

QR factorization for the Cell Broadband Engine

Jakub Kurzak^{a,*} and Jack Dongarra^{a,b,c}

^a *Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN, USA*

^b *Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA*

^c *School of Mathematics and School of Computer Science, University of Manchester, Manchester, UK*

Abstract. The QR factorization is one of the most important operations in dense linear algebra, offering a numerically stable method for solving linear systems of equations including overdetermined and underdetermined systems. Modern implementations of the QR factorization, such as the one in the LAPACK library, suffer from performance limitations due to the use of matrix–vector type operations in the phase of panel factorization. These limitations can be remedied by using the idea of updating of QR factorization, rendering an algorithm, which is much more scalable and much more suitable for implementation on a multi-core processor. It is demonstrated how the potential of the cell broadband engine can be utilized to the fullest by employing the new algorithmic approach and successfully exploiting the capabilities of the chip in terms of single instruction multiple data parallelism, instruction level parallelism and thread-level parallelism.

Keywords: Cell broadband engine, multi-core, numerical algorithms, linear algebra, matrix factorization

1. Introduction

State of the art, numerical linear algebra software utilizes *block algorithms* in order to exploit the memory hierarchy of traditional cache-based systems [1–4]. Public domain libraries such as LAPACK [5] and ScaLAPACK [6] are good examples. These implementations work on square or rectangular submatrices in their inner loops, where operations are encapsulated in calls to *Basic Linear Algebra Subroutines* (BLAS) [7], with emphasis on expressing the computation as level 3 BLAS (*matrix–matrix* type) operations.

The fork-and-join parallelization model of these libraries has been identified as the main obstacle for achieving scalable performance on new processor architectures. The arrival of multi-core chips increased the demand for new algorithms, exposing much more thread-level parallelism of much finer granularity. This paper presents an implementation of the QR factorization based on the idea of updating the QR factorization. The algorithm, referred to as *tile QR*, processes the input matrix by small square blocks of fixed size, providing for great data locality and fine granularity of parallelization. In the case of the *Cell Broadband Engine* (Cell BE), it also readily solves the problem of limited

size of private memory associated with each computational core.

Section 2 provides a brief discussion of related work. Section 3 presents a short description of the algorithm. Section 4 gives a quick overview of processor architecture, followed by a discussion of vectorization and parallelization of the code. Sections 5–7 follow with the presentation of the performance results, conclusions and possibilities for future developments.

This article focuses exclusively on the aspects of efficient implementation of the algorithm and makes no attempts at discussing the issues of numerical quality of the results related to the use of single precision with truncation rounding, and lack of support for NaNs and denorms (which is the way the Cell BE implements single precision floating point operations).

2. Related work

The first experiences with implementing dense matrix operations on the Cell BE were reported by Chen et al. [8]. Performance results were presented for single precision matrix multiplication and the solution of dense systems of linear equations in single precision using LU factorization. The authors of this article refined this work by using the LU factorization in single

*Corresponding author. E-mail: kurzak@eecs.utk.edu.

precision along with the technique of iterative refinement to achieve double precision accuracy of the final solution [9].

Cholesky factorization was identified as an algorithm rendering itself easily to formulation as *algorithm by tiles*. It was subsequently implemented, delivering parallel scaling far superior to that of LU (in its classic form). The mixed-precision iterative refinement technique was used to solve symmetric positive definite systems of equations, producing results with double precision accuracy while exploiting the speed of single precision operations [10].

Other developments worth noting were further refinements of the work on optimizing the matrix multiplication, first by Hackenberg [11,12] and then by Alvaro et al. [13]. It is also worthwhile to note that impressive performance was achieved by Williams et al. for sparse matrix operations on the Cell BE [14].

The most important issue in performance optimization of orthogonal transformations is aggregation of transformations leading to efficient use of the memory system. The idea was first demonstrated by Dongarra et al. [15], later by Bischof and van Loan [16], and yet later by Schreiber and van Loan [17], resulting in the compact *WY* technique for accumulating Householder reflectors.

Elmroth and Gustavson [18–20] generalized this work to produce high performance recursive QR factorization. In this work the problem of reducing the amount of extra floating point operations was addressed by the introduction of *mini blocking/register blocking*, referred to as *inner blocking* in this article. Serial implementation was presented as well as parallel implementation with dynamic scheduling of tasks on symmetric multiprocessors.

One of the early references discussing methods for updating matrix factorizations is the paper by Gill et al. [21]. Berry et al. successfully applied the idea of using orthogonal transformations to annihilate matrix elements by tiles, in order to achieve a highly parallel distributed memory implementation of matrix reduction to the block upper-Hessenberg form [22].

It is crucial to note that the technique of processing the matrix by square blocks only provides performance in tandem with data organization by square blocks, a fact initially observed by Gustavson [23,24] and recently investigated in depth by Gustavson et al. [25]. The layout is referred to as *Square Block (SB)* format by Gustavson et al. and as *Block Data Layout (BDL)* in this work. The paper by Elmroth et al. [26] gives an excellent introduction to many of the important issues

concerning deep memory hierarchies and the use of recursion and hybrid data structures and also contains a section on the QR factorization.

The idea of inner blocking was exploited by Gunter and van de Geijn [27] to achieve high performance for an *Out-Of-Core* (OOC) implementation of QR factorization. Recently, the combination of processing the matrix by tiles and storing the matrix by tiles was applied to achieve high performance for matrix factorization on “standard” (x86 and alike) multi-core architectures. Buttari et al. [28,29] reported initial results for QR factorization without inner blocking [28] and then results for QR and LU factorizations with inner blocking and also for Cholesky factorization [29] (inner blocking does not apply here).

