



International Conference on Computational Science, ICCS 2012

A class of communication-avoiding algorithms for solving general dense linear systems on CPU/GPU parallel machines

Marc Baboulin^a, Simplicé Donfack^a, Jack Dongarra^b, Laura Grigori^a, Adrien Rémy^a, Stanimire Tomov^b^a*Inria and University Paris-Sud, France*^b*University of Tennessee, USA*

Abstract

We study several solvers for the solution of general linear systems where the main objective is to reduce the communication overhead due to pivoting. We first describe two existing algorithms for the LU factorization on hybrid CPU/GPU architectures. The first one is based on partial pivoting and the second uses a random preconditioning of the original matrix to avoid pivoting. Then we introduce a solver where the panel factorization is performed using a communication-avoiding pivoting heuristic while the update of the trailing submatrix is performed by the GPU. We provide performance comparisons and tests on accuracy for these solvers on current hybrid multicore-GPU parallel machines.

Keywords: hybrid multicore/GPU computing, dense linear algebra libraries, linear system solvers, communication-avoiding algorithms, LU factorization

1. Introduction

There have been several main changes in the development of dense linear algebra libraries over the years. These changes have always been triggered by major hardware developments. For example, LINPACK [1] in the 70's targeted the vector machines at the time for which cache reuse was not essential, and as a result LINPACK had relied on just Level 1 BLAS. In the 80's LINPACK had to be rewritten, leading to LAPACK [2], that would rely on Level 3 BLAS for cache based machines. In the 90's it was extended to ScaLAPACK [3] for parallel platforms, relying on the PBLAS [4] message passing. Now, in the 00's, with the explosion in parallelism and heterogeneity as well as the ever increasing data-communication costs, the old libraries had to be redesigned once again. An example is the MAGMA¹ library - a collection of this next generation linear algebra libraries [5, 6, 7].

MAGMA, similarly to LAPACK, is being build as a community effort, incorporating the newest developments in hybrid algorithms and scheduling, and aiming at minimizing synchronizations and communication in these algorithms. The goal of these efforts is to redesign the dense linear algebra algorithms in LAPACK to fully exploit the power of current heterogeneous systems of multi/manycore CPUs and accelerators, and deliver the fastest possible time to an

Email addresses: marc.baboulin@inria.fr (Marc Baboulin), simplice.donfack@lri.fr (Simplicé Donfack), dongarra@eecs.utk.edu (Jack Dongarra), laura.grigori@inria.fr (Laura Grigori), adrien.remy@lri.fr (Adrien Rémy), tomov@eecs.utk.edu (Stanimire Tomov)

¹Matrix Algebra on GPU and Multicore Architectures, <http://icl.cs.utk.edu/magma/>

accurate solution within given energy constraints. Indeed, the algorithms included so far in MAGMA 1.1 manage to overcome bottlenecks associated with just multicore or GPUs, to significantly outperform corresponding packages for any of these components taken separately. MAGMA's one-sided factorizations for example (and linear solvers) on a single Fermi GPU (and a basic CPU host) can outperform state-of-the-art CPU libraries on high-end multi-socket, multicore nodes (e.g., using up to 48 modern cores). In parallel to the success stories in the development of hybrid algorithms, there have been a number of new developments related to minimizing the communication in one-sided factorizations (e.g. [8]). Such improvements have become essential due to the increasing gap between communication and computation costs.

For the linear system solvers on current multicore or GPU architectures, a bottleneck in terms of communication cost and parallelism comes from the pivoting, a technique used to prevent divisions by too-small numbers in the Gaussian Elimination (GE) process. The commonly used method of Gaussian Elimination with partial pivoting (GEPP) is implemented in current linear algebra libraries for solving square linear systems $Ax = b$ resulting in very stable algorithms. These systems are in general solved using the well-known LU factorization that decomposes the input matrix A into the product $L \times U$, where L is a lower triangular matrix and U is an upper triangular matrix. Current libraries like LAPACK implement GE using a block algorithm, which factors the input matrix by iterating over its blocks of columns (panels). At each iteration, the LU factorization of the current panel is computed, and then the trailing submatrix is updated. In the LAPACK implementation, pivoting is used during the factorization of each column of a panel, and leads to swapping rows. Pivoting not only requires communication (or synchronization in a shared memory environment), but it also limits the exploitation of asynchronicity between block operations. This is because the update of the trailing submatrix can be performed only when the panel factorization is completed. We can find in [9] an evaluation of the communication overhead due to partial pivoting using MAGMA on a given CPU/GPU architecture. This cost can represent on some hybrid architectures up to 40% of the global factorization time, depending on the matrix size. The communication cost of GEPP is asymptotically larger than the lower bounds on communication [10]. Other classical pivoting strategies can be used, as rook pivoting or complete pivoting (see [11] for a comprehensive review), but they always require between $O(n^2)$ and $O(n^3)$ comparisons to search for the pivot. In this paper we consider two alternative strategies to these pivoting techniques, that have the property of reducing the communication in the LU factorization while providing a satisfying accuracy.

