arXiv:2008.00701v2 [cs.DC] 4 Dec 2020

Memory Optimal Dispersion by Anonymous Mobile Robots

Archak Das
Department of Mathematics, Jadavpur University, India
archakdas.math.rs@jadavpuruniversity.in

Kaustav Bose
Department of Mathematics, Jadavpur University, India
kaustavbose.rs@jadavpuruniversity.in

Buddhadeb Sau
Department of Mathematics, Jadavpur University, India
buddhadeb.sau@jadavpuruniversity.in

Abstract

Consider a team of $k \leq n$ autonomous mobile robots initially placed at a node of an arbitrary graph G with n nodes. The *dispersion* problem asks for a distributed algorithm that allows the robots to reach a configuration in which each robot is at a distinct node of the graph. If the robots are anonymous, i.e., they do not have any unique identifiers, then the problem is not solvable by any deterministic algorithm. However, the problem can be solved even by anonymous robots if each robot is given access to a fair coin which they can use to generate random bits. In this setting, it is known that the robots require $\Omega(\log \Delta)$ bits of memory to achieve dispersion, where Δ is the maximum degree of G. On the other hand, the best known memory upper bound is $min\{\Delta, max\{\log \Delta, \log D\}\}$ ($D =$ diameter of G), which can be $\omega(\log \Delta)$, depending on the values of Δ and D. In this paper, we close this gap by presenting an optimal algorithm requiring $O(\log \Delta)$ bits of memory.

1 Introduction

1.1 Background and Motivation

A considerable amount of research has been devoted in recent years to the study of distributed algorithms for autonomous multi-robot system. A multi-robot system consists of a set of autonomous mobile computational entities, called $robots$, that coordinate with each other to achieve some well defined goals, such as forming a given pattern, exploration of unknown environments etc. The robots may be operating on continuous space or graph-like environments. The most fundamental tasks in graphs are Gathering [3, 6, 7, 10, 13–17, 22–24, 28, 29] and Exploration [4, 5, 8, 9, 11, 12, 27]. A relatively new problem which has attracted a lot of interest recently is Dispersion, introduced by Augustine and Moses Jr. [2]. The problem asks $k \leq n$ robots, initially placed arbitrarily at the nodes of an n-node anonymous graph, to reposition themselves to reach a configuration in which each robot is at a distinct node of the graph. The problem has many practical applications, for example, in relocating self-driven electric cars to recharge stations where finding new recharge stations is preferable to multiple cars queuing at the same station to recharge. The problem is also interesting because of its relationship to other well-studied problems such as Exploration, Scattering and Load Balancing [2].

It is easy to see that the problem cannot be solved deterministically by a set of anonymous robots. Since all robots execute the same deterministic algorithm and initially they are in the same state, the

co-located robots will perform the same moves. This is true for each round and hence they will always mirror each other's move and will never do anything different. Hence, throughout the execution of the algorithm, they will stick together and as a result, dispersion cannot be achieved. Using similar arguments, it can be shown that the robots need to have $\Omega(\log k)$ bits of memory each in order to solve the problem by any deterministic algorithm [2]. However, it has been recently shown in [25] that if we consider randomized algorithms, i.e., each robot is given access to a fair coin which can be used to generate random bits, then DISPERSION can be solved by anonymous robots with possibly $o(\log k)$ bits of memory. In [25], two algorithms are presented for DISPERSION from a *rooted configuration*, i.e., a configuration in which all robots are situated at the same node. The first algorithm requires each robot to have $O(max\{\log \Delta, \log D\})$ bits of memory, where Δ and D are respectively the maximum degree and diameter of G. The second algorithm requires each robot to have $O(\Delta)$ bits of memory. In [25], it is also shown that the robots require $\Omega(\log \Delta)$ bits of memory to achieve dispersion in this setting. Notice that while the memory requirement of the second algorithm is clearly $\omega(\log \Delta)$, that of the first algorithm too can be $\omega(\log \Delta)$ depending on the values of Δ and D. In this paper, we close this gap by presenting an asymptotically optimal algorithm that requires $O(\log \Delta)$ bits of memory.

1.2 Related Works

DISPERSION was introduced in [2] where the problem was considered in specific graph structures such as paths, rings, trees as well as arbitrary graphs. In [2], the authors assumed $k = n$, i.e., the number of robots k is equal to the number of nodes n. They proved a memory lower bound of $\Omega(\log k)$ bits at each robot and a time lower bound of $\Omega(\log D)$ rounds for any deterministic algorithm to solve the problem in a graph of diameter D. They then provided deterministic algorithms using $O(\log n)$ bits of memory at each robot to solve DISPERSION on lines, rings and trees in $O(n)$ time. For rooted trees they provided an algorithm requiring $O(\Delta + \log n)$ bits of memory and $O(D^2)$ rounds and for arbitrary graphs, they provided an algorithm, requiring $O(n \log n)$ bits of memory and $O(m)$ rounds (m is the number of edges in the graph). In [18], a $\Omega(k)$ time lower bound was proved for $k \leq n$. In addition, three deterministic algorithms were provided in [18] for arbitrary graphs. The first algorithm requires $O(k\log\Delta)$ bits of memory and $O(m)$ time, ($\Delta =$ the maximum degree of the graph), the second algorithm requires $O(D \log \Delta)$ bits of memory and $O(\Delta^D)$ time, and the third algorithm requires $O(\log(\max(k, \Delta)))$ bits of memory and $O(mk)$ time. Recently, a deterministic algorithm was provided in [19] that runs in $O(\min(m, k\Delta) \log k)$ time and uses $O(\log n)$ bits of memory at each robot. In [21], the problem was studied on grid graphs. The authors presented two deterministic algorithms on anonymous grid graphs that achieve simultaneously optimal bounds with respect to both time and memory complexity. For the first algorithm, the authors considered the local communication model where a robot can only communicate with other robots that are present at the same node. Their second algorithm works in global communication model where a robot can communicate with other robots present anywhere on the graph. In the local communication model, they showed that the problem can be solved in an n-node square grid graph in $O(\min(k, \sqrt{n}))$ time with $O(\log k)$ bits of memory at each robot. In the global communication model, the authors showed that it can be solved in $O(\sqrt{k})$ time with $O(\log k)$ bits of memory at each robot. In [20], the authors extended the work in global communication model to arbitrary graphs. They gave three deterministic algorithms, two for arbitrary graphs and one for trees. For arbitrary graphs, their first algorithm is based on DFS traversal and has time complexity of $O(\min(m, k\Delta))$ and memory complexity of $\Theta(\log(\max(k, \Delta)))$. The second algorithm is based on BFS traversal and has time complexity $O(\max(D, k)\Delta(D + \Delta))$ and memory complexity $O(\max(D, \Delta \log k))$. The third algorithm in arbitrary trees is a BFS based algorithm that has time and memory complexity $O(D \max(D, k))$ and $O(\max(D, \Delta \log k))$ respectively. In [1], the problem was studied on dynamic rings. Fault-tolerant DISPERSION was considered for the first time in [26] where the authors studied the problem on a ring in presence of Byzantine robots. In [25], randomization was used to break the $\Omega(\log k)$ memory lower bound for deterministic algorithms. In particular, the authors considered anonymous robots that can generate random bits and gave two deterministic algorithms that achieve dispersion from rooted configurations on an arbitrary graph. The

memory complexity of the algorithms are respectively $O(max\{\log \Delta, \log D\})$ and $O(\Delta)$. For arbitrary initial configurations, they gave a random walk based algorithm that requires $O(\log \Delta)$ bits of memory, but the robots do not terminate.

