

Efficient Communication Using Message Prediction for Cluster of Multiprocessors

Ahmad Afsahi

Nikitas J. Dimopoulos

Department of Electrical and Computer Engineering
Queen's University
Kingston, Canada K7L 3N6

Department of Electrical and Computer Engineering
University of Victoria
P.O. Box 3055, Victoria, Canada V8W 3P6

Email: ahmad@ee.queensu.ca
Tel.: (613) 533-3068
Fax: (613) 533-6615

Email: nikitas@ece.uvic.ca
Tel.: (250) 721-8902
Fax: (250) 721-6052

Abstract. *With the increasing uniprocessor and SMP computation power available today, interprocessor communication has become an important factor that limits the performance of cluster of workstations. Many factors including communication hardware overhead, communication software overhead, and the user environment overhead (multithreading, multiuser) affect the performance of the communication subsystems in such systems. A significant portion of the software communication overhead belongs to a number of message copying. Ideally, it is desirable to have a true zero-copy protocol where the message is moved directly from the send buffer in its user space to the receive buffer in the destination without any intermediate buffering. However, due to the fact that message-passing applications at the send side do not know the final receive buffer addresses, early arrival messages have to be buffered at a temporary area. In this paper, we show that there is a message reception communication locality in message-passing applications. We have utilized this communication locality and devised different message predictors at the receiver sides of communications. In essence, these message predictors can be efficiently used to drain the network and cache the incoming messages even if the corresponding receive calls have not been posted yet. The performance of these predictors, in terms of hit ratio, on some parallel applications are quite promising and suggest that prediction has the potential to eliminate most of the remaining message copies. We also show that the proposed predictors do not have sensitivity to the starting message reception call, and that they perform better than (or at least equal to) our previously proposed predictors in [3].*

1.0 Introduction

With the increasing uniprocessor and SMP computation power available today, interprocessor communication has become an important factor that limits the performance of workstation clusters. Essentially, communication overhead is one of the most important factors affecting the performance of parallel computers. Many factors affect the performance of communication subsystems in parallel systems. Specifically, communication hardware and its services, communication software, and the user environment (multiprogramming, multiuser) are the major sources of the communication overhead.

Communication software overhead currently dominates communication time in clusters of workstations. In the current generation of parallel computer systems, the software overheads are tens of microseconds [17]. This is worse in cluster of workstations. Even with high performance networks [11, 23, 40] available today, there is still a gap between what the network can offer and what the user application can see. The communication software overhead cost comes mainly from three different sources; crossing protection boundaries several times between the user space and the kernel space, passing several protocol layers, and involving a number of memory copying.

Several researchers are working to minimize the cost of crossing protection boundaries, and using simpler protocol layers by utilizing *user-level messaging* techniques such as *Active Messages (AM)* [44], *Fast Messages (FM)* [35], *Virtual Memory-Mapped Communications (VMMC-2)* [19], *U-Net* [45], *LAPI* [39], *Basic Interface for Parallelism (BIP)* [36], *Virtual Interface Architecture (VIA)* [20], and *PM* [43]. A significant portion of the software communication overhead belongs to a number of message copying. Ideally, message protocols should transfer messages in a single copy (this is usually called a true zero-copy). In other words, the protocol should copy the message directly from the send buffer in its user space to the receive buffer in the destination without any intermediate buffering. However, applications at the send side do not know the final receive buffer addresses and, hence, the communication subsystems at the receiving end still copy messages unnecessarily from the network interface to a system buffer, and then from the system buffer to the user buffer when the receiving application posts the receive call.

Some researchers have tried to avoid memory copying [19, 29, 37, 8, 42, 41]. While they have been able to remove the memory copying between the application buffer space and the network interface at the send side by using user-level messaging techniques, they haven't been able to remove the memory copying at the receiver sides completely. They may achieve a zero-copy messaging at the receiver sides only if the receive call is already posted, a rendez-vous type communication is used for large messages, or the destination buffer address is already known by a pre-communication. Note, however, that MPI-2 [32] supports a remote memory access (RMA) operation but this is mostly suitable for receiver-initiated communications arising from the shared-memory paradigm.

We are interested in bypassing the memory copying at the destination in the general case, synchronous or asynchronous, eager or rendez-vous and for sender-initiated communications as in MPI [31, 32]. In this paper, we argue that it is possible to address the message copying problem at the receiving side by speculation. We support our claim by showing that messages display a form of locality at the receiving ends of communications.

This paper, for the first time introduces the notion of message prediction for the receiving side of message-passing systems [3, 4]. By predicting the next receive communication call, and hence the next destination buffer address, before the receiving call is posted we will be able to copy the message directly into the CPU cache speculatively before it is needed so that an effect of a zero-copy can be achieved. As a matter of fact, communication operations at both the sender and the receiver sides may be issued at the earliest possible, while the predictive algorithms will ensure that they will be scheduled and the data made available at the optimal place and time. We are interested in utilizing similar predictors as in [1, 2], but this time at the receiver sides to predict the next consumable message and drain the network as soon as the message arrives.

The first contribution of this paper is that we show evidence that there exists message communication locality at the receiver sides of message-passing parallel applications. The second contribution of this work is the introduction and evaluation of different message predicting techniques for the receiving side of message-passing systems.

This paper concentrates on message predictions at the destinations in message-passing systems using MPI in isolation. This is analogous to branch prediction, and coherence activity prediction [34] in isolation. Our tools are not ready for measuring the effectiveness of our predictors on the application run-time yet. Our preliminary evaluation measures the accuracy of the predictors in terms of hit ratio. The results are quite promising and suggest that prediction has the potential to eliminate most of the remaining message copies.

In Section 2.0 of this paper, we explain the motivation behind this work and mention related works. We elaborate on how prediction would help eliminate the message copies at the receiving side of communications, in Section 3.0. Our experimental methodologies to gather communication traces of our parallel applications are explained in Section 4.0. In Section 5.0, we show communication frequency and unique message identifier distributions in the applications, and present evidence of message locality at the receiver sides. In Section 6.0, we propose our message predictors and present their performance on the applications. In section Section 7.0, we briefly discuss the integration of the predictors with the network interface. Finally, we conclude our paper in Section 8.0.

2.0 Motivation and Related Work

High performance computing is increasingly concerned with efficient communication across the interconnect due to the availability of high-speed highly-advanced processors. Modern switched networks, called *System Area Networks* (SAN), such as Myrinet [11] and ServerNet [23], provide high communication bandwidth and low communication latency. However, because of high processing overhead due to communication software including network interface control, flow control, buffer management, memory copying, polling and interrupt handling, users cannot see much difference compared to traditional local area networks.

Fortunately, several user-level messaging techniques have been developed to remove the operating system kernel and protocol stack from the critical path of communications [44, 35, 19, 45, 20, 36, 39, 43]. This way, applications can send and receive messages without operating system intervention which often greatly reduces the communication latency.

Data transfer mechanisms and message copying, control transfer mechanisms, address translation mechanisms, protection mechanisms, and reliability issues are the key factors for the performance of a user-level communication system. In this paper, we are particularly interested to avoid message copying at the receiver sides of communications.

A significant portion of the software communication overhead belongs to a number of message copying. With the traditional software messaging layers, there are usually four message copying operations from the send buffer to the receive buffer, as shown in Figure 1. These copies are namely from the send buffer to the system buffer (1), from the system buffer to the *network interface* (NI) (2), and at the other end of communication from the network interface to the system buffer (3), and from the system buffer to the receive buffer (4) when the receive call is posted. Note that, we haven't considered data transfer from the network interface at the sending process to the network interface at the receiving process as a separate copy. Also, the network interface's place can be either on the I/O bus or on the memory bus.

At the send side, some user-level messaging layers use programmed I/O to avoid system buffer copying. FM uses programmed I/O while AM-II and BIP do so only for small messages. Some other user-messaging layers use DMA. VMMC-2, U-Net, and PM use DMA to bypass the system buffer copy while AM-II and BIP do so only for large messages. In systems that use DMA, applications or a library dynamically pins and unpins pages in the user space that contain the send and the receive buffers. Address translation can be done using a kernel module as in BIP, or by caching a limited number of address translations for the pinned pages as in VMMC-2, U-Net/MM [9], and PM. Some network interfaces also permit bypassing message copying at the network interface by directly writing into the network.

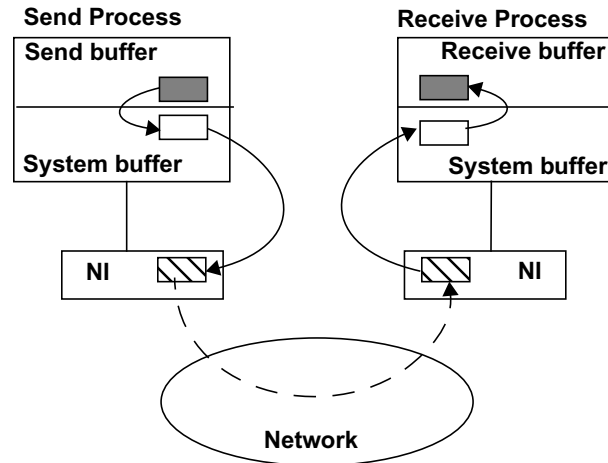


Figure 1. Data transfers in a traditional messaging layer

Contrary to the send side, bypassing the system buffer copying at the receiving side may not be achievable. Processes at the sending sides do not know the destination buffer addresses. Therefore, when a message arrives at the receiving side it has to be buffered if the receive call has not been posted yet. VMMC [10] for the SHRIMP multicomputer is a communication model that provides direct data transfer between the sender's and receiver's virtual address space. However, it can achieve zero-copy transfer only if the sender knows the destination buffer address. Therefore, the receiver exports its buffer address by scouting a message to the sender before the actual transmission can take place. This leads to a 2-phase rendez-vous protocol which adds to the network traffic, and network latency especially for short messages.