This article combines the idea of storing and processing the matrix by tiles to achieve a high performance implementation of the QR factorization, while relying on inner blocking to keep the overhead extra floating point operations negligible.

3. Algorithm

The tile QR algorithm is very well documented in the literature [27,29]. The algorithm produces the same R factor as the classic algorithm, e.g., the implementation in the LAPACK library (elements may differ in sign). However, a different set of Householder reflectors is produced and a different procedure is required to build the Q matrix. Whether the Q matrix is actually needed depends on the application.

The algorithm relies on four basic operations implemented by four computational kernels (Fig. 1). Here the LAPACK-style naming convention, introduced by Buttari et al. [29], is followed. The capital letter S at the beginning indicates the use of single precision.

SGEQRT: The kernel performs the QR factorization of a diagonal tile of the input matrix and produces an upper triangular matrix R and a unit lower triangular matrix V containing the Householder reflectors. The kernel also produces the upper triangular matrix T as defined by the compact *WY* technique for accumulating Householder reflectors [16,17]. The R factor overrides the upper triangular portion of the input and the reflectors override the lower triangular portion of the input. The T matrix is stored separately.




Fig. 1. Basic operations of the tile QR factorization.

STSQRT: The kernel performs the QR factorization of a matrix built by coupling an R factor, produced by SGEQRT or a previous call to STSQRT, with a tile below the diagonal tile. The kernel produces an updated R factor, a square matrix V containing the Householder reflectors and the matrix T resulting from accumulating the reflectors V . The new R factor overrides the old R factor. The block of reflectors overrides the square tile of the input matrix. The T matrix is stored separately.

SLARFB: The kernel applies the reflectors calculated by SGEQRT to a tile to the right of the diagonal tile, using the reflectors V along with the matrix T .

SSSRFB: The kernel applies the reflectors calculated by STSQRT to two tiles to the right of the tiles factorized by STSQRT, using the reflectors V and the matrix T produced by STSQRT.

LAPACK-style block QR factorization relies on the compact WY technique for accumulating Householder reflectors in order to express computation in terms of level 3 BLAS (matrix–matrix) operations. The technique requires calculation of a square matrix T per each panel of the input matrix, where the size of T is equal to the width of the panel and, most of the time, much smaller than the height of the panel. In this case, the overhead associated with manipulating the T matrices is negligible.




Fig. 2. Inner blocking of the tile operations.

In a naive implementation of the tile QR factorization, a T matrix is produced for each square tile of the panel and used in updating the tiles to the right. This approach results in 25% more operations than the standard QR algorithm.

It can be observed, however, that in principle the updating algorithm can be implemented relying on level 2 BLAS (matrix–vector) operations, without the use of the T matrices and associated overheads. Interestingly, in such case, the updating algorithm results in the same number of floating point operations as the standard QR algorithm ($2MN^2 - (2/3)N^3$). Obviously, such implementation has to perform poorly due to the memory-bound nature of level 2 BLAS.

The key to achieving performance is to find the right trade-off between extra operations and memory intensity. This can be achieved by implementing the tile operations using the block algorithms within the tile. With internal block size much smaller than tile size, resulting T matrices are not “full” upper triangular matrices, but instead consist of upper triangular blocks along the diagonal of size equal to the inner block size (Fig. 2).

4. Implementation

The process of implementing the algorithm on the Cell BE included a few design choices (some of them arbitrary), which the authors would like to discuss here.

The tile size of 64×64 is a common practice for implementing dense matrix operations in single precision on the Cell BE. It has been shown that at this size matrix multiplication kernels can achieve over 99% of the SPE peak [13]. At the same time, the DMA transfer of a single tile fully utilizes the memory system consisting of 16 banks interleaved on a cache line boundary of 128 bytes. For these reasons the tile size of 64×64 is chosen here. Also, since the code is only a *proof-of-concept* prototype, only problem sizes divisible by 64 are handled.

Typically, the inner block size is chosen using some method of auto-tuning. In this case the inner block size of 4 has been chosen arbitrarily, mostly for coding simplicity stemming from the size of the SIMD vector of four single precision floating point elements. It turns out, however, that even at such a small size, the code does not become memory-bound thanks to the small, flat latency of the Local Store. It also introduces an acceptable amount of extra floating point operations. It is very unlikely that a different choice would yield significantly better results.

Finally, it has been chosen to implement the SGEQRT and STSQRT kernels using LAPACK-style block algorithm internally within the kernels. Potentially, the tile algorithm could also be used inside the kernels. Such approach would, however, dramatically complicate the application of the reflectors. The update operation could not be implemented efficiently.

4.1. Cell BE architecture overview

The Cell BE has been available since 2005 and is well known to the numerical computing community. It is not, however, a main-stream solution and is often perceived as a special-purpose accelerator device. As a result, the authors restrain from an extensive overview of the architecture, but do introduce the basic Cell BE vocabulary, and the highlights of the chip computing core design.

The Cell BE is an innovative multi-core architecture consisting of a standard processor, the *Power Processing Element* (PPE), and eight short-vector, *Single Instruction Multiple Data* (SIMD) processors, referred to as the *Synergistic Processing Elements* (SPEs). The SPEs are equipped with *scratchpad memory* referred to as the *Local Store* (LS) and a *Memory Flow Controller* (MFC), to perform *Direct Memory Access* (DMA) transfers of code and data between the system memory and the Local Store.

