

Approximation algorithms for information dissemination problems.

Pierre Fraigniaud, Sandrine Vial

► **To cite this version:**

Pierre Fraigniaud, Sandrine Vial. Approximation algorithms for information dissemination problems.. [Research Report] LIP RR-1995-15, Laboratoire de l'informatique du parallélisme. 1995, 2+23p. hal-02101953


HAL Id: hal-02101953

<https://hal-lara.archives-ouvertes.fr/hal-02101953>

Submitted on 17 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Laboratoire de l'Informatique du Parallélisme


Ecole Normale Supérieure de Lyon
Unité de recherche associée au CNRS n° 1398

Approximation algorithms for information dissemination problems

Pierre Fraigniaud
Sandrine Vial

September 1995

Research Report N° 95-15



Ecole Normale Supérieure de Lyon

46 Allée d'Italie, 69364 Lyon Cedex 07, France

Téléphone : (+33) 72.72.80.00 Télécopieur : (+33) 72.72.80.80

Adresse électronique : lip@lip.ens-lyon.fr

Approximation algorithms for information dissemination problems

Pierre Fraigniaud
Sandrine Vial

September 1995

Abstract

Broadcasting and gossiping are known to be NP-hard problems. This paper deals with approximation algorithms for such problems. We consider both round-complexity and step-complexity in the telephone model. After an overview of previously derived approximation algorithms, we present new strategies for broadcasting and gossiping in any graphs. Broadcasting strategies are based on the construction of edge-disjoint spanning trees. Gossiping strategies are based on on-line computation of matchings along with the gossiping process. Our approximation algorithms for broadcasting offer almost optimal complexity when the number of messages to be broadcasted is large. We show that our best approximation algorithm for gossiping performs optimally in many cases. We also show experimentally that it can perform faster than the best known handmade algorithms in some particular cases.

Keywords: Broadcasting, Gossiping, Approximation algorithms

Résumé

La diffusion et l'échange total sont des problèmes NP-durs. Cet article traite d'algorithmes d'approximation pour ces problèmes. Nous considérons à la fois la complexité en temps et en taille de messages dans le modèle "téléphone". Après un bref aperçu des algorithmes d'approximation existants, nous présentons de nouvelles stratégies pour la diffusion et l'échange total dans un graphe quelconque. Les stratégies pour la diffusion sont basées sur la construction d'arbres de recouvrement arêtes-disjoints. Celles pour l'échange total sont basées sur le calcul de couplages tout au long du processus d'échange total. Nos algorithmes d'approximation pour la diffusion ont une complexité presque optimale lorsque le nombre de messages diffusés est grand. Nous montrons que notre meilleur algorithme d'approximation pour l'échange total s'exécute optimalement dans de nombreux cas. Nous montrons, expérimentalement, que cet algorithme peut dans certains cas, produire des algorithmes d'échange total plus rapides que les meilleurs algorithmes connus.

Mots-clés: Diffusion, Échange Total, Algorithmes d'approximation

Approximation algorithms for information dissemination problems

Pierre Fraigniaud and Sandrine Vial *

Laboratoire de l'Informatique du Parallélisme – CNRS
Ecole Normale Supérieure de Lyon
69364 Lyon Cedex 07, France

Abstract

Broadcasting and gossiping are known to be NP-hard problems. This paper deals with approximation algorithms for such problems. We consider both round-complexity and step-complexity in the telephone model. After an overview of previously derived approximation algorithms, we present new strategies for broadcasting and gossiping in any graphs. Broadcasting strategies are based on the construction of edge-disjoint spanning trees. Gossiping strategies are based on on-line computation of matchings along with the gossiping process. Our approximation algorithms for broadcasting offer almost optimal complexity when the number of messages to be broadcasted is large. We show that our best approximation algorithm for gossiping performs optimally in many cases. We also show experimentally that it can perform faster than the best known handmade algorithms in some particular cases.

*Both authors are supported by the research programs PRS and ANM of the CNRS.

1 Introduction

Information dissemination problems

Given a connected and undirected graph G , an information dissemination problem on G consists in realizing a communication pattern known in advance between the vertices of G . *Broadcasting* and *gossiping* are two famous information dissemination problems. In broadcasting, a single node knows a piece of information that it wants to send to all the other nodes of the network. In gossiping, all nodes know some piece of information, a gossip, and they want to exchange their pieces so that, at the end, all the nodes will be aware of the whole information.

Information dissemination algorithms are not only interesting as graph theoretical parameters, but also as software kernels for high performance communication networks. They are helpful for the design of parallel and/or distributed algorithms [30]. Because of that, information dissemination has received a lot of interest from researchers in the past. These problems were approached following several different manners.

The first one consists, given a particular graph G selected because of its interest in other domains (as network architecture for parallel computer), to compute the *broadcast time* and the *gossip time* of G . We refer to [14, 17, 19] for surveys of results obtained in this direction.

Another approach consists, given a particular problem, to find networks allowing to solve this problem in a minimum time with minimum communication facilities (typically the number of edges). This yields to the construction of *minimum broadcast* or *minimum gossip networks*. We refer to [4, 15, 27, 41] for recent results in this domain.

Approximation algorithms

It is only recently that a third approach has been considered: given a specific information dissemination problem \mathcal{P} (as the broadcasting problem or the gossiping problem), deriving a general algorithm \mathcal{A} that returns, for any graph G , an algorithm $\mathcal{A}[G]$ that solves \mathcal{P} on G .

Unfortunately, but not surprisingly, the problem of finding the minimum number of rounds necessary to broadcast an atomic message from one processor to all the others is NP-hard [16](see [38] for his proof), even for particular class of graphs, see [34]. Therefore, the general broadcast problem is hard to attack (as well as for the gossip problem), and one cannot hope to find optimal information dissemination algorithms in a reasonable (i.e. polynomial) amount of time. This is why \mathcal{A} generally only produces approximate solutions $\mathcal{A}[G]$, for all the graphs G .

The efficiency of an approximation algorithm \mathcal{A} is measured by its complexity (we are looking for polynomial approximation algorithms) and by the efficiency of the communication algorithms generated. In particular, one usually tries to express the complexity of $\mathcal{A}[G]$ as a function of the complexity of the considered information dissemination problem on G , or as function of static characteristics of G (diameter, degree, bisection,...).

Telephone model and complexity measurements

Approximation algorithms were derived under the following communication constraints (usually called telephone model): communications proceed by a serial-parallel sequence of *calls*. Many calls

can be performed in parallel, but a given call involves exactly two neighboring vertices, and a vertex can participate to at most one call at a time. Two vertices participating to the same call can exchange all the information they know. The total number of calls necessary to complete some specified information dissemination problem is called the *call-complexity* of that problem. Even if the call-complexity gives interesting information about the difficulty of solving a problem (see [17]), one often prefers other measures directly related to several notions of time. The most usual consists in counting the number of *rounds*.

Definition 1 *A round is composed of the set of calls performed at a same given time. The round-complexity of an information problem \mathcal{P} on a graph G is the minimum number of rounds necessary to solve \mathcal{P} on G .*

In particular, the round-complexity of broadcasting from node x of graph G is denoted by $b_{\mathcal{R}}(G, x)$, and the round-complexity of broadcasting in G is denoted by $b_{\mathcal{R}}(G)$, that is $b_{\mathcal{R}}(G) = \max_x b_{\mathcal{R}}(G, x)$. Similarly, the round-complexity of gossiping in G is denoted by $g_{\mathcal{R}}(G)$. Clearly, for any graph of order n ,

$$\log_2 n \leq b_{\mathcal{R}}(G) \leq g_{\mathcal{R}}(G) \leq 2b_{\mathcal{R}}(G) - 1 \leq 2n - 3 \quad (1)$$

Hypercubes satisfy $\log_2 n = b_{\mathcal{R}}(G) = g_{\mathcal{R}}(G)$, and the star of n vertices (one center and $n - 1$ leaves) satisfies $g_{\mathcal{R}}(G) = 2b_{\mathcal{R}}(G) - 1 = 2n - 3$. We have also: for any graph of maximum degree Δ and diameter D ,

$$D \leq b_{\mathcal{R}}(G) \leq \Delta D \quad (2)$$

because it is possible to broadcast in at most ΔD rounds in any shortest paths spanning tree of G of depth D and maximum degree Δ .

