On the Impact of Sense of Direction on Message Complexity *

Paola Flocchini[†] Bernard Mans[‡] Nicola Santoro [§]

Abstract

In this paper, we prove a general result on the impact of sense of direction. We show that, in arbitrary graphs, *any* sense of direction has a dramatic effect on the communication complexity of several important distributed problems: *Broadcast*, *Depth-First Traversal*, *Election*, and *Spanning Tree Construction*.
In systems with n nodes and e communication links, the solution for the Depth First Traversal and the Broadcast problems require $\Omega(e)$ messages with arbitrary labelings; we show that, with *any* sense of direction, they can be solved exchanging only $\Theta(n)$ messages, even if the system is *anonymous*. The problems of Election and of Spanning-Tree Construction require $\Omega(e + n \log n)$ messages with arbitrary labelings; on the other hand, we show that they can be solved with $\Theta(n \log n)$ messages with any sense of direction.
The results presented here completely explain and generalize the existing results which now follow as corollaries for specific labelings.

Key Words Distributed algorithms, Complexity, Graph algorithms.

1 Introduction

A *distributed system* is a collection of processing entities (e.g., agents, processes, processors) connected by a communication network. The network will be viewed as an edge-labeled graph where nodes correspond to entities and edges correspond to bidirectional links between them. Each entity has a local, non-shared, memory and can communicate with its neighbors by exchanging messages; moreover, every entity associates a distinct label (e.g., port number) to each of its incident links so that it can distinguish among its neigbours.

The properties of the chosen edge-labeling can be used in the solution of distributed problems. *Sense of direction* is a global consistency property of edge-labelings which has a strong impact on both *computability* and *complexity*. For instance, the presence of sense of direction empowers an anonymous system (the least powerful distributed system where the entities have no identifiers) with a unique global view of the entire system; it is, thus, possible to route a

*Research supported in part by N.S.E.R.C, grant #A2415, by a Carleton University Research Achievement Award, by a CNR International Fellowship, and by a Merit Research Grant (MRG#A128) from James Cook University of North Queensland, Australia.

†Département d' Informatique et de Recherche Opérationnelle, Université de Montréal, Montréal, Quebec, H3C 3J7, Canada. flocchini@iro.umontreal.ca

‡Dept. of Computing, School of Mathematics, Physics, Computing and Electronics, Macquarie University, Sydney, NSW2109, Australia, bmans@mpce.mq.edu.au

§Center for Parallel and Distributed Computing, School of Computer Science, Carleton University, Ottawa K1S 5B6 Canada. santoro@scs.carleton.ca

message optimally to its destination even though there are no global names, i.e., neither the source nor the destination have an identity [7].

As for the impact of sense of direction on message complexity, there have been many results, but obtained for special labelings. The first result is due to [14]: the *chordal* labeling of a *complete graph* allows the message complexity of the Election problem to be reduced from $\Omega(n \log n)$ [12] to $\Theta(n)$ messages, where n is the number of entities in the system (see also [17, 23]).

Since this first result, the evidence of the impact of sense of direction for specific labelings in particular graphs has been accumulating in the recent years. For example, in *circulant graphs* with arbitrary labelings, the Election problem requires $\Omega(n \log n)$ messages; with the *chordal* labeling there exist algorithms whose complexity depends on the chord structure and can be linear [2, 10, 19]. In these graphs, the properties of the labeling have been efficiently used for the Wake-up problem [9] and for Fault-Tolerant Election (e.g., [18]). However, in a very particular case of unlabelled circulant graphs (a ring chord associated with a chord of length approximately $\sqrt{N}$), the Election problem can be solved using $\Theta(N)$ messages [15].

Similarly, $\Theta(n)$ Election algorithms exist for a *hypercube* with the traditional *dimensional* labeling (e.g., [4, 20]). An even simpler $\Theta(n)$ technique has been found if the hypercube has the *chordal* labeling [4].

In systems where the topology is unknown to the entities (the arbitrary graph case), the availability of the *neighboring* labeling reduces the complexity of the Election problem from $\Omega(e + n \log n)$ [1, 21] to $\Theta(e)$ messages, where e is the number of links in the network [13]. With the same labeling, the message complexity of Depth First Traversal drops from $O(e)$ to $\Theta(n)$ [22]. The availability of the *commutative group* labeling reduces the complexity of the Broadcast problem from $\Omega(e)$ to $\Theta(n)$ [24]. The same reductions for the Election, the Broadcast, and Depth First Traversal problems can be obtained also with the simpler *chordal* labeling [5, 16].

