ANNALES DE L'INSTITUT FOURIER

JEAN-CLAUDE BERMOND

A. BONNECAZE

T. KODATE

STÉPHANE PÉRENNES

PATRICK SOLÉ

Symmetric flows and broadcasting in hypercubes

SYMMETRIC FLOWS AND
BROADCASTING IN HYPERCUBES

by J.-C. BERMOND, A. BONNECAZE,
T. KODATE, S. PÉRENNES and P. SOLÉ

---

1. Introduction.

The problem we consider here is motivated by communications in interconnection networks. *Broadcasting* (also called One to All) is a communication scheme in which a given node (called initiator) sends its information to all the other nodes of the network. We consider here broadcasting in hypercubes under a circuit switched model.

Several multiprocessors with hypercube or hypercube-like topology have been designed. This topology is widely accepted as it has a logarithmic diameter and regular structure and offers high communication bandwidth.

DEFINITION 1 (Hypercube). — *The Hypercube $H(n)$ of dimension n is defined as the graph whose vertices are words of length n on the alphabet $\{0, 1\}$ and where two vertices are adjacent if and only if they differ exactly in one coordinate.*

Hypercube $H(5)$ is displayed in Figure 1.

In the circuit-switched model, a node x sends its information to a node y via a directed path (called *"circuit"* in the telecommunications terminology). There exist different ways of implementing such a model like wormhole routing; they mainly differ in the manner the circuit is established and released and how the acknowledgments are done.

Here we consider a generic model in which the communication protocol consists of rounds (or steps). A new round starts only when the preceding one is completely finished. During a round, vertices which have the information can send it to as many vertices they want (model called Δ-port, all-port or F_*), but *all the paths used for communication should be arc disjoint.* Figure 1 shows a broadcast scheme in some spanning subgraph of the hypercube of dimension 5 in two rounds. In the first round, the initiator 00000 sends the message to the five other black nodes; five disjoint paths can easily be found as $H(5)$ is 5 edge-connected. In the second round the six informed vertices inform all the others, the paths used are shown on the figure.

Figure 1. Scheme for $H(5)$ in two communication rounds

Note that other models exist like store-and-forward model (where a vertex can only communicate with a neighbor at distance 1) or 1-port models (in which a vertex can send only one message): see the book [25] or the surveys [10], [16], [18]. In first approximation the number of rounds of a broadcasting scheme represents the time needed to realize the scheme. In a more precise model, one should also consider the size of the message and the length of the communicating dipaths (as some time is needed to set up the switches in the intermediate nodes).

Broadcasting schemes under this model have been considered by many authors who studied various topologies; surveys on communication schemes under circuit switched model can be found in [6], [9].

Let us define $b_{F_*}(G)$ as the minimum number of rounds needed to complete broadcasting in the network G. Note that in a graph with maximum degree Δ, an informed vertex can inform, via edge-disjoint paths, at most Δ other vertices. Hence, a lower bound on the number of steps necessary to broadcast in a graph G with maximal degree Δ is at least $\log_{\Delta+1}|V(G)|$. This bound has been proved to be attained for cycles (see [25]), 2-dimensional toroidal grids in [23] and almost attained for n-dimensional grids in [7]. For $H(n)$, the lower bound becomes

$$b_{F_*}(H(n)) \geq \left\lceil \frac{n}{\log_2(n+1)} \right\rceil.$$

Different algorithms have been given for broadcasting in $H(n)$: McKinley and Trefftz [22] presented an algorithm of $\lceil \frac{1}{2}n \rceil$ rounds, based on edge-disjoint spanning trees. The best known algorithm has been proposed by Ho and Kao [17]. It is recursive and uses special routing (called e-routing); the algorithm is asymptotically optimal but the broadcasting time does not match the lower bound. Another problem such as multicast (where a node has to send a message to some subset of nodes) in hypercubes under the same model is examined in [24]. We prove here that:

THEOREM 2. — *For every n :*

$$b_{F_*}(H(n)) \leq \left\lceil \frac{n}{\lfloor \log_2(n+1) \rfloor} \right\rceil.$$

Our result improves all the preceding ones (see Table 1 for numerical values), is optimal for $n = 2^k - 1$, and is the best in a class of natural schemes.

The rest of the paper is organized as follows. First, we show that the problem can be formulated in the context of undirected graphs. In Section 2, the problem is related to the design of a *good sequence of multi-broadcasts* and to *flow networks*. In Section 3 we show how symmetries can be used to simplify the study of flow networks. This part is strongly related to previous work on connectivity in symmetric networks. In Section 4 the results obtained are applied in order to derive a simple condition insuring that a sequence of multi-broadcasts can be done. In Section 5 we construct two different schemes, their validity is proven by using the condition derived in Section 4.

We make use of classical tools from graph theory (see [1], [3], [25] for an introduction).

Undirected and directed models.

Usually, we model a communication network as a symmetric digraph $G = (V, E)$ where the vertex set V represents the nodes of the network and the arc set E the links between nodes. Note that communicating dipaths used during one round are arc-disjoint. But we can for simplicity consider undirected graphs and undirected communicating paths. Indeed in a broadcasting scheme useful communications are only from an informed vertex to a non-informed vertex and so a vertex can be either sending or receiving during one round but not both. Furthermore, suppose that there exists two dipaths, one from x_1 to y_1 and the other from x_2 to y_2 using two arcs in opposite directions say (u, v) and (v, u) (the dipaths cross the same edge but in opposite direction). Let the first dipath be $P_1 = P_1[x_1, u](u, v)P_1[v, y_1]$ and the second be $P_2 = P_2[x_2, v](v, u)P_2[u, y_2]$. Then x_1 can inform y_2 along $P_1[x_1, u]P_2[u, y_2]$ and x_2 can inform y_1 along $P_2[x_2, v]P_1[v, y_1]$. So, we obtain a scheme with one less pair of opposite arcs. By repeating it we obtain dipaths which do not contain opposite arcs. Hence, it can be assumed that we use edge-disjoint paths on the undirected graph.

