

# Parallel Tiled QR Factorization for Multicore Architectures

Alfredo Buttari<sup>1</sup>, Julien Langou<sup>2</sup>, Jakub Kurzak<sup>1</sup>, and Jack Dongarra<sup>1,3,4</sup>

<sup>1</sup>*Department of Electrical Engineering and Computer Science, University Tennessee, Knoxville, Tennessee*

<sup>2</sup>*Department of Mathematical Sciences, University of Colorado at Denver and Health Sciences Center, Colorado*

<sup>3</sup>*Oak Ridge National Laboratory, Oak Ridge, Tennessee*

<sup>4</sup>*University of Manchester, Manchester UK*

January 6, 2008

## Abstract

As multicore systems continue to gain ground in the High Performance Computing world, linear algebra algorithms have to be reformulated or new algorithms have to be developed in order to take advantage of the architectural features on these new processors. Fine grain parallelism becomes a major requirement and introduces the necessity of loose synchronization in the parallel execution of an operation. This paper presents an algorithm for the QR factorization where the operations can be represented as a sequence of small tasks that operate on square blocks of data (referred to as “tiles”). These tasks can be dynamically scheduled for execution based on the dependencies among them and on the availability of computational resources. This may result in an out of order execution of the tasks which will completely hide the presence of intrinsically sequential tasks in the factorization. Performance comparisons are presented with the LAPACK algorithm for QR factorization where parallelism can only be exploited at the level of the BLAS operations and with vendor implementations.

## 1 Introduction

In the last twenty years, microprocessor manufacturers have been driven towards higher performance rates only by the exploitation of higher degrees of *Instruction Level Parallelism* (ILP). Based on this approach, several generations of processors have been built where clock frequencies were higher and higher and pipelines were deeper and deeper. As a result, applications could

benefit from these innovations and achieve higher performance simply by relying on compilers that could efficiently exploit ILP. Due to a number of physical limitations (mostly power consumption and heat dissipation) this approach cannot be pushed any further. For this reason, chip designers have moved their focus from ILP to *Thread Level Parallelism* (TLP) where higher performance can be achieved by replicating execution units (or *cores*) on the die while keeping the clock rates in a range where power consumption and heat dissipation do not represent a problem. It is easy to imagine that multicore technologies will have a deep impact on the High Performance Computing (HPC) world where high processor counts are involved and, thus, limiting power consumption and heat dissipation is a major requirement. The Top500 [1] list released in June 2007 shows that the number of systems based on dual-core Intel Woodcrest processors grew in six months (i.e. from the previous list) from 31 to 205 and that 90 more systems are based on dual-core AMD Opteron processors.

Even though many attempts have been made in the past to develop parallelizing compilers, they proved themselves efficient only on a restricted class of problems. As a result, at this stage of the multicore era, programmers cannot rely on compilers to take advantage of the multiple CPUs available on a processor. All the applications that were not explicitly coded to be run on parallel architectures must be rewritten with parallelism in mind. Also, those applications that could exploit parallelism may need considerable rework in order to take advantage of the fine-grain parallelism features provided by multicores.

The current set of multicore chips from Intel and AMD are for the most part multiple processors glued together on the same chip. There are many scalability issues to this approach and it is unlikely that this type of architecture will scale up beyond 8 or 16 cores. Even though it is not yet clear how chip designers are going to address these issues, it is possible to identify some properties that algorithms must have in order to match high degrees of TLP:

**fine granularity:** cores are (and probably will be) associated with relatively small local memories (either caches or explicitly managed memories like in the case of the STI Cell [22] architecture or the Intel Polaris[4] prototype). This requires splitting an operation into tasks that operate on small portions of data in order to reduce bus traffic and improve data locality.

**asynchronicity:** as the degree of TLP grows and granularity of the operations becomes smaller, the presence of synchronization points in a parallel execution seriously affects the efficiency of an algorithm.

The LAPACK [6] and ScaLAPACK [10] software libraries represent a *de facto* standard for high performance dense Linear Algebra computations and have been developed, respectively, for shared-memory and distributed-memory architectures. In both cases exploitation of parallelism comes from the availability of parallel BLAS. In the LAPACK case, a number of BLAS libraries can be used to take advantage of multiple processing units on shared memory systems; for example, the freely distributed ATLAS [27] and GotoBLAS [16] or other vendor BLAS like Intel MKL [2] or AMD ACML [3] are popular choices.

These parallel BLAS libraries use common techniques for shared memory parallelization like pThreads [21] or OpenMP [12]. This is represented in Figure 1 (left).

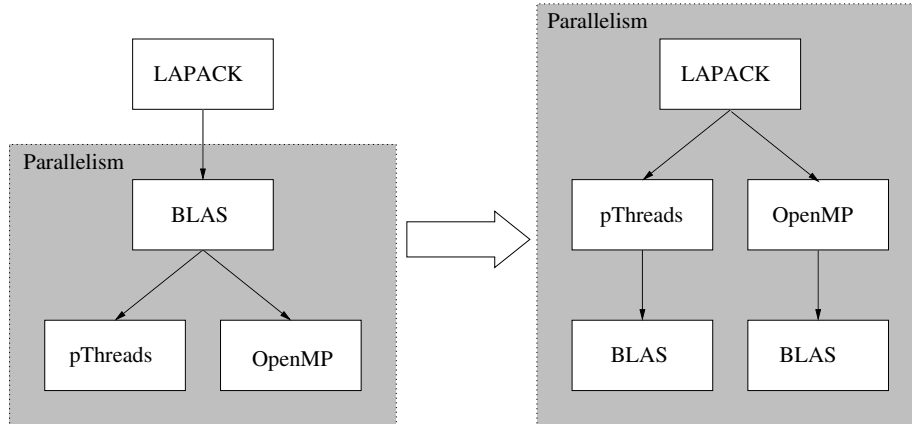


Figure 1: Transition from sequential algorithms that rely on parallel BLAS to parallel algorithms.

In the ScaLAPACK case, parallelism is exploited by PBLAS [11] which is a parallel BLAS implementation that uses the Message Passing Interface [14] (MPI) for communications on a distributed memory system. Substantially, both LAPACK and ScaLAPACK implement sequential algorithms that rely on parallel building blocks (i.e., the BLAS operations). As multicore systems require finer granularity and higher asynchronicity, considerable advantages may be obtained by reformulating old algorithms or developing new algorithms in a way that their implementation can be easily mapped on these new architectures. This transition is shown in Figure 1. An approach along these lines has already been proposed in [8, 9, 20] where operations in the standard LAPACK algorithms for some common factorizations were broken into sequences of smaller tasks in order to achieve finer granularity and higher flexibility in the scheduling of tasks to cores. The importance of fine granularity algorithms is also shown by the authors of this paper in earlier works [19]. The usage of recursive factorization techniques [13] is also relevant to the topics discussed in the rest of the paper.

