Efficient Collective Communication on Heterogeneous Networks of Workstations [1]

Mohammad Banikazemi, Vijay Moorthy, and Dhabaleswar K. Panda

Dept. of Computer and Information Science
The Ohio State University, Columbus, OH 43210-1277
Tel: (614)-292-5199, Fax: (614)-292-2911
Email: {banikaze, moorthy, panda}@cis.ohio-state.edu

Contact Author: Mohammad Banikazemi

Abstract:

Networks of Workstations (NOW) have become an attractive alternative platform for high performance computing. Due to the commodity nature of workstations and interconnects and due to the multiplicity of vendors and platforms, the NOW environments are being gradually redefined as *Heterogeneous Networks of Workstations* (HNOW) environments. This paper presents a new framework for implementing collective communication operations (as defined by the Message Passing Interface (MPI) standard) efficiently for the emerging HNOW environments. We first classify different types of heterogeneity in HNOW and then focus on one important characteristic: *communication capabilities of workstations*. Taking this characteristic into account, we propose two new approaches (*Speed-Partitioned Ordered Chain* (SPOC) and *Fastest-Node First* (FNF)) to implement collective communication operations with reduced latency. We also investigate methods for deriving *optimal* trees for broadcast and multicast operations. However, generating such trees is found to be computationally intensive. It is shown that the FNF approach, in spite of its simplicity, can deliver performance within 1% of the performance of the optimal trees. Finally, these new approaches are compared with the approach used in the MPICH implementation on experimental as well as on simulated testbeds. On a 24-node existing HNOW environment with SGI workstations and ATM interconnection, our approaches reduce the latency of broadcast and multicast operations by a factor of up to 3.5 compared to the approach used in the existing MPICH implementation. On a 64-node simulated testbed, our approaches can reduce the latency of broadcast operation by a factor of up to 4.9. Thus, these results demonstrate that there is significant potential for our approaches to be applied towards designing scalable collective communication libraries for current and future generation HNOW environments.

Keywords: Parallel Architecture, Network of Workstations (NOW), Heterogeneous Network of Workstations (HNOW), Collective Communication, Broadcast, Multicast.

[1] This research is supported in part by NSF Career Award MIP-9502294 and NSF Grant CCR-9704512.

Contents

1 Introduction **3**

2 Characterizing Heterogeneous Networks of Workstations **5**
2.1 Major Characteristics 5
2.2 Overhead of Point-to-Point Communication under Heterogeneity 6

3 A New Framework for Collective Communication **6**

4 Speed-Partitioned Ordered Chain (SPOC) Approach **7**
4.1 Basic Idea 8
4.2 Broadcast 8
4.3 Multicast (Single and Multiple) 9

5 Optimal Trees **10**

6 Fastest-Node First (FNF) Approach **12**
6.1 Broadcast 12
6.2 Multicast (Single and Multiple) 13
6.3 Comparison With Optimal Trees 13
6.4 Comparisons between FNF-based and SPOC-based Trees 14

7 Performance Evaluation **14**
7.1 Experimental Results 15
7.1.1 Experimental Setup 15
7.1.2 Broadcast 15
7.1.3 Multicast 16
7.2 Simulation Results 17
7.2.1 Simulation Setup 17
7.2.2 Broadcast 18
7.2.3 Multicast 18

8 Conclusions and Future Research **21**

9 Acknowledgment **22**

1 Introduction

Networks of Workstations (NOW) are becoming increasingly popular for providing cost-effective and affordable parallel computing for day-to-day computational needs [1, 11, 34]. Such environments consist of clusters of workstations connected by *Local Area Networks* (LANs). Hardware and software LAN technology was not initially developed for parallel processing, and thus the communication overhead between workstations can be quite high. In order to achieve performance comparable to MPP systems, many research projects are currently being undertaken in academia and industry to provide fast communication and synchronization in NOW systems [32]. However, most of these research projects focus on *homogeneous* NOW systems where similar kinds of workstations (nodes) are connected over a single network architecture. Some popular network architectures used in the current NOW environments are Ethernet, ATM [2], FDDI, Fiber-Channel, and cut-through routed networks such as SP-2 [33] and Myrinet [4].

Due to the commodity nature of workstations and networking equipment, LAN environments are gradually becoming *heterogeneous*. The capability of a LAN environment to incrementally expand by incorporating new generations of workstations and network architectures over a period of time is also forcing this trend. Such heterogeneity may get reflected in terms of varying speed and communication capability of workstations, coexistence of multiple network architectures, availability of alternative communication protocols, and availability of specialized support for communication and synchronization over a set of workstations. Thus, in today's networked high performance computing environment, heterogeneity is common and its extent will continue to grow over the years. This is forcing the NOW environments to be gradually redefined as *Heterogeneous Networks of Workstations* (HNOW) environments.

A portable parallel programming environment is key to the success of the NOW/HNOW paradigm. Over the last few years, researchers have developed software packages like PVM [34] and *Message Passing Interface* standards like MPI [7, 24] to provide such portability. Even though these softwares and standards do not force an application developer to understand the intricate details of the hardware, software, and network characteristics, the performance of an application in a NOW/HNOW environment heavily depends on the latency and bandwidth required for interprocessor communication and synchronization across the workstations (nodes).

The need for *collective communication* operations such as *broadcast, multicast, global reduction, scatter, gather, complete exchange*, and *barrier synchronization* arises frequently in parallel applications [3, 5, 6, 13, 18, 22, 25, 29, 35]. Thus, it is critical that the collective communication operations be implemented in the best possible manner (scalable as well as high performance) in a HNOW system. Recently, some projects have emphasized issues related to collective communication in NOW systems. These projects have been centered around the following interconnects: ATM [14], Ethernet [37], and Myrinet [36]. Performance of collective communication operations in NOW environments have been evaluated in [20, 27]. However, all these studies focus on only one type of interconnect in a NOW system. They also do not consider heterogeneity in workstation speeds, communication protocols, etc. Thus, the solutions derived in these research projects cannot be directly applied to HNOW systems to obtain maximum performance.

To the best of our knowledge, the ECO [21] package has been the only effort made to consider the heterogeneity of workstations in NOW environments. ECO is built on top of PVM. It proposes heuristics

to partition the participating workstations of a collective communication operation into subnetworks based on pair-wise round-trip latencies. Next, it divides the required communication steps into two major phases: inter-subnetwork and intra-subnetwork. Different trees are used for performing collective communication operations in each of these phases. However, the proposed partitioning approach based on pair-wise round-trip latencies may result in incorrect partitioning in the presence of many factors such as background traffic and workstations with different communication capabilities. This may cause inefficient implementation of collective communication operations. This framework also does not consider other types of heterogeneity. This leads to the following challenges: 1) how to characterize the heterogeneity of a HNOW environment and 2) how to implement efficient collective communication on HNOW environments by exploiting one or more of the heterogeneous characteristics.

In this paper we take on these challenges. We first classify different types of heterogeneity that can exist in HNOW environments and characterize them. Then, we focus on one major characteristic, *communication capabilities of workstations*. We study the impact of this characteristic on the communication overhead of MPI point-to-point communication in a typical LAN environment consisting of heterogeneous workstations. The experimental results indicate that the communication overhead among workstations in a HNOW environment may vary as much as 5:1. Using this observation, we propose a generalized framework to implement efficient collective communications on HNOW systems. We first introduce a new *Speed-Partitioned Ordered Chain* (SPOC) approach to order the participating nodes of a collective communication based on their communication capabilities. Using this framework, broadcast and multicast collective communication operations are implemented with reduced latency using binomial trees [23]. Then, we argue that using binomial trees might not be the best approach for implementing these collective operations. We propose a method for finding optimal trees for broadcast and multicast operations. We show that deriving an optimal tree for a set of nodes with arbitrary communication capabilities is however a computationally intensive operation. Next, we introduce a more efficient approach called *Fastest-Node First* (FNF) for implementing these operations. A quick performance comparison of the FNF scheme with that of optimal trees suggests that the FNF scheme (with its low complexity) can deliver performance within 1% of the performance of optimal trees.

Finally, the SPOC and FNF approaches are evaluated on an experimental testbed consisting of a cluster of 24 SGI workstations and compared with the existing approach used in the MPICH implementation of the MPI standard. Furthermore, these approaches are evaluated on a 64-node simulated HNOW environment with different architectural characteristics. It is shown that latency of these collective communication operations can be reduced by a factor of up to 4.9 using the proposed algorithms. These results show that considerable benefits can be obtained by using the proposed approaches for implementing collective communication operations in HNOW environments.