2. Valid sequences and flows.

Given a graph G, designing a broadcast scheme ending in T communication rounds is equivalent to define a *valid sequence of sets*

DEFINITION 3. — *A sequence $\mathcal{S} = C_0, C_1, \ldots, C_T$ of sets of vertices is valid if and only if*

1) $C_0 = \{0\}$;

2) $C_t \subset C_{t+1}$;

3) $C_T = V(G)$;

4) *vertices in C_t can inform the vertices in $C_{t+1} \setminus C_t$ in one communication round, via edge-disjoint paths.*

Note that it is always easy to fulfill conditions 1), 2), 3) from construction, and we will always construct sequences $\mathcal{S}$ such that these

conditions are ensured. Hence, the problem reduces to ensure that C_t can inform $C_{t+1} \setminus C_t$ in a single communication round. When the sequence S is known, the question is indeed exactly a *multi-broadcast* problem.

DEFINITION 4 (multi-broadcast). — *In a graph G, given a set of originators O and a set of destinations D such that $D \cap O = \emptyset$, the multi-broadcast problem (O, D) consists in finding $|D|$ edge-disjoint paths from O to D ending at different nodes in D. In this case we will say shortly that O can inform D.*

According to this definition, we have to solve T successive multi-broadcast problems $(C_t, C_{t+1} \setminus C_t)$, for $0 \leq t \leq T - 1$.

Note that at the first round, C_0 reduces to the initiator; and the initiator can inform any set of λ vertices or less, if λ is the edge-connectivity of G. Indeed, in a λ edge-connected graph, there exists by Menger's theorem (or flow theorem) λ edge-disjoint paths from any vertex to any set of λ vertices.

Multi-broadcast and flows.

The multi-broadcast problem (O, D) is easily reduced to a flow problem. By a *flow network N* we will always mean a triple $N = (H, s, t)$ where H is a capacited[1] graph, and s (resp. t) a specific vertex of H called the source (resp. the sink).

Let us recall some terminology and properties regarding graphs, and flow networks.

- Given a graph G and $S \subset V(G)$, the border $\delta_G(S)$ of S denotes the set of edges between S and $\bar{S} = V(G) \setminus S$ (in [1] it is denoted $m(S, \bar{S})$).

- Given a flow network H, a *cut* is a set $F \subset V(H)$ such that $s \in F$ and $t \notin F$. The *border* of the cut F is the set of edges $\delta_H(F)$. The capacity of the cut F, denoted $c(F)$, is the sum of the capacities of the edges belonging to the border $\delta(F)$. If the capacities are all equal to 1, it is simply the number $|\delta(F)|$ of edges between F and $\bar{F}$.

THEOREM 5 (Ford-Fulkerson). — *In a flow network N the maximum value of a flow from s to t is equal to the minimum capacity of a cut.*

---

[1] To each edge is associated a positive integer called *capacity of the edge*.

DEFINITION 6. — *Given a multi-broadcast problem (O, D) in a graph G, the flow network associated $N(O, D)$ is defined as follows :*

- *the vertex set of $N(O, D)$ is $V(G) \cup \{s\} \cup \{t\}$;*
- *to each edge of G we associate an edge of capacity 1 in $N(O, D)$;*
- *for each vertex $o \in O$ we add the edge $[s, o]$ with capacity $+\infty$;*
- *for each vertex $d \in D$ the edge $[d, t]$ with capacity 1.*

LEMMA 7. — *The multi-broadcast (O, D) is possible if and only if the maximal flow in $N(O, D)$ is $|D|$.*

Proof. — The value of the flow is at most $|D|$ (which corresponds to the cut $V(G) \setminus \{t\}$), and this value is clearly achieved if and only if there exists $|D|$ edge-disjoint paths in G starting in O and ending at each vertex of D. ☐

According to Lemma 7, given a sequence $\mathcal{S}$, deciding if $\mathcal{S}$ is valid and if so finding the paths that allow to perform the multi-broadcast in T communication rounds, takes at most T times the maximum flow complexity $(O(ne \log(n^2/e))$ in a graph with n vertices and e edges with Goldberg and Tarjan algorithm [12]).

However, that does not tell us how to construct the sequence $\mathcal{S}$; one way would be to use a non-constructive method and to consider random subsets of $V(G)$. When G has nice symmetry properties there exists a better constructive approach. As example, when G is a toroidal mesh one can use a sequence $\mathcal{S}$ made up of linear codes over a vector space Z_k^n (for broadcasting in the k-dimensional torus [7], for gossiping in respectively the 2 and 3-dimensional tori [8], [5]). Here, in the case of the cube, the sequence will be built from linear binary codes. The high symmetry and the algebraic properties of the sequence will enable us to reduce condition 4) so that it becomes easy to check.

3. Symmetric flow graph, symmetric cut.

In this part, our aim is to show that in a flow network having symmetries there exists a symmetric minimum cut. In a graph G an *automorphism* is a one to one mapping $V(G) \xrightarrow{\phi} V(G)$ which preserves the edges (*i.e.* $[x, y] \in E(G)$ if and only if $[\phi(x), \phi(y)] \in E(G)$, see [2]). In a flow network $N = (H, s, t)$ a *symmetry* is simply an automorphism ϕ of the capacited graph H (*i.e.* ϕ preserves also the capacities) fixing both s and t.

Note that in [27] Watkins studied connectivity properties of a transitive graph. In particular, he proved that, when a graph is both vertex and edge transitive then it is superconnected (*i.e.* its edge-connectivity equals its degree). Since that many related results have been derived (see [11], [21], [15]). This question is very related to our multi-broadcast problem (O, D), which is a connectivity problem between the two sets O and D.

Atoms and fragments.

The following definitions and lemmas are exact counterparts of the ones introduced by Watkins [27], Mader [21] and Hamidoune [15] to study connectivity, the only difference is that we consider flow networks that is a graph labeled with a source and a sink. For a comprehensive treatment on connectivity we refer to the work of Hamidoune [14].

