Sense of Direction in Distributed Computing

Paola Flocchini[1], Bernard Mans[2], Nicola Santoro[3]

[1] Département d'Informatique Université du Quebec à Hull
(flocchini@uqah.uquebec.ca)
[2] Department of Computing, School of MPCE, Macquarie University, Sydney,
(bmans@mpce.mq.edu.au)
[3] School of Computer Science, Carleton University, Ottawa
(santoro@scs.carleton.ca)

1 Introduction

1.1 Distributed Model

There are numerous models for distributed systems, differing from one another on a large number of important factors and parameters. We shall restrict ourselves to systems based on the point-to-point message-passing model: there is no common memory, and a node may communicate directly only with its direct neighbors by exchanging messages. The communication topology of the system can be described by an edge labeled graph where nodes correspond to entities and edges correspond to direct communication links between entities [32]. Let $G = (V, E)$ be the graph; let $E(x)$ denote the set of edges incident to node $x \in V$, and $d(x) = |E(x)|$ the degree of x. Each node has a local label (port number) associated to each of its incident edges. Given a set Σ of labels, a *local orientation* of $x \in V$ is any injective function $\lambda_x : E(x) \to \Sigma$ which associates a distinct label $l \in \Sigma$ to each edge $e \in E(x)$. The set $\lambda = \{\lambda_x : x \in V\}$ of local labeling functions will be called a *labeling* of G, and by (G, λ) we shall denote the corresponding *(edge-)labeled graph*.

Any property of the labeling λ can be exploited to improve the performance of the system, e.g., by reducing the amount of communication required to perform some distributed tasks. The most basic property is *Local Orientation*: the capacity at each node to distinguish between the incident links; by definition, we only consider labelings having this property. Another interesting property is *Edge Symmetry*: there exists a bijection $\psi : \Sigma \to \Sigma$ such that for each $\langle x, y\rangle \in E$, $\lambda_y(\langle y, x\rangle) = \psi(\lambda_x(\langle x, y\rangle))$; ψ will be called the edge-symmetry function. The particular case of edge symmetry called *Coloring* where the edge-symmetry function is the identity (i.e., the labels on the two sides of each edge are the same) is also of interest. For specific labelings in specific topologies, some other properties have been extensively studied; e.g., *Orientation* in ring networks with "left-right" labeling or in tori with the compass labeling.

Without any doubt, the property of labeled graphs which has been shown to have a definite impact on computability and complexity, and whose applicability ranges from the analysis of graph classes to distributed object systems, is *Sense of Direction*. In the rest of the paper, we will provide some introduction and pointers to the relevant results and literature.

1.2 Sense of Direction

Informally, in a labeled a graph (G, λ). the labeling is a (weak) sense of direction if it is possible to understand, from the labels associated to the edges, whether different walks from a given node x end in the same node or in different nodes. A *walk* π 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 the walks starting from $x \in V$, $P[x, y]$ the set of walks starting from $x \in V$ and ending in $y \in V$. Let $\Lambda_x : P[x] \to \Sigma^+$ and $\Lambda = \{\Lambda_x : x \in V\}$ denote the extension of λ_x and λ, respectively, from edges to walks; let $\Lambda[x] = \{\Lambda_x(\pi) : \pi \in P[x]\}$, and $\Lambda[x, y] = \{\Lambda_x(\pi) : \pi \in P[x, y]\}$.

A *coding function* f of a graph (G, λ) be any function such that: $\forall x, y, z \in V$, $\forall \pi_1 \in P[x, y]$, $\pi_2 \in P[x, z]$ $f(\Lambda_x(\pi_1)) = f(\Lambda_x(\pi_2))$ iff $y = z$.

Definition 1 *[15]* - Weak Sense of Direction
A system (G, λ), has weak sense of direction *iff there exists a coding function f.*

A *decoding function* h for f is a function that, given a label and a coding of a string (sequence of labels), returns the coding of the concatenation of the label and the string. More precisely, given a coding function f, a decoding function h for f is such that $\forall x, y, z \in V$, such that $\langle x, y\rangle \in E(x)$ and $\pi \in P[y, z]$, $h(\lambda_x(\langle x, y\rangle), f(\Lambda_y(\pi)) = f(\lambda_x(\langle x, y\rangle) \circ \Lambda_y(\pi))$, where $\circ$ is the concatenation operator. We can now define sense of direction:

Definition 2 *[15]* - Sense of Direction
A system (G, λ), has a sense of direction *($\mathcal{SD}$) iff the following conditions hold:*
1) there exists a coding function f,
2) there exists a decoding function h for f.

We shall also say that (f, h) is a sense of direction in (G, λ).

Several instances of sense of direction (*contracted, chordal, neighboring, cartographic*) have been described in [15]; a class of sense of direction called *group* $\mathcal{SD}$ that comprises contracted, chordal, and cartographic $\mathcal{SD}$s has been defined in [48].

