

# Towards Automatic Translation of OpenMP to MPI

Ayon Basumallik  
basumall@purdue.edu

Rudolf Eigenmann  
eigenman@purdue.edu

School of Electrical and Computer Engineering  
Purdue University, West Lafayette, IN 47907-1285  
<http://www.ece.purdue.edu/ParaMount>\*

## Abstract

We present compiler techniques for translating OpenMP shared-memory parallel applications into MPI message-passing programs for execution on distributed memory systems. This translation aims to extend the ease of creating parallel applications with OpenMP to a wider variety of platforms, such as commodity cluster systems. We present key concepts and describe techniques to analyze and efficiently handle both regular and irregular accesses to shared data.

We evaluate the performance achieved by our translation scheme on seven representative OpenMP applications, two from SPEC OMPM2001 and five from the NAS Parallel Benchmarks suite, on two different platforms. The average scalability (execution time relative to the serial version) achieved is within 12% of that achieved by corresponding hand-tuned MPI applications. We also compare our programs with versions deployed for a Software Distributed Shared Memory (SDSM) system and find that the direct translation to MPI achieves up to 30% higher scalability. A comparison with High Performance Fortran (HPF) versions of two NAS benchmarks indicates that our translated OpenMP versions achieve 12% to 89% better performance than the HPF versions.

**Keywords:** Compiler Techniques, OpenMP, MPI, Performance, Commodity Clusters.

## 1 Introduction

OpenMP [33] has established itself as an important method and language extension for programming shared-memory parallel computers. On these platforms, OpenMP offers an easier programming model than the currently widely used message-passing paradigm. Programs written in OpenMP can usually be parallelized stepwise, starting from a sequential program.

---

\*This material is based upon work supported in part by the National Science Foundation under Grants No. 9974976-EIA, 0103582-EIA, and 0429535-CCF. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

OpenMP programs often resemble their original, sequential versions, the main difference being the inserted OpenMP directives.

This approach contrasts with using the Message Passing Interface (MPI) [16], which generally requires programs to be translated into parallel form as a whole, often causing them to look drastically different from their original version. Additionally, MPI places the burden of data-to-process mapping and remote access on the programmer, since the only mechanism for data management is the judicious use of explicit messages. Management of logically shared-data, coupled with the task of parallelizing the application as a whole, can make programming in MPI very effort-intensive.

While OpenMP has clear advantages in terms of ease of programming on shared-memory platforms, message-passing is today still the predominant programming paradigm available for distributed-memory computers, including clusters and most of the highly-parallel systems. In this paper, we explore the suitability of OpenMP for distributed systems as well. Our approach is to automatically translate *standard* OpenMP programs directly to MPI programs. To this end, we have investigated compiler techniques for the translation of OpenMP programs to MPI.

To achieve good performance, our translation scheme includes efficient management of shared data as well as advanced handling of irregular accesses. These issues were also considered by related approaches to compile shared-address-space languages onto distributed memory machines. Compiler techniques for High Performance Fortran (HPF) [20] and those for shared-memory programming on clusters using Software Distributed Shared Memory (DSM) Systems are important representatives. A key difference that distinguishes our approach from these others is the concept of *partial replication*. In HPF-like approaches, the focus was on the availability of data distribution directives. Data partitioning, computation partitioning, and message insertion was derived from these directives. Most often, data had a single owner and computation was performed on the owning node (a.k.a owner-computes rule). Software DSM systems, such as TreadMarks [2], take a contrasting approach. They not only replicate all shared program data but also allocate management data structures (such as shadow copies of shared data) on all processors. In our scheme, shared data is allocated on all nodes. However, no shadow copies or management structures need to be allocated. Furthermore we envision that for future work, arrays with fully-regular accesses will be distributed between nodes by the compiler. Therefore, we refer to our scheme as *partial-replication*. Data replication comes at the cost of limiting the data scalability of programs – that is, programs will not be able to handle  $n$ -times larger data sets on  $n$  processors. While this is a limitation of our scheme,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

JCS '05 June 20-22, Boston, MA, USA.

Copyright © 2005, ACM 1-59593-167-8/06/2005...\$5.00

by *partially replicating* data, our method simplifies data management and communication to the point where our techniques outperform HPF programs substantially. At the same time, our system is able to run the large data sets of our benchmarks, which was not possible on software DSMs due to their extensive replication. We have studied a number of OpenMP programs and identified baseline techniques for such a translation as well as optimizations that can improve performance significantly. This paper makes the following contributions:

- We present compiler techniques for translating OpenMP programs to MPI under the *partial replication* model, including
  - the discussion of OpenMP semantics under the new model.
  - an algorithm for computing message sets,
  - techniques for statically handling irregular accesses, and
  - optimizations based on collective communication.
- We evaluate seven representative OpenMP programs from the NAS and SPEC OMP benchmarks on two platforms. We compare the performance of our new translation scheme with corresponding hand-tuned MPI program versions and with programs optimized for SDSM platforms. We also compare the performance of available HPF program versions and add qualitative comparisons with related paradigms (such as UPC [15]).

The rest of the paper is organized as follows. Section 2 presents the OpenMP to MPI baseline translation and discusses OpenMP semantics and message set computation under the partial replication model. It also presents basic optimizations. Section 3 presents techniques for handling irregular accesses at compile-time. Section 4 presents a set of compiler optimizations that significantly improve the performance of the translated applications. Section 5 presents an evaluation of the performance achieved by a set of representative OpenMP applications translated using these techniques on two different platforms. Section 6 places this paper in the context of related work. Section 7 concludes the paper.