The first alternative is tournament pivoting, introduced in the last years in the context of CALU, a communication-avoiding LU factorization algorithm [12]. It is shown in [10] that tournament pivoting is as stable as partial pivoting in practice and that CALU minimizes communication. With this strategy, the panel factorization, referred to as TSLU (Tall Skinny LU), can be efficiently parallelized as follows. The panel is partitioned into P_r blocks. From each block, a set of local pivots is selected in parallel using GEPP. A tournament is used on the P_r local sets to select a set of global pivots. These global pivots are moved to the diagonal positions, and then the LU factorization with no pivoting of the entire panel is performed. The tournament is implemented as a reduction operation, with GEPP being the operator used at each step of the reduction.

The second alternative is proposed in [9] where the communication overhead due to pivoting is completely removed by considering a randomization technique referred to as Partial Random Butterfly Transformation (PRBT). This method revisits the approach initially proposed in [13] where the original matrix A is transformed into a matrix that would be sufficiently "random" so that, with a probability close to 1, pivoting is not needed. It consists of a multiplicative preconditioning U^TAV where the matrices U and V are chosen among a particular class of random matrices called *recursive butterfly matrices*. Then Gaussian Elimination with No Pivoting (GENP) is performed on the matrix U^TAV . It has been shown in [9] that a small number of recursions (2 in practice) is sufficient to obtain a satisfying backward error. Moreover, by exploiting the particular structure of the recursive butterflies, this technique of randomization is attractive in terms of Gflop/s ($O(n^2)$ operations) and data storage. Note that since in this approach we know in advance that we are not going to pivot, GENP that follows randomization is implemented as a very efficient fully BLAS 3 algorithm. Note also that when the initial matrix is randomized, we systematically add iterative refinement in the working precision for better stability, as indicated in [11, p. 232].

We show in this paper that the usage of these techniques in the context of hybrid CPU/GPU architectures allows us to take advantage of each computational unit.

The paper is organized as follows. First we describe in Section 2 two existing algorithms for LU factorization on hybrid CPU/GPU machines. The first one is implemented in the parallel library MAGMA and is based on the partial pivoting technique. The second algorithm uses randomization to avoid pivoting and performs GENP on the

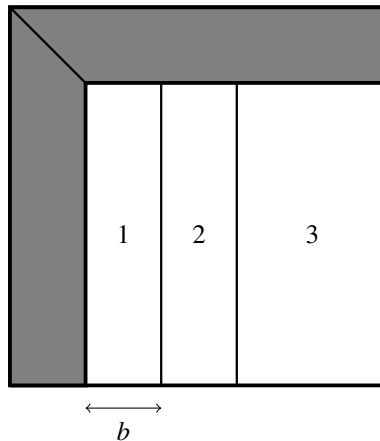


Figure 1: Block splitting in hybrid LU factorization

randomized matrix. Then in Section 3 we introduce tournament pivoting, a strategy based on CALU that we adapt specifically for CPU/GPU architectures. In this new implementation, the panel is factored on the CPU using a modified CALU factorization while the update of the trailing submatrix is performed on the GPU. The resulting solver is called H-CALU solver. Finally we propose in Section 4 performance and accuracy results where we compare H-CALU with LU solvers using partial pivoting (MAGMA) and PRBT. Concluding remarks are given in Section 5.

2. Two existing hybrid algorithms for LU factorization

2.1. LU factorization based on partial pivoting

Let us illustrate how the hybrid multicore + GPU approach can be applied to the LU factorization by describing the algorithm as it is implemented in the MAGMA library. The method is based on splitting the computation as shown in Figure 1 that represents a current matrix factored via a right looking block LU factorization [14, p. 85], where the dark part has been already factored. The initial matrix has been downloaded to the GPU and we describe here a current iteration:

1. The current panel (1) containing b columns is downloaded to the CPU.
2. (1) is factored by the CPU and the result is sent back to the GPU.
3. The GPU updates (2) that corresponds to the next panel (first b columns of the trailing submatrix).
4. The updated panel (2) is sent to the CPU and asynchronously factored on the CPU while the GPU updates the rest of the matrix (3).

