Efficient Algorithms for Global Data Communication on the Multidimensional Torus Network

Paraskevi Fragopoulou and Selim G. Akl
(fragopou, akl)@qucis.queensu.ca
Department of Computing & Information Science, Queen's University
Kingston, Ontario, Canada, K7L 3N6

Abstract

Efficient interprocessor communication is crucial to increasing the performance of parallel multiprocessors. In this paper, a special framework is developed on the multidimensional torus, a network that is currently receiving considerable attention. Using this framework as the basic tool, a spanning graph with special properties, to fit various communication needs, is constructed on the network. The importance of this graph is demonstrated with the development of optimal algorithms for three fundamental communication problems, namely the multinode broadcasting and the single-node and multinode scattering under the store-and-forward, all-port communication model.

1 Introduction

In this paper we concentrate on three fundamental communication problems, namely the *multinode broadcasting*, and the *single-node* and *multinode scattering*, on the popular *multidimensional torus* network. *Broadcasting* is the distribution of the same group of messages from a source processor to all other processors, and *scattering* is the distribution of distinct groups of messages from a source processor to all other processors. We consider scattering from a single processor of the network (single-node scattering). Furthermore, we consider broadcasting and scattering simultaneously from all processors of the network (multinode broadcasting and scattering). Algorithms for the same communication problems using a similar approach have been previously constructed on the binary hypercube [1, 7], the generalized hypercube [4], and the star networks [3].

All of the communication problems are studied under the *store-and-forward*, *all-port* communication model, meaning that in one time step a processor can exchange messages of fixed length with all of its neighbors simultaneously. The communication is *bidirectional*, meaning that an edge can be used for message transmission in both directions at each time step. Each message requires unit time to be transmitted on an edge, i.e. the *unit* cost model is used. Finally, each processor wishes to transmit a number of M messages to each one of its destination processors.

All of the algorithms presented in this paper are based on the construction of a spanning graph with special properties on the multidimensional torus network. A special framework is developed to facilitate the construction of the spanning graph and the development of the communication algorithms. The multinode broadcasting and scattering problems are of special interest. A special technique is developed on the multidimensional torus so that messages originating at individual nodes are interleaved in such a manner that no two messages contend for the same edge at any given time during the execution of the algorithm. For the single-node scattering problem, where the edges incident to the source node constitute a bottleneck for the transmission of the messages, the spanning graph offers the capability to transmit an equal number of messages over each edge incident to the source node.

2 Notations and Definitions

An n-dimensional, k-ary, multidimensional torus network, denoted by $MT_{n,k}$, is an undirected graph of k^n nodes, each one labeled by an n digit number in radix k arithmetic [2, 8]. Each node v is connected to $2n$ other nodes with which it differs in only one digit by $j \bmod k$, $j \in \{-1, 1\}$, i.e.,

$$v_{n-1} \ldots v_{i+1} v_i v_{i-1} \ldots v_0, \;\; v_{n-1} \ldots v_{i+1} v_i' v_{i-1} \ldots v_0$$

is an edge for all $0 \leq i \leq n-1$, and $v_i' = (v_i + j) \bmod k$, $j \in \{-1, 1\}$, Fig. 1. The network is edge and node symmetric with degree $2n$ and diameter $n\lfloor \frac{k}{2} \rfloor$. $MT_{n,k}$

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)

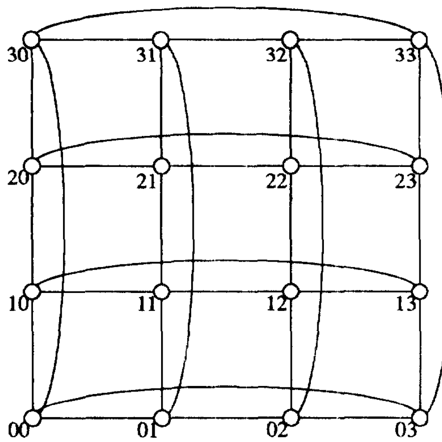

Figure 1: The $MT_{2,4}$ network.

