at the border of b , neighbor of b_1 and engaged to move toward it. This implies that either b_3 or b'_3 (or both) have merged with b_1 and thus b_1 is of size at least 6 (while b is of size at most 4). At this point, there is a unique largest block and all engaged robots are neighbors of it and engaged to move toward it. Thus after finite time all smaller blocks will join the largest one and the configuration changes to type C or D (depending on the current interdistance), with no robots engaged. This proves the lemma in this case. Consider now the remaining case, i.e. when there are at most two blocks of size larger than 2 and none of them contains engaged robots. Moreover

all engaged robots are neighbors of one of these blocks and are engaged to move toward it. Thus after finite time all blocks of size 2 join the block(s) of size larger than 2 and no robots are engaged. If the configuration is not yet of type C or D , then the smaller block is progressively dismantled (type $B2$) and all its robots join the largest block. This concludes the proof of the lemma. \square

Lemma 3.4 *Assume that at some time t the configuration is symmetric, of type $B1$, with blocks of size $s \geq 3$, and that no robots are engaged. Then after finite time, the configuration is of type $B2$, C or D , of the same interdistance as in time t , no robots are engaged, and there are fewer blocks than at time t .*

Proof: Assume that at some time t the configuration is symmetric (two leaders), of type $B1$, with blocks of size $s \geq 3$, and that no robots are engaged. Let us consider two cases.

Case 1. The players are in the same block b .

After the first player has moved, the block b is of size smaller than all the other blocks. Therefore no isolated robot will ever go toward it and no robot in a block different from b can be a player, and thus engaged, until the block b has been completely dismantled and all its robots have joined other blocks. Thus after finite time the configuration is of type $B2$, C or D , no robots are engaged, and there is one block less than at time t .

Case 2. The players \mathcal{R}_1 and \mathcal{R}_2 are in different blocks b_1 and b_2 .

By definition of the players in such a configuration, the blocks b_1 and b_2 are neighbors but \mathcal{R}_1 , resp. \mathcal{R}_2 , is not a neighbor of b_2 , resp. b_1 , but of another block b'_1 , resp. b'_2 . Note that it may be possible that $b'_1 = b'_2$. Let \mathcal{R}_1 be the first player to move, at time t_1 . If the second player \mathcal{R}_2 is not engaged at that time, then the situation is as in Case 1. Thus we may assume that \mathcal{R}_2 is engaged at time t_1 . Let t_2 be the time when \mathcal{R}_2 moves.

If block b_1 is not completely dismantled at time t_2 , then any engaged robot must be either isolated or in block b_1 , and in both cases engaged to move toward block b'_1 . Indeed in the period between time t_1 and time t_2 , the block b_1 is the only block of size smaller than s . Moreover, if \mathcal{R}_1 is still isolated then it is the only robot possibly engaged at time t_2 . Otherwise \mathcal{R}_1 joined block b'_1 , which becomes larger than any other block, and thus any player between time t_1 and time t_2 must be either isolated or in block b_1 , and in both cases engaged to move toward block b'_1 . Starting from time t_2 , the isolated robots and the robots from blocks b_1 and b_2 get engaged (if not already) and move toward blocks b'_1 and b'_2 . When there are no more isolated robots nor blocks b_1 and b_2 , the configuration is of type B , C or D , no robots are engaged, and there are two blocks less than at time t .

If block b_1 is completely dismantled at time t_2 , then consider the first time $t_3 \geq t_2$ when the configuration is of a type different from A . At time t_3 the block b'_1 has size at least $2s$ because all robots from b_1 joined b'_1 (they could not have joined b_2 , of size at most s). Moreover, no other blocks than b'_1 are of size larger than $s + 1$ and all engaged robots are neighbors of b'_1 engaged to move toward b'_1 . After finite time, all robots join block b'_1 . Thus the configuration is of type C or D , no robots are engaged, and there are fewer blocks than at time t . \square

Lemma 3.5 *Assume that at some time t the configuration is of type $B2$ and that no robots are engaged. Then after finite time, the configuration is of type B , C or D , of the same interdistance as in time t , no robots are engaged, and there are fewer blocks than at time t .*

Proof: Assume that at some time t the configuration is of type $B2$, with minimum block size s , and that no robots are engaged. Consider the situation after time t but before the first time $t' > t$ where there are no isolated robots and no blocks of size at most s , if such a time t' exists.