The remaining part of the paper is organized as follows. Different types of heterogeneity are characterized in section 2. Experimental results on message initiation cost on a set of workstations are also presented. The basic idea behind the development of the proposed algorithms is presented in section 3. The new SPOC framework for collective communication and its corresponding algorithms are proposed in section 4, and the method for finding optimal trees is introduced in section 5. The FNF approach, its application for implementing efficient broadcasts and multicasts, and its comparison with the SPOC and optimal tree approaches are presented in section 6. Experimental and simulation results are presented in section 7. Finally, we conclude the paper with conclusions and future research directions.

2 Characterizing Heterogeneous Networks of Workstations

In this section we characterize factors leading to heterogeneity in HNOW systems. We show how overhead for MPI point-to-point communication can vary significantly in a HNOW system by considering only one factor - communication capabilities of the nodes.

2.1 Major Characteristics

A typical HNOW system can be characterized by the following four factors. These factors are orthogonal to each other. A typical HNOW environment can have one or more of these characteristics.

Communication capabilities of workstations (nodes): The commodity nature of a LAN environment forces workstations with different processor speed, memory speed, and network interface support to be interconnected. These parameters reflect the hardware capabilities to send or receive a message. A node with faster processor and memory can run a communication routine faster and prepare a message to be injected in less time. Similarly, a faster network interface support can also inject a message into the network at a faster speed. Similar arguments also hold good for receiving a message. These capabilities have direct impact on communication overheads for point-to-point communication as well as collective communication.

Network architectures: A typical LAN system can have multiple networks (intra-networks) connected to each other through an inter-network. Different network architectures such as Ethernet, Fast Ethernet, ATM, Myrinet, and Fiber Channel can coexist in a LAN. Different parameters which can affect the overall performance of a LAN are: network characteristics, overall topology, connectivity (inter-network and intra-network), and communication bandwidth (intra-network and inter-network). Thus, algorithms for collective communication need to take care of these characteristics to achieve reduced latency.

Communication protocols: Current generation LANs typically support different communication protocols (such as TCP/IP, ATM [2], and FM [28]) which workstations (nodes) use to communicate. These protocols have their own characteristics and affect the latency of communication operations. Furthermore, some nodes may have access to more than one set of protocols. In this case it is crucial to choose an appropriate protocol to implement a collective communication operation with reduced latency.

Dedicated support for communication and synchronization: Since a typical HNOW environment consists of multiple subnets, there might be dedicated support for collective communication within some subnets. Examples include hardware multicasting within ATM subnet, broadcasting capability within Ethernet and Fast Ethernet, fast synchronization within SMP systems, etc. Availability of such dedicated support for communication and synchronization needs to be taken into account to implement collective communication operations with reduced latency.

All of the above factors have significant impact on the implementation of collective communication operations on HNOW systems. To illustrate this significance, in this paper, we limit our scope to the first characteristic only. Similar approach can also be used for other characteristics and we are currently working along these directions.

2.2 Overhead of Point-to-Point Communication under Heterogeneity

We present experimental results to show the effect of communication capabilities of workstations on the latency of MPI point-to-point communication. We measured round trip latency between four different pairs of workstations in a heterogeneous environment. The workstations were connected via Ethernet and used MPICH communication library [12] to communicate. Table 1 shows these results. Since the results are symmetric, values are shown for only the upper triangle entries. These values indicate how processor speed affects the time taken to transmit a message from one workstation to another. The fastest workstation we used was an HP 735 and the slowest one was a Sun 4. It can be observed that the communication startup time for a Sun 4 is around 5 times that of an HP 735.

Table 1: Roundtrip times in microseconds between different types of workstations in a heterogeneous network. (Entries with * indicate that the sender and receiver were in different clusters.)

	HP 735	HP 715/100	HP 715/64	Sun4
HP 735	871	973	2491 *	5806 *
HP 715/100		1020	2538 *	5871 *
HP 715/64			1869	6050 *
Sun4				4196

These experimental results demonstrate that workstation speeds can have direct impact on the communication latency. Since collective communication operations involve more than one workstations the question arises whether we can use the heterogeneity to our advantage to implement the operations faster. We propose such a framework in the following section.

3 A New Framework for Collective Communication

In this section we propose a framework to take advantage of heterogeneity in communication capabilities to implement a collective communication operation faster. First, we provide the basic idea behind such a framework and then formalize the problem. In the following sections, we provide alternative approaches to solve the problem.

As we observed earlier, various factors such as processor speed, memory speed, and network interface support affect the communication capability of a node. The above parameters can be combined together to a single parameter known as message initiation cost, t_{ini}. Let us consider a broadcast operation on an example HNOW environment consisting of eight workstations $(1, 2, \cdots, 8)$. Let node 1 be the source node. Let six of these workstations (2, 3, 4, 5, 7, and 8) be slow ones having a message initiation cost, t^s_{ini}. Similarly, let the other two (1 and 6) be fast ones and have lower message initiation cost, t^f_{ini}. Based on our experimental data (shown in Table 1) , it can be observed that the t^s_{ini}/t^f_{ini} can be as large as 5. For an example quantitative evaluation, let us consider $t^f_{ini} = 100.0$ microsec and $t^s_{ini} = 300.0$ microsec, leading to $t^s_{ini}/t^f_{ini} = 3$. Because of the high value of t_{ini} (which is typical of NOW systems) let us ignore the time required for transmitting the messages in our example.

Consider a naive and simple binomial-tree-based scheme [23], based on node numbering. This leads to the broadcast tree as shown in Fig. 1(a). Using this scheme the broadcast can be completed in 700.0 microsec in this HNOW environment. However, a more efficient scheme, as shown in Fig. 1(b), can be designed by considering the differences in the communication capabilities of the workstations. Here

although we are still using a binomial tree, the fast workstations are used as intermediate nodes to broadcast the message faster. This scheme takes only 500.0 microsec to implement the broadcast and demonstrates a 29% improvement in broadcast latency. However, this is not the best way of implementing broadcast in this system. Fig. 1(c) illustrates the optimal implementation with a latency of 400.0 microsec (43% improvement over the naive implementation). This example shows that binomial trees which are optimal trees for implementing broadcast on homogeneous systems might not be optimal in heterogeneous environments.

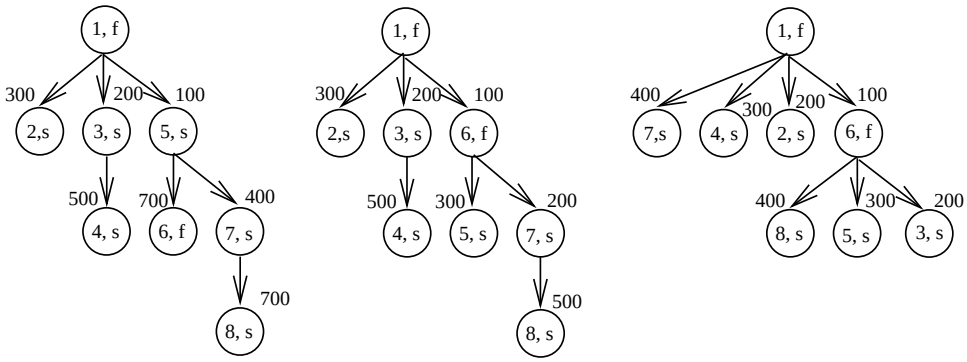

(a) naive scheme (b) better scheme (c) optimal scheme

Figure 1: Different ways to implement broadcast in an eight workstation HNOW environment: (a) simple binomial tree-based scheme using node numbering, (b) a better implementation by considering communication capabilities of the workstations, and (b) the optimal implementation.

It can be observed from this example that several alternatives exist to implement a broadcast communication faster by considering the communication capabilities of the participating nodes and the tree used for implementing the collective operations. In this example we assumed that there are only two types of workstations. The problem becomes more complex if we have multiple levels of communication capabilities. In general, the communication capabilities of the nodes in an N-node HNOW environment can be totally different. This leads to the following problem:

Problem Statement: *How to design efficient algorithms for collective communication operations on an N-node HNOW environment with each node having arbitrary communication capability: $t^i_{ini}, 0 \leq i \leq N-1$.*

In the following sections, we provide alternative solutions to solve this problem. We first present a generalized approach to design efficient binomial broadcast trees by considering multiple levels of communication capabilities across nodes, and develop algorithms for broadcast and multicast collective communication operations. Then, we propose a method to derive optimal trees for performing broadcast and multicast operations. Since deriving such trees is computationally expensive, we finally propose a simple approach which can deliver performance close to those of the optimal trees.

4 Speed-Partitioned Ordered Chain (SPOC) Approach

This approach uses information about the message initiation cost (t_{ini}) of nodes which are participating in a collective communication operation. Using this information, an efficient ordering of participating nodes

is derived to implement the operation faster. In this section, we first present the basic idea behind this approach. Then, we present algorithms for broadcast and multicast (single and multiple) operations.

