

SENSE OF DIRECTION AND COMMUNICATION COMPLEXITY IN DISTRIBUTED NETWORKS

Nicola Santoro*, Jorge Urrutia⁺ and Shmuel Zaks⁺⁺

*Distributed Computing Group, Carleton University, Ottawa, Canada

⁺Department of Computer Science, Ottawa University, Ottawa, Canada

⁺⁺Department of Computer Science, Technion, Haifa, Israel

ABSTRACT

The connection between sense of direction and communication complexity in distributed complete networks is studied for two basic problems: finding a minimum-weight spanning tree (*MST*) and finding a spanning tree (*SP*). Several models of the complete network are defined, the difference being the amount (and type) of sense of direction available, forming a hierarchy which includes the models previously studied in the literature. It is shown that to move up in the hierarchy might require $\Omega(n^2)$ messages in the worst case. It is shown that the existing $O(n)$ bound for *SP* can still be achieved at a lower level in the hierarchy; and that the $\Omega(n^2)$ bound for *MST* still holds at a higher level in the hierarchy.

1. INTRODUCTION

Consider a network of n processors. Each processor has a distinct identity of which it alone is aware, and has available some labeled direct communication lines to other (possibly, all) processors; it also knows the (non-negative) cost associated with each such line. The network can be viewed as an undirected graph $G=(V,E)$ where $|V|=n$. Communication is achieved by sending messages along the communication lines. It is assumed that messages arrive, with no error, after a finite but otherwise arbitrary delay, and are kept in order of arrival in a queue until processed. All processors execute the same algorithm, which involves local processing as well as sending a message to a neighbor and receiving a message from a neighbor. Any non-empty set of processors may start the algorithm. In this context, two basic problems are finding a minimum-weight spanning-tree (*MST*) and distinguishing a unique processor (a *leader*); the latter is very closely related to the problem of finding a spanning tree (*SP*). These problems have been mostly studied in the literature for circular networks and complete networks.

The interest in the circular network derives from the fact that it is the simplest symmetrical graph. In this graph both problems are equivalent and have been extensively studied; different bounds

have been established depending on whether the lines are unidirectional [DKR, P] or bidirectional, and whether (in the latter case) there exists a strong and consistent sense of direction (i.e., "left" means the same to all processors) [F, KRS]; in all cases the message complexity is $\Theta(n \log n)$, and the difference is felt only in the multiplicative constant.

The interest in the complete network derives from the fact that it is the 'densest' symmetrical graph, each processor having a direct connection to all others. The $O(n \log n + |E|)$ bounds of [GHS] for minimum-weight spanning-tree construction in arbitrary graph implies an $O(n^2)$ bound on both *SP* and *MST* in the case of complete networks. By exploiting the 'density' of the complete graph, in [KMZ₁] it has been shown that $\Theta(n \log n)$ messages are sufficient and in the worst case necessary to solve *SP* in a complete network. On the other hand, in [KMZ₂], it has been shown that solving *MST* in a complete network might require $\Omega(n^2)$ messages. Thus, the two problems are not equivalent in a complete network.

These results for complete networks do not assume any sense of direction, where (informally) sense of direction refers to the knowledge that processors have about the labeling of their incident lines. However, it has been shown in [LMW] that, if the processors have a strong and consistent sense of direction (to be defined later), then *SP* can be solved in $\Theta(n)$ messages; that is, presence of sense of direction drastically reduces the message complexity of the *SP* problem in the complete graph. On this basis, the following questions naturally arise (e.g., see [S]): does sense of direction influence also the complexity of *MST*? how much sense of direction is actually needed to achieve the $O(n)$ bound for *SP*?

