

HPC programming on Intel Many-Integrated-Core hardware with MAGMA Port to Xeon Phi

Jack Dongarra^{a,b,c,*}, Mark Gates^a, Azzam Haidar^a, Yulu Jia^a, Khairul Kabir^a, Piotr Luszczyk^a and Stanimire Tomov^a

^a *University of Tennessee, Knoxville, TN, USA*

^b *Oak Ridge National Laboratory, Oak Ridge, TN, USA*

^c *University of Manchester, Manchester, UK*

Abstract. This paper presents the design and implementation of several fundamental dense linear algebra (DLA) algorithms for multicore with Intel Xeon Phi Coprocessors. In particular, we consider algorithms for solving linear systems. Further, we give an overview of the MAGMA MIC library, an open source, high performance library that incorporates the developments presented here, and, more broadly, provides the DLA functionality equivalent to that of the popular LAPACK library while targeting heterogeneous architectures that feature a mix of multicore CPUs and coprocessors. The LAPACK-compliance simplifies the use of the MAGMA MIC library in applications, while providing them with portably performant DLA. High performance is obtained through the use of the high-performance BLAS, hardware-specific tuning, and a hybridization methodology whereby we split the algorithm into computational tasks of various granularities. Execution of those tasks is properly scheduled over the heterogeneous hardware by minimizing data movements and mapping algorithmic requirements to the architectural strengths of the various heterogeneous hardware components. Our methodology and programming techniques are incorporated into the MAGMA MIC API, which abstracts the application developer from the specifics of the Xeon Phi architecture and is therefore applicable to algorithms beyond the scope of DLA.

Keywords: Numerical linear algebra, Intel Xeon Phi processor, Many Integrated Cores, hardware accelerators and coprocessors, dynamic runtime scheduling using dataflow dependences, communication and computation overlap

1. Introduction and background

Solving linear systems of equations and eigenvalue problems is fundamental to scientific computing. The popular LAPACK library [4], and in particular its vendor optimized implementations such as Intel's MKL [15] or AMD's ACML [2], have been the software of choice to provide solver routines for dense matrices on shared memory systems. This paper considers a redesign of the LAPACK algorithms and their implementation to add efficient support for heterogeneous systems of multicore processors with Intel Xeon Phi coprocessors. This is not the first time that DLA libraries have needed a redesign to be efficient on new architectures. Notable examples being the transition from LINPACK [10] to LAPACK [4] in the 1980s to make al-

gorithms cache-friendly. Also, ScaLAPACK [8] in the 1990s added support for distributed memory systems. And at present time, the PLASMA and MAGMA libraries [1] target efficiency on, respectively, multicore and heterogeneous architectures.

The Intel Xeon Phi coprocessor is a hardware accelerator that made its debut in the late 2012 as a platform for high-throughput technical computing. It is sometimes known under an alternative name of Many Integrated Cores (MIC). For the purposes of this paper, the common mode of operation for the device is called off-load. However, the stand-alone and reverse off-load modes are also valid possibilities. When in off-load mode, the device receives work from the host processor and reports back as soon as the computational task completes. Any such assignment of work proceeds and completes without the host device being involved. In a typical scenario, the host is an Intel x86 CPU such as Sandy Bridge, Ivy Bridge or even more recent Haswell

*Corresponding author: Jack Dongarra, University of Tennessee, Knoxville, TN, USA. E-mail: dongarra@cs.utk.edu.

1 and Ivy Town. The CPU may monitor the activity of
2 communication and/or computation through an event-
3 based interface and can also pursue its own compu-
4 tational activities between events. This is very simi-
5 lar to the operation of hardware accelerators based on
6 throughput-oriented GPUs and compute-capable FP-
7 GAs that are specialized for certain types of workloads
8 beyond what could be achieved on standard multicore
9 CPUs. In fact, Xeon Phi is often considered to be an al-
10 ternative to the hardware accelerators from AMD and
11 NVIDIA despite the fact that there exist many technical
12 differences between the three.

13 The development of new high-performance numer-
14 ical libraries is a complex endeavor, which requires
15 meticulous accounting for the extreme levels of par-
16 allelism, heterogeneity, and wide variety of accelera-
17 tors and coprocessors available in the current architec-
18 tures. Challenges vary from new algorithmic designs to
19 choices of programming models, languages and frame-
20 works that ease the development, future maintenance
21 and portability. This paper addresses these issues while
22 presenting our approach and algorithmic designs in the
23 development of the MAGMA MIC [23] library. Spec-
24 ific differences between the GPU-based MAGMA [1]
25 and the MIC version are elaborated upon in Section 3.

26 To provide a uniform portability across a variety of
27 coprocessors/accelerators, we developed an API that
28 abstracts the application developer from the low level
29 specifics of the architecture. In particular, we use low
30 level vendor libraries, like SCIF for Intel Xeon Phi (see
31 Section 5), to define API for memory management and
32 off-loading computations to coprocessors and/or accel-
33 erators.

