

The Complexity of Leader Election: A Chasm at Diameter Two

Soumyottam Chatterjee
University of Houston
Department of Computer Science
Houston, TX, USA
schatterjee4@uh.edu

Gopal Pandurangan*
University of Houston
Department of Computer Science
Houston, TX, USA
gopalpandurangan@gmail.com

Peter Robinson
Royal Holloway, University of London
Department of Computer Science
Egham, Surrey, UK
peter.robinson@rhul.ac.uk

ABSTRACT

Leader election is one of the fundamental problems in distributed computing. In its *implicit* version, only the leader must know who is the elected leader. This paper focuses on studying the message complexity of leader election in synchronous distributed networks, in particular, in networks of *diameter two*. Kutten et al. [JACM 2015] showed a fundamental lower bound of $\Omega(m)$ (m is the number of edges in the network) on the message complexity of (implicit) leader election that applied also to Monte Carlo randomized algorithms with constant success probability; this lower bound applies for graphs that have *diameter at least three*. On the other hand, for complete graphs (i.e., diameter 1), Kutten et al. [TCS 2015] established a tight bound of $\tilde{O}(\sqrt{n})^1$ on the message complexity of randomized leader election (n is the number of nodes in the network). For graphs of diameter two, the complexity was not known.

In this paper, we settle this complexity by showing a tight bound of $\tilde{O}(n)$ on the message complexity of leader election in diameter-two networks. We first give a simple *randomized* Monte-Carlo leader election algorithm that with high probability (i.e., probability at least $1 - n^{-c}$, for some positive constant c) succeeds and uses $O(n \log^3 n)$ messages and runs in $O(1)$ rounds; this algorithm works without knowledge of n (and hence needs no global knowledge). We then show that any algorithm (even Monte Carlo randomized algorithms with large enough constant success probability) needs $\Omega(n)$ messages (even when n is known), regardless of the number of rounds. We also present an $O(n \log n)$ messages *deterministic* algorithm that takes $O(\log n)$ rounds (but needs knowledge of n); we show that this message complexity is tight for deterministic algorithms.

Our results show that leader election can be solved in diameter-two graphs in (essentially) linear (in n) message complexity and thus the $\Omega(m)$ lower bound does not apply to diameter-two graphs. Together with the two previous results of Kutten et al., our results fully characterize the message complexity of leader election vis-à-vis the graph diameter.

*Supported, in part, by NSF grants CCF-1527867, CCF-1540512, IIS-1633720, and CCF-1717075.

¹Notation \tilde{O} hides a $\frac{1}{\text{polylog}(n)}$ factor; \tilde{O} and $\tilde{\Theta}$ hide a $\text{polylog}(n)$ factor.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICDCN '18, January 4–7, 2018, Varanasi, India
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-6372-3/18/01...\$15.00
<https://doi.org/10.1145/3154273.3154308>

KEYWORDS

Distributed Algorithm, Leader Election, Randomized Algorithm, Message Complexity, Time Complexity, Lower Bounds

ACM Reference Format:

Soumyottam Chatterjee, Gopal Pandurangan, and Peter Robinson. 2018. The Complexity of Leader Election: A Chasm at Diameter Two. In *ICDCN '18: 19th International Conference on Distributed Computing and Networking, January 4–7, 2018, Varanasi, India*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3154273.3154308>

1 INTRODUCTION

Leader election is a classical and fundamental problem in distributed computing. The leader election problem requires a group of processors in a distributed network to elect a unique leader among themselves, i.e., exactly one processor must output the decision that it is the leader, say, by changing a special *status* component of its state to the value *leader* [13]. All the rest of the nodes must change their status component to the value *non-leader*. These nodes need not be aware of the identity of the leader. This *implicit* variant of leader election is quite standard (cf. [13]), and has been extensively studied (see e.g., [10] and the references therein) and is sufficient in many applications, e.g., for token generation in a token ring environment [12]. In this paper, we focus on this implicit variant.²

The complexity of leader election, in particular, its message and time complexity, has been extensively studied both in general graphs as well as in special graph classes such as rings and complete networks, see e.g., [10, 11, 13, 16, 18, 19]. While much of the earlier work focused on deterministic algorithms, recent works have studied randomized algorithms (see e.g., [10, 11] and the references therein). Kutten et al. [10] showed a fundamental lower bound of $\Omega(m)$ (m is the number of edges in the network) on the message complexity of (implicit) leader election that applied even to Monte Carlo randomized algorithms with (large-enough) constant success probability; this lower bound applies for graphs that have *diameter at least three*. We point that the $\Omega(m)$ lower bound applies even for algorithms that have knowledge of n, m, D (throughout, n denotes the number of nodes, m the number of edges, and D the network diameter). The lower bound proof involves constructing a “dumb-bell” graph G which consists of two regular subgraphs G_1 and G_2 (each having approximately $\frac{m}{2}$ edges) joined by a couple of “bridge” edges (the bridge edges are added so that the regularity is preserved). Note that (even) if G_1 and G_2 are cliques (in particular, they can be any

²In another variant, called *explicit* leader election, all the non-leaders change their status component to the value *non-leader*, and moreover, every node must also know the identity of the unique leader. In this variant, $\Omega(n)$ messages is an obvious lower bound (throughout, n denotes the number of nodes in the network) since every node must be informed of the leader’s identity. Clearly, any lower bound for implicit leader election applies to explicit leader election as well.

2-connected graph) then G will be of diameter (at least) three. This is the smallest diameter that makes the lower bound proof work; we refer to [10] for details.

On the other hand, for complete graphs (i.e., diameter one), Kutten et al. [11] established a tight bound of $\tilde{\Theta}(\sqrt{n})$ on the message complexity of randomized leader election (n is the number of nodes in the network). In other words, they showed an $\tilde{O}(\sqrt{n})$ messages algorithm that elects a (unique) leader with high probability. To complement this, they also showed that any leader election algorithm in a complete graph requires $\tilde{\Omega}(\sqrt{n})$ messages to succeed with (large-enough) constant probability.

For graphs of diameter two, the message complexity was not known. In this paper, we settle this complexity by showing a tight bound of $\tilde{\Theta}(n)$ on the message complexity of leader election in diameter-two networks. In particular, we present a simple randomized leader election algorithm that takes $O(n \log^3 n)$ messages and $O(1)$ rounds that works *even when n is not known*. In contrast, we show that any randomized algorithm (even Monte Carlo algorithms with constant success probability) needs $\Omega(n)$ messages. Our results show that leader election can be solved in diameter-two graphs in (essentially) linear (in n) message complexity which is optimal (up to a polylog(n) factor) and thus the $\Omega(m)$ message lower bound does not apply to diameter-two graphs. Together with the previous results [10, 11], our results fully characterize the message complexity of leader election vis-à-vis the graph diameter (see Table 1).

1.1 Our Results

This paper focuses on studying the message complexity of leader election (both randomized and deterministic) in synchronous distributed networks, in particular, in networks of *diameter two*.

For our algorithms, we assume that the communication is *synchronous* and follows the standard *CONGEST* model [17], where a node can send in each round at most one message of size $O(\log n)$ bits on a single edge. We assume that the nodes have unique IDs. We assume that all nodes wake up simultaneously at the beginning of the execution. (Additional details on our distributed computation model are given in Section 1.3.)

We show the following results:

(1) Algorithms: We show that the message complexity of leader election in diameter-two graphs is $\tilde{O}(n)$, by presenting a randomized (implicit) leader election algorithm (cf. Section 2), that takes $O(n \log^3 n)$ messages and runs in $O(1)$ rounds with high probability (whp).³ This algorithm works even without knowledge of n . While it is easy to design an $O(n \log n)$ messages randomized algorithm with knowledge of n (Section 1.2), not having knowledge of n makes the analysis more involved. We also present a *deterministic* algorithm that uses only $O(n \log n)$ messages, but it takes $O(\log n)$ rounds. Also this algorithm needs knowledge (or a constant factor upper bound) of n (or $\log n$) (cf. Section 4). It is not difficult to convert this algorithm (under the same bounds) to solve explicit leader election, where the identity of the leader is broadcast to all nodes. Thus *broadcast*, another fundamental problem, can be solved in diameter-two graphs in $O(n \log n)$ messages and $O(\log n)$ rounds if n is known (in contrast we note that $\Omega(m)$ is a lower bound for broadcast on

graphs of diameter at least three, even if n is known and even for randomized algorithms [10]). (We note that all our algorithms will work seamlessly for complete networks as well.)