It has been shown that $\mathcal{WSD}$ does not imply $\mathcal{SD}$; in fact there are labeled graphs which have $\mathcal{WSD}$ but do not have any $\mathcal{SD}$ [4].

1.3 Sense of Direction and Local Names

The definition of both coding and decoding (and, thus, of sense of direction) can be restated in terms of "translation" capability of local names.

Each node x refers to the other nodes using local names from a finite set $\mathcal{N}$ called name space; let $\beta_x(y)$ be the name associated by x to y. Let us stress that these local names are *not* necessarily identities (i.e., unique global identifiers); in fact, the system could be *anonymous*. The family of injective functions $\beta = \{\beta_x : V \to \mathcal{N} : x \in V\}$ will be called a *local naming* of G.

Let us restate the definition of coding function: a *coding function* f of a graph (G, λ) endowed with local naming β is any function such that: $\forall x, y \in V$,

$\forall \pi \in P[x, y]$, $f(\Lambda_x(\pi)) = \beta_x(y)$. In other words, a coding function translates the sequence of labels of a path from x to y into the local name that x gives to y. Note that while the resulting name is *local* (i.e. x and z might choose different local names for the same node y), the coding function is *global* (i.e., the same for all the nodes).

Let us restate the definition of decoding in terms of local naming. Given a coding function f, a decoding function h for f is any map such that $\forall x, y, z \in V$, with $\langle x, y\rangle \in E(x)$, $h(\lambda_x(\langle x, y\rangle), \beta_y(z)) = \beta_x(z)$. To understand the capabilities of the decoding function in terms of names, consider the situation of node y sending to its neighbor x a message containing information about a node z (see Figure 1). Node z is known at y as $\beta_y(z)$, thus, the message sent by y will

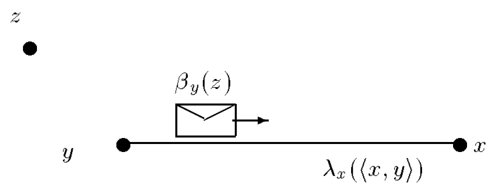

Fig. 1. Communication of information about node z from node y to x.

contain information about a node called "$\beta_y(z)$". The decoding function allows x to "translate" such a name into its own local name for z $\beta_x(z)$, knowing only the label of the link $\lambda_x(\langle x, y\rangle)$ on which the message is received.

1.4 Sense of Direction and Structural Knowledge

The study of sense of direction is part of the larger investigation on structural knowledge. Informally, the term "structural knowledge" refers to any knowledge about the structure of the system; in general we can distinguish between four distinct types, depending on whether the knowledge is about: 1. the communication topology (T); 2. the labeling (L); 3. the (input) data (D); 4. the status of the entities and of the system (S).

Clearly, sense of direction is of type L. In most cases, its impact, existence and properties will be strongly affected also by the other forms of structural information available in the system. It is thus important to both distinguish among different types and within each type.

Examples of type-T knowledge are: *metric information* (TM): "numeric" information about the network; e.g., size of the network (exact or approximate), number of edges, maximum degree, length of the diameter, girth, etc. *topological properties* (TP): knowledge of some properties of the topology; e.g., the fact that the topology of the network is a regular graphs, a Cayley graph, a (edge-, vertex-, cycle-) symmetric graph, etc. *topological classes* (TC): information about the "class" to which the topology belongs; for example, the fact that the topology is a ring, a tree, a mesh, a tori,etc. *topological awareness* (TA): a knowledge equivalent to the knowledge of the adjacency matrix that describes the topology. *complete knowledge of the topology* (TK): each node knows the adjacency matrix and its position in it (this knowledge is partial when truncated at distance d)

Examples of type-D knowledge are: *unique identifiers*: all input values are distinct; *anonymous network*: all input values are identical; *size*: number of distinct values.

Examples of type-S knowledge are: *system with leader*: there is a unique entity in state "leader"; *reset* all nodes are in the same state.

2 Impact of Sense of Direction on Computability

A large amount of research has been devoted to the study of computability in anonymous systems; i.e., the study of what problems can be solved when there are no distinct identifiers associated to the nodes (e.g., [2, 29, 30, 51]). Clearly, which problems can be solved depend on many factors including the structural properties of the system as well as the amount and type of structural knowledge available to the system entities (the nodes).

The computational power of sense of direction in anonymous systems has been studied in [19] and shown to be linked to the notion of surrounding introduced in [18]. The *surrounding* $N(u)$ of a node u in (G, λ) is a labeled graph isomorphic to (G, λ) where the isomorphism maps each node $v \in V$ to the set $\Lambda[u, v]$ of strings corresponding to all walks from u to v. It has been proven in [19] that, in anonymous distributed systems with sense of direction, the surrounding of a node represents the maximum information that an entity can obtain by message transmissions. Moreover, what is computable in such systems depends on the number of distinct surroundings as well as on their multiplicity (i.e., how many nodes have a given surrounding) [19]. It also depends on the amount of type-T knowledge available.