1.3 Our Results

We study Dispersion from a rooted configuration on arbitrary graphs by a set of anonymous robots with random bits. In [25], two algorithms with memory complexity $O(max\{\log \Delta, \log D\})$ and $O(\Delta)$ were reported. The question of whether the problem can be solved with $O(\log \Delta)$ bits of memory at each robot was left as an open problem. In this paper, we answer this question affirmatively by presenting an algorithm with memory complexity $O(\log \Delta)$. The lower bound result presented in [25] implies that the algorithm is asymptotically optimal with respect to memory complexity.

1.4 Organization of the Paper

In Section 2, we describe the model and introduce notations that will be used in the paper. In Section 3, we describe the main algorithm. In Section 4, we prove the correctness of our algorithm and establish the time and memory complexity.

2 Technical Preliminaries

Graph. We consider a connected undirected graph G of n nodes, m edges, diameter D and maximum degree Δ. For any node v, its degree is denoted by $\delta(v)$ or simply δ when there is no ambiguity. The nodes are anonymous, i.e., they do not have any labels. For every edge connected to a node, the node has a corresponding port number for the edge. For every node, the edges incident to the node are uniquely identified by port numbers in the range $[0, \delta - 1]$. There is no relation between the two port numbers of an edge. If u, v are two adjacent nodes then port(u, v) denotes the port at u that corresponds to the edge between u and v.

Robots. Robots are anonymous, i.e., they do not have unique identifiers. Each robot has $O(\log \Delta)$ bits of space or memory for computation and to store information. Each robot has a fair coin which they can use to generate random bits. Each robot can communicate with other robots present at the same node by message passing: a robot can broadcast some message which is received by all robots present at the same node. The size of a message is no more than its memory size because it cannot generate a message whose size is greater than its memory size. Therefore, the size of a message must be $O(\log \Delta)$. Also, when there are many robots (co-located at a node) broadcasting their messages, it is not possible for a robot to receive all of these messages due to limited memory. When there is not enough memory to receive all the messages, it receives only a subset of the messages. The view of a robot is local: the only things that a robot can 'see' when it is at some node, are the edges incident to it. The robots have access to the port numbers of these edges. It cannot 'see' the other robots that may be present at the same node. The only way it can detect the presence of other robots is by receiving messages that those robots may broadcast. The robots can move from one node to an adjacent node. Any number of robots are allowed to move via an edge. When a robot moves from a node u to node v, it is aware of the port through which it enter v.

Time Cycle. We assume a fully synchronous system. The time progresses in rounds. Each robot knows when a current round ends and when the next round starts. Each round consists of the following.

- The robots first performs a series of synchronous computations and communications. These are called *subrounds*. In each subround, a robot performs some local computations and then broadcasts some messages. The messages received in the ith subround are read in the $(i + 1)$th

subround. The local computations are based on its memory contents (which contains the messages that it might have received in the last subround and other information that it had stored) and a random bit generated by the fair coin.

- Then robots move through some port or remains at the current node.

Problem Definition. A team of k ($\leq n$) robots are initially at the same node of the graph G. The DISPERSION problem requires the robots to re-position themselves so that i) there is at most one robot at each node, and ii) all robots have terminated within a finite number of rounds.

3 The Algorithm

3.1 Local Leader Election

Before presenting our main algorithm, we give a brief description of the LEADERELECTION() subroutine. We adopt this subroutine from [25]. When $k \geq 1$ robots are co-located together at a node, LEADERELECTION() subroutine allows exactly one robot to be selected as the leader within one round. Formally, 1) if $k = 1$, the robot finds out that it is the only robot at the node, 2) if $k > 1$, after finitely many rounds (with high probability), i) exactly one robot is elected as leader, ii) all robots can detect when the process is completed. Each robot starts off as a candidate for leader. In the first subround, every robot broadcasts 'start'. If a robot finds that it has received no message, it then concludes that it is the only robot at the node. Otherwise, it concludes that there are multiple robots at the node and does the following. In each subsequent subround, each candidate flips a fair coin. If heads, it broadcasts 'heads', otherwise it does not broadcast anything. If a robot gets tails, and receives at least one ('heads') message, it stops being a candidate. This process is repeated until exactly one robot, say r, broadcasts in a given subround. In this subround, r broadcasts 'heads', but receives no message, while all other non-candidate robots have not broadcasted, but received exactly one message. So r elects itself as the leader, and all robots detect that the process is completed. The process requires $O(1)$ bits of memory at each robot and terminates in $O(\log k)$ subrounds with high probability.

3.2 Overview of the Algorithm

In this subsection, we present a brief overview of the algorithm. The execution of our algorithm can be divided into three stages. In the first stage, the robots, together as a group, perform a DFS traversal in the search of empty nodes, starting from the node where they are placed together initially. We shall call this node the $root$ and denote it by v_R. Whenever the group reaches an empty node, they perform the LEADERELECTION() subroutine to elect a leader. The leader settles at that node, while the rest of the group continues the DFS traversal. Note that the settled robot does not terminate. This is because when the robots that are performing the DFS return to that node, they need to detect that the node is occupied by a settled robot. Recall that a robot cannot distinguish between an empty node and a node with a terminated robot. Therefore, the active settled robot helps the travelling robots to distinguish between an occupied node and an empty node, and also provides them with other information that are required to correctly execute the DFS. The size of the travelling group decreases by one, each time the DFS traversal reaches an empty node. The first stage completes when each robot has found an empty node for itself. Let r_L denote the last robot that finds an empty node, v_L, for itself. Although dispersion is achieved, this robot will not terminate. The other settled robots do not know that dispersion is achieved and will remain active. Therefore r_L needs to revisit those nodes and ask the settled robots to terminate. First r_L will return to the root v_R via the *rootpath* which is the unique path in the DFS tree from v_L to v_R. This is the second stage of the algorithm. In the third stage, r_L performs a second DFS traversal and asks the active settled robots to terminate. Since the active settled robots play crucial in the DFS traversal, r_L needs to be careful about the order in which it should ask the settled robots to terminate. Finally, r_L terminates after it returns to v_L.

