

Partial Multinode Broadcast and Partial Exchange Algorithms for d -Dimensional Meshes¹

EMMANOUEL A. VARVARIGOS* AND DIMITRI P. BERTSEKAS†

* Department of Electrical and Computer Engineering, University of California, Santa Barbara, California 93106; and † Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

In this paper we consider the partial multinode broadcast and the partial exchange communication tasks in d -dimensional meshes. The partial multinode broadcast in an N -processor network is the task in which each of $M \leq N$ arbitrary nodes broadcasts a packet to all the remaining $N - 1$ nodes. Correspondingly, in the partial exchange there are $M \leq N$ nodes that wish to send a separate, personalized packet to each of the other nodes. We propose algorithms for the d -dimensional mesh network that execute the partial multinode broadcast and the partial exchange communication tasks in near-optimal time. No assumption is made concerning the locations of the M source nodes. The communication algorithms proposed are "on line" and distributed. We further look at a dynamic version of the broadcasting problem, where broadcast requests are generated at random times. In particular, we assume that the broadcast requests are generated at each node of the mesh according to a Poisson distribution with rate λ . Based on the partial multinode broadcast algorithm, we propose a dynamic decentralized scheme to execute the broadcasts in this dynamic environment. We find an upper bound on the average delay required to serve each broadcast. We prove that the algorithm is stable for network utilization ρ close to 1, and the average delay is of the order of the diameter for any load in the stability region. © 1994 Academic Press, Inc.

1. INTRODUCTION

The processors of a multiprocessor system, when doing computations, often have to communicate intermediate results. The interprocessor communication time may be substantial relative to the time needed exclusively for computations, so it is important to carry out the information exchange as efficiently as possible.

Two of the most frequent communication tasks are the *multinode broadcast* (MNB) and the *total exchange* (TE). The first task involves broadcasting a packet (the same packet) from every node to all the other nodes. It arises, for example, in iterations of the form

$$x = f(x), \quad (1)$$

where each processor computes an entry (or some entries) of the vector x . At the end of each iteration it is

¹ Research supported by the NSF under Grant NSF-DDM-8903385 and by the ARO under Grant DAAL03-86-K-0171.

necessary that each processor broadcast the updated value of the component that it computes to all other processors in order for it to be used at the next iteration.

The total exchange is the communication task where each node has to send a *personalized* (different) packet to each one of the other nodes. An example where the total exchange arises is the transposition of a matrix, when each processor stores, say, a column of the matrix. Then every processor i has to send the (i, k) th entry of the matrix to processor k , for all k , which is a total exchange.

Algorithms to perform a MNB or TE have been studied by several authors under a variety of assumptions on the communication network connecting the processors. Saad and Schultz [18, 19] were the first to consider these problems and to propose corresponding routing algorithms. Johnsson and Ho [8] have developed minimum and nearly minimum completion time algorithms for routing problems similar to those of Saad and Schultz but using a different communication model and a hypercube network. Bertsekas *et al.* [2], and Bertsekas and Tsitsiklis [1] have used the communication model of Saad and Schultz to derive minimum completion time algorithms for a MNB or TE in a hypercube. Varvarigos and Bertsekas [24] considered a class of communication tasks, called isotropic tasks, in hypercubes and d -dimensional wraparound meshes, and devised algorithms which are optimal jointly with respect to completion time, average packet delay, and storage requirements. This class of tasks includes the TE as a special case. The same authors in [26] proved that the multinode broadcast task when packets have random lengths can be executed in near-optimal time with high probability. Several other works deal with various communication problems and network architectures related to those discussed in the present paper; see [4-7, 10, 12, 13, 17, 22].

In iterations of the kind given in (1) it is very probable that only some of the components of the vector x change appreciably during an iteration. As these iterations approach their convergence point, fewer and fewer of the processors need to broadcast the updated values of the components of x that they compute. This gives rise to a task in which a strict (but unpredictable) subset of the processors have to broadcast a packet. We call this task a *partial multinode broadcast* (or PMNB for brevity). The

PMNB task, aside from being important on its own merit, is also a critical subroutine of the dynamic broadcast schemes that we propose in Section 5. The PMNB task arises also in clustering algorithms (see [16, Chap. 5], where the M nodes that store the coordinates of the centers of the clusters broadcast them after each iteration) and other problems. Because of its many applications we believe that the PMNB deserves a position among the prototype tasks of a communication library.

Similarly, during the transposition of a matrix that has both sparse and dense columns, it is more efficient if the nodes storing sparse columns do not participate in the TE, but send instead their packets as ordinary traffic through the 1-1 routing algorithm used by the machine. Since most large problems involve sparse matrices one can see that this situation arises frequently, giving rise to the *partial exchange tasks* (PE), where only M nodes send a (separate) packet to every other node. A task which is related to the PE is the *partial gather* (PG) task. In this task, M arbitrary nodes have to receive a (different) packet from every other node of the network (combining packets originated at different nodes is not allowed). Note that the PG task is *dual* to the PE task; if we find an algorithm to execute the PE we immediately get an algorithm of the same time complexity that executes the PG. In the transposition of a matrix stored by columns in a multiprocessor network, a PG arises when the matrix has only M dense rows. By combining a PE and a PG algorithm we get an algorithm that transposes a matrix which has M_1 dense columns and M_2 dense rows. The dense rows and columns can be arbitrary. This sparseness pattern arises very frequently in applications. The smaller M is, the less efficient a full MNB or TE algorithm would be and the more necessary it becomes to employ algorithms that are specially designed for partial tasks.

The main focus of this paper is to propose optimal and near-optimal communication algorithms for the partial multinode broadcast and the partial exchange tasks in d -dimensional meshes, with or without wraparound. These problems are considered here for the first time. PMNB algorithms have previously been studied for hypercubes in [20, 25]. The partial exchange problem was also considered for hypercube networks in [27]. In what follows, to avoid confusion, we call a (d -dimensional) mesh *with* wraparound a *torus* and a mesh *without* wraparound an *array*.

An algorithm that executes a task in the minimum number of steps is called *optimal*. We will say that an algorithm is *near-optimal* if the potential loss of optimality with respect to completion time is of strictly smaller order of magnitude than the optimal completion time itself. We will say that an algorithm is of *optimal order* if its worst case time complexity is asymptotically within a constant factor of the optimal value.

One of the main contributions of the paper is the development of near-optimal algorithms for a partial multinode

broadcast in a d -dimensional torus and in a d -dimensional array. We propose algorithms for two different communication models. In the first model, packets can be split and recombined at the destination without any overhead. In the second model the splitting is not allowed, and messages are always transmitted as one packet. We also present the first partial exchange algorithm of optimal order for the 2-dimensional array.

The PMNB communication task is a *static* broadcasting task, that is, it assumes that at time $t = 0$ some nodes have to broadcast a packet. In this paper, we also consider the dynamic version of this problem. We assume that broadcast requests are generated at each node according to a Poisson process with rate λ , independently of the other nodes. We propose routing schemes for d -dimensional tori and arrays that work under such a dynamic environment, and we evaluate their performance. The performance criterion used is the average packet delay, that is, the time that elapses on the average between the arrival of a packet to be broadcast at a node and the completion of the broadcast of the packet. Dynamic broadcasting schemes are also studied in the companion paper [25], and in [20] for hypercube networks. The scheme to be proposed for d -dimensional meshes is similar to the one given in [25] for hypercubes. It is stable for load asymptotically equal to the maximum possible, and its average delay is of the order of the diameter of the mesh for any load in the stability region.

The organization of the paper is the following. Section 2 shows how a mesh without wraparound can simulate a mesh with wraparound, and presents the first (strictly) optimal multinode broadcast algorithm for 2-dimensional meshes without wraparound. We also present a theorem concerning arbitrary broadcasts in rings and linear arrays, and two useful results by Nassimi and Sahnhi [14]. In Section 3 we present near-optimal algorithms to execute a partial multinode broadcast in a d -dimensional mesh. In particular, in Subsection 3.1 we give an algorithm where packets can be split, while in Subsection 3.2 we give an algorithm that does not require packet splitting. In Section 4 we present an algorithm for a partial exchange in a 2-dimensional array. In Section 5 we give the dynamic broadcasting scheme and evaluate its performance.