4.1 Basic Idea

Consider a collective communication operation, say broadcast, which can be implemented by using a binomial-tree [23]. The binomial-tree algorithm is well suited for homogeneous networks. However, nodes that are higher up in the binomial tree send more messages and therefore incur more message startup overhead. Thus, if faster nodes appear higher up in the binomial tree, then the overall startup overhead incurred during the collective communication operation can be reduced. For example, in Figure 1(a) node 5 is a slow node and has 3 descendents while node 6 is a fast node and has no descendent. In Figure 1(b) node 6 has 3 descendent and node 5 has none. Other faster nodes have also been moved to tree nodes having more children. As a result, the latency of the broadcast in Figure 1(b) is 200 microsec lower than the latency of the tree in Figure 1(a).

The above observation leads to the following: The problem of finding the fastest way to implement a binomial tree-based algorithm for a collective communication operation is just finding a way to assign faster workstations to nodes in the binomial tree which have to send more messages. Let us assume that the number of participating workstations in this operation is N. We can always construct a binomial-tree consisting of N tree nodes and having a depth of $\lceil \log_2 N \rceil$. The source node of the broadcast operation remains fixed as the root of the tree, but we have the flexibility of assigning the rest of the participating nodes to any of the other tree nodes as we please.

In order to find an efficient assignment we construct the binomial tree for the given number of participating nodes and find the total number of descendents at each tree node. The number of descendents of a tree node simply specifies the number of nodes that should receive the message from that node, directly or indirectly. We sort the tree nodes from the node with the highest number of descendents to the node with the lowest number of descendents. Obviously, the root of the tree will have the highest number of descendents ($N-1$), but the source node is already assigned to it. Thus, we have flexibility of assigning other participating nodes to tree nodes. We can sort participating nodes in ascending order of message initiation cost and then assign faster participating nodes to tree nodes with greater number of descendents from these two sorted lists.

It can be observed that the dimension-ordered chain [23] resulting from such an assignment will use faster nodes in the initial steps of the communication and will lead to reduced broadcast latency. We call this ordered chain of participating nodes a *Speed-Partitioned Ordered Chain* (SPOC).

In the following two subsections we use the SPOC approach to develop new algorithms for implementing frequently used collective communication operations such as broadcast, multicast, and multiple multicast.

4.2 Broadcast

First we consider the situation where only two classes of workstations are present in a HNOW environment. Then we extend it to include multiple classes.

Consider a HNOW system in which N workstations are participating in a broadcast. Let these work-

stations be divided into two classes: C_1 and C_2, with N_1 and N_2 workstations in each class, respectively. Workstations in each class have similar communication capabilities. Let t^1_{ini} and t^2_{ini} be the message initiation cost of workstations in classes C_1 and C_2. Let us assume that $t^1_{ini} < t^2_{ini}$. As outlined in the previous section, a binomial-tree with N nodes can be created and an ordering based on the number of their descendents, in descending order, can be constructed. The source node of the broadcast is assigned to be the root of the tree. Let us consider the scenario where the source belongs to class C_1. In this case, the rest of the $N_1 - 1$ workstations are assigned to the first $N_1 - 1$ tree nodes appearing in the ordered list of the nodes right after the root. Then N_2 workstations of the class C_2 are assigned to the rest of the tree nodes. If the source node belongs to class C_2 then all N_1 workstations of class C_1 are first assigned to the tree nodes coming immediately after the root in the ordered list. Next $N_2 - 1$ nodes of C_2 are assigned to the following tree nodes in the ordered list. To implement the broadcast operation it is enough to make each workstation (except the source) in the binomial tree receive the message from its parent and forward it to all of its children.

Let us consider the case where more than two classes of workstations exist. Let the workstations be partitioned into a collection of workstation classes (C). The number of workstations belonging to each class ($N_i, 1 \leq i \leq C$) can also vary. Let t^i_{ini} be the message initiation cost for the workstations in class C_i. A binomial tree with N nodes can be created and all tree nodes can be sorted based on the number of their descendents. Let's call this ordered list as $SL_t : \{T_0, T_1, \cdots, T_{N-1}\}$. We also sort all participating nodes (excluding the source of the broadcast) based on their t_{ini}. Let's call this list as $SL_n : \{W_1, W_2, \cdots, W_{N-1}\}$ where, $t^{W_i}_{ini} \leq t^{W_{i+1}}_{ini}$. Then we can assign workstation W_i to node T_i in the binomial tree. Workstations participating in the broadcast receive the message from their parent and forward it to their children (if they have any). The most general form of SPOC-based tree construction for broadcast, where each node has a different t_{ini}, is presented in Fig. 2.

The completion time of broadcast, where W_0 is the source and $\{W_1, W_2, \cdots, W_{N-1}\}$ is the list of other participating nodes in a nondecreasing order with respect to their message initiation cost, can be expressed as:

$$T^{SPOC}_{broadcast} = \sum_{i=0}^{\lceil \log N \rceil - 1} \max(t^0_{ini}, t^{2^i - 1}_{ini}) \quad (1)$$

4.3 Multicast (Single and Multiple)

A multicast [8, 26, 15, 16, 31, 36] operation involves only a subset of workstations in the HNOW environment. The above approach for broadcast can also be used for efficient implementation of single multicast. A binomial tree is first created with the participating nodes in the multicast. Workstations belonging to the multicast group are assigned to the nodes of the binomial tree based on their message initiation cost. Faster nodes are assigned to higher levels of the tree so that they can pass the message to their descendents with reduced latency.

In a multiple multicast operation, more than one group of nodes participate in the multicast [17]. Each group has its own source and destination nodes. Each node can be the source of at most one multicast while, a node can appear in more than one group as one of the destinations. An algorithm similar to the algorithm used for multicast can be used to implement the multiple multicast with the only difference

SPOC-based tree construction for broadcast

Input:

root: the source node of broadcast;
L_n: The list of participating nodes with their respective initiation costs (t_{ini}).

Tree construction steps:

1. $N \leftarrow number_of_nodes(L_n)$
2. $T_b \leftarrow construct_binomial_tree(N)$
3. $SL_t \leftarrow sort(list(T_b), Descending, Num_Descendents)$
   $// \; SL_t = \{T_i \mid 0 \leq i < N,\; num_decs(T_i) \leq num_decs(T_j)\; for\; i < j)$
4. $SL_n \leftarrow sort(L_n, Ascending, T_ini)$
   $// \; SL_n = \{W_i \mid 0 \leq i < N,\; t_ini(W_i) \leq t_ini(W_j)\; for\; i < j)$
5. $SL_n \leftarrow remove(SL_n, root)$
6. $assign_node_tree(item(SL_t, 0), root)$
7. for $i = 1$ to N do
   $assign_node_tree(item(SL_t, i), item(SL_n, i-1))$

Figure 2: Outline of the SPOC-based tree construction for broadcast.

being the fact that different trees are used for each multicast.

5 Optimal Trees

In Section 3 we showed that binomial trees might not be the best tree for implementing broadcast (or multicast) operations on a HNOW. Consider the construction of an optimal tree for performing broadcast (or multicast) among N nodes, $\{W_0, W_1, \cdots, W_{N-1}\}$, where W_0 is the source node. In the first step W_0 sends the message to W_i. Then, W_0 and W_i will be responsible for sending the message to the rest of the participating nodes through two subtrees. First, the Node W_i must be chosen such that the overall tree is *optimal*. Second, the two subtrees must themselves be *optimal*. The same procedure can be applied recursively to each of the multicast subtrees. Therefore, this optimization problem which exhibits optimal substructure and overlapping subproblems properties can be solved by using the dynamic programming technique [9]. This technique can be used for the current problem as follows.

Let $L_{i,A}$ be the minimum latency required to multicast the message from node W_i to all nodes in the set A. If A is an empty set, the latency will be equal to zero. Otherwise, in the first step, message is sent to a node in A (say W_j), and the latency would be the maximum of latencies associated with the two obtained subtrees (where W_i and W_j are the roots of these subtrees and descendents of these two nodes will be all nodes in $A - \{j\}$). Therefore, the overall latency can be obtained through the following recurrence:

$$L_{i,A} = \begin{cases} 0 & \text{if } |A| = 0 \\ min\{t^i_{ini} + max(L_{i,B}, L_{j,C})\} \quad \text{, where } j \epsilon A, \text{ and } B + C = A - \{j\} & \text{otherwise} \end{cases}$$

It is to be noted that the total running time for finding the optimal tree for broadcasting a message among N nodes will be $O(N^2 2^{2N})$. However, this method for finding the optimal tree is too computa-

tionally intensive to be useful in any practical system. In the following section, we propose a near-optimal algorithm which runs in polynomial time. Before describing this near-optimal algorithm, let us look at two important properties of optimal trees presented as the following two lemmas.