The contribution of this paper is to shed some light on the relationship between sense of direction and communication complexity for these two problems in complete networks. Several models of the complete network are defined, the difference being the amount (and type) of sense of direction available; in these models not all the lines connected to a processor are undistinguishable. These models form a hierarchy which include the models previously studied in the literature; and it is shown that to acquire more sense of direction (i.e., to move up in the hierarchy) $\Omega(n^2)$ messages might be required in the worst case. The impact of sense of direction on the communication complexity of *SP* and *MST* is then investigated. In particular, it is shown that the $O(n)$ bound for *SP* can still be achieved with some degree of ambiguity in the sense of direction (i.e., at a lower level in the hierarchy); and that the $\Omega(n^2)$ bound for *MST* still holds in spite of additional sense of direction (i.e., at a higher level in the hierarchy).

2. THE MODELS

The communication network is a complete graph $G = (V,E)$. Each node has a distinct identity of which it alone is aware; without loss of generality and unless otherwise specified, assume that the identities are integers in $[1,n]$, where $n=|V|$. Furthermore, a node has a label associated with each incident edge. Let $l(i,j)$ denote the label associated at node i with edge (i,j) ; again, without loss of generality, assume that the labels used at node i are integers in $[1,n-1]$. Depending on which property is assumed to exist on the labels, different models can be defined and have been considered. Let $N_i = \{1,2,\dots,n\} - \{i\}$, and let $\{\{\}\}$ denote a multiset.

The following models of the complete network are considered here:

[G] *The general model*: for every processor i

$$\{l(i,j) \mid j \in N_i\} = N_n.$$

That is, all the edges incident to a node have different labels; however, no relationship is known between $l(i,j)$ and $l(j,i)$. This is the model discussed in [KM1, KMZ2].

[CS] *The consistent strong model*: for every i and j

$$\{l(i,j) \mid j \in N_i\} = N_n$$

$$l(i,j) = (j-i) \bmod n$$

$$l(i,j) + l(j,i) = n.$$

That is, at each node the labels of the incident edges are distinct and denote the distance between that node and its neighbours with respect to a predefined directed Hamiltonian circuit (the same for all nodes). This is the model discussed in [LMW].

[CW] *The consistent weak model*: for every i and j

$$\{l(i,j) \mid j \in N_i\} \subseteq \{\{1,1,2,2,\dots,n-1,n-1\}\}$$

$$l(i,j) = \min \{(i-j) \bmod n, (j-i) \bmod n\}$$

$$l(i,j) = l(j,i).$$

That is, at each node the labels of the incident edges denote the minimum distance between that node and its neighbours with respect to a predefined Hamiltonian circuit. It also follows that, at every node, each label (except one, if n is even) is associated to two incident edges; furthermore, each edge has the same label at both end nodes.

[IW] *The inconsistent weak model*: for every i and j

$$\{l(i,j) \mid j \in N_i\} \subseteq \{\{1,1,2,2,\dots,n-1,n-1\}\}$$

$$l(i,j) = l(j,i).$$

That is, it is possible to have the same label associated to two different edges incident on the same node; however, each edge has the same label at both end nodes.

Depending on whether the model is G , CS , CW , or IW , we shall say that there exist a G -labeling, a CS -labeling, a CW -labeling, or an IW -labeling of the communication lines, respectively.

3. HIERARCHY

It is clear that every CS -labeling is also a CW -labeling, every CW -labeling is an IW -labeling, and all of them are also G -labelings. The situation is depicted in the following diagram:

$$CS \text{ -----} > CW \text{ -----} > IW \text{ -----} > G$$

which also express the hierarchy of sense of direction defined by the labelings. The amount of information in any two levels of the hierarchy is significantly different, and moving from one labeling to another is quite costly from a computational point of view. In fact, we show that $\Omega(n^2)$ messages might be needed to move upward in the hierarchy.

Theorem 1 In the worst case, $\Omega(n^2)$ messages are required to transform a G -labeling to an IW -labeling, an IW -labeling to a CS -labeling, and a CW -labeling to a CS -labeling, with the minimum number of label changes.