All these labelings differ greatly from each other: the "chordal" labeling in circulant graphs, the "dimensional" labeling in hypercubes, the "neighboring" and the "commutative group" labeling in arbitrary graphs, etc. What they have in common is the property of sense of direction [6]. In other words, all those results are evidence of the impact that specific instances of sense of direction has on the communication complexity of distributed problems (in specific graphs).

In this paper, we prove a general result on the impact of sense of direction. We show that, in arbitrary graphs, the availability of *any* sense of direction has a dramatic effect on the communication complexity of several important distributed problems: *Broadcast*, *Depth-First Traversal*, *Election*, *Spanning Tree Construction*, and *Minimum Finding*.

In fact, in the case of Broadcasting and Depth First Traversal, we show that the presence of any sense of direction reduces the message complexity from $\Omega(e)$ to $\Theta(n)$; furthermore, this result holds even if the system is anonymous.

We then show that, with any sense of direction we can construct Election protocols which require only $O(n \log n)$ messages reducing the $\Omega(e+n \log n)$ message complexity for this problem. This results implies a similar bound and reduction for Spanning Tree Construction.

The results presented here completely explain and generalize the results of [13, 16, 22] which now follow as corollaries for specific labelings.

The paper is organized as follows. The framework is described in Section 2 where the basic definitions and properties of sense of direction are given. We prove lower-bounds on the communication complexity in arbitrary labeled networks in Section 3; the upper-bounds obtainable with sense of direction are established in Section 4.

2 The Framework

2.1 Labelings and Sense of Direction

Let $G = (V, E)$ be a graph where nodes correspond to entities and edges correspond to direct bidirectional communication links between entities. Let $E(x)$ denote the set of edges incident to node x.

Every node associates a label, from a finite set of labels $\mathcal{L}$, to each incident edge; moreover, each node associates a local name, from a finite set of names $\mathcal{N}$, to the other nodes in the systems. Notice that these names are not global identifiers.

A distributed system can, thus, be described as a triple (G, λ, β), where: G is the graph representing the communication topology, $\lambda = \{\lambda_x : x \in V\}$ is the set of labeling functions, where $\lambda_x(\langle x, z\rangle)$ denotes the label associated by x to the edge $\langle x, z\rangle \in E(x)$; $\beta = \{\beta_x : x \in V\}$ is the set of naming functions, where $\beta_x(y)$ denotes the name that node x locally associates to node y. We say that a labeling λ is a *Local Orientation* when each node can distinguish among its incident edges; that is when, $\forall x \in V$, $\forall e_1, e_2 \in E(x)$, $\lambda_x(e_1) = \lambda_x(e_2)$ iff $e_1 = e_2$. The *local view* of a node x is the set of names $\{\beta_x(y) : y \in V\}$ used by x to denote the other nodes.

A *path* in G is a sequence of edges in which the endpoint of one edge is the starting point of the next edge. Let $P[x]$ denote the set of all paths with $x \in V$ as a starting point, and let $P[x, y]$ denote the set of paths starting from node $x \in V$ and ending in node $y \in V$. Let Λ be the extension of the labeling function λ from edges to paths.

A *consistent coding function* f of a graph (G, λ, β) is a function that maps the sequences of labels associated to any paths from x to y, to the local name $\beta_x(y)$ used by x to refer to y. More precisely, a coding function f is *consistent* in (G, λ, β) iff $\forall x, y \in V$, $\pi \in P[x, y]$, $f(\Lambda_x(\pi)) = \beta_x(y)$.

A *consistent decoding function* h for f is a function that allows nodes to translate the local views of their neighbors. More precisely, given a consistent coding function f, a decoding function h for f is consistent iff $\forall x$, $\forall \langle x, y\rangle \in E(x)$, $\forall \pi \in P[y, z]$, $h(\lambda_x(\langle x, y\rangle), f(\Lambda_y(\pi))) = \beta_x(z)$.

A system (G, λ, β), has a *Sense of Direction* ($\mathcal{SD}$) iff the following conditions hold [6]:
1) λ is a Local Orientation,
2) there exists a consistent coding function f,
3) there exists a consistent decoding function h for f.
Notice that (2) implies (1).
We shall also say that (f, h) is a sense of direction in (G, λ, β).

2.2 Instances: Chordal and Neighboring $\mathcal{SD}$

To illustrate some of these concepts, we will describe here the two types of labeling which have been studied with respect to solving distributed problems in arbitrary graphs: the *chordal* labeling and the *neighboring* one.

