Interconnection Networks: Graph- and Group Theoretic Modelling

Christian Lavault

Interconnection Networks: Graph- and Group-Theoretic Modelling

Christian Lavault*

E-mail: lavault@lipn.univ-paris13.fr

Abstract

The present paper surveys the most recent and promising results about graph-theoretic and group-theoretic modelling, from the viewpoint of relationships between *Structural Information* (e.g. *Sense of Direction*) and communication complexity of distributed leader election. The specific behaviour of various classes of networks (Cayley networks, de Bruijn and Kautz networks, etc.) is studied in terms of usual efficiency requirements, such as computability, routing, symmetry and algebraic structure.

1 Introduction

One of the main topics under investigation in distributed computing concerns the study and design of network topologies which have optimal efficiency with regard to several specific parameters, such as communication complexity of leader election, spanning tree construction, or broadcasting, ease of routing and message transmission, fault-tolerance, etc. To optimize the communication complexity of distributed algorithms, one introduces labellings on the network links in order to give the network *Structural Information*, and more precisely a *"Sense of Direction"* (or *"Orientation"*).

Above all, this paper is a survey on the present "state of the art" in graph-theoretic and group-theoretic modelling of interconnection networks in terms of sense of direction. Since it has developed to a "bench-mark", we consider the Leader Election Problem (*LEP*) to study the effect of structural information on the communication complexity.

The paper is organized as follows. In the Introduction, the model of distributed network, the notion of sense of direction and preliminaries are presented. In Section 2, we survey the properties generated from the very rich algebraic structure of the Cayley networks. In Section 3, we study the specific properties of another important family of networks, viz. de Bruijn and Kautz networks, in terms of various efficiency requirements. The impact of orientation on the communication complexity of the leader election problem is addressed in Section 4 for several network topologies. Finally, Section 5 offers some conclusions and raises open problems in this domain of research.

1.1 The Model

The model is a standard point-to-point asynchronous network $\mathcal{N}$ of N processes connected by m bidirectional communication links. As usual, the network topology is described by an undirected,

*LIPN-CNRS UPRES A 7030 (Institut *Galilée*) University of Paris-Nord 99, Av. J.B. Clément 93430 Villetaneuse. France.

connected graph (V, E) (devoid of multiple edges and loop-free): (V, E) is defined on a set V of vertices representing the processes of $\mathcal{N}$, and E is a set of edges representing the bidirectional communication links of $\mathcal{N}$ operating between neighbouring vertices. In the sequel, $|V| = N$ is the *order* of the graph and $|E| = m$ is its number of edges (or its *size*. In order to simplify notation, we also denote $\mathcal{N} = (V, E)$, .

Given a message driven algorithm $\mathcal{A}$ on $\mathcal{N}$, it is assumed that the messages are transferred on links in FIFO order, without error, and in a finite but unbounded delay (*asynchronously*). The worst-case message complexity of $\mathcal{A}$ (for a given input size N) is the maximum over all networks $\mathcal{N}$ of order N of the largest number of messages sent in any execution of $\mathcal{A}$ on $\mathcal{N}$.

1.2 Sense of Direction

The notion of sense of direction refers to this capability of a processor (or a process) to distinguish between its adjacent communication links (or its ports), according to a *globally consistent scheme* [18, 19]. In order to give a network a sense of direction, one introduces labellings on (a subset) of its links.

In an *arbitrary* distributed network $\mathcal{N} = (V, E)$, a natural globally consistent labelling on the links of the network is defined as follows in [16]. Fix a cyclic ordering of all the processors. $\mathcal{N}$ has a global sense of direction if at each processor each incident link is labelled according to the distance in the above cycle to the other nodes reached by this link. In particular, if a link, between two processors P and Q, is labelled by distance d at processor P, this link is labelled by $N - d$ at the other incident processor Q, where $N = |V|$.

Note that such a definition requires the knowledge of the order N of the network.

1.3 Preliminaries

For a given small degree, we are interested primarily in *dense networks*. A dense network $\mathcal{N}$ is one of large order N for a given *diameter* D, defined as the maximum *distance* between all nodes pairs in $\mathcal{N}$. Here, the distance between two nodes refers to the smallest number of hops between these two nodes. Obviously, a dense network $\mathcal{N}$ allows the interconnection of a large number of processing elements with relatively small communication delay.

