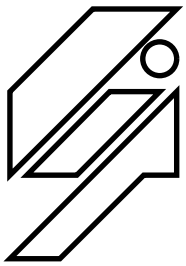

Laboratoire de l'Informatique du Parallélisme

Ecole Normale Supérieure de Lyon
Unité de recherche associée au CNRS n°1398

Optimization of the ScaLAPACK LU Factorization Routine Using Communication/Computation Overlap

F. Desprez, S. Domas and
B. Tourancheau

July 12, 1996

Research Report N^{O} 96-17

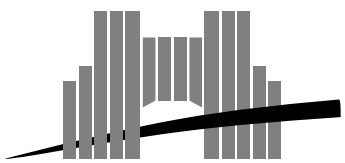

Ecole Normale Supérieure de Lyon
46 Allée d'Italie, 69364 Lyon Cedex 07, France
Téléphone : (+33) 72.72.80.00 Télécopieur : (+33) 72.72.80.80
Adresse électronique : lip@lip.ens–lyon.fr

Optimization of the ScaLAPACK LU Factorization Routine Using Communication/Computation Overlap

F. Desprez, S. Domas and
B. Tourancheau

July 12, 1996

Abstract

This paper presents some works on the ScaLAPACK LU factorization. First, a complexity analysis is given. It allows to compute the optimal block size for the block scattered distribution used in ScaLAPACK LU. It also gives the communication phases that are interesting to overlap. Second, two optimizations based on computations/communications overlap are given with experimental results on Intel Paragon system.

Keywords: LU, ScaLAPACK, communication/computation overlapping

Résumé

Ce papier présente différents travaux sur la factorisation LU de ScaLAPACK. En premier lieu, une analyse de complexité est donnée. Elle permet de calculer la taille de bloc optimale pour la distribution bloc cyclique utilisée dans ScaLAPACK. Elle donne également les phases de communications intéressantes à recouvrir. En second lieu, deux optimisations utilisant les recouvrements calculs/communications sont données ainsi que les résultats expérimentaux obtenus sur Intel Paragon.

Mots-clés: LU, ScaLAPACK, recouvrement calcul/communication

Optimization of the ScaLAPACK LU Factorization Routine Using Communication/Computation Overlap

F. Desprez*, S. Domas† and B. Tourancheau*

LIP, CNRS URA 1398, INRIA Rhône-Alpes
Ecole Normale Supérieure de Lyon
46, Allée d'Italie
69364 LYON cedex 07
(desprez,sdomas,btouranc)@lip.ens-lyon.fr

1 Introduction

The *LU* factorization is the kernel of many applications. Thus, the importance of optimizing this routine has not to be proven because of the increasing demand of applications dealing with large matrices. Its efficient parallel implementation can bring real improvements in the execution speed of the whole application. The speed-up depends greatly on the kind of supercomputer chosen. Vector machines have very high performances, but a prohibitive cost. Distributed memory machines seem to be a good balance between performances and cost.

Portability is one of the key issue of computer programming. Many libraries have been designed to ensure portability and performances across multiple architectures. The BLAS [4, 5, 10] and LAPACK [6] are available on many platforms, provided by computers vendors. *LU* factorization was released in the LAPACK package, using levels 1, 2 and 3 BLAS. ScaLAPACK [1] is the parallel version of a subset of LAPACK. ScaLAPACK has been designed to ensure portability, performances and ease of use across many parallel machines. Matrices are distributed in a block scattered way. Parallelism is hidden in a parallel version of the BLAS called PBLAS [9]. Communications between processors on a virtual grid are done using the BLACS package [7].

The aim of this paper is to show that improvements can be obtained in the existing ScaLAPACK *LU* factorization routine by the use of overlap techniques.

The first section presents the parallel block *LU* decomposition. The second section presents a recent bibliography on the parallel *LU* factorization. It contains a selection of articles that reflects the different solutions for this problem. The third section presents a complexity study of the ScaLAPACK *LU* routines, and two possible optimizations based on communication / computation overlap. The main goal of the complexity analysis is to provide the optimal block size for the block scattered decomposition, but it also gives the communication phases that are interesting to overlap.

*This work has been supported by the INRIA Rhône-Alpes and the EUREKA-EUROTOPS project.
†This work has been supported by the EUREKA-EUROTOPS project.

2 Parallel block LU decomposition

In ScaLAPACK, the parallel LU factorization uses a block scattered decomposition of matrix A on a $P \times Q$ processors grid. The $M \times N$ matrix is divided in square blocks $(r \times r)$. Thus, each processor owns a local matrix with $\lceil \frac{M}{P \times r} \rceil \times \lceil \frac{N}{Q \times r} \rceil$ blocks. This distribution is really suitable for the block LU decomposition and it provides a good load balancing between processors.

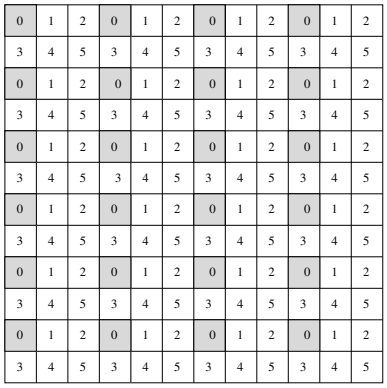

Figure 1: Block scattered decomposition of a 12r $\times$ 12r matrix on a 2×3 grid.

Figure 1 shows how to distribute the $r \times r$ blocks of a $12r \times 12r$ matrix on a 2×3 grid in a block scattered way. Squares marked with a number represent a single block of the global matrix. For example, all grey blocks belong to processor 0.

The block LU decomposition consists in three phases, repeated as many times as there are block columns to be factorized in the global matrix: instead of working on a single column of the global matrix A at a time, r columns are factored at each step. And for convenience, the local L and U matrices are stored in place of the matrix to be decomposed.