A *chordal* labeling of a graph $G = (V, E)$, with $|V| = n$, is defined by fixing a cyclic ordering of the nodes and labeling each incident link $\langle x, y\rangle$ of x by the distance $\delta(x, y)$ in the above cycle (see Figure 1).

Let λ be a chordal labeling and, $\forall x, y$, let $\beta_x(y) = \delta(x, y)$; note that $\mathcal{N} = \mathcal{L} = \mathbf{Z}_n^+$. This labeling has a sense of direction in (G, λ, β) [6]. For example, a consistent coding function f is

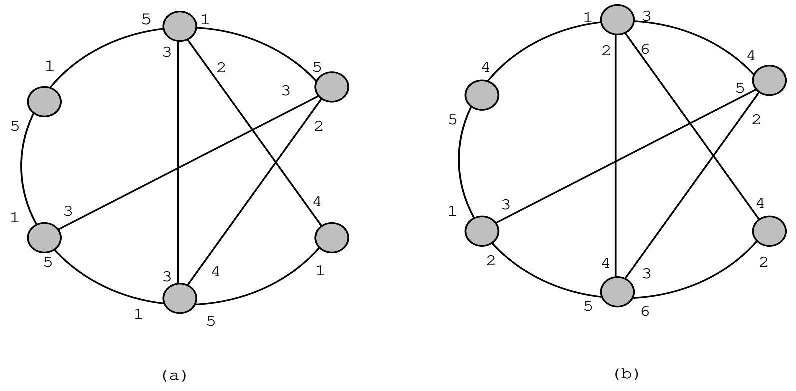

Figure 1: Graph with (a) Chordal and with (b) Neighboring labeling

the following: for any path $\pi = (\langle x_0, x_1\rangle, \langle x_1, x_2\rangle, \ldots, \langle x_{m-1}, x_m\rangle) \in P[x_0]$

$$f(\Lambda_{x_0}(\pi)) = \sum_{i=0}^{m-1} \lambda_{x_i}(\langle x_i, x_{i+1}\rangle) \bmod n$$

It is easy to verify that $f(\Lambda_{x_0}(\pi)) = \delta(x_0, x_m) = \beta_{x_0}(x_m)$, and thus, f is consistent.

The corresponding consistent decoding function is the following: $\forall \langle x_0, y_0\rangle \in E(x_0)$, $\forall \pi \in P[y_0]$

$$\begin{aligned} h(\lambda_{x_0}(\langle x_0, y_0\rangle), f(\Lambda_{y_0}(\pi))) &= \lambda_{x_0}(\langle x_0, y_0\rangle) + f(\Lambda_{y_0}(\pi)) \\ &= \lambda_{x_0}(\langle x_0, y_0\rangle) + \sum_{i=1}^{m-1} \lambda_{y_i}(\langle y_i, y_{i+1}\rangle) \bmod n = \beta_{x_0}(x_m) \end{aligned}$$

Sometimes called *Distance* $\mathcal{SD}$, its impact on the Election problem in arbitrary graphs has been studied in [16], where a $O(n \log n)$ solution algorithm has been shown.

In a *neighboring* labeling, all the links ending in the same node x are labeled with the same label which we shall denote by $l_{(x)}$ (see Figure 1). If, $\forall x, y \in V$, $\beta_x(y) = l_{(y)}$, then there is a $\mathcal{SD}$ in (G, λ, β) [6]. In fact, both the coding function, $\forall \pi \in P[x_0]$, $\pi = (\langle x_0, x_1\rangle, \ldots, \langle x_{m-1}, x_m\rangle)$

$$f(\Lambda_x(\pi)) = \lambda_{x_{m-1}}(\langle x_{m-1}, x_m\rangle)$$

and the decoding function, $\forall \langle x_0, y_0\rangle \in E(x_0)$, $\forall \pi \in P[y_0]$, $\pi = (\langle y_0, y_1\rangle, \ldots, \langle y_{m-1}, y_m\rangle)$

$$h(\lambda_{x_0}(\langle x_0, y_0\rangle), f(\Lambda_{y_0}(\pi))) = f(\Lambda_{y_0}(\pi))$$

are consistent.

This type of labeling is available for example in systems where every node has a distinct global identifier (the *identity*) and knows the identities of all its neighbors (hence the name "neighboring").

It has been shown that, in arbitrary graphs with Neighboring $\mathcal{SD}$, the *Election* problem can be solved with $O(n \log n)$ messages [13], and Depth-First Traversal can be done in $O(n)$ messages [22]. Since distinct global identifiers are easily recoverable from a Neighboring labeling, systems with such labelings are *not* considered anonymous [6].