Let

$$c_{\min} = \min\{c(F), \ F \text{ a cut of } N\}.$$

DEFINITION 8 (Fragments). — *In a flow network a cut F such that $c(F) = c_{\min}$ is called a fragment.*

LEMMA 9. — *Let F_1 and F_2 be two fragments of a flow network. Then, $F_1 \cap F_2$ and $F_1 \cup F_2$ are also two fragments.*

Proof. — We have

$$c(F_1 \cup F_2) \leq c(F_1) + c(F_2) - c(F_1 \cap F_2),$$

and then

$$c(F_1 \cup F_2) + c(F_1 \cap F_2) \leq 2c_{\min}.$$

Furthermore, $F_1 \cup F_2$ and $F_1 \cap F_2$ both contain s and not t, hence they are cuts. It follows that $c(F_1 \cup F_2) \geq c_{\min}$ and $c(F_1 \cap F_2) \geq c_{\min}$, which implies $c(F_1 \cup F_2) = c(F_1 \cap F_2) = c_{\min}$.   □

DEFINITION 10 (Atoms). — *An atom is a fragment F of minimum size (i.e. $|F|$ is minimum).*

LEMMA 11. — *In a flow network there exists a unique atom.*

Proof. — Let A be the intersection of all the fragments; from Lemma 9, A is a fragment. Moreover it is contained in any fragment so it is both minimal and minimum. □

PROPOSITION 12. — *Let ϕ be a symmetry of a flow network N, then the atom A of N is invariant by ϕ (i.e. $\phi(A) = A$).*

Proof. — By definition, a symmetry maps an atom on another atom. As the atom of N is unique, $\phi(A) = A$. □

The above lemma shows that in order to find the minimal cut in a flow network N it is enough to consider the capacity of a symmetric cut. Indeed there exists a minimum cut which is symmetric. In the next section we will show that, instead of considering the flow problem on N, it is possible to consider a reduced flow network, obtained by quotienting by its symmetries.

4. Sequences for the cube.

In the case of the hypercube the sequence

$$\mathcal{S} = C_0 = \{0\} \subset C_1 \subset \cdots \subset C_T = V(H_n)$$

will be made of linear codes of length n over $\mathbb{Z}_2$ (*i.e.* linear spaces of the vector space $\mathbb{Z}_2^n$ [(2)]). We will note $\text{Span}\{\mathcal{F}\}$ the linear space generated by a family $\mathcal{F}$ of vectors. Given two independent linear spaces A and B, $A \oplus B$ denotes their sum (that is combinations of vectors in A and in B). We will denote

$$e_i = \underbrace{00\ldots0}_{i-1}10\ldots0$$

the i-th vector of the natural basis of $\mathbb{Z}_2^n$.

As our codes are nested, we have

$$C_{t+1} = C_t \oplus V_{t+1} \quad \text{and} \quad |C_{t+1}| = |V_{t+1}| \cdot |C_t|\,;$$

moreover, as the number of informed nodes is multiplied by at most $n + 1$ during a round, we must have

$$\dim(V_{t+1}) \leq \lfloor \log_2(n + 1) \rfloor.$$

[(2)] A linear space of dimension k contains 2^k elements which are the linear combinations of k independent elements.

Let $\kappa = \lfloor \log_2(n+1) \rfloor$. We will always choose V_{t+1} with maximum dimension; that is $\dim(V_{t+1}) = \kappa$ for any $t < T-1$ and $\dim(V_T) = n - \lfloor n/\kappa \rfloor \kappa \leq \kappa$.

As an example for $n = 9$, let

$$f_1 = 110000000, \; f_2 = 011000000, \; \ldots, f_8 = 000000011.$$

We will use the codes

$$C_1 = \{0\} \oplus V_1 = V_1, \qquad V_1 = \text{Span}\{f_1, f_4, f_7\},$$
$$C_2 = C_1 \oplus V_2, \qquad\qquad V_2 = \text{Span}\{f_2, f_5, f_8\}.$$

So, $C_2 = V_1 \oplus V_2$ contains the 2^6 combinations of $f_1, f_2, f_4, f_5, f_7, f_8$. Note that this is the set of vectors abc where a, b, c are three words of length 3 having an even number of 1. The last code is

$$C_3 = C_2 \oplus V_3, \qquad\qquad V_3 = \text{Span}\{e_1, e_4, e_7\} \, ;$$

note that $C_3 = V_1 \oplus V_2 \oplus V_3 = \mathbb{Z}_2^9$.

We consider the multi-broadcast $(C_t, C_{t+1} \setminus C_t)$ in $H(n)$, and the associated flow network $N(C_t, C_{t+1} \setminus C_t)$. We want to derive a condition ensuring that there exists a flow with value

$$|C_{t+1} \setminus C_t| = (|V_{t+1}| - 1)|C_t|$$

in $N(C_t, C_{t+1} \setminus C_t)$.

The hypercube $H(n)$ is a *Cayley graph* and this structure is the key one.

DEFINITION 13. — *Given an Abelian group $\mathcal{G}$ and a multi-set $S \subset \mathcal{G}$ the Cayley multi-graph on $\mathcal{G}$ with generators S, denoted $\text{Cay}(\mathcal{G}, S)$, is defined by :*

- *the vertices are elements of $\mathcal{G}$;*

- *the neighborhood of x is the set $x + S$.*

In order to get an undirected graph we must have $-S = S$. Note that if we choose for S a multi-set (repeating some generators) we get a Cayley multi-graph. Note that in $\text{Cay}(\mathcal{G}, S)$ the mapping $\phi_y : x \mapsto x + y$ is an automorphism. For simplicity we restrict ourselves to the Abelian case,

but similar notions and results both exist for non-Abelian groups. For an overview of applications of Cayley graphs to interconnection networks we refer to the chapter of Heydemann in [13].

The hypercube $H(n)$ is indeed $\mathrm{Cay}(\mathbb{Z}_2^n, \{e_1, e_2, \ldots, e_n\})$, where $\{e_1, e_2, \ldots, e_n\}$ is the basis of $\mathbb{Z}_2^n$ previously defined. The application $\phi_c : x \mapsto x + c$, $c \in C_t$ is an automorphism of the cube. Moreover it lets both C_t and $C_{t+1} \setminus C_t$ invariant. Hence, ϕ_c, $c \in C_t$ is a symmetry of $N(C_t, C_{t+1} \setminus C_t)$. Now, the set $\{\phi_c \mid c \in C_t\}$ is an Abelian subgroup of the group of automorphisms of $H(n)$. According to Proposition 12, there exists a minimal cut of $N(C_t, C_{t+1} \setminus C_t)$ which is invariant by any mapping ϕ_c, $c \in C_t$, that is there exists a minimal cut invariant modulo C_t. This property is indeed equivalent to a connectivity property of the quotient multi-graph $H(n)/C_t$ defined below.