The technique consisting of factoring (2) while still updating (3) is often referred to as *look-ahead* [15]. In the current implementation of MAGMA, the panel factorization is performed using GEPP but this algorithm is general enough to be applicable to many forms of LU factorizations, where the distinction can be made based on the form of pivoting that they employ. In Section 3.2 we will use a different pivoting strategy that turns out to be very efficient for factoring the panel due to its particular “tall and skinny” structure. Depending on the problem size n and on the hardware used, MAGMA proposes a default value for the parameter b (width of the panel). Note that the design of the hybrid LU in MAGMA avoids communicating by having only panels transferred between CPU and GPU ($O(n * b)$ data vs $O(n * n * b)$ computation in the updates), enabling also the total overlap of the panel computation by the updates for n large enough.

2.2. PRBT solver

Using the PRBT solver, we solve the general linear system $Ax = b$ by the following steps:

1. Compute the randomized matrix $A_r = U^T AV$, with U and V recursive butterflies.
2. Factorize A_r with GENP.
3. Solve $A_r y = U^T b$.
4. Solution is $x = Vy$.

In this section, we present an implementation of a PRBT solver for hybrid CPU/GPU architectures. The randomization by recursive butterflies uses 2 recursions and the resulting computational cost for randomization is $4n^2$ flops (see [9] for more details on the computational kernels that exploit the particular structure of the recursive butterflies). The PRBT solver for hybrid CPU/GPU architectures performs the following tasks:

- Random generation and packed storage of U and V on the CPU, while sending A to the device memory.
- The packed representation of U and V are sent from the host memory to the device memory.
- Randomization is performed on the GPU, updating A in-place (no additional memory needed) on the device memory.
- The randomized matrix is factorized with GENP on the GPU, the panel factorization being performed on the CPU host.
- We compute $U^T b$ on the GPU, $A_r y = U^T b$ is solved on the GPU, followed by the solution $x = Vy$.
- The solution is sent to the host memory, followed if necessary by iterative refinement on the CPU.

When using the PRBT solver, the bulk of the computation corresponds to the GENP factorization (the randomization represents less than 2% of the global computational time). Then we expect that the Gflop/s performance of the PRBT solver will provide us with an upper bound for other LU solvers on hybrid CPU/GPU architectures.

3. Hybrid communication-avoiding LU factorization

3.1. Communication-avoiding LU

The poor evolution of latency and memory bandwidth that we observed over recent years for parallel architectures is a major bottleneck for algorithms that require communication like GEPP. New approaches have been recently proposed to minimize the amount of communication due to pivoting. Among them are the communication-avoiding algorithms introduced for distributed-memory machines [12] and for multicore architectures [16]. These algorithms are effective in the sense that they reduce significantly the communication while being stable in practice. CALU is an algorithm that partitions the input matrix into block of columns (panels), iteratively factors the panel and updates the trailing submatrix. The factorization of the panel is one of the most important task in the LU factorization since it is part of the critical path in the diagram of tasks and its effective execution influences the performance of the algorithm. In CALU, the panel factorization is performed by the TSLU algorithm which factors a block column of size b with the following steps:

1. Partition the block column in P_r threads where P_r is the number of threads participating in the panel factorization. Let A_i , the block assigned to each thread.
2. Each thread applies GEPP to its block and we have $\Pi_i A_i = L_i U_i$. We denote by $C_i = (\Pi_i A_i)(1 : b, 1 : b)$ the block containing the potential pivots, where Π_i is the permutation matrix resulting from the factorization of A_i .
3. For each step of the reduction tree (see [16] for more details on the most adapted reduction tree):
If j is the thread number exchanging data with thread i and $C'_i = [C_i; C_j]$, then we compute $\Pi_i C'_i = L_i U_i$ and $C_i = (\Pi_i C'_i)(1 : b, 1 : b)$ is the new pivot candidate.
4. Apply the latest permutation computed in 3) to the original panel in order to compute the block column of L and to compute the block row of U .

Once the panel is factored using TSLU then we update the trailing submatrix. Following the approach presented in [17, 18], the CALU algorithm can be represented as a Directed Acyclic Graph (DAG) where nodes are elementary tasks that operate on one or several $b \times b$ blocks and where edges represent the dependencies among them. A dependency occurs when a task must access data that is the output of another task, either to further update or just read data. In Figure 2 we represent an example of LU factorization with CALU as a sequence of DAGs using 2 threads. The panel is partitioned into 3 column blocks. Red tasks represent the factorization of the panel via TSLU, the green tasks represent the update of the trailing submatrix, and the blue tasks are those that are not executed so far.

