



HAL
open science

Two-dimensional Block Partitionings for the Parallel Sparse Cholesky Factorization: the Fan-in Method

Bogdan Dumitrescu, Mathias Doreille, Jean-Louis Roch, Denis Trystram

► **To cite this version:**

Bogdan Dumitrescu, Mathias Doreille, Jean-Louis Roch, Denis Trystram. Two-dimensional Block Partitionings for the Parallel Sparse Cholesky Factorization: the Fan-in Method. [Research Report] RR-3156, INRIA. 1997. inria-00073533

HAL Id: inria-00073533

<https://inria.hal.science/inria-00073533v1>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Two-dimensional Block Partitionings for the
Parallel Sparse Cholesky Factorization : the
Fan-in Method***

B. Dumitrescu, M. Doreille, J.-L. Roch and D. Trystram

N° 3156

April 21, 1997

_____ THÈME 1 _____



***rapport
de recherche***

Two-dimensional Block Partitionings for the Parallel Sparse Cholesky Factorization : the Fan-in Method

B. Dumitrescu*, M. Doreille, J.-L. Roch and D. Trystram †

Thème 1 — Réseaux et systèmes
Projet APACHE

Rapport de recherche n° 3156 — April 21, 1997 — 29 pages

Abstract: This paper presents a discussion on 2D block mappings for the sparse Cholesky factorization on parallel MIMD architectures with distributed memory. It introduces the fan-in algorithm in a general manner and proposes several mapping strategies. The grid mapping with row balancing, inspired from Rothberg's work [21, 22] proved to be more robust than the original fan-out algorithm. Even more efficient is the proportional mapping, as show the experiments on a 32 processors IBM SP1 and on a Cray T3D. Subforest-to-subcube mappings are also considered and give good results on the T3D.

Key-words: sparse Cholesky factorisation, parallel algorithms, fan-in communication, 2D block partitioning, proportional mapping

(Résumé : *tsvp*)

* Politehnica University of Bucharest, Department of Automatic Control and Computers, 313, Splaiul Independenței, 77206 Bucharest, Romania. This work was performed while the author was visiting the Laboratoire de Modélisation et Calcul de Grenoble, supported by a TEMPRA grant.

† Research supported by the (CNRS - INRIA - INPG - UJF) joint project *Apache*. IMAG-LMC, 100 rue des Mathématiques, 38041 Grenoble cedex France.

Partitionnements par blocs bi-dimensionnels pour la factorisation parallèle creuse de Cholesky : la méthode fan-in

Résumé : Ce rapport étudie les partitionnements par blocs bi-dimensionnels pour la factorisation parallèle creuse de Cholesky sur des machines MIMD à mémoire distribuée. Nous introduisons l'algorithme fan-in dans un cadre général et étudions différentes stratégies de placement. Le placement sur grille avec équilibrage de charge sur les lignes, inspiré des travaux de Rothberg [21, 22], s'avère plus robuste que l'algorithme fan-out original. Le placement proportionnel est encore plus efficace, comme le montrent les expérimentations sur un IBM SP1 à 32 processeurs et sur un Cray T3D. Le placement sous-forêt vers sous-cube est également étudié et donne de bons résultats sur le Cray T3D.

Mots-clé : factorisation creuse de Cholesky, algorithmes parallèles, communication fan-in, partitionnement blocs 2D, placement proportionnel.

1 Introduction

Many problems in scientific and engineering computation request to solve a linear system $Ax = b$, where A is a sparse symmetric positive definite matrix. To solve the system, the Cholesky factorization $A = LL^T$ is the most time consuming step.

Although a classic problem, the factorization continues to request interest due to the effort to find algorithms well adapted to actual computer architectures. The class of parallel algorithms is especially targeted, since many approaches are possible and promising. We do not intend to begin with a brief history of the domain, but rather to spread it through the paper.

To give only the general characteristic, our contributions concern the class of two-dimensional block algorithms, using fan-in communication, on MIMD architectures with distributed memory.

An outline of the paper follows. We remind in section 2 the main stages in computing the sparse Cholesky factorization. Section 3 describes the operations preceding a 2D partitioning of the sparse matrix, such that BLAS 3 routines be used efficiently. In section 4, we present the general principles of fan-out and fan-in algorithms and give an outline of the previous work of Rothberg and Gupta [21]. In section 5, the fan-in algorithm for 2D mapping is introduced in a general manner. Section 6 contains the key to efficiency: specific mapping strategies, which combine known heuristics in a new way. Section 7 is devoted to experiments, which show the good behavior of the new fan-in algorithms. Finally, in section 8, we discuss some promising perspective issues.

2 General lines for Cholesky factorization

The nonzero structure of the Cholesky factor L includes the structure of A (supposing that no exact cancellation occurs); the other nonzeros are generally named fill-in. Since a sparse matrix is stored in a compressed format, containing only the nonzero elements and the index information required to access them, the process of factorization cannot be performed directly on the initial matrix A , but is split, for efficiency reasons, in three distinct stages.

First, a permutation (ordering) of the initial matrix is computed, in order to reach several objectives; the main is to reduce the fill-in occurring during factorization, and hopefully the number of operations required to compute L . For a parallel implementation, a tradeoff between minimizing fill-in and enhancing the natural parallelism offered by matrix sparsity is desired. In a raw classification, there are two classes of ordering algorithms, using minimum degree and nested dissection strategies, respectively. The purpose of this paper is not to discuss such algorithms; however, we will present some comparisons in section 7.

Second, the symbolic factorization of A is computed, i.e. the structure of the Cholesky factor L . This is crucial for efficiency, since storage requirements are solved and all the elements of L have a fixed position in memory, after this stage but before their actual computation.




Figure 1: Structure of a Cholesky factor and its elimination tree.

Third, the numerical factorization is performed. This is the most time consuming stage. Sequentially, there are several ways to organize the computation, all based on a Gaussian elimination process. If we mention only column oriented algorithms, there are left-looking and right-looking algorithms, the latter class including the multifrontal method. This paper is concerned only with the numerical factorization stage.

An important tool for both sequential and parallel algorithms is the elimination tree associated with the Cholesky factor. In this tree, each node represents a column; the father of node j is node i , where l_{ij} is the first subdiagonal nonzero on column j (i.e. the smallest $i > j$ such that $l_{ij} \neq 0$). An example of sparse matrix and its elimination tree is given in figure 1. For many algorithms, including some of this paper, it is important that the elimination tree be numbered in postorder, i.e. a father be numbered immediately after its sons; usually, ordering algorithms give a postorder numbering or at least offer all the information to perform it. For an extensive study of elimination trees, see Liu [16].

The elimination tree illustrates the intrinsic parallelism of the sparse Cholesky factorization; the computation of a column j of L may be done after all columns in the subtree rooted in j are factorized (by contrast to the dense Cholesky factorization, where all columns $k < j$ must be factorized). Intuitively, an elimination tree with small height and large width is more appropriate to parallel computation than a high thin tree. The elimination tree offers a column view on parallelism, but we must stress that parallelism exists also for operations within a column.