Besides density, *vertex symmetry* (or *vertex-transitivity*) is another desirable attribute of an efficient interconnection network topology. This notion of symmetry implies that for any two nodes $u, v \in V$ there exists a label preserving automorphism $\varphi \in \mathrm{Aut}(G)$ such that $\varphi(u) = v$: informally, a vertex symmetric (or vertex-transitive) network looks the same from any node. This property allows the use of identical routing algorithms at every node, and makes it possible to define a natural labelling which provides the network with a sense of direction. Many well-known interconnection networks, such as complete networks, Rings, tori, hypercubes, cube-connected cycles, n-stars, etc., are examples of such vertex symmetric networks. Most of them belong to the class of *Cayley graphs* which are connected graphs constructed from a *group* and a *set of generators* as defined in section 2. (See also [1, 7, 12, 13, 20, 21].)

In Section 3 we deal with another class of network, whose characteristic is to have the *largest number of vertices for given maximum degree* Δ *and diameter* D, viz. de Bruijn and Kautz networks. Though they are *not* vertex symmetric, these networks enjoy very interesting properties, such as having an optimal number of nodes (for small value of D or Δ), easy routings, an optimal fault-tolerance, the feasibility for designing efficient "consensus protocols", symmetry of extensions, possibilities of quasi-optimal generalizations in all respects, etc.
Notation.

$\mathcal{N}$	Group	Gen.	Order
Or. Ring	C_n	$(1, 2, \ldots, n)$	n
2-Ring	D_n	$(1, 2, \ldots, n), \rho_n$	$2n$
Torus	$(\mathrm{C}_n)^d$	direct product	n^d
n-cube	Γ_n	$\varphi_{\{1\}}, \ldots, \varphi_{\{n\}}$	2^n
Star	S_n	$(1, k)$ $1 < k \leq n$	$n!$
Bubble	S_n	$(k-1, k)$ $1 < k \leq n$	$n!$
Pancake	S_n	$\rho_k, \overline{\rho}_k$ $(1 < k \leq n)$	$n!$

Table 1: Examples of Cayley networks.

In the sequel, we use the usual terminology of group theory and graph theory. Since we only consider finite groups, the groups are mainly represented as *permutation groups.* The following notation is used:

Z_q for the ring of integers $0, 1, \ldots, q-1$ (modulo q), and $(\mathrm{Z}_p)^n$ for the n-dimensional vector space over Z_p (p being a prime); S_n for the symmetric group on n symbols; $\langle S \rangle$ for the group generated by the set S of generators; I_n for the trivial identity group consisting of the identity permutation in S_n; e for the identity element of a group; and $|u|$ for the order of $u \in G$, i.e. the smallest positive integer k such that $u^k = e$. Given a (finite) group G and $H \leq G$ (H is a subgroup of G), $(G : H)$ denotes the (finite) index of H in G, i.e. the (finite) number of cosets of H in G.

2 Cayley Networks

2.1 Definition of Cayley Networks

Definition 2.1 *Let G be a group and let $S \subseteq G$ be a set of generators of G. The Cayley network $\mathcal{N}_S$ of G with set of generators S is $\mathcal{N}_S = (V, E)$, where $V = G$ and $E = \{(u, v) \ / \ u^{-1}v \in S\}$. We assume that $S = S^{-1}$, where S^{-1} is the set of g^{-1} such that $g \in S$, so that $\mathcal{N}_S$ can be viewed as undirected. To avoid loops in the network $\mathcal{N}_S$, we assume $e \notin S$; further, if $g = g^{-1}$ then we identify the edges g and g^{-1}. The Cayley network $\mathcal{N}_S$ has $|G| = N$ nodes and the degree of each node is $|S|$, denoted by $\delta(S)$.*

For each $g \in G$, let $[g]_S$ denote the least number of generators from S needed to represent g (with possible repetitions). The diameter of $\mathcal{N}_S$ is $D(S) = \max\{[g]_S \ / \ g \in G\}$.

The resulting Cayley network depends on the set S of generators. On the one hand we can choose $S = G$, in which case $\mathcal{N}_S$ is the complete network K_N with $|G| = N$ nodes. On the other hand, we are usually interested in "small" sets of generators with "not too big" diameter. In fact, as pointed out in [12], we would want sets S of generators minimizing the quantity $D^2(S) \cdot \sum_{g \in S} |g|^2$. It is well known that such small S do exist, since every finite group G has a set of generators of order $O(\log |G|)$. (We refer to [1, 2, 20, 21] for more indications on the importance of such groups.)