34 To deal with the extreme level of parallelism and
35 heterogeneity in the current architectures, MAGMA
36 MIC uses a hybridization methodology, described in
37 Section 6, where we split the algorithms of interest into
38 computational tasks of various granularities, and prop-
39 erly schedule those tasks' execution over the hetero-
40 geneous hardware. Thus, we use a Directed Acyclic
41 Graph (DAG) approach to parallelism and schedul-
42 ing that has been developed and successfully used for
43 dense linear algebra libraries such as PLASMA and
44 MAGMA [1], as well as in general task-based ap-
45 proaches to parallelism, such as runtime systems like
46 StarPU [5] and SMPs [6].

47 Obtaining high performance depends on a combi-
48 nation of algorithmic and hardware-specific optimiza-
49 tions, discussed in Section 6.4. This is in addition to the
50 use of high-performance low-level libraries, which we
51 address in Section 5. This has implications on the re-

52 sulting software: in order to maintain the performance
53 portability across hardware, it is necessary to provide
54 in the library a number of algorithmic variations that
55 are tunable, e.g., at installation time. This is the basic
56 premise of autotuning – a prominent example of these
57 kinds of advanced optimization techniques.

58 A performance study is presented in Section 7. Be-
59 sides verifying our approach and confirming the ap-
60 peal of the Intel Xeon Phi coprocessors for high-
61 performance DLA, the results open up a number of
62 future work opportunities discussed in Section 8 that
63 concludes the paper.

2. Related work

64
65
66
67
68 Intel Xeon Phi [14,16] is a family of Intel copro-
69 cessors known before under the MIC (Many Inte-
70 grated Cores) moniker. Knights Corner (KNC) is the
71 first official product accelerator card in a series that
72 will be followed by Knights Landing (KNL). Phi is a
73 hardware platform based on x86 instruction set with
74 modifications for throughput-oriented workloads. In
75 some sense, Phi may be regarded as an alternative
76 to NVIDIA's compute GPU cards that require CUDA
77 programming [19] or AMD's compute GPU cards that
78 are programmed with OpenCL [17] and the AMD's
79 GPU libraries [3].

80 Phi's use for scientific applications that require solu-
81 tion to PDEs (Partial Differential Equations) was stud-
82 ied and under some scenarios revealed opportunities
83 and advantages [28,29].

84 There is a rich area of work on execution environ-
85 ments that begin with serial code and result in paral-
86 lel execution, often using task superscalar techniques,
87 for example Jade [22], Cilk [9], Sequoia [12], OmpSS
88 [20], Habanero [7], StarPU [5] or the DepSpawn [13]
89 project.

3. Differences between GPU and MIC versions of MAGMA

90
91
92
93
94
95 We mostly focus on the CUDA-based version of
96 MAGMA for the comparison because it is the basis for
97 functional interface and, in terms of the feature set, it is
98 our aim to reproduce it on the Intel MIC coprocessor.

99 Fundamentally, hardware accelerators require refac-
100 toring of the existing code base to accommodate the
101 new compute device and include it harmoniously into
102 the mix with the CPU so that the performance gains

1 may be fully realized. In terms of raw performance
 2 across a broad spectrum of applications, the most ef-
 3 ficient programming language is CUDA [19]. Our ex-
 4 periments show that it easily outperforms portable
 5 standards-based APIs such as OpenCL [11]. While it
 6 might be tempting to include CUDA in the family of
 7 languages derived from C and C++, it is worth not-
 8 ing that the clear syntactic differences from the base
 9 language (mostly C++ and its 1998 standard) form an
 10 easily distinguishable delineation of the computational
 11 spaces of the CPU and the GPU. At the CPU code
 12 level the triple-chevron launch notation, for exam-
 13 ple: `<<<blocks, threadsPerBlock >>>gpu_kernel(args)`,
 14 launches GPU kernels by means of incompatible syn-
 15 tax that requires NVIDIA’s own `nvcc` compiler. This
 16 divergent syntax has spurred over the years a num-
 17 ber of ways to simplify the coding with the use of
 18 directive-based code and as of lately, these efforts
 19 have coalesced into the OpenACC initiative [21,25] –
 20 directive-based approach that hides some of the CUDA
 21 complexity behind compiler’s pragma syntax.

22 The directive-based approach is what Intel MIC fea-
 23 tured from the beginning and this is what MAGMA’s
 24 port to the coprocessor used. However, the MIC port
 25 of MAGMA accommodated changes in the interfaces,
 26 feature set and performance levels. Thus, the end user
 27 was shielded from the effects of the growth of the plat-
 28 form and the flux of the software ecosystem. A particu-
 29 lar example of such an underlying change was the early
 30 use of SCIF (see Section 5) which was essential for ex-
 31 changing non-contiguous memory regions between the
 32 host and the device with a very low overhead. This has
 33 been progressively phased as the need for SCIF dimin-
 34 ished with richer functionality available through the di-
 35 rectives and improvements in the Linux kernel drivers
 36 and runtime overheads. From the user perspective, this
 37 change was transparent for programming on Xeon Phi
 38 while the recent changes in event-driven APIs of CU-
 39 DAs had to percolated to MAGMA’s publicly visible
 40 interface.

