

Irregular computations in Fortran – expression and implementation strategies

Jan F. Prins *, Siddhartha Chatterjee
and Martin Simons

*Department of Computer Science, The University of
North Carolina, Chapel Hill, NC 27599-3175, USA
E-mail: {prins,sc,simons}@cs.unc.edu*

Modern dialects of Fortran enjoy wide use and good support on high-performance computers as performance-oriented programming languages. By providing the ability to express nested data parallelism, modern Fortran dialects enable irregular computations to be incorporated into existing applications with minimal rewriting and without sacrificing performance within the regular portions of the application. Since performance of nested data-parallel computation is unpredictable and often poor using current compilers, we investigate *threading* and *flattening*, two source-to-source transformation techniques that can improve performance and performance stability. For experimental validation of these techniques, we explore nested data-parallel implementations of the sparse matrix-vector product and the Barnes–Hut n -body algorithm by hand-coding thread-based (using OpenMP directives) and flattening-based versions of these algorithms and evaluating their performance on an SGI Origin 2000 and an NEC SX-4, two shared-memory machines.

1. Introduction

Modern science and engineering disciplines make extensive use of computer simulations. As these simulations increase in size and detail, the computational costs of naive algorithms can overwhelm even the largest parallel computers available today. Fortunately, computational costs can be reduced using sophisticated modeling methods that vary model resolution as needed, coupled with sparse and adaptive solution techniques that vary computational effort in time and space as needed. Such techniques have been developed and are routinely employed in sequential computation, for example, in cosmological simulations (using adap-

tive n -body methods) and computational fluid dynamics (using adaptive meshing and sparse linear system solvers).

However, these so-called irregular or unstructured computations are problematic for parallel computation, where high performance requires equal distribution of work over processors and locality of reference within each processor. For many irregular computations, the distribution of work and data cannot be characterized *a priori*, as these quantities are input-dependent and/or evolve with the computation itself. Further, irregular computations are difficult to express using performance-oriented languages such as Fortran, because there is an apparent mismatch between data types characteristic of irregular computations such as trees, graphs, and nested sequences, and the statically analyzable rectangular multi-dimensional arrays that are the core data types of Fortran.

In Fortran 90 and Fortran 95 [1,24], irregular data types can be introduced using the data abstraction facilities, with a representation exploiting pointers. However, compilers have difficulty generating high-performance parallel code for irregular computations expressed using such data types. In this paper we illustrate how irregular computations can be expressed well in Fortran 9X, and how additional source-level techniques can be used to achieve high performance execution of such computations on parallel processors. While HPF [21] focuses on the irregular distribution of regular data structures, our approach is based on the (regular) distribution of irregular data structures.

These modern Fortran dialects enjoy increasing use and good support as mainstream performance-oriented programming languages. By providing the ability to express irregular computations as Fortran modules, and by preprocessing these modules into a form that current Fortran compilers can successfully optimize, we enable irregular computations to be incorporated into existing applications with minimal rewriting and without sacrificing performance within the regular portions of the application.

*Corresponding author.

For example, consider the NAS Conjugate Gradient (CG) benchmark, which solves an unstructured sparse linear system using the method of conjugate gradients [3]. Within the distributed sample sequential Fortran solution, 79% of the lines of code are standard Fortran 77 concerned with problem construction and performance reporting. The next 16% consist of scalar and regular vector computations of the BLAS 2 variety [22], while the final 5% of the code is the irregular computation of the sparse matrix-vector product. Clearly to obtain high performance on the irregular computation we want to rewrite only this 5% of the code (which performs 97% of the work in the NAS CG class B benchmark), while the remainder should be left intact for the Fortran compiler. This is not just for convenience. It is also critical for performance reasons; following Amdahl's Law, as the performance of the irregular computation improves, the performance of the regular component becomes increasingly critical for sustained high performance overall. Fortran compilers provide good compiler/annotation techniques to achieve high performance for the regular computations in the problem and thus constitute an effective platform for building an efficient and seamless interface between the regular and irregular portions of the computation.

We manually applied the implementation techniques described in Section 4 to the irregular computation in the NAS CG problem. (Ultimately, these techniques are intended to be mechanized for inclusion in a compiler system.) The resultant Fortran program achieved a performance on the NAS CG class B benchmark of 13.5 GFLOPS using a 32 processor NEC SX-4 [32]. We believe this to be the highest performance achieved for this benchmark to date (April 1998). It exceeds, by a factor of 2.6, the highest performance reported in the last NAS Parallel Benchmark (1.0) Results [34], and is slightly faster than the 12.9 GFLOPS recently achieved using a 1024 processor Cray T3E-900 [23]. These encouraging initial results support the thesis that high-level expression and high-performance for irregular computations can be supported simultaneously in a production Fortran programming environment.

2. Expressing irregular computations using nested data parallelism

We adopt the data-parallel programming model of Fortran as our starting point. The data-parallel programming model has proven to be popular because of

its power and simplicity. Data-parallel languages are founded on the concept of collections (e.g., arrays) and a means to allow programmers to express parallelism through the application of an operation independently to all elements of a collection (e.g., the elementwise addition of two arrays). Sipelstein and Blleloch survey "collection-oriented" languages in [39].

Most of the common data-parallel languages, such as Fortran 9X, offer restricted data-parallel capabilities: they limit collections to multidimensional rectangular arrays, limit the elements of a collection to scalar and record types, and limit the operations that can be applied in parallel. These limitations simplify compile-time analysis and partitioning of the work and communication for parallel execution, but make it difficult to express irregular computations in this model.