3 Block algorithms: basic ingredients

The architecture of current processors and their greater speed in performing arithmetic operations, compared to memory accesses, imposed the use of BLAS 3 [7] as a crucial condition to the effectiveness of linear algebra programs. BLAS 3 routines operate at block

level; blocks may be kept in the fast memory (cache) while performing several operations using the same matrix elements; this way, the floating point unit of the processor may be fed at full speed.

For dense matrices, the block approach is natural and more or less straightforward for most problems; the short interval at which LAPACK [2] followed BLAS 3 is a consequence of this state of things.

For sparse matrices there are several difficulties. We will focus on the solutions appropriated to the Cholesky factorization.

The main constraint is that sparsity is opposed to arbitrary grouping (consecutive) columns, as for the dense case. A group of columns $j, j + 1, \dots, j + s - 1$, for which the diagonal block $L(j : j + s - 1, j : j + s - 1)$ is full lower triangular and which share the same structure of rows with index $i \geq j + s$, is called *supernode*; in the elimination tree, these columns are each one a son of the next one; (if the son is unique, then the supernode is called fundamental;) if the elimination tree is not in postorder, the columns are not necessarily consecutive; this is an important reason to use postorder. For the matrix in figure 1, there are 5 supernodes, made up of columns $\{1\}$, $\{2,3\}$, $\{4\}$, $\{5,6\}$, $\{7,8\}$.

The columns of a supernode can be factorized together as for a dense matrix, allowing block operations; this approach was used, to give only few examples, in [18] or [9], but goes back to [5] and even earlier.

However, the supernode structure is specific to each matrix, so the supernodes may be too small or too large (the size of a supernode is s , the number of its columns). The latter case is benign; if required, large supernodes can be divided into *panels*, i.e. groups of consecutive columns which obviously have the same properties as supernodes.

Small supernodes are indeed a problem because their behavior is similar to that of isolated columns. It is worth mentioning that most supernodes have small size. A compromise is in order: several supernodes can be considered as a greater supernode, but with the sacrifice of treating some zeros as nonzeros, such that the new *amalgamated* supernode fit the definition of a true supernode, i.e. its columns have the same row structure. For example, in figure 1, if the element (3,1) is treated as a nonzero, then columns 1, 2 and 3 form a larger supernode. Even if more floating point operations are necessary to factorize the matrix after amalgamation, one can hope a reduction of the execution time due to greater block sizes. Ashcraft and Grimes [4] proposed an effective and simple algorithm for supernode amalgamation, depending on a single parameter – the maximal number of nonzeros that may be added to an amalgamated supernode. We coded and used this algorithm for all our tests. (It should be mentioned that supernodes selected for amalgamation are not necessarily consecutive; thus, a reordering of matrix columns is required in order to obtain real block columns in the physical memory; however, rows are not affected by reordering: indices are modified, but not the relative positions of elements.)

From now on we assume a given block column structure, obtained with the techniques listed above. As for storage, the use of BLAS 3 imposes that a block column is indeed a matrix, i.e. zeros corresponding to the upper part of the submatrix $L(j : j + s - 1, j : j +$




Figure 2: Block sparse Cholesky factor and its elimination tree.

$s - 1$) are explicitly stored. The extra memory required is however relatively not important, moreover when large supernodes are split into panels.

So far, the treatment applied to the matrix is common for many sequential and parallel methods for the Cholesky factorization. We turn our attention to parallel methods.

All parallel algorithms working by columns presented in the early review [13] can be successfully extended for block columns. There are several implementations for distributed memory multiprocessors; the recent review of Duff [8] offers some bibliographical references.

The principle of allocating parts of the matrix to processors is a first sensible point. Schreiber [24] was the first to explain that a one-dimensional mapping (i.e. a whole block column to one processor) has poor scalability, and thus two-dimensional (2D) mappings are required for parallel efficiency.

At the present time, two classes of algorithms seem to take great advantage of the 2D mapping: the block fan-out algorithm of Rothberg and Gupta [21], based on a classic right-looking Gaussian elimination, and the multifrontal parallel algorithm proposed by Gupta, Karypis and Kumar [12].

Since our study is in the same framework, let us start describing the basic lines of Rothberg's approach, which is very simple. After splitting the matrix vertically in N block columns, as described above, the same split is applied by rows and thus a 2D partitioning is obtained. A block sparse matrix results, i.e. sparsity is now thought in terms of blocks L_{IJ} (we will use capital letters for block indices); however, nonzero blocks are not necessarily full; all we can say is that if a row of a block is nonzero, then all its elements are nonzero; only diagonal blocks are surely full lower triangular. We present in figure 2 the 2D block structure of the matrix from figure 1, after amalgamation (zeros (3, 1) and (5, 4) considered as nonzeros); there are 3 block columns (rows); block L_{21} is zero; blocks L_{31} , L_{32} have only one nonzero row.

```

BLOCK GAUSSIAN ELIMINATION ( $L$  is initialized with  $A$ )
1.   for  $K = 1 : N$ 
2.       compute in place the Cholesky factor of  $L_{KK}$ 
3.       for  $I \in \text{col}(K)$ 
4.            $L_{IK} \leftarrow L_{IK} L_{KK}^{-T}$ 
5.       for  $J \in \text{col}(K)$ 
6.           for  $I \in \text{col}(K)$  and  $I \geq J$ 
7.                $L_{IJ} \leftarrow L_{IJ} - L_{IK} L_{JK}^T$ 

```

Figure 3: 2D block-level Gaussian elimination.

We work now with a block column elimination tree, defined similarly to the column elimination tree, as seen in figure 2. Its signification is the same: a block column J can be computed only after all columns in the subtree rooted in J are computed.

The sequential Gaussian elimination algorithm at block level is similar to the element level one. We present in figure 3 the *kij*, or "right-looking" version; this algorithm should be viewed only as a general framework. We denote by $\text{col}(K)$ the set of row indices of subdiagonal blocks of column K , i.e. $\text{col}(K) = \{I \mid L_{IK} \neq 0, I > K\}$.

The Cholesky factorization of a diagonal block may be performed with the DPOTRF routine from LAPACK. BLAS 3 routines may be used for the other operations; in statement 4, DTRSM is used directly for the computation of L_{IK} , since a block contains full rows; the matrix multiplication $L_{IK} L_{JK}^T$ is performed with DGEMM; the result of this multiplication usually has not the same structure as L_{IJ} , neither on rows or columns; an insert-add operation completes the update 7 of L_{IJ} . The techniques outlined here are detailed in [21]; we note also the use of relative (row) indices (Schreiber [23] introduced this notion). Adapting the traditional notation, we will denote by $\text{bdiv}(K, I)$ and $\text{bmod}(K, I, J)$ the operations 4 and 7, respectively (standing for "block division" and "block modification").

4 Fan-out, fan-in algorithms

We describe here in general terms the possible organization of parallel algorithms based on the 2D block structure of a sparse matrix.

This 2D structure allows the mapping of a block column K to several processors. The set of processors sharing K will be called $\text{group}(K)$, or simply group , when the context is clear. We will denote by $\text{owner}(I, K)$ the processor holding the block L_{IK} . From now on we will use "panel" instead of "block column".