(2) Lower Bounds: We show that, in general, it is not possible to improve over our algorithm substantially, by presenting a lower bound for leader election that applies also to randomized (Monte Carlo) algorithms. We show that $\Omega(n)$ messages are needed for any leader election algorithm (regardless of the number of rounds) in a diameter-two network which succeeds with any constant probability that is strictly larger than $\frac{1}{2}$ (cf. Section 3). This lower bound holds even in the *LOCAL* model [17], where there is no restriction on the number of bits that can be sent on each edge in each round. To the best of our knowledge, this is the first non-trivial lower bound for randomized leader election in diameter-two networks.

We also show a simple deterministic reduction that shows that any super-linear message lower bound for complete networks also applies to diameter-two networks as well (cf. Section 5). It can be shown that $\Omega(n \log n)$ messages is a lower bound for deterministic leader election in complete networks [1, 9] (under assumption that the number of rounds is bounded by some function of n)⁴ and by our reduction this lower bound also applies for diameter-two networks. (We point out that lower bounds for complete networks do not directly translate to diameter-two networks.)

Diameter	RANDOMIZED		DETERMINISTIC	
	Time	Messages	Time	Messages
$D = 1$: [1, 11]				
Upper Bound	$O(1)$	$O(\sqrt{n} \log^{\frac{3}{2}} n)$	$O(1)^\dagger$	$O(n \log n)^\ddagger$
Lower Bound	$\Omega(1)$	$\Omega(\sqrt{n})$	$\Omega(1)$	$\Omega(n \log n)$
$D = 2$: Our Results				
Upper Bound	$O(1)$	$O(n \log^3 n)$	$O(\log n)^\ddagger$	$O(n \log n)^\S$
Lower Bound	$\Omega(1)$	$\Omega(n)$	$\Omega(1)$	$\Omega(n \log n)$
$D \geq 3$: [10]				
Upper Bound	$O(D)$	$O(m \log \log n)$	$O(D \log n)$	$O(m \log n)$
Lower Bound	$\Omega(D)$	$\Omega(m)$	$\Omega(D)$	$\Omega(m)$

[†] Note that attaining $O(1)$ time requires $\Omega(n^{1+\Omega(1)})$ messages in cliques, whereas achieving $O(n \log n)$ messages requires $\Omega(\log n)$ rounds; see [1].

[§] Needs knowledge of n .

^{‡†} Note that it is easy to show a $O(1)$ round deterministic algorithm that takes $O(m)$ messages.

Table 1: Message and time complexity of leader election.

1.2 Technical Overview

All our algorithms exploit the following simple “neighborhood intersection” property of diameter-two graphs: Any two nodes (that are non-neighbors) have at least one neighbor in common (please refer to Observation 1). However, note that unlike complete networks (which have been extensively studied with respect to leader election — cf. Section 1.4), in diameter-two networks, nodes generally don’t have knowledge of n , the network size (in a complete graph, this is

⁴Afek and Gafni[1] show the $\Omega(n \log n)$ message lower bound for complete networks under the non-simultaneous wakeup model in synchronous networks. The same message bound can be shown to hold in the simultaneous wake-up model as well under the restriction that the number of rounds is bounded by a function of n [9].

³Throughout, “with high probability” means with probability at least $1 - 1/n^c$, for some constant $c > 0$.

trivially known by the degree). This complicates obtaining sublinear in m (where m is the number of edges) message algorithms that are fully localized (don't have knowledge of n). Indeed, if n is known, the following is a simple randomized algorithm: each node becomes a candidate with probability $\Theta(\frac{\log n}{n})$ and sends its ID to all its neighbors; any node that gets one or more messages acts as a "referee" and notifies the candidate that has the smallest ID (among those it has received). The neighborhood intersection property implies that at least one candidate will be chosen uniquely as the leader with high probability.

If n is not known, the above idea does not work. However, we show that if each node v becomes a candidate with probability $\frac{1+\log d(v)}{d(v)}$, (where $d(v)$ is the degree of v) then the above idea can be made to work. The main technical difficulty is then showing that at least one candidate is present (cf. Section 2.1) and in bounding the message complexity (cf. Section 2.2). We use Lagrangian optimization to prove that on expectation at least $\Theta(\log n)$ candidates will be selected and then use a Chernoff bound to show a high probability result.

Our $\Omega(n)$ randomized lower bound is inspired by the *bridge crossing* argument of [10] and [15]. In this argument, we construct a "dumbbell" graph G which is done by taking two identical regular graphs G_1 and G_2 , removing an edge from each and adding them as bridge edges between G_1 and G_2 (so that regularity is preserved). The argument is that any leader election algorithm should send at least one message across one of the two bridge edges (bridge crossing); otherwise, it can be shown that the executions in G_1 and G_2 are identical leading to election of two leaders which is not valid. The argument in [10] shows that $\Omega(m)$ messages are needed for bridge crossing. As pointed out earlier in Section 1, this construction makes the diameter of G at least three and hence does not work for diameter-two graphs. To overcome this, we modify the construction that takes two complete graphs and add a set of bridge edges (as opposed to just two); see Fig 1. This creates a diameter-two graph; however, the large number of bridge edges requires a different style of argument and results in a bound different compared to [10]. We show that $\Omega(n)$ messages (in expectation) are needed to send a message across at least one bridge.

We also present a deterministic algorithm that uses $O(n \log n)$ messages, but takes $O(\log n)$ rounds. Note that, in a sense, this improves over the randomized algorithm that sends $O(n \log^3 n)$ messages (although, we did not strive to optimize the log factors). However, the deterministic algorithm is slower by a $\log(n)$ -factor and is more involved compared to the very simple randomized algorithm (although its analysis is a bit more complicated). Our deterministic algorithm uses ideas similar to Afek and Gafni's [1] leader election algorithm for complete graphs; however, the algorithm is a bit more involved. Our algorithm assumes knowledge of n (this is trivially true in complete networks, since every node can infer n from its degree) which is needed for termination. It is not clear if one can design an $O(n \log n)$ messages algorithm (running in say $O(\log n)$ rounds) that does not need knowledge of n , which is an interesting open question (cf. Section 6).

Finally, we present a simple reduction that shows that superlinear (in n) lower bounds in complete networks also imply lower bounds

for diameter-two networks, by showing how using only $O(n)$ messages and in $O(1)$ rounds, a complete network can be converted to a diameter-two network in a distributed manner. This shows that our deterministic algorithm (cf. Section 4) is message optimal.

1.3 Distributed Computing Model

The model we consider is similar to the models of [1, 3, 5, 7, 8], with the main addition of giving processors access to a private unbiased coin. We consider a system of n nodes, represented as an undirected graph $G = (V, E)$. In this paper, we focus on graphs with diameter $D(G) = 2$, where $D(G)$ is the diameter of $G = (V, E)$. An obvious consequence of this is that G is connected, therefore $n - 1 \leq m \leq \frac{n(n-1)}{2}$, where $m = |E|$ and $n = |V|$.

Each node has a unique identifier (ID) of $O(\log n)$ bits and runs an instance of a distributed algorithm. The computation advances in synchronous rounds where, in every round, nodes can send messages, receive messages that were sent in the same round by neighbors in G , and perform some local computation. Every node has access to the outcome of unbiased private coin flips (for randomized algorithms). Messages are the only means of communication; in particular, nodes cannot access the coin flips of other nodes, and do not share any memory. Throughout this paper, we assume that all nodes are awake initially and simultaneously start executing the algorithm. We note that initially nodes have knowledge only of themselves, in other words we assume the *clean network model* — also called the *KTO model* [17] which is standard and most commonly used. On the other hand, if one assumes the *KTI model*, where nodes have an initial knowledge of the IDs of their neighbors, there exists a trivial algorithm for leader election in a diameter-two graph that uses only $O(n)$ messages.

