Sense of Direction in Distributed Computing

Paola Flocchini
School of Information Technology and Engineering, University of Ottawa,
(flocchin@site.uottawa.ca)

Bernard Mans
Department of Computing, Macquarie University
(bmans@ics.mq.edu.au)

Nicola Santoro
School of Computer Science, Carleton University
(santoro@scs.carleton.ca)

Abstract

Sense of Direction is a property of labeled graphs which has been shown to have a definite impact on computability and complexity in systems of communicating entities, and whose applicability ranges from the analysis of graph classes to distributed object systems. The full consequences of this property are still not known; in fact, the ongoing investigations continue to bring new (often surprising) results, to establish unsuspected links with other research and/or application areas, and to pose more questions than they answer. The aim of this paper is to provide a view of the current status of research, describing some of the relevant results, and providing pointers to future research directions.

1 Introduction

In its more general formulation, a *distributed system* is a collection of computational entities communicating by exchanging finite amounts of information, which we shall call messages. The exact nature of the entities (i.e., processors, processes, network nodes, agents, objects, etc.,) will depend on the particular instance; it is however assumed that each entity has local memory and processing capabilities.

Different models exist, depending on how the communication is achieved. We shall restrict ourselves to systems based on the *point-to-point* model [46]. In such systems, each entity can communicate directly and individually with a subset of the other entities, called its neighbors, and can distinguish between them. For such systems, the communication topology can be represented by an edge-labeled undirected graph $(G = (V, E), \lambda)$, where each node $x \in V$ corresponds to a system entity, each edge $\langle x, y \rangle \in E$ corresponds to a pair of neighboring entities x and y, and each entity x has a distinct label (e.g., port number) $\lambda_x(\langle x, y \rangle)$ associated to each of its incident links. Note that for each edge $\langle x, y \rangle$ there are two labels, $\lambda_x(\langle x, y \rangle)$ and $\lambda_y(\langle y, x \rangle)$, one for each of its incident nodes.

Any property of the labeling $\lambda = \{\lambda_x : x \in V\}$ 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 ability at each node to distinguish among the incident links; by definition of point-to-point, we only consider labelings having this property. Another interesting property is *Edge Symmetry*: the ability to understand from the local label what is the label at the other end; this is for example the case of the "left-right" labeling in rings or of the "north-south, east-west" compass labeling of meshes and tori. Also of interest is 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); an example is the traditional "dimensional" labeling of hypercubes.

Without any doubt, one of the properties 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*: the ability to understand, from the labels associated to the edges, whether different walks from any given node x end in the same node or not. The consequences and the impact of this simple property are still not totally clear; in fact, as new results are established, more and more open questions are posed and unsuspected links with other research area discovered. The aim of this paper is to provide a view of the current status of research, describing some of the relevant results, and providing pointers to future research directions.

2 Definitions and Terminology

2.1 Labeled Graphs

Let $G = (V, E)$ be a simple undirected graph; let $E(x)$ denote the set of edges incident to node $x \in V$, and $d(x) = |E(x)|$ the degree of x.

Given $G = (V, E)$ and a set Σ of labels, a *local orientation* of $x \in V$ is any injective function $\lambda_x : E(x) \rightarrow \Sigma$ which associates a distinct label to each edge. A set $\lambda = \{\lambda_x : x \in V\}$ of local orientations will be called a *labeling with local orientation* (or, simply, *labeling*) of G, and by (G, λ) we shall denote the corresponding *(edge-)labeled graph.*

A labeling λ is *minimal* if it uses $d(G) = max\{d(x) : x \in V\}$ labels. It is *symmetric* if there exists a bijection $\psi : \Sigma \rightarrow \Sigma$ 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. A symmetric labeling is a *coloring* if the edge symmetry function is the Identity.

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 non empty 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 $\Lambda_x : P[x] \rightarrow \Sigma^+$ 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 $\Lambda[x, y] = \{\Lambda_x(\pi) : \pi \in P[x, y]\}$.

2.2 Sense of Direction

Given a labeled graph (G, λ), the system is said to have *Sense of Direction* when it is possible to understand, from the labels associated to the edges, whether different walks from any given node x end in the same node or not. More precisely, sense of direction involves the existence of a consistent coding and a consistent decoding function.

Given (G, λ), a *consistent coding function* (or, simply, *coding function*) f for λ is any function with domain Σ^+ such that walks originating from the same node are mapped to the same value (called *local name*) if and only if they end in the same node; that is, $\forall x, y, z \in V$, $\forall \pi_1 \in P[x, y]$, $\pi_2 \in P[x, z]$, $f(\Lambda_x(\pi_1)) = f(\Lambda_x(\pi_2)) \quad \Leftrightarrow \quad y = z$. We shall denote by $\mathcal{N}$ the codomain of f.

Definition 1 *[27]* - Weak Sense of Direction
A system (G, λ), has weak sense of direction *iff there exists a coding function f for λ.*

Given a coding function f, a decoding function h for f is any function $h : \Sigma \times \mathcal{N} \rightarrow \mathcal{N}$ such that $\forall x, y, z \in V$, with $\langle x, y\rangle \in E(x)$ and $\pi \in P[y, z]$: $h(\lambda_x(\langle x, y\rangle), f(\Lambda_y(\pi)) = f(\lambda_x(\langle x, y\rangle) \cdot \Lambda_y(\pi))$, where $\cdot$ is the concatenation operator.

Definition 2 *[27]* - 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 for λ,
2) there exists a decoding function h for f.

We shall also say that (f, h) is a sense of direction in (G, λ).
Notice that $\mathcal{SD}$ is a stronger notion than $\mathcal{WSD}$; in fact, there exist labeled graphs with $\mathcal{WSD}$ but without $\mathcal{SD}$ (see, for example, [11]).

2.3 Examples of Sense of Direction

Some types of sense of direction exist only for specific classes of topologies, while others are "universal", in the sense that they can be constructed in any graph. In the following, we briefly review some of the universal instances.

Chordal $\mathcal{SD}$. Given any graph $G = (V, E)$, with $|V| = n$, a chordal labeling can be obtained by fixing an arbitrary cyclic ordering of the nodes and labeling a link $\langle x, y\rangle$ by the distance (modulo n) between x and y in the predefined ordering (see Figures 1 and 2 (a)). In this case the set of labels Σ is the set of positive integers modulo n.
If λ is a chordal labeling, the system (G, λ) has *chordal sense of direction*; the coding function $f : \Sigma^+ \to \Sigma$ is the function that maps a sequence of labels into their sum modulo n. More precisely, for any sequence of labels $a_1, a_2, \ldots a_k$ with $a_i \in \Sigma$, $f(a_1, \ldots a_k) = \sum_{i=1}^{k} a_i \bmod n$

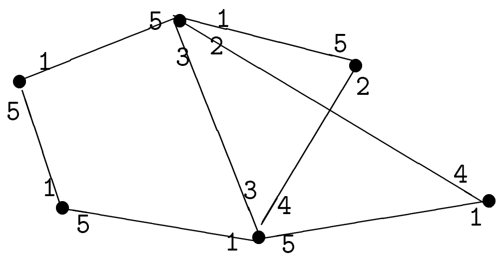

Figure 1: Chordal sense of direction.