The core of the SPE is the *Synergistic Processing Unit* (SPU). The SPU is a RISC-style SIMD processor featuring 128 general purpose registers and 32-bit fixed-length instruction encoding. An SPU implements instructions to perform single and double precision floating point arithmetics, integer arithmetics, logicals, loads and stores, compares and branches. SPU's nine execution units are organized into two pipelines, referred to as the odd and even pipeline. Instructions are issued in-order, and two independent instructions can be issued simultaneously if they belong to different pipelines (what is referred to as *dual-issue*).

SPU executes code from the Local Store and operates on data residing in the Local Store, which is a fully pipelined, single-ported, 256 kB of *Static Random Access Memory* (SRAM). Load and store instructions are performed within local address space, which is untranslated, unguarded and noncoherent with respect to the system address space. Loads and stores transfer 16 bytes of data between the register file and the Local Store and complete with fixed six-cycle delay and without exception.

4.2. SIMD vectorization

The keys to maximum utilization of the SPEs are highly optimized implementations of the computational kernels, which rely on efficient use of the short-vector SIMD architecture. For the most part, the kernels are developed by applying standard loop optimization techniques, including tiling, unrolling, reordering, fusion, fission, and sometimes also collapsing of loop nests into one loop spanning the same iteration space with appropriate pointer arithmetics. Tiling and unrolling are mostly dictated by Local Store latency and the size of the register file, and aim at hiding memory references and reordering of vector elements, while balancing the load of the two execution pipelines. Due to the huge size of the SPU's register file, unrolling is usually quite extensive.

Most of the techniques used to build the tile QR kernels are similar to those used to build the Cholesky factorization kernels [10] and the high performance SGEMM (matrix multiplication) kernels [13]. The main difference here is the use of inner blocking, which substantially narrows down the design choices. Most importantly, the use of inner blocking imposes the structure of nested loops, where a single iteration of the outermost loop implements a single block operation. For instance, one iteration of the outermost loop of SGEQRT and STSQRT produces a block of four re-

flectors and the associated 4×4 upper triangular block of T ; one iteration of the outermost loop of SLARFB and SSSRFB applies a block of four reflectors and utilizes a 4×4 block of T .

All kernels are written in C using mostly SIMD language extensions (intrinsics) and sometimes inline assembly. Table 1 shows the size of the C source code, assembly code and object code of the kernels. Since all the code is hand-written, it gives some idea of its complexity. Table 2 reports the performance of the kernels in terms of Gflop/s and percentage of the peak (of a single SPE). The authors are only able to achieve this performance while compiling the kernels with SPU GCC 3.4.1. The paragraphs that follow briefly discuss technicalities related to each of the kernels.

The SSSRFB kernel, being the most performance-critical (contributing the most floating point operations), is optimized the most. This kernel actually allows for the most extensive optimizations, since all loops have fixed boundaries. Therefore, the technique of collapsing loop nests into one loop is used here, along with double-buffering, where odd and even iterations overlap each other's arithmetic operations with loads, stores and vector element permutations. Also, input arrays are constrained with 16 kB alignment, and

pointer arithmetic is implemented by calculating data offsets from the iteration variable (loop counter) by using bit manipulation. It turns out that all these operations can be implemented using quadword shifts, rotations and shuffles, and placed in the odd pipeline, where they can be hidden behind floating point arithmetics. Interestingly, it also turns out that for some loops, mostly rearranging vector elements, shuffles can be replaced with bit select operations to yield more balanced odd and even pipeline utilization.

In principle, the SSSRFB kernel shares many properties with the SGEMM kernel, and one could expect performance similarly close to the peak (99.8% was reported for SGEMM [13]). This is not the case for a few reasons. The main contributor of performance loss of the SSSRFB kernel is the prologue and epilogue code of the inner loops, which cannot be hidden behind useful work. Also, the reported performance of the SSSRFB kernel cannot reach the peak because of the extra operations, related to the application of the T matrix, which are not accounted for in the standard formula for operation count, $4b^3$. The actual number of operations is $4b^3 + sb^2$, where s is the size of internal blocking.

The STSQRT kernel has been identified as the second most critical for performance. STSQRT produces data, which is consumed by many SSSRFB kernels in parallel. As a result, it is important, in the context of parallel scheduling, that the STSQRT kernel executes in a shorter time than the SSSRFB kernel. This task proved quite difficult and the STSQRT kernel took significant coding effort and results in the longest code. One fact that is taken advantage of is that, at each step (each outer loop iteration), a block of reflectors of the same size is produced (64×4). This allows for performing the panel factorization (production of four reflectors) to be executed entirely in the register file, using 64 registers. First, the panel is loaded, then four steps are performed, each producing one reflector and applying it to the rest of the panel, then the panel is stored. The whole procedure is completely unrolled to one block of *straight-line* code.

As extreme as it might seem, this step alone proves to be insufficient to deliver the desired performance. The operations applying the panel to the remaining submatrix have to also be extensively optimized by heavy unrolling and addressing of special cases (e.g., different treatment of odd and even loop boundaries). It took significant effort to accomplish execution time slightly below the one of the SSSRFB kernel at an execution rate of less than half of the peak.