Similarly, given an approximation algorithm \mathcal{A} for the broadcasting problem (resp. the gossiping problem), the number of rounds necessary to complete $\mathcal{A}[G]$ on a graph G is denoted by $b_{\mathcal{R}}(\mathcal{A}[G])$ (resp. $g_{\mathcal{R}}(\mathcal{A}[G])$). The general objective is to derive approximation algorithms for the broadcasting and the gossiping problem such that, for any graph G , $b_{\mathcal{R}}(\mathcal{A}[G])$ (resp. $g_{\mathcal{R}}(\mathcal{A}[G])$) differs from $b_{\mathcal{R}}(G)$ (resp. $g_{\mathcal{R}}(G)$) by a small multiplicative or additive factor.

Previous works

The problem of finding heuristic algorithms for broadcasting in point-to-point networks was initiated by Scheuermann and Wu. In [37], they present a formulation of the broadcast problem based on maximum matching: a broadcast from a vertex u of a graph $G = (V, E)$ is simply a sequence

$$\{u\} = V_0, E_1, V_1, E_2, V_2, \dots, E_k, V_k = V$$

such that, for any $i \in \{1, 2, \dots, k\}$,

1. $V_i \subset V$ and $E_i \subset E$,
2. E_i is a (maximum) matching between V_{i-1} and $\Gamma(V_{i-1})$ where, for any set S of vertices, $\Gamma(S)$ denotes the set of all the neighbors of the vertices in S but S , and
3. $V_i = V_{i-1} \cup E_i(V_{i-1})$, where, for any set S of vertices and any matching M between S and $\Gamma(S)$, $M(S)$ denotes the set of the vertices of $\Gamma(S)$ that are matched by M to vertices in S .

This formulation allows Scheuermann and Wu to derive a recursive formula

$$b_{\mathcal{R}}(G, S) = 1 + \min_M \{b_{\mathcal{R}}(G, S \cup M(S))\}$$

where the minimum is taken over all the maximum matchings M between S and $\Gamma(S)$. They present several heuristics to solve this recurrence. However, although they give some experimental results to illustrate the efficiency of their approach, no bounds were derived on the broadcast complexity of their solutions. In fact, it was shown in [24] that even the best heuristics can be $\Omega(\sqrt{n})^1$ times worse than optimal for some particular graphs.

Kortsarz and Peleg have then presented in [24] an approximation algorithm \mathcal{A}_{KP} that polynomially computes a scheme that completes a broadcast in minimum time within a constant factor of the optimal plus an additive factor of $O(\sqrt{n})$. More precisely,

$$b_{\mathcal{R}}(\mathcal{A}_{KP}[G]) \in O(b_{\mathcal{R}}(G) + \sqrt{n})$$

for any graph G of order n . This approximation algorithm is based on a decomposition of G in clusters of size \sqrt{n} and on the use of a solution for a restricted version of the so called *minimum weight cover* problem. However, as it was pointed out by the authors, the algorithm in [24] is not efficient when applied to graphs of small broadcast time, namely when $b_{\mathcal{R}}(G) \in o(\sqrt{n})$.

Ravi has recently presented in [36] a polynomial approximation algorithm \mathcal{A}_R that appears to be particularly efficient for fast broadcast graphs. In particular, his algorithm satisfies

$$b_{\mathcal{R}}(\mathcal{A}_R[G]) \in O\left(\frac{\log^2 n}{\log \log n} b_{\mathcal{R}}(G)\right)$$

for any graph G of order n . This algorithm is based on the so called *poise* of a graph. The poise $P(T)$ of a tree T is the quantity: maximum degree of T plus diameter of T . The poise $P(G)$ of a graph G is the minimum poise of any of its spanning tree. Ravi showed that the poise of a graph is strongly related to its broadcast time. More precisely, he showed that

$$b_{\mathcal{R}}(G) \in \Omega(P(G)) \text{ and } b_{\mathcal{R}}(G) \in O\left(P(G) \frac{\log n}{\log \log n}\right)$$

for any graph G of order n . Since computing the poise of a graph is NP-hard, Ravi developed a polynomial algorithm for computing a tree T such that $P(T) \in O(\log n P(G) + \log^2 n)$. His approximation algorithm directly follows this result.

Feige, Peleg, Raghavan and Upfal have derived in [11] a simple but quite efficient randomized broadcast algorithm \mathcal{A}_{RB} . The algorithm \mathcal{A}_{RB} performs as follows:

```

repeat
  for all  $v \in V$ , in parallel do
    if  $v$  has already received the message then
       $v$  sends the message to a randomly and uniformly chosen neighbor;
end.
```

To study the performances of their algorithm, Feige et al. consider the *almost sure coverage time* $\hat{b}_{\mathcal{C}}(\mathcal{A}_{RB}[G])$ of a network G , that is defined by

$$Prob \left[\begin{array}{l} \text{all the vertices have received} \\ \text{the message after } \hat{b}_{\mathcal{C}}(\mathcal{A}_{RB}[G]) \\ \text{rounds of } \mathcal{A}_{RB}[G] \end{array} \right] \geq 1 - \frac{1}{n}.$$

¹All along this paper, $O(f)$, $\Omega(f)$, and $o(f)$ denote the sets of functions that are asymptotically upper bounded, lower bounded, and negligible in front of a function that differs from f by a constant, respectively.

They show that for every network G of order n , maximum degree Δ and diameter D :

$$\hat{b}_{\mathcal{C}}(\mathcal{A}_{RB}[G]) \in O(\Delta(D + \log n))$$

but that \mathcal{A}_{RB} can performs better is some particular case; for instance,

$$\hat{b}_{\mathcal{C}}(\mathcal{A}_{RB}[Q_d]) \in \Theta(\log n)$$

where Q_d is the d -dimensional hypercube. Moreover, they show that for almost all random graphs G of order n and whose edges are present independently with probability p , we have

$$b_{\mathcal{R}}(\mathcal{A}_{RB}[G]) \in \Theta(\log n)$$

if $p \geq (1 + \epsilon) \frac{\log n}{n}$ for some fixed constant $\epsilon > 0$.

Step-complexity

From Equation 1, we have $g_{\mathcal{R}}(G) \in \Theta(b_{\mathcal{R}}(G))$ for any graph G . Moreover, a gossiping can be performed in two phases by first accumulating (a reverse broadcast) all the pieces of information in some particular vertex, and then broadcasting the single message obtained by a concatenation of all the pieces of information from this vertex to all the others. For any approximation algorithm B for broadcasting, we denote by B^2 the corresponding two-phases approximation algorithm for gossiping. Thus, all the algorithms B for the broadcasting problem that we have previously described give adapted solutions B^2 for the gossiping problem. The main criticism about this is that the number of exchanged messages increases exponentially during the first phase, and the message broadcasted during the second phase has size n . Counting 1 when manipulating messages of size $\Theta(n)$ might not reflect the real complexity of the problem.

One can refine the study of the complexity of information dissemination problems under the telephone model by counting not only the number of rounds, but also the sum of the maximum number of pieces of information exchanged during each successive round. During a call between two vertices x and y , if they exchange k_x pieces of information (a piece is supposed to be atomic) from x to y , and k_y pieces of information from y to x , then this call requires $k = \max\{k_x, k_y\}$ steps. A round requires the maximum, over all the calls of that round, of the number of steps of these calls. The rounds being synchronized, r rounds require $\sum_{i=1}^r k_i$ steps where k_i is the number of steps of round i , for all i .

