



Faster Leader Election and Its Applications for Mobile Agents with Parameter Advice

Ajay D. Kshemkalyani¹ , Manish Kumar² , Anisur Rahaman Molla³ ,
and Gokarna Sharma⁴

¹ University of Illinois Chicago, Chicago, USA
ajay@uic.edu

² Indian Institute of Technology Madras, Chennai, India
manishsky27@gmail.com

³ Indian Statistical Institute, Kolkata, India

⁴ Kent State University, Kent, USA
sharma@cs.kent.edu

Abstract. Leader election is a critical and extensively studied problem in distributed computing. This paper introduces the study of leader election using mobile agents. Consider n agents initially placed arbitrarily on the nodes of an arbitrary, n -node, m -edge graph G . These agents move autonomously across the nodes of G and elect one agent as the leader such that the leader is aware of its status as the leader, and the other agents know they are not the leader. The goal is to minimize both time and memory usage.

We study the leader election problem in a synchronous setting where each agent performs operations simultaneously with the others, allowing us to measure time complexity in terms of rounds. We assume that the agents have prior knowledge of the number of nodes n and the maximum degree of the graph Δ .

We first elect a leader deterministically in $O(n \log^2 n + D \Delta \log n)$ rounds with each agent using $O(\log n)$ bits of memory, where D is the diameter of the graph. Leveraging this leader election result, we then present a deterministic algorithm for constructing a minimum spanning tree of G in $O(m + n \log n)$ rounds, with each agent using $O(\Delta \log n)$ bits of memory.

Finally, using the same leader election result, we improve time and memory bounds for other key distributed graph problems, including gathering, maximal independent set, and minimal dominating set. For all the aforementioned problems, our algorithms remain memory-optimal.

Keywords: Distributed algorithms · mobile agents · local communication · leader election · MST · MIS · gathering · minimal dominating set · time and memory complexity · graph parameters

M. Kumar—Supported by CyStar at IIT Madras.

A. R. Molla—Supported, in part, by ANRF-SERB Core Research Grant, file no. CRG/2023/009048 and R. C. Bose Centre's internal research grant.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025

Q. Bramas et al. (Eds.): ICDCIT 2025, LNCS 15507, pp. 108–123, 2025.

https://doi.org/10.1007/978-3-031-81404-4_9

1 Introduction

Leader election is one of the fundamental and well-studied problems in distributed computing due to its significance in various applications, including resource allocation, reliable replication, load balancing, synchronization, membership management, and crash recovery. It can also be viewed as a form of symmetry breaking, where a single designated process or node (the leader) is responsible for making the critical decisions. In this paper, we explore the graph-level task of leader election in a distributed network under the agent-based model. In the agent-based model, the leader election problem involves a group of agents operating within a distributed network, where the goal is for exactly one agent to declare itself as the leader.

The agent-based model is particularly valuable in scenarios like private military networks or sensor networks in remote areas, where direct access to the network may be challenging, but small, battery-powered, relocatable devices can traverse the environment to gather information about the network’s structure for management purposes. Significant applications of the agent-based model in network management can be found in areas such as underwater navigation [6], network-centric warfare in military systems [15], social network modeling [25], and the study of social epidemiology [8]. Moreover, limited storage capacity reduces the risk of sensitive network information being exposed if a device is compromised. Additionally, mobility in these devices enhances communication security by mitigating concerns over message interception.

Recent applications of the agent-based model have emerged in various areas. A notable example is the work by Martinkus et al. [16], which introduces **AgentNet**—a graph neural network (GNN) architecture where a group of relocatable (neural) devices, referred to as neural agents, ‘walk’ across the graph and collaboratively classify graph-level tasks such as detecting triangles, cliques, and cycles. In this model, neural agents can gather information from the nodes they occupy, neighboring nodes they visit, and other co-located devices. Moreover, a recent study [4] demonstrated that a fundamental graph-level task can be solved in the agent-based model using a deterministic algorithm within $O(\Delta \log n)$ rounds, with each device requiring only $O(\Delta \log n)$ bits of storage.

We present a deterministic algorithm for leader election that offers provable guarantees on two key performance metrics in the agent-based model: the solution’s *time complexity* and the storage requirements for each agent. We focus specifically on deterministic algorithms, as they may be more appropriate for relocatable devices. In the agent-based model, storage requirements are treated as a primary performance metric alongside time complexity. The objective in the agent-based model is to minimize storage usage, ideally keeping it as small as the size of a device identifier, typically $O(\log n)$ bits per device. The limited storage prevents relocatable devices from first traversing the graph to learn the topology and then performing graph computation as a separate step. Our goal is to develop an algorithm that is storage-optimal.

Building on our deterministic leader election algorithm, which guarantees both time and storage efficiency, we present improved algorithms for several fun-

Table 1. Summary of Our Results - Leader Election, MST, MIS, MDS and Gathering.