41 Another departure from the CUDA-based MAGMA
 42 was the device- and software-specific tuning and opti-
 43 mization (described in more detail in Section 6). There

52 is very little commonality between the targeted sys-
 53 tems, both in terms of hardware and software. The
 54 Xeon Phi implementation has to balance the perfor-
 55 mance sensitivity of the BLAS calls in MKL, custom
 56 kernels, and their mapping onto the much different
 57 hardware substrate. Similarly, the levels and layer-
 58 ing of parallelism nesting (software threads, hardware
 59 threads, versus BLAS threads) is anything but what
 60 is presented to the CUDA programmer. Despite the
 61 differences, however, MAGMA’s external interface re-
 62 mains almost indistinguishable.

4. Compiler support for off-load

66 In this paper, we consider the off-load mode as the
 67 primary mode of operation for the Xeon Phi coproces-
 68 sor. The device receives work from the host proces-
 69 sor and reports back upon completion of the assign-
 70 ment without the host being involved in between these
 71 two events. This is very similar to the operation of net-
 72 work off-load engines, specifically, the TCP Off-load
 73 Engines (TOEs) that feature an optimized implemen-
 74 tation of the TCP stack that handles the majority of the
 75 network traffic to lessen the burden of the main pro-
 76 cessor, which handles other operating system and user
 77 application tasks.

78 The off-load mode for the Xeon Phi devices has
 79 direct support from the compiler in that it is possi-
 80 ble to issue requests to the device and ascertain the
 81 completion of tasks directly from the user’s C/C++
 82 code. The support for this mode of operation is offered
 83 by the Intel compiler through Phi-specific pragma di-
 84 rectives: `offload`, `offload_attribute`, `offload_transfer` and
 85 `offload_wait` [14]. This is very closely related to the off-
 86 load directives now included in the OpenMP 4 stan-
 87 dard. In fact, the two are syntactically and semantically
 88 equivalent, barring the difference in the “omp” prefix
 89 for the OpenMP syntax. A similar standard for GPUs is
 90 called OpenACC. A summary of various programming
 91 methods on Xeon Phi is provided in Table 1. From our
 92 rudimentary experiments we concluded that the com-
 93 piler directive overhead is very close that of the Com-
 94 mon Offload Interface (COI) library.

Table 1

Programming models for the Intel Xeon Phi coprocessors and their current status and properties				
Programming model/API	Status	Portability	Overhead	Language support
SCIF	Mature	No	None	No
COI	Mature	Yes	Minimal	Yes
OpenMP 4.0	Early	Yes	Varies	Yes
OpenCL	Experimental	Yes	Minimal	No

5. Programming model: Host-device with a server based on LLAPI

For many scientific applications, the off-load model offered by the Intel compiler, described in Section 4, is sufficient. This is not the case for a fully equivalent port of MAGMA to the Xeon Phi because of the very rich functionality that MIC MAGMA inherits from both its CUDA and OpenCL ports. We had to use the LLAPI (Low-Level API) based on Symmetric Communication InterFace (SCIF) that offers, as the name suggests, a very low level interface to the host and device hardware. The use of this API is discouraged for most workloads as it tends to be error-prone and offers very little abstraction on top of the hardware interfaces. What motivated us to use it for the port of our library was: (1) the asynchronous events capability that allows low-latency messaging between the host and the device to notify about completion of kernels on Xeon Phi; (2) the possibility of hiding the cost of data transfer between the host and the device which requires the transfer of submatrices to overlap with the computation. The direct access to the DMA (Direct Memory Access) engine allowed us to maximize the bandwidth of data transfers over the PCI Express bus. The only requirement was that the memory regions for transfer be page-aligned and pinned to guarantee their fixed location in the physical memory. Figure 1(a) shows the interaction between the host and the server running on the Xeon Phi and responding to requests that are remote invocations of numerical kernels on data that have already been transferred to the device.




Fig. 1. (a) MIC MAGMA programming model with a LLAPI server mediating requests between the host CPU and the Xeon Phi device. (b) DLA algorithm as a collection of BLAS-based tasks and their dependencies. The algorithm’s critical path is, in general, scheduled on the CPUs, and large data-parallel tasks on the Xeon Phi. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140404>.)

6. Hybridization methodology and optimization strategies

The hybridization methodology used in CUDA MAGMA [26], adopted for MIC MAGMA, is an extension of the task-based approach for parallelism and developing DLA on homogeneous multicore systems [1]. In particular,

- the computation is split into BLAS-based tasks of various granularities, with their data dependencies, as shown in Fig. 1(b);
- small, non-parallelizable tasks with significant control-flow are scheduled on the CPUs;
- large, parallelizable tasks are scheduled on Xeon Phi.