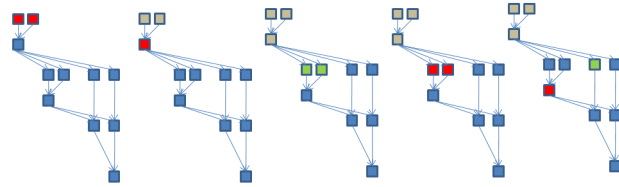


Figure 2: Example of asynchronous LU factorization using multithreaded CALU (2 threads).

3.2. Hybrid version of CALU (H-CALU)

As explained in Section 2.1, the LU algorithm implemented in MAGMA factors each block of columns iteratively. Each step is essentially decomposed into two distinct phases: the factorization of the panel by the CPU followed by the update of the trailing submatrix by the GPU. The algorithm's design minimizes the CPU-GPU communications. In the following we describe an algorithm that further improves the algorithm in MAGMA by minimizing the communication associated with the panel factorizations.

At each step, a panel of size B is factored on the CPU by applying CALU to a rectangular matrix and the update of the trailing submatrix is performed by the GPU. CALU factors the panel by splitting the initial block of columns into smaller blocks containing b columns that are factored iteratively using TSLU. Thus, the factorization of the panel is considered as a variant of the algorithm at the first level where we factor a rectangular matrix using only the CPU. The use of this second level of blocking is important for performance on hybrid CPU/GPU architectures because the CPU and GPU processors have different size of cache. The block size B is chosen in order to optimize the performance of the matrix-matrix product on the GPU and to ensure a good grain for increasing parallelism. Then the block size b is tuned in order to optimize the utilization of the multicore cache. This decomposition of the algorithm into small tasks allows us to operate on block of data that fit into the cache. It results in an asynchronous and dynamic execution of the panel factorization on the CPU, yielding to good performance on multicore machines [16]. This asynchronous execution keeps busy most of the CPU threads. When $b = B$, CALU behaves simply as TSLU. If B is large enough (which will be the case for our hybrid implementation), the panel is factored using CALU rather than TSLU because CALU can be executed asynchronously [16]. Our approach also uses the well known technique referred to as *look-ahead* [15] but adapted here so that the CPU and the GPU can work together while minimizing the number of memory transfers. In this approach, we start factoring the next panel as soon as possible.

Figure 3 shows an example of the factorization of a matrix at the top level. We consider that the matrix is initially stored on the GPU. Red tasks represent the factorization of the panel using multithreaded CALU and the green tasks represent the update of the trailing submatrix in the GPU. At each step of the factorization, the block corresponding to the panel is transferred to the CPU and factored using CALU. Once a panel is factored, it is transferred again to the GPU to update the trailing submatrix. The GPU updates in priority the column block corresponding to the next panel. Note that, similarly to [6], the data transfer between CPU and GPU is overlapped by computation.

4. Numerical experiments

In this section we present performance results for the H-CALU algorithm described in Section 3.2. These numerical experiments were carried out using a hybrid CPU/GPU system where:

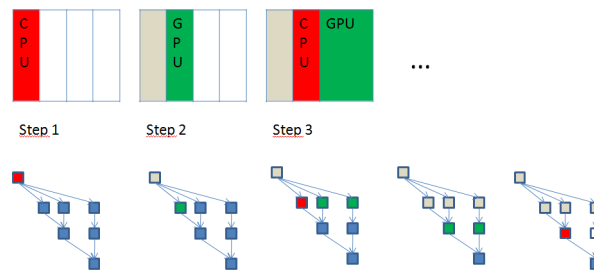


Figure 3: Hybrid CALU factorization (4 panels).

- The GPU device is an NVIDIA Fermi Tesla S2050 with 448 CUDA cores running at 1.15 GHz and 2687 MB memory.
- The multicore host is a 48 cores system (4 sockets \times 12 cores) AMD Opteron 6172 (2.1 GHz).

For experiments involving only the multicore host (panel factorization), comparisons are made against the MKL [19] multithreaded library. For experiments involving both CPU and GPU (factorization of the whole matrix), comparisons are made against version 1.1 of the MAGMA library. All computations are performed on random matrices and in double precision arithmetic.

4.1. Panel factorization

As described in Section 3.2, the panel factorization is performed by the CPU while the update of the trailing submatrix is executed on the GPU. Let us evaluate specifically the performance for the panel factorization phase in an LU factorization. This performance is measured by summing the total number of flops executed in factoring successively each panel throughout the factorization and dividing it by the time spent during these steps. This performance (expressed in Gflop/s) is plotted in Figure 4 for the factorization of four square matrices, each associated with a given panel size (parameter B defined in Section 3.2, corresponding to the number of columns for the panel). For factoring the panel, we consider different number of threads (one CPU core being used for each thread) varying from 1 to 26. Note that using more than 26 threads does not provide us with better performance, due to the too-large amount of communication involved in the panel factorization. The panel size B considered in Figure 4 for each matrix size corresponds to a value empirically tuned in order to provide the best global factorization time for each matrix when using a hybrid implementation.