The corresponding decoding function is the following: $\forall a, b \in \Sigma$, $h(a, b) = (a + b) \bmod n$,

Neighboring $\mathcal{SD}$. A graph (G, λ) has *neighboring sense of direction* when λ is as follows: $\forall \langle x, y\rangle \in E[x]$, $\langle z, w\rangle \in E[z]$, $\lambda_x(\langle x, y\rangle) = \lambda_z(\langle z, w\rangle)$ iff $y = w$.
That is, in a neighboring sense of direction, all the links ending in the same node are labeled with the same label (see Figure 2 (b)).
Let Σ be the set of labels; the coding function $f : \Sigma^+ \to \Sigma$ is the following: for any sequence of labels $a_1, a_2, \ldots a_k$ with $a_i \in \Sigma$, $f(a_1, \ldots a_k) = a_k$. The corresponding decoding function is such that $\forall a, b \in \Sigma$, $h(a, b) = b$

A property of neighboring sense of direction that sets it apart from most classes of sense of directions (including the others described here) is that, when it is present in an *anonymous network*, it always destroys anonymity. An anonymous network is a network where the nodes are totally indistinguishable; in such systems, neighboring sense of direction allows the construction of distinct identifiers for the nodes.

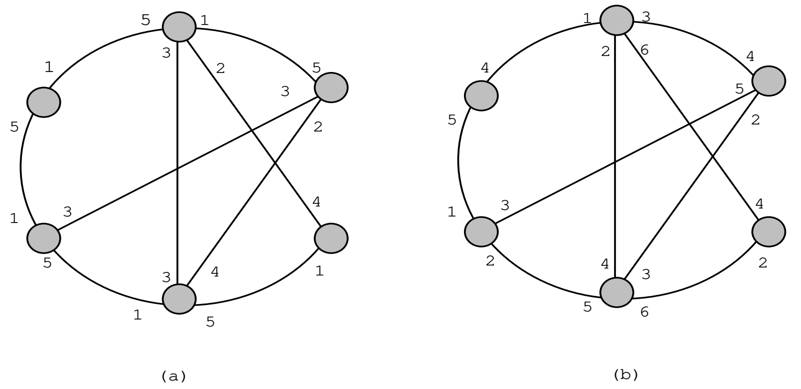

Figure 2: A system with (a) Chordal and with (b) Neighboring sense of direction

Coordinate $\mathcal{SD}$. Given an embedding of G in the plane, λ is a coordinate labeling iff: $\forall\langle u,v\rangle \in E[u]$, $\lambda_u(\langle u,v\rangle) = (x_1 - x_0, y_1 - y_0)$ where (x_0,y_0) and (x_1,y_1) are the coordinates of u and v, respectively, in the embedding (see Figure 3).

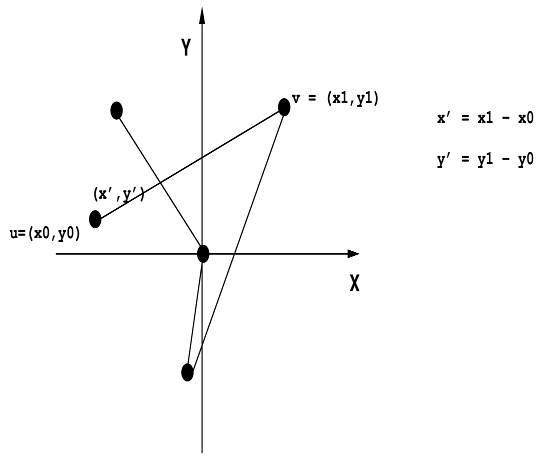

Figure 3: Coordinate Sense of Direction.

With such a labeling (G,λ) has a *coordinate sense of direction.* Let Σ be the set of labels; the coding function $f : \Sigma^+ \rightarrow \Sigma$ is defined as follows: for any sequence of labels $(x_1,y_1),\ldots(x_m,y_m)$ with $(x_i,y_i) \in \Sigma$,$f((x_1,y_1),\ldots(x_m,y_m)) = (\sum_{i=1}^m x_i, \sum_{i=1}^m y_i)$ and the corresponding decoding function h: $\forall(x,y),(x',y') \in \Sigma$, $h((x,y),(x',y')) = (x+x', y+y')$.

Other examples of labelings with $\mathcal{SD}$, not necessarily universal, can be found in [27]; an elegant and compact definition of a large class can be found in [65].

2.4 Translation of Local Views

The notion of consistent coding, decoding, and sense of direction, can be rephrased in terms of local names.

Each entity 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; the set $\{\beta_x(y) : y \in V\}$ is called the *local view* of x. Let us stress that these local names are *not* necessarily identities (i.e., unique global identifiers); in fact,

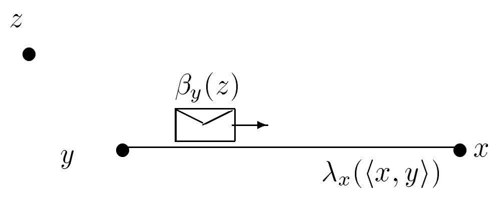

Figure 4: Communication of information about node z from node y to x.

the system could be *anonymous*. The family of injective functions $\beta = \{\beta_x : V \rightarrow \mathcal{N} : x \in V\}$ will be called a *local naming* of G.

Let us restate the definition of coding function in terms of local names: 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 4). Node z is known at y as $\beta_y(z)$, thus, the message sent by y will 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, namely $\beta_x(z)$, knowing only the label of the link $\lambda_x(\langle x, y\rangle)$ on which the message is received.

2.5 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 (G, λ); 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 distinguish both 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 size or an upper bound of the size), 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 graph, 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 "labeled adjacency matrix" of (G, λ). *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. *Multiset*: input values are not necessarily 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. *Unique initiator*: there is a unique entity in state "initiator".

3 Impact of Sense of Direction on Complexity

3.1 Problems

The evidence of the impact that specific labelings with sense of direction have on the communication complexity of several problems, in particular network topologies, has been accumulating in the recent years.

In general, the investigations have focused on a small set of typical problems that recur in many applications, for example: broadcast ($\mathcal{B}$), spanning tree construction ($\mathcal{SPT}$), depth-first traversal ($\mathcal{DFT}$), topology recognition ($\mathcal{TR}$), minimum finding ($\mathcal{MF}$), leader election ($\mathcal{LE}$), edge election ($\mathcal{EE}$), wakeup ($\mathcal{W}$), etc.

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*. Notice that, unless otherwise stated, no assumption is made on the number of *initiators*.

The *leader election* ($\mathcal{LE}$) 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*. The *wake-up problem* ($\mathcal{W}$) consists in arriving from an initial configuration where some entities are in the *initiator* state while the others are in a different state (say *asleep*), to a configuration where all the entities are *awake*.

For some problems it is also required that, when the final state is reached, some knowledge has been acquired or constructed. Examples of these problems are: The *broadcast* problem ($\mathcal{B}$), which consists in arriving from an initial configuration where exactly one entity (the *unique initiator*) has an information, to a configuration where all the entities have the information. The *minimum finding* problem ($\mathcal{MF}$) which consists in moving to a final configuration where every entity knows the minimum of the input values. The *depth-first traversal* problem ($\mathcal{DFT}$), which consists in moving to a final configuration in which a depth-first spanning-tree of the graph rooted at the *unique initiator* has been identified (i.e., every entity knows which of its neighbors belong to the tree). The *spanning tree construction* problem ($\mathcal{SPT}$), which consists in moving to a final configuration in which a spanning-tree of the graph rooted at the *unique initiator* has been identified (i.e., every entity knows which of its neighbors belong to the tree). The *topology recognition* problem ($\mathcal{TR}$), which consists in moving from an initial configuration where entities do not know (G, λ), to a final configuration where they know the "labeled adjacency matrix" of (G, λ). The *complete topology recognition* problem ($\mathcal{CT}$), in the final configuration of which, besides knowing the "labeled adjacency matrix" of (G, λ), each entity must also know its position in the matrix.