Variable	Description
role	It indicates the role that the robot is playing in the algorithm. It takes values from {`explore`, `settled`, `return`, `acknowledge`, `done`}. Initially, $role \leftarrow$ `explore`.
entered	It indicates the port through which the robot has entered the current node. Initially, $entered \leftarrow \emptyset$. For simplicity, assume that it is automatically updated when the robot entered a node.
received	It indicates the message(s) received by the robot in the current subround. After the end of each subround, the messages are erased, i.e., it is reset to $\emptyset$. Initially, $received \leftarrow \emptyset$.
direction	It indicates the direction of movement of a robot during a DFS traversal. It takes values from {`forward`, `backward`}. Initially, $role \leftarrow$ `forward`.
parent	For a settled robot on some node, it indicates the port number towards the parent of that node in the DFS tree. Initially, $parent \leftarrow \emptyset$.
child	For a settled robot on some node that is on the rootpath, it indicates the port number towards the child of that node in the DFS tree that is on the rootpath. Initially, $child \leftarrow \emptyset$.
visited	For a settled robot on some node, it indicates whether the node where the robot is settled has been visited by r_L in the third stage. Initially, $visited \leftarrow 0$.

Table 1: Description of the variables used by the robots

A pseudocode description of the algorithm is given in Algorithm 1. In Table 1, we give details of the variables used by the robots. If $variable_name$ is some variable, then we shall denote the value of the variable stored by r as $r.variable_name$.

3.3 Detailed Description of the Algorithm

In the starting configuration, all robots are present at the root node v_R. Initially, *role* of each robot is `explore`. In the first stage, the robots have to perform a DFS traversal together as a group. This group of robots is called the *exploring group*. Whenever the exploring group reaches an empty node (a node with no settled robot), one of the robots will settle at that node, i.e., it will change its *role* to `settled` and remain at that node. For the rest of the algorithm, it does not move. However, it stays active and checks for any received messages. A settled robot can receive three types of messages:

- it may receive a query about the contents of its internal memory
- it may receive an to instruction to change the value of some variable
- it may be asked to terminate

When queried about its memory, it broadcasts a message containing its *role*, *parent*, *child* and *visited*. If it is asked to change the value of some variable or terminate, then it does so accordingly. Any robot with role `explore`, `return` or `acknowledge`, in the first subround of any round, broadcasts a message querying about internal memory of any settled robot at the node. If it receives no message in the second subround, then it concludes that there is no settled robot at that node. Whenever the robots find that there is no settled robot at the node, during the first stage, they start the LeaderElection() subroutine to elect a leader. For any robot r, LeaderElection() results in one of the following outcomes:

- it is elected as the leader

Algorithm 1: Dispersion

```
Procedure DISPERSION()
    r ← myself
    if r.role = settled then
        if I am queried then
            BROADCAST(role = settled, parent = r.parent, child = r.child, visited = r.visited)
        else if r.received = "Set child x" then
            r.child ← x
        else if r.received = "Set visited = 1" then
            r.visited ← 1
        else if r.received = "terminate" then
            TERMINATE()
    else if r.role = explore then
        QUERY()
        if r.received = ∅ then
            LEADERELECTION()
            if I am alone then
                r.role ← return
                Move through r.entered
            else if I am elected as leader then
                r.role ← settled
                r.parent ← r.entered
            else if I am not elected as leader then
                if r.entered = ∅ then
                    Move via port 0
                else
                    if (r.entered + 1 = r.entered) mod δ then
                        r.direction ← backward
                    Move via port (r.entered + 1)modδ
        else if r.received = "role = settled, parent = x, child = ∅, visited = 0" then
            if r.direction = forward then
                r.direction ← backward
                Move via port r.entered
            else if r.direction = backward then
                if (r.entered + 1)modδ = x then
                    Move via port (r.entered + 1)modδ
                else
                    r.direction ← forward
                    Move via port (r.entered + 1)modδ
    else if r.role = return then
        QUERY()
        if r.received = "role = settled, parent = x, child = ∅, visited = 0" then
            if x ≠ ∅ then
                BROADCAST("Set child r.entered")
                Move via port x
            else
                BROADCAST("Set child r.entered")
                r.direction ← forward
                r.role ← acknowledge
                r.entered ← ∅
    else if r.role = acknowledge then
        QUERY()
        if r.received = "role = settled, parent = x, child = y, visited = 0" then
            BROADCAST("Set visited = 1")
            if r.entered = ∅ then
                Move via port 0
            else
                if r.entered = (r.entered + 1)modδ then
                    r.direction ← backward
                    BROADCAST("terminate")
                else if y = (r.entered + 1)modδ then
                    BROADCAST("terminate")
                Move via port (r.entered + 1)modδ
        else if r.received = "role = settled, parent = x, child = y, visited = 1" then
            if r.direction = forward then
                r.direction ← backward
                Move via port r.entered
            else if r.direction = backward then
                if x = (r.entered + 1)modδ then
                    BROADCAST("terminate")
                else if y = (r.entered + 1)modδ then
                    r.direction ← forward
                    BROADCAST("terminate")
                else
                    r.direction ← forward
                Move via port (r.entered + 1) mod δ
        else if r.received = ∅ then
            if r.direction = forward then
                r.direction ← backward
            else if r.direction = backward then
                r.role = done
            Move via port r.entered
    else if r.role = done then
        TERMINATE()
```

- it is not elected as the leader
- it finds that it is the only robot at that node

In the first case, it changes $r.role$ to `settled` and sets $r.parent$ equal to $r.entered$. Recall that $r.entered$ is the port through which it entered the current node and in the beginning, $r.entered$ is set to $\emptyset$. We shall call $r.parent$ the *parent port* of the node where r resides. We shall refer to a robot that has set its $role$ to `settled` as a *settled robot*. In the second case, it will continue the DFS: if $r.entered = \emptyset$, it leaves via port 0 and if $r.entered \neq \emptyset$, it leaves via port $(r.entered + 1)$ mod δ. If $(r.entered + 1) = r.entered$ mod δ, it changes its variable *direction* to `backward` before exiting the node. Recall that the variable *direction* is used to indicate the direction of the movement during a DFS traversal. In the third case, it changes $r.role$ to `return`.