belongs to the class of Cayley graphs [8]. The $2n$ generators that define the edges of $MT_{n,k}$ are denoted by g_i^j, $0 \leq i \leq n-1$, $j \in \{-1,1\}$. Generator g_i^j connects node $v = v_{n-1}...v_{i+1}v_iv_{i-1}...v_0$ to node $v' = v_{n-1}...v_{i+1}((v_i + j) \bmod k)v_{i-1}...v_0$, which results by adding $j \bmod k$ to the i^{th} digit of v. In this case we say that edge (v, v') is of dimension g_i^j, or $dim(v, v') = g_i^j$. In what follows, node 00...0 of $MT_{n,k}$ is referred to as node 0^n. In order to avoid tedious details, in what follows, we consider only multidimensional torus networks with k odd.

We now define two automorphisms on the multidimensional torus network,namely, the *translation* and the *rotation* operations, that will be of primary importance for the construction of the spanning graph and the development of the communication algorithms.

Definition 1 *The translation of a node v with respect to node s, denoted by $Tr_s(v)$, is defined to be node $t = Tr_s(v)$, so that $t_i = (v_i + s_i) \bmod k$, $0 \leq i \leq n-1$.*

The translation operation is an automorphism on the multidimensional torus that preserves the dimension of each edge [5]. This automorphism allows the network to be viewed identically from all nodes. The translation operation on $MT_{n,k}$ is analogous to the exclusive-OR operation on nodes of the binary hypercube [7, 1].

As emphasized in the introduction, the multinode broadcasting and scattering algorithms are designed so that messages originating at individual nodes are interleaved in such a manner that no two messages contend for the same edge at any given time. The properties of the rotation operation, as explained below, will help achieve this attribute. The rotation operation on nodes of the multidimensional torus has properties similar to those of the right cyclic shift operation on nodes of the binary hypercube [7, 1].

Definition 2 *Given function $r(i) = (k - i) \bmod k$ from the set $\{0, 1, ..., k-1\}$ to itself, the rotation of a node $v = v_{n-1}...v_{i+1}v_iv_{i-1}...v_0$, denoted by $Ro(v)$, is defined to be node $r(v_0)v_{n-1}v_{n-2}...v_{i+1}v_iv_{i-1}...v_1$. This can be viewed as a right cyclic shift of the digits of v with the wraparound digit being mapped through function r.*

The rotation operation has the following properties:

1. If (v, u) is an edge of dimension g_i^j, $0 \leq i \leq n-1$, $j \in \{-1, 1\}$, then edge $(Ro(v), Ro(u))$ is of dimension $g_{i'}^{j'}$ so that:
$$i' = (i-1) \bmod n, \qquad j' = \begin{cases} -j, & \text{if } i = 0, \\ j, & \text{otherwise.} \end{cases}$$
2. As a result of property 1, the $2n$ edges derived from edge (v, u) by consecutive applications of the rotation operation are all of different dimensions.
3. The rotation operation preserves the distance of each node from node 0^n.

The nodes of $MT_{n,k}$ are grouped into equivalence classes under the action of the rotation operation as follows:

Definition 3 *An ordered group of nodes, each one derived from its preceding one cyclically, by the application of a rotation is called a necklace.*

Necklaces have the following properties [5]:

1. A necklace contains *at most* $2n$ nodes.
2. The size of a necklace always divides $2n$ [6].
3. All nodes of a necklace are at the same distance form node 0^n.

The nodes of $MT_{n,k}$ at each distance from node 0^n are collections of necklaces. Below, the necklaces of $MT_{3,3}$ at each distance d, $0 \leq d \leq n\lfloor \frac{k}{2} \rfloor$, from node 0^n are given enclosed in parentheses [5].