Definition 2 *A step is the operation consisting of transmitting an atomic piece of information during a call between two vertices. The step-complexity of an information problem \mathcal{P} on a graph G is the minimum number of consecutive steps necessary to solve \mathcal{P} on G .*

In particular, the step-complexity of a gossiping in G is denoted by $g_{\mathcal{S}}(G)$, and, given an approximation algorithm \mathcal{A} for the gossiping problem, the number of steps necessary to complete $\mathcal{A}[G]$ on a graph G is denoted by $g_{\mathcal{S}}(\mathcal{A}[G])$. Clearly, for any graph of order n ,

$$g_{\mathcal{S}}(G) \in \Omega(n + b_{\mathcal{R}}(G)) \tag{3}$$

since every vertex must receive at least $n - 1$ pieces of information, and a gossiping includes at least one broadcasting. Therefore, the two-phases gossiping algorithm obtained by accumulating and broadcasting yields a quite inefficient algorithm when considering the number of steps:

$$g_{\mathcal{S}}(\mathcal{B}^{*2}[G]) \in O(n \times b_{\mathcal{R}}(\mathcal{B}^*[G]))$$

where \mathcal{B}^* is a polynomial approximation algorithm for broadcasting under the telephone model with the minimum number of rounds. This is at least $b_{\mathcal{R}}(G)$ times worse than the lower bound of Equation 3.

Also, when broadcasting a message that is not an atomic piece of information but composed of k pieces of information, then again algorithms derived to minimize the number of rounds might be inefficient. Indeed, let $b_{\mathcal{S}}(G, x)$ be the step-complexity of a broadcasting from a node x of a graph G , and let $b_{\mathcal{S}}(G) = \max_x b_{\mathcal{S}}(G, x)$ be the step-complexity of a broadcasting in G . We have,

$$b_{\mathcal{S}}(G) \in \Omega(k + b_{\mathcal{R}}(G)) \tag{4}$$

when broadcasting a message of size k . On the other hand, \mathcal{B}^* satisfies

$$b_{\mathcal{S}}(\mathcal{B}^*[G]) \in O(k \times b_{\mathcal{R}}(\mathcal{B}^*[G]))$$

which is at least $b_{\mathcal{R}}(G)$ times worse than the bound of Equation 4.

Our results

In this paper, we focus on both measurements: round and step-complexity. More precisely, we present a polynomial approximation algorithm for broadcasting that almost matches the bound of Equation 4 for large k . This result is based on the construction of λ edge-disjoint spanning trees in the considered graph (λ denotes the edge-connectivity of the graph). Pieces of information move without conflict along the edges of the tree.

Concerning gossiping, we present several strategies to construct efficient approximation algorithms. All strategies are based on matchings of different kinds. Off-line matchings are computed in advance, and the gossiping proceeds by a succession of exchanges along the edges of the matchings. On-line matchings are more efficient. In this case, the matching used at round t of an approximation algorithm for gossiping depends on the way the information flows during the $t - 1$ previous rounds. In particular, we present the Dynamic Weights algorithm that is shown to work asymptotically quite efficiently on particular graphs, and is experimentally shown to work well on about all the usual examples of graphs used to interconnect processors of a parallel computer. In fact, we have sometimes observed that the DW approximation algorithm performs faster than the best known handmade algorithms for some particular graphs.

Extensions

We extend our results in two directions.

First, we consider the so-called store-and-forward routing model, and show that the approximation algorithms derived in Section 2 and 3 apply in this case, and can offer good performances in practice.

Our second direction of investigation consists in studying the so-called telegraph model. In this abstract model, information cannot traverse an edge in both directions simultaneously. This strongly reduces the performances of the broadcasting and gossiping algorithms derived to satisfy this constraint. We show that our best approximation algorithm for gossiping does not suffer too much of this restriction, and applies in the telegraph model with almost the same efficiency as in the telephone model.

Organization of the paper

The next section is devoted to broadcasting. Section 3 deals with the gossiping problem. Finally, section 4 presents results obtained in both store-and-forward and telegraph models.

2 Broadcasting

In all this section, the size of the message to be broadcasted is denoted by k , and we recall that the best known polynomial approximation algorithm for broadcasting when counting only the number of rounds is denoted by \mathcal{B}^* .

The following result shows that it is easy to design efficient algorithms that minimize the number of steps from efficient algorithms designed to minimize the number of rounds. This is particularly true when $k \in O(\frac{b_{\mathcal{R}}(\mathcal{B}^*[G])}{\Delta})$ where Δ is the maximum degree of the considered graph G .

Theorem 1 *There exists a polynomial time approximation algorithm \mathcal{A} for broadcasting such that, for any graph G of maximum degree Δ , $b_{\mathcal{S}}(\mathcal{A}[G]) \in O(k\Delta + b_{\mathcal{R}}(\mathcal{B}^*[G]))$.*

Proof. Let G be any graph of maximum degree Δ . The algorithm $\mathcal{A}[G]$ proceeds as follows. Consider a broadcast tree T induced by \mathcal{B}^* on G . The edges of T are labeled by the time at which they are used in \mathcal{B}^* . Let $\ell(e)$ be the label of an edge e of the tree T . The i -th piece of information of the message is sent through e at time $(i-1)\Delta_T + \ell(e)$ where Δ_T is the maximum degree of T . The first piece is known by all the vertices after a time $b_{\mathcal{R}}(\mathcal{B}^*[G])$. Every Δ_T steps, a new piece is received by all the vertices. Thus this algorithm takes a total time $b_{\mathcal{R}}(\mathcal{B}^*[G]) + (k-1)\Delta_T$. \square

However, for long messages ($k \gg b_{\mathcal{R}}(G)$), the result of Theorem 1 is Δ times worse than the bound of Equation 4. The next result shows that one can be much more efficient for long messages. Before stating this result, recall that the *chromatic index* [2] of a graph G is the minimum number q such that it is possible to label all the edges of G with numbers in $\{0, \dots, q-1\}$, each node having all its incident edges labeled with different numbers. Vizing's Theorem (see [2]) says that for any graph of maximum degree Δ , the chromatic index q of G satisfies $\Delta \leq q \leq \Delta + 1$. The bad news is that knowing whether or not $q = \Delta$ is NP-complete [16]. The good news however is that finding a coloring of the edges of a graph G with at most $\Delta + 1$ colors is polynomial (from the proof of Vizing's Theorem). This result is the base of Theorem 2.

The telephone model implies that when two nodes x and y of a graph G communicate during a call, they *exchange* their information, that is links are *full-duplex* and the underlying communication network could be represented by a symmetric digraph G^* where G^* denotes the digraph obtained from G by replacing every edge (x, y) of G by two symmetric arcs x, y and y, x .

Theorem 2 *There exists a polynomial time approximation algorithm \mathcal{A} for broadcasting such that, for any graph G of maximum degree Δ and edge-connectivity λ , $b_{\mathcal{S}}(\mathcal{A}[G]) = (\Delta + 1)\lceil \frac{k}{\lambda} \rceil + h$, where h is independent of k .*

Proof. Following Menger's Theorems (see [33]), Edmonds [8] (see also [32]) has shown that for any vertex r of a finite digraph H , there exists λ arc-disjoint spanning trees of H rooted at r (λ is the arc-connectivity of the digraph H). If the edge-connectivity of a graph G is λ , then λ is also

the arc-connectivity of G^* . Thus, for any vertex r of a finite digraph G^* , there exists λ arc-disjoint spanning trees of G^* rooted at r , that is spanning trees of G where any edge (x, y) appears at most in two trees, each using (x, y) in a distinct direction. Let us call $T_1, T_2, \dots, T_\lambda$ a family of such trees.