Summarizing Results of This Paper			
Algorithm	Knowledge	Time	Memory (optimal)
Leader Election	n, Δ	$O(n \log^2 n + D\Delta \log n)$	$O(\log n)$
MST	Leader, Dispersed	$O(m + n \log n)$	$O(\Delta \log n)$
MIS	Leader, Dispersed	$O(n \log \Delta)$	$O(\log n)$
MDS	Leader, Dispersed	$O(n \log \Delta)$	$O(\log n)$
Gathering	Leader, Dispersed	$O(n \log \Delta)$	$O(\log n)$

damental distributed graph problems, including minimum spanning tree (MST), maximal independent set (MIS), minimal dominating set (MDS), and gathering. A summary of these results, detailing their time and memory complexities along with the required knowledge, is provided in Table 1. Our leader election algorithm requires prior knowledge of the number of nodes/agents, denoted as n , and the maximum degree of the graph, Δ . We assume that the number of agents equals the number of nodes in the graph, which directly applies to graph problems such as MST, MIS, MDS, and gathering. Furthermore, the results presented in Table 1, aside from the leader election algorithm, assume that the leader is known—that is, each agent is aware of whether they are the leader. Additionally, our leader election process disperses the agents, so we assume each agent is positioned at a distinct node. The aforementioned problem has been studied by Kshemkalyani et al. [11], where no prior knowledge of global parameters was assumed. In contrast, our approach, with prior knowledge of the number of nodes and the maximum degree of the graph, achieves faster and memory-optimal results [1]. A comparative analysis of our results is presented in Table 2.

1.1 Our Results

We examine the leader election problem in a synchronous setting where all agents operate simultaneously, allowing us to measure time complexity in terms of rounds. We assume that the agents have prior knowledge of both the number of nodes n and the maximum degree Δ of the graph. We show the following main results.

1. Leader Election - Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that elects a leader in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits of memory per agent for a known value of n and Δ .

2. Minimum Spanning Tree (MST) - Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MST of G in $O(m + n \log^2 n + D\Delta \log n)$ rounds with $O(\Delta \log n)$ bits per agent for a known value of n and Δ .

3. Maximal Independent Set (MIS) - Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MIS of G in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent for a known value of n and Δ .

4. Minimal Dominating Set (MDS) - Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MDS of G in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent for a known value of n and Δ .

5. Gathering - Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that collects all n agents to a node in G not fixed a priori in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent for a known value of n and Δ .

The rest of the paper is organized as follows.

Paper Organization: Sect. 2 states our distributed computing model. Section 3 discusses the closely related work. Section 4 presents the leader election algorithm, which is the main contribution of the paper. Section 5 explores some graph-related applications of the leader election with improved time/memory complexity. Lastly, Sect. 6 concludes the paper with some interesting future work.

2 Computational Model

We model the network as a connected, undirected graph $G = (V, E)$ with $|V| = n$ nodes and $|E| = m$ edges. Each node $v_i \in V$ has δ_i ports corresponding to each incident edge, labeled $[1, \dots, \delta_i]$. We consider a set $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$ of n agents, initially positioned on the nodes of G ¹. The agents have unique IDs within the range $[1, n^c]$, for some constant, c and agents know the value of c .

Initially, a node may host zero, one, or multiple agents. An agent at a node can communicate with other agents present at the same node but cannot communicate with agents at different nodes (this is the *local communication* model). An agent can move from node v to a neighboring node u via the edge e_{vu} . Following standard message-passing models, such as in [10], we assume that an agent can traverse an edge within a single round, regardless of the edge's weight, even in weighted graphs. When an agent moves from v to u along port p_{vu} , it learns the corresponding port p_{uv} upon arrival at u . Furthermore, at any node v , the agent is aware of the weight $w(e)$ (if G is weighted) of the edge e_{vu} connecting v to its neighbor u . We assume no correlation between the two port numbers of an edge. Multiple agents can traverse an edge simultaneously, as the agent-based model imposes no restrictions on how many agents may move across an edge at the same time.

The agents operate in a synchronous setting, similar to the standard *CONGEST* model: In each round, an agent r_i at node v_i can perform local

¹ Certain graph problems may be solvable with $k < n$ agents, but not all, such as the Minimum Spanning Tree (MST), where computing the MST with $k < n$ agents will yield the MST of a subgraph of G .

Table 2. Summary of previous and our results in the agent-based model. M is the memory required for the Universal Exploration Sequence (UXS) [24] and γ is the number of clusters of agents in the initial configuration. ‘–’ means no prior knowledge of any parameter is required.

Comparative Analysis of the Closely Related Work			
Algorithm	Knowledge	Time	Memory
Leader Election			
Sect. 4	n, Δ	$O(n \log^2 n + D\Delta \log n)$	$O(\log n)$
Kshemkalyani et al. [11]	–	$O(m)$	$O(n \log n)$
MST			
Sect. 5	n, Δ	$O(m + n \log^2 n + D\Delta \log n)$	$O(\Delta \log n)$
Kshemkalyani et al. [11]	–	$O(m + n \log n)$	$O(n \log n)$
MIS			
Sect. 5	n, Δ	$O(n \log^2 n + D\Delta \log n)$	$O(\log n)$
Pattanayak et al. [19]	n, Δ	$O(n\Delta \log n)$	$O(\log n)$
Kshemkalyani et al. [11]	–	$O(n\Delta)$	$O(n \log n)$
MDS			
Sect. 5	n, Δ	$O(n \log^2 n + D\Delta \log n)$	$O(\log n)$
Chand et al. [5]	n, Δ, m, γ	$O(\gamma\Delta \log n + n\gamma + m)$	$O(\log n)$
Kshemkalyani et al. [11]	–	$O(m)$	$O(n \log n)$
Gathering			
Sect. 5	n, Δ	$O(n \log^2 n + D\Delta \log n)$	$O(\log n)$
Molla et al. [17]	n	$O(n^3)$	$O(M + m \log n)$
Kshemkalyani et al. [11]	–	$O(m)$	$O(n \log n)$

computations based on its stored information and the (unique) port labels at its current node, and then decide to either remain at the node or move to a neighboring node via a port. Before leaving, the agent may write information to the storage of another agent at the same node. An agent exiting a node always reaches another node by the end of the round (agents are never positioned on edges at the end of a round).

