

# Gossiping and Broadcasting versus Computing Functions in Networks

Martin Dietzfelbinger  
Fachbereich Informatik, Lehrstuhl II  
Universität Dortmund  
D-44221 Dortmund  
Germany

email: [dietzf@ls2.informatik.uni-dortmund.de](mailto:dietzf@ls2.informatik.uni-dortmund.de)

September 25, 1996

## Abstract

The fundamental assumption in the classical theory of dissemination of information in interconnection networks (gossiping and broadcasting) is that atomic pieces of information are communicated. Messages in the network always consist of a set of such pieces. The purpose of this paper is to make results from this theory applicable to the problem of computing functions in synchronous networks built from processors that are connected by (bidirectional) links, across which arbitrary messages may be sent. We show that, under suitable assumptions about the way processors may communicate, computing an  $n$ -ary function that has a “critical input” and distributing the result to all processors on an  $n$ -processor network takes exactly as long as performing gossiping in the graph that defines the network structure. A similar relation exists between the problem of computing functions with the output appearing at only one processor and the complexity of broadcasting in the corresponding graph.

The restrictions we have to impose on the way processors communicate are, e. g., in the case of 1-way communication across the links, that in one step, a processor may either send one message or receive one message or do nothing, and that

- (\*) “*Predictable reception*”: a processor must know at the beginning of the step whether it is to receive a message across one of its links or not.

As an example, we consider the problem of computing the OR of  $n$  bits and distributing the result, which is essentially the same as solving a certain type of synchronization problem. Our results show that on a network that

obeys the aforementioned restrictions, for this task it is optimal to use the communication pattern of a gossiping scheme, while for computing the OR of  $n$  bits with the result appearing at one processor it is optimal to use the pattern given by a broadcasting, or, equivalently, by an accumulation scheme. It is obvious that gossiping and accumulation schemes can be used to compute functions, as long as the length of messages is unrestricted and computation is given for free; however, it seems that until now there was no formal study of the opposite direction. All results are sharp, that means, no constant factors are involved.

We further study the complexity of broadcasting one bit in a synchronous network that does not necessarily satisfy restriction (\*). In fact, the only restriction is that in one step a processor can send only one message. We show that in this model the number of steps needed for broadcasting a bit equals the number of rounds needed by so-called 2-broadcasting protocols, in which in one step a processor can inform two of its neighbors of the piece of information that is to be broadcasted. The method can be applied to parallel random-access machines (PRAMs) and distributed memory machines (DMMs) with the ARBITRARY access resolution rule to obtain lower bounds for broadcasting. The result also implies that in networks without restriction (\*) computing the OR function and distributing the result can be done faster than gossiping, but at most by a factor of 4.

## 1 Introduction

The purpose of this paper is to demonstrate that the well-established theory of gossiping and broadcasting in interconnection networks can directly be applied to obtain lower bounds for algorithms that compute functions in synchronous networks consisting of processors connected by bidirectional links. In this introductory section, we informally discuss the relevant structures and describe the results. Formal definitions will be given in Section 2.

### 1.1 Gossiping and broadcasting

In *gossiping* theory one deals with the following basic situation: each node of a network initially has an atomic piece of information; the purpose of a gossiping scheme is to distribute all the information to all the nodes. For this, in rounds, the nodes send each other messages consisting of an arbitrary number of pieces of information. The standard restriction is that in one round a node can communicate with only one of its neighbors. One distinguishes 1-way (or half-duplex) mode, where in a round information can be sent through a link in only one direction, and 2-way (or full-duplex) mode, where in a round two nodes may exchange all their information through a link that connects them. The most intensively studied efficiency criterion in this theory is the number of rounds needed for

disseminating all pieces of information to every node. The *broadcasting* problem is similar, excepting that only one piece of information, initially located at one node, is to be spread to all others. The *accumulation* problem is the converse of the broadcast problem: the aim is to collect all pieces of information initially located at the single nodes in one distinguished node. For an account of the history of the problem area and the intensive research devoted to it see the survey [17], the more recent surveys [14, 20], or other articles in the special issue of *Discrete Applied Mathematics* (vol. **53**, 1994). In the fundamental paper [8], which also contains many lower bound arguments, the relevance of the gossip model for real multiprocessor systems is discussed.

## 1.2 Computing functions in processor networks

We consider the problem of computing functions in networks. The computational model is a network of  $n$  processors,  $P_1, \dots, P_n$ , that are connected by bidirectional links, according to a network graph  $G = (V, E)$ . The processor network is to compute an  $n$ -ary function  $f : A_1 \times \dots \times A_n \rightarrow A$ , for arbitrary sets  $A_1, \dots, A_n$ , and  $A$ , in the following way. Initially, processor  $P_i$  knows the  $i$ th component  $a_i$  of the input  $a = (a_1, \dots, a_n)$ ; at the end, all processors know the result  $f(a)$  (“global output”). (For example, for realizing a synchronization barrier in the network, i. e., in order to find out whether all processors have finished certain subtasks from a set  $\mathcal{T}_1$  given to them and inform them whether to proceed to another set  $\mathcal{T}_2$  of subtasks that requires that all activities for  $\mathcal{T}_1$  are finished, the Boolean function OR must be computed with global output.) We will also consider the situation where only one processor has to know the result (“local output”). The processors work synchronously in lock-step, i. e., in global steps  $t = 1, \dots, T$ . In one step, a processor may communicate with at most one of its neighbors. Different models are obtained by allowing only 1-way traffic on a link in a step (half-duplex mode) or 2-way-traffic (full-duplex-mode). There is one essential restriction, which turns out to be necessary for our results to be true, which can be phrased informally as follows.

- (\*) “*Predictable reception*”: a processor must know at the beginning of a step whether it is to receive a message across one of its links in this step or not.

That means, a processor must not, in step  $t$ , wait for a message that may or may not arrive and proceed in different ways when a message arrives and otherwise. This restriction makes it impossible that information is transferred by *not* sending a message in a certain round. The relevance of such a possibility in the context of computations on the parallel random-access machine (PRAM) was observed in [7], and investigated in more depth in [3]. Apart from this restriction, the model is quite general, e. g., a processor may wait for a message from any one of its neighbors without specifying the sender or may make a message available to

any neighboring processor that wants to receive it. As will be discussed later, algorithms that observe restriction (\*) are suitable to be executed on asynchronous networks as well (Remark 2.8).

### 1.3 Results

It is almost obvious that if a network has a  $T$ -round gossiping scheme then it can solve the synchronization problem in  $T$  steps with 1-bit messages and trivial computation, and, under the assumption that messages may be arbitrarily long and computation is for free, it can compute any function with global output in  $T$  steps. The main result of this paper essentially is that this is best possible. More specifically, we show the following for processor networks with “predictable reception”.

- If an  $n$ -ary function  $f$  that has a *critical input*  $a$ , i. e., an input with the property that each of the  $n$  components of the input can be changed in such a way that the value of the function changes, can be computed on a processor network  $G$  with global output in  $T$  steps, then  $G$  has a  $T$ -round gossiping scheme.
- If an  $n$ -ary function  $f$  that has a critical input can be computed on a processor network  $G$  in  $T$  steps with local output at processor  $P_{i_0}$ , or a function  $f$  that depends on component  $a_{i_0}$  can be computed on  $G$  in  $T$  steps with global output, then  $G$  has a  $T$ -round broadcasting scheme.

These results hold if both in the processor network and in the gossiping network 1-way resp. 2-way communication is assumed. Of course, the idea of the proof is to show that any communication pattern that the network produces on the critical input in fact essentially is a gossiping scheme respectively a broadcasting scheme. Although this is quite intuitive, the result does not seem to be obvious, due to the fact that the communication in the processor network is not oblivious, i. e., communication patterns for different inputs  $a$  may be different, in contrast to gossiping schemes, which have one fixed communication pattern. The role played by the restriction “predictable reception” becomes apparent only in the proof. In contrast, e. g. in [8] it is stated, but not proved formally, that lower bounds for gossiping carry over directly to the synchronization problem or, even more generally, to computing any multiple-output function in which all output components depend on all input components, like matrix inversion or computing the discrete Fourier transform of a vector, on arbitrary processor networks. In Sections 4 and 5 we will consider networks without restriction (\*). As for computing functions with global or local output, we shall see that the possibility of using different communication patterns on different inputs actually makes it possible to compute the OR faster than by a gossiping respectively an accumulation scheme. Still, it can be shown that in no network the speedup can be more than

a factor of 4. In Section 5, we will exactly characterize the complexity of broadcasting a bit in a synchronous network with very general communication rules in terms of a variant of broadcasting schemes in the gossiping type model. The only remaining restriction is that in one step a processor may send only one message that may not be duplicated by the communication system.

## 1.4 Applications

The consequences of our results for networks with “predictable reception” are threefold. First, as corollaries we obtain a host of lower-bound results for computing functions in networks of different topologies with algorithms that obey restriction (\*), since all lower bounds proved for gossiping and broadcasting carry over. As is common in gossiping theory, many of these bounds are tight. E. g., we obtain the following.

- Computing the OR in 1-way mode with global output on a complete network of  $n$  processors takes  $1.44\dots \log n \pm O(1)$  steps<sup>1</sup> [8, 12, 23, 27].
- Computing the OR in 1-way mode with global output in a ring of  $n$  processors takes time  $n/2 + \sqrt{2n} \pm O(1)$  [19, 8].
- Computing the OR in 1-way or 2-way mode with global or local output on a butterfly network with  $n = k \cdot 2^k$  nodes takes at least  $1.7417k$  steps [22].
- Computing the OR in a complete  $k$ -ary tree of depth  $d$  with local output at the root takes exactly  $kd$  steps; with global output,  $2kd$  steps are necessary and sufficient [20].

Second, it follows that the so-called minimum broadcast subgraphs (cf. [17]) of a network or the trees and schedules used to achieve optimal broadcast times are also the optimal communication pattern for computing functions with a critical input on the corresponding networks, with local output, under restriction (\*). The analogous statement for minimum gossip graphs (cf. [17]) also holds. Third, it has been an objection to upper bounds described in gossiping theory that the model used is somewhat unrealistic in that it allows to send messages containing an arbitrary number of atomic pieces of information from one node to another in one round. Thus, claims to the effect that a lower bound was optimal by giving the matching upper bound were not completely satisfying. (See [1] for an investigation of gossiping with bounded packet sizes.) In our network model, the corresponding upper and lower bounds hold for the problem of computing the OR function, which can be solved by sending messages consisting of one bit. Thus, both upper and lower bounds hold for a quite realistic model.