VMMC-2 [19], uses a *transfer redirection* mechanism instead. It uses a default, redirectable receive buffer for a sender who does not know the address of the receive buffer. When a message arrives at the receiving network interface, the redirection mechanism checks to see if the receiver has already posted its buffer address. If the receive buffer has been posted earlier than the message arrival, the message will be directly transferred to the user buffer. Thus it achieves a zero-copy transfer. If the buffer address is not posted, the message must be buffered in the default buffer. It will then be transferred when the receive buffer is posted. Thus, it achieves a one-copy transfer. However, if the receiver posts its buffer address when the message arrives, part of the message is buffered at the default buffer and the rest is transferred to the user buffer.

Fast sockets [37] has been built using active messages. It uses a mechanism at the receiver side called *receive posting* to avoid the message copy in the fast socket buffer. If the message handler knows that the data's final memory destination is already known upon message arrival the message is directly moved to the application user space. Otherwise, it has to be copied into the fast socket buffer.

FM 2.x [29] uses a similar approach as fast sockets, namely *layer interleaving*. FM collaborates with the handler to direct the incoming messages into the destination buffer if the receive call has already been posted.

MPI-LAPI [8] is an implementation of MPI on top of LAPI [39] for the IBM SP machines. In the implementation of the eager protocol, the header handler of the LAPI returns a buffer pointer to LAPI which tells LAPI where the packets of the message must be reassembled. If a receive call has been posted, the address of the user buffer is returned to LAPI. If the header handler doesn't find a matching receive, it will return the address of an *early arrival buffer* and hence a one-copy transfer is accomplished. Meanwhile, messages with larger sizes than the eager size are transferred using 2-phase rendez-vous protocol.

Some research projects have proposed solutions to multi-protocol message-passing interfaces on *cluster of multiprocessors* (Clumps) using both shared-memory for intra-node communications and message-passing for inter-node communications [41, 21, 30].

MPICH-PM/CLUMP [41] is an MPI library implemented on a cluster of SMPs. It uses a message-passing only model where each process runs on a processor of an SMP node. For inter-node communications, it uses *eager* and *rendez-vous* protocols. For short messages, it achieves one-copy using eager protocol as the message is copied into a temporary buffer if the MPI receive primitive has not been issued. For large message, it uses rendez-vous protocol to achieve zero-copy by using a remote write operation but it needs an extra communication. For intra-node communications, it achieves a one-copy using a kernel primitive that allows to copy messages from the sender to the receiver without the involvement of a communication buffer.

BIP-SMP [21], for intra-node communications, uses shared memory for small messages with two memory copy, and direct copy for large messages with a kernel overhead. For inter-node communications, it works like MPI-BIP which is a port of MPICH [22].

TOMPI [15] is a threaded implementation of MPI on a single SMP node. It copies a message only once by utilizing multiple threads on an SMP node. Unfortunately, it is not scalable to cluster of SMP machines.

Other techniques to bypass extra copying are the *re-mapping*, and *copy-on-write* techniques [13, 18]. Both techniques require to switch to the supervisor mode, acquire necessary locks to virtual memory data structure, and change virtual memory mapping at several levels for each page, and then perform *Translation Lookaside Buffer* (TLB)/cache consistency actions, and finally return to the user mode. This limits the performance of the page re-mapping, and copy-on-write techniques. A zero-copy TCP stack is implemented in Solaris by using copy-on-write pages and re-mapping to improve communication performance [13]. It achieves a relatively high throughput for large messages. However, it does not have a good performance for small messages. This work is also solely dedicated to the SUN Solaris virtual memory system.

fbufs [18] is also using the re-mapping technique to avoid the penalty of copying large messages across different layers of protocol stack. However, *fbufs* allows re-mapping only for a limited range of user virtual memory.

It is quite clear that the user-level messaging techniques may not achieve a zero-copy communication all the time at the receiver side of communications. Meanwhile, the major problem with all page re-mapping techniques is their poor performance for short messages which is extremely important for parallel computing.

Prediction techniques have been proposed in the past to predict the future accesses of sharing patterns and coherence activities in distributed shared memory (DSM) by looking at their observed behavior [34, 28, 25, 48, 14, 38]. These techniques assume that memory accesses and coherence activities in the near future will follow past patterns. Sakr and his colleagues have used time series and neural networks for the prediction of the next memory sharing requests [38]. Dahlgren and his colleagues devised hardware regular stride techniques to prefetch several blocks ahead of the current data block [14]. More elaborate hardware-based irregular stride prefetching approaches have been proposed by Zhang and Torrellas [48]. Kaxiras and Goodman have recently proposed an instruction-based approach which maintains the history of load and store instructions in relation to cache misses and predicting their future behavior [25]. This is in contrast to address-based techniques that keep data-access history for the predictions. Mukherjee and Hill proposed a general pattern-based predictor, *cosmos*, to learn and predict the coherence activity for a memory block in a DSM [34]. *Cosmos* makes a prediction in two steps. First, it uses a cache block address to index into a message history table to obtain the <processor and message-type> tuples of the last few coherence messages received for that cache block. Then it uses these <processor, message-type> tuples to index a pattern history table to obtain a <processor, message-type> tuple prediction. In a recent paper, Lai and Falsafi proposed a new class of pattern-based predictors, *memory sharing predictors*, to eliminate the coherence overhead on a remote access latency by just predicting the memory request messages, those primary messages that invoke a sequence of protocol actions [28]. It improves prediction accuracy over *cosmos* by eliminating the acknowledgments messages from the pattern tables. It also reduces memory overhead and perturbation in the tables due to message re-ordering. Both works in [34, 28] are adaptations of Yeh and Patt's two-level *PAP* branch predictor [47]. *PAP* is a two-level adaptive branch predictor based on the past behavior of the same branch.

In software-controlled prefetching, the programmer or compiler decides when and what to prefetch by analyzing the code and inserting *prefetch* instructions. Mowry and Gupta [33] have used software-controlled prefetching, and multithreading to hide and reduce the latency in shared memory multiprocessors.

As stated above, many prediction techniques have been proposed to reduce or hide the latency of a remote memory access in shared memory systems. Recently, Afsahi and Dimopoulos proposed some heuristics to predict the destination target of subsequent communication requests at the send

side of communications in message-passing systems [1, 2]. However, to the best of our knowledge, no prediction technique has been proposed for the receive side of communications in message-passing systems to reduce the latency of a message transfer.

This paper reports on an innovative approach for removing message copying at the receiving ends of communications for message-passing systems [3, 4]. We argue that it is possible to address the message copying problem at the receiving sides by speculation. We introduce message prediction techniques such that messages can be directly transferred to the cache even if the receive calls have not been posted yet.

3.0 Using Message Predictions

In this section, we analyze the problem with the early arrival of messages at the destinations in message-passing systems. In such systems, a number of messages arrive in arbitrary order at the destinations. The consuming process or thread will consume one message at a time. If we know which message is going to be consumed next, then we can move the message upon its arrival to near the place that it is to be consumed (e.g. a staging cache), or we could schedule which thread to execute next preferably at the same processor as the consuming thread to enhance the chances that the data will be in the processor cache when it is accessed by the consumer.

For this, we have to consider three different issues. First, deciding which message is going to be consumed next. This can be done by devising receive call predictors, history-based predictors that predict subsequent receive calls by a given process in a message-passing program, Second, deciding where and how this message is to be moved in the cache. Third, efficient cache re-mapping and late binding mechanisms need to be devised for when the receive call is posted.

In this work, we are addressing the first problem. That is, devising message predictors and evaluating their performance. We are working on several methods to address the remaining issues. We shall report on these issues in the future.

4.0 Experimental Methodology

In exploring the effect that different heuristics have in predicting the next receive call, we utilized a number of parallel benchmarks, and extracted their communication traces on which we applied our predictors. We have used some well-known parallel benchmarks from the *NAS parallel benchmarks* (NPB) suite [7], and the *Parallel Spectral Transform Shallow Water Model* (PSTSWM) application [46]. We used the MPI [31] implementation of the NPB suite (version 2.3), and version 6.2 of the PSTSWM application. Specifically, we used the, *block tridiagonal* (BT), *scalar pentadiagonal* (SP), and *conjugate gradient* (CG) benchmarks from NPB suite [7]. We didn't use the *multigrid* (MG) and *lower-upper diagonal* (LU) benchmarks from the NPB suite because these benchmarks use

MPI_ANY_SOURCE in some of their receive calls (*MPI_Recv* and *MPI_Irecv*). This means that the applications may receive a particular message from different sources depending on the order of arrival.

We are only interested in the patterns of the point-to-point communications between pair-wise processes in our applications. For this, we executed these applications on an IBM SP2 machine. We wrote our own profiling code using the wrapper facility of the MPI to gather the communication traces. We did this by inserting monitor operations in the profiling MPI library for the communication related activities. These operations include arithmetic operations for the calculation of the desired characteristics. Collecting communication traces does not affect the communication patterns of the applications.

We considered different system sizes and problem sizes for our applications to evaluate the performance of our prediction heuristics. Specifically, we experimented with the workstation class “W”, and the large class “A” of the NPB suite, and the default problem size for the PSTSWM application. The NPB results are almost the same for “W” and “A” classes. Hence, we report only for the “A” class here. Note that we also removed the initialization part from the communication traces of the PSTSWM application. Although the results presented in this paper are for the above parallel applications, these applications have been widely used as benchmarks representing the computations in scientific and engineering parallel applications.