If the elements of a collection are themselves permitted to have arbitrary type, then arbitrary functions can be applied in parallel over collections. In particular, by operating on a collection of collections, it is possible to specify a parallel computation, each simultaneous operation of which in turn involves (a potentially different-sized) parallel subcomputation. This programming model, called *nested data parallelism*, combines aspects of both data parallelism and control parallelism. It retains the simple programming model and portability of the data-parallel model while being better suited for describing algorithms on irregular data structures. The utility of nested data parallelism as an expressive mechanism has been understood for a long time in the LISP, SETL [36], and APL [25] communities, although always with a sequential execution semantics and implementation.

Nested data parallelism occurs naturally in the succinct expression of many irregular scientific problems. Consider the general sparse matrix-vector product at the heart of the NAS CG benchmark. In the popular compressed sparse row (CSR) format of representing general sparse matrices, the nonzero elements of an $m \times n$ sparse matrix A are represented as a sequence of m rows $A = [R_1, \dots, R_m]$. Each row R is, in turn, represented by a sequence of (v, c) pairs that combine a nonzero value v with the column c , $1 \leq c \leq n$, in which it occurs. With the dense vector x represented as a simple sequence of n values, the sparse matrix-vector product $y = Ax$ can be visualized graphically as shown in Fig. 1 and may be written using the NESL notation [6] as follows:

$$y = \{\text{sparse_dot_product}(R, x) : R \text{ in } A\}.$$




Fig. 1. Nested parallelism contained in the sparse matrix-vector product. We write `dot_product` instead of `sparse_dot_product` for brevity.

This expression specifies the application of `sparse_dot_product` (grey boxes), in parallel, to each row of A to yield the m element result sequence y . The sequence constructor $\{ \dots \}$ serves a dual role: it specifies parallelism (for each R in A), and it establishes the order in which the result elements are assembled into the result sequence, i.e., for $1 \leq i \leq m$,

$$y_i = \text{sparse_dot_product}(R_i, x).$$

We obtain nested data parallelism if the body expression `sparse_dot_product(R, x)` itself specifies the parallel computation of the dot product of row R with x as the sum-reduction of a sequence of nonzero products:

```
function sparse_dot_product(R, x)
    = sum({v * x[c]: (v, c) in R})
```

More concisely, the complete expression could also be written as follows:

$$y = \{\text{sum}(\{v * x[c]: (v, c) \text{ in } R\}) : R \text{ in } A\}$$

where the nested parallelism is now visible as nested sequence constructors in the source text.

Nested data parallelism thus provides a succinct and powerful notation for specifying parallel computation, including irregular parallel computations. Many more examples of efficient parallel algorithms expressed using nested data parallelism have been described in [6].

3. Nested data parallelism in Fortran

If we consider the expression of nested data parallelism in standard imperative programming languages,

we find that they either lack a data-parallel control construct (C, C++) or else lack a nested collection data type (Fortran). A nested data-parallel control construct can be added to C [13] or C++ [37], but the pervasive pointer semantics of these languages complicate its meaning. There is also incomplete agreement about the form parallelism should take in these languages.

The `FORALL` construct, originated in HPF [21] and later added into Fortran 95 [1,24], specifies data-parallel evaluation of expressions and array assignments. To ensure that there are no side effects between these parallel evaluations, functions that occur in the expressions must have the `PURE` attribute. Fortran 90 lacks a construct that specifies parallel evaluations. However, many compilers infer such an evaluation if specified using a conventional `DO` loop, possibly with an attached directive asserting the independence of iterations. Fortran 95 `FORALL` constructs (or Fortran 90 loops) may be nested. To specify nested data-parallel computations with these constructs, it suffices to introduce nested aggregates, which we can do via the data abstraction mechanism of Fortran 9X.

3.1. Sparse matrix-vector product

As a consequence of these language features, it is entirely possible to express nested data-parallel computations in modern Fortran. For example, we might introduce the types shown in Fig. 2 (top) to represent a sparse matrix. `sparse_element_t` is the type of a sparse matrix element, i.e., the (v, c) pair of the NESL example. `sparse_row_t` is the type of vectors (1-D arrays) of sparse matrix elements, i.e., a row of the matrix. A sparse matrix is characterized by the number of rows and columns, and by the nested sequence of sparse matrix elements.

Using these definitions, the sparse matrix-vector product can be succinctly written as shown in Fig. 2

```

MODULE Sparse_matrices

  IMPLICIT none

  TYPE Sparse_element_t
    REAL                :: val
    INTEGER              :: col
  END TYPE Sparse_element_t

  TYPE Sparse_row_t
    TYPE (Sparse_element_t), DIMENSION (:), POINTER :: elts
  END TYPE Sparse_row_t

  TYPE Sparse_matrix_t
    INTEGER              :: nrow, ncol
    TYPE (Sparse_row_t), DIMENSION (:), POINTER    :: rows
  END TYPE Sparse_matrix_t

CONTAINS

  SUBROUTINE smvp(a, x, y)

    TYPE (Sparse_matrix_t), INTENT(IN) :: a
    REAL, DIMENSION(:), INTENT(IN)   :: x
    REAL, DIMENSION(:), INTENT(OUT)  :: y

    FORALL (i = 1:a%nrow)
      y(i) = SUM(a%rows(i)%elts%val * x(a%rows(i)%elts%col))
    END FORALL

  END SUBROUTINE smvp

END MODULE Sparse_matrices

```

Fig. 2. Fortran 9X definition of a nested sequence type for sparse matrices (top) and its use for computing a sparse matrix-vector product (bottom).