Although the 2D mapping was proved to be superior to 1D mapping, there is still interest in allocating all the blocks of some panels to the same processor. If $|\text{group}(K)| > 1$, communication is needed; for example, $\text{owner}(K, K)$ must send L_{KK} to all other processors in $\text{group}(K)$ in order that bdiv operations be possible. The benefit is greater for bmod ; let us assume that a whole subtree of the elimination tree is mapped to a processor, as in figure




Figure 4: An elimination tree: local subtrees and distributed panels.

4; then, at least for panels K and J belonging to this subtree, where K is a descendant of J , the operation $bmod(K, I, J)$ is local for any I . Greater (higher) the local subtrees, smaller the communication volume.

Certainly, there must be a tradeoff between the size of local subtrees and balancing processor load. Deferring the details to section 6, we assume now that some algorithm has been used and the elimination tree is separated in several local subtrees and panels that are distributed among processors. In figure 4, a dotted curve separates the local subtrees and the distributed panels. Let us remark that several subtrees may be mapped to the same processor and that a distributed panel is not mapped to all processors, generally.

Let us further study how a parallel algorithm could be described, given the above distribution of blocks. The strategies to be used are two, well known as fan-out and fan-in, depending on the processor computing $bmod(K, I, J)$.

In the fan-out strategy, $bmod(K, I, J)$ is computed by $owner(I, J)$ ("at destination", as sometimes said); to this purpose, L_{IK} and L_{JK} must be sent to $owner(I, J)$.

In the fan-in strategy, $bmod(K, I, J)$ is computed either by $owner(I, K)$ or $owner(J, K)$ ("at source"); we assume that $owner(I, K)$ performs the mentioned update, recalling that $I \geq J$; this could be somehow an arbitrary assumption, but simplifies the presentation and the practical behavior was satisfactory (however, more complicated schemes may be tried). Each processor involved in the computation of L_{IJ} holds its own contribution in a local block; only $owner(I, J)$ initializes this block with A_{IJ} , while the other processors with zero. When a processor has finished all updates of its contribution, it sends it to $owner(I, J)$ which adds it to L_{IJ} ; in fact this is rather an insert-add operation, since contributions structure may differ of that of L_{IJ} . To conclude, in the fan-in strategy, for each $bmod(K, I, J)$ (here K is variable), the block L_{JK} must be sent to $owner(I, K)$, but each processor sends at most one aggregated contribution to $owner(I, J)$.

Generally, the fan-in strategy implies a smaller number of communicated blocks. It is difficult to make a similar assertion about communication volume; supernodes tend to

		Distributed step	
		Fan-out	Fan-in
Local step	Fan-out	2D block mapping: costly column mapping: [11]	
	Fan-in	2D block mapping: Rothberg [21] column mapping: [14]	2D block mapping: this paper column mapping: [3]

Figure 5: Communication strategies and possible combinations.

became larger for greater indices, and thus destination blocks (L_{IJ}) usually have greater size than source blocks. However, our experiments showed a smaller communication volume of fan-in methods.

Costs modeling depend on the communication strategy; the cost of $bmod(K, I, J)$ is associated with block L_{IJ} for fan-out, and with L_{IK} for fan-in.

We must distinguish now between the local computation step and the distributed one. Part of the local computation potentially affects panels that are distributed; to be more precise, let assume that K is a local panel and L_{JK} , L_{IK} are subdiagonal blocks; if panel J is also local (and necessarily mapped to the same processor), then $bmod(K, I, J)$ may be performed locally, as we already remarked; if panel J is distributed, then some blocks must be sent to $owner(I, J)$, depending on the strategy. An essential observation is that, no matter the strategy, the communication could be postponed until all local subtrees are computed.

A 2D block parallel algorithm for the Cholesky factorization will thus have three steps :

1. Factorization of panels in local subtrees (only local computation).
2. Communication of blocks (or updates) computed at step 1 and affecting distributed panels.
3. Factorization of distributed panels (involving computation and communication).

It is not necessary that the same communication strategy (fan-out or fan-in) be used in the local and distributed steps. Figure 5 presents the four possible combinations. Three of them have "ancestors", i.e. column oriented algorithms; the fourth, which is fairly unnatural, was never tried; despite the age of about a decade of these algorithms, we never saw such a classification.

For 2D block mappings, the only existent algorithm belongs to Rothberg and Gupta [21] and is of "fan-in fan-out" type. The "fan-out fan-out" algorithm is presumably less efficient due to large communication costs, as for the column case; we implemented this algorithm and its efficiency is indeed lower; however, less communication is required than for a 1D based algorithm, since only few blocks of local panels are sent. The "fan-in fan-in" algorithm is much more appealing; we will present it in the next sections.

For the sake of brevity, Rothberg’s algorithm will be called ”fan-out”, while our algorithm ”fan-in”.

We present now an outline of the fan-out algorithm distributed step. The underlying architecture is supposed to be a $p_r \times p_c$ grid; a cyclic mapping is used, i.e. a block L_{IJ} is mapped to processor $(I \bmod p_r, J \bmod p_c)$. Thus, the communication is reduced: a computed block L_{IK} must be sent only to processors owning row I or column I of L , to participate if necessary to updates $bmod(K, I, J)$ or $bmod(K, J, I)$. The algorithm is data driven; after receiving a block, a processor looks at what updates this block can participate, perform these updates and, if it is the case, diagonal block factorizations and *bdiv* operations on local blocks; finally, it sends any newly factorized block to appropriate processors. A very clear description of the algorithm is presented in [21].

Since the cyclic mapping is rigid and may cause load imbalancing, Rothberg and Schreiber [22] proposed an improved mapping scheme; matrix block rows and columns are still mapped to processor grid rows and columns, but not cyclically; instead, a balancing of grid rows (columns) load is searched. (The problem reduces to bin packing.) For rows, the idea was indeed effective; for columns, it seems that usually the cyclic mapping is fairly satisfactory.

5 The fan-in algorithm: general presentation

We present in this section the details of the fan-in algorithm, excepting matrix blocks mapping, which is the key of efficiency, but is not affecting the validity of the algorithm; we need only the notion of group of processors owning blocks in a certain panel K , denoted by $group(K)$. The reader familiar with Rothberg’s algorithm can imagine that $group(K)$ is a column of the processor grid. Only the distributed step is detailed, since the others are straightforward.

The main structure of the algorithm is presented in figure 6, and two important functions are listed in figures 7 and 8.

The general communication pattern was described in the previous section. We recall that once a block L_{IK} of the Cholesky factor is computed, it is broadcast by its owner to all processors in $group(K)$. By the other hand, each processor accumulates all his updates for non local blocks and sends them only once to the owner.

Similarly to Rothberg’s fan-out algorithm, the number of updates a block must suffer is computed initially; we denote it by $nmod(I, J)$; this computation may take place in the symbolic factorization stage. During the factorization, $nmod(I, J)$ is decremented at each update.

Let us explain the data structures appearing in the algorithm; $wait(K)$ is a list containing row indices of local blocks from panel K which need only the *bdiv* operation to be factorized (i.e. the arrival of block L_{KK}); $ready(K)$ is a list of local or received blocks of panel K of the Cholesky factor (i.e. already computed); $diag(K)$ is a flag indicating if the block L_{KK} was factorized and is present in the local memory; $queue$ is a list of local blocks that

can be factorized with local information; their factorization was postponed for a simpler organization of the algorithm. All these variables have an initial value set to zero or empty.