Example. To simplify notation the elements of S are listed without their inverses and multiplication of permutations is considered to the left. Throughout, we assume that n is arbitrary but fixed.

The first four Cayley networks of Table 1 are arising from cyclic, abelian and dihedral groups. With cyclic groups, we obtain a variety of tori: the oriented Ring, the double Ring (with $n \neq 2$) and the d-dimensional Torus.

The group of automorphisms of the n-dimensional Hypercube is denoted by Γ_n; it is the group of bit-complement automorphisms, and $\Gamma_n \cong (\mathrm{Z}_2)^n$. The generators $\varphi_{\{i\}}$ of Γ_n are such that $\varphi_{\{i\}}(b_1, \ldots, b_n)$ is the sequence of bits obtained from $(b_1, \ldots, b_n)$ by complementing the ith bit, while leaving the others unchanged.

The last three Cayley networks are examples arising from the symmetric group S_n.

In Table 1, ρ_k and $\overline{\rho}_k$ are the reflection permutations:

$$\rho_k = (k\ k-1\ \cdots\ 1),$$

and

$$\overline{\rho}_k = \begin{pmatrix} n-k+1 & n-k+2 & \cdots & n \\ n & n-1 & \cdots & n-k+1 \end{pmatrix}.$$

2.2 Properties of Cayley Networks

2.2.1 Basic Properties

Definition 2.2 *Let $\mathcal{N}_S$ be a Cayley network. A natural labelling $\mathcal{L}_S$ on $\mathcal{N}_S$ is such that the label of the edge $(u, v) \in E$ is $\mathcal{L}_S(u, v) = u^{-1}v$. The resulting labelled Cayley network is thus denoted by $\mathcal{N}_S[\mathcal{L}_S]$. Any automorphism φ of $\mathcal{N}_S[\mathcal{L}_S]$ is such that the edge-labels are preserved under φ:*

$$(\forall (u, v) \in E) \quad \mathcal{L}_S(u, v) = \mathcal{L}_S(\varphi(u), \varphi(v)) \tag{1}$$

The group of automorphisms of $\mathcal{N}_S$ satisfying (1) is denoted by $\mathrm{Aut}(\mathcal{N}_S[\mathcal{L}_S])$.

Theorem 2.1 [7, 13] *Every Cayley network is vertex symmetric (or vertex-transitive), in the sense that*

$$(\forall u, v \in G)(\exists \varphi \in \mathrm{Aut}(\mathcal{N}_S)) \quad \varphi(u) = v.$$

The automorphism φ is thus label preserving and such that

$$(\forall x \in G) \quad x \mapsto vu^{-1}x.$$

This automorphism is uniquely determined from u and v, which makes the action of Aut($\mathcal{N}_S[\mathcal{L}_S]$) on the vertices of $\mathcal{N}_S$ regular. □

Theorem 2.2 [1] *Consider a Cayley network $\mathcal{N}_S$ defined by a set S of generators on n symbols (the "permutation group representation" of $\mathcal{N}_S$). $\mathcal{N}_S$ is edge symmetric iff for every pair of generators $(g_1, g_2) \in S$ there exists a permutation of the n symbols that maps S into S, and, in particular, maps g_1 into g_2.* □

A labelled Cayley network $\mathcal{N}_S[\mathcal{L}_S]$ is said *strongly symmetric* if it is both vertex and edge symmetric. Examples of strongly symmetric Cayley networks are the Complete network K_N, the n-Hypercube, the n-Star, the n-Bubble Sort and the n-Pancake Sort networks.

Theorem 2.3 [13] *The group of automorphisms of* $\mathrm{Aut}(\mathcal{N}_S[\mathcal{L}_S])$ *is isomorphic to G.* □

Note that Cayley network may thus also be characterized as those transitive networks whose automorphism group has a regular transitive subgroup (cf. [7]).