2.3 The Power of Sense of Direction

In sense of direction, the existence of a consistent decoding function is clearly the crucial property which allows the nodes to solve global problems while working solely and truly in a local mode. In particular, it provides the ability to consistently *decode* the information received from another node.

As an example, consider the situation of a node x wanting to send to its neighbor z information about a node y. Node y is known at x as $\beta_x(y)$; thus, the message sent by x will contain information about a node called "$\beta_x(y)$". Suppose that this information is received by z along the edge locally labeled with $\lambda_z(\langle z, x\rangle)$. Informally, if there is sense of direction, node z, based on the label $\lambda_z(\langle z, x\rangle)$ and on the name $\beta_x(y)$, can deduce that the received information is about the node locally called $\beta_z(y)$. In other words, when there is a sense of direction, each node can consistently *translate the local views* of its neighbors. This can be done even is the system is anonymous.

It is by exploiting the existence of this property in the neighboring, the group, and chordal labelings that [13, 22], [24], and [5, 16], respectively, where able to obtain the reduction in complexity in arbitrary graphs for several important problems In fact, as we show in this paper, it is possible to obtain the same results with *any* sense of direction.

3 Problems and Lower Bounds

3.1 Problems

In this paper we consider a variety of interrelated distributed problems, briefly described below. Each problem can be described as a transformation of the global system configurations; an entity will spontaneously and independently start the transformation process if in state *initiator*.

The *Election* or *Leader Finding* ($\mathcal{LF}$) is the process to transform an initial system configuration, where one or more of the entities are in state *initiator* and the others in a different state (say, *asleep*), into a final system configuration where exactly one entity is in a *leader* state and all the other entities are in state *lost*. Notice that there are no a priori restrictions on which entity can become *leader*, nor on the number of *initiators*. The Election problem is *solvable* (i.e., there exists a deterministic protocol which will always correctly solve the problem) if and only if each entity has a unique identity (e.g., [1]). In this case, the Election problem is computationally equivalent to the problems of *Spanning-Tree Construction* ($\mathcal{SPT}$) and *Minimum Finding* ($\mathcal{M}$) with an arbitrary number of initiators (e.g., [21]).

The *Depth-First Traversal* ($\mathcal{DFT}$) problem consists to arrive from an initial configuration where all the entities except one (the *initiator*) are in the same state to a final configuration in which a depth-first spanning-tree of the graph rooted at the initiator has been identified.

The *Broadcast* ($\mathcal{B}$) problem consists to arrive from an initial configuration where exactly one entity (the *initiator*) has an information, to a configuration where all the entities have the information. The Broadcast problem is a particular instance of the *Wake-up* ($\mathcal{W}$) problem which consists to arrive from an initial configuration where some entities are in the same state (say *awake*) while the others are in a different state (say *asleep*), to a configuration where all the entities are *awake*. Namely, a Broadcast is a Wake-up with only one entity initially *awake*.

3.2 Complexity in Arbitrary Labeled Graphs

Given a distributed problem P, let $\mathcal{C}(P)$ denote the message complexity of P in arbitrary labeled (or unlabeled) graphs. We summarise the following results in Table 1.

	Broadcast ($\mathcal{B}$), Depth First Traversal ($\mathcal{DFS}$),	Election ($\mathcal{LF}$), Spanning Tree ($\mathcal{SPT}$), Min Finding ($\mathcal{M}$)
arbitrary labeling	$\Omega(e)$	$\Omega(e + n\log n)$
sense of direction	$2n-2$	$3n\log n + O(n)$

Table 1: Reductions in message complexity using any sense of direction.

Lemma 1 *[21] If the entities have unique identities,* $\mathcal{C}(\mathcal{LF}) = \mathcal{C}(\mathcal{SPT}) = \mathcal{C}(\mathcal{M})$

Lemma 2 $\mathcal{C}(\mathcal{DFT}) \geq \mathcal{C}(\mathcal{B})$

Proof Trivially, since any traversal performs a broadcast. □

Lemma 3 $\mathcal{C}(\mathcal{LF}) \geq \mathcal{C}(\mathcal{B})$

Proof Any Election protocol solves the Wake-up problem: just map the Election states *initiator*, *leader* and *lost* into the Wake-up state *awake*. Thus, $\mathcal{C}(\mathcal{LF}) \geq \mathcal{C}(\mathcal{W})$. Since a broadcast is just a wake-up with only one initiator, we have $\mathcal{C}(\mathcal{W}) \geq \mathcal{C}(\mathcal{B})$. It follows that $\mathcal{C}(\mathcal{LF}) \geq \mathcal{C}(\mathcal{B})$. □

Theorem 1 *Let the entities have unique identities, let* $P \in \{\mathcal{LF}, \mathcal{SPT}, \mathcal{M}, \mathcal{DFT}, \mathcal{B}\}$*, then, in arbitrary networks of unknown size* $\mathcal{C}(P) \geq \Omega(|E|)$.