An important question is when a processor sends to $owner(I, J)$ its contribution to L_{IJ} , since this communication must be performed only after all possible local updates were aggregated. The solution we adopted is to count, initially, the number of updates a processor must perform on each block; the cost of this operation similar to a partial symbolic factorization, is negligible, moreover when block size is large; we use $nmod(I, J)$ to keep the number of updates; note that $nmod$ has different significations for local and non local blocks.

We will denote by $bmod_send(K, I, J)$ (see figure 7) the operation $bmod(K, I, J)$, followed, if L_{IJ} is not local, by a check of the number of updates and by a send, if all updates were performed; note that the number of updates must be transmitted to $owner(I, J)$, in order to adjust $nmod(I, J)$ in line 15 of the main algorithm; if L_{IJ} is local, and is completely updated, then it is added to the list $wait(J)$ or to the *queue*.

Although the algorithm is detailed, there are several aspects to be made more clear. Statements 24 in the main routine, 8 and 13 in *last_block_op* are not quite correct; when scanning $ready(K)$ for indices J , there are three situations when the current processor performs $bmod_send(K, I, J)$: if L_{IK} and L_{JK} are both local; if L_{IK} is local, L_{JK} is a received block (in the sense that $owner(J, K) \neq me$), and $I > J$; if L_{IK} is a received block, L_{JK} is local and $I < J$. In a real program, when calling $bmod_send(K, I, J)$, second argument value must be greater than (or equal to) third argument value; our description doesn't actually respect this rule.

The condition in statement 8 of the main routine is implemented by initially counting the number of blocks of the Cholesky factor to be received by a processor and updating this number after each receive of such a block; another counter with the number of local blocks is decremented when a block is factorized.

Another small remark is that when *last_block_op* is called for a block L_{IK} extracted from the *queue*, this block is always factorized, i.e. one of the following conditions is true: $I = K$ or $diag(K) = 1$.

6 Mappings for the fan-in algorithm

We present now several methods to map the blocks of a matrix such that the fan-in algorithm be efficient. There are two objectives (recall figure 4): to find large enough local subtrees that can be mapped to processors such that the load imbalance is reasonable; and to map the blocks of the distributed panels. These objectives may be attained by the same, or by distinct algorithms.

Such algorithms are based on costs associated with the computation performed on a panel; these costs can be computed with low overhead in the stage of symbolic factorization; the cost associated with a panel is the sum of its blocks costs; the cost of a block operation may be modeled as the sum of the number of floating point operations and a constant representing other operations (function calls, block address calculation, etc.) [22]; the costs of block L_{KK} Cholesky factorization and of $bdiv(I, K)$ are associated naturally with blocks

```

(me is the number of current processor)
// local step and involved communication
1. for  $K = 1 : N$ 
2.   if  $K$  is a local panel
3.     factorize all blocks in  $col(K)$ 
4.     for all  $I, J \in col(K)$  with  $I \geq J$ ,  $bmod\_send(K, I, J)$ 
5.   else if  $me \in group(K)$ 
6.     performed all possible computation on local data
7.     initialize lists  $wait(K)$  and  $ready(K)$ 
// distributed step
8. while there are blocks to be received and local blocks to compute
9.   while  $queue$  is not empty
10.    take block  $L_{IK}$  from  $queue$ 
11.     $last\_block\_op(I, K)$ 
12.   if a block  $L_{IJ}$  was received
13.     if  $L_{IJ}$  is an update block
14.       insert-add it to local block  $L_{IJ}$ 
15.       subtract from  $nmod(I, J)$  the number of updates of the received block
16.       if  $nmod(I, J) = 0$ 
17.          $last\_block\_op(I, J)$ 
18.     else ( $L_{IJ}$  is a block of the Cholesky factor)
19.       if  $I = J$ 
20.          $diag(J) \leftarrow 1$ 
21.          $bdiv(I', J)$  for all  $I' \in wait(J)$ 
22.       else
23.         put  $I$  in list  $ready(J)$ 
24.         perform all my  $bmod\_send(J, I, I')$ , for  $I' \in ready(J)$ 

```

Figure 6: The fan-in 2D block algorithm.

```

function bmod_send( $K, I, J$ )
1.   bmod( $K, I, J$ )
2.    $nmod(I, J) \leftarrow nmod(I, J) - 1$ 
3.   if  $nmod(I, J) = 0$ 
4.     if  $me = owner(I, J)$ 
5.       if  $I = J$  or  $diag(J) = 1$ 
6.         put  $L_{IJ}$  in the queue
7.       else put  $I$  in list wait( $J$ )
8.     else
9.       send  $L_{IJ}$  to owner( $I, J$ )

```

Figure 7: Function *bmod_send*.

```

function last_block_op( $I, K$ )
1.   if  $I = K$ 
2.     compute in place the Cholesky factor of  $L_{KK}$ 
3.      $diag(K) \leftarrow 1$ 
4.     broadcast  $L_{KK}$  to processors in group( $K$ )
5.     for all  $I' \in wait(K)$ 
6.       bdiv( $I', K$ )
7.       put  $I'$  in list ready( $K$ )
8.     perform all my bmod_send( $K, I', J$ ), for  $J \in ready(K)$ 
9.   else if  $diag(K) = 1$ 
10.    bdiv( $I, K$ )
11.    broadcast  $L_{IK}$  to processors in group( $K$ )
12.    put  $I$  in list ready( $K$ )
13.    perform all my bmod_send( $K, I, J$ ), for  $J \in ready(K)$ 
14.   else
15.     put  $I$  in list wait( $K$ )

```

Figure 8: Function *last_block_op*.

L_{KK} and L_{IK} , respectively; the cost of $bmod(K, I, J)$ is associated either with L_{IK} or L_{IJ} , for the fan-in or fan-out methods, respectively. The cost associated with a subtree is the sum of its panels costs.

It is not very easy to evaluate a mapping; there are two criteria, usually in conflict: communication volume and load balance. For evaluation, the relative importance of these criteria depends on the properties of the parallel computer, especially on the ratio communication vs. computation speed. In this section we will thus present qualitative comparisons between the mappings presented below, with respect to the two criteria. In the next section, the experiments will give a more precise insight.

6.1 Grid mappings

Geist and Ng [10] algorithm is used to map local subtrees. They proposed to keep a list of subtrees, initialized with the root, and to try a mapping by the means of a bin-packing algorithm (i.e. giving subtrees in decreasing cost order to the currently least loaded processor); if the workload imbalance is unacceptable, the heaviest subtree is deleted from the list, and its sons are added to the list.

Distributed panel blocks are mapped following the idea of Rothberg and Gupta [21], used by them for the fan-out algorithm. Processors are supposed to be connected in a $p_r \times p_c$ grid; a block row (column) of the matrix is mapped to a row (column) of the grid; the algorithm from [22] is used to balance grid rows work; this mapping algorithm is adapted to the fan-in strategy by only changing costs associated with blocks, as $bmod$ operations are performed at source and not at destination.