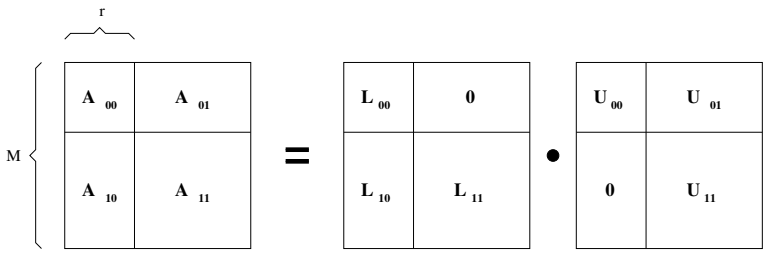

Figure 2: First step of the block LU factorization.

According to Figure 2, the three steps are the followings:

- in order to obtain (L_{00}, L_{10}) and U_{00}, a simple Gaussian elimination is computed on $\binom{A_{00}}{A_{10}}$.

$$L_{00}.U_{00} = A_{00} \tag{1}$$

$$L_{10}.U_{00} = A_{10} \tag{2}$$

- the U_{01} block is obtained by a triangular solve (see equation 3).

```
pcol = 0
prow = 0
for k = 0 to min(M_b, N_b) − 1 by step r do
    for i = 0 to r − 1 do                                          ⎫
       if (my_col = pcol) then                                     ⎪
          find pivot and its position                              ⎪
       end if                                                      ⎪
       broadcast the two values to all processors                  ⎪
       exchange pivot rows                                         ⎬ phase 1
       if (my_col = pcol) then                                     ⎪
         divide under-diag.  elts.  of col.  i by pivot            ⎪
         update col.  i+1 to r-1 /* _GER */                       ⎪
       end if                                                      ⎪
    end for                                                        ⎭

    if (my_row = prow) then                                        ⎫
       broadcast L_00 to all processors of the prow row            ⎬ phase 2
       solve L_00.U_01 = A_01 /* _TRSM */                         ⎪
    end if                                                         ⎭

    broadcast L_10 on processors rows                              ⎫
    broadcast U_01 on processor columns                            ⎬ phase 3
    update A_11 ← A_11 − L_10.U_01 /* _GEMM */                     ⎭

    pcol = (pcol + 1) mod Q
    prow = (prow + 1) mod P
end for
```

Figure 3: Parallel block LU factorization using a block scattered data distribution.

$$L_{00}.U_{01} = A_{01} \tag{3}$$

- $L_{11}.U_{11}$ is given by equation 4. A matrix product is needed $(A_{11} - L_{10}.U_{01})$.

$$L_{10}.U_{01} + L_{11}.U_{11} = A_{11} \tag{4}$$

The three steps are recursively computed on $L_{11}.U_{11}$ to obtain L_{11} and U_{11}.

The general parallel algorithm is given in Figure 3. `_GER` is a level 2 BLAS, `_TRSM` is the level 3 BLAS triangular solve routine and `_GEMM` is the general matrix product routine.

The ScaLAPACK routines `PDGETRF` and `PDGETF2` execute the block LU factorization on a matrix distributed in a block scattered way. Figures 4 and 5 represent the execution schemes of these two routines :

- `PDGETF2` performs the factorization of a block column to compute L_{00},L_{10}, and U_{00} (phase 1 of the general algorithm in figure 3).
- `PDGETRF` calls `PDGETF2`, then updates the remaining blocks of the matrix to compute $L_{11}.U_{11}$, U_{01} (phase 2 and 3 of the general algorithm in figure 3).

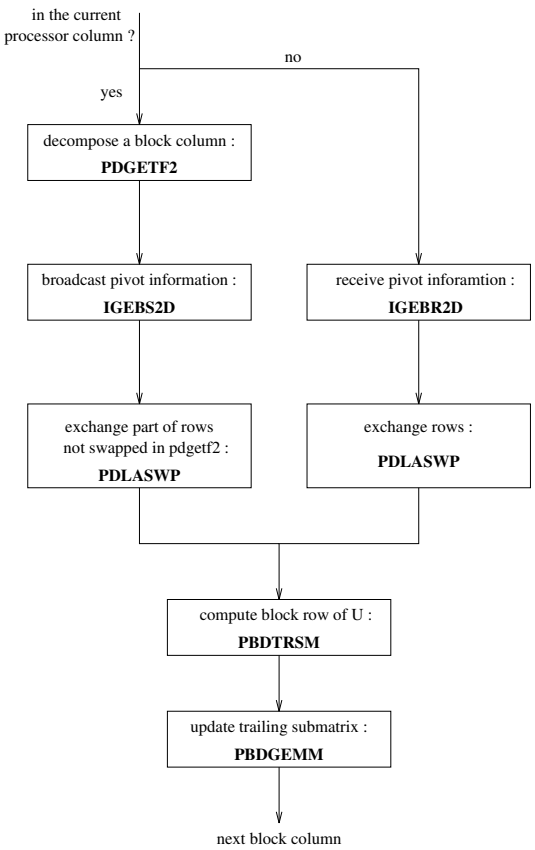

Figure 4: Execution scheme of the `PDGETRF` routine.