DEFINITION 14. — *The quotient multi-graph $H(n)/C_t$ is defined as follows :*

• *The vertex set is $\mathbb{Z}_2^n/C_t \sim \mathbb{Z}_2^{n-\dim(C_t)}$ and each vertex corresponds to a coset of C_t ;*

• *$\tilde{x} = x + C_t$ will denote the coset of x. ;*

• *for all $i \in \{1, 2, \ldots, n\}$, we add one edge from $\tilde{x}$ to $\tilde{y}$ along "dimension i" if $e_i \in \tilde{y} - \tilde{x}$.*

The quotient $H(n)/C_t$ is a Cayley multi-graph on $\mathbb{Z}_2^{n-\dim(C_t)}$ with generators $\tilde{e}_i$, $i = \{1, 2, \ldots, n\}$. Indeed there exist a edges between $\tilde{x}$ and $\tilde{y}$ if there are $a|C_t|$ edges in $H(n)$ between the two sets $x + C_t$ and $y + C_t$. Note that we can have $\tilde{e}_i = \tilde{e}_j$ for some $i \neq j$; so $H(n)/C_t$ is an Abelian Cayley multi-graph. Note also that if $e_i \in C_t$ then $H(n)/C_t$ contains a loop.

In our example, the quotient $H(9)/C_1$ is a multi-hypercube $H(6)$ with some edges having multiplicity 2: $\tilde{e}_1 = \tilde{e}_2$ as $e_1 + e_2 = f_1$, similarly $\tilde{e}_4 = \tilde{e}_5$ and $\tilde{e}_7 = \tilde{e}_8$.

The multiset of generators is $\{\tilde{e}_1, \tilde{e}_1, \tilde{e}_3, \tilde{e}_4, \tilde{e}_4, \tilde{e}_6, \tilde{e}_7, \tilde{e}_7, \tilde{e}_9\}$.

The quotient $H(9)/C_2$ is simply the hypercube $H(3)$ where each edge is repeated three times, as now,

$$\{\tilde{e}_1 = \tilde{e}_2 = \tilde{e}_3 \ ; \ \tilde{e}_4 = \tilde{e}_5 = \tilde{e}_6 \ ; \ \tilde{e}_7 = \tilde{e}_8 = \tilde{e}_9\}.$$

Note that in general $H(n)/C_t$ is not necessary a multi-hypercube. As example, if $C_t = (00 \ldots 0, 11 \ldots 1)$, we obtain for $H(n)/C_t$ the "Halved cube" by identifying antipodal vertices; this graph is not a multi-hypercube (see also [4] for other examples of quotient).

Both $H(9)/C_1$ and $H(9)/C_2$ have edge-connectivity 9.

LEMMA 15. — *If $H(n)/C_t$ has edge-connectivity at least $|V_{t+1}| - 1$, then C_t can inform $C_{t+1} \setminus C_t$.*

Proof. — By Lemma 7 and the Ford-Fulkerson Theorem (Theorem 5), we simply need to check that the capacity of the border of a minimum cut of $N(C_t, C_{t+1} \setminus C_t)$ is at least $|C_{t+1}| - |C_t| = (|V_{t+1} - 1)|C_t|$, whenever the edge-connectivity of $H(n)/C_t$ is at least $|V_{t+1}| - 1$.

According to Proposition 12, we can consider the unique symmetric atom A of $N(C_t, C_{t+1} \setminus C_t)$, having by definition minimum border capacity (*i.e.* the capacity of the border of A equals the maximum value of flow in $N(C_t, C_{t+1} \setminus C_t)$). Note that $A = U \cup \{s\}$, where U is a subset of $V(H(n))$ such that $\phi_c(U) = U$, for all $c \in C_t$. This means that $\phi_c(U) = U + c = U$ for all $c \in C_t$. Hencefore, U is invariant by translation in C_t, and is a union of cosets: $U = \widetilde{U} + C_t$.

Note also that U must contain all the vertices in C_t, otherwise, as edges $[s, c]$ for $c \in C_t$ have infinite capacity, the border of A would have infinite capacity. If U contains a vertex $c' \in C_{t+1} \setminus C_t$ then we find the edge $[c', t]$ in the border of A. We also find in the border of A the border $\delta_{H(n)}(U)$ of U in the hypercube. On the total

(1) $c_{N(C_t, C_{t+1} \setminus C_t)}(A) = |U \cap (C_{t+1} \setminus C_t)| + |\delta_{H(n)}(U)|.$

As U and C_{t+1} both contain C_t, we have

$$|U \cap (C_{t+1} \setminus C_t)| = |U \cap C_{t+1}| - |C_t|.$$

Now let us consider $H(n)/C_t$. We have

$$|\delta_{H(n)}(U)| = |\delta_{H(n)/C_t}(\widetilde{U})| \cdot |C_t|.$$

Hence Equation (1) can be written:

$$c_{N(C_t, C_{t+1} \setminus C_t)}(A) = |C_t| \cdot \left(|\widetilde{U} \cap \widetilde{C}_{t+1}| - 1 + |\delta_{H(n)/C_t}(\widetilde{U})| \right).$$

As, U contains C_t, $\widetilde{U}$ contains the vertex 0 in $H(n)/C_t$, we have $|\widetilde{U} \cap \widetilde{C}_{t+1}| - 1 \geq 0$, and it follows that:

$$c_{N(C_t, C_{t+1})}(A) \geq |C_t| \cdot |\delta_{H(n)/C_t}(\widetilde{U})|.$$

Now, the condition on the connectivity of $H(n)/C_t$ implies that $|\delta_{H(n)/C_t}(\widetilde{U})|) \geq |V_{t+1} - 1|$ implying the lemma.   □

The lemma above allows us to reduce the condition C_t can inform C_{t+1} to a simple condition on the connectivity of the quotient graph $H(n)/C_t$.

We will now construct sequences of linear codes over $\mathbb{Z}_2$ such that $H(n)/C_t$ has high enough edge-connectivity to ensure the condition of Lemma 15.