The rest of this document shows how this can be achieved for the QR factorization. Section 2 describes the algorithm for block QR factorization used in the LAPACK library; Section 3 describes the tiled QR factorization that provides both fine granularity and high level of asynchronicity; performance results for this algorithm are shown in Section 4. A comment on the usage of recursive techniques is given in Section ???. Finally future working directions are illustrated in Section 6.

## 2 Block QR Factorization

### 2.1 Description of the block QR Factorization

The QR factorization is a transformation that factorizes an  $m \times n$  matrix  $A$  into its factors  $Q$  and  $R$  where  $Q$  is a unitary matrix of size  $n \times n$  and  $R$  is a triangular matrix of size  $m \times m$ . This factorization is operated by applying  $\min(m, n)$  Householder reflections to the matrix  $A$ . Since Householder reflections are orthogonal transformations, this factorization is stable as opposed to the LU one; however, stability comes at the price of a higher flop count: QR requires  $2n^2(m - n/3)$  as opposed to the  $n^2(m - n/3)$  needed for LU. A detailed discussion of the QR factorization can be found in [15, 25, 26]. LAPACK uses a particular version of this algorithm which achieves higher performance on architectures with memory hierarchies thanks to blocking. This algorithm is based on accumulating a number of Householder transformations in what is called a *panel factorization* which are, then, applied all at once by means of high performance Level 3 BLAS operations. The technique used to accumulate Householder transformation was introduced in [24] and is known as the compact WY technique.

The LAPACK subroutine that performs the QR factorization is called DGEQRF and is explained below. Consider a matrix  $A$  of size  $m \times n$  that can be represented as

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

where  $A_{11}$  is of size  $b \times b$ ,  $A_{12}$  of size  $b \times (n - b)$ ,  $A_{21}$  of size  $(m - b) \times b$  and  $A_{22}$  of size  $(m - b) \times (n - b)$ .

The LAPACK algorithm for QR factorization can be described as a sequence of steps where, at each step, the transformation in Equation (1) is performed.

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \implies \begin{pmatrix} V_{11} \\ V_{21} \end{pmatrix}, \begin{pmatrix} R_{11} & R_{12} \\ 0 & \tilde{A}_{22} \end{pmatrix} \quad (1)$$

The transformation in Equation (1) is obtained in two steps:

1. **Panel Factorization.** At this step a QR transformation of the panel ( $A_{*1}$ ) is performed as in Equation (2).

$$\begin{pmatrix} A_{11} \\ A_{21} \end{pmatrix} \implies \begin{pmatrix} V_{11} \\ V_{21} \end{pmatrix}, (T_{11}), (R_{11}) \quad (2)$$

This operation produces  $b$  Householder reflectors ( $V_{*,1}$ ) and an upper triangular matrix  $R_{11}$  of size  $b \times b$ , which is a portion of the final  $R$  factor, by means of the DGEQR2 LAPACK subroutine; also, at this step, a triangular matrix  $T_{11}$  of size  $b \times b$  by means of the DLARFT LAPACK subroutine<sup>1</sup>. Please note that  $V_{11}$  is a unit lower triangular matrix of size  $b \times b$ . The arrays  $V_{*,1}$  and  $R_{11}$  do not need extra space to be stored since they overwrite  $A_{*1}$ . A temporary workspace is needed to store  $T_{11}$ .

<sup>1</sup>for the meaning of the matrix  $T_{11}$  please refer to [24]

2. **Trailing submatrix update.** At this step, the transformation that was computed in the panel factorization is applied to the rest of the matrix, also called *trailing submatrix* as shown in Equation (3).

$$\begin{pmatrix} R_{12} \\ \tilde{A}_{22} \end{pmatrix} = \begin{pmatrix} I - \begin{pmatrix} V_{11} \\ V_{21} \end{pmatrix} \cdot (T_{11}) \cdot \begin{pmatrix} V_{11}^T & V_{21}^T \end{pmatrix} & \\ & \end{pmatrix} \begin{pmatrix} A_{12} \\ A_{22} \end{pmatrix} \quad (3)$$

This operation, performed by means of the DLARFB LAPACK subroutine, produces a portion  $R_{12}$  of the final  $R$  factor of size  $b \times (n - b)$  and the matrix  $\tilde{A}_{22}$ .

The QR factorization is continued by applying the transformation (1) to the submatrix  $\tilde{A}_{22}$  and, then, iteratively, until the end of the matrix  $A$  is reached. The value of  $b \ll m, n$  (the so called *block size*) is set by default to 32 in LAPACK-3.1.1 but different values may be more appropriate, and provide higher performance, depending on the architecture characteristics.

## 2.2 Scalability limits of the LAPACK implementation

The LAPACK algorithm for QR factorization can use any flavor of parallel BLAS to exploit parallelism on a multicore, shared-memory architecture. This approach, however, has a number of limitations due to the nature of the transformation in Equation (2), i.e., the panel factorization. Both the DGEQR2 and the DLARFT are rich in Level 2 BLAS operations that cannot be efficiently parallelized on currently available shared memory machines. To understand this, it is important to note that Level 2 BLAS operations can be, generally speaking, defined as all those operations where  $O(n^2)$  floating-point operations are performed on  $O(n^2)$  floating-point data; thus, the speed of Level 2 BLAS computations is limited by the speed at which the memory bus can feed the cores. On current multicores architectures, there is a vast disproportion between the bus bandwidth and the speed of the cores. For example the Intel Clovertown processor is equipped with four cores each capable of a double precision peak performance of 10.64 Gflop/s (that is to say a peak of 42.56 Gflop/s for four cores) while the bus bandwidth peak is 10.64 GB/s which provides 1.33 GWords/s (a word being a 64 bit double precision number). As a result, since one core is largely enough to saturate the bus, using two or more cores does not provide any significant benefit. The LAPACK algorithm for QR factorization is, thus, characterized by the presence of a sequential operation (i.e., the panel factorization) which represents a small fraction of the total number of flops performed ( $\mathcal{O}(n^2)$  flops for a total of  $\mathcal{O}(n^3)$  flops) but limits the scalability of the block QR factorization on a multicore system when parallelism is only exploited at the level of the BLAS routines. This approach will be referred to as the *fork-join* approach since the execution flow of the QR factorization would show a sequence of sequential operations (i.e. the panel factorizations) interleaved to parallel ones (i.e., the trailing submatrix updates).