Proof By Lemmas 1, 2, 3, it suffices to prove $\mathcal{C}(\mathcal{B}) \geq \Omega(|E|)$.
By contradiction. Assume that there is an execution of a broadcast protocol in some graph $G = (V, E)$ such that no messages were transmitted between two neighbors x and y.
Consider the graph $G' = (V', E')$ defined as follows: $V' = V \cup \{z\}$, where $z \notin V$, and $E' = E \cup \{\langle x, z\rangle, \langle y, z\rangle\} - \{\langle x, y\rangle\}$ (see Figure 2).

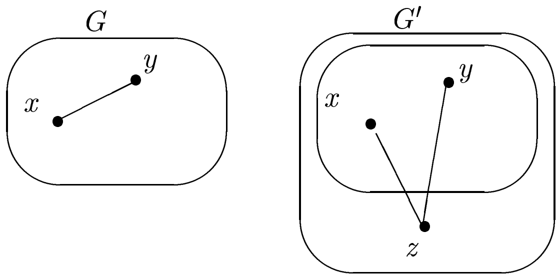

Figure 2: Proof of Theorem 1.

Since the network size is unknown to any solution protocol by hypothesis, and since x and y did not communicate, the protocol cannot distinguish between G and G'; that is, the same execution is possible in G'. But, in this execution, z will not receive any message and, thus, the broadcast will never be correctly completed. □

This establishes an $\Omega(e)$ lower bound on the message complexity of $\mathcal{LF}, \mathcal{SPT}, \mathcal{M}, \mathcal{DFT}$. Notice that the $\Omega(e)$ bound for broadcast and depth first traversal holds even if the entities have unique identities.

Since the proof of the Theorem 1 relies on the fact that $|V|$ is not known to the entities, a natural question is whether knowledge of the number of nodes can be used to escape from the bound. We give now a negative partial answer to this question; in fact, we prove that, for any given size e there exist graphs of almost that size which require $\Omega(e)$ messages, even if the topology of the network (including the number of nodes) is *common knowledge* to all entities.

Given a positive integer n, any integer m such that $n-1 \le m \le \frac{n(n-1)}{2}$ is said to be an *edge-value* for n (or, simply, edge-value).

Theorem 2 *For any n and edge-value m, for any problem $P \in \{\mathcal{B}, \mathcal{LF}, \mathcal{SPT}, \mathcal{M}, \mathcal{DFT}\}$, there exists a graph $G(V,E)$ with n nodes and e edges, $m \le e \le m+n-2$, such that $\mathcal{C}(P)[G] \ge e$, even if n is common knowledge.*

Proof By Lemmas 1, 2, and 3, it is sufficient to show that the theorem holds for $P = \mathcal{B}$. Let n be even and let m be an edge-value for n. Let G' be a k-regular graph on $n' = n-2$ nodes, where $k = \left\lceil \frac{2(m-1)}{n-2} \right\rceil$; clearly, G' has $e' = \frac{k(n-2)}{2}$ edges. Consider now the graph G obtained from G' by adding two new nodes u and v, removing an edge (x,y), and adding the new edges (x,u) and (y,v) (see Figure 3). The new graph G has n nodes and $e = e'+1 = \frac{k(n-2)}{2}+1$ links. Notice that, because of the ceiling operator in the definition of k, we have $m \le e \le m+n-2$. It is easy to see that in G broadcasting requires $\Theta(e)$ messages even if n is common knowledge: the two leaf nodes must receive the message; however, every internal node has the same degree and it is incapable of distinguishing whether or not it is connected to a leaf. A similar construction and argument holds when n is odd. □

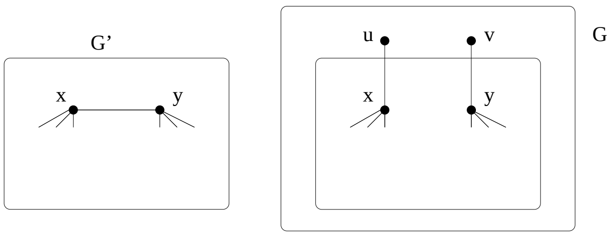

Figure 3: Proof of Theorem 2.

That is, for any given size e there exist graphs of almost that size which require $\Omega(e)$ messages, even if the topology of the network is common knowledge.