Let $\mathcal{W}_K(P)$ denote the set of graphs with sense of direction where, in presence of type-T knowledge K, problem P can be solved. In [19] it has been shown that, with weak sense of direction, for every type-T knowledge K and for every problem P, $\mathcal{W}_K(P) \supseteq \mathcal{W}_{TK}(P)$. Since $\mathcal{W}_K(P)$ do not depend on K we will denote it as $\mathcal{W}(P)$.

A powerful implication of this result is the formal proof that, *with sense of direction, it is possible to do shortest path routing even in anonymous networks.*

This result also implies that the obvious knowledge-computability hierarchy

$$\mathcal{W}_{ni}(P) \subseteq \mathcal{W}_{ub}(P) \subseteq \mathcal{W}_{size}(P) \subseteq \mathcal{W}_{TA}(P) \subseteq \mathcal{W}_{TK}(P)$$

collapses, where ni, ub, $size$, TA, TK denote no knowledge, knowledge of an upper bound on the size of the network, knowledge of the exact size, topological awareness, and complete knowledge of the topology, respectively.

Another interesting result is about the relationship between the strongest type-T knowledge and the lowest with sense of direction (i.e., $\mathcal{W}(P)$). Let $D_K(P)$

denote the set of graphs with just local orientation where, in presence of type-T knowledge K, problem P can be solved; clearly, $\mathcal{D}_K(P) \subseteq \mathcal{W}_K(P)$ for every type-T knowledge K. In [19] it has been shown that sense of direction and complete knowledge of the topology have the same computational power; in fact $\mathcal{W}(P) = \mathcal{D}_{TK}(P)$ for every problem P.

The relation between *topological awareness* and sense of direction has been studied with respect to several problems including leader and edge election, and spanning tree construction. It has been shown that, for each of these problems, there exist graphs in which they are not solvable with topological awareness (TA); on the other hand, they can be solved with *any* weak sense of direction; that is, *sense of direction is strictly more powerful than topological awareness.*

3 Impact of Sense of Direction on Complexity

The evidence of the impact that specific sense of direction have on the communication complexity of several problems in particular network topologies has been accumulating in the recent years. Most of the distributed algorithms presented in the literature assume and exploit specific, and mostly implicit, instances of sense of direction within their pre-defined set of assumptions on the structural knowledge (held by the entities) and the attributes of the network. Techniques developed for most of the existing applications are *ad hoc*, network- or application-specific; they do not provide seamless mechanisms for allowing processors to inter-operate with all intended assumptions on knowledge.

A large amount of investigations has been dedicated to the study of a small but typical set that recur in most applications: broadcast ($\mathcal{B}$), depth-first traversal ($\mathcal{DFT}$), spanning tree construction ($\mathcal{SPT}$), minimum finding ($\mathcal{MF}$), election (or leader finding)($\mathcal{LF}$). The next sections provide a schematic view of the most important results.

3.1 Arbitrary Network

A network topology is said to be arbitrary if it is unknown to the entities. This can be understood as a lack of topological awareness. The results on the impact of sense of direction in arbitrary networks are listed below.

Arbitrary Networks

	Broadcast Depth First Traversal	Election Spanning Tree Min Finding
local orientation	$\Theta(e)$	$\Theta(e + n \log n)$ [24, 43]
neighboring $\mathcal{SD}$	$\Theta(\min(e, n^{1+\Theta(1)}))$ [3, 24] Theta(n)	$O(e + n \log n)$
chordal $\mathcal{SD}$	$\Theta(n)$ [36]	$O(e + n \log n)$
any $\mathcal{SD}$	$2n - 2$ [17]	$3n \log n + O(n)$ [17, 36, 35]

Here and in the following n denotes the number of nodes of the network and e the number of edges.

Note that Depth First Traversal can be performed in $O(n)$ with *any sense of direction* even if the graph is *anonymous*. In arbitrarily labeled graphs, this problem requires $\Omega(e)$ messages even if the system is not anonymous; this bound can be easily achieved (e.g., [7]). 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 [44]. This result implies that, in presence of neighboring sense of direction, $\mathcal{DFT}$ 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 [16]. It has been proven in [3] that a similar complexity can be obtained for broadcasting without sense of direction only if the size of messages is unbounded.