5.0 Receiver-side Locality Estimation

Our applications use blocking and nonblocking standard MPI receive primitives, namely *MPI_Recv* and *MPI_Irecv* [31]. *MPI_Recv* (*buf*, *count*, *datatype*, *source*, *tag*, *comm*, *status*) is a standard blocking receive call. When it returns, data is available at the destination buffer. The PSTSWM application uses this type of receive call. *MPI_Irecv* (*buf*, *count*, *datatype*, *source*, *tag*, *comm*, *request*) is a standard nonblocking receive call. It immediately posts the call and returns. Hence, data is not available at the time of return. It needs another call to complete the call. All applications in our study use this type of receive call.

One of the communication characteristics of any parallel application is the frequency of communications. Figure 2 illustrates the minimum, average, and maximum number of receive communication calls in the applications under different system sizes. We ran our applications once for each different system size and counted the number of receive calls for each process of the applications. Hence, in Figure 2, by average, minimum, and maximum, we mean the average, minimum, and maximum number of receive calls taken over all processes of each application. It is clear that all processes in the BT, SP, and CG applications have the same number of receive communication calls. While processes in the PSTSWM application have different number of receive communication calls.

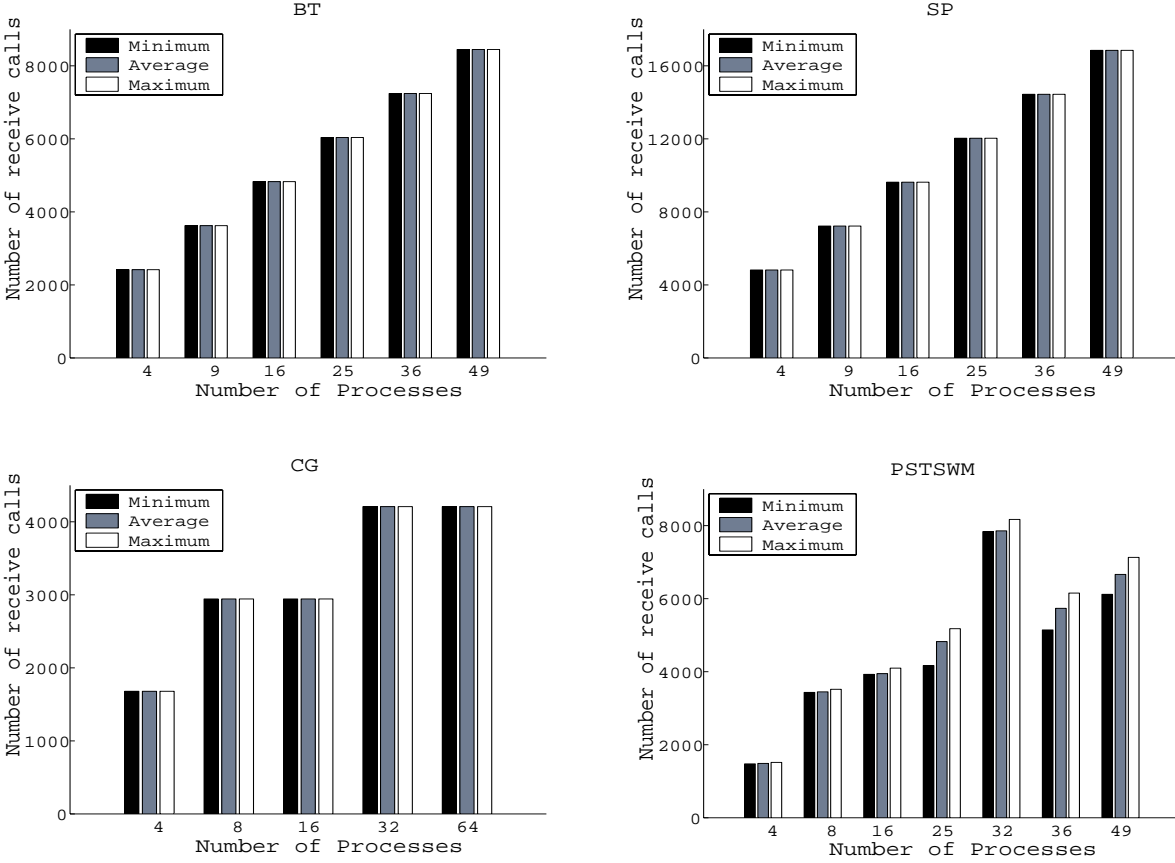


Figure 2. Number of receive calls per process in the applications under different system sizes

As stated earlier, *MPI_Recv* and *MPI_Irecv* calls have a 7-tuple set consisting of *source*, *tag*, *count*, *datatype*, *buf*, *comm*, and *status* or *request*. In order to choose precisely one of the received messages at the network interface and transfer it to the cache, our predictors need to consider all the details of a message envelop. That is, *source*, *tag*, *count*, *datatype*, *buf*, and *comm* (we don't consider *status* and *request* as they are just a handle when the calls return). We did not rely only on the buffer address, *buf*, of a receive call as many processes may send their messages to the same buffer address of a particular destination process. Neither we could depend only on the sender, *source*, of a message, or on the length, *count*, of a message. Therefore, we assigned a different identifier for each unique 6-tuple found in the communication traces of the applications. Figure 3 shows the number of *unique message identifiers* in our applications under different system sizes. By average, minimum, and maximum, we mean the average, minimum, and maximum number of unique identifiers taken over all processes of each application. It is evident that all processes in the BT, and CG applications have the same number of unique message identifiers while processes in the SP, and PSTSWM applications have different number of unique message identifiers (except when the number of processes is four for the SP benchmark).

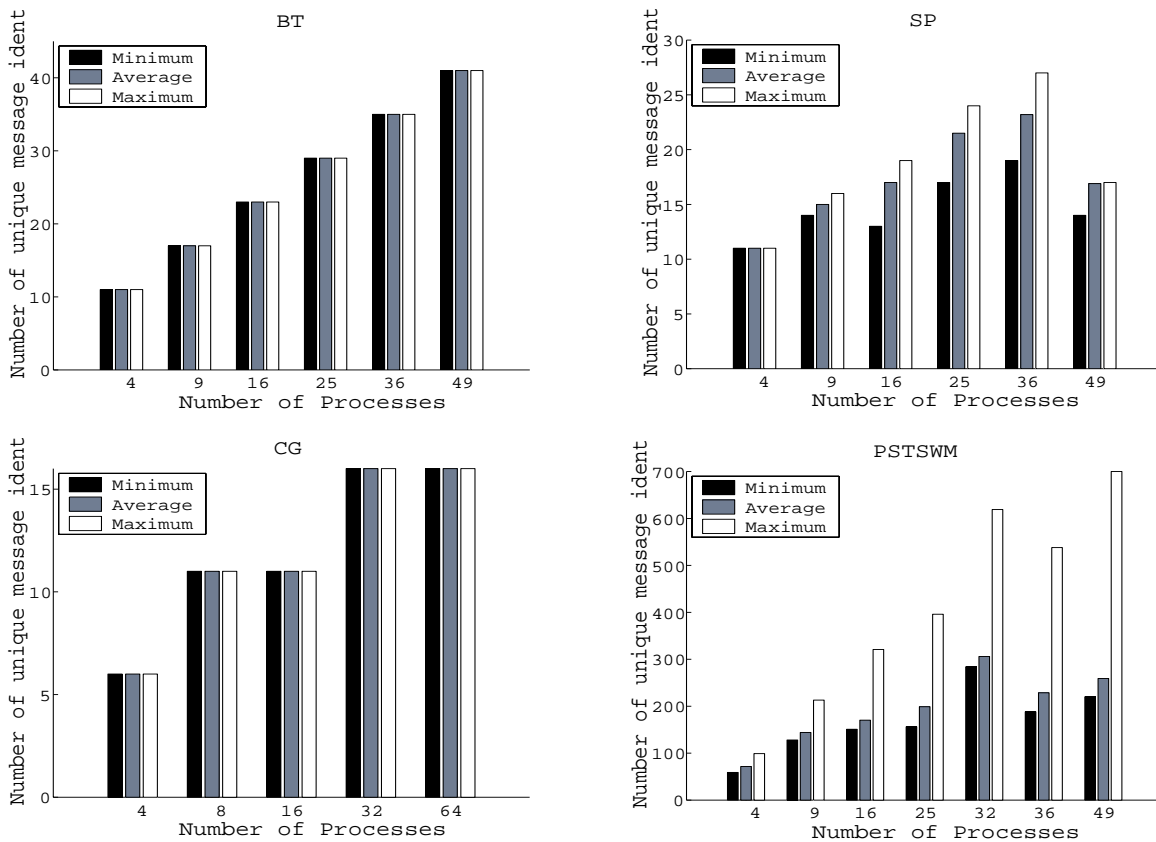


Figure 3. Number of unique message identifiers per process in the applications under different system sizes

Figure 4 shows the distribution of each unique message identifier for process zero of the applications when the number of processes is 64 for CG and 49 for the other applications. We chose process zero because this process has the largest number of unique message identifiers among all processes and is also responsible for distributing data and verifying the results of the computation. As it is shown in Figure 3, the message identifiers are evenly distributed in BT. However, the distribution of the message identifiers in CG and PSTSWM are almost bimodal with two separated peaks. The SP benchmark shows four different peaks for the message identifiers.