An important aspect of these problems is that their solutions are usually modular, and the solutions for the simpler problems can be used as building blocks for solving the complex ones; for instance, a solution to $\mathcal{SPT}$ can be constructed starting from any network traversal protocol (e.g., a solution to $\mathcal{DFT}$) [40]. Furthermore, large-scale problems and applications are likely to be (and are) modularly decomposed in terms of these basic problems; hence, the intrinsic cost of solving such basic problems has a direct impact on the cost of larger ones. Finally, there are known complexity equivalences and relationships among some of these problems. For example, in absence of any type-T knowledge (and assuming distinct identities), the problems $\mathcal{LE}$, $\mathcal{SPT}$ and $\mathcal{MF}$ are equivalent [59]; the communication complexity of $\mathcal{B}$ is not greater than that of $\mathcal{LE}$ [26]; and, clearly, any solution to $\mathcal{DFT}$ is also a solution to $\mathcal{B}$.

consistency	Broadcast Depth First Traversal	Election Spanning Tree Minimum Finding
local orientation	$\Omega(e)$ [26]	$\Omega(e + n\log n)$ [58]
neighboring $\mathcal{SD}$	$\Theta(\min(e, n^{1+\Theta(1)}))$ [3] $\Theta(n)$ [60]	$O(n\log n)$ [44]
chordal $\mathcal{SD}$	$\Theta(n)$ [25]	$O(n\log n)$ [48]
c-group $\mathcal{SD}$	$\Theta(n)$ [65]	
arbitrary $\mathcal{SD}$	$2n-2$ [26]	$3n\log n + O(n)$ [26]

Table 1: Universal Protocols

The studies have concentrated both on the inherent limits (lower bounds) and on the achievable results (upper bounds) for solving these problems in presence of sense of direction. In systems with point-to-point communication, the main cost parameter is the total number of messages transmitted during the execution of a protocol, and thus the communication complexity of a problem is measured mainly in terms of its message complexity.

The rest of this section provides a schematic view of the significant and important results; they will be listed in chronological order according to the level of consistency (from local orientation to sense of direction) and of knowledge required.

In the following, the problems $\mathcal{LE}$, $\mathcal{SPT}$ and $\mathcal{MF}$ will be considered only in networks with unique identifiers; $n = |V|$ will denote the number of nodes and $e = |E|$ the number of edges.

3.2 Universal Protocols

A protocol is *universal* if it can be executed on any system (G, λ); obviously, to be universal, a protocol should *not* require a priori knowledge of the network topology (type-T knowledge) at the nodes. In the literature, this situation is referred to as the "arbitrary topology" case, or as "topological ignorance". The advantages of universal protocols are clear and many, including portability and robustness to network changes.

The salient results on the impact of sense of direction for the complexity of universal protocols are the following.

- Both $\mathcal{B}$ and $\mathcal{DFT}$ can be performed with $O(n)$ messages with *any* sense of direction even if the system is anonymous. This represents a dramatic improvement from the $\Omega(e)$ lower bound for these problems for arbitrary labelings; note that this lower bound holds even if the entities have distinct identities.
- All three problems $\mathcal{LE}$, $\mathcal{SPT}$ and $\mathcal{MF}$ can be solved with $O(n\log n)$ messages with *any* sense of direction; this represents a significant improvement on the $\Omega(e + n\log n)$ lower bound for these problems in the case of arbitrary labelings.

A more detailed discussion of the known results (summarized in Table 1) follows.

3.2.1 Broadcast and Depth First Traversal

In arbitrarily labeled graphs (i.e., with local orientation alone), and in absence of additional type-T knowledge, both the broadcast and the depth first traversal problems require $\Omega(e)$ messages even if the entities have distinct identities [26]; such a bound can be easily achieved (e.g., see [13]).

An improvement in the message complexity had been shown to exist if each entity has a distinct identity *and* knows the identities of all its neighbors; in this case, $O(n)$ messages suffice to solve

$\mathcal{DFT}$ (and, thus, $\mathcal{B}$) [60]. These results imply that, in presence of *neighboring* sense of direction, $\mathcal{B}$ and $\mathcal{DFT}$ can be performed in $O(n)$ messages. Recall that in graphs with neighboring labelings anonymity can be broken.

A similar reduction to $O(n)$ was then shown to be possible in the presence of the less powerful *chordal* sense of direction even in anonymous systems [25]. This result has then been extended to the larger class of labelings with *commutative group* sense of direction [65].

Finally, it has been shown that *any* sense of direction allows these two problems to be solved with only $O(n)$ messages [26].

On a related note, it has been shown that, without sense of direction, a $O(n)$ broadcast protocol exists only if the size of messages is *unbounded* [3].

3.2.2 Election, Spanning Tree and Minimum Finding

In arbitrarily labeled graphs and in absence of additional type-T knowledge, the three problems $\mathcal{LE}$, $\mathcal{SPT}$, and $\mathcal{MF}$ are equivalent and require $\Omega(e+n\log n)$ messages [58]; such a bound can be achieved (e.g., see [35]).

An improvement was shown to be possible if each entity knows the identities of all its neighbors, that is in presence of the *neighboring* sense of direction; in this case, $O(n\log n)$ messages suffice [44]. A similar reduction to $O(n\log n)$ messages was then shown to be possible also in the presence of the less powerful *chordal* sense of direction [25].

These two labeling-dependent results have been recently fully generalized; in fact, *any* sense of direction allows for $O(n\log n)$ solutions to these three problems [26].

3.3 Network-specific Protocols

Most of the original research on sense of direction has focused on specific network topologies. Clearly, the solution algorithms, to be executable, require at each entity type-T knowledge in addition to local orientation. All the results below assume the presence of *topological awareness* $\mathcal{TA}$. The known results are summarized in Tables 1, 2, 3, and 4.

3.3.1 Complete Networks

In complete networks, the *broadcast* process is trivially and optimally achieved with $n-1$ messages regardless of sense of direction.

The situation is quite different for the $\mathcal{DFT}$ problem. In fact, in absence of sense of direction, $\Omega(n^2)$ messages are required, *even if there is a leader*; this follows from the lower bound on the construction of degree-restricted spanning-trees of the complete graph of [41]. On the other hand, with *any* sense of direction, this problem can be solved in $O(n)$ messages using the universal protocol of [26], *even if the network is anonymous.*

Historically, the *election* problem in the complete graph has provided one of the major motivations for the studies on the relationship between consistency of the labeling and communication complexity. In fact, it is known that the election process requires $\Omega(n\log n)$ messages in complete networks [41]. Because the same lower bound holds also for ring networks, this implies that, as far as the election process is concerned, there is no complexity difference between the two types of networks. But this in turns implies that, for the election task, it is totally useless to have a dedicated direct communication links between every pair of entities ($O(n^2)$ links in all) since the same bounds can be achieved connecting each entity to only two others in a ring network (with n links in all). This last implication went mostly unnoticed; those who did notice were puzzled. The solution to the puzzle was provided by the result that, with *chordal* $\mathcal{SD}$, the election process can be carried out with only $O(n)$ messages [45]; this was the first clear indication of the impact of $\mathcal{SD}$. This result

	Depth First Traversal	Election	Fault-Tolerant Election
local orientation with topological awareness	$O(n^2)$	$\Theta(n \log n)$ [41]	$\Omega(n \log k + kf)$ [38]
chordal $\mathcal{SD}$	$\Theta(n)$ [25]	$\Theta(n)$ [45, 49, 62]	$\Theta(n + kf)$ [50]
arbitrary $\mathcal{SD}$	$\Theta(n)$ [26]	$\Theta(n)$ [18]	?

Table 2: Complete Networks.

Wake-Up	Time $\times$ Bits
local orientation with topological awareness	$\Theta(n^2)$
chordal $\mathcal{SD}$	$O(n \log^2 n)$ [37]

Table 3: Synchronous Complete Network.

implies that consistency of the labeling is the key to exploit the greater hardware capabilities of the complete network with respect to the ring; these implications were first observed and explained in [59]. Recently, a long standing open problem has been closed by showing that the same reduction in complexity is achievable in presence of *any* $\mathcal{SD}$ [18].