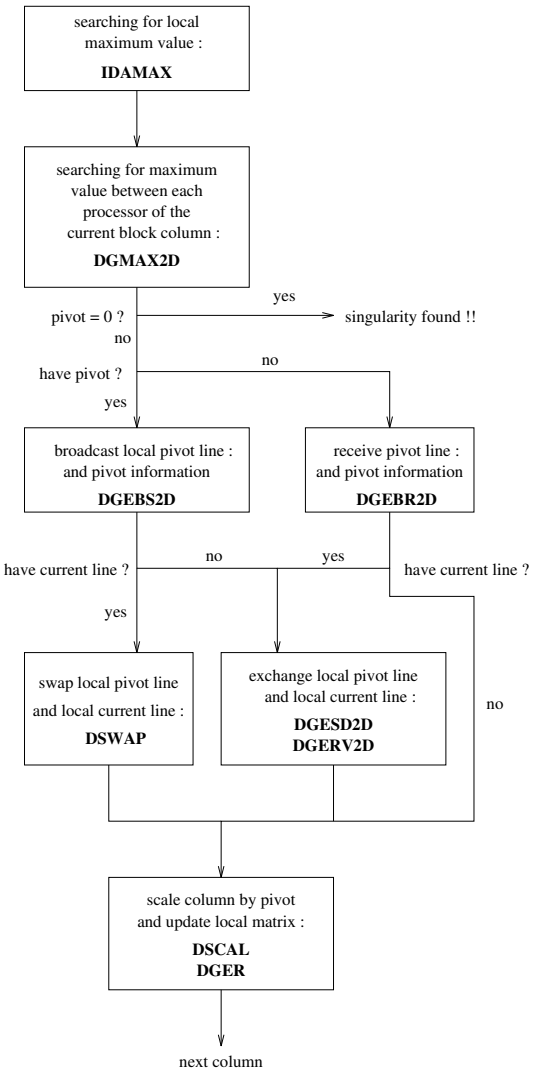

Figure 5: Execution scheme of the `PDGETF2` routine.

3 Related work

Various methods have been proposed to improve the parallel LU factorization. The corresponding papers present experiments which are sometimes confirmed by complexity studies.

In [8], a row oriented method is presented with an efficient pivot selection. It uses a reverse spanning tree broadcast to choose the real pivot among all the local pivots, owned by different processors. In ScaLAPACK, it is now achieved by the BLACS operation `_GMAX2D`. Furthermore, in [2], a load balancing strategy for the choice of the pivot is added to the row oriented method with a low cost row interchange.

Meanwhile, a row distribution supposes that the real pivot is often chosen on a different processor than the owner of the current row. Thus, there are some communications to exchange and to broadcast the pivot row.

In [11], a column-scattered distribution is chosen with the well-known pipelined ring algorithm. A complete performance estimation is provided for distributed memory parallel machines, and compared to experimental results.

In [3], a column-scattered distribution is chosen and asynchronous communications are used to overlap computations (columns updates). A theoretical model is given for the complexity on a complete network

and a ring topology.

The two last methods are quite efficient because choosing the pivot on a single processor (instead of all, like in [8, 2]) is very fast and exchanging rows occurs in local memory. Therefore, it uses a fine grain parallelism.

ScaLAPACK LU is a block LU decomposition with a block scattered distribution of the data. In [9], a complexity study and different assumptions on the choice of the block size and the grid size are proposed. It appears that a rectangular grid with few rows of processors is optimal (for the reasons exposed upper). Thus, it can use the best of the precedent methods but with coarse grain computations.

The block scattered distribution used in the ScaLAPACK LU factorization seems to be the most interesting solution for distributed memory computers such as the Paragon system, as far as the block and the grid sizes are well chosen.

Consequently, this paper presents an optimized version of LU factorization with a block scattered distribution and an overlap between communications and computations. It also gives a method to compute an optimal block size without an "enquiry routine" (as previewed in [9]).

4 Analysis and optimization of the ScaLAPACK LU factorization

4.1 Complexity analysis

A prediction of the execution time is important to confirm the experimental results. The value of the parameters which influence the execution time can be deduced form the model. Furthermore, it gives the execution time of the communication phases. Such predictions are very interesting to choose what communications to overlap.

In the ScaLAPACK LU factorization, performances depend greatly on the block size used for the block scattered decomposition. Thus, it is interesting to compute the optimal block size to avoid a number of tests. It is also interesting to compute automatically the best data distribution (for parallel compilers, for example).

Figure 6 shows the impact of the block size on the performance of the ScaLAPACK LU factorization. On the Intel Paragon at the University of Lyon, a size of 64 is less efficient than 8, which is the optimal value.

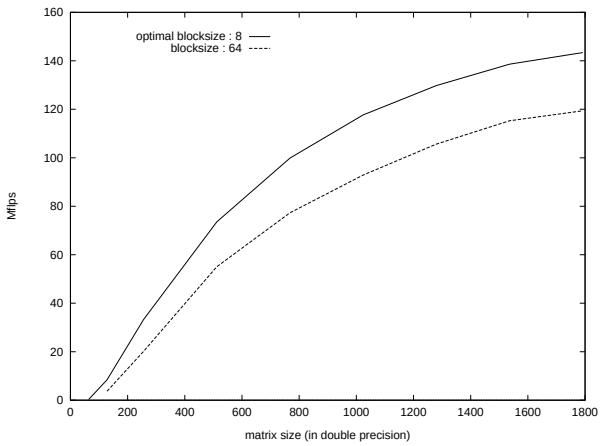

Figure 6: Performance (In Mflops) of LU factorization on 4 processors with two different block sizes.

4.1.1 The model

In this section, the theoretical optimal block size S_b is obtained by an interpolation, based on experimental measurements (for a given supercomputer) of the LU subroutines. For all subroutines function of S_b (see below), the execution time is expressed literally, and then derived to find the optimal S_b.

The complexities are given below with S_b as the block size, B_c the number of block columns (or rows) in a single processor (i.e. $\frac{M}{P_c \times S_b}$), P_c the number of processor columns (or rows) on the grid (the grid is assumed to be square (P_c^2 processors) in order to simplify the equations below) and M the matrix size. The subroutines names are the BLAS or LAPACK names expressed in Figures 4 and 5.

