

HPF-2 Support for Dynamic Sparse Computations^{*}

R. Asenjo¹, O. Plata¹, J. Touriño², R. Doallo², and E.L. Zapata¹

¹ Dept. Computer Architecture, University of Málaga, Spain
{`asenjo`,`oscar`,`ezapata`}@ac.uma.es

² Dept. Electronics and Systems, University of La Coruña, Spain
{`juan`,`doallo`}@udc.es

Abstract. There is a class of sparse matrix computations, such as direct solvers of systems of linear equations, that change the fill-in (nonzero entries) of the coefficient matrix, and involve row and column operations (pivoting). This paper addresses the problem of the parallelization of these sparse computations from the point of view of the parallel language and the compiler. Dynamic data structures for sparse matrix storage are analyzed, permitting to efficiently deal with fill-in and pivoting issues. Any of the data representations considered enforces the handling of indirections for data accesses, pointer referencing and dynamic data creation. All of these elements go beyond current data-parallel compilation technology. We propose a small set of new extensions to HPF-2 to parallelize these codes, supporting part of the new capabilities on a runtime library. This approach has been evaluated on a Cray T3E, implementing, in particular, the sparse LU factorization.

1 Introduction

Irregular computations, where data-access patterns and workload are not known at compile time, appear profusely on scientific and engineering applications. An approach to handle such computations is based on extending a data-parallel language with new constructs suitable to express non-structured parallelism. With this information, the compiler can perform at compile time a number of optimizations, usually embedding the rest of them into a runtime library. In Fortran D [12], the programmer can specify a mapping of array elements to processors using another array. Vienna-Fortran [22] lets programmers define functions to specify irregular distributions. HPF-2 [15,16] provides a generalized block distribution (GEN-BLOCK), where the contiguous array partitions may be of different sizes,

^{*} This work was supported by the Ministry of Education and Science (CICYT) of Spain (TIC96-1125-C03), by the Xunta de Galicia (XUGA20605B96), by the European Union (BRITE-EURAM III BE95-1564), by the Human Capital and Mobility programme of the European Union (ERB4050P1921660), and by the Training and Research on Advanced Computing Systems (TRACS) at the Edinburgh Parallel Computing Centre (EPCC)

and an indirect distribution (INDIRECT), where a mapping array is defined to specify an arbitrary assignment of array elements to processors.

A different approach is based on runtime techniques, that is, the non-structured parallelism is captured and managed fully at runtime. These techniques automatically manage programmer-defined data distributions, partition loop iterations, remap data and generate optimized communication schedules. Most of these solutions are based on the inspector-executor paradigm [18,7].

Current language constructs and the supportive runtime libraries are insufficiently developed, leading to low efficiencies when they are applied to a wide set of irregular codes. In the context of sparse computations, we found useful to inform the compiler not only about the data distribution, but also about how these data are stored in memory. We will call *distribution scheme* the combination of these two aspects (data structure + data distribution). We have developed and extensively tested a number of pseudo-regular distribution schemes for sparse problems, which combines natural extensions of regular data distributions with compressed data storages [2,19,20,21]. These distribution schemes can be incorporated to a data-parallel language (HPF) in a simple way.

The above mentioned distribution schemes are faced to static sparse problems. In this paper we discuss data structures and distributions in the context of dynamic sparse matrix computations, involving fill-in and pivoting operations. Direct methods for solving sparse systems of linear equations, for instance, present this kind of computations. Factorization of the coefficient matrix may produce new nonzero values (fill-in), so that data structures must consider the inclusion of new elements at runtime. Also, row and/or column permutations of the coefficient matrix are usually accomplished to assure numerical stability and limit fill-in. All these features make such sparse computations hard to parallelize.

The rest of the paper is organized as follows. Section 2 discusses the dynamic data distributions schemes we have tested to implement efficient parallel sparse codes involving pivoting and fill-in. Specifically, a direct method for the LU factorization is taken as a working example. Section 3 describes our proposal to extend HPF-2 for considering the above dynamic distributions. Experimental results validating our approach are presented in Section 4.

2 Sparse Data Structures and Distributions

2.1 Sparse Data Structures

Two different approaches may be considered to represent a sparse matrix: static and dynamic data structures. Static data structures are the most used in Fortran codes. Common examples are Compressed Row and Column Storages (CRS and CCS) [4]. If the computation includes fill-in and/or pivoting operations, it may be preferably to use some more complex and flexible data structures (dynamic). We have experimented with linked lists and hybrid (semi-)dynamic data structures, depending on the type of data accesses we have to deal with.

To simplify the discussion, hereafter we will consider as a working example the LU factorization of a sparse matrix, computed using a general method [1,10].

```

do k = 1, n
  Find pivot = Aij
  if (i ≠ k)
    swap A(k, 1 : n) and A(i, 1 : n)
  endif
  if (j ≠ k)
    swap A(1 : n, k) and A(1 : n, j)
  endif
  A(k + 1 : n, k) = A(k + 1 : n, k) / A(k, k)
  do j = k + 1, n
    do i = k + 1, n
      A(i, j) = A(i, j) - A(i, k)A(k, j)
    enddo
  enddo
enddo
    
```

Fig. 1. LU algorithm (General approach, right-looking version)