Assume the edges of G are colored with at most $\Delta + 1$ colors $c_1, c_2, \dots, c_{\Delta+1}$. The broadcasting algorithm proceeds as follows: split the message of k pieces in λ packets $P_1, P_2, \dots, P_\lambda$, each composed of at most $\lceil \frac{k}{\lambda} \rceil$ pieces. Packet P_i is broadcasted using T_i . Pieces of information are sent in a pipelined fashion, and edges colored c_j are used every $\Delta + 1$ steps, at time $j, j + \Delta + 1, j + 2(\Delta + 1), \dots$

Every $\Delta + 1$ rounds, at least one piece of each packet has been sent by the root. Therefore, after $(\Delta + 1)\lceil \frac{k}{\lambda} \rceil$ rounds, the root has completed its broadcast. In time at most $(\Delta + 1)h$, where h is the maximum depth of the λ spanning trees, the remaining packets still in the networks will reach their destinations. This time additional is independent of k . \square

This result yields two remarks. First, if $\lambda = \Delta$ and if (by chance) the coloring obtained in polynomial time has Δ colors, then the approximation algorithm described in Theorem 2 is asymptotically optimal when k tends to infinity, and the lower bound of Equation 4 is reached (replace λ by Δ , and $\Delta + 1$ by Δ in the statement of Theorem 2). Note also that the term corresponding to $o(k)$ in Theorem 2 is strongly dependent on the maximum depth of λ disjoint spanning trees of a graph G . Unfortunately, minimizing this parameter has been shown to be NP-hard by N. Alon (see [3] and [13]).

3 Gossiping

Any gossiping algorithm under the telephone model can be described by a sequence of matchings, each matching corresponding to a set of calls performed during the same round. Conversely, any sequence of matching M_1, M_2, \dots, M_r yields a communication algorithm described by: at round i , each pairs of extremities of the edges of M_i exchange all the information they are aware of. It is therefore natural to describe gossiping algorithms by sequences of matchings.

Off-line matchings

People were recently interested in studying the structure of the matching sequence, and see what happens if one forces this sequence to be composed of the repetition of a same subsequence:

$$\begin{aligned} M_1, M_2, \dots, M_r &= M_1, \dots, M_s, M_1, \dots, M_s, \dots, M_1, \dots, M_s \\ &= (M_1, M_2, \dots, M_s)^t \end{aligned}$$

This approach gives rise to a class of approximation algorithms that (1) look for some sequence of matchings M_1, M_2, \dots, M_s in a graph G , and (2) apply the corresponding communication algorithm by repetition of this sequence.

Liestman and Richards were the first to consider such an approach (see [31]). They assume that the subsequence M_1, M_2, \dots, M_s is obtained by a coloration of the edges of the graphs with s colors. They study this strategy in paths, cycles, trees and meshes to produce the best adapted gossip algorithm. They also show the following lemma:

Lemma 1 (Liestman and Richards)

If there exists a coloring of a graph G of diameter D with k colors, then $g_{\mathcal{R}}(G) \leq D(k-1) + 1$.

Proof. Each edge is colored. Therefore any edge is used every k rounds, and any piece of information will be routed on a path of length D in less than $k + (k-1)(D-1) = D(k-1) + 1$ rounds. \square

In this model, Labhan, Hedetniemi and Laskar [28] have proposed approximation algorithms for gossiping in trees. In [21, 20], Hromkovic et al. consider the so called *systolic gossiping*. In this case, the choice of the subsequence M_1, M_2, \dots, M_s is enlarged (and is not necessarily chosen among chromatic matchings). In particular two matchings M_i and M_j might have common edges. As Liestman and Richards, they consider particular cases like cycles and trees but did not give a general way of finding the appropriate subsequence M_1, M_2, \dots, M_s for all graphs.

This approach yields the following result that improves by a factor of two the gossiping time resulting of an algorithm that (1) accumulates all the information in some node in at most ΔD rounds, and (2) broadcasts the whole information from that node in at most ΔD rounds (see Equation 2).

Theorem 3 *There exists a polynomial time approximation algorithm \mathcal{A}_χ for gossiping such that, for any graph G of maximum degree Δ and diameter D , $b_{\mathcal{R}}(\mathcal{A}_\chi[G]) \leq \Delta D + 1$.*

Proof. As we said in Section 2, it is possible to polynomially find a coloring of the edges of G with at most $\Delta + 1$ colors. The result then directly follows Lemma 1 \square

Although the approximation algorithm \mathcal{A}_χ of Theorem 3 is quite attractive because of its simplicity, it can be quite inefficient. For instance, let us take a similar example as in [24], and consider the *sparse wheel* SW_n of n vertices and \sqrt{n} rays: a center node x_0 , and $n-1$ other nodes organized in a cycle x_1, x_2, \dots, x_{n-1} ; the center is connected to vertices $x_1, x_{\sqrt{n}+1}, x_{2\sqrt{n}+1}, \dots, x_{(\sqrt{n}-1)\sqrt{n}+1}$. It is easy to check that $b_{\mathcal{R}}(SW_n) \in \Theta(\sqrt{n})$. Therefore, $g_{\mathcal{R}}(SW_n) \in \Theta(\sqrt{n})$. However, there exists a $(\sqrt{n} + 1)$ -coloring of the edges of SW_n such that $g_{\mathcal{R}}(\mathcal{A}_\chi[SW_n]) \in \Omega(n)$: color the edges of the cycle with $O(1)$ colors, and each ray with a different color. Indeed, the information of x_0 needs $\Omega(n)$ rounds to reach a node of the form $x_{k\sqrt{n} + \frac{\sqrt{n}}{2}}$ because the edges of the cycle are used every $\Omega(\sqrt{n})$ rounds. Thus the approximation algorithm of Theorem 3 can be $O(\sqrt{n})$ time worse than the optimal.

In fact, \mathcal{A}_χ is particularly inefficient when a node of high degree implies a large chromatic index and let therefore many edges are often inactive during the gossiping. In the following, we will focus on a more dynamic approach, where the matchings are chosen on-line and not fixed statically before the gossiping algorithm starts.

On-line matching

Our gossiping algorithms will proceed as follows:

Generic gossiping algorithm

```

repeat
  compute a matching  $M$ ;
  exchange messages along the edges of  $M$ ;
until all the vertices get everything

```

This strategy depends on the way the matching M is chosen at each phase. In the following, we will explore two main directions.

In the first case, each edge has a fixed weight at the beginning and weights evolve as the communication algorithm goes on: the matching M_i , chosen at phase i , is a matching of maximum weight and the weights of the edges not in M_i are increased following some given rule so that the next matching M_{i+1} selects other edges.

The other approach consists in computing new weights for the edges after each phase. For instance the weight of an edge can be defined by the number of pieces of information that knows one of its extremities but that the other does not know.

Before analyzing these two approaches, let us recall that a matching of k edges is *maximum* if there does not exist any other matching with more than k edges. A matching is *maximal* if it is not possible to add an edge to that matching. (Of course, a maximal matching is not necessarily maximum.) A matching is of *maximum weight* if the sum of the weights of its edges is maximum. (Again, a matching of maximum weight is not necessarily maximum, but if all the weights are positive, then a matching of maximum weight is maximal.)

In the following, the weight of an edge e is denoted by $w(e)$, and the matching chosen at phase i is denoted by M_i . We say that a generic gossiping algorithm converges if, for any graph G , the gossiping completes in G in a finite time.

Increasing weights

Our first approach is formalized as follows: at the beginning, all the weights are set to 1. During the algorithm, all matchings are chosen to be of maximum weight. It means that M_1 is a maximum matching. The weights evolve by successive multiplication by a fixed constant $\alpha > 1$. At phase $i > 1$, if $e \notin M_{i-1}$ then $w(e) \leftarrow \alpha w(e)$, otherwise $w(e)$ is not modified. We denote by α -IW the corresponding approximation algorithm (for α -Increasing Weights).