The difference with multicore algorithms is the task splitting, which here is of various granularities to make different tasks suitable for particular architectures, and the scheduling itself. Specific algorithms using this methodology, and covering the main classes of DLA, are described in the subsections below.

6.1. Design and functionality

The MIC MAGMA interface is similar to LAPACK. For example, compare LAPACK’s LU factorization interface vs. MIC MAGMA’s:

```
lapackf77_dgetrf(&M, &N, hA,
&llda, ipiv, &info)
magma_dgetrf_mic(M, N, dA, 0,
ldda, ipiv, &info, queue)
```

Here, hA is the typical CPU pointer (double *) to the matrix of interest in the CPU memory and dA is a pointer in the Xeon Phi memory (its type is magmaDouble_ptr). The last argument in every MIC MAGMA call is an Xeon Phi queue, through which the computation will be streamed on the Xeon Phi device (its type is magma_queue_t).

To abstract the user away from knowing the low-level directives, library functions (such as BLAS), CPU-Phi data transfers, and memory allocations and deallocations are redefined in terms of MIC MAGMA data types and functions. This design allows us to more easily port the MIC MAGMA library to many devices as was the case for the GPU accelerators that either use CUDA [19] or OpenCL [11,17] and eventually to merge them in order to maintain a single source code tree with conditional compilation options that allow seamless targeting of specific hardware. Also, the MIC MAGMA wrappers provide a complete set of func-

52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102

tions for programming hybrid high-performance numerical libraries. Thus, not only users but application developers as well can opt to use the MIC MAGMA wrappers. MIC MAGMA provides the four standard floating-point arithmetic precisions – single and double precision real as well as single and double precision complex. It has routines for the so called one-sided factorizations (LU, QR and Cholesky), and recently we are developing the two-sided factorizations (Hessenberg, bi- and tridiagonal reductions), linear system and least squares solvers, matrix inversions, symmetric and non-symmetric standard eigenvalue problems, SVD and orthogonal transformation routines.

6.2. Task distribution based on hardware capability

Programming models that raise the level of abstraction are of great importance for reducing software development efforts. A traditional approach has been to organize algorithms in terms of BLAS calls, where hardware specific optimizations would be hidden in BLAS implementations such as Intel’s MKL or AMD’s ACML. This is still valid and used but has shown some drawbacks on new architectures. In particular, parallelization is achieved using a fork–join approach since each BLAS call, e.g., a matrix–matrix multiplication, can be performed in parallel (fork) but a synchronization is needed before performing the next call (join). The number of synchronizations thus can become a prohibitive bottlenecks for performance on highly parallel devices such as the MICs. This type of programming has been popularized under the Bulk Synchronous Processing name [27].

Instead, the algorithms (like matrix factorizations) are broken into computational tasks (e.g., panel factorizations followed by trailing submatrix updates) and pipelined for execution on the available hardware components (see below). Moreover, particular tasks are scheduled for execution on the hardware components that are best suited for them. Thus, this task distribution based on *hardware capability* allows the user for the efficient use of each hardware component. In the case of DLA factorizations, the less parallel panel tasks are scheduled for execution on multicore CPUs, and the parallel updates mainly on the MICs. We illustrate this in Algorithm 1.

6.3. LU, QR and Cholesky factorizations for Intel Xeon Phi

The one-sided factorization routines implemented and currently available through MIC MAGMA are:

Algorithm 1. Two-phase (first: panel, two: update) factorization of $A = [P_1, P_2, \dots]$ with lookahead of depth 1. Matrix A and the result are assumed to reside on the MIC memory

```

1 PanelStartReceivingon CPU( $P_1$ );
2 for  $P_i = P_1, P_2, \dots$  do
3   PanelFactorizeon CPU( $P_i$ );
4   PanelSendto MIC( $P_i$ );
5   TrailingMatrixUpdateon MIC( $P_{i+1}$ );
6   PanelStartReceivingon CPU( $P_{i+1}$ );
7   TrailingMatrixUpdateon MIC( $P_{i+2}, \dots$ );

```




Fig. 2. Typical computational pattern for the hybrid one-sided factorizations in MIC MAGMA. (The colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140404>.)

magma_zgetrf_mic computes an LU factorization of a general M -by- N matrix A using partial pivoting with row interchanges;

magma_zgeqrf_mic computes a QR factorization of a general M -by- N matrix A ;

magma_zpotrf_mic computes the Cholesky factorization of a complex Hermitian positive definite matrix A .

Routines in all standard four floating-point precision arithmetics are available, following LAPACK’s naming convention. Namely, the first letter of the routine name (after the prefix **magma_**) indicates the precision – **z**, **c**, **d** or **s** for double complex, single complex, double real or single real, respectively. The suffix **_mic** indicates that the input and the output matrices are in the Xeon Phi memory.