5. Constructions.

5.1. Case of cyclic codes.

For some values of n it is easy to construct an optimal scheme by using a sequence made of cyclic linear codes of appropriate dimension. *A cyclic code is simply a subspace of $\mathbb{Z}_2^n$ which is invariant by left-shift.* Note that our first proof, in the case $n = 2^p - 1$ was made using more coding theory and special properties of BCH codes (see [19], [20]). With the technique used here the result is more straightforward.

LEMMA 16. — *If C_t is a cyclic code and $|C_{t+1}| \leq (n+1)|C_t|$ then C_t can inform C_{t+1}.*

Proof. — Consider the quotient graph $H(n)/C_t$, it is a Cayley multigraph. Consider the left shift, σ defined as the linear mapping of $\mathbb{Z}_2^n$ such that

$$\sigma(e_i) = e_{i+1 \bmod n}.$$

As C_t is cyclic, $\sigma(C_t) = C_t$, and the mapping σ acts also on the vertices of $H(n)/C_t$. Now, if $\tilde{x}$ is adjacent to $\tilde{y}$ in $H(n)/C_t$ along dimension i (that is $\tilde{x} = \tilde{y} + e_i$) then $\sigma(\tilde{x})$ is adjacent to $\sigma(\tilde{y})$ along dimension $i + 1 \bmod n$. So σ is an automorphism of $H(n)/C_t$. It follows that $H(n)/C_t$ is edge transitive. According to the result of Watkins [27], $H(n)/C_t$ is superconnected, or its edge-connectivity equals its degree. Note that if the degree of $H(n)/C_t$ is less than n, then, for some i, $e_i \in C_t$; since the code is cyclic it implies $C_t = \mathbb{Z}_2^n$. So $H(n)/C_t$ has degree n and edge-connectivity n. $\square$

Cyclic Codes.

To complete we need to find some good sequence of codes. For this we briefly recall some basics of the theory of error correcting codes. The reader will find more information in [20] or [26]. Cyclic codes are very convenient

since they can be viewed as *principal ideal* of the ring $R = F_2[x]/(x^n - 1)$. In other words, a cyclic code C consists of all multiples of a polynomial $g(x)$, called generator polynomial. Note that this polynomial must be a factor of $x^n - 1$. The code C is then totally defined by $g(x)$ and we write: $C = \langle g(x) \rangle$. One can find all the possible cyclic codes by considering the factorization of $x^n - 1$ over F_2.

Decomposition of $x^n - 1$ over F_2. — We consider only binary codes of length $n = 2^p - 1$. Let α be a primitive n-th root of unity over a suitable extension of $\mathbb{Z}_2$. We have:

$$x^n - 1 = \prod_{i=0}^{n-1} (x - \alpha^i).$$

Furthermore, it is well known that over $\mathbb{Z}_2$

$$(2) \qquad x^n - 1 = (x - 1)f_1 \cdots f_{T-1} = \prod_t f_t,$$

where f_t denotes the minimal polynomials corresponding to the cyclotomic classes. Then we consider the cyclic codes generated by f_t.

• The Hamming code can be generated by anyone of the primitive polynomials f_t. Its parameters are

$$[N_H = 2^p - 1, k_H = 2^p - 1 - p, d_H = 3],$$

where N_H denotes the length, k_H the dimension (seen as a linear space) and d_H the minimum Hamming distance of the code.

• The simplex code is defined as the (algebraic) dual of the Hamming code. Then its parameters are $[n, p, 2^{p-1}]$.

• A BCH code of length $n = 2^p - 1$ and designed distance d' is a cyclic code generated by the product (without repetition of factors) of the minimum polynomials of $\alpha^r, \alpha^{r+1}, \ldots, \alpha^{r+d'-2}$, where r is a non negative integer.

Note that all these codes are defined up to equivalence.

Nested BCH codes. — Let $n = 2^p - 1$, that is $\kappa = p$. If p is a prime number, one can use a family of BCH codes that have the required dimensions (*i.e.* a sequence of nested codes with dimensions

$0, \kappa, 2\kappa, \ldots, \lfloor n/\kappa \rfloor \kappa, n)$. In this case, the partition of $H(n)$ can be done as follows: We set

$$C_0 = \{0^n\},$$
$$C_1 = \langle (x-1)f_1 \cdots f_{T-2} \rangle,$$
$$C_2 = \langle (x-1)f_1 \cdots f_{T-3} \rangle,$$
$$\vdots$$
$$C_{T-2} = \langle f_1 f_2 \rangle,$$
$$C_{T-1} = \langle f_1 \rangle,$$
$$C_T = F_2^n.$$

The polynomials f_i in (2) are sorted in a way so that we obtain the Hamming code for C_{T-1}, the simplex code for C_1 and BCH codes for the other codes. We have $C_0 = \{0\}$, $|C_{t+1}| = (n+1)|C_t|$ and $C_T = \mathbb{Z}_2^n$. So, this sequence of codes gives an optimal scheme for $n = 2^p - 1$ (see in [19] for more details).

PROPOSITION 17. — For p a prime,

$$b_{F_*}\left(H(2^p - 1)\right) = \left\lceil \frac{2^p - 1}{p} \right\rceil.$$

5.2. A connectivity condition on the quotient graph.

In many other cases one can prove that the connectivity of $H(n)/C_t$ is high enough. As the graph $H(n)/C_t$ is a Cayley multi-graph the following result makes the determination of its edge-connectivity relatively easy.

PROPOSITION 18 (Hamidoune). — In a Cayley graph $\mathrm{Cay}(\mathcal{G}, \mathcal{S})$ the atom containing $\{0\}$ is a subgroup of $\mathcal{G}$.

Proof. — Let A be the atom containing $\{0\}$, and note that $\phi_a : x \mapsto x + a$ is an automorphism of G. Hence $A + a$ is an atom of G, and it contains $a + 0 = a$. So $a + A$ and A are two atoms with a non empty intersection, they must be equal. So $a + A = A$ for all $a \in A$. In the same way, ϕ_{-a} maps A on A so $A - a = A$. It follows that A is a subgroup of $\mathcal{G}$. $\square$

DEFINITION 19. — Let W be a subgroup of $\mathbb{Z}_2^n/C_t$, and let denote $\widetilde{e_i}$ the C_t coset associated to e_i (indeed the set $e_i + C_t$), then the number of $\widetilde{e_i} \notin W$ is denoted $\ell(W)$.