2.2.2 Computability of Boolean Functions

Most of the results of this section are in [12, 13, 14]. Let $\mathrm{S}(f)$ denote the group of permutations in S_n that leave f invariant on all inputs. If a Boolean function $f \in B_N$ is computable on a distributed network $\mathcal{N}$, then $\mathrm{S}(f) \geq \mathrm{Aut}(\mathcal{N}_S[\mathcal{L}_S])$. The converse is not necessarily true, in general. However, it is true for the class of distributed Cayley networks:

Theorem 2.4 [12] *For any Boolean function $f \in B_N$, f is computable in the distributed network $\mathcal{N}_S[\mathcal{L}_S]$ iff $\mathrm{S}(f) \geq \mathrm{Aut}(\mathcal{N}_S[\mathcal{L}_S])$.* □

In general, oriented distributed Cayley networks are more powerful than unoriented distributed Cayley networks, in the sense that the former can compute more Boolean functions than the latter. A characterization is now given of those abelian groups G which have a canonical set of generators S such that the network $\mathcal{N}$ computes more Boolean functions than the network $\mathcal{N}_S[\mathcal{L}_S]$.

Definition 2.3 *The labelling $\mathcal{L}_S$ is said to be* strong iff *there is a Boolean function on $|G| = N$ variables which is computable in the network $\mathcal{N}_S[\mathcal{L}_S]$ but not computable in $\mathcal{N}$.*

Theorem 2.5 [13] *For any automorphism $\varphi \in \mathrm{Aut}(G)$, let*

$$S_\varphi = S \cup \varphi(S) \cup \varphi^2(S) \cup \cdots.$$

If $\varphi(S) \notin \langle g \rangle$ for some $g \in S$, then $\mathcal{L}_{S_\varphi}$ is strong.

If S is a canonical set of generators for the cyclic group $G = \mathrm{C}_n$, then $\mathcal{L}_S$ is strong exactly when $n \neq 2, 3, 4, 5$. □

Note that the 11 abelian groups $\oplus_{n \in A} \mathrm{C}_n$, where $A \subseteq \{2, 3, 4, 5\}$ and $|A| \geq 2$ have an interesting behaviour. Although it is proven in [13] that the networks $\mathcal{N}_S$ and $\mathcal{N}_S[\mathcal{L}_S]$ cannot "distinguish" the Boolean functions they can compute from their automorphism groups alone, it is also shown that the labelled network $\mathcal{N}_S[\mathcal{L}_S]$ can actually compute more Boolean functions that $\mathcal{N}_S$. In particular, for the 11 abelian groups cited above, there exist Boolean functions which are computable on $\mathcal{N}_S$ but such that $\mathrm{S}(f) \geq \mathrm{Aut}(\mathcal{N}_S)$.

2.3 Extensions of Cayley Networks

While most symmetric networks considered in the literature can be viewed as Cayley networks, there remain certain vertex symmetric graphs that cannot be represented as Cayley graphs. A prime example is the Petersen graph.

Theorem 2.6 [1] *Every vertex symmetric graph can be represented as the quotient of two Cayley graphs.* □

The following conjecture constitutes a very important open problem.

Conjecture 2.1 *Every Cayley graph is Hamiltonian, i.e., has a Hamiltonian cycle. Furthermore, every vertex symmetric graph has a Hamiltonian path.*

A weaker form is Alspach's conjecture: every connected $2k$-regular Cayley graph on a finite abelian group can be partitioned into k Hamiltonian cycles. □

For example, the Petersen graph again, is a vertex symmetric graph but not a Cayley graph; it has a Hamiltonian path but has no Hamiltonian cycle. Besides, it is the only graph of order ≤ 10 without a Hamiltonian cycle such that the deletion of one vertex yields a Hamiltonian graph.The following two results may be regarded as a step towards the proof of Alspach's conjecture.

Theorem 2.7 [5] *Every connected 4-regular Cayley graph on a finite abelian group can be decomposed into two edge-disjoint Hamiltonian cycles.* □

Most symmetric interconnection networks that are found in the literature can be represented as Cayley networks. Along these lines, the new Star networks and Pancake networks (as designed in [1] from the group-theoretic model of Cayley networks) feature relevant directions for extensions and representations of Cayley networks.