(bottom). The `FORALL` loop specifies parallel evaluation of the inner products for all rows. Nested parallelism is a consequence of the use of parallel operations such as sum and elementwise multiplication, projection, and indexing.

3.2. Barnes–Hut force calculation algorithm

An illustration of a more complex form of nested data-parallel computation in Fortran 95 is the hierarchical force calculation algorithm of Barnes and Hut [4], as used, for example, in the simulation of self-gravitating systems. This algorithm exploits the fact that, at a distance, the combined potential of a group of particles can be *approximated* by the potential of the center of mass of that group. The algorithm makes use of a hierarchical quad-tree (or oct-tree in 3D) de-

composition of the space containing the particles, and associates with each region its center of mass. Fig. 3 (top) shows the Fortran types for this structure in the 2D case. The attribute `ntype` distinguishes between a particle and an inner cell of the tree that has up to four subtrees.

The force calculation traverses the tree top down for each particle to accumulate the total force on the particle. Subregions are explored only if the region's center of mass is not sufficiently far away from the particle to be used as an approximation. The function `tree_force` shown in Fig. 3 (bottom) encodes this algorithm in Fortran 95. Note that it is qualified as being `RECURSIVE` and, because it is called in a parallel context, `PURE`. The function `zero_force` returns the zero force vector; and the function `sum_forces` sums an array of force vectors.

```

MODULE Barnes_Hut

  IMPLICIT NONE

  TYPE Node_t
    INTEGER                :: ntype
    REAL                   :: mass
    REAL, DIMENSION(2)    :: pos
    TYPE (Body_t), POINTER :: body
    TYPE (Node_p), DIMENSION(4) :: subtree
  END TYPE Node_t

  TYPE Node_p
    TYPE (Node_t), POINTER :: node
  END TYPE Node_p

  TYPE Body_t
    REAL, DIMENSION(2)    :: vel
    REAL, DIMENSION(2)    :: acc
  END TYPE Body_t

CONTAINS

  PURE RECURSIVE FUNCTION tree_force(p,q) RESULT(f)

    TYPE (Node_t), POINTER :: p, q
    REAL, DIMENSION(2)    :: f

    INTEGER                :: I
    REAL, DIMENSION(4,2)  :: fs

    IF ("q is a body" THEN
      f = <body(p)-body(q) interaction>
    ELSE
      IF ("p is distant enough from q" ) THEN
        f = <body(p)-cell(q) interaction>
      ELSE
        fs = zero_force()
        FORALL (I = 1:4, ASSOCIATED(q%subtree(i)%node))
          fs(I,:) = tree_force(p,q%subtree(i)%node)
        END FORALL
        fs = sum_forces(fs)
      END IF
    END IF

  END SUBROUTINE tree_force

END MODULE Barnes_Hut

```

Fig. 3. Fortran 95 definition of a quad-tree for hierarchical force calculation in 2D (top) and its use in a hierarchical force calculation (bottom).

Since the different force-computations are independent of each other, they can be performed in parallel for all particles. For a given tree rooted at `root` and an array of particles `ps`, we can write

```
FORALL (I=1:SIZE(ps))
  CALL tree_force(ps(I),root)
END FORALL
```

to compute the forces on all particles in parallel.

This `FORALL` loop expresses nested parallelism because `tree_force` specifies parallelism in its recursive application. The parallelism specified in `tree_force` is dynamic since the available parallelism increases with the depth of the recursion. It is irregular because the degree of parallelism specified depends on the distribution of the particles. The realization of this kind of parallelism is complex and approaches to it are discussed in Section 4.

3.3. Nested data parallelism in an imperative setting

Earlier experiments with nested data parallelism in imperative languages include V [13], Amelia [37], and F90V [2]. For the first two of these languages the issues of side-effects in the underlying notation (C++ and C, respectively) were problematic in the potential introduction of interference between parallel iterations, and the efforts were abandoned. Fortran 95 finesses this problem by requiring procedures used within a `FORALL` construct to be `PURE`, an attribute that can be verified statically. This renders invalid those constructions in which side effects (other than the nondeterministic orders of stores) can be observed, although such a syntactic constraint is not enforced in Fortran 90.

The specification of nested data parallelism in Fortran and NESL differ in important ways, many of them reflecting differences between the imperative and functional programming paradigms. While V and F90V are imperative languages, they introduce nested parallelism in a functional fashion, following the approach found in NESL.

First, a sequence is formally a function from an index set to a value set. The NESL sequence constructor specifies parallelism over the value set of a sequence while the Fortran `FORALL` statement specifies parallelism over the index set of a sequence. This makes explicit the shape of the common index domain shared by several collections participating in a `FORALL` construct, and, in some cases, allows a more concise syntax.

Second, the NESL sequence constructor implicitly specifies the ordering of result elements, while this ordering is explicit in the `FORALL` statement. One consequence is that the restriction clause has different semantics. For instance, the NESL expression

$$v = \{i: i \text{ in } [1:n] \mid \text{oddp}(i)\}$$

yields a result sequence v of length $\lfloor n/2 \rfloor$ of odd values while the Fortran statement

```
FORALL (i = 1:n, odd(i)) v(i) = i
```

replaces the elements in the odd-numbered positions of v .

Third, the Fortran `FORALL` construct provides explicit control over memory. Explicit control over memory can be quite important for performance. For example, if we were to multiply the same sparse matrix repeatedly by different right hand sides (which is in fact exactly what happens in the CG benchmark), we could reuse a single temporary instead of freeing and allocating each time. Explicit control over memory also gives us a better interface to the regular portions of the computation.