2. SOME PRELIMINARY RESULTS

The d -dimensional mesh, denoted by M_d , consists of $N = p^d$ processors arranged along the points of a d -dimensional space that have integer coordinates numbered from 0 to $p - 1$. Along the i th dimension, obtained by fixing coordinates $x_{d-1}, \dots, x_{i+1}, x_{i-1}, \dots, x_0$, there are p processors with identities $(x_{d-1}, \dots, x_i, \dots, x_0)$, $x_i = 0, 1, \dots, p - 1$. Two processors $(x_{d-1}, \dots, x_i, \dots, x_0)$ and $(y_{d-1}, \dots, y_i, \dots, y_0)$ are connected by a (two-directional) link if and only if for some i we have $|x_i - y_i| = 1$ and $x_j = y_j$ for all $j \neq i$. In addition to these links in the d -dimensional

mesh with wraparound (also called a torus), all links of the type

$$((x_{d-1}, \dots, x_{i+1}, 0, x_{i-1}, \dots, x_0), (x_{d-1}, \dots, x_{i+1}, p - 1, x_{i-1}, \dots, x_0))$$

are present. The latter links do not exist in the d -dimensional mesh without wraparound (also called an array). The set of nodes of an array (or torus) whose identities differ from the identity of node $x = (x_{d-1}, \dots, x_{i+1}, x_{i-1}, \dots, x_0)$ only in the i th digit is called the i -level linear array (or ring, respectively) of node x , and is denoted by $(x_{d-1}, \dots, x_{i+1}, *, x_{i-1}, \dots, x_0)$. The node with identity $(x_{d-1}, x_{d-2}, \dots, x_0)$ is also represented by the base p number of the form $x = x_{d-1}x_{d-2} \dots x_0$. The 0th digit is considered the least significant digit of the above representation. For two nodes x and y , we will write $x < y$ when the identity of x (viewed as a base p number) is smaller than the identity of y . A link connecting two nodes which differ only in the i th digit is called a link of dimension i .

Packets can be simultaneously transmitted along a link in both directions. Only one packet can travel along a link in each direction at any one time; thus, if more than one packet are available at a node and are scheduled to be transmitted on the same incident link of the node, then only one of these packets can be transmitted at the next time period, while the remaining packets must be stored at the node while waiting in queue. Each node is assumed to have infinite storage space. All incident links of a node

can be used simultaneously for packet transmission and reception. Each packet requires one unit of time for transmission over a link. We consider both a model where packets can be split at the origin, and be recombined at the destination, and a model where packets cannot be split; in the first model if a packet is split into d parts, each of them requires $1/d$ units of time to be transmitted over a link.

We start by describing how a torus can be simulated by an array. A linear (i.e., one-dimensional) array can simulate a ring of the same size with a slowdown factor of 2. This can be done as indicated in Fig. 1. By using this fact, a torus of any dimension can be simulated by an array of the same size and dimension with a slowdown factor of 2, as shown again in Fig. 1.

The optimal time T_{MNB}^t to execute a (full) multinode broadcast in a $p \times p$ torus was found in [15] (see also [1, pp. 81–82]) to be equal to

$$T_{MNB}^t = \frac{p^2}{4} = \frac{N}{4}$$

if p is even and

$$T_{MNB}^t = \frac{p^2 - 1}{4} = \frac{N - 1}{4}$$

if p is odd.

The following theorem gives a corresponding result for the MNB task in a 2-dimensional array.

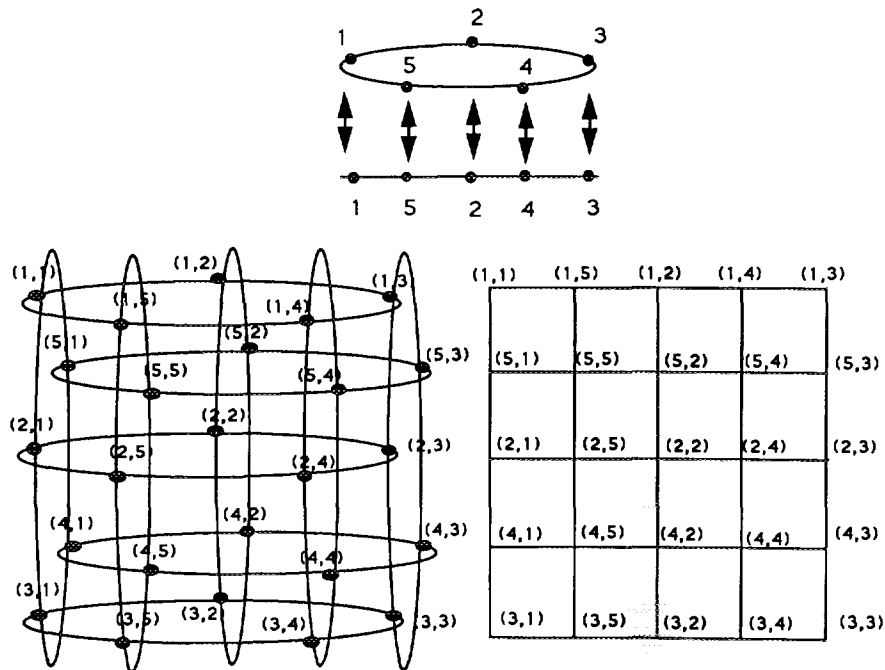


FIG. 1. The upper part of the figure shows how a ring can be simulated by a linear array with a factor of 2 slowdown. The nodes of the linear array are renamed as shown in the figure. Any communication between neighbors in the ring can be performed in two steps between the corresponding nodes in the linear array. This idea is easily extended to the simulation of a d -dimensional torus by a d -dimensional array, as can be seen from the lower part of the figure. The nodes along dimension i of the d -dimensional array form a linear array. The i th digit of the representations of these nodes is defined as in the simulation of a ring by a linear array (top of the figure).

THEOREM 1. *The minimum time T_{MNB}^a required to execute a (full) multinode broadcast in a 2-dimensional array is exactly twice the minimum time T_{MNB}^l required to execute a multinode broadcast in a 2-dimensional torus of the same size, that is,*

$$T_{\text{MNB}}^a = 2T_{\text{MNB}}^l = \left\lfloor \frac{N}{2} \right\rfloor. \quad (2)$$

Proof. As we indicated earlier, a mesh without wrap-around can simulate with a slowdown factor of 2 a mesh with wraparound of the same size. Each step of a torus can be simulated in two steps by an array even if all the links of the torus are simultaneously used. This gives the inequality $T_{\text{MNB}}^a \leq 2T_{\text{MNB}}^l$. Since node $(0, 0)$ has only two neighbors and receives $N - 1$ packets we have $T_{\text{MNB}}^a \geq (N - 1)/2$. This together with the fact that T_{MNB}^a has to be integer proves that

$$T_{\text{MNB}}^a = 2T_{\text{MNB}}^l = \left\lfloor \frac{N}{2} \right\rfloor.$$

Q.E.D.

The next theorem deals with arbitrary broadcasts in rings and arrays.

THEOREM 2. *Consider a linear (one-dimensional) array of p nodes, where each node has a certain (not necessarily the same) number of packets to broadcast to all other nodes. Let K be the total number of packets in the array. Then the broadcasts can be completed in time less than or equal to*

$$K + p - 1.$$

In a ring of the same size, the task requires half this time, provided that packets can be split into two parts without additional overhead.