3.2 Regular Networks

Complete Networks

	Broadcast	Depth First Traversal	Election
local orientation	$(n-1)$	$O(n^2)$	$\Theta(n^2)$ [24]
local orientation with topological awareness	$(n-1)$	$O(n^2)$	$\Theta(n \log n)$ [28]
chordal $\mathcal{SD}$	$(n-1)$	$(n-1)$	$\Theta(n)$ [31, 37, 45]

The results for the Leader Election problem show explicitly that, despite maximum connectivity, there are strict (reachable) lower bounds on what optimal solutions can achieve. A strict and proper classification of the structural information can be deduced. In complete networks also the problem of fault-tolerant leader election has been studied with chordal $\mathcal{SD}$ [35, 38].

Hypercube Network

	Broadcast	Election
local orientation	$O(n)$ [9, 11]	$O(n \log\log n)$ [12]
local orientation and Hamming node-labeling	$O(n)$ [41]	$O(n \log\log n)$
neighboring $\mathcal{SD}$	$O(n)$ [3]	$O(n \log\log n)$
dimensional $\mathcal{SD}$	$O(n)$ (Folklore)	$\Theta(n)$ [14, 42, 47, 50]
chordal $\mathcal{SD}$	$O(n)$ (Folklore)	$O(1)$ [14]

Note that for both Broadcast and Election, a similar complexity can be obtained in all cases (beside neighboring) even if the graph is *anonymous*.

All solutions provided in [14, 47, 42, 50] exploit the implicit region partitioning of the topology and an efficient and implicit scheme to compute and represent shortest paths. Recently, several Broadcast algorithms succeed to exploit the highly symmetric topology without considering processor or link names. Nevertheless, in this case, the best known message complexity for the Election problem is $O(n \log\log n)$ (for large Hypercubes and with a large asymptotic constant) [12].

Chordal Rings

	Broadcast	Election
$\langle 1, 2, .., k\rangle_n$ local orientation chordal $\mathcal{SD}$	 $O(n)$ [34] $n-1$ (Folklore)	 $O(n \log n)$ $O(n \log n)$
$\langle S\rangle_n$ $\|S\| < \log n$ local orientation chordal $\mathcal{SD}$	 $O(\|S\|n)$ $n-1$ (Folklore)	 $O(n \log n)$ $O(n)$ [1]
$\langle S\rangle_n$ $\|S\| < \log\log n$ local orientation chordal $\mathcal{SD}$	 $O(\|S\|n)$ $n-1$ (Folklore)	 $O(n \log n)$ $O(n)$ [27]
$\langle S\rangle_n$ $\|S\| < \log^* n$ local orientation chordal $\mathcal{SD}$	 $O(\|S\|n)$ $n-1$ (Folklore)	 $O(n \log n)$ $O(n \log^* n)$ [40]

The results on chordal rings give information on the situation of a complete graph where, at each node, the chordal labeling is known only for a subset (S) of the edges. In this view must also be interpreted the result on double loop [33].

Wake-up

	Time $\times$ Bits
local orientation	$\Theta(n^2)$ [17]
chordal $\mathcal{SD}$	$O(n \log^2 n)$ [25]

In all previous results, the communication complexity is measured in terms of number of messages. In the case of synchronous systems, the complexity measure is the trade-off between completion time and amount of transmitted information. A strong example of impact of sense of direction in synchronous systems is the *wake-up* (or, *weak unison*) problem in complete graphs:

3.3 Open Problems

The interplay between symmetry of the graph and consistency of the labels is an important open question. As illustrated in the previous results, there is an empirical evidence that highly regular topologies are difficult to exploit efficiently when the edges are not labeled with consistent structural information.

4 Constructing and Testing Sense of Direction

4.1 Deciding and Testing Sense of Direction

Given a labeled graph (G, λ), how can we verify (decide) whether there is sense of direction?

In [4], the authors have shown that there exist polynomial algorithms for deciding both $\mathcal{WSD}$ and $\mathcal{SD}$. Furthermore, they have shown that deciding weak sense of direction can be done very efficiently in parallel. In fact, considering as model of computation a CRCW PRAM, weak sense of direction is in AC^1 for all graphs, where AC^1 denotes the class of problems which are solvable in time $O(\log n)$ using a polynomial number of processors (in this case n^6 processors). This result is based on a characterization of weak sense of direction in purely combinatorial terms. In fact weak sense of direction is shown to be equivalent to a combinatorial condition on the labeling of a graph called *uniformity*, and checking such a condition has then been proved to be in AC^1. Deciding sense of direction is not as efficient in parallel, in fact it is in AC^1 only for some classes of graphs (k-reducible graphs).

4.2 Constructing Sense of Direction

Since Sense of Direction is known to improve the communication complexity of distributed algorithms, computing $\mathcal{SD}$ as a preprocessing phase in unlabeled topology has been studied in [46, 48] showing that any algorithm computing the common sense of direction for cliques, hypercubes, arbitrary graphs, or tori, exchanges at least $\Omega(e - \frac{1}{2}n)$ messages in a network with n nodes and e edges.

