Unified Compilation of Fortran 77D and 90D

ALOK CHOUDHARY*, GEOFFREY FOX*, SEEMA HIRANANDANI†, KEN KENNEDY†,
CHARLES KOELBEL†, SANJAY RANKA* and CHAU-WEN TSENG†

*Northeast Parallel Architectures Center, 111 College Place, Syracuse University, Syracuse, NY 13244-4100

†Center for Research on Parallel Computation, Rice University, P.O. Box 1892, Houston, TX 77251-1892

We present a unified approach to compiling Fortran 77D and Fortran 90D programs for efficient execution on MIMD distributed-memory machines. The integrated Fortran D compiler relies on two key observations. First, array constructs may be *scalarized* into FORALL loops without loss of information. Second, *loop fusion*, *partitioning*, and *sectioning* optimizations are essential for both Fortran D dialects.

Categories and Subject Descriptors: C.1.2 [**Processor Architectures**]: Multiple Data Stream Architectures (Multiprocessors) – *multiple-instruction-stream, multiple-data-stream processors (MIMD), parallel processors*; D.3.2 [**Programming Languages**]: Language Constructs – *concurrent programming structures*; D.3.4 [**Programming Languages**]: Processors – *code generation, compilers, optimization, preprocessors*

General Terms: Languages, Performance

Additional Keywords and Phrases: parallel languages; Fortran D; parallel programming

---

1. INTRODUCTION

Parallel computing on distributed-memory machines is scalable and cost-effective; however, it is hindered by both the difficulty of parallel programming and lack of portability of the resulting programs. We propose to solve this problem using Fortran D, a version of Fortran extended with data decomposition specifications that can be applied to Fortran 77 and Fortran 90 [6] to produce Fortran 77D and Fortran 90D, respectively. Fortran D has contributed to the development of High Performance Fortran (HPF), an informal Fortran standard for programming massively-parallel machines [22].

We are developing the compiler technology needed to automate translation of Fortran D to different parallel architectures. Our goal is to establish a machine-independent programming model for data-parallel programs that is easy to use, yet performs with acceptable efficiency on different parallel architectures. In this paper, we describe a unified strategy for compiling both Fortran 77D and Fortran 90D into efficient *single-program, multiple-data* (SPMD) message-passing programs. In particular, we concentrate on the design of a prototype Fortran 90D compiler for the Intel iPSC/860 and Thinking Machines CM-5, two

This research was supported by the Center for Research on Parallel Computation (CRPC), a National Science Foundation Science and Technology Center. CRPC is funded by NSF through Cooperative Agreement Number CCR-9120008. Additional support was provided by DARPA under contract # DABT63-91-C-0028. Alok Choudhary is also supported by an NSF Young Investigator Award CCR-9357840.

multiple-instruction, multiple-data (MIMD) distributed-memory machines.

The principal issues involved in compiling Fortran 90D are *partitioning* the program across multiple nodes and *scalarizing* it for execution on each individual node. Previous work has described both the scalarization [5] and partitioning process [20, 24, 25]. The contributions of this paper are to:

- utilize the FORALL loop to preserve the semantics of Fortran 90 array constructs in a common intermediate form,
- provide a compilation framework for integrating program partitioning with Fortran 90 scalarization,
- show that an efficient portable run-time library can ease the task of compiling Fortran D,
- validate the effectiveness of a unified compilation approach through empirical case studies.

Because our main interest is in the interaction between scalarization and partitioning optimizations, in this paper we only focus on a subset of Fortran 90, namely array operations, array constructs, and the FORALL loop.

The remainder of this paper presents a brief overview of the Fortran D language and compilation strategy, then describes the Fortran 90D and 77D front ends and the common Fortran D back end. The design of the run-time library is discussed, and an example is used to illustrate the compilation process. We conclude with a discussion of related work.

2. FORTRAN D LANGUAGE

We briefly overview aspects of Fortran D relevant to this paper. These extensions can be added to either Fortran 77 or Fortran 90. The complete language is described elsewhere [19].

2.1 DATA ALIGNMENT AND DISTRIBUTION

In Fortran D, the DECOMPOSITION statement declares an abstract problem or index domain. The ALIGN statement maps each array element onto one or more elements of the decomposition. This provides the minimal requirement for reducing data movement for the program given an unlimited number of processors. The DISTRIBUTE statement groups decomposition elements, mapping them and any array elements aligned with them to the finite resources of the physical machine. Each dimension of the decomposition is distributed in a block, cyclic, or block-cyclic manner; the symbol ":" marks dimensions that are not distributed. Because the alignment and distribution statements are executable, dynamic data decomposition is possible.

2.2 FORALL

Fortran D provides FORALL loops to permit the user to specify difficult parallel loops in a deterministic manner. Its semantics are borrowed from the Myrias PARDO and previous FORALL constructs [3, 10, 33]. In a FORALL loop, each iteration *uses only values defined before the loop or within the current iteration.* When a statement in an iteration of the FORALL loop accesses a memory location, it will not get any value written by a different iteration of the loop. Instead, it will get the *old* value at that memory location (*i.e.,* the value at that location before the execution of the FORALL loop) or it will get some new value written on the current iteration. Similarly, a merging semantics ensures that a deterministic value is obtained after the FORALL if several iterations assign to the same memory location.

Another way of viewing the FORALL loop is that it has copy-in/copy-out semantics. In other words, each iteration gets its own copy of the entire data space that exists before the execution of the loop, and writes its results to a new data space at the end of the loop. Since no values depend on other iterations, the FORALL loop may be executed in parallel without synchronization. However, communication may still be required before the loop to acquire non-local values, and after the loop to update or merge non-local values. Single-statement Fortran D FORALL loops are identical to those supported in CM FORTRAN [2] and HPF [22], but multi-statement Fortran D FORALL loops are different from those found in HPF.

3. FORTRAN D COMPILATION STRATEGY

3.1 OVERALL STRATEGY

Our strategy for parallelizing Fortran D programs for distributed-memory MIMD computers is illustrated in Figure 1. In brief, we transform both Fortran 77D and Fortran 90D to a common intermediate form, which is then compiled to code for the individual nodes of the machine. We have several pragmatic and philosophical reasons for this strategy:

- Sharing a common back end for both the Fortran 77D and Fortran 90D avoids duplication of effort.
- Decoupling the Fortran 77D and Fortran 90D front ends allows them to become machine independent.
- Providing a common intermediate form helps us experiment with defining an efficient compiler/programmer interface for programming the nodes of a massively parallel machine.