Proof. ($G \text{ ---} > IW$)

Let A be an algorithm which always correctly transforms a G -labeling into an IW -labeling. Execute A on a complete graph G with a G -labeling which is already a IW -labeling of G ; thus A will terminate without changing any label. Assume that there are two edges (x,y) and (x,z) , $y \neq z$ and $l(x,y) \neq l(x,z)$, along which no messages were transmitted during the execution of A . Consider now the graph G' obtained by exchanging $l(x,y)$ and $l(x,z)$ in the IW -labeling of G ; note that this labeling of G' is no longer an IW -labeling. Since everything is the same except for the exchange (at a single node) of the labels of two edges along which no message was transmitted, the two graphs G and G' are undistinguishable for algorithm A ; thus, the same execution is possible in both G and G' terminating, in the latter case, with a labeling which is not IW : a contradiction. Therefore, at least one out of every couple of edges (x,y) and (x,z) must carry a message; that is, $\Omega(n^2)$ messages must be transmitted.

($IW \text{ ---} > CW$)

Let A be an algorithm which always correctly transforms a IW -labeling to a CW -labeling. Execute A on a complete graph G with $n=2k$ (for an odd k) nodes and with a IW -labeling which is already a

CW -labeling of G ; thus, A will terminate without changing any label. Assume that no messages were transmitted along the edges $(x, x+i)$, $(x+i, x+k)$, $(x+k, x+k+i)$ and $(x+k+i, x)$ for a node $x \in \{0, 1, \dots, n-1\}$ where $i \in \{1, 2, \dots, k-1\}$ and all arithmetic is modulo n . Consider now the graph G' obtained by exchanging the labels of those edges in the IW -labeling of G as follows: $l(x, x+i) = l(x+k, x+k+i) = k-i$ and $l(x+i, x+k) = l(x+k+i, x) = i$. Since k is odd, it follows that $k-i \neq i$; hence this labeling of G' is no longer a CW -labeling. Since everything is the same except for the exchange (at two nodes only) of the labels of edges along which no message was transmitted, the two graphs G and G' are undistinguishable for algorithm A ; thus, the same execution is possible in both G and G' terminating, in the latter case, with a labeling which is not CW : a contradiction. Therefore, at least one out of every four edges (called a 'square') of the form $(x, x+i)$, $(x+i, x+k)$, $(x+k, x+k+i)$ and $(x+k+i, x)$ must carry a message; since there are $O(k^2)$ such squares for $x \in \{0, 1, \dots, k-1\}$ and $i \in \{1, 2, \dots, k-1\}$, and since $k = n/2$, it follows that $\Omega(n^2)$ messages must be transmitted.

($CW \rightarrow CS$)

Let A be an algorithm which always correctly transforms a CW -labeling to a CS -labeling. Let G have $n = 2k+1$ nodes and a CW -labeling where $\{l(i, j) \mid j \in N_i\} = \{\{1, 1, 2, 2, \dots, k, k\}\}$. The execution of A on G will obviously change only k labels at each node; namely, at each node x , of the two edges labeled i ($i=1, \dots, k$) one will be relabeled $n-i$ and the other will be unchanged. Assume that, for some node x , no message is transmitted on the edges $\mu = (x, x+i)$ and $\partial = (x-i, x)$, where $i \in \{1, \dots, k\}$ and all arithmetic is modulo n . Consider the graph G' obtained from G by exchanging μ and ∂ ; the identical CW -labeling is obviously possible in both G and G' . In this case, since everything is the same except for the exchange of those two edges along which no messages were transmitted, the two graphs G and G' are undistinguishable for algorithm A ; thus, the same execution is possible in both G and G' terminating, in the latter case, with a labeling which is not CS : a contradiction. Therefore, at least one out of every couple of edges $(x, x+i)$ and $(x, x-i)$ must carry a message; that is, $\Omega(n^2)$ messages must be transmitted.

□

4. COMMUNICATION COMPLEXITY

In this section we study the complexity of the problems of constructing a spanning tree (SP) and a minimum-weight spanning tree (MST) in a complete network with respect to the hierarchy of models defined above. In figure 1 at the end of the paper, the existing bounds for these two problems are shown; we strengthen both results by showing that the $\Omega(n)$ bound for SP can be achieved with

less sense of direction, and that the $\Omega(n^2)$ bound for *MST* still holds in spite of more available sense of direction.