4.1.2 Subroutines used in `PDGETF2`

- `IDAMAX:` is executed S_b times for each `PDGETF2` call. There are B_c `PDGETF2` calls for each processor (Figures 1 and 4), and thus $B_c \times P_c \times S_b$ in total. The execution time of a single `IDAMAX` call is linear.

$$T_{idamax} = \sum_{i=1}^{B_c} P_c.S_b.[a_{idamax}(S_b \times i) + b_{idamax}]$$

- `DCOPY:` is executed $2 \times S_b$ times for each `PDGETF2` call, with a fixed data size S_b. Its first use is to copy the current pivot segment to broadcast it. Its second use is to copy the real pivot segment into local matrix. The execution time of a single `DCOPY` call is linear.

$$T_{copy} = 2 \sum_{i=1}^{B_c} P_c.S_b.[a_{copy}S_b + b_{copy}]$$

- `DGEBS2D,DGEBR2D:` is executed S_b times for each `PDGETF2` call, with a fixed data size S_b. The execution time of a single broadcast is linear. It is considered to be an hypercube broadcast (see the BLACS users' guide [7]), proportional to the basic communication time: $t_{broadcast}(L) = \lfloor log_2 P_c \rfloor.[a_{sendd}L + b_{sendd}]$ [1]

$$T_{broadcastd} = \sum_{i=1}^{B_c} P_c.S_b.[a_{broadcastd}S_b + b_{broadcastd}]$$ [2]

- `DGESD2D,DGERV2D:` is executed S_b times for each `PDGETF2` call, with a fixed data size S_b. The execution time of a single send is linear.

$$T_{send} = \sum_{i=1}^{B_c} P_c.S_b.[a_{sendd}S_b + b_{sendd}]$$

- `DSCAL:` is executed S_b times for each `PDGETF2` call. The execution time of a single `DSCAL` call is linear.

$$T_{dscal} = \sum_{i=1}^{B_c} P_c.S_b.[a_{dscal}(S_b \times i) + b_{dscal}]$$

- `DGER:` is executed S_b times for each `PDGETF2` call, with a decreasing data size. The execution time of a single `DGER` call is quadratic. It is a level 2 BLAS which computes $A \leftarrow \alpha x.y^t + A$ with $A_{(m \times n)}$:

$$t_{dger}(i,n) = (\alpha_{dger} \times S_b).i.n + \gamma_{dger} \times n + \lambda_{dger}$$

$$T_{dger} = \sum_{i=1}^{B_c} P_c. \sum_{j=1}^{S_b - 1} [(\alpha_{dger} \times j)(S_b \times i) + (\gamma_{dger} \times j + \lambda_{dger})]$$

[1] a_{sendd} is expressed for double precision buffer send. For integers, the a_{sendi} notation is used.

[2] $a_{broadcastd}$ is expressed for double precision buffer broadcast. For integers, the $a_{broadcasti}$ notation is used.

4.1.3 Subroutines used in PDGETRF

- `IGEBS2D,IGEBR2D:` is executed $B_c \times P_c$ time, with a fixed data size S_b.
  $T_{broadcasti} = B_c.P_c.[a_{broadcasti}S_b + b_{broadcasti}]$

- `PBDTRSM:` is executed $B_c \times P_c$ time, with a variable data size, multiple of S_b. It must be decomposed in two parts:

  1. **broadcasting:** the factored column can be broadcast in one chunk in order to distribute the current L_{00} block (see figure 2) for `DTRSM` computation, and L_{10} for later `DGEMM` computation. Before broadcasting, the whole column is copied in a working buffer.

$$\begin{aligned} T_{broadcastd} &= \sum_{i=1}^{B_c} P_c.[(a_{broadcastd} + a_{copy})(i \times S_b^2) + (b_{broadcastd} + b_{copy})] \\ &- (a_{broadcastd} + a_{copy})S_b^2 + (b_{broadcastd} + b_{copy}) \end{aligned}$$

  2. `DTRSM:` The execution time of a single `DTRSM` call is cubic. It is a level 3 BLAS which computes $B \leftarrow \alpha.A^{-1}.B$ with $A_{(m\times m)}$ and $B_{(m\times n)}$:
     $t^{S_b}_{dtrsm}(i) = a_{dtrsm}.(i \times S_b) + b_{dtrsm}$
     with:
     - $a_{dtrsm} = \alpha_{dtrsm} \times S_b^2 + \beta_{dtrsm} \times S_b + \delta_{dtrsm}$
     - $b_{dtrsm} = \gamma_{dtrsm} \times S_b + \lambda_{dtrsm}$

$$\begin{aligned} T_{DTRSM} = &\sum_{i=1}^{B_c} P_c.[a_{dtrsm}.(i \times S_b) + b_{dtrsm}] \\ &-[a_{dtrsm}.(B_c \times S_b) + b_{dtrsm}] \ ^{3} \end{aligned}$$

- `PBDGEMM:` executed $B_c \times P_c$ time, with a variable data size, multiple of S_b. It must be decomposed in two parts:

  1. **broadcasting:** the last factored row (see `DTRSM`) must be broadcast in order to achieve `DGEMM` computation. Before broadcasting, the whole row is copied in a working buffer.

$$\begin{aligned} T_{broadcastd} = &\sum_{i=1}^{B_c} P_c.[(a_{broadcastd} + a_{copy})(i \times S_b^2) + (b_{broadcastd} + b_{copy})] \\ &-[(a_{broadcastd} + a_{copy})(B_c \times S_b^2) + (b_{broadcastd} + b_{copy})] \ ^{4} \end{aligned}$$

  2. `DGEMM:` The execution time of a single `DGEMM` call is cubic. It is a level 3 BLAS which computes $C \leftarrow \alpha.A.B + \beta.C$ with $A_{(m\times k)}$, $B_{(k\times n)}$ and $C_{(m\times n)}$:
     $t^{S_b}_{dgemm}(i) = a_{dgemm}.(i \times S_b)^2 + b_{dgemm}.(i \times S_b) + c_{dgemm}$
     with:
     - $a_{dgemm} = \alpha_{dgemm} \times S_b$
     - $b_{dgemm} = \beta_{dgemm} \times S_b + \delta_{dgemm}$
     - $c_{dgemm} = \gamma_{dgemm} \times S_b + \lambda_{dgemm}$

[3] `DTRSM` is executed one less time at maximal size because of block scattered distribution and LU algorithm.
[4] `broadcastd` is executed one less time at maximal size because of block scattered distribution and LU algorithm.

$$T_{DGEMM} = \sum_{i=1}^{B_c} P_c.[a_{dgemm}.(i \times S_b)^2 + b_{dgemm}.(i \times S_b) + c_{dgemm}]$$
$$-[a_{dgemm}.(B_c \times S_b)^2 + b_{dgemm}.(B_c \times S_b) + c_{dgemm}]$$ [5]