1.4 Other Related Works

The complexity of the leader election problem and algorithms for it, especially deterministic algorithms (guaranteed to always succeed), have been well-studied. Various algorithms and lower bounds are known in different models with synchronous/asynchronous communication and in networks of varying topologies such as a cycle, a complete graph, or some arbitrary topology (e.g., see [4, 10, 11, 13, 16, 18, 19] and the references therein).

The study of leader election algorithms is usually concerned with both message and time complexity. We discuss two sets of results, one for complete graphs and the other for general graphs. As mentioned earlier, for complete graphs, Kutten et al. [11] showed that $\tilde{\Theta}(\sqrt{n})$ is the tight message complexity bound for randomized (implicit) leader election. In particular, they presented an $O(\sqrt{n} \log^{3/2} n)$ messages algorithm that ran in $O(1)$ rounds; they also showed an almost matching lower bound for randomized leader election, showing that $\Omega(\sqrt{n})$ messages are needed for any leader election algorithm that succeeds with a sufficiently large constant probability.

For deterministic algorithms on complete graphs, it is known that $\Theta(n \log n)$ is a tight bound on the message complexity [1, 9]. In particular, Afek and Gafni [1] presented an $O(n \log n)$ messages algorithm for complete graphs that ran in $O(\log n)$ rounds. For complete graphs, Korach et al. [6] and Humblet [3] also presented $O(n \log n)$ message algorithms. Afek and Gafni [1] presented asynchronous and synchronous algorithms, as well as a tradeoff between the message

and the time complexity of synchronous *deterministic* algorithms for complete graphs: the results varied from a $O(1)$ -time, $O(n^2)$ -messages algorithm to a $O(\log n)$ -time, $O(n \log n)$ -messages algorithm. Afek and Gafni [1], as well as [6, 8] showed a lower bound of $\Omega(n \log n)$ messages for *deterministic* algorithms in the general case.⁵

For general graphs, the best known bounds are as follows. Kutten et al. [10] showed that $\Omega(m)$ is a very general lower bound on the number of messages and $\Omega(D)$ is a lower bound on the number of rounds for any leader election algorithm. It is important to point out that their lower bounds applied for graphs with *diameter at least three*. Note that these lower bounds hold even for randomized Monte Carlo algorithms that succeed even with (some large enough, but) constant success probability and apply even for implicit leader election. Earlier results, showed such lower bounds only for deterministic algorithms and only for the restricted case of comparison algorithms, where it was also required that nodes may not wake up spontaneously and that D and n were not known. The $\Omega(m)$ and $\Omega(D)$ lower bounds are *universal* in the sense that they hold for all universal algorithms (namely, algorithms that work for all graphs), apply to every $D \geq 3$, m , and n , and hold even if D , m , and n are known, all the nodes wake up simultaneously, and the algorithms can make any use of node's identities. To show that these bounds are tight, they also present an $O(m)$ messages algorithm (this algorithm is not time-optimal). An $O(D)$ time leader election algorithm is known [16] (this algorithm is not message-optimal). They also presented an $O(m \log \log n)$ messages randomized algorithm that ran in $O(D)$ rounds (where D is the network diameter) that is simultaneously almost optimal with respect to both messages and time. They also presented an $O(m \log n)$ and $O(D \log n)$ deterministic leader election algorithm for general graphs.

2 A RANDOMIZED ALGORITHM

In this section, we present a simple randomized Monte Carlo algorithm that works in a constant number of rounds. Algorithm 1 is entirely local, as nodes do not require any knowledge of n . Nevertheless, we show that we can sub-sample a small number of candidates (using only local knowledge) that then attempt to become leader. In the remainder of this section, we prove the following result.

THEOREM 2.1. *There exists a Monte Carlo randomized leader election algorithm that, with high probability, succeeds in n -node networks of diameter at most two in $O(1)$ rounds, while sending $O(n \log^3 n)$ messages.*

2.1 Proof of Correctness: Analyzing the number of candidates selected

We use the following property of diameter-2 graphs crucially in our algorithm.

⁵This lower bound assumes non-simultaneous wakeup though. If nodes are assured to wake up at the same time in synchronous complete networks, there exists a trivial algorithm: if a node's identity is some i , it waits i time before it sends any message then leader election could be solved (deterministically) in $O(n)$ messages on complete graphs in synchronous networks. Recently, Kutten [9] shows that the $\Omega(n \log n)$ lower bound holds for simultaneous wakeup as well, if the number of rounds is bounded.

Algorithm 1 Randomized leader election in $O(1)$ rounds and $O(n \log^3 n)$ message complexity

- 1: Each node $v \in V$ selects itself to be a “candidate” with probability $\frac{1+\log(d_v)}{d_v}$, where d_v is the degree of v .
 - 2: **if** v becomes a candidate **then** v sends its ID to all its neighbors.
 - 3: Each node acts as a “referee node” for all its candidate neighbors (including, possibly itself).
 - 4: If a node w receives ID's from its neighbors v_1, v_2, \dots, v_j (say), then w computes the minimum ID of those and sends it back to those neighbors. That is, w sends $\min\{ID(v_1), ID(v_2), \dots, ID(v_j)\}$ back to each of v_1, v_2, \dots, v_j .
 - 5: A node v decides that it is the leader if and only if it receives its own ID from *all* its neighbors. Otherwise v decides that it is not the leader.
-

Observation 1. Let $G = (V, E)$ be a graph of diameter 2. Then for any $u, v \in V$, either $(u, v) \in E$ or $\exists w \in V$ such that $(u, w) \in E$ and $(v, w) \in E$, i.e., u and v has at least one common neighbor w (say).

We note that if one or more candidates are selected, then only the candidate node with the minimum ID is selected as the leader. That is, the leader is unique, and therefore the algorithm produces the correct output. The only case when the algorithm may be wrong is if no candidates are selected to begin with, in which case no leader is selected. In this section, we show that, with high probability, at least two candidates are ⁶

LEMMA 2.2. *Let $f(x_1, x_2, \dots, x_n)$ be a function of n variables x_1, x_2, \dots, x_n , where x_1, x_2, \dots, x_n are positive reals. f is defined as*

$$f(x_1, x_2, \dots, x_n) \stackrel{\text{def}}{=} \sum_{i=1}^n \frac{1 + \log x_i}{x_i}.$$

Let C be a constant $\geq n\sqrt{2}$. Then $f(x_1, x_2, \dots, x_n)$ is minimized, subject to the constraint $\sum_{i=1}^n x_i = C$, when $x_i = \frac{C}{n}$, for all $1 \leq i \leq n$. The minimum value that $f(x_1, x_2, \dots, x_n)$ takes is at the point

$$\left(\frac{C}{n}, \frac{C}{n}, \dots, \frac{C}{n}\right), \text{ and is given by}$$

$$f^{\min} = f\left(\frac{C}{n}, \frac{C}{n}, \dots, \frac{C}{n}\right) = \frac{n^2}{C} \left(1 + \log\left(\frac{C}{n}\right)\right).$$

LEMMA 2.3. *Let X be a random variable that denotes the total number of candidates selected in Algorithm 1. Then the expected number of selected candidates is lower-bounded by $E[X] > 2 + \frac{1}{2} \log n$.*

PROOF. Let X_v be an indicator random variable that takes the value 1 if and only if v becomes a candidate. Then $E[X_v] = \Pr[X_v = 1] = \frac{1+\log d_v}{d_v}$. Thus if X denotes the total number of candidates selected, then

$$E[X] = \sum_{v \in V} E[X_v] = \sum_{v \in V} \frac{1 + \log(d_v)}{d_v}.$$

Since G is connected, $m \geq n - 1 \implies 2m \geq 2n - 2 > n\sqrt{2}$. Thus by Lemma 2.2, $E[X]$ is minimized subject to the constraint $\sum_{v \in V(G)} d_v = 2m$ when $d_v = \frac{2m}{n}$, i.e., when G is regular.