The agents are assumed to wake up simultaneously at the beginning of the execution. The time complexity is measured in the number of rounds required to reach a solution, while the memory complexity is the number of bits stored by each agent during the execution.

At any round, the agent distribution on G may satisfy one of the following configurations:

- *Dispersed*: each of the n agents occupies a distinct node,
- *Rooted*: all n agents are gathered at a single node,
- *General*: the configuration is neither dispersed nor rooted.

3 Related Work

The leader election problem was first stated by Le Lann [14] in the context of token ring networks, and since then it has been central to the theory of distributed computing. Gallager, Humblet, and Spira [9] provided a deterministic algorithm for any n -node graph G with time complexity $O(n \log n)$ rounds and message complexity $O(m \log n)$. Awerbuch [2] provided a deterministic algorithm with time complexity $O(n)$ and message complexity $O(m + n \log n)$. Peleg [20] provided a deterministic algorithm with optimal time complexity $O(D)$ and message complexity $O(mD)$. Recently, an algorithm is given in [12] with message complexity $O(m)$ but no bound on time complexity, and another algorithm with $O(D \log n)$ time complexity and $O(m \log n)$ message complexity. Additionally, it was shown in [12] the message complexity lower bound of $\Omega(m)$ and time complexity lower bound of $\Omega(D)$ for deterministic leader election in graphs. Leader election was not studied in the agent-based model before, except for [11]. See Table 2.

For MST, the algorithm in Gallager, Humblet, and Spira [9] is the first deterministic algorithm in the message-passing model with time complexity $O(n \log n)$ and message complexity $O(m \log n)$. Time was improved to $O(n)$ in [2] and to $O(\sqrt{n} \log^* n + D)$ in [10, 13]. Furthermore, a time lower bound of $\Omega(\sqrt{n}/\log n + D)$ was given in [21]. MST was not studied in the agent-based model before, except for [11]. See Table 2.

For MIS in the message-passing model, the best-known deterministic distributed algorithms have time complexity $O(2^{\sqrt{\log n}})$ [3, 18]. For MDS, Deurer *et al.* [7] gave two algorithms with an optimal approximation ratio of $(1 + \epsilon)(1 + \log(\Delta + 1))$ running respectively in $O(2^{O(\sqrt{\log(n) \log(\log(n))})})$ and $O(\Delta \text{polylog}(\Delta) + \text{polylog}(\Delta) \log^*(n))$ rounds for $\epsilon > \frac{1}{\text{polylog}(\Delta)}$. MIS and MDS were solved in the agent-based model in [5, 19] with time and memory complexities reported in Table 2 assuming n, Δ (additionally m, γ for MDS) are known to agents a priori. [11] improved these results w.r.t. time/memory complexities as well as lifted the parameter assumptions, see Table 2.

Gathering is a very old problem and has been studied extensively in the agent-based model. The recent results are [17, 24] (detailed literature in [17]). [24] provided a $\tilde{O}(n^5 \log \beta)$ time solution to gather $k \leq n$ agents in arbitrary graphs, where \tilde{O} hides polylog factors and β is the smallest label among agents. Molla *et al.* [17] provided improved time bounds for large values of k assuming n is known but not k : (i) $O(n^3)$ rounds, if $k \geq \lfloor \frac{n}{2} \rfloor + 1$ (ii) $\tilde{O}(n^4)$ rounds, if $\lfloor \frac{n}{2} \rfloor + 1 \leq k < \lfloor \frac{n}{3} \rfloor + 1$, and (iii) $\tilde{O}(n^5)$ rounds, if $\lfloor \frac{n}{3} \rfloor + 1 > k$. Each agent requires $O(M + m \log n)$ bits of memory, where M is the memory required to implement the universal traversal sequence (UXS) [24]. [11] lifted the assumption

of known n and improved time bound for $k = n$ case from $O(n^3)$ to $O(m)$ and additionally the memory bound.

Finally, there is a model, called *whiteboard* [22], related to the agent-based model. The whiteboard model considers (limited) storage at the network nodes in addition to (limited) storage per device. Other aspects remain the same as in the agent-based model, for example, devices have identifiers, computation ability, and communication through relocation. It is immediate that any solution designed in the agent-based model works in the whiteboard model without any change but the opposite may not be true.

4 Leader Election

In this section, we present our deterministic leader election algorithm which, starting from any initial configuration (dispersed, rooted, or arbitrary) of n agents on an n -node graph G , ensures that one agent is elected as a leader. Initially, we disperse the agents on the graph with the help of the algorithm discussed in [23]. After dispersing the agents, we elect the leader. The running time of our algorithm depends on the starting configuration, increasing from the dispersed configuration to the rooted configuration, and finally to the general configuration. Our primary goal is to minimize the round complexity by keeping the optimal memory (in bits) per agent.