This result is not attractive for these dense topologies (e.g., $\Omega(n^2)$ for cliques and $\Omega(n \log n)$ for the hypercube); although natural algorithms matching the lower bounds have been proposed [46], most of the solutions proposed cannot avoid a complete flooding of the communication links in the network.

The interest is more relevant for topologies with a linear number of edges such as tori or chordal rings of constant degree, as shown in the following table presenting the known message complexity.

	Cost	Type of $\mathcal{SD}$
Complete graph	$\Theta(n^2)$ [46]	chordal $\mathcal{SD}$
Hypercube	$\Theta(n \log n)$ [46]	dimensional $\mathcal{SD}$
Torus ($\sqrt{n} \times \sqrt{n}$)	$\Theta(n)$ [33]	compass SD
Double Loop $\langle 1, \sqrt{n} \rangle_n$	$\Theta(n)$ [33]	chordal $\mathcal{SD}$

Interesting constructions of sense of direction have been given in [5], where the authors have concentrated on the problem of constructing senses of direction that use a small number of labels. Such a number always lies between the maximum degree of the graph (in that case it is minimum) and the number of vertices. By exploiting compositional, algebraic and geometrical techniques, they have shown some constructions of large graphs with sense of direction using few labels; the idea of their constructions is to build large networks by appropriately connecting smaller ones. The authors show that with their approach they obtain minimal senses of direction in Cayley graphs. Their result is a first attempt of developing a set of techniques for constructing sense of direction in graphs by using primitive constructions.

4.3 Coping with Lack of Sense of Direction

The results on constructing sense of direction suggested that it may be more efficient to cope with the lack of sense of direction than to use a pre-processing phase which constructs it. Indeed, some problems in unlabeled networks require only a partial orientation to achieve an optimal solution. In tori and chordal rings, the *handrail* technique, introduced in [33], uses the complete knowledge of the topology truncated at distance 2 (or 2-surrounding) to provide any node with the ability to pass a message in the same dimension or to make the message "turn" in the same consistent direction. In the torus, the message-passing scheme will be consistent (clockwise, or counter-wise, squares) in the world according to the leader.

For hypercubes, the *mask* paradigm, introduced in [11], allows a source node to build a small dominating set; the mask coverage is built by including consistently (according to the source) appropriate nodes in each layer of the hypercube. Such a set, once available, can be used efficiently to solve problems such as broadcasting even when the traditional dimensional $\mathcal{SD}$ is not available.

4.4 Open Problems.

The problem of constructing graphs have a sense of direction with the minimum number of labels has been extensively studied in regular graphs (e.g., [6, 13, 18]), but still very little is known in the general case. This problem will be discussed in more details in Section 5.

The paradigms of handrail and mask could be exploited for other topologies with relevant partial structural properties. Both schemes demonstrate that a *local* information of a distance-2 surrounding (although, combined with non-trivial algorithms) may be sufficient to solve particular problems optimally. This raises the interesting and more general question of characterizing the minimal structural properties that a network must possess to solve particular classes of problems within a specific message complexity range (see also [10]).

5 Minimal Sense of Direction in Regular Graphs

In this Section we consider labeled graphs (G, λ) where λ is minimal; i.e., it uses $d(G) = max\{d(x) : x \in V\}$ labels. In particular, we focus on d-regular graphs. If λ is minimal and there is (weak) sense of direction, (G, λ) is said to have *minimal (weak) sense of direction.*

5.1 Minimal $\mathcal{SD}$ and Cycle Symmetry

A graph has *cycle symmetry* if each node belongs to the same number of cycles of the same length [13].

In [13], the author has shown that cycle symmetry is a *necessary condition* for having a minimal sense of direction in a regular graph. For example, the graph of Figure 2 *b*) does not have a minimal $\mathcal{SD}$ since it is not cycle symmetric. On the other hand, figure 2 *a*) shows a graph with cycle symmetry where there exists a minimal $\mathcal{SD}$. Such a condition, however, is *not sufficient*; for example, Peterson's graph has cycle symmetry but it cannot have minimal sense of direction [13].