Lemma 1 *Let W_0 be the source node of a broadcast (or multicast) operation and $\{W_1, W_2, \cdots, W_{N-1}\}$ be the set of other participating nodes in the order of the time they have received the message. There exists an optimal tree for performing the broadcast (or multicast) operation such that W_k $(1 \leq K \leq N-1)$ receives the message from one of the nodes in the set $\{W_0, W_1, \cdots, W_{k-1}\}$ and the time at which it receives the message is the earliest possible time.*

Proof: For any given tree which implements broadcast or multicast, before W_k receives the message, $\{W_0, W_1, \cdots, W_{k-1}\}$ is the set of nodes which have already received the message. Let us call the earliest time at which node W_i can deliver the message to another node, R_i $(0 \leq i \leq (k-1))$. Now we use proof by contradiction to prove the lemma.

Let us assume there does not exist an optimal tree in which every node receives the message at the earliest possible time from the nodes which have received the message before it. Therefore, in the optimal tree (T) there should exist a node (W_k) which receives the message from W_i where $0 \leq i \leq k-1$ and there exists W_j such that $0 \leq j \leq k-1$, $j \neq i$ and $R_j < R_i$ when W_k receives the message. Now there are two possibilities: 1) W_j sends the message to at least one node W_l where $l > k$. In this case, W_l would have received the message in a time earlier than the time W_k has received the message contradicting the fact that W_k is the *kth* node to receive the message. 2) W_j will never send a message to a node W_l such that $l > k$. In this case we can build another tree T', which is similar to T with the only difference being that the last node to receive the message in T receives the message from W_j in T'. The latency of T' would be less than or equal to that of T. If the latency of T' is less than the latency of T, it means T could not have been optimal (contradiction), and if the latency of T' is equal to that of T, it means that there is an optimal tree in which the *kth* node to receive the message receives it at the earliest possible time. This procedure can be applied repeatedly to obtain an optimal tree in which the *kth* node $(0 \leq k \leq N-1)$ receives the message at the earliest possible time. ∎

Lemma 2 *Let W_0 be the source node of a broadcast (or multicast) operation and $\{W_1, W_2, \cdots, W_{N-1}\}$ be the set of other participating nodes. Let t^i_{ini} be the message initiation cost of node W_i. There exists an optimal tree for performing the broadcast (or multicast) operation in which the message initiation cost of any node other than the source node is less than or equal to that of its children.*

Proof: Again we use proof by contradiction to show the correctness of this lemma. Let us assume that there does not exist an optimal tree in which the message initiation of any node other than the source node is less than or equal to that of its children. Furthermore, let us assume that T is an optimal tree. Therefore, there is a node W_j in T such that it receives the message from W_i $(i \neq 0)$ and $t^j_{ini} < t^i_{ini}$. Let t_1 be the time at which W_j receives the message from W_i. Both W_i and W_j will be ready to send the message to other nodes at t_1. Now consider the tree T' which is similar to T with the only difference being that the positions of W_i and W_j in the tree have been exchanged. Let t_2 be the time at which W_i receives the message from W_j. W_i and W_j will be ready to send the message to other nodes at t_2 and obviously, $t_2 < t_1$. On the other hand, all other nodes which should have received the message before t_1 in T would have received the message earlier (if they receive it from W_i in T) or in the worst case at the

same time in T'. In other words, if T' is used all of the first k nodes (and therefore the rest of the nodes) receive the message and are ready to send out the message either at the same time or earlier than the time they would have, if T was used. Thus, the completion time of the operation while T' is used is at least as low as that if T is used. This process can be executed again and again until there is no node (other than possibly the source node) that has a bigger initiation cost than that of its children, contradicting the assumption that such an optimal tree does not exist. ■

In the next section we use these two properties to propose a near-optimal algorithm which runs in polynomial time with respect to the number of participating nodes.

6 Fastest-Node First (FNF) Approach

Using the above properties of an optimal tree, we propose a greedy algorithm called Fastest-Node First (FNF). In each iteration of this algorithm, one node which has not received the message is added to the tree. Obviously at each instance, we need to make two decisions. First, we need to decide which node is going to send the message to the new node. From Lemma 1 we can easily find the node which should deliver the message to the new node. The second decision to be made is selecting the new node among the nodes which have not been added to the tree yet. To make sure that the property presented in Lemma 2 is preserved, we select the fastest node among the nodes not in the tree. This way, we can generate trees through which multicast and broadcast operations can be implemented. Now, we describe how FNF can be specifically applied for implementing broadcast and multicast operations.

6.1 Broadcast

Consider a HNOW system in which N workstations are participating in a broadcast. Let the workstations be partitioned into C workstation classes $(C_0, C_1, \cdots, C_{C-1})$, where $1 \leq C \leq N$ (note that when $C = 1$ we actually have a homogeneous NOW with respect to the message initiation cost). The number of workstations in each class can also vary. Let t^i_{ini} be the message initiation cost for the workstations in class C_i. The FNF algorithm whose outline is presented in Figure 3, starts by creating a list of the nodes which have a copy of the message (*Senders*), and a list of the nodes which are participating in broadcast but have not received the message yet (*Receivers*). Obviously, at the beginning the first list contains only the source node and the second list contains all other nodes. A variable (R) is associated with each node indicating the earliest time when the node can send out a message to another node. Since at the beginning of the operations none of the nodes except the source node has the message, infinity is assigned to this variable of all nodes except that of the source which is set to zero. The FNF tree which is presented as a set of (parent, child) two-tuples (*Solution*) is also initialized. Then, in $N-1$ successive iterations, the node from which the message is supposed to be received at one of the participating nodes (except the root) is found. In each iteration, the best candidate for sending the message to one of the nodes which have not received the message is found by minimizing the time by which this message can be delivered (R). On the other hand, the fastest node among those which have not received the message yet is chosen as the receiver. After both sender and receiver are selected, the times at which these nodes can send out a new message are each adjusted. During this step, the receiver is taken out of the *Receivers* list and is added to the *Senders* list.

FNF-based tree construction for broadcast

Input:

$Tini$: the initiation costs of all nodes(t_{ini});
root: the source node of broadcast;
L_n: The list of other participating nodes.

Tree construction steps:

1. $N \leftarrow number_of_nodes(L_n) + 1$
2. $R_0 \leftarrow 0$
3. for $i = 1$ to $N - 1$ do
   $R_i \leftarrow \infty$
4. $Senders \leftarrow \{root\}$
5. $Receivers \leftarrow L_n$
6. $Solution \leftarrow \{\}$
7. for $i = 1$ to $N - 1$ do
   $sender \leftarrow k$,where $k \, \epsilon \, Senders$ and $\forall \, l \, \epsilon \, Senders \;\; (R_k + Tini_k) \leq (R_l + Tini_l)$
   $receiver \leftarrow k$,where $k \, \epsilon \, Receivers$ and $\forall \, l \, \epsilon \, Receivers \;\; Tini_k \leq Tini_l$
   $Solution \leftarrow concat(Solution, (sender, receiver))$
   $R_{sender} \leftarrow R_{sender} + Tini_{sender}$
   $R_{receiver} \leftarrow R_{sender}$
   $Receivers \leftarrow remove(Receivers, receiver)$
   $Senders \leftarrow concat(Senders, receiver)$
   endfor

Figure 3: Outline of the FNF algorithm.

It should be noted that the order of the tuples in $Solution$ is important. For any two-tuples whose first items (or sender nodes) are the same, the sender node will send the message to the receiver mentioned in the first two-tuples before sending the message to the other receiver. In other words, the obtained tree is an ordered tree. It should also be noted that for an efficient implementation of the algorithm, nodes can be sorted based on their message initiation costs. This algorithm can be easily implemented with $O(N^2)$ complexity where N is the number of participating nodes.

6.2 Multicast (Single and Multiple)

Similar to the broadcast operation, multicast can be implemented in a more efficient manner under this approach. The FNF algorithm can be directly applied for implementing multicast operation by limiting the list of the participating nodes to the nodes participating in the multicast operation. Similarly, FNF can also be used for implementing multiple multicast by constructing one tree for each multicast operation.

