Gossiping in Cayley Graphs by Packets

Jean-Claude Bermond, Takako Kodate and Stephane Perennes

Laboratoire I3S, Université de Nice - Sophia Antipolis
Bât. Essi, 650 Route des Colles, B.P.145
06903 Sophia Antipolis Cedex, France
emails: {bermond, kodate, sp}@unice.fr

Abstract. *Gossiping* (also called *total exchange* or *all-to-all communication*) is the process of information diffusion in which each node of a network holds a packet that must be communicated to all other nodes in the network. We consider here gossiping in the *store-and-forward*, *full-duplex* and *Δ-port* (or *shouting*) model. In such a model, the protocol consists of a sequence of rounds and during each round, each node can send (and receive) messages from all its neighbors. The great majority of the previous works on gossiping problems allows the messages to be freely concatenated and so messages of arbitrary length can be transmitted during a round. Here we restrict the problem to the case where at each round communicating nodes can exchange exactly one packet. We give a lower bound of $\frac{N-1}{\delta}$, where δ is the minimum degree, and show that it is attained in Cayley symmetric digraphs with some additional properties. That implies the existence of an optimal gossiping protocol for classical networks like hypercubes, k-dimensional tori, and star-graphs.

1 Introduction

Gossiping (also called total exchange or all–to–all communication) in distributed systems is the process of distribution of information known to each processor to every other processor of the system. This process of information dissemination is carried out by means of a sequence of message transmissions between adjacent nodes in the network.

The gossiping problem was originally introduced by the community of discrete mathematicians, to which it owes most of its terminology, as a combinatorial problem in graphs. Nonetheless, it was soon realized that, once cast in more realistic models of communication, gossiping is a fundamental primitive in distributed memory multiprocessor system. There are a number of situations in multiprocessor computation, such as global processor synchronization, where gossiping occurs. Moreover, the gossiping problem is implicit in a large class of parallel computation problems, such as linear system solving, Discrete Fourier Transform, and sorting, where both input and output data are required to be distributed across the network [3]. Due to the interesting theoretical questions it poses and its numerous practical applications, gossiping has been widely studied under various communication models. Hedetniemi, Hedetniemi and Liestman [6]

provide a survey of the area. Two more recent surveys paper collecting the latest results are [5, 8]. The reader can also profitably see the book [4].

The great majority of the previous work on gossiping has considered the case in which the packets known to a processor at any given time during the execution of the gossiping protocol can be freely concatenated and the resulting (longer) message can be transmitted in a constant amount of time, that is, it has been assumed that the time required to transmit a message is independent from its length. While this assumption is reasonable for short messages, it is clearly unrealistic when the size of the messages becomes large. Notice that most of the gossiping protocols proposed in the literature require the transmission, in the last rounds of the execution of the protocol, of messages of size $\Theta(N)$, where N is the number of nodes in the network. Here we consider the gossiping problem under the restriction that communicating nodes can exchange exactly one packet at each round.

1.1 Communication Models

We model a communication network by a symmetric digraph $G = (V, E)$ where the node set V represents the set of processors and the arc set E represents the communication links between processors. Here we suppose that if a node x can directly communicate with a node y, then the converse is true. So we have a symmetric digraph where if $(x, y) \in E$ then $(y, x) \in E$. Initially, each node holds a packet that must be transmitted to all the other nodes of the network by a sequence of calls between adjacent processors. The restriction considered in the introduction implies that during each call, communicating nodes can exchange only one packet. We therefore see the gossiping protocol as a sequence of *rounds*. During each round, we suppose that each processor can communicate *with all its neighbors*, more exactly it can send a packet to all its neighbors and receive a packet from all its neighbors. Furthermore we allow each node to send *different packets* to *different neighbors* at each round.

We will denote by $g_{F_*}(1, G)$ the minimum gossip time, that is the minimum number of rounds to complete the gossiping process in the network G subject to the above condition. This model is often referred as F_* or *full-duplex Δ-port* or *shouting model*. Other models also popular restrict a node to communicate only with one of its neighbors during each round (F_1 or *telephone model*) or do not allow both emission and reception (*half-duplex model* denoted H_* or H_1 according one node can send to all its neighbors or only one.) The problem of estimating $g_{F_1}(1, G)$ and more generally $g_{F_1}(p, G)$ where one allows to exchange up to a fixed number p of packets at each round has been considered in [2]. The similar problem for $g_{H_1}(p, G)$ has been considered in [1]. Analogous problems on bus networks have been considered in [7]. In [9], a similar problem is considered for the toroidal mesh as they limit the size of the buffers. However they use a linear time model (of the form $\beta + L\tau$) and allow pipelining.

1.2 Graph Definitions

We use the definitions of the book [4].

In what follows:

- N will denote the number of vertices of G.
- D (or $D(G)$) will denote the *diameter* of a graph G, that is the maximum of all over the minimum distances between every pair of vertices.
- δ (or $\delta(G)$) will denote *minimum degree* of a graph G, that is the minimum over the degrees of all vertices of G.

Let us now present some of the classical networks for which we want to determine $g_{F_*}(1, G)$. We give their definitions in terms of graphs, but we recall that in our model, we will always work with the symmetric digraphs associated obtained by replacing each edge $[x, y]$ by the two opposite arcs (x, y) and (y, x).