5.1 Communication Locality

In the context of message passing programming, many parallel algorithms are built from loops consisting of computation and communication phases. Therefore, communication patterns may be repetitive. This has motivated researchers to find or use the *communications locality* properties of parallel applications [1, 2, 26, 24, 27, 16, 12]. Kim and Lilja [26] have shown that there is a locality in message destination, message sizes, and consecutive runs of send/receive primitives in parallel algorithms. They have proposed and expanded the concept of memory access locality based on the *Least Recently Used*, LRU, stack model to determine these localities. In [1, 2], Afsahi and Dimopoulos have

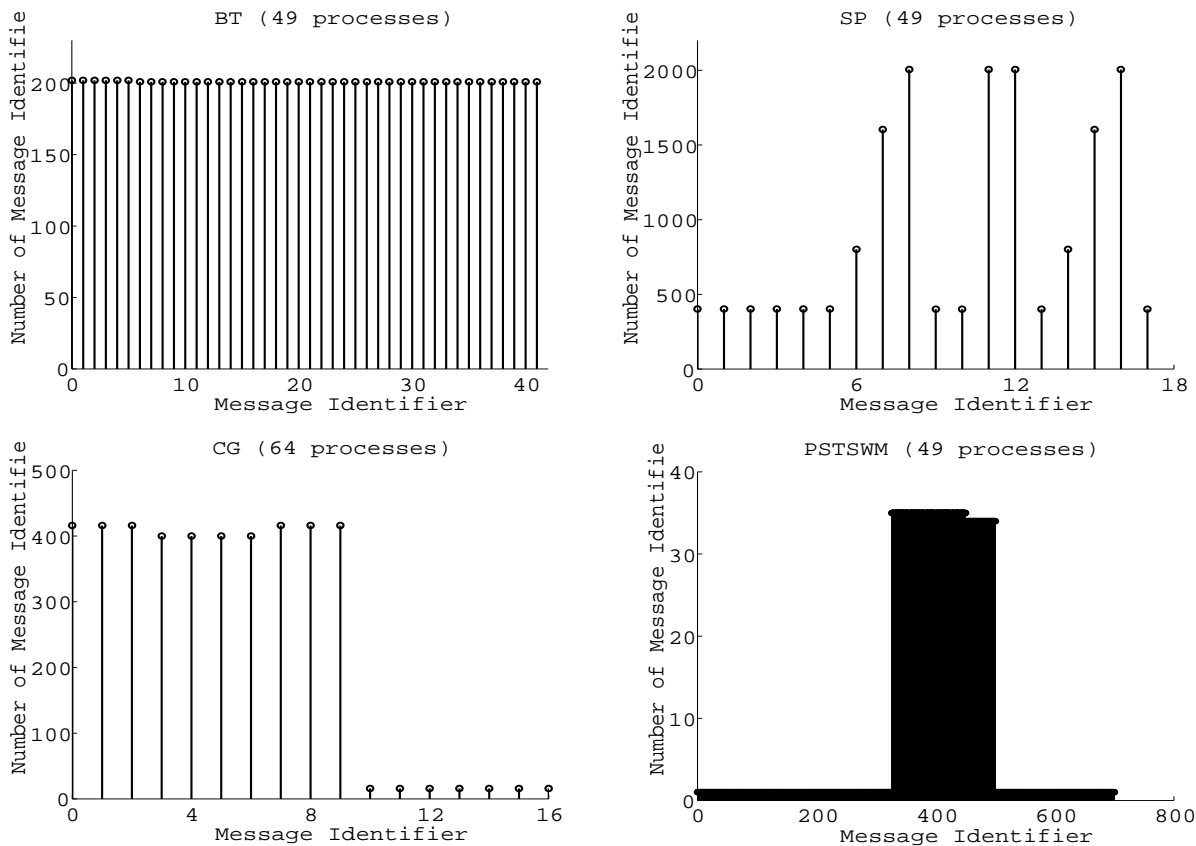


Figure 4. Distribution of the unique message identifiers for process zero in the applications

shown the communication locality of message-passing application in terms of message destination locality. Karlsson and Brorsson [24] have compared the communication properties of parallel applications in message-passing systems using MPI, and shared memory systems using TreadMarks [6].

We define the terms *message reception locality* in conjunction with this work. By message reception locality we mean that if a certain message reception call has been used it will be re-used with high probability by a portion of code that is “near” the place that was used earlier, and that it will be re-used in the near future.

In the following subsection, we present the performance of the classical LRU, LFU, and FIFO heuristics on the applications to see the existence of locality or repetitive receive calls. We use the *hit ratio* to establish and compare the performance of these heuristics. As a hit ratio, we define the percentage of times that the predicted receive call was correct out of all receive communication requests.

5.2 The LRU, FIFO and LFU Heuristics

The *Least Recently Used* (LRU), *First-In-First-Out* (FIFO), and *Least Frequently Used* (LFU) heuristics, all maintain a set of k (k is the window size) unique message identifiers. If the next message identifier is already in the set, then a hit is recorded. Otherwise, a miss is recorded and the new message identifier replaces one of the identifiers in the set according to which of the LRU, FIFO or LFU strategies is adopted.

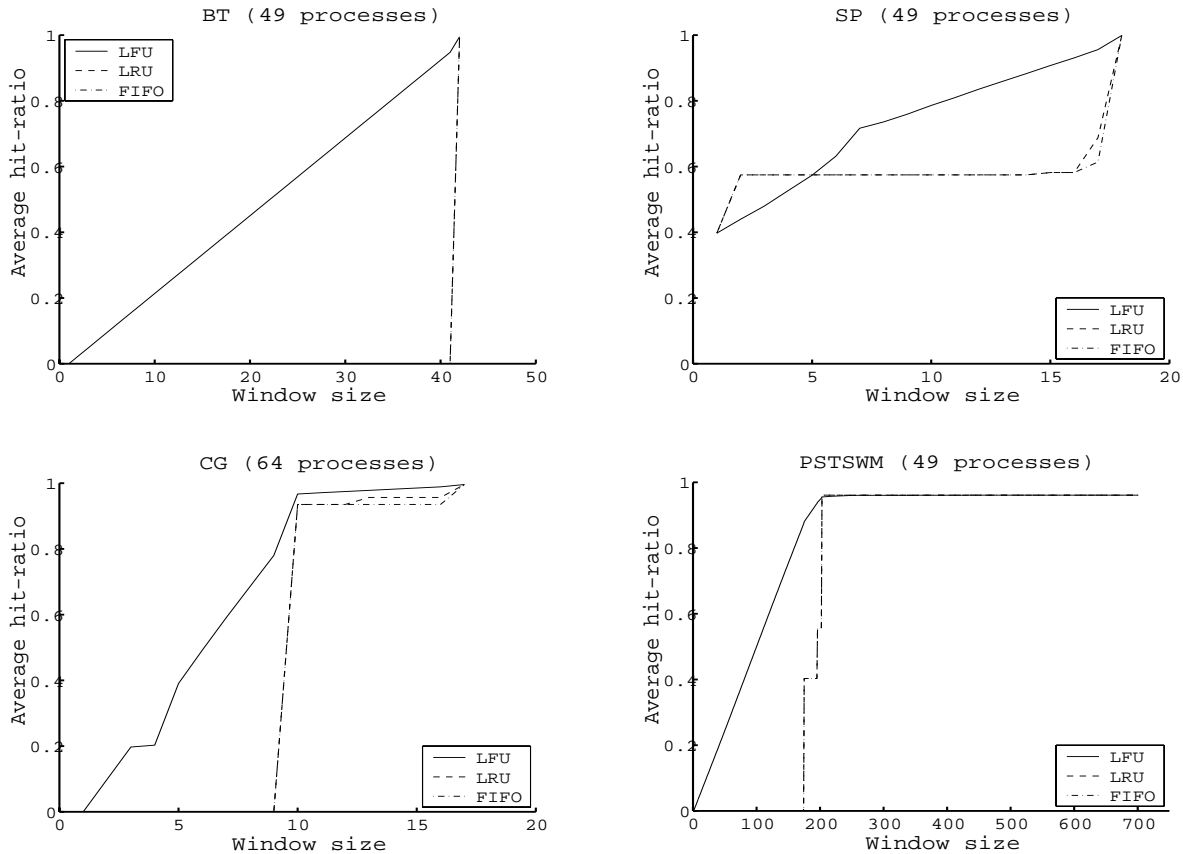


Figure 5. Effects of the LRU, FIFO, and LFU heuristics on the applications

Figure 5 shows the results of the LRU, FIFO, and LFU heuristics on the application benchmarks when the number of processes is 64 for CG and 49 for all other applications. It is clear that the hit-ratios in all benchmarks approach 1 as the window size increases. The performance of the FIFO algorithm is the same as the LRU for BT, and PSTSWM benchmarks, and almost the same for the SP and CG benchmarks. The LFU algorithm consistently has a better performance than the LRU and FIFO heuristics on the BT, CG, and PSTSWM applications. It also has a better performance than the LRU and FIFO heuristics on the SP benchmark for window sizes of greater than five. It is interesting to see that a real application like PSTSWM needs window sizes of greater than 150 to achieve a good performance (hit ratios above 80%) under the LFU policy. Similar performance results for the LRU, FIFO, and LFU heuristics on other system sizes can be found in [5].

Essentially, the LRU, FIFO and LFU heuristics do not predict exactly the next receive call but show the probability that the next receive call might be in the set. For instance, the SP benchmark shows nearly 60% hit ratio for a window size of five under the LRU heuristic. This means that 60% of the time one of the five most recently issued call will be issued next. These heuristics perform better when the window size k is sufficiently large. However, this large window adds to the hardware/software implementation complexity as one need to move all messages in the set to the cache in the likelihood that one of them is going to be used next. This is prohibitive for large window sizes.

We are interested in having predictors that can predict the next receive call with a high probability. In Section 6.0, we introduce our novel message predictors employing different heuristics and evaluate their performance.