6.3 Comparison With Optimal Trees

To compare the performance of the FNF-based algorithms with that of optimal algorithms in a qualitative fashion, we evaluated the latency of broadcast and multicast operations obtained from these algorithms. We considered a system with 9 nodes. (We could not go further because of the complexity of the optimal

algorithm which is exponential with respect to the number of participating nodes). The message initiation cost of each node was randomly chosen from the $\{100, 200, 300, \cdots, 800\}$ set. For each operation, the source node and other participating nodes were chosen randomly. We recorded the latency of the FNF-based algorithm (L_{FNF}) and that of the optimal algorithm ($L_{optimal}$) for $10,000$ cases for each particular number of participating nodes. We then calculated the average latency of the FNF-based and the optimal algorithms. Table 2 shows these latencies. It can be observed that the latency of FNF-based algorithms is equal to that of the optimal algorithms upto 5 participating nodes. Beyond that, the FNF algorithm produces latencies which are within 1% of the latency produced by the optimal algorithm. Furthermore, trees generated by the FNF-based algorithm were found to be identical to those generated by the optimal algorithm for 90-100% of the cases. Considering the very minor difference between the latency of the FNF-based and the optimal algorithms, and the very low complexity of the FNF-based algorithms, we conclude that FNF-based algorithms will be more practical to be incorporated in future HNOW environments. Thus, for the remaining part of the paper, we does not consider the optimal algorithm any further. now on we will only consider FNF-based algorithms.

Table 2: Comparison between FNF-based and optimal algorithms for implementing broadcast and multicast on a system with 9 nodes.

Number of Participating Nodes	L_{FNF}	$L_{optimal}$	$\frac{(L_{FNF}-L_{optimal})}{L_{optimal}}$
2	453.37	453.37	0.00%
3	705.60	705.60	0.00%
4	805.70	805.70	0.00%
5	871.01	871.01	0.00%
6	914.90	913.15	0.19%
7	947.56	942.26	0.56%
8	976.90	967.27	0.99%
9	984.29	977.12	0.73%

6.4 Comparisons between FNF-based and SPOC-based Trees

Let us compare the trees produced by the FNF-based and SPOC-based trees. In general, they will be different because SPOC-based trees are always binomial in nature. It is to be noted that when all nodes participating in broadcast or multicast have the same message initiation startup (i.e. the participating nodes are homogeneous), the tree obtained from the FNF algorithm will be the same as that obtained from the SPOC algorithm and the tree produced by the naive binomial tree algorithm (such as the one used in the existing MPI implementations). It is also to be noted that where 50% of the participating nodes (including the source node) belong to the fastest group of participating nodes, the trees obtained from the SPOC and FNF algorithms are identical.

7 Performance Evaluation

In this section, the performance of the algorithms developed in the previous sections are compared. First, we present experimental results obtained from a cluster of SGI workstations connected by an ATM network. Due to the limitations of our experimental testbed (limited number of nodes and only two levels of speed),

we also carried out simulation results for larger number of nodes with greater variation in speed. We present these simulation results next.

7.1 Experimental Results

We used an ATM network of 24 SGI workstations at the Ohio Supercomputer Center (OSC) to implement and evaluate the proposed algorithms. We compared the performance of our algorithms with those of MPICH v1.1. In the following subsections, the setup used in the experiments is explained and then the results are presented.

7.1.1 Experimental Setup

The testbed used in our experiments consisted of 24 SGI workstations with two different speeds. There were 16 slow and 8 fast nodes. This allowed us flexibility to take measurements on three different 16-node configurations with 12.5% (2), 25% (4), and 50% (8) fast nodes, respectively. Using MPICH point-to-point communication, the roundtrip latency for a short message (4-byte long) was measured to be 1380 microsec between two fast nodes and 2960 microsec between two slow nodes. The ratio of the communication capabilities of slow and fast nodes is therefore 2.15:1.0.

7.1.2 Broadcast

For measuring the broadcast latency we followed a method similar to the one used in [19]. A broadcast operation starts when the source node initiates it. It is said to be complete when all the other nodes have received the broadcast message. The broadcast latency is defined as the time elapsed between the source node initiating it and the last recipient receiving it. Measurement of broadcast latency was done in the following way. For an N node system, $N-1$ broadcasts were performed. Each time, after the broadcast, one of the $N-1$ recipients sent back an acknowledgment (instead of issuing another broadcast as in [19]). At the source node the time between initiation of the broadcast and receipt of the acknowledgment was measured. The maximum of these $N-1$ time readings corresponds to the last broadcast recipient sending back the acknowledgment. Therefore, this maximum is the sum of the broadcast latency and half of the round trip latency between the source node and the last recipient. Since the roundtrip time can be easily measured, the broadcast latency is obtained by subtracting half of the roundtrip time from the maximum reading. Furthermore, each experiment was repeated 100 times for a given source node and the minimum broadcast latency was taken into account.

Figure 4 shows the results. In the graphs, latencies of the MPICH implementation are referred to as MPI-DEF (MPI with default implementation) latencies. The modified version of MPICH routines in which SPOC-based algorithms are used, is called MPI-SPOC. We refer to the modified version of MPICH in which FNF-based algorithms are employed to perform MPI-Bcast as MPI-FNF. It can be observed that the proposed algorithms always perform better than the naive ordering scheme used in MPICH. Factors of improvement over MPI-DEF are shown in Figure 4(d). As the fraction of fast nodes in the system increases, so does the factor of improvement. A factor of improvement of 1.7 is achieved in a system with 50% faster nodes by using the SPOC-based algorithm. Factors of improvement up to 2.3 are achievable when FNF-based algorithms are used. As predicted in earlier sections, when 50% of the nodes are fast nodes, SPOC-based algorithms perform as good as FNF-based algorithms. It can also be observed that

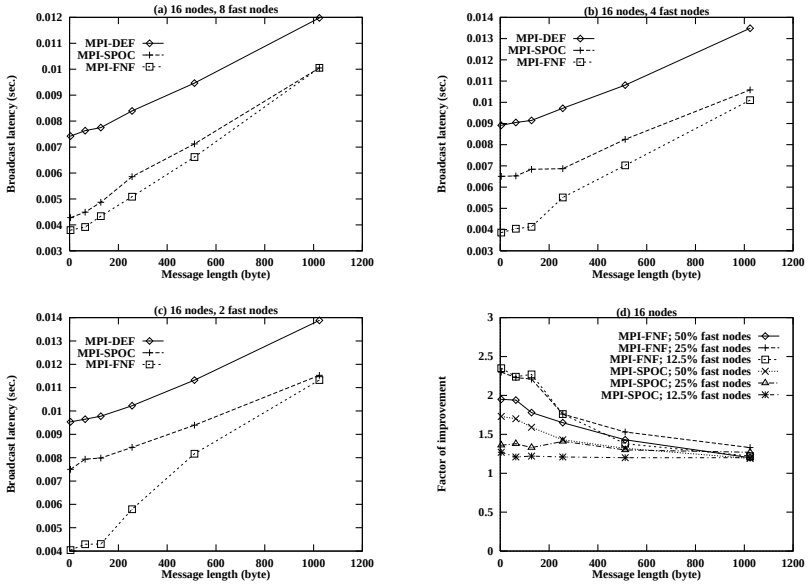

Figure 4: Broadcast latency and factor of improvement on a HNOW system with 16 nodes and different number of fast nodes.

for different fractions of faster nodes, the factor of improvement curves almost coincide beyond a certain message length due to the fact that with increasing message length, the message transmission time begins to dominate the latency. For shorter messages, where the startup overhead dominates, we get higher factors of improvement.

7.1.3 Multicast

In a multicast, the set of recipients is a subset of the set of nodes. The same procedure was therefore used to measure multicast latencies, that is, each experiment for multicast was repeated $100 \times (R-1)$ times, where R is the number of recipients.

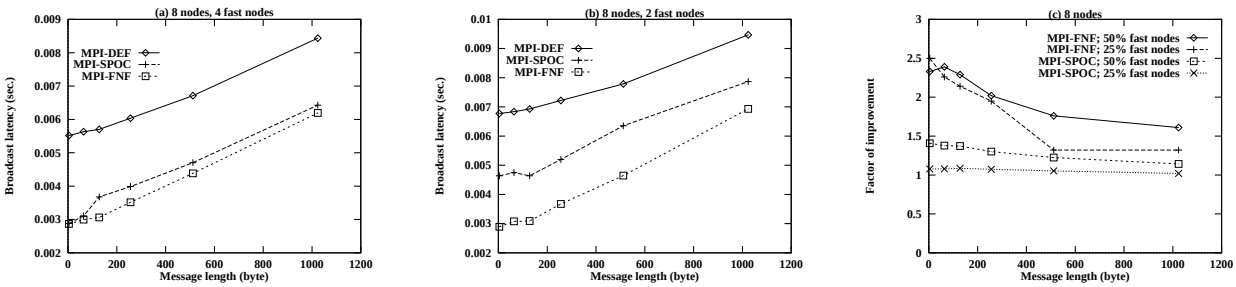

Figure 5: Multicast latency and factor of improvement on a HNOW system with 8 nodes and different number of fast nodes.