## 2 OpenMP to MPI Baseline Translation

The objective of the baseline translation scheme is to perform a source-to-source translation of a shared-memory OpenMP program to an MPI program. This is accomplished by two categories of transformations – (1) transformations that interpret the OpenMP directives and (2) transformations that identify and satisfy producer-consumer relationships for shared data. We have implemented these compiler transformations using the Cetus compiler infrastructure [26].

### 2.1 Interpretation of OpenMP Semantics Under Partial Replication

OpenMP directives fall into three categories –

- (1) directives that specify work-sharing (*omp parallel for*, *omp parallel sections*) – the compiler interprets these to partition work between processes,
- (2) directives that specify data properties (*omp shared*, *omp private* etc.) – these are used for constructing the set of *shared* variables in the program. By default, data is shared as per the OpenMP standard.
- (3) synchronization directives (*omp barrier*, *omp critical*, *omp flush* etc.) – the compiler incorporates these into the control flow graph.

The execution model under partial replication is SPMD [12] with the following characteristics:

- All participating processes redundantly execute serial regions and parallel regions demarcated by *omp master* and *omp single* directives. Iterations of OpenMP parallel for loops are partitioned between processes using *block-scheduling*.
- Shared data, is allocated on all processes. There is no concept of an *owner* for any shared data item. There are only producers and consumers of shared data.
- At the end of parallel constructs, each participating process communicates the shared data it has produced that other processes *may* use in the future.

The method of producers forwarding data to all *potential* consumers ensures that processes always make local accesses when they read shared data. This method may result in redundant communication where accesses are not completely analyzable. Section 5 will show that our techniques perform well for several representative applications. Furthermore, we present a technique for refining irregular accesses in Section 3.

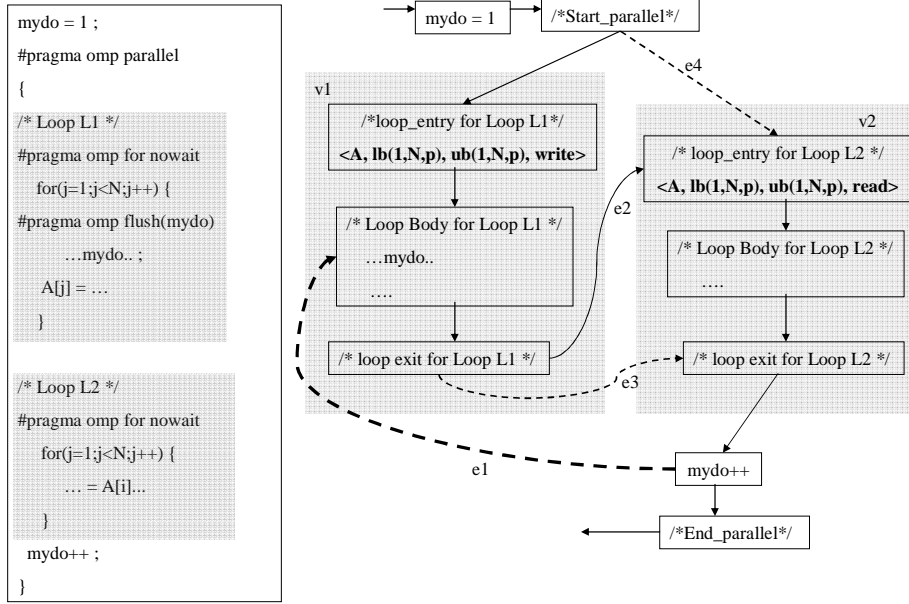
An exception to redundant execution of the serial regions is file I/O. Reading from files is redundantly done by all processes (we assume a file-system visible to all processes). Writing to file is done by only one process (the process with the smallest MPI rank).

Our compiler converts OpenMP *section* constructs to OpenMP loops using a *switch-case* construct so that each iteration executes a different OpenMP section. Henceforth, the term “work-sharing construct” will only refer to OpenMP loops. OpenMP has four scheduling strategies that may be specified for parallel loops – static, dynamic, guided and runtime. Currently, we support only static scheduling. Additionally, at this step, if the OpenMP loop contains a *reduction* clause, then the compiler replaces operations on the reduction variable with operations on a temporary variable and inserts an *MPI\_Allreduce* function call to reduce the value of the temporaries into the reduction variable.

In the first translation step, the compiler thus converts the OpenMP application to SPMD form, interprets the OpenMP directives and performs the requisite work partitioning. After work partitioning, the compiler constructs the set of *shared* data used in the program, using an algorithm described in previous work [31]. Additionally, at this stage, the compiler also does computation repartitioning. Computation repartitioning, discussed previously in the context of deploying OpenMP on Software DSM systems [31], is also important in translating OpenMP applications to MPI. It transforms patterns where the OpenMP application has loops partitioned in such a way that multi-dimensional arrays are accessed along different subscripts in different loops. Computation repartitioning then maintains data affinity by changing the level at which parallelism is specified and partitions work using the parallelism of the inner loops.

### 2.2 Generation of MPI Messages using Array Dataflow

In the second step, the compiler inserts MPI calls to communicate data from producers to *potential* future consumers. To resolve producer-consumer relationships, the compiler has to perform precise array-dataflow analysis. Several schemes such as Linearized Memory Access Descriptors [34] and Regular Section Descriptors [9] have been proposed to characterize array accesses. Our compiler constructs bounded regular section descriptors [19] to characterize accesses to shared arrays. The compiler constructs a control flow graph (with each vertex corresponding to a program statement) and records the Regular Section Descriptors (RSDs) by annotating the vertices of the control flow graph. The starting points of OpenMP work-



**Figure 1. Creation of the Producer-Consumer Flow Graph:** Loop entry vertices  $v1$  and  $v2$  are annotated with Regular Section Descriptors summarizing accesses to shared array  $A$ . Edge  $e1$  is added since the flush directive asserts that writes to  $mydo$  by one thread after loop  $L2$  will be visible to reads of  $mydo$  by other threads in loop  $L1$ . The `nowait` clause in loop  $L1$  implies that writes to  $A$  by one thread in loop  $L1$  are not guaranteed to be visible to other threads in loop  $L2$ . Therefore, edge  $e2$  is deleted and edges  $e3$  and  $e4$  are added to break the path from the producer vertex  $v1$  to the consumer vertex  $v2$ .

sharing constructs are also recorded by annotating the loop entry vertices. For each shared array accessed within a loop, the loop entry vertex is annotated with an RSD summarizing accesses within the loop. For accesses to shared-arrays outside loops, the corresponding program statement vertex is annotated with an RSD describing the access.