---

<sup>1</sup>All logarithms are to the base 2, unless noted otherwise.

The model without “predictable reception” is general enough to make it possible to draw conclusions about the complexity of broadcasting a bit in some other models, viz., PRAMs and distributed memory machines (DMMs), with communication that obeys the EXCLUSIVE READ rule. Not surprisingly, the mode of writing is not relevant in this context. (See [13] for information on the complexity of PRAM computations, and [11] for a specification of the DMM model with various access conflict resolution rules.)

## 1.5 Related work

In [8], the relevance of results for the gossiping problem for real multiprocessor systems was discussed in depth, but informally, i. e., without making the model for the multiprocessor system explicit. To the best of the knowledge of the author, the problem considered here has not been studied before on a comparable technical level, with the exception of work done by Belting in his diploma thesis [4], who used a combination of the lower-bound method for the gossip problem on the complete network from [8, 12, 23, 27] with the degree technique for proving lower bounds for CREW PRAMs from [10] to show that computing the OR function in networks that obey restriction (\*) and the concurrent-write property (it is forbidden that in one step, more than one message are sent to one processor) takes exactly as long as gossiping in this network. The proof method used in the present paper is completely different. Our technique is, however, related to the method introduced in [30] for analyzing computations on concurrent-read concurrent-write PRAMs with bounded communication width. The method for analyzing broadcasting algorithms in networks with unrestricted communication is a new and more general formalization of the idea of analyzing PRAM computations by keeping track of those cells and processors that are “affected” by some input bit, which has been used in [3, 7].

In [2], the problem of computing and distributing the value of a commutative and associative function in a complete network is considered, with the variant of communication mode that allows a processor to send one message *and* receive one message in one step. The lower bound  $\lceil \log n \rceil$  claimed in that paper for the broadcasting problem is called “obvious”, and it seems that a fixed communication pattern is assumed. Similarly, the lower bound mentioned in [5] for the more general “ $k$ -port” model seem to be based on this assumption.

Also, it should be mentioned that in the context of *asynchronous* communication in networks, Tel studied a family of algorithms that were all recognized to be equivalent (called “normal algorithms” in [28] and “wave algorithms” in [29]). They would correspond to the problem of computing functions that have a critical input with “local output”. However, because of the absence of a notion of time in these models, Tel’s results do not have direct applications in our setting.

## 1.6 Structure of paper

The structure of the rest of this paper is as follows. In Section 2 we recall the relevant notions from gossiping theory and formally describe the network model. Section 3 contains the proof of the main theorem; Section 4 contains variants of the main result, in particular, extensions to functions that do not have a critical input, to multiple-valued functions, and to networks without restriction (\*). Finally, in Section 5 we consider the problem of broadcasting one bit in unrestricted networks, and mention applications to other parallel computational models (exclusive-read PRAMs, and distributed memory machines (DMMs) with the ARBITRARY read and write conflict resolution rule).

## 2 Preliminaries

In this section, we first recall some basic definitions and facts from gossiping and broadcasting theory. Next, we give a precise description of the type of processor networks we will consider. For completeness, we also recall the definition of some well-known complexity measures for  $n$ -ary functions.

### 2.1 A formal view of gossiping and broadcasting

For the rigorous definitions of gossiping, broadcasting, and accumulation schemes we partly follow the survey paper [20]. (See that paper as well as [14] for more information on the subject, in particular for lower and upper bounds for specific networks, and for variations of the models.) Throughout this paper, by a graph  $G = (V, E)$  we mean an undirected graph without loops or multiple edges, where, without loss of generality,  $V = \{1, \dots, n\}$  for some  $n \in \mathbf{N}$ . An undirected edge between nodes  $i$  and  $j$  is denoted by  $\{i, j\}$ , a directed edge from  $i$  to  $j$  by  $(i, j)$ .

**Definition 2.1** *Let  $G = (V, E)$  be a graph.*

- (a) *A directed matching in  $G$  is a set  $M \subseteq V \times V - \{(i, i) \mid i \in V\}$  such that
  - (i) *if  $(i, j) \in M$  then  $\{i, j\} \in E$ ;*
  - (ii) *if  $(i, j)$  and  $(i', j')$  are distinct elements of  $M$  then  $\{i, j\} \cap \{i', j'\} = \emptyset$ .**
- (b) *An undirected matching of  $G$  is a set  $M \subseteq E$  consisting of node-disjoint edges.*

**Definition 2.2** *A communication protocol for a graph  $G = (V, E)$  in 1-way [resp. 2-way] mode is a sequence  $\mathcal{M} = (M_1, \dots, M_T)$  of directed [resp. undirected] matchings of  $G$ . The number  $T$  is called the number of rounds of protocol  $\mathcal{M}$ . Inductively, we define a sequence  $\mathcal{K}(\mathcal{M}) = (K_0, K_1, \dots, K_T)$  of mappings  $K_t: V \rightarrow \mathcal{P}(V)$ , where  $\mathcal{P}(V)$  denotes the power set of  $V$ , as follows.*

- (i)  $K_0(i) = \{i\}$ , for  $i \in V$ ;
- (ii) For  $1 \leq t \leq T$ ,  $i \in V$ :

$$K_t(i) = \begin{cases} K_{t-1}(l) \cup K_{t-1}(i) & \text{for the (unique) } l \in V \text{ such that} \\ & (l, i) \in M_t \quad [\text{resp. } \{l, i\} \in M_t], \\ & \text{if such an } l \text{ exists.} \\ K_{t-1}(i), & \text{otherwise.} \end{cases}$$

This definition is the straightforward formalization of the description of the flow of pieces of information in the graph that is induced by the communication protocol  $\mathcal{M}$  in the obvious way: in step  $t$ , processor  $P_i$  sends a message consisting of all pieces of information it has collected so far to  $P_j$  if and only if  $(i, j) \in M_t$  [resp.  $\{i, j\} \in M_t$ ]. The statement “ $j \in K_t(i)$ ” is interpreted as “by the end of round  $t$ , processor  $P_i$  has received (‘knows’) the piece of information that was initially located at  $P_j$ ”.

**Definition 2.3** *Let  $G$  be a graph,  $\mathcal{M}$  a communication protocol for  $G$  in 1-way [2-way] mode. Let  $\mathcal{K}(\mathcal{M})$  be the sequence associated with  $\mathcal{M}$  as in the previous definition.*

- (a)  $\mathcal{M}$  is a gossip protocol for  $G$  in 1-way [2-way] mode if  $K_T(i) = V$  for all  $i \in V$ ;
- (b)  $\mathcal{M}$  is a broadcast protocol for  $G$  in 1-way mode with source node  $i_0 \in V$  if  $i_0 \in K_T(i)$  for all  $i \in V$ ;
- (c)  $\mathcal{M}$  is an accumulation protocol for  $G$  in 1-way mode with target node  $i_0 \in V$  if  $K_T(i_0) = V$ .

**Definition 2.4** *Let  $G = (V, E)$  be a graph.*

- (a) *The 1-way [2-way] gossip complexity  $r(G)$  [ $r_2(G)$ ] of  $G$  is the minimum  $T$  such that there is a gossip protocol for  $G$  in 1-way [2-way] mode with  $T$  rounds.*
- (b) *The (1-way) broadcast complexity  $b(G, i_0)$  and the (1-way) accumulation complexity  $a(G, i_0)$  are defined analogously.*

One can also define broadcast and accumulation complexity in 2-way mode. However, it is easily seen that these do not differ from their 1-way counterparts. In fact, we even have the following. (For the simple proofs of these claims, see, e. g., [20].)

**Fact 2.5**  $a(G, i_0) = b(G, i_0)$  for all  $G = (V, E)$  and  $i_0 \in V$ . ■



## 2.2 Specification of the network model

In this section, we describe the machine model that will be the basis of our considerations. The description will be slightly informal; a fully rigorous definition (involving abstract state sets, transition functions, read-address, write-address, and write-value functions, as well as input and output functions) can quite easily be constructed, e. g., along the lines of the formal description of a CREW PRAM in [7, 10]. We consider a network consisting of  $n$  processors,  $P_1, \dots, P_n$ , and of a set of bidirectional links, each of which joins two of the processors. The topology of the network is described by a graph  $G = (V, E)$  with  $V = \{1, \dots, n\}$ , the edges  $\{i, j\} \in E$  representing the links. Assume the network is to compute a function  $f : A_1 \times \dots \times A_n \rightarrow A$ . At the beginning of the computation, input  $a_i$  is given to processor  $P_i$ , for  $i \in V$ , that means, the initial state of  $P_i$  depends on  $a_i$ .

The computation proceeds synchronously in steps  $t = 1, \dots, T$ . The actions of the network in one step are as follows. We first describe the 1-way case (half-duplex use of links). Processor  $P_i$ ,  $i \in V$ , on the basis of its state after step  $t - 1$ , chooses one of the following two possibilities.

- (S): Choose a message  $m_{i,t}$  and a set  $V_{i,t} \subseteq \{j \in V \mid \{i, j\} \in E\}$  representing possible recipients of the message. We say that  $P_i$  **SENDS** a message in this step. (For notational convenience, we also allow that  $V_{i,t} = \emptyset$ , which means that  $P_i$  does nothing in this step.)
- (R): Choose a nonempty set  $W_{i,t} \subseteq \{l \in V \mid \{l, i\} \in E\}$  of neighbors, representing possible senders from which  $P_i$  wishes to **RECEIVE** a message.

Now, messages are transmitted, so as to satisfy the choices made by the processors, as described in the following. Consider the set

$$E_t = \{(i, j) \mid j \in V_{i,t} \text{ and } i \in W_{j,t}\},$$

representing all edges across which information may flow. It will be important that we may choose from  $E_t$  disjoint edges to be used for communication in a greedy manner, and still are guaranteed that all processors that want to receive a message actually get one. For this, condition (\*) is formulated technically as follows:

- (\*)<sup>1</sup> “*Predictable reception (1-way)*”: For any directed matching  $M' \subseteq E_t$  there is a directed matching  $M$  with  $M' \subseteq M \subseteq E_t$  that covers all recipients, i. e., if  $P_j$  has chosen to receive a message in step  $t$  then  $(i, j) \in M$  for some  $i \in V$ .