Table 1

Complexity characteristics of tile QR SPE kernels

Kernel name	Lines of code in C ^a	Lines of code in ASM ^b	Object size (KB) ^c
SSSRFB	1600	2200	8.8
STSQRT	1900	3600	14.2
SLARFB	600	600	2.2
SGEQRT	1600	2400	9.0
Total	5700	8800	34.2

Note: Bold font indicates the most complex kernel.

^aSize of code in C before or after preprocessing, whichever is smaller; ^bsize of code in assembly after removing the `.align` statements; ^csum of `.text` and `.rodata` sections (not size of the `.o` file).

Table 2

Performance characteristics of tile QR SPE kernels

Kernel name	Exec. time (μ s) ^a	Flop count formula ^b	Exec. rate (Gflop/s) ^c	Fraction of peak (%) ^{a,d}
SSSRFB	47	$4b^3$	22.20	87
STSQRT	46	$2b^3$	11.40	45
SLARFB	41	$2b^3$	12.70	50
SGEQRT	57	$(4/3)b^3$	6.15	24

Note: Bold font indicates the most performance-critical kernel.

^aValues are rounded; ^btile size $b = 64$; ^cvalues are truncated; ^dsingle SPE.

There is less to be said about the two remaining kernels, SLARFB and SGEQRT. The SLARFB kernel turns out to deliver very good performance without much effort. On the other hand, SGEQRT does not deliver good performance despite efforts similar to the STSQRT kernel. This is to be expected, however, since none of the loops have fixed boundaries. This kernel is executed the least and its poor performance does not affect the overall performance much. The situation is analogous to the SPOTRF kernel of the Cholesky factorization, for which similar performance is reported (roughly 6 Gflop/s [10]).

The last technical detail, which has not been revealed so far, is that the T factors are stored in a compact format. Each element of T is pre-splatted across a 4-element vector; each 4×4 triangular block of T is stored in a column of 10 vectors and the T array contains 16 such columns of overall size of 2560 bytes.

4.3. Parallelization – single Cell BE

For the distribution of work for parallel execution on the SPEs, static 1D cyclic partitioning is used, shown in Fig. 3. The effect of “wrapping” the SPEs assignment from one step to another results in pipelining of factorization steps, basically implementing the technique known in linear algebra as the *lookahead*. Following Fig. 3, one can observe that SPE 5 can start factorizing the second panel as soon as SPE 1 finishes the first SSSRFB operation.

Static work partitioning makes the synchronization extremely straightforward. With all the work predetermined, each SPE can proceed on its own, and only needs to check if dependencies are satisfied for each operation. Figure 4 shows the dependencies between

tasks of the tile QR algorithm expressed as a Direct Acyclic Graph (DAG).

Before fetching a tile for an operation in a given step, the SPE needs to check if the preceding step has completed on that tile. The SPE does that by looking up a progress table in its Local Store. The progress table contains the global progress information and is replicated on all SPEs. The progress table holds one entry (byte) for each tile of the input matrix, indicating the number of the step which has completed on that tile. At the completion of an operation, an SPE broadcasts the progress information to all progress tables with an LS-to-LS DMA.

As one can see, the scheme implements the right-looking (aggressive) variant of the algorithm. Although different scenarios can be easily imagined, this version makes sense from the standpoint of ease of implementation. In this arrangement, an SPE factorizing the panel can hold the diagonal tile in place, while streaming the tiles below diagonal through Local Store. Similarly, an SPE updating a column of the trailing submatrix can hold the topmost tile in place, while streaming the tiles below it through Local Store. Data reuse is accomplished this way, which minimizes the traffic to main memory. It needs to be pointed out, though, that this is absolutely not necessary from the standpoint of memory bandwidth. The tile QR factorization is so compute intensive that all memory traffic can easily be hidden behind computation with data reuse or without it.

At each step, the tiles of the input matrix are exchanged between the main memory and local store. Important aspect of the communication is double-buffering. Since work partitioning is static, upcoming operations can be anticipated and the necessary data




Fig. 3. Cyclic partitioning of work to eight SPEs in the five consecutive steps of factorizing a 5×5 block matrix.