4 Impact of Sense of Direction

4.1 Distributed Depth First Traversal

In this Section, we consider the distributed problem of constructing a Depth First spanning tree of the graph G representing the system.

In arbitrarily labeled graphs, this problem requires $\Omega(|E|)$ messages (see Lemma 2 and Theorem 1) even if the system is not anonymous; this bound can be easily achieved (e.g., [3]). An improvement in the communication complexity has been shown to exist if each entity has a distinct identity *and* knows the identities of all its neighbors; in this case, this problem can be solved with $O(n)$ messages [22]. This result implies that, in presence of neighboring sense of direction (discussed in Section 2.2), $\mathcal{DFT}$ can be performed in $O(n)$ messages. Recall that graphs with neighboring labelings are *not* anonymous. A similar reduction has been shown to exist in the presence of chordal sense of direction (discussed in section 2.2) [5].

In the following we show that Depth First Traversal can be performed in $O(n)$ messages with *any* sense of direction, even if the graph is *anonymous*.

Algorithm DFT

The algorithm is token-based; at any point in time, at most one node has the "token". Initially, the token is held by the initiator node. During each step of the algorithm, the node holding the "token" transmits it to one of its "untraversed" neighbors; if all the neighbors have been already traversed, the node "backtracks" the token to the node from which the token was first received. The algorithm terminates when the owner of the "token" detects that all the nodes have been traversed. With respect to a DF-traversal of G, the edges of G can be partitioned into "tree edges" and "back edges". To achieve the desired complexity, the "token" must be transmitted only along the tree edges. However, a node does not know a priori which of its *unused* edges (i.e., on which no message has been transmitted or received) is a "tree edge" or a "back edge". The properties of sense of direction are used to determine exactly this, without message transmissions on "back edges".

Let (G, λ, β) be a system with sense of direction. Given a set $S \subseteq V$ of nodes, let $V_x(S) = \{\beta_x(s) : s \in S\}$ be S as viewed by node x; given a neighbor y of x, let $H_y(S) = \{h(\lambda_y(\langle y, x\rangle), q) : q \in V_x(S)\}$, where h is the consistent decoding function.

Property 1 *For every $S \subseteq V, x \in V, y \in E(x)$, $H_y(S) = V_y(S)$.*

Proof It directly follows from the definition of consistent coding and decoding function. □

In other words, if a node x transmits its view (i.e., its local names) of a set S of nodes to a neighbor y, the neighbor is capable of correctly constructing its view of the same set. This implies that, even in absence of distinct identities, if a node sends a set of local names, the receiver node can understand to which nodes the sender is referring to.

Furthermore, $F(x) = \{f(\lambda_x(\langle x, y\rangle)) : y \in E(x)\} = \{\beta_x(y) : y \in E(x)\}$, where f is the consistent coding function. Thus, for any $V' \subseteq V_x$ subset of the local view of $x \in V$ it is possible to correctly determine whether or not a neighbor is in V'.

Summarising, if a node transmits its view of a set of nodes, the receiver is capable of correctly determining whether or not an incident link is connected to a node in that set. The algorithm is shown in Figure 4.

Theorem 3 *In an anonymous system with sense of direction, Depth First Traversal can be performed using at most $2n - 2$ messages and time.*

Proof Using Property 1, messages in the algorithm described above are transmitted only along the tree edges. In particular, a "token" traverses each tree edge twice; the first time when the edge is first "traversed", the second time in the backtracking phase. □

It is crucial to note that this bound can be reached even in absence of distinct node identities. This must be confronted with the algorithm of [22] which assumed distinct node identities as well as a priori knowledge at each entity of the identity of the neighbors; the size of the messages in the two algorithms is the same.

4.2 Broadcast

With an arbitrary labeling, the complexity of the broadcast problem is $\Omega(|E|)$ when the number of nodes n is unknown (by Theorem 1). If n is common knowledge, for any given size e there exists graphs which require $\Omega(e)$ messages, (by Theorem 2).

When the labeling is a chordal sense of direction, there exists an algorithm for solving the Broadcast problem in $O(n)$ [5] and the same complexity is achievable when the labeling is a commutative group sense of direction [24]. We can now generalize the existing results which hold for specific sense of directions.

Theorem 4 *In an anonymous network with sense of direction, the broadcast can be performed using at most* $2n - 2$ *messages.*

Proof It follows from Theorem 3. □

Notice that the size of the messages used is the same as the one in the algorithms for the specific senses of direction of [16, 24].