Figures 5 and 6 shows results for multicast. Characteristics similar to broadcast results are observed. The factor of improvement increases with increasing system size. Again, as message length increases, transmission time dominates and the factor of improvement decreases. For an 8 node multicast with 50% fast participating nodes factors of improvement of up to 2.5 are observed.

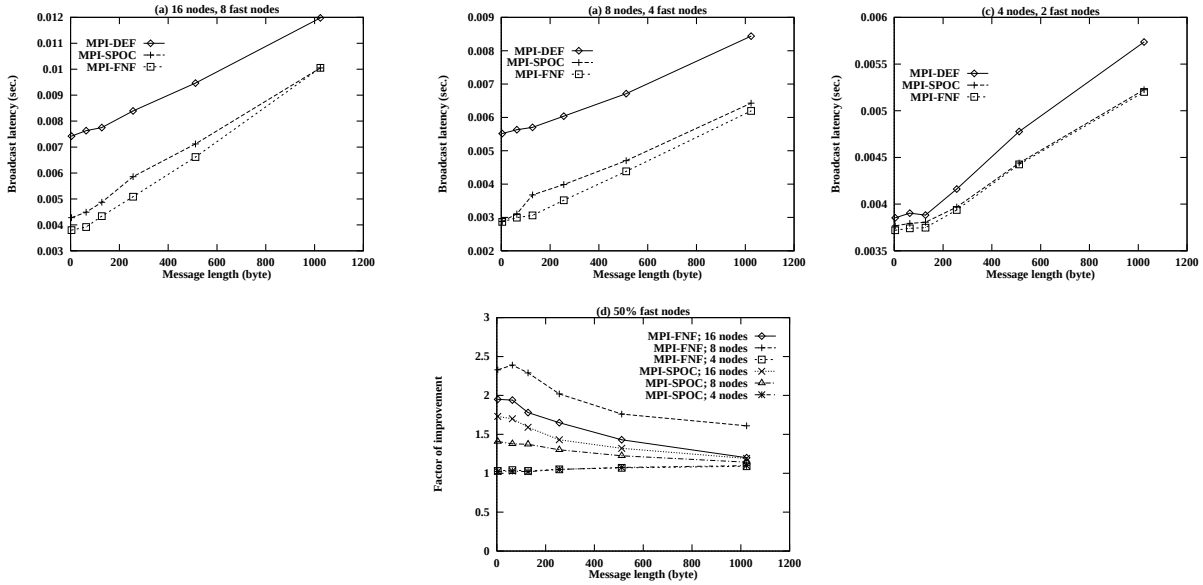

Figure 6: Multicast latency and factor of improvement on configurations with half of the nodes being fast nodes.

7.2 Simulation Results

In this section we present our simulation results. Several experiments were performed to measure the impact of different workstation speeds, number of workstations and fraction of faster nodes on the performance of our algorithms. In the following subsections, first the simulation setup is described in detail. Next, the results for broadcast, single multicast, and multiple multicast are presented.

7.2.1 Simulation Setup

We modeled a representative HNOW system where workstations are interconnected with Myrinet switches. A detailed flit level simulator (built using CSIM [30]) was used to model irregular topologies and the wormhole switching technique. A 64-node HNOW system was considered. Based on the experimental results, presented in Section 2.2, we considered the following communication startup times. Two classes of workstations were considered. The communication startup time for faster class was kept constant at t^f_{ini} = 400.0 microsec. The time for slower class (t^s_{ini}) was varied (800, 1600 and 2400 microsec). These values lead to *speed factors* (ratio of communication cost between the slow class to the fast class) of 2, 4, and 6, respectively. Such a variation helps to study a wide range of HNOW systems.

With respect to the interconnection network [10], the following parameters, representing current generation systems, were used: t_{phy} (link propagation time per byte) = 12.5 nanoseconds, t_{route} (routing delay at switch) = 500 nanoseconds, t_{sw} (switching time across the router crossbar for a flit) = 12.5 nanoseconds, t_{inj} (time to inject a flit into network) = 12.5 nanoseconds and t_{cons} (time to consume a flit from network) = 12.5 nanoseconds. For all experiments we assumed the following default system configuration: a 64 workstation system interconnected by 16 eight-port switches and a network having 50% connectivity [2].

[2]Connectivity is defined as the fraction of ports in a switch which are used for interconnection with other switches [16].

We performed several experiments to study the impact of message length, speed factor and fraction of faster nodes. We evaluated our new algorithms (indicated as SPOC and FNF in the graphs) with the default ordering algorithm (indicated as DEF) which is used by the current MPI implementations. Participating nodes for a given collective communication and the network configurations were generated randomly. Latency value for each data point in the graphs was averaged over 100 experiments (10 different sets of participating nodes for each of 10 different network configurations).

7.2.2 Broadcast

For the broadcast operation, we simulated the FNF-based algorithms, the SPOC-based algorithms, and the default ordering scheme (DEF) for system configurations with 10%, 25% and 50% faster nodes, and for message lengths of 64 bytes, 1 KBytes and 4 KBytes. Figure 7 shows the results for different percentages of fast nodes and speed factors while the message size is fixed to 1 KBytes. It can be observed that the new algorithms always perform better than the naive schemes being currently used in the MPI implementations. The factor of improvement increases with increase in percentage of faster nodes and with increase in differences in the startup times. Furthermore, it can be observed that by increasing the speed factor, the FNF-base algorithm outperform the SPOC-based algorithm more strongly. Overall, a speed factor of up to 4.9 can be achieved. Again it should be noted that when half of the nodes are fast nodes, the SPOC-based and FNF-based algorithms perform equally well (when the source node is also one of the fast nodes). Figure 8 shows the broadcast latencies when message size is changed. It can be observed that the new algorithms outperform the naive scheme in all cases.

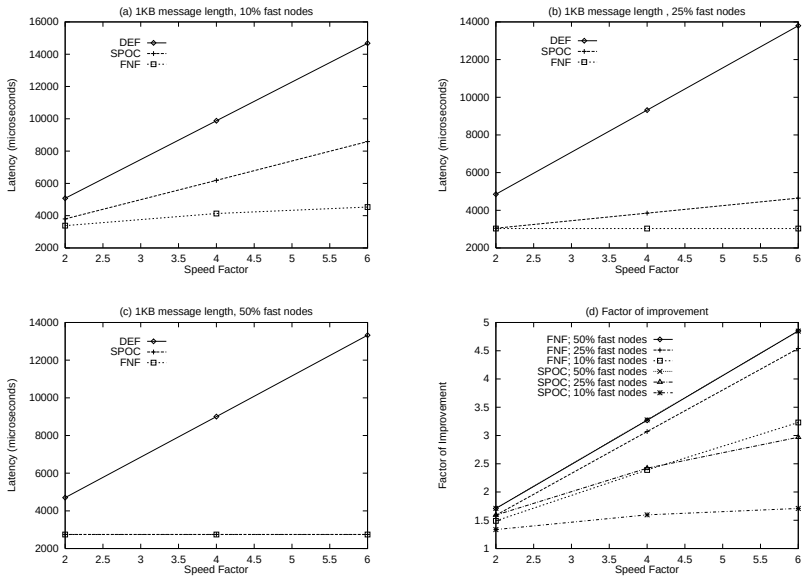

Figure 7: Impact of percentage of faster nodes on broadcast latency: (a) 10% fast nodes, (b) 25% fast nodes, (c) 50% fast nodes, and (d) factor of improvement for (a) - (c).

7.2.3 Multicast

Figures 9, 10, and 11 show the impact of speed factor, percentage of faster nodes, and message length on the latency of single multicast with varying sizes of destination sets. The factor of improvement of the

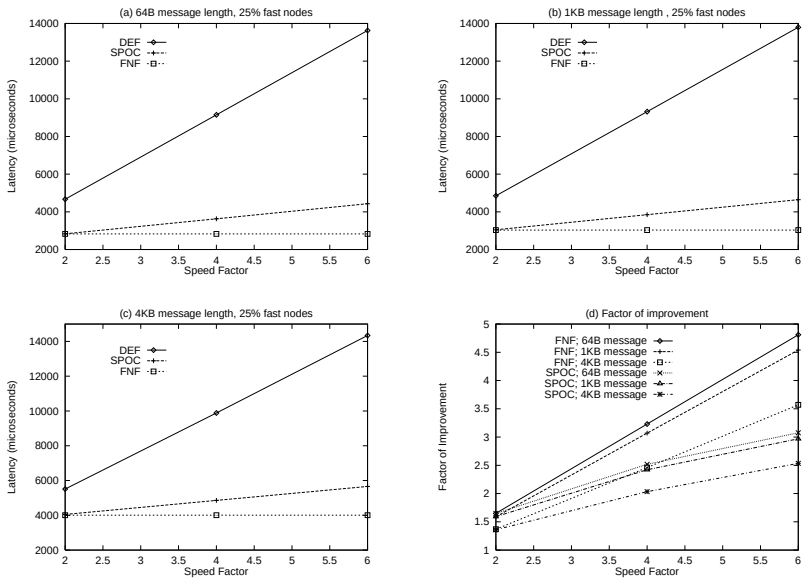

Figure 8: Impact of message length on broadcast latency: (a) 64 Bytes, (b) 1 KBytes, (c) 4 KBytes, and (d) factor of improvement for (a) - (c).