During this period, any player in a configuration of type $B2$ is a robot in some block of size at most s , neighbor of some block of size larger than s , and its destination is its adjacent node in the direction of this large block. Indeed, either the smallest size of a block is s and then there must exist a block of such a size that is a neighbor of a larger block, or the smallest size is smaller than s and then any smallest block is a neighbor of a block larger than s since it shrank by having at least one of its robots move to such a large block.

Similarly, between time t and t' , any player \mathcal{R} in a configuration of type A is a robot neighbor of some block of size larger than s , and its destination is its adjacent node in the direction of this large block. Indeed, if this robot \mathcal{R} got isolated by moving outside a block then it is a neighbor of a block of size larger than s , by the previous argument. Otherwise it got isolated after another robot \mathcal{R}' left the block of size 2 to which \mathcal{R} belonged. In this case, \mathcal{R}' was a neighbor of a block of size larger than s and \mathcal{R} cannot be a player before \mathcal{R}' joins the large block and thus before \mathcal{R} becomes a neighbor of this large block, except if \mathcal{R} is itself already a neighbor of a block of size larger than s .

Since the configuration is necessarily of type A or $B2$ before time t' , we have that, while there is at least one isolated robot or one block of size s , the only engaged robots are robots not belonging to a block larger than s but neighbors of one of them and engaged to move toward it. Since there is at least one player in each configuration of type A or $B2$, this implies that the above defined time t' exists and that no robots are engaged at that time. Moreover no blocks have been created during the process and some (the ones of size at most s) have been dismantled. This concludes the proof of the lemma. \square

Lemma 3.6 *Assume that at some time t the configuration is of type C , of interdistance $d \geq 2$, and that no robots are engaged. Then at some time $t' > t$ one of the two following situations occurs:*

- *The configuration is of type A , of interdistance $d - 1$, and the only engaged robots are isolated robots engaged to move toward a block.*
- *The configuration is of type B , of interdistance $d - 1$, and no robots are engaged.*

Proof: Assume that at some time t the configuration is of type C , of interdistance $d \geq 2$, and no robots are engaged. The players are the two leaders. After finite time, at least one will move. Consider the moment t_1 where the first moves. At this moment the configuration changes to type A .

If the other player moved at the same time or is not engaged, we are done because the configuration is of type A , of interdistance $d - 1$, and no robots are engaged.

Thus we assume that the other player \mathcal{R} is engaged at time t_1 . By the definition of type C configurations and the fact that $k \geq 17$, it is engaged to move toward an isolated robot \mathcal{R}' that is at distance exactly d . Note that until \mathcal{R} moves, there is only one block (of interdistance $d - 1$) and \mathcal{R} is a neighbor of it (its other neighbor is \mathcal{R}'). Moreover, \mathcal{R} is isolated and no robots will move toward it to make a block because no other robot is engaged at the moment and because a player in a configuration of type A never moves toward an isolated robot. Therefore, the configuration will remain of type A while \mathcal{R} does not move.

Consider now the time t_2 where \mathcal{R} makes its move. If this move does not make it belong to a block, then the configuration is of type A , of interdistance $d - 1$, and the only engaged robots are isolated robots engaged to move toward a block. Assume now that the move of \mathcal{R} makes it belong to a block. Then necessarily it is a new block, of size two, and formed with robot \mathcal{R}' . If \mathcal{R}' is not engaged then we are at a time t' satisfying the lemma. Indeed if there are isolated robots then the first situation occurs, and if there are not then the other block is larger, of size $k - 2$, and thus the second situation occurs. If \mathcal{R}' is engaged at time t_2 , then there are no isolated robots because there is none between \mathcal{R} and the other block (in the segment excluding \mathcal{R}') and there is none between \mathcal{R}' and the other block (in the segment excluding \mathcal{R}) since \mathcal{R}' got engaged as an isolated robot and thus was engaged to move toward a block. Therefore, we are in the following situation: there are two blocks of sizes $k - 2$ and 2 , at distances at least $d + 1$ (on both sides); exactly one robot of the smaller block is engaged to move toward the other block, and no other robots are engaged. Thus the configuration is of type $B2$ and after some finite time, one of the two robots of the smaller block will move. At this moment, the first situation occurs, which concludes the proof of the lemma. \square