Definition 1. The *cartesian sum* (also called *cartesian product* or *box product*) denoted by $G \Box G'$ of two graphs $G = (V, E)$ and $G' = (V', E')$, is the graph whose vertices are the pairs (x, x') where x is a vertex of G and x' is a vertex of G'. Two vertices (x, x') and (y, y') of $G \Box G'$ are adjacent if and only if $x = y$ and $[x', y']$ is an edge of G', or if $x' = y'$ and $[x, y]$ is an edge of G.

Definition 2. *The k-dimensional toroidal mesh* (or *the torus*) is the cartesian sum of k cycles of orders $p_1, p_2, \ldots, p_k$ and is denoted by $TM(p_1, p_2, \ldots, p_k) = C_{p_1} \Box C_{p_2} \Box \cdots \Box C_{p_k}$.

If all $p_i \geq 3$, it is a regular graph of degree $2k$. Its order is $p_1 \times p_2 \times \cdots \times p_k$ and its diameter is $\sum_{i=1}^{k} \lfloor \frac{p_i}{2} \rfloor$. When $p_1 = p_2 = \cdots = p_k$, we will use the abbreviated notation $TM(p)^k$ and suppose in what follows that $p \geq 3$.

Definition 3. When $p = 2$, $TM(2)^k = \underbrace{K_2 \Box K_2 \Box \cdots \Box K_2}_{k \text{ times}}$ is known as *the hypercube of dimension k*, denoted by $H(k)$. $H(k)$ can be therefore defined as the graph whose vertices are words of length k over the two-letter alphabet $\{0, 1\}$ and whose edges connect two words which differ in exactly one coordinate. A vertex $x_0 x_1 \cdots x_i \cdots x_{k-1}$ is thus joined to the vertices $x_0 x_1 \cdots \overline{x_i} \cdots x_{k-1}$ with $i = 0, 1, \ldots, k-1$. An edge between two vertices which differ in the ith coordinate will be called an edge of dimension i, or of type e_i.

In fact, these graphs are a particular case of a more general class of graphs.

Definition 4. Let $\mathcal{G}$ be a group and $\mathcal{S} = (s_0, s_1, \ldots, s_{d-1})$ be a set of generators of $\mathcal{G}$ not containing the identity. *The associated Cayley digraph* is the graph whose vertices are the elements of $\mathcal{G}$ and whose arcs are the couples $(x, s_i x)$ for $x \in \mathcal{G}$ and $s_i \in \mathcal{S}$.

As we restrict our attention to *symmetric digraphs*, we will always suppose that $\mathcal{S} = \mathcal{S}^{-1}$, that is if $s \in \mathcal{S}$, then $s^{-1} \in \mathcal{S}$. Some authors call them "Cayley graphs" as they identify the arcs $(x, s_i x)$ and $(s_i x, x)$ with the edge $[x, s_i x]$.

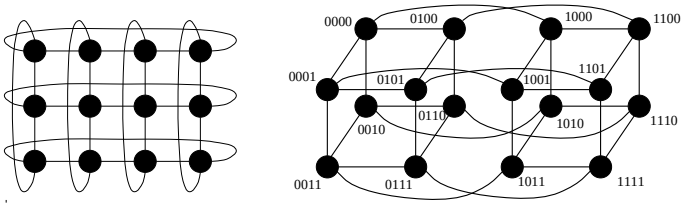

Fig. 1. Toroidal mesh $TM(3,4)$ and hypercube $H(4)$

Example 1. The symmetric k-dimensional torus $TM(p)^k$ is a Cayley digraph. The group $\mathcal{G}$ is $\mathbf{Z}_p^k$ with elements $(x_0 x_1 \cdots x_{k-1})$ and the $2k$ generators are the "canonical basis" $\pm e_i$ where $e_i = (0 \cdots 1 \cdots 0) = (x_0 x_1 \cdots x_{k-1})$ with $x_j = 0$ for $j \neq i$, $x_i = 1$. In the case $k = 2$, we have the 4 generators $s_0 = e_0 = (1,0)$, $s_1 = e_1 = (0,1)$, $s_2 = -e_0 = (-1,0)$, and $s_3 = -e_1 = (0,-1)$.

Example 2. In the case of $p = 2$, that is for the hypercube $\mathcal{G} = \mathbf{Z}_2^k$ with only k generators as $e_i = -e_i$.

Example 3. The symmetric "star-graph", denoted by $S(k)$, is the Cayley digraph whose vertices are the permutations of a k-element set, and where the generators are the $k-1$ transposition exchanging respectively 1 and i, for $2 \leq i \leq k$. We will denote a permutation π by the word $(\pi(1)\pi(2)\ldots\pi(k))$. For example in figure 2, the permutation identity $e = (1234)$ is joined by an arc to the permutations (2134), (3214) and (4231). $S(k)$ is regular of degree $k-1$, its order is $k!$ and its diameter is $\lfloor \frac{3(k-1)}{2} \rfloor$.