4.3 Election, Spanning Tree Construction and Minimum Finding

Consider now non-anonymous systems; i.e., where each entity has a unique value (or identity). It is well known that in arbitrary networks the Election problem requires $\Omega(e+n \log n)$ messages (e.g., [21]) and such a bound is achievable (e.g., [8]).

With Chordal $\mathcal{SD}$ [16] and Neighboring $\mathcal{SD}$ [13, 16], this problem can be solved in $O(n \log n)$ messages. In the following, we prove that the availability of *any* $\mathcal{SD}$ allows to achieve the same bound.

Lemma 4 *[11] The Election problem can be solved using* $(t(n) + n)(\log n + 1)$ *or* $(t(e) + n)$ *messages where* $t(n)$ *(resp.,* $t(e)$*) is the message complexity of graph traversal with* n *nodes (resp.,* e *edges).*

Theorem 5 *Let* $P \in \{\mathcal{LF}, \mathcal{SPT}, \mathcal{M}\}$*, With sense of direction, the* P *can be solved using* $3n \log n + O(n)$ *messages.*

Proof By Lemma 4 and Theorem 1.

□

Because of Lemma 1, the same bounds hold also for $\mathcal{SPT}$ and for $\mathcal{M}$. Notice that the size of the messages is the same as in [13, 16].

Local information:
•$Unused_x$: list of labels of edges at node x, initialized to $\{\lambda_x(\langle x,y\rangle) : y \in E(x)\ \}$.
•$Tree$: list of local node names which represents the set of nodes reached by now.
•$Parent_x$: label of the edge to the parent node in the rooted tree being constructed.

procedure *INITIATE*
begin
/* entity initiator */
  $Tree := \{x\}$
  SEND *(TOKEN,Tree)* on edge $\lambda_x(\langle x,y\rangle) \in Unused_x$
  $Unused_x := Unused_x - \{\lambda_x(\langle x,y\rangle)\}$
end INITIATE

Upon RECEIPT of *(TOKEN,Tree)* on edge labeled $\lambda_x(\langle x,y\rangle)$ by node x
  $Parent_x := \lambda_x(\langle x,y\rangle)$
  $Unused_x := Unused_x - \{\lambda_x(\langle x,y\rangle)\}$
  DECODE($Tree, \lambda_x(\langle x,y\rangle)$)
  $Tree := Tree \cup \{x\}$
  for each edge $\lambda_x(\langle x,z\rangle) \in Unused_x$
    if $\beta_x(z) \in Tree$ **then**
      $Unused_x := Unused_x - \{\lambda_x(\langle x,z\rangle)\}$
  if $Unused_x \neq \emptyset$ **then**
    SEND *(TOKEN,Tree)* on an edge labeled $\lambda_x(\langle x,z\rangle) \in Unused_x$
    $Unused_x := Unused_x - \{\lambda_x(\langle x,z\rangle)\}$
  else /* backtrack */
    if $|Tree| = n$ **then**
      Terminate
    else
      SEND *(BACKTRACK,Tree)* on edge labeled $Parent_x$
    fi
  fi
end TOKEN

Upon RECEIPT of *(BACKTRACK,Tree)* on edge labeled $\lambda_x(\langle x,y\rangle)$ by node x
  DECODE($Tree, \lambda_x(\langle x,y\rangle)$)
  for each edge $\lambda_x(\langle x,z\rangle) \in Unused_x$
    if $\beta_x(z) \in Tree$ **then**
      $Unused_x := Unused_x - \{\lambda_x(\langle x,z\rangle)\}$
  if $Unused_x \neq \emptyset$ **then**
    SEND *(TOKEN,Tree)* on an edge labeled $\lambda_x(\langle x,z\rangle) \in Unused_x$
    $Unused_x := Unused_x - \{\lambda_x(\langle x,z\rangle)\}$
  else /* backtrack again */
    SEND *(BACKTRACK,Tree)* on edge labeled $Parent_x$
  fi
end BACKTRACK

procedure DECODE($Tree, \lambda_x(\langle x,y\rangle)$)
begin
  foreach $r = \beta_y(z) \in Tree$ **do** $r := h(\lambda_x(\langle x,y\rangle), \beta_y(z))$
end DECODE

Figure 4: DFT Algorithm

References

[1] D. Angluin. Local and global properties in networks of processors. In *Proc. of 12th A.C.M. Symposium on Theory of Computing*, pages 82–93, 1980.

[2] H. Attiya, J. van Leeuwen, N. Santoro, and S. Zaks. Efficient elections in chordal ring networks. *Algorithmica*, 4:437–446, 1989.