⁶Please refer to the full paper for the detailed proofs.

Case 1 ($n-1 \leq m \leq n^{\frac{3}{2}}$): The minimum value that $E[X]$ takes is given by

$$\begin{aligned} E[X]_{\min} &= \frac{n^2}{2m} \left(1 + \log\left(\frac{2m}{n}\right)\right) \\ &> \frac{n^2}{2m} \quad (\text{since } 1 + \log\left(\frac{2m}{n}\right) > 1) \\ &\geq \frac{\sqrt{n}}{2} \quad (\text{since } m \leq n^{\frac{3}{2}}) \end{aligned}$$

Case 2 ($n^{\frac{3}{2}} < m \leq \binom{n}{2}$): The minimum value that $E[X]$ takes is given by

$$\begin{aligned} E[X]_{\min} &= \frac{n^2}{2m} \left(1 + \log\left(\frac{2m}{n}\right)\right) \\ &> 1 + \log\left(\frac{2n^{\frac{3}{2}}}{n}\right) \quad (\text{since } \frac{n^2}{2m} > 1 \text{ and } m > n^{\frac{3}{2}}) \\ &= 1 + \log 2 + \log\left(n^{\frac{1}{2}}\right) = 2 + \frac{1}{2} \log n. \end{aligned}$$

□

We use the following variant of Chernoff Bound [14] to show *concentration*, i.e., to show that the number of candidates selected is not too less than its expected value.

THEOREM 2.4 (CHERNOFF BOUND). *Let X_1, X_2, \dots, X_n be independent indicator random variables, and let $X = \sum_{i=1}^n X_i$. Then the following Chernoff bound holds: for $0 < \delta < 1$,*

$$\Pr[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^\mu, \text{ where } \mu \stackrel{\text{def}}{=} E[X].$$

LEMMA 2.5. *If X denotes the number of candidates selected, then $\Pr[X \leq 1] < n^{-\frac{1}{3}}$.*

2.2 Computing the message complexity

Note that the expected total message complexity of the algorithm can be bounded as follows. Let random variable M^{entire} denote the total messages sent during the course of the algorithm. Let M_v be the number of messages sent by node v . Thus $M^{\text{entire}} = \sum_{v \in V} M_v$. A node v becomes a candidate with probability $\frac{1 + \log d_v}{d_v}$ and, if it does, it sends $d(v)$ messages (the referees reply to these, but this increases the total number of messages by a factor of 2). Hence by linearity of expectation, it follows that $E[M^{\text{entire}}] = \sum_{v \in V} E[M_v] = \sum_{v \in V} 2 \frac{1 + \log d_v}{d_v} d_v = 2 \sum_{v \in V} (1 + \log d_v) \leq 2 \sum_{v \in V} (1 + \log n) \leq 2n + 2n \log n$. To show concentration, we cannot directly apply a standard Chernoff bound that works for 0-1 random variables, since M_v s are not 0-1 random variables (they take values either 0 or d_v). To handle this, we bucket the degrees into (at most) $\log n$ categories based on their value then use a Chernoff bound as detailed below.

We use the following variant of Chernoff Bound [14] in the following analysis.

THEOREM 2.6 (CHERNOFF BOUND). *Let X_1, X_2, \dots, X_n be independent indicator random variables, and let $X = \sum_{i=1}^n X_i$. Then the following Chernoff bound holds: for $R \geq 6E[X]$, $\Pr[X \geq R] \leq 2^{-R}$.*

Definition 2.7. Let k be a positive integer such that $2^{k-1} < n \leq 2^k$. For $0 \leq j \leq k$, let $V_j \subset V$ be the set of vertices with degree in $(2^{j-1}, 2^j]$, i.e., if $v \in V_j$, then $2^{j-1} < d_v \leq 2^j$.

Remark 1. We note that $\sum_{j=0}^k n_j = n$, where $n_j = |V_j|$ for $0 \leq j \leq k$. In particular, $n_j \leq n$ for all $j \in [0, k]$.

Analyzing vertices with degree ≤ 2 : We recall that X_v is an indicator random variable that takes the value 1 if and only if v becomes a candidate. Then $\Pr[X_v = 1] = 1$ if $v \in V_0 \cup V_1$, i.e., every vertex with degree 1 or degree 2 selects itself to be a candidate, deterministically.

For $v \in V$, let m_v denote the number of messages that v sends. So $m_v = d_v$ if v becomes a candidate, and $m_v = 0$ otherwise. Let M_j be the total number of messages that members of V_j send, i.e.,

$$\begin{aligned} M_j &\stackrel{\text{def}}{=} \sum_{v \in V_j} m_v \leq \sum_{v \in V_j} d_v \leq \sum_{v \in V_j} 2^j = n_j \cdot 2^j \leq n \cdot 2^j \\ &\implies M_0 \leq n \text{ and } M_1 \leq 2n. \end{aligned}$$

Analyzing vertices with degree > 2 : We recall that for $v \in V$, X_v is an indicator random variable that takes the value 1 if and only if v becomes a candidate. Let i be an integer in $[2, k]$ and let $v \in V_i$. Then

Observation 2. $\frac{i}{2^i} < E[X_v] < \frac{3i}{2^i}$.

PROOF. For $v \in V_i$, $2^{i-1} < d_v \leq 2^i$. So

$$\begin{aligned} E[X_v] &= \Pr[X_v = 1] \quad (\text{since } X_v \text{ is an indicator random variable}) \\ &= \frac{1 + \log d_v}{d_v} \\ &\implies \frac{1 + \log(2^{i-1})}{2^i} < E[X_v] < \frac{1 + \log(2^i)}{2^{i-1}} \\ &[\text{since } 2^{i-1} < d_v \leq 2^i] \\ \text{or, } \frac{i}{2^i} < E[X_v] < \frac{i+1}{2^{i-1}} &\leq \frac{3i}{2^i} \\ [\text{since } i \geq 2 \implies \frac{3i}{2} &\geq i+1] \end{aligned}$$

□

For $0 \leq j \leq k$, let Y_j be a random variable that denotes the total number of candidates selected from V_j .

Observation 3. For $2 \leq i \leq k$, $\frac{in_i}{2^i} < E[Y_i] < \frac{3in_i}{2^i}$.

PROOF.

$$\begin{aligned} Y_i &= \sum_{v \in V_i} X_v \implies E[Y_i] = E\left[\sum_{v \in V_i} X_v\right] \\ &= \sum_{v \in V_i} E[X_v] \quad [\text{by linearity of expectation}] \\ &\implies \sum_{v \in V_i} \frac{i}{2^i} < E[Y_i] < \sum_{v \in V_i} \frac{3i}{2^i} \\ &\implies \frac{in_i}{2^i} < E[Y_i] < \frac{3in_i}{2^i} \end{aligned}$$

□

Remark 2. $\forall u, v \in V(G)$, $u \neq v$, X_u and X_v are independent, and for $0 \leq j \leq k$, we define Y_j as $Y_j = \sum_{v \in V_j} X_v$, i.e., Y_j is a sum of independent indicator random variables. Hence we can use Theorem 2.6 to show that Y_j is concentrated around its mean, i.e., its expectation.