Proof. Consider the following algorithm. Each node immediately transmits over its left (right) neighbor every packet that it receives from its right (left) neighbor. Whenever a node does not receive anything from its left (right) neighbor, it sends one of its own packets to the right (left). In other words, each node passes in the same direction the packets that come to it, and inserts a packet of its own whenever it sees an empty slot. Note that a packet is never delayed after it starts getting transmitted. In order to evaluate the time complexity we can focus on one direction, say the one going from left to right. Since there are K packets in the linear array, the packet can be delayed at most K times before starting transmission in this direction, and after at most $p - 1$ slots it will have arrived to all the nodes in that direction. To prove the result about the ring, we can split each packet in two parts, each requiring 0.5 units of time. The ring can be viewed as two edge-disjoint unidirectional linear arrays,

and by similar arguments, applied to each direction, the result follows. Q.E.D.

Remark. In the case where each node of the linear array has at most one packet, all the nodes can broadcast their packet in time $p - 1$ (see [1]). In the case of a ring the same task requires time $\lceil p/2 \rceil$, if the packets cannot be split, and $(p - 1)/2$, if the packets can be split into two parts without overhead.

The following two theorems are proved in [14]. The first deals with a problem called the packing routing problem.

THEOREM 3 [14] (Mesh Packing Routing). *Let $s^{(i)}$, $i = 0, 1, \dots, K - 1$, be nodes of a d -dimensional array such that $s^{(0)} < s^{(1)} < \dots < s^{(K-1)}$. Consider the communication task where each node $s^{(i)}$ sends a packet to processor i . This can be done without conflicts through a greedy scheme in time $d(p - 1)$. This greedy scheme uses only links of dimension j during steps $jp, jp + 1, \dots, jp + p - 1$, $j = 0, \dots, d - 1$.*

The next theorem, also proved in [14], treats a more general routing problem, which is called the *mesh monotone routing problem*.

THEOREM 4 [14] (Mesh Monotone Routing). *Let $s^{(i)}$ and $v^{(i)}$, $i = 0, 1, \dots, K - 1$, be nodes of a d -dimensional array such that $s^{(0)} < s^{(1)} < \dots < s^{(K-1)}$ and $v^{(0)} < v^{(1)} < \dots < v^{(K-1)}$. Consider the communication task where each node $s^{(i)}$ has to send a packet to processor $v^{(i)}$. This can be performed through a greedy scheme, without conflicts, in time $2d(p - 1)$. The greedy scheme uses only links of dimension j during steps $jp, jp + 1, \dots, jp + p - 1$, with $0 \leq j \leq d - 1$, and only links of dimension $2d - j$ during steps $jp, jp + 1, \dots, jp + p - 1$, with $d \leq j \leq 2d - 1$.*

Theorems 3 and 4 assume that packets $s^{(i)}$ know their rank i . The rank can be computed in time $2(p - 1)dt_p$, where t_p is the time required for a single parallel prefix step, through a parallel prefix operation as explained in various references (see, e.g., [11, pp. 37–44]), and described briefly in Phase 1 of the PMNB algorithm given in the next section.

3. PARTIAL MULTINODE BROADCAST IN D -DIMENSIONAL TORI AND ARRAYS

In this section we consider the problem where M arbitrary nodes of a d -dimensional mesh with $N = p^d$ nodes want to broadcast a packet to all the other nodes. We call these M nodes *active nodes*. Let T_{PMNB}^l be the optimal time required for the partial multinode broadcast in a d -dimensional torus, and let T_{PMNB}^a be the corresponding time for a d -dimensional array. T_{PMNB}^l and T_{PMNB}^a may actually depend on the identities of the M nodes that want to broadcast. A lower bound, however, is always

$$T_{\text{PMNB}}^t \geq \frac{M-1}{2d} \quad (3a)$$

and

$$T_{\text{PMNB}}^a \geq \frac{M-1}{d} \quad (3b)$$

where d is the dimension of the mesh. To see that, note that in a d -dimensional array (or torus) node $00 \dots 0$ has only d input ports (or $2d$ input ports, respectively), and has to receive at least $M-1$ packets.

Nassimi and Sahni [14] considered the more general broadcasting problem where some nodes of the network broadcast a packet to a set of nodes of the network, and each node receives at most M packets. If we use the algorithm of [14] to execute the PMNB task, its time complexity is (using our notation) $O(d^2p + Mdp)$, which is not of optimal order.

One way to execute the partial multinode broadcast is to perform a full multinode broadcast (with dummy packets for the nodes that have nothing to broadcast). A full MNB in a d -dimensional torus with $N = p^d$ nodes, when each packet requires one time unit (or slot) to be transmitted over a link, will require at least

$$T_{\text{MNB}}^t \geq \frac{N-1}{2d} \quad (4a)$$

time slots (every node has $2d$ links and has to receive $N-1$ packets), while a MNB in a d -dimensional array will require at least

$$T_{\text{MNB}}^a \geq \frac{N-1}{d} \quad (4b)$$

time slots. The gap between the right hand sides of Eqs. (3a) and (4a), and between the right hand sides of Eqs. (3b) and (4b), suggests that the use of a MNB algorithm to execute the PMNB task will be inefficient when $M \ll N$.

In this section we present communication algorithms that execute the PMNB task in d -dimensional meshes with or without wraparound in near-optimal time. In Subsection 3.1 we present an algorithm which assumes that packets can be split at the origin, and be recombined at the destination without any overhead. This algorithm executes the PMNB task in time at most

$$\frac{M}{2d} \frac{N-1}{N} + 2d(p-1)t_p + 1.5(p-1) \quad (5a)$$

for a d -dimensional torus with $N = p^d$ nodes, and in time at most

$$\frac{M}{d} \frac{N-1}{N} + 2d(p-1)t_p + 2(p-1) \quad (6a)$$

for a d -dimensional array of the same size, where t_p is the time required for a single parallel prefix step. For the case where the splitting of packets is undesirable (because of the overhead introduced and the cost of packet reassembling), we will present in Subsection 3.2 an algorithm that avoids the splitting of packets, and executes the PMNB task in time less than

$$\left\lceil \frac{M}{d} \right\rceil \frac{1}{p-1} \left\lceil \frac{p-1}{2} \right\rceil \frac{N-1}{N} + (p-1)d + d \left\lceil \frac{p-1}{2} \right\rceil + 4(p-1)t_p \quad (5b)$$

for a d -dimensional torus and in time less than

$$\left\lceil \frac{M}{d} \right\rceil + 2(p-1)d - 1 + 4(p-1)t_p \quad (6b)$$

for a d -dimensional array. Comparing Eqs. (5a)–(5b) and Eqs. (6a)–(6b) with the lower bounds (3a) and (3b), respectively, we see that the leading terms of the corresponding right hand sides have the same coefficient. So, the algorithms to be proposed are near-optimal.

3.1. A Near-Optimal PMNB Algorithm with Splitting of Packets

The algorithm in this section assumes that packets can be split at the origin, and recombined at the destination without any overhead. Each packet requires one time slot for transmission over a link. If a packet is split in d parts, each of these parts requires $1/d$ time units to be transmitted over a link.

Let s_1, s_2, \dots, s_M , $M \leq N$, be the active nodes. The rank of a packet located at node s is defined as

$$r_s = \sum_{t < s} x_t - 1,$$

where x_t is equal to one if processor t has a packet to broadcast and zero otherwise.

We will first present a suboptimal partial multinode broadcast algorithm for the d -dimensional mesh, with or without wraparound. This algorithm will not make full use of the links of a mesh. We will then modify the algorithm to achieve efficient link utilization and near-optimal completion time. The suboptimal algorithm consists of three phases:

Phase 1 (Rank Computation Phase). The rank r_s ($0 \leq r_s \leq M-1$) of each active node s is computed. This can be done in $2(p-1)d$ steps for a d -dimensional array or a torus by performing a *parallel prefix operation* (see [11, pp. 37–44]) on a tree P , called a *parallel prefix tree*, embedded in the mesh. The i th leaf of the tree from the leaf is the i th node of the mesh. The operation is described in Fig. 2 for a linear array and a mesh with $p = 3$ and $d = 2$.