3.2 INTERMEDIATE FORM

To compile both dialects of Fortran D using a single back end, we must select an appropriate intermediate form. In addition to standard computation and control flow information, the intermediate form must capture three important aspects of the program:

- Data decomposition information, telling how data is aligned and distributed among processors.
- Parallelization information, telling when operations in the code are independent.
- Communication information, telling what data must be transferred between processors.

In addition, we believe that the primitive operations of the intermediate form should be relatively low-level operations that can be translated simply for single-processor execution.

We have chosen Fortran 77 with data decompositions, FORALL, and intrinsic functions to be the intermediate form for the Fortran D compiler. We show later that this form preserves all of the information available in a Fortran 90 program, but maintains the flexibility of Fortran 77. Parallelism and communication can be determined by the compiler for simple computations, and specified by the user using FORALL and intrinsic functions for complex computations.

3.3 NODE INTERFACE

Another topic of interest in the overall strategy is the node interface—the node program produced by the Fortran D compiler. It must be both portable and efficient. In addition, the level of the node interface should be neither so high that efficient translation to object code is impossible, nor so low that its workings are completely opaque to the user. We have selected Fortran 77 with calls to communication and run-time libraries based on Express, a collection of portable message-passing primitives [36]. Evaluating our experiences with this node interface is the first step towards defining an "optimal" level of support for programming individual nodes of a parallel machine.

4. FORTRAN D COMPILER

The Fortran D compiler thus consists of three parts. The Fortran 90D and 77D front ends process input programs into the common intermediate form. The Fortran D back end then compiles this to the SPMD message-passing node program. The Fortran D compiler is implemented in the context of the ParaScope programming environment [18].

4.1 FORTRAN 90D FRONT END

The function of the Fortran 90D front end is to *scalarize* the Fortran 90D program, translating it to an equivalent Fortran 77D program. This is necessary because the underlying machine executes computations sequentially, rather than on entire arrays at once as specified in Fortran 90. For the Fortran D compiler we find it useful to view scalarization as three separate tasks:

- **Scalarizing Fortran 90 Constructs.** Many Fortran 90 features are not present in our intermediate form. They must be translated into equivalent Fortran 77D statements.

- **Fusing Loops.** Simple scalarization results in many small loop nests. Fusing these loop nests can improve the locality of data accesses, simplify partitioning, and enable other program transformations.

- **Sectioning.** Fortran 90 array operations allow the programmer to access and modify entire arrays atomically, even if the underlying machine lacks this capability. The Fortran D compiler must divide array operations into *sections* that fit the hardware of the target machine [5].

We defer both loop fusion and sectioning to the Fortran D back end. Loop fusion is deferred because even hand-written Fortran 77 programs can benefit significantly [15, 28, 34]. Sectioning is needed in the back end because FORALL loops may also be present in Fortran 77D.

We assign to the Fortran 90D front end the remaining task, scalarizing Fortran 90 constructs that have no equivalent in the Fortran 77D intermediate form. There are three principal Fortran 90 language features that must be scalarized: array constructs, WHERE statements, and intrinsic functions [6].

ARRAY CONSTRUCTS Fortran 90 *array constructs* allow entire arrays to be manipulated atomically, enhancing the clarity and conciseness of the program. Previous research has shown that efficiently implementing array constructs for scalar machines may be difficult [5]. One problem is that when Fortran 90 array constructs are used in assignment statements, the entire right-hand side (*rhs*) must be evaluated before storing the results in the left-hand side (*lhs*). Without adequate analysis, *rhs* array elements would need to be stored in temporary buffers to ensure that they are not overwritten before their values are used. The Fortran 90 front end can defer this problem by relying on a key observation—the FORALL loop possesses copy-in/copy-out semantics identical to Fortran 90 assignment statements utilizing array constructs. Such statements may thus be translated into equivalent FORALL loops with no loss of information.

WHERE STATEMENT Another Fortran 90 feature is the WHERE statement. It takes a boolean argument that is used to *mask* array operations, inhibiting assignments to array elements whose matching boolean flag has the value *false*. Fortunately, the WHERE statement may be easily translated into equivalent IF and FORALL statements. Consider the following example where A is assumed to be an 1D N-element array. Because of FORALL copy-in/copy-out semantics, it is unnecessary at this point to explicitly store the value of the boolean argument to prevent it from being overwritten.

```
WHERE (A .EQ. 0)         FORALL i = 1,N
  A = 1.0                  IF (A(i) .EQ. 0) THEN
ELSEWHERE        ==>         A(i) = 1.0
  A = 0.0                  ELSE
ENDWHERE                     A(i) = 0.0
                           ENDIF
                         ENDFOR
```

INTRINSIC FUNCTIONS Intrinsic functions are fundamental to Fortran 90. They not only provide a concise means of expressing operations on arrays, but also identify parallel computation patterns that may be difficult to detect automatically. Fortran 90 provides intrinsic functions for operations such as shift, reduction, transpose, and matrix multiplication. Additional intrinsics are described in Table 1. To avoid excessive complexity and machine-dependence in the Fortran D compiler, we convert most Fortran 90 intrinsics into calls to customized run-time library functions.

The strategy used by the Fortran 90D front end is thus to preserve all intrinsic functions, passing them to the Fortran D compiler back end. However, some processing is necessary. Like the WHERE statement, some intrinsic functions accept a mask expression that restricts execution of the computation. The Fortran 90D front end may need to evaluate the expression and store it in a temporary boolean array before performing the computation, so the mask can be passed as an argument to the run-time library.

For example, consider the following reduction operation, where X is a scalar and A, B are arrays:

```
X = MAXVAL(A, A .EQ. B)
```

It should return the value of the element of A that is the maximum of all elements for which element of A is equal to the corresponding element of B. The Fortran 90D front end translates this to:

```
FORALL i = 1,N
  TMP(i) = A(i) .EQ. B(i)
ENDFOR
X = MAXVAL(A, TMP)
```

TMP can then be passed as an argument to the run-time routine MAXVAL. Temporary arrays may also be introduced when intrinsic functions return a value that is part of a Fortran 90 expression.

TEMPORARY ARRAYS When the Fortran 90D front end needs to create temporary arrays, it must also generate appropriate Fortran D data decomposition statements. A temporary array is usually aligned and distributed in the same manner as its master array. For example, in the previous example the temporary logical array TMP is aligned and distributed in the same manner as A and B. If A and B are distributed differently, automatic data decomposition techniques must be applied to select efficient data distributions for compiler-generated temporary arrays [16, 27, 29].