In these experiments, we compare the performance of the panel factorization for the following routines:

- CALU factorization routine modified for the H-CALU solver and linked with the sequential version of MKL for the required BLAS and LAPACK routines.
- MKL implementation of the LAPACK routine `dgetrf`, used in the MAGMA implementation of LU for factoring the panel.
- A recursive routine for GEPP `rgetf2` (linked with MKL multithreaded BLAS) described in [20] and known to give good performance on “tall and skinny” matrices.
- PRBT where the panel is factored using the GENP routine `dgetrf_nopiv` (no pivoting).

The routines compared in this section have been selected on the fact that they can be used as kernels for our hybrid CPU/GPU implementation. If we use only multicore machines without GPU, then other solvers can be considered (see e.g. recursive tile version in [21]). The performance of PRBT, based on a GENP routine can be considered here as a “peak” performance for the panel factorization. In this respect, we observe that, in percentage of this peak performance and depending on the matrix size n , CALU achieves between 36 % ($n = 5120$) and 48 % ($n = 21504$), `dgetrf` achieves between 35 % ($n = 5120$) and 23 % ($n = 21504$), and `rgetf2` achieves between 31 % ($n = 5120$) and

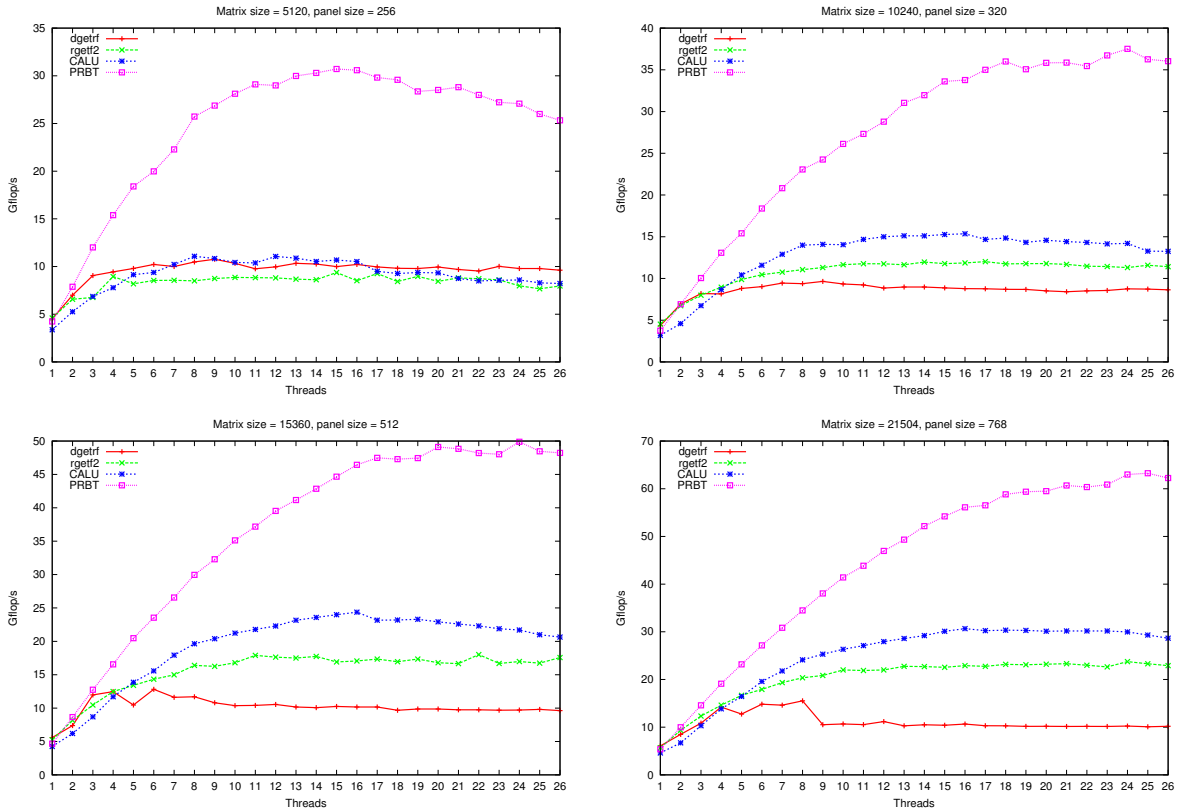


Figure 4: Comparison of CPU multi-threaded panel factorizations.