3 De Bruijn and Kautz Networks

The problem of constructing large networks of a given degree and diameter—known as the (Δ, D) graph problem—is a well-studied extremal graph theory problem. The maximum number of vertices $N(\Delta, D)$ of a network of given maximum degree Δ and diameter D is bounded by the following relation known as the *Moore bound* (see also inequality **??**):

$$\begin{aligned} \text{if} \quad \Delta \geq 3 \qquad N(\Delta, D) &\leq \frac{\Delta(\Delta-1)^D - 2}{\Delta - 2}, \\ \text{if} \quad \Delta = 2 \qquad N(\Delta, D) &\leq 2D + 1. \end{aligned} \tag{2}$$

It is known that this bound is unachievable except in the following cases: N cliques ($D = 1$), $(2D + 1)$ cycles ($\Delta = 2$), Petersen graph ($D = 2$, $\Delta = 3$), Hoffman-Singleton graph ($D = 2$, $\Delta = 7$), and, possibly, $D = 2, \Delta = 57$ (no known construction yet). (Refer to [15] for more information on Moore graphs.) As far as the the Moore bound is concerned, The best of the general known classes of networks are de Bruijn or Kautz networks, which are defined by using words over alphabets. (See [6].)

3.1 Definitions of de Bruijn and Kautz Networks

Definition 3.1 (de Bruijn Networks $B(q, D)$)

The vertices of a directed de Bruijn network are the words of length D constructed on an alphabet of q letters. Let $(x_1, \ldots, x_D)$, with $x_i \in \{0, \ldots, q-1\}$ $(1 \leq i \leq D)$, denote a vertex. There is an arc between $(x_1, \ldots, x_D)$ and all vertices $(x_2, \ldots, x_D, \alpha)$, where α is any letter from the alphabet. This digraph is thus q-regular, it has $N = q^D$ vertices, and its diameter is D.

The associated undirected de Bruijn network $B(q, D)$ is obtained from the above digraph by forgetting the orientations of the arcs, removing the self-loops, and replacing each double edge with a single edge. Now the vertex $(x_1, \ldots, x_D)$ is adjacent to all the vertices $(x_2, \ldots, x_D, \alpha)$ and $(\alpha, x_1, \ldots, x_{D-1})$, where α is any letter of the alphabet. Therefore, the undirected de Bruijn network $B(q, D)$ has maximum degree $\Delta = 2q$ and it is not regular. Indeed, $B(q, D)$ has $N = q^D$ vertices, and it easily seen that $B(q, D)$ has $N - q$ vertices of degree $2q$, $q^2 - q$ vertices of degree $2q - 1$, and q vertices of degree $2q - 2$. (See Figure 1.)

Kautz networks $K(q, D)$ are defined in a similar way: the vertices are the word of length D constructed over an alphabet of $q+1$ letters. $K(q, D)$ has maximum degree $\Delta = 2q$ and it is not regular. $K(q, D)$ is indeed of order $N = q^D + q^{D-1}$; the graph has $N - q^2 - q$ vertices of degree $2q$ and $q^2 + q$ vertices of degree $2q - 1$.

It is interesting to compare (as in Table 2) the order of $B(q, D)$, and n-Hypercube, for given values of Δ and D. Since the n-Hypercube is defined only for $\Delta = D = n$, Table 2 is given for $\Delta = D = 4, 6, 8, 10, 12$.

From Table 2, it is clear that, for the same values of Δ and D, $B(q, D)$ and $K(q, D)$ have many more vertices than the n-Hypercube. Furthermore, for a same given number of vertices,

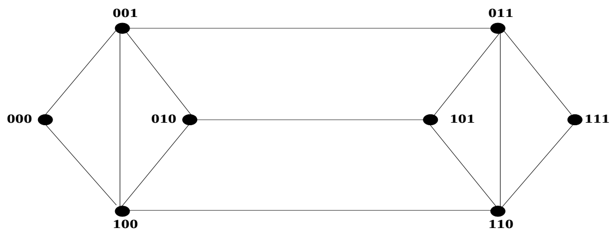

Figure 1: The Undirected de Bruijn Graph $B(2,3)$.

$\Delta = D$	n-Hypercube	de Bruijn
4	16	16
6	64	72
8	256	65,536
10	1,024	9,765,625
12	4,093	2,176,782,336

Table 2: Order of $B(q,D)$ and n-Hypercube.

different values of Δ and D are possible. For example, with 256 vertices, there is a unique 8-Hypercube of degree and diameter 8. By contrast, the de Bruijn networks with 256 vertices provide four choices: $B(4,8)$, $B(2,8)$, $B(4,4)$, and $B(16,2)$.