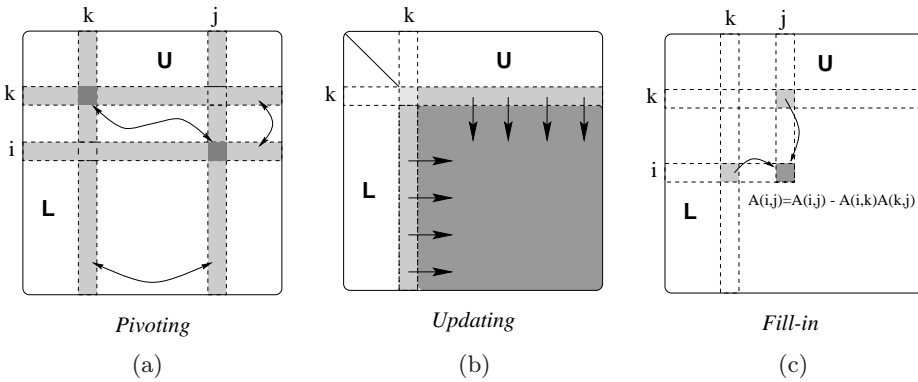


Fig. 2. Pivoting (a) and updating (b) operations, and fill-in (c) in right-looking LU

Fig. 1 shows an in-place code for the direct right-looking LU algorithm, where an n -by- n matrix A is factorized. The code includes a row and column pivoting operation (full pivoting) to provide numerical stability and preserve sparsity. Fig. 2 depicts the access patterns for the pivoting and updating operations on both matrices, L and U , and the generation of new entries. Note that efficient data accesses both by rows and columns are required.

The sparse coefficient matrix may be structured as a two-dimensional doubly linked list (see Fig. 3 (c)), to make efficient data accesses both by rows and columns. Each item in such a dynamic structure stores not only the value and the local row and column indices, but also pointers to the previous and next nonzero element in its row and column.

The complexity of this list can be reduced if full pivoting is replaced by partial pivoting, where only columns (or rows) are swapped. This may imply large memory and computation savings as we can use a simple list of packed vectors, or a one-dimensional doubly linked list structure, to store the sparse

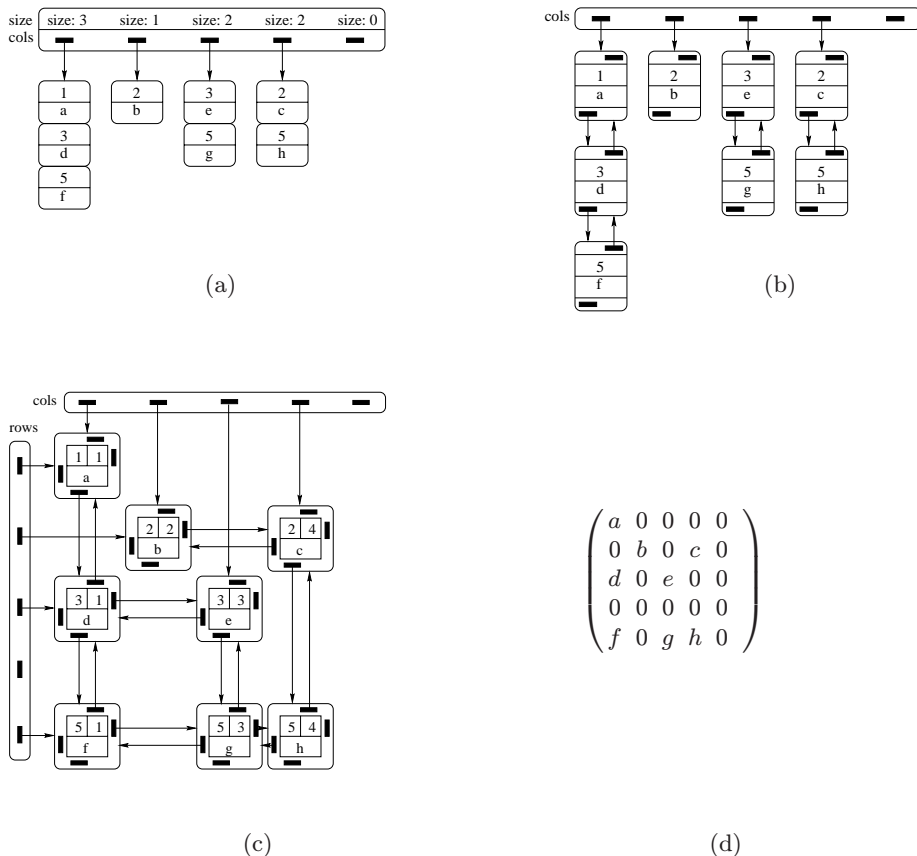


Fig. 3. Packed vectors and linked lists as efficient data structures for direct methods: (a) List of packed vectors; (b) one-dimensional doubly linked list; (c) two-dimensional doubly linked list; (d) local sparse matrix

matrix. As depicted in Fig. 3 (b), each linked list represents one column of the sparse matrix, where its nonzero entries are arranged in growing order of the row index. Each item of the list stores the row index, the matrix entry and two pointers. A simplification of the linked list is shown in Fig. 3 (a), where columns are stored as packed vectors, and they are referenced by means of an array of pointers. The list of packed vectors do not have pointers and, therefore, this mixed structure requires much less memory space than the doubly linked list.

Compressed formats and lists of packed vectors are very compact and allow fast accesses by rows or by columns to the matrix entries (but not both at the same time). Linked lists are useful when more flexible data accesses are needed. Two-dimensional lists, for instance, allow accesses to both rows and columns with the same overhead. The fill-in and pivoting issues are easily managed when doubly linked lists are used, as they make easy the entry insertion and dele-

```

!Doubly LLRS, LLCS (one-dimensional)
TYPE entry
  INTEGER:: index
  REAL:: value
  TYPE (entry), POINTER:: prev, next
END TYPE entry

!Doubly LLRCS (two-dimensional)
TYPE entry
  INTEGER:: indexi, indexj
  REAL:: value
  TYPE (entry), POINTER:: previ, prevj
  TYPE (entry), POINTER:: nexti, nextj
END TYPE entry

```

Fig. 4. Fortran 90 derived types for the items of LLRS, LLCS and LLRCS storage schemes, and a definition of an array of pointers (pex) to these items

tion operations. In the case of compressed formats (CRS, CCS ...) or a list of packed vectors, the fill-in problem is more difficult to solve. Compressed formats also have the inconvenience of not allowing the pivoting operation (column/row swapping) in an easy way. This can be overcome by using some mixed data structure, such as the list of packed vectors, or a linked list structure. Column pivoting is then implemented by just interchanging pointer values.