Table 1 shows the scalability limits of the panel factorization and how this affects the scalability of the whole QR factorization on an 8-socket dual-core AMD Opteron system with MKL-9.1 and GotoBLAS-1.15 parallel BLAS libraries.

# cores	AMD ACML-4.0.0		GotoBLAS-1.15	
	DGEQR2	DGEQRF	DGEQR2	DGEQRF
	Gflop/s	Gflop/s	Gflop/s	Gflop/s
1	0.1218	2.9	0.4549	3.31
2	0.1577	5.4	0.4558	5.51
4	0.2083	9.0	0.4557	9.69
8	0.5055	12.8	0.4549	10.58
16	0.4496	8.7	0.4558	13.01

Table 1: Scalability of the fork-join parallelization on a 8-socket dual Opteron system (sixteen cores total).

In [8, 9, 20], a solution to this scalability problem is presented. The approach described in [8, 20] consists of breaking the trailing submatrix update into smaller tasks that operate on a block-column (i.e., a set of  $b$  contiguous columns where  $b$  is the block size). The algorithm can then be represented as a Directed Acyclic Graph (DAG) where nodes represent tasks, either panel factorization or update of a block-column, and edges represent dependencies among them. The execution of the algorithm is performed by asynchronously scheduling the tasks in a way that dependencies are not violated. This asynchronous scheduling results in an out-of-order execution where slow, sequential tasks are hidden behind parallel ones. This approach can be described as a dynamic lookahead technique. Even if this approach provides significant speedup, as shown in [20], it is exposed to scalability problems. In fact, due to the relatively high granularity of the tasks, the scheduling of tasks may have a limited flexibility and the parallel execution of the algorithm may be affected by an unbalanced load. These problems become a major limitation when the degree of parallelism grows (see [20] for more details).

The following sections describe the application of this idea of dynamic scheduling and out of order execution to an algorithm for QR factorization where finer granularity of the operations and higher flexibility for the scheduling can be achieved. Performance results in Section 4 will show how a good scalability and significant speedup can be achieved using the proposed algorithm in combination with graph-driven, dynamic scheduling of the tasks.

### 3 Tiled QR Factorization

The idea of dynamic scheduling and out of order execution can be applied to a class of algorithms for common Linear Algebra operations. Previous work in this direction shows how the Symmetric Rank-K update (SYRK), Cholesky factorization, block LU factorization and block QR factorization [8, 9, 20] can be parallelized with these techniques. Fine granularity algorithms for SYRK and the Cholesky factorization can be easily derived from those used in the BLAS and LAPACK libraries (respectively) by “tiling” the elementary operations that

they are made of (see [9] for details). In the case of the block LAPACK algorithms for LU and QR factorizations, however, the panel reduction involves tall and narrow portions of the matrix (i.e. a block-column) and cannot be reformulated as a sequence of tile-operations; this represents the major limitation of the approach presented in [8, 20] (see also the previous section for more details). In order to overcome these limitations, a major algorithmic change is necessary.

The algorithmic change proposed is actually well-known and takes its roots in updating factorizations [15, 25]. Using updating techniques to tile the algorithms have first<sup>2</sup> been proposed by Yip [28] for LU to improve the efficiency of out-of-core solvers, and were recently reintroduced in [17, 23] for LU and QR, once more in the out-of-core context. A similar idea has also been proposed in [7] for Hessenberg reduction in the parallel distributed context.

The originality of this paper is to study this techniques in the multicore context, where they can be used to formulate fine granularity algorithms and achieve high flexibility for the dynamic scheduling of tasks.

### 3.1 A Fine-Grain Algorithm for QR Factorization

The tiled QR factorization will be constructed based on the following four elementary operations:

**DGEQT2.** This subroutine was developed to perform the unblocked factorization of a diagonal tile  $A_{kk}$  of size  $b \times b$ . This operation produces an upper triangular matrix  $R_{kk}$ , a unit lower triangular matrix  $V_{kk}$  that contains  $b$  Householder reflectors and an upper triangular matrix  $T_{kk}$  as defined by the WY technique for accumulating the transformations. Note that both  $R_{kk}$  and  $V_{kk}$  can be written on the memory area that was used for  $A_{kk}$  and, thus, no extra storage is needed for them. A temporary work space is needed to store  $T_{kk}$ .

Thus,  $\text{DGEQT2}(A_{kk}, T_{kk})$  performs

$$A_{kk} \leftarrow V_{kk}, R_{kk} \quad T_{kk} \leftarrow T_{kk}$$

**DLARFB.** This LAPACK subroutine will be used to apply the transformation  $(V_{kk}, T_{kk})$  computed by subroutine **DGEQT2** to a tile  $A_{kj}$ .

Thus,  $\text{DLARFB}(A_{kj}, V_{kk}, T_{kk})$  performs

$$A_{kj} \leftarrow (I - V_{kk}T_{kk}V_{kk}^T)A_{kj}$$

**DTSQT2.** This subroutine was developed to perform the unblocked QR factorization of a matrix that is formed by coupling an upper triangular tile  $R_{kk}$  with a square tile  $A_{ik}$ . This subroutine will return an upper triangular matrix  $\tilde{R}_{kk}$  which will overwrite  $R_{kk}$  and  $b$  Householder reflectors where  $b$  is the tile size. Note that, since  $R_{kk}$  is upper triangular, the resulting Householder reflectors can be represented as an identity tile  $I$  on top

---

<sup>2</sup>to our knowledge

of a square tile  $V_{ik}$ . For this reason no extra storage is needed for the Householder vectors since the identity tile need not be stored and  $V_{ik}$  can overwrite  $A_{ik}$ . Also a matrix  $T_{ik}$  is produced for which storage space has to be allocated.

Then,  $\text{DTSQT2}(R_{kk}, A_{ik}, T_{ik})$  performs

$$\begin{pmatrix} R_{kk} \\ A_{ik} \end{pmatrix} \leftarrow \begin{pmatrix} I \\ V_{ik} \end{pmatrix}, \tilde{R}_{kk} \quad T_{ik} \leftarrow T_{ik}$$

**DSSRFB.** This subroutine was developed to apply the transformation computed by  $\text{DTSQT2}$  to a matrix formed coupling two square tiles  $A_{kj}$  and  $A_{ij}$ .