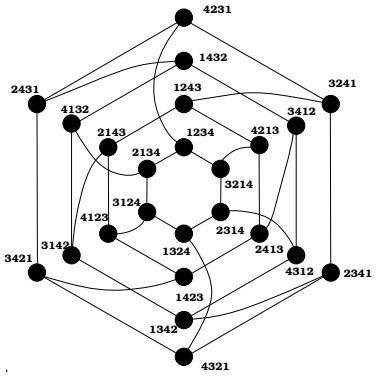

Fig. 2. Star-graph $S(4)$

1.3 Lower Bounds

If we have no restriction on the size of the messages, that is if we can exchange an arbitrary number of packets during each call, then it is well known that in the model F_*, one can gossip in $D(G)$ rounds. It is sufficient to use the greedy protocol where at each round each node send all the messages it has just received during the precedent round to all its neighbors.

Theorem 5. *For a graph G,*

$$g_{F_*}(\infty, G) = D(G) \tag{1}$$

So we have a first immediate lower bound.

Proposition 6.

$$g_{F_*}(1, G) \geq D(G) \tag{2}$$

The following lower bound is often more appropriate.

Proposition 7. *For a graph G having N vertices and minimum degree δ,*

$$g_{F_*}(1, G) \geq \lceil \frac{N-1}{\delta} \rceil \tag{3}$$

Proof. Let u be a vertex of minimum degree δ. During each round, u can receive only one packet from each of its neighbors and so at most δ packets. At the end of the protocol, u should have received $N-1$ packets, so the protocol needs at least $\lceil \frac{N-1}{\delta} \rceil$ rounds.

The same proof gives:

Proposition 8. *For a digraph G with N vertices and minimum in-degree d,*

$$g_{F_*}(1, G) \geq \lceil \frac{N-1}{d} \rceil \tag{4}$$

Remark. We can sometimes obtain a better lower bound in digraphs by considering an edge cut, separating a set S from its complement $V \setminus S$. Let $m^+(S, V \setminus S)$ be the size of this cut; we need at least $\frac{|S|}{m^+(S, V \setminus S)}$ rounds. The bound of (3) corresponds to the particular case where $V \setminus S = \{x\}$ and $m^+(S, V \setminus S) = d$.

1.4 Results

In this paper, we shall show that for k-dimensional tori (and in particular 2-dimensional tori), hypercubes and star-graphs, the lower bound is attained. That will follow from a more general property valid for Cayley symmetric digraphs with special automorphisms.

2 Gossiping in Cayley symmetric digraphs

In what follows G will always denote a Cayley symmetric digraph of in and out-degree d associated to a group $\mathcal{G}$ with identity element denoted e and a set of d generators $\mathcal{S} = (s_0, s_1, \ldots, s_{d-1})$. The arc $(x, s_i s)$ will be said of dimension i.

If A and B are two subsets of G, the set $\{ab \mid a \in A, b \in B\}$ will be denoted by AB; if $x \in G$ the set $\{ax, a \in A\}$ will be denoted by Ax.

2.1 Balanced sequence of sets

Definition 9. A sequence of sets $\{S_t\}_{t=0,\ldots,T-1}$ containing elements of G of length T is said to be *balanced* if there exist d vertices $x_0^t, x_1^t, \ldots, x_{d-1}^t$ (not necessarily distinct) in S_t such that $S_{t+1} \subset S_t \cup (\bigcup_{i=0,\ldots,d-1} s_i x_i^t)$.

The balanced property means that from S_t we can reach the vertices of $S_{t+1} \setminus S_t$ by using at most one arc in each dimension. In term of communications, suppose that S_0 is reduced to one vertex x (the initiator of a broadcast), then S_t represents the set of vertices which know the information of x at time t. The information is then forwarded to vertices of $S_{t+1} \setminus S_t$ by using at most once each generator. We can then associate to the sequence of sets $\{S_t\}_{t=0,\ldots,T-1}$ a *broadcast tree* by considering at round t only the arcs from S_{t-1} to S_t. If we label these arcs with t, the balanced property means that for any t, there exists at most one arc of dimension i.

Figure 3 and figure 4 give such sequences of sets for different tori, the sequence of sets are in fact implicitely defined by their associated broadcast tree. The label t of a vertex x indicates that x belongs to $S_t \setminus S_{t-1}$. So S_t consists of all the vertices with label less than or equal to t.

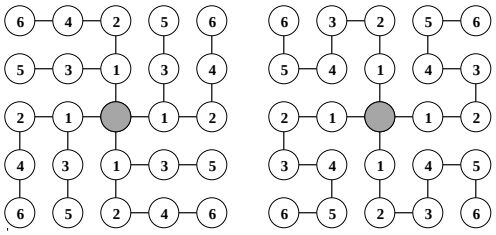

Fig. 3. Balanced set and broadcast tree in toroidal mesh $TM(5)^2$

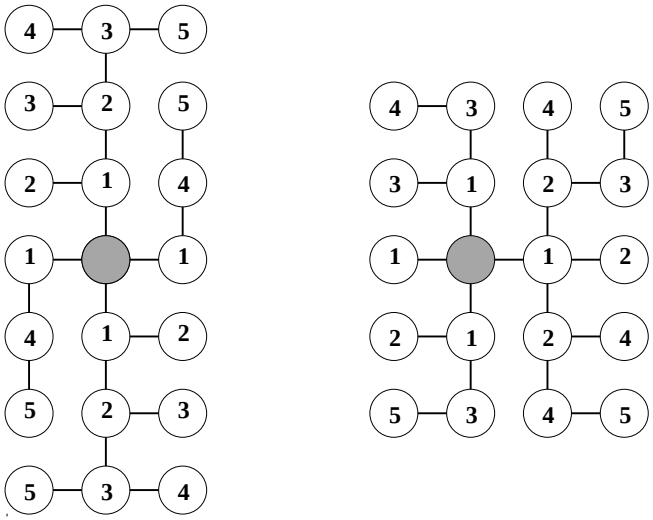

Fig. 4. Balanced set and broadcast tree in $TM(3,7)$ and $TM(4,5)$

Lemma 10. *Let G be a Cayley symmetric digraph with a balanced sequence of sets $\{S_t\}_{t=0,\ldots,T-1}$ with $S_0 = \{e\}$, then we can broadcast concurrently the packet of x to $S_T x$ in T rounds.*