PROPOSITION 20. — *If $|W|\ell(W) \geq \lambda$ for every subgroup W of $\mathbb{Z}_2^n/C_t$, then $H(n)/C_t$ has edge-connectivity at least λ.*

Proof. — Consider a subgroup W of $\mathbb{Z}_2^n/C_t$. If $\widetilde{e}_i \notin W$, then for any $w \in W$, $\widetilde{e}_i + w \notin W$, so each edge $[w, w + \widetilde{e}_i]$ is in the border of W. If $\widetilde{e}_i \in W$ then all the edges $[w, w + \widetilde{e}_i]$, $w \in W$ are inside W. Consequently the border of W contains $|W|\ell(W)$ edges. Note that if some e_i belongs to C_t, then $\widetilde{e}_i = 0$ and any subgroup W contains $\widetilde{e}_i$. $\square$

5.3. An ad-hoc construction.

We propose here a sequence of nested linear codes of length $\lceil n/\kappa \rceil$; such a sequence is optimal among sequences made of linear codes.

PROPOSITION 21. — *One has*

$$b_{F_*}\left(H(n)\right) \leq \left\lceil \frac{n}{\kappa} \right\rceil, \quad \text{where} \quad \kappa = \left\lfloor \log_2(n+1) \right\rfloor.$$

Note that for $n = 2^k - 1$, $\lfloor \log_2(n+1) \rfloor$ is the integer k and our scheme is optimal (*i.e.* $b_{F_*}(H(2^k - 1)) = \lceil (2^k - 1)/k \rceil$) this generalizes Proposition 17 valid for k prime. The proposition will follow from the construction of valid sequences of length $\lceil n/\kappa \rceil$.

The sequence.

Let $n = (p+1)\kappa - x, 0 \leq x < \kappa$, so that $p + 1 = \lceil n/\kappa \rceil$; we have $n = (p+1)(\kappa - x) + px$.

As in our example of Section 4, let

- $f_1 = e_1 + e_2$, $f_2 = e_2 + e_3, \ldots$, $f_{n-1} = e_{n-1} + e_n$;

- $f_1' = f_{px+1}$, $f_2' = f_{px+2}, \ldots$, $f_{(\kappa-x)(p+1)-1}' = f_{n-1}$;

- $e_i' = e_{i+px}$.

Our sequence of codes is defined as follows (see the example after

for $n = 17$):

$$
\begin{cases}
V_1 = \mathrm{Span}\{f_1, f_{1+p}, f_{1+2p}, \ldots, f_{1+(x-1)p}\} \\
\qquad \oplus \mathrm{Span}\{f'_1, f'_{1+(p+1)}, \ldots, f'_{1+(\kappa-x-1)(p+1)}\}, \\
V_2 = \mathrm{Span}\{f_2, f_{2+p}, f_{2+2p}, \ldots, f_{2+(x-1)p}\} \\
\qquad \oplus \mathrm{Span}\{f'_2, f'_{2+(p+1)}, \ldots, f'_{2+(\kappa-x-1)(p+1)}\}, \\
\qquad \cdots \\
V_i = \mathrm{Span}\{f_i, f_{i+p}, f_{i+2p}, \ldots, f_{i+(x-1)p}\} \\
\qquad \oplus \mathrm{Span}\{f'_i, f'_{i+(p+1)}, \ldots, f'_{i+(\kappa-x-1)(p+1)}\} \, ;
\end{cases}
$$

rounds $1, 2, \ldots, p-1$

round $p \quad \begin{cases} V_p = \mathrm{Span}\{f'_p, f'_{p+(p+1)}, \ldots, f'_{p+(\kappa-x-1)(p+1)} = f_{n-1}\} \end{cases}$

round $p+1 \quad \begin{cases} V_{p+1} = \mathrm{Span}\{e_1, e_{1+p}, \ldots, e_{1+(x-1)p}\}, \\ \qquad \oplus \mathrm{Span}\{e'_1, e'_{1+(p+1)}, \ldots, e'_{1+(\kappa-x-1)(p+1)}\}. \end{cases}$

To prove that the sequence is valid, we write $\mathbb{Z}_2^n$ as a sum of κ independent subspaces, x of dimension p and $\kappa - x$ of dimension $p + 1$. We will call these spaces *blocks*. Let:

$$
\begin{aligned}
A_1 &= \mathrm{Span}\{e_1, e_2, \ldots, e_p\}, \\
A_2 &= \mathrm{Span}\{e_{1+p}, e_{2+p}, \ldots, e_{2p}\}, \\
&\quad \cdots \\
A_x &= \mathrm{Span}\{e_{1+(x-1)p}, e_{2+(x-1)p}, \ldots, e_{xp}\}, \\
A'_1 &= \mathrm{Span}\{e'_1, e'_2, \ldots, e'_{p+1}\}, \\
A'_2 &= \mathrm{Span}\{e'_{1+(p+1)}, e'_{2+(p+1)}, \ldots, e'_{2(p+1)}\}, \\
&\quad \cdots \\
A'_{\kappa-x} &= \mathrm{Span}\{e'_{1+(\kappa-x-1)(p+1)} e'_{2+(\kappa-x-1)(p+1)}, \ldots, e'_{(k-x-1)(p+1)+(p+1)}\}.
\end{aligned}
$$

From definition

$$
\mathbb{Z}_2^n = A_1 \oplus A_2 \oplus \cdots \oplus A_x \oplus A'_1 \oplus A'_2 \oplus \cdots \oplus A'_{\kappa-x}.
$$

Example.