Subsequent research has focused on improvements of the constant factor of the message complexity and on improvements of the time complexity [49, 62], and adding fault-tolerance to the requirements [48, 50, 52, 61]. In the results for fault-tolerant election, f is the number of faulty processors, k is the number of initiators, and the faults are "fail-stop".

In the case of *synchronous* complete networks, a good example of the impact of sense of direction is the *wake-up* (or, *weak unison*) problem. In an arbitrarily labeled complete network, the worst case *message complexity* is $\Omega(n^2)$, and the bound is trivially achievable. In presence of *chordal* $\mathcal{SD}$, this problem can be easily solved with $O(n)$ *bits*; this solution however has a high time complexity [37]. Interestingly enough, even the composite cost measure *Time* $\times$ *Bits*, is asymptotically less than the message complexity without sense of direction; in fact, the *Time* $\times$ *Bits* complexity is no more than $O(n \log^2 n)$ [37].

3.3.2 Chordal Rings

In the case of chordal rings, the chord structure S has clearly a critical impact on the achievable bounds. It is also important to understand that the results on a chordal ring with structure S give information on the situation of a complete graph where, at each node, the chordal labeling is known only for the subset S of the edges.

The first result was the observation that the chord structure $S = \{1, 2, \ldots, k\}$ with $k = O(\log n)$ suffice to obtain a $\Theta(n)$ election algorithms [1]; other chord structures achieving the same bound with the same number of chords were also identified. Subsequent investigations have reduced the *size* of the structure necessary to achieve linear election algorithms from $O(\log n)$ to $O(\log \log n)$ [39], to $O(\log \log \log n)$ [54], to $O(1)$ in [47, 65].

3.3.3 Hypercube Network

The results for *broadcasting* in the hypercube show a general *insensitivity* of this topology to sense of direction. In fact, broadcasting can be performed with $O(n)$ messages in presence of neighboring sense of direction [3], and it is well known that the same bound may be achieved with the dimensional

$\langle S \rangle_n$	Election local orient.	Election chordal $\mathcal{SD}$
$S = \{1, 2, \ldots, k\}$	$O(nk + n \log n)$	$O(\frac{n}{k} \log \frac{n}{k})$ [1]
$\|S\| = O(\log \log n)$	$O(n \log n)$	$O(n)$ [39]
$\|S\| = O(\log^* n)$	$O(n \log n)$	$O(n \log^* n)$ [54]
$S = \{1, \sqrt{n}\}$	$O(n)$ [47]	$O(n)$ Fabri&Tel95 in [65]
arbitrary $S, \|S\| = k$	$O(nk + n \log n)$	?

Table 4: Chordal Ring.

	Broadcast	Election
local orientation	$O(n)$ [15, 22]	$O(n \log \log n)$ [19]
neighboring $\mathcal{SD}$	$O(n)$ [3]	$O(n \log \log n)$
dimensional $\mathcal{SD}$	$O(n)$	$\Theta(n)$ [24, 57, 64, 66]
chordal $\mathcal{SD}$	$O(n)$	$O(1)$ [24]
arbitrary $\mathcal{SD}$	$O(n)$	$O(n)$ [18]

Table 5: Hypercube.

sense of direction. However, the same bound has been obtained assuming only local orientation and topological awareness [15, 22].

For the *election* problem, in *absence* of sense of direction, a solution has been developed which uses $O(n \log \log n)$ messages [19]; this is an improvement on the $O(n \log n)$ bound implied by the universal protocol. In presence of the *dimensional* sense of direction (traditional for hypercubes), several $\Theta(n)$ election algorithms have been presented [24, 57, 64, 66]. Most of these solutions exploit the implicit region partitioning of the topology and an efficient and implicit scheme to compute and represent shortest paths. These results can be obtained also by using the more recent modular technique of [21].

An interesting additional result is that the *chordal* sense of direction would break the symmetry of the hypercube so to allow a surprising $O(1)$ solution [24].

Recently, the question of whether a $O(n)$ solution exists with *any* sense of direction has been positively answered [18].

3.4 Insensitivity to Sense of Direction

Some graphs have particular structural properties which can be exploited, for some problems, to obtain solutions which are efficient even in absence of sense of direction. These graphs will then be said to be *insensitive* to sense of direction for those problems. The identification of the properties which renders a graph insensitive is a fascinating research area. The knowledge to date is limited to some vertex-symmetric regular graphs of constant degree.

In particular, the *election* problem can be solved with the exchange of $\Theta(n)$ messages in *tori* [55, 47], in the class of *chordal rings* $\langle 1, \sqrt{n} \rangle_n$ [47], as well as in *butterflies* and *cube-connected-cycles* [21].

Broadcast can be easily performed with $O(n)$ messages in all those graphs, since they are sparse; thus, the research has focused on the multiplicative constant. In *tori*, the $3n + 1$ trivial upperbound has been reduced first to $2n + O(\sqrt{n})$ [16], and subsequently to $\frac{10}{7}n + O(\sqrt{(n})$ [20]; the current lower-bound is $\frac{8}{7}n + O(1)$[20].

It is interesting to note that there exists graphs in which broadcasting without sense of direction can be done with the absolutely minimum cost: n-1 messages. A characterization of regular networks having this property can be found in [17]; the general problem to decide whether a graph has this property is NP-hard [17].

For some classes of regular networks with *non-constant* degree, it is clearly possible that their structural properties might be exploited to (partially) cope with the absence of sense of direction. However, extremely little is known. As mentioned before, the hypercube is insensitive to sense of direction in regards to the broadcast problem. There is widespread suspicion that it might be insensitive also in the case of the election problem. It is interesting to note that all algorithms for hypercubes without $\mathcal{SD}$ [15, 19, 22] exploit the well-known fact that the hypercube of size n has a dominating set of size $O(\frac{n}{\log n})$; active nodes build such a dominating set to efficiently solve problems such as broadcast and election. The method has been further refined to reduce the overall bit-complexity by using the *mask* paradigm [22].

3.5 Open Problems

The development of generic protocols which exploit sense of direction is an ongoing research effort. Any such result would provide not only an efficient and portable solution but also further indications on the nature of sense of direction. Of particular interest would be investigations on problems other than the ones already being explored. Similarly, lower-bounds would give precious information on limits and boundaries.

Most of the topology dependent results have been established only for particular sense of directions. The achievability of those bounds by larger classes of (possibly, all) labelings with sense of direction is a foremost open problem.

For some classes of graphs (e.g. chordal rings) results are known only for particular chord structures. An outstanding open problem is to determine an election algorithm which would effectively exploit the properties of chordal sense of direction in *every* graph in the class.

The discovery of classes of graphs insensitive to sense of direction is an important research goal; in particular, the identification of which graph properties are powerful enough to substitute for sense of direction would have a direct practical effect for the design of protocols for systems with such a communication topology. A particularly intriguing set of question is about the "level" of insensibility of a graph; e.g., are there graphs which are insensitive for *every* problem? are there graphs that become insensitive when endowed with additional knowledge?

An outstanding open question is whether the hypercube is insensitive to sense of direction also for the election problem; the development of an $O(n)$ election algorithm in presence of an arbitrary local orientation would provide a positive answer.

Finally, for *synchronous* networks there are no known results except the ones for the wake-up problem.

4 Impact of Sense of Direction on Computability

4.1 Knowledge and 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, 9, 5, 42, 43, 51, 68]). Clearly, which problems can be solved depends 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 [28] and shown to be linked to the notion of surrounding, introduced in [29].

Before describing the notion of surrounding, we need the following definition of labeled graph isomorphism. Given two labeled graphs $(G = (V, E), \lambda)$ and $(G' = (V', E'), \lambda')$, a *labeled graph isomorphism* is a bijection $\alpha : (G, \lambda) \to (G', \lambda)$ such that: $(\alpha(u), \alpha(v)) \in E'$ iff $(u, v) \in E$, and $\lambda_{\alpha(u)}(\alpha(u), \alpha(v)) = \lambda_u(u, v)$.