Finally, the base types of a nested aggregate in Fortran are drawn from the Fortran data types and include multidimensional arrays and pointers. In NESL, we are restricted to simple scalar values and record types. Thus, expressing a sparse matrix as a collection of supernodes would be cumbersome in NESL. Another important difference is that we may construct nested aggregates of heterogeneous depth with Fortran, which we used in the representation of the Barnes–Hut tree.

4. Implementation issues

Expression of nested data-parallelism in Fortran is of limited interest and of no utility if such computations can not achieve high performance. Parallel execution and tuning for the memory hierarchy are the two basic requirements for high performance. Since the locus of activity and amount of work in a nested data-parallel computation can not be statically predicted, run-time techniques are generally required.

In the following, we will discuss two general strategies for the parallel execution of nested data parallelism, both consisting of a compile-time and a run-time component. The two strategies are illustrated for a nested data-parallel computation shown in Fig. 4(a).




Fig. 4. (a) Nested data-parallel program. (b) The associated dependence graph. (c) Thread decomposition of the graph. (d) Data-parallel decomposition of the graph.

Its associated dependence graph¹ is shown in Fig. 4(b). Here G and H denote assignment statements that can not introduce additional dependences, since there can be no data dependences between iterations of FORALL loops.

4.1. The thread-based approach

This technique conceptually spawns a different thread of computation for every parallel evaluation within a FORALL construct. The compile-time component constructs the threads from the nested loops. A run-time component dynamically schedules these threads across processors. Scheduling very fine-grained threads (e.g., a single multiplication) is impractical, hence compile-time techniques are required to increase thread granularity, although this may result in lost parallelism and increased load imbalance. Recent work has resulted in run-time scheduling techniques that minimize completion time and memory use of the generated threads [11,8,27]. The automatic construction of threads of appropriate granularity is currently being investigated by several researchers [26,19].

In Fig. 4(c) we show a decomposition of the total work into four parallel threads T_1, \dots, T_4 . In this decomposition the body of the inner FORALL loop has been serialized to increase the grain size of each thread.

¹We are using HPF INDEPENDENT semantics for the control dependences of a FORALL loop.

As a result the amount of work in each thread is quite different. On the other hand, since each thread executes a larger portion of the sequential implementation, it can exhibit good locality of reference.

Applying this strategy to the the sparse matrix-vector product code shown in Fig. 2 results in a parallelization of the outer loop and a serialization of the dot-product inner loop. This is not always optimal, since the distribution of work over outermost iterations depends on the input matrix and may be uneven or there may be insufficient parallelism in the outer iterations.

Similarly, applying the strategy to the Barnes–Hut force computation shown on page 6,

```

FORALL (I=1:SIZE(ps))
  fs(I) = tree_force(ps(I),root)
END FORALL

```

would serialize the parallelism within the `tree_force` function, but apply `tree_force` to each particle in parallel. In execution, each thread might evaluate the force on one or more particles, leading to potential load imbalances depending on the particle distribution. Some of the techniques that have been developed for the parallelization of treecodes in general and the Barnes–Hut algorithm in particular, have application in this setting. These include geometric orderings of the particles to improve locality, and cost-zoning to improve load balance [43,38].

4.2. The flattening approach

This technique replaces nested loops by a sequence of steps, each of which is a simple data-parallel operation. The compile-time component of this approach is a program transformation that replaces `FORALL` constructs with “data-parallel extensions” of their bodies and restructures the representation of nested aggregate values into a form suitable for the efficient implementation of the data-parallel operations [10,33]. The runtime component is a library of data-parallel operations closely resembling `HPFLIB`, the standard library that accompanies `HPF`.

In Fig. 4(d) we show a decomposition of the work into sequential steps S_1, \dots, S_3 , each of which is a simple data-parallel operation. The advantage of this approach is that we may partition the work in each operation very easily since it is all independent. Thus we can partition the work over processors and create parallel slack at each processor to hide network or memory latencies. In this example, the dependence structure permits the parallel execution of steps S_2 and S_3 , although this increases the complexity of the run time scheduler.

A nested data-parallel loop that has been flattened may perform a small multiplicative factor of additional work compared with a sequential implementation. However, full parallelism and optimal load balance are easily achieved in this approach. Compile-time techniques to fuse data-parallel operations can reduce the number of barrier synchronizations, decrease space requirements, and improve reuse [14,31,19].

To flatten the sparse matrix-vector product `smvp`, we replace the nested sequence representation of A with a linearized (flattened) representation (A', s) . Here A' is an array of r pairs, indexed by `val` and `col`, partitioned into rows of A by s , i.e., s is an array of n integers equal to the length of each row and whose sum is r . This datatype transformation is part of the flattening transformation rules. Application of the flattening transformations to the loop in Fig. 2 yields the single statement

```
y = segmented_sum(A'%val * x(A'%col), s),
```

where `segmented_sum` is a data-parallel operation with efficient parallel implementations [5]. By substituting `A'%val * x(A'%col)` for the first argument in the body of `segmented_sum`, the sum and product may be fused into a *segmented dot-product*.

To flatten the Barnes–Hut force computation is more complex and requires an extension of the conventional

flattening transformation to handle tree-like recursive datatypes. This extension has recently been proposed by Keller and Chakravarty in a functional setting [20, 19]. In this transformation, each level of the Barnes–Hut tree is represented by the sequence of the nodes at that level together with a segment descriptor that partitions the nodes according to their parents at the preceding level. This representation permits all invocations of `tree_force` to be evaluated in parallel at a given level of the tree.