$d = 0$: ($\underline{000}$)
$d = 1$: ($\underline{200}$, 020, 002, 100, 010, 001)
$d = 2$: ($\underline{220}$, 022, 102, 110, 011, 201)
($\underline{210}$, 021, 202, 120, 012, 101)
$d = 3$: ($\underline{222}$, 122, 112, 111, 211, 221)
($\underline{212}$, 121)

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)
1063-7133/95 $10.00 © 1995 IEEE

Definition 4 *The period of a node v, denoted by $P(v)$, is defined to be the number of nodes contained in the necklace to which it belongs.*

Definition 5 *An unfolded necklace is an ordered group of exactly $2n$ nodes, not necessarily distinct, each one obtained from its preceding one, cyclically, by the application of a rotation.*

Each necklace has a corresponding unfolded necklace. For necklaces that contain $2n$ nodes, the corresponding unfolded necklace is the necklace itself. For necklaces that contain $P < 2n$ nodes, the corresponding unfolded necklace is the necklace repeated $\frac{2n}{P}$ times. This is possible since the size of a necklace is always a divisor of $2n$. Below, the unfolded necklaces of $MT_{3,3}$ at each distance from node 0^n, are given [5].

$$
\begin{aligned}
d=0:&\quad (\underline{000},\ 000,\ 000,\ 000,\ 000,\ 000)\\
d=1:&\quad (\underline{200},\ 020,\ 002,\ 100,\ 010,\ 001)\\
d=2:&\quad (\underline{220},\ 022,\ 102,\ 110,\ 011,\ 201)\\
&\quad (\underline{210},\ 021,\ 202,\ 120,\ 012,\ 101)\\
d=3:&\quad (\underline{222},\ 122,\ 112,\ 111,\ 211,\ 221)\\
&\quad (\underline{212},\ 121,\ 212,\ 121,\ 212,\ 121)
\end{aligned}
$$

The following definition aims to distinguish one particular node of each necklace.

Definition 6 *The binary correspondent of a node v of $MT_{n,k}$ is the binary number obtained if we substitute each nonzero digit in v with 1. One particular node of each necklace is now distinguished as follows:*

1. *Select the nodes of the necklace that have the largest binary correspondent.*
2. *Choose the largest among the nodes selected in step 1, if the k digits that are used to label the nodes of $MT_{n,k}$ are ordered as follows: $0 < 1 < k-1 < \ldots i < k-i < \ldots < \lceil \frac{k}{2} \rceil$. This ordering of the digits is adopted in order to reflect how each digit contributes to the distance of a node from node 0^n.*

The node selected in step 2 is defined to be the leader of the necklace.

The property of the rotation operation that $2n$ edges each of which is obtained as a rotation of its preceding one are all of different dimensions, along with the property of edge dimension preservation of the translation operation will help guarantee that messages originating at individual nodes during a multinode broadcasting or scattering algorithms are interleaved in such a manner that no two messages contend for the same edge at any given time. In a multinode broadcasting or scattering algorithm, all nodes of the network are source of messages. Under the all-port communication model $2nk^n$ edges are available on $MT_{n,k}$ for message transmission at each time step. The algorithm proceeds symmetrically from all nodes and as a consequence messages originating at each one of the k^n nodes of $MT_{n,k}$ are transmitted through at most $2n$ edges at each time step. Let us denote by $E_i(0^n)$ the set of $2n$ edges on which messages originating at node 0^n are transmitted at time step i of the algorithm. Since a multinode algorithm proceeds symmetrically from each node of the network, the $2n$ edges on which messages originating at node s are transmitted at time step i, denoted by $E_i(s)$, is obtained from $E_i(0^n)$ using the operation of translation with respect to s, i.e., if $(v, u) \in E_i(0^n)$ then $(Tr_s(v), Tr_s(u)) \in E_i(s)$. The following lemma is enough to guarantee that no conflicts arise during the execution of an algorithm.

Lemma 1 *At each time step i, if the $2n$ edges in $E_i(0^n)$ are all of different dimensions, then the sets of $2n$ edges $E_i(s)$, where s ranges over all nodes of $MT_{n,k}$, are disjoint [5].*

3 The Spanning Graph

We construct a spanning graph rooted at node 0^n of $MT_{n,k}$, and denoted by BSG_{0^n}. This is a special type of graph, which is composed of $2n$ subtrees, denoted by T_i, $0 \le i \le 2n$.