Proof. The result is proved by induction on t. The proposition is true for $t = 0$. For an arbitrary vertex x, we suppose that vertices in $S_t x$ have received the packet of x at time t. Then the nodes $x_0^t x, x_1^t x, \ldots, x_{d-1}^t x$ which are in $S_t x$ send the packet of x to the vertices $s_i x_i^t x$, hence vertices of $S_{t+1} x = S_t x \cup (\bigcup_{i=0,\ldots,d-1} s_i x_i^t x)$ will have received the packet of x at time $t+1$.

Now we have to check that there is no communication conflict. As a vertex y sends the information of x along the arc s_i at time t if and only if $y = x_i^t x$, y uses the arc s_i at time t only to send the information originated from $(x_i^t)^{-1} y$. Hence the different calls can be performed in parallel.

Proposition 11. *Let G be a Cayley digraph and $\{S_t\}_{t=0,\ldots,T-1}$ be a balanced sequence of sets with $S_0 = \{e\}$, and $S_T = V$ then the gossip time of G is at most T.*

Proof. Just note that $xV = V$ and that gossiping corresponds to do concurrently broadcast to all the vertices from all the vertices.

Example 4. The examples of figure 3 and figure 4 show respectively that we can gossip in 6 rounds in $TM(5)^2$ and in 5 rounds in $TM(3,7)$ and $TM(4,5)$. These protocols are optimal as by (3): $g_{F_*}(1, TM(p_1,p_2)) \geq \lceil \frac{p_1 p_2 - 1}{4} \rceil$.

For the torus $TM(2p+1)^2$, we can easily exhibit a balanced sequence of sets S_t such that $|S_t| = 4t+1$ and so $g_{F_*}(1,G) \leq p^2 + p$ as for $T = p^2 + p$, $|S_T| = 4p^2 + 4p + 1 = (2p+1)^2$. As by (3): $g_{F_*}(1,G) \geq \lceil \frac{N-1}{4} \rceil = p^2 + p$, we conclude that : $g_{F_*}(1, TM(2p+1)^2) = p^2 + p$.

A very simple way to find the sets S_t consists in dividing the vertices of $TM(2p+1)^2$ different from $(0,0)$ into 4 symmetric subsets T_0, T_1, T_2 and T_3, where T_0 consists of the vertices (i,j) with $1 \leq i \leq p$, $0 \leq j \leq p$ and T_1 is

obtained from T_0 by a rotation ω of $\frac{\pi}{2}$, T_2 by the rotation ω^2 of π, and T_3 by the rotation ω^3 of $\frac{3\pi}{2}$ (see figure 5). Then the sequence S_t can be defined as follows: $S_0 = (0,0)$, S_1 consists of the 4 neighbors of $(0,0)$. They are $(1,0)$ $(0,1)$ $(-1,0)$ and $(0,-1)$ and belong respectively to T_0, T_1, T_2 and T_3. Then suppose that we have defined S_t by induction. To construct S_{t+1}, choose in T_0 a vertex y_0 connected to some x_0^t of $S_t \cap T_0$ by an arc in some dimension i_0. Then we put in S_{t+1} the vertex y_i of T_i, $0 \leq i \leq 3$, obtained by applying the rotation ω^i to y_0. Then y_i is connected to $x_i^t = \omega^i(x_0^t)$ by an arc in dimension $i_0 + i$. We will now extend that idea to Cayley symmetric digraph by using the notion of "*complete rotation*".

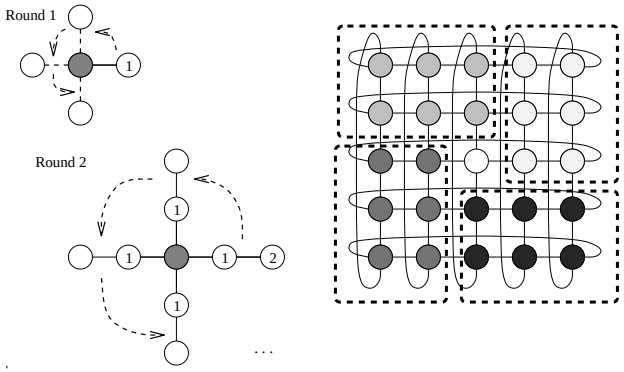

Fig. 5. Algorithm for the toroidal square mesh $TM(5)^2$ and its four same parts.

2.2 Complete rotation

Definition 12. An automorphism ω of a Cayley symmetric digraph $G = (\mathcal{G}, \mathcal{S})$ is *a complete rotation* if $\omega(e) = e$ and $\forall x \in G$ and $\forall i = 0 \cdots d-1$, $\omega(s_i x) = s_{i+1}\omega(x)$ (where the indices are taken modulo d).