Albeit their flexibility, linked lists may have some drawbacks. The dynamic memory allocation for each new entry, as well as the list traversing, are time-consuming operations. Additionally, they consume more space memory than packed vectors. Finally, memory fragmentation due to allocation/deallocation of items may arise, as well as spatial data locality loss.

2.2 Dynamic Sparse Distribution Schemes

Four data storage schemes will be considered: LLCS (Linked List Column Storage), LLRS (Linked List Row Storage), LLRCS (Linked List Row-Column Storage) and CVS (Compressed Vector Storage), the first three schemes to represent sparse matrices, and the last one to represent sparse (one-dimensional) arrays. The LLCS storage scheme corresponds to the structure shown in Fig. 3 (b). In this figure the lists are doubly linked, but it can also be defined as singly linked, in order to save memory overhead. The LLRS storage scheme is similar to the LLCS scheme but considering linking by rows instead of columns. A combination of compressed columns and rows representation, interlinked among themselves, can be declared using the LLRCS storage scheme, as shown in Fig. 3 (c). As well as with the other two schemes, the entries can be singly or doubly linked. Finally, CVS scheme represents a sparse vector as two arrays and one scalar: the index array, containing the indices of the nonzero entries of the sparse array; the value array, containing the nonzero entries themselves; and the size scalar, containing the number of nonzero entries.

Fig. 4 displays the Fortran 90 derived types which may define the items of each kind of linked list. The first type corresponds to the LLRS and LLCS schemes

```

<sparse-directive> ::= <datatype>, SPARSE (<sparse-content>) :: <array-objects>
<datatype> ::= REAL | INTEGER
<sparse-content> ::= LLRS (<ll-spec>)
                    | LLCS (<ll-spec>)
                    | LLRCS (<ll2-spec>)
                    | CVS (<cv-spec>)
<ll-spec> ::= <pointer-array-name>, <pointer-array-name>,
             <size-array-name>,
             <link-spec>
<ll2-spec> ::= <pointer-array-name>, <pointer-array-name>,
             <pointer-array-name>, <pointer-array-name>,
             <size-array-name>, <size-array-name>,
             <link-spec>
<cv-spec> ::= <index-array-name>, <value-array-name>, <size-scalar-name>
<link-spec> ::= SINGLY | DOUBLY
<array-objects> ::= <sized-array>{,<sized-array>}
<sized-array> ::= <array-name>(<subscript>[,<subscript>])

```

Fig. 5. Syntax for the proposed HPF-2 SPARSE directive with dynamic data structures

(doubly linked), indistinctly, and the second one to the LLRCS scheme, doubly linked. The singly linked versions for these data types are equivalent but without the *prev* pointers. The list itself is declared also through a derived type, *pex*, which defines an array (or two) of pointers to the above items.

Once storage schemes have been defined, we can use the SPARSE directive to specify that a sparse matrix (or sparse array) is stored using a particular linked list scheme. This directive was previously introduced [2,19] in the context of static sparse applications. Fig. 5 shows the BNF syntax for the dynamic SPARSE directive. The first two data structures, LLRS and LLCS, are defined by two arrays of pointers (<pointer-array-name>), which point to the beginning and to the end, respectively, of each row (or column) list, and a third array (<size-array-name>), containing the number of elements per row (for LLRS) or per column (for LLCS). The option <link-spec> specifies the type of linking of the list data structure (singly or doubly). Regarding the LLRCS data structure, we have four arrays of pointers which point to the beginning and to the end of each row and each column of the sparse matrix, and two additional arrays storing the number of elements per row and per column, respectively.

As an example, the following statement,

```
!HPF$ REAL, DYNAMIC, SPARSE (CVS(vi, vv, sz)):: V(10)
```

declares *V* as a sparse vector compressed using the CVS format. *V* will work in the code as a place holder of the sparse vector, which occupies no storage. What is really stored are the nonzero entries of the sparse array in *vv*, the corresponding array indices in *vi*, and the number of nonzero entries in *sz*. The place holder *V* actually provides an abstract object with which other data objects can be aligned and which can then be distributed. The DYNAMIC keyword means that the contents of the three arrays, *vi*, *vv* and *sz*, are determined dynamically, as a result of executing a DISTRIBUTE statement.

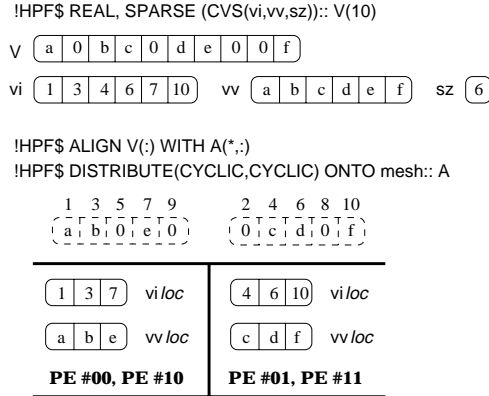


Fig. 6. Alignment and distribution of a sparse array on a 2x2 processor mesh

The HPF directives `DISTRIBUTE` and `ALIGN` can be applied to sparse place holders with the same syntax as in the standard. Distributing a sparse place holder is equivalent to distributing it as if it was a dense array (matrix). For instance, the statement,

```
!HPF$ DISTRIBUTE(CYCLIC) ONTO mesh:: V
```

considers `V` as a dense array (not compressed), mapping this array on the processors using the standard `CYCLIC` data distribution, and representing the distributed (local) sparse arrays using the `CVS` compressed format.