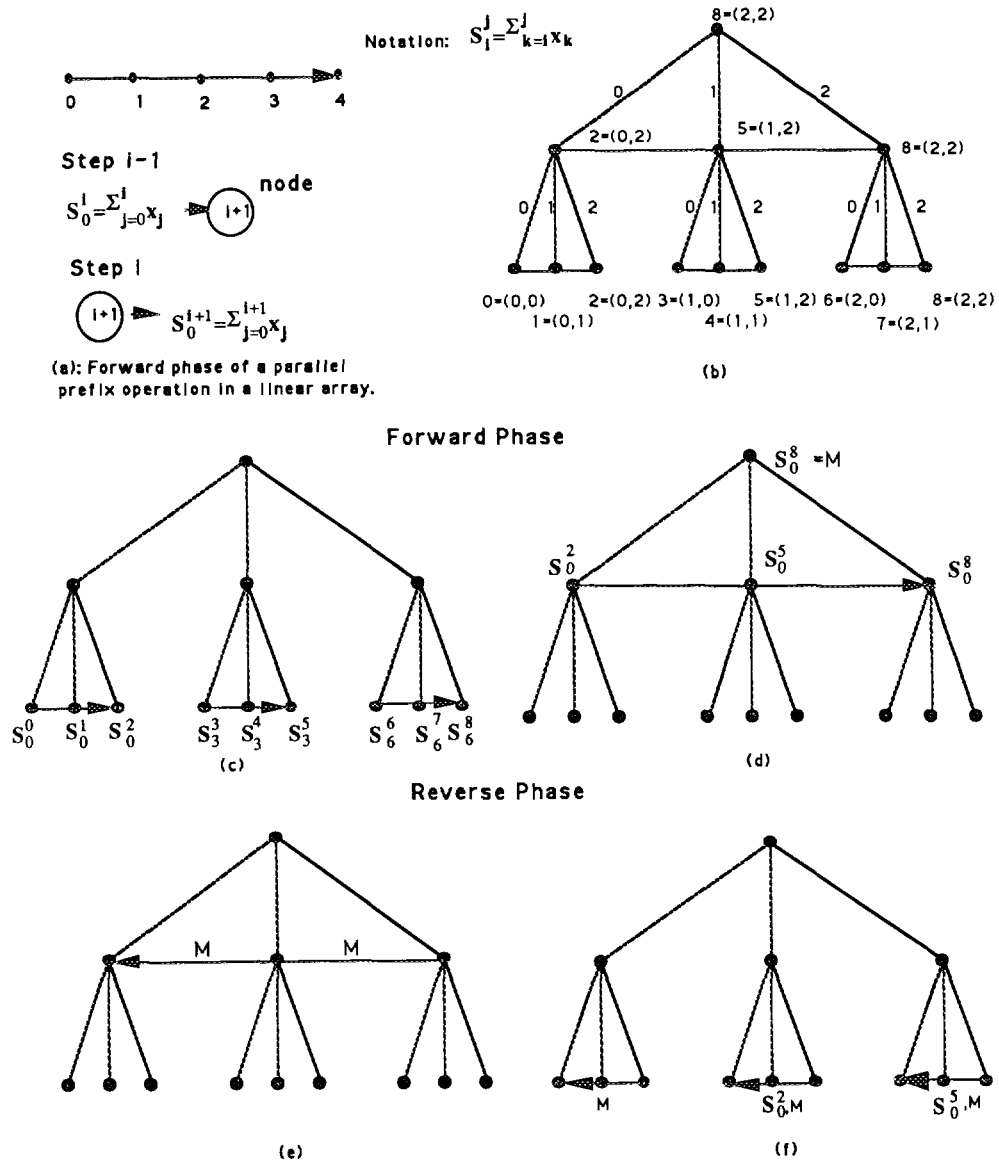


FIG. 2. (a) illustrates the operation of each node during a (forward) parallel prefix operation in a linear array. The partial sums $\sum_{k=0}^i x_k$ are obtained at each node i in time $p - 1$. (b)–(f) illustrate the parallel prefix operation in a mesh with $d = 2$ and $p = 3$. It consists of two phases (forward and reverse), each of which consists of d subphases. Each subphase is a parallel prefix operation in a linear array and requires $p - 1$ steps. The total duration of the operation is $2d(p - 1)$ steps. More precisely, (b) illustrates what we call tree representation of a mesh. An intermediate node is a root of a subtree whose leaves form a submesh of the original mesh. At the end of subphase l of the forward phase a node of level l from the bottom forms the partial sum of the values of the leaves under it. During the forward phase information moves from the bottom to the top, and from the left to the right. The notation S_i^j stands for $S_i^j = \sum_{k=i}^j x_k$. In the reverse phase, information moves from the top to the bottom and from the right to the left.

Note that during each step only links of a particular dimension are used. The packets involved in a parallel prefix operation are small (one byte of information), and require only t_p time units to be transmitted over a link. Thus it is reasonable to assume that $t_p \leq 1$, where one time unit is the time required to transmit a whole packet over a link; in fact it is reasonable to expect that in many parallel machines we have $t_p \ll 1$. Thus Phase 1 takes $2(p - 1) dt_p$ time units to be completed.

Phase 2 (Packing Phase). The packet of node s and rank r_s is sent to processor r_s , where r_s is interpreted as a p -ary number. This is a mesh packing problem, and can

be performed in $(p - 1)d$ time units according to Theorem 4.

Phase 3 (Broadcast Phase). The broadcast phase consists of d subphases $l = 1, 2, \dots, d$. During each subphase l , every node $r = r_{d-1}r_{d-2} \dots r_0$ broadcasts (in any order) to all the nodes in the ring or linear array (depending on whether we are considering a mesh with or without wraparound) $(r_{d-1} \dots r_{d-l+1} * r_{d-l-1} \dots r_0)$, the packets that were located at the node at the beginning of Phase 3 plus the packets that the node has received during all the previous subphases. The broadcast algorithms used are those described in Theorem 2.

During subphase 0 the nodes have (at most) one packet and this is the only one they broadcast. The current subphase l is easily known to each processor.

To prove that the algorithm delivers the packets to all the nodes, it is useful to introduce some new notation. Let $\beta = \beta_{d-1}\beta_{d-2}\cdots\beta_0$ be a p -ary number of length d . We denote by $S_l(\beta) = (*^l\beta_{d-l-1}\beta_{d-l-2}\cdots\beta_0)$ the submesh of the nodes whose $d-l$ less significant digits are equal to the $d-l$ less significant digits of β .

The next theorem proves that the previous algorithm actually executes the PMNB task.

THEOREM 5. *For each $\beta \in \{0, 1, \dots, p-1\}^d$, at the end of subphase l of Phase 3, $l = 1, 2, \dots, d$, each node in submesh $S_l(\beta)$ has received a copy of every packet located at the beginning of Phase 3 at some node in $S_l(\beta)$, completing a PMNB within each of these submeshes.*

Proof. The proof will be done by induction on l . For $l = 0$ (i.e., at the beginning of Phase 3 of the algorithm) it holds trivially since every node has its own (if any) packet. Assume it is true for some l . Every submesh $S_l(\beta)$ is composed of the p submeshes $S_{l-1}(\beta_{d-1}\cdots\beta_{d-l+1}0\beta_{d-l-1}\cdots\beta_0)$, $S_{l-1}(\beta_{d-1}\cdots\beta_{d-l+1}1\beta_{d-l-1}\cdots\beta_0)$, ..., $S_{l-1}(\beta_{d-1}\cdots\beta_{d-l+1}(p-1)\beta_{d-l-1}\cdots\beta_0)$. During subphase l every node in one of these submeshes broadcasts to all nodes in its $(d-l)$ -level linear array (or ring) all the packets it has received during the previous subphases, together with its own packet. This together with the induction hypothesis proves the theorem. Q.E.D.

Letting $l = d$ we find that at the end of subphase d each packet has been broadcast to all the nodes, and therefore, the PMNB has been completed.

The next lemma calculates the time complexity of Phase 3.

LEMMA 1. *Phase 3 of the algorithm requires at most*

$$\frac{N-1}{N} \frac{M}{\gamma} + \frac{(p-1)d}{\gamma}$$

time units, where $\gamma = 1$ for the d -dimensional array, and $\gamma = 2$ for the d -dimensional torus.