4.1 Challenges and High-Level Idea for Leader Election

We are given a set \mathcal{R} of n agents with unique IDs positioned initially arbitrarily on the nodes of an n -node, m -edge graph G such that n and Δ is known to the agents. We also assume that the range $[1, n^c]$ of the IDs is known to the agents for any constant $c > 1$. We aim to elect a unique leader; the pseudocode is given in Algorithm 1. Initially, we disperse all the agents (if not dispersed already) with a well-known algorithm [23]. Then each agent iteratively passes the information of a higher known ID to its neighbors, and finally, the agent with the highest ID agent becomes the leader. The algorithm (after dispersion) uses $O(D\Delta \log n)$ rounds and $O(\log n)$ bits of memory per agent. Our leader election following dispersion presents three key challenges to the aforementioned complexity: (i) The graph diameter is unknown, making it difficult to terminate the algorithm. (ii) Since agents are mobile entities, therefore, meeting neighboring agents is challenging. (iii) The algorithm assumes each agent has only $O(\log n)$ bits of memory, making it difficult to track all the information passed by neighboring agents.

On a high level, we deal with all the above challenges in the following ways (in order): (i) Based on a voting mechanism, our algorithm ensures that the agents have passed the information (about the highest ID) from one end of the graph to another. (ii) Based on ID, our algorithm makes sure that each neighbor meets all of its neighbors in a definite interval of time/rounds. (iii) Based on the child-parent relationship, our algorithm makes sure only each child keeps track of its parent rather than the parent keeps track of all its children which requires $\Delta \log n$ memory.

4.2 Detailed Description of the Algorithm

We disperse the agents in $O(n \log^2 n)$ rounds² with the help of the dispersion algorithm for the general configuration, given in [23]. Their algorithm does not have the termination time, therefore, with the help of a known parameter of n all the agents get to know the dispersion time. After dispersing all the agents, the algorithm operates in phases of $2\Delta \log n$ rounds each. In a phase, each agent meets all its neighbors in $O(\Delta \log n)$ time. Each agent tries to figure out the

Algorithm 1: Leader Election

Input : A set \mathcal{R} of n agents with unique IDs positioned initially arbitrarily on the nodes of an n -node, m -edge graph G such that n and Δ is known to the agents.

Output: An agent in \mathcal{R} is elected as a leader with status *leader*.

- 1 Disperse the agents in $O(n \log^2 n)$ rounds with the help of an arbitrary dispersion algorithm in [23]. ▷ Takes $O(n \log^2 n)$ rounds.
 - 2 A phase consists of $O(\Delta \log n)$ rounds, in which agent r_u meets all its neighbors, follows from Lemma 1.
 - 3 Each agent r_u does nothing for a phase of $c\Delta \log n$ rounds.
 - 4 **while** agent r_u is unaware about the leader **do**
 - 5 **if** agent r_u has not sent its highest ID known from the previous phase **then**
 - 6 Agent r_u sends its highest known ID to all the neighbors
 - 7 **if** r_u receives higher ID from neighbour r_v **then**
 - 8 r_u considers r_v as its parent and becomes the child of r_v
 - 9 **if** r_v remains parent of r_u for next two phases **then**
 - 10 r_u sends “undecided” vote to r_v for next two phases
 // Initialization of “undecided” vote.
 - 11 **if** r_u received “yes” vote and did not receive the vote “undecided” from all its children (if any), in the previous phase, after two phases of proposal **then**
 - 12 r_u sends the vote “yes” to its parent (if any)
 - 13 **else**
 - 14 r_u sends the vote “undecided” to its parent (if any)
 - 15 **if** r_u does not have child **then**
 - 16 r_u keeps sending a vote “yes” to its parent until it gets to know about the leader // r_u did not receive a vote from any neighbor.
 - 17 **if** r_u received its own proposed ID from all its children with a “yes” vote in the previous phase **then**
 - 18 r_u becomes the leader and informs its ID to all neighbors
 - 19 **if** r_u knows the leader ID from the previous phase **then**
 - 20 r_u informs the leader ID to all its neighbors
-

² Our approach uses the algorithm as it is and considers the constant as in [23] throughout the paper.

highest ID agent—in phase k the highest ID in a k -hop neighborhood is identified. The highest ID agent eventually becomes the leader.

Meeting with Neighboring Agents: The approach is based on the ID of the agent, whose size is $O(\log n)$, i.e., $c \log n$ since the highest ID is n^c (n is known parameter for the Algorithm 1) and the known value of Δ . Each agent r_u keeps the ID length $c \log n$; if ID is shortened by x bits then x number of bits with value 0 is prepended as MSB (Most Significant Bit). Starting from the LSB (Least Significant Bit) of an agent r_u 's ID, each agent r_u decides whether it would visit its neighbor or not. If LSB is 1 then the agent r_u visits its neighboring nodes one by one based on port numbering (in increasing order) which takes $2\delta_{r_u}$ rounds (two rounds to visit a neighboring node—back and forth) to explore all the neighbors, where δ_{r_u} is the degree for the agent r_u . If LSB is 0 for an agent r_u then agent r_u waits for 2Δ rounds at its node. After 2Δ rounds, r_u decides based on the next bit in its ID whether it would explore its neighboring nodes or would stay at its node. Notice that if $\delta_{r_u} < \Delta$ then r_u completes visiting its neighboring agents in less than 2Δ rounds for bit value 1 in its ID. Therefore, r_u stays at its node for the remaining rounds $(2\Delta - 2\delta_{r_u})$. Lemma 1 shows that agent r_u meets with all its neighbors at least once in $O(\Delta \log n)$ rounds.