The *surrounding* $N(u)$ of a node u in (G, λ) is a labeled graph isomorphic to (G, λ) through a labeled graph isomorphism χ, together with the image of node u, $\chi(u)$. The surrounding of a node represents the maximum information that an entity can obtain by message transmissions in anonymous distributed systems with sense of direction. In particular, what is computable in these systems depends on the number of distinct surroundings, as well as on their multiplicity (i.e., how many nodes have a given surrounding) [28]. An interesting property is that the classes of nodes having the same surrounding have the same cardinality $\rho_{(G,\lambda)}$, called *surrounding symmetricity* of (G, λ). Notice that the availability at node u of its surrounding is equivalent to availability at u of the complete topology of the system (TK).

Let $\mathcal{G}$ denote the set of all labeled graphs, $\mathcal{W}$ the set of labeled graphs with weak sense of direction, and $\mathcal{D}$ the set of labeled graphs with sense of direction. Given $\mathcal{C} \subseteq \mathcal{G}$ and a set of graphs $\mathcal{S}$, let $\mathcal{C}/\mathcal{S} = \{(G, \lambda) \in \mathcal{C} : G \in \mathcal{S}\}$ denote the restriction of $\mathcal{C}$ to $\mathcal{S}$; e.g., $\mathcal{W}/\{G_1, G_2\}$ is the set of labeled graphs obtained by labeling G_1 and G_2 with a weak sense of direction. For simplicity, denote $\mathcal{W}/\{G\}$ by $\mathcal{W}/G$.

A problem P is said to be *solvable* in $\mathcal{C}/\mathcal{S}$ if it is solvable in all labeled graphs $(G, \lambda) \in \mathcal{C}/\mathcal{S}$; it is K-solvable in $\mathcal{C}/\mathcal{S}$ if it is solvable in all $(G, \lambda) \in \mathcal{C}/\mathcal{S}$ when all nodes are empowered with a-priori knowledge K.

Notice that there are problems which, although not solvable in $\mathcal{G}/\mathcal{S}$, might be K-solvable in $\mathcal{G}/\mathcal{S}$ for some additional a-priori knowledge K. Consider, for example, the problem of computing the AND of the input values in a graph. It is easy to see that there are graphs (e.g., rings) where the AND of the input values can or cannot be distributively computed, depending on the edge labeling. However, if some type-T knowledge is available (e.g., the number of nodes $|V|$, the number of edges $|E|$), the problem becomes solvable in those graphs with any labeling.

4.2 General Results.

The main property of sense of direction is the following: Given a graph G and a type-T knowledge K, if a problem P is K-solvable in $\mathcal{G}/G$ then P is solvable in $\mathcal{W}/G$ (where knowledge of the coding function is assumed) [28]. In other words, *with weak sense of direction, no other knowledge is needed.*

This result is based on the fact that, in a labeled graph with weak sense of direction (and without any other information except the coding function), every node can construct its surrounding even if the system is anonymous (and, thus, no leader can be elected). Note that, without sense of direction (i.e., with an arbitrary labeling), the problem of constructing a surrounding in anonymous networks is *unsolvable*.

A powerful and surprising implication of this result is that, *with weak sense of direction, it is possible to do shortest path routing in anonymous networks*, i.e., even if there are no global identifiers for neither source nor destination (nor for any other node in the graph).

Another interesting consequence of this result is that from a computational point of view, weak sense of direction and sense of direction are equivalent; this is not so from a complexity viewpoint.

4.3 Results on Specific Problems.

The computational power of anonymous systems with $\mathcal{WSD}$ has been studied in [28] considering some specific, typical problems: *leader election* $\mathcal{LE}$, *edge election* $\mathcal{EE}$, *spanning tree construction* $\mathcal{SPT}$, *topology recognition* $\mathcal{TR}$, *complete topology recognition* $\mathcal{CT}$.

Solvability of these problem with sense of direction has been shown to depend on the level of symmetricity of the surrounding; in particular:

- $\mathcal{LE}$ is solvable in $(G, \lambda) \in \mathcal{W}/G$, iff $\rho_{(G,\lambda)} = 1$; i.e., all surroundings are different.
- $\mathcal{EE}$ and $\mathcal{SPT}$ are solvable in $(G, \lambda) \in \mathcal{W}/G$ iff $\rho_{(G,\lambda)} \leq 2$ and, if $\rho_{(G,\lambda)} = 2$ there must be two adjacent nodes with the same surrounding.
- $\mathcal{TR}$ and $\mathcal{CT}$ are solvable for every labeled graph in $\mathcal{W}$.

Let $\mathcal{W}(\mathcal{LE})$, $\mathcal{W}(\mathcal{EE})$, $\mathcal{W}(\mathcal{ST})$, $\mathcal{W}(\mathcal{TR})$, and $\mathcal{W}(\mathcal{CT})$ denote the classes of graphs on which such problems are solvable, respectively. The relationship between these classes is as follows:

$$\mathcal{W}(\mathcal{LE}) \subset \mathcal{W}(\mathcal{EE}) = \mathcal{W}(\mathcal{SPT}) \subset \mathcal{W}(\mathcal{TR}) = \mathcal{W}(\mathcal{CT}) = \mathcal{W}$$

where the inclusions have been shown to be strict.

4.4 Results on Specific Forms of Knowledge.

From a computational point of view, weak sense of direction is strictly *stronger* than topological awareness. In fact, there exists a class of graphs $\mathcal{S}$ where some problems are solvable in $\mathcal{W}/\mathcal{S}$ but not in $\mathcal{G}/\mathcal{S}$, even with the additional knowledge of the topology. Consider, for example, the graph G of Figure 5 (also called minimum identity graph) and let us assume that (labeled) topological awareness ($\mathcal{TA}$) is available at every node. Even with this additional knowledge, there exist labelings (for example, the one shown in the Figure) for which the leader election cannot be solved That is, $\mathcal{LE}$ can be solved with knowledge $\mathcal{TA}$ only on a strict subset of $\mathcal{G}/G$. However, even without topological awareness, $\mathcal{LE}$ can be solved on G with *any* weak sense of direction. (i.e., in every $(G, \lambda) \in \mathcal{W}/G$).

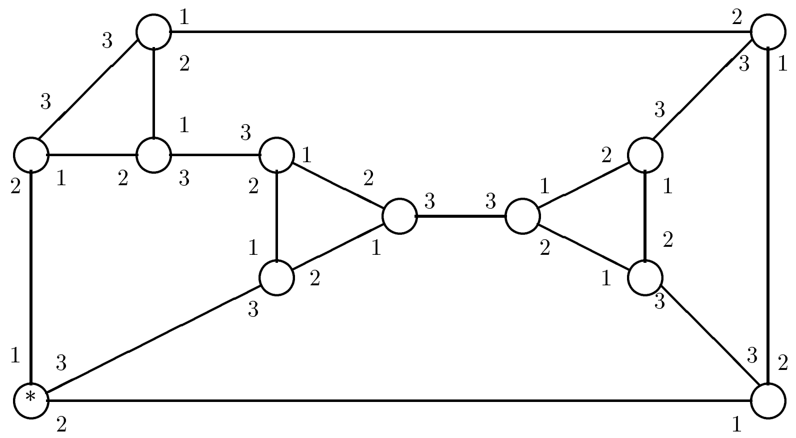

Figure 5: A symmetric labeling of the minimum regular identity graph.

4.5 Open Problems.

From a computational point of view, with (weak) sense of direction no other form of knowledge is necessary. A crucial question is whether sense of direction is the only property with such a computational power, or, in other words, there exists some other knowledge $\mathcal{K}$ with the same property. If such a form of knowledge exists, it must lie between (labeled) topological awareness and complete knowledge of the topology.