SPOC and FNF approaches over the DEF increases with increasing number of destinations, increasing speed factor, and decrease in message length. For a multicast involving 1 KBytes message on a system with 25% faster nodes and a speed factor of 4, SPOC and FNF algorithms give factors of improvement of 2.3 and 3.3, respectively.

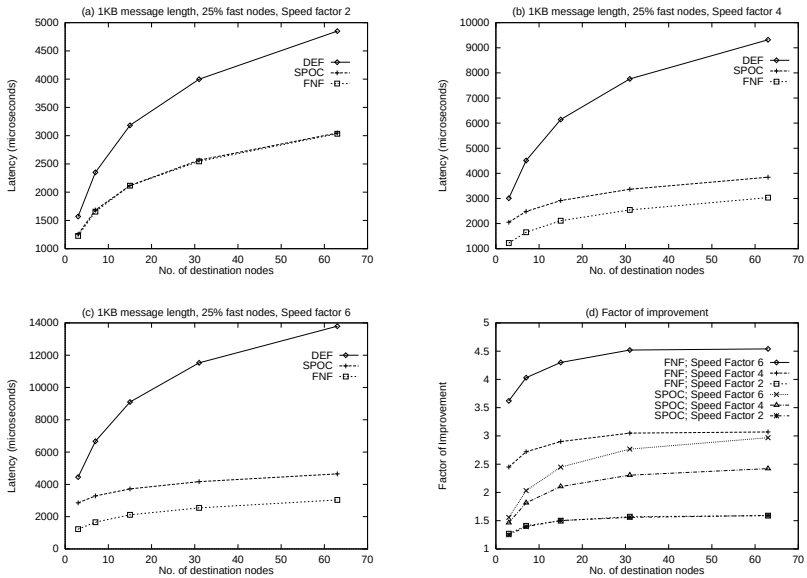

Figure 9: Impact of speed factor on single multicast latency with 25% fast nodes: (a) speed factor 2, (b) speed factor 4, (c) speed factor 6, and (d) factor of improvement for (a) - (c).

Results for multiple multicasts with varying number of sources and destinations are presented in Figure 12. It can be observed that the factor of improvement increases as more and more sources are involved in the multiple multicast. It can be observed that for a system with 50% faster nodes and a speed

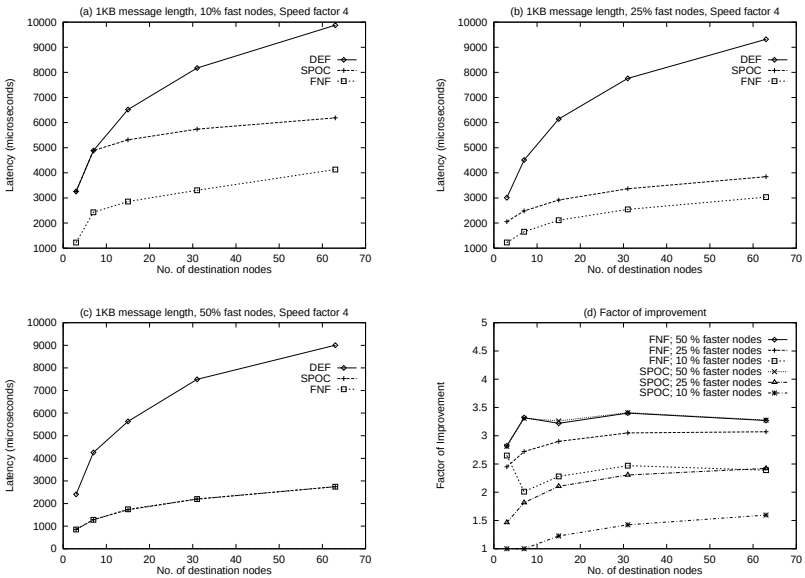

Figure 10: Impact of percentage of faster nodes on single multicast latency with speed factor 4: (a) 10% fast nodes, (b) 25% fast nodes, (c) 50% fast nodes, and (d) factor of improvement for (a) - (c).

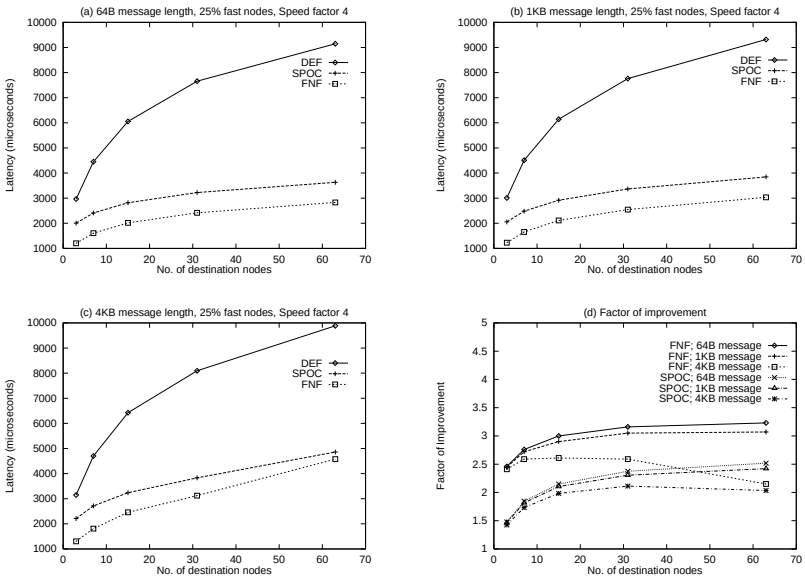

Figure 11: Impact of message length on single multicast latency with 25% fast nodes and speed factor 4: (a) 64 Bytes, (b) 1 KBytes, (c) 4 KBytes, and d) factor of improvement for (a) - (c).

factor of 4, multiple multicast latency can be reduced by a factor of 2.4 using our new algorithms.

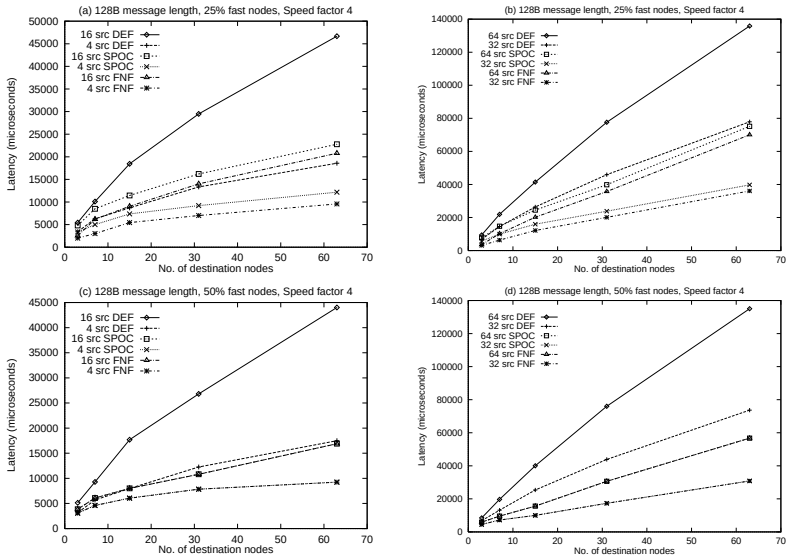

Figure 12: Impact of percentage of faster nodes on latency for multiple multicast: (a) 25% fast nodes with 4 and 16 sources, (b) 25% fast nodes with 32 and 64 sources, (c) 50% fast nodes with 4 and 16 sources, and (d) 50% fast nodes with 32 and 64 sources.

8 Conclusions and Future Research

In this paper, we have presented three new approaches to implement fast collective communication in the emerging HNOW systems. Major factors in HNOW systems have been characterized. A new SPOC-based framework has been introduced to order the participating nodes of a collective communication based on their communication capabilities. Using this framework, algorithms for frequently used collective communication operations (broadcast, single multicast, and multiple multicast) have been developed. An algorithm for generating optimal trees have been proposed. Furthermore, another new approach (FNF) with a low complexity has been introduced in which near-optimal trees are used for implementing collective communications efficiently. Performance evaluation of these new algorithms on a 24-node experimental testbed and a 64-node simulated testbed indicates that latency of collective communication operations can be reduced by a factor up to 4.9 compared to the naive algorithms used in current MPI implementations. As the commodity LAN systems go towards heterogeneity, these new approaches provide a mean to harness maximum performance benefits from these systems.