Election of the Highest ID Agent as Leader: After dispersing every agent such that each node possesses only one agent, each agent r_u does nothing for a phase of $O(\Delta \log n)$ rounds. We consider *Meeting with Neighboring Agents* as a phase that takes $2\Delta \log n$ rounds. In this phase, we keep our algorithm deliberately ideal to simplify the flow of the algorithm, i.e., to pass the updated information in the next phase that was received in the prior phase.

The following statements and conditions are executed by agent r_u until r_u gets to know the leader and informs all its neighbors:

In the very first phase, each agent already has its own ID which would be passed during the next phase. Each agent r_u passes its ID to neighboring agents if that is the highest ID known to r_u . Therefore, if agent r_u has not sent (earlier) its highest ID known from the prior phase then agent r_u sends its highest known ID to all the neighbors.

If r_u receives a higher ID from neighbor r_v then r_u considers r_v as its parent and becomes the child of r_v . This is required since due to memory limitation r_u can not keep track of all its neighbors to whom it passes the higher ID. Therefore, each receiving agent (child) keeps the track of sender (parent) which is not more than 1 in numbers. In case of a tie—(i) more than one agent sends the same ID which is not less than the earlier known ID, an agent r_u gives priority to the agent that was the parent in the prior phase. Notice that this condition is more than breaking a tie since in another case—giving priority to the latest sending agent, an agent might keep changing its parent forever. (ii) If there is more than one agent that proposes the same ID in the same phase then priority (to be considered as parent) is given to the higher ID proposer. Additionally, if r_v remains the parent of r_u for the next two phases then r_u sends an “undecided” vote to r_v for the next two phases. In this way, the initialization of the “undecided” vote takes place. The vote “undecided” is sent to the parent

by its child if it received the “undecided” vote from any of its children or r_v becomes the parent two phases before. In these two phases, r_u ’s children (if any) receive the ID and send their vote. Therefore, r_u sends an “undecided” vote in the first two phases irrespective of its children’s vote. From the next phase onwards, r_u receives its children’s vote and considers them for voting to its parent.

In case, if r_u received “yes” vote from all its children (if any) and did not receive the vote “undecided”, in the previous phase, after two phases of the proposal then r_u sends the vote “yes” to its parent (if any). An agent r_u considers a neighbor its child only if it sends the vote either “yes” or “undecided”. An agent r_u sends a “yes” vote to its parent only if it has received a “yes” vote from all its children (if any) and it has not received any “undecided” vote from its neighbors. Otherwise, r_u sends the vote “undecided” to its parent (if any). Notice that if agent r_u does not possess any child, in that case, r_u does not receive any “yes” or “undecided” vote. This implies that r_u does not have any child and r_u initiates the vote “yes” to its parent after two phases of passing the highest ID known to its neighbors. A child always responds with a vote of “yes” or “undecided”. If r_u does not have a child then r_u keeps sending a vote “yes” to its parent until it gets to know about the leader.

Whenever r_u receives its own proposed ID from all its children with a “yes” vote in the prior phase then r_u becomes the leader and informs its ID to all neighbors. If r_u knows the leader ID from the prior phase, in that case, r_u informs the leader ID to all its neighbors.

Now, we present the above claims formally. In Lemma 1, we show the number of rounds required to meet any two neighbors. Lemma 2 shows the time and memory complexity required by Algorithm 1. Further, Lemma 3 shows the correctness of the Algorithm 1 such that it elects a unique leader.

Lemma 1. *After dispersion, every agent requires at most $O(\Delta \log n)$ rounds to meet all its neighbors.*

Proof. To prove our lemma, we prove the results for agent r_u and any neighboring agent r_v . Each agent r_u has a unique ID with length $O(\log n)$ bits. Therefore, there exists at least one index between r_u and any neighboring agent r_v such that their bits are different. Furthermore, each agent explores its neighboring nodes in the interval of 2Δ rounds based on the next bit (starting from the LSB). This implies during the different bit for a certain index either r_u would visit r_v or vice versa. Henceforth, each r_u meets all its neighbors in $O(\Delta \log n)$ rounds. Hence, the lemma. \square

Lemma 2. *Algorithm 1 takes $O(n \log^2 n + D\Delta \log n)$ rounds and $O(\log n)$ memory to elect the leader.*

Proof. For round complexity, Line 1 takes $O(n \log^2 n)$ rounds to disperse all the n agents [23]. Furthermore, Lemma 1 takes $O(\Delta \log n)$ rounds to meet its neighbor. Therefore, to pass the information one hop away takes $O(\Delta \log n)$ rounds. The farthest agent is D hops away from any other agent which takes overall $O(D\Delta \log n)$ rounds to know the highest ID agent in the graph. Then it

takes $O(D\Delta \log n)$ rounds for the “yes” votes to propagate back to the highest ID agent. After that, it takes $O(D\Delta \log n)$ rounds required to inform all the agents in the graph about the highest ID agent, i.e., the leader. This shows the overall round complexity for leader election is $O(n \log^2 n + D\Delta \log n)$. In case of memory complexity, Line 1 takes $O(\log n)$ memory [23]. On the other hand, after dispersion, knowing the highest ID also takes $O(\log n)$ bits to store the information of ID, parent, highest ID known, and Leader ID. Hence, the lemma. \square

Lemma 3. *Algorithm 1 correctly elects a unique agent as the leader.*

Proof. For the sake of contradiction, let us suppose Algorithm 1 elects more than one leader, namely, ℓ_1 and ℓ_2 such that $\ell_1 > \ell_2$, i.e., ℓ_1 ’s ID is higher than ℓ_2 . In such case, ℓ_2 ’s children voted “yes” for ℓ_2 and so did their children and their children’s children, and so on. Therefore, ℓ_1 which is an agent situated at the graph also voted for ℓ_2 which is not possible. Hence, there exists a unique leader with the highest ID. \square

From the above discussion, we have the following result.