Thus,  $\text{DSSRFB}(A_{kj}, A_{ij}, V_{ik}, T_{ik})$  performs

$$\begin{pmatrix} A_{kj} \\ A_{ij} \end{pmatrix} \leftarrow \begin{pmatrix} I - \begin{pmatrix} I \\ V_{ik} \end{pmatrix} \cdot (T_{ik}) \cdot \begin{pmatrix} I & V_{ik}^T \end{pmatrix} \end{pmatrix} \begin{pmatrix} A_{kj} \\ A_{ij} \end{pmatrix}$$

All of this elementary operations rely on BLAS subroutines to perform internal computations.

Assuming a matrix  $A$  of size  $pb \times qb$

$$\begin{pmatrix} A_{11} & A_{12} & \dots & A_{1q} \\ A_{21} & A_{22} & \dots & A_{2q} \\ \vdots & & \ddots & \vdots \\ A_{p1} & A_{p2} & \dots & A_{pq} \end{pmatrix}$$

where  $b$  is the tile size and each  $A_{ij}$  is of size  $b \times b$ , the QR factorization can be performed as in Algorithm 1.

---

**Algorithm 1** The tiled algorithm for QR factorization.

---

```

1: for  $k = 1, 2, \dots, \min(p, q)$  do
2:    $\text{DGEQT2}(A_{kk}, T_{kk});$ 
3:   for  $j = k + 1, k + 2, \dots, q$  do
4:      $\text{DLARFB}(A_{kj}, V_{kk}, T_{kk});$ 
5:   end for
6:   for  $i = k + 1, k + 1, \dots, p$  do
7:      $\text{DTSQT2}(R_{kk}, A_{ik}, T_{ik});$ 
8:     for  $j = k + 1, k + 2, \dots, q$  do
9:        $\text{DSSRFB}(A_{kj}, A_{ij}, V_{ik}, T_{ik});$ 
10:    end for
11:  end for
12: end for

```

---

Figure 2 gives a graphical representation of one repetition (with  $k = 1$ ) of the outer loop in Algorithm 1 with  $p = q = 3$ . The red, thick borders show what tiles in the matrix are being read and the light blue fill shows what tiles are being written in a step. The  $T_{kk}$  tiles are not shown in this figure for clarity purposes.



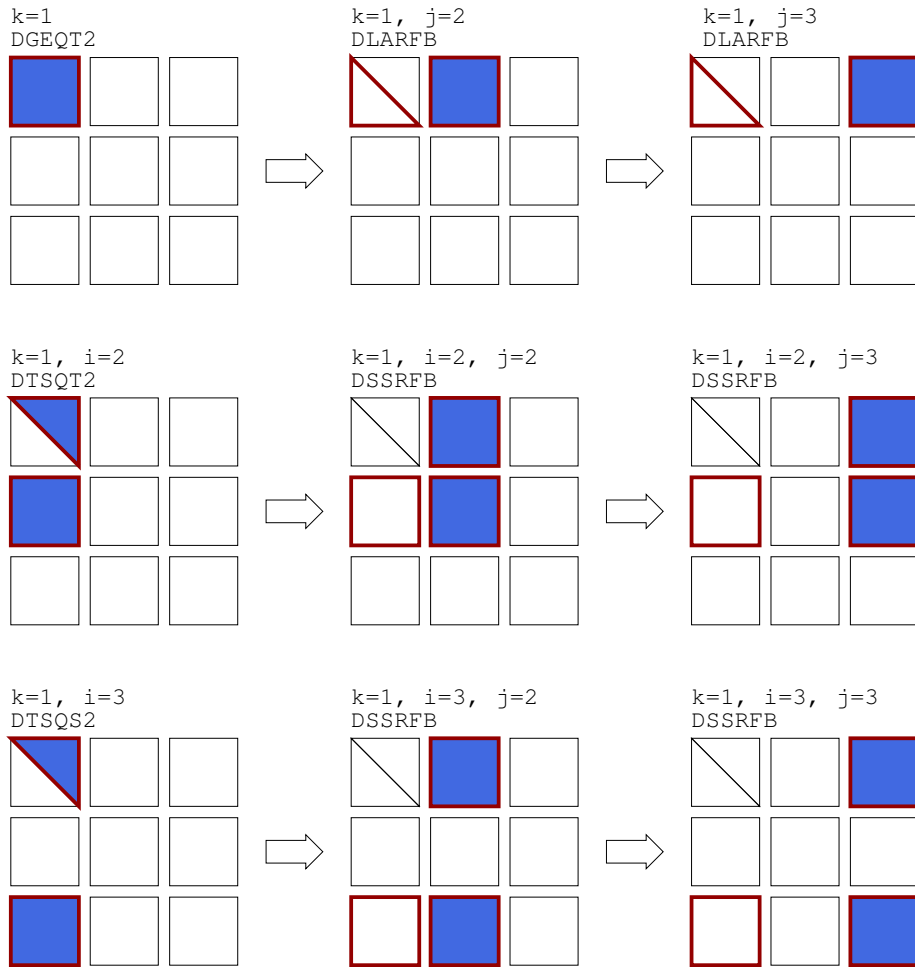


Figure 2: Graphical representation of one repetition of the outer loop in Algorithm 1 on a matrix with  $p = q = 3$ . As expected the picture is very similar to the out-of-core algorithm presented in [17].

### 3.2 Operation count

This section shows that Algorithm 1 has a higher operation count than the LAPACK algorithm discussed in Section 2. Performance results in Section 4 will demonstrate that it is worth paying this cost for the sake of scaling. The operation count of the tiled algorithm for QR factorization can be derived starting from the operation count of each elementary operation; assuming that  $b$  is the tile size:

**DGEQT2:** this operation is a standard non blocked QR factorization of a  $b \times b$  tile where, in addition, the  $T_{kk}$  triangular tile is computed. Thus,  $4/3b^3$  floating point operations are performed for the factorization plus  $2/3b^3$  for computing  $T_{kk}$ . This subroutine accounts for  $2b^3$  floating point operations total.

**DLARFB:** since both  $V_{kk}$  and  $T_{kk}$  are triangular tiles,  $3b^3$  floating point operations are done in this subroutine.