In the case of the `ALIGN` directive, however, the semantics is slightly different. From the next example code,

```

REAL, DIMENSION(10,10):: A
INTEGER, DIMENSION(10):: vi
REAL, DIMENSION(10):: vv
INTEGER:: sz
!HPF$ PROCESSORS, DIMENSION(2,2):: mesh
!HPF$ REAL, DYNAMIC, SPARSE (CVS(vi, vv, sz)):: V(10)
!HPF$ ALIGN V(:) WITH A(*,:)
!HPF$ DISTRIBUTE(CYCLIC,CYCLIC) ONTO mesh:: A
  
```

the (nonzero) entries of `V` (that is, `vv`) are aligned with the columns of `A` depending on the positions stored in the array `vi`, and not in the corresponding positions in their own `vv` array (which is the standard semantics). Now, the `DISTRIBUTE` directive replicates the `V` array over the first dimension of the processor array `mesh`, and distributes it over the second dimension in the same way as the second dimension of the `A` matrix. Observe that in this distribution operation, `vi` is taken as the index array for the entries stored in `vv`. Fig. 6 shows the combined effect of alignment/distribution for a particular case.

The combination of the directives `SPARSE` and `DISTRIBUTE` defines the distribution scheme of a sparse matrix. The variable `V` in the example code above

```

INTEGER, PARAMETER:: n=1000, dim=8
INTEGER:: k, i, j
REAL:: maxpiv, pivot, amul, product
INTEGER:: actpiv, pivcol
TYPE (entry), POINTER:: aux

TYPE (ptr), DIMENSION(n):: first, last, vpiv
INTEGER, DIMENSION(n):: vsize

REAL, DIMENSION(n):: vcolv, vmaxval
INTEGER, DIMENSION(n):: vcoli
INTEGER:: size

!HPF$ PROCESSORS, DIMENSION(dim):: linear
!HPF$ REAL, DYNAMIC, SPARSE(LLCS(first, last, vsize, DOUBLY)):: A(n,n)
!HPF$ REAL, DYNAMIC, SPARSE(CVS(vcoli, vcolv, size)):: VCOL(n)
!HPF$ ALIGN iq(:) WITH A(*,:)
!HPF$ ALIGN vpiv(:) WITH A(*,:)
!HPF$ ALIGN vmaxval(:) WITH A(*,:)
!HPF$ ALIGN VCOL(:) WITH A(:,*)
!HPF$ DISTRIBUTE (*,CYCLIC) ONTO linear:: A

```

Fig. 7. Declarative section of the extended HPF-2 parallel sparse LU code

really works as a place holder for the sparse array. The `SPARSE` directive establishes the connection between the logical entity `V` and its actual representation (compressed format). The benefit of this approach is that we can use the standard HPF `DISTRIBUTE` and `ALIGN` directives applied to the array `V` and, at the same time, store the array itself using a compressed format. In the rest of the code, the sparse matrix is operated using directly its compressed format.

3 Parallel Dynamic Sparse Computations

The `SPARSE` directive establishes a link between the sparse matrix (or array) and its storage structure. From this point on, we can choose to hide the storage scheme to programmers, and allow them to write the parallel sparse code using dense matrix notations. The compiler will be in charge of translating these dense notations into parallel sparse codes taking into account the storage schemes specified. However, this approach supposes a great effort in compiler implementation, as well as the possibility of mixing in the same code place holders (dense notations) with *real* arrays. Bik and Wijshoff [6] and Kotlyar and Pingaly [17] propose a similar approach, based on the automatic transformation of a dense program, annotated with sparse directives, into a semantically equivalent sparse code. The design of such compiler is, however, very complex, in such a way that no implementation of it is available for general and real problems.

A different approach is based on forcing programmers to use explicitly the compressed storage structures common in sparse codes, and allow them to use the place holders (dense notations) only for aligning and distributing purposes. Parallelism is constrained to the directives. If the parallel code is sequentially compiled, the resulting code would run properly.

3.1 Parallel Sparse LU Code

A direct right-looking LU factorization with partial pivoting (column swapping) will be considered in this section. In most cases, partial pivoting leads to similar

numerical error results than full pivoting, but at a lower cost. However, a matrix reordering stage (*analyze* stage) should be added before the factorization. This is in charge of updating the permutation vectors so as sparsity and numerical stability are preserved in the subsequent factorization stage. A partial numerical pivoting is however retained in the factorization stage to cover the case that the selected pivot in the analyze stage turns to be unstable during factorization.

Despite pivoting, the sparsity of the matrix usually decreases during the factorization. In such case, a switch to a dense LU factorization may be advantageous at some point of the computation. This dense code is based on Level 2 BLAS, and includes numerical partial pivoting. At the switch point, the reduced sparse submatrix is scattered to a dense array. The overhead of the switch operation is negligible (as the analyze stage) and the reduced dense submatrix appears distributed in a regular cyclic manner. A 15% sparsity threshold value was used in our experiments to switch from the sparse to the dense code.

Fig. 7 shows the declarative section of the parallel sparse LU code, using the proposed extensions to HPF-2. Matrix **A** is defined as sparse and stored using the LLCS data structure. The arrays of pointers **first** and **last** indicate the first and the last nonzero entry, respectively, of each column of **A**. The array **vsize** stores the number of nonzero entries on each column of **A**. The sparse array **VCOL** is also defined, stored using the CVS format. This array contains the normalized pivot column of **A**, calculated in each outer iteration of the algorithm.