38 % ($n = 21504$). We also observe that CALU is even faster for larger ratios rows/columns. Moreover, CALU and PRBT have better scalability properties. This can be explain by the fact that CALU minimizes communication thanks to its pivoting strategy and PRBT does not pivot at all. The plateau observed for each curve after a certain number of threads corresponds to cases where the volume of communication becomes too large and cannot be overlapped by computation. For $n = 5120$, CALU, dgetrf and rgetf2 give similar performance. However, when the matrix size increases and then the panel becomes more “tall and skinny”, CALU outperforms the two other solvers and achieves a reasonable fraction of the PRBT rate. This good behavior of CALU for factoring the panel was already mentioned in [16]. In particular this better scalability of CALU enables us to use more CPU threads in factoring the panel and then to improve the overall performance of a hybrid LU solver.

4.2. Performance of hybrid LU implementations

In this section we study the performance of LU factorization routines that utilize resources from multicore (16 threads) and one GPU. We compare in Figure 5a the following routines, applied to square matrices of various sizes:

- The MAGMA routine magma_dgetrf, where the panel is factored using the MKL routine dgetrf.
- H-CALU, where the panel is factored using the CALU routine mentioned in Section4.1.
- The PRBT solver (randomization + GENP).

As expected, PRBT outperforms the routines because it does not pivot and the randomization time is negligible. We can also observe that in the range 1024-5120, H-CALU gives similar performance as MAGMA but it is slightly faster for matrix sizes larger than 5120. This trend can be explained by the fact that, for matrix sizes smaller than 5120, the panels are not enough “tall and skinny” to take advantage of the CALU algorithm. We notice that the difference

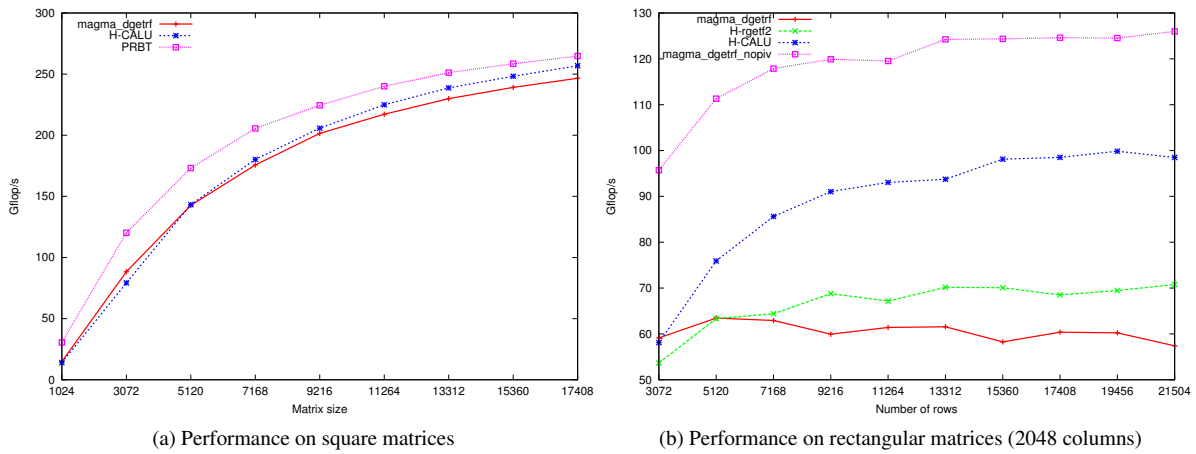


Figure 5: Performance on AMD + Tesla 2050, 16 threads

of performance observed for the panel in Section 4.1 has a moderate impact on the whole factorization since the update phase performed on the GPU represents the bulk of the computation. In particular we do not mention in these results the experiments performed with a MAGMA routine modified so that the panel is factored by the routine `rgetf2` mentioned in Section 4.1 because we obtained performance results similar to that of `magma_dgetrf`. Note that asymptotically, the performance of the three routines should be close because communication becomes negligible compared to the $O(n^3)$ computations for large dimensions.

In Figure 5b we compare the performance of hybrid LU factorization routines for rectangular matrices of size $m \times n$ with $m > n$, using 16 threads. Such an LU factorization exists when $A(1 : k; 1 : k)$ is nonsingular for $k = 1 : n$ (see [22, p. 102]). In our experiments $n = 2048$ and m varies from 3072 to 21504. Comparisons are made against MAGMA routines `magma_dgetrf` and `magma_dgetrf_nopiv` (instead of PRBT since the latter has no implementation for rectangular matrices). We also compare with the MAGMA routine `magma_dgetrf` modified by factoring the panel using the recursive GEPP kernel (this routine is named H-`rgetf2` in our graph). On this type of matrices, H-CALU outperforms `magma_dgetrf` and H-`rgetf2`. Indeed, for rectangular matrices, the proportion of computation performed during the panel factorization is bigger. Hybrid factorization on rectangular matrices could be for instance useful in a future hybrid factorization with multiple GPUs where the (rectangular) panel could be factored using CPU and a GPU.