**DTSQT2:** taking advantage of the triangular structure of  $R_{kk}$ , the factorization can be computed in  $2b^3$  floating point operations. The computation of the triangular  $T_{ik}$  tile can also be performed exploiting the structure of the Householder vectors built during the factorization (remember that the  $b$  reflectors can be represented as an identity tile on top of a square full tile). Since  $4/3b^3$  are needed to compute  $T_{ik}$ , the DTSQT2 accounts for  $10/3b^3$  floating point operations.

**DSSRFB:** exploiting the structure of the Householder reflectors and of the  $T_{ik}$  tile computed in DTSQT2, this subroutine needs  $5b^3$  floating point operations.

For each repetition of the outer loop in Algorithm 1, one DGEQT2,  $q - k$  DLARFB,  $p - k$  DTSQT2 and  $(p - k)(q - k)$  DSSRFB are performed for a total of  $2b^3 + 3(q - k)b^3 + 10/3(p - k)b^3 + 5(p - k)(q - k)b^3$ . Assuming, without loss of generality, that  $q < p$  and integrating this quantity over all the  $q$  repetitions of the outer loop in Algorithm 1, the total operation count for the QR factorization is

$$\begin{aligned} & \sum_{k=1}^q (2b^3 + 3(q - k)b^3 + \frac{10}{3}(p - k)b^3 + 5(p - k)(q - k)b^3) \\ & \simeq \frac{5}{2}q^2(p - \frac{q}{3})b^3 \\ & = \frac{5}{2}n^2(m - \frac{n}{3}). \end{aligned} \tag{4}$$

Equation (4) shows that the tiled algorithm for QR factorization needs 25% more floating point operations than the standard LAPACK algorithm.

Since the extra flops are due to the formation and application of multiple  $T_{ik}$  tiles at each step, using unblocked transformations in the tiled algorithm would lead to exactly the same operation count of the block LAPACK algorithm but this drastically affects the performance on a system with memory hierarchy.

### 3.3 Graph driven asynchronous execution

Following the approach presented in [8, 20], Algorithm 1 can be represented as a Directed Acyclic Graph (DAG) where nodes are elementary tasks that operate on  $b \times b$  blocks and where edges represent the dependencies among them. Figure 3 show the DAG when Algorithm 1 is executed on a matrix with  $p = q = 3$ . Note that the DAG has a recursive structure and, thus, if  $p_1 \geq p_2$  and  $q_1 \geq q_2$  then the DAG for a matrix of size  $p_2 \times q_2$  is a subgraph of the DAG




Figure 3: The dependency graph of Algorithm 1 on a matrix with  $p = q = 3$ .

for a matrix of size  $p_1 \times q_1$ . This property also holds for most of the algorithms in LAPACK.

Once the DAG is known, the tasks can be scheduled asynchronously and independently as long as the dependencies are not violated. A critical path can be identified in the DAG as the path that connects all the nodes that have the higher number of outgoing edges. Based on this observation, a scheduling policy can be used, where higher priority is assigned to those nodes that lie on the critical path. Clearly, in the case of our block algorithm for QR factorization, the nodes associated to the DGEQT2 subroutine have the highest priority and then three other priority levels can be defined for DTSQT2, DLARFB and DSSRFB in descending order.

This dynamic scheduling results in an out of order execution where idle time is almost completely eliminated since only very loose synchronization is required between the threads. Figure 4 shows part of the execution flow of Algorithm 1 on a 16-cores machine (8-socket Dual Opteron) when tasks are dynamically scheduled based on dependencies in the DAG. Each line in the execution flow shows which tasks are performed by one of the threads involved in the factorization.

Figure 4 shows that all the idle times, which represent the major scalability limit of the fork-join approach, can be removed thanks to the very low synchronization requirements of the graph driven execution. The graph driven execution also provides some degree of adaptivity since tasks are scheduled to threads depending on the availability of execution units.

### 3.4 Block Data Layout

The major limitation of performing very fine grain computations, is that the BLAS library generally have very poor performance on small blocks. This sit-




Figure 4: The execution flow for dynamic scheduling, out of order execution of Algorithm 1.

uation can be considerably improved by storing matrices in Block Data Layout (BDL) instead of the Column Major Format that is the standard storage format for FORTRAN arrays.




Figure 5: A comparison of Column Major storage format (*left*) and Block Data Layout (*right*).

Figure 5 compares Column Major Format (*left*) and Block Data Layout (*right*). In BDL a matrix is split into blocks and each block is stored into contiguous memory locations. Each block is stored in Column Major Format and blocks are stored in Column Major Format with respect to each other. As a result the access pattern to memory is more regular and BLAS performance is considerably improved. The benefits of BDL have been extensively studied in the past, for example in [18], and recent studies like [9] demonstrate how fine-granularity parallel algorithms can benefit from BDL.

## 4 Performance Results

The performance of the tiled QR factorization with dynamic scheduling of tasks has been measured on the systems listed in Table 2 and compared to the performance of the fork-join approach, i.e., the standard algorithm for block QR factorization of LAPACK associated with multithreaded BLAS.

	8-socket dual Opteron	2-socket quad Clovertown
Architecture	Dual-Core AMD Opteron™8214	Intel®Xeon®CPU X5355
Clock speed	2.2 GHz	2.66 GHz
# cores	$8 \times 2 = 16$	$2 \times 4 = 8$
Peak performance	70.4 Gflop/s	85.12 Gflop/s
Memory	62 GB	16 GB
Compiler suite	Intel 9.1	Intel 9.1
BLAS libraries	GotoBLAS-1.15 ACML-4.0.0	GotoBLAS-1.15 MKL-9.1

Table 2: Details of the systems used for the following performance results.

Figures 6, 7, 8, 9 report the performance of the QR factorization for both the block algorithm with dynamic scheduling, the LAPACK subroutine linked to multithreaded BLAS and a vendor implementation of the QR factorization (this last is only provided by the MKL and ACML libraries but no the GotoBLAS one). A block size of 200 has been used for the tiled algorithm while the block size for the LAPACK algorithm<sup>3</sup> has been tuned in order to achieve the best performance for all the combinations of architecture and BLAS library.

In each graph, two curves are reported for the block algorithm with dynamic scheduling; the solid curve shows its relative performance when the operation count is assumed equal to the one of the LAPACK algorithm reported in Section 2 while the dashed curve shows its “raw” performance, i.e. the actual flop rate computed with the exact operation count for this algorithm (given in Equation (4)). As already mentioned, the “raw performance” (dashed curve) is 25% higher than the relative performance (solid curve).