Proof. We denote by T_l the duration of subphase l , and we let $m = \lceil \log_p M \rceil$. At the beginning of Phase 3 only nodes $0, 1, \dots, M-1$ have a packet. From Theorem 5 we know that just before the beginning of phase l , node $s = s_{d-1}s_{d-2}\cdots s_0$ has received all the packets originally located at nodes in the submesh $(*^{l-1}s_{d-1}s_{d-l-1}\cdots s_0)$. The number of these packets is equal to the cardinality of the set

$$\mathcal{W}_l(s) = \{w = w_{d-1}w_{d-2}\cdots w_0 \mid 0 \leq w \leq M-1, \\ w_{d-l} = s_{d-l}, w_{d-l-1} = s_{d-l-1}, \dots, w_0 = s_0\}.$$

During subphase l , node s will broadcast these packets to the nodes in its $(d-l)$ -level linear array or ring. Since a multinode broadcast in a linear array requires $p-1$

steps, while in a ring it requires $(p-1)/2$ steps (see the remark following Theorem 5), we have

$$T_l \leq \frac{p-1}{\gamma} \max_s |\mathcal{W}_l(s)|,$$

where $\gamma = 1$ for d -dimensional arrays, $\gamma = 2$ for d -dimensional tori, and $|\cdot|$ denotes the cardinality of a set. Let $s' = s_{d-l}p^{d-l} + s_{d-l-1}p^{d-l-1} + \cdots + s_0$. The cardinality of $\mathcal{W}_l(s)$ is equal to the cardinality of the set

$$\mathcal{W}_l(s) - s' = \{v \mid -s' \leq v \leq M-1-s', \\ v \text{ divisible by } p^{d-l+1}\},$$

or equivalently (since $s' < p^{d-l+1}$) equal to the number of integers between 0 and $M-1-s'$ which are divisible by p^{d-l+1} . Thus

$$\max_s |\mathcal{W}_l(s)| \leq \max_s \left\lceil \frac{M-s'}{p^{d-l+1}} \right\rceil \leq \left\lceil \frac{M}{p^{d-l+1}} \right\rceil.$$

The total duration of Phase 3 satisfies

Duration of Phase 3

$$\begin{aligned} &= \sum_{l=1}^d T_l \leq \frac{p-1}{\gamma} \sum_{l=1}^d \left\lceil \frac{M}{p^{d-l+1}} \right\rceil \\ &\leq \frac{p-1}{\gamma} \left(d + M \sum_{l=1}^d \frac{1}{p^l} \right) \\ &= \frac{(p-1)d}{\gamma} + \frac{M}{\gamma} \left(1 - \frac{1}{p^d} \right). \end{aligned}$$

Q.E.D.

Adding up the duration of Phases 1, 2, and 3, we obtain the following lemma:

LEMMA 2. *The partial multinode broadcast task can be executed in a d -dimensional torus with $N = p^d$ processors in*

$$T_{\text{PMNB}}^1 \leq \frac{M}{2} \frac{N-1}{N} + 2d(p-1)t_p + 1.5(p-1)d$$

time units, where M is the number of active nodes. Similarly, the PMNB task can be executed in a d -dimensional array with $N = p^d$ processors in

$$T_{\text{PMNB}}^a \leq M \frac{N-1}{N} + 2d(p-1)t_p + 2(p-1)d$$

time units.

The PMNB algorithm that we described so far is not of optimal order, as the gap between the lower bounds of Eqs. (3) and (4) and the results of Lemma 2 indicates. In fact, they are suboptimal by a factor of roughly d . This is

due to the fact that at each step only links of a particular dimension are used. In the next theorem we modify the algorithms so that all dimensions are used at the same time, and near-optimal completion time is achieved.

THEOREM 6. *The partial multinode broadcast task can be executed in a d -dimensional torus with $N = p^d$ processors in*

$$T_{\text{PMNB}}^l \leq \frac{M}{2d} \frac{N-1}{N} + V_l \quad (7)$$

time units, where M is the number of active nodes, and

$$V_l = 2d(p-1)t_p + 1.5(p-1).$$

Similarly, the PMNB task can be executed in a d -dimensional array with $N = p^d$ processors in

$$T_{\text{PMNB}}^a \leq \frac{M}{d} \frac{N-1}{N} + V_a \quad (8)$$

time units, where

$$V_a = 2d(p-1)t_p + 2(p-1).$$

Proof. We call the PMNB algorithm analyzed in Lemmas 1 and 2 Algorithm \mathcal{A}_0 . At each step of Phases 1, 2, and 3 of \mathcal{A}_0 , only links of a particular dimension are used. Indeed, it can be seen from Fig. 2 that during each step of the parallel prefix phase only links of a particular dimension are used. Similarly, in the packing phase, only links of a particular dimension are used at each step, as indicated in Theorem 4. Finally, during subphase l of the broadcast phase only links of dimension $d-l$ are used.

For any c , consider now another PMNB algorithm, referred to as algorithm \mathcal{A}_c . According to \mathcal{A}_c a packet is transmitted over the link of dimension $(l+c) \bmod d$ of its current location, whenever the same packet would be transmitted under the \mathcal{A}_0 algorithm over the l th-dimensional link of its current location. Since \mathcal{A}_c is identical to \mathcal{A}_0 after appropriately renaming the mesh dimensions (and the nodes), and since \mathcal{A}_0 performs the PMNB independently of the location of the M active nodes, we conclude that \mathcal{A}_c also executes the PMNB task, and requires the same amount of time as \mathcal{A}_0 .²

² In the algorithm \mathcal{A}_c the rank of an active node is defined in the following way. On the p -ary numbers of length d , we first define the order with respect to class c , $c \in \{0, 1, \dots, d-1\}$ (denoted by \prec_c), as follows: $s \prec_c t$ iff the right shift of s by c positions is less (with the usual order) than the right shift of t by c positions. The rank with respect to class c of a packet located at node s is then defined as

$$r_s^c = \sum_{\{t: t \prec_c s\}} x_t - 1,$$

where x_t is equal to one if processor t has a packet to broadcast and zero otherwise. The parallel prefix tree P^c used in the calculation of r_s^c is the same as P , but with the digits of the nodes shifted by c positions.

Using simultaneously all the algorithms $\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_{d-1}$ we can find a new algorithm which requires the amount of time claimed in the theorem. In particular, each packet is split into d parts, called *mini packets*. Each mini packet is assigned a distinct integer c between 0 and $d-1$, called class. The mini packets of class c are routed according to algorithm \mathcal{A}_c . Packets of different classes use different mesh dimensions at any time. According to our communication model, a mini packet requires $1/d$ time units for transmission over a link. Therefore, the theorem follows from Lemma 2. Q.E.D.

The terms V_l and V_a in Eqs. (7) and (8), respectively, are growing linearly with the dimension d . In practice, however, $2d(p-1)t_p$ is small, since t_p is very small. Indeed, at each step of a parallel prefix operation only one byte has to be transmitted between neighbors. Some parallel computers, such as the Connection Machine model CM-2 of Thinking Machines Corporation, the IBM/RP-3, and the NYU Supercomputer, have very efficient implementation of the parallel prefix, otherwise called "scan" operation [3, 23]. Theoretically, however, the parallel prefix operation takes time proportional to the diameter.

No upper ceilings are needed in Eqs. (7) and (8), since we allow fragmented slots. Note also that under the communication model used in this section (which allows the splitting of packets in d parts), a broadcast from a single node requires $\Theta(p)$ time units, instead of $\Theta(dp)$ which is the diameter. A near-optimal PMNB algorithm which does not use the splitting of packets is presented in the next subsection.

3.2. A Near-Optimal PMNB Algorithm without Splitting of Packets

In this subsection we modify the previous algorithms in order to avoid the potential drawbacks of packet splitting. This is done at the expense of a slight increase in the complexity. Messages in this section require one time slot in order to be transmitted over a link, and are always transmitted as one packet.