[3] T.Y. Cheung. Graph traversal techniques and the maximum flow problem in distributed computation. *I.E.E.E. Transactions on Software Engineering*, 9:504–512, 1983.

[4] P. Flocchini and B. Mans. Optimal election in labeled hypercubes. *Journal of Parallel and Distributed Computing*, 33(1):76–83, 1996.

[5] P. Flocchini, B. Mans, and N. Santoro. Distributed traversal and broadcasting in arbitrary network with distance sense of direction. In *Proc. of 9th International Symposium on Computer and Information Sciences*, pages 196–203, 1994.

[6] P. Flocchini, B. Mans, and N. Santoro. Sense of direction: formal definition and properties. In *Proc. of 1st Colloquium on Structural Information and Communication Complexity*, International Informatic Series 1, pages 9–34, Ottawa, Canada, 1994. Carleton University Press.

[7] P. Flocchini and A. Roncato. Computing on anonymous networks with sense of direction. In *Proc. of 3rd Colloquium on Structural Information and Communication Complexity*, International Informatic Series 6, pages 32–46, Siena, 1996.

[8] R.G. Gallager, P.A. Humblet, and P.M. Spira. A distributed algorithm for minimum spanning tree. *A.C.M. Transactions on Programming Languages and Systems*, 5(1):66–77, 1983.

[9] A. Israeli, E. Kranakis, D. Krizanc, and N. Santoro. Time-message trade-offs for the weak unison problem. In *Conference on Algorithms and Complexity*, Lecture Notes in Computer Science 778, pages 167–178. Springer-Verlag, 1994.

[10] T.Z. Kalamboukis and S.L. Mantzaris. Towards optimal distributed election on chordal rings. *Information Processing Letters*, 38:265–270, 1991.

[11] E. Korach, Kutten, and S. Moran. A modular technique for the design of efficient distributed leader finding algorithms. *A.C.M. Transactions on Programming Languages and Systems*, 12(1):84–101, January 1990.

[12] E. Korach, S. Moran, and S. Zaks. Tight lower and upper bounds for a class of distributed algorithms for a complete network of processors. In *Proc. of 3rd Symposium on Principles of Distributed Computing*, pages 199–207, Vancouver, 1984.

[13] I. Lavallée and G. Roucairol. A fully distributed (minimal) spanning tree algorithm. *Information Processing Letters*, 23:55–62, Aug 1986.

[14] M.C. Loui, T.A. Matsushita, and D.B. West. Election in complete networks with a sense of direction. *Information Processing Letters*, 22:185–187, 1986. see also Information Processing Letters, vol.28, p.327, 1988.

[15] B. Mans. Optimal distributed algorithms in unlabeled tori and chordal rings. In *Proc. of the 3rd International Colloquium on Structural Information and Communication Complexity,*, International Informatics Series 6, pages 17–31, Siena, Italy, June 1996. Carleton Press.

[16] B. Mans and N. Santoro. On the impact of sense of direction in arbitrary networks. In *Proc. of 14th International Conference on Distributed Computing Systems*, pages 258–265, Poznan, 1994.

[17] G.H. Masapati and H. Ural. Effect of preprocessing on election in a complete network with a sense of direction. In *Proc. of I.E.E.E. International Conference on Systems, Man and Cybernetics*, volume 3, pages 1627–1632, 1991.

[18] T. Masuzawa, N. Nishikawa, K. Hagihara, and N. Tokura. A fault-tolerant algorithm for election in complete networks with a sense of direction. *Systems and Computers in Japan*, 22(12):11–22, 1991.

[19] Yi Pan. An improved election algorithm in chordal ring networks. *International Journal of Computer Mathematics*, 40(3-4):191–200, 1991.

[20] S. Robbins and K.A. Robbins. Choosing a leader on a hypercube. In *Proc. of International Conference on Databases, Parallel Architectures and their Applications*, pages 469–471, Miami Beach, 1990.

[21] N. Santoro. On the message complexity of distributed problems. *Journal of Computing and Information Science*, 13:131–147, 1984.

[22] M.B. Sharma, S.S. Iyengar, and N.K. Mandyam. An efficient distributed depth-first-search algorithm. *Information Processing Letters*, 32:183–186, September 1989. see also Information Processing Letters, vol.35, p.55, 1990.

[23] G. Singh. Leader election in complete networks. In *Proc. of 11th Symposium on Principles of Distributed Computing*, pages 179–190, 1992.

[24] G. Tel. Sense of direction in processor networks. In *Theory and Practice of Informatics*, Lecture Notes in Computer Science 1012, pages 50–82. Springer-Verlag, 1995.