In this paper, we have used only the heterogeneous communication capabilities of nodes to implement collective communication efficiently. We are also developing schemes to take advantage of heterogeneity in network architectures, communication protocols, and dedicated communication/synchronization units to obtain further improvements. Finally, we plan to propose a combined framework which can take advantage of all these factors and build a scalable collective communication library for HNOW systems. We are also investigating the load balancing issues in HNOW.

9 Acknowledgment

The authors are grateful to The Ohio Supercomputing Center for making their SGI cluster available for obtaining experimental results for this paper.

Additional Information: A number of related papers and technical reports are available electronically through the home page of *Parallel Architecture and Communication* (PAC) research group. The URL is `http://www.cis.ohio-state.edu/~panda/pac.html`. Alternatively, the papers can be obtained by anonymous ftp from: `ftp.cis.ohio-state.edu, cd pub/communication`.

References

[1] T. Anderson, D. Culler, and Dave Patterson. A Case for Networks of Workstations (NOW). *IEEE Micro*, pages 54–64, Feb 1995.

[2] ATM Forum. *ATM User-Network Interface Specification, Version 3.1*, September 1994.

[3] M. Barnett, S. Gupta, D. G. Payne, L. Shuler, R. van de Geijn, and J. Watts. Interprocessor Collective Communication Library (Intercom). In *Scalable High Performance Computing Conference*, pages 357–364, 1994.

[4] N. J. Boden, D. Cohen, and et al. Myrinet: A Gigabit-per-Second Local Area Network. *IEEE Micro*, pages 29–35, Feb 1995.

[5] R. V. Boppana, S. Chalasani, and C. S. Raghavendra. On Multicast Wormhole Routing in Multicomputer Networks. In *Symposium on Parallel and Distributed Processing*, pages 722–729, 1994.

[6] J. Bruck, R. Cypher, P. Elustando, A. Ho, C.T. Ho, V. Bala, S. Kipnis, and M. Snir. CCL: A Portable and Tunable Collective Communication Library for Scalable Parallel Computers. In *Proceedings of the International Parallel Processing Symposium*, 1994.

[7] J. Bruck et al. Efficient Message Passing Interface (MPI) for Parallel Computing on Clusters of Workstations. *Journal of Parallel and Distributed Computing*, pages 19–34, Jan 1997.

[8] G. T. Byrd, N. P. Saraiya, and B. A. Delagi. Multicast Communication in Multiprocessor Systems. In *International Symposium on Computer Architecture*, pages I:196–200, 1989.

[9] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction To Algorithms*. The MIT Press, 1990.

[10] J. Duato, S. Yalamanchili, and L. Ni. *Interconnection Networks: An Engineering Approach*. The IEEE Computer Society Press, 1997.

[11] E. W. Felten, R. A. Alpert, A. Bilas, M. A. Blumrich, D. W. Clark, S. N. Damianakis, C. Dubnicki, L. Iftode, and K. Li. Early Experience with Message-Passing on the SHRIMP Multicomputer. In *International Symposium on Computer Architecture (ISCA)*, pages 296–307, 1996.

[12] W. Gropp, E. Lusk, N. Doss, and A. Skjellum. A High-Performance, Portable Implementation of the MPI, Message Passing Interface Standard. Technical report, Argonne National Laboratory and Mississippi State University.

[13] C.-T. Ho and M.-Y. Kao. Optimal Broadcast in All-Port Wormhole-Routed Hypercubes. *IEEE Transactions on Parallel and Distributed Systems*, 6(2):200–204, 1995.

[14] C. Huang and P. K. McKinley. Communication Issues in Parallel Computing Across ATM Networks. *IEEE Parallel and Distributed Technology*, Winter:73–86, 1994.

[15] Y. Huang, C. C. Huang, and P. K. McKinley. Multicast Virtual Topologies for Collective Communication in MPCs and ATM Clusters. In *Proceedings of the Supercomputing*, 1995.

[16] R. Kesavan, K. Bondalapati, and D. K. Panda. Multicast on Irregular Switch-based Networks with Wormhole Routing. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA-3)*, pages 48–57, February 1997.

[17] R. Kesavan and D. K. Panda. Minimizing Node Contention in Multiple Multicast on Wormhole k-ary n-cube Networks. In *Proceedings of the International Conference on Parallel Processing*, pages I:188–195, Chicago, IL, Aug 1996.

[18] V. Kumar, A. Grama, A. Gupta, and G. Karypis. *Introduction to Parallel Computing: Design and Analysis of Algorithms*. Benjamin/Cummings, 1994.

[19] M. Lauria. High Performance MPI Implementation on a Network of Workstations. Master's thesis, Department of Computer Science, University of Illinois at Urbana-Champaign , Oct 1996.

[20] M. Lin, J. Hsieh, D. H. C. Du, J. P. Thomas, and J. A. MacDonald. Distributed Network Computing over Local ATM Networks. *IEEE Journal on Selected Areas in Communications*, 13(4), May 1995.

[21] B. Lowekamp and A. Beguelin. ECO: Efficient Collective Operations for Communication on Heterogeneous Networks. In *Int'l Parallel Processing Symposium*, pages 399–405, 1996.

[22] P. K. McKinley and D. F. Robinson. Collective Communication in Wormhole-Routed Massively Parallel Computers. *IEEE Computer*, pages 39–50, Dec 1995.

[23] P. K. McKinley, H. Xu, A.-H. Esfahanian, and L. M. Ni. Unicast-based Multicast Communication in Wormhole-routed Networks. *IEEE Transactions on Parallel and Distributed Systems*, 5(12):1252–1265, Dec 1994.

[24] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard*, Mar 1994.

[25] P. Mitra, D. G. Payne, and et al. Fast Collective Communication Libraries, Please.

[26] Lionel Ni. Should Scalable Parallel Computers Support Efficient Hardware Multicasting? In *ICPP Workshop on Challenges for Parallel Processing*, pages 2–7, 1995.

[27] N. Nupairoj and L. M. Ni. Performance Evaluation of Some MPI Implementations on Workstation Clusters. In *Proceedings of the SPLC Conference*, 1994.

[28] S. Pakin, M. Lauria, and A. Chien. High Performance Messaging on Workstations: Illinois Fast Messages (FM). In *Proceedings of the Supercomputing*, 1995.

[29] D. K. Panda. Issues in Designing Efficient and Practical Algorithms for Collective Communication in Wormhole-Routed Systems. In *ICPP Workshop on Challenges for Parallel Processing*, pages 8–15, 1995.

[30] D. K. Panda, D. Basak, D. Dai, R. Kesavan, R. Sivaram, M. Banikazemi, and V. Moorthy. Simulation of Modern Parallel Systems: A CSIM-based approach. In *Proceedings of the 1997 Winter Simulation Conference (WSC'97)*, pages 1013–1020, December 1997.

[31] J. Y. L. Park, H. A. Choi, N. Nupairoj, and L. M. Ni. Construction of Optimal Multicast Trees Based on the Parameterized Communication Model. In *Proceedings of the International Conference on Parallel Processing*, Chicago, IL, Aug 1996.

[32] F. Silla, M. P. Malumbres, A. Robles, P. Lopez, and J. Duato. Efficient Adaptive Routing in Networks of Workstations with Irregular Topology. In *Proceedings of the First International Workshop on Communication and Architectural Support for Network-Based Parallel Computing (CANPC '97)*, pages 46–60, February 1997. Available as `Lecture Notes in Computer Science #1199, Springer-Verlag`.

[33] Craig B. Stunkel, D. G. Shea, B. Abali, et al. The SP2 High-Performance Switch. *IBM System Journal*, 34(2):185–204, 1995.

[34] V. S. Sunderam. PVM: A Framework for Parallel and Distributed Computing. *Concurrency: Practice and Experience*, 2(4):315–339, December 1990.

[35] Y.-C. Tseng, D. K. Panda, and T.-H. Lai. A Trip-based Multicasting Model in Wormhole-routed Networks with Virtual Channels. *IEEE Transactions on Parallel and Distributed Systems*, 7(2):138–150, February 1996.

[36] K. Verstoep, K. Langendoen, and H. Bal. Efficient Reliable Multicast on Myrinet. In *Proceedings of the International Conference on Parallel Processing*, pages III:156–165, Aug 1996.

[37] T. von Eicken, A. Basu, V. Buch, and W. Vogels. U-Net: A User-level Network Interface for Parallel and Distributed Computing. In *ACM Symposium on Operating Systems Principles*, 1995.