The last sentence in the declaration section distributes the columns of the sparse matrix **A** cyclically over a one-dimensional arrangement of abstract processors (the one-dimensional characteristic is not essential). Previously, three dense arrays, **iq**, **vpiv** and **vmaxval**, were aligned with the columns of **A**, while **VCOL** was aligned with the rows of **A**. Hence, after distributing **A**, **VCOL** is replicated over all the processors. At each iteration of the main loop of the algorithm (loop *k* in Fig. 1), the owner of the column *k* of **A** selects and updates on **VCOL** the pivot column, which is consistently broadcast to the rest of processors to enable the subsequent parallel submatrix update. Fig. 8 shows an example of this declaration.

Fig. 9 presents the rest of the parallel LU code. The first action corresponds to the initialization of the array **vpiv**, which should point to the row that includes the pivot. This loop is parallel and no communications are required, as both arrays, **vpiv** and **first**, were aligned. Next, loop *k* starts. **SwitchIter** was calculated by the analyze stage, value from which the sparse code switches to an equivalent dense one.

The first action inside the main loop corresponds to column pivoting pivoting, in which we look for a stable pivot and, if possible, in agreement with the recommended permutation vector *iq* (obtained in the analyze stage). To fulfill the first condition, the pivot should be greater than the maximum absolute value of the pivot row times an input parameter called *u* ($0 \leq u \leq 1$). The maximum absolute value is calculated using the Fortran 90 **MAXVAL()** intrinsic function, evaluated over **vmaxval** vector. The update of **vmaxval** takes place on the second **INDEPENDENT** loop which traverses the pivot row storing the absolute

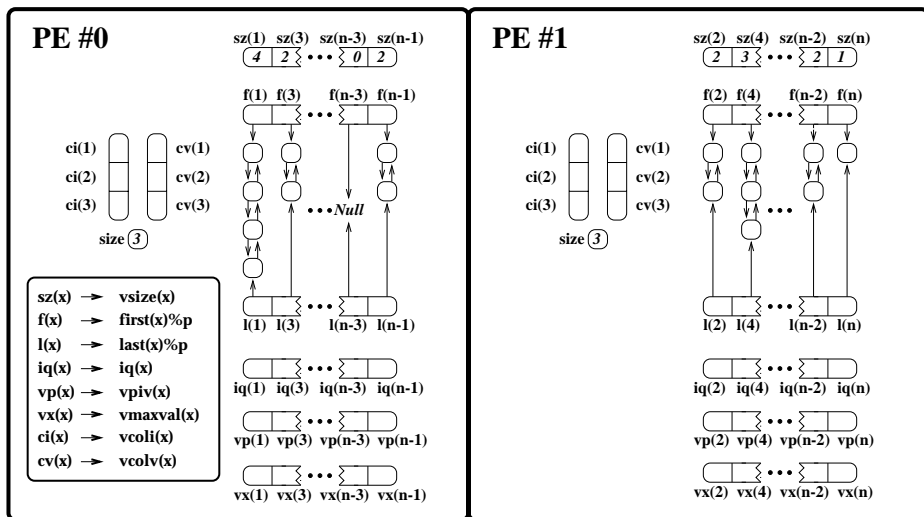


Fig. 8. Partitioning of most LU arrays/matrices on two processors, according to the HPF declaration of Fig. 7 (an even number of columns for A, and that the outer loop of the LU algorithm is in the fourth iteration, are assumed)

value of each entry on `vmaxval`. These entries are candidates for pivot. The `ON HOME (vpiv(j))` directive tells the compiler that the processor owning `vpiv(j)` will be encharged of iteration `j`. The `RESIDENT` annotation points out to the compiler that all variables referenced inside the directive's body are local. Thus, the compiler analysis is simplified and more optimized code may be generated.

Once the threshold `maxpiv` is obtained, the pivot is chosen in such a way that its value is greater than the above threshold, and, on the other hand, sparsity is preserved by following the `iq` recommendations. This computation is, in fact, a reduction operation, and consequently we annotate the corresponding `INDEPENDENT` loop with such directive. This user-defined reduction operation is indeed not considered by the HPF-2 standard, but its inclusion would not add any significant complexity to the compiler implementation. Finally, after selecting the pivot, the `swap()` routine is called to perform the permutation of the current column `k` and the pivot column of matrix `A`.

After the pivoting operation, the pivot column is updated and packed into the sparse `VCOL` array. This is computed by the owner of such column (`ON HOME` directive). As `VCOL` is a replicated array, any update made on it is communicated to the rest of processors. Finally, the submatrix $(k + 1 : n, k + 1 : n)$ of `A` is updated. Loop `j` runs over the columns of the matrix, and it is parallel. The `NEW` directive prevents the compiler from considering inexistent data dependences due to variables that are actually private to each iteration.

The code also contains the user-defined routines `append()` and `insert()` for list management, which are included in a Fortran 90 module. The `append()`

```

! --> Initialization
!HPF$ INDEPENDENT
DO j = 1, n
    vpiv(j)%p => first(j)%p
END DO

! --> Main loop LU
main: DO k = 1, SwitchIter

! --> Pivoting
! --> Candidates for pivot are selected and ...
!HPF$ INDEPENDENT
DO j = k, n
!HPF$ ON HOME (vpiv(j)), RESIDENT BEGIN
    IF (.NOT.ASSOCIATED(vpiv(j)%p)) CYCLE
    IF (vpiv(j)%p%index /= k) CYCLE
    vmaxval(j) = ABS(vpiv(j)%p%value)
!HPF$ END ON
END DO

! --> ... the maximum value is calculated
maxpiv = MAXVAL(vmaxval(k:n))
maxpiv = maxpiv*u

! --> The pivot is chosen from the candidates
! --> (reduction operation)
actpiv = 0
pivcol = 0
!HPF$ INDEPENDENT, REDUCTION(actpiv,pivcol)
DO j = k, n
    IF (vmaxval(j) > maxpiv .AND.
        iq(pivcol) > iq(j)) THEN
        actpiv = vmaxval(j)
        pivcol = j
    END IF
END DO
IF(pivcol == 0) pivcol=k
IF(pivcol /= k) THEN
! ----> Columns are swapped
CALL swap(k,pivcol,first,last,vpiv,vsize,iq)
END IF

! --> Pivot column is updated and packed
!HPF ON HOME (vpiv(k)), RESIDENT BEGIN
    aux => vpiv(k)%p
    pivot = 1/(aux%value)
    aux%value = pivot
    aux => aux%next
    size = vsize(k)-1

    DO i = 1, size
        aux%value = aux%value*pivot
        vcolv(i) = aux%value
        vcoli(i) = aux%index
        aux => aux%next
    END DO
!HPF END ON