Theorem 4 *For any $\alpha > 1$, the approximation algorithm α -IW converges.*

Proof. Let $\alpha > 1$ and let $G = (V, E)$ be any graph. Let us show that there does not exist t such that there is an edge $e \in E$ that is never selected in M_i for any $i > t$.

The proof is by contradiction. Let S be the set of edges e such that there exist t_e satisfying $e \notin M_i, \forall i > t_e$. If $S \neq \emptyset$ then let $t_0 = \max\{t_e, e \in S\}$. For any edge $e \in E$ and any time $t \geq 1$, let

$$\nu(e, t) = |\{i \leq t, e \notin M_i\}|.$$

For any $t \geq t_0$ and any $e \in S$, we have $w(e) = \alpha^{\nu(e, t_0) + t - t_0}$. For any $e \notin S$, we have

$$\forall r \geq 1, \exists t \geq r \mid \nu(e, t) \leq t - r$$

that is e has been selected r times before time t . Let $e \in S$, let

$$r > \log_\alpha |E| - \nu(e, t_0) + t_0$$

and let $t > t_0$ such that

$$\nu(e, t) \leq t - r, \forall e \notin S.$$




Figure 1: Evolution of the weights during the first four steps of a gossiping algorithm that does not converge. (Edges in the matching are represented by dashed lines.)

Then $\forall e \in S, w(e) > |E|\alpha^{t-r} \geq \sum_{f \notin S} w(f)$. Therefore at least an edge of S must be selected in $M_t, t > t_0$, a contradiction.

The proof directly follows since

$$\forall e \in E, \forall t \geq 1, \exists t' \geq t \mid e \in M_{t'}.$$

This implies that any path of G is traversed in a finite time. \square

Note that the fact that α is a scalar is important in the proof of Theorem 4. For instance, consider the case where $\vec{\alpha}$ is a vector of $|E|$ components $\alpha_e, e \in E$. In Figure 1, such a case is illustrated with $\alpha_e = 10$ for all the edges but the central edge e_c , and $\alpha_{e_c} = 2$. Since $\log_2 10 > 2$, the edge e_c is never selected and the information cannot cross from one side of the network to the other side.

The efficiency of the approximation algorithm α -IW is strongly related to the value of α . Consider for instance the path $P_4 = \{x_1, x_2, x_3, x_4\}$ of 4 vertices. The edge (x_2, x_3) will be selected not before round t such that $\alpha^{t-1} > 2$, and then α -IW $[P_4]$ will be very slow if α is too close to 1. Of course, choosing a large α solves the problem in this case. In fact, it is possible to show that if $\alpha \in \Omega(|E|)$, then for any graph $G = (V, E)$ of maximum degree Δ and diameter D , $\mathfrak{g}_{\mathcal{R}}(\alpha$ -IW $[G]) \in O(\Delta D \log n)$. This bound is obtained by estimating the frequency at which an edge can be selected in a matching. However this bound is much worse than the one derived in Theorem 3, and does not reflect the real behavior of this algorithm. For instance, let us consider again the *sparse wheel* SW_n of n vertices and \sqrt{n} rays. If α is large enough, a cycle edge is selected once every $O(1)$ rounds and the information of all the vertices $x_{k\sqrt{n}+1}, \dots, x_{(k+1)\sqrt{n}+1}$ is concentrated in $x_{k\sqrt{n}+1}$ in $O(\sqrt{n})$ rounds. Again, if α is large enough, a ray is selected every $O(\sqrt{n})$ rounds. Therefore the center x_0 is aware of all the information in $O(\sqrt{n})$ rounds. By symmetry, x_0 broadcasts the whole information in $O(\sqrt{n})$ rounds. Hence the gossiping completes in $O(\sqrt{n})$ rounds. Therefore, α -IW can be $\Omega(\sqrt{n})$ times faster than the chromatic gossiping \mathcal{A}_χ .

However, even if computing the matching on-line as in the α -IW approximation algorithm tends to limit the time during when an edge is not used, the fact that this time strongly depends on α makes α -IW very difficult to tune. Moreover, the choice of the matching in both \mathcal{A}_χ and α -IW depends mainly on the topology of the network, and not really on the way the information flows inside the network. We present below another approach that takes into account the knowledge of each node at any time during the gossiping.

Dynamic weights

Our second approach is formalized as follows. Again, all along the algorithm, all matchings are chosen to be of maximum weight. At each phase $i \geq 1$, the weights are computed as follows. For every neighboring nodes x and y , let $\mathcal{I}_{x \rightarrow y}$ be the set of pieces of information that knows x but that does not know y . For any edge $e = (x, y)$, we set $w(e) = |\mathcal{I}_{x \rightarrow y}| + |\mathcal{I}_{y \rightarrow x}|$. The weight $w(e)$ represents the number of pieces of information that will cross e if e is selected in the matching of the current phase. We denote by DW the corresponding approximation algorithm (for Dynamic Weights).

Theorem 5 *The approximation algorithm DW converges.*

Proof. Let $G = (V, E)$ be any connected graph. Let \mathcal{I}_x^t be the set of pieces of information known by $x \in V$ after phase t . (Of course, $|\mathcal{I}_x^t| \leq n, \forall x \in V, \forall t \geq 1$.) If there exists x such that $|\mathcal{I}_x^t| < n$ then $\sum_{x \in V} |\mathcal{I}_x^{t+1}| > \sum_{x \in V} |\mathcal{I}_x^t|$ because otherwise $\mathcal{I}_x^t = \mathcal{I}_{x'}^t, \forall x, x' \in V$. Therefore there exists t such that $|\mathcal{I}_x^t| = n, \forall x \in V$. \square

To give a flavor of the efficiency of the DW approximation algorithm, we have derived the following results:

Property 1 *Let C_n and P_n be respectively the cycle and the path of n vertices.*

$$\begin{aligned} g_{\mathcal{R}}(DW[P_n]) &= g_{\mathcal{R}}(P_n) = \begin{cases} n & \text{if } n \text{ is odd} \\ n - 1 & \text{if } n \text{ is even} \end{cases} \\ g_{\mathcal{R}}(DW[C_n]) &= g_{\mathcal{R}}(C_n) = \lceil \frac{n}{2} \rceil. \end{aligned}$$

Note that the chromatic approximation algorithm does not reach such performances since it directly follows the results in [31] that $g_{\mathcal{R}}(\mathcal{A}_\chi[C_n]) = \frac{3}{4}n + O(1)$.

Proof. The complexity of the DW algorithm is easy to compute for the path of any length, and the cycle of even size. In both cases, two matching alternate and allow to complete a gossiping in the specified time. The cycle $C_n = \{x_0, x_1, \dots, x_n\}$ of odd size must be considered separately. Indeed, at each round, only $\frac{n}{2}$ edges are selected. Assume that the two consecutive edges not selected at the first round are (x_0, x_1) and (x_1, x_2) . At the second round, one of these two edges is selected. Assume without loss of generality that it is (x_0, x_1) . Then (x_1, x_2) and (x_2, x_3) are the two consecutive edges that are not selected at round 2. It is not difficult to see that, at round r , (x_{r-1}, x_r) and (x_r, x_{r+1}) are the two consecutive edges that are not selected. Thus the gossiping time of $DW[C_n]$ is $\lceil \frac{n}{2} \rceil$. (The optimality of the results can be checked easily [10].) \square

Similar results can be obtained for the step-complexity:

Property 2 *Let C_n and P_n be respectively the cycle and the path of $n \geq 4$ vertices.*

$$\begin{aligned} g_{\mathcal{S}}(DW[P_n]) &= \begin{cases} 2n - 2 & \text{if } n \text{ is odd} \\ 2n - 3 & \text{if } n \text{ is even} \end{cases} \\ g_{\mathcal{S}}(DW[C_n]) &= \begin{cases} n & \text{if } n \text{ is odd} \\ n - 1 & \text{if } n \text{ is even} \end{cases} \end{aligned}$$

In all cases (but P_n, n odd), these results are optimal.