Now consider the case where the robots find that there is a settled robot at the node. If the *direction* is set to `forward` when they encounter the settled robot, it indicates the onset of a cycle. So the robots change the *direction* to `backward` and leave the node via the port through which they entered it. Now suppose that the *direction* is set to `backward` when they encounter the settled robot. Recall that the robots have received from the settled robot, say a, a message which contains $a.parent$. The robots check if $a.parent$ is equal to the port number through which it entered, say z, plus 1 (modulo the degree of the node). If yes, it implies that the robots have moved through all edges adjacent to the node, and hence they leave the node via $a.parent$ which is the port through which they entered for the first time. If no, then it means that they have not moved through the port $(z + 1)$modδ before. So they change the direction to `forward` and leave via $(z + 1)$modδ.

The DFS traversal in the first stage ends when a robot, say r_L, with *role* set to `explore`, finds that it is the only robot at a node, say v_L. Recall that when this happens, r_L changes its *role* to `return`. At this point, the first stage ends, and the second stage starts. It then leaves v_L via the port through which it entered. In each of the following rounds where the *role* of r_L is `return`, it does the following. In the first subround, it broadcasts a query. In the next subround, it receives a message from the settled robot at that node which contains its *parent*. If the obtained value of *parent*, say x, is not $\emptyset$, it means that r_L is yet to reach the root v_R. Then r_L broadcasts an instruction for the settled robot to change the value of its *child* to the port via which r_L entered the node. This value of $child$ will be called the *child port* of the node. After broadcasting the instruction, r_L leaves through the port x. If $x = \emptyset$, then it means that r_L has reached the root v_R. In this case, r_L broadcasts the same instruction and then changes the values of $r_L.role$, $r_L.direction$ and $r_L.entered$ to respectively `acknowledge`, `forward` and $\emptyset$. At this point the second stage ends, and the third stage starts.

In the following rounds, r_L with *role* `acknowledge` does the following. In the first subround, it broadcasts a query. It either receives a reply or does not. If it receives a message, then it contains the values of *parent*, say x, and *child*, say y, and *visited* of the settled robot at that node. Now, the value of variable *visited* can be 0 or 1. If the value of *visited* is 0, it denotes that the settled robot is visited for the first time in the third stage. The robot r_L then broadcasts a message instructing the settled robot to change the value of its variable $visited$ to 1. Now $(r_L.entered + 1)$modδ can be equal to $r_L.entered$ (the case of one degree node) or y or neither of them. In the former case, it changes its variable $direction$ to `backward`. In the first two cases, it broadcasts a message instructing the settled robot to terminate and leaves through port $(r_L.entered + 1)$modδ. If $(r_L.entered + 1)$modδ is neither equal to $r_L.entered$, nor equal to y, r_L just exits through $(r_L.entered + 1)$modδ without broadcasting any message for termination. If the value of *visited* is 1, it denotes that the settled robot has been visited before in the third stage. If the value of variable $direction$ of r_L is `forward`, it changes the value of $direction$ to `backward` and exits through the port through which it entered the node at the previous round. Otherwise the value of variable $direction$ of r_L is `backward`. In this case, three sub-cases arise. If $(r_L.entered + 1)$modδ is equal to x, then r_L broadcasts a message instructing the settled robot to terminate, and then r_L exits through the port $(r_L.entered + 1)$modδ. Otherwise if, $(r_L.entered + 1)$modδ is equal to y, then also r_L broadcasts a message instructing the settled robot to terminate, changes the variable $direction$ to `forward` and then r_L exits through the port $(r_L.entered + 1)$modδ. If $(r_L.entered + 1)$modδ is neither equal to x, nor y, then r_L changes

$direction$ to `forward` and exits through $(r_L.entered + 1)\text{mod}\delta$. Now, we consider the case where r_L in third stage does not receive any answer to its query. If its $direction$ is set to `forward`, it changes its $direction$ to `backward` and then exits through the same port by which it entered the node in the previous round. If its direction is set to `backward`, then it means that r_L was at v_L in the previous round. So r_L changes its *role* to `done` and leaves the node through the port via which it entered. Then it will reach v_L in the next round and it will find that its *role* is `done` and terminate.

4 Correctness Proof and Complexity Analysis

The first stage of our algorithm is the same as that of [25]. The robots simply perform a DFS traversal. Whenever a new node is visited, one of the robots settle there. The DFS continues until k distinct nodes are visited. To see that the DFS traversal can be correctly executed in our setting, it suffices to verify that the robots can correctly ascertain 1) if a node is previously visited and 2) if all neighbors of a node have been visited. For 1), observe that the presence of settled robot at a node indicates that the node has already been visited. So, when the robots with *direction* `forward` go to a node which has a settled robot, it backtracks, i.e., it changes its $direction$ to `backward` and leaves the node via the port through which it had entered. For 2), observe that the port p through which robots first enters a node v is set as its parent port, i.e., the robot settled at v sets its variable $parent$ to p. Then the robots will move through all other ports with *direction* `forward` in the order $p+1, p+2, \ldots, \delta-1, 0, 1, \ldots, p-1$ (unless the DFS is stopped midway for k distinct nodes have been visited). This is because if the robots leaves via a port q (with *direction* `forward`), it re-enters v via the same port q after some rounds (with *direction* `backward`) and then leaves via $(q+1)\text{mod}\delta(v)$ (with *direction* `forward`) in the next round if $(q+1)\text{mod}\delta(v) \neq p$. Clearly, when $(q+1)\text{mod}\delta(v) = p$, it indicates that the robots have moved through all ports other than p with *direction* `forward`, i.e., all neighbors of v have been visited. The robots can check if $(q+1)\text{mod}\delta(v) = p$ because their variable *entered* is equal to q and the variable *parent* of the robot settled at v is equal to p. Inspecting the pseudocode of Algorithm 1, it is easy to see that these are correctly implemented in the algorithm. Therefore, have the following result.

Theorem 1. *There is a round t_1, at the beginning of which*

1. *each node of G has at most one robot*
2. role *of exactly one robot r_L is `explore` and the role of the remaining $k-1$ robots is `settled`*
3. *if $V' \subseteq V$ is the set of nodes occupied by robots, then $G[V']$ (the subgraph of G induced by V') is connected*
4. *if $E' \subseteq E$ is the set of edges corresponding to the variable* parent *of robots in $\mathcal{R} \setminus \{r_L\}$ and variable* entered *of r_L, then the graph $T = T(V', E')$ is a DFS spanning tree of $G[V']$,*
5. *r_L is at a leaf node v_L of T.*

In the following lemmas, we present some observations regarding the execution of DFS traversal in the first stage.