Theorem 3.1 *Any initial configuration is transformed after finite time into a configuration of type D (i.e. no-towers-final) without engaged robots.*

Proof: Let Φ be the property that the only engaged robots (if any) in a given configuration are isolated ones and they are engaged to move toward a neighboring block. For any configuration c of type A , B or C define the triple $T(c) = (d, t, x)$, where d is the interdistance of c , t is the type of c , i.e., t is A , B or C , and x is the number of blocks in c . Order all triples lexicographically, assuming that $C < B < A$. Lemmas 3.1 – 3.6 imply that any configuration c of type A , B or C satisfying property Φ is transformed after finite time either in a configuration c' of type A , B or C satisfying property Φ , such that $T(c') < T(c)$, or in a configuration of type D with no robots engaged. Since any initial configuration satisfies property Φ , this concludes the proof.

Figure 1 illustrates the progress of configurations toward type D .

\square




Figure 1: Progress toward type D. Dashed arrows correspond to transitions where the interdistance decreases. The loop (in bold) corresponds to a transition where the number of blocks decreases.

3.3 Tower-Creation Phase

In the following, as a slight abuse of notation, we will use “(a segment of) a consecutive robots” with the meaning “a segment of a consecutive nodes, each occupied by a robot of multiplicity one, i.e., a single robot”. Similarly, “a tower” may stand for “a node occupied by a tower”.

The second phase of the protocol is *Tower-Creation*. This phase begins with a configuration of type D , i.e., one of the configurations *no-towers-final* which consists of either a single block of consecutive robots or two such blocks of the same size. The goal of this phase is to create one or two towers in each block (depending on the parity of the number of robots per block). More precisely, in a block of odd size there will be one tower, and in a block of even size there will be two towers. In a block of odd size the tower is formed by the central robot (the player) moving to the neighbor (the destination) containing the robot with the larger view (in the case when both have the same view, an arbitrary adjacent node is chosen by the adversary). In a block of even size the two towers are formed by the two central robots (the players) moving to their other neighbors (their destination). The obtained configuration is called *towers-completed*. More precisely:

- *towers-completed*: Four types of such configurations exist:
 - (1) a segment of a consecutive robots followed by an empty node, followed by a tower, followed by a segment of $a - 1$ consecutive robots, with $2a + 1 = k$;
 - (2) a segment of $a - 2$ consecutive robots followed by a tower, followed by two empty nodes followed by a tower, followed by a segment of $a - 2$ consecutive robots, with $2a = k$;
 - (3) two occurrences of segments of type (2); that is: a segment of $a - 2$ consecutive robots followed by a tower, followed by two empty nodes followed by a tower, followed by a segment of $a - 2$ consecutive robots, followed by b_1 empty nodes, again followed by $a - 2$ consecutive robots, a tower, two empty nodes, a tower, $a - 2$ consecutive robots, and b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (4) two occurrences of segments of type (1); that is: a segment of a consecutive robots followed by an empty node, followed by a tower, followed by a segment of $a - 1$ consecutive robots, followed by b_1 empty nodes, again followed by a consecutive robots, a tower, an empty node, $a - 1$ consecutive robots, and b_2 empty nodes, with $4a + 2 = k$ and $k + b_1 + b_2 = n$.

Configurations *towers-completed* are easily recognizable as each block of a *no-towers-final* configuration is transformed as follows (see Figure 2). A single block or two blocks of odd size $2a + 1$, with $a \geq 2$, are transformed into *towers-completed* of type (1). A single block or two blocks of even size $2a$ each, with $a \geq 3$, are transformed into *towers-completed* of type (2).