Proof. Bermond, Gargano, Rescigno and Vaccaro have shown in [5] that $g_{\mathcal{S}}(P_n) = 2n - 3$. In the case n even, this bound is reached by $DW[P_n]$. Indeed, as we saw in the proof of Property 1, there are two selected matching that alternate during the gossiping: one with $\frac{n}{2}$ edges, and another with $\frac{n}{2} - 1$ edges. These matching allow to gossip in $2n - 3$ steps. Unfortunately, $DW[P_n]$ is not optimal when n is odd and differs from the optimal gossiping by one. Indeed, if n is odd, there are two matching, both with $\frac{n-1}{2}$ edges, that alternates. This yields a gossiping time of $2n - 2$ steps.

Also, Bermond, Gargano, Rescigno and Vaccaro have shown in [5] that, for any Hamiltonian graph G of n vertices,

$$g_{\mathcal{S}}(G) = \begin{cases} n & \text{if } n \text{ is odd} \\ n - 1 & \text{if } n \text{ is even} \end{cases}$$

During the gossiping $DW[C_n]$ on the even cycle, the first round consists of exchanging one packet, and, during the $\frac{n}{2} - 1$ other rounds, two packets are exchanged at each round. That gives $n - 1$ steps in total. When n is odd, the matching chosen at each round is composed of $\frac{n-1}{2}$ edges. Moreover, at each round, two consecutive edges are not selected in the matching. It is easy to see, as in the proof of Property 1, that these pairs of edges “rotate” (say clockwise) around the cycle. This phenomenon slows down the information flow going counter clockwise by 1 step, but does not slow down the information flow going clockwise. This yields a gossiping time of $1 + 2\frac{n-1}{2} = n$ steps. \square

In the following, we experimentally study the DW approximation algorithm for gossiping.

Experimental comparisons

We have experimented the DW approximation algorithm on several graphs as, in particular, the complete graph on n vertices K_n , the d -dimensional hypercube Q_d , the d -dimensional *cube-connected-cycle* CCC_d [35], the d -dimensional shuffle-exchange SE_d [30], the two dimensional meshes and tori ($P_p \times P_q$ and $C_p \times C_q$), and the d -dimensional star-graph S_d [1]. All these graphs are often considered as efficient interconnection networks for parallel computers. We comment the results of our experiments below. We compare these results with the best ones known in the literature [14].

Figures 2 and 3 present results on the round-complexity as a function of characteristic parameters of the graph (number of nodes, or dimension). Figure 4 present results on the step-complexity, again as a function of characteristic parameters of the graphs.

Note that the execution of the DW approximation algorithm depends on the chosen matching algorithm: two different matching algorithms may produce different results using the DW approximation algorithm.

The complete graph K_n

Figures 2(a) and (b), and Figure 4(a) give the results obtained on complete graph. The DW algorithm performs quite well. In particular, it is easy to see that, for $n = 2^k$, the DW approximation algorithm is optimal:

Property 3 *Let K_{2^k} be the complete graph on 2^k vertices,*

$$\begin{aligned} g_{\mathcal{R}}(DW[K_{2^k}]) &= g_{\mathcal{R}}(K_{2^k}) = k \\ g_{\mathcal{S}}(DW[K_{2^k}]) &= g_{\mathcal{S}}(K_{2^k}) = 2^k - 1 \end{aligned}$$

The first even value of n for which the $DW[K_n]$ is not round-optimal is 14 (Knodel [23] has shown that for n even, it is possible to gossip in $\lceil \log_2 n \rceil$ rounds). Our experiment shows that the DW algorithm performs well on K_{14} during the 3 first rounds (that is the global amount of information known by each vertex double at each round). Unfortunately, the way in which the information flows in the fourth round is not enough to complete the gossiping in K_{14} . For n odd, the general structure of the $DW[K_n]$ consists in a gossiping among $n - 1$ vertices plus a broadcast from the non informed vertex to all the other. This doubles at least the usual gossiping time of $\lceil \log_2 n \rceil + 1$ for n odd [23] (this is what happened for $n = 9$).

The hypercube Q_d

Figure 2(c) gives the experimental results obtained on the d -dimensional hypercube Q_d . These results strongly depends on the labeling of the vertices. For instance, for the “usual” labeling \mathcal{L} (two nodes are linked by an edge if and only if their binary representations differ in exactly one bit), the DW gossiping algorithm has performed optimally in term of both round and step. However, Figure 2(c) shows that, for other labelings, the DW gossiping algorithm may be less efficient, even with a slight modification of the labels: $\mathcal{L}'(x) = \mathcal{L}(x) + 13 \bmod n$ and $\mathcal{L}''(x) = \mathcal{L}(x) + 7 \bmod n$. This property must be kept in mind when implementing the DW algorithm for a given graph: the labeling of the vertices has a great influence on its performances.

Meshes $P_p \times P_q$

Figures 2(d) and (e) present experimental results on two dimensional meshes. Compared to the optimal, the performances (in terms of rounds) are quite good: always less than the optimal time plus 2 in our experiments.

Torus $C_p \times C_q$

Figures 3(a) and (b), and Figure 4(b) present experimental results on two dimensional tori. In our experiments (that may have been helped by the usual labeling of the torus), we have always obtained optimal results (in term of both number of rounds and steps) when p and q are both even. In the general case, the round-complexity of $DW[C_p \times C_q]$ is always between the known lower and upper bounds of the round-complexity of the gossiping in $C_p \times C_q$.

Cube-connected-cycle CCC_d

Figures 2(f) and 4(c) present experimental results on d -dimensional CCC . The round-complexity of $DW[CCC_d]$ is about the same as the best known upper bound for this graph (it can even perform faster in some cases). The step-complexity of $DW[CCC_d]$ is close to the general lower bound $n - 1$ for even order, and n for odd order.

Star-graph S_d

Figures 3(c) and 4(d) present experimental results on star-graphs. The round-complexity of $DW[S_d]$, $d = 3, 4, 5$, is much closer to the optimal than the best known upper bound. The re-

sult on the step complexity are also quite good. In fact, they are optimal for $d = 3, 4, 5$.

Shuffle-exchange SE_d

Figure 3(d) presents experimental results on shuffle-exchange. Gossiping in this graph is still an open problem. However our results shows that gossiping and broadcasting in the shuffle-exchange must have similar round-complexities. Indeed, the round-complexity of $DW[SE_d]$ is quite close to the round-complexity of broadcasting in SE_d .

4 Extensions

Communications in store-and-forward routed parallel computers

In a store-and-forward routed parallel computer, a message proceeds along a path $P = \{x_0, x_1, \dots, x_\ell\}$ of length ℓ in ℓ successive phases. At phase i , $1 \leq i \leq \ell$, the message is sent by x_{i-1} to x_i , and stored in the memory of x_i . At the next phase, x_i will forward the message to x_{i+1} .

Many experiments [6, 12] have shown that the time of a neighbor-to-neighbor communication in many computers can be modeled by $\beta + L\tau$ where β is the *start-up* time necessary to set up the communication, and $L\tau$ is the *propagation* time proportional to the size L of the message ($\frac{1}{\tau}$ is the *bandwidth* of each communication link). Though it can be more difficult to handle (see [14]), this model reflects, in some sense, both step-complexity (by looking at the dominating term in τ) and round-complexity (by looking at the dominating term in β).