Remark. A complete rotation is a particular case of what is called in the literature an *Adam automorphism* (automorphism such that $\omega(xy) = \omega(x)\omega(y)$ for any couple of elements x and y). It has the additional property of acting cyclically on the set of generators.

Definition 13. A vertex $x \neq e$ will be called a *fixed point* for a complete rotation ω if there exists some i, $1 \leq i \leq d-1$ such that $\omega^i(x) = x$. The set of fixed points of ω will be denoted by F_ω.

Remark. The orbit of $x \neq e$ under the action of ω is of length d except when x is a fixed point of ω in which case the orbit is degenerated.

Example 5. In the torus $TM(p)^2$, consider the automorphism ω which associates to x_0x_1 the vertex $(-x_1)x_0$ (which corresponds to a rotation of $\frac{\pi}{2}$ in the plane). Then ω is a complete rotation and we have $\omega(s_0) = \omega(1,0) = (0,1) = s_1$, $\omega(s_1) = (-1,0) = s_2$ and $\omega(s_2) = (0,-1) = s_3$.

If p is odd, then $F_\omega = \emptyset$.
If p is even, then $F_\omega = \{(0,\frac{p}{2}),(\frac{p}{2},0),(\frac{p}{2},\frac{p}{2})\}$ as $\omega(\frac{p}{2},\frac{p}{2}) = (-\frac{p}{2},\frac{p}{2}) = (\frac{p}{2},\frac{p}{2})$, $\omega^2(0,\frac{p}{2}) = (0,-\frac{p}{2}) = (0,\frac{p}{2})$ and $\omega^2(\frac{p}{2},0) = (-\frac{p}{2},0) = (\frac{p}{2},0)$.

More generally, in the torus $TM(p)^k$, the automorphism ω which associates to $x_0x_1\cdots x_{k-1}$ the vertex $(-x_{k-1})x_0\cdots x_{k-2}$ is a complete rotation. Indeed $\omega^i(e_0) = e_i$ for $0 \leq i \leq k-1$. Then $\omega^k(e_0) = \omega(e_{k-1}) = -e_0$ and $\omega^{k+i}(e_0) = -e_i$. One can check that ω satisfies the definition if we rank the generators of $TM(p)^k$ in the order $s_i = e_i$, $s_{i+k} = -e_i$ for $0 \leq i \leq k-1$. We will characterize the set of fixed points later on, but we can note that if p is odd, then $F_\omega = \emptyset$.

Example 6. For the hypercube $H(k)$, we consider similarly the automorphism ω which associates to $x_0x_1\cdots x_{k-1}$ the vertex $x_{k-1}x_0\cdots x_{k-2}$ (recall that $-x_i = x_i$). ω is clearly a linear one to one mapping of $\mathbf{Z}_2^k$ onto itself. Moreover if $s_i = e_i$, we have $\omega(s_i) = s_{i+1}$.

F_ω consists of the vertices such that the associated word is periodic, that is x can be written as $uu\cdots u$ where u is repeated r times with r dividing k and u being a word of length $q = \frac{k}{r}$.
So if d is a prime number, then F_ω is reduced to the vertex $111\cdots1$.

Lemma 14. *Suppose that there exists in G a complete rotation ω and let $\Gamma^*_{G\setminus F_\omega}(e)$ be the connected component of $G \setminus F_\omega$ containing e, then G admits a balanced sequence of sets $\{S_t\}_{t=0,\ldots,T-1}$, with $S_0 = \{e\}$ and $S_T = \Gamma^*_{G\setminus F_\omega}(e)$ with $T = \frac{|\Gamma^*_{G\setminus F_\omega}(e)-1|}{d}$.*

Proof. We construct inductively a sequence of sets using the following algorithm, which generalizes the construction given for the torus in example 4.

- $S_0 = \{e\}$.
- If S_t does not contain all the vertices of $\Gamma^*_{G\setminus F_\omega}(e)$, choose a vertex x in $\Gamma^*_{G\setminus F_\omega}(e)$ adjacent to S_t, and let $S_{t+1} = S_t \cup \{\omega^j(x);\ j = 0,\ldots,d-1\}$.
- If $S_t = \Gamma^*_{G\setminus F_\omega}(e)$, then stop.

Clearly, the algorithm stops when $S_T = \Gamma^*_{G\setminus F_\omega}(e)$. The construction implies that $\forall t : \omega(S_t) = S_t$, as $S_0 = \omega(S_0)$ and $\omega(S_{t+1}) = \omega(S_t \cup \{\omega^j(x);\ j = 0,\ldots,d-1\}) = S_t \cup \{\omega^j(x);\ j = 0,\ldots,d-1\} = S_{t+1}$. S_t being invariant under rotation none of the vertices $\omega^j(x)$ can belong to S_t when $x \notin S_t$. Hence $|S_{t+1}| = |S_t| + |\{\omega^j(x);\ j = 0,\ldots,d-1\}|$. As we always choose $x \notin F_\omega$, $|\{\omega^j(x);\ j = 0,\ldots,d-1\}| = d$, and $|S_{t+1}| = |S_t| + d$. The cardinality of S_T is now clearly $1 + dT$, and we get $T = \frac{|\Gamma^*_{G\setminus F_\omega}(e)|-1}{d}$.