The compiler then uses this annotated control flow graph to create a *producer-consumer flow graph* which is used to resolve producer-consumer relationships for shared data. This graph is created by modifying the annotated control-flow graph to conform to the relaxed memory consistency model of OpenMP. OpenMP specifies implicit and explicit memory synchronization points. Writes to shared data by one thread are not guaranteed to be visible to other threads till a memory synchronization point is reached. OpenMP loops have a synchronization point only at the end, thus writes to shared data in one iteration on one thread are not guaranteed to be visible to other iterations on other threads. Thus, in creating the producer-consumer flow graph, the compiler removes the loop back-edge for OpenMP loops to eliminate producer-consumer paths between different iterations of the same loop. Similarly, edges are deleted to account for `nowait` clauses. For explicit synchronization clauses (such as `flush`), producer-consumer edges are added. This is illustrated by the example in Figure 1.

#### Algorithm to Compute Message Sets:

The compiler then uses this producer-consumer flow graph for computing message sets using the inter-procedural algorithm shown in Figure 2. The objective of computing the message sets is to identify the minimal subset of shared data that a producer  $p$  needs to communicate to a potential future remote consumer  $q$ . The vertex  $v_{end}$  marks the end of the program. The graph is traversed starting from annotated Regular Section Descriptor nodes that specify writes to shared-arrays within OpenMP loops. For each shared array written to inside a work-sharing construct, the set of reachable uses is com-

puted. For each reachable use, the corresponding message set is computed by an intersection operation of the regular section descriptors of the current array-write and the reachable use and by subtracting from this the set of intervening productions and consumptions for the array. The computed message sets are communicated using non-blocking *send/receive* and blocking *wait* calls. For affine accesses, the algorithm produces affine expressions for  $UW_{ij}$ . If the accesses are regular but not affine, the compiler inserts code to evaluate the symbolic expressions for  $UW_{ij}$  at runtime. In case of irregular accesses, the compiler may not be able to precisely determine the sets  $PROD_A$  and  $NEED_A$  ( $PROD_A$ ,  $NEED_A$  and  $UW_{ij}$  are defined in Figure 2). The compiler then conservatively treats these to be the entire array  $A$ . Methods for further refining this conservative estimate are discussed in Section 3.

Techniques for generating communication using RSDs have been proposed in the context of HPF compilers [23, 18]. These techniques have presented communication generation in the context of user specified data-distribution and the owner-computes rule for computation partitioning. Our technique differs in two ways. Firstly, instead of data distribution and owner-computes, our scheme is based on partial replication of shared data and user-specified parallelism (with static block scheduling of parallel loops). Therefore, our compiler looks ahead from every production point of shared data to identify potential future consumers and formulates the communication accordingly. Secondly, serial regions are redundantly executed in our scheme and thus our compiler needs to generate communication only for data produced in the parallel regions.

### 3 Translation of Irregular Accesses

A major challenge in translating OpenMP applications directly to message passing is the handling of irregular accesses. Irregular accesses are the reads and writes that cannot be analyzed precisely at compile-time. Existing techniques for dealing with irregular accesses have relied on runtime meth-

### Algorithm *COMPUTE\_MSG\_SETS\_FOR\_PROGRAM*

#### Input :

The producer-consumer flow graph for the program ,  
N - the number of participating processes.

#### Output :

$UW_{ij}$  - The set of data elements that must be communicated  
from process  $i$  to process  $j$  at the end of work sharing construct  $W$   
for each work-sharing construct in the program and  $\forall i, j, 1 \leq i, j \leq N$ .

#### Operations :

We define the following operations -  
REACH-ANY-USE( $v_1, v_2, A$ ) returns the set of RSD nodes, encountered along *any* (interprocedural) path  
between  $v_1$  and  $v_2$  (excluding  $v_1$  and  $v_2$ ), that describe reads to the array  $A$ .  
REACH-ALL-USE( $v_1, v_2, A$ ) returns the set of RSD nodes, encountered along *all* (interprocedural) paths  
between  $v_1$  and  $v_2$  (excluding  $v_1$  and  $v_2$ ), that describe reads to the array  $A$ .  
REACH-ALL-DEF( $v_1, v_2, A$ ) returns the set of RSD nodes, encountered along *all* (interprocedural) paths  
between  $v_1$  and  $v_2$  (excluding  $v_1$  and  $v_2$ ), that describe writes to the array  $A$ .  
SEC-READ( $v, A$ ) returns the RSD for reads to array  $A$  contained in the annotation of node  $v$ .  
SEC-WRITE( $v, A$ ) returns the RSD for writes for array  $A$  contained in the annotation of node  $v$ .  
EVAL( $e, p$ ) returns the accesses of RSD  $e$  made by process  $p$ .