6.0 Message Predictors

The set of predictors proposed in this section predict the subsequent receive calls based on the past history of communication patterns on a per process basis. We propose two sets of predictors in this work: *Single-cycle* predictor, which is purely a dynamic predictor, and *Tag-based* predictors, which are static/dynamic predictors. In the *Single-cycle* predictor, predictions are done dynamically at the network interface without any help from the programmer or compiler. In the *Tag-based* predictors, *Tagging*, *Tag-cycle*, and *Tag-bettercycle*, predictions are done dynamically at the network interface as well, but they require some information to be passed from the program to the network interface. This can be done with the help of the programmer and/or the compiler through inserting instructions such as *pre-receive (tag)* in the program. The *Tag-based* predictors can be pure dynamic predictors if another level of prediction is done on the tag themselves at the network interface. This way, there is no need for the program to pass pre-connect (tag) information to the network interface. We leave this approach for the future research. These heuristics were originally proposed in [1, 2] to predict the destination target of subsequent communication requests at the send side of communications. However in this paper, we have added an “initialization phase” to the *Single-cycle*, *Tag-cycle*, and *Tag-bettercycle* predictors to remove the initialization patterns of communications, if any [4]. We explain this initialization phase in Section 6.1.1.

It is worth mentioning that the message re-ordering effect [28] (messages from different processes may arrive out-of-order even if messages from the same processes may arrive in-order in most networks) has no effect on the predictions as the predictors predict the next receive calls based on the patterns of the receive calls in the program that runs on the same process and not on the arriving messages unless the order of receive calls depends on the order of message arrival. Note that in the following figures, by average, minimum, and maximum, we mean the average, minimum, and maximum hit ratio taken over all processes of each application.

6.1 The Single-cycle Predictor

The *Single-cycle* predictor is consisted of an “initialization phase” followed by a “prediction phase”. The “prediction phase” is based on the fact that if a group of receive calls are issued repeatedly in a cyclical fashion, then we can predict the next request one step ahead. Figure 6 illustrates an example for the operation of the single-cycle predictor. The top trace represents the sequence of requested receive calls, while the bottom trace represents the predicted sequence. The arrows with the cross represent misses, while the ones with the circle represent hits. The “dash” in place of a predicted request indicates that a cycle is being formed, and therefore no prediction is offered (note that this is also added to the misses).




Figure 6. Operation of the Single-cycle predictor on the sample request sequence

This predictor implements a simple cycle discovery algorithm. Starting with a *cycle-head* receive call (this is the first receive call that is requested at start-up, or the receive call that causes a miss), we log the sequence of requests until the cycle-head receive call is requested again. Note that during cycle formation, the previously requested message destination is offered as the predicted message destination. This stored sequence constitutes a cycle, and can be used to predict the subsequent requests. If the predicted receive call coincides with the subsequent requested one, then we record a hit. If the requested receive call does not coincide with the predicted one, then we record a miss and the cycle formation stage commences with the cycle-head being the receive call that caused the miss.

6.1.1 Initialization Phase

We have added an initialization phase to the *Single-cycle*, *Tag-cycle*, and *Tag-bettercycle* predictors proposed in [1, 2, 3]. As stated earlier, the prediction phase of the Single-cycle predictor (and all other predictors) starts with the first message reception call as the cycle-head. However, it is possible that we never find this cycle-head again to establish a cycle. Thus, the predictor fails with a zero prediction hit ratio. This may also happen if we skip the first receive call and start forming the cycles with the second receive call as the cycle-head. In other words, applications may have a sequence of receive calls at the beginning that may never come back to them and hence no cycle is formed. This is illustrated in the example in Figure 7. It is clear that no cycle will be formed with the message cycle-heads 1, 3, 5, and 6. However, if we skip ahead we see that a cycle is formed with the message cycle-

head 4. This cycle consists of message calls 4, 7, 8, 9, 10, 11, and 12. It is clear that the first four message reception calls, 1, 3, 5, and 6 are considered the initialization pattern. Therefore, the new *Single-cycle* predictor (along with the *Tag-cycle*, and *Tag-bettercycle* predictors) proposed in this paper starts with the “initialization phase”, and then after forming the first cycle switches to the “prediction phase”. It is clear that misses are added until the first cycle is formed. Note that in the “initialization phase” of the predictors, we looked for a cycle with a length of greater than 5. This was chosen arbitrarily so as not to form a small cycle.




Figure 7. Message reception call sequence (from left to right)

The performance of the Single-cycle predictor is shown in Figure 8. It is evident that its performance is consistently very high (hit ratios of more than 0.9).

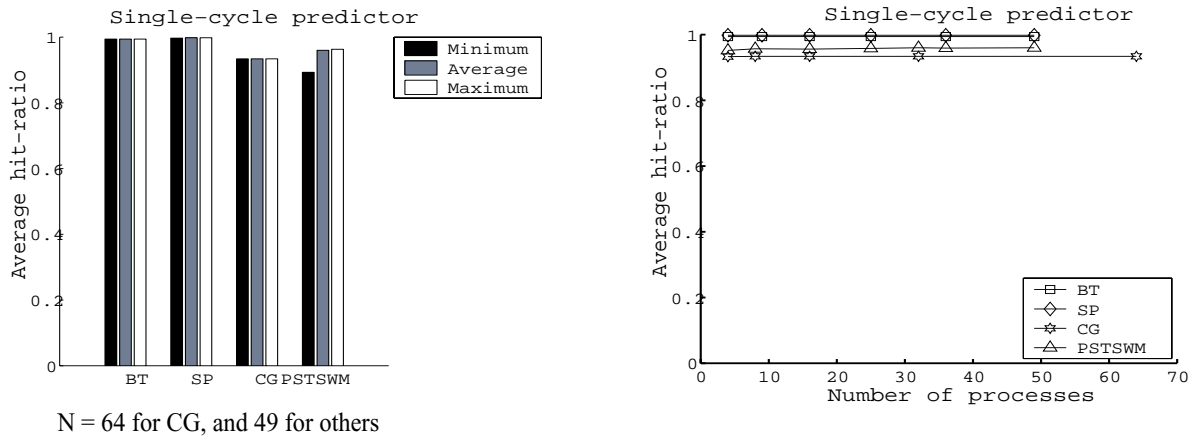


Figure 8. Effects of the Single-cycle predictor on the applications

6.2 The Tagging Predictor

The *Tagging* predictor assumes a static communication environment in the sense that a particular communication receive call in a section of code, will be the same one with a large probability. We attach a different *tag* (this is different than the tag in an MPI communication call; It may be a unique identifier or the program counter at the address of the communication call) to each of the receive calls found in the applications. This can be implemented with the help of the compiler or by the programmer through a *pre-receive (tag)* operation which will be passed to the communication subsystem to predict the next receive call before the actual receive call is issued.

To this tag and at the communication assist, we assign this receive call. A hit is recorded if in subsequent encounters of the tag, the requested communication is the same as the receive call already associated with the tag. Otherwise, a miss is recorded and the tag is assigned the newly requested

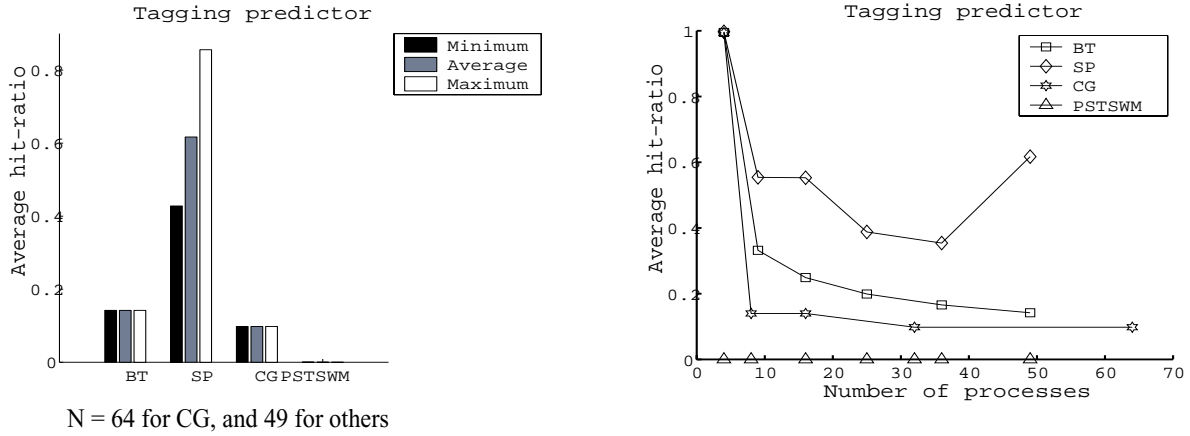


Figure 9. Effects of the Tagging predictor on the applications

receive call. The performance of the Tag predictor is shown in Figure 9. It is evident that this predictor doesn't have a good performance on the applications. It cannot predict the communication patterns of PSTSWM at all, and has a degrading performance for all other applications when the number of processes increases.

6.3 The Tag-cycle Predictor

The Tag predictor didn't have a good performance on the applications while the Single-cycle predictor had a very good performance. We would like to see the impact of the cycle algorithm on the Tag predictor. Therefore, we combine the Tag algorithm with the Single-cycle algorithm and call it the *Tag-cycle* predictor.

In the Tag-cycle predictor, we attach a different tag to each of the communication requests found in the benchmarks and do a Single-cycle discovery algorithm on each tag. To this tag and at the communication assist, we assign the requested receive call, to be called *tagcycle-head* (this is the first receive call that is requested at this tag, or the call that causes a miss). We log the sequence of the requests at this tag until the tagcycle-head is requested again. This stored sequence constitutes a cycle at each tag, and can be used to predict the subsequent requests. The performance of the Tag-cycle predictor is shown in Figure 10. The Tag-cycle predictor performs well on all benchmarks. Its performance is the same as the Single-cycle predictor on BT and PSTSWM. However, it has a better performance on CG and a lower performance on SP.