Now let us prove that the sequence is balanced. Suppose that the added vertex x is joined to some vertex x_i^t of S_t along dimension i, that is $x = s_i x_i^t$. Then $\omega^j(x) = s_{i+j}\omega^j(x_i^t)$. As S_t is invariant under ω, the element $\omega^j(x_i^t)$ is in S_t. Let us call it x_{i+j}^t. Then $\bigcup_{j=0,\ldots,d-1}\omega^j(x) = \bigcup_{j=0,\ldots,d-1} s_{i+j}\, x_{i+j}^t$. Therefore the sequence is balanced.

Corollary 15. *If a Cayley symmetric digraph G admits a complete rotation ω such that $F_\omega = \emptyset$, then $g_{F_*}(1, G) = \frac{N-1}{d}$.*

Proof. As $F_\omega = \emptyset$, $\Gamma^*_{G\setminus F_\omega}(e) = V$ and so by lemma 14, $S_T = V$ with $T = \frac{N-1}{d}$. The result follows from proposition 11.

Proposition 16. *For the k-dimensional torus $TM(2p+1)^k$, we have*

$$g_{F_*}(1, TM(2p+1)^k) = \frac{N-1}{k} \tag{5}$$

Proof. That follows from (4) and corollary 15.

Example 7. Lemma 14 enables us also to conclude in many other cases. For example, consider the hypercube $H(k)$ with k prime. F_ω is reduced to the vertex $111\cdots1$.
By lemma 14, we can construct a balanced sequences of sets $\{S_t\}_{t=0,\ldots,T-1}$ with $S_T = V - \{111\cdots1\}$ and $T = \frac{2^k-2}{k}$. We can easily extend this sequence by adding $\{111\cdots1\}$ in S_{T+1} to obtain $S_{T+1} = V$. So we have a gossiping protocol in $\frac{2^k-2}{k}+1$ rounds. This result is optimal as, k being a prime, $2^k - 2 \equiv 0 \ mod\ k$ so $\lceil\frac{N-1}{k}\rceil = \frac{2^k-2}{k}+1$ and therefore $g_{F_*}(1, H(k)) = \lceil\frac{N-1}{k}\rceil$ for k prime.

Example 8. Similarly let us consider the 2-dimensional torus $TM(2p)^2$. By lemma 14, we have a balanced sequence of sets $\{S_t\}_{t=0,\ldots,T-1}$, with $S_T = V \setminus F_\omega$ with $T = \frac{4p^2-4}{4} = p^2 - 1$ as $|V| = 4p^2$ and $|F_\omega| = 3$ (see example 5). But again one can easily extend this sequence by adding the 3 vertices of F_ω, which have all their 4 neighbors in $S_T = V \setminus F_\omega$. We can for example join one of them say (p, p) via generator $s_0 = (1, 0)$; another say $(p, 0)$ via generator $s_1 = (0, 1)$ and the last one $(0, p)$ via generator $s_2 = (-1, 0)$. So we obtain $S_{T+1} = V$ and have an optimal gossip protocol in $p^2 = \lceil\frac{N-1}{4}\rceil$ rounds (see figure 6 where the elements of F_ω are in light grey and the corrections as above).

In fact, these results are particular cases of a more general proposition.

Let us recall that a subset A of vertices is said to be *independent* (or *stable*) if there is no arc between any couple of vertices of A. A is said to be *separating* (or a *vertex cut set*) if the deletion of vertices of A disconnects G.

Lemma 17. *If a Cayley symmetric digraph G admits a complete rotation such that F_ω is an independent and not separating set of G, then $g_{F*}(1, G) = \lceil\frac{N-1}{d}\rceil$.*

Proof. As F_ω is not a separating set, $\Gamma^*_{G\setminus F_\omega}(e) = V \setminus F_\omega$. So we can construct from lemma 14 a sequence with $S_0 = \{e\}$, $S_{t_0} = V \setminus F_\omega$, and $|S_{t_0}| = 1 + t_0 d$. After this step any vertex in $V \setminus S_{t_0+t} \subset F_\omega$ will be adjacent to $V \setminus F_\omega \subset S_{t_0+t-1}$ along any direction as F_ω is an independent set. So we can continue the construction by including d vertices in S_{t_0+t}, as long as there exist d vertices in $V \setminus S_{t_0+t-1}$. If there are less than d vertices in $V \setminus S_{t_0+t-1}$, we add in S_{t_0+t} all the remaining vertices. Let T be the value such that $S_T = V$. Therefore for $t < T$, $S_t = 1 + td$ and so $T = \lceil \frac{N-1}{d} \rceil$.

We have reduced the problem to check if F_ω is an independent, not separating set. Clearly in the digraph of examples 7 and 8, the set F_ω was independent and not separating set.