Some matching  $M \subseteq E_t$  that covers all recipients is chosen (by “the system”), and for each pair  $(i, j) \in M$  message  $m_{i,t}$  is delivered to  $P_j$ . Messages that are not delivered are discarded. We require that no matter which decision is made here

by “the system”, the output produced at the end of the computation is always correct. (This rule is analogous to the ARBITRARY write-conflict resolution rule for PRAMs, cf. [21].) Processors  $P_i$  that send a message in step  $t$  change their state only by noting that step  $t$  is finished; those processors  $P_j$  that have received a message  $m_{i,t}$  assume a new state that also depends on this message. After the last step,  $T$ , the result  $f(a)$  is known to all processors (in “global output” mode) or to one designated processor  $P_{i_0}$  (in “local output” mode), that means,  $f(a)$  is a function of the state that is finally reached by each processor resp. by processor  $P_{i_0}$ .

**Remark 2.6** We sketch two special combinations of possible choices for the sets  $V_{i,t}$  and  $W_{j,t}$  that may make it clearer which variety of possibilities is covered by this model.

(a) All senders  $P_i$  might specify a set  $V_{i,t}$  with  $|V_{i,t}| = 1$  or  $V_{i,t} = \emptyset$ , and all recipients  $P_j$  might specify  $W_{j,t} = \{i \mid \{i,j\} \in E\}$ : this corresponds to the situation where each receiving processor has a “write window” into which other processors may write; conflicts are resolved by the ARBITRARY rule known from PRAMs. Requirement (\*)<sup>1</sup> simplifies to the condition that for each recipient  $P_j$  there must be at least one message actually addressed to it.

(b) All recipients  $P_j$  might specify a set  $W_{j,t}$  with  $|W_{j,t}| = 1$ , i. e., they are interested in seeing a message from a specific neighbor, and all senders  $P_i$  might specify  $V_{i,t} = \{j \mid \{i,j\} \in E\}$ , i. e., make their respective message available to all their neighbors: this corresponds to a situation in which senders offer their information via a “read window” accessible to all their neighbors. Requirement (\*)<sup>1</sup> turns into the EXCLUSIVE READ rule known from PRAMs.

In the 2-way (full-duplex) variant of the model, the basic structure of the computation is similar. However, here at the beginning of step  $t$  processor  $P_i$  fixes a message  $m_{i,t}$  and a set  $V_{i,t} \subseteq \{j \in V \mid \{i,j\} \in E\}$  of neighbors that are possible partners. (If  $V_{i,t} = \emptyset$ , the processor does not want to communicate.) Consider the graph  $G_t = (V_t, E_t)$ , where  $V_t = \{i \mid V_{i,t} \neq \emptyset\}$  and  $E_t = \{\{i,j\} \mid i,j \in V_t \text{ and } i \in V_{j,t} \text{ and } j \in V_{i,t}\}$ . The condition “predictable reception” here takes on the following form.

(\*)<sup>2</sup> “Predictable reception (2-way)”: Any partial matching  $M' \subseteq E_t$  can be extended to a perfect matching  $M$  for  $G_t$ , i. e., to a matching in  $G_t$  that covers all nodes in  $V_t$  and includes  $M'$ .

The “system” (arbitrarily) chooses one such perfect matching  $M$ , and the processors communicate according to this matching, i. e., for every edge  $\{i,j\} \in M$  message  $m_{i,t}$  is delivered to  $P_j$  and message  $m_{j,t}$  is delivered to  $P_i$ . Processors that have received a message change their state accordingly.

**Definition 2.7** Let  $G = (V, E)$  be a graph specifying a processor network, and let  $f$  be an  $n$ -ary function.

- (a) A network algorithm (in 1-way or 2-way mode) is said to compute  $f$  with “global output” (“g”) if for all inputs  $a$ , after  $T$  steps have been performed, all processors know  $f(a)$ .
- (b) The 1-way network complexity  $T_G^{1,g}(f)$  is the minimum number  $T$  of steps of a 1-way algorithm with global output that computes  $f$ . The 2-way complexity  $T_G^{2,g}(f)$  is defined analogously.
- (c) A network algorithm in 1-way or 2-way mode is said to compute  $f$  with “local output” (“l”) at node  $P_{i_0}$  if after step  $T$  of the algorithm processor  $P_{i_0}$  knows the result  $f(a)$ .
- (d) The network complexities  $T_{G,i_0}^{1,l}(f)$  and  $T_{G,i_0}^{2,l}(f)$  are defined for local output at processor  $P_{i_0}$  in analogy to the case of global output.

**Remark 2.8** While restriction  $(*)^1$  (respectively  $(*)^2$  in the 2-way case) may not seem to be a natural assumption to make in fully synchronous networks, it arises quite naturally in connection with the problem of performing synchronous algorithms on fully asynchronous networks. Assume that in an asynchronous network internal computations of processors and delivery of messages may be delayed for indefinite but finite periods of time. Still, we want to perform an algorithm written for a synchronous network in which processors can either receive or send one message in one step. The obvious idea is to have each processor keep an internal (virtual) step counter, and to keep the exchange of messages synchronized by the use of time stamps. Conceptually, we assume that there is a “communication system” that manages a global message buffer. If  $P_i$  wants to send a message  $m_{i,t}$  to one of the processors  $P_j$ ,  $j \in V_{i,t}$ , it places the message, extended by a time stamp  $t$ , and the list  $V_{i,t}$  into this buffer, and proceeds to step  $t + 1$  no matter what happens to the message. If  $P_j$  wants to receive a message in step  $t$  from one of the processors  $P_i$ ,  $i \in W_{j,t}$ , it submits a corresponding request to the system, again with time stamp  $t$ . The system searches the buffer for a message that has time stamp  $t$  and lists  $P_j$  among its possible recipients. If one is found, it is delivered to  $P_j$  and removed from the buffer;  $P_j$  proceeds to step  $t + 1$ . Otherwise,  $P_j$  is blocked until a message appears in the buffer that can be delivered to it. Synchronous algorithms that obey restriction  $(*)^1$  are suited for being run in this way on an asynchronous network, in the following sense: no matter in which order messages are submitted or delivered, every receive request for step  $t$  will finally be satisfied; it never happens that a processor waits indefinitely for a message that will not arrive. (In the terminology of asynchronous distributed systems, this is a “liveness” property; see [29].)

### 2.3 Complexity measures for functions

We recall some definitions concerning functions  $f : A_1 \times \cdots \times A_n \rightarrow A$ . Originally, these notions were formulated for Boolean functions, but they can readily be

generalized to other domains (cf. [13]). The concept of a *critical input* will be central for the main result. For discussion and applications of all these notions in the context of parallel computing on PRAMs see, e. g., [6, 7, 25, 30], and in particular the survey [13].

**Definition 2.9** Fix a function  $f : A_1 \times \cdots \times A_n \rightarrow A$ .

- (a) An input  $a = (a_1, \dots, a_n)$  is *critical* for  $f$  if for every  $i \in \{1, \dots, n\}$  there is an input  $b$  that differs from  $a$  only in the  $i$ th component and satisfies  $f(a) \neq f(b)$ .
- (b) The *critical complexity*  $c(f)$  is the maximal  $k$  such that there is an input  $a$  with the property that for  $k$  different indices  $i \in \{1, \dots, n\}$  there is an input  $b$  that differs from  $a$  only in the  $i$ th component and satisfies  $f(a) \neq f(b)$ .
- (c) For each input  $a$ , the *sensitive complexity* (or *certificate complexity*)  $s(f, a)$  of  $f$  at  $a$  is the minimal  $k$  such that there is a set  $I_a \subseteq \{1, \dots, n\}$  of cardinality  $k$  with the property that all inputs  $b = (b_1, \dots, b_n)$  with  $\forall i \in I_a: a_i = b_i$  satisfy  $f(a) = f(b)$ . The *sensitive complexity*  $s(f)$  of  $f$  is  $\max\{s(f, a) \mid a \text{ is an input}\}$ .
- (d) The function  $f$  depends on input bit  $i$  if there are inputs  $a$  and  $b$  that differ only in the  $i$ th component and satisfy  $f(a) \neq f(b)$ .

### 3 Computing functions versus gossiping

This section contains the formulation and the proof of the main results concerning the relationship between gossip and accumulation protocols on the one hand and computing functions in networks with global respectively local output on the other hand.

Throughout this section, let  $G = (V, E)$  be a graph with  $V = \{1, \dots, n\}$  for some  $n \in \mathbb{N}$ , and let  $i_0 \in V$  be fixed. We assume alternately that  $G$  represents a gossiping network and a processor network. The following is well known and an almost immediate consequence of the definitions.

**Observation 3.1** If  $f : A_1 \times \cdots \times A_n \rightarrow A$  is an arbitrary  $n$ -ary function, then

- (a)  $T_G^{1,g}(f) \leq r(G)$ ;
- (b)  $T_G^{2,g}(f) \leq r_2(G)$ ;
- (c)  $T_{G,i_0}^{1,1}(f), T_{G,i_0}^{2,1}(f) \leq b(G, i_0)$ , for  $i_0 \in V$ .

If  $f$  is the OR of  $n$  bits, it can be computed within these time bounds by using 1-bit messages only.

*Proof.* We only consider (a); the other statements are proved similarly. Let  $\mathcal{M} = (M_1, \dots, M_T)$  be a 1-way gossip protocol, and let  $a = (a_1, \dots, a_n)$  be the input. The sequence  $\mathcal{K}(\mathcal{M}) = (K_0, K_1, \dots, K_T)$  is defined as in Definition 2.2. In step  $t$ ,  $1 \leq t \leq T$ , the processors send each other messages according to the communication pattern given by  $M_t$ . If  $(i, j) \in M_t$ , processor  $P_i$  sends a message  $m_{i,t}$  to  $P_j$ , where  $m_{i,t}$  is a representation of the input fragment  $(a_l)_{l \in K_i(t-1)}$ . From the definition of  $\mathcal{K}(\mathcal{M})$  it follows by induction that the required information is available to  $P_i$  in this step; since  $\mathcal{M}$  is a gossip protocol, after step  $t$  each processor knows the complete input  $a$  and can compute  $f(a)$ . Note that the longest message can be at most as long as  $a$ . — If  $f$  is the OR function, it is sufficient to send as message  $m_{i,t}$  the single bit  $\bigvee_{l \in K_i(t-1)} a_l$ . (Note that the only properties of the OR function used here are that it is commutative, associative, and idempotent. In complete networks, idempotency is not required, see [2, 5].) ■

The main result of this paper is essentially that, under the restriction “predictable reception”, these algorithms for computing  $f$  are optimal if global respectively local output is required and  $f$  has a critical input. If global output is required and  $f$  is nonconstant, lower bounds for broadcast protocols apply.