The graphs on the left part of each figure show the performance measured using the maximum number of cores available on each system with respect to the problem size. The graphs on the right part of each figure show the weak scalability, i.e. the flop rates versus the number of cores when the local problem size is kept constant ( $n_{loc}=5,000$ ) as the number of cores increases.

Figures 6, 7, 8, 9 show that, despite the higher operation count, the block algorithm with dynamic scheduling is capable of completing the QR factorization in less time than the LAPACK algorithm when the parallelism degree is high enough that the benefits of the asynchronous execution overcome the penalty

<sup>3</sup>the block size in the LAPACK algorithm sets the width of the panel.

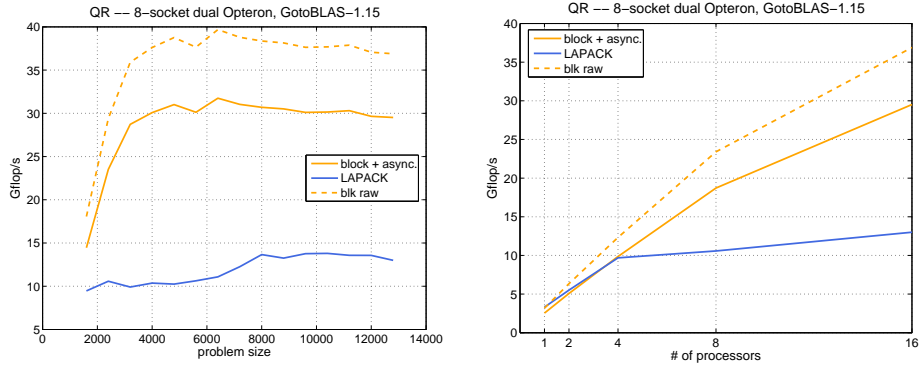


Figure 6: Performance of the tiled algorithm with dynamic scheduling using GotoBLAS-1.15 on an 8-socket Dual Opteron system. The dashed curve reports the raw performance of the tiled algorithm with dynamic scheduling, i.e., the performance as computed with the true operation count in Equation (4).




Figure 7: Performance of the tiled algorithm with dynamic scheduling using GotoBLAS-1.15 on a 2-socket Quad Clovertown system. The dashed curve reports the raw performance of the tiled algorithm with dynamic scheduling, i.e., the performance as computed with the true operation count in Equation (4).

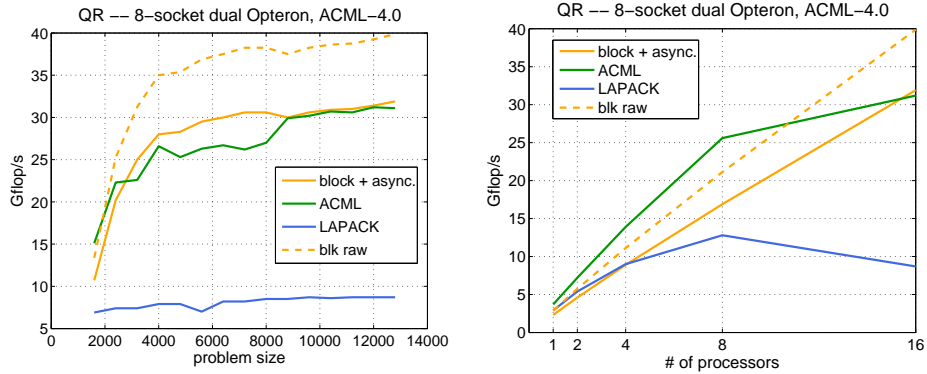


Figure 8: Performance of the tiled algorithm with dynamic scheduling using ACML-4.0.0 on an 8-socket Dual Opteron system. The dashed curve reports the raw performance of the tiled algorithm with dynamic scheduling, i.e., the performance as computed with the true operation count in Equation (4).

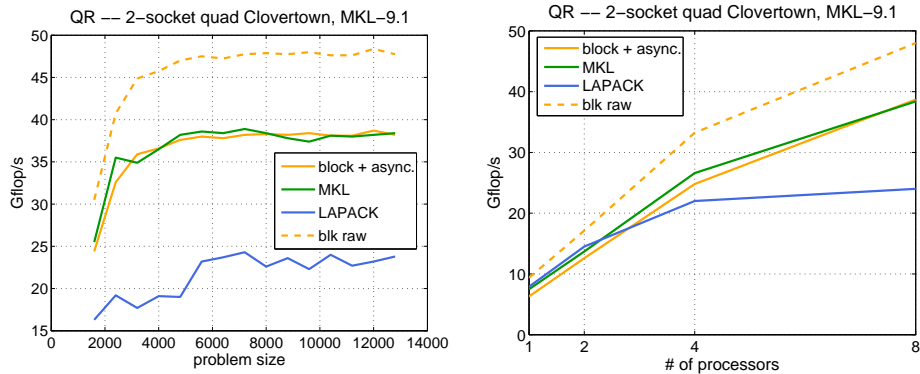


Figure 9: Performance of the tiled algorithm with dynamic scheduling using MKL-9.1 on an 2-socket Quad Clovertown system. The dashed curve reports the raw performance of the tiled algorithm with dynamic scheduling, i.e., the performance as computed with the true operation count in Equation (4).

of the extra flops. For lower numbers of cores, in fact, the fork-join approach has a good scalability and completes the QR factorization in less time than the block algorithm because of the lower flop count. Note that the actual execution rate of the block algorithm for QR factorization with dynamic scheduling (i.e., the dashed curves) is always higher than that of the LAPACK algorithm with multithreaded BLAS even for low numbers of cores.

The scalability data reported on the right side of Figures 6, 7, 8 and 9 clearly show how the fork-join approach has scalability limits that become more evident for higher degree of parallelism leading also to slowdown in some cases (see Figure 8). The scalability of the tiled algorithm with dynamic scheduling resembles that of vendors implementations. However, the data in the graphs suggest that the tiled algorithm may have an advantage over the vendor implementations when the degree of parallelism is higher than what was possible to achieve on the systems used.