The typical hybrid computation and communication pattern for the one-sided factorizations (LU, QR and Cholesky) is shown in Fig. 2. At a given iteration, panel i is copied to the CPU and factored using LAPACK, and the result is copied back to Xeon Phi. The trailing matrix, consisting of the next panel $i + 1$ and the rest of the matrix, is updated on the Xeon Phi. After receiving panel i back from the CPU, panel $i + 1$ is updated first

Algorithm 2. Cholesky factorization in MIC MAGMA

```

1  for  $j = 0, nb, 2nb, 3nb, \dots, n - 1$  do
2
3  1   $jb = \min(nb, n - j)$ ;
4  magma_zherk_mic( MagmaUpper, MagmaConjTrans, jb, j, m_one,
5  dA(0,j), ldda, one, dA(j,j), ldda, queue );
6  magma_zgetmatrix_async_mic(jb, jb, dA(j,j), ldda, work, 0, jb,
7  queue, &event );
8  if  $j + jb < n$  then
9  magma_zgemm_mic( MagmaConjTrans, MagmaNoTrans, jb,
10 n - j - jb, j, mz_one, dA(0,j), ldd, dA(0, j + jb), ldda, queue );
11 magma_event_sync_mic(event );
12 lapackf77_zpotrf( MagmaUpperStr, &jb, work, &jb, info );
13 if  $*info \neq 0$  then
14 *info += j;
15 magma_zsetmatrix_async_mic(jb, jb, work, 0, jb, dA(j,j), ldda,
16 queue, &event );
17 if  $j + jb < n$  then
18 magma_event_sync_mic(event );
19 magma_ztrsm_mic(MagmaLeft, MagmaUpper,
20 MagmaConjTrans, MagmaNonUnit, jb, n - j - jb, z_one, dA(j,j),
21 ldda, dA(j, j + jb), ldda, queue );

```

using panel i and the result is sent to the CPU (as being the next panel to be factored there). While the CPU starts the factorization of i , the rest of trailing matrix, panels $i + 1, i + 2, \dots$, is updated on the Xeon Phi device in parallel with the CPU factorization of panel $i + 1$. In this pattern, only data to the right of the current panel is accessed and modified, and the factorizations that use it are known as right-looking. The computation can be organized differently – to access and modify data only to the left of the panel – in which case the factorizations are known as left-looking.

An example of a left-looking factorization, demonstrating a hybrid implementation, is given in Algorithm 2 for the Cholesky factorization. The algorithm introduces a notion of a *blocking factor* denoted as nb , which is the algorithm-level entity that defines the number of columns in the panel and the inner dimension of the outer-product update to the trailing submatrix. Copying the panel to the CPU, in this case just a square block on the diagonal, is done on line 4. The data transfer is asynchronous, so before we factor it on the CPU (line 8), we synchronize on line 7 to enforce that the data has arrived. Note that the CPU work from line 8 is overlapped with the Xeon Phi work on line 6. This is indeed the case because line 6 is an asynchronous call/request from the CPU to Xeon Phi to start a ZGEMM operation. Thus, the control is passed to lines 7 and 8 while Xeon Phi is performing the ZGEMM. The resulting factored panel from the CPU work is sent to Xeon Phi on line 11 and used on line 14, after making sure that it has arrived through the sync command on line 13.

6.4. Hybrid implementation and optimization techniques

In order to explain our hybrid methodology and the optimization that we have developed, let us give a detailed analysis for the QR decomposition algorithm. While the description below only addresses the QR factorization, it is straightforward to derive with the same ideas the analysis for both the Cholesky and LU factorizations. For that we start briefly by recalling the description of the QR algorithm.

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 $m \times m$ and R is an upper trapezoidal matrix of size $m \times n$. The QR algorithm can be described as a sequence of steps where, at each step, a QR of a panel is performed 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 in what is called the “*trailing matrix update*”. Despite that this approach can exploit the parallelism of the Level 3 BLAS during the trailing matrix update, it has a number of limitations when implemented on massively multithreaded system such as the Intel Xeon Phi coprocessor due to the nature of its operations. On the one hand, the panel factorization relies on Level 2 BLAS operations that cannot be efficiently parallelized on either Xeon Phi or any accelerator such as GPU-based architectures, and thus it can be considered to be close to sequential operations that limit the scalability of the algorithm. On the other hand, this algorithm is referred as the *fork-join approach* since the execution flow will show a sequence of sequential operations (panel factorizations) interleaved with parallel ones (trailing matrix updates). In order to take advantage of the high execution rate of the massively multithreaded system, in particular, the Phi coprocessor we redesigned the standard algorithm in a way to perform the Level 3 BLAS operations (trailing matrix update) on the Xeon Phi while performing the Level 2 BLAS operations (panel factorization) on the CPU. We also proposed an algorithmic change to remove the fork-join bottleneck and to minimize the overhead of the panel factorization by hiding its costs behind the parallel trailing matrix update. This approach can be described as the *scalable look ahead techniques* [24]. Our idea is to split the trailing matrix update into two phases, the update of the lookahead panel (panel of step $i + 1$, i.e., dark blue portion of Fig. 2) and the update of the remaining trailing sub-

