Modeling the LU Factorization for SMP Clusters

Jack Dongarra
ICL, U. Tennessee

Emmanuel Jeannot
ICL, U. Tennessee, INRIA

Julien Langou
ICL, U. Tennessee

1 Introduction

In this paper we target the problem of modeling the LU factorization in the context of a 2-D block-cyclic distribution.

Modeling the LU factorization is an important challenge as it can help to understand the scalability of the algorithm and helps the user to compute both the best block size and the optimal processor grid-size.

Modeling the LU factorization has been addressed in several prior works [1, 2, 3].

In [1], a very crude model is proposed. It is simply used for showing the scalability of ScaLapack. It does not use the processor grid shape whereas it is well known that it has a great impact on the performance.

In [2], the authors propose a very fine model with multiple parameters for the environment (more than 30!). They use this model to compute the optimal grid size. However, they do not show how to compute these parameters.

In [3], the authors propose a simple model with only three parameters for the environment (the bandwidth and the latency of the network and the flop rate of the processor). However, our experiments show that the value of these parameters depends on the phase of the algorithm. For instance a better floprate is achieved when updating the trailing sub-matrix than when updating the panel.

It is important to note that none of the above models take into account the SMP clusters architecture where each processor can be grouped on a multi-processor node.

In order to address the above problem we propose a model that has three environment parameters per subroutine. As it takes into account the grid-size, it is able to derive the optimal grid shape in most of the cases. Moreover, we propose a simple and fast method to compute the model parameters. Finally we show how to enhance our model for multiprocessor clusters.

2 A model of the LU factorization for multiprocessor architecture

The LU factorization is decomposed in steps, which are repeatedly executed up to the end of the decomposition. Each step is composed of several substeps. In ScaLapack, each of these substeps is implemented by a parallel routine. We study and model the cost of each of these substeps in order to build a model of the whole LU factorization. Hence we will use the following parameters: N is the number of rows and columns of the matrix. We assume that the blocks of the distribution are squared. NB is the block-size. The matrix is distributed among a processor grid of size $P \times Q = NP$ For each substep `xxx`, we have three environments parameters: γ_{xxx} is the time to perform one operation within the subroutine that implements this substep; we call β_{xxx} the latency of the network for the subroutine and α_{xxx} the time to transfer one matrix element for this subroutine[1].

3 Computing the parameters of the model

For each substep we have three parameters. Computing this efficiently and quickly is an important issue. In [3], the authors suggest using the characteristics of the machine (flop rate) and of the network (latency and bandwidth). However, our experiments show that for each substep the parameters have different values. Our methodology for computing these parameters consists i-of three steps. First we perform several runs of the LU factorization on the target architecture. Second we time each subroutine during these runs and sum the duration of each[2]. Third, based on these timings, we perform a parameter evaluation using a least square method. We have been able to limit the number of runs to five and still have a precise parameter evaluation as shown in the next section.

4 Results

The experiments were performed on a dual-processor Opteron with Giga-ethernet interconnect. Computing the grid size for N=30000 and up to 64 processors with our model shows that we are able to find the optimal grid-size in 68% of the cases. For the other cases, the proposed grid-size always has a running time lower than 6% of the optimal one. The figure on the right shows the histogram of error of our model (log scale) for LU factorization with NB=32, $N \in [10000, 20000, 30000]$ and up to 64 processors. The average of est.time/time is 1.07 and the variance is 1.17.

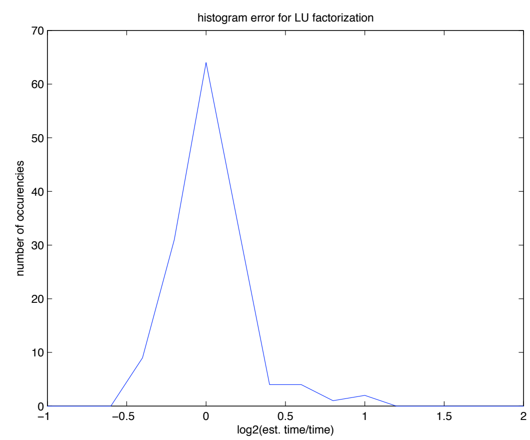

References

[1] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*. SIAM, 1997.

[2] S. Domas, F. Desprez, and B. Tourancheau. Optimization of the ScaLAPACK LU Factorization Routine Using Communication/Computation Overlap. In *Europar'96 Parallel Processing*, volume 1124 of *Lecture Notes in Computer Science*, pages 3–10. Springer Verlag, August 1996.

[3] J. Dongarra, R van de Geijn, and D Walker. A Look at Scalable Dense Linear Algebra Libraries. LAPACK Working Note 43, University of Tennessee, may 1992.

[1] Due to lack of space, we cannot show how to compute the model and even the model itself, it will be available soon as a technical report.

[2] Since the computation is not synchronized on the processors summing the duration is tricky and will be detailed in the technical report.