6.4 The Tag-bettercycle Predictor

In the Tag-cycle predictor, as soon as a receive call breaks a cycle we remove the cycle and form a new cycle. In the *Tag-bettercycle* predictor, we keep the last cycle associated with each tagcycle-head encountered in the communication patterns of each process. This means that when a cycle breaks we

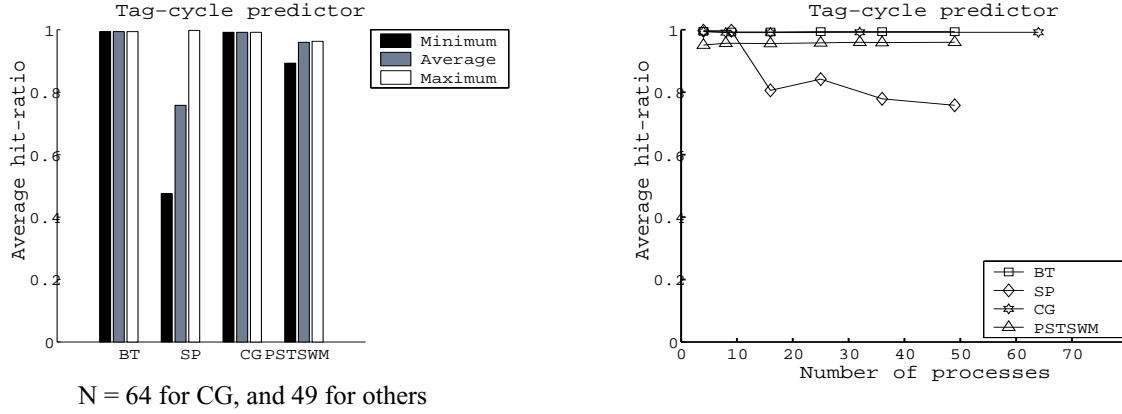


Figure 10. Effects of the Tag-cycle predictor on the applications

maintain this cycle in memory for later references. If we haven't already seen the new tagcycle-head then we form a new cycle for it, otherwise we predict the next communication call based on the member of the cycle associated with this new tagcycle-head that we have from the past in memory.

The performance of the Tag-bettercycle predictor is shown in Figure 11. The Tag-bettercycle predictor performs well on all benchmarks. Its performance is the same as the Single-cycle and Tag-cycle predictors on BT and PSTSWM. However, it has a better performance on the CG and a lower performance on SP relative to the Single-cycle predictor. The Tag-bettercycle predictor has a better performance on SP compared to the Tag-cycle predictor. We also found that the applications have very small number of tagbettercycle-heads (at most 2) under the Tag-bettercycle predictor and different system sizes.

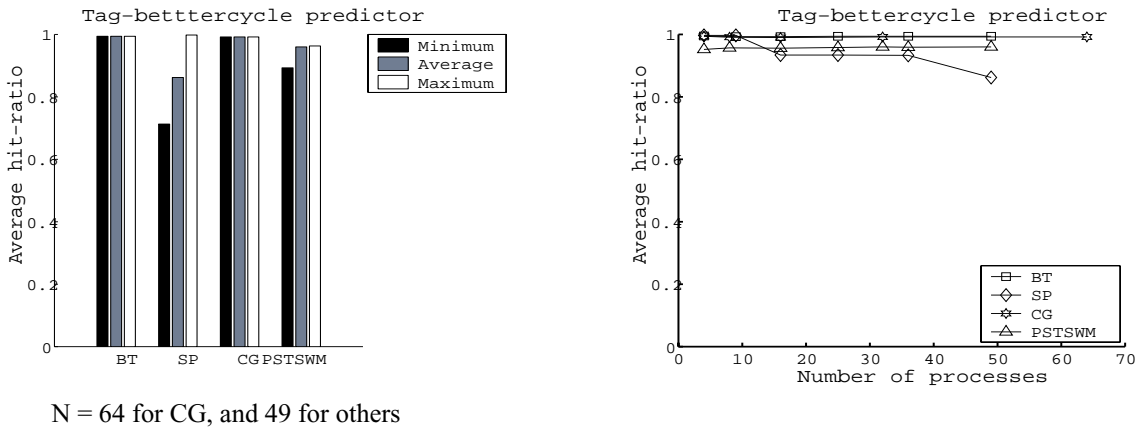


Figure 11. Effects of the Tag-bettercycle predictor on the applications

6.5 Message Predictors' Comparison

Figure 12 presents a comparison of the performance of the predictors presented in this paper when the number of processes is 64 for CG and 49 for the other benchmarks. As we have seen so far, Single-cycle, Tag-cycle and Tag-bettercycle all perform well on the benchmarks. However, the perfor-

mance of the Single-cycle is better on the SP benchmark while Tag-cycle and Tag-bettercycle have better performance for the CG benchmark. Similar results have been found for the other system sizes [5].

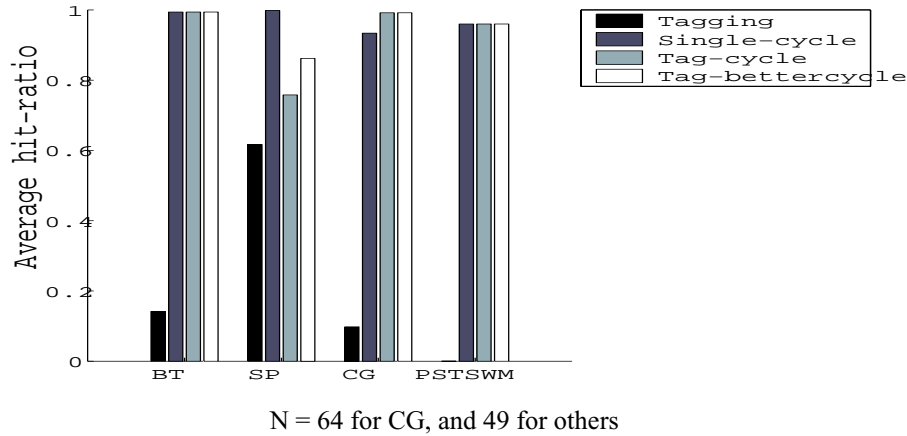


Figure 12. Comparison of the performance of the predictors on the applications

6.5.1 Predictor’s Memory Requirements

Table 1 compares the maximum memory requirement of the message predictors on the application benchmarks when the number of processes is 64 for CG, and 49 for BT, SP, and PSTSWM. We have found that the memory requirement of the predictors decrease gradually when the number of processes decreases. The numbers in the table are the multiplication factor for the amount of storage needed to maintain the message 6-tuple sets. It is quite clear that the memory requirements of the predictors is low. This makes them very attractive for the implementation at the network interface. Comparatively, predictors (Single-cycle, Tag-cycle, and Tag-bettercycle) need higher memory requirement for the PSTSWM application. Although, the classical LRU, LFU, and FIFO heuristics need less memory requirements, but as stated earlier, the beauty of the predictors lies on the fact that they predict with high accuracy and transfer only one message to the cache which should dramatically reduce the cache pollution effect, if any. This should also bring down the software cost of the implementation.

Table 1: Memory requirements (in 6-tuple sets) for the predictors when $N = 64$ for CG, and $N = 49$ for BT, SP, and PSTSWM

	BT	SP	CG	PSTSWM
Tagging	12	12	10	7
Single-cycle	43	43	138	204
Tag-cycle	60	72	40	693
Tag-bettercycle	60	108	40	693

The storage requirement of the predictors have been found using the following formulae:

$$Mem_{Single-cycle} = \text{Maximum cycle length} \quad (1)$$

$$Mem_{Tagging} = \text{Maximum number of tags} \quad (2)$$

$$Mem_{Tag-cycle} = Mem_{Tagging} \times \text{Maximum cycle length of each tags} \quad (3)$$

$$Mem_{Tag-bettercycle} = Mem_{Tag-cycle} \times \text{Maximum number of tagbettercycle-heads} \quad (4)$$

6.6 Sensitivity Analysis of the Message Predictors

In this section, we are interested in discovering the prediction sensitivity of the *Single-cycle*, *Tag-cycle*, and *Tag-bettercycle* message predictors to the varying starting message reception call. That is, what would be the difference in the predictor's performance if we begin the predictors with the first message reception call, second message reception call, and so on. Actually, we applied the predictors on the benchmarks starting with each of the first 100 message reception calls. Figure 13 compares the average hit ratios of the proposed predictors in this paper (that is with the initialization phase) over the starting first 100 message reception calls with the predictors without the initialization phase as proposed in [3]. It is clear that the revised predictors proposed in this work perform better than the previously proposed predictors in [3] for the PSTSWM application. However, they have the same average performance on BT, SP, and CG benchmarks. Meanwhile, the average performance of the predictors, with varying starting message reception calls, are more than 95% on the BT, CG, and PSTSWM benchmarks, and more than 75% on the SP benchmark.

7.0 Integrating Message Predictors with the Network Interface

In this section, we briefly discuss how a message predictor can be used and integrated into the network interface. Predictors would reside beside the communication assist or network interface. They monitor the message reception call patterns of their host process and make a prediction according to their prediction algorithms. Then, the network interface uses the predictions to move the early arrival messages into the cache.