This grid mapping is attractive for its simplicity and for limiting communication: since $group(K)$ is a grid column, a processor will broadcast a factorized block only to $p_r - 1$ processors; on the other hand, $owner(I, J)$ will receive updates only from processors on its row because $bmod(K, I, J)$ is always performed by $owner(I, K)$ (recall that $I \geq J$).

For further reference, we will call *fan-in on grid* (FLGRID) the fan-in algorithm with grid mapping with row and column locality and row balancing.

6.2 Proportional mappings

The grid mapping is somehow rigid; that is, always allocating a grid column to a panel is a restrictive scheme. As noticed in the previous section, we are free to choose processors in $group(K)$ upon desire without affecting the correctness of the fan-in algorithm. Let us suppose that the cost of a communication between any two processors is the same, no matter the physical connectivity of the architecture (in fact this is a fair assumption for many actual parallel distributed-memory computers).

Looking again at figure 4, let imagine that a distributed panel J have two local subtrees as descendants, like in the right side of the figure. Let suppose that two different processors are in charge with the local computation of the two subtrees. Since all updates $bmod(K, I, J)$ are performed by the two processors, it is natural to map panel J to them; all updates targeted to panel J will be communicated only between these processors; moreover, the

factorization of panel J will also imply the same pattern. Finally, global load balance seems to be preserved, with two conditions; first, that work on local subtrees has been balanced, which is supposed to be true; second, that the work on panel J is evenly distributed between the two processors forming $group(J)$.

This mapping idea can be immediately generalized for the whole elimination tree. More precisely

$$group(J) = \bigcup_{K \in sons(J)} group(K),$$

i.e. a panel is mapped to all processors owning its sons.

Geist and Ng [10] mapping of local subtrees may be used, but it has the drawback of mapping several subtrees to the same processor; presumably, more communication will be necessary then in the case of one subtree per processor, because large groups of processors tend to appear (every processor may be allocated to any subtree rooted in a distributed panel).

The proportional mapping of Pothén and Sun [19], meant by the authors for the multifrontal method, seems more promising. Although the one local subtree per processor condition is not guaranteed, this mapping is very effective in providing disjoint groups. We present the basic lines of a slightly modified version of the algorithm in figure 9. The principle is simple: each son K (in the elimination tree) of a panel J is mapped to a subset of $group(J)$ of size proportional to the cost $w(K)$ of the subtree rooted in K ; the algorithm is recursively applied; the recursion is stopped when a group has size 1, i.e. a local subtree was obtained. The calling arguments for *prop_map* are the root of the elimination tree (argument J), a list of all processors ($g \equiv group(J)$) and a list of length p initialized with zeros, standing for processor total work (pw). (If the elimination tree is in fact a forest, its trees are mapped in decreasing order of costs; a more subtle approach could be taken, but it is not worth since in practice only one tree is large, while the others, if any, are rather isolated nodes.)

Since $|group(K)|$ is an integer, there are problems caused by rounding, e.g. sons for which less than i processors can be allocated; the original paper [19] shows how to avoid all these problems. We didn't address the problem of how the i processors are selected in line 6; it seems difficult to propose a heuristic; so that we simply took the first i processors from the list g .

The blocks of a panel J are mapped after the mapping of all subtrees rooted in sons of J , i.e. a bottom-top approach was preferred. The explanation is natural; the mapping in line 15 may be based on the most recent processor workload information, since the Gaussian elimination evolution is from bottom to top in the elimination tree.

We intentionally left unaddressed the details of how blocks in a panel are mapped. Any processor in $group(K)$ may be a candidate to own any block. A mapping heuristic has two requirements which are rather disharmonic: good load balancing and small communication volume. We will distinguish two classes of mappings.

Greedy mapping. The simpler idea is to give, in line 15 of *prop_map*, the current block to the least loaded processor of the group. The loop 14 goes in decreasing order

```

function prop_map(J, g, pw)
1.   sons(J) is the list of panel J sons, ordered upon decreasing cost
2.    $m \leftarrow |g|$ , i.e. the number of processors in group(J)
3.    $w \leftarrow \sum_{K \in \text{sons}(J)} w(K)$ 
   // allocate processors in g to sons of J
4.   for  $K \in \text{sons}(J)$  (in decreasing cost order)
5.      $i \leftarrow \lfloor w(K)/w * m + 0.5 \rfloor$ 
6.     group(K)  $\leftarrow$  a set of i processors in g
7.      $g \leftarrow g \setminus \text{group}(K)$ 
   // recursive proportional mapping
8.   for  $K \in \text{sons}(J)$ 
9.     if  $|\text{group}(K)| > 1$ 
10.      prop_map(K, group(K), pw)
11.    else
12.      map the subtree rooted in K to P (group(K) = {P})
13.       $pw(P) \leftarrow pw(P) + w(K)$ 
   // map current panel blocks to group(K)
14.   for  $I \in \text{col}(J) \cup \{J\}$ 
15.     map  $L_{IJ}$  to a processor  $P \equiv \text{owner}(I, J) \in \text{group}(J)$ 
16.     add to  $pw(P)$  the cost associated with  $L_{IJ}$ 

```

Figure 9: An outline of the proportional mapping algorithm.

of row indices, usually close to a decreasing block cost order. A good load balancing is assured; a possible drawback is that there is no attempt to limit communication (if we don't count the limitations provided by the proportional mapping itself); a factorized block will be broadcast to all processors in the group, while updates may be received from all processors in the group.

Subgrid mappings. Let imagine the case of a supernode divided in several panels; in the elimination tree, these panels form a chain; if the supernode is fundamental – and, if not, in most cases – all the panels will be mapped to the same group of processors. Structuring the group becomes interesting, as a way to further reduce communication; if the group is split in subgroups and each panel is mapped to a subgroup, broadcasts will occur only inside subgroups.

The most natural (to the block Cholesky factorization) is the (sub)grid structure. A group of m processors may be thought as a $m_r \times m_c$ grid. A panel will be mapped to a column of the subgrid.

Of course, there are some precautions to observe. Subgrid dimensions m_r and m_c may be chosen among the divisors of m ; we take m_r and m_c such that the subgrid be as close of a square as possible. Since the proportional mapping algorithm doesn't offer a control of group sizes, m may have few divisors. We adopted the following heuristic necessary conditions to use subgrid mapping instead of the greedy one: $m_r \leq m_c$, $m_c/m_r \leq \alpha$, $m_r \neq 1$; the constant α is used to limit the "distance" to a square grid; we used $\alpha = 3$, but this choice is purely intuitive.

The problem is now how to map panels to columns of subgrids. We identified two appealing techniques:

- mapping a panel to the currently least loaded subgrid column;
- wrapping the panels of a supernode on subgrid columns (starting with the least loaded column).

Further on panel blocks must be mapped; we can favorize load balancing or communication reduction. Two strategies result, respectively

- mapping a block to the least loaded processor on current subgrid column;
- wrap mapping nonzero blocks of the panel on the subgrid column.

The subgrid column and row strategies are independent so that four combinations result from the above discussion. (We should say however that other possibilities exist, for example the trivial cyclic mapping. We chose those techniques that seemed offering more robustness.)

For further reference, we will use the following notations for the variants of the fan-in algorithm with proportional mapping: FL_PROP_G for the greedy mapping and FL_PROP_SG_xx for the subgrid mappings, where the first x is for the column strategy and the second for the row strategy; we use the letter L for the "least loaded" approach and W for the wrap mapping.