```

⋮

Fig. 9. Outline of an extended HPF-2 specification of the parallel right-looking partial pivoting LU algorithm (first part)

routine adds an entry at the end of a list, while the `insert()` routine adds an element at the beginning or in the middle of a list.

4 Evaluating Results

A parallel sparse right-looking partial pivoting LU algorithm was implemented using the Cray T3E Fortran 90 and the SHMEM library [5]. The columns of

⋮

```

! --> Submatrix of A is Updated
!HPF$ INDEPENDENT, NEW (aux,i,amul,product)
loopj: DO j = k+1, n
!HPF$ ON HOME (vpiv(j)), RESIDENT BEGIN
    aux => vpiv(j)%p
    IF (.NOT.ASSOCIATED(aux)) CYCLE
    IF (aux%index /= k) CYCLE
    amul = aux%value
    vsize(j) = vsize(j)-1
    vpiv(j)%p => aux%next
    aux => aux%next
loopi: DO i = 1, size
    product = -amul*vcolv(i)
    DO
        IF (.NOT.ASSOCIATED(aux)) EXIT
        IF (aux%index >= vcoli(i)) EXIT
        aux => aux%next
    END DO
outer_if: IF (ASSOCIATED(aux)) THEN
    IF (aux%index == vcoli(i)) THEN
        aux%value = aux%value + product
    ELSE
! ----> First or middle position insertion
        CALL insert(aux,vcoli(i),product,first(j)%p,
            vsize(j))
        IF (vpiv(j)%p%index >= aux%prev%index)
            vpiv(j)%p => aux%prev
        END IF
    ELSE outer_if
! ----> End position insertion
        CALL append(vcoli(i),product,first(j)%p,last(j)%p,
            vsize(j))
        IF (.NOT.ASSOCIATED(vpiv(j)%p)) vpiv(j)%p => last(j)%p
    END IF outer_if
    END DO loopi
!HPF$ END ON
    END DO loopj
END DO main

```

Fig.9 (cont.). Outline for an extended HPF-2 specification of the parallel right-looking partial pivoting LU algorithm (last part)

the sparse matrix A were cyclically distributed over the processors (linearly arranged), and stored in the local memories using one-dimensional doubly linked lists. This parallel algorithm is similar to the sequential version, but with local indices instead of the global ones, and Cray SHMEM routines performing communication/synchronization operations. All these operations were encapsulated into calls to the DDLY (Data Distribution Layer) runtime library [20]. The parallel code was designed in such a way that it could be the output of a hypothetical extended HPF-2 compiler (extended with the directives for the proposed distribution schemes). That is, it should be not considered as an optimized hand-coded program.

Fig. 10 shows execution times and speed-up for the parallel LU algorithm. Test sparse matrices were taken from the Harwell-Boeing suite and University of Florida Sparse Matrix Collection [8] (see Table 1). The efficiency of the parallel code is high when the size of the input matrix is significantly large. We also carried out experiments considering meshes of processors instead of linear arrays, but the best times were obtained in the latter case and when the matrices were distributed by columns. Load imbalances due to fill-in (cyclic distribution) were not a problem for any matrix (see Fig. 11).

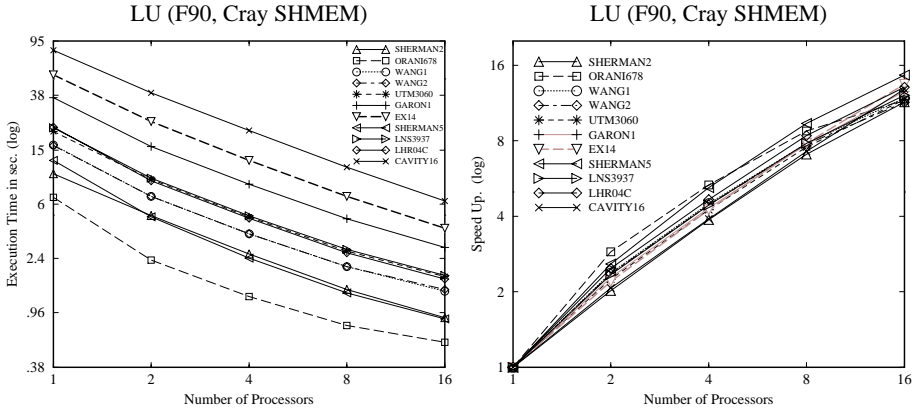


Fig. 10. Parallel sparse LU execution times and speed-up for different sparse matrices, using F90 linked lists and Cray T3E SHMEM