matrix (clear blue portion of Fig. 2). Thus, during the submatrix update the CPU can receive asynchronously the panel $i + 1$ and performs its factorization. As a result, our MIC MAGMA implementation of the QR factorization can be described by a sequence of the three phases described below. Consider a matrix A that can be represented as:

$$A = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{bmatrix}. \quad (1)$$

- Phase 1, the panel factorization:** at a step i , this phase consists of a QR transformation of the panel $A_{i:n,i}$ as in Eq. (2). This operation consists of calling two routines. The DGEQR2 that factorizes the panel and produces nb Householder reflectors (V_{*i}) and an upper triangular matrix R_{ii} of size $nb \times nb$, which is a portion of the final R factor, and the DLARFT that generates the triangular matrix T_{ii} of size $nb \times nb$ used for the trailing matrix update. This phase is performed on the CPU,

$$\begin{bmatrix} A_{11} \\ A_{21} \\ A_{31} \end{bmatrix} \Rightarrow \begin{bmatrix} V_{11} \\ V_{21} \\ V_{31} \end{bmatrix}, [R_{1,1}], [T_{1,1}]. \quad (2)$$

- Phase 2, the look ahead panel update:** the transformation that was computed in the panel factorization needs to be applied to the rest of the matrix (trailing matrix, i.e., the blue portion of Fig. 2). This phase consists into updating only the next panel (dark blue portion of Fig. 2) in order to let the CPU start its factorization as soon as possible while the update of the remaining portion of the matrix is performed in phase 3. The idea is to hide the cost of the panel factorization. This operation presented in Eq. (3), is performed on the Phi coprocessor and involves the DLARFB routine which has been redesigned as a sequence of DGEMM's to better take advantage of the Level 3 BLAS operations,

$$\begin{bmatrix} R_{12} \\ \tilde{A}_{22} \\ \tilde{A}_{32} \end{bmatrix} = [I - V_{*i} T_{ii}^T V_{*i}^T] \begin{bmatrix} A_{12} \\ A_{22} \\ A_{32} \end{bmatrix}. \quad (3)$$

- Phase 3, the trailing matrix update:** Similarly to phase 2, this phase consists of applying the Householder reflectors generated during the panel




Fig. 3. Effect of the blocking factor on performance of MAGMA MIC factorizations. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140404>.)

factorization of step i according to Eq. (3), to the remaining portion of the matrix (the trailing submatrix i.e., the clear blue portion of Fig. 2). This operation is also performed on the Phi coprocessor, while in parallel to it, the CPU performs the factorization of the panel $i + 1$ that has been computed in phase 2.

This hybrid technique of distribution of tasks between CPU-Phi allows us to hide the memory bound operations occurred during the panel factorization (phase 1) by performing such operation on the CPU in parallel with the trailing submatrix update (phase 3) on the Phi coprocessor. However, one of the key parameters to performance tuning is the blocking size as the performance and the overlap between the CPU-Phi will be solely guided by it. Figure 3 illustrates the effect of the blocking factor on the performance. It is obvious that, a small nb will reduce the cost of the panel factorization phase 1, but it decreases the efficiency of the Level 3 BLAS kernel of phase 2 and phase 3 and thus resulting a bad performance. As opposed, a large nb will dramatically affect the panel factorization phase 1 which becomes slow and thus the CPU-Phi computation cannot be overlapped, providing a deterioration in the performance as shown in Fig. 3. As a consequence, the challenging problem is the following: on the one hand, the blocking size nb needs to be large enough to extract high performance from the Level 3 BLAS phase 3 and on the other hand, it has to be small enough to extract efficiency (thanks to the cache speed up) from the Level 2 BLAS phase 1 and overlap CPU/Phi computation. Figure 3 shows the performance obtained for different blocking sizes and we

can see a trade-off between small and large nb 's. Either $nb = 480$ or $nb = 960$ can be considered as a good choice because MKL Phi BLAS is optimized for multiples of 240. Moreover, to extract the maximum performance and allow the maximum overlap between both of the CPU and the Xeon Phi coprocessor, we developed a new variant that can use a variable nb during the steps of the algorithm. The flexibility of our implementation allows an efficient task execution overlap between the CPU host and the Phi coprocessor which enables the implementation to scale almost linearly with the number of cores on the Phi coprocessor, which we see (below) from very good performance that is close to the practical peak obtained on such a system from matrix-matrix multiply and related dense linear algebra operations, which achieve over 70% of the theoretical peak performance. Our tuned variable implementation is represented by the red curve of Fig. 3 where we can easily observe its advantages over the other variants.