Theorem 2 The message complexity of the *SP* problem for the *CW* (and, thus, for the *CS*) model is $\Theta(n)$.

Proof (sketch). Since $\Omega(n)$ is an obvious lowerbound and since once a leader is elected $O(n)$ messages suffice to construct a spanning tree, to prove the theorem is sufficient to design an algorithm for finding a leader in a complete network with a *CW*-labeling using at most $O(n)$ messages. First note that the edges labeled 1 form a bidirectional ring without a global sense of direction (i.e., processors know that they are on such a ring, but cannot distinguish their left from their right). We modify the election algorithm of [F] for such rings so to take advantage of the fact that more communication lines are available. In this algorithm, each initiator (there could be one or more initiator, possibly all nodes) becomes active and sends a message carrying its identity in both directions; upon reception of such a message, a non-initiator becomes passive and acts as a relay; the messages travel along the ring until they encounter another active node. An active node waits until it receives a message from both directions; it will stay active iff both received identities are smaller than its own; if the identities are its own, it becomes elected; otherwise it becomes passive, disregards the message and acts as a relay of future messages. In case it remains active, a node will start the process again (i.e., send a message in both directions, etc.). This algorithm, when executed on a ring will exchange $O(n \log n)$ messages; this is due to the fact that, although the cumulative sum of all active nodes in all iterations of the algorithm is just linear, all nodes (including the passive ones) will transmit a message in each stage. In a complete network with a *CS*-labeling, by adding a counter to each message (describing the distance travelled by the message), it is possible for the active node receiving the message to determine the direct link connecting it to the originator of the message; this link (called *cord*) can then be used as a shot-cut in the next stage. This is essentially the technique of the algorithm described in [LMW]. In presence of a *CW*-labelling, the active node receiving a message with such a counter has two edges which could possibly be the correct cord. This ambiguity is easily resolved by having the node sending a *check* message (containing the identity of the node which should be at the other end of the cord) over the two candidate edges; the correct node will then reply with an *acknowledgment* which will break the ambiguity. It is easy to show that the edges carrying 'normal' messages in this modified algorithm necessarily form a planar graph (hence their number is of $O(n)$), and each is carrying a constant number of messages, which amounts to a total of $O(n)$ messages. Furthermore, the number of *check* and *acknowledgment* messages in each stage is linear in the number of nodes active in that

stage; since the cumulative sum of all active nodes in all iterations of the algorithm is $O(n)$, the bound stated in the theorem follows. \square

It should be pointed out that the described algorithm can easily be modified so to have a smaller multiplicative constant in the bound; further improvement can be obtained by using as a basis a more efficient election algorithm for bidirectional rings without sense of orientation (e.g., [KRS]).

Theorem 3 The worst-case message complexity of the *MST* problem for the *IW* model is $\Omega(n^2)$.