We recall that for $0 \leq j \leq k$, M_j is the total number of messages that members of V_j send, i.e., for $2 \leq i \leq k$,

$$M_i = \sum_{v \in V_i} m_v = \sum_{v \in V_i, X_v=1} d_v$$

LEMMA 2.8. *For any integer $i \in [2, k]$, it holds that $\Pr[M_i \geq 24n \log^2 n] \leq \frac{1}{n^4}$.*

PROOF.

$$\begin{aligned} M_i &= \sum_{v \in V_i} m_v = \sum_{v \in V_i, X_v=1} d_v \\ &\implies \sum_{v \in V_i, X_v=1} 2^{i-1} < M_i \leq \sum_{v \in V_i, X_v=1} 2^i \text{ [since } 2^{i-1} < d_v \leq 2^i \text{]} \\ &\implies 2^{i-1} \cdot Y_i < M_i \leq 2^i \cdot Y_i \end{aligned}$$

Case 1 ($E[Y_i] = 0$): $E[Y_i] = 0$ if and only if $n_i = 0$, i.e., if and only if $\nexists v \in V$ such that $2^{i-1} < d_v \leq 2^i$. But $n_i = 0 \implies V_i = \emptyset$, the empty set. Therefore, $M_i = 0$.

Case 2 ($0 < E[Y_i] < 1$): Assuming $n \geq 3$, $4 \log n > 6 > 6E[Y_i]$. Therefore, by Theorem 2.6,

$$\begin{aligned} \Pr[Y_i \geq 4 \log n] &\leq 2^{-4 \log n} = n^{-4} \\ &\implies \Pr[M_i \geq 2^i \cdot 4 \log n] \leq n^{-4} \text{ [since } M_i \leq 2^i \cdot Y_i \text{]} \\ &\implies \Pr[M_i \geq 8n \log n] \leq n^{-4} \text{ [since } i \leq k < \log n + 1 \text{]} \end{aligned}$$

Case 3 ($E[Y_i] \geq 1$): We have shown before that $E[Y_i] \leq \frac{3in_i}{2^i}$. But $n_i \leq n$ for all $2 \leq i \leq k$. Hence $E[Y_i] \leq \frac{3ni}{2^i}$. Assuming $n \geq 3$, $4 \log n > 6$. Therefore, by Theorem 2.6,

$$\begin{aligned} \Pr[Y_i \geq 12n \log n \cdot \frac{i}{2^i}] &\leq \Pr[Y_i \geq 4 \log n E[Y_i]] \\ &\leq 2^{-4 \log n E[Y_i]} = n^{-4E[Y_i]} \leq n^{-4} \text{ [since } E[Y_i] \geq 1 \text{]} \\ &\implies \Pr[M_i \geq 12in \log n] \leq n^{-4} \text{ [since } M_i \leq 2^i \cdot Y_i \text{]} \\ &\implies \Pr[M_i \geq 24n \log^2 n] \leq \Pr[M_i \geq 12in \log n] \leq n^{-4} \\ &\text{ [since } i \leq k < \log n + 1 < 2 \log n \text{]} \end{aligned}$$

□

LEMMA 2.9. *If M denotes the total number of messages sent by the candidates (in the first round only), then $\Pr[M \geq 27n \log^3 n] < \frac{1}{n^3}$.*

PROOF.

$$\begin{aligned} M &\stackrel{\text{def}}{=} \sum_{i=0}^k M_i = M_0 + M_1 + \sum_{i=2}^k M_i \\ &= n + 2n + \sum_{i=2}^k M_i \text{ [since } M_0 = n \text{ and } M_1 = 2n \text{]} \\ &= 3n + \sum_{i=2}^k M_i \end{aligned}$$

But for $2 \leq i \leq k$, $\Pr[M_i \geq 24n \log^2 n] \leq \frac{1}{n^4}$. Taking the union bound over $2 \leq i \leq k$,

$$\begin{aligned} \Pr[M_{i'} \geq 24n \log^2 n] \text{ for some } i' \in [2, k] &\text{ is } \leq \frac{\log n}{n^4} < \frac{1}{n^3} \\ &\implies \Pr\left[\sum_{i=2}^k M_i \geq 24n \log^3 n\right] < \frac{1}{n^3} \\ &\implies \Pr\left[3n + \sum_{i=2}^k M_i \geq 3n + 24n \log^3 n\right] < \frac{1}{n^3} \\ &\implies \Pr[M \geq 3n + 24n \log^3 n] < \frac{1}{n^3} \text{ [since } M = 3n + \sum_{i=2}^k M_i \text{]} \end{aligned}$$

But $3n \leq 3n \log^3 n$ for $n \geq 2$, or, $3n + 24n \log^3 n \leq 27n \log^3 n$. Hence

$$\Pr[M \geq 27n \log^3 n] \leq \Pr[M \geq 3n + 24n \log^3 n] < \frac{1}{n^3}.$$

□

LEMMA 2.10. *If M^{entire} denotes the total number of messages sent during the entire run of Algorithm 1, then $\Pr[M^{\text{entire}} \geq 54n \log^3 n] < \frac{1}{n^3}$.*

PROOF. Let M' denote the number of messages sent by the “referee” nodes in the *second* round of the algorithm. We recall that M is the number of messages sent by the “candidate” nodes in the *first* round of the algorithm. Then $M' \leq M$, and $M^{\text{entire}} = M + M' \leq 2M$, and the result follows. □

This completes the proof of Theorem 2.1.

3 A LOWER BOUND FOR RANDOMIZED ALGORITHMS

In this section we show that $\Omega(n)$ is a lower bound on the message complexity for solving leader election with any randomized algorithm in diameter-two networks. Notice that [11] show a lower bound of $\Omega(\sqrt{n})$ for the special case of diameter 1 networks, and we know from [10] that, for the message complexity becomes $\Omega(m)$ for (most) diameter 3 networks. Thus, Theorem 3.1 completes the picture regarding the message complexity of leader election when considering networks according to their diameter.

THEOREM 3.1. *Any algorithm that solves implicit leader election with probability at least $\frac{1}{2} + \epsilon$ in any n -node network with diameter at most 2, for any constant $\epsilon > 0$, sends at least $\Omega(n)$ messages in expectation. This holds even if nodes have unique IDs and know the network size n .*

In the remainder of this section, we prove Theorem 3.1. Assume towards a contradiction, that there is an algorithm that elects a leader with probability $\frac{1}{2} + \epsilon$ that sends $o(n)$ messages with probability approaching 1. In other words, we assume that the event where the algorithm sends more than $o(n)$ messages (of arbitrary size) happens with probability at most $o(1)$.

Unique IDs vs Anonymous. Before describing our lower bound construction, we briefly recall a simple reduction used in [11] that shows that assuming unique IDs does not change the success probability of the algorithm by more than $\frac{1}{n}$: Since we assume that nodes have knowledge of n , it is straightforward to see that nodes

can obtain unique IDs (whp) by choosing a random integer in the range $[1, n^c]$, for some constant $c \geq 4$. Thus, we can simulate an algorithm that requires unique IDs in the anonymous case and the simulation will be correct with high probability. Suppose that there is an algorithm A that can break the message complexity bound of Theorem 3.1 while succeeding with probability $\geq \frac{1}{2} + \epsilon$, for some constant $\epsilon > 0$, when nodes have unique IDs. Then, the above simulation yields an algorithm A' that works in the case where nodes are anonymous with the same message complexity bound as algorithm A and succeeds with probability at least $(\frac{1}{2} + \epsilon - \frac{1}{n}) \geq \frac{1}{2} + \epsilon'$, for some constant $\epsilon' > 0$. We conclude that proving the lower bound for the anonymous case is sufficient to imply a lower bound for the case where nodes have unique IDs.

The Lower Bound Graph. Our lower bound is inspired by the bridge crossing argument of [10] and [15]. For simplicity, we assume that $\frac{n}{4}$ is an integer. Consider two cliques C_1 and C_2 of $\frac{n}{2}$ nodes each and let G' be the n -node graph consisting of the two (disjoint) cliques. The *port numbering* of an edge $e = (u_i, v_j) \in E(G')$ refers to the port number at u_i and the respective port number at v_j that connects e . The port numberings of the edges defines an *instance of G'* .