**Theorem 3.2** *Let  $G = (V, E)$  be a graph with node set  $V = \{1, \dots, n\}$ , which represents a network of processors, and let  $f$  be an  $n$ -ary function. If  $f$  has a critical input, then*

- (a)  $T_G^{1,g}(f) = r(G)$ ;
- (b)  $T_G^{2,g}(f) = r_2(G)$ ;
- (c)  $T_{G,i_0}^{1,1}(f) = T_{G,i_0}^{2,1}(f) = b(G, i_0)$ , for  $i_0 \in V$ .

Further,

- (d) if  $f$  depends on the  $i_0$ th input component for some  $i_0 \in V$ , then  $T_G^{1,g}(f), T_G^{2,g}(f) \geq b(G, i_0)$ .

**Corollary 3.3** *The complexity of computing the OR of  $n$  bits on  $G$  (under restriction  $(*)^1$  resp.  $(*)^2$ ) is given by*

- (a)  $r(G)$  [ $r_2(G)$ ] for 1-way [2-way] communication and global output, and
- (b)  $b(G, i_0)$  for 1-way or 2-way communication and local output at processor  $P_{i_0}$ .

(In the upper bounds, 1-bit messages are sufficient.) ■

*Proof of Theorem 3.2.* We deal with part (a) in detail; the proofs of the other parts, being similar, will only be sketched. In view of Observation 3.1, only the inequality  $T_G^{1,g}(f) \geq r(G)$  has to be proved. For this, assume that a 1-way algorithm for computing  $f$  on  $G$  in  $T$  steps is given, and that  $a^* = (a_1^*, \dots, a_n^*)$  is a critical input for  $f$ . Obviously, it is sufficient to construct a 1-way gossip protocol for  $G$  that has  $T$  rounds.

The construction splits into two parts. First, we eliminate ambiguities from computations according to the algorithm, i. e., for each input  $a$  we fix a computation  $C_a$ , which essentially corresponds to a communication pattern. In the second part, we show that  $C_{a^*}$  induces a gossip protocol for  $G$ .

*Part 1: Fix computations.* Consider the network algorithm that computes  $f$  in  $T$  steps. First, we arbitrarily fix a computation  $C_{a^*}$  for input  $a^*$ . Consider steps  $t = 1, \dots, T$  one after the other. Let  $E_t(a^*)$  be the set of all possible pairs of senders and recipients determined on the basis of step  $t - 1$  (cf. Section 2.2). Then an arbitrary matching  $M_t^* \subseteq E_t(a^*)$  is chosen that covers all recipients, and messages are delivered according to  $M_t^*$ . Now, consider some input  $a \neq a^*$ . We proceed by induction on  $t$ . Assume  $C_a$  has been fixed up to step  $t - 1$ , and consider the graph  $(V, E_t(a))$  induced by the communication requests of the processors for step  $t$ . Let  $M_t'(a) = M_t^* \cap E_t(a)$  be the set of those edges in  $E_t(a)$  that are used in step  $t$  of  $C_{a^*}$ . By restriction  $(*)^1$ , we may extend  $M_t'(a)$  to some directed matching  $M_t(a) \subseteq E_t(a)$  that covers all recipients. Deliver messages according to  $M_t(a)$ .

*Part 2: Identify a gossip protocol.* Consider the communication protocol  $\mathcal{M}^* = (M_1^*, \dots, M_T^*)$ . The proof of part (a) of the theorem is finished once the following assertion is established.

**Lemma 3.4 (Main Lemma)**  $\mathcal{M}^*$  is a 1-way gossip protocol for  $G$ .

*Proof.* Let the sequence  $\mathcal{K}(\mathcal{M}^*) = (K_0, K_1, \dots, K_T)$  be as in Definition 2.2. We must show that  $K_T(i) = V$  for all  $i \in V$ . For this, it is sufficient to establish the following assertion  $(\mathcal{A}_t)$  for  $t = T$ :

$(\mathcal{A}_t)$  for all inputs  $a = (a_1, \dots, a_n)$  and all  $i \in V$ :  
 $(\forall j \in K_t(i): a_j = a_j^*) \Rightarrow P_i$  is in the same state in  $C_{a^*}$  and  $C_a$  after step  $t$ .

Indeed, if  $K_T(i) \neq V$ , e. g.,  $j \notin K_T(i)$ , then by  $(\mathcal{A}_T)$  processor  $P_i$  is in the same state after step  $T$  on input  $a^*$  and each input that differs from  $a^*$  only in the  $j$ th component, which contradicts the assumption that  $a^*$  is critical for  $f$  and that the network computes the function  $f$  with global output.

We prove  $(\mathcal{A}_t)$  by induction on  $t$ . For  $t = 0$ , the claim follows from the definitions:  $K_0(i) = \{i\}$ , for  $i \in V$ , and the initial state of  $P_i$  only depends on the  $i$ th input bit. Now assume  $t > 0$ , and that  $(\mathcal{A}_{t-1})$  is true. Let  $i \in V$  and  $a$  be an input so that  $a_j = a_j^*$  for all  $j \in K_t(i)$ . There are two cases.

*Case 1:* There is no  $l$  such that  $(l, i) \in M_t^*$ . Then, by definition,  $K_t(i) = K_{t-1}(i)$ , in particular  $a_j = a_j^*$  for all  $j \in K_{t-1}(i)$ . By the induction hypothesis,

$P_i$  is in the same state after step  $t - 1$  in both  $C_{a^*}$  and  $C_a$ . Since no  $l$  satisfies  $(l, i) \in M_t^*$ , no processor  $P_l$  delivers a message to  $P_i$  in step  $t$  of  $C_{a^*}$ , i. e., by restriction  $(*)^1$  (“predictable reception”), in step  $t$  of  $C_{a^*}$  processor  $P_i$  is a sender. Since this decision is based on the state at the end of step  $t - 1$ ,  $P_i$  decides in the same way in computation  $C_a$ , thus does not receive a message in  $C_a$  in step  $t$  either, and enters the same state in  $C_a$  as in  $C_{a^*}$ .

*Case 2:* There is some (unique)  $l$  such that  $(l, i) \in M_t^*$ . In this case, by definition,  $K_t(i) = K_{t-1}(l) \cup K_{t-1}(i)$ . We apply the induction hypothesis  $(\mathcal{A}_{t-1})$  to  $a$  with respect to both  $K_{t-1}(l)$  and  $K_{t-1}(i)$  to conclude that  $P_l$  is in the same state after step  $t - 1$  of  $C_a$  and of  $C_{a^*}$ , and that the same is true for  $P_i$ . Since  $(l, i) \in M_t^*$ , the message  $m_{l,t}$  sent by processor  $P_l$  in this step is delivered to  $P_i$  in step  $t$  of computation  $C_{a^*}$ ; in particular,  $P_l$  is a sender with  $i \in V_{l,t}$  and  $P_i$  is a recipient with  $l \in W_{i,t}$ . Since both  $P_i$  and  $P_l$  are in the same state after step  $t - 1$  in  $C_{a^*}$  and  $C_a$ , this will also be true in computation  $C_a$ ; thus, edge  $(l, i)$  is in  $E_t(a)$  and in  $M_t^*$ , and  $P_l$  sends message  $m_{l,t}$  in step  $t$  of  $C_a$  as well. By the construction of  $C_a$  described in Part 1, edge  $(l, i)$  will be chosen to be in  $M_t(a)$ , and message  $m_{l,t}$  will be delivered to  $P_i$  in  $C_a$ . This implies that  $P_i$  receives identical messages in step  $t$  of  $C_{a^*}$  and of  $C_a$ ; hence  $P_i$  will enter the same state at the end of step  $t$  in these two computations. This finishes the induction step, and the proof of the main lemma.  $\blacksquare$

We will not prove parts (b) and (c) of the theorem in detail, since the arguments are essentially the same as in part (a). We only point out the changes to be made. In part (b), undirected matchings are used in place of directed matchings. We fix a computation  $C_{a^*}$  arbitrarily, and show exactly as in the 1-way case that the sequence  $\mathcal{M} = (M_1(a^*), \dots, M_T(a^*))$  of matchings that describe the pairs of processors that communicate in each step is a 2-way gossip protocol. Restriction  $(*)^2$  is formulated in such a way that the argument from (a) carries over without any difficulties. In part (c), the construction of the communication protocol is exactly the same as in part (a) respectively (b). The only difference in the proof is that since it is only required that  $P_{i_0}$  knows the result at the end, we can only conclude that  $K_T(i_0) = V$ , which means that we have constructed an accumulation protocol for  $G$ . Then Fact 2.5 and the remark preceding it yield (c).

Finally, we sketch the proof of part (d), for 1-way mode. Choose  $a^*$  and  $a$  such that  $f(a^*) \neq f(a)$  and  $a^*$  and  $a$  differ only in the  $i_0$ th component. Computation  $C_{a^*}$  and (afterwards) computation  $C_a$ , each with  $T$  steps, are fixed exactly as in the proof of part (a); the sequence  $\mathcal{M} = (M_1(a^*), \dots, M_T(a^*))$  is defined as above. Consider the resulting sequence  $\mathcal{K}(\mathcal{M}) = (K_0, \dots, K_T)$  (Definition 2.2). The assertion

$$(\mathcal{A}'_t) \quad \forall i \in V [ i_0 \notin K_t(i) \Rightarrow P_i \text{ is in the same state after step } t \text{ of } C_{a^*} \text{ and } C_a ]$$

is proved by induction on  $t$ . Since all processors  $P_i$  must be able to distinguish

$a^*$  and  $a$  after step  $T$ , we must have  $i_0 \in K_T(i)$  for all  $i \in V$ , thus,  $\mathcal{M}^*$  is a broadcast protocol, which establishes the inequality  $T \geq b(G, i_0)$ . ■

## 4 Extensions and limitations

In this section, we consider possible extensions of the main result. First, we discuss functions that do not necessarily have a critical input; next, multiple-output functions are considered; finally, the situation of processor networks that do not necessarily satisfy restriction  $(*)$  is discussed. The complexity of broadcasting a bit in networks with an even more general communication mode is treated in detail in Section 5.