The actual performance of the tiled algorithm, even if considerably higher than that of the fork-join one, is still far from the peak performance of the systems used for the measures. This is mostly due to two factors. First the nature of the BLAS operations involved; the DGEQR2 and the DLARFT in the LAPACK algorithm and the DGEQT2 and DTSQT2 in the block algorithm are based on Level 2 BLAS operations that, being memory bound, represent a limit for performance. Second, the performance of BLAS routines on small size blocks. The block size used in the experiments reported above is 200; this block size represents a good compromise between flexibility of the scheduler and performance of the BLAS operations but it is far from being ideal. Such a block size, in fact, does not allow a good task scheduling for smaller size problems and still the performance of BLAS operations is far from what can be achieved for bigger size blocks.

## 5 On the use of recursion in the block LAPACK algorithm

As pointed out in Section 2, the block QR algorithm in LAPACK suffers scalability limitations due to the intrinsically sequential nature of the panel reduction. Recursive techniques to perform the panel reduction as proposed by Elmroth et al. [13] can be used to weaken these limitations. Even if recursive panel factorization was introduced as a cache oblivious technique to improve data locality, it must be noted that the usage of recursion allows to perform some of the operations in the panel reduction by means of Level-3 BLAS subroutines which means that the panel can be parallelized to some extent on shared memory architectures. However, part of the panel reduction has still to be performed in Level-2 BLAS operations which represents a limitation to the scalability of the algorithm for the reasons discussed in Section 2. Due to the unavailability of source code for the algorithm presented in [13], a variant of this algorithm was implemented and its performance compared to that of the traditional LAPACK



block algorithm and the tiled algorithm described in Section 3. Specifically, in the implemented variant, recursion in the panel is stopped at a given point (set through parameter that must be tuned according to the architecture characteristics) after which conventional unblocked code (i.e. the LAPACK DGEQR2) is used to perform the remainder of the computations. The results of this comparison on the 8-socket Dual Opteron system in Table 2 are reported in Figure 10. Figure 10 shows that, even if the recursive panel provides much better per-

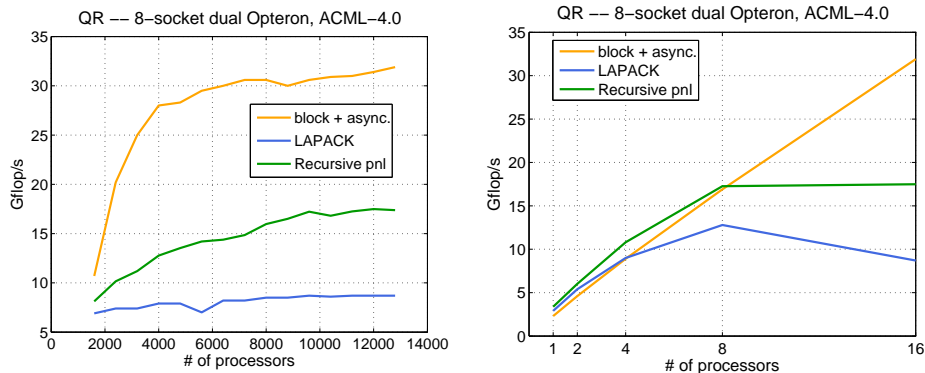


Figure 10: Comparison between the performance of the tiled algorithm with dynamic scheduling using ACML-4.0.0 on an 8-socket Dual Opteron system, the traditional LAPACK block algorithm and the block algorithm with recursive panel.

formance with respect to the traditional LAPACK block algorithm, it is still considerably slower than the proposed tiled algorithm for higher parallelism degrees.

## 6 Conclusion

By adapting known algorithms for updating the QR factorization of a matrix, we have derived an implementation scheme of the QR factorization for multicore architectures based on dynamic scheduling and block data layout. Although the proposed algorithm is performing 25% more flops than the regular algorithm, the gain in flexibility allows an efficient dynamic scheduling which enables the algorithm to scale almost perfectly when the number of cores increases.

While this paper only addresses the QR factorization, it is straightforward to derive with the same ideas the two important computational routines that consists in applying the Q factor to a set of vectors (see DORMQR in LAPACK) and constructing the Q-factor (see DORGQR in LAPACK).

The ideas behind this work can be extended in many directions:

**Explore techniques to reduce the extra flops.** It can be noted that the 25% overhead can be reduced by using nonsquare tiles. For example, us-

ing  $2b \times b$  tiles, the overhead reduces to 12.5%. Even though this may seem an effective solution to the problem of extra computations, using rectangular tiles yields a coarser granularity which limits the flexibility of the dynamic scheduling and results in poorer parallelization of the elementary operations. A potentially effective technique for reducing the amount of extra flops consists in accumulating the transformations computed inside the DTSQT2 operation in subsets of size  $s$  where  $s \ll b$  (this method is explained in [17]). The efficient implementation of this “internal blocking” technique is not easy to accomplish since it has to face the limitations of BLAS subroutines on small portions of data. The internal blocking is currently under investigation.

**Implement other linear algebra operations.** The LU factorization can be performed with an algorithm that is analogous to the QR one described in Section 3. This algorithm has been discussed in [23, 28] as a way of improving the out-of-core LU factorization. Even though the only difference between the tiled algorithms for the LU and QR factorizations is in the elementary operations, in the LU case the cost of the tiled algorithm is 50% higher than the LAPACK algorithm. For this reason, the benefits of the improved scalability may be visible only at very high processor counts or may not be visible at all. Techniques must be investigated to eliminate or reduce the extra cost.

The same idea of tiled operations may also be applied to other two sided transformations like Hessenberg reduction, Tridiagonalization and Bidiagonalization. In these transformations, Level 2 BLAS operations are predominant and panel reductions account for almost 50% of the time of a sequential execution. Breaking the panel into smaller tasks that can be executed in parallel with other tasks may yield considerable performance improvements.

**Enforcing data locality.** The results presented in [9] show that enforcing data locality and CPU affinity may provide considerable benefits. It must be noted that the improvements that can be expected on non-multicore SMPs are higher than on currently available multicore systems and this is due to the fact that on multicores, some of the higher level memories are shared between multiple cores. Moreover enforcing data locality has a major drawback in the fact that it seriously limits the scheduling of tasks since each core can only be assigned tasks that operate on data that resides on the memory associated with it. Preliminary results show that enforcing data locality and CPU affinity provides a slight speedup on the 8-socket Dual Opteron system which is a NUMA architecture. These techniques require further investigation.

**Implement the same algorithms in distributed memory systems.** The fact that the block algorithms for QR and LU factorizations only require loose synchronization between tasks make them also good candidates for

the implementation on distributed memory systems based on MPI communications.