Given an instance of G' , we will now describe how to obtain an instance of graph G that has the same node set as G' . Fix some arbitrary enumeration $u_1, \dots, u_{\frac{n}{2}}$ of the nodes ⁷ in C_1 and similarly let $v_1, \dots, v_{\frac{n}{2}}$ be an enumeration of the nodes in C_2 . To define the edges of G , we randomly choose a maximal matching M_1 of $\frac{n}{4}$ edges in the subgraph C_1 . Consider the set of edges $M'_2 = \{(v_i, v_j) \mid \exists (u_i, u_j) \in M_1\}$, which is simply the matching in C_2 corresponding to M_1 in C_1 . We define M_2 to be a randomly chosen maximal matching on C_2 when using only edges in $E(G') \setminus M'_2$. Then, we remove all edges in $M_1 \cup M_2$ from G' . So far, we have obtained a graph where each node has one *unwired port*.

The edge set of G consists of all the remaining edges of G' in addition to the set $M = \{(u_1, v_1), \dots, (u_{\frac{n}{2}}, v_{\frac{n}{2}})\}$, where we connect these *bridge edges* by using the unwired ports that we obtained by removing the edges as described above. We say that an edge is an *intra-clique edge* if it has both endpoints in either C_1 or C_2 . Observe that the intra-clique edges of G are a subset of the intra-clique edges of G' . Figure 1 gives an illustration of this construction.

LEMMA 3.2. *Graph G is an n -node network of diameter 2 and the port numbering of each intra-clique edge in G is the same as of the corresponding edge in G' .*

A *state* σ of the nodes in C_1 is a $\frac{n}{2}$ -size vector of the local states of the $n/2$ nodes in C_1 . Since we assume that nodes are anonymous, a state σ that is reached by the nodes in C_1 , can also be reached by the nodes in C_2 . More formally, when executing the algorithm on the disconnected network G' , we can observe that every possible state σ (of $\frac{n}{2}$ nodes) has the same probability of occurring in C_1 as in C_2 . Thus, a state where there is exactly one leader among the $\frac{n}{2}$ nodes of a clique in G' , is reached with some specific probability q depending on the algorithm. By a slight abuse of notation, we also use G' and G to denote the event that the algorithm executes on G' respectively G . For the probability of the event One, which occurs




Figure 1: The lower bound graph construction used in Theorem 3.1 for $n = 12$, with cliques C_1 and C_2 , where $V(C_1) = \{u_1, \dots, u_6\}$ and $V(C_2) = \{v_1, \dots, v_6\}$. The dotted red edges are the edges in M_1 and M_2 that are removed from C_1 and C_2 when constructing G and the blue dashed inter-clique edges are given by the maximal matching M between C_1 and C_2 . Each blue edge incident to some node u_i is connected by using the port number of u_i 's (removed) red edge.

when there is exactly 1 leader among the n nodes, we get

$$\Pr[\text{One} \mid G'] = 2q(1 - q) \leq \frac{1}{2}, \quad (1)$$

which holds for any value of q . Since G' is disconnected, the algorithm does not need to succeed with nonzero probability when being executed on G' . However, below we will use this observation to obtain an upper bound on the probability of obtaining (exactly) one leader in G .

Now consider the execution on the diameter 2 network G (obtained by modifying the ports of G' as described above) and let $C_1 \leftrightarrow C_2$ be the event that no message is sent across the bridges between C_1 and C_2 . Since we assume the port numbering model where nodes are unaware of their neighbors initially, it follows by Lemma 3.2 that

$$\Pr[\text{One} \mid C_1 \leftrightarrow C_2, G] = \Pr[\text{One} \mid G']. \quad (2)$$

Let M be the event that the algorithm sends $o(n)$ messages. Recall that we assume towards a contradiction that $\Pr[M \mid G] = 1 - o(1)$.

LEMMA 3.3. $\Pr[C_1 \leftrightarrow C_2 \mid G, M] = o(1)$.

PROOF. The proof is inspired by the guessing game approach of [2] and Lemma 16 in [15]. Initially, any node $u \in C_1$ has $\frac{n}{2} - 1$ ports that are all equally likely (i.e., a probability $p = \frac{1}{\frac{n}{2}-1}$) to be connected to the (single) bridge edge incident to u . As u sends messages to other nodes, it might learn about some of its ports connecting to non-bridge edges and hence this probability can increase over time. However, we condition on event M , i.e., the algorithm sends at most $o(n)$ messages in total and hence at least $\frac{n}{4}$ ports of each node u remain unused at any point.

It follows that the probability of some node u to activate a (previously unused) port that connects a bridge edge is at most $\frac{4}{n}$ at any point of the execution. Let X be the total number of ports connecting bridge edges that are activated during the run of the algorithm and

⁷This enumeration is used solely for the description of the lower bound construction and is unbeknownst to the nodes.

let X_u be the indicator random variable that is 1 iff node u sends a message across its bridge edge. Let S_u be the number of messages sent by node u . It follows by the hypergeometric distribution that

$$E[X_u | G, M] = S_u \frac{1}{\Theta(n)},$$

for each node u and hence,

$$E[X | G, M] = \sum_{u \in V(G)} \frac{S_u}{\Theta(n)} = \frac{1}{\Theta(n)} \sum_{u \in V(G)} S_u = o(1)$$

where we have used the fact that $\sum_{u \in V(G)} S_u = o(n)$ due to conditioning on event M . By Markov's Inequality, it follows that the event $C_1 \leftrightarrow C_2$, i.e., $X \geq 1$, occurs with probability at most $o(1)$. \square

We now combine the above observations to obtain

$$\begin{aligned} \Pr[\text{One} | G, M] &= \Pr[\text{One} | C_1 \leftrightarrow C_2, G, M] \Pr[C_1 \leftrightarrow C_2 | G, M] \\ &\quad + \Pr[\text{One} | C_1 \not\leftrightarrow C_2, G, M] \Pr[C_1 \not\leftrightarrow C_2 | G, M] \\ &\leq \Pr[\text{One} | C_1 \leftrightarrow C_2, G, M] + o(1) \quad (\text{by Lem. 3.3}) \\ &\leq \frac{1}{2} + o(1), \end{aligned} \quad (3)$$

where the last inequality follows by first using (2) and noting that the upper bound (1) still holds when conditioning on the event M . Finally, we recall that the algorithm succeeds with probability at least $\frac{1}{2} + \epsilon$ and $\Pr[M | G] \geq 1 - o(1)$, which yields

$$\frac{1}{2} + \epsilon \leq \Pr[\text{One} | G] \leq \Pr[\text{One} | G, M] + o(1) \leq \frac{1}{2} + o(1),$$

which is a contradiction, since we have assumed that $\epsilon > 0$ is a constant.

4 A DETERMINISTIC ALGORITHM

Our algorithm (Algorithm 2) is inspired by the solution of Afek and Gafni [1] for the n -node clique. However, there are some complications that we explain below, since we cannot rely on all nodes to be connected by an edge. Note that our algorithm assumes that n (or a constant factor upper bound for $\log n$) is known to all nodes.

For any node $v \in V$, we denote the degree of v by d_v and the ID of v by ID_v . At any time-point in the algorithm, L_v denotes the highest ID that v has so far learned (among all the probe messages it has received, in the current round or in some previous round).

The algorithm proceeds as a sequence of $\Theta(\log n)$ phases. Initially every node is a “candidate” and is “active”. Each node v numbers its neighbors from 1 to d_v , denoted by $w_{v,1}, w_{v,2}, \dots, w_{v,d_v}$ respectively. In phase i , if a node v is active, v sends probe-messages containing its ID to its neighbors $w_{v,2^{i-1}}, \dots, w_{v,k}$, where $k = \min\{d_v, 2^i - 1\}$. Each one of them replies back with the highest ID it has seen so far. If any on those ID's is higher than ID_v , then v stops being a candidate and becomes inactive. Node v also becomes inactive if it has finished sending probe-messages to all its neighbors. After finishing the $\Theta(\log n)$ phases v becomes leader if it is still a candidate.