Definition 7 *Subtree T_0 of BSG_{0^n}, contains all the leader nodes (definition 6) of $MT_{n,k}$ and is defined through the following parent function. For leader node v, let p be the position of its rightmost nonzero digit.*

$$
parent(v) = \begin{cases} \emptyset, & \textit{if } v = 0^n, \\ v_{n-1} \ldots v_{p+1} v_p^- v_{p-1} \ldots v_0, & \textit{otherwise.} \end{cases}
$$

$$
\textit{where} \quad v_p^- = \begin{cases} v_p - 1, & \textit{if } k \le \lfloor \frac{k}{2} \rfloor, \\ (v_p + 1) \bmod k, & \textit{otherwise.} \end{cases}
$$

Any other subtree T_i, $0 < i < 2n$, of BSG_{0^n} is derived as a rotation of subtree T_{i-1}.

BSG_{0^n} has the following properties [5]:

1. The paths from the source node to any other node through BSG_{0^n} are all of shortest length.
2. All the subtrees contain the same number of nodes.
3. Nodes with period P that belong to nonfull necklaces have $\frac{2n}{P}$ paths to node 0^n through BSG_{0^n}. These nodes are always leaf nodes of the subtrees of BSG_{0^n}.

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)
1063-7133/95 $10.00 © 1995 IEEE

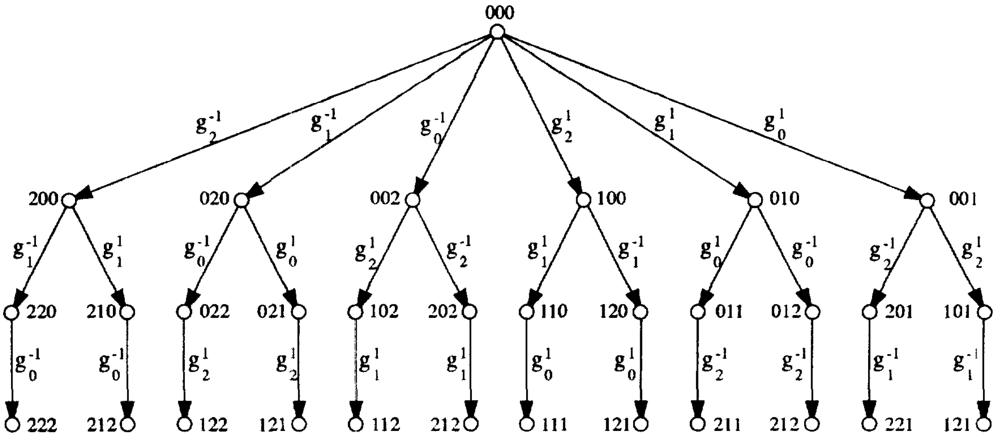

Figure 2: The BSG_0^n spanning tree on the $MT_{3,3}$ network.

4. Corresponding nodes of the $2n$ subtrees form unfolded necklaces and corresponding edges of the $2n$ subtrees are of different dimensions.

Using the BSG_{0^n} graph we can easily derive a BSG_s, rooted at any other node s of $MT_{n,k}$ by simply applying the operation of translation with respect to s, on BSG_{0^n}. The BSG_{0^n} spanning graph on the $MT_{3,3}$ network can be seen in Fig. 2.

The importance of the BSG_s graph lies in several different properties it possesses. The fact that each of the $2n$ subtrees of BSG_s contains the same number of nodes is used in the single-node and multinode scattering algorithms in order for each source node to transmit an equal number of its messages over each one of its incident edges. Furthermore, messages originating at individual nodes in a multinode broadcasting or scattering algorithm are interleaved in such a manner that no two messages contend for the same edge at any time during the execution of the algorithm. The $2n$ subtrees of BSG_{0^n} are rotations of each other, and as a consequence $2n$ corresponding edges of the subtrees of BSG_{0^n} are all of different dimensions. According to lemma 1 it is enough to use $2n$ corresponding edges of the subtrees of BSG_{0^n} in order to avoid message conflicts.