3.2 Properties of de Bruijn and Kautz Networks

De Bruijn and Kautz Networks are those networks which have the largest number of vertices for a given degree and diameter.

For any fixed D, the number of vertices of these digraphs is of the same asymptotic order as the *directed* Moore bound: $1 + q + \cdots + q^D$ (which is different from inequation 2). Whenever $D = 2$, the order of $K(q,2)$ is $q+q^2$, and $K(q,2)$ is an *optimal graph*, since it has been shown that the directed Moore bound cannot be achieved if $q > 1$ and $D > 1$ de Bruijn and Kautz networks are neither vertex symmetric, nor edge symmetric; in this sense, Cayley networks, and the n-hypercube for example, are better. However, de Bruijn and Kautz networks enjoy interesting algebraic properties.

$B(q,D)$ gives asymptotically optimal protocols for computing simple functions that do not require *metric regularity* (MAX, MIN, AND, OR, etc.). (See [14, 15].) For the communication complexity of certain classes of consensus protocols such as the leader election, optimality is achieved in $B(2,D)$ or $K(2,D)$ ($\Delta = 4$ and $D = \lg N$).

4 The Impact of Sense of Direction

Among the distributed problems considered, the leader election problem is certainly the most significant, and it has developed to a "bench-mark" to study the impact of structural information on the communication complexity for several network topologies.

In this section, we focus on the leader election problem in *named and anonymous* distributed networks. For most distributed network topologies, the availability of sense of direction has been shown to have positive impact on the communication complexity of the election problem

$\mathcal{N}$	Order	Degree	Diam.
Ori. Ring	n	2	$n-1$
Unor. Ring	$2n$	3	n
d-Torus	n^d	$2d$	$d(n-1)$
$BF(n)$	$n2^n$	4	$\lfloor 3n/2 \rfloor$
Chord. Ring	n	4	$\Theta(\log n)$ (a.e.)
Hypercube	2^n	n	n
n-Bubble	$n!$	$n-1$	$n(n-1)/2$
n-star	$n!$	$n-1$	$\lfloor 3(n-1)/2 \rfloor$
n-Pancake	$n!$	$n-1$	$O(n)$
K_N	N	$N-1$	1
$B(q,D)$	q^D	$[2q-2, 2q]$	D
δ-Regular	N	δ	$\Omega(\log N)$ (a.e.)
Arbitr.	N	$\Delta \leq N-1$	$D \geq 1$

Table 3: Characteristics of some significant networks.

(Complete network, Circulant Graphs, Chordal Ring , etc.). As opposed to almost all networks, the existence of an orientation does not help in two important classes of networks, namely rings and d-tori.

The question whether sense of direction has positive impact on the message complexity of the election on a n-Hypercube remains open. Election can be performed in $2N \log N$ messages in the oriented hypercube, but, disappointingly, since the hypercube has only $O(N \log N)$ edges, an $O(N \log N)$ complexity is also achieved with the standard algorithm of Gallager *et al.* Recent results [10, 23] achieve an $O(N)$ message complexity for the election on the oriented hypercube. It is commonly believed that, for the unoriented case, the $O(N \log N)$ message complexity is optimal.

In the following Tables 3 and 4, n is defined as in Section 2, N denotes the order of the network $\mathcal{N}$, d is the dimension of the Torus, and δ and D the degree and the diameter of $\mathcal{N}$, respectively. In addition, general δ-regular and arbitrary networks are also quoted in Table 3 and 4. Table 3 exhibits some characteristics of relevant networks we are especially concerned with from the communication complexity point of view. Table 4 provides the communication complexity of the Leader Election in the significant *named* networks of Table 3, in the unoriented and the oriented case.

Remarks.

- The case of chordal rings and δ-regular networks is specific, since randomized bounds are given for the diameter in Tables 3 and 4. Indeed, we know from [8] that the diameter of *almost every* (a.e.) δ-regular graph of order N is at least

$$\lfloor \log_{\delta-1} N \rfloor + \left\lceil \log_{\delta-1} \ln N - \log_{\delta-1}(\frac{6\delta}{\delta-2}) \right\rceil + 1.$$

- The particular network $B(2,D)$ (or $K(2,D)$), for example, achieves a $\Theta(N \log N)$ message complexity in the unoriented case, and a $\Theta(N)$ message complexity in the oriented one. Note that $B(2,D)$ or $K(2,D)$ is only defined for $D > 1$.