4.1.4 Results on Intel Paragon

In order to compute the optimal block size, all the precedent equations have been derivated over S_b. The total complexity optimum is obtain when the sum of the derivatives equals zero. Hence, the optimal block size is given by the resolution of a third degree equation. We computed the three roots for different grid sizes and matrix sizes. Only one solution was positive each time. Results are given in table 1. It gives the optimal block size as a function of different grid sizes and matrix sizes (1000 signifies a 1000×1000 matrix in double precision).

grid size	**matrix size**							
	1000	2000	3000	4000	5000	6000	7000	8000
4×4	8.31	9.28	9.64	9.83	9.94	10.0	10.1	10.1
8×8	7.38	8.38	8.79	9.03	9.17	9.27	9.35	9.41
16×16	7.26	8.0	8.38	8.61	8.77	8.88	8.96	9.03

Table 1: Optimal block size.

Table 2 gives a comparison between experimental and theoretical optimal block size, on a 4×4 grid.

type	**matrix size**			
	1000	2000	3000	4000
experimental	8 - 9	8	8	10
theoretical	8.31	9.28	9.64	9.83

Table 2: Comparison between theoretical and experimental block sizes on a 4×4 grid.

Six important points have to be noticed:

- Theoretical optimal block sizes are close to the experimental ones. They range between 7.2 and 10.1 and actual tests give an optimal block size of 8 or 10 on a 4×4 grid[6].
- The theoretical optimal block size is a function of the matrix size, and it raises up to a top value around 10. But this size is a real. A "good" block size is the nearest even value from the theoretical one.
- The optimal block size hardly depends on the number of processors. Asymptotic values are 9 for 256 processors and 10 for 16. Again, experimental results confirm this point.
- The grid is assumed to be squared for convenience. According to experimental results, the grid shape has no real influence on the optimal block size, though it greatly influences the execution time [9]. A rectangular grid with few rows of processors works faster than a square grid. There are less communications since pivoting is achieved more often in local memory.

[5] `DGEMM` is executed one time less at maximal size because of block scattered distribution and *LU* algorithm.

[6] A size of 16 was found on previous tests. It has been done on the same machine but with an older version of the operating system. Results depend greatly on the machine and its system.

- Results are identical if only the subroutines `DGER`, `DTRSM`, and `DGEMM` are used for block size computation (they represent 95% of total computation time).
- All these results are only applying to Intel Paragon machines. Recent tests on an IBM SP2 show that optimal block size is varying with the matrix size.

4.2 Optimizations

The ScaLAPACK version of LU has been implemented in order to be scalable: each subroutine call is a BLAS or BLACS call. These two libraries are already fully optimized for a wide range of computers. Furthermore, the minimal number of operations to achieve a LU factorization is well-known ($\frac{2}{3}n^3 + 2n^2$). Thus, only communication phases can be optimized since the computation time is fixed. Asynchronous messages are used to overlap communications with computations.

4.2.1 Broadcast overlap

By looking closely at the algorithm, we can see that processors are often waiting results from other processors ! During the block column decomposition, only a processor column is working. And only one processor row is working during the triangular solve. This brings us to the first optimization solution:

"Instead of broadcasting (L_{00}, L_{10}) panel before the triangular solve (`_TRSM`), do it at the same time."

This means that general synchronous BLACS broadcasting routines (`_GEBS2D` and `_GEBR2D`) are used on processors that do not compute `_TRSM`, and asynchronous communications are used to send (L_{00}, L_{10}) during `_TRSM` on processors that compute it.

Therefore, the single block L_{00} must be broadcast on the current processor row to perform `_TRSM`. But it takes less time than broadcasting (L_{00}, L_{10}).

This solution seems interesting since we gain at each step i of factorization the broadcast time of $P_b - i - 1$ blocks, compared to the original version. But, unfortunately, that gain represents only 1 to 2 percents of total factorization time on the Paragon system. This is due to the number of times this broadcast is done (only $P_b - 1$ times). Moreover, this deceiving result can be predicted by the complexity analysis. The sum of broadcasting time for all steps never exceeds 2 percents of total execution time.

4.2.2 Rows pivoting overlap

It appears that a good speed-up cannot be obtained unless a communication phase is overlapped most of the time. Thus, it is interesting to overlap the pivoting time since it is executed for each row of the matrix :

"Instead of broadcasting pivot informations and then exchanging rows after the (A_{00}, A_{10}) decomposition (`PDGETF2`), do it at the same time."

This means that we can use the `DSCAL` time of the processors column which decompose (A_{00}, A_{10}), to exchange current and real pivot rows using asynchronous communications. Then we use the `DGER` time to send asynchronously the pivot informations to the next processor in the processors row. Thus, as soon as this processor receives pivot informations, it can exchange the rows for pivoting and send the information to the next processor in the row. And so on, until the last processor on the pseudo-ring receives its data. During this step, the block column decomposition continues on the processor column.

Figure 7 represents the different steps of these operations for an 64×64 matrix distributed on a 4×4 grid using a 8×8 block size. We explain this example in the following:

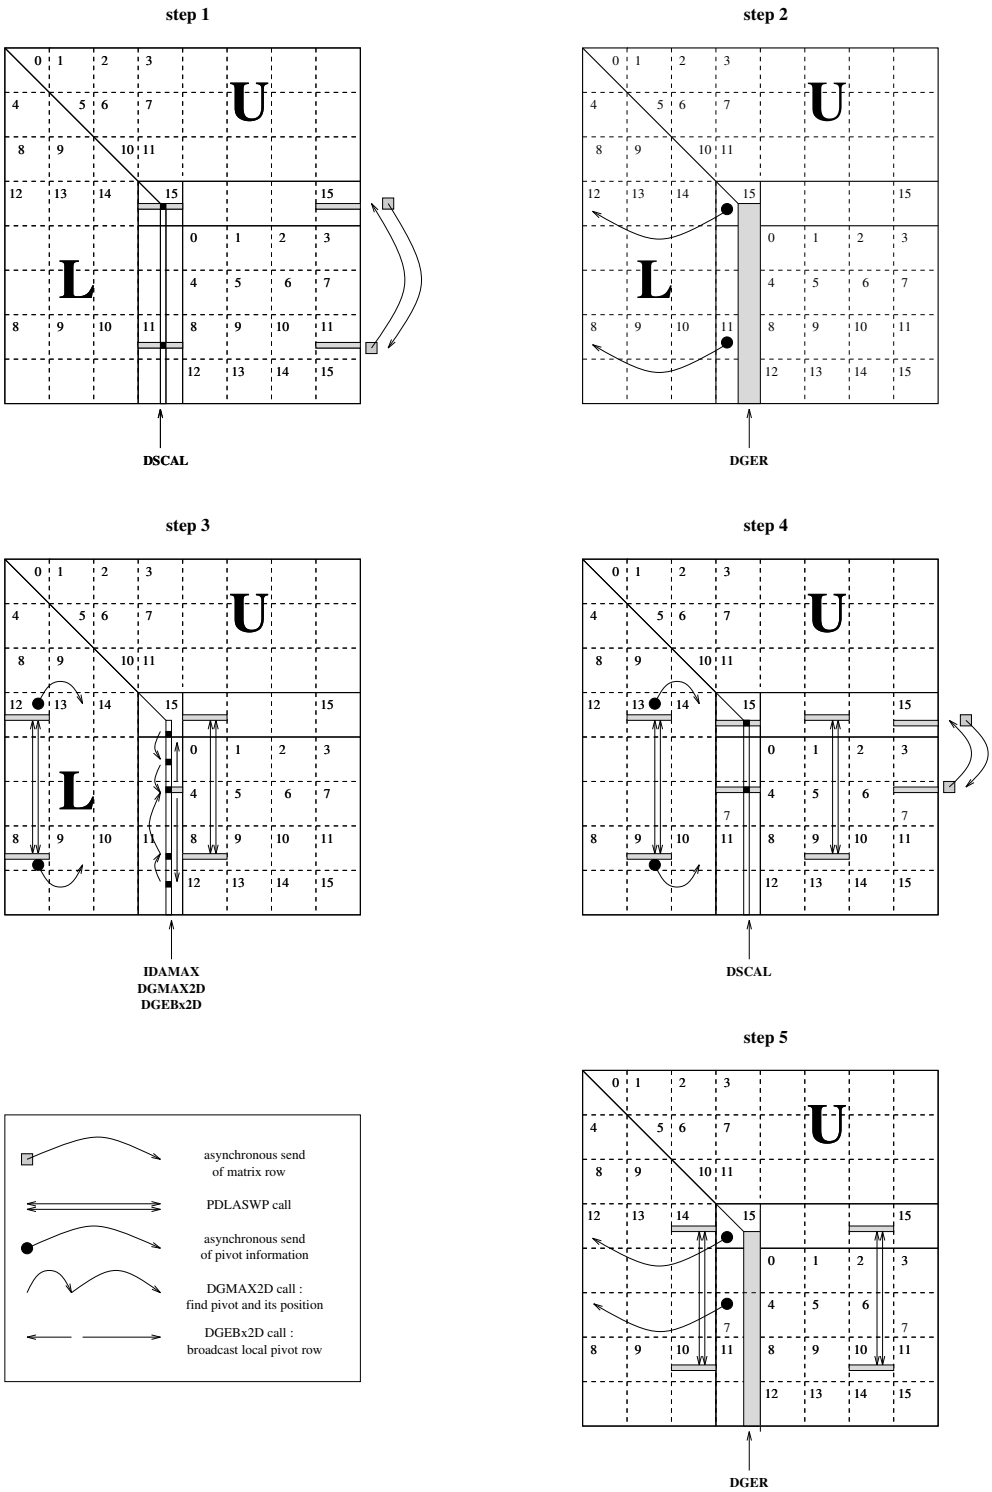

Figure 7: One iteration of optimized `PDGETF2` routine.

- **step** 1: This is the k^{th} iteration of `PDGETF2`. The real pivot row has been found on processor 11. Then, processor 15, that owns the current pivot row k, divides the current pivot by the real pivot and

asynchronously send the whole k^{th} row to processor 11. After, it proceeds with `DSCAL` and after, waits for the completion of the asynchronous send and receives the real pivot row.

Identically, processor 11 asynchronously sends the whole pivot row to processor 15 and computes the `DSCAL` at the same time. After completion, it receives the current pivot row.

In the best case, processors 15 and 11 do not need to wait for send completion since the communication is already over when `DSCAL` ends. In fact, a wait state appears with a very large matrix, when their size reaches memory size limits. In this case, the `DSCAL` time does not completely overlap the communication time, since `DSCAL` domain size decreases each step.

Other processors in the pivot processor column just execute the `DSCAL` routine.

- **step** 2: `DSCAL` and pivoting is done. Now, the local sub-matrix must be updated with `DGER`. Processors 15 and 11 use this time to send the pivot information to their right neighbor on a pseudo-ring made from the processors row. In the figure 7, a pseudo-ring is $(12, 13, 14, 15)$, and the right neighbor of 15 is 12.

  Processors 12 and 8 are just waiting for pivot information using a blocking receive.

- **step** 3: processors 12 and 8 have just received the pivot information. They can now exchange the k^{th} row and the pivot row, and send the pivot informations to their right neighbor. Meanwhile, processors in the pivot column continue the decomposition, finding the pivot and its position, broadcasting the local pivot row, ...

- **step** 4: as in step 1, `DSCAL` is overlapped by an asynchronous exchange of current and real pivot row. But now, processors 13 and 9 are working, exchanging rows, instead of waiting for the completion of `PDGETF2` to work.

- **step** 5: as in step 2, but with two more processors working.