The idea behind the algorithm is to exploit the *neighborhood intersection property* (cf. Observation 1) of diameter-2 networks. Since for any $u, v \in V$, there is an $x \in V$ that is connected to both u and v (unless u and v are directly connected via an edge) and acts as a “referee” node for candidates u and v . This means that x serves to inform u and v who among them is the winner, i.e., has the higher

Algorithm 2 Deterministic Leader Election in $O(\log n)$ rounds and with $O(n \log n)$ messages: Code for a node v

```

1:  $v$  becomes a “candidate” and “active”.
2:  $L_v \leftarrow ID_v$ .
3:  $N_v \leftarrow ID_v$ .
4:  $v$  numbers its neighbors from 1 to  $d_v$ , which are called
    $w_{v,1}, w_{v,2}, \dots, w_{v,d_v}$  respectively.
5: for phase  $i = 1$  to  $\Theta(\log n)$  do
6:   if  $v$  is active then
7:      $v$  sends a “probe” message containing its ID to its neigh-
       bors  $w_{v,2^{i-1}}, \dots, w_{v,\min\{d_v, 2^i - 1\}}$ .
8:     if  $d_v \leq 2^i - 1$  then
9:        $v$  becomes inactive.
10:    end if
11:   end if
12:   Let  $X$  be the set (possibly empty) of neighbors of  $v$  from
       whom  $v$  receives messages in this round.
13:   Let  $\mathcal{ID}$  be the set of ID's sent to  $v$  by the members of  $X$ .
14:   Let  $ID_u$  be the highest ID in  $\mathcal{ID}$ .
15:   if  $ID_u > L_v$  then
16:      $v$  sends  $ID_u$  to  $N_v$ .
17:      $L_v \leftarrow ID_u$ .  $\triangleright v$  stores the highest ID seen so far in  $L_v$ .
18:      $N_v \leftarrow x$ .  $\triangleright v$  remembers neighbor who told  $v$  about  $L_v$ .
19:      $v$  becomes “inactive” and “non-candidate”.
20:   end if
21:    $v$  tells every member of  $X$  about  $L_v$ , i.e., the highest ID it
       has seen so far.
22: end for
23: if  $v$  is still a candidate then
24:    $v$  elects itself to be the leader.
25: end if

```

ID. Thus at the end of the algorithm, every node except the one with the highest ID should know that he is not a leader. We present the formal analysis of Theorem 4.1 in Sections 4.1 and 4.2.

THEOREM 4.1. *There exists a deterministic leader election algorithm for n -node networks with diameter at most 2 that sends $O(n \log n)$ messages and terminates in $O(\log n)$ rounds.*

In the pseudocode and the subsequent analysis we use v and ID_v interchangeably to denote the node v .

4.1 Proof of Correctness

Define v^{\max} to be the node with the highest ID in G .

LEMMA 4.2. v^{\max} becomes a leader.

PROOF. Since v^{\max} has the highest ID in G , the if-clause of Line 15 of Algorithm 2 is never satisfied for v^{\max} . Therefore v^{\max} never becomes a non-candidate, and hence becomes a leader at the end of the algorithm. \square

LEMMA 4.3. No other node except v^{\max} becomes a leader.

PROOF. Consider any $u \in V$ such that $u \neq v^{\max}$.

- **Case 1 (v^{\max} and u are connected via an edge):** Since v^{\max} has the highest ID in G , the if-clause of Line 15 of Algorithm

2 is never satisfied for v^{\max} . Therefore v^{\max} becomes inactive only if it has already sent probe-messages to all its neighbors (or v^{\max} never becomes inactive). In particular, u always receives a probe-message from v^{\max} containing $ID_{v^{\max}}$. Since $ID_{v^{\max}} > ID_u$, u becomes a non-candidate at that point (if u was still a candidate until that point) and therefore does not become a leader.

- **Case 2 (v^{\max} and u do not have an edge between them):** By Observation 1, there is some $x \in V$ such that both v^{\max} and u have edges going to x . And we have already established that v^{\max} will always send a probe-message to x at some point of time or another.

– **Case 2(a) (u does not send a probe-message to x):** This implies that u became inactive before it could send a probe-message to x . But then u could have become inactive only if the if-clause of Line 15 of Algorithm 2 got satisfied at some point. Then u became a non-candidate too at the same time and therefore would not become a leader.

– **Case 2(b) (u sends a probe-message to x before v^{\max} does):** Suppose u sends a probe-message to x at round i and v^{\max} sends a probe-message to x at round i' , where $1 \leq i < i' \leq \log n$. If x had seen an ID higher than ID_u up until round i , then x immediately informs u and u becomes a non-candidate.

So suppose not. Then, after round i , x sets its local variables L_x and N_x to ID_u and u respectively. Let $j > i$ be the smallest integer such that x receives a probe-message from a neighbor u' at round j , where $ID_{u'} > ID_u$. Note that v^{\max} will always send a probe-message to x , therefore such a u' exists. Then, after round j , x sets its local variables L_x and N_x to $ID_{u'}$ and u' respectively, and informs u of this change. u becomes a non-candidate at that point of time.

– **Case 2(c) (u and v^{\max} each sends a probe-message to x at the same time):** Since $ID_{v^{\max}}$ is the highest ID in the network, L_x is assigned the value $ID_{v^{\max}}$ at this point, and x tells u about $L_x = ID_{v^{\max}} > ID_u$, causing u to become a non-candidate.

– **Case 2(d) (u sends a probe-message to x after v^{\max} does):** Suppose v^{\max} sends a probe-message to x at round i and u sends a probe-message to x at round i' , where $1 \leq i < i' \leq \log n$. Then x sets its local variables L_x and N_x to $ID_{v^{\max}}$ and v^{\max} , respectively, after round i . So when u comes probing at round $i' > i$, x tells u about $L_x = ID_{v^{\max}} > ID_u$, causing u to become a non-candidate. \square

4.2 Message Complexity

LEMMA 4.4. *At the end of round i , there are at most $\frac{n}{2^i}$ “active” nodes.*

PROOF. Consider a node v that is active at the end of round i . This implies that the if-clause of Line 15 of Algorithm 2 has not so far been satisfied for v , which in turn implies that $ID_v > ID_{w_{v,j}}$ for $1 \leq j \leq 2^i - 1$, therefore none of $w_{v,1}, w_{v,2}, \dots, w_{v,2^i-1}$ is active after round i . Thus for every active node at the end of round i , there are at least $2^i - 1$ inactive nodes. We call this set of inactive nodes, together with v itself, the “kingdom” of v , i.e.,

$$KINGDOM(v) \stackrel{\text{def}}{=} \{v\} \cup \{w_{v,1}, w_{v,2}, \dots, w_{v,2^i-1}\}$$

$$\text{and } |KINGDOM(v)| = 2^i.$$

If we can show that these kingdoms are disjoint for two different active nodes, then we are done.

Proof by contradiction: Suppose not. Suppose there are two nodes u and v such that

$$u \neq v \text{ and } KINGDOM(u) \cap KINGDOM(v) \neq \emptyset$$

(after some round i , $1 \leq i \leq \log n$). Let x be such that $x \in KINGDOM(u) \cap KINGDOM(v)$. Since an active node obviously cannot belong to the kingdom of another active node, this x equals neither u nor v , and therefore

$$x \in \{w_{v,1}, w_{v,2}, \dots, w_{v,2^i-1}\} \cap \{w_{u,1}, w_{u,2}, \dots, w_{u,2^i-1}\},$$

that is, both u and v have sent their respective probe-messages to x . Without loss of generality, let $ID_v > ID_u$.

• **Case 1 (u sends a probe-message to x before v does):** Suppose u sends a probe-message to x at round j and v sends a probe-message to x at round j' , where $1 \leq j < j' \leq i$. If x had seen an ID higher than ID_u up until round j , then x immediately informs u and u becomes inactive. Contradiction.

So suppose not. Then, after round j , x sets its local variables L_x and N_x to ID_u and u respectively. Let $k > j$ be the smallest integer such that x receives a probe-message from a neighbor u' at round k , where $ID_{u'} > ID_u$. Note that v sends a probe-message to x at round j' , where $j < j' \leq i$, and $ID_v > ID_u$. Therefore such a u' exists. Then, after round k , x sets its local variables L_x and N_x to $ID_{u'}$ and u' respectively, and informs u of this change. u becomes inactive at that point of time, i.e., after round k , where $k \leq j' \leq i$. Contradiction.