The Phi-specific techniques had to be employed in order to reap the benefits of the above design in the presence of particular constraints and opportunities present on the Intel hardware. One opportunity is to choose the best one out of a number of interfaces for transferring data between the CPU and the coprocessor – refer to Table 1 for details. The Phi implementation of MAGMA seeks to minimize the latency and maximize the bandwidth of the PCIe transfers while maintaining a good computational load of both the host and the device. If the proper API for the right size of data transfer is chosen, the DMA hardware can take over and off-load the transfer logistics so that the compute components can remain the busy computing on matrix elements and not polluting their cache hierarchy with spurious messaging data. In particular, SCIF offers the lowest latency but the large data transfers create complexity burden of dealing of many smaller transfer requests. Higher level mechanisms, such as COI and virtual shared memory regions, carry a larger overhead but allow the handling of large volumes of data in a much more automated fashion. The switching between these interfaces occurs seamless behind the familiar MIC MAGMA functions.

6.5. Task-based runtime model

The scheduling of tasks for execution can be static or dynamic. In either case, the small and not easy to parallelize tasks from the critical path (e.g., panel factorizations) are executed on CPUs, and the large and

highly parallel task (like the matrix updates) mostly on the MICs.

The use of multiple coprocessors complicates the development using static scheduling. Instead, the use of a light-weight runtime system is preferred as it can keep scheduling overhead low, while enabling the expression of parallelism through sequential-like code. The runtime system relieves the developer from keeping track of the computational activities that, in the case of heterogeneous systems, are further exacerbated by the separation between the address spaces of the main memory of the CPU and the MICs. Our runtime model is build on the QUARK [30] superscalar execution environment that has been originally used with great success for linear algebra software on just multi-core platforms [18]. The conceptual work though could be replicated within other models such as StarPU [5], OmpSS [20], Cilk [9] and Jade [22], to just mention a few.

Dynamic runtime scheduling plays an important role in translating dependences annotated at the source code level and discovered at runtime when the execution traverses the Direct Acyclic Graph of computational tasks. For example, one of the symbolic dependences of tasks in Algorithm 1 could be:

$$\begin{aligned} & \text{PanelFactorize}_{\text{CPU}}(P_i) \\ & \rightarrow \text{TrailingMatrixUpdate}_{\text{MIC}}(P_{i+1}) \end{aligned}$$

At runtime, this dependence formula is repeatedly applied to form a sequence of tasks:

$$\begin{aligned} & \text{PanelFactorize}_{\text{CPU}}(P_1) \\ & \rightarrow \text{TrailingMatrixUpdate}_{\text{MIC}}(P_2) \\ & \text{PanelFactorize}_{\text{CPU}}(P_2) \\ & \rightarrow \text{TrailingMatrixUpdate}_{\text{MIC}}(P_3) \\ & \dots \end{aligned}$$

The runtime environment for scheduling maintains the current set of tasks and the future set of tasks. The completed tasks enable execution of their dependent tasks and are discarded from the system.

6.6. Improving off-load mode communication

It is well known that the off-load transfer mode copies only continuous chunks of data from and to the coprocessors. However most of the scientific application algorithms require to exchange data with 2D or

1 3D storage and thus this may create an issue when using
2 the off-load transfer mode. In particular, the one-
3 sided factorizations (Cholesky, LU and QR) require to
4 send the panel to the CPU and then receive it later af-
5 ter being factorized by the CPU. A simple implemen-
6 tation loops over one direction and calls the off-load
7 section to send and receive a contiguous vector. Such
8 an implementation behaves poorly and as a result the
9 communication will become expensive and will slow
10 down the algorithm. Indeed, another alternative is to
11 copy the 2D panel to a contiguous temporary space on
12 the MIC, and then to send it and vice versa. Hence,
13 there are two points that need to be taken into consid-
14 eration. Firstly, the copy needs to be implemented as a
15 multi-threaded operation in order to hide its cost. For
16 that, we implemented a parallel copy that uses all of
17 the 240 hardware threads of the MIC to perform the
18 copy. This might be against the common wisdom that
19 multi-threading is of little help for bandwidth-limited
20 operations such as a memory copy. This is not the ex-
21 perience on the MIC, where the clock frequency of the
22 compute cores is twice as low as that of the memory –
23 the exact opposite of what is the case in Intel x86 mul-
24 ticore processors. In addition to the low frequency, the
25 current MIC hardware is to a large degree an in-order
26 architecture with dual-pipeline execution and single-
27 issue fetch/decode units [16] which poses constraints
28 on the amount of bandwidth that can be utilized by a
29 single core. These can be overcome in multiple ways,
30 including the use of streaming loads and have the mul-
31 tiple threads requesting data. Secondly, when the MIC
32 copies data to or from the temporary space, it should
33 be the only kernel running, otherwise, it will run sim-
34 ultaneously with another executing kernel and this may
35 slow down both of the kernels. To that end, we repre-
36 sented the copy kernel as a task with high priority and
37 the scheduler is responsible for executing it as soon
38 as possible and to handle the dependencies such as no
39 other kernel will be running at the same time. Xeon Phi
40 requires multiple cores driving a single FPU, which is
41 similar to Hyperthreading in the recent Intel x86 pro-
42 cessors. In fact, the core-to-FPU ratio must be two-to-
43 one to satisfy the data rate that a single FPU can sus-
44 tain. If the ratio is lower, the FPU goes largely under-
45 utilized because the data request rate from memory is
46 too low.