Proof. Let G be the complete graph on $n=2^{p+1}$ nodes; partition the nodes into two sets $A=\{a_1, a_2, \dots, a_m\}$ and $B=\{b_1, b_2, \dots, b_m\}$ where $m=n/2$, and construct a symmetric latin square L of size m with the first column being the identity permutation (i.e., each row and each column of L is a permutation of $\langle 1, 2, \dots, m \rangle$, $L(i,k)=j$ iff $L(j,k)=i$, and $L(i,1)=i$). Define four edges of the form (a_i, a_j) , (b_i, b_j) , (a_i, b_j) , (a_j, b_i) to be a r -square if $L(i, 2r)=j$ and $L(j, 2r)=i$, $r=1, 2, \dots, m-1$; it is not difficult to see that, by the definition of L , there is a total of exactly $m(m-1)/2$ such squares and they are all edge-disjoint. Construct a *IW*-labeling of G as follows: for each r -square ($r=1, \dots, m-1$) assign label r to each edge in the square at both end-points; to the remaining edges (i, i) assign label m . Assign now weight 0 to all edges connecting nodes in the same partition, and weight 1 to all other edges; obviously, the weight of the minimum spanning tree must be 1. Let A be an algorithm which always correctly construct the minimum-weight spanning tree of a complete graph; execute A on G . Assume that there is an r -square (a_i, a_j) , (b_i, b_j) , (a_i, b_j) , (a_j, b_i) where no messages were transmitted during this execution, $r \in \{1, \dots, m-1\}$. Consider the graph G' obtained from G by changing the weights of the edges on that square; that is, by setting $w(a_i, a_j)=w(b_i, b_j)=1$ and $w(a_i, b_j)=w(a_j, b_i)=0$. Obviously, the weight of the minimum spanning tree of G' is 0. Since everything is the same except for the exchange of the weights of four edges along which no message was transmitted, the two graphs G and G' are undistinguishable for algorithm A ; thus, the same execution is possible in both G and G' terminating, in the latter case, with a spanning tree of weight 1: a contradiction. Therefore, at least one out of the four edges in a square must carry a message; since there are $m(m-1)/2=(n^2-2n)/8$ such squares, the theorem follows. \square

A summary of the above results (designated with a $*$) as well as the existing ones is shown in the following figure.

Model	SP	MST
CS	$\Omega(n)$ [LMW]	?
CW	$\Omega(n)$ [*]	?
IW	?	$\Omega(n^2)$ [*]
G	$\Omega(n \log n)$ [KMZ1]	$\Omega(n^2)$ [KMZ2]

Figure 1. New and existing bounds for the different models

CONCLUDING REMARKS

The problem considered in this paper is the relationship between sense of direction and communication complexity in distributed networks. This problem seems to be definitely important in that drastic reduction in complexity can be associated to presence of sense of direction; on the other hand, as argued in [S], 'insensitivity' of a problem to sense of direction seems to indicate the presence of an inherent complexity of the problem.

The results presented here should be seen as preliminary insights in this problem. A main obstacle to a deeper understanding seems to be the lack of an accurate definition of the notion of 'sense of direction'. The definitions given here for complete networks (the four models) have been sufficient to shed some light at least for the two problems considered (see the above figure); however, much is still unknown even within these models (e.g., the '?' entries in the above figure).

REFERENCES

- [DKR] D. Dolev, M. Klawe and M. Rodeh, "An $O(n \log n)$ unidirectional distributed algorithm for extrema finding in a circle", *Journal of Algorithms* 3, (1982), 245-260.
- [F] W.R. Franklin, "On an improved algorithm for decentralized extrema-finding in a circular configuration of processes", *Communications of the ACM* 25, 5 (May 1982), 336-337.
- [GHS] R.G. Gallager, P.A. Humblet and P.M. Spira, "A distributed algorithm for minimum spanning tree", *ACM Transactions on Programming Languages and Systems* 6, 1 (1983), 66-77.
- [KMZ1] E. Korach, S. Moran and S. Zaks, "Tight lower and upper bounds for distributed algorithms for a complete network of processors", *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, Vancouver, August 1984, 199-207.
- [KMZ2] 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", *Proc. 4th ACM Symposium on Principles of Distributed Computing*, Minaki, August 1985, 277-286.
- [KRS] E. Korach, D. Rotem, N. Santoro, "Distributed election in a circle without a global sense of orientation", *Int. J. of Computer Mathematics* 16, (1984), 115-124.
- [LMW] M. Loui, T.A. Matsushita, D. West, "Elections in a complete network with a sense of direction", *Information Processing Letters*, to appear.
- [P] G.L. Peterson, "An $O(n \log n)$ unidirectional algorithm for the circular extrema problem", *ACM Transactions on Programming Languages and Systems* 4, 4 (Oct. 1982), 758-762.
- [S] N. Santoro, "Sense of direction, topological awareness and communication complexity", *SIGACT News* 16, 2 (1984), 50-56.