Theorem 1. *Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that elects a leader in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits of memory per agent for a known value of n and Δ .*

Remark 1. For the rooted configuration, dispersion takes $O(n \log n)$ rounds in Line 1 of Algorithm 1 as shown in [23]. Therefore, round complexity becomes $O(n \log n + D\Delta \log n)$. On the other hand, for the case of already dispersed configuration, round complexity is $O(D\Delta \log n)$. Furthermore, memory complexity remains unchanged in both configurations, i.e., $O(\log n)$ bits per agent.

5 Applications to Additional Graph Problems

We leverage the leader election result (Theorem 1) to enhance existing solutions for MST, MIS, MDS, and Gathering problems in the agent-based model. The improvements over previous results (discussed in Table 2) are achieved either by optimizing time or memory. All our solutions are memory-optimal, with most offering significant reductions in round complexity as well.

5.1 Minimum Spanning Tree

A **Minimum Spanning Tree (MST)** of a graph $G = (V, E)$ with n nodes is a subgraph $T = (V, E')$ such that $w(e)$ represents the weight of an edge e and for the MST T :

$$T = \arg \min_{T' \subseteq G} \sum_{e \in T'} w(e)$$

subject to the constraints that T' is connected, acyclic, and spans all n nodes in G .

Kshemkalyani et al. [11] constructed the MST in $O(m + n \log n)$ rounds and $O(\Delta \log n)$ bits per agent from a given dispersed configuration in the presence of a leader. From Sudo et al. [23], we can disperse all the n agents in $O(n \log^2 n)$ rounds and $O(\log n)$ memory. Their protocol does not have termination [23], therefore, knowledge of n provides the dispersed configuration in $O(n \log^2 n)$ rounds. From the above discussion and Theorem 1, we have the following result.

Theorem 2 (MST). *Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MST of G in $O(m + n \log^2 n + D \Delta \log n)$ rounds with $O(\Delta \log n)$ bits per agent.*

Remark 2. In an MST, an agent might have to track all its neighboring edges as the MST edge that requires $O(\Delta \log n)$ bits of memory. Therefore, Theorem 2 is memory optimal.

5.2 Maximal Independent Set (MIS)

Suppose n agents are initially located arbitrarily on the nodes of an n -node anonymous graph $G = (V, E)$. The goal in the maximal independent set (MIS) problem is to relocate the agents autonomously to find a subset $S \subset V$ of nodes such that S forms an MIS of G .

Description of the MIS Algorithm: We construct the MIS from a dispersed configuration and known leader in Algorithm 2. In this, the leader becomes the MIS agent and converts all its neighbors non-MIS in $O(\log \Delta)$ rounds with the help of the Helping-Each-Other (HEO) technique introduced by [23]. In this, each agent r_u having the token visits its neighbor by port 1 and comes back with the agent situated across port 1. Now, these two agents visit ports 2 and 3 individually and come back with one agent each. Similarly, after every 2 rounds the number of agents becomes double, and for the degree δ_{r_u} , it takes $2 \log \delta_{r_u}$ rounds to gather all the neighboring agents at r_u 's node. During the MIS algorithm, we run a variant of Depth-First-Search (DFS) with the help of a token. We run the MIS algorithm until some neighbor of agent r_u remains unexplored. We call an agent r_u as unexplored if it never held the token. Specifically, if the leader's neighbor remains unexplored, we run the following protocol for any agent r_u . If agent r_u has the token and it is neither an MIS nor a non-MIS agent, in that case, r_u collects all the agents using HEO technique and figures out the status of all other agents. If none of the neighboring agents is an MIS agent only then r_u becomes the MIS agent. Otherwise, r_u becomes a non-MIS agent. If agent r_u has the token and it is either an MIS or a non-MIS agent and there exists any unexplored port then r_u passes the token to the next unexplored port (in increasing order) and becomes the parent of the node across the port. Otherwise, r_u sends the token to the parent node. In this protocol, each agent stores its own ID and parent's port number. It requires $O(\log n)$ memory. Furthermore, the

token reaches all the n nodes and takes at most $O(\log \Delta)$ rounds to run HEO. Therefore, the overall time complexity to run the MIS algorithm is $O(n \log \Delta)$ rounds and $O(\log n)$ memory.

Below, we show the correctness of Algorithm 2.

Lemma 4. *Algorithm 2 generates an MIS.*

Proof. We prove the lemma by considering two cases: (i) The token is passed to each agent r_u . (ii) Either the agent r_u is part of the MIS or one of its neighbors is.

We prove both cases by contradiction. For case (i), let us assume the token is not passed to agent r_u . Then, it is also not passed to r_u 's parent, grandparent, and so forth, due to Line 13 of Algorithm 2. This implies that the token was not passed by the leader, contradicting Line 2. For case (ii), given that the token reaches all agents, suppose there exists an agent at a node that is neither part of the MIS nor has any neighbors in the MIS. This would mean that Line 8 in Algorithm 2 was not executed, contradicting the assumption. \square

From the above discussion and Theorem 1 of leader election, we have the following result.