48 Experiments showed that when using these opti-
49 mizations the performance of the off-load communica-
50 tion mode is comparable to both the SCIF and the COI
51 mode with a variance of less than 5%.

6.7. Trading extra computation for higher execution rate

The optimization discussed here is MIC-specific but is often valid for any hardware architecture with multi-layered memory hierarchy. The `dlarfb` routine used by the QR decomposition consists of two `dgemms` and one `dtrmm`. Since coprocessors are better at handling compute-bound tasks, for computational efficiency, we replace the `dtrmm` by `dgemm`, yielding 5–10% performance improvement. For the Cholesky factorization, the trailing matrix update requires a `dsyrk`. Due to uneven storage, the multi-device `dsyrk` cannot be assembled purely from regular `dsyrk` calls on each device. Instead, each block column must be processed individually. The diagonal blocks require special attention. One can use a `dsyrk` to update each diagonal block, and a `dgemm` to update the remainder of each block column below the diagonal block. The small `dsyrk` operations have little parallelism and therefore their execution is inefficient on MICs. This can be improved to some degree by using `pragma` to run several `dsyrk`'s simultaneously. Nevertheless, because we have copied the data to the device, we can consider the space above the diagonal to be a scratch workspace. Thus, we update the entire block column, including the diagonal block, writing extra data into the upper triangle of the diagonal block, which is subsequently ignored. We do extra computation for the diagonal block, but gain efficiency overall by launching fewer BLAS kernels on the device and using the more efficient `dgemm` kernels, instead of small `dsyrk` kernels.

The per-kernel improvement in performance exceeds 20% and for the entire factorization a 5–10% improvement levels may be observed.

7. Performance results

This section presents the performance results obtained by our hybrid CPU–Xeon Phi implementation in the context of the development of the state-of-the-art numerical linear algebra libraries.

7.1. Experimental environment

Our experiments were performed on a system equipped with Intel Xeon Phi formerly known as Knights Corner. It is representative of a vast class of servers and workstations commonly used for computationally intensive workloads. We benchmarked all im-

52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102

1 plementations on an Intel multicore system with dual-
 2 dual-socket, 8 core Intel Xeon E5-2670 (Sandy Bridge) pro-
 3 cessors, each running at 2.6 GHz. Each socket has a
 4 24 MB shared L3 cache, and each core has a private
 5 256 KB L2 and 64 KB L1. The system is equipped with
 6 52 Gbytes of memory. The theoretical peak for this archi-
 7 tecture in double precision is 20.8 Gflop/s per core,
 8 giving 332 Gflops in total. The system is also equipped
 9 with an Intel Xeon Phi cards with 7.7 Gbytes per card
 10 running at 1.09 GHz, and giving a double precision
 11 theoretical peak of 1046 Gflops.

12 There are a number of software packages available.
 13 On the CPU side we used the MKL (Math Kernel Li-
 14 brary) [15] which is a commercial software package
 15 from Intel that is a highly optimized numerical library.
 16 On the Intel Xeon side, we used the MPSS 2.1.5889-
 17 16 as the software stack, icc 13.1.1 20130313 which
 18 comes with the composer_xe_2013.3.163 suite as the
 19 compiler and the Level 3 BLAS routine GEMM from
 20 MKL 11.00.03.

22 7.2. Experimental results

24 Figure 4 reports the performance of the three linear
 25 algebra factorization operations, the Cholesky, QR and
 26 LU factorizations with our hybrid implementation and
 27 compare it to the performance of the CPU implementa-
 28 tion of the MKL libraries. For our implementation, the
 29 blocking factor has been chosen to be flexible in order
 30 to achieve the best performance. A detailed descrip-
 31 tion of how to choose this factor is included in Sec-
 32 tion 6.4 and in the results presented in this section we
 33 choose the factor to be in the range between 480 and
 34 960. As a general rule, we use smaller blocking fac-
 35 tors for smaller matrices and larger ones for the larger
 36 matrices. The graphs show the performance measured
 37 using all the cores available on the system (i.e., 60 for
 38 the Intel Phi and 16 for the CPU) with respect to the
 39 problem size. In order to reflect the time to completion,
 40 for each algorithm the operation count is assumed to be
 41 the same as that of the LAPACK algorithm, i.e., $\frac{1}{3}N^3$,
 42 $\frac{2}{3}N^3$ and $\frac{4}{3}N^3$ for the Cholesky factorization, the LU
 43 factorization and the QR decomposition, respectively.

44 Figure 4(a), (b) and (c) provides the common type
 45 of information that is characteristic of dense linear al-
 46 gebra computations. Clearly, our algorithms from the
 47 MIC MAGMA library, that employ hybrid techniques,
 48 deliver higher execution rates than their CPU coun-
 49 terparts optimized by the vendor. This is in correspon-
 50 dence with the difference of the peak performance rates
 51 between the two hardware components. It should be