5. Nested parallelism using current Fortran compilers

To gain some insight into the performance of nested data parallel Fortran programs, we examined the NAS CG benchmark and the sparse matrix-vector product contained within it. We started by presenting the program in Fig. 2 to be compiled for parallel execution by current Fortran compilers. For Fortran 90 compilers we replaced the `FORALL` with a corresponding `DO` statement.

For shared-memory multiprocessors we examined two auto-parallelizing Fortran 90 compilers: the SGI F90 V7.2.1 compiler (beta release, March 1998) for SGI Origin class machines and the NEC FORTRAN90/SX R7.2 compiler (release 140, February 1998) for the NEC SX-4. Since the nested parallel loops in `smvp` do not define a polyhedral iteration space, many classical compiler techniques for parallelization do not apply. However, both compilers recognize that iterations of the outer loop (over rows) are independent and, in both cases, these iterations are distributed over processors. The dot-product inner loop is compiled for serial execution (SGI F90) or vectorized (NEC F90).

For distributed memory multiprocessors we examined one `HPF` compiler. This compiler failed to compile `smvp` because it had no support for pointers in Fortran 90 derived types. Our impression is that this situation is representative of `HPF` compilers in general, since the focus has been on the parallel execution of programs operating on rectangular arrays. The data distribution issues for the more complex derived types with pointers are unclear. Instead, `HPF 2.0` supports the non-uniform distribution of arrays over processors. This requires the programmer to embed irregular data structures in an array and determine the appropriate mapping for the distribution.

We concluded that current Fortran compilers do not sufficiently address the problems of irregular nested

data parallelism. The challenge for irregular computations is to achieve uniformly high and predictable performance in the face of dynamically varying distribution of work.

Our approach is therefore use the threading and flattening techniques to transform nested data parallel constructs into simpler Fortran 90, providing integration with regular computations, and leveraging the capabilities of current Fortran compilers. This source-to-source translation restricts our options somewhat for the thread scheduling strategy. Since threads are not part of Fortran 90, the only mechanism for their (implicit) creation are loops, and the scheduling strategies we can choose from are limited by those offered by the compiler/run-time system. In this regard, standardized loop scheduling directives like the OpenMP directives [30] improve portability. Dynamic scheduling can be used to tolerate variations in progress among threads.

5.1. Sparse matrix-vector product

Consider a sparse $m \times n$ matrix A with a total of r nonzeros. Implementation of the simple nested data parallelism in the procedure `smvp` of Fig. 2 must address many of the problems that may arise in irregular computations:

- Uneven units of work: A may contain both dense and sparse rows.
- Small units of work: A may contain rows with very few nonzeros.
- Insufficient units of work: if n is less than the number of processors and r is sufficiently large, then parallelism should be exploited within the dot products rather than between the dot products.

We constructed two implementations of `smvp`. The *nested* implementation is obtained by direct compilation of the program in Fig. 2 using loop-level parallelization directives. This results in a parallelized outer loop, in which the dot products for different rows are statically or dynamically scheduled across processors. The *flat* implementation is obtained by flattening `smvp` as outlined in the previous section.

For the SGI Origin 2000, the flattened representation A' of A is divided into $p\sigma$ sections of length $r/(p\sigma)$ where p is the number of processors and $\sigma \geq 1$ is a factor to improve the load balance in the presence of multiprocessing and operating system overhead on the processors. Sections are processed independently and dot products are computed sequentially within each

section. Sums for segments spanning sections are adjusted after all sections are summed.

For the NEC SX-4, A' is divided into $pq\sigma$ sections where q is the vector length required by the vector units [7]. Section i , $0 \leq i < pq$, occupies element $i \bmod q$ in a vector of length q in thread $\lfloor i/p \rfloor$. Prefix dot-products are computed independently for all sections using a sequence of $r/(pq)$ vector additions on each processor. Segment dot-products are computed from the prefix dot-products and sums for segments spanning sections are adjusted after all sections are summed [32]. On the SX-4, typically $\sigma = 1$ since the operating system performs gang-scheduling and the threads experience very similar progress rates.

5.2. Experimental results

The SGI Origin 2000 used is a 32 processor cache-based shared memory multiprocessor. The processors are 250MHz R10000s with 32KB data cache and 4MB L2 unified instruction/data cache per processor. The NEC SX-4 used is a 16 processor shared-memory parallel vector processor with vector length 256. Each processor has a vector unit that can perform 8 or 16 memory reads or writes per cycle. The clock rate is 125 MHz. The memory subsystem provides sufficient sustained bandwidth to simultaneously service independent references from all vector units at the maximum rate.

The performance on square sparse matrices of both implementations is shown for 1, 2, 4, and 8 processors for the Origin 2000 in Fig. 5 and for 1, 2, and 4 processors for the SX-4 in Fig. 6. The top graph of each figure shows the performance as a function of problem size in megaflops per second, where the number of floating point operations for the problem is $2r$. Each row contains an average of 20 nonzeros and the number of rows is varied between 1000 and 175000. The bottom graph shows the influence of the average number of nonzeros per row (r/n) on the performance of the code. To measure this, we chose a fixed total number of non-zeros (12 million for the Origin 2000 and 3.6 million for the SX-4) and varied the average number of nonzeros on each row. In each case, the performance reported is averaged over 50 different matrices.