Algorithm 2: MAXIMAL INDEPENDENT SET

Input : A set \mathcal{R} of n agents with unique IDs dispersed on a graph of n nodes with known leader's ID.

Output: MIS configuration.

- 1 Helping-Each-Other (HEO) technique to gather all the neighboring agents at a node as discussed in [23]. A HEO phase consists of $O(\log \Delta)$ rounds, in which, agent r_u meets all its neighbors.
 - 2 The leader becomes the MIS agent and converts all its neighbors in non-MIS in $\log \Delta$ rounds with the help of HEO, passes the token via port 1, and becomes the parent of the across-the-port node.
 - 3 **while** agent r_u 's neighbors are unexplored **do**
 - 4 **if** agent r_u has the token **then**
 - 5 **if** r_u is neither an MIS nor a non-MIS agent **then**
 - 6 r_u collects all the agents using HEO technique and figures out the status of all other agents
 - 7 **if** none of the neighboring agents is an MIS agent **then**
 - 8 r_u becomes the MIS agent
 - 9 **else**
 - 10 r_u becomes a non-MIS agent
 - 11 **if** r_u is either MIS or Non-MIS **then**
 - 12 **if** any unexplored port exists **then**
 - 13 r_u passes the token to the next unexplored port (in increasing order) and becomes the parent of the node across the port
 - 14 **else**
 - 15 r_u sends the token to the parent node
-

Theorem 3 (MIS). *Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MIS of G in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent.*

5.3 Minimal Dominating Set (MDS)

Suppose n agents are initially located arbitrarily on the nodes of an n -node anonymous graph $G = (V, E)$. A *dominating set* of G is a subset $DS \subset V$ of nodes such that for any $v \notin DS$, v has a neighbor in DS .

For the sake of completeness, we prove the following lemma to show every MIS is an MDS. Therefore, the result of MIS holds for MDS as it is.

Lemma 5. *Every MIS is an MDS.*

Proof. We prove this by contradiction. Let us suppose there exists a neighbor in MIS that is not covered by any of the MIS agents then there exists an agent r_u which is neither MIS nor its neighbor which contradicts the assumption that it is an MIS. Hence, the proof. \square

From the Lemma 5 and Theorem 3, we have the following result.

Theorem 4 (MDS). *Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that finds an MDS of G in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent.*

5.4 Gathering

Suppose n agents are initially located arbitrarily on the nodes of an n -node anonymous graph $G = (V, E)$. The goal of the gathering problem is to relocate the agents autonomously to position all of them to a node in G not fixed a priori (n agents at a node of G).

Algorithm 2 can be adapted to gather all the agents at the leader's node. In Algorithm 2 (Line 15), when an agent passes the token to its parent, the parent agent keeps the account of the number of children. Therefore, after $O(n \log \Delta)$ rounds all the agents are traversed, and MIS is formed. The agent without children gathers at its parent node. When the parent finds all its children in its place, the parent moves to its parent with all of its children. This process occurs recursively and eventually all the agents are gathered at the leader's node. In the worst case, an agent can be at most n hop away from its leader. Therefore, it takes at most $O(n)$ rounds to gather all the agents. Notice that an agent only counts its number of children other than the memory used for MIS, therefore, the memory complexity remains unchanged, i.e., $O(\log n)$.

From the above discussion and Leader Election Algorithm 1, we have the following result.

Theorem 5 (Gathering). *Given any configuration of n agents positioned initially arbitrarily on the nodes of a n -node graph G , there is a deterministic algorithm that collects all n agents to a node in G not fixed a priori in $O(n \log^2 n + D\Delta \log n)$ rounds with $O(\log n)$ bits per agent.*

6 Conclusions

We solved leader election in the agent-based model in $O(n \log^2 n + D\Delta \log n)$ time with optimal memory of $O(\log n)$ bits per agent for general initial configurations. Furthermore, the memory complexities of the existing results on MST, MIS, MDS, and Gathering on the agent-based model were improved using the leader election result, and we achieved the optimal memory complexity bits per agent for all the aforementioned problems. Our result also solved near-linear time complexity for MIS, MDS, and Gathering problems after leader election in a dispersed setting in $O(n \log \Delta)$ rounds (see Table 1). It would be interesting to improve the time complexity of our leader election, which would directly help to improve the complexity of the other graph-related problems. It would also be interesting to see whether these results can be achieved without any prior knowledge of the parameter.