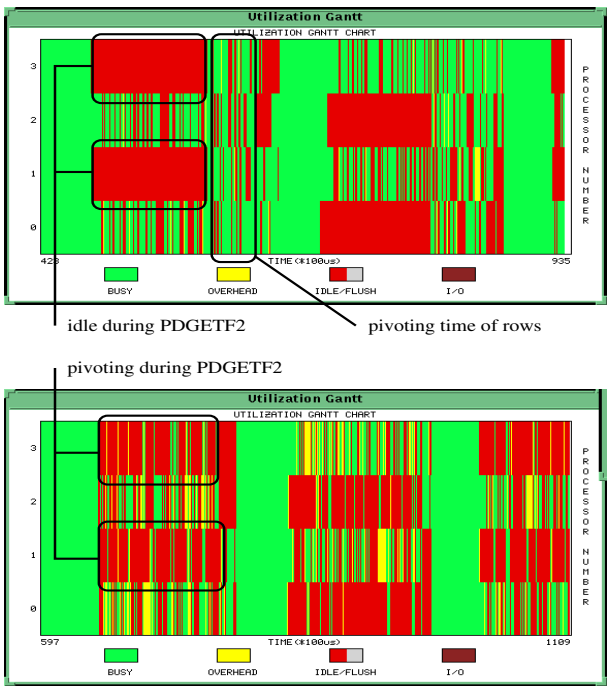

Figure 8: Comparison of the Gantt diagram for optimized and non-optimized versions.

The impact of this optimization is clearly shown on Figure 8. On the top of the figure, we can see that processors 1 and 3 are completely idle during the non-optimized `PDGETF2` execution. After, there is a communication phase to exchange the pivot row. On the bottom of the figure, processors 1 and 3 are no more idle : they are receiving the pivot information and exchanging rows during the optimized `PDGETF2` execution. Thus, there are no more communication post-phase. They are overlapped by `PDGETF2`.

4.3 Experimental results

All experimental results have been produced on Paragon systems, with various grid size, on two different machines (ORNL, Tennessee USA, and Lyon, France). These two machines have not the same operating system and it appeared that it leads to different experimental results.There are four different grid sizes: 4×4, 6×5 (Lyon), 8×8, and 16×16 (ORNL).

All the results have been successfully run with the test driver included in the ScaLAPACK LU factorization package in order to test the correctness of the optimizations. This driver completes the LU factorization on a matrix A and then solves the system $LUx = B$. It also provides different possibilities to vary the execution parameters such as block size, number of processors, number of columns of the right-hand-side (RHS) matrix B, ...

A block size of 16×16 appears to be the optimal value for the LU factorization on the ORNL system, and 8×8 on the Lyon system. So, all Mflops are given for these block sizes and one RHS vector.

Figure 9 shows a comparison between optimized and non-optimized results for 16×16 grid. The optimized version grows faster (in Mflops) than the non-optimized, before becoming almost parallel. Indeed, the optimization works better for small matrix sizes ($< 375 \times P_c$ with P_c, the number of column of processors) because the pivoting communications are all executed during `PDGETF2`. After this limit, the pivoting time becomes more important and communications cannot be completely overlapped by the `PDGETF2` execution time. This is confirmed by the complexity analysis. The Figure 10 shows the predicted and the experimental gain for a 4×4 grid. For each matrix size, the predicted gain is given by the minimum of the total time of `PDGETF2` and the total communication time for pivoting.

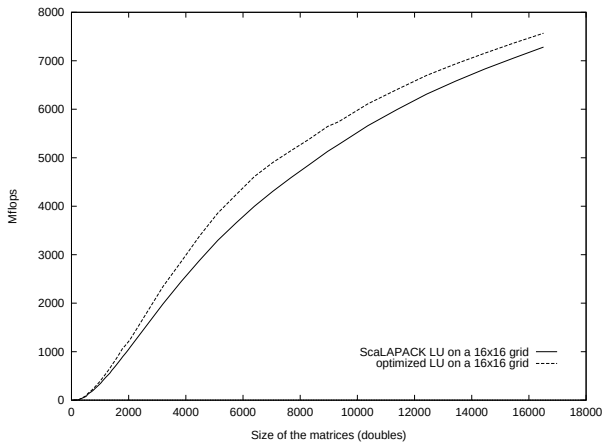

Figure 9: Experimental results on a 16×16 grid.

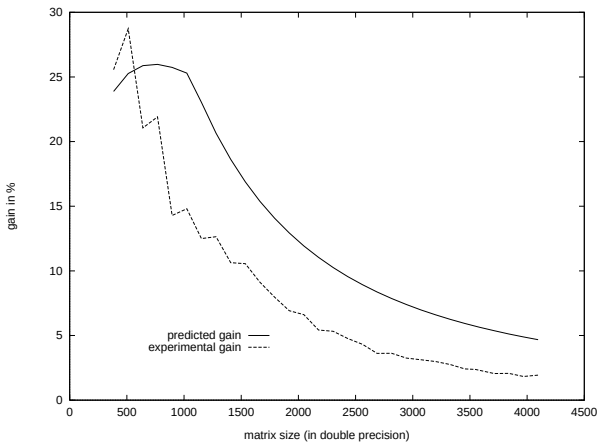

Figure 10: Predicted and experimental gain on a 4×4 grid.

Figure 11 shows the gain in percents over non-optimized version for 8×8 and 16×16 grids. It can reach 15% for small matrix sizes, and stay above 4% for the largest matrix size that can be allocated. This figure confirms that the gain progressively decreases for a matrix size greater than $375 \times P_c$.

Figure 12 is the same as figure 11 but the x-coordinates are expressed as a function of the sub-matrix size for a single processor. It shows that the gain does not depend on the total matrix size but only on the size of the sub-matrix that one processor owns. This is due to the block scattered distribution which provides a good load balancing.