#### Start *COMPUTE\_MSG\_SETS\_FOR\_PROGRAM*

do for each OpenMP work-sharing construct  $W$

Compute  $S_W = \{\text{Set of shared-arrays written within } W\}$

Initialize  $UW_{ij} = \emptyset, 1 \leq i, j \leq N$

do for each array  $A, A \in S_W$ ,

Let  $v_{AW}$  be the entry node for  $W$

Compute  $PROD_A = \text{SEC-WRITE}(v_{AW}, A)$

//Array elements produced (written) in  $W$

Compute  $U_A = \text{REACH-ANY-USE}(v_{AW}, v_{end}, A)$

Set  $WR_A = \emptyset, UA_A = \emptyset$

do for each RSD node  $v_R \in U_A$ ,

Compute  $WR_A = \text{REACH-ALL-DEF}(v_{AW}, v_R, A)$

//Set of intervening producers

Compute  $UA_A = \text{REACH-ALL-USE}(v_{AW}, v_R, A)$

//Set of intervening consumers

Compute  $NEED_A = S1_A - (S2_A \cup S3_A)$

where  $S1_A = \text{SEC-READ}(v_R)$ ,

//Array elements consumed in future read

$S2_A = \bigcup_{v \in WR_A} \text{SEC-WRITE}(v, A)$

//Array elements produced by intervening writes

$S3_A = \bigcup_{v \in UA_A} \text{SEC-READ}(v, A)$

//Array elements consumed by intervening reads

Compute  $UW_{ij} = UW_{ij} \cup \{ \text{EVAL}(PROD_A, i) \cap \text{EVAL}(NEED_A, j) \}, 1 \leq i, j \leq N, i \neq j$

enddo

enddo

enddo

End *COMPUTE\_MSG\_SETS\_FOR\_PROGRAM*

Figure 2. Algorithm to Compute Message Sets for Resolving Remote Accesses

ods [13, 37]. In this section, we present a technique for handling irregular accesses to shared arrays at compile-time. Section 5 will show that this technique performs well in practice.

To handle such irregular accesses, we make use of the property of *monotonicity*. Monotonicity states  $\forall i, \forall j, i \leq j \Leftrightarrow A[i] \leq A[j]$ . Consider an irregular access of the form  $A[B[i]]$  where  $A$  is a one dimensional shared array,  $B$  is an indirection array and  $i$  is the loop index. Process  $p$  and  $q$  execute iterations  $[l_p, u_p]$  and  $[l_q, u_q]$  of the loop respectively. If the array  $B$  can be proved to be monotonic, then the compiler asserts that  $p$  and  $q$  access the intervals  $[B[l_p], B[u_p]]$  and  $[B[l_q], B[u_q]]$  respectively.

In the baseline case for irregular reads, the compiler conservatively overestimates irregular accesses to an array as accesses to the entire array and processes inter-communicate all their preceding writes to the array. This can be done efficiently in our translation scheme since the entire shared data space is allocated on all processes and thus, place-holders for shared data already exist on potential consumers. This is different from common HPF implementations, where temporary storage has to be allocated for non-owned shared data, potentially introducing overheads. However, conservatively forwarding all writes may result in redundant communication. So, as a compile-time optimization, the compiler checks if the future irregular read has a monotonicity property and accordingly refines the accesses.

For irregular writes, it is not possible to exactly determine which process writes which element of an array. In such a case it is generally incorrect to treat them as writes to the entire

array. The analysis of irregular writes may be refined using the monotonicity property. For the applications that we have encountered, all instances of irregular writes have been analyzable using the monotonicity property.

However, there may be applications where this compile-time scheme is not applicable, either because the indirection function is not provably monotonic or because the irregular write is not in the form of *ARRAY[indirection function]*. In such cases, we use a runtime mechanism as a fallback. In this mechanism, the compiler inserts code to record the writes at runtime. The inserted code allocates a buffer on each process and puts in, at runtime, (location,value) pairs for the elements that are written. At the end of the loop that contains the irregular write, processes intercommunicate this buffer.

Testing for monotonicity in irregular accesses has been discussed in the context of parallelizing compilers [7, 29] for analyzing properties of index arrays. We have manually applied these techniques to our benchmarks.

## 4 Optimizations Using Collective Communications

In the previous two sections, we have described techniques for the baseline translation of OpenMP applications to MPI and for handling irregular accesses, respectively. In this section, we describe two additional compile-time optimizations which have been found to improve performance considerably. These are motivated by the observation that most MPI implementations optimize collective communication routines such

as MPI\_Reduce and MPI\_Alltoall. While a detailed evaluation of the impact of each optimization is beyond the scope of this paper, Section 5 indicates the applicability of these optimization techniques for our benchmarks.

#### 4.1 Recognition of Transformed Reduction Idioms

The objective of this optimization is to recognize that the programmer is using a series of OpenMP constructs to perform a reduction operation and to substitute it with a MPI reduction operation. Recognition of reduction idioms have been discussed in the context of parallelizing compilers [35]. We apply this technique in a way that is the converse of the transformations done by parallelizing compilers. OpenMP does have a reduction clause which can be specified for OpenMP loops. However, this is most commonly used for scalar reductions. Array reduction and other complex forms of reduction operations are often implemented using an OpenMP *for* loop, followed by an OpenMP *critical* section that updates global memory. Instead of a critical section, a parallel loop may also be used where threads update disjoint parts of the shared array in each iteration. In the baseline translation from OpenMP to MPI, an OpenMP critical section is implemented as a section which processes enter in order. Each process leaving the critical section communicates its updates of shared data to the next process entering the critical section. The last process leaving the critical section communicates its updates to all other processes. However, this baseline translation scheme has a large performance overhead, which grows linearly with the number of processes involved. This overhead can be reduced considerably if the compiler recognizes that this pattern represents a transformed reduction.