The algorithm makes use of the algorithm \mathcal{A}_c described in the proof of Theorem 6 (recall that algorithm \mathcal{A}_c consists of three phases: the rank computation, the packing, and the broadcast phases), and consists of two parts.

Class Computation Part. The rank r_s , $0 \leq r_s \leq M-1$, $s \in \{s_1, s_2, \dots, s_M\}$, of each packet is computed through a parallel prefix operation. This requires $2d(p-1)t_p$ time units. The packet of node s is assigned a class number $c = r_s \bmod d$.

Main Part. The packets of class c are routed according to algorithm \mathcal{A}_c . Only packets of class c take part in the rank computation phase, or in any other phase of \mathcal{A}_c .

To estimate the complexity of the algorithm, we first note that each class has at most $\lceil M/d \rceil$ packets. Lemma 1 has been proved under the assumption that a packet can

be split into two parts. When packets cannot be split, it can be shown (by a proof similar to that in Lemma 1) that algorithm \mathcal{A}_c requires time less than or equal to

$$M' \frac{1}{p-1} \left\lceil \frac{p-1}{\gamma} \right\rceil \frac{N-1}{N} + d \left\lceil \frac{p-1}{\gamma} \right\rceil,$$

where M' is the number of packets that participate in \mathcal{A}_c , $\gamma = 1$ for d -dimensional arrays, and $\gamma = 2$ for d -dimensional tori. Substituting $\lceil M/d \rceil$ instead of M' , and adding the time required for the parallel prefix operations and the packing routing phase, we get

$$T_{\text{PMNB}} \leq \left\lceil \frac{M}{d} \right\rceil \frac{1}{p-1} \left\lceil \frac{p-1}{\gamma} \right\rceil \frac{N-1}{N} + (p-1)d + d \left\lceil \frac{p-1}{\gamma} \right\rceil + 4(p-1) dt_p.$$

The algorithm just presented for the PMNB task gives rise to an efficient algorithm for the MNB task. Indeed, a multinode broadcast can be treated as a partial multinode broadcast with $M = N$. The class computation part and the rank computation phase are not necessary any more, since the class number and the rank of each packet are known in advance. The packing and the broadcast phases alone can execute the MNB in time less than or equal to

$$\left\lceil \frac{N}{d} \right\rceil \frac{1}{p-1} \left\lceil \frac{p-1}{\gamma} \right\rceil \frac{N-1}{N} + (p-1)d + d \left\lceil \frac{p-1}{\gamma} \right\rceil,$$

which is near-optimal for tori with p odd or arrays, and of optimal order for tori with p even. This MNB algorithm is apparently new.

4. PARTIAL EXCHANGE IN 2-DIMENSIONAL ARRAYS

In this section we present an algorithm to execute the partial exchange task in a $p \times p$ array (and, therefore also in a $p \times p$ torus). In particular, we initially assume that there are M nodes, called *active*, that want to send a personalized packet to each of the other nodes. The algorithm to be presented has time complexity which is of optimal order.

We first present lower bounds on the minimum time required to execute the partial exchange in 2-dimensional array. Let us define

$$m = \lfloor M^{1/2} \rfloor.$$

First, consider the case (Fig. 3a) where m^2 of the M active nodes are in the $m \times m$ subarray $M_{0,0}$, where

$$M_{0,0} = \{(i, j) \mid 0 \leq i \leq m-1, 0 \leq j \leq m-1\}.$$

Then $m^2(N - m^2)$ packets have to cross the $2m$ links connecting $M_{0,0}$ to the rest of the array (since there are no wraparound links). This gives

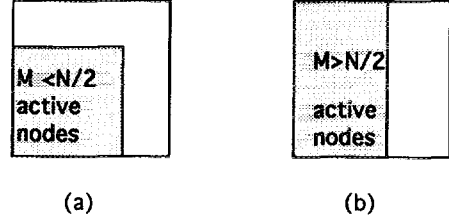


FIG. 3. Position (a) of the active nodes corresponds to the first lower bound, while position (b) corresponds to the second lower bound.

$$T_{\text{PE}} \geq \frac{m(N - m^2)}{2}.$$

Thus,

$$T_{\text{PE}} \geq \frac{mN}{4}, \quad \text{for } M \leq N/2.$$

We next consider the case where $M \geq N/2$, and all the active nodes are at the left side of the mesh (see Fig. 3b). Consider the packets that pass from left to right through the cut that bisects the mesh. At least $N^2/4$ packets cross the p links of this cut. Thus,

$$T_{\text{PE}} \geq \frac{Np}{4} \geq \frac{mN}{4}, \quad \text{for } M \geq N/2.$$

The previous bounds show that

$$T_{\text{PE}} \geq \frac{mN}{4} = \Omega(M^{1/2}N).$$

Before describing the algorithm, we introduce some notation. The nodes of the mesh are represented as pairs (i, j) with $0 \leq i, j \leq p-1$. The *row sum* (or *column sum*) of node (i, j) is defined as the number of active nodes of row i (or column j) and is denoted by r_i (or c_j , respectively).

We now describe the algorithm. It consists of two phases:

Phase 1 (Parallel Prefix Phase). The row sums r_i and column sums c_j are computed. All r_i 's and c_j 's can be found in p steps by concurrently performing a parallel prefix operation within each row and column. A row or column is a linear array, and can be viewed as a tree rooted at a median node of that linear array of depth $\lceil p/2 \rceil$. The parallel prefix operation is performed with value equal to one for active nodes and zero for the other nodes. Phase 1 requires pt_p time units, where $t_p \leq 1$ is the time required for a single parallel prefix step.

Phase 2 (Exchange Phase). The set of active nodes is partitioned into the two sets R and C , where

$$R = \{(i, j) \text{ active} \mid r_i \leq c_j\}, \quad C = \{(i, j) \text{ active} \mid r_i > c_j\}.$$

The nodes in R (or in C) send each of their packets along the unique shortest path that first crosses horizontal (respectively, vertical) links exclusively, and then crosses vertical (respectively, horizontal) links exclusively. The order of transmission of the packets at each link is arbitrary subject to two restrictions:

- (a) Packets originating at nodes of R (or of C) have priority on the horizontal (respectively, vertical) links over packets originating at nodes in C (respectively, R).
- (b) Transmission is nonwasting in the sense that no link remains idle if there is a packet waiting at the queue of the link.

We have the following lemma:

LEMMA 3. *The number of nodes of R that belong to the same row is at most m .*

Proof. Our proof is by contradiction. Suppose that for some i , the nodes $(i, j_1), (i, j_2), \dots, (i, j_x)$ belong to R , and $x > m$. Then by the definition of the set R , $c_{j_k} \geq r_i \geq x > m$ for all $k = 1, 2, \dots, x$. This implies that there are $x \geq m + 1$ columns, each of which has at least x active nodes. This is a contradiction since there are only $M < (m + 1)^2$ active nodes. Q.E.D.

We claim that Phase 2 requires at most $2m(N - p) + 2(p - 1)$ time units. To see this, we first note that the number of packet originating at nodes of R that must cross at least one horizontal link of any given row i is at most $m(N - p)$. The reason is that by Lemma 3 there are at most m active nodes from R in row i and each of these nodes has a total of $N - p$ packets to send to nodes that belong to a different column. Since each packet increases the delay of another packet by at most one unit along the horizontal path, we see that the time required for all the packets originating at nodes in R to traverse completely the horizontal portion of their path is at most $m(N - p) + (p - 1)$. Similarly, the time required for all the packets originating at nodes in C to traverse completely the vertical portion of their path is also at most $m(N - p) + (p - 1)$.

Let us make the worst-case assumption that packets originating at nodes of R (or C) are delayed after completing their horizontal (respectively, vertical) transmissions so that their transmission starts after exactly $m(N - p) + (p - 1)$ time units. We will show that at most $m(N - p) + (p - 1)$ additional time units are needed to complete Phase 2. Indeed, at the end of the first $m(N - p) + (p - 1)$ time units, each node has at most $m(p - 1)$ packets originating at nodes in R to send over the vertical links. Therefore, at most $m(p - 1)p$ such packets remain to traverse the links of its column. This requires at most an additional $m(N - p) + (p - 1)$ time units.