### 4.1 Boolean functions without critical input

What can be said about functions that do not have a critical input? We can use the approach taken in the proof of Theorem 3.2 to obtain the following. Assume a network  $G = (V, E)$  computes a function  $f$  with global output, and  $a^* = (a_1^*, \dots, a_n^*)$  is an arbitrary input. Define  $\mathcal{M}^* = (M_1^*, \dots, M_t^*)$  as before, in accordance with an arbitrarily chosen computation on this input. Then it is easily seen that assertion  $(\mathcal{A}_T)$  still holds: for all inputs  $a$  and all  $j \in V$  we have that if  $a_i = a_i^*$  for all  $i \in K_T(j)$ , then  $f(a) = f(a^*)$ . This means that for each processor  $P_j$  the set  $K_T(j)$  must be a “certificate set”  $I_{a^*}$  for the function value  $f(a^*)$ , as in Definition 2.9(c). If we choose  $a^*$  to be an input with  $s(f) = s(f, a^*)$ , we obtain that  $|K_T(j)| \geq s(f)$  for all  $j \in V$ .

**Definition 4.1** For  $1 \leq s \leq n$ , let  $r(G, s)$  denote the minimum number of rounds of a communication protocol for  $G$  that satisfies  $|K_T(j)| \geq s$  for all  $j \in V$ .

While the message complexity of the problem to obtain at least  $s$  different pieces of information was defined and investigated already in [26], its round complexity  $r(G, s)$  does not seem to have received much attention, with the following exception: in the case of the complete network, in [4] it was established that the technique of [8, 12, 23, 27] can quite easily be adapted to prove a lower bound of about  $\log_\rho(s)$  for this case with  $\rho = \frac{1}{2}(1 + \sqrt{5})$ , which is about  $1.44 \dots \log(s)$ .

Alternatively, one can consider the critical complexity of  $f$  (Definition 2.9(b)). Let  $a^*$  be an input and  $I \subseteq V$  be a set of size  $c(f)$  such that for every  $i \in I$  there is an input  $a$  that differs from  $a^*$  only in the  $i$ th component and satisfies  $f(a) \neq f(a^*)$ . Then from  $(\mathcal{A}_T)$  it is clear that the communication protocol defined in the proof of Theorem 3.2 satisfies  $I \subseteq K_T(j)$  for all  $j \in V$ . This relates to the problem of broadcasting from  $k = c(f)$  fixed sources. Results for this problem for specific networks and specific placement of these sources have been obtained by Höltring in her diploma thesis [18].



## 4.2 Functions with multiple outputs

In applications, often functions  $f : A_1 \times \cdots \times A_n \rightarrow B_1 \times \cdots \times B_n$ ,  $a \mapsto (f_1(a), \dots, f_n(a))$ , must be computed in such a way that the  $j$ th component  $f_j(a)$  appears at processor  $P_j$  at the end. In [8], matrix inversion, discrete Fourier transform, and sorting are listed as examples. Our methods apply to this situation *directly* if there is one input  $a^*$  that is critical for all functions  $f_1, \dots, f_n$ . This is the case, e. g., for the problem of sorting  $n$  numbers  $a_1, \dots, a_n$  from the set  $\{1, \dots, m\}$  with  $m \geq n + 2$ , so that the number with rank  $j$  appears at processor  $P_j$  at the end, since the input  $(2, \dots, n + 1)$  is critical for all output positions. In many important cases, though, no single input can be found that is critical for all components of the output, e. g., for the problem of sorting  $n$  bits. For the special case of the complete network, we still can obtain lower bounds in terms of the sensitive complexity of  $f_1, \dots, f_n$  on specific inputs, in the following way.

**Proposition 4.2** *If the multiple-output function  $f : a \mapsto (f_1(a), \dots, f_n(a))$  is computed on the complete network in 1-way mode in  $T$  steps, under restriction  $(*)^1$ , and  $a$  is any input, then*

$$T \geq \frac{1}{2} \cdot \log_\rho \left( \left( \sum_{1 \leq j \leq n} s(f_j, a)^2 \right) / n \right) \geq 0.72\dots \cdot \log \left( \left( \sum_{1 \leq j \leq n} s(f_j, a)^2 \right) / n \right),$$

where  $\rho = (1 + \sqrt{5})/2$ .

*Proof.* (Sketch.) Fix any input  $a^*$ . Just as in Section 4.1 we can see that the communication protocol  $\mathcal{M}^*$  induced by  $a^*$  and the corresponding sequence  $\mathcal{K}(\mathcal{M}^*) = (K_1, \dots, K_T)$  must satisfy  $|K_T(j)| \geq s(f_j, a)$ , for  $1 \leq j \leq n$ . The proof method from [8, 12, 23, 27] can be applied to show that  $T$  must satisfy

$$\rho^T \cdot \sqrt{n} \geq \left( \sum_{1 \leq j \leq n} |K_T(j)|^2 \right)^{1/2}.$$

Combining these inequalities and taking logarithms yields the result. ■

**Corollary 4.3** *Sorting  $n$  bits in 1-way mode on a complete network with restriction  $(*)^1$  takes at least  $1.44\dots \log n - O(1)$  steps.*

*Proof.* Assume the problem of sorting  $n$  bits can be solved in  $T$  steps. The function  $f_j$  is the output bit that appears at processor  $P_j$ . As input, we choose  $a^* = (0, \dots, 0)$ . Since  $f_j(a) = 0$  if and only if  $a$  contains  $j$  zeroes, it is clear that  $s(f_j, a^*) = j$ . By Proposition 4.2,  $T \geq \frac{1}{2} \log_\rho((\sum_{1 \leq j \leq n} j^2)/n) = \log_\rho(n) - O(1)$ . ■

We leave it to the reader to find inputs for the other problems mentioned (discrete Fourier transform or matrix inversion) that make it possible to show that these problems have high complexity in complete processor networks.

In order to apply results from the gossiping framework to the computation of multiple-output functions in other, sparse, networks, we would have to have more information on partial gossiping in these networks, in particular, the complexity of collecting any  $k$  input pieces or a predetermined set of  $k$  input pieces in every node.

### 4.3 Networks without “predictable reception”

In this section, we show that Theorem 3.2 becomes false if the network algorithm does not have to satisfy restriction (\*) (“predictable restriction”) but still is fully synchronous. In this case, we may collect information with a large fan-in, and distribute information faster using the phenomenon of “transmitting information by *not* sending a message”, cf. [3, 7, 9]. Then, we show that even such networks cannot compute functions with a critical input in global output mode more than four times faster than gossip protocols.

Let us start with two examples, the tree network and the complete network.

**Proposition 4.4** *Let  $d \geq 1$ , and let  $G$  be a full binary tree of depth  $d$  with root  $P_1$ . (Thus,  $G$  has  $2^{d+1} - 1$  nodes.) With synchronous algorithms in 1-way mode (which do not satisfy restriction  $(*)^1$ ), in  $G$*

- (a) *the OR of  $n$  bits with local output at  $P_1$  can be computed in  $d$  steps;*
- (b) *a bit  $b$  located initially at  $P_1$  can be broadcasted to all nodes in  $d$  steps;*
- (c) *the OR of  $n$  bits with global output can be computed in  $2d$  steps.*

*Proof.* (a) It is sufficient to note that the OR of 3 bits  $a_1, a_2$ , and  $a_3$  can be computed in 1 step, because then we can proceed iteratively, starting from the leaves. In one step,  $P_2$  [ $P_3$ ] sends a message to  $P_1$  if and only if  $a_2 = 1$  [ $a_3 = 1$ ]. If  $P_1$  gets any messages, the result is 1, otherwise it is equal to  $a_1$ . — (b) Again, we only have to show how this can be done for  $d = 1$ ; in larger trees, we iterate, starting from the root.  $P_1$  informs both of its sons in one step as follows: if  $b = 0$  then it sends a message to  $P_2$ ; if  $b = 1$  then it sends a message to  $P_3$ . In each case, the processor that did not get a message can derive  $b$  from just this fact. (The contents of the message are irrelevant. Cf. [3, 9] for this trick.) — (c) Apply (a) and (b) one after the other. ■

The resulting running times should be contrasted with the fact that if  $G$  is the binary tree of depth  $d$  with root  $P_1$  then  $b(G, P_1) = 2d$  and  $r(G) = 4d$  (cf. [20]). In the complete network, the situation is slightly different.

**Proposition 4.5** *Let  $G$  be the complete network of size  $n$ . With synchronous algorithms in 1-way mode (which do not satisfy restriction  $(*)^1$ ), in  $G$*

- (a) the OR of  $n$  bits with local output can be computed in one step;
- (b) a bit  $c$  located initially at some node can be broadcasted to all nodes in  $\lceil \log_3 n \rceil \approx 0.63 \log n$  steps;
- (c) the OR of  $n$  bits with global output can be computed in  $1 + \lceil \log_3 n \rceil$  steps.

*Proof.* (a) For  $2 \leq i \leq n$ , processor  $P_i$  sends a message to  $P_1$  if and only if  $a_i = 1$ . If  $P_1$  gets a message, the result is 1, otherwise it is  $a_1$ . — (b) Using the same trick as in the proof of the previous proposition, in one step a processor that knows  $c$  can inform two other processors of this value; thus, the number of processors that know  $c$  can be tripled in one step. — (c) Combine (a) and (b). ■

The running times from this proposition should be compared with the values  $a(G, i_0) = b(G, i_0) = \lceil \log n \rceil$  and  $r(G) = 1.44\dots \log n \pm O(1)$  for  $G$  the complete network. Next, we note that for computing nonconstant functions with global output in any network  $G$ , at least  $\frac{1}{4}r(G)$  steps are needed. This bound is tight, as the example of the binary tree shows.