The basis of recognizing such patterns is to first identify critical sections (or parallel loops) where shared arrays are being updated using commutative operations with private arrays, then verify that these private arrays are produced in an OpenMP loop immediately preceding the critical section. In such a case, we translate the array reduction directly using MPI\_Allreduce functions.

#### 4.2 Optimization of Sends and Receives for Shared Data

The global dataflow analysis, described in Section 2, identifies the production and consumption points of the message sets. To overlap computation with communication, messages are initiated early using *non-blocking sends/receives* and completed just before the consumption point at the receiver with a *wait* operation. A similar scheme has been used by some HPF compilers [18]. This scheme has limited benefit when the distance between the production and consumption points is insufficient. In such a case, if the producer-consumer relationships are of the many-to-many type (as computed by the algorithm in Figure 2), our technique re-casts the sends and receives to collective communications using the MPI\_Alltoallv function.

### 5 Performance Evaluation

We have implemented the translation steps, described in Section 2, in the Cetus compiler infrastructure [26]. We then manually applied the techniques described in Sections 3 and 4 to seven representative OpenMP applications. As our benchmarks, we have selected five out of the eight benchmarks from the NAS Parallel Benchmarks suite [4, 22], namely CG, IS, EP, FT and LU; and two benchmarks from the SPEC OMPM2001 suite [3], namely ART and EQUAKE. Our compiler infrastructure handles C programs. Thus for SPEC OMPM2001 we have used two out of the three available C benchmarks and for the NAS benchmarks, we have used the NPB-2.3 OpenMP C versions created by the Real World Computing Partnership

(<http://phase.hpcc.jp/Omni/benchmarks/NPB/>). In this section, we evaluate the performance of these translated OpenMP versions (henceforth referred to simply as the OpenMP versions). For our experiments, we have selected two different platforms – (1) a cluster of sixteen PIII 800 MHz Linux nodes, with 256 MB memory per node, connected by a commodity 100 Mbps Ethernet network and (2) sixteen IBM SP2 WinterHawk-II nodes connected by a high-performance switch. We expect the scaling behavior of the benchmarks on these systems to be representative of a wide class of high-performance computing platforms. The MPI libraries used for these platforms is MPICH version 1.2.5 on Linux and IBM MPI libraries on the IBM SP2. The back-end compilers are gcc-3.3 on Linux and xlc version 6 on IBM, both at optimization level O3.

The main objective of this performance evaluation is to compare the performance of the OpenMP versions with the performance of their hand-coded MPI counterparts. The hand-coded MPI versions for CG, IS, EP, FT and LU are part of the NAS Benchmarks NPB2.4 suite. They have been used to evaluate MPI implementations on several platforms and represent an expected upper bound on performance. The SPEC OMPM2001 benchmarks EQUAKE and ART do not have any such existing MPI versions. We created the MPI versions, using reasonable programming efforts.

Figures 3,4,5,6,7,8 and 9 show the performance of the OpenMP and the MPI versions on the Linux cluster and the IBM SP2. The performance is shown both in terms of the execution time as well as scalability. We define scalability of a version  $x$  as  $Scalability = T_{serial}/T_x$  where  $T_x$  is the execution time of version  $x$  and  $T_{serial}$  is the execution time of the corresponding serial version. The serial versions used are NPB-3.0-SER for the NAS benchmarks and OpenMP versions with the OpenMP directives removed for ART and EQUAKE. Thus we have a common baseline for both OpenMP and MPI versions in measuring scalability. On average, the scalability of our translated OpenMP versions is 12% less and the execution times are 14% more than that of their hand-coded MPI counterparts. Next, we individually examine the performance of each benchmark.

In the case of EP, the MPI version uses several large auxiliary arrays, which are not present in the serial version [22]. Since the OpenMP version is derived from this serial version, its execution time is about 10% less than the hand-tuned MPI version. The OpenMP to MPI translation of EP benefited from the recognition of an array reduction.

The hand-tuned MPI version of CG benefits from the programmer's knowledge of exact reduction and data exchange patterns in the algorithm. In our translation of the OpenMP version, the message set computation in Figure 2 recognizes the access affinity in the regular accesses for several arrays in the *conj\_grad* subroutine. However, there are irregular reads to two arrays. Analysis of the indirection array provides loose bounds for these irregular reads. For these arrays, each process has to broadcast its writes to all other processes. In spite of this difference, our version achieves performance close to that achieved by the hand-coded MPI version.

For FT, the hand-tuned MPI version employs a sophisticated strategy to map the problem into a virtual network topology depending on the problem set size and the number of processes. Furthermore, the resulting transpose communication patterns are also hand-optimized. Owing to these differences, the translated OpenMP version is about 15% slower than its hand-coded MPI counterpart. The dataset used for FT on PIII/Ethernet is CLASS A because the CLASS B problem size was too large (also for the serial and the MPI versions) for nodes with 256 MB main memory.





Figure 3. Performance of EP (CLASS B).





Figure 4. Performance of CG (CLASS B).

IS is a NAS benchmark that performs integer sorting. The hand-tuned MPI version of IS uses a *bucket-sort* algorithm. The OpenMP version does not use *bucket-sort*, in effect performing the ranking using only one *bucket*. The version with buckets has been found to have better cache affinity and the MPI version thus has significantly lower execution times than the OpenMP version (and the serial version without buckets). However, the OpenMP version still achieves scalability to within 14% of the MPI version on both platforms. The translation of IS benefited from the recognition of reduction patterns discussed in Section 4.