On the Origin 2000 the flattened implementation performed at least as well as the nested version over most inputs. The absolute performance of neither implementation is particularly impressive. The sparse matrix-vector problem is particularly tough for processors with limited memory bandwidth since there is




Fig. 5. Performance measurements for the nested and the flattened implementations of `smvp` on the SGI Origin 2000.

no temporal locality in the use of A (within a single matrix-vector product), and the locality in reference to x diminishes with increasing n . While reordering may mitigate these effects in some applications, it has little effect for the random matrices used here. The Origin 2000 implementations also do not exhibit good parallel scaling. This is likely a function of increased latency to remote memory with increasing processor count, coupled with limited latency-hiding capabilities in the pro-

cessors. Higher performance can be obtained with further tuning. For example, the current compiler does not perform optimizations to map the `val` and `col` components of A into separate arrays. When applied manually, this optimization increases performance by 25% or more.

On the SX-4 the flattened implementation performs significantly better than the nested implementation over all inputs. This is because the flattened implemen-




Fig. 6. Performance measurements for the nested and the flattened implementations of `smvvp` on the NEC SX-4. (Note the logarithmic scale of the y -axis.)

tation always operates on full-sized vectors (provided $r \geq pq$), while the nested implementation performs vector operations whose length is determined by the number of nonzeros in a row. Hence the nested implementation is insensitive to problem size but improves with average row length. For the flattened implementation, absolute performance and parallel scaling are good primarily because the memory system has

sufficient bandwidth and the full-sized vector operations fully amortize the memory access latencies.

5.3. Summary

This experiment provides evidence that the flattening technique can be used in an implementation to improve the performance stability for irregular problems

while maintaining or improving on the performance of the simple thread-based implementation. The flattening techniques may be particularly helpful in supporting the instruction-level and memory-level parallelism required for high performance in modern processors. The experiments also demonstrate that, for the matrices generated by the NAS benchmark, the scheduling strategies provided by the Fortran runtime system suffice for good load balance.

It is possible to construct sparse matrices with large variations in the number of nonzeros per row for which the dynamic thread scheduling techniques, in the simple form generated by Fortran compilers, can not solve the load imbalance problems. While these irregular matrices may not be representative of typical problems, the basic characteristic of large amounts of work in small portions of the iteration space is not unusual. For example, it can arise with data structures for the adaptive spatial decomposition of a highly clustered n -body problem, or with divide-and-conquer algorithms like quicksort or quickhull [6].

The flattened sparse matrix-vector product described in this section was incorporated into the NAS CG benchmark code to obtain the results quoted in the introduction for the NEC SX-4.

6. Related work

The facilities for data abstraction and dynamic aggregates are new in Fortran 9X. Previously, Norton et al. [28], Deczyk et al. [16], and Nyland et al. [29] have experimented with these advanced features of Fortran 90 to analyze their impact on performance.

HPF 2.0 provides a `MAPPED` irregular distribution to support irregular computations. This is a mechanism and not a policy, and leaves the user responsible for developing a coherent policy for its use. Further, the ramifications of this distribution on compilation are not yet fully resolved. Our approach is fundamentally different in attempting to support well a smaller class of computations with an identifiable policy (nested data parallelism) and by preprocessing the irregular computation to avoid reliance on untested strategies in the HPF compiler. While HPF focuses on the irregular distribution of regular data structures, our approach is based on the (regular) distribution of irregular data structures.

The FX programming model [17,42] integrates data and task parallelism using directives in a Fortran environment. The data-parallel features are similar to those of High Performance Fortran. The task parallelism di-

rectives are based on the notions of task partitions and processor subgroups, and thus represents a processor-centric approach. Recent versions of this model allow task regions to be dynamically nested, allowing the representation of algorithms such as quicksort and Barnes–Hut.

Split-C [15] also provides a number of low-level mechanisms for expressing irregular computations. We are attempting to provide a higher level of abstraction while providing the same level of execution efficiency of low-level models.

The Chaos library [35] is a runtime library based on the inspector/executor model of executing parallel loops involving irregular array references. It is a suitable back end for the features supporting irregular parallelism in HPF 2.0. The library does not provide obvious load balancing policies, particularly for irregularly nested parallel loops. Recent work on Chaos is looking at compilation aspects of irregular parallelism.

Flattening transformations have been implemented for the languages NESL [9], Proteus [33], Amelia [37], and V [13], differing considerably in their completeness and in the associated constant factors. There has been little work on the transformation of imperative constructs such as sequential loops within a `FORALL`, although there do not appear to be any immediate problems. The flattening techniques are responsible for several hidden successes. Various high performance implementations are really hand-flattened nested data-parallel programs: FMA [18], radix sort [44], as well as the NAS CG implementation described in the introduction. Furthermore, the set of primitives in HPFLIB itself reflects a growing awareness and acceptance of the utility of the flattening techniques.

The mainstream performance programming languages Fortran and SISAL [12,40] can express nested data parallelism, but currently do not address its efficient execution in a systematic way. Languages that do address this implementation currently have various disadvantages: they are not mainstream languages (NESL, Proteus); they subset or extend existing languages (Amelia, V, F90V); they do not interface well with regular computations (NESL, Proteus); they are not imperative, hence provide no control over memory (NESL, Proteus); and they are not tuned for performance at the level of Fortran (all).

7. Conclusions

Nested data parallelism in Fortran is attractive because Fortran is an established and important language

for high-performance parallel scientific computation and has an active community of users. Many of these users are now facing the problem of implementing irregular computations on parallel computers. They have substantial investments in existing code and depend on Fortran or HPF to achieve high performance on the regular portions of their computations. For them it is highly desirable to stay within the Fortran framework.