In fact, given any approximation algorithm \mathcal{A} for gossiping, one can hope that its completion time in a store-and-forward routed parallel computer of topology G will be of the form

$$g_{\mathcal{R}}(\mathcal{A}[G])\beta + g_{\mathcal{S}}(\mathcal{A}[G])\tau.$$

However, there exist graphs G for which the store-and-forward gossiping time is larger than $g_{\mathcal{R}}(G)\beta + g_{\mathcal{S}}(G)\tau$. This is typically the case when the round-optimal gossiping algorithm on G is different from the step-optimal gossiping algorithm in G . The cycle C_n , n odd, is an example of such a graph: $g_{\mathcal{R}}(C_n) = \frac{n+1}{2}$, $g_{\mathcal{S}}(C_n) = n - 1$, but we do not know if it is possible to gossip in $\frac{n+1}{2}\beta + (n - 1)\tau$ (see [14]). However, choosing an efficient approximation algorithm depends on the tradeoff between β and τ . The DW gossiping algorithm based on the differences of knowledge between neighbors of the networks is quite appropriate to that problem since it tends to minimize simultaneously the total start-up times and the total propagation times.

In the broadcasting case, the dominating term in the complexity of the algorithm described in Theorem 2 corresponds to the propagation time of the broadcasting algorithm. However, the $o(k)$ term cannot be neglected in store-and-forward routed networks since it depends on the start-up time, and β is usually much larger than τ . The adapted solution is *pipelining* [22, 39] where, instead on sending the information contained in each packet piece after piece, pieces are first grouped in sub-packets. Then the sub-packets play the role of the pieces, being sent one after the other through the network. Computing the number of pieces that must form each sub-packet is critical and is described in detail in [14].

Telegraph model

In the telephone model (also called 2-way mode), communication links are *full-duplex*, that is they can be used simultaneously in both directions. In the *telegraph* model (also called 1-way mode), links are *half-duplex*, that is they are bidirectional but they cannot be used simultaneously in both directions.

The telegraph model appears to be much more tricky than the telephone model. For instance, in the telegraph model, for any graph G of order n ,

$$2 + \lceil \log_\rho \frac{n}{2} \rceil \leq g_{\mathcal{R}}(G)$$

where $\rho = \frac{1+\sqrt{5}}{2}$ [7, 9, 29, 40]. This bound must be compared to Equation 1. To give an idea of the difficulty, the complexity of gossiping in the complete graph K_n is not known for all the values of n , and the complexity of gossiping in hypercube is totally unknown (it belongs to the interval $[1.44 \log_2 n, 1.88 \log_2 n]$, see [25]).

However, our approach applies also in the telegraph model. For instance, Theorem 1 directly applies. Theorem 2 does not directly apply because the proof is based on trees that use edges in both directions simultaneously. Nevertheless, this theorem can be adapted to the half-duplex constraint by using $\lfloor \frac{\lambda}{2} \rfloor$ edge-disjoint spanning trees (the existence of such trees has been shown by Kundu [26]).

For gossiping, the same approximation algorithm as for the telephone model applies. The rule is that for any edge $e = (x, y)$, $w(e) = \max\{|\mathcal{I}_{x \rightarrow y}|, |\mathcal{I}_{y \rightarrow x}|\}$ and the information goes from x to y if $|\mathcal{I}_{x \rightarrow y}| > |\mathcal{I}_{y \rightarrow x}|$, and from y to x otherwise.

The experimental results show that this approach is efficient. We experiment the DW algorithm in the telegraph model on the complete graph, the hypercube, and the cycle of even size (quite a few results are known about other topologies [14]), and consider only the round-complexity (the step-complexity of gossiping has not been yet investigated for the usual topologies in the half-duplex model [14]). Figure 5(a) presents results obtained on the complete graph. Even if the number of rounds is larger than the optimal (about $1.44 \log_2 n$ [7, 9, 29, 40]), it stays not larger than a constant factor times this optimal. Same remarks hold for the hypercube (Figure 5(b)). In the case of the cycle of even size, the number of rounds of the DW algorithm is n whereas the optimal gossiping algorithm performs in $\frac{n}{2} + \lceil \sqrt{2n} \rceil - 1$ rounds [18] (Figure 5(c)).


References

- [1] S. Akers, D. Harel, and B. Krishnamurthy. The star graph: an attractive alternative to the n -cube. Proceedings of the International Conference on Parallel Processing, Pennsylvania, pages 393-400, 1987.
- [2] C. Berge. *Graphes et Hypergraphes*. Dunod, Paris, 1973.
- [3] J.-C. Bermond and P. Fraigniaud. Broadcasting and NP-Completeness. *Graph Theory Notes of New York*, XXII:8–14, 1992.
- [4] J.-C. Bermond, P. Fraigniaud, and J. Peters. Antepenultimate broadcasting. *Networks*, (To appear). Technical Report CMPT TR 92-03, SFU, Vancouver.


- [5] J-C. Bermond, L. Gargano, A. Rescigno, and U. Vaccaro. Optimal gossiping by short messages. Dipartimento di Informatica ed Applicazioni, Salerno, Italy.
- [6] L. Bomans and D. Roose. Communication benchmarks for the iPSC/2. In North-Holland, editor, *Hypercube and Distributed Computers*, pages 311–328, 1989.
- [7] G. Cybenko, D.W. Krumme, and K.N. Venkataraman. Gossiping in minimum time. *SIAM Journal on Computing*, 21(1):111–139, 1992.
- [8] J. Edmonds. Edge-disjoint branchings, combinatorial algorithms. In R.Rustin, editor, *Combinatorial Algorithms 91-96*. Algorithmics press, New York, 1972.
- [9] S. Even and B. Monien. On the number of rounds necessary to disseminate information. ACM Symposium on Parallel Algorithms and Architectures, Santa Fe, June 1989.
- [10] A. Farley and A. Proskurowski. Gossiping in grid graphs. *Journal of Combinatorics Information and Systems Science*, 5:161–172, 1980.
- [11] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. In *International Symposium SIGAL'90*, volume 450 of *Lecture Notes in Computer Science*, pages 128–137, 1990.
- [12] P. Fraigniaud. Performance analysis of broadcasting in hypercubes with restricted communication capabilities. *Journal of Parallel and Distributed Computing*, 16(1):15–26, 1992.
- [13] P. Fraigniaud and C. Laforest. Disjoint spanning trees of small depth. In *Parallel Computing: Trends and Applications*. Elsevier Science, 1994.
- [14] P. Fraigniaud and E. Lazard. Methods and Problems of Communication in Usual Networks. *Discrete Applied Mathematics*, 53:79–133, 1994.
- [15] P. Fraigniaud and J. Peters. Minimum linear gossip graphs and maximal linear (Δ, k) -gossip graphs. Technical report TR 94-05, Simon Fraser University, 1994.
- [16] M. Garey and D. Johnson. *Computers and Intractability: a guide to theory of NP-completeness*. Freeman, San Fransisco, 1979.
- [17] S.M. Hedetniemi, S.T. Hedetniemi, and A.L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18:319–349, 1986.
- [18] J. Hromkovic, C. Jeschke, and B. Monien. Optimal algorithms for dissemination of information in some interconnection networks. Technical Report 71, Universitat Paderborn, 1990.
- [19] J. Hromkovic, R. Klasing, B. Monien, and R. Peine. Dissemination of information in interconnection networks (broadcasting and gossiping). Technical report, University of Paderborn, Germany, 1993.
- [20] J. Hromkovic, R. Klasing, D. Pardubska, W. Unger, J. Waczulic, and H. Wagener. Effective systolic algorithms for gossiping in cycles and two-dimensional grids. In *10-th International Conference on Fundamentals of Computation Theory (FCT95)*, Lecture Notes in Computer Science. Springer-Verlag, (to appear).