Table 1. Harwell-Boeing and Univ. of Florida test matrices

Matrix	Origin	n	# entries	sparsity
STEAM2	Oil reservoir simulation	600	13760	3.82%
JPWH991	Circuit physics modeling	991	6027	0.61%
SHERMAN1	Oil reservoir modeling	1000	3750	0.37%
SHERMAN2	Oil reservoir modeling	1080	23094	1.98%
ORANI678	Economic modeling	2529	90158	1.41%
WANG1	Discretized electron continuity	2903	19093	0.22%
WANG2	Discretized electron continuity	2903	19093	0.22%
UTM3060	Uedge test matrix	3060	42211	0.45%
GARON1	2D FEM, Navier-Stokes, CFD	3175	88927	0.88%
EX14	2D isothermal seepage flow	3251	66775	0.63%
SHERMAN5	Oil reservoir modeling	3312	20793	0.19%
LNS3937	Compressible fluid flow	3937	25407	0.16%
LHR04C	Light hydrocarbon recovery	4101	82682	0.49%
CAVITY16	Driven cavity problem	4562	138187	0.66%

The sequential efficiency of the Fortran 90 implementation of the sparse LU algorithm was also tested. Table 2 presents comparison results from this implementation and the Fortran 77 MA48 routine [10]. We observe that the MA48 routine is significantly faster than our algorithm for many matrices, but it should be considered the fact that the Cray Fortran 90 compiler is not efficient generating code for managing lists. However, the resulting computing errors are practically the same for both algorithms. The main advantage of our approach is its ease to be parallelized, as opposite to the MA48 routine, which is inherently sequential, as corresponds to a left-looking algorithm.

The analyze and solve (forward and backward substitution) stages of the LU algorithm were also implemented, but they are not presented here as no

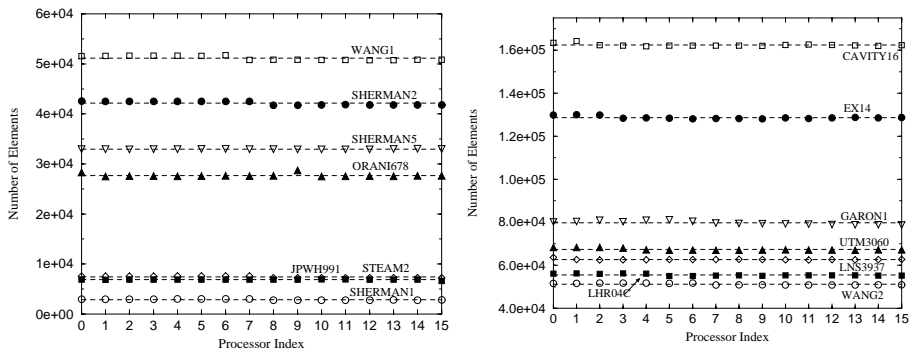


Fig. 11. Workload (non-null matrix values) on each processor after executing the parallel sparse LU factorization on a 16-processor system

Table 2. Comparison between Fortran 90 LU and MA48 (times in sec.)

Matrix	Times		Errors	
	F90	MA48	ratio	F90 - MA48
STEAM2	.572	.373	1.53	.15E-11 - .13E-11
JPWH991	1.039	.563	1.84	.44E-13 - .82E-13
SHERMAN1	.574	.148	3.87	.14E-12 - .16E-12
SHERMAN2	10.05	9.77	1.02	.14E-05 - .15E-05
ORANI678	6.74	3.52	1.91	.70E-13 - .74E-13
WANG1	16.33	18.76	0.87	.11E-12 - .97E-13
WANG2	16.19	16.54	0.97	.53E-13 - .52E-13
UTM3060	20.57	22.68	0.90	.53E-08 - .58E-08
GARON1	36.39	28.80	1.26	.53E-09 - .21E-08
EX14	53.68	62.78	0.85	.28E+01 - .93E+01
SHERMAN5	12.58	6.09	2.06	.75E-12 - .59E-12
LNS3937	21.88	15.09	1.44	.15E-02 - .13E-02
LHR04C	22.15	10.42	2.12	.22E-03 - .10E-03
CAVITY16	81.23	88.48	0.91	.39E-09 - .49E-09

additional relevant aspect is contributed. Both execution time and fill-in are comparable with those of the MA48 routine (they do not differ more than 10%).

5 Related Work

There are many parallel sparse LU factorization designs in the literature. From the loop-level parallelism point of view, the parallel pivot approach allows the extraction of an additional parallelism due to the sparsity of the matrix, besides the obvious one coming from the independences on the loops traversing rows and columns [3]. Coarser parallelism level can be exploited thanks to the elimination tree, which can be used to schedule parallel tasks in a multifrontal [11] code. It is also possible to use a coarse matrix decomposition to obtain an ordering to

bordered block triangular form, as is done in the MCSPARSE package [14]. The supernodal [9] approach is also a parallelizable code [13].

Some of the above parallel solutions can be implemented using the approach described in this paper. Loop-level LU approaches can be implemented using the LLRCS data storage (in addition to LLCs), and some other complex reductions to choose a good parallel pivot set, but loosing some of the performance due to the semi-automatic implementation. The multifrontal approach, however, is not suitable to the linked list sparse directive, due to the use of different data storage schemes. However, they could be implemented using the basic BCS or BRS sparse distributions [2,19]. The implementation of the supernodal code in [9] uses some sort of column compressed storage, but it would be necessary to simplify the memory management and the data access patterns to consider a data-parallel implementation of this code.

6 Conclusions

This paper presented a solution to the parallelization of dynamic sparse matrix computations (applications suffering from fill-in and/or involving pivoting operations) in a HPF-2 environment. The programmer is allowed to specify a particular sparse data storage representation, in addition to a standard data distribution. Sparse computations are specified by means of the storage representation constructs, while the (dense) matrix notation is reserved to declare alignments and distributions. Our experiments (a parallel sparse direct LU solver, *emulating* the output of an extended HPF-2 compiler) show that we can obtain high efficiencies using that strategy.