4 Communication Algorithms

4.1 Lower bounds

In a multinode broadcasting problem on $MT_{n,k}$, each node receives a total of $M(k^n-1)$ messages, and a lower bound for the number of message transmissions is $M(k^n-1)k^n$. Since each node of $MT_{n,k}$ has $2n$ incident edges, a lower bound for the number of time steps required for this problem is $\lceil \frac{M(k^n-1)}{2n} \rceil$.

In a single-node scattering problem on $MT_{n,k}$, the source node transmits a total of $M(k^n-1)$ messages. Since each node of $MT_{n,k}$ has $2n$ incident edges, a lower bound for the number of time steps required for this problem is $\lceil \frac{M(k^n-1)}{2n} \rceil$. A lower bound for the number of message transmissions required is the sum of the shortest distances of all nodes to the source node, multiplied by M. If we denote by N_d the number of nodes at distance d from the source node, A lower bound for the number of message this lower bound is derived as follows:

$$M \sum_{d=1}^{n\lfloor \frac{k}{2} \rfloor} dN_d = Mk^n \left(\frac{\sum_{d=1}^{n\lfloor \frac{k}{2} \rfloor} dN_d}{k^n} \right).$$

The expression in parentheses is the average diameter of the $MT_{n,k}$ network which is equal to $\frac{n(k^2-1)}{4k}$ [5].

The multinode scattering problem is equivalent to nk^n scattering problems, one from each node of $MT_{n,k}$. A lower bound for the number of message transmissions $\lceil \frac{Mn(k^2-1)k^{2n-1}}{4} \rceil$. At most $2nk^n$ message transmissions can be performed at each time step and a lower bound for the number of time steps is $\lceil \frac{M(k^2-1)k^{n-1}}{8} \rceil$.

4.2 Multinode broadcasting

In a multinode broadcasting problem, each node of the network transmits the same M messages to all the

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)
1063-7133/95 $10.00 © 1995 IEEE

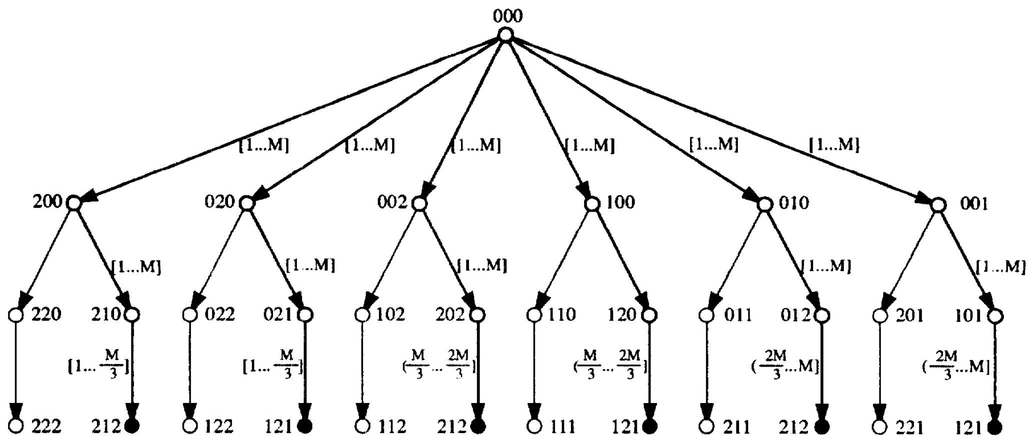

Figure 3: Multinode broadcasting on the $MT_{3,3}$ network using $BSG_0{}^n$.

other nodes. Each source node s uses BSG_s for the transmission of its messages.