For $n = 9$, we find the example of section 4, in this case $\kappa = 3, p = 2, x = 0$ and $f'_i = f_i, e'_i = e_i$; we find

$$
\begin{aligned}
V_1 &= \mathrm{Span}\{f'_1, f'_4, f'_7\} = \mathrm{Span}\{f_1, f_4, f_7\}, \\
V_2 &= \mathrm{Span}\{f'_2, f'_5, f'_8\} = \mathrm{Span}\{f_2, f_5, f_8\}, \\
V_3 &= \mathrm{Span}\{e'_1, e'_4, e'_7\} = \mathrm{Span}\{e_1, e_4, e_7\}.
\end{aligned}
$$

We also give the example for $H(17)$:

- $n = 17$, $\kappa = 4$, $p = \lceil n/\kappa \rceil - 1 = 4$, so $x = 3$ and $n = 5 \cdot 1 + 4 \cdot 3$;

- the vector space $\mathbb{Z}_2^{17}$ is the direct sum of the following subspaces:

  ▷ three spaces of dimension 4:

$$A_1 = \text{Span}\{e_1, e_2, e_3, e_4\},$$
$$A_2 = \text{Span}\{e_5, e_6, e_7, e_8\},$$
$$A_3 = \text{Span}\{e_9, e_{10}, e_{11}, e_{12}\} \, ;$$

  ▷ one space of dimension 5: $A_1' = \text{Span}\{e_{13}, e_{14}, e_{15}, e_{16}, e_{17}\}$;

- sequence of codes:

Codes	added	A_1	A_2	A_3	A_1'
$C_1 = C_0 \oplus V_1$	V_1	f_1	f_5	f_9	f_{13}
$C_2 = C_1 \oplus V_2$	V_2	f_2	f_6	f_{10}	f_{14}
$C_3 = C_2 \oplus V_3$	V_3	f_3	f_7	f_{11}	f_{15}
$C_4 = C_3 \oplus V_4$	V_4				f_{16}
$C_5 = C_4 \oplus V_5$	V_5	e_1	e_5	e_9	e_{13}

Connectivity of the quotient graphs.

For $t \leq p$, the codes C_t have the following property: *if a subspace $W \subset \mathbb{Z}_2^n/C_t$ contains some set of vectors $\widetilde{e}_i$ located in d different blocks then $\dim(W) \geq d$.*

LEMMA 22. — *The edge-connectivity of $H(n)/C_t$ is at least $2^\kappa - 1$.*

Proof. — According to Proposition 20, we only need to check that for a linear space W of $\mathbb{Z}_2^n/C_t$, $\ell(W)|W| \geq 2^\kappa - 1$. Let $d(W)$ denote the dimension of the linear subspace W, and note that if $d(W) \geq \kappa$ then $\ell(W)|W| \geq 2^\kappa$ and the condition is ensured. So we can restrict ourselves to the case $d(W) \leq \kappa$.

Rounds $1, 2, \ldots p$. — According to the property of the sequence, if $d(W)$ is less than d, W can contain only the e_i of d distinct blocks. Hence, if $d(W) \leq \kappa$, the worst case space W is clearly obtained by picking vectors of the basis in $d(W)$ different blocks as a generating set for W. So doing we add at most p vectors (the ones of the bloc) of the basis in W when we

pick a new block. Hence, we find at most $pd(W)$ vectors of the basis in W. Consider the function

$$\ell(W)|W| = 2^{d(W)}\left(n - pd(W)\right).$$

It is easy to check that the minimum of such a function is attained either for $d(W) = \kappa$, then $d(W) = \kappa$ and $|W| \geq 2^{\kappa}$, or for $d(W) = 0, |W| = 1$ and then $\ell(W)|W| = n \geq 2^{\kappa} - 1$ (by definition $\kappa = \lfloor \log_2(n+1)\rfloor$).

In both cases $\ell(W)|W| \geq 2^{\kappa}$. We conclude that the edge-connectivity of $H(n)/C_t$ is at least 2^{κ} for $t \leq p$.

Round $p+1$. — Note that, H_n/C_p is a multi-hypercube of dimension κ with x dimensions with multiplicity p and the $\kappa - x$ others with multiplicity $p + 1$. Due to symmetry, the minimal cut contain all the dimensions of capacity p (or $p + 1$). So the minimum cut has value $2^{\kappa-x}xp$ (or $2^x(p+1)(\kappa - x)$). This value must be larger than $|V_{p+1}| - 1 = 2^{\kappa-x} - 1$. So, if there exists some contradiction, it is in one of the two extremal cases: $x = 0$ (in this case dimensions are all equal and the only cut is $\{0\}$ and has capacity $(p + 1)\kappa$) or $x = \kappa - 1$. In both cases the cuts are large enough. □

Distance 3 codes.

One can easily use other kind of schemes. As an example, if C_t has minimal distance at least 3 one can show, using results of Sections 3 and 4 that the connectivity of $H(n)/C_t$ is n. It follows that it is possible to inform any distance 3 codes in $(n - \Theta(\log(n)))/\kappa$ communication rounds. As distance 3 codes with dimension $n - \Theta(\log^2(n))$ do exist, a scheme first informs such a code, at that point the broadcast is almost completed, and one has then to add a few rounds $(\Theta(\log(n)^2/\log(n)) \sim \log(n))$ to inform the whole cube.

6. Conclusion.

In this paper we have derived some efficient, and sometimes optimal scheme to broadcast information in the cube in wormhole like models (see Table 1). It turns out that high symmetry of our solution made the proof of our scheme possible. Our scheme does not use the e-cube routing but we can mention that it defines implicitly a routing function which is not too complicated. Our scheme uses a non immediate sequence of nested codes; let us point out that in their algorithm Ho and Kao used such a sequence,

but the simplest possible one: the set of informed nodes at a given round was a sub-cube (that is a very simple linear code). Hence, the vertices informed were packed in the same area of the cube, and the algorithm was not very efficient. However the analysis was simple and the scheme was using e-routing. Our scheme is more efficient as vertices informed at round t are better spread in the cube, but it is also more complex. As an example, for $n = 31$ our scheme uses 7 rounds and their scheme 10.

n	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	...	31
Low. Bound	2	2	2	2	3	3	3	3	3	4	4	4	4	4	4	...	5
MT [22]	2	2	2	3	3	4	4	5	5	6	6	7	7	8	8	...	16
HK [17]	2	2	2	3	3	3	4	4	4	5	5	5	6	6	6	...	10
Our	2	2	2	2	3	3	3	3	4	4	4	5	5	4	4	...	7

Table 1. Comparison of the different algorithms

Note that our scheme is optimal among schemes such that at a given round the set of informed nodes is a linear code. Hence to improve it one would need to have informed at round t a set having a structure more complex than a linear one, the analysis would certainly be complicated. The example given for $H(5)$ in Figure 1 uses non-linear set of vertices and two rounds, and any scheme using linear sets takes at least three rounds.