All intermediate configurations (e.g., an even block with one tower only, or a pair of even blocks with a tower formed in one) are called *uncompleted-towers*. More precisely:

- *uncompleted-towers* (to be read either clockwise or counterclockwise):
 - (1) a segment of $a - 2$ consecutive robots followed by a tower, followed by one empty node, followed by a segment of $a - 1$ consecutive robots, with $2a = k$;
 - (2) a segment of $2a$ consecutive robots followed by b_1 empty nodes, followed by a segment of $a - 2$ consecutive robots, a tower, two empty nodes, a tower, $a - 2$ consecutive robots, and finally b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (3) a segment of $a - 2$ consecutive robots followed by a tower, followed by one empty node, followed by a segment of $a - 1$ consecutive robots, followed by b_1 empty nodes, followed by a segment of $a - 2$ consecutive robots, a tower, two empty nodes, a tower, $a - 2$ consecutive robots, and finally b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (4) a segment of $a - 1$ consecutive robots followed by one empty node, a tower, a segment of $a - 2$ consecutive robots, b_1 empty nodes, $a - 1$ consecutive robots, one empty node, a tower, a segment of $a - 2$ consecutive robots, b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (5) a segment of $a - 1$ consecutive robots followed by one empty node, a tower, a segment of $a - 2$ consecutive robots, b_1 empty nodes, $a - 2$ consecutive robots, a tower, one empty node, $a - 1$ consecutive robots, b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (6) a segment of $2a$ consecutive robots followed by b_1 empty nodes, followed by a segment of $a - 2$ consecutive robots, followed by a tower, one empty node, a segment of $a - 1$ consecutive robots, and b_2 empty nodes, with $4a = k$ and $k + b_1 + b_2 = n$;
 - (7) a segment of $2a + 1$ consecutive robots followed by b_1 empty nodes, followed by a segment of $a - 1$ consecutive robots, followed by a tower, one empty node, a consecutive robots, and b_2 empty nodes, with $4a + 2 = k$ and $k + b_1 + b_2 = n$.

Notice that the players and their destination are unambiguously determined both in *no-towers-final* and in *uncompleted-towers* configurations.




Figure 2: Transformed blocks (a) of odd size (b) of even size

Theorem 3.2 *Let $k \geq 17$. The Tower Creation phase transforms in finite time the system from a no-towers-final configuration with no engaged robots to a towers-completed configuration with no engaged robots.*

Proof: Let the configuration become *no-towers-final* at time t and let no robot be engaged at that time. First notice that, by definition, an even block of a *towers-completed* configuration requires at least 6 robots (odd blocks require 5), and such a size is always guaranteed by the condition $k \geq 17$. By definition of players, the players and their destination are uniquely determined. Since no robot is engaged at time t , only the players will be allowed to move. Let $Play(t)$ denote the set of players at time $t' \geq t$, and let $Still(t')$ denote the set of those that at time $t' \geq t$ have not reached their destination yet. Let $t_1 > t$ be the first time $|Still(t_1)| < |Play(t)|$. First observe that if $|Still(t_1)| = 0$ (i.e., all players have reached their destination) the configuration is *towers-completed*, otherwise it is *uncompleted-towers* and the robots that moved at this time have reached their destination. Further notice that, if the configuration is *uncompleted-towers*, $Play(t_1) = Still(t_1)$; i.e., the set of players coincides with $Still(t_1)$ and their destination is unchanged. As a consequence, all engaged robots at time t_1 are a subset of $Still(t_1)$; in other words, only robots in $Still(t_1)$ will move and their destination is the same as at time t . Hence, at some time $t'' > t$ all robots in $Play(t)$ will be in their final position forming a *towers-completed* configuration, with no engaged robots. \square

3.4 Exploration Phase

Exploration starts when towers in the preceding phase are created. Note that the empty nodes adjacent to towers have already been explored, so the segments of empty nodes between the transformed blocks are the only ones possibly not yet explored. Each of these segments is explored in the current phase using one or two robots closest to the segment. If k is even, such a segment must lie between two segments of consecutive robots of equal size, and it is explored by the two border robots that meet in the middle of the segment (either at the extremities of the central edge, or in the central node). The obtained configuration is called *terminal*. More precisely:

- *terminal* (to be read either clockwise or counterclockwise): There are several types of such configurations:
 - (EC_1) $a + 1$ consecutive robots followed by an empty node, followed by a tower, followed by a segment of $a - 2$ consecutive robots, with $2a + 1 = k$;
 - (EC_2) $a - 3$ consecutive robots, followed by a tower, two empty nodes, a tower, $a - 3$ consecutive robots, x empty nodes, a tower, x empty nodes, with $2a = k$ and $k + 2x - 1 = n$;
 - (EC_3) $a - 3$ consecutive robots, followed by a tower, two empty nodes, a tower, $a - 3$ consecutive robots, x empty nodes, a tower, x empty nodes, $a - 3$ consecutive robots, a tower, two empty nodes, a tower, $a - 3$ consecutive robots, y empty nodes, two consecutive robots, y empty nodes, with $4a = k$ and $k + 2x + 2y - 1 = n$;
 - (EC_4) $a - 1$ consecutive robots, followed by one empty node, a tower, $a - 2$ consecutive robots, x empty nodes, a tower, x empty nodes, $a - 2$ consecutive robots, a tower, one empty node, $a - 1$ consecutive robots, y empty nodes, two consecutive robots, y empty nodes, with $4a + 2 = k$ and $k + 2x + 2y - 1 = n$;
 - (EC_5) $a - 2$ consecutive robots, followed by a tower, one empty node, $a - 1$ consecutive robots, x empty nodes, a tower, x empty nodes, $a - 1$ consecutive robots, one empty node, a

tower, $a - 2$ consecutive robots, y empty nodes, two consecutive robots, y empty nodes, with $4a + 2 = k$ and $k + 2x + 2y - 1 = n$.

In the case of odd k , the configuration starting the *Exploration* phase is a single block of odd size with one tower, with $a = (k - 1)/2$ (hence in particular $a \geq 3$), see Figure 2 (a). The unique player is the robot in the segment of $a - 1$ consecutive robots, farthest from the tower. The destination of this robot is its unique empty neighbor node. In a resulting configuration with a single isolated robot, the player is this robot and its destination is the neighbor node toward the segment of a consecutive robots. When the configuration contains $a + 1$ consecutive robots followed by an empty node, followed by a tower, followed by a segment of $a - 2$ consecutive robots, all robots remain idle. At this point we reach *terminal* configuration (EC_1) and the exploration is completed.

In the case of even k , the configuration starting the *Exploration* phase is composed by either a single block or two blocks with towers, see Figure 2. In this case two players are identified at the border of each transformed block and their destination is their neighbor node towards the center of the empty segment towards the neighboring block (or towards the other side of the same block). The players keep being isolated robots until they either become neighbors in the middle of the segment (configurations (EC_3), (EC_4) and (EC_5) in the case of two blocks) or they form another tower (configuration (EC_2) in the case of a single block, and configurations (EC_3), (EC_4) and (EC_5) in the case of two blocks). In all these cases they form a *terminal* configuration and the exploration is completed.

Theorem 3.3 *The Exploration phase transforms the system from a towers-completed configuration with no engaged robots to a terminal configuration with no engaged robots.*

Proof: Let the configuration become *completed-towers* at time t and let no robot be engaged at that time. By definition, the players and their destination are uniquely determined. Since no robot is engaged at time t , only the players will be allowed to move. Let $Play(t')$ denote the set of players at time $t' \geq t$, and let $Still(t')$ denote the set of those that at time $t' \geq t$ have not reached their destination yet. Let $t'_1 > t$ be the first time $|Still(t'_1)| < |Play(t)|$. First observe that if $|Still(t_1)| = 0$ (i.e., all players have reached their destination) the configuration is *terminal*, otherwise the robots that moved at this time are closer to their destination and the players are still unambiguously determined. Further notice that, if the exploration is not terminated, $Play(t_1) = Still(t_1)$; i.e., the set of players coincides with $Still(t_1)$ and their destination is unchanged. As a consequence, all engaged robots at time t_1 are a subset of $Still(t_1)$; in other words, only robots in $Still(t_1)$ will move and their destination is the same as at time t . Hence, at some time $t'' > t$ all robots in $Play(t)$ will be in their final position forming a *terminal* configuration, with no engaged robots. \square

From Theorem 3.1 describing the conclusion of the *Set-Up* phase, from Theorem 3.2 guaranteeing the termination of the *Tower-Creation* phase and from Theorem 3.3, we get the following result.