The multinode broadcasting algorithm proceeds as follows:

1. Each source node s transmits the M messages it wishes to broadcast to all of its neighbors simultaneously.

2. When an intermediate node v of a BSG_s receives a group of M messages originating at node s, it stores a copy, and performs the following procedures. The messages have to be forwarded to the first child of node v, node u, in BSG_s. If node $Tr_s^{-1}(u)$ has period $P = 2n$, then node v transmits the M messages it received from its parent to node u. However if node $Tr_s^{-1}(u)$ has period $P < 2n$, then a message splitting technique is employed. In this case, the group of M messages is split into $\frac{2n}{P}$ subgroups of $\frac{MP}{2n}$ messages each. Node v in subtree T_i, $0 \leq i < 2n$, of BSG_s transmits the $(i \text{ div } P)^{th}$ subgroup of messages to its first child node in BSG_s.

3. When a leaf node of BSG_s receives a group of messages broadcast by node s, it transmits an acknowledgment to its parent node in BSG_s.

4. When an intermediate node v receives an acknowledgment from one of its children nodes in BSG_s, it forwards the messages it received in the past from node s to its next child in BSG_s following the message splitting technique described above. When an acknowledgment is received from the rightmost child node of v in BSG_s, node v transmits an acknowledgment to its parent node in BSG_s.

The algorithm terminates when each source node receives acknowledgments from all its neighbors. In this algorithm, the transmission of messages in each BSG_s corresponds to a simultaneous depth first traversal of its subtrees. At each time step of the algorithm, messages originating at node 0^n of the network are transmitted over $2n$ corresponding edges of BSG_{0^n}. From the properties of BSG_{0^n}, these edges are all of different dimensions, and the requirement of lemma 1 for message conflict avoidance is satisfied. An example of a multinode broadcasting algorithm on the $MT_{2,4}$ network can be seen in Fig. 3. This figure helps illustrate the technique of message splitting performed by the algorithm.

The number of message transmissions and the number of time steps required by this algorithm are $M(k^n - 1)k^n$ and $\lceil \frac{M(k^n-1)}{2n)} \rceil$, respectively. Consequently, the algorithm is optimal with respect to both of these measures.

4.3 Single node scattering

In a single-node scattering problem, a source node s transmits distinct groups of M messages to each other node. Node s uses BSG_s for the transmission of its messages. Each source node keeps a table of $O(\frac{k^n}{2n})$ nodes, which correspond to the nodes in subtree T_0 of BSG_{0^n}, sorted in reverse ordering of their distance from node 0^n. Each one in this table is accompanied by a number to indicate its period P. Recall that nodes with period P have $\frac{2n}{P}$ paths to node 0^n through BSG_{0^n}. The interesting property of this algorithm is that nodes with period P receive $\frac{MP}{2n}$ of their M messages through each one of the $\frac{2n}{P}$ paths from node 0^n.

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)

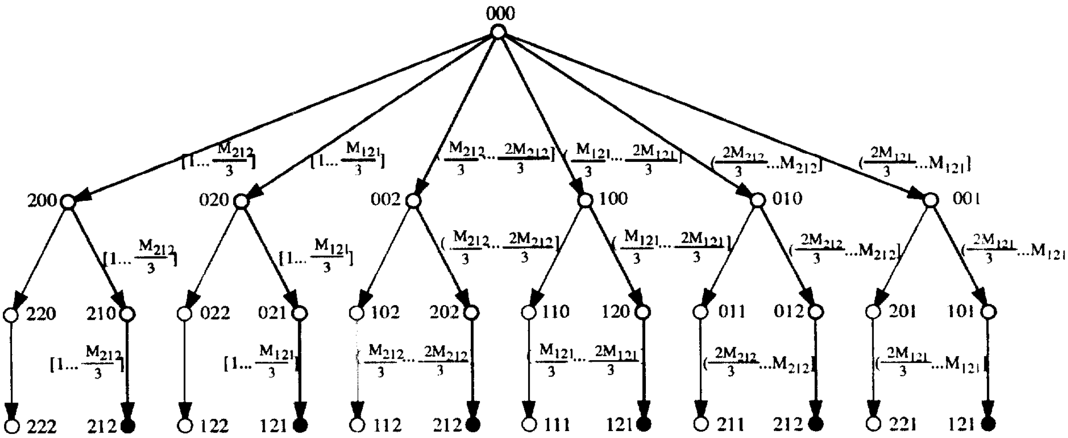

Figure 4: Single-node scattering on the $MT_{3,3}$ network using BSG_0^n.