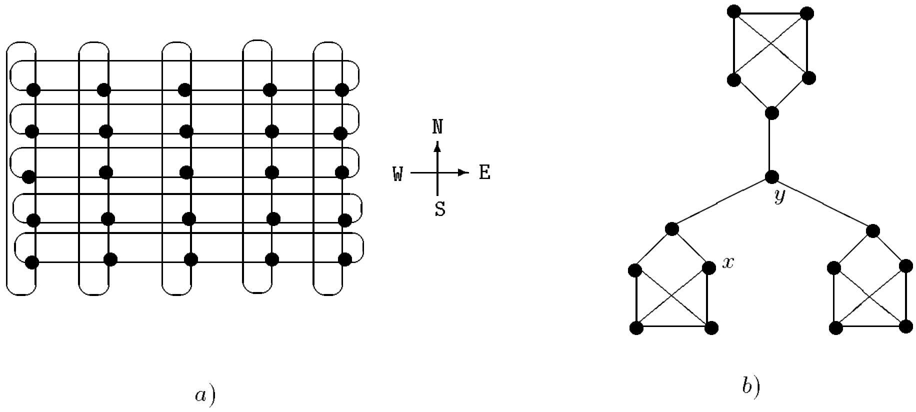

Fig. 2. Regular graphs for which *a*) there exists a minimal $\mathcal{SD}$, *b*) there exists no minimal $\mathcal{SD}$.

5.2 Minimal $\mathcal{SD}$ with Symmetric Labeling.

A labeling λ is *symmetric* if there exists a bijection $\psi : \Sigma \to \Sigma$ such that for each $\langle x, y\rangle \in E$, $\lambda_y(\langle y, x\rangle) = \psi(\lambda_x(\langle x, y\rangle))$; ψ will be called the edge-symmetry function.

In the case of symmetric labelings, minimality is linked to the notion of surrounding (discussed in section 2). A labeled graph is *surrounding symmetric* (S-symmetric) when every node has the same surrounding.

A necessary and sufficient condition has been given in [18] where the authors have shown that a regular graph with symmetric labeling has a minimal $\mathcal{SD}$ iff it is surrounding symmetric.

Hence, we have a characterization of minimal $\mathcal{SD}$ in terms of Cayley graphs: a regular graph has a minimal $\mathcal{SD}$ iff it is a Cayley graph with a Cayley labeling [6, 18]. Note that in [6] the theorem is more general since it holds for directed graphs.

5.3 Minimal $\mathcal{SD}$ with non-Symmetric Labeling.

Given a regular graph with a non-symmetric labeling (G, λ), a necessary and sufficient condition for (G, λ) to have minimal sense of direction has been given in [23]. A regular graph has a minimal $\mathcal{SD}$ iff it is the graph of a semi-group which the direct product of a group and a left-zero semi-group.

5.4 Open Problems

Cycle symmetry is not yet perfectly understood; in fact, its relationship with other well known notions of symmetries (e.g., vertex symmetry) is not always clear. For example, it is evident that a vertex symmetric graph is also cycle symmetric, but it is not known whether a cycle symmetric graph is always vertex symmetric.

The general characterization of minimal sense of direction in arbitrary graphs is still completely open.

An other interesting problem is the following: given a topology which does not have minimal sense of direction, what is the minimum number of labels necessary for having a sense of direction in that topology ? If it cannot be done with $d(G)$ labels (where $d(G)$ is the maximum degree) is it always possible to do it with $d(G)+1$ labels ?

Also the problem: given a graph G and an integer k, can G have a sense of direction using at most k labels ? is still open and suspected to be NP-complete.

6 Topological Constraints for Sense of Direction

The interplay between topology of the system and the properties that a labeling must satisfy for having sense of direction have been studied in [21, 22, 20].

In some graphs, all the labelings with the same simple properties have sense of direction. It is clearly important to identify which property for which graphs.

The authors have considered the following question: *Given a system (G, λ), where the labeling λ has certain properties, under what topological conditions (G, λ) has sense of direction ?*

In particular, they have considered a hierarchy of labelings: *local orientation*, local orientations with *edge symmetry* , and *colorings*. For each class, it has been shown for what graphs every labeling has a sense of direction (*feasible graphs*) and for what other graphs the property of the labeling does not suffice for having sense of direction (they contain *forbidden subgraphs*). This characterization is *complete*: in fact it has been shown that every graph is either feasible or contains a forbidden subgraph.

Local orientation suffices for having sense of direction only in trees and in some trivial graphs; the simultaneous presence of local orientation and edge symmetry guarantees sense of direction in a larger class of graphs which, in addition to trees, includes rings and particular types of "spiked rings"; finally, a coloring suffices for having sense of direction in a even larger class of graphs (some of them are shown in Figure 3).

7 An Application: Object Naming

Sense of direction has been used to implement an efficient naming scheme in systems of distributed objects [49].

An object system consists of a collection of objects and their relations; each object has a state (e.g., local variables) and a behavior (set of actions it may execute) and the global behavior of a system is described in terms of interactions between its objects.