- As mentioned above, a lower bound (or an upper bound) on the complexity of the election in a n-Hypercube is not yet known (see $^{(*)}$ for the unoriented n-hypercube). Hence, it is not

$\mathcal{N}$	Unoriented	Oriented
Ori. Ring	$\Theta(N \log N)$	$\Theta(N \log N)$
Unor. Ring	$\Theta(N \log N)$	$\Theta(N \log N)$
d-Torus	$\Theta(N)$	$\Theta(N)$
Chord. Ring	$\Theta(N \log N)$	$\Theta(N)$
Hypercube	$O(N \log N)^{(*)}$	$\Theta(N)$
n-Bubble	$O(nN + N \log N)$	$O(N \log N)$
n-Star	$O(nN + N \log N)$	$O(N \log N)$
n-Pancake	$O(nN + N \log N)$	$O(N \log N)$
K_N	$\Theta(N \log N)$	$\Theta(N)$
$B(2, D)$	$\Theta(N \log N)$	$\Theta(N)$
δ-Regular	$\Theta(\delta \cdot N + N \log N)$	$\Theta(N \log N)$
Arbitrary	$\Theta(\delta \cdot N + N \log N)$	$\Theta(N \log N)$

Table 4: Message Complexity of Leader Election.

clear whether the orientation helps or not in this case. Along the same lines, notice that, up to our knowledge, the message complexity of leader election on the Butterfly network $BF(n)$ is not yet known. Consequently, deciding whether orientation helps or not in this case remains an interesting open problem.

Nevertheless, Table 4 shows that the behaviour of rings, d-tori and double rings regarding the message complexity of leader election is not at all a coincidence. Also notice that the behaviour of de Bruijn networks $B(2, D)$ (or Kautz networks $K(2, D)$) is very similar to the one of Chordal Rings (at least to the behaviour of *almost every* Chordal Ring), since orientation does help from $\Theta(N \log N)$ to $\Theta(N)$ (for a.e. Chordal Ring).

It is important to emphasize that similar results do hold, comparatively, in the anonymous case.

5 Conclusions and Open Problems

In summary, the present paper offers a medley of ideas arising from the notion of sense of direction. It seems obvious that algebraic group-theoretic and graph-theoretic models are a fertile field on which structural information problems can be solved. We are thus provided with theoretical tools which allow us to take full advantage of the rich algebraic structures of large classes of networks.

Despite the very promising results of Kranakis and Krizanc (e.g. in [12, 13, 14]), we still know little about the possibilities and potentialities of Cayley networks. Borel Cayley networks, the study of dense networks as well as of the diameter of finite groups [2] and of distance-regular graphs [7, 14], etc. also open wide and fruitful directions of research. On the other hand, the class of de Bruijn and Kautz networks and their generalizations provides a rich alternative among the networks enjoying many properties and efficiency requirements of massively distributed computing. Yet, very little is also known about these networks, and especially about the generalized de Bruijn and Kautz networks [6, 15].

The notions of symmetry and density of networks are fundamental for the solution of the above questions and should be investigated in detail from a static and dynamic point of view. Along the same lines, the behaviour of networks should be considered as a whole, and not only from such or such isolated point of view. We need to take into account objective parameters such as symmetry and algebraic structure of networks, ease of routing, broadcasting, fault-tolerance (which does not only mean the connectivity of the underlying graphs), extendability and generalization of networks, and the structure of the graphs.

References

[1] **S.B. Akers and B. Krishnamurthy.** A Group-Theoretic Model for Symmetric Interconnection Networks. *IEEE Transactions on Computers, Vol. 38, No. 4*, April 1989, 555-566.

[2] **L. Babai, G. Hetyei, W.M. Kantor, A. Lubotzky and Á. Seress.** On the Diameter of Finite Groups. *Proc. of the 31st IEEE Conf. on Foundation of Computer Science*, 1990.

[3] **E. Bannai and T. Ito.** Algebraic Combinatorics I. Association Schemes. *Mathematics Lecture Note Series*, Cummings Pub. Comp., 1984.