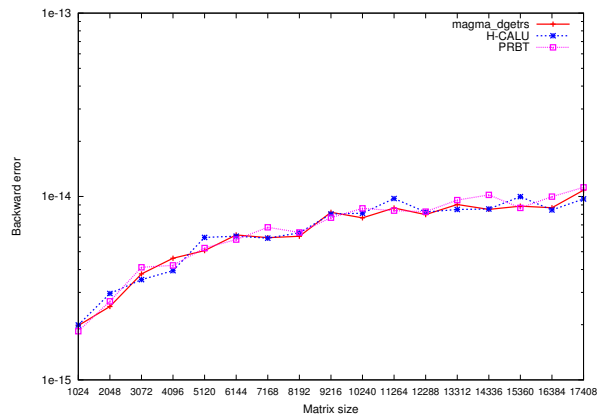


Figure 6: Comparison of componentwise backward error

4.3. Tests on accuracy

First we study the backward error obtained for the linear system solution computed with the solvers LU MAGMA (based on the factorization routine `magma_dgetrf`), H-CALU and PRBT, using the random matrices considered in the previous experiments. The quantity plotted in Figure 6 corresponds to the componentwise backward error given in [2, p. 78] and expressed by

$$\omega = \max_i \frac{|Ax - b|_i}{(|A| \cdot |x| + |b|)_i},$$

where x is the computed solution. The number of threads used in factoring the panel in H-CALU is 16 (we mention this because, as explained in [10], this might affect the accuracy). We observe that the backward errors are very similar for the three hybrid solvers.

Now we present additional tests using LAPACK test cases given in [23]. Table 1 describes the 11 matrices used in our experiments (size 512, all in double precision). In this table, ε denotes the machine precision and κ is the infinity-norm condition number of the matrix. Table 2 shows the componentwise backward error obtained for the three hybrid solvers LU MAGMA, H-CALU and PRBT. Iterative refinement (in the working precision) is added if necessary using the LAPACK routine `dgerfs` on the multicore host. The iterative refinement is based on the stopping criterion given [24] ($\omega > (n + 1)\varepsilon$), with a maximum of 5 iterations. Matrices 5 to 7 are singular and have at least one row and column zero. These are used in [23] to test the error return codes. Nevertheless, for these matrices the random transformation of A allows the LU factorization to continue while the pivoting is not capable of removing the zero pivots. Matrix 9 is scaled to near underflow and the three solvers give comparable results but less accurate than with the other matrices. For the ill-conditioned matrix 9, the backward is slightly less accurate than the other solvers. For the other matrices, the three solvers give the same accuracy. Tests on accuracy for specific matrix collections can be found in [9] and [10] respectively for PRBT and CALU.

Table 1: Test matrices

1	Diagonal	7	Last $n/2$ columns zero
2	Upper triangular	8	Random, $\kappa = \sqrt{0.1/\varepsilon}$
3	Lower triangular	9	Random, $\kappa = 0.1/\varepsilon$
4	Random, $\kappa = 2$	10	Scaled near underflow
5	First column zero	11	Scaled near overflow
6	Last column zero		

Table 2: Componentwise Backward Error

Matrix Type	MAGMA LU (<code>magma_dgetrf</code>)	H-CALU	PRBT
1	0.0	0.0	2.10145e-16
2	1.31539e-16	1.31539e-16	2.18841e-16
3	184697e-16	184697e-16	2.06543e-16
4	2.1647e-16	2.75832e-16	1.92510e-16
5	-	-	2.66472e-16
6	-	-	2.14281e-16
7	-	-	1.97144e-16
8	2.10408e-16	3.76095e-16	1.55625e-16
9	2.70036e-16	6.36540e-16	1.08967e-13
10	7.59577e-14	7.40225e-14	7.54745e-14
11	2.27295e-16	2.11000e-16	2.42990e-16

5. Conclusion

In this paper we presented different LU factorization routines using a multicore machine accelerated with one GPU. The difference between these approaches comes from the pivoting strategy chosen for factoring the panel. We proposed a new hybrid communication-avoiding solver H-CALU where the panel is factored on the CPU while the update is performed by the GPU. In our experiments, this solver turns out to be faster than the classical GEPP implementation in MAGMA for square matrices larger than 5120 and in all cases for rectangular matrices. H-CALU is scalable, allowing us to use larger panels and thus to limit the amount of transfer between the CPU and the GPU memory. We point out that further optimizations are possible with e.g. additional tuning and scheduling but our experiments give a general trend for the performance of algorithms as dictated by the amount of communication that they perform. However, the solver based on randomization always outperforms other solvers since we do not pivot at all and the computational cost for randomizing is small.