6.3 Forest-to-subcube mappings

As we mentioned, a difficulty of the proportional mapping is the lack of control on group sizes. Gupta, Karypis and Kumar [12] proposed the *subforest-to-subcube* mapping, which ensures that groups are always subcubes of a hypercube (the underlying architecture being a hypercube). The basic idea is to try to split a list of unassigned subtrees, initialized with the root, into two parts of roughly equal costs, and to map each part to a half of the current group (initially, the whole hypercube); the algorithm is recursively applied for the resulted sublists and subgroups (which are always subcubes of the hypercube); if, when splitting the list, the imbalance is unacceptable, then the root of the heaviest subtree is mapped to the whole group, the root is deleted from the list, its child subtrees are added to the list and the algorithm is applied to the new list and the same group.

We can introduce in this algorithm, which is only panel oriented, the same techniques for mapping the blocks of a panel as in the previous subsection. Let us name the algorithm FL_CUBE (suffixes may be added like for FL_PROP). Since group sizes are a power of 2, the advantage over the proportional mapping is that subgrids are always a square or a rectangle with $m_c = 2m_r$. A possible drawback is that there may be more local subtrees per processor; moreover, like Geist and Ng's mapping, a bound of the accepted imbalance must be input to the algorithm; the proportional mapping has no parameters. We do not insist anymore because the experiments showed that usually the proportional mapping is better.

6.4 Heuristic comparisons

When evaluating algorithms, we insisted above on communication volume; it is presumable that the proportional mapping is the best for this criterion. Load balancing is harder to evaluate in general, but we can appreciate that FL_PROP and FL_CUBE are better than FL_GRID, due to their flexibility. Another feature of interest is the intrinsic parallelism offered by the algorithm, besides load balancing; even if globally the processors have the same amount of work, idle times may occur for some processors due to the lack of received blocks. The proportional mapping seems more subject to idle times; a processor has assigned blocks in panels that form a chain in the elimination tree, i.e. has no other work alternative if idle. For the other mappings, a processor may have assigned blocks in panels of the whole elimination tree; a temporary lack of work on one subtree may be compensated by available work on another.

7 Experiments

We implemented in C all methods described in previous sections, using MPI [17] for communication. BLAS 3 routines were used whenever possible. Portable and (hopefully) efficient programs resulted.

Let us recall the abbreviations for our methods; FO_GRID and FL_GRID are the fan-out and fan-in with blocks mapped as for a processor grid, with global row balancing;

Matrix name	Size	Nonzeros in A	Nonzeros in L	Mflop for L
B15_K	3,948	117,816	574,104	122.48
B15_A			627,763	155.45
B16_K	4,884	290,378	754,734	150.53
B16_A			812,183	186.42
B17_K	10,974	428,650	1,188,305	207.13
B17_A			1,055,927	162.99
B18_K	11,948	149,090	673,070	111.58
B18_A			645,717	131.67
B25_K	15,439	252,241	1,842,860	491.96
B25_A			1,479,108	316.32
G150	22,500	111,900	721,862	62.51

Table 1: Test matrices.

FI_PROP_G is the fan-in method with proportional mapping of panels and greedy mapping of blocks inside a panel; FI_PROP_SG_xx are fan-in methods using a proportional mapping scheme combined with subgrid mappings. FI_CUBE_G and FI_CUBE_SG_xx use a subforest-to-subcube general mapping, combined with greedy and subgrid mapping of panel blocks, respectively.

We tested our programs on several matrices of the Harwell-Boeing collection presented in table 1. For our use, we abbreviated their names; B15 stands for BCSSTK15, etc. The last matrix, G150, is a 5-point grid discretization. We added a suffix indicating the method used for ordering; K stands for the method used in METIS [15], using nested dissection principle (K is from Karypis, the first author); A stands for the approximate (external) minimum degree of Amestoy, Davis and Duff [1] (which is faster than exact minimum degree and gave similar results for our matrices). Matrix G150 is ordered with optimal nested dissection.

At least for these matrices, minimum degree methods seem less suited for parallel methods; we will report thus mainly for METIS orderings; a comparison will be finally given for B17 and B25, the only matrices where minimum degree offered better sequential performance.

As we described in section 3, supernodes were identified, amalgamated and split into panels. To give an idea, let say that even after amalgamation supernode size is rather small, about 10 in average; however, there are few large supernodes (among them, the last is the largest) with sizes in the hundreds. Amalgamation implies a greater number of operations for factorization; table 1 gives the flop count without amalgamation; however we will report our Mflops performances to this flop count; as a consequence, the actual figures are 5-10% better.

Panel size is an important issue. Intuitively, large panels favorize BLAS 3 routines, while small panels offer a greater intrinsic parallelism; small panels are slowing not only computation, but communication also, due to the increased effect of latency. A compromise

is necessary between these antagonistic tendencies. In our experiments, we varied panel size in order to find the best choice; we will present some recommendations which cannot be generalized without care. When no explicit mention, we report results for the best panel size.

We experimented our programs on two parallel computers. The first is a 32 processors IBM SP1, located at LMC-IMAG. Its processors peak performance for BLAS 3 routines is about 100 Mflops. The communication rate can go to 30 Mbytes/s, while the start-up time for a message is about 60 μ s. Each node has 64 Mbytes of local memory. For sparse matrix computations, we cannot hope at maximal computation and communication speed. For block sizes resulted for the test matrices, fair figures are 60 Mflops and 10 Mbytes/s. This means a ratio of about 50 flop for one transmitted double precision floating point, which is rather high; i.e. communication is slow with respect to computation. More than that, decreasing panel size implies a greater degradation of this ratio.

The second is a Cray T3D located at CEA Grenoble. Processors peak performance using BLAS 3 is about 110 Mflops, i.e. similar to the SP1; communication rate with MPI is at most 35 Mbytes/s; however, compared to the SP1, communication is roughly three times faster for current block sizes.

In the SP1, any pair of processors can be physically connected by the means of a switch. In the T3D, the communication network is a three-dimensional torus. On both computers, communication speed is the same between any two processors if we neglect possible delays due to conflicts on reserving switch channels on the SP1 or communication paths on the T3D. Since programs execution times are affected by paging effects and possibly by communication contention, we always report the best of four successive executions. Usually, there were not significant variations.

7.1 Sequential performance

The sequential version of our programs attained usually 50-60 Mflops on the SP1, which may be considered satisfactory for test matrix sizes, but only 25-30 Mflops on the T3D. We must stress that the performance varies enough function of panel size. An example is given in figure 10, for three of our matrices. On the SP1, it can be seen that for a block size of 16, the computation speed is little more than one half of the maximal speed (notice that for panels larger than 128, performance is still slowly growing); this fact is clearly limiting parallel speed-up, since small panels enhance problem parallelism. On the T3D, the higher performance is attained for smaller blocks than on SP1; more than that, for very large panels, performance becomes worse. These remarks, together with the better communication vs. computation speed ratio, allow us to anticipate better speed-ups on the T3D.