[4] **D. Barth and A. Raspaud.** Two Edge-Disjoint Hamiltonian Cycles in the Butterfly Graph. *Information Processing letters*, to appear.

[5] **J.-C. Bermond, O. Favaron and M. Maheo.** Hamiltonian Decomposition of Cayley Graphs of Degree 4. *Journal of Combinatorial Theory*, Series B 46, 1989, 142-153.

[6] **J.-C. Bermond and C. Peyrat.** de Bruijn and Kautz Networks: a Competitor for the Hypercube. *Hypercube and Distributed Computers*, F. André and J.-P. Verjus eds., North-Holland, 1989, 279-293.

[7] **N. Biggs.** *Algebraic Graph Theory.* Cambridge University Press, 1974.

[8] **B. Bollobás and W. Fernandez de la Vega.** The Diameter of Random Regular Graphs. *Combinatorica*, **2** (2), 1982, 125-134.

[9] **L. Bougé.** *Modularité et symétrie pour les systèmes répartis ; application au langage CSP.* Thèse de Doctorat d'Etat, Université Paris VII, mars 1987.

[10] **P. Flocchini and B. Mans.** Optimal Election in Labeled Hypercubes. *Technical Report TR-231, School of Computer Science*, Carleton University, November 1993.

[11] **C.M. Hoffmann.** Group-Theoretic Algorithms and Graph Isomorphism. *Lecture Notes in Computer Science 136*, Springer-Verlag, 1982.

[12] **E. Kranakis and D. Krizanc.** Distributed Computing on Cayley Networks. *4th IEEE Symp. on Parallel and Distributed Processing*, Arlington, 1992, 222-229.

[13] **E. Kranakis and D. Krizanc.** Labeled versus Unlabeled Distributed Cayley Networks. *1st Colloquium on Structural Information and Communication Complexity*, (SICC-1), Carleton Un. Press, May 1994, This issue.

[14] **E. Kranakis, D. Krizanc and J. van den Berg.** Computing Boolean Functions on Anonymous Networks. *17th Int. Colloquium on Automata, Languages, and programming (ICALP'90)*, Lecture Notes in Computer Science 443, Springer-Verlag, 1990, 254-267.

[15] **T. Lakshman and V.K. Wei.** Distributed Computing on Regular Networks with Anonymous Nodes. *IEEE Transactions on Computers, Vol. 43, No. 2*, February 1994, 211-218.

[16] **B. Mans and N. Santoro.** On the Impact of Sense of Direction in Arbitrary Networks. *14th Int. Conf. on Distributed Computing Systems (ICDCS'94)*, June 1994, to appear.

[17] **S. Moran and Y. Wolfstahl.** Extended Impossibility Results for Asynchronous Complete Networks. *Information Processing Letters*, **26**, 1987, 145-151.

[18] **N. Santoro.** Sense of Direction, Topological Awareness and Communication Complexity. *SIGACT NEWS, 2 (16)*, 1984, 50-56.

[19] **N. Santoro, J. Urrutia and S. Zaks.** Sense of Direction and Communication Complexity in Distributed Networks. *Proc. of the 1st International Workshop on Distributed Algorithms and Graphs*, (WDAG-1), Carleton Un. Press, 1985, 123-132.

[20] **K.W. Tang and B.W. Arden.** Vertex-Transitivity and Routing for Cayley Graphs in GCR Representations. *Proc. of the 1992 ACM Symposium on Applied Computing, Vol. 2*, 1992, 1180-1187.

[21] **K.W. Tang and B.W. Arden.** Representations of Borel Cayley Graphs. *SIAM J. Discrete Math., Vol. 6, No. 4*, November 1993, 665-676.

[22] **G. Tel.** Orientation of Regular Processor Networks. *École d'informatique sur l'algorithmique répartie : calculs répartis et causalité (6-10 septembre 1993)*, Ed. IRISA, 1993.

[23] **G. Tel.** Linear Election for Oriented Hypercubes. *Research Report RUU-CS-93-39*, Utrecht University, December 1993.

[24] **G. Tel.** *Introduction to Distributed Algorithms.* Cambridge University Press, 1994.

[25] **M. Yamashita and T. Kameda.** Computing on Anonymous Networks. *Proc. of the 7th Symposium on Principles of Distributed Computing (PODC)*, 1988, 117-130.