As stated above, the predictors would execute on the communication assist of each node of the parallel machine, and predict the message reception calls based on the past history of communications. The Single-cycle predictor does not need any help from the compiler or programmer. However, the Tag-based predictors (Tagging, Tag-cycle, and Tag-bettercycle) require an interface to pass some information from the program to the network interface. With a simple help from the programmer or compiler, this can be done through inserting *pre-recv* (*tag*) instructions in the program well above

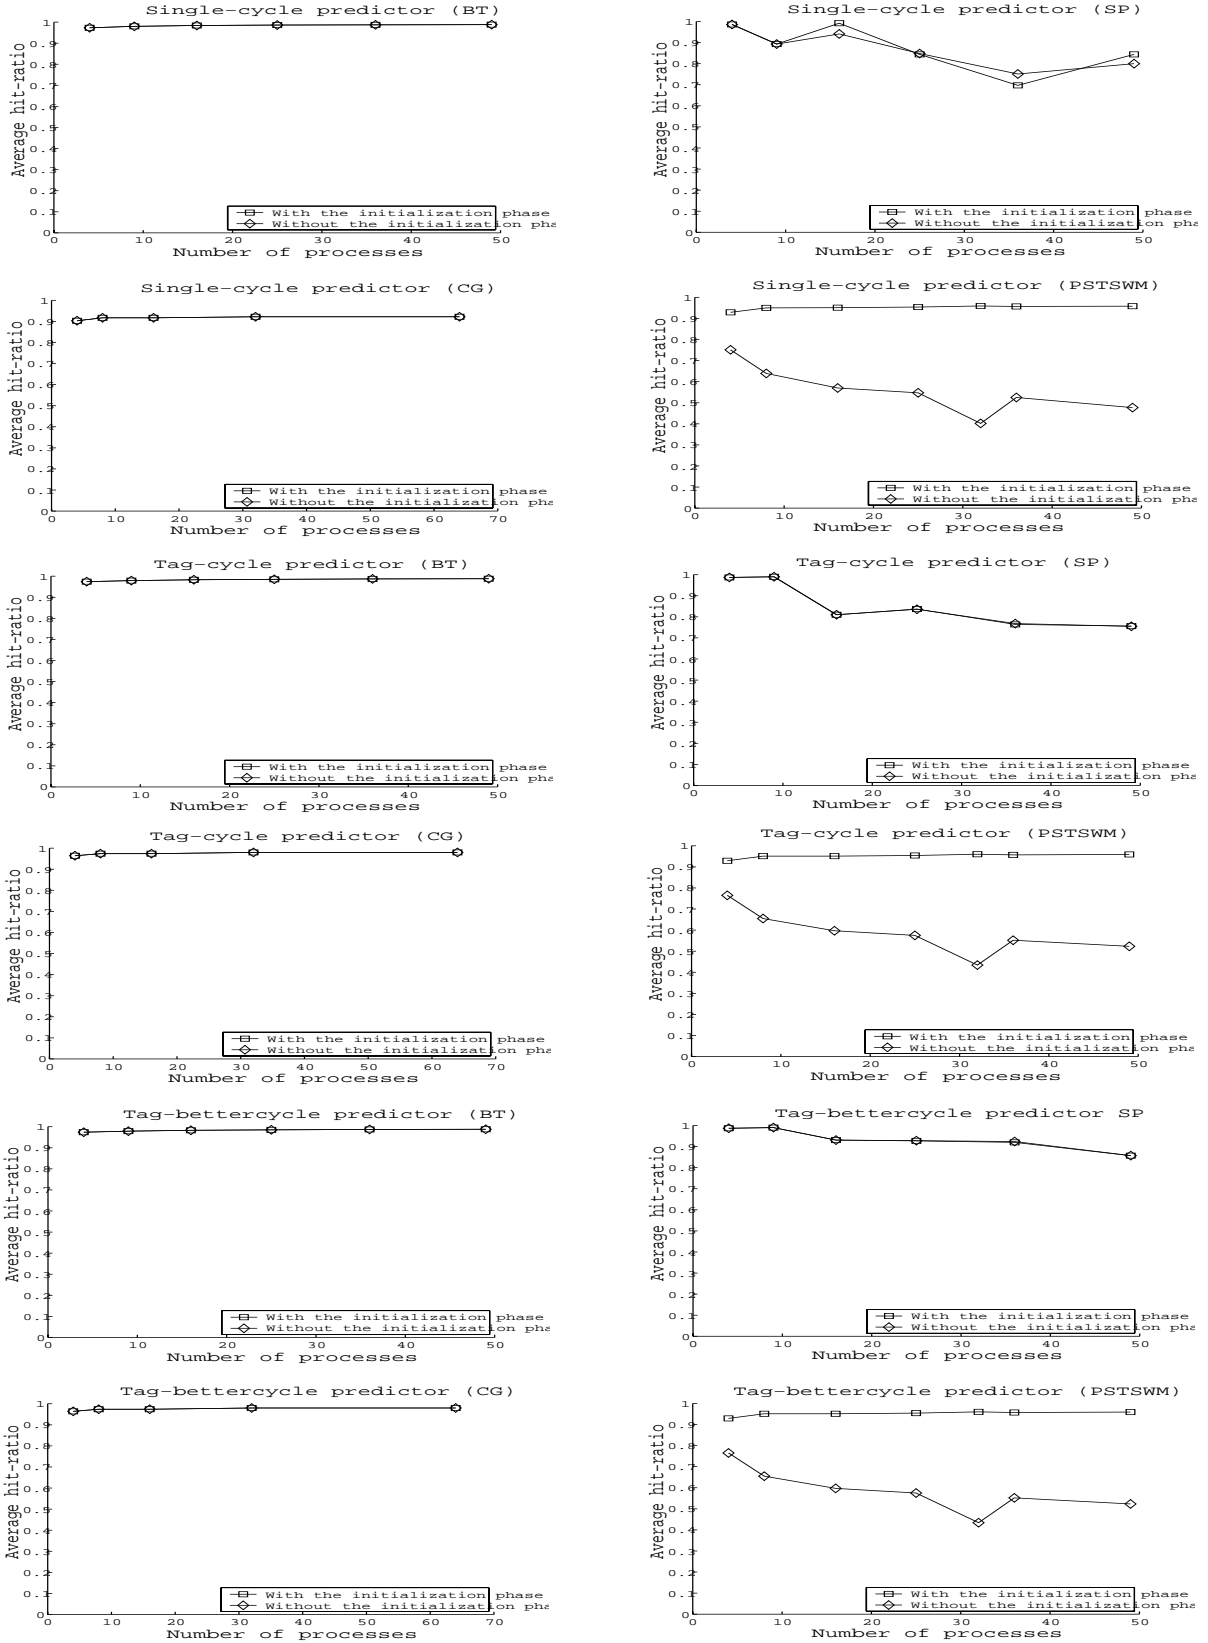


Figure 13. Comparison of the performance of the predictors with the initialization phase with the predictors without the initialization phase over the starting first 100 message reception calls

each specific receive communication operation but evidently after the previous receive communication operation. The Tag-based predictors can be pure dynamic predictors if another level of prediction is done on the tag themselves at the network interface. This way, there is no need for the program to pass pre-receive (tag) information to the network interface. However, the performance of these *2-level Tag-based* prediction techniques has not been evaluated yet.

8.0 Conclusion

Communication latency adversely affects the performance of networks of workstations. A significant portion of the software communication overhead belongs to a number of message copying operations. Ideally, it is very desirable to have a true zero-copy protocol where the message is moved directly from the send buffer in its user space to the receive buffer in the destination without any intermediate buffering. However, this is not always possible as a message may arrive at the destination where the corresponding receive call has not been issued yet. Hence, the message has to be buffered in a temporary buffer.

In this paper, we have shown that there is a message reception communication locality in message-passing applications. We have utilized this communication locality and devised different message predictors for the receiver sides of communications. By predicting receive calls early, a node can perform the necessary data placement upon message reception and move the message directly into the cache. We presented the performance of these predictors on some parallel applications. The performance results are quite promising and justify more work in this area. We also presented that these predictors are not sensitive to the starting message reception call and can capture the patterns of communications for the correct prediction.

We envision these predictors to be used to drain the network and place the incoming messages in the cache in such a way so as to increase the probability that the messages will still be in cache when the consuming thread needs to access them.

Further issues we are presently investigating include mechanisms for in-the-cache late binding and thread scheduling to guarantee that the consuming thread finds the message in the cache of the processor it executes on. We shall report on these issues in the future.

Acknowledgments

This work was supported by grants from Natural Sciences and Engineering Research Council of Canada (NSERC) and the University of Victoria. We would like to thank Dr. Murray Campbell at the IBM T J Watson Research Center for his kind help in accessing the IBM Deep Blue machine, and the computer staff at the University of Victoria for the access to the University of Victoria IBM SP2.

References

- [1] A. Afsahi and N. J. Dimopoulos, "Hiding Communication Latency in Reconfigurable Message-Passing Environments", *Proceedings of the of IPPS/SPDP 1999, 13th International Parallel Processing Symposium and 10th Symposium on Parallel and Distributed Processing*, April 1999, pp. 55-60.
- [2] A. Afsahi and N. J. Dimopoulos, "Communication Latency Hiding in Reconfigurable Message-Passing Environments: Quantitative Studies", *13th Annual International Symposium on High Performance Computing Systems and Applications, HPCS'99*, June, 1999, pp. 111-126.
- [3] A. Afsahi and N. J. Dimopoulos, "Efficient Communication Using Message Prediction for Cluster of Multiprocessors", *Proceedings of the CANPC'00, Fourth Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing, held in conjunction with the 6th International Symposium on High-Performance Computer Architecture, HPCA-6*, January 2000.
- [4] A. Afsahi and N. J. Dimopoulos, "Communication Prediction in Message-Passing systems", *14th Annual International Symposium on High Performance Computing Systems and Applications, HPCS'2000*, June, 2000.
- [5] A. Afsahi, "Design and Evaluation of Communication Latency Hiding/Reduction Techniques for Message-Passing Environments", *Department of Electrical and Computer Engineering, University of Victoria*, April 2000.
- [6] C. Amza, A. L. Cox, S. Dwarkadas, P. Keleher, H. Lu, R. Rajamony, W. Yu and W. Zwaenepoel, "TreadMarks: Shared Memory Computing on Networks of Workstation", *IEEE Computer*, Volume 29, no. 2, February 1996, pp. 18-28.
- [7] D. H. Bailey, T. Harsis, W. Saphir, R. V. der Wijngaart, A. Woo and M. Yarrow, "The NAS Parallel Benchmarks 2.0: Report NAS-95-020", *Nasa Ames Research Center*, December 1995.
- [8] M. Banikazemi, R. K. Govindaraju, R. Blackmore and D. K. Panda, "Implementing Efficient MPI on LAPI for IBM RS/6000 SP Systems: Experiences and Performance Evaluation", *Proceedings of the of IPPS/SPDP 1999, 13th International Parallel Processing Symposium and 10th Symposium on Parallel and Distributed Processing*, April 1999, pp. 183-190.
- [9] A. Basu, M. Welsh, T. V. Eicken, "Incorporating Memory Management into User-Level Network Interfaces", *Hot Interconnects V*, August 1997.
- [10] M. Blumrich, K. Li, R. Alpert, C. Dubnicki, E. Felten, and J. Sandberg, "A Virtual Memory Mapped Network Interface for the SHRIMP Multicomputer", *Proceedings of the 21st Annual International Symposium on Computer Architecture*, 1994, pp. 142-153.
- [11] N. J. Boden, D. Cohen, R. E. Felderman, A. E. Kulawik, C. L. Seitz, J. N. Seizovic and W-K. Su, "Myrinet: A Gigabit-per-Second Local Area Network", *IEEE Micro*, February 1995.
- [12] S. Chodnekar, V. Srinivasan, A. Vaidya, A. Sivasubramaniam and C. Das, "Towards a Communication Characterization Methodology for Parallel Applications", *Proceedings of the Third International Symposium on High Performance Computer Architecture*, 1997.