References

- [1] J. J. Dongarra, C. B. Moler, J. R. Bunch, G. W. Stewart, LINPACK Users' Guide, SIAM, 1979.
- [2] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, LAPACK Users' Guide, SIAM, 1999, third edition.
- [3] L. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, R. Whaley, ScaLAPACK Users' Guide, SIAM, 1997.
- [4] J. Choi, J. Dongarra, L. Ostrouchov, A. Petitet, D. Walker, R. Whaley, A proposal for a set of parallel basic linear algebra subprograms, Tech. rep., LAPACK Working Note 100 (1995).
- [5] R. Nath, S. Tomov, J. Dongarra, An improved MAGMA GEMM for Fermi GPUs, *International Journal of High Performance Computing Applications* 24 (4) (2010) 511–515.
- [6] S. Tomov, J. Dongarra, M. Baboulin, Towards dense linear algebra for hybrid GPU accelerated manycore systems, *Parallel Computing* 36 (5&6) (2010) 232–240.
- [7] S. Tomov, R. Nath, J. Dongarra, Accelerating the reduction to upper Hessenberg, tridiagonal, and bidiagonal forms through hybrid GPU-based computing, *Parallel Computing* 36 (12) (2010) 645–654.
- [8] M. Anderson, G. Ballard, J. Demmel, K. Keutzer, Communication-Avoiding QR decomposition for GPUs LAPACK Working Note 240, proceedings of IPDPS'11.
- [9] M. Baboulin, J. Dongarra, J. Herrmann, S. Tomov, Accelerating linear system solutions using randomization techniques LAPACK Working Note 246.
- [10] L. Grigori, J. Demmel, H. Xiang, CALU: a communication optimal LU factorization algorithm, *SIAM J. Matrix Anal. and Appl.* 32 (2011) 1317–1350.
- [11] N. J. Higham, Accuracy and Stability of Numerical Algorithms, SIAM, 2002, second edition.
- [12] L. Grigori, J. Demmel, H. Xiang, Communication avoiding Gaussian elimination, in: Proceedings of the 2008 ACM/IEEE conference on Supercomputing, IEEE Press, 2008, p. 29.
- [13] D. S. Parker, Random butterfly transformations with applications in computational linear algebra, Technical Report CSD-950023, Computer Science Department, UCLA (1995).
- [14] J. Dongarra, I. Duff, D. Sorensen, H. van der Vorst, Numerical Linear Algebra for High-Performance Computers, SIAM, 1998.
- [15] J. Kurzak, J. Dongarra, Implementing linear algebra routines on multi-core processors with pipelining and a look ahead LAPACK Working Note 178.
- [16] S. Donfack, L. Grigori, A. K. Gupta, Adapting communication-avoiding LU and QR factorizations to multicore architectures, in: Parallel & Distributed Processing (IPDPS), 2010 IEEE International Symposium on, IEEE, 2010, pp. 1–10.
- [17] A. Buttari, J. Dongarra, J. Kurzak, J. Langou, P. Luszczek, S. Tomov, The impact of multicore on math software In Proceedings of PARA 2006, Workshop on state-of-the-art in scientific computing.
- [18] A. Buttari, J. Langou, J. Kurzak, J. Dongarra, Parallel tiled QR factorization for multicore architectures, *Concurr. Comput. : Pract. Exper.* 20 (2007) 1573–1590.
- [19] Intel, Math Kernel Library (MKL), <http://www.intel.com/software/products/mkl/>.
- [20] F. G. Gustavson, Recursion leads to automatic variable blocking for dense linear-algebra algorithms, *IBM Journal of Research and Development* 41 (6) (1997) 737–755.
- [21] J. Dongarra, M. Faverge, H. Ltaief, P. Luszczek, Achieving numerical accuracy and high performance using recursive tile LU factorization- LAPACK Working Note 259.
- [22] G. H. Golub, C. F. van Loan, Matrix Computations, The Johns Hopkins University Press, 1996, third edition.
- [23] S. Blackford, J. Dongarra, Installation Guide for LAPACK LAPACK Working Note 41, revised version 3.0.
- [24] R. D. Skeel, Iterative refinement implies numerical stability for Gaussian elimination, *Math. Comput.* 35 (1980) 817–832.