Adding up the times required for each phase, and taking into account that $N = p^2$ and $m = \lfloor M^{1/2} \rfloor$, we find that the time T_{PE} required for the partial exchange in a 2-

dimensional array satisfies

$$T_{PE} \leq 2\lfloor M^{1/2} \rfloor(N - p) + 2(p - 1) + pt_p.$$

Comparing this inequality with the lower bound found earlier, we see that our algorithm is of optimal order (within a factor of roughly 8 of being optimal).

The PE algorithm that we presented for 2-dimensional arrays does not seem to be easily extendable to d -dimensional arrays. If a similar idea were followed in the d -dimensional case, the M active nodes would have to be partitioned in d disjoint sets, and the packets of set i should traverse dimensions in the order $i \bmod d, i + 1 \bmod d, \dots, i + d - 1 \bmod d$. The partition should be done in a way that would guarantee that the congestion on the links would be small. It is not clear, however, how to find such a partition. Therefore, the problem of finding a PE algorithm of optimal order for a d -dimensional array is still an open problem.

5. DYNAMIC BROADCASTING SCHEMES

The PMNB task considered in Section 3 is *static* in the sense that it is executed only once, starting at time $t = 0$. In this section we consider the *dynamic* version of this task. We assume that broadcast requests arrive at each node of a mesh according to a Poisson process with rate λ . We propose an algorithm that works well in such a dynamic environment, and evaluate its performance. The performance criterion used is the average packet delay, that is, the time between the arrival of a packet to be broadcast at a node and the completion of the broadcast of the packet. The same problem has been studied in [20] and in a companion paper [25] where dynamic schemes are proposed for the hypercube case, and their performance is analyzed. The dynamic scheme that we propose for meshes is simple and stable for network utilization very close to the maximum possible. Furthermore, for each fixed utilization in the stability region, the average delay is of the order of the diameter of the mesh, which is the best we could hope for.

Our scheme is essentially a repetition of partial multinode broadcasts, each starting when the previous one has finished. The PMNB algorithm that we will assume throughout this section is the one of Subsection 3.1 where packets are split; if the algorithm of Subsection 3.2, where packets are not split, is used we get corresponding results (slightly worse, however, especially in the case where p is even). The time axis is divided into PMNB intervals (see Fig. 4). Within each PMNB interval, a PMNB is executed, involving exactly one packet from each of the nodes that have a packet to broadcast at the start of the interval. Each PMNB interval is divided into two parts. The first part is called *reservation interval*, and consists of the parallel prefix and the packing phases of the PMNB algorithm of Section 3. Its duration

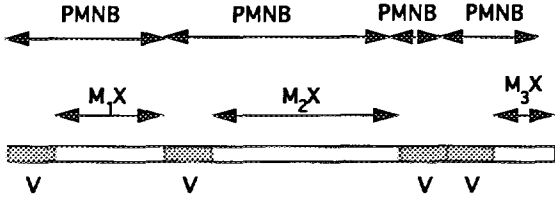


FIG. 4. The dynamic broadcasting scheme. Each PMNB interval consists of two intervals: a reservation interval (marked by gray) of duration V , and a broadcast interval of duration MX , where M is the number of active nodes at the start of the PMNB interval.

can be upper bounded by a known constant that depends only on the size of the network, and is independent of the number of active nodes M . The second part of a PMNB interval is called *broadcast interval*. Its duration is known once M is known. Thus, even though the duration of each partial multinode broadcast is random (because packet arrivals are random), it is known to all the nodes of the network, because each node learns during the broadcast interval the number M of active nodes and, from there, the duration of the following broadcast interval. Therefore, if the nodes initiate the dynamic broadcast scheme at the same time, no further synchronization is needed.

It is important for the performance of the dynamic scheme that the duration of the PMNB algorithms given in Section 4 is *linear* in the number of active nodes M , with the constant of proportionality being the smallest possible. In particular, the duration of the PMNB algorithm for a d -dimensional mesh with $N = p^d$ processors when we use the algorithm that allows the splitting of packets was found in Section 3 to satisfy

$$T_{\text{PMNB}} \leq XM + V,$$

where

$$X = \frac{1}{\gamma d} \frac{N-1}{N}, \quad (9a)$$

$$V = 2d(p-1)t_p + \left(1 + \frac{1}{\gamma}\right)(p-1), \quad (10a)$$

and $\gamma = 1$ for the d -dimensional array and $\gamma = 2$ for the d -dimensional torus. If the splitting of packets is not allowed, then X and V satisfy the following expressions:

$$X = \frac{1}{d} \frac{1}{p-1} \left\lceil \frac{p-1}{\gamma} \right\rceil \frac{N-1}{N}, \quad (9b)$$

$$V = (p-1)d + d \left\lceil \frac{p-1}{\gamma} \right\rceil + 4(p-1)dt_p + 1. \quad (10b)$$

In [25] the following theorem was proved (in fact a slightly stronger theorem was proved there, but the following simplified version is adequate for our purposes):

THEOREM 7 [25]. *Let the arrivals of broadcast requests at a node of an N -processor network be Poisson with rate λ . Assume also that there exists a PMNB algorithm for that network, which executes the PMNB task in time at most*

$$XM + V,$$

where M is the number of nodes that have a packet to broadcast and X, V are independent of M (they may depend on the size of the network). Let $\rho = \lambda NX$, and suppose that

$$1 - \rho - \lambda V > 0.$$

Then if this PMNB algorithm is used in a dynamic broadcasting scheme, in the way described earlier, the system is stable in the sense that the average packet delay is finite and satisfies

$$T \leq (1 + \rho) \left(\frac{\rho X}{2(1 - \rho - \lambda V)} + \frac{(1 - \rho)V}{2(1 - \rho - \lambda V)} + \frac{(1 - \lambda V)V}{1 - \rho - \lambda V} \right) + X. \quad (11)$$

For the d -dimensional mesh we have found algorithms that satisfy the conditions of Theorem 7. Thus, the average packet delay is bounded as in Eq. (11), with X and V given by Eqs. (9a) and (10a), when the splitting of packets is allowed, and by Eqs. (9b) and (10b), when the splitting of packets is not allowed.

The scalar

$$\rho = \lambda NX = \frac{\lambda(N-1)}{\gamma d} \quad (12)$$

is called the *mesh utilization factor*, for reasons that will become evident soon. To find necessary conditions for stability for any broadcasting scheme, consider a d -dimensional array or torus, the outgoing links of node $(00 \dots 0)$, and the traffic that passes through them. There are $2d$ such links for the torus, and d for the array. Thus, for stability we must have

$$\lambda(N-1) \leq \gamma d$$

or

$$\rho \leq 1, \quad (13)$$

no matter what broadcasting scheme we use. For the torus ($\gamma = 2$) and a given load, ρ is equal to the ratio of the average number of transmissions per unit of time necessary to execute the broadcasts (each broadcast requires $N-1$ transmissions), over the total number of links of the network. For the array ($\gamma = 1$), which is not a sym-

metric network, ρ equals the average fraction of time during which the links of node $(00 \dots 0)$ have to be used under any broadcasting scheme.

Our algorithm is guaranteed to be stable for $\rho < 1 - \lambda V$. Using Eq. (12) the stability condition becomes

$$\rho < 1 - \rho \frac{\gamma d V}{N - 1}.$$

If the splitting of packets is allowed, so that Eq. (9a) is used, the preceding expression becomes

$$\rho < 1 / \left(1 + \frac{d\gamma(2d(p-1)t_p + (1+1/\gamma)(p-1))}{p^d - 1} \right). \quad (14)$$

The right hand side of Eq. (14) is very close to the maximum possible load, given by Eq. (13), that a d -dimensional array or torus could sustain. As the number of nodes p^d tends to infinity, $d^2 p/p^d$ tends to zero. Thus, the right hand side of Eq. (14) tends to one, which is the maximum utilization that can be accommodated by the network.