Lemma 1. *If v is a non-rootpath node, then the exploring group leaves it via its parent port once. If v is a rootpath node other than v_L, then the exploring group leaves it via its child port once.*

Lemma 2. *Suppose that v is a non-rootpath node and the exploring group leaves v through its parent port at round t with $direction$ `backward`. If the exploring group returns to v at some round $t', t < t' \leq t_1$, then its $direction$ must be `forward`.*

Lemma 3. *Suppose that v is a rootpath node and the exploring group leaves v through its child port at round t with $direction$ `forward`. If the exploring group returns to v at some round $t', t < t' \leq t_1$, then its $direction$ must be `forward`.*

There is a unique path, i.e., the rootpath $v_R = v_1, v_2, \ldots, v_s = v_L$ in $T(V', E')$ from v_R to v_L. Furthermore, for any consecutive vertices v_i, v_{i+1} on the path 1) if $i + 1 < s$, the variable *parent* of the settled robot at v_{i+1} is set to port(v_{i+1}, v_i) and 2) if $i + 1 = s$, the variable *entered* of robot r_L at v_{i+1} is set to port(v_{i+1}, v_i). So, according to our algorithm, r_L will move along this path to reach v_R. For each node $v_i, i < s$, on the rootpath, when r_L reaches v_i along its way to v_R, it instructs the settled robot at v_i to set its variable *child* to port(v_i, v_{i+1}). Therefore, we have the following result.

Theorem 2. *There is a round t_2, at the beginning of which*

1. *r_L is at v_R with $r_L.role =$ `return`*
2. *each node of $T(V', E') \setminus \{v_L\}$ has a settled robot*
3. *if $v_R = v_1, v_2, \ldots, v_s = v_L$ is the rootpath and r_i is the settled robot at $v_i, i < s$, then $r_i.child =$ port(v_i, v_{i+1})*
4. *if r is a settled robot on a non-rootpath node, then $r.child = \emptyset$.*

From round $t_2 + 1$, r_L will start a second DFS traversal. This DFS traversal is trickier than the earlier one because the settled robots will one by one terminate during the process. Recall that the settled robots played important role in the first DFS. We shall prove that r_L will correctly execute the second DFS traversal. In fact, we shall prove that the DFS traversal in the first stage is exactly same as the DFS traversal in the third stage in the sense that if the exploring group is at node v at round $i < t_1$ (in the first stage), then r_L is at node v at round $t_2 + i$ (in the third stage).

Let us first introduce a definition. In the following definition, whenever we say 'at round', it is to be understood as 'at the beginning of round'. Round i in the first stage is said to be *identical* to round j in the third stage if the exploring group at round i and r_L at round j are at the same node, say u and one of the following holds:

I1 At round i, there is no settled robot at u, the exploring group contains more than one robot and the variable $direction$ for each robot in the exploring group is set to `forward`. At round j there is a settled robot at u with its variable $visited$ set to 0. The variables $direction$ and $entered$ of r_L at round j are equal to those of each robot in the exploring group at round i.

I2 At round i, there is a settled robot at u, the variable $direction$ for each robot in the exploring group is set to `forward` and the variable $entered$ for each robot in the exploring group is $\neq \emptyset$. At round j, either there is a terminated robot at u, or there is a settled robot at u with its variable $visited$ set to 1. The variables $direction$ and $entered$ of r_L at round j are equal to those of each robot in the exploring group at round i .

I3 At round i, there is a settled robot at u, the variable $direction$ for each robot in the exploring group is set to `backward` and the variable $entered$ for each robot in the exploring group is $\neq \emptyset$. At round j, there is an active settled robot at u with its variable $visited$ set to 1. The variables $direction$ and $entered$ of r_L at round j are equal to those of each robot in the exploring group at round i .

Lemma 4. *Round i is identical to $t_2 + i$ for all $1 \leq i < t_1$.*

Proof. We prove this by induction. At the beginning of round 1, the exploring group (consisting of more than one robots) is at the root node v_R, there is no settled robot at v_R, and the variables $direction$ and $entered$ of each robot in the exploring group are set to `forward` and $\emptyset$ respectively. Note that when the robot r_L enters the root node v_R at round t_2, it sets its $direction$ and $entered$ to `forward` and $\emptyset$ respectively, and does not move (See line 46-49 in Algorithm 1). Hence at the beginning of round $t_2 + 1$, r_L is at node v_R, with its $direction$ and $entered$ set to `forward` and $\emptyset$ respectively. Furthermore, there is a settled robot at v_R at round $t_2 + 1$. This is the robot which was elected as the

leader in round 1 and had settled there. Since the variable $visited$ of this robot was initially set to 0 and it has been not instructed to change it thus far, it is still set to 0 at the beginning of t_2+1. These observations imply that I1 holds and round 1 is identical to t_2+1.

Now assume that round j is identical to round t_2+j for all $1 \leq j \leq i-1, i < t_1$. We shall prove that round i is identical to round t_2+i. Let the robot r_L be at node u at round t_2+i-1. Then it implies that the exploring group was at node u at round $i-1$. Now we can have following three cases.

Case 1. Suppose that at round $i-1$, 1) there is no settled robot at u, 2) the exploring group contains more than one robot and 3) the variable $direction$ for each robot in the exploring group is set to `forward`. Suppose that their variable $entered$ are set to p. Then the robots will perform the $\textsc{LeaderElection}()$ protocol and one of them will settle. The remaining robots will move via 0 if $p = \emptyset$ or otherwise, will move via $(p+1)$ mod δ. In the later case, the robots will change their $direction$ to `backward` iff $(p+1)$ mod $\delta = p$. Since we assumed that round $i-1$ is identical to round t_2+i-1, at the beginning of round t_2+i-1, 1) r_L is at u, 2) there is a settled robot at u with its variable $visited$ set to 0, and 3) the variables $direction$ and $entered$ of r_L are equal to `forward` and p. According to our algorithm, r_L will move via 0 if $p = \emptyset$ or otherwise, will move via $(p+1)$modδ. In the later case, r_L will change its $direction$ to `backward` iff $(p+1)$mod$\delta = p \Leftrightarrow \delta(u) = 1$. Hence the exploring group at round i and r_L at t_2+i must be at the same node, say v, and have same values of $direction$ and $entered$. So now it remains to show that one of I1, I2 or I3 is true for round i and t_2+i. For this, consider the following two cases.

Case 1a. First consider the case where the exploring group and r_L exit u with $direction$ `forward` at round $i-1$ and t_2+i-1 respectively. At the beginning of round i, the exploring group has more than one robots as $i < t_1$. Now, at the beginning of round i, v either has no settled robot or has a settled robot. In the first case, it implies that exploring group is entering v for the first time at round i. The robots will then perform the $\textsc{LeaderElection}()$ protocol and one of them, say a, will settle. By the induction hypothesis, it implies that r_L will enter v for the first time in the third stage at round t_2+i. Hence the settled robot a must be there with the value of $visited$ set to 0. So I1 holds and hence round i is identical to round t_2+i. In the second case, it implies that the exploring group had entered v at some previous round $j < i$. Then by the induction hypothesis, r_L had visited v at round $t_2+j < t_2+i$. Hence, if the robot a is active, then variable $visited$ of a is 1, or otherwise a has terminated. So I2 holds and hence round i is identical to round t_2+i.