The advanced features of modern Fortran, such as derived data types, modules, pointers, and the `FORALL` construct, together constitute a sufficient mechanism to express complex irregular computations as nested data parallel computations. This makes it possible to express both irregular and regular computations within a common framework and in a familiar programming style.

How to achieve high performance from such high-level specifications of irregular computation is a more difficult question. The flattening technique can be effective for machines with very high and uniform shared-memory bandwidth, as that found in current parallel vector processors from NEC and SGI/Cray or the parallel multithreaded Tera machine [41]. For cache-based shared-memory processors, the improved locality of the threading approach is a better match. Here, the flattening techniques may help to extract threads from a nested parallel computation that, on the one hand, are sufficiently coarse grain to obtain good locality of reference and amortize scheduling overhead, and, on the other hand, are sufficiently numerous and regular in size to admit good load balance with run-time scheduling.

We have illustrated that irregular computations can be efficiently executed through a combination of source-to-source preprocessing, leveraging of the Fortran compilers, and runtime support. The source-to-source preprocessing is presently beyond the capabilities of automatic parallelization techniques and hence has been tediously performed manually [29,18]. However, we believe that compiler support for flattening and threading, possibly guided by additional directives, is achievable, and will provide a simple route to high-performance irregular computation.

Acknowledgments

This research was supported in part by NSF Grants #CCR-9711438 and #INT-9726317. Chatterjee is supported in part by NSF CAREER Award #CCR-9501979. Simons is on leave from TU Berlin and is

supported by a scholarship from the German Research Foundation (DFG). Use of the NEC SX-4 has been sponsored by the CSCS, Manno, Switzerland.

References

- [1] J.C. Adams, W.S. Brainerd, J.T. Martin, B.T. Smith and J.L. Wagener, *Fortran 95 Handbook – Complete ISO/ANSI Reference*, MIT Press, 1997.
- [2] P. Au, M. Chakravarty, J. Darlington, Y. Guo, S. Jähnichen, G. Keller, M.Köhler, M. Simons and W. Pfannenstiel, Enlarging the scope of vector-based computations: Extending Fortran 90 with nested data parallelism, in: *International Conference on Advances in Parallel and Distributed Computing*, W. Giloi, ed., IEEE, 1997, pp. 66–73.
- [3] D.H. Bailey, E. Barszcz, J.T. Barton, D.S. Browning, R.L. Carter, L. 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, *Internat. J. Supercomputer Appl.* **5**(3) (Fall 1991), 63–73.
- [4] J.E. Barnes and P. Hut, A hierarchical $O(N \log N)$ force calculation algorithm, *Nature* (1986), 446–449.
- [5] G.E. Blelloch, *Vector Models for Data-Parallel Computing*, The MIT Press, Cambridge, MA, 1990.
- [6] G.E. Blelloch, Programming parallel algorithms, *Commun. ACM* **39**(3) (March 1996), 85–97.
- [7] G.E. Blelloch, S. Chatterjee and M. Zagha, Solving linear recurrences with loop raking, *J. Parallel Distrib. Computing* **25**(1) (Feb. 1995), 91–97.
- [8] G.E. Blelloch, P.B. Gibbons and Y. Matias, Provably efficient scheduling for languages with fine-grained parallelism, in: *Proceedings of the 7th Annual ACM Symposium on Parallel Algorithms and Architectures*, Santa Barbara, CA, June 1995, pp. 1–12.
- [9] G.E. Blelloch, J.C. Hardwick, J. Sipelstein, M. Zagha and S. Chatterjee, Implementation of a portable nested data-parallel language, *J. Parallel Distrib. Computing* **21**(1) (Apr. 1994), 4–14.
- [10] G.E. Blelloch and G.W. Sabot, Compiling collection-oriented languages onto massively parallel computers, *J. Parallel Distrib. Computing* **8**(2) (Feb. 1990), 119–134.
- [11] R.D. Blumofe, C.F. Joerg, B.C. Kuszmaul, C.E. Leiserson, K.H. Randall and Y. Zhou, Cilk: An efficient multithreaded runtime system, in: *Proceedings of the Fifth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Santa Barbara, CA, July 1995, pp. 207–216.
- [12] D. Cann and J. Feo, SISAL versus FORTRAN: A comparison using the livermore loops, in: *Proceedings of Supercomputing '90*, New York, NY, Nov. 1990, pp. 626–636.
- [13] M.M.T. Chakravarty, F.-W. Schröer and M. Simons, V – Nested parallelism in C, in: *Programming Models for Massively Parallel Computers*, W.K. Giloi, S. Jähnichen and B.D. Shriver, eds., IEEE Computer Society, 1995, pp. 167–174.
- [14] S. Chatterjee, Compiling nested data-parallel programs for shared-memory multiprocessors, *ACM Trans. Prog. Lang. Syst.* **15**(3) (July 1993), 400–462.