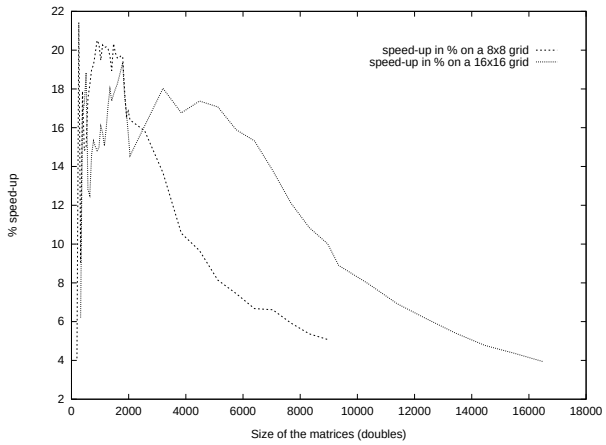

Figure 11: Gain on the 8×8 and 16×16 grids

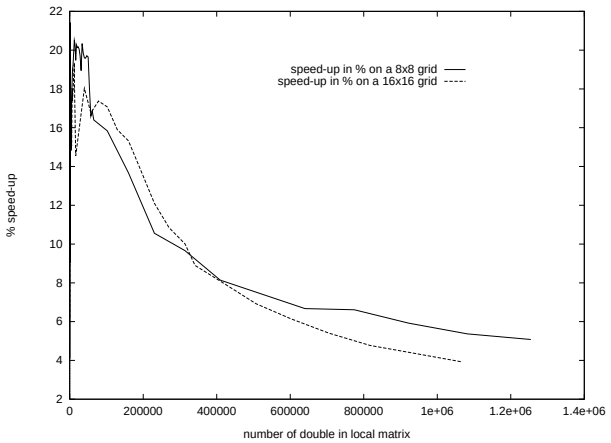

Figure 12: Gain on the 8×8 and 16×16 grids, as a function of the matrix size on a single processor.

5 Conclusion and future work

After a description of the LU algorithm in ScaLAPACK, a complete analysis of complexity has been presented. This theoretical model allows the computation of the optimal block size for the block scattered decomposition. Thus, it is possible to have the best performance with a simple pre-computation.

The second part has presented two optimizations based on communication / computation overlap. In the two cases, our aim was to "hide" the time of some big communication phases. Furthermore, it allows to reduce the idle time of some processors that are waiting results from other to continue the execution.

All experimentations have been done on Paragon systems but the methods presented in this paper are general. Thus, they can be applied to any supercomputer.

Meanwhile, some hints can be given as future work:

- All the updates have been done using Paragon system calls syntax because asynchronous BLACS do not exist. Consequently, our code is not fully portable. But the modifications are simple enough to be rewritten on any supercomputer. For example, the code has been recently ported on an IBM SP2 in about five minutes, and it will soon be ported on a CRAY T3D.

- the gain decreases as the matrix size grows: the `PDGETF2` time becomes not big enough to overlap the communications due to pivoting. Thus, another routine could be used to overlap more communications (like `DGEMM`).

- the optimal block size computation is also interesting as input for parallel compilers because it gives the best data distribution for a given matrix size, and number of processors. It would be interesting to include such a computation in a HPF compiler.

References

[1] E. Anderson, A. Benzoni, J. Dongarra, S. Moulton, S. Ostrouchov, B. Tourancheau, and R. Van de geijn. Lapack for distributed memory architecture progress report. In *Fifth SIAM Conference on Parallel Processing for Scientific Computing,*, 1991.

[2] E. Chu and A. George. Gaussian Elimination with Partial Pivoting and Load Balancing on a Multiprocessor. *Parallel Computing*, 5:65–74, 1987.

[3] F. Desprez, J.J. Dongarra, and B. Tourancheau. Performance Complexity of LU Factorization with Efficient Pipelining and Overlap on a Multiprocessor. *Parallel Processing Letters*, 5-II, 1995.

[4] J.J. Dongarra, J. Du Croz, I. Duff, and S. Hammarling. A Set of Level 3 Basic Linear Algebra Subprograms. *ACM Transaction on Mathematical Software*, 16(1):1–17, 1990.

[5] J.J. Dongarra, J. Du Croz, S. Hammarling, and R. Hanson. An Extended Set of Fortran Basic Linear Algebra Subroutines. *ACM Transaction on Mathematical Software*, 14(1):1–17, March 1988.

[6] J.J. Dongarra, R. Van De Geijn, and D.W. Walker. A Look at Dense Linear Algebra Libraries. Technical Report ORNL/TM-12126, Oak Ridge National Laboratory, July 1992.

[7] J.J. Dongarra, R.A. Van De Geijn, and R.C. Whaley. A User's Guide to the BLACS, March 1993.

[8] G.A. Geist and M.T. Heath. Matrix Factorization on a Hypercube Multiprocessor. Technical report, Oak Ridge National Laboratory, 1985.

[9] J.Choi, J. Demmel, I. Dhillon, J. Dongarra, S. Ostrouchov, A. Petitet, K. Stanley, D. Walker, and R.C. Whaley. LAPACK Working Note: ScaLAPACK: A Portable Linear Algebra Library for Distributed Memory Computers - Design Issues and Performances. Technical Report 95, Department of Computer Science - University of Tennessee, 1995.

[10] C. Lawson, R. Hanson, D. Kincaid, and F. Krogh. Basic Linear Algebra Subprograms for Fortran Usage. *ACM Transaction on Mathematical Software*, 5:308–323, 1979.

[11] B.V. Purushotham, A. Basu, P.S. Kumar, and L.M. Patnaik. Performance Estimation of LU Factorisation on Message Passing Multiprocessors. *Parallel Processing Letters*, 2(1):51–60, 1992.

Contents

1 Introduction **1**

2 Parallel block LU decomposition **2**

3 Related work **4**

4 Analysis and optimization of the ScaLAPACK LU factorization **5**
4.1 Complexity analysis 5
4.1.1 The model 6
4.1.2 Subroutines used in `PDGETF2` 6
4.1.3 Subroutines used in `PDGETRF` 7
4.1.4 Results on Intel Paragon 8
4.2 Optimizations 9
4.2.1 Broadcast overlap 9
4.2.2 Rows pivoting overlap 9
4.3 Experimental results 12

5 Conclusion and future work **13**