**Proposition 4.6** *If a nonconstant function  $f$  can be computed in  $T$  steps with global output in the network  $G$  (without restriction  $(*)^1$ ), then  $b(G, i_0) \leq 2T$  for some  $i_0$ , and  $r(G) \leq 4T$ .*

*Proof.* The main part of the proof is postponed to Section 5, where we characterize the complexity of broadcasting a bit in synchronous networks that are even stronger than those considered here in terms of a variant of the broadcast complexity of the underlying graph. Clearly, if a network  $G$  can compute a function  $f$  that depends on input position  $i_0$  in  $T$  steps with global output, then it can broadcast one bit from source  $P_{i_0}$  in  $T$  steps. In Section 5, we show that broadcasting one bit from a node  $i_0$  in  $G$  takes at least  $b_2(G, i_0)$  steps, which is the number of rounds in an optimal “2-broadcast” protocol, in which one node may send a message to two of its neighbors in one step. It is easy to see that  $b(G, i_0) \leq 2b_2(G, i_0)$ , hence  $b(G, i_0) \leq 2T$ . By the trivial fact that  $r(G) \leq 2b(G, i_0)$  (cf. [20]), we obtain  $r(G) \leq 4T$ . ■

## 5 Broadcasting a bit in a synchronous network

In this section, we study the complexity of broadcasting one bit in a synchronous network of processors that communicate by message passing, with hardly any restriction on the communication mode excepting that a processor must not send more than one message in one step. We show, in analogy to Theorem 3.2, that a well-known complexity measure from the theory of broadcast protocols for atomic pieces of information is appropriate for describing the complexity of this problem.

Note that the problem of broadcasting a bit captures the essence of all situations in which differences in the state of a single processor finally influence all others.

We treat communication in 1-way mode and in 2-way mode separately (Sections 5.3 and 5.4, respectively). The results of this section can be applied to other parallel models like EXCLUSIVE READ PRAMs and 1-ARBITRARY distributed memory machines, as will be indicated in Section 5.4.

## 5.1 The general network model in 1-way mode

As before, we consider a network of  $n$  processors, connected by bidirectional links according to a graph  $G = (V, E)$ . First, we focus on 1-way communication. One node  $P_{i_0}$  is distinguished as the source of the broadcasting process, i. e., this processor can be in either one of two different states initially, representing inputs 0 and 1; the other processors are in an initial state that is independent of the input. In step  $t$ , a processor may send a message *and* receive several messages, as specified by the following rules. Depending on its state after step  $t - 1$ , processor  $P_i$  fixes a message  $m_{i,t}$  that it wishes to send, and a set  $V_{i,t} \subseteq \{j \mid \{i, j\} \in E\}$  of possible recipients. (As in Section 2.2, the choice  $V_{i,t} = \emptyset$  indicates that  $P_i$  does not send a message at all; with  $|V_{i,t}| = 1$  a unique recipient can be specified.) Further,  $P_i$  specifies a set  $W_{i,t} \subseteq \{j \mid \{i, j\} \in E\}$  of processors from which it wishes to receive a message. Let  $E_t = \{(i, j) \mid j \in V_{i,t} \text{ and } i \in W_{j,t}\}$ . Some set  $E'_t \subseteq E_t$  is selected arbitrarily such that for each  $i$  that occurs as a first component in  $E_t$  there is exactly one  $j$  such that  $(i, j) \in E'_t$ . Then, for each  $(i, j) \in E'_t$ , message  $m_{i,t}$  is delivered to  $P_j$ . Thus, a processor  $P_j$  that has specified  $W_{j,t} \neq \emptyset$  may receive messages from none, some, or all  $P_i$  with  $j \in W_{j,t}$ . On the basis of *all* messages received  $P_j$  changes its state. After step  $T$ , all processors must know whether the input bit was 0 or 1. (The reader is invited to check against his or her intuition that alternative rules for dealing with surplus messages like buffering for later delivery, combining, discarding, etc. are at most as strong as this scheme.)

**Definition 5.1** *The broadcast complexity  $\hat{T}_{G,i_0}$  of a processor network  $G$  in 1-way mode with source  $P_{i_0}$  is the smallest number of steps an algorithm for broadcasting a bit from  $P_{i_0}$  can have.*

## 5.2 2-broadcast protocols

Turning now to communication networks, we generalize broadcast protocols (see Section 2) to *2-broadcast protocols*, in which one node may pass the information it has to two of its neighbors in one step. (This modification is discussed as “DMA-bound model”  $H_2$ , an abbreviation for “half-duplex with outdegree 2”, in [8, 14].) Our definition of 2-broadcast protocols differs only formally from the standard definition.

**Definition 5.2** *Let  $G = (V, E)$  be a graph, and  $v_0 \in V$ .*

- (a) A 2-broadcast protocol in 1-way mode with source node  $i_0$  is a sequence  $\mathcal{M} = (M_1, \dots, M_T)$  of sets  $M_t$  of directed edges with the following properties:
- (i) if  $(i, j) \in M_t$  then  $\{i, j\} \in E$ , for  $1 \leq t \leq T$ ;
  - (ii)  $i$  occurs at most twice as a first component in  $M_t$  and at most once as a second component in  $M_t$ , for each  $i \in V$  and  $1 \leq t \leq T$ .
  - (iii) if  $\mathcal{K}(\mathcal{M}) = (K_1, \dots, K_T)$  is as in Definition 2.2 then  $i_0 \in K_T(i)$  for all  $i \in V$ .
- (b) The 2-broadcast complexity of  $G$  with source node  $i_0$  is the minimum  $T$  such that there is a 2-broadcast protocol as in (a) with  $T$  rounds.

A 2-broadcast protocol can be used to broadcast a piece of information, initially located at node  $i_0$ , in the following obvious way: if edge  $(i, j)$  is in  $M_t$  and  $i_0 \in K_t(i)$ , then in round  $t$  the piece of information is sent from  $i$  to  $j$ . The definition makes sure that the node that has to send the piece of information actually has received it before and that every processor sends the piece of information to at most two of its neighbors in one step. As mentioned before, it is obvious that  $b(G, i_0) \leq 2b_2(G, i_0)$  for arbitrary networks  $G$  and nodes  $i_0$  in  $G$ .

### 5.3 Processor networks versus 2-broadcast networks

By using the trick described in the proof of Proposition 4.4(b) for sending one bit to two recipients in one step, we can transform any 2-broadcast protocol for a graph  $G$  into a broadcast algorithm in 1-way mode for the processor network with topology given by  $G$ .

**Observation 5.3**  $\hat{T}_{G, i_0} \leq b_2(G, i_0)$ .

*Proof.* Assume that a 2-broadcast protocol for  $G$  with  $T$  rounds is given, and define  $\mathcal{K}(\mathcal{M}) = (M_1, \dots, M_T)$  as before. We obtain an algorithm for the network as follows. If edges  $(i, j_1)$  and  $(i, j_2)$  are in  $E_t$  and  $i_0 \in K_t(i)$ , then in step  $t$  processor  $P_i$  informs  $P_{j_1}$  and  $P_{j_2}$  of the value of  $b$  with the trick described in the proof of Proposition 4.4(b). If there is only one edge leaving  $i$  in  $E_t$ , the trick can be applied nonetheless. We note that in this way all processors can use the same message in all steps. ■

The purpose of this section is to show that this algorithm is optimal. This may be intuitively plausible, but to the best of the knowledge of the author a formal proof, especially for a network model as general as considered here, has not been available so far.

**Theorem 5.4**  $\hat{T}_{G, i_0} = b_2(G, i_0)$ .

*Proof.* In view of Observation 5.3 we must only prove that  $\hat{T}_{G,i_0} \geq b_2(G, i_0)$ . Let a broadcast algorithm for the processor network  $G$  in 1-way mode with source  $i_0$  be given that runs in  $T$  steps. We construct a 2-broadcast protocol with at most  $T$  steps.

We proceed similarly as in the proof of Theorem 3.2. However, here computations  $C_0$  and  $C_1$  on inputs 0 and 1 are fixed simultaneously by induction on  $t$ . Thus, assume that  $C_0$  and  $C_1$  have been fixed up to step  $t - 1$ . For  $b = 0, 1$ , the communication requests of the processors (as represented by the sets  $V_{i,t}(b)$  and  $W_{i,t}(b)$ , for  $i \in V$ ) induce a set  $E_t(b)$  of directed edges across which messages may flow. We let  $E_t'' = E_t(0) \cap E_t(1)$  and choose a subset  $E_t' \subseteq E_t''$  so that if  $i$  appears as a first component in  $E_t''$  then  $(i, j) \in E_t'$  for exactly one  $j$ . Then, for  $b = 0, 1$  separately, a set  $E_t'(b)$  with  $E_t' \subseteq E_t'(b) \subseteq E_t(b)$  is chosen with the property that if  $i$  appears as a first component in  $E_t(b)$  then  $(i, j) \in E_t'(b)$  for exactly one  $j$ , and in step  $t$  of  $C_b$  message  $m_{i,t}$  is delivered to  $P_j$  for all pairs  $(i, j) \in E_t'(b)$ .

The intuition behind the next definition is that we try to “peel off” inessential parts of the two computations, i. e., to identify “meaningless” messages and eliminate them. For this, we define a set  $E_{0,1}$  of labeled, directed edges that run along some of the edges of  $G$ , and identify a 2-broadcast protocol as part of this set.

- Edge  $(i, j)$ , with label  $t$ , for  $1 \leq t \leq T$ , is in  $E_{0,1}$  if and only if there is some  $b \in \{0, 1\}$  such that in computation  $C_b$  the message  $m_{i,t}$  sent by processor  $P_i$  is delivered to  $P_j$ , but in the other computation  $C_{\bar{b}}$  no message or a different one is sent from  $P_i$  to  $P_j$  in step  $t$ .

Note that parallel edges are possible, involving different time steps. Note also that identical messages that are sent across the same edge in both  $C_0$  and  $C_1$  are ignored in this definition, corresponding to the intuition that they are “meaningless”. The following simple observations are crucial.

*Claim 1:* Let  $i \neq i_0$  and  $j \neq i_0$ . If  $(i, j)$  with label  $t$  is in  $E_{0,1}$ , and no edge  $(i', j)$  with a label  $t' < t$  is in  $E_{0,1}$ , then there must be some  $l \in V$  such that edge  $(l, i)$  with label  $t''$  is in  $E_{0,1}$  for some  $t'' < t$ .

(Intuitively spoken, in order to send a meaningful message in step  $t$ , a processor  $P_i$  with  $i \neq i_0$  must either

- have received a meaningful message in an earlier step in the same computation, or
- have noticed that a meaningful message that would have been due to arrive in the other computation has not turned up, or
- be treated differently by the intended recipient of the message in the two computations.)

*Claim 2:* If  $i \neq i_0$  then there is at least one labeled edge that enters node  $i$ .

(Intuitively, if  $P_i$ ,  $i \neq i_0$ , never gets a meaningful message in either computation, it cannot know the input bit at the end.)