Case 1b. Now consider the case where the exploring group and r_L exit u with $direction$ `backward` at round $i-1$ and t_2+i-1 respectively. This implies that u is a one degree node. Also, the exploring group and r_L were at v at round $i-2$ and t_2+i-2 respectively. Therefore, there must be a settled robot, say b, present at v , when the exploring group visits it at round i, as the node was visited earlier. So when r_L enters v at round t_2+i, if b is still active, its $visited$ value is set to 1 and I3 holds. We prove that this is the only case. For the sake of contradiction, let us assume that b has terminated. Consider the following cases.

- First let v be a non-rootpath node. Since b has terminated, it implies that r_L had exited v via its parent port with $direction$ `backward` at some round $t_2+l < t_2+i-1$. Then by the induction hypothesis, the exploring group exited v via its parent port with $direction$ `backward` at some round $l < i-1$. But then the fact that the exploring group returns to v with $direction$ `backward` at round i contradicts Lemma 2.

- Now let v be a rootpath node. Since b has terminated, it implies that r_L had exited v via its child port with $direction$ `forward` at some round $t_2+l < t_2+i-1$. Then by the induction hypothesis, the exploring group exited v via its child port with $direction$ `forward` at some round $l < i-1$. But then the fact that the exploring group returns to v with $direction$ `backward` at round i contradicts Lemma 3.

Case 2. At the beginning of round $i-1$, 1) there is a settled robot at u, 2) the variable $direction$ for each robot in the exploring group is set to `forward` and 3) the variable $entered$ for each robot in the exploring group is $p \neq \emptyset$. According to our algorithm, the exploring group will change their

$direction$ to backward and leave the node via p. Since we assumed that round $i-1$ is identical to round t_2+i-1, at the beginning of round t_2+i-1, r_L is at u, there is either an active settled robot at u with its variable $visited$ set to 1 or a terminated robot, and the variables $direction$ and $entered$ of r_L are equal to forward and p respectively. According to our algorithm, r_L will change it $direction$ to backward and leave the node via p. Hence the exploring group at round i and r_L at t_2+i must be at the same node, say v, both having $direction$ set to backward and $entered$ set to port(v,u).

It is clear that the exploring group and r_L was at v at rounds $i-2$ and t_2+i-2 respectively. Hence, there is a settled robot at v, say a, at round i, as it had been visited at least once before. Also, a is still situated at v at round t_2+i, either active or terminated. If it is active, then the value of its variable $visited$ is 1, as r_L had been at v at round t_2+i-2. So, in this case I3 holds. Using the same arguments as in Case 1b, we can show that this is the only possible case.

Case 3. At the beginning of round $i-1$, 1) there is a settled robot at u, say a, 2) the variable $direction$ for each robot in the exploring group is set to backward and 3) the variable $entered$ for each robot in the exploring group is $p \neq \emptyset$. Let the $parent$ of a be equal to x. According to our algorithm, the robots will not change their $direction$ if $(p+1)\text{mod}\delta = x$ (Case 3a), and otherwise, it will change it to forward (Case 3b). In any case, they will leave the node via port $(p+1)\text{mod}\delta$. Since we assumed that round $i-1$ is identical to round t_2+i-1, at the beginning of round t_2+i-1, 1) r_L is at u, 2) there is a settled robot at u with its variable $visited$ set to 1, and 3) the variables $direction$ and $entered$ of r_L are equal to backward and p respectively. Since a settled robot does not move or change its $parent$, the settled robot at u at round t_2+i-1 is a and its $parent$ is set to x. According to our algorithm, r_L will not change its $direction$ if $(p+1)\text{mod}\delta = x$ (Case 3a), and otherwise, it will change it to forward (Case 3b). In any case, they will leave the node via port $(p+1)\text{mod}\delta$. Hence at round i and t_2+i, the exploring group and r_L must be at the same node, say v, with same $direction$ and $entered$. So now it remains to show that one of I1, I2 or I3 is true for round i and t_2+i.

Case 3a. Clearly in this case v must have been visited at least once before round i in the first stage. Hence a robot, say c, is already settled there at round i. Hence, at round t_2+i, either c is active with $visited$ set to 1, or is terminated. In the first case, I3 holds and hence we are done. We can prove that the later case is impossible using the same arguments as in Case 1b.

Case 3b. If there is no settled robot at v at round i, then the exploring group is visiting v for the first time and one of them, say b, will settle there. It implies from our induction hypothesis that r_L is also visiting v for the first time in the third stage at round t_2+i. Hence it will find b with its $visited$ set to 0. So I1 holds. If there is a settled robot at v at round i, say c, then at round t_2+i, v has c which is either terminated or its variable $visited$ set to 1. So we see that I2 holds. □

Theorem 3. *By round t_2+t_1 all the settled robots have terminated.*

Proof. A settled robot at a non-rootpath node will terminate if r_L moves from that node to its parent and a settled robot at a rootpath node will terminate if r_L moves from that node to its child. Recall that if v is a non-rootpath node, then in Phase 1 the exploring group leaves it via its parent port once. Also, if v is a rootpath node other than v_L, then in Phase 1 the exploring group leaves it via its child port once. So the result follows from Lemma 4. □

Theorem 4. *At round t_2+t_1+2, r_L terminates at v_L.*

Proof. It follows from Lemma 4 that r_L will be at v_{s-1} at round t_2+t_1-1. It is easy to see that it will then move to $v_s = v_L$ with $direction$ forward. So at round t_2+t_1, r_L is at v_L with $direction$ set to forward. Since v_L has no settled robot, r_L will change its $direction$ to backward and exit through the port through which it entered v_L. But moving through this port leads to v_{s-1}. The settled robot at this node is already terminated by Lemma 3. Hence at round t_2+t_1+1, r_L enters with $direction$ backward a node where it does not receive any message. So it will then change its $role$ to done and exit the port through which it entered the node. Therefore, at round t_2+t_1+2, r_l again returns to v_L, this time with $role$ done, and hence will terminate. □

Theorem 5. *Algorithm 1 is correct.*

Proof. It follows from the above results that at round $t_2 + t_1 + 2$, dispersion accomplished and all robots have terminated. Hence Algorithm 1 solves DISPERSION from any rooted configuration. □

Theorem 6. *Algorithm 1 requires $O(\log \Delta)$ bits of memory at each robot and this is optimal in terms of memory complexity.*

Proof. The LEADERELECTION() subroutine costs $O(1)$ bits of memory for each robot. Among the variables, $role$, $visited$ and $direction$ costs $O(1)$ bits of memory, and the variables $entered$, $parent$, $child$, $received$ costs $O(\log \Delta)$ bits of memory for each robot. Hence the algorithm requires $O(\log \Delta)$ bits of memory at each robot. The optimality follows from the lower bound result proved in [25]. □

Theorem 7. *Algorithm 1 requires $\Theta(k^2)$ rounds in the worst case (assuming that LEADERELECTION() terminates each time).*

Proof. Exactly k distinct nodes are visited in our algorithm. In the first stage, movement of the exploring group takes $O(m')$ rounds where m' is the number of edges in the subgraph of G induced by these k vertices. Clearly $m' = O(k^2)$. So the first stage requires $O(k^2)$ rounds. The third stage requires (k^2) rounds as well since apart from the last two rounds, it is exactly identical to the first stage. Clearly the second round takes $O(k)$ rounds. So the overall round complexity is $O(k^2)$.