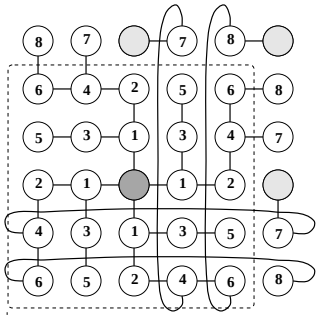

Fig. 6. Balanced sets in $TM(6)^2$

Lemma 18. *Let G be a Cayley symmetric digraph with complete rotation ω; If any pair of vertices (y, y') such that $d(y, y') = 2$, has at most two common neighbors, then F_ω is an independent set.*

Proof. a) If x and y are two adjacent vertices and if ω^t fixes both x and y, then t is a multiple of d. Indeed let $y = s_i x$, then $\omega^t(y) = s_{i+t}\omega^t(x) = s_{i+t}x = y = s_i x$. So $s_{i+t} = s_i$ which implies that t is a multiple of d.

b) Now let x and y be two adjacent vertices in F_ω. As x and $y \in F_\omega$, there exists p and q such that $\omega^p(x) = x$, $\omega^q(y) = y$, with p and q proper dividers of d. As x and y are adjacent, $p \neq q$ by a).

c) Let $y' = \omega^p(y)$ and $x' = \omega^q(x)$. As $p \neq q$ and p and q are proper dividers of d by a), $x' \neq x$ and $y' \neq y$.
Furthermore y' is a neighbor of x as $y' = \omega^p(y) = \omega^p(s_i x) = s_{i+p}\omega^p(x) = s_{i+p}x$ and so $d(y, y') = 2$. Similarly x' is a neighbor of y and y'.

d) By the hypothesis of the lemma, x and x' are the only neighbors of y and y'. So repeating the argument of c) with x' and y', we obtain that $\omega^q(x')$ is a neighbor of y and y' and so $\omega^q(x') = x$ and similarly $\omega^p(y') = y$.

e) Therefore $\omega^{2q}(x) = x$ and $\omega^{2p}(y) = y$. By b), $\omega^{2p}(x) = x$ and $\omega^{2q}(y) = y$. So by a), $2p$ and $2q$ are multiple of d. As p and q are proper dividers of d, the only possibility is that d is even and $p = q = \frac{d}{2}$ contradicting $p \neq q$.

Corollary 19. *For any complete rotation ω of $TM(p)^k$, $H(k)$ and $S(k)$, F_ω is an independent set.*

Proof. For $TM(p)^k$ and $H(k)$, that follows from the fact that there is a unique C_4 containing any pair of vertices at distance 2.

For $S(k)$, that follows from the fact that there is no C_4 as the girth of $S(k)$ is 6.

Proposition 20. *For the hypercube $H(k)$,*

$$g_{F_*}(1, H(k)) = \lceil \frac{2^k - 1}{k} \rceil \tag{6}$$

Proof. By lemma 17 and lemma 18, it remains to prove that F_ω is not a separating set of $H(k)$.

a) First let us prove that if $x \notin F_\omega$, x has at most two neighbors in F_ω. For that, let f be the application which associates to a vertex x of $H(k)$ the complex $f(x) = \sum_{i=o}^{k-1} x_i \theta^i$ where θ is a primitive root of 1 of order k. If $x \in F_\omega$, then $f(x) = 0$. Indeed x is of the form $uu \cdots u$ where u is repeated r times with r dividing k and u being a word of length $q = \frac{k}{r}$. Then $f(x) = \sum_{i=0}^{r-1} \theta^{iq} \varphi(u) = 0$ where $\varphi(u) = \sum_{i=0}^{q-1} u_i \theta^i$. Note also that f is nearly linear as $f(s_i x) = \varepsilon \theta^i + f(x)$ where $\varepsilon \in \{-1, 1\}$ depending on x_i.
Now let $x \notin F_\omega$ and $s_i x$ and $s_j x$ be two neighbors of x in F_ω. Then $f(s_i x) = f(s_j x) = 0$ which implies $\varepsilon \theta^i + f(x) = \varepsilon \theta^j + f(x)$, that is $\theta^i = \pm \theta^j$. So either $i = j$ or $\theta^{i-j} = -1$ which implies k even and $i = j + \frac{k}{2}$. In summary either k is odd and x has at most one neighbor in F_ω, or k is even and x has at most two neighbors in F_ω.

b) We prove now that $V \setminus F_\omega$ is connected. For this we show inductively on j that from e one can reach all the vertices in $V \setminus F_\omega$ at distance at most j. We start from vertex e, then as the vertices at distance 1 are not in F_ω we can reach these vertices. Vertices at distance 2 which are not in F_ω can be reached as they have one neighbor at distance 1, vertices at distance $j > 2$ have at least three neighbors at distance $j - 1$; so by a) one of them is not in F_ω and can be reached from e.

Lemma 21. *For the k-dimensional torus $TM(p)^k$,*

$$g_{F_*}(1, TM(p)^k) = \lceil \frac{p^k - 1}{2k} \rceil \tag{7}$$

Proof. By lemma 17 and lemma 18, it remains to prove that F_ω is not separating. That is clear if p is odd as $F_\omega = \emptyset$. If p is even, we found only a tedious and technical proof and so we skip it. In some cases (k even), we can do a

proof identical to that of the hypercube. The difficulty comes from the fact that the fixed points are not "*fully periodic*". Indeed if $\omega^q(x) = x$, then we have $x_i = x_{i+q} = \cdots = x_{i+rq}$ for $0 \leq i \leq q - i$ and r such that $i + rq < k$ plus $x_i = -x_{i+(r_0+1)q}$ where $i + r_0 q < k \leq i + (r_0 + 1)q$. So the fixed points are of the form $u_1u_2u_1u_2 \cdots u_1u_2u_1$ with $(-u_2)(-u_1) = u_1u_2$. For example, for $k = 5$, $x = a(-a)a(-a)a$ is a fixed point as $\omega^2(x) = x$ (here $u_1 = a$ and $u_2 = -a$).