It would be interesting to use non constructive approach (using random subsets of the cube) to improve our bound, but this would certainly be only an existence result.

At least we believe that our result demonstrates once again that symmetries can be used in order to simplify graph problems (for an overview of symmetry technique: see [13]), the key point being that symmetric flow-problems admit symmetric minimum cuts.

BIBLIOGRAPHY

[1]   C. BERGE, Graphes et hypergraphes, Dunod, Paris, 1970.

[2]   N. BIGGS, Algebraic Graph Theory, Cambridge University Press, 1974.

[3]   J.A. BONDY, U.S.R. MURTY, Graph theory with applications, McMillan Press, 1976.

[4]   A.E. BROUWER, A.M. COHEN, A. NEUMAIER, Distance regular graphs, Springer Verlag, 1989.

[5]   C. CALVIN, S. PÉRENNES, D. TRYSTRAM, Gossiping in torus with wormhole-like routing, in Proceedings of the 7-th IEEE Symposium on Parallel and Distributed Processing, San-Antonio, 1995.

[6]   O. DELMAS, Communications par commutation de circuits dans les réseaux d'interconnexion, Thèse, Université de Nice, Sophia Antipolis, 1997.

[7]   O. DELMAS, S. PÉRENNES, Diffusion par commutation de circuits dans les tores de dimension k/Circuit switched broadcasting in the k-th dimentional torus networks, Technique et science informatique, RAIRO, AFCET, 16 (5) (1997), 563–581.

[8]   O. DELMAS, S. PÉRENNES, Circuit-Switched Gossiping in 3–Dimensional Torus Networks, in Proceedings of the Euro-Par'96 Parallel Processing, Second International EURO-PAR Conference, Lyon, Lecture Notes in Computer Science, Springer Verlag, vol. 1123, 1996, 370–373.

[9]   E. FLEURY, Communication, routage et architectures des machines à mémoires distribuées – autour du routage wormhole, Thèse, École normale supérieure de Lyon, 1996.

[10]   P. FRAIGNIAUD, E. LAZARD, Methods and problems of communication in usual networks, Discrete Applied Math., 53 (1994), 79–133.

[11]   C.D. GODSIL, Connectivity of minimal Cayley graphs, Arch. Math., 37 (1981), 437–476.

[12]   A.V. GOLDBERG, R.E. TARJAN, A new approach to the maximum flow problem, J. ACM, 35 (1988), 921–940.

[13]   G. HAHN, G. SABIDUSSI, Graph symmetry, Nato ASI Series, vol. 497, Kluwer Academic Publishers, 1996.

[14]   Y.O. HAMIDOUNE, Quelques problèmes de connexité dans les graphes orientés, Thèse, Université Pierre et Marie Curie, Paris VII, 1978.

[15]   Y.O. HAMIDOUNE, On the connectivity of Cayley digraphs, Eur. J. Comb., 5 (1984), 309–312.

[16]   S.M. HEDETNIEMI, S.T. HEDETNIEMI, A.L. LIESTMAN, A survey of gossiping and broadcasting in communication networks, Networks, 18 (1988), 319–349.

[17]   C.-T. HO, M.-Y. KAO, Optimal broadcast in all-port wormhole-routed hypercubes, IEEE Transactions on Parallel and Distributed Systems, 6 (2) (1995), 200–204.

[18]   J. HROMKOVIC, R. KLASING, B. MONIEN, R. PEINE, Combinatorial Network Theory, Chap. 'Dissemination of information in interconnecting networks (Broadcasting and Gossiping)', 125–212, Kluwer Academic Publishers, 1995.

[19]   T. KODATE, Communications structurées dans les réseaux d'interconnection, Thèse, Université de Nice-Sophia Antipolis, 1996.

[20]   F.J. MACWILLIAMS, N.J.A. SLOANE, The Theory of Error-Correcting Codes, North-Holland, 1977.

[21]   W. MADER, K minimal n-fach kantensuzammenhangenden, Math. Annal., 191 (1971), 21–28.

[22]   P.K. MCKINLEY, C. TREFFTZ, Efficient broadcast in all-port wormhole-routed hypercubes, in 'International Conference on Parallel Processing (ICPP'98)', vol. II, St. Charles, IL, USA, 1993.

[23]   J.G. PETERS, M. SYSKA, Circuit switched broadcasting in torus networks, IEEE Transactions on Parallel and Distributed Processing, 7 (3) (1996), 246–255.

[24]   D.F. ROBINSON, D. JUDD, P.K. MCKINLEY, B.H.C. CHENG, Effective collective data distribution in all-port wormhole-routed hypercubes, in 'Proceedings Super-computing'93', 1993, 792–780.

[25]   J. DE RUMEUR, Communication dans les réseaux de processeurs, Masson, Paris, 1994.

[26]   J.H. VAN LINT, Introduction to Coding Theory, Springer Verlag, Berlin, 1982.

[27]   M.E. WATKINS, Connectibity of transitive graphs, J. Comb. Theory, 8 (1970), 23–29.

J.-C. BERMOND,
SLOOP, joint project CNRS/INRIA/UNSA
2004 route des Lucioles
BP 93
06 902 Sophia Antipolis (France).
Jean–Claude.Bermond@inria.fr

&

A. BONNECAZE,
Université de Toulon-Var
GECT
BP 132
83 857 La Garde Cedex (France).
bonnecaz@univ–tln.fr

&

T. KODATE,
Laboratoire I3S
UPRESA 6070 CNRS-UNSA
650 route des Colles
BP 145
06 903 Sophia-Antipolis (France).
kodate@twcu.ac.jp

&

S. PÉRENNES,
SLOOP, joint project CNRS/INRIA/UNSA
2004 route des Lucioles
BP 93
06 902 Sophia Antipolis (France).
Stephane.Perennes@inria.fr

&

P. SOLÉ,
Laboratoire I3S
UPRESA 6070 CNRS-UNSA
650 route des Colles
BP 145
06 903 Sophia-Antipolis (France).
ps@essi.fr