For any fixed ρ in the stability region, Eq. (11) gives

$$T = O(V). \quad (15)$$

If the splitting of packets is allowed then, using Eq. (10a), we get

$$T = O(pdt_p + p).$$

If the splitting of packets is not allowed then V is given by Eq. (10b), and Eq. (15) becomes

$$T = O(pdt_p + pd).$$

The diameter of a d -dimensional mesh is $\Theta(pd)$. For any fixed ρ in the stability region, the average delay of the dynamic scheme is of optimal order of magnitude when packets are not split, in which case the diameter is a lower bound on any broadcast. Under the model where packets can be split in d parts, the diameter is no longer a lower bound on T (in this case $T = \Omega(p)$), and the average delay is suboptimal (unless t_p is so small that it can be neglected). Note that for very light load ($\lambda \approx 0$, $\rho \approx 0$) we get from Eq. (11) that

$$T \leq 1.5V + X \quad (\rho \approx 0).$$

REFERENCES

- Bertsekas, D. P., and Tsitsiklis, J. N. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- Bertsekas, D. P., Ozveren, C., Stamoulis, G. D., Tseng, P., and Tsitsiklis, J. N. Optimal communication algorithms for hypercubes. *J. Parallel Distrib. Comput.* **11**, 4 (Apr. 1991), 263-275.
- Blelloch, G. E. Scands as primitive parallel operations. *Proc. Int'l Conf. Parallel Processing*. Aug. 1986, pp. 355-362.
- Edelman, A. Optimal matrix transposition and bit reversal on hypercubes: All-to-all personalised communication. *J. Parallel Distrib. Comput.* **11**, 4 (Apr. 1991), 328-331.
- Fraigniaud, P. Complexity analysis of broadcasting in hypercubes with restricted communication capabilities. *J. Parallel Distrib. Comput.* **16**, 1 (Sept. 1992), 15-26.
- Hedetniemi, S. M., Hedetniemi, S. T., and Liestman, A. L. A survey of gossiping and broadcasting in communication networks. *Networks* **18**, pp. 319-349, 1988.
- Ho, C. T. Full bandwidth communications on folded hypercubes. Research Report RJ 7434 (69605), IBM Almaden Research Center, Apr. 1990.
- Johnsson, S. L., and Ho, C. T. Optimum broadcasting and personalized communication in hypercubes. *IEEE Trans. Comput.* **C-38** (1989), 1249-1268.
- Kermani, P., and Kleinrock, L. Virtual cut-through: A new computer communicating switching technique. *Comput. Networks* **3** (1979), 267-286.
- Krumme, D. W., Venkataraman, K. N., and Cybenko, G. The token exchange problem. Tufts University, Technical Report 88-2, 1988.
- Leighton, F. T. *Introduction to Parallel Algorithms and Architectures: Arrays—Trees—Hypercubes*. Morgan Kaufmann, San Mateo, CA, 1992.
- Lan, Y., Esfahanian, A.-H., and Ni, L. Multicast in hypercube multiprocessors. *J. Parallel Distrib. Comput.* **8**, 1 (Jan. 1990), 30-41.
- McBryan, O. A., and Van de Velde, E. F. Hypercube algorithms and their implementations. *SIAM J. Sci. Statist. Comput.* **8** (1987), 227-287.
- Nassimi, D., and Sahni, S. Data broadcasting in SIMD computers. *IEEE Trans. Comput.* **C-30** (1981), 101-106.
- Ozveren, C. Communication aspects of parallel processing. Laboratory for Information and Decision Systems Report LIDS-P-1721, Massachusetts Institute of Technology, Cambridge, MA, 1987.
- Ranka, S., and Sahni, S. *Hypercube Algorithms with Applications to Image Processing and Pattern Recognition*. Springer-Verlag, New York, 1990.
- Saad, Y., and Schultz, M. H. Topological properties of hypercubes. *IEEE Trans. Comput.* **37** (July 1988), 867-872.
- Saad, Y., and Schultz, M. H. Data communication in hypercubes. *J. Parallel Distrib. Comput.* **6**, 1 (Feb. 1989), 115-135.
- Saad, Y., and Schultz, M. H. Data communication in parallel architectures. *Parallel Comput.* **11** (1989), 131-150.
- Stamoulis, G. D. Routing and performance evaluation in interconnection networks. Ph.D. Thesis, Massachusetts Institute of Technology, Report LIDS-TH-2035, May 1991.
- Stout, Q. F., and Wagar, B. Intensive hypercube communication, prearranged communication in link-bound machines. *J. Parallel Distrib. Comput.* **10** (1990), 167-181.
- Topkis, D. M. Concurrent broadcast for information dissemination. *IEEE Trans. Software Engrg.* **13** (1983), 207-231.
- Tucker, L. W., and Robertson, G. G. Architectures and applications of the Connection Machine. *IEEE Comput.* 26-38, August 1988.
- Varvarigos, E. A., and Bertsekas, D. P. "Communication Algo-

- rithms for Isotropic Tasks in Hypercubes and Wraparound Meshes," *Parallel Computing*, Vol. 18, pp. 1233–1257, 1992.
25. Varvarigos, E. A., and Bertsekas, D. P. Dynamic broadcasting in parallel computing. Laboratory for Information and Decision Systems, Report LIDS-P-2111, Massachusetts Institute of Technology, Cambridge, MA, May 1992; submitted for publication.
 26. Varvarigos, E. A., and Bertsekas, D. P. Multinode broadcast in hypercubes and rings with randomly distributed length of packets. *IEEE Trans. Parallel Distrib. Systems* 4, 2 (Feb. 1993).
 27. Varvarigos, E. A. Static and dynamic communication in high speed parallel computing. Ph.D. Thesis, Massachusetts Institute of Technology, Report LIDS-TH-2126, September 1992.

EMMANOUEL (MANOS) VARVARIGOS was born in Athens, Greece, in 1965. He received a Diploma (1988) in electrical engineering and computer science from the National Technical University of Athens, Greece, and the M.S. (1990), Electrical Engineer (1991), and Ph.D. (1992) degrees in electrical engineering and computer science from the Massachusetts Institute of Technology. In 1990, he conducted research at Bell Communications Research, Morristown, NJ. He is currently an assistant professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. His research interests are in the areas of data networks, parallel and distributed computation, algorithms, and protocols. Dr. Varvarigos received the

Received June 5, 1992; revised July 29, 1993; accepted September 1, 1993

first panhellenic prize in the Greek Mathematic Olympiad in 1982, and has received the Technical Chamber of Greece award four times. He is a member of the Technical Chamber of Greece.

DIMITRI P. BERTSEKAS received a combined B.S.E.E. and B.S.M.E. from the National Technical University of Athens, Greece in 1965, the M.S.E.E. degree from George Washington University in 1969, and the Ph.D. degree in system science from the Massachusetts Institute of Technology in 1971. Dr. Bertsekas has held faculty positions with the Engineering–Economic Systems Department at Stanford University (1971–1974) and the Electrical Engineering Department at the University of Illinois, Urbana (1974–1979). He is currently Professor of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. He consults regularly with private industry and has held editorial positions with several journals. He was elected Fellow of the IEEE in 1983. Professor Bertsekas has done research in the areas of estimation and control of stochastic systems, linear, nonlinear and dynamic programming, data communication networks, and parallel and distributed computation, and has written numerous papers in each of these areas. He is the author of *Dynamic Programming and Stochastic Control* (Academic Press, 1976); *Constrained Optimization and Lagrange Multiplier Methods* (Academic Press, 1982); *Dynamic Programming: Deterministic and Stochastic Models* (Prentice–Hall, 1987); *Linear Network Optimization: Algorithms and Codes* M.I.T. Press, 1991); and co-author of *Stochastic Optimal Control: The Discrete-Time Case* (Academic Press, 1978); *Data Networks* (1987); and *Parallel and Distributed Computation: Numerical Methods* (Prentice–Hall, 1989).