Proposition 22. *For the star-graph $S(k)$,*

$$g_{F_*}(1, S(k)) = \lceil \frac{N-1}{k} \rceil \tag{8}$$

Proof. a) First let us prove that the star-graph admits a complete rotation.
Recall that we denote a permutation π by the word $\pi(1)\pi(2)\cdots\pi(k)$. Let σ be the permutation $(1, 3, 4, \ldots, k, 2)$; σ fixes 1 and acts cyclically on the elements $2, \ldots, n$. Let define ω as the automorphism which associates to a permutation π the permutation $\sigma^{-1}\pi\sigma$. We have $\omega(e) = e$ and $\omega(\pi\pi') = \omega(\pi)\omega(\pi')$.
Now let the generator s_i, $0 \leq i \leq k-2$ denote the transposition exchanging 1 and $i + 2$. Then $\omega(s_i) = \sigma^{-1}s_i\sigma$.
$s_i\sigma = (i + 3, 3, 4, \ldots, i + 2, 1, i + 4, \ldots, k, 2)$ for $i \leq k - 3$
$s_i\sigma = (2, 3, 4, \ldots, k, 1)$ for $i = k - 2$
Then $\sigma^{-1}s_i\sigma = s_{i+1}$.

b) By corollary 19, F_ω is an independent set. Now let us note that if a permutation $\pi \in F_\omega$, then necesarily $\pi(1) = 1$. But one can check that the set U of permutations such that $\pi(1) = 1$ is itself not a separating set. Indeed let x be any vertex such that $\pi(1) \neq 1$ and let j be such that $\pi(j) \neq 1$. We first do the s_{j-2} (exchanging of 1 and j) and then do other generators different from s_{j-2} so we will never use a vertex in U.

3 Conclusion

In this paper, we have exhibited a way how to achieve optimal gossiping with packet size limited to 1 in Cayley symmetric digraphs having a complete rotation. It will be nice to characterize this class of digraphs; it does not include all the Cayley symmetric digraphs as they should be *arc-transitive*, but it contains many interesting networks like hypercubes, star-graphs, k-dimensional tori. The technics developped in this paper can be easily extend to the case where we allow p packets to be send (and receive) during each round and it can be shown that in the networks considered in the article, one can optimally gossip in $\lceil \frac{N-1}{pd} \rceil$ rounds.

We have also been able to construct balanced sequences of sets for any 2-dimensional torus $TM(p, q)$, so we can optimally gossip in them. We conjecture that optimal gossip exists in any k-dimensional torus $TM(p_1, p_2, \ldots, p_k)$ although the construction of balanced sequence of sets can be tedious. We have also considered the case of 2-dimensional meshes, which are not Cayley graphs

and on which we have obtained partial results.

Finally let us note that many interesting questions remain open. For example, it would be interesting to determine the computational complexity of computing $g_{F_*}(1,G)$ or $g_{F_*}(p,G)$ for a fixed p; it is very likely that it is NP-hard. It will be also interesting to study the same problem with the model H_* (half-duplex) although in that case, the case $p=\infty$ is already difficult.
Another direction will be to use the technics of this paper for other communication problems.

References

1. A. Bagchi, E.F. Schmeichel and S.L. Hakimi.: Parallel Information Dissemination by Packets. SIAM J.on Computing, **23** (1994) 355–372
2. J.-C. Bermond, L.Gargano, A. Rescigno and U. Vaccaro.: Fast Gossiping by Short Messages. *in ICALP 95,* Hungary, (1995)
3. D. P. Bertsekas and J. N. Tsitsiklis.: Parallel and Distributed Computation: Numerical Methods. Prentice–Hall, Englewood Cliffs, NJ, (1989)
4. J. de RUMEUR.: Communications dans les réseaux de processeurs. Masson Paris, (1994)
5. P. Fraigniaud and E. Lazard.: Methods and Problems of Communication in Usual Networks. Disc.Appl.Math **53** (1994) 79–134
6. S.M. Hedetniemi, S.T. Hedetniemi and A.L. Liestman.: A Survey of Gossiping and Broadcasting in Communication Networks. NETWORKS, **18** (1988) 319–349
7. A. Hily and D. Sotteau.: Communications in Bus Networks. Parallel and Distributed Computing, Lectures Notes in Computer Science, Springer–Verlag, **805** (1994) 197–206
8. J. Hromkovič, R. Klasing, B. Monien and R. Peine.: Dissemination of Information in Interconnection Networks (Broadcasting and Gossiping). To appear in: F. Hsu, D.-Z. Du (Eds.) Combinatorial Network Theory, Science Press & AMS
9. J.-C. König, P.S. Rao, and D. Trystram.: Analysis of Gossiping Algorithms in Torus with Restricted Bufferization Capabilities. Technical Report IMAG Grenoble, (1994)
10. M. Mahéo and J.-F. Saclé.: Note on the Problem of Gossiping in Multidimensional Grids. Disc.Apple.Math **53** (1994) 287–290

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