• **Case 2 (u and v each sends a probe-message to x at the same time):** Suppose that u and v each sends a probe-message to x at the same round j , where $1 \leq j \leq i$. Since $ID_v > ID_u$, x has at least one neighbor u' such that $ID_{u'} > ID_u$. Therefore x would not set L_x to ID_u (or N_x to u), and x would inform u about that after round j , causing u to then become inactive. Contradiction.

• **Case 3 (u sends a probe-message to x after v does):** Suppose v sends a probe-message to x at round j and u sends a probe-message to x at round j' , where $1 \leq j < j' \leq i$. Then x sets its local variables L_x and N_x to ID_v and v , respectively, after round j . So when u comes probing at round $j' > j$, x tells u about $L_x \geq ID_v > ID_u$, causing u to become inactive. Contradiction. \square

LEMMA 4.5. *In round i , at most $3n$ messages are transmitted.*

PROOF. In round i , each active node sends exactly 2^{i-1} probe messages, and each probe-message generates at most two responses (corresponding to Lines 16 and 21 of Algorithm 2). Thus, in round i , each active node contributes to, directly or indirectly, at most $3 \cdot 2^{i-1}$ messages. The result immediately follows from Lemma 4.4. \square

Since the algorithm runs for $\log n$ rounds, Theorem 4.1 immediately follows.

5 A DETERMINISTIC LOWER BOUND

We will show a lower bound of $\Omega(n \log n)$ message complexity by reducing the problem of “leader election in complete graphs” to that of “leader election in graphs of diameter two”. This reduction itself would take two rounds and $O(n)$ messages. Then, since the former is known to have $\Omega(n \log n)$ message complexity, the latter would have the same lower bound too (cf. Section 1.1).

Suppose \mathcal{A} is a leader election algorithm that works for any graph of diameter two. Let $G = (V, E)$ be our input instance for the problem of “leader election in complete graphs”, i.e., G is the complete graph on n nodes, say.

The reduction: G sparsifies itself into a diameter-two graph (G' , say, where $G' = (V, E')$, where $E' \subseteq E$) on which \mathcal{A} works thereafter. This sparsification takes $O(n)$ messages and a constant number of rounds (two, to be exact) and is done as follows.

- **Round 1:** Each node v chooses one of its neighbours (any arbitrary one) and asks its ID. If this neighbour’s ID is larger than its own ID, then v will “drop” that edge, i.e., it won’t use that for communication in the subsequent simulation of \mathcal{A} . Otherwise v will keep that edge.
For $v \in V$, if v has $\lceil \frac{n}{2} \rceil$ or more edges removed, then v makes itself a “candidate”.
- **Round 2:** The candidates from the previous round send their ID’s to all the nodes in the network using edges of G . By Lemma 5.1, there can be at most two such nodes. Thus the total number of messages sent is still $O(n)$. Then each node (including the candidates themselves) receives the ID’s of up to two candidates and chooses the highest of them to be the ID of the leader.

If no such node exists which has had $\lceil \frac{n}{2} \rceil$ or more edges removed, then G' has diameter two (please refer to Lemma 5.2), and we run \mathcal{A} on G' . \mathcal{A} returns a leader on G' which makes itself the leader of G too, and informs all its neighbors. This takes $O(n)$ messages.

5.1 Proof of Correctness

Observation 4. E has at most $n - 1$ edges more than E' .

PROOF. Each node except the node with the highest ID drops at most one edge. The node with the highest ID drops no edge. \square

LEMMA 5.1. *For $n \geq 3$, there can be at most two nodes in G' that has had $\lceil \frac{n}{2} \rceil$ or more edges removed.*

LEMMA 5.2. *If no node exists in G' which has had $\lceil \frac{n}{2} \rceil$ or more edges removed, then G' has diameter two.*

6 CONCLUSION

We settle the message complexity of leader election throughout the diameter spectrum, by presenting almost tight bounds (tight upto polylog(n) factors) for diameter-two graphs which were left open by previous results [10, 11]. Several open problems arise from our work. Is it possible to show an high probability bound of $O(n)$ messages for randomized leader election that runs in $O(1)$ rounds?

This will match the lower bounds, by closing the polylog(n) factor. It might be possible to improve the analysis of our randomized algorithm to show $O(n \log n)$ messages. Another very interesting question is whether explicit leader election (i.e., where all nodes should also know the identity of the leader) can be performed in $\tilde{O}(n)$ messages in diameter-two graphs (this is true for complete graphs, but not for diameter three and beyond). Also removing the assumption of the knowledge of n (or showing that it is not possible) for deterministic algorithms with $\tilde{O}(n)$ message complexity and running in $\tilde{O}(1)$ rounds is open.

REFERENCES

- [1] Yehuda Afek and Eli Gafni. Time and message bounds for election in synchronous and asynchronous complete networks. *SIAM Journal of Computing*, 20(2):376–394, 1991.
- [2] Seth Gilbert, Peter Robinson, and Suman Sourav. Brief announcement: Gossiping with latencies. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25–27, 2017*, pages 255–257, 2017.
- [3] P Humblet. Electing a leader in a clique in $O(n \log n)$ messages. Intern. Memo., Laboratory for Information and Decision Systems, M.I.T., Cambridge, Mass, 1984.
- [4] Maleq Khan, Fabian Kuhn, Dahlia Malkhi, Gopal Pandurangan, and Kunal Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing, PODC '08*, pages 263–272, New York, NY, USA, 2008. ACM.
- [5] E. Korach, S. Kutten, and S. Moran. A modular technique for the design of efficient distributed leader finding algorithms. *ACM Trans. Program. Lang. Syst.*, 12(1):84–101, January 1990.
- [6] E. Korach, S. Moran, and S. Zaks. Tight lower and upper bounds for some distributed algorithms for a complete network of processors. In *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing, PODC '84*, pages 199–207, New York, NY, USA, 1984. ACM.
- [7] E. Korach, S. Moran, and S. Zaks. The optimality of distributive constructions of minimum weight and degree restricted spanning trees in a complete network of processors. *SIAM Journal on Computing*, 16(2):231–236, 1987.
- [8] E. Korach, S. Moran, and S. Zaks. Optimal lower bounds for some distributed algorithms for a complete network of processors. *Theoretical Computer Science*, 64(1):125 – 132, 1989.
- [9] Shay Kutten. Private communication. 2017.
- [10] Shay Kutten, Gopal Pandurangan, David Peleg, Peter Robinson, and Amitabh Trehan. On the complexity of universal leader election. *Journal of the ACM*, 62(1):7:1–7:27, March 2015. Invited paper from ACM PODC 2013.
- [11] Shay Kutten, Gopal Pandurangan, David Peleg, Peter Robinson, and Amitabh Trehan. Sublinear bounds for randomized leader election. *Theoretical Computer Science*, 561, Part B:134 – 143, 2015. Special Issue on Distributed Computing and Networking.
- [12] Gérard Le Lann. Distributed systems - towards a formal approach. In *IFIP Congress*, pages 155–160, 1977.
- [13] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.
- [14] Michael Mitzenmacher and Eli Upfal. *Probability and Computing - Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
- [15] Shreyas Pai, Gopal Pandurangan, Sriram V. Pemmaraju, Talal Riaz, and Peter Robinson. Symmetry breaking in the CONGEST model: Time- and message-efficient algorithms for ruling sets. In *DISC 2017, Vienna*, to appear.
- [16] David Peleg. Time-optimal leader election in general networks. *Journal of Parallel and Distributed Computing*, 8(1):96 – 99, 1990.
- [17] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [18] Nicola Santoro. *Design and Analysis of Distributed Algorithms (Wiley Series on Parallel and Distributed Computing)*. Wiley-Interscience, 2006.
- [19] Gerard Tel. *Introduction to distributed algorithms*. Cambridge University Press, New York, NY, USA, 1994.