LU benefited from the computation repartitioning optimization mentioned in Section 2. The optimized nearest neighbor communication strategy in the hand-tuned MPI version still enables it to achieve better performance and scalability. In LU, we also observed a significantly higher execution time for the OpenMP version on one processor compared to the MPI and serial versions. We attribute a part of this to the difference in code produced by C and Fortran compilers (the OpenMP version is in C while the MPI and serial versions are in Fortran).

For EQUAKE and ART, the translated versions do not differ significantly from the hand-crafted versions and thus achieve similar performance.

## 5.1 Comparison with OpenMP deployed using Software Distributed Shared Memory

As discussed in Section 1, recent schemes proposed for the deployment of OpenMP on clusters utilize an underlying layer of Software Distributed Shared Memory (SDSM) [21, 32]. In

this section, we compare the performance of the OpenMP versions (translated to MPI using our proposed techniques) with the performance of the OpenMP versions deployed on the TreadMarks SDSM system. For this set of experiments, we have used the sixteen node Linux cluster. Figure 10 shows the performance of the two schemes.

Apart from EQUAKE, the versions translated to MPI always perform better compared to the corresponding SDSM versions. EQUAKE has two large arrays that are accessed along different subscripts in the original partitioning scheme. Furthermore, it contains a sparse matrix vector product where it expands the shared data size by allocating work-arrays whose size is proportional to the number of threads. An MPI implementation would need to heavily restructure the existing application to achieve performance. In the SDSM version of EQUAKE, an additional optimization was that iterations for the main loop in the *smvp* (sparse-matrix-vector-product) subroutine were dynamically scheduled, which serialized most accesses and eliminated communication. Similar performance would be achieved if this loop were serialized for the MPI version as well. For the other benchmarks, the MPI versions mainly profit from the implicit aggregation of messages, which has been cited by others [30] as a principal advantage of message-passing systems.

An important shortcoming of the SDSM versions not evident from the performance comparison, is the constraint on the size of shared data. Each TreadMarks SDSM process creates a *twin* for every page of shared data that it writes to and tracks writes to shared data at runtime by comparing current pages with their original *twins*. Furthermore, memory also needs to





Figure 5. Performance of FT (CLASS A on Ethernet Cluster and CLASS B on IBM SP2 ).

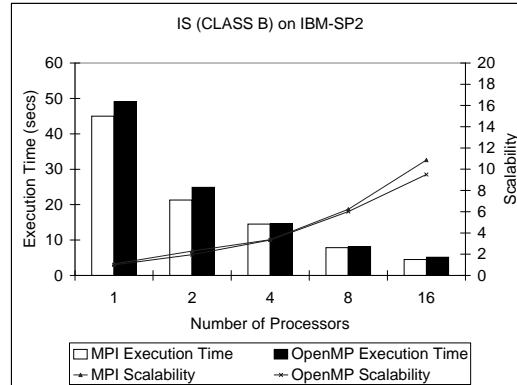



Figure 6. Performance of IS (CLASS B).

be allocated for storing the list of writers to each page of shared data in a program interval. As a result, the maximum size of shared data that can be allocated is constrained to be a fraction of the total available address space. Due to this, one cannot use large datasets such as SPEC's *ref* datasets (which have sizes close to 2 GB) of ART and EQUAKE for the SDSM versions of these benchmarks. By contrast, our translation to MPI does not introduce such size constraints, despite the replication of shared data on all nodes, and we can run the largest problem sizes allowed by the architecture.

## 6 Related Work

An important contribution towards shared-memory programming support for distributed-memory machines was the development of High Performance Fortran (HPF) [20, 24]. There are important differences between our approach and that taken by HPF. Even though, like OpenMP, HPF provided directives to specify parallel loops, HPF's focus was on the use of the data distribution directives. Data and computation partitioning was derived from these directives. Most often, data had a single owner and computation was performed on the owning node (a.k.a owner-computes rule). Input operands to the computations were received via messages from their owners. This execution scheme could also add significant overhead to serial sections, as these needed to be executed on multiple nodes owning parts of the data. Handling irregular data was difficult and usually employed runtime schemes [13].

In contrast to HPF implementations, our execution model starts from the available parallelism specified through OpenMP directives. Partial replication allows serial regions to be executed intact and input operands of parallel computations are

locally available. Communication happens primarily at the end of parallel loops, facilitated by collective communication. Partial replication also obviates the need for data partitioning techniques [25], even though data distribution information is not provided by the user. In all our benchmarks, irregular access patterns are successfully handled at compile-time.

Figures 11 and 12 compares the performance of our programs with those available in HPF form, compiled with the *pghpf* compiler<sup>1</sup>. Of these programs, we have chosen two extremes: in CG, HPF performs well, while in LU, HPF performs poorly. In both cases, our translation scheme comes close to the MPI performance. In CG our techniques perform 10% better than the HPF versions whereas in LU, our techniques outperform the HPF version 10-fold. The CLASS A dataset has been used for LU, instead of the CLASS B dataset used in comparing OpenMP and MPI, because the available 256 MB memory per node is insufficient to run the HPF version. (We observed the main source of slowdown in the HPF version of LU to be the subroutines *jacld*, *blts*, *buts* and *jacu* which have execution times that are more than ten times slower than the serial versions.) These trends are consistent with those reported by the authors of the HPF versions for the NAS benchmarks [17].

Split-C [11] and Unified Parallel C (UPC) [15] are other programming paradigms that have been proposed to ease programming effort by providing a global address space that is logically partitioned between threads. The user specifies affinity

<sup>1</sup>While other HPF compilers may perform better, we believe the *pghpf* compiler [8] (the only one we could still obtain) to be representative





Figure 7. Performance of LU (CLASS B).





Figure 8. Performance of EQUAKE (REF Dataset).

between threads and data. In a comparative study for the NAS benchmarks [14], the authors have compared hand optimized UPC versions (hand optimized to privatize local accesses) of the NAS benchmarks with MPI versions. Their results indicate that the UPC versions achieve 50% to 75% the performance of the MPI versions. They find that an important reason for the better scalability of the MPI versions is the optimized implementation of collective operations provided by MPI vendors. Our translation scheme and our optimizations based on collective communication leverages the availability of efficient MPI libraries on multiple HPC platforms.

Recently, several research efforts have suggested the use of OpenMP on clusters. Some have proposed language extensions and data distribution directives [36, 1, 10] similar to HPF. Others have proposed architectural support and page placement techniques to map data to the most suitable processing nodes [6]. A common scheme suggested for deploying standard OpenMP (without language extensions) on clusters has been to use an underlying Software Distributed Shared Memory (SDSM) system [21, 31]. We have already evaluated the advantage of our approach over deployment using SDSM in Section 5.1.

Some parallelizing compilers have targeted distributed-memory systems and have attempted to directly generate parallel message-passing versions of serial applications [38, 5]. By contrast, our starting point is code parallelized by the programmer. Some research efforts have presented a general framework for matching syntactic reference patterns with appropriate aggregate communication routines [27, 28]. To the best of our knowledge, ours is the only effort that comprehensively

addresses the direct, source-level translation of shared-memory programs in standard OpenMP to a portable message-passing form in MPI.

## 7 Conclusion

We have presented compiler techniques for translating standard OpenMP shared-memory programs into MPI message passing variants, based on a novel model of *partial replication*. The contributions include a new algorithm to compute message sets, techniques for statically handling irregular accesses, and optimizations based on collective communication. Partial replication contrasts with data distribution models, as used in HPF-like languages, in which a single node owns a (partition of a) shared array. Partial replication simplifies the generation of messages, which holds for irregular accesses. Exploiting monotonicity properties of index arrays, our techniques were able to handle all message generation for irregular accesses in our program suite statically. Our methods, although less data scalable than data distribution schemes, can handle much larger data sets than the extensive data replication schemes employed in SDSM systems such as TreadMarks. Our experiments used the full *ref* data sets (up to 2 GB) of the SPEC OMPM2001 benchmarks.

Our measurements show that the presented techniques significantly outperform both HPF programs (using the pghpf compiler) and SDSM programs (using TreadMarks). Perhaps more importantly, they come close to or match hand-tuned MPI programs in all of our measured cases. This fact, combined with the greater ease of programming that OpenMP is generally attributed with, indicates a promising new path toward higher programming productivity on clusters and highly paral-






Figure 9. Performance of ART (REF Dataset).




Figure 10. Performance Comparison of the OpenMP versions translated to MPI with the OpenMP versions deployed on the TreadMarks DSM system. Datasets are *train* for ART and EQUAKE; CLASS A for LU, FT and CLASS B for CG, IS, EP.

lel computer platforms.

## 8 References

- [1] T. Abdelrahman and T. Wong. Compiler support for data distribution on NUMA multiprocessors. *Journal of Supercomputing*, 12(4):349–371, October 1998.
- [2] 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 Workstations. *IEEE Computer*, 29(2):18–28, 1996.
- [3] V. Aslot, M. Domeika, R. Eigenmann, G. Gaertner, W. B. Jones, and B. Parady. SPEComp: A New Benchmark Suite for Measuring Parallel Computer Performance. In *Proc. of the Workshop on OpenMP Applications and Tools (WOMPAT2001)*, *Lecture Notes in Computer Science*, 2104, pages 1–10, July 2001.
- [4] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, D. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrishnan, and S. K. Weeratunga. The NAS Parallel Benchmarks. *The International Journal of Supercomputer Applications*, 5(3):63–73, Fall 1991.
- [5] P. Banerjee, J. A. Chandy, M. Gupta, J. G. Holm, A. Lain, D. J. Palermo, S. Ramaswamy, and E. Su. The PARADIGM Compiler for Distributed-Memory Message Passing Multicomputers. In *The First International Workshop on Parallel Processing*, pages 322–330, Bangalore, India, Dec. 1994.
- [6] J. Bircsak, P. Craig, R. Crowell, Z. Cvetanovic, J. Harris, C. Nelson, and C. Offner. Extending OpenMP for NUMA Machines. In *Proc. of the IEEE/ACM Supercomputing'2000: High Performance Networking and Computing Conference (SC2000)*, November 2000.
- [7] W. Blume and R. Eigenmann. The range test: a dependence test for symbolic, non-linear expressions. In *Proceedings of the 1994 ACM/IEEE conference on Supercomputing*, pages 528–537. ACM Press, 1994.
- [8] Z. Bozkus, L. Meadows, S. Nakamoto, V. Schuster, and M. Young. Pghpf—an optimizing high performance fortran compiler for distributed memory machines. *Sci. Program.*, 6(1):29–40, 1997.
- [9] D. Callahan and K. Kennedy. Analysis of interprocedural side effects in a parallel programming environment. *J. Parallel Distrib. Comput.*, 5(5):517–550, 1988.
- [10] B. Chapman, P. Mehrotra, and H. Zima. Enhancing OpenMP with features for Locality Control. Technical Report TR99-02, Inst. for Software Technology and Parallel Systems, U. Vienna, www.par.univie.ac.at., 1999.
- [11] D. E. Culler, A. C. Arpacı-Dusseau, S. C. Goldstein, A. Krishnamurthy, S. Lumetta, T. von Eicken, and K. A. Yelick. Parallel Programming in Split-C. In *Supercomputing*, pages 262–273, 1993.
- [12] F. Darema, D. A. George, V. A. Norton, and G. F. Pfister. A single-program-multiple-data computational model for




Figure 11. Performance of CG (CLASS A).




Figure 12. Performance of LU (CLASS A).

- epex/fortran. *Parallel Computing*, 7(1):11–24, 1988.
- [13] R. Das, M. Uysal, J. Saltz, and Y.-S. S. Hwang. Communication optimizations for irregular scientific computations on distributed memory architectures. *Journal of Parallel and Distributed Computing*, 22(3):462–478, 1994.
  - [14] T. El-Ghazawi and F. Cantonnet. "upc performance and potential: a npb experimental study". In *Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, pages 1–26. IEEE Computer Society Press, 2002.
  - [15] T. El-Ghazawi, W. Carlson, and J. Draper. UPC Language Specifications V1.0, Feb. 2001.
  - [16] M. P. I. Forum. MPI: A Message-Passing Interface Standard. Technical Report UT-CS-94-230, 1994.
  - [17] M. Frumkin, H. Jin, and J. Yan. Implementation of NAS Parallel Benchmarks in High Performance Fortran. Technical Report NAS-98-009.
  - [18] M. Gupta, E. Schonberg, and H. Srinivasan. A unified framework for optimizing communication in data-parallel programs. *IEEE Transactions on Parallel and Distributed Systems*, 7(7):689–704, 1996.
  - [19] P. Havlak and K. Kennedy. An implementation of interprocedural bounded regular section analysis. *IEEE Transactions on Parallel and Distributed Systems*, 2(3):350–360, 1991.
  - [20] High Performance Fortran Forum. High Performance Fortran language specification, version 1.0. Technical Report CRPC-TR92225, Houston, Tex., 1993.
  - [21] Y. Hu, H. Lu, A. Cox, and W. Zwaenepoel. OpenMP for Networks of SMPs. *Journal of Parallel and Distributed Computing*, 60(12):1512–1530, December 2000.
  - [22] H. Jin, M. Frumkin, and J. Yan. The OpenMP Implementation of NAS Parallel Benchmarks and Its Performance. Technical Report NAS-99-011.
  - [23] K. Kennedy and N. Nedeljković. Combining dependence and data-flow analyses to optimize communication. In *Proceedings of the 9th International Parallel Processing Symposium*, Santa Barbara, CA, 1995.
  - [24] C. Koelbel, D. Loveman, R. Schreiber, G. S. Jr., and M. Zosel. *The High Performance Fortran Handbook*. MIT Press, 1994.
  - [25] U. Kremer. Automatic data layout for distributed memory machines. Technical Report TR96-261, 14, 1996.
  - [26] S.-I. Lee, T. A. Johnson, and R. Eigenmann. Cetus - An Extensible Compiler Infrastructure for Source-to-Source Transformation. In *Proc. of the Workshop on Languages and Compilers for Parallel Computing (LCPC'03)*, pages 539–553. (Springer-Verlag Lecture Notes in Computer Science), Oct. 2003.
  - [27] J. Li and M. Chen. Generating explicit communication from shared-memory program references. In *Proceedings of the 1990 ACM/IEEE conference on Supercomputing*, pages 865–876. IEEE Computer Society, 1990.
  - [28] J. Li and M. Chen. Compiling communication-efficient programs for massively parallel machines. *IEEE Trans. Parallel Distrib. Syst.*, 2(3):361–376, 1991.
  - [29] Y. Lin and D. Padua. Compiler analysis of irregular memory accesses. In *Proceedings of the ACM SIGPLAN 2000 conference on Programming language design and implementation*, pages 157–168. ACM Press, 2000.
  - [30] H. Lu, S. Dwarkadas, A. L. Cox, and W. Zwaenepoel. Quantifying the performance differences between PVM and TreadMarks. *Journal of Parallel and Distributed Computing*, 43(2):65–78, 1997.
  - [31] S.-J. Min, A. Basumallik, and R. Eigenmann. Optimizing OpenMP programs on Software Distributed Shared Memory Systems. *International Journal of Parallel Programming*, 31(3):225–249, June 2003.
  - [32] S.-J. Min, A. Basumallik, and R. Eigenmann. Supporting realistic OpenMP applications on a commodity cluster of workstations. In *OpenMP Shared Memory Parallel Programming: International Workshop on OpenMP Applications and Tools, WOMPAT 2003, Toronto, Canada, June 26-27, 2003. Proceedings Editors: M.J. Voss (Ed.)*, pages 170 – 179, 2003.
  - [33] OpenMP Forum. OpenMP: A Proposed Industry Standard API for Shared Memory Programming. Technical report, October 1997.
  - [34] Y. Paek, J. Hoeflinger, and D. Padua. Efficient and precise array access analysis. *ACM Trans. Program. Lang. Syst.*, 24(1):65–109, 2002.
  - [35] W. M. Pottenger and R. Eigenmann. Idiom recognition in the Polaris Parallelizing Compiler. In *International Conference on Supercomputing*, pages 444–448, 1995.
  - [36] V. Schuster and D. Miles. Distributed OpenMP, Extensions to OpenMP for SMP Clusters. In *Proc. of the Workshop on OpenMP Applications and Tools (WOMPAT2000)*, July 2000.
  - [37] R. von Hanxleden, K. Kennedy, C. H. Koelbel, R. Das, and J. H. Saltz. Compiler analysis for irregular problems in Fortran D. In *1992 Workshop on Languages and Compilers for Parallel Computing*, number 757, pages 97–111, New Haven, Conn., 1992. Berlin: Springer Verlag.
  - [38] J. Zhu and J. Hoeflinger. Compiling for a Hybrid Programming Model Using the LMAD Representation. In *Proc. of the 14th annual workshop on Languages and Compilers for Parallel Computing (LCPC2001)*, August 2001.