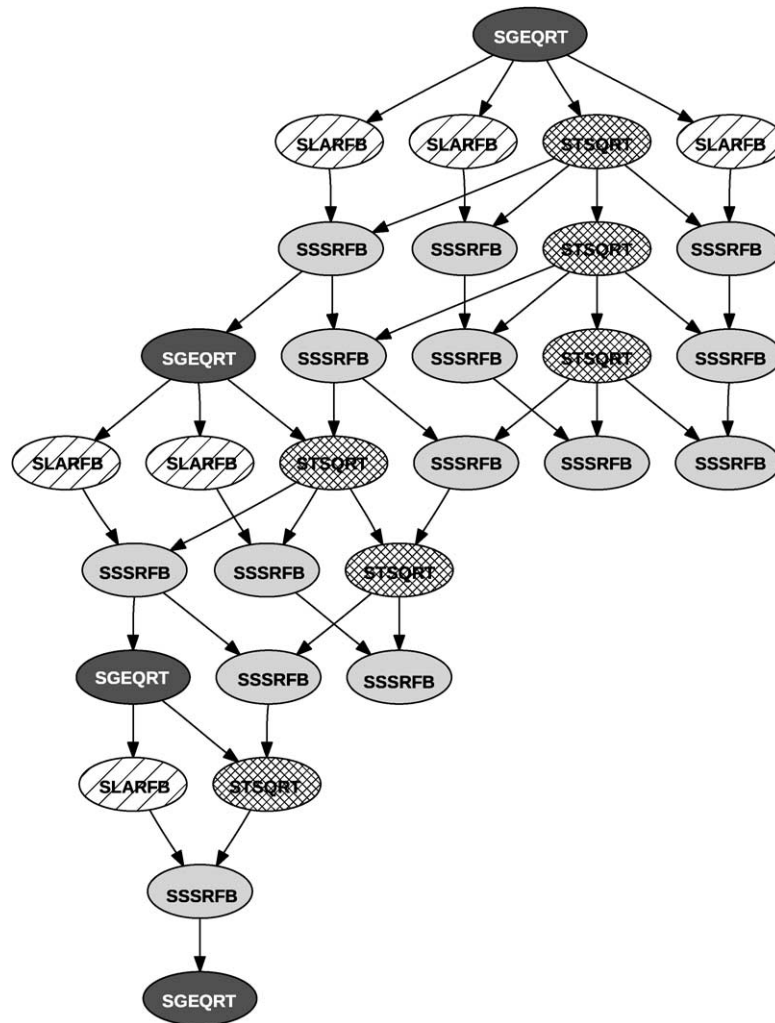


Fig. 4. The direct acyclic graph of a tile QR factorization of a 4×4 block matrix (some dependencies are omitted for clarity).

fetches. In fact all data buffers are duplicated and, at each operation, a prefetch of data is initiated for the following operation (subject to dependency check). If the prefetch fails for dependency reasons, data is fetched in a blocking mode right before the operation. Algorithm 1 shows the mechanism of double buffering in the tile QR implementation. Figure 5 shows the execution trace of factorizing a 512×512 matrix using all the eight SPEs.

Alternatively to hand-coded communication and synchronization, a generic DAG scheduling framework could be used. To the best of the authors' knowledge, the only such framework available for the Cell BE today is the *Cell Superscalar* (CellSs) project from the Barcelona Supercomputer Center [30,31]. However, currently the software is not competitive, in terms of

Algorithm 1. Double buffering of communication in the tile QR implementation.

- 1: **while** more work to do **do**
 - 2: **if** data not prefetched **then**
 - 3: wait for dependencies
 - 4: fetch data
 - 5: **end if**
 - 6: **if** more work to follow **then**
 - 7: **if** dependencies met **then**
 - 8: prefetch data
 - 9: **end if**
 - 10: **end if**
 - 11: compute
 - 12: swap buffers
 - 13: **end while**
-




Fig. 5. Execution trace of a factorization of a 512×512 matrix. (Total time: 1645 μ s, execution rate: 109 Gflop/s.)

performance, with the approach presented here due to the handling of task scheduling on the PPE.

4.4. Parallelization – Dual Cell BE

Given single-Cell BE implementation, extension to a dual-Cell BE system, like the IBM QS20 blade, is relatively straightforward. A single PPE process can launch 16 SPE threads, eight on each Cell BE. The single-Cell BE code is going to run correctly on a dual-Cell BE system by simply increasing the number of SPEs to 16.

The only problem is the one of performance of the memory system. The QS20 blade is a *Non-Uniform Memory Access* (NUMA) system. Each Cell BE is associated with a separate memory node. Peak bandwidth to the local node is 25.6 GB/s. Cross-traffic, however, is handled at a much lower bandwidth (roughly half of that number). It is important, then, that each SPE satisfies its data needs mostly from the local memory node.

This situation is addressed by duplicating the input matrix in both memory nodes (*libnuma* is used for correct memory placement). Each SPE reads data only from the local node, but writes data to both nodes. From the perspective of the shared memory model, it can be viewed as a manual implementation of the write-back memory consistency protocol. From the perspective of a distributed memory model, it can be viewed as non-blocking collective communication (broadcast) or as one-sided communication. The obvious limitation is that the approach would not be scalable to larger NUMA systems. As of today, however, larger Cell BE-based NUMA systems do not exist.

One technical detail to be mentioned here are the acknowledgment DMAs implementing the synchronization protocol between SPEs. When 16 SPEs are used, each SPE needs to send 16 acknowledgment messages following a write of data to the system memory. The acknowledgment DMA is fenced with the data DMA and the SPE also sends such message to its own

progress table (since 16 messages are sent and not 15). The DMA request queue is, however, only 16 entries deep and issuing 16 acknowledgment requests at the same time stalls data transfers until some requests clear the queue. A simple remedy is the use of a DMA list with 16 elements, where the elements point to appropriate Local Store locations of the other SPEs. The code alternates between two such lists in the double-buffered communication cycle.

5. Results

Results presented in this section are produced on one and two 3.2 GHz Cell BEs of the QS20 dual-socket blade running Fedora Core 7 Linux and on a PlayStation 3 running Fedora Core 7 Linux. The code is cross-compiled using x86 SDK 3.0, although, as mentioned before, the kernels are cross-compiled with an old x86 SPU GCC 3.4.1 cross-compiler, since this compiler yields the highest performance. The results are checked for correctness by comparing the R factor produced by the algorithm to the R factor produced by a call to the LAPACK routine SGEQRF ran on the PPE. The tile algorithm produces a valid R factor, where elements may differ in sign from the elements of the R factor produced by the block algorithm. For that reason, absolute values of elements are compared.