7.2 Parallel performance

We present here some significant results issued from many timings for the different methods, matrices and panel sizes. There are still parameters we didn't vary; the most important is the imbalance bound in Geist and Ng's algorithm, used for local subtrees mapping in FO_GRID




Figure 10: Sequential performance (Mflops) for variable panel size.

		B15_K	B16_K	B17_K	B18_K	B25_K	G150	B17_A	B25_A
on SP1	$p = 16$	0.7%	4.9%	0.4%	0.9%	0%	10.8%	0.7%	0%
	$p = 32$	10.2%	9.9%	4.3%	15.5%	6.5%	18.0%	7.5%	7.1%
on T3D	$p = 16$	0.6%	3.6%	-2.6%	6.7%	2.0%	7.9%	0%	1.1%
	$p = 32$	-0.2%	5.2%	1.8%	4.3%	3.5%	7.2%	-4.2%	-1.9%

Table 2: Performance improvement: FL_GRID with respect to FO_GRID.

and FL_GRID; we used the value 1.4 (the ratio between the largest load of a processor and the average load), which gave good results for $p = 16$ processors; for $p = 2$ we used the bound 1.2; it is clear that for some given matrix, p and panel size a better value can be found, but we appreciate that the improvement is minor.

Since FO_GRID was implemented by Rothberg on a iPSC/860 [20] having a communication vs. computation speed ratio of the same order as the SP1 we compared his and our speed-ups. They are similar, but we must remark that Rothberg used only minimum degree orderings, to which the comparison was thus limited.

Grid mappings. Let compare first the two grid algorithms. For a small number of processors (i.e. $p \leq 8$), FO_GRID and FL_GRID have very similar performance. On the SP1, for $p = 16$, there is a small advantage for FL_GRID; for $p = 32$, FL_GRID is always better. On the T3D, it is difficult to say which method to choose, although FL_GRID seems slightly better. See table 2. The general performance of FL_GRID on the test matrices is presented in figure 11.




Figure 11: Performance (Mflops) for the fan-in on grid (FLGRID) algorithm.

Although FLGRID brings an improvement over FO_GRID, the scalability of the algorithm is poor on the SP1; in fact, the matrices used for experiments offer a small amount of work, compared to processor power; even the largest, B25_K, is factorized in little more than one second on 32 processors. We can thus expect much better results for larger matrices. It is also to say that the test matrices are also irregular, excepting G150.

Proportional mappings. The proportional mapping reveals its superiority over the grid mappings, for nested dissection orderings. Again, for $p \leq 8$, it is difficult to distinguish a better method; for the proportional mapping it is intuitive that a small number of processors is not favorable, because rounding errors in group allocation may be important (line 5 in figure 9). For $p = 16$, $p = 32$, FLPROP_G is largely superior, as seen in table 3; compare also to table 2. However, let remark that the improvement on the T3D is smaller than on the SP1.

The performance of FLPROP_G is presented in figure 12, which show a better scalability of this algorithm, compared to grid ones. The best result is again for B25_K, letting us estimate that larger matrices will furnish even better performance.

However, minimum degree orderings give advantage to grid algorithms. The explanation is simple; the elimination tree is high and thin, at least in its upper part; the proportional mapping produces groups with slowly decreasing size; a greater communication volume results. In this case, grid algorithms are more robust, with their general and simple communication pattern.

Experiments with FLPROP_SG_XX algorithms didn't show an improvement over the greedy variant; among them, the "wrap" scheme (FLPROP_SG_WW) seems the most promising; on the SP1, it gives better results than FLPROP_G on G150, but slightly lower on

		B15_K	B16_K	B17_K	B18_K	B25_K	G150
on SP1	$p = 16$	1.9%	21.8%	4.2%	14.4%	14.0%	16.3%
	$p = 32$	18.6%	43.1%	18.0%	27.5%	27.0%	40.2%
on T3D	$p = 16$	-6.4%	3.6%	-5.6%	8.1%	6.8%	6.3%
	$p = 32$	13.7%	32.7%	8.4%	28.4%	9.1%	13.8%

Table 3: Performance improvement: FI_PROP_G with respect to FO_GRID.




Figure 12: Performance (Mflops) for the fan-in proportional greedy (FI_PROP_G) algorithm.




Figure 13: Relative performance: FLPROP_G vs. FLCUBE_SG_WW on the T3D.

the other matrices; on the T3D, it is on the same level as FLPROP_G. However, further investigation is needed, since there are several parameters to be tuned before stating a firm conclusion.

Subforest-to-subcube mappings gave poor results on the SP1; the cause is the great communication volume; we will exemplify later this issue. On the T3D, these mappings were clearly the best for a small number of processors and left the advantage to FLPROP_G only for $p = 32$. FLCUBE_SG_WW is the best in the family; we give in figure 13 a comparison between FLPROP_G and FLCUBE_SG_WW on the T3D.

7.3 The effect of panel size

We present here some results and recommendations concerning panel size. In our experiments we varied panel size from 16 to 128 on the SP1 and to 80 on the T3D.

A simple and general rule is obvious: as the number of processors increases, the best panel size is decreasing. On the SP1, a panel of 128 is still the best for $p = 2$ and $p = 4$, for all algorithms, while for $p = 32$ the best size is smaller and depends on the method. It is interesting to remark that the fan-out algorithm requires a smaller panel size to reach the optimum, than fan-in methods (going down to 16, for certain matrices); the same situation occurs when fan-in on grid is compared to proportional mapping. Figure 14 presents the execution times of four algorithms for B25_K, $p = 32$; we may affirm that curves shape is representative. FO_GRID and FL_GRID have similar behavior for small panels, but FL_GRID is less affected by large panel size. FLPROP_G has also a good behavior for large panel size (due probably to the greedy strategy). On the contrary, FLPROP_SG_WW is not so much affected by small panels, when the wrap mapping preserves load balancing; for large




Figure 14: Execution times (in seconds) for B25_K, $p = 32$ and variable panel size, on the SP1.

panels, the same mapping may cause "accidents" (however, the degradation is not usually as important as in figure 14).

On the T3D, the variation of optimal panel size is not so large, when the number of processors is increased. This is natural, if we remember the sequential performance curve from figure 10. For small number of processors, the best panel size is from 40 to 64, depending on the matrix; for $p \geq 16$, the best panel size is 32 or 40. For a fixed p , FL_CUBE_G and FL_CUBE_SG_WW have similar behaviors to FL_PROP_G and FL_PROP_SG_WW, respectively.

7.4 Communication volume

We argumented some of our strategies with intuitive evaluations of communication volume. Using the same example – B25_K and $p = 32$ – we present in figure 15 the number of messages and the communication volume of the six discussed algorithms.

It results – and the fact is true for the other matrices and other number of processors – that fan-in generally implies less communication than fan-out; the exception is FL_CUBE_G. Other remarks are that proportional mapping reduces communication with respect to grid mapping and that the subgrid wrapping schemes reduce communication with respect to greedy mappings. Moreover, FL_GRID, FL_PROP_G and FL_CUBE_G manifest a decrease of communication when panel size is increased, which partially motivates the good behavior of these methods for large panels.




Figure 15: Average number of messages and communication volume (in Kbytes) per processor for B25_K, $p = 32$ and variable panel size.