A particularly interesting open problem is the characterization of the computability relationship between sense of direction and other forms of consistency; some results have already been recently obtained in [30].

5 Impact of Sense of Direction on Systems: Object Naming

Sense of direction can be used to construct an efficient naming scheme in systems of distributed objects [4].

A *distributed object system* consists of a collection of objects and relations between objects; each object has a state (for example, expressed by 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.

A *naming scheme* is a mechanism used for naming objects and manipulating them 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., [36, 53, 67]) 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 based on the location of the objects (e.g., internet, email addresses). On the other hand, locality of names and independence on the location are two very desirable characteristics of a naming scheme.

In [4], a naming scheme in which a name of an object depends neither on its state nor on its characteristics, not even on its location, has been proposed based on sense of direction as described in the following.

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 (G, λ) 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 an arbitrarily 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 correct communication* of names and is *location independent.*

Sense of direction has also been shown to be an interesting approach to the problem of naming in the context of object composition [56].

6 Testing and Constructing Sense of Direction

6.1 Testing Sense of Direction

Given a labeled graph (G, λ), how can we verify (decide) whether there is sense of direction? In [11], this question has been answered by showing that there exist polynomial algorithms for testing both $\mathcal{WSD}$ and $\mathcal{SD}$. In this section we briefly describe some of their results.

Given a labeled graph (G, λ) with n nodes, let M_λ and Q_λ be the $n^2 \times n^2$ matrices defined as follows:

$$M_\lambda(\langle x, x' \rangle, \langle y, y' \rangle) = 1 \Leftrightarrow \langle x, y \rangle, \langle x', y' \rangle \in E \wedge \lambda_x(\langle x, y \rangle) = \lambda_{x'}(\langle x', y' \rangle)$$

and

$$Q_\lambda(\langle x, y \rangle, \langle x', y' \rangle) = M^*_\lambda(\langle x, x' \rangle, \langle y, y' \rangle)$$

where $*$ denotes the transitive closure. Finally let $T_\lambda = Q^*_\lambda$.
It has been shown that a system (G, λ) has $\mathcal{WSD}$ iff $\forall x, y, z$ with $y \neq z$, $T_\lambda(\langle x, y \rangle, \langle x, z \rangle) = 0$; that is, (G, λ) has $\mathcal{WSD}$ iff the $n \times n$ matrices along the diagonal of T_λ are identities. The algorithm based on the above property (the best-known so far) gives a sequential time complexity of $O(n^{4.752} \log n)$. The most time-consuming steps are the transitive closures of matrices whose complexity depends on that of matrix multiplication. The sequential complexity for testing $\mathcal{WSD}$ can be further reduced in practice by considering special properties of M_λ and Q_λ.

Interestingly, it has also been shown that deciding weak sense of direction can be done efficiently in parallel. In fact, considering as model of computation a CRCW PRAM, weak sense of direction is in AC^1 for all graphs using n^6 processors.

The algorithm for deciding $\mathcal{SD}$ is more complicated and its time complexity is $O(n^{14.256} logn)$. Unfortunately, deciding sense of direction is not as efficient in parallel as deciding weak sense of direction: in fact it is in AC^1 only for some classes of graphs.

6.2 Topological Constraints for Sense of Direction

Since the algorithm for testing sense of direction on a given labeled graph (G, λ) has a high polynomial complexity, an interesting research direction is the study of how to exploit the topological property of G to find simpler testing algorithms. This is one of the motivation for studying the interplay between the topology of the system and the properties that a labeling must satisfy for having sense of direction. Such an interplay has been investigated in [31, 32], where the following questions have been considered and answered:

- In which graphs, *every* labeling guarantees the existence of sense of direction?
- For which graphs, the fact that the labeling is *symmetric* suffices for the existence of sense of direction?
- In which class of graphs every proper *edge-coloring* guarantees sense of direction

In each case, a complete characterization of the corresponding graph class has been given. Not surprisingly, an arbitrary labeling suffices for having sense of direction only in few trivial graphs; the simultaneous presence of local orientation and edge symmetry guarantees sense of direction in a larger class of graphs which includes trees and rings; finally, a coloring suffices for having sense of direction in a larger class of graphs which includes particular types of "spiked rings".

As a consequence, testing for sense of direction becomes very easy for graphs in those classes; for example, if G is a tree or a ring, the test consists in checking whether λ is symmetric and has local orientation which can be done in $O(n)$.

6.3 Constructing Sense of Direction

Since the presence of sense of direction decreases the communication costs of many distributed algorithms, the problem of its construction (e.g., in preprocessing phase) is of obvious relevance.

The problem of constructing $\mathcal{SD}$ in an unlabeled network has been studied in [63, 65]. It has been shown that any algorithm constructing the traditional sense of direction for cliques, hypercubes,

	Cost	Type of $\mathcal{SD}$
Complete Network	$\Theta(n^2)$ [63]	chordal $\mathcal{SD}$
Hypercube	$\Theta(n \log n)$ [63]	dimensional $\mathcal{SD}$
Torus ($\sqrt{n} \times \sqrt{n}$)	$\Theta(n)$ [47]	compass SD
Double Loop $\langle 1, \sqrt{n} \rangle_n$	$\Theta(n)$ [47]	chordal $\mathcal{SD}$

Table 6: Constructing $\mathcal{SD}$.

and tori, exchanges at least $\Omega(e - \frac{1}{2}n)$ messages in a network with n nodes and e edges. The same result hold for constructing a chordal *calSD* in arbitrary graphs. Algorithms matching the lower bounds have also been proposed [63]; most of the solutions cannot avoid a complete flooding of the communication links in the network.

These results are not attractive for dense topologies, like cliques and hypercubes, while they are more relevant for sparse graphs such as tori or chordal rings of constant degree (see Table 6).

In systems with sense of direction, the number of labels lies between the maximum, degree of the graph (in which case it is *minimum*) and the number of links; it is considered *small* if it is not grater than the number of vertices. The problem of constructing graphs having a sense of direction with the minimum number of labels has been extensively studied for regular graphs, and will be discussed in Section 7. The problem of constructing senses of direction that use a small number of labels has been addressed in [8, 7], where techniques are presented for constructing large labeled graphs with sense of direction and a small number of labels; these techniques are based on the idea of constructing large networks by appropriately connecting smaller ones. As a consequence of this approach, minimal senses of direction can be obtained in Cayley graphs. This result is the first attempt of developing a set of techniques for constructing sense of direction in graphs by using primitive constructs.

6.4 Open Problems.

There is a clear need for a more efficient testing technique for weak sense of direction and, especially so, for sense of direction. An interesting open problem is to determine the class of graphs for which every labeling with a given "easy-to-test" property P (other that local orientation, symmetry and colorings), guarantees the existence of sense of direction; here, "easier" means less costly than the general technique for testing sense of direction. Obviously important is also the determination of such labeling properties P.

The cost of constructing sense of direction is known only for very specific topologies and very specific labelings. The development of more general techniques is an open research direction; in particular, it would be interesting to explore more efficient techniques for sense of direction in arbitrary topologies. The modular bottom-up approach to the construction of sense of direction with a small number of labels is a promising research direction; where it can lead it is still unknown.

7 Minimal Sense of Direction and Regular Graphs

The properties of sense of direction with the minimum number of labels has been extensively studied for regular graphs. In this section we consider only 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.*

An interesting and important property of regular graphs with minimal sense of direction is that, for every string $\alpha \in \Sigma^+$ of labels, there exists a walk with that sequence of labels; actually, there is one starting from *every* node.

7.1 Cycle Symmetry and Minimality.

In [23] it has been shown that the existence of a minimal sense of direction is related to the notion of *cycle symmetry* of a graph: a graph is cycle symmetric if each node belongs to the same number of cycles of the same length. More precisely, it has been shown that cycle symmetry is a *necessary* condition for the existence of a minimal sense of direction in a regular graph.