It also needs to be mentioned that the implementation utilizes *Block Data Layout* (BDL) [32,33], where each tile is stored in a continuous 16 kB portion of the main memory, which can be transferred in a single DMA, what puts an equal load on all 16 memory banks. Tiles are stored in the row-major order, and also data within tiles is arranged in the row-major order, a common practice on the Cell BE. Translation from standard, (FORTRAN) layout to BDL can be implemented very efficiently on the Cell BE [9]. Here the translation is not included in timing results. Also, in order to avoid the problem of TLB misses, all the memory is allocated in huge TLB pages and “faulted in” at

Table 3

Selected performance points of the tile QR factorization in single precision on a single 3.2 GHz Cell BE of the Sony PlayStation 3 (6 SPEs) dual-socket blade

Matrix size ^a	Execution rate (Gflop/s) ^b	Fraction of peak (%) ^c	Fraction of SSSRFB peak (%) ^c
128	12	8	10
256	40	27	31
384	74	48	56
448	80	52	60
640	101	66	76
960	114	75	86
1280	120	78	90
1536	123	80	93
3072	128	83	96
4096	128	84	97

Note: Bold font indicates the point of exceeding half of the processor peak.

^aSquare matrices were used; ^bvalues are truncated; ^cvalues are rounded.

Table 4

Selected performance points of the tile QR factorization in single precision on a single 3.2 GHz Cell BE (8 SPEs) of the IBM QS20 dual-socket blade

Matrix size ^a	Execution rate (Gflop/s) ^b	Fraction of peak (%) ^c	Fraction of SSSRFB peak (%) ^c
128	12	6	7
256	40	20	23
384	81	40	46
512	109	53	62
768	137	67	77
1024	150	73	85
1280	157	77	89
1536	162	79	91
2048	166	81	94
4096	171	84	97

Note: Bold font indicates the point of exceeding half of the processor peak.

^aSquare matrices were used; ^bvalues are truncated; ^cvalues are rounded.

initialization. As a result, an SPE never incurs a TLB miss during the run. For single-Cell BE runs as well as dual-Cell BE run correct memory placement is enforced using the *libnuma* library.

Tables 3–5 and Fig. 6 show the performance of the algorithm in Gflop/s, while using the standard formula, $2MN^2 - (2/3)N^3$, for operation count. Tables 3–5 also show percentage of system peak and percentage of the SSSRFB kernel performance times the number of SPEs.

Table 5

Selected performance points of the tile QR factorization in single precision on two 3.2 GHz Cell BE of the IBM QS20 dual-socket blade (16 SPEs)

Matrix size ^a	Execution rate (Gflop/s) ^b	Fraction of peak (%) ^c	Fraction of SSSRFB peak (%) ^c
128	12	3	3
256	38	9	11
384	81	20	23
512	137	34	39
768	212	52	60
1024	266	65	75
1536	307	75	87
2048	322	79	91
3072	335	82	95
4096	340	83	96

Note: Bold font indicates the point of exceeding half of the system peak.

^aSquare matrices were used; ^bvalues are truncated; ^cvalues are rounded.

The presented implementation crosses half of the peak performance for problems of size 512×512 on a single Cell BE and 768×768 on two Cell BEs. The performance of 150 Gflop/s is reached for a problem of size 1024×1024 using single Cell BE. The performance of 300 Gflop/s is crossed for a problem of size 1536×1536 using dual Cell BE system. For a problem of size 4096×4096 performance of 171 Gflop/s is reached using single Cell BE and 340 Gflop/s using two Cell BEs.

The code used to produce the reported results is freely available through one of the author's web site, <http://www.cs.utk.edu/~kurzak/>.

6. Conclusions

The presented implementation of tile QR factorization on the Cell BE allows for factorization of a 4000×4000 dense matrix in single precision in exactly half a second. To the authors' knowledge, at present, it is the fastest reported time of solving such problem by any semiconductor device implemented on a single semiconductor die.

It has been demonstrated that a complex dense linear algebra operation, such as the QR factorization, can be very efficiently implemented on a modern multi-core processor, such as the Cell BE, through the use of appropriate algorithmic approaches. Specifically, fine granularity of parallelization and loose model of synchronization allow for achieving high performance.