7.5 Ordering effects

We will present here the behavior of FLGRID on the matrices for which METIS ordering is sequentially less efficient than the approximate minimum degree one (in terms of flop count for the Cholesky factorization), i.e. B17 and B25. As we mentioned, for minimum degree, the proportional mapping is not very effective; however, FLGRID remains better than FO_GRID; on the T3D, FL_CUBE_SG_WW is also a good alternative to the fan-out algorithm.

Figure 16 presents the execution times (Mflops performance is not relevant, since the sequential flop counts depend on the ordering) on the SP1. For both matrices, METIS allows better execution times for $p \geq 8$. A similar situation occurs on the T3D, but for $p \geq 16$.

For all other test matrices, the METIS ordering is better for any number of processors.

8 Conclusions and future work

We presented a fan-in algorithm for computing the sparse Cholesky factorization, using several 2D block mappings. The algorithm FLGRID, using a grid mapping, is usually better than Rothberg’s fan-out algorithm. Even better is the FI_PROP_G, based on proportional mapping of columns and a greedy mapping inside a column.

Experiments on a IBM SP1 showed that mapping design aiming communication volume reduction is successful, since proportional mapping was proved to be the best. On a Cray T3D, where communication is cheaper, the differences between the proposed algorithms are attenuated; while the fan-in principle remain better than fan-out, the subforest-to-subcube




Figure 16: Comparisons of FL_GRID execution times (in seconds) for different ordering algorithms.

mappings become an alternative for the proportional mapping, especially for a small number of processors.

Although our algorithms proved their efficiency, there is still place to improvements and comparisons. We will also continue to investigate the possibility of finding better mappings and to test those already proposed on larger matrices and on a greater number of processors.

We are also working in the implementation of the sparse Cholesky factorization with the ATHAPASCAN 1 programming interface [6], currently under development in the Apache projet. This programming interface makes a clear distinction between the application description and the run time scheduling policy. Additionally, it is possible to chose the better policy for a specific application and a target machine architecture. In the ATHAPASCAN 1, we intend to integrate some of the various mapping techniques presented in this report for the sparse Cholesky bi-dimensional factorization algorithm in order to compare this approach with the MPI programming model used in this paper.

References

- [1] P.R. Amestoy, T.A. Davis, and I.S. Duff. An Approximate Minimum Degree Ordering Algorithm. *SIAM J.Matrix Anal.Appl.*, 17(4):886–905, October 1996.
- [2] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide, Second Edition*. SIAM, 1995.
- [3] C. Ashcraft, S.C. Eisenstat, and J.W.H. Liu. A Fan-in Algorithm for Distributed Sparse Numerical Factorization. *SIAM J.Sci.Stat.Comput.*, 11(3):593–599, May 1990.

-
- [4] C. Ashcraft and R. Grimes. The Influence of Relaxed Supernode Partitions on the Multifrontal Method. *ACM Trans.Math.Soft.*, 15(4):291–309, December 1989.
 - [5] C. Ashcraft, R. Grimes, J. Lewis, B. Peyton, and H. Simon. Progress in Sparse Matrix Methods for Large Linear Systems on Vector Supercomputers. *Internat.J.Supercomput.Appl.*, 1:10–29, 1987.
 - [6] Gerson Cavalheiro and Mathias Doreille. ATHAPASCAN: A C++ library for parallel programming. In *Stratagem'96*, page 75, Sophia Antipolis, France, July 1996. INRIA.
 - [7] J.J. Dongarra, J. Du Croz, S. Hammarling, and I. Duff. A Set of Level-3 Basic Linear Algebra Subprograms. *ACM Trans.Math.Software*, 16:1–17,18–28, 1990.
 - [8] I.S. Duff. Sparse numerical linear algebra: direct methods and preconditioning. Technical Report TR/PA/96/22, CERFACS, 1996.
 - [9] L. Facq and J. Roman. Distribution par bloc pour une factorisation parallèle de Cholesky. In Authié, G. and al., editor, *Parallélisme et applications irrégulières*, pages 135–147. Hermès, 1995.
 - [10] G.A. Geist and E. Ng. Task Scheduling for Parallel Sparse Cholesky Factorization. *Internat. J. Parallel Programming*, 18:291–314, 1989.
 - [11] A. George, M.T. Heath, J. Liu, and E. Ng. Sparse Cholesky Factorization on a Local-Memory Multiprocessor. *SIAM J.Sci.Stat.Comput.*, 9(2):327–340, March 1988.
 - [12] A. Gupta, G. Karypis, and V. Kumar. Highly Scalable Parallel Algorithms for Sparse Matrix Factorization. Technical Report 94-63, Department of Computer Science, University of Minnesota, Minneapolis, 1994.
 - [13] M.T. Heath, E. Ng, and B.W. Peyton. Parallel Algorithms for Sparse Linear Systems. *SIAM Review*, 33(3):420–460, September 1991.
 - [14] L. Hulbert and E. Zmijewski. Limiting Communication in Parallel Sparse Cholesky Factorization. *SIAM J.Sci.Stat.Comput.*, 12(5):1184–1197, September 1991.
 - [15] G. Karypis and V. Kumar. METIS – Unstructured Graph Partitioning and Sparse Matrix Ordering System, version 2.0. Technical report, Department of Computer Science, University of Minnesota, Minneapolis, 1995.
 - [16] J.W.H. Liu. The Role of Elimination Trees in Sparse Factorization. *SIAM J.Matrix Anal. Appl.*, 11(1):134–172, January 1990.
 - [17] Message Passing Interface Forum. *MPI: A Message Passing Interface Standard*. 1995.
 - [18] E. Ng and B.W. Peyton. Block Sparse Cholesky Algorithms on Advanced Uniprocessor Computers. *SIAM J.Sci.Comput.*, 14(5):1034–1056, September 1993.

- [19] A. Pothen and C. Sun. A Mapping Algorithm for Parallel Sparse Cholesky Factorization. *SIAM J.Sci.Comput.*, 14(5):1253–1257, September 1993.
- [20] E. Rothberg. Performance of Panel and Block Approaches to Sparse Cholesky Factorization on the iPSC/860 and Paragon Multicomputers. *SIAM J.Sci.Comput.*, 17(3):699–713, May 1996.
- [21] E. Rothberg and A. Gupta. An Efficient Block-oriented Approach to Parallel Sparse Cholesky Factorization. *SIAM J.Sci.Comput.*, 15(6):1413–1439, November 1994.
- [22] E. Rothberg and R. Schreiber. Improved Load Distribution in Parallel Sparse Cholesky Factorization. In *Supercomputing '94*, pages 783–792, 1994.
- [23] R. Schreiber. A New Implementation of Sparse Gaussian Implementation. *ACM Trans. Math. Software*, 8(3):256–276, September 1982.
- [24] R. Schreiber. Scalability of Sparse Direct Solvers. In A. George, J.R. Gilbert, and J.W.H. Liu, editors, *Graph Theory and Sparse Matrix Computation*, pages 191–209. The IMA Volumes in Mathematics and its Applications, Volume 56, 1993.



Unit e de recherche INRIA Lorraine, Technop le de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS L ES NANCY
Unit e de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unit e de recherche INRIA Rh ne-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unit e de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unit e de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

 diteur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399