References

1. Augustine, J., Moses Jr, W.K.: Dispersion of mobile robots: a study of memory-time trade-offs. In: ICDCN, pp. 1:1–1:10 (2018)
2. Awerbuch, B.: Optimal distributed algorithms for minimum weight spanning tree, counting, leader election and related problems (detailed summary). In: Aho, A.V. (ed.) STOC, pp. 230–240. ACM (1987). <https://doi.org/10.1145/28395.28421>
3. Awerbuch, B., Goldberg, A.V., Luby, M., Plotkin, S.A.: Network decomposition and locality in distributed computation. In: FOCS, vol. 30, pp. 364–369. Citeseer (1989)
4. Chand, P.K., Das, A., Molla, A.R.: Agent-based triangle counting and its applications in anonymous graphs. CoRR abs/2402.03653 (2024). <https://doi.org/10.48550/arxiv.2402.03653>
5. Chand, P.K., Molla, A.R., Sivasubramaniam, S.: Run for cover: dominating set via mobile agents. In: ALGOWIN, pp. 133–150. Springer, Cham (2023)
6. Cong, Y., Changjun, G., Zhang, T., Gao, Y.: Underwater robot sensing technology: a survey. Fundam. Res. **1** (2021). <https://doi.org/10.1016/j.fmre.2021.03.002>
7. Deurer, J., Kuhn, F., Maus, Y.: Deterministic distributed dominating set approximation in the congest model. In: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, pp. 94–103 (2019)
8. El-Sayed, A., Scarborough, P., Seemann, L., Galea, S.: Social network analysis and agent-based modeling in social epidemiology. Epidemiol. perspect. Innov. EP+I **9**, 1 (2012). <https://doi.org/10.1186/1742-5573-9-1>
9. Gallager, R.G., Humblet, P.A., Spira, P.M.: A distributed algorithm for minimum-weight spanning trees. ACM Trans. Program. Lang. Syst. **5**(1), 66–77 (1983). <https://doi.org/10.1145/357195.357200>

10. Garay, J.A., Kutten, S., Peleg, D.: A sub-linear time distributed algorithm for minimum-weight spanning trees (extended abstract). In: FOCS, pp. 659–668. IEEE Computer Society (1993). <https://doi.org/10.1109/SFCS.1993.366821>
11. Kshemkalyani, A.D., Kumar, M., Molla, A.R., Sharma, G.: Brief announcement: agent-based leader election, MST, and beyond. In: DISC 2024. LIPIcs, vol. 319, pp. 50:1–50:7 (2024). <https://doi.org/10.4230/LIPICS.DISC.2024.50>
12. Kutten, S., Pandurangan, G., Peleg, D., Robinson, P., Trehan, A.: On the complexity of universal leader election. *J. ACM* **62**(1), 7:1–7:27 (2015)
13. Kutten, S., Peleg, D.: Fast distributed construction of small k -dominating sets and applications. *J. Algorithms* **28**(1), 40–66 (1998)
14. Lann, G.L.: Distributed systems - towards a formal approach. In: Gilchrist, B. (ed.) *Information Processing, Proceedings of the 7th IFIP Congress 1977*, Toronto, Canada, 8–12 August 1977, pp. 155–160. North-Holland (1977)
15. Lee, J., Shin, S., Park, M., Kim, C.: Agent-based simulation and its application to analyze combat effectiveness in network-centric warfare considering communication failure environments. *Math. Probl. Eng.* **2018**, 1–9 (2018). <https://doi.org/10.1155/2018/2730671>
16. Martinkus, K., Papp, P.A., Schesch, B., Wattenhofer, R.: Agent-based graph neural networks. In: *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, 1–5 May 2023*. OpenReview.net (2023). <https://openreview.net/pdf?id=8WTAh0tj2jC>
17. Molla, A.R., Mondal, K., Moses Jr, W.K.: Fast deterministic gathering with detection on arbitrary graphs: the power of many robots. In: *IEEE International Parallel and Distributed Processing Symposium, IPDPS*, pp. 47–57. IEEE (2023)
18. Panconesi, A., Srinivasan, A.: On the complexity of distributed network decomposition. *J. Algorithms* **20**(2), 356–374 (1996)
19. Pattanayak, D., Bhagat, S., Chaudhuri, S.G., Molla, A.R.: Maximal independent set via mobile agents. In: *ICDCN*, pp. 74–83. ACM (2024)
20. Peleg, D.: Time-optimal leader election in general networks. *J. Parallel Distrib. Comput.* **8**(1), 96–99 (1990). [https://doi.org/10.1016/0743-7315\(90\)90074-Y](https://doi.org/10.1016/0743-7315(90)90074-Y)
21. Peleg, D., Rubinfeld, V.: A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. *SIAM J. Comput.* **30**(5), 1427–1442 (2000). <https://doi.org/10.1137/S0097539700369740>
22. Sudo, Y., Baba, D., Nakamura, J., Ooshita, F., Kakugawa, H., Masuzawa, T.: A single agent exploration in unknown undirected graphs with whiteboards. In: *Proceedings of the 3rd International ACM Workshop on Reliability, Availability, and Security, WRAS 2010*, vol. E98.A (2010). <https://doi.org/10.1145/1953563.1953570>
23. Sudo, Y., Shibata, M., Nakamura, J., Kim, Y., Masuzawa, T.: Near-linear time dispersion of mobile agents. In: Alistarh, D. (ed.) *38th International Symposium on Distributed Computing, DISC 2024, October 28 to November 1, 2024, Madrid, Spain*. LIPIcs, vol. 319, pp. 38:1–38:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2024). <https://doi.org/10.4230/LIPICS.DISC.2024.38>
24. Ta-Shma, A., Zwick, U.: Deterministic rendezvous, treasure hunts, and strongly universal exploration sequences. *ACM Trans. Algorithms* **10**(3), 12:1–12:15 (2014). <https://doi.org/10.1145/2601068>
25. Zhuge, C., Shao, C., Wei, B.: An agent-based spatial urban social network generator: a case study of Beijing, China. *J. Comput. Sci.* **29** (2018). <https://doi.org/10.1016/j.jocs.2018.09.005>