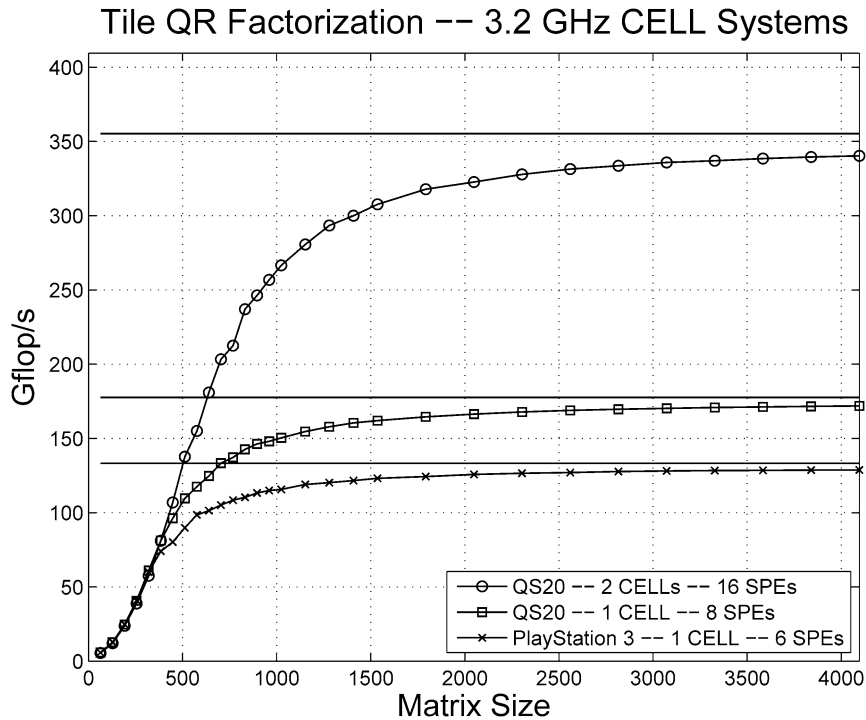


Fig. 6. Performance of the tile QR factorization in single precision on Sony PlayStation 3 (6 SPEs), IBM QS20 blade using one Cell BE (8 SPEs) and IBM QS20 blade using two Cell BEs (16 SPEs). Square matrices were used. The solid horizontal lines mark performance of the SSSRFB kernel times the number of SPEs.

It has been shown that a short-vector SIMD architecture, such as the one of the SPE, can handle complex operations very efficiently, although, at this moment, significant programming effort by an experienced programmer is required.

7. Future work

Experiences with solutions of linear systems of equations using LU and Cholesky factorizations show that the technique of mixed-precision, iterative refinement can be used to achieve double precision accuracy, while exploiting the speed of single precision. It would be straightforward to apply the same approach to solve linear systems of equations or least squares problems using QR factorization. In fact, due to the higher cost, in terms of floating point operations of the QR factorization, the overhead of the iterative process will be much smaller than for the other cases.

Finally, it should be pointed out that LU factorization can be implemented in the same manner, yielding the tile LU algorithm that is bound to produce scaling similar to the QR and Cholesky algorithms, which is

superior to the LU implementation reported so far. Although, it needs to be pointed out that the tile LU algorithm has different properties in terms of numerical stability.

Acknowledgment

The authors thank Alfredo Buttari and Julien Langou for their insightful comments, which helped immensely to improve the quality of this article.

References

- [1] J.J. Dongarra, I.S. Duff, D.C. Sorensen and H.A. van der Vorst, *Numerical Linear Algebra for High-Performance Computers*, SIAM, Philadelphia, PA, 1998; ISBN: 0898714281.
- [2] J.W. Demmel, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, PA, 1997; ISBN: 0898713897.
- [3] G.H. Golub and C.F. van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1996; ISBN: 0801854148.
- [4] L.N. Trefethen and D. Bau, *Numerical Linear Algebra*, SIAM, Philadelphia, PA, 1997; ISBN: 0898713617.