The single-node scattering algorithm proceeds as follows: For each node v in the table of $O(\frac{k^n}{2n})$ entries do the following:

1. Each source node s transmits messages destined to nodes $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, simultaneously. If node v has period $P = D$, then these $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, are distinct and the M messages destined to node $Tr_s(Ro^i(v))$ are transmitted through subtree T_i of BSG_s, for $0 \leq i < 2n$, respectively. However, if node v has period $P < D$, then nodes $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, are not distinct, but they are a group of P nodes, repeated $\frac{2n}{P}$ times. In this case, a message splitting technique is employed by the source node. Each one of the P groups of M messages that node s has to transmit is split into $\frac{2n}{P}$ subgroups, each containing $\frac{MP}{2n}$ messages. The i^{th}, $0 \leq i < \frac{2n}{P}$, subgroup of the j^{th}, $0 \leq j < P$, group of messages is transmitted through subtree T_{iP+j} of BSG_s. As a consequence, each of the P nodes receives $\frac{MP}{2n}$ of its M messages through each of the $\frac{2n}{P}$ paths from node s through BSG_s.

2. As soon as an intermediate node v receives a new message, it performs the following procedures. If node v is the destination of the message it stores a copy and removes it from the network. If v is not the destination of the message, the identity of the child node to which the message will be forward has to be determined. Node v in subtree T_i, $0 \leq i < 2n$, of BSG_s identifies the first digit to the left of digit $(n-1-i) \bmod n$ in its label that is not equal to the corresponding digit of the destination node. The message is forwarded to the child node of v that has this digit updated.

3. As soon as a source node has transmitted the messages to nodes $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, through its incident edges, it starts transmitting messages to nodes $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, for the next entry u in the table.

An instance of the single-node scattering algorithm on the $MT_{3,3}$ network is shown in Fig. 4.

BSG_s is a shortest path spanning graph and the number of message transmissions performed is $\lceil \frac{Mn(k^2-1)k^{n-1}}{4} \rceil$, which is optimal. Each source node transmits an equal number of its messages over each one of its incident edges and the the number of time steps required is $\lceil \frac{M(k^n-1)}{2n} \rceil$, which is also optimal.

4.4 Multinode scattering

In a multinode scattering problem each node transmits distinct groups of M messages to each other node. Each source node s uses BSG_s for the transmission of its messages. The method used for the multinode scattering algorithm is similar to the one used for the single-node scattering algorithm, but simultaneously executed from all nodes of the network. Each node keeps a table of $O(\frac{k^n}{2n})$ entries that correspond to the nodes in subtree T_0 of BSG_{0^n}. Each node in this table is accompanied by a number to indicate its period P.

The multinode scattering algorithm from each node of the network proceeds as follows: For each node v in the table of $\frac{k^n}{2n}$ entries do the following:

1. Source node s determines the destination of the messages to be transmitted over its i^{th}, $0 \leq i <$

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)
1063-7133/95 $10.00 © 1995 IEEE

$2n$, port as $Tr_s(Ro^i(v))$. For node v with period P, each of the P groups of M messages that have to be transmitted by the source node is split into $\frac{2n}{P}$ subgroups of $\frac{MP}{2n}$ messages each. The i^{th}, $0 \leq i < \frac{2n}{P}$, subgroup of the j^{th}, $0 \leq j < P$, group of messages is transmitted over subtree T_{iP+j} of BSG_s.