Theorem 3.4 *Let $17 \leq k < n$. Algorithm RING EXPLORATION allows a team of k robots to explore an n -node ring and enter a terminal state within finite time, provided $\gcd(n, k) = 1$.*

4 On the size of the minimum team

In this section we show that the minimum number of robots that can explore an n -node ring regardless of their initial position, is logarithmic in n . More precisely, we have the following result.

Theorem 4.1 *The minimum number $\rho(n)$ of robots that can explore an n -node ring has the following properties:*

1. $\rho(n) \in O(\log n)$;
2. there exists a constant c such that, for infinitely many n , we have $\rho(n) \geq c \log n$.

Proof: Let p_j denote the j -th prime, and let $p_j\#$ denote the p_j -primorial, that is

$$p_j\# = \prod_{i=1}^j p_i \tag{1}$$

An important property of the primorial is the following [40]:

$$\lim_{j \rightarrow \infty} (p_j\#)^{\frac{1}{p_j}} = e. \tag{2}$$

We will now prove each part of the theorem separately.

Part 1.

Let $f(n)$ be the smallest integer coprime with n and larger than 16. Thus, by Theorem 3.4, exploration is possible with $f(n)$ agents. Hence, $\rho(n) \leq f(n)$.

Take j such that $\frac{p_j\#}{13\#} \leq n < \frac{p_{j+1}\#}{13\#}$. We have $f(n) \leq p_{j+1}$. (Otherwise, all primes in $\{17, \dots, p_{j+1}\}$ divide n and hence $n \geq \frac{p_{j+1}\#}{13\#}$, contradiction.) By property (2) we have $2 \leq (p_j\#)^{\frac{1}{p_j}}$, for sufficiently large j . Hence $2^{p_j} \leq p_j\#$, and thus $p_j \leq \log(p_j\#)$. Hence $p_{j+1} \leq \log(p_{j+1}\#) = \log(p_j\#) + \log p_{j+1}$. Since $p_{j+1} \leq p_j\# + 1 \leq 2 \cdot 13\# \cdot n$, we have $\rho(n) \leq f(n) \leq p_{j+1} \leq \log n + \log(2 \cdot 13\# \cdot n)$, which is at most $3 \log n$, for sufficiently large n .

Part 2.

Let n be the least common multiple of integers $1, 2, \dots, m$. Let $g(n)$ be the smallest integer not dividing n . By Lemma 2.1 we have $\rho(n) \geq g(n)$. We have $g(n) \geq m + 1$. The Prime Number Theorem implies $\frac{\ln n}{m} \rightarrow 1$. Hence $\ln n \leq 2m$, for sufficiently large m . This implies the existence of a constant c such that $\rho(n) \geq g(n) \geq m + 1 > m \geq \frac{\ln n}{2} \geq c \log n$. \square

It should be noted that for some specific values of n , the number $\rho(n)$ is constant. For example, if $n > 17$ is prime, then Theorem 3.4 shows that 17 robots can explore the n -node ring, hence $\rho(n) \leq 17$.

5 Conclusions

In this paper we have analyzed the exploration problem in rings by asynchronous oblivious robots that can see the environment but cannot communicate. This is a further step in the understanding of how these robots' capabilities, standard in continuous universes, can be exploited in the discrete ones. Our results open several interesting problems.

First, the complete characterization of couples (n, k) for which exploration of the ring is solvable remains open. In particular, what happens if the number k of robots is less than 17? A partial answer has been given in [35], where it has been shown that no deterministic exploration is feasible with fewer than five robots, whenever the size of the ring is even, and that five robots are sufficient for any n that is co-prime with five.

Next, the problem of exploring other topologies and arbitrary graphs is a natural extension of this work. Moreover, since the robots cannot communicate, they have to be able to observe the environment; an immediate question is what happens if the robots can only see within a fixed distance. Accuracy of vision as well as fault-tolerance are issues that should be also addressed by future research.