- [15] D.E. Culler, A. Dusseau, S.C. Goldstein, A. Krishnamurthy, S. Lumetta, T. von Eicken and K. Yelick, Parallel programming in Split-C, in: *Proceedings of Supercomputing '93*, Nov. 1993, pp. 262–273.
- [16] V.K. Decyk, C.D. Norton and B.K. Szymanski, High performance object-oriented programming in Fortran 90, *ACM Fortran Forum* **16**(1) (Apr. 1997).
- [17] T. Gross, D. O'Hallaron and J. Subhlok, Task parallelism in a High Performance Fortran framework, in: *IEEE Parallel and Distributed Technology*, Vol. 2, 1994, pp. 16–26.
- [18] Y.C. Hu, S.L. Johnsson and S.-H. Teng, High Performance Fortran for highly irregular problems, in: *Proceedings of the Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Las Vegas, NV, June 1997, pp. 13–24.
- [19] G. Keller, Transformation-based implementation of nested data-parallelism for distributed memory machines, Dissertation, Technische Universität Berlin, 1999.
- [20] G. Keller and M.T. Chakravarty, Flattening trees, in: *Euro-Par'98, Parallel Processing*, D. Pritchard and J. Reeve, eds., Lecture Notes Comput. Sci., Vol. 1470, Springer, 1998, pp. 709–719.
- [21] C.H. Koelbel, D.B. Loveman, R.S. Schreiber, G.L. Steele Jr. and M.E. Zosel, *The High Performance Fortran Handbook*, Scientific and Engineering Computation, The MIT Press, Cambridge, MA, 1994.
- [22] C.L. Lawson, R.J. Hanson, D.R. Kincaid and F.T. Krogh, Basic linear algebra subprograms for Fortran usage, *ACM Trans. Math. Softw.* **5**(3) (Sept. 1979), 308–323.
- [23] J. McCalpin, Personal communication, Apr. 1998.
- [24] M. Metcalf and J. Reid, *Fortran 90/95 Explained*, Oxford Univ. Press, 1996.
- [25] T. More, The nested rectangular array as a model of data, in: *APL '79 Conference Proceedings*, ACM, 1979, pp. 55–73.
- [26] G.J. Narlikar, Pthreads for dynamic and irregular parallelism, in: *Supercomputing '98: High Performance Networking and Computing*, 1998.
- [27] G.J. Narlikar and G.E. Blelloch, Space-efficient implementation of nested parallelism, in: *Proceedings of the Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Las Vegas, NV, June 1997, pp. 25–36.
- [28] C.D. Norton, B.K. Szymanski and V.K. Decyk, Object-oriented parallel computation for plasma simulation, *Commun. ACM* **38**(10) (Oct. 1995), 88–100.
- [29] L.S. Nyland, S. Chatterjee and J.F. Prins, Parallel solutions to irregular problems using HPF, in: *First HPF UG Meeting*, Santa Fe, NM, Feb. 1997.
- [30] OpenMP Group, OpenMP: A proposed standard API for shared memory programming. White paper, OpenMP Architecture Review Board, Oct. 1997.
- [31] D.W. Palmer, J.F. Prins, S. Chatterjee and R.E. Faith, Piecewise execution of nested data-parallel programs, in: *Languages and Compilers for Parallel Computing*, C.-H. Huang, P. Sadayappan, U. Banerjee, D. Gelernter, A. Nicolau and D. Padua, eds., Lecture Notes Comput. Sci., Vol. 1033, Springer, 1996, pp. 346–361.
- [32] J. Prins, M. Ballabio, M. Boverat, M. Hodous and D. Maric, Fast primitives for irregular computations on the NEC SX-4, *Crosscuts* **6**(4) (1997), 6–10.
- [33] J.F. Prins and D.W. Palmer, Transforming high-level data-parallel programs into vector operations, in: *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, San Diego, CA, May 1993, pp. 119–128.
- [34] S. Saini and D. Bailey, NAS Parallel Benchmark (1.0) results 1-96, Technical Report NAS-96-018, NASA Ames Research Center, Moffett Field, CA, Nov. 1996.
- [35] J. Saltz, R. Ponnusammy, S. Sharma, B. Moon, Y.-S. Hwang, M. Uysal and R. Das, A manual for the CHAOS runtime library, Technical Report CS-TR-3437, Department of Computer Science, University of Maryland, College Park, MD, March 1995.
- [36] J.T. Schwartz, R.B.K. Dewar, E. Dubinsky and E. Schonberg, *Programming with Sets: An Introduction to SETL*, Springer, New York, 1986.
- [37] T.J. Sheffler and S. Chatterjee, An object-oriented approach to nested data parallelism, in: *Proceedings of the Fifth Symposium on the Frontiers of Massively Parallel Computation*, McLean, VA, Feb. 1995, pp. 203–210.
- [38] J.P. Singh, J.L. Hennessy and A. Gupta, Implications of hierarchical N -body methods for multiprocessor architecture, *ACM Transact. Computer Syst.* (2) (1995), 141–202.
- [39] J.M. Sipelstein and G.E. Blelloch, Collection-oriented languages, *Proceedings of the IEEE* **79**(4) (Apr. 1991), 504–523.
- [40] S.K. Skedzelewski, Sisal, in: *Parallel Functional Languages and Compilers*, B.K. Szymanski, ed., ACM Press, New York, NY, 1991, pp. 105–157.
- [41] A. Snively and L. Carter, Multi-processor performance on the Tera MTA, in: *Supercomputing '98 (SC'98)*, 1998.
- [42] J. Subhlok and B. Yang, A new model for integrated nested task and data parallel programming, in: *Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'97)*, 1997, pp. 1–12.
- [43] M.S. Warren and J.K. Salmon, A parallel hashed oct-tree N -body algorithm, in: *Supercomputing '93*, IEEE Computer Society, 1993, pp. 12–21.
- [44] M. Zagha and G.E. Blelloch, Radix sort for vector multiprocessors, in: *Proceedings of Supercomputing'91*, Albuquerque, NM, Nov. 1991, pp. 712–721.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