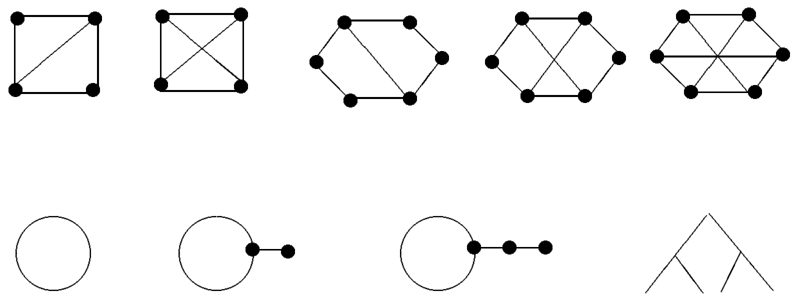

Fig. 3. Some feasible graphs for colorings.

A *naming scheme* is a mechanism used for naming objects and manipulating objects through their names. In object systems, some of the main objectives of a naming scheme are: to *designate* an object, to distinguish between objects *independently of their state*; to *communicate* to other objects the knowledge about some objects. Furthermore, a naming scheme should be *stable* and *reliable*, that is, names must remain valid, and keep denoting the same objects when the system changes its state.

In order to achieve these goals, the existing naming schemes (e.g., [26, 39, 8]) either use an approach based on *global* names (e.g., each object uses as its name an intrinsic property which distinguishes it from the other objects) or they use *hierarchical* names often based on the location of the objects (e.g., email addresses). On the other hand, locality of names and independence on the location are two very desirable characteristics of a naming scheme.

In [49], a naming scheme in which a name of an object does not depend neither on its state nor to its characteristics, nor to its location has been proposed based on sense of direction. A distributed object system is here described by a labeled graph (G, λ) where nodes are objects, links represent references between objects and the labeling λ is chosen in such a way that the system has sense of direction. The name that an object uses to refer to another object is either the label of the link between them (if there is such a direct link), or the coding of a (arbitrarily chosen) sequence of labels leading to it. More precisely,

Given a coding function f, the local object naming scheme β_x for x is constructed as follows: $\forall x, y$,

$$\beta_x(y) = \begin{cases} \lambda_x(x, y) & \text{if } y \in E(x) \\ f(\alpha) & \text{otherwise} \end{cases}$$

where α is the sequence of labels corresponding to an arbitrary path between

x and y. The collection of all local object naming schemes $\beta = \{\beta_x : x \in V\}$ constitutes the (global) *object naming scheme*.

This approach based on sense of direction constitutes the first naming scheme that uses *local names*, provides *mechanisms for the correct communication* of names and is *location independent*.

Acknowledgements. This work has been partially supported by NSERC and FCAR.

References

1. H. Attiya, J. van Leeuwen, N. Santoro, and S. Zaks. Efficient elections in chordal ring networks. *Algorithmica*, 4:437–446, 1989.
2. H. Attiya, M. Snir, and M.K. Warmuth. Computing on an anonymous ring. *Journal of the A.C.M.*, 35(4):845–875, 1988.
3. B. Awerbuch, O. Goldreich, D. Peleg, and R. Vainish. A trade-off between information and communication in broadcast protocols. *Journal of the A.C.M.*, 37(2):238–256, April 1990.
4. P. Boldi and S. Vigna. On the complexity of deciding sense of direction. In *Proc. of 2nd Colloquium on Structural Information and Communication Complexity*, pages 39–52, Olympia, 1995.
5. P. Boldi and S. Vigna. Good fibrations and other constructions which preserve sense of direction. In *3rd International Colloquium on Structural Information and Communication Complexity*, number 6, pages 47–58, Siena, 1996.
6. P. Boldi and S. Vigna. Minimal sense of direction and decision problems for Cayley graphs. *Information Processing Letters*, 64:299–303, 1997.
7. 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.
8. DCE. Guide to writing distributed computing environement applications, digital equipment corporation, 1994.
9. K. Diks, S. Dobrev, E. Kranakis, A. Pelc, and P. Ružička. Broadcasting in unlabeled hypercubes with linear number of messages. *Information Processing Letters*. to appear.
10. S. Dobrev. An alternative view on structural information. In *5th International Colloquium on Structural Information and Communication Complexity*, Amalfi, 1998. to appear.
11. S. Dobrev, P. Ružička, and G. Tel. Time and bit optimal broadcasting in anonymous unoriented hypercubes. In *Proc. of 5th Colloquium on Structural Information and Communication Complexity*, Amalfi, 1998. to appear.
12. S. Dobrev and P. Ruzicka. Linear broadcasting and $n \log \log n$ election in unoriented hypercubes. In *Proc. of the 4th International Colloquium on Structural Information and Communication Complexity*, pages 53–68, 1997.
13. P. Flocchini. Minimal sense of direction in regular networks. *Information Processing Letters*, 61:331–338, 1997.
14. P. Flocchini and B. Mans. Optimal election in labeled hypercubes. *Journal of Parallel and Distributed Computing*, 33(1):76–83, 1996.