- [21] J. Hromkovic, R. Klasing, D. Pardubska, W. Unger, and H. Wagerer. The complexity of systolic dissemination of information in interconnection networks. In *4-th Scandinavian Workshop on Algorithm Theory (SWAT94)*, number 824 in Lecture Notes in Computer Science, pages 219–230. Springer-Verlag, 1994.
- [22] S.L. Johnsson and Ching-Tien Ho. Optimum broadcasting and personalized communication in hypercubes. *IEEE Transaction on Computers*, 38(9):1249–1268, 1989.
- [23] W. Knodel. New gossips and telephones. *Discrete Mathematics*, 13:95, 1975.
- [24] G. Kortsarz and D. Peleg. Approximation algorithms for minimum time broadcast. In Springer-Verlag, editor, *Proceedings of the 1992 Israel Symposium on Theory of Computer Science*, number 601 in Lecture Notes in Computer Science, 1992.
- [25] D.J. Krumme. Fast gossiping for the hypercube. *SIAM Journal on Computing*, 21(2):365–380, 1992.
- [26] S. Kundu. Bounds on the number of disjoint spanning trees. *Journal of Combinatorial Theory*, 17:199–203, 1974.
- [27] R. Labahn. Some minimum gossip graphs. *Networks*, 23:333–341, 1993.
- [28] R. Labahn, S. Hedetniemi, and R. Laskar. Periodic gossiping on trees. *Discrete Applied Mathematics*, 53:235–245, 1994.
- [29] R. Labahn and I. Warnke. Quick gossiping by telegraphs. *Discrete Mathematics*, 126:421–424, 1994.
- [30] T. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*. Morgan Kaufmann, 1992.
- [31] A. Liestman and D. Richards. Network communication in edge-colored graphs: gossiping. *IEEE Transaction on Parallel and Distributed Systems*, 4(4):438–445, 1993.
- [32] L. Lovasz. On two minimax theorems in graph theory. *J. Combinatorial Theory*, B(21):96–103, 1976.
- [33] W. McCuaig. A simple proof of menger’s theorem. *Journal of Graph Theory*, 8:427–429, 1984.
- [34] M. Middendorf. Minimum broadcast time is NP-complete for 3-regular planar graphs and deadline 2. *Information Processing Letters*, 46:281–287, 1993.
- [35] F. Preparata and J. Vuillemin. The cube connected cycles: a versatile network for parallel computation. *Communications of the ACM*, 24(5):300–309, 1981.
- [36] R. Ravi. Rapid rumor ramification: approximating the minimum broadcast time. In IEEE Computer Society Press, editor, *35-th Symposium on Foundations of Computer Science*, pages 202–213, 1994.
- [37] P. Scheuermann and G. Wu. Heuristic algorithms for broadcasting in point-to-point computer networks. *IEEE Transaction on Computers*, C-33(9):804–811, 1984.
- [38] P. Slater, E Cockayne, and S Hedetniemi. Information dissemination in trees. *SIAM Journal on Computing*, 10(4):692–701, 1981.


- [39] Q.F. Stout and B. Wagar. Intensive hypercube communication, prearranged communication in link-bound machines. *Journal of Parallel and Distributed Computing*, 10:167–181, 1990.
- [40] V.S. Sunderam and P. Winkler. Fast information sharing in a distributed system. *Discrete Applied Mathematics*, 42:75–86, 1993.
- [41] J. Ventura and X. Weng. A new method for constructing minimal broadcast networks. *Networks*, 23:481–497, 1993.




(c) Hypercube



(d) Meshes



(e) Meshes



(f) Cube Connected Cycle




Figure 2: Experimental round-complexity of the DW approximation algorithm (to be continued)

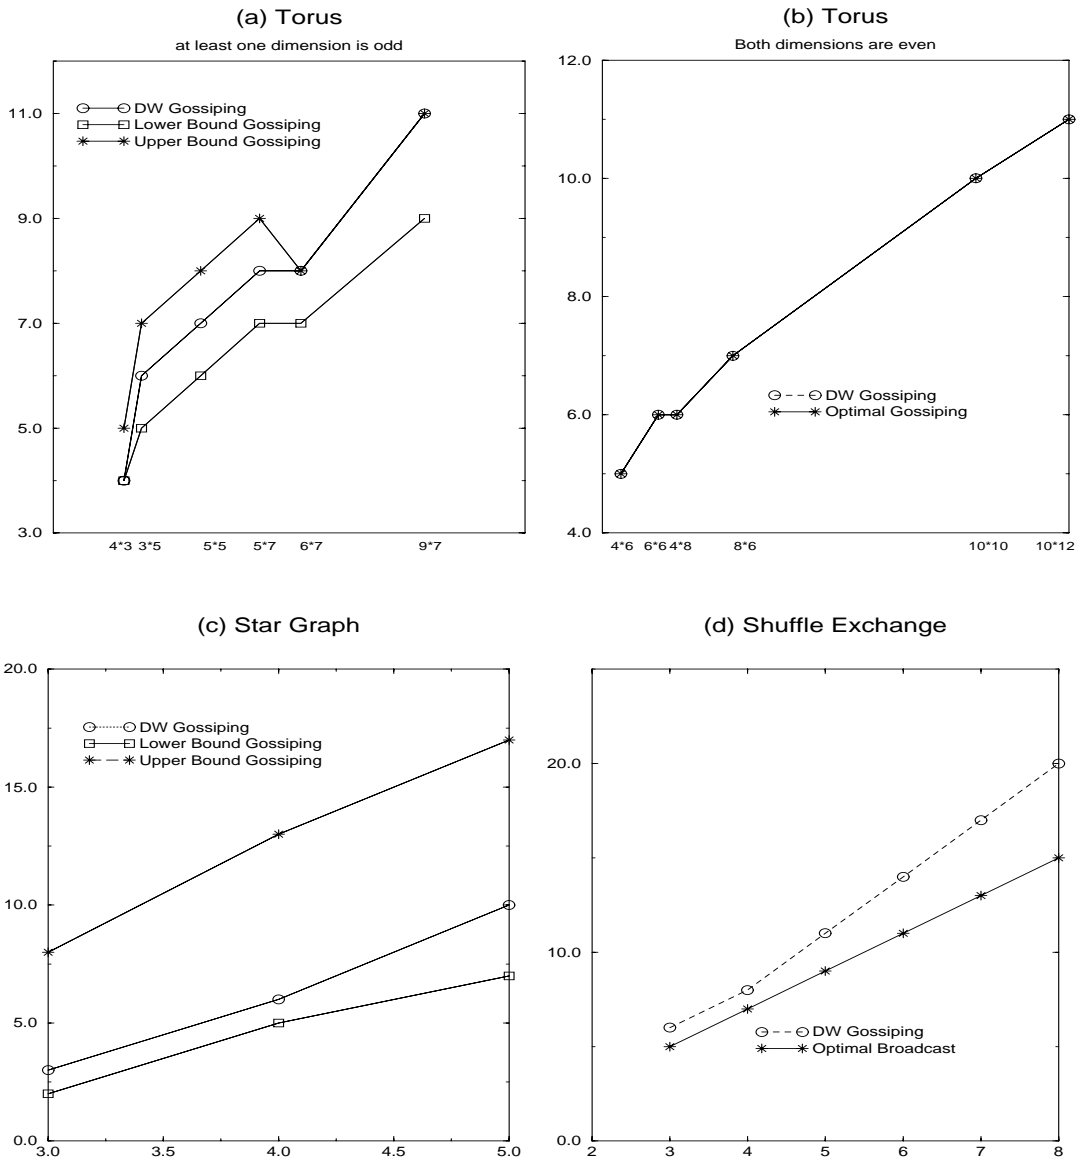


Figure 3: (continued) Experimental round-complexity of the DW approximation algorithm




Figure 4: Experimental step-complexity of the DW approximation algorithm




Figure 5: Experimental round-complexity of the DW approximation algorithm in telegraph model.