- [5] E. Anderson, Z. Bai, C. Bischof, L.S. Blackford, J.W. Demmel, J.J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney and D. Sorensen, *LAPACK Users' Guide*, SIAM, Philadelphia, PA, 1992; <http://www.netlib.org/lapack/lug/>.
- [6] L.S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J.J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker and R.C. Whaley, *ScaLAPACK Users' Guide*, SIAM, Philadelphia, PA, 1997; <http://www.netlib.org/scalapack/slug/>.
- [7] Basic Linear Algebra Technical Forum, *Basic Linear Algebra Technical Forum Standard*, August 2001; <http://www.netlib.org/blas/blast-forum/blas-report.pdf>.
- [8] T. Chen, R. Raghavan, J.N. Dale and E. Iwata, Cell Broadband Engine architecture and its first implementation – A performance view, *IBM J. Res. & Dev.* **51**(5) (2007), 559–572; DOI: 10.1147/rd.515.0559.
- [9] J. Kurzak and J.J. Dongarra, Implementation of mixed precision in solving systems of linear equations on the CELL processor, *Concur. Comput.: Pract. Exp.* **19**(10) (2007), 1371–1385; DOI: 10.1002/cpe.1164.
- [10] J. Kurzak, A. Buttari and J.J. Dongarra, Solving systems of linear equation on the CELL processor using Cholesky factorization, *Trans. Parallel Distrib. Syst.* **19**(9) (2008), 1175–1186; DOI: TPDS.2007.70813.
- [11] D. Hackenberg, Einsatz und Leistungsanalyse der Cell Broadband Engine, Institut für Technische Informatik, Fakultät Informatik, Technische Universität Dresden, Großer Beleg, February 2007.
- [12] D. Hackenberg, Fast matrix multiplication on CELL systems, http://tu-dresden.de/die_tu_dresden/zentrale_einrichtungen/zih/forschung/architektur_und_leistungsanalyse_von_hochleistungsrechnern/cell/matmul/, July 2007.
- [13] W. Alvaro, J. Kurzak and J.J. Dongarra, Fast and small short vector SIMD matrix multiplication kernels or the synergistic processing element of the CELL processor, in: *Computational Science – ICCS 2008, 8th International Conference*, Kraków, Poland, June 23–25, 2008, Lecture Notes in Computer Science, Vol. 5101, pp. 935–944; DOI: 10.1007/978-3-540-69384-0_98.
- [14] S. Williams, L. Oliker, R. Vuduc, J. Shalf, K. Yelick and J. Demmel, Optimization of sparse matrix–vector multiplication on emerging multicore platforms, in: *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, Reno, NV, USA, November 10–16, 2007, ACM, New York, NY; DOI: 10.1145/1362622.1362674.
- [15] J.J. Dongarra, L. Kaufman and S. Hammarling, Squeezing the most out of eigenvalue solvers on high-performance computers, *Lin. Alg. Applic.* **77** (1986), 113–136.
- [16] C. Bischof and C. van Loan, The WY representation for products of householder matrices, *J. Sci. Stat. Comput.* **8** (1987), 2–13.
- [17] R. Schreiber and C. van Loan, A storage-efficient WY representation for products of householder transformations, *J. Sci. Stat. Comput.* **10** (1991), 53–57.
- [18] E. Elmroth and F.G. Gustavson, Applying recursion to serial and parallel QR factorization leads to better performance, *IBM J. Res. & Dev.* **44**(4) (2000), 605–624.
- [19] E. Elmroth and F.G. Gustavson, New serial and parallel recursive QR factorization algorithms for SMP systems, in: *Applied Parallel Computing, Large Scale Scientific and Industrial Problems, 4th International Workshop, PARA'98*, Umeå, Sweden, June 14–17, 1998, Lecture Notes in Computer Science, Vol. 1541, pp. 120–128; DOI: 10.1007/BFb0095328.
- [20] E. Elmroth and F.G. Gustavson, High-performance library software for QR factorization, in: *Applied Parallel Computing, New Paradigms for HPC in Industry and Academia, 5th International Workshop, PARA 2000*, Bergen, Norway, June 18–20, 2000, Lecture Notes in Computer Science, Vol. 1947, pp. 53–63; DOI: 10.1007/3-540-70734-4_9.
- [21] P.E. Gill, G.H. Golub, W.A. Murray and M.A. Saunders, Methods for modifying matrix factorizations, *Math. Comput.* **28**(126) (1974), 505–535.
- [22] M.W. Berry, J.J. Dongarra and Y. Kim, LAPACK working note 68: A highly parallel algorithm for the reduction of a nonsymmetric matrix to block upper-Hessenberg form. Technical Report UT-CS-94-221, Computer Science Department, University of Tennessee, 1994; <http://www.netlib.org/lapack/lawns/pdf/lawn68.pdf>.
- [23] F.G. Gustavson, Recursion leads to automatic variable blocking for dense linear-algebra algorithms, *IBM J. Res. & Dev.* **41**(6) (1997), 737–756; DOI: 10.1147/rd.416.0737.
- [24] F.G. Gustavson, New generalized matrix data structures lead to a variety of high-performance algorithms, in: *Proceedings of the IFIP TC2/WG2.5 Working Conference on the Architecture of Scientific Software*, Ottawa, Canada, October 2–4, 2000, Kluwer Academic Publishers, Norwell, MA, 2000, pp. 211–234; ISBN: 0792373391.
- [25] F.G. Gustavson, J.A. Gunnels and J.C. Sexton, Minimal data copy for dense linear algebra factorization, in: *Applied Parallel Computing, State of the Art in Scientific Computing, 8th International Workshop, PARA 2006*, Umeå, Sweden, June 18–21, 2006, Lecture Notes in Computer Science, Vol. 4699, pp. 540–549; DOI: 10.1007/978-3-540-75755-9_66.
- [26] E. Elmroth, F.G. Gustavson, I. Jonsson and B. Kågström, Recursive blocked algorithms and hybrid data structures for dense matrix library software, *SIAM Rev.* **46**(1) (2004), 3–45; DOI: 10.1137/S0036144503428693.
- [27] B.C. Gunter and R.A. van de Geijn, Parallel out-of-core computation and updating the QR factorization, *ACM Trans. Math. Software* **31**(1) (2005), 60–78; DOI: 10.1145/1055531.1055534.
- [28] A. Buttari, J. Langou, J. Kurzak and J.J. Dongarra, Parallel tiled QR factorization for multicore architectures, *Concur. Comput.: Pract. Exp.* **20**(13) (2008), 1573–1590; DOI: 10.1002/cpe.1301.
- [29] A. Buttari, J. Langou, J. Kurzak and J.J. Dongarra, LAPACK working note 191: A class of parallel tiled linear algebra algorithms for multicore architectures, Technical Report UT-CS-07-600, Electrical Engineering and Computer Science Department, University of Tennessee, 2007; <http://www.netlib.org/lapack/lawns/pdf/lawn191.pdf>.
- [30] J.M. Perez, P. Bellens, R.M. Badia and J. Labarta, CellSs: Making it easier to program the Cell Broadband Engine processor, *IBM J. Res. & Dev.* **51**(5) (2007), 593–604; DOI: 10.1147/rd.515.0593.
- [31] P. Bellens, J.M. Perez, R.M. Badia and J. Labarta, CellSs: A programming model for the Cell BE architecture, in: *Proceedings of the 2006 ACM/IEEE Conference on Supercomput-*

- ing, Tampa, FL, USA, November 11–17, 2006, ACM; DOI: 10.1145/1188455.1188546.
- [32] N. Park, B. Hong and V.K. Prasanna, Tiling, block data layout, and memory hierarchy performance, *IEEE Trans. Parallel Distrib. Syst.* **14**(7) (2003), 640–654; DOI: 10.1109/TPDS.2003.1214317.
- [33] N. Park, B. Hong and V.K. Prasanna, Analysis of memory hierarchy performance of block data layout, in: *International Conference on Parallel Processing*, Vancouver, Canada, August 2002; DOI: 10.1109/ICPP.2002.1040857.

Copyright of Scientific Programming is the property of IOS Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.