2. As soon as an intermediate node u receives a group of messages it performs the following procedures. If node v is the destination node of the messages, it stores a copy and removes them from the network. If node v is not the destination node of the messages, it has to identify the child node to which the messages have to be forwarded. Node v in subtree T_i of BSG_s, locates the first digit to the left of digit $(n-1-i) \bmod n$ in its label that is not equal to the corresponding digit of the destination node. The messages are forwarded to the child node of v that has this digit updated.

3. When the messages transmitted from a source node s have reached their destination nodes $Tr_s(Ro^i(v))$, $0 \leq i < 2n$, then s can transmit messages to nodes $Tr_s(Ro^i(u))$, $0 \leq i < 2n$, for the next entry u in the table.

From the properties of BSG_{0^n}, we know that the $2n$ paths that lead to nodes $Ro^i(v)$, $0 \leq i < 2n$, through its subtrees T_i, $0 \leq i < 2n$, respectively, are rotations of each other, and as a consequence, the $2n$ edges at each level of these paths are of different dimensions. Each node in a path receives all the messages from its parent node before it starts transmitting them to the next node down the path. As a consequence, at each time step, $2n$ edges that are all at the same level of the paths are used. Since these edges are all of different dimensions the requirement of lemma 1 is satisfied, and no two messages contend for the same edge during the execution of the algorithm.

Each message follows a shortest path to its destination node and the minimum number of message transmissions, $\lceil \frac{Mn(k^2-1)k^{2n-1}}{4} \rceil$, is performed. Furthermore, an equal number of the $M(k^n-1)$ messages that each source node has to transmit are transmitted over each one of its incident edges and the minimum number of time steps, $\lceil \frac{M(k^2-1)k^{n-1}}{8} \rceil$, is achieved.

5 Conclusions

A general framework was developed on the multidimensional torus network, that led to the construction of a spanning graph that possesses some special properties. The application of this graph to the development of optimal communication algorithms was demonstrated by describing algorithm for the multinode broadcasting, and the single-node and multinode scattering problems, under the store-and-forward, all-port communication model. Our method is mostly useful for communication problems that require a group or all nodes of the network to be sources of messages, such as the multinode broadcasting and scattering problems. The property that corresponding edges of the subtrees are of different dimensions, along with lemma 1, give the necessary condition for conflict avoidance.

Acknowledgments

This research was supported by the Telecommunications Research Institute of Ontario and the Natural Sciences and Engineering Research Council of Canada.

References

[1] D.P. Bertsekas, C. Ozveren, G.D. Stamoulis, P. Tseng, and J.N. Tsitsiklis, "Optimal Communication Algorithms for Hypercubes", *J. Parall. Distr. Comput.*, 11(4), 263-275, 1991.

[2] L.N. Bhuyan and D.P. Agrawal, "Generalized Hypercube and Hyperbus Structures for a Computer Network", *IEEE Trans. Comput.*, 33(4), 323-333, 1984.

[3] P. Fragopoulou and S.G. Akl, "Optimal Communication Algorithms on Star Graphs Using Spanning Tree Constructions", *J. Parall. Distr. Comput.*, 24(1), 55-71, 1995.

[4] P. Fragopoulou, S.G. Akl, and H. Meijer, "Optimal Communication Primitives on the Generalized Hypercube Network", submitted 1994.

[5] P. Fragopoulou and S.G. Akl, "A Framework for Optimal Communication on the Multidimensional Torus Network", submitted 1994.

[6] L. Garding and T. Tambour, *Algebra for Computer Science*, Springer-Verlag, 1988.

[7] S.L. Johnson and C.T. Ho, "Optimum Broadcasting and Personalized Communication in Hypercubes", *IEEE Trans. Comput.*, 38(9), 1249-1268, 1989.

[8] S. Lakshmivarahan, J.S. Jwo, and S.K. Dhall, " Symmetry in Interconnection Networks based on Cayley Graphs of Permutation Groups: A Survey", *Parall. Comput.*, 19(4), 361-407, 1993.

Proceedings of the 9th International Parallel Processing Symposium (IPPS '95)
1063-7133/95 $10.00 © 1995 IEEE