Acknowledgment. This work was done during the stay of David Ilcinkas at the Research Chair in Distributed Computing at the Université du Québec en Outaouais and at the University of Ottawa, as a postdoctoral fellow. Andrzej Pelc was partially supported by the Research Chair in Distributed Computing at the Université du Québec en Outaouais, Paola Flocchini was partially supported by the University Research Chair of the University of Ottawa. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada under Discovery grants.

References

- [1] N. Agmon, D. Peleg, Fault-tolerant gathering algorithms for autonomous mobile robots, *SIAM Journal on Computing* 36 (2006), 56–82.
- [2] S. Albers and M. R. Henzinger, Exploring unknown environments, *SIAM Journal on Computing* 29 (2000), 1164–1188.
- [3] C. Ambühl, L. Gasieniec, A. Pelc, T. Radzik, X. Zhang, Tree exploration with logarithmic memory, *ACM Transactions on Algorithms* 7 (2011), article 17.
- [4] I. Averbakh and O. Berman, A heuristic with worst-case analysis for minimax routing of two traveling salesmen on a tree, *Discrete Applied Mathematics* 68 (1996), 17–32.
- [5] I. Averbakh and O. Berman, $(p - 1)/(p + 1)$ -approximate algorithms for p -traveling salesmen problems on a tree with minmax objective, *Discrete Applied Mathematics* 75 (1997), 201–216.

- [6] B. Awerbuch, M. Betke, R. Rivest and M. Singh, Piecemeal graph learning by a mobile robot, Proceedings of the 8th Annual Conference on Computational Learning Theory (COLT 1995), 321–328.
- [7] M.A. Bender, A. Fernandez, D. Ron, A. Sahai and S. Vadhan, The power of a pebble: exploring and mapping directed graphs, Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC 1998), 269–278.
- [8] M.A. Bender and D. Slonim, The power of team exploration: Two robots can learn unlabeled directed graphs, Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS 1994), 75–85.
- [9] M. Betke, R. Rivest and M. Singh, Piecemeal learning of an unknown environment, Machine Learning 18 (1995), 231–254.
- [10] J. Chalopin, P. Flocchini, B. Mans, and N. Santoro, Network exploration by silent and oblivious robots, Proceedings of the 36th Int. Workshop on Graph Theoretic Concepts in Computer Science (WG 2010), LNCS 6410, 208–219.
- [11] M. Cieliebak, P. Flocchini, G. Prencipe, N. Santoro, Solving the robots gathering problem, Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP 2003), LNCS 2719, 1181–1196.
- [12] R. Cohen, D. Peleg, Convergence properties of the gravitational algorithm in asynchronous robot systems. SIAM Journal on Computing 34 (2005), 1516–1528.
- [13] R. Cohen, D. Peleg, Robot convergence via center-of-gravity algorithms, Proceedings of the 11th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2004), LNCS 3104, 79–88.
- [14] C. Cooper, R. Klasing, and T. Radzik, Searching for black-hole faults in a network using multiple agents. Proceedings of the 10th International Conference on Principles of Distributed Systems (OPODIS 2006), LNCS 4288, 320–332.
- [15] J. Czyzowicz, L. Gasieniec, A. Pelc, Gathering few fat mobile robots in the plane, Theoretical Computer Science 410 (2009), 481–499.
- [16] J. Czyzowicz, D. Kowalski, E. Markou, and A. Pelc, Searching for a black hole in synchronous tree networks. Combinatorics, Probability & Computing 16 (2007), 595–619.
- [17] S. Das, P. Flocchini, S. Kutten, A. Nayak, N. Santoro, Map construction of unknown graphs by multiple agents, Theoretical Computer Science 385 (2007), 34–48.
- [18] X. Défago and S. Souissi. Non-uniform circle formation algorithm for oblivious mobile robots with convergence toward uniformity. Theoretical Computer Science 396 (2008), 97–112.
- [19] X. Deng and C. H. Papadimitriou, Exploring an unknown graph, Journal of Graph Theory 32 (1999), 265–297.