To see that our analysis is tight, we show an instance where $\Omega(k^2)$ rounds will be required. Consider the graph of size k in Fig. 1. Here port$(v_R, v_1) = 0$, port$(v_1, v_R) = 0$, port$(v_1, v_2) = 1$, port$(v_1, v_L) = 2$. Here, after reaching v_1, the exploring group will go v_2. It is easy to see that $\Theta(k^2)$ rounds will be spent inside the $(k-3)-$clique. Finally the last robot will return to v_1 and then move to v_L. □

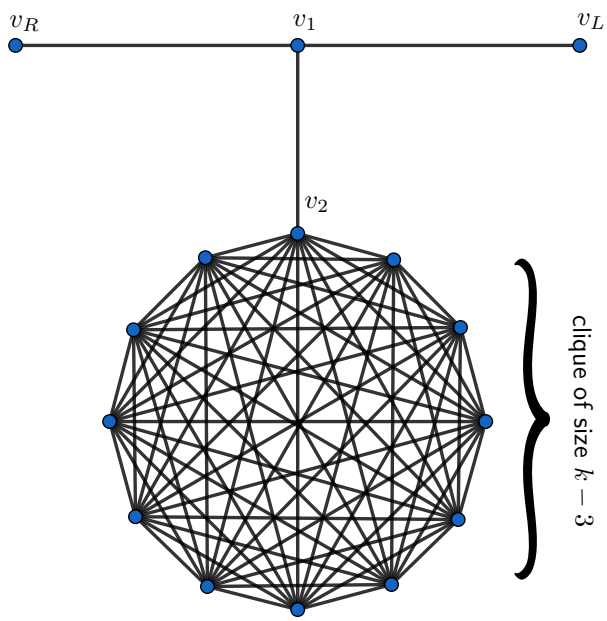

Figure 1: An example where Algorithm 1 requires $\Theta(k^2)$ rounds to complete.

5 Concluding Remarks

We have presented a memory optimal randomized algorithm for DISPERSION from rooted configuration by anonymous robots. This resolves an open problem posed in [25]. Time complexity of our algorithm

is $\Theta(k^2)$ rounds in the worst case (assuming that $\textsc{LeaderElection}()$ terminates each time). Any algorithm that solves the problem requires $\Omega(k)$ rounds in the worst case. To see this, consider a path with $n \geq k$ nodes with all robots initially at one of its one degree nodes. An interesting open problem is to close this gap.

For arbitrary configuration, the random walk based algorithm presented in [25] requires the robots to stay active indefinitely. Therefore an interesting open question is whether it is possible to solve the problem by anonymous robots from non-rooted configurations without requiring robots to stay active indefinitely.

Acknowledgement. We would like to thank Pritam Goswami for valuable discussions. The first two authors are supported by UGC, Govt. of India, and NBHM DAE, Govt. of India respectively. We would like to thank the anonymous reviewers for their valuable comments which helped us to improve the quality and presentation of the paper.

References

[1] Ankush Agarwalla, John Augustine, William K. Moses Jr., Sankar Madhav K., and Arvind Krishna Sridhar. Deterministic dispersion of mobile robots in dynamic rings. In Paolo Bellavista and Vijay K. Garg, editors, *Proceedings of the 19th International Conference on Distributed Computing and Networking, ICDCN 2018, Varanasi, India, January 4-7, 2018*, pages 19:1–19:4. ACM, 2018. doi:10.1145/3154273.3154294.

[2] John Augustine and William K. Moses Jr. Dispersion of mobile robots: A study of memory-time trade-offs. In *Proceedings of the 19th International Conference on Distributed Computing and Networking, ICDCN 2018, Varanasi, India, January 4-7, 2018*, pages 1:1–1:10, 2018. doi:10.1145/3154273.3154293.

[3] Kaustav Bose, Manash Kumar Kundu, Ranendu Adhikary, and Buddhadeb Sau. Optimal gathering by asynchronous oblivious robots in hypercubes. In *Algorithms for Sensor Systems - 14th International Symposium on Algorithms and Experiments for Wireless Sensor Networks, ALGOSENSORS 2018, Helsinki, Finland, August 23-24, 2018, Revised Selected Papers*, volume 11410 of *Lecture Notes in Computer Science*, pages 102–117. Springer, 2018. doi:10.1007/978-3-030-14094-6_7.

[4] Peter Brass, Flavio Cabrera-Mora, Andrea Gasparri, and Jizhong Xiao. Multirobot tree and graph exploration. *IEEE Trans. Robotics*, 27(4):707–717, 2011. doi:10.1109/TRO.2011.2121170.

[5] Reuven Cohen, Pierre Fraigniaud, David Ilcinkas, Amos Korman, and David Peleg. Label-guided graph exploration by a finite automaton. *ACM Trans. Algorithms*, 4(4):42:1–42:18, 2008. doi:10.1145/1383369.1383373.

[6] Jurek Czyzowicz, Adrian Kosowski, and Andrzej Pelc. How to meet when you forget: log-space rendezvous in arbitrary graphs. *Distributed Comput.*, 25(2):165–178, 2012. doi:10.1007/s00446-011-0141-9.

[7] Gianlorenzo D'Angelo, Gabriele Di Stefano, and Alfredo Navarra. Gathering on rings under the look-compute-move model. *Distributed Comput.*, 27(4):255–285, 2014. doi:10.1007/s00446-014-0212-9.

[8] Shantanu Das, Dariusz Dereniowski, and Christina Karousatou. Collaborative exploration of trees by energy-constrained mobile robots. *Theory Comput. Syst.*, 62(5):1223–1240, 2018. doi:10.1007/s00224-017-9816-3.

[9] Dariusz Dereniowski, Yann Disser, Adrian Kosowski, Dominik Pajak, and Przemyslaw Uznanski. Fast collaborative graph exploration. *Inf. Comput.*, 243:37–49, 2015. doi:10.1016/j.ic.2014.12.005.