**Implement the same algorithms on the STI Cell architecture.** In the STI Cell processor, no caches are present but a small, explicitly managed memory is associated to each core. Due to the small size of these local memories (only 256 KB), the LAPACK algorithms for LU and QR factorizations cannot be efficiently implemented. The block algorithms for LU and QR factorizations represent ideal candidates for the STI Cell architecture since they can be parallelized with a very fine granularity.

**Explore the usage of parallel programming environments.** The task of implementing Linear Algebra operations with dynamic scheduling of tasks on multicore architectures can be considerably simplified by the use of graph driven parallel programming environments. One such environment is SMP Superscalar [5] developed at the Barcelona Supercomputing Center. SMP Superscalar addresses the automatic exploitation of the functional parallelism of a sequential program in multicore and SMP environments. The focus is on the portability, simplicity and flexibility of the programming model. Based on a simple annotation of the source code, a source to source compiler generates the necessary code and a runtime library exploits the existing parallelism by building at runtime a task dependency graph. The runtime takes care of scheduling the tasks and handling the associated data. Besides, a temporal locality driven task scheduling can be implemented.

## Acknowledgment

We would like to thank John Gilbert from University of California at Santa Barbara for sharing with us its 8-socket dual Opteron system.

## References

- [1] <http://top500.org>.
- [2] <http://www.intel.com/cd/software/products/asmo-na/eng/307757.htm>.
- [3] <http://developer.amd.com/acml.jsp>.
- [4] Teraflops research chip. <http://www.intel.com/research/platform/terascale/teraflops.htm>.
- [5] SMP Superscalar (SMPSs) User's Manual, July 2007. [www.bsc.es/media/1002.pdf](http://www.bsc.es/media/1002.pdf).

- [6] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK User's Guide*. SIAM, Philadelphia, 3. edition, 1999.
- [7] M. W. Berry, J. J. Dongarra, and Y. Kim. A parallel algorithm for the reduction of a nonsymmetric matrix to block upper-hessenberg form. *Parallel Comput.*, 21(8):1189–1211, 1995.
- [8] A. Buttari, J. Dongarra, J. Kurzak, J. Langou, P. Luszczek, and S. Tomov. The impact of multicore on math software. In *PARA 2006, Umeå Sweden*, June 2006.
- [9] E. Chan, E. S. Quintana-Orti, G. Quintana-Orti, and R. van de Geijn. Supermatrix out-of-order scheduling of matrix operations for SMP and multi-core architectures. In *SPAA '07: Proceedings of the nineteenth annual ACM symposium on Parallel algorithms and architectures*, pages 116–125, New York, NY, USA, 2007. ACM Press.
- [10] J. Choi, J. Demmel, I. Dhillon, J. Dongarra, S. Ostrouchov, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. ScaLAPACK: A portable linear algebra library for distributed memory computers - design issues and performance. *Computer Physics Communications*, 97:1–15, 1996. (also as LAPACK Working Note #95).
- [11] J. Choi, J. Dongarra, S. Ostrouchov, A. Petitet, D. W. Walker, and R. Clinton Whaley. A proposal for a set of parallel basic linear algebra subprograms. In *PARA '95: Proceedings of the Second International Workshop on Applied Parallel Computing, Computations in Physics, Chemistry and Engineering Science*, pages 107–114, London, UK, 1996. Springer-Verlag.
- [12] L. Dagum and R. Menon. OpenMP: An industry-standard API for shared-memory programming. *IEEE Computational Science and Engineering*, 05(1):46–55, 1998.
- [13] E. Elmroth and F. G. Gustavson. Applying recursion to serial and parallel QR factorization leads to better performance. *IBM Journal of Research and Development*, 44(4):605–624, July 2000.
- [14] Message Passing Interface Forum. MPI: A Message-Passing Interface Standard. *The International Journal of Supercomputer Applications and High Performance Computing*, 8, 1994.
- [15] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 3rd edition, 1996.
- [16] K. Goto and R. van de Geijn. High-performance implementation of the level-3 blas. Technical Report TR-2006-23, The University of Texas at Austin, Department of Computer Sciences., 2006. FLAME Working Note 20.

- [17] B. C. Gunter and R. A. van de Geijn. Parallel out-of-core computation and updating of the QR factorization. *ACM Trans. Math. Softw.*, 31(1):60–78, 2005.
- [18] F. G. Gustavson. New generalized data structures for matrices lead to a variety of high performance algorithms. In *PPAM '01: Proceedings of the th International Conference on Parallel Processing and Applied Mathematics-Revised Papers*, pages 418–436, London, UK, 2002. Springer-Verlag.
- [19] J. Kurzak, A. Buttari, and J. Dongarra. Solving systems of linear equations on the CELL processor using Cholesky factorization. Technical Report UT-CS-07-596, Innovative Computing Laboratory, University of Tennessee Knoxville, April 2007.
- [20] J. Kurzak and J. Dongarra. Implementing linear algebra routines on multi-core processors with pipelining and a look ahead. LAPACK Working Note 178, September 2006. Also available as UT-CS-06-581.
- [21] F. Mueller. Pthreads library interface, 1993.
- [22] D. Pham, S. Asano, M. Bolliger, M. N. Day, H. P. Hofstee, C. Johns, J. Kahle, A. Kameyama, J. Keaty, Y. Masubuchi, M. Riley, D. Shippy, D. Stasiak, M. Suzuoki, M. Wang, J. Warnock, S. Weitzel, D. Wendel, T. Yamazaki, and K. Yazawa. The design and implementation of a first-generation CELL processor. In *IEEE International Solid-State Circuits Conference*, pages 184–185, 2005.
- [23] E. Quintana-Orti and R. van de Geijn. Updating an LU factorization with pivoting. Technical Report TR-2006-42, The University of Texas at Austin, Department of Computer Sciences, 2006. FLAME Working Note 21.
- [24] R. Schreiber and C. van Loan. A storage-efficient WY representation for products of Householder transformations. *SIAM J. Sci. Stat. Comput.*, 10(1):53–57, 1989.
- [25] G. W. Stewart. *Matrix Algorithms*, volume 1. SIAM, Philadelphia, 1. edition, 1998.
- [26] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997.
- [27] R. C. Whaley, A. Petitet, and J. Dongarra. Automated empirical optimization of software and the ATLAS project. *Parallel Computing*, 27(1–2):3–25, 2001.
- [28] E. L. Yip. FORTRAN Subroutines for Out-of-Core Solutions of Large Complex Linear Systems. Technical Report CR-159142, NASA, November 1979.