Cycle symmetry has been studied in relations to other notions of symmetries in graphs; in particular it has been compared to *vertex transitivity.* A graph G is vertex transitive if, $\forall x, y \in V$, there exists an automorphism ρ such that $(\rho(x), \rho(y)) \in E$ iff $(x, y) \in E$.

Since a vertex transitive graph is also cycle symmetric, a weaker necessary condition for the existence of minimal sense of direction is given by vertex transitivity [23]. Notice that the two notions are not equivalent; in fact it has been recently shown that a cycle symmetric graph is not necessarily vertex transitive [34].

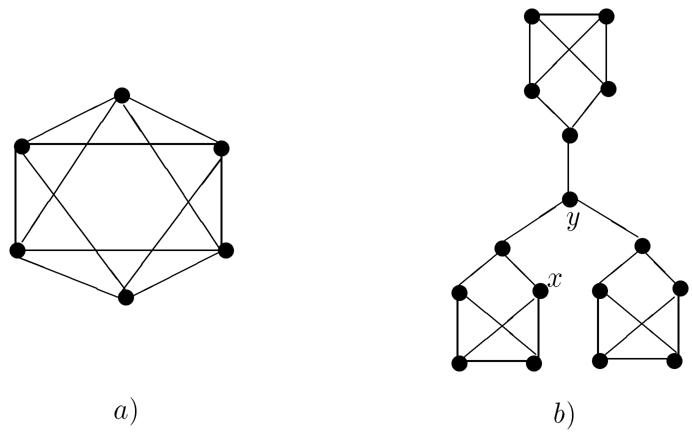

Figure 6: Regular graphs for which a) there exists a minimal $\mathcal{SD}$, b) there exists no minimal $\mathcal{SD}$.

An example of graph that is not cycle symmetric is given in Figure 6 (b) (e.g., node x belongs to a cycle of length five, while node y does not belong to any cycle of such a length): thus, this graph cannot be labeled using three labels so to have a sense of direction.

On the other hand, the graph of figure 6 (a) is cycle symmetric; in this case it is easy to see that there exists a minimal $\mathcal{SD}$.

However, cycle symmetry is *not* a *sufficient condition* for the existence of minimal sense of direction. As an example consider the Petersen's graph G (see Figure 7); this graph is vertex transitive and, thus, cycle symmetric, but it has been shown in [23] that for any labeling λ which uses three labels, (G, λ) does not have sense of direction. In order to have a sense of direction a fourth label is necessary.

7.2 Surrounding Symmetry, Cayley Graphs, and Minimality.

Necessary and sufficient conditions for a labeled graph (G, λ) to have minimal sense of direction have been given in [29] for the case of symmetric labelings and in [33] for the case of non-symmetric labelings.

We remind that a labeling λ is *symmetric* if there exists a bijection $\psi : \Sigma \rightarrow \Sigma$ such that for each $\langle x, y\rangle \in E$, $\lambda_y(\langle y, x\rangle) = \psi(\lambda_x(\langle x, y\rangle))$.

In the case of symmetric labelings, minimality is linked to the notion of surrounding (discussed in

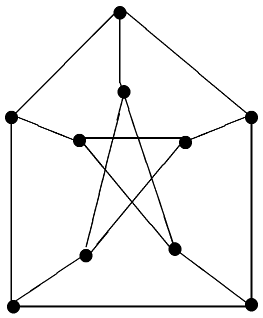

Figure 7: Petersen's graph.

Section 4): a regular graph with symmetric labeling has a minimal $\mathcal{SD}$ iff it is *surrounding symmetric* (i.e., if every node has the same surrounding) [29].

On the other hand, a labeled graph (G, λ) is surrounding symmetric iff it is a Cayley graph with a Cayley labeling. A Cayley graph is a graph where nodes correspond to the elements of a group and edges correspond to the action of the generators; a Cayley graph has a Cayley labeling when labels on the edges correspond to the generators of the group.

As a consequence, we have a characterization of minimal $\mathcal{SD}$ in terms of Cayley graphs: a regular graph with symmetric labeling has a minimal $\mathcal{SD}$ iff it is a Cayley graph with a Cayley labeling [29]; this result holds also for directed graphs [10].

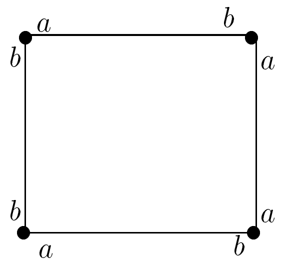

Figure 8: A minimal $\mathcal{SD}$ with a non-symmetric labeling.

If the graph is not symmetric, this necessary and sufficient condition does not hold. Consider, for example, the labeled graph (G, λ) shown in in Figure 8. In this case, the graph G is a Cayley graph but the labeling λ does not correspond to the set of generators; however, it is easy to verify that (G, λ) has minimal sense of direction.

When the labeling is not symmetric, a complete characterization has been given in [33]: A regular labeled graph (G, λ) has a minimal $\mathcal{SD}$ iff G is the graph of a semigroup S which is the direct product of a group and a (particular type of semigroup called) left-zero semigroup, *and* λ corresponds to the generators of S. The graph of a semigroup is the graph where nodes correspond to the elements of the semigroup and the edges correspond to the action of the generators.

7.3 Open Problems.

The general characterization of minimal sense of direction in non-regular graphs is still a challenging open question.

Another 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 of the graph) is it always possible to do it with $d(G)+1$ labels ? This property has been conjectured and it has been experimentally observed to hold in all undirected graphs which have been considered so far. In *undirected* graphs, the conjecture has been recently disproved by showing a graph G where the minimum number of labels for having a sense of direction is $d(G)+2$, where $d(G)$ is the out-degree of the graph [6].

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

8 Concluding Remarks

The range of the applications of sense of direction continues to grow; this is not surprising as the research is only starting and the evidence of the positive impact is already quite strong, as this paper hopes to have shown. For example, the notion of sense of direction has just recently found applications in the area of self-stabilization [12, 14].

The scope of the investigations on sense of direction and on the consistency of the labelings continues to expand. For example, research has recently focused to consider cases where the communication is not point-to-point, i.e., the labeling λ is not injective. These cases model systems which use more advanced communication and interconnection technology such as *buses*, *optical* networks, *wireless* communication media, etc. Preliminary investigations have identified the notion of *backward sense of direction* as the relevant consistency property and results are very promising [30].

The variety in the nature and scope of the problems still open, several indicated throughout the paper, offers a unique opportunity for researchers with very distinct background.

A portal on Structural Information, and in particular on sense of direction, is currently under construction; the interim home page of sense of direction can be accessed at:
http://www.site.uottawa.ca/∼flocchin/sdlist.html.

Acknowledgements. The authors would like to thank the Pontignano Research School for providing the initial motivation to write this survey. They would also like to thank the anonymous referees whose comments have contributed to improve the presentation of the material.

This work has been partially supported by ARC, 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] G. v. Bochmann, P. Flocchini, and D. Ramazani. Distributed objects with sense of direction. In *1st Workshop on Distributed Data and Structures*, pages 1–15, Orlando, 1998.

[5] P. Boldi, B. Codenotti, P. Gemmell, S. Shammah, J. Simon, and S. Vigna. Symmetry breaking in anonymous networks: Characterizations. In *Proc. of 4th Israeli Symposium on Theory of Computing and Systems*, pages 16–26, Jerusalem, 1996.

[6] P. Boldi and S. Vigna. Personal Communication.

[7] P. Boldi and S. Vigna. Coverings that preserve sense of direction. *Information Processing Letters*. To appear.

[8] P. Boldi and S. Vigna. On some constructions which preserve sense of direction. In *Proc. of 3rd International Colloquium on Structural Information and Communication Complexity*, pages 47–57, Siena, 1996.