- [20] A. Dessmark and A. Pelc, Optimal graph exploration without good maps, *Theoretical Computer Science* 326 (2004), 343–362.
- [21] S. Devismes, Optimal exploration of small rings, *Proceedings of the 3rd International ACM SIGOPS/SIGACT Workshop on Reliability, Availability, and Security (WRAS 2010)*, pages 9:1–9:6.
- [22] S. Devismes, F. Petit, and S. Tixeuil, Optimal probabilistic ring exploration by semi-synchronous oblivious robots, *Proceedings of the 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2009)*, LNCS 5869, 195–208.
- [23] K. Diks, P. Fraigniaud, E. Kranakis and A. Pelc, Tree exploration with little memory, *Journal of Algorithms* 51 (2004), 38–63.
- [24] C.A. Duncan, S.G. Kobourov and V.S.A. Kumar, Optimal constrained graph exploration, *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2001)*, 807–814.
- [25] R. Fleischer, G. Trippen, Exploring an unknown graph efficiently, *Proceedings of the 13th European Symposium on Algorithms (ESA 2005)*, LNCS 3669, 11–22.
- [26] P. Flocchini, D. Ilcinkas, A. Pelc, N. Santoro, Remembering without memory: tree exploration by asynchronous oblivious robots, *Theoretical Computer Science* 411 (2010), 1544–1557.
- [27] P. Flocchini, D. Ilcinkas, A. Pelc, N. Santoro, Computing without communicating: ring exploration by asynchronous oblivious robots, *Proceedings of the 11th International Conference on Principles of Distributed Systems (OPODIS 2007)*, LNCS 4878, 105–118.
- [28] P. Flocchini, D. Ilcinkas, N. Santoro, Ping-Pong in dangerous graphs: Optimal black hole search with pure tokens, *Algorithmica* (2011), in press.
- [29] P. Flocchini, G. Prencipe, N. Santoro, Computing by Mobile Robotic Sensors, Chapter 3 of [37], 2011.
- [30] P. Flocchini, G. Prencipe, N. Santoro, P. Widmayer, Arbitrary pattern formation by asynchronous, anonymous, oblivious robots, *Theoretical Computer Science* 407 (2008), 412–447.
- [31] P. Fraigniaud, L. Gasieniec, D. Kowalski, A. Pelc, Collective tree exploration, *Networks* 48 (2006), 166–177.
- [32] G. N. Frederickson, M. S. Hecht and C. E. Kim, Approximation algorithms for some routing problems. *SIAM Journal on Computing* 7 (1978), 178–193.
- [33] R. Klasing, A. Kosowski, A. Navarra, Taking advantage of symmetries: Gathering of many asynchronous oblivious robots on a ring. *Theoretical Computer Science* 411 (2010), 3235–3246.
- [34] R. Klasing, E. Markou, A. Pelc, Gathering asynchronous oblivious mobile robots in a ring, *Theoretical Computer Science* 390 (2008), 27–39.

- [35] A. Lamani, M. Gradinariu Potop-Butucaru, and S. Tixeuil, Optimal deterministic ring exploration with oblivious asynchronous robots, Proceedings of the 17th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2010), LNCS 6058, 183–196.
- [36] S. Kamei, A. Lamani, F. Ooshita, and S. Tixeuil, Asynchronous mobile robot gathering from symmetric configurations without global multiplicity detection, Proceedings of the 18th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2011), LNCS 6796, 150–161.
- [37] S. Nikolettseas and J. Rolim (Editors), *Theoretical Aspects of Distributed Computing in Sensor Networks*, Springer, 2011.
- [38] P. Panaite and A. Pelc, Exploring unknown undirected graphs, *Journal of Algorithms* 33 (1999), 281–295.
- [39] G. Prencipe, Impossibility of gathering by a set of autonomous mobile robots. *Theoretical Computer Science* 384 (2007), 222–231.
- [40] S.M. Ruiz, A result on prime numbers, *The Mathematical Gazette* 81 (1997), 269.
- [41] S. Souissi, X. Défago and M. Yamashita, Gathering asynchronous mobile robots with inaccurate compasses. Proceedings of the 10th International Conference on Principles of Distributed Systems (OPODIS 2006), LNCS 4305, 333–349.
- [42] I. Suzuki, M. Yamashita, Distributed anonymous mobile robots: formation of geometric patterns. *SIAM Journal on Computing* 28 (1999), 1347–1363.