[10] Yoann Dieudonné and Andrzej Pelc. Anonymous meeting in networks. *Algorithmica*, 74(2):908–946, 2016. doi:10.1007/s00453-015-9982-0.

[11] Krzysztof Diks, Pierre Fraigniaud, Evangelos Kranakis, and Andrzej Pelc. Tree exploration with little memory. *J. Algorithms*, 51(1):38–63, 2004. doi:10.1016/j.jalgor.2003.10.002.

[12] Christian A. Duncan, Stephen G. Kobourov, and V. S. Anil Kumar. Optimal constrained graph exploration. *ACM Trans. Algorithms*, 2(3):380–402, 2006. doi:10.1145/1159892.1159897.

[13] Tomoko Izumi, Taisuke Izumi, Sayaka Kamei, and Fukuhito Ooshita. Mobile robots gathering algorithm with local weak multiplicity in rings. In Boaz Patt-Shamir and Tínaz Ekim, editors, *Structural Information and Communication Complexity, 17th International Colloquium, SIROCCO 2010, Sirince, Turkey, June 7-11, 2010. Proceedings*, volume 6058 of *Lecture Notes in Computer Science*, pages 101–113. Springer, 2010. doi:10.1007/978-3-642-13284-1_9.

[14] Sayaka Kamei, Anissa Lamani, Fukuhito Ooshita, Sébastien Tixeuil, and Koichi Wada. Gathering on rings for myopic asynchronous robots with lights. In *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland*, volume 153 of *LIPIcs*, pages 27:1–27:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.OPODIS.2019.27.

[15] Ralf Klasing, Adrian Kosowski, and Alfredo Navarra. Taking advantage of symmetries: Gathering of many asynchronous oblivious robots on a ring. *Theor. Comput. Sci.*, 411(34-36):3235–3246, 2010. doi:10.1016/j.tcs.2010.05.020.

[16] Ralf Klasing, Euripides Markou, and Andrzej Pelc. Gathering asynchronous oblivious mobile robots in a ring. *Theor. Comput. Sci.*, 390(1):27–39, 2008. doi:10.1016/j.tcs.2007.09.032.

[17] Dariusz R. Kowalski and Adam Malinowski. How to meet in anonymous network. *Theor. Comput. Sci.*, 399(1-2):141–156, 2008. doi:10.1016/j.tcs.2008.02.010.

[18] Ajay D. Kshemkalyani and Faizan Ali. Efficient dispersion of mobile robots on graphs. In *Proceedings of the 20th International Conference on Distributed Computing and Networking, ICDCN 2019, Bangalore, India, January 04-07, 2019*, pages 218–227, 2019. doi:10.1145/3288599.3288610.

[19] Ajay D. Kshemkalyani, Anisur Rahaman Molla, and Gokarna Sharma. Fast dispersion of mobile robots on arbitrary graphs. In *Algorithms for Sensor Systems - 15th International Symposium on Algorithms and Experiments for Wireless Sensor Networks, ALGOSENSORS 2019, Munich, Germany, September 12-13, 2019, Revised Selected Papers*, pages 23–40, 2019. doi:10.1007/978-3-030-34405-4_2.

[20] Ajay D. Kshemkalyani, Anisur Rahaman Molla, and Gokarna Sharma. Dispersion of mobile robots in the global communication model. In Nandini Mukherjee and Sriram V. Pemmaraju, editors, *ICDCN 2020: 21st International Conference on Distributed Computing and Networking, Kolkata, India, January 4-7, 2020*, pages 12:1–12:10. ACM, 2020. doi:10.1145/3369740.3369775.

[21] Ajay D. Kshemkalyani, Anisur Rahaman Molla, and Gokarna Sharma. Dispersion of mobile robots on grids. In M. Sohel Rahman, Kunihiko Sadakane, and Wing-Kin Sung, editors, *WALCOM: Algorithms and Computation - 14th International Conference, WALCOM 2020, Singapore, March 31 - April 2, 2020, Proceedings*, volume 12049 of *Lecture Notes in Computer Science*, pages 183–197. Springer, 2020. doi:10.1007/978-3-030-39881-1_16.

[22] Giuseppe Antonio Di Luna, Paola Flocchini, Linda Pagli, Giuseppe Prencipe, Nicola Santoro, and Giovanni Viglietta. Gathering in dynamic rings. *Theor. Comput. Sci.*, 811:79–98, 2020. doi:10.1016/j.tcs.2018.10.018.

[23] Avery Miller and Andrzej Pelc. Fast rendezvous with advice. *Theor. Comput. Sci.*, 608:190–198, 2015. doi:10.1016/j.tcs.2015.09.025.

[24] Avery Miller and Andrzej Pelc. Time versus cost tradeoffs for deterministic rendezvous in networks. *Distributed Comput.*, 29(1):51–64, 2016. doi:10.1007/s00446-015-0253-8.

[25] Anisur Rahaman Molla and William K. Moses Jr. Dispersion of mobile robots: The power of randomness. In T. V. Gopal and Junzo Watada, editors, *Theory and Applications of Models of Computation - 15th Annual Conference, TAMC 2019, Kitakyushu, Japan, April 13-16, 2019, Proceedings*, volume 11436 of *Lecture Notes in Computer Science*, pages 481–500. Springer, 2019. doi:10.1007/978-3-030-14812-6_30.

[26] Anisur Rahaman Molla, Kaushik Mondal, and William K. Moses Jr. Efficient dispersion on an anonymous ring in the presence of weak byzantine robots. In Cristina Maria Pinotti, Alfredo Navarra, and Amitabha Bagchi, editors, *Algorithms for Sensor Systems - 16th International Symposium on Algorithms and Experiments for Wireless Sensor Networks, ALGOSENSORS 2020, Pisa, Italy, September 9-10, 2020, Revised Selected Papers*, volume 12503 of *Lecture Notes in Computer Science*, pages 154–169. Springer, 2020. doi:10.1007/978-3-030-62401-9_11.

[27] Petrisor Panaite and Andrzej Pelc. Exploring unknown undirected graphs. *J. Algorithms*, 33(2):281–295, 1999. doi:10.1006/jagm.1999.1043.

[28] Gabriele Di Stefano and Alfredo Navarra. Optimal gathering of oblivious robots in anonymous graphs and its application on trees and rings. *Distributed Comput.*, 30(2):75–86, 2017. doi:10.1007/s00446-016-0278-7.

[29] Amnon Ta-Shma and Uri Zwick. Deterministic rendezvous, treasure hunts, and strongly universal exploration sequences. *ACM Trans. Algorithms*, 10(3):12:1–12:15, 2014. doi:10.1145/2601068.