- [13] H. Chu, “Zero-copy TCP in Solaris”, *Proceedings of the USENIX Annual Technical Conference*, 1996, pp. 253-263.
- [14] F. Dahlgren, M. Dubois and P. Stenström, “Sequential Hardware Prefetching in Shared-Memory Multiprocessors”, *IEEE Transactions on Parallel and Distributed Systems*, 6(7), 1995.
- [15] E. D. Demaine, “A Threads-Only MPI Implementation for the Development of Parallel Programs”, *Proceedings of the 11th International Symposium on High Performance Computing Systems and Applications, HPCS’97*, 1997, pp. 153-163.
- [16] B. V. Dao, Sudhakar Yalamanchili, and Jose Duato, “Architectural Support for Reducing Communication Overhead in Multiprocessor Interconnection Networks”, *Proceedings of the Third International Symposium on High Performance Computer Architecture*, 1997, pp. 343-352.
- [17] J. J. Dongarra and T. Dunigan, “Message-Passing Performance of Various Computers”, *Concurrency: Practice and Experience*, Volume 9, Issue 10, 1997, pp. 915-926.
- [18] P. Druschel and L. L. Peterson, “Fbufs: A High-bandwidth Cross-domain Transfer Facility”, *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, 1993, pp. 189-202.
- [19] C. Dubnicki, A. Bilas, Y. Chen, S. Damianakis and K. Li, “VMMC-2: Efficient Support for Reliable, Connection-Oriented Communication”, *Proceedings of the Hot Interconnect’97*, 1997.
- [20] D. Dunning, G. Regnier, G. McAlpine, D. Cameron, B. Shubert, F. Berry, A. M. Merritt, E. Gronke and C. Dodd, “The Virtual Interface Architecture”, *IEEE Micro*, March-April, 1998, pp. 66-76.
- [21] P. Geoffray, L. Prylli, and B. Tourancheau, “BIP-SMP: High Performance Message Passing Over a Cluster of Commodity SMPs”, *SC99: High Performance Networking and Computing Conference*, November, 1999.
- [22] W. Gropp and E. Lusk, “User’s Guide for MPICH, a Portable Implementation of MPI”, *Argonne National Laboratory, Mathematics and Computer Science Division*, June, 1999.
- [23] R. W. Horst and D. Garcia, “ServerNet SAN I/O Architecture”, *Proceedings of the Hot Interconnects V*, 1997.
- [24] S. Karlson and M. Brorsson, “A Comparative Characterization of Communication Patterns in Applications Using MPI and Shared Memory on an IBM SP2”, *Proceedings of the Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing, International Symposium on High Performance Computer Architecture*, February 1998.
- [25] S. Kaxiras and J. R. Goodman, “Improving CC-NUMA Performance Using Instruction-Based Prediction”, *International Symposium on High Performance Computer Architecture*, 1999.
- [26] J. Kim and D. J. Lilja, “Characterization of Communication Patterns in Message-Passing Parallel Scientific Application Programs”, *Proceedings of the Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing, International Symposium on High Performance Computer Architecture*, February 1998, pp. 202-216.
- [27] D. G. de Lahaut and C. Germain, “Static Communications in Parallel Scientific Programs”, *Proceedings of PARLE’94, Parallel Architecture and Languages*, July 1994.

- [28] A.-C. Lai and B. Falsafi, "Memory Sharing Predictor: The Key to a Speculative Coherent DSM", *Proceedings of the 26th Annual International Symposium on Computer Architectures*, 1999, pp. 172-183.
- [29] M. Lauria, S. Pakin and A. A. Chien, "Efficient Layering for High Speed Communication: Fast Messages 2.x", *Proceedings of the 7th High Performance Distributed Computing, HPDC7, Conference*, 1998.
- [30] S. S. Lumetta, A. M. Mainwaring, and D. E. Culler, "Multi-Protocol Active Messages on a Cluster of SMPs", *SC97: High Performance Networking and Computing Conference*, November, 1997.
- [31] *Message Passing Interface Forum: MPI: A Message-Passing Interface Standard*. Version 1.1 (June 1995).
- [32] *Message Passing Interface Forum: MPI-2: Extensions to the Message-Passing Interface*, (July 1997).
- [33] T. Mowry and A. Gupta, "Tolerating Latency Through Software-Controlled Prefetching in Shared-Memory Multiprocessors", *Journal of Parallel and Distributed Computing*, 12(2), 1991, pp. 87-106.
- [34] S. S. Mukherjee and M. D. Hill, "Using Prediction to Accelerate Coherence Protocols", *Proceedings of the 25th Annual International Symposium on Computer Architecture*, 1998.
- [35] S. Pakin, M. Lauria, and A. Chien, "High Performance Messaging on Workstation: Illinois Fast Messages (FM) for Myrinet," *Proceedings of the Supercomputing '95*, Nov., 1995.
- [36] L. Prylli and B. Tourancheau, "BIP: A New Protocol Designed for High Performance Networking on Myrinet", *Proceedings of the PC-NOW98: International Workshop on Personal Computer based Networks Of Workstations, in conjunction with IPPS/SPDP '98*, 1998.
- [37] S. H. Rodrigues, T. E. Anderson and D. E. Culler, "High-Performance Local Area Communication with Fast Sockets", *USENIX 1997 Annual Technical Conference*, January 1997.
- [38] M. F. Sakr, S. P. Levitan, D. M. Chiarulli, B. G. Horne, and C. L. Giles, "Predicting Multiprocessor Memory Access Patterns with Learning Models", *Proceedings of the Fourteenth International Conference on Machine Learning*, 1997, pp. 305-312.
- [39] G. Shah, J. Nieplocha, J. Mirza and C. Kim, R. Harrison, R. K. Govindaraju, K. Gildea, P. DiNicola, and C. Bender, "Performance and Experience with LAPI -- a New High-Performance Communication Library for the IBM RS/6000 SP", *First Merged Symposium IPPS/SPDP 1998 12th International Parallel Processing symposium & 9th Symposium on Parallel and Distributed Processing*, 1998.
- [40] R. Sheifert, "Gigabit Ethernet", *Addison-Wesley*, 1998.
- [41] T. Takahashi, F. O'Carrol, H. Tezuka, A. Hori, S. Sumimoto, H. Harada, Y. Ishikawa, P.H. Beckman, "Implementation and Evaluation of MPI on an SMP Cluster", *Proceedings of the PC-NOW99: International Workshop on Personal Computer based Networks Of Workstations, in conjunction with PPS/SPDP '99*, 1999.
- [42] Y. Tanaka, M. Matsuda, M. Ando, K. Kubota and M. Sato, "COMPaS: A Pentium Pro PC-based SMP Cluster and its Experience", *Proceedings of the PC-NOW98: International Workshop on Personal Computer based Networks Of Workstations, in conjunction with IPPS/*

SPDP'98, 1998.

- [43] H. Tezuka, F. O'Carroll, A. Hori, and Y. Ishikawa, "Pin-down Cache: A Virtual Memory Management Technique for Zero-copy Communication", *First Merged Symposium IPPS/SPDP 1998 12th International Parallel Processing symposium & 9th Symposium on Parallel and Distributed Processing*, 1998.
- [44] T. V. Eicken, D. E. Culler, S. C. Goldstein, and K. E. Schauser, "Active Messages: A Mechanism for Integrated Communication and Computation", *Proceedings of the 19th Annual International Symposium on Computer Architecture*, May 1992, pp. 256-265.
- [45] T. V. Eicken, A. Basu, V. Buch and W. Vogels, "U-Net: A User-Level Network Interface for Parallel and Distributed Computing", *Proceedings of the 15th ACM Symposium on Operating Systems Principles*, December, 1995.
- [46] P. H. Worley and I. T. Foster, "Parallel Spectral Transform Shallow Water Model: A Runtime-tunable parallel benchmark code", *Proceedings of the Scalable High Performance Computing Conference*, 1994, pp. 207-214.
- [47] T-Y Yeh and Y. Patt, "Alternative Implementation of Two-Level Adaptive Branch Prediction", *Proceedings of the 19th Annual International Symposium on Computer Architecture*, 1992, pp.124-134.
- [48] Z. Zhang and J. Torrellas, "Speeding Up Irregular Applications in Shared-Memory Multiprocessors: Memory Binding and Group Prefetching", *Proceedings of the 22nd Annual International Symposium on Computer Architectures*, 1995, pp. 188-199.