The research discussed in this paper gives new in-depth understanding in the semi-automatic parallelization of irregular codes dealing with dynamic data structures (list based), in such a way that the parallel code becomes a generalization of the original sequential code. An efficient parallel sparse code can be obtained by annotating the corresponding sequential version with a few number of HPF-like directives. The techniques described in this paper are not only useful to deal with the fill-in and pivoting problems, but they can also be applied to many other applications where the same or similar data structures are in use.

Acknowledgements

We gratefully thank Iain Duff and all members in the parallel algorithm team at CERFACS, Toulouse (France), for their kindly help and collaboration. We also thank the CIEMAT (Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas), Spain, for giving us access to the Cray T3E multiprocessor.

References

1. R. Asenjo. Sparse LU Factorization in Multiprocessors. Ph.D. Dissertation, Dept. Computer Architecture, Univ. of Málaga, Spain, 1997. 231

2. R. Asenjo, L.F. Romero, M. Ujaldón and E.L. Zapata. Sparse Block and Cyclic Data Distributions for Matrix Computations. in *NATO Adv. Res. Works. on High Performance Computing: Technology, Methods and Applications*, Cetraro, Italy, 1994. (Elsevier Science B.V., The Netherlands, pp. 359–377, 1995). 231, 235, 244
3. R. Asenjo and E.L. Zapata. Sparse LU Factorization on the Cray T3D. *Int'l. Conf. on High-Performance Computing and Networking (HPCN'95)*, Milan, Italy, pp. 690–696, 1995. 243
4. R. Barret, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine and H. van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. Siam Press, 1994. 231
5. R. Barriuso, A. Knies. SHMEM User's Guide for Fortran, Rev. 2.2. Cray Research, Inc, 1994. 240
6. A. Bik. Compiler Support for Sparse Matrix Computations. Ph.D. Dissertation, University of Leiden, The Netherlands, 1996. 237
7. P. Brezany, K. Sanjari, O. Cheron and E. Van Konijnenburg. Processing Irregular Codes Containing Arrays with Multi-Dimensional Distributions by the PRE-PARE HPF Compiler. *Int'l. Conf. on High-Performance Computing and Networking (HPCN'95)*, Milan, Italy, pp. 526–531, 1995. 231
8. T. Davis, University of Florida Sparse Matrix Collection. *NA Digest*, 92(42), 1994, 96(28), 1996, 97(23), 1997. See <http://www.cise.ufl.edu/~davis/sparse/>. 241
9. J.W. Demmel, S.C. Eisenstat, J.R. Gilbert, X.S. Li and J.W.H. Liu. A Supernodal Approach to Sparse Partial Pivoting. Tech. Report UCB/CSD-95-883, Computer Science Division, Univ. of California at Berkeley, CA, 1995. 244, 244
10. I.S. Duff and J.K. Reid. MA48, a Fortran Code for Direct Solution of Sparse Unsymmetric Linear Systems of Equations. Tech. Report RAL-93-072, Rutherford Appleton Lab., UK, 1993. 231, 242
11. I.S. Duff and J.A. Scott. The Design of a New Frontal Code of Solving Sparse Unsymmetric Systems. *ACM Trans. on Mathematical Software*, 22(1):30–45, 1996. 243
12. G. Fox, S. Hiranandani, K. Kennedy, C. Koelbel, U. Kremer, C-W. Tseng and M. Wu. Fortran D Language Specification. Tech. Report COMP TR90-141, Computer Science Dept., Rice University, 1990. 230
13. C. Fu and T. Yang,. Run-time Compilation for Parallel Sparse Matrix Computations. *10th ACM Int'l Conf. on Supercomputing*, Philadelphia, pp. 237–244, May 1996. 244
14. K. Gallivan, B.A. Marsolf and H.A.G. Wijshoff. Solving Large Nonsymmetric Sparse Linear Systems Using MCSPARSE. *Parallel Computing*, 22(10):1291–1333, 1996. 244
15. High Performance Fortran Forum. High Performance Language Specification, Ver. 1.0. *Scientific Programming*, 2(1–2):1–170, 1993. 230
16. High Performance Fortran Forum. High Performance Language Specification, Ver. 2.0". Rice University, Houston, TX, February 1997. 230
17. V. Kotlyar and K. Pingali. Sparse Code Generation for Imperfectly Nested Loops with Dependences. *11th ACM Int'l Conf. on Supercomputing*, Vienna, Austria, 188–195, July 1997. 237
18. R. Ponnusamy, Y.-S. Hwang, R. Das, J. Saltz, A. Choudhary and G. Fox. Supporting Irregular Distributions Using Data-Parallel Language. *IEEE Parallel and Distributed Technology: Systems and Applications*, 3(1):12–24, 1995. 231
19. L.F. Romero and E.L. Zapata. Data Distributions for Sparse Matrix Vector Multiplication. *Parallel Computing*, 21(4):583–605, 1995. 231, 235, 244

20. G.P. Trabado and E.L. Zapata. Exploiting Locality on Parallel Sparse Matrix Computations. *3rd EUROMICRO Works. on Parallel and Distributed Processing*, San Remo, Italy, pp. 2–9, 1995. [231](#), [241](#)
21. M. Ujaldón, E.L. Zapata, B. Chapman and H.P. Zima. Vienna-Fortran/HPF Extensions for Sparse and Irregular Problems and their Compilation. *IEEE Trans. on Parallel and Distributed Systems*, 8(10):1068–1083, 1997. [231](#)
22. H. Zima, P. Brezany, B. Chapman, P. Mehrotra and A. Schwald. Vienna Fortran – A Language Specification. Tech. Report ACPC–TR92–4, Austrian Center for Parallel Computation, University of Vienna, Austria, 1992. [230](#)