15. P. Flocchini, B. Mans, and N. Santoro. Sense of direction: definition, properties and classes. *Networks.* To appear. Preliminary version in *Proc. of 1st Colloquium on Structural Information and Communication Complexity*, 9-34,1994.
16. 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.
17. P. Flocchini, B. Mans, and N. Santoro. On the impact of sense of direction on message complexity. *Information Processing Letters*, 63(1):23–31, 1997.
18. P. Flocchini, A. Roncato, and N. Santoro. Complete symmetries and minimal sense of direction in labeled graphs. *Discrete Applied Mathematics.* to appear. Preliminary version in *Proc. of 27th SE Conference on Combinatorics, Graph Theory and Computing*; Congressus Numerantium 121, 3-18, 1996.
19. P. Flocchini, A. Roncato, and N. Santoro. Computing on anonymous networks with sense of direction. *Theoretical Computer Science.* To appear. (Preliminary version in *Proc. of 3rd Colloquium on Structural Information and Communication Complexity*, 32–46, 1996).
20. P. Flocchini and N. Santoro. Proper coloring and sense of direction. Submitted.
21. P. Flocchini and N. Santoro. Topological constraints for sense of direction. *International Journal on Foundations of Computer Science.* To appear. Preliminary version in *Proc. of 2nd Colloquium on Structural Information and Communication Complexity*, 27-38, 1995.
22. P. Flocchini and N. Santoro. Topological constraints for sense of direction. In *Proc. of 2nd Colloquium on Structural Information and Communication Complexity*, pages 27–38, Olympia, 1995.
23. S. Foldes. Personal communication.
24. 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.
25. A. Israeli, E. Kranakis, D. Krizanc, and N. Santoro. Time-message trade-offs for the weak unison problem. *Nordic Journal of Computing*, 4:317–329, 1997.
26. ISO/IEC JTC1. Information technology - open distributed processing - naming framework, iso/iec dis147771, july 1997.
27. T.Z. Kalamboukis and S.L. Mantzaris. Towards optimal distributed election on chordal rings. *Information Processing Letters*, 38:265–270, 1991.
28. 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.
29. E. Kranakis and D. Krizanc. Distributed computing on anonymous hypercubes. *Journal of Algorithms*, 23:32–50, 1997.
30. E. Kranakis and N. Santoro. Distributed computing on anonymous hypercubes with faulty components. In *Proc. of 6th Int. Workshop of Distributed Algorithms*, Lecture Notes in Computer Science 647, pages 253–263. Springer-Verlag, 1992.
31. 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.
32. N. Lynch. *Distributed Algorithms.* Morgan-Kaufmann, 1995.
33. B. Mans. Optimal distributed algorithms in unlabeled tori and chordal rings. *Journal on Parallel and Distributed Computing*, 1997. To appear. Preliminary version in *Proc. of the 3rd International Colloquium on Structural Information and Communication Complexity*, 17-31,1996.

34. B. Mans and D. Peleg. Broadcasting in unlabelled chordal rings, 1997. Personal Communication.
35. B. Mans and N. Santoro. Optimal fault-tolerant leader election in chordal rings. *IEEE Transactions on Computers*. To appear.
36. 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.
37. 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.
38. 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.
39. OMG. Naming service specification clause 3, corba services, omg, march 199.
40. Yi Pan. A near-optimal multi-stage distributed algorithm for finding leaders in clustered chordal rings. *Information Sciences*, 76(1-2):131–140, 1994.
41. D. Peleg. Message-optimal broadcast on a hypercubes with no neighbour knowledge, 1997. Personal Communication.
42. 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.
43. N. Santoro. Sense of direction, topological awareness and communication complexity. *SIGACT NEWS*, 2(16):50–56, Summer 1984.
44. 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.
45. G. Singh. Efficient leader election using sense of direction. *Distributed Computing*, 10:159–165, 1997.
46. G. Tel. Network orientation. *International Journal of Foundations of Computer Science*, 5(1):1–41, 1994.
47. G. Tel. Linear election in hypercubes. *Parallel Processing Letters*, 5(1):357–366, 1995.
48. G. Tel. Sense of direction in processor networks. In *Proc. of Conference on Theory and Practice of Informatics*, Lecture Notes in Computer Science 1012, pages 50–82. Springer-Verlag, 1995.
49. G. v. Bochmann, P. Flocchini, and D. Ramazani. Distributed objects with sense of direction. In *1st Workshop on Distributed Data and Structures*, Orlando. To appear.
50. A.M. Verweij. Linear-message election in hypercubes. Manuscript.
51. M. Yamashita and T. Kameda. Computing on anonymous networks, part I: characterizing the solvable cases. *IEEE Transaction on Parallel and Distributed Computing*, 7(1):69–89, 1996.

This article was processed using the LaTeX macro package with LLNCS style