*Proof of Claim 1:* Assume for a contradiction that  $(i, j)$  with label  $t$  is in  $E_{0,1}$ , but that in  $E_{0,1}$  there is neither an edge  $(i', j)$  nor an edge  $(l, i)$  with a label  $t'' < t$ . By symmetry, we may assume without loss of generality that in computation  $C_0$  processor  $P_i$  sends a message  $m_{i,t}$ , which is delivered to  $P_j$ , but that this message does not flow from  $P_i$  to  $P_j$  in  $C_1$ . This means that  $(i, j) \in E_t(0)$ , hence we have  $j \in V_{i,t}(0)$  and  $i \in W_{j,t}(0)$ . The last two assumptions imply that processor  $P_i$ , which does not know  $b$  initially since  $i \neq i_0$ , has received exactly the same messages in each of the steps  $1, \dots, t-1$  in both computations, and that the same is true for  $P_j$ . This entails that  $P_i$  and  $P_j$  both are in the same state after step  $t-1$  in  $C_0$  and  $C_1$ , hence  $j \in V_{i,t}(1)$  and  $i \in W_{j,t}(1)$ , which implies  $(i, j) \in E_t(1)$ , and  $P_i$  sends message  $m_{i,t}$  in  $C_1$  as well. By the definition of  $E_t''$ , edge  $(i, j)$  appears in this set. In computation  $C_0$  edge  $(i, j)$  is used for delivering  $m_{i,t}$ , which implies  $(i, j) \in E_t'$ . Hence,  $(i, j) \in E_t'(1)$  as well, which means that message  $m_{i,t}$  is delivered to  $P_j$  in  $C_1$  as well, a contradiction.

*Proof of Claim 2:* If no labeled edges enter  $i$ , then all messages received by  $P_i$  in steps  $t$ ,  $1 \leq t \leq T$ , are identical in  $C_0$  and  $C_1$ . Thus, since  $i \neq i_0$ , in both computations  $P_i$  will be in the same state at the end of step  $T$ . This contradicts the requirement that at the end all processors must know  $b$ .

Next, we eliminate some of the labeled edges, as follows. All labeled edges that enter  $i_0$  are removed; all labeled edges that enter  $i \neq i_0$  excepting that one with the smallest label  $t$  are also removed. (This operation corresponds to the intuition that once a processor has received a meaningful message, or noted that it did not arrive at its due time, it knows  $b$ , and later messages are irrelevant.) The resulting edge set is called  $E_{0,1}^*$ . The digraph  $G^* = (V, E_{0,1}^*)$  has the following two properties.

- (i) All nodes excepting  $i_0$  have indegree 1 in  $G^*$ .
- (ii) If  $(i, j)$  with label  $t$  is in  $E_{0,1}^*$ , then either  $i = i_0$  or the (unique) edge that enters  $i$  in  $G^*$  has a label  $t' < t$ .

(Property (i) is immediate from Claim 2 and the definition of  $E_{0,1}^*$ . Property (ii) is a consequence of Claim 1 and the definition of  $E_{0,1}^*$ . Indeed, if  $(i, j)$  with label  $t$  is in  $E_{0,1}^*$ , then there cannot be an edge  $(i', j)$  in  $E_{0,1}$  with an “earlier” label  $t' < t$ , and  $j \neq i_0$ . We may assume that  $i \neq i_0$ , and apply Claim 1 to conclude that there is an edge in  $E_{0,1}$  that enters  $i$  and has a label  $t'' < t$ . Again by the definition of  $E_{0,1}^*$ , the label of the unique edge  $(l, i) \in E_{0,1}^*$  that enters  $i$  must have a step number which is at most  $t''$ , thus, smaller than  $t$ .)

Properties (i) and (ii) taken together say that  $G^*$  is a directed spanning tree for  $V$  with root  $i_0$ , and that the labels along directed paths in  $G^*$  are strictly increasing with respect to their  $t$ -parts. Moreover, for each node and each  $t$  there can be at most two edges leaving  $i$  that are labeled with  $t$ . From this it easily follows that we obtain a 2-broadcast protocol for  $G$  with at most  $T$  rounds by defining  $E_t = \{(i, j) \mid (i, j) \in E_{0,1}^* \text{ and } (i, j) \text{ has label } t\}$ , for  $1 \leq t \leq T$ . ■

## 5.4 Broadcasting in 2-way mode and in other parallel models

The complexity of broadcasting a bit on EREW PRAMs has been determined in [3]. Let us recall briefly how such a parallel model works. It consists of processors  $Q_1, \dots, Q_p$  and cells  $C_1, \dots, C_r$ . In one step, a processor reads from a cell, performs some internal computation, and writes to a cell. It is forbidden that in any step more than one processor reads from the same cell or more than one processor writes to the same cell. Initially, one cell contains a bit, the others have a neutral content. EREW algorithms can be regarded as algorithms for processor networks with  $n = p + r$  processors. Each PRAM processor and each PRAM cell are represented as a network processor. A write phase is represented as a step in which all PRAM processors are senders (with 0 or 1 possible recipients) and all cells are recipients (specifying all PRAM processors as possible senders). A read phase is represented as a step in which all PRAM processors are recipients (specifying 0 or 1 cell as possible sender). If we apply the technique from the previous section, we obtain a 2-broadcast tree in which levels alternately consist of PRAM cells and PRAM processors. It is not hard to show that the numbers  $q_t$  of the cumulative size of the levels 1 through  $t$  of processors and the numbers  $c_t$  of the cumulative size of the level 1 through  $t$  of cells satisfy the recurrence inequalities

$$\begin{aligned} c_0 &= 1, \\ p_0 &= 0, \\ p_t &\leq p_{t-1} + c_{t-1}, \\ c_t &\leq c_{t-1} + 2p_t \end{aligned}$$

(cf. [3]). This implies that for all  $p$  processors to be informed of the bit,  $\log_{2+\sqrt{3}} p \pm O(1)$  steps must be made. Incidentally, we note that the EXCLUSIVE WRITE rule is not needed in this argument; rather concurrent writing with any conflict resolution rule may be permitted.

We want to generalize this lower bound to a PRAM in which concurrent read accesses or write accesses are not forbidden but rather are resolved by the ARBITRARY rule.

**ARBITRARY READ:** If in a step several processors try to read from the same cell, an arbitrary one of them is given the contents of the cell, the other ones receive a negative acknowledgement (“reading failed”);

**ARBITRARY WRITE:** If in a step several processors write to the same cell, an arbitrary one of them succeeds *and* is given an acknowledgement (“writing successful”), the other ones are given a negative acknowledgement (“writing failed”).



The reader should note that for reading, this is an unusual rule in the context of PRAMs. It has only been proposed in the context of distributed memory machines (DMMs) as a relaxation of the COLLISION access rule, see [11, 24]. The DMM with COLLISION access rule, also known as OCPC, has been under intensive investigation in recent years, and it has been shown that such machines are very strong in a randomized setting; in particular they are able to perform routing and simulate PRAMs in sublogarithmic time [15, 16, 24]. Here, we show that the deterministic version of the model has a significant weakness: the elementary task of broadcasting one bit to  $p$  processors takes  $\Omega(\log p)$  steps.

If we try to translate the rules for such ARBITRARY-READ ARBITRARY-WRITE PRAMs into communication rules for a network as above for the EREW PRAM, we note that this time we must take 2-way communication into account, because of the acknowledgements received by the processors that write. Thus, we must describe communication rules for processor networks without restriction  $(*)^2$  but allowing full-duplex use of links in one step. The rules for step  $t$  in such a network are as follows. Depending on its state after step  $t - 1$ , each processor  $P_i$  fixes a message  $m_{i,t}$  and a set  $V_{i,t} \subseteq \{j \mid \{i, j\} \in E\}$  of possible partners. Let  $V_t = \{i \in V \mid V_{i,t} \neq \emptyset\}$  and  $E_t = \{\{i, j\} \mid i \in V_{j,t} \text{ and } j \in V_{i,t}\}$ , and consider the graph  $G_t = (V_t, E_t)$ . A maximal (i.e., not extendible) matching  $M_t \subseteq E_t$  is chosen arbitrarily, and for each edge  $\{i, j\} \in M_t$  message  $m_{i,t}$  is delivered to  $P_j$  and  $m_{j,t}$  is delivered to  $P_i$ . Processors that have received a message change their state based on this message; processors that have not received anything change their state based on this knowledge. At the beginning, i.e., before step 1, processor  $P_{i_0}$  knows the bit  $b$ ; after step  $T$ , all processors must know bit  $b$  no matter in which way the sets  $M_t$  were chosen.

**Definition 5.5** *The 2-way broadcast complexity  $\hat{T}_{G, i_0}^2$  of a processor network  $G$  with source  $P_{i_0}$  is the smallest number of steps a 2-way algorithm for broadcasting a bit from  $P_{i_0}$  can have.*

**Remark 5.6** In analogy to Remark 2.8 we note that this kind of algorithm can be performed on an asynchronous network without deadlocks and without processors waiting indefinitely for messages that will not arrive. As in the case of algorithms that obey restriction  $(*)^2$ , in its local step  $t$  processor  $P_i$  places its message  $m_{i,t}$  and the list  $V_{i,t}$  of its possible partners into a global buffer, together with a time stamp  $t$ . Whenever the buffer contains a matching pair of processors (i.e.,  $i \in V_{j,t}$  and  $j \in V_{i,t}$ ), after some finite delay the system delivers the messages  $m_{i,t}$  and  $m_{j,t}$  and removes the communication requests. Ties are broken arbitrarily. Whenever it happens that  $P_i$  has submitted a list  $V_{i,t}$ , but all  $P_j$  with  $j \in V_{i,t}$  either have specified  $V_{j,t} = \emptyset$  or have already found other partners for communicating in step  $t$ , the system informs  $P_i$  that its communication attempt in step  $t$  has failed and  $P_i$  proceeds to step  $t + 1$ . It is easily seen that the set  $\{\{i, j\} \mid P_i \text{ exchanges messages with } P_j \text{ in step } t\}$  is a maximal matching  $M_t$  in  $E_t$ , as required.

The 1-way algorithm of Observation 5.3 can also be run in 2-way mode, since no concurrent writing is used. Thus, we have the following.

**Observation 5.7**  $\widehat{T}_{G,i_0}^2 \leq b_2(G, i_0)$ . ■

However, one can also prove the (again intuitively obvious) fact that for the broadcast problem 2-way communication does not help.