[9] P. Boldi and S. Vigna. Computing vector functions on anonymous networks. In *Proc. of 4th International Colloquium on Structural Information and Communication Complexity*, pages 201–214, Ascona, 1997.

[10] P. Boldi and S. Vigna. Minimal sense of direction and decision problems for Cayley graphs. *Information Processing Letters*, 64:299–303, 1997.

[11] P. Boldi and S. Vigna. On the complexity of deciding sense of direction. *SIAM Journal on Computing*, 29(3):779–789, 2000.

[12] A. Bui, A.K. Datta, F. Petit, and V. Villain. Snap-stabilizing pif algorithm in the tree networks without sense of direction. In *Proc. of 6th International Colloquium on Structural Information and Communication Complexity*, pages 32–46, Lacanau, 1999.

[13] 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.

[14] AK Datta, S Gurumurthy, F Petit, and V Villain. Self-stabilizing network orientation algorithms in arbitrary networks. In *ICDCS 2000, 20th International Conference on Distributed Computing Systems*, pages 576–583, 2000.

[15] K. Diks, S. Dobrev, E. Kranakis, A. Pelc, and P. Ružička. Broadcasting in unlabeled hypercubes with linear number of messages. *Information Processing Letters*, 66:181–186, 1998.

[16] K. Diks, E. Kranakis, and A. Pelc. Broadcasting in unlabeled tori. *Parallel Processing Letters*, 8:177–188, 1998.

[17] K. Diks, E. Kranakis, and A. Pelc. Perfect broadcasting in unlabeled networks. *Discrete Applied Mathematics*, 87:33–47, 1998.

[18] S. Dobrev. Leader election using any sense of direction. In *Proc. of 6th International Colloquium on Structural Information and Communication Complexity*, pages 93–104, Lacanau, 1999.

[19] S. Dobrev and P. Ružička. 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, Ascona, 1997.

[20] S. Dobrev and P. Ružička. Broadcasting in anonymous unoriented torus. In *Proc. of 24rd International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 50–62, Smolenice-Castle, 1998.

[21] S. Dobrev and P. Ružička. Yet another modular technique for efficient leader election. In *Proc. of 25th Annual Conference on Current Trends in Theory and Practice of Informatics*, pages 312–321, Slovakia, 1998. Springer-Verlag.

[22] S. Dobrev, P. Ružička, and G. Tel. Time and bit optimal broadcasting in anonymous unoriented hypercubes. In *Proc. of 5th International Colloquium on Structural Information and Communication Complexity*, pages 173–187, Amalfi, 1998.

[23] P. Flocchini. Minimal sense of direction in regular networks. *Information Processing Letters*, 61:331–338, 1997.

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

[25] 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, Antalya, 1994.

[26] 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.

[27] P. Flocchini, B. Mans, and N. Santoro. Sense of direction: definition, properties and classes. *Networks*, 32(3):165–180, 1998.

[28] P. Flocchini, A. Roncato, and N. Santoro. Computing on anonymous networks with sense of direction. *Theoretical Computer Science*. To appear.

[29] P. Flocchini, A. Roncato, and N. Santoro. Symmetries and sense of direction in labeled graphs. *Discrete Applied Mathematics*, 87:99–115, 1998.

[30] P. Flocchini, A. Roncato, and N. Santoro. Backward consistency and sense of direction in advanced distributed systems. In *Proc. of the 18th A.C.M. Symposium on Principles of Distributed Computing*, pages 189–198, Atlanta, 1999.

[31] P. Flocchini and N. Santoro. Topological constraints for sense of direction. In *Proc. of 2nd International Colloquium on Structural Information and Communication Complexity*, pages 27–38, Olympia, 1995.

[32] P. Flocchini and N. Santoro. Topological constraints for sense of direction. *International Journal on Foundations of Computer Science*, 9(2):179–198, 1998.

[33] S. Foldes and J. Urrutia. Sense of direction, semigroups, and cayley graphs. Manuscript, 1998.

[34] J-L Fouquet and G. Hahn. Cycle regular graphs need not be transitive. Manuscript, 1999.

[35] 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.

[36] International Standard Organization. ISO/IEC JTC1, information technology - open distributed processing - naming framework, ISO/IEC DIS147771, July 1997.

[37] 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.

[38] A. Itai, S. Kutten, Y. Wolfstahl, and S. Zaks. Optimal distributed t-resilient election in complete networks. *I.E.E.E. Transactions on Software Engineering*, 16(1):415–420, 1990.

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

[40] 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.

[41] E. Korach, S. Moran, and S. Zaks. Optimality of distributed constructions of minimum weight and degree restricted spanning trees in a complete network of processors. *SIAM Journal on Computing*, 16(2):231–236, 1987.

[42] E. Kranakis and D. Krizanc. Distributed computing on anonymous hypercubes. *Journal of Algorithms*, 23:32–50, 1997.

[43] E. Kranakis and N. Santoro. Distributed computing on anonymous hypercubes with faulty components. In *Proc. of 6th International Workshop on Distributed Algorithms*, pages 253–263, Haifa, 1992.

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

[45] 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.

[46] N. Lynch. *Distributed Algorithms*. Morgan-Kaufmann, 1995.

[47] B. Mans. Optimal distributed algorithms in unlabeled tori and chordal rings. *Journal on Parallel and Distributed Computing*, 46(1):80–90, 1997.

[48] B. Mans and N. Santoro. Optimal elections in faulty loop networks and applications. *IEEE Transactions on Computers*, 4(3):286–297, 1998.

[49] 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.

[50] 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.

[51] Y. Metivier, A. Muscholl, and P.A. Wacrenier. About the local detection of termination of local computations in graphs. In *Proc. of the 4th International Colloquium on Structural Information and Communication Complexity*, pages 188–200, Ascona, 1997.

[52] N. Nishikawa, T. Masuzawa, and N. Tokura. Fault-tolerant distributed algorithm in complete networks with link and processor failures. *IEICE Transactions on Information and Systems*, J74D-I(1):12–22, Jan 1991.

[53] Object Management Group. Naming service specification clause 3, CORBA services, March 1995.

[54] Yi Pan. A near-optimal multi-stage distributed algorithm for finding leaders in clustered chordal rings. *Information Sciences*, 76(1-2):131–140, 1994.

[55] G.L. Peterson. Efficient algorithms for elections in meshes and complete networks. Technical Report TR-140, Dept. of Computer Science, Univ. of Rochester, Rochester, NY-14627, 1985.

[56] D. Ramazani, G. v. Bochmann, and P. Flocchini. Object naming and object composition. Technical Report 1135, Université de Montréal, 1998.

[57] 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.

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

[59] N. Santoro. Sense of direction, topological awareness and communication complexity. *SIGACT NEWS*, 2(16):50–56, Summer 1984.

[60] 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.

[61] G. Singh. Leader election in the presence of link failures. *IEEE Transaction on Parallel and Distributed Systems*, 7:231–236, 1996.

[62] G. Singh. Efficient leader election using sense of direction. *Distributed Computing*, 10:159–165, 1997.

[63] G. Tel. Network orientation. *International Journal of Foundations of Computer Science*, 5(1):1–41, 1994.

[64] G. Tel. Linear election in hypercubes. *Parallel Processing Letters*, 5(1):357–366, 1995.

[65] 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.

[66] A.M. Verweij. Linear-message election in hypercubes. Manuscript.

[67] R. Wieringa and W. de Jonge. Object identifiers, keys, and surrogates - objects identifiers revisited. *Theory and Practice of Object Systems*. To appear.

[68] M. Yamashita and T. Kameda. Computing on anonymous networks, part I: characterizing the solvable cases. *I.E.E.E. Transaction on Parallel and Distributed Computing*, 7(1):69–89, 1996.