4.2 FORTRAN 77D FRONT END

The Fortran 77D front end does not need to perform much work since Fortran 77D is very close to the intermediate form. Its only task is to detect complex high-level parallel computations, replacing or annotating them by their equivalent Fortran 90 intrinsics. These intrinsic functions help the compiler recognize complex computations such as reductions and scans that are supported by the run-time library. With advanced program analysis, some operations such as DOTPRODUCT, SUM, TRANSPOSE, or MATMUL can be detected automatically with ease. Others computations such as COUNT or PACK may require user assistance.

4.3 FORTRAN D BACK END

The Fortran D back end performs two main functions—it partitions the program onto the nodes of the parallel machine and completes the scalarization of Fortran D into Fortran 77. We find that the desired order for compilation phases is to apply loop fusion first, followed by partitioning and sectioning.

Loop fusion is performed first because it simplifies partitioning by reducing the need to consider inter-loop interactions. It also enables optimizations such as *strip-mining* and *loop interchange* [4, 41]. In addition, loop fusion does not increase the difficulty of later compiler phases. On the other hand, sectioning is performed last because it can significantly disrupt the existing program structure, increasing the difficulty of partitioning analysis and optimization.

4.3.1 LOOP FUSION Loop fusion is particularly important for the Fortran D back end because scalarized Fortran 90 programs present many single-statement loop nests. Fusing such loops simplifies the partitioning process and enables additional optimizations.

Data dependence is a concept developed for vectorizing and parallelizing compilers to characterize memory access patterns at compile time [4, 31, 41]. A true dependence indicates definition followed by use, while an anti-dependence shows use before definition. Data dependences may be either loop-carried or loop-independent. Loop fusion is legal if it does not reverse the direction of any data dependence between two loop nests [5, 41].

The current Fortran D back end fuses all adjacent loop nests where legal, if no loop-carried true dependences are introduced. This heuristic does not adversely affect the parallelism or communication overhead of the resulting program, and should perform well for the simple cases found in practice. More sophisticated algorithms are discussed elsewhere [34].

Loop fusion also has the added advantage of being able to improve memory reuse in the resulting program. Modern high-performance processors are so fast that memory latency and bandwidth limitations become the performance bottlenecks for most scientific programs. Transformations such as loop fusion promote memory reuse and can significantly improve program efficiency for both scalar and vector machines [5, 31, 34, 39]. For instance, consider the following example.

```
FORALL i = 1,N           FORALL i = 1,N
  A(i) = i                 A(i) = i
ENDFOR            ==>      B(i) = A(i)*A(i)
FORALL i = 1,N           ENDFOR
  B(i) = A(i)*A(i)
ENDFOR
```

The occurrences of $A(i)$ in separate loops means that the memory location referenced by $A(i)$ in the first loop is likely to have been flushed from the cache by the reference in the second loop. If the two loops are fused, all accesses to $A(i)$ occur in the same loop iteration, allowing the value to be reused in a register or cache. For this example, we measured improvements of up to 30% for some problem sizes on an Intel i860, as shown in Figure 2. Additional transformations to enhance memory reuse and increase unit-stride memory accesses are also quite important; they are described elsewhere [28, 34].

4.3.2 PROGRAM PARTITIONING The major step in compiling Fortran D for MIMD distributed-memory machines is to partition the data and computation across processors, introducing communication where needed. We present a brief overview of the Fortran D compilation process below, details are discussed elsewhere [20, 24, 25, 40].

- **Analyze Program.** Symbolic and data dependence analysis is performed.
- **Partition data.** Fortran D data decomposition specifications are analyzed to determine the decomposition of each array in a program.
- **Partition computation.** The compiler partitions computation across processors using the "owner computes" rule—where each processor only computes values of data it owns [14, 37, 43].
- **Analyze communication.** Based on the work partition, references that result in nonlocal accesses are marked.
- **Optimize communication.** Nonlocal references are examined to determine optimization opportunities. The key optimization, message vectorization, uses the level of loop-carried true dependences to combine element messages into vectors [8, 43].
- **Manage storage.** "Overlaps" [43] or buffers are allocated to store nonlocal data.
- **Generate code.** Information gathered previously is used to generate the SPMD program with explicit message-passing that executes directly on the nodes of the distributed-memory machine.

Two extensions are needed in the Fortran D back end to handle FORALL loops and intrinsics. During communication optimization, the Fortran D compiler treats all true dependences carried by FORALL loops as anti-dependences. This reflects the semantics of the FORALL loop and ensures that the message vectorization algorithm will place all communication outside the loop. In addition, during code generation intrinsic functions are translated into calls to the run-time library. Parameters are added where necessary to provide necessary data partitioning information.

4.3.3 SECTIONING

The final phase of the Fortran D back end completes the scalarization process. After partitioning is performed, the compiler applies *sectioning* to convert FORALL loops into DO loops [5] in the node program. The Fortran D back end detects cases where temporary storage may be needed using data dependence analysis. True dependences carried on the FORALL loop represent instances where values are defined in the loop and used on later iterations; they point out where the copy-in/copy-out semantics of the FORALL loop is being violated.

During simple translation of Fortran 90 array constructs or FORALL loops, arrays involved in loop-carried true dependences must be saved in temporary buffers to preserve their old values. For instance, consider the translation of the following concise Fortran 90 formulation of the Jacobi algorithm:

```
A(2:N-1) = 0.5 * (A(1:N-2) + A(3:N))
      ⇓
FORALL i = 2,N-1
  A(i) = 0.5 * (A(i-1) + A(i+1))
ENDFOR
      ⇓
DO i = 1,N-2
  TMP(i) = A(i-1)
ENDDO
DO i = 2,N-1
  A(i) = 0.5 * (TMP(i) + A(i+1))
ENDFOR
```

A loop-carried true dependence exists between the definition to $A(i)$ and the use of $A(i-1)$. A temporary array TMP is needed so that the old values of $A(i-1)$ are not overwritten before they are used. The values of $A(i+1)$ do not need to be buffered since they are used before being redefined.

The previous example is problematic because temporary storage is required for the values of $A(i-1)$. In some cases, the Fortran D compiler can eliminate buffering through program transformations such as *loop reversal*. In other cases, the compiler can reduce the amount of temporary storage required through *data prefetching* [5, 13]. For instance, in the Jacobi example a more efficient translation would result in:

```
X = A(1)
DO i = 2,N-1
  Y = 0.5 * (X + A(i+1))
  X = A(i)
  A(i) = Y
ENDFOR
```

This reduces the temporary memory required significantly, from an entire array to two scalars. For this version of Jacobi, we measured improvements of up to 50% for certain problem sizes on an Intel i860, as shown in Figure 2.

5. RUN-TIME LIBRARY

Fortran 90 intrinsic functions represent computations (such as TRANSPOSE and MATMUL) that may have complex communication patterns. It is possible to support these functions at compile time, but we have chosen to implement these functions in the run-time library instead to reduce the complexity and machine-dependence of the compiler. The Fortran D compiler translates intrinsics into calls to run-time library routines using a standard interface. Additional information is passed describing bounds, overlaps, and partitioning for each array dimension. The run-time library developed at Syracuse is built on top of the Express communication package to ensure portability across different architectures [36].

Table 2 presents some sample performance numbers for a subset of the intrinsic functions on an iPSC/860, details are presented elsewhere [1]. The times in the table include both the computation and communication times for each function. These measurements are also displayed in Figure 3. Timings in seconds are plotted logarithmically along the Y-axis. The number of processors is plotted along the X-axis. Different lines in the graph correspond to timings of individual functions in the run-time library for different problem sizes. For large problem sizes, we were able to obtain almost linear speedups. In the case of TRANSPOSE function, going from one processor to two or four degrades execution time due to increased communication. However, speedup improves as the number of processors increases.

6. FORTRAN 90D COMPILATION EXAMPLE

6.1 COMPILATION

Figure 4 shows a code fragment implementing one sweep of ADI integration on a 2D mesh, a typical (if short) numerical algorithm. Conceptually, the code is solving a tridiagonal system (represented by the arrays A and B) along each row of the matrix X. The tridiagonal systems are solved by a sequential method, but separate columns are independent and may be solved in parallel. The full version of ADI integration sweeps each dimension of the mesh, preventing completely parallel execution for any static data decomposition.

In the example, Fortran D data decomposition statements are used to partition the 2D array into blocks of columns. For clarity, we declare the number of processors (N$PROC) to be 32 at compile time. The Fortran 90D example is concise and convenient for the user, since it can be written for a single address space without requiring explicit communication. However, additional compilation techniques are required to generate efficient code. First, the Fortran 90D front end translates the program into intermediate form as shown in Figure 5, converting all array constructs into FORALL loops. Since no true dependences are carried on the FORALL loops, they may be directly replaced with DO loops.

The compilation process for the Fortran D back end merits closer examination. First, array bounds are reduced to the local sections plus overlaps. The local processor number is determined using *myproc*(), a library function; it is used to compute expressions for reducing loop bounds. Analysis determines that both I and J are *cross-processor* loops—loops carrying true dependences that sequentialize the computation across processors. To exploit pipeline parallelism, the Fortran D compiler interchanges such loops inward. We call this technique *fine-grain pipelining* [24, 25].

For this version of ADI integration, data dependences permit the Fortran D compiler to interchange the J loop inwards. However, if loop fusion is not performed, the imperfectly nested K loops inhibit loop interchange for loop I, forcing it to remain in place. During code generation, true dependences for nonlocal references carried on the I and J loop cause calls to *send* and *recv* to be inserted to provide communication and synchronization. Figure 6 shows the resulting program. Unfortunately, the computation in the I loop has been sequentialized, since each processor has to wait for its predecessor to complete. Note that this is not due to communication placement; the values needed by the succeeding processor are simply computed last.

If loop fusion is enabled, the Fortran D back end will fuse the two inner K loops. This is legal because the dependence between the definition and use of B is carried on the I loop and is thus unaffected. Fusion is also conservative because it does not introduce any true dependences carried by the K loop. Fusing the K loops promotes reuse of A and B, but its main benefit is to enable the Fortran D back end to interchange the I and K loops, exposing pipeline parallelism. The resulting program is displayed in Figure 7. For simplicity, only the first loop is shown. The remaining loops are compiled in a similar manner as before. Note that *loop distribution* applied to the I loop can enable loop interchange, but is prevented by the recurrence for B.

To reduce communication overhead, we can also apply strip-mining in conjunction with loop interchange to adjust the granularity of pipelining. We call this technique *coarse-grain pipelining* [24, 25]. In the ADI example, we strip-mine the K loop by four, then interchange the resulting loop outside the I loop. For this study, the desirable strip size was derived by hand; we describe elsewhere how it may be automatically calculated for a given architecture based on communication and computation costs [25]. After strip-mining, messages inserted outside the K loop allow each processor to reduce communication costs at the expense of some parallelism, resulting in Figure 8. Except for coarse-grain pipelining, all these versions of ADI integration were generated automatically by the Fortran D compiler.

6.2 PERFORMANCE RESULTS

To validate these methods, we executed these codes on an iPSC/860. The programs were compiled under -O4 using Release 3.0 of *if77*, the iPSC/860 compiler. Timings were taken for three double-precision problem sizes using *dclock*() on a 32 node Intel iPSC/860 with 8 Meg of memory per node. Results for three problem sizes are tabulated in Table 3. Timings are not provided where problem size exceeds available memory. We also graphically display the timings in Figure 9. Execution times in seconds are plotted logarithmically along the Y-axis. The number of processors used is plotted logarithmically along the X-axis. Numbers for perfect or ideal speedup (sequential execution time divided by number of processors) are provided for comparison.

The original version of ADI (Figure 6) exploits pipeline parallelism in the J loop, but shows limited speedup, since the I loop is sequentialized. Fusing the K loops to improve memory reuse provides very little improvement in this case, yielding nearly identical results. Applying loop interchange after fusion to enable fine-grain pipelining (Figure 7) parallelizes the I loop as well, yielding significant speedup. Strip-mining to apply coarse-grain pipelining can improve efficiency an additional 10–50% (Figure 8). Pipelining comes closest to perfect speedup for large problems on a small number of processors.

We also compared the efficiency of pipelining versus dynamic data decomposition. By changing the distribution of data at run-time from columns to rows, all dependences in each sweep of ADI may be internalized, enabling completely parallel execution. Data must be redistributed twice, once to achieve the desired distribution, then a second time to return it to its original configuration. The cost of redistributing is approximated by the performance of the TRANSPOSE routine shown in Table 2.

Our results show that on the iPSC/860, dynamic data decomposition for this formulation of ADI integration achieves speedup. However, the resulting program is significantly slower than pipelining, even for small problems distributed over large numbers of processors, the expected best case for dynamic data decomposition. Our experiences show that some common algorithms, such as ADI integration, require significant amounts of optimization to compete with hand-crafted code.

7. RELATED WORK

The Fortran D compiler is a second-generation distributed-memory compiler that integrates and extends many previous analysis and optimization techniques. Many distributed-memory compilers reduce communication overhead by aggregating messages outside of parallel loops [26, 30] or parallel procedures [21, 38], while others rely on functional language [32] or single assignment semantics [37]. In comparison, the Fortran D compiler uses dependence analysis to automatically exploit parallelism and extract communication even from sequential loops such as those found in ADI integration.

Several other projects are also developing Fortran 90 compilers for MIMD distributed-memory machines. ADAPT [35] and ADAPTOR [12] propose to scalarize and partition Fortran 90 programs using a run-time library for Fortran 90 intrinsics. The CM FORTRAN compiler compiles Fortran 90 with alignment and layout specifications directly to the physical machine, and can optimize floating point register usage [2, 13]. The FORTRAN-90-Y compiler uses formal specification techniques to generate efficient code for the CM-2 and CM-5 [17]. FORGE90, formerly MIMDIZER, is an interactive parallelization system for MIMD shared and distributed-memory machines from Applied Parallel Research [7]. It performs data-flow and dependence analyses, and also supports loop-level transformations. Our compiler resembles the VIENNA FORTRAN 90 compiler [11] derived from SUPERB [43]. It has also been influenced by a proposal by Wu & Fox that discussed program generation and optimization using a test-suite approach [42].

A number of researchers have studied techniques to reduce storage and promote memory reuse [5, 28, 31, 34, 39]. These optimizations have proved useful for both scalar and parallel machines. The goal of the Fortran 90D compiler is to integrate these scalarization techniques with advanced communication and parallelism optimizations.

8. CONCLUSIONS

This paper presents an integrated approach to compiling both Fortran 77D and 90D based on a few key observations. First, using FORALL preserves information in Fortran 90 array constructs. Dividing the scalarization process into translation, loop fusion, and sectioning allows it to be easily integrated with the partitioning performed by the Fortran D compiler. A portable run-time library can also reduce the complexity and machine-dependence of the compiler. All optimizations except data prefetching have been implemented in the current Fortran D compiler prototype.

Compiling for MIMD distributed-memory machines is only a part of the Fortran D project. We also are working on Fortran 77D and Fortran 90D compilers for SIMD machines, translations between the two Fortran dialects, support for irregular computations, and environmental support for static performance estimation and automatic data decomposition [8, 9, 23, 27].

9. ACKNOWLEDGEMENTS

We are grateful to the ParaScope and Fortran D research groups for their assistance, and to Parasoft for providing the Fortran 90 parser and Express. Use of the Intel iPSC/860 was provided by the CRPC under NSF Cooperative Agreement Nos. CCR-8809615 and CDA-8619893 with support from the Keck Foundation. The content of this information does not necessarily reflect the position or policy of the government, and no official endorsement should be inferred.

REFERENCES

[1] Ahmad, I., Choudhary, A., Fox, G., Parasuram, K., Ponnusamy, R., Ranka, S., and Thakur, R. Implementation and scalability of Fortran 90D intrinsic functions on distributed memory machines. Technical Report SCCS-256, NPAC, Syracuse University, Mar. 1992.

[2] Albert, E., Knobe, K., Lukas, J., and Steele, Jr., G. Compiling Fortran 8x array features for the Connection Machine computer system. In *Proceedings of the ACM SIGPLAN Symposium on Parallel Programming: Experience with Applications, Languages, and Systems (PPEALS)* (New Haven, CT, July 1988).

[3] Albert, E., Lukas, J., and Steele, Jr., G. Data parallel computers and the forall statement. *Journal of Parallel and Distributed Computing 13*, 2 (Oct. 1991), 185–192.

[4] Allen, J. R., and Kennedy, K. Automatic translation of Fortran programs to vector form. *ACM Transactions on Programming Languages and Systems 9*, 4 (Oct. 1987), 491–542.

[5] Allen, J. R., and Kennedy, K. Vector register allocation. *IEEE Transactions on Computers 41*, 10 (Oct. 1992), 1290–1317.

[6] ANSI X3J3/S8.115. Fortran 90, June 1990.

[7] Applied Parallel Research. *Forge 90 Distributed Memory Parallelizer: User's Guide*, version 8.0 ed. Placerville, CA, 1992.

[8] Balasundaram, V., Fox, G., Kennedy, K., and Kremer, U. An interactive environment for data partitioning and distribution. In *Proceedings of the 5th Distributed Memory Computing Conference* (Charleston, SC, Apr. 1990).

[9] Balasundaram, V., Fox, G., Kennedy, K., and Kremer, U. A static performance estimator to guide data partitioning decisions. In *Proceedings of the Third ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Williamsburg, VA, Apr. 1991).

[10] Beltrametti, M., Bobey, K., and Zorbas, J. The control mechanism for the Myrias parallel computer system. *Computer Architecture News 16*, 4 (Sept. 1988), 21–30.

[11] Benkner, S., Chapman, B., and Zima, H. Vienna Fortran 90. In *Proceedings of the 1992 Scalable High Performance Computing Conference* (Williamsburg, VA, Apr. 1992).

[12] Brandes, T. Automatic translation of data parallel programs to message passing programs. In *Proceedings of AP'93 International Workshop on Automatic Distributed Memory Parallelization, Automatic Data Distribution and Automatic Parallel Performance Prediction* (Saarbrücken, Germany, Mar. 1993).

[13] Bromley, M., Heller, S., McNerney, T., and Steele, Jr., G. Fortran at ten gigaflops: The Connection Machine convolution compiler. In *Proceedings of the SIGPLAN '91 Conference on Program Language Design and Implementation* (Toronto, Canada, June 1991).

[14] Callahan, D., and Kennedy, K. Compiling programs for distributed-memory multiprocessors. *Journal of Supercomputing 2* (October 1988), 151–169.

[15] Carr, S., Kennedy, K., McKinley, K. S., and Tseng, C. Compiler optimizations for improving data locality. Tech. Rep. TR92-195, Dept. of Computer Science, Rice University, Nov. 1992.

[16] Chatterjee, S., Gilbert, J., Schreiber, R., and Teng, S. Automatic array alignment in data-parallel programs. In *Proceedings of the Twentieth Annual ACM Symposium on the Principles of Programming Languages* (Charleston, SC, Jan. 1993).

[17] Chen, M., and Cowie, J. Prototyping Fortran-90 compilers for massively parallel machines. In *Proceedings of the SIGPLAN '92 Conference on Program Language Design and Implementation* (San Francisco, CA, June 1992).

[18] Cooper, K., Hall, M. W., Hood, R. T., Kennedy, K., McKinley, K. S., Mellor-Crummey, J. M., Torczon, L., and Warren, S. K. The ParaScope parallel programming environment. *Proceedings of the IEEE 81*, 2 (Feb. 1993), 244–263.

[19] Fox, G., Hiranandani, S., Kennedy, K., Koelbel, C., Kremer, U., Tseng, C., and Wu, M. Fortran D language specification. Tech. Rep. TR90-141, Dept. of Computer Science, Rice University, Dec. 1990.

[20] Hall, M. W., Hiranandani, S., Kennedy, K., and Tseng, C. Interprocedural compilation of Fortran D for MIMD distributed-memory machines. In *Proceedings of Supercomputing '92* (Minneapolis, MN, Nov. 1992).

[21] Hatcher, P., Quinn, M., Lapadula, A., Seevers, B., Anderson, R., and Jones, R. Data-parallel programming on MIMD computers. *IEEE Transactions on Parallel and Distributed Systems 2*, 3 (July 1991), 377–383.

[22] High Performance Fortran Forum. High Performance Fortran language specification, version 1.0. Tech. Rep. CRPC-TR92225, Center for Research on Parallel Computation, Rice University, Houston, TX, Jan. 1993.

[23] Hiranandani, S., Kennedy, K., Koelbel, C., Kremer, U., and Tseng, C. An overview of the Fortran D programming system. In *Languages and Compilers for Parallel Computing, Fourth International Workshop* (Santa Clara, CA, Aug. 1991), U. Banerjee, D. Gelernter, A. Nicolau, and D. Padua, Eds., Springer-Verlag.

[24] Hiranandani, S., Kennedy, K., and Tseng, C. Compiling Fortran D for MIMD distributed-memory machines. *Communications of the ACM 35*, 8 (Aug. 1992), 66–80.

[25] Hiranandani, S., Kennedy, K., and Tseng, C. Evaluation of compiler optimizations for Fortran D on MIMD distributed-memory machines. In *Proceedings of the 1992 ACM International Conference on Supercomputing* (Washington, DC, July 1992).

[26] Ikudome, K., Fox, G., Kolawa, A., and Flower, J. An automatic and symbolic parallelization system for distributed memory parallel computers. In *Proceedings of the 5th Distributed Memory Computing Conference* (Charleston, SC, Apr. 1990).

[27] Kennedy, K., and Kremer, U. Automatic data alignment and distribution for loosely synchronous problems in an interactive programming environment. Technical Report TR91-155, Dept. of Computer Science, Rice University, Apr. 1991.

[28] Kennedy, K., and McKinley, K. S. Optimizing for parallelism and data locality. In *Proceedings of the 1992 ACM International Conference on Supercomputing* (Washington, DC, July 1992).

[29] Knobe, K., Lukas, J., and Steele, Jr., G. Data optimization: Allocation of arrays to reduce communication on SIMD machines. *Journal of Parallel and Distributed Computing 8*, 2 (Feb. 1990), 102–118.

[30] Koelbel, C., and Mehrotra, P. Compiling global name-space parallel loops for distributed execution. *IEEE Transactions on Parallel and Distributed Systems 2*, 4 (Oct. 1991), 440–451.

[31] Kuck, D., Kuhn, R., Padua, D., Leasure, B., and Wolfe, M. J. Dependence graphs and compiler optimizations. In *Conference Record of the Eighth Annual ACM Symposium on the Principles of Programming Languages* (Williamsburg, VA, Jan. 1981).

[32] Li, J., and Chen, M. Compiling communication-efficient programs for massively parallel machines. *IEEE Transactions on Parallel and Distributed Systems 2*, 3 (July 1991), 361–376.

[33] Lundstrom, S., and Barnes, G. Controllable MIMD architectures. In *Proceedings of the 1980 International Conference on Parallel Processing* (St. Charles, IL, Aug. 1980).

[34] McKinley, K. S. *Automatic and Interactive Parallelization.* PhD thesis, Dept. of Computer Science, Rice University, Apr. 1992.

[35] Merlin, J. ADAPTing Fortran-90 array programs for distributed memory architectures. In *First International Conference of the Austrian Center for Parallel Computation* (Salzburg, Austria, Sept. 1991).

[36] Parasoft Corporation. *Express User's Manual*, 1989.

[37] Rogers, A., and Pingali, K. Process decomposition through locality of reference. In *Proceedings of the SIGPLAN '89 Conference on Program Language Design and Implementation* (Portland, OR, June 1989).

[38] Rosing, M., Schnabel, R., and Weaver, R. The DINO parallel programming language. *Journal of Parallel and Distributed Computing 13*, 1 (Sept. 1991), 30–42.

[39] Sarkar, V., and Gao, G. Optimization of array accesses by collective loop transformations. In *Proceedings of the 1991 ACM International Conference on Supercomputing* (Cologne, Germany, June 1991).

[40] Tseng, C. *An Optimizing Fortran D Compiler for MIMD Distributed-Memory Machines.* PhD thesis, Dept. of Computer Science, Rice University, Jan. 1993.

[41] Wolfe, M. J. *Optimizing Supercompilers for Supercomputers.* The MIT Press, Cambridge, MA, 1989.

[42] Wu, M., and Fox, G. A test suite approach for Fortran 90D compilers on MIMD distributed memory parallel computers. In *Proceedings of the 1992 Scalable High Performance Computing Conference* (Williamsburg, VA, Apr. 1992).

[43] Zima, H., Bast, H.-J., and Gerndt, M. SUPERB: A tool for semi-automatic MIMD/SIMD parallelization. *Parallel Computing 6* (1988), 1–18.

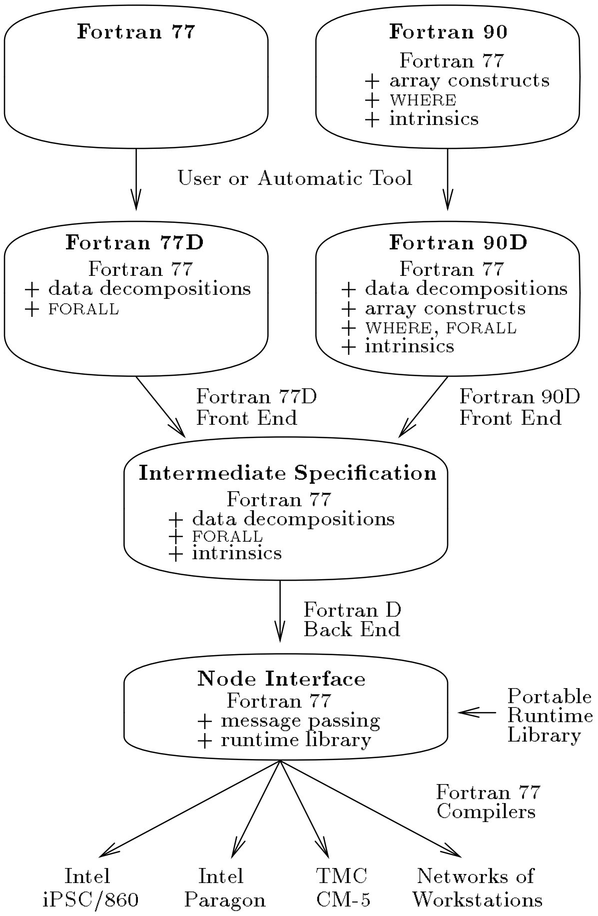

Figure 1 Fortran D Compilation Strategy

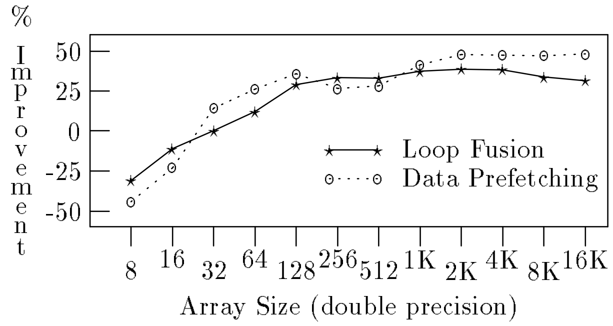

Figure 2 Effect of Scalarization Optimizations

Data Movement	Reductions	Irregular Operations	Special Routines
CSHIFT EOSHIFT SPREAD RESHAPE TRANSPOSE	DOTPRODUCT ALL, ANY, COUNT MAXVAL, MINVAL SUM, PRODUCT MAXLOC, MINLOC	PACK UNPACK	MATMUL

Table 1 Representative Intrinsic Functions of Fortran 90D

Processor	*Time* (milliseconds)						
	ALL	ANY	MAXVAL	PRODUCT	TRANSPOSE		
	$1K \times 1K$	$1K \times 1K$	$1K \times 1K$	$256K$	256×256	512×512	$1K \times 1K$
1	580.6	606.2	658.8	90.1	58	299	-
2	291.0	303.7	330.4	50.0	118	575	-
4	146.2	152.6	166.1	25.1	87	395	-
8	73.84	77.1	84.1	13.1	61	224	1039
16	37.9	39.4	43.4	7.2	41	140	539
32	19.9	20.7	23.2	4.2	36	85	316

Table 2 Performance of Some Fortran 90 Intrinsic Functions

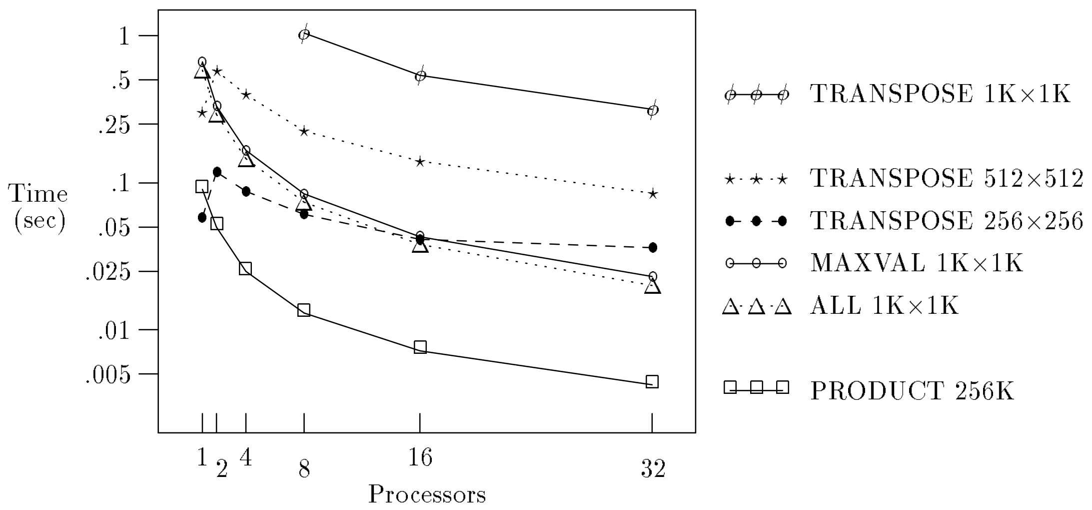

Figure 3 Performance of Run-time Library

```
PARAMETER (N = 512, N$PROC = 32)
REAL X(N,N), A(N,N), B(N,N)
DECOMPOSITION DEC(N,N)
ALIGN X, A, B WITH DEC
DISTRIBUTE DEC(:,BLOCK)
DO I = 2,N
  X(1:N,I) = X(1:N,I) - X(1:N,I-1)*A(1:N,I)/B(1:N,I-1)
  B(1:N,I) = B(1:N,I) - A(1:N,I)*A(1:N,I)/B(1:N,I-1)
ENDDO
X(1:N,N) = X(1:N,N) / B(1:N,N)
DO J = N-1,1,-1
  X(1:N,J) = (X(1:N,J)-A(1:N,J+1)*X(1:N,J+1))/B(1:N,J)
ENDDO
```

Figure 4 ADI integration in Fortran 90D

```
PARAMETER (N = 512)
REAL X(N,N), A(N,N), B(N,N)
DO I = 2,N
  FORALL K = 1,N
    X(K,I) = X(K,I) - X(K,I-1)*A(K,I)/B(K,I-1)
  ENDFOR
  FORALL K = 1,N
    B(K,I) = B(K,I) - A(K,I)*A(K,I)/B(K,I-1)
  ENDFOR
ENDDO
FORALL K = 1,N
  X(K,N) = X(K,N)/B(K,N)
ENDFOR
DO J = N-1,1,-1
  FORALL K = 1,N
    X(K,J) = (X(K,J)-A(K,J+1)*X(K,J+1))/B(K,J)
  ENDFOR
ENDDO
```

Figure 5 ADI in Intermediate Form

```
REAL X(512,0:17), A(512,17), B(512,0:16)
MY$P = myproc()      {* 0...31 *}
LB1 = MAX((MY$P*16)+1,2) - MY$P*16
UB1 = MIN((MY$P+1)*16,511) - MY$P*16
IF (MY$P .GT. 0) recv(X(1:N,0),B(1:N,0),MY$P-1)
DO I = LB1, 16
  DO K = 1,N
    X(K,I) = X(K,I) - X(K,I-1)*A(K,I)/B(K,I-1)
  ENDDO
  DO K = 1,N
    B(K,I) = B(K,I) - A(K,I)*A(K,I)/B(K,I-1)
  ENDDO
ENDDO
IF (MY$P .LT. 31) send(X(1:N,16),B(1:N,16),MY$P+1)
IF (MY$P .EQ. 31) THEN
  DO K = 1,N
    X(K,16) = X(K,16)/B(K,16)
  ENDDO
ENDIF
IF (MY$P .GT. 0) send(A(1:N,1),MY$P-1)
IF (MY$P .LT. 31) recv(A(1:N,17),MY$P+1)
DO K = 1,N
  IF (MY$P .LT. 31) recv(X(K,17),MY$P+1)
  DO J = UB1,1,-1
    X(K,J) = (X(K,J)-A(K,J+1)*X(K,J+1))/B(K,J)
  ENDDO
  IF (MY$P .GT. 0) send(X(K,1),MY$P-1)
ENDDO
```

Figure 6 ADI without Loop Fusion

```
REAL X(512,0:17), A(512,17), B(512,0:16)
MY$P = myproc()      {* 0...31 *}
LB1 = MAX((MY$P*16)+1,2) - MY$P*16
DO K = 1,N
  IF (MY$P .GT. 0) recv(X(K,0),B(K,0),MY$P-1)
  DO I = LB1,16
    X(K,I) = X(K,I) - X(K,I-1)*A(K,I)/B(K,I-1)
    B(K,I) = B(K,I) - A(K,I)*A(K,I)/B(K,I-1)
  ENDDO
  IF (MY$P .LT. 31) send(X(K,16),B(K,16),MY$P+1)
ENDDO
```

Figure 7 ADI with Fine-grain Pipelining (first loop only)

```
REAL X(512,0:17), A(512,17), B(512,0:16)
MY$P = myproc()      {* 0...31 *}
LB1 = MAX((MY$P*16)+1,2) - MY$P*16
DO KK = 1,N,4
  IF (MY$P .GT. 0) THEN
     recv(X(KK:KK+3,0),B(KK:KK+3,0),MY$P-1)
  ENDIF
  DO I = LB1,16
    DO K = KK,KK+3
      X(K,I) = X(K,I) - X(K,I-1)*A(K,I)/B(K,I-1)
      B(K,I) = B(K,I) - A(K,I)*A(K,I)/B(K,I-1)
    ENDDO
  ENDDO
  IF (MY$P .LT. 31) THEN
     send(X(KK:KK+3,16),B(KK:KK+3,16),MY$P+1)
  ENDIF
ENDDO
```

Figure 8 ADI with Coarse-grain Pipelining (first loop only)

Problem Size	P	ADI w/o Loop Fusion	Fine-grain Pipelining	Coarse-grain Pipelining	Data Redistribution	Perfect Speedup
	1	1.22	1.45	1.32	-	1.22
256	2	1.13	0.78	0.69	1.32	0.61
$\times$	4	1.02	0.45	0.38	0.83	0.30
256	8	0.96	0.28	0.21	0.52	0.15
	16	0.93	0.20	0.12	0.32	0.08
	32	0.96	0.17	0.08	0.25	0.04
	1	5.44	6.26	5.93	-	5.44
512	2	4.79	3.18	2.98	6.17	2.72
$\times$	4	4.29	1.72	1.53	3.72	1.36
512	8	4.04	0.97	0.81	2.02	0.68
	16	3.94	0.59	0.44	1.18	0.34
	32	3.94	0.41	0.26	0.68	0.17
	1	21.74	-	-	-	21.74
1K	4	17.13	6.44	5.98	-	5.44
$\times$	8	16.09	3.42	3.07	8.95	2.72
1K	16	15.61	1.91	1.62	4.59	1.36
	32	15.47	1.17	0.89	2.58	0.68

Table 3 Performance of ADI Integration (in seconds)

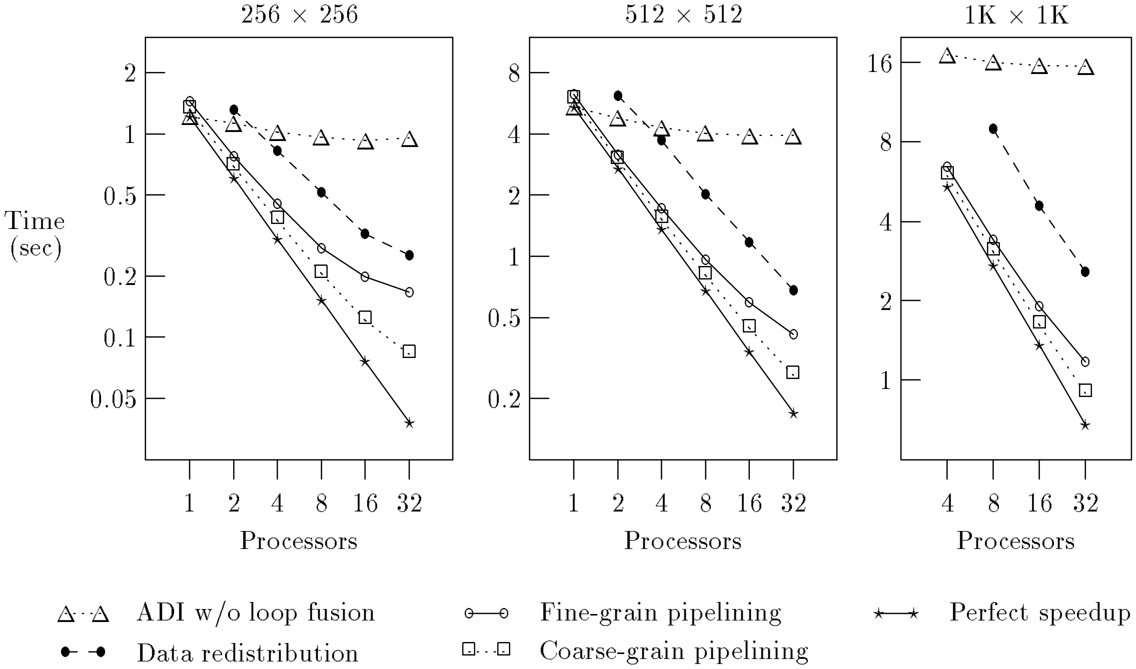

Figure 9 Execution Times for ADI Integration (double precision)