**Theorem 5.8**  $\widehat{T}_{G,i_0}^2 = b_2(G, i_0)$ .

*Proof.* (Sketch) We must only prove that  $b_2(G, i_0) \leq \widehat{T}_{G,i_0}^2$ . Let a 2-way algorithm be given that performs broadcast from  $P_{i_0}$ . We proceed similarly as in the proof of Theorem 5.4, and only indicate the changes. First, we consider undirected edges. Computations  $C_0$  and  $C_1$  are fixed simultaneously by induction on  $t$ . If  $E_t(b) = \{\{i, j\} \mid i \in V_{i,t}(b) \text{ and } j \in V_{i,t}(b)\}$ , for  $b \in \{0, 1\}$ , is the collection of possible pairs of processors to communicate in step  $t$ , we first fix a maximal matching  $E'_t$  in  $E''_t = E_t(0) \cap E_t(1)$  and then let  $E'_t(b) \supseteq E''_t$  be a maximal matching in  $E_t(b)$ , for  $b = 0, 1$ . Next, we choose labeled edges. Edge  $\{i, j\}$  is put into  $E_{0,1}$  with label  $t$  if there is some  $b \in \{0, 1\}$  such that in  $C_b$  this edge is used for exchanging messages  $m_{i,t}(b)$  and  $m_{j,t}(b)$  but in the other computation  $C_{\bar{b}}$  this edge either is not used or different messages flow across it. We need the following observations.

*Claim 1:* If  $i \neq i_0$  and  $j \neq i_0$ , and  $\{i, j\}$  has label  $t$ , then there is some  $l \in V$  such that either  $\{l, i\}$  or  $\{l, j\}$  is in  $E_{0,1}$  with a label  $t' < t$ .

*Claim 2:* If  $i \neq i_0$ , then node  $i$  is incident with some labeled edge.

The proofs of these claims are similar as in the 1-way case. Now, we eliminate some labeled edges as follows. For each node  $i \neq i_0$ , choose the unique incident labeled edge with the smallest label, and direct this edge towards  $i$ . The resulting set of  $n - 1$  directed edges is called  $E_{0,1}^*$ . Using Claim 1, it is easily seen that if  $(i, j) \in E_{0,1}^*$  with label  $t$ , then either  $i = i_0$  or there is an edge  $(l, i) \in E_{0,1}^*$  with label  $t' < t$ . This implies that  $E_{0,1}^*$  forms a directed spanning tree of  $G$  with root  $i_0$ . Again, by the definition of  $E_{0,1}$ , at most two directed edges with the same time stamp can leave any node. This means that from  $E_{0,1}^*$  we obtain a 2-broadcast protocol by defining  $E_t = \{(i, j) \mid (i, j) \in E_{0,1}^* \text{ and } (i, j) \text{ has label } t\}$ . ■

We leave it to the reader to work out the details of the following argument. We wish to apply Theorem 5.8 to PRAMs with ARBITRARY READ and ARBITRARY WRITE with acknowledgement. To this end, we model both processors and cells as processors in a network. The aim is to inform all  $p$  processors of one bit initially located in one cell. If we carry out the construction of the edge set  $E_{0,1}^*$  as in the previous proof, it turns out that the resulting tree splits into levels, which alternately correspond to PRAM processors only and to PRAM cells only. Due to the special properties of the read operation there can be at most one edge

with time stamp  $t$  running from a “cell node” to a “processor node” at a higher-numbered level, whereas from a “processor node” two edges with the same time stamp may emanate. This implies that the recurrence inequalities mentioned at the beginning of this section also hold here. In this way we obtain the following generalization of the main result from [3].

**Theorem 5.9** *Broadcasting a bit from one cell to all  $p$  processors in a PRAM with the ARBITRARY READ rule and the ARBITRARY WRITE rule with acknowledgement takes  $\log_{2+\sqrt{3}} p \pm O(1)$  steps. The same bound holds for broadcasting a bit on a DMM with  $p$  processors and  $p$  memory modules. ■*

## References

- [1] A. BAGCHI, E. F. SCHMEICHEL, AND S. L. HAKIMI, Parallel information dissemination by packets, *SIAM J. Comput.* **23** (1993) 355–372.
- [2] A. BAR-NOY, S. KIPNIS, AND B. SCHIEBER, An optimal algorithm for computing census functions in message-padding systems, *Parallel Processing Letters* **3** (1993) 19–23.
- [3] P. BEAME, M. KUTYŁOWSKI, AND M. KIK, Information Broadcasting by exclusive-read PRAMs, *Parallel Processing Letters* **1 & 2** (1994) 159–169.
- [4] G. BELTING, *Untere Schranken für die Berechnung von Booleschen Funktionen in vollständigen Prozessornetzwerken im Telefon- und Telegraf-Modus*, Diplomarbeit, Universität-Gesamthochschule-Paderborn, Paderborn, 1994.
- [5] J. BRUCK AND C.-T. HO, Efficient global combine operations in multi-port message-passing systems, *Parallel Processing Letters* **3** (1993) 335–346.
- [6] S. BUBLITZ, U. SCHÜRFELD, I. WEGENER, AND B. VOIGT, Properties of complexity measures for PRAMs and WRAMs, *Theoret. Comput. Science* **48** (1986) 53–73.
- [7] S. COOK, C. DWORK, AND R. REISCHUK, Upper and lower time bounds for parallel random access machines without simultaneous writes, *SIAM J. Comput.* **15** (1986) 87–97.
- [8] D. W. KRUMME, G. CYBENKO, AND K. N. VENKATAMARAN, Gossiping in minimal time, *SIAM J. Comput.* **21** (1992) 111–139.
- [9] M. DIETZFELBINGER, M. KUTYŁOWSKI, AND R. REISCHUK, Exact time bounds for computing Boolean functions on PRAMs without simultaneous writes, in: Proc. 2nd Annual ACM Symp. on Parallel Algorithms and Architectures, 1990, pp. 125–135.

- [10] M. DIETZFELBINGER, M. KUTYŁOWSKI, AND R. REISCHUK, Exact lower bounds for computing Boolean functions on CREW PRAMs, *J. Comput. Syst. Sci.* **48** (1994) 1231–254.
- [11] M. DIETZFELBINGER AND F. MEYER AUF DER HEIDE, Simple, efficient shared memory simulations, in: Proc. 5th ACM Symp. on Parallel Algorithms and Architectures, 1993, pp. 110–118.
- [12] S. EVEN AND B. MONIEN, On the number of rounds necessary to disseminate information, in: Proc. ACM Symp. on Parallel Algorithms and Architectures, 1989, pp. 318–327.
- [13] F. FICH, The complexity of computation on the parallel random access machine, in J. H. Reif (ed.), *Synthesis of Parallel Computation*, Morgan Kaufmann, San Mateo, 1994, pp. 843–899.
- [14] P. FRAIGNIAUD AND E. LAZARD, Methods and problems of communication in usual networks, *Discrete Applied Math.* **53** (1994) 79–134.
- [15] L. A. GOLDBERG, M. JERRUM, T. LEIGHTON, AND S. RAO, A doubly logarithmic communication algorithm for the completely connected optical communication parallel computer, in: Proc. 5th Annual ACM Symp. on Parallel Algorithms and Architectures, 1993, pp. 300–309.
- [16] L. A. GOLDBERG, Y. MATIAS, AND S. RAO, An optical simulation of shared memory, in: Proc. 6th Annual ACM Symp. on Parallel Algorithms and Architectures, 1994, pp. 257–267.
- [17] S. M. HEDETNIEMI, S. T. HEDETNIEMI, AND A. L. LIESTMAN, A survey of gossiping and broadcasting in communication networks, *Networks* **18** (1986) 319–349.
- [18] I. HÖLTRING, *Broadcast und Gossip in parallelen Netzwerken*, Diplomarbeit, Universität–Gesamthochschule–Paderborn, Paderborn, 1994.
- [19] J. HROMKOVIČ, C.-D. JESCHKE, AND B. MONIEN, Optimal algorithms for dissemination of information in some interconnection networks, *Algorithmica* **10** (1993) 24–40.
- [20] J. HROMKOVIČ, R. KLASING, B. MONIEN, AND R. PEINE, Dissemination of information in interconnection networks (broadcasting & gossiping), in: D.-Z. Du and D. F. Hsu (eds.), *Combinatorial Network Theory*, Kluwer Academic Publishers, Amsterdam, 1996, pp. 125–212.

- [21] R. M. KARP AND V. RAMACHANDRAN, Parallel algorithms for shared-memory machines, in J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science*, Vol. A, Algorithms and Complexity, Elsevier, Amsterdam, 1990, pp. 869–941.
- [22] R. KLASING, B. MONIEN, R. PEINE, AND E. STÖHR, Broadcasting in butterfly and deBruijn networks, *Discrete Applied Math.* **53** (1994) 183–197.
- [23] R. LABAHN AND I. WARNKE, Quick gossiping by multi-telegraphs, in: R. Bodendiek and R. Henn (eds.), *Topics in Combinatorics and Graph Theory*, Physica-Verlag, Heidelberg, 1990, pp. 451–458.
- [24] F. MEYER AUF DER HEIDE, C. SCHEIDELER, AND V. STEMANN, Exploiting storage redundancy to speed up randomized shared memory simulations, in: E. W. Mayr and C. Puech (eds.), Proc. 12th Annual Symposium on Theoretical Aspects of Computer Science (STACS 95), Lecture Notes in Computer Science 900, Springer, Berlin, 1995, pp. 267–278.
- [25] N. NISAN, CREW PRAMs and decision trees, *SIAM J. Comput.* **20** (1991) 999–1007.
- [26] D. RICHARDS AND A. L. LIESTMAN, Generalizations of broadcasting and gossiping, *Networks* **18** (1988) 125–138.
- [27] V. S. SUNDERAM AND P. WINKLER, Fast information sharing in a complete network, *Discrete Applied Math.* **42** (1991) 75–86.
- [28] G. TEL, *Topics in Distributed Algorithms*, Vol. 1 of *Cambridge International Series on Parallel Computation*, Cambridge University Press, Cambridge, 1991.
- [29] G. TEL, *Introduction to Distributed Algorithms*, Cambridge University Press, Cambridge, 1994.
- [30] U. VISHKIN AND A. WIGDERSON, Trade-offs between depth and width in parallel computation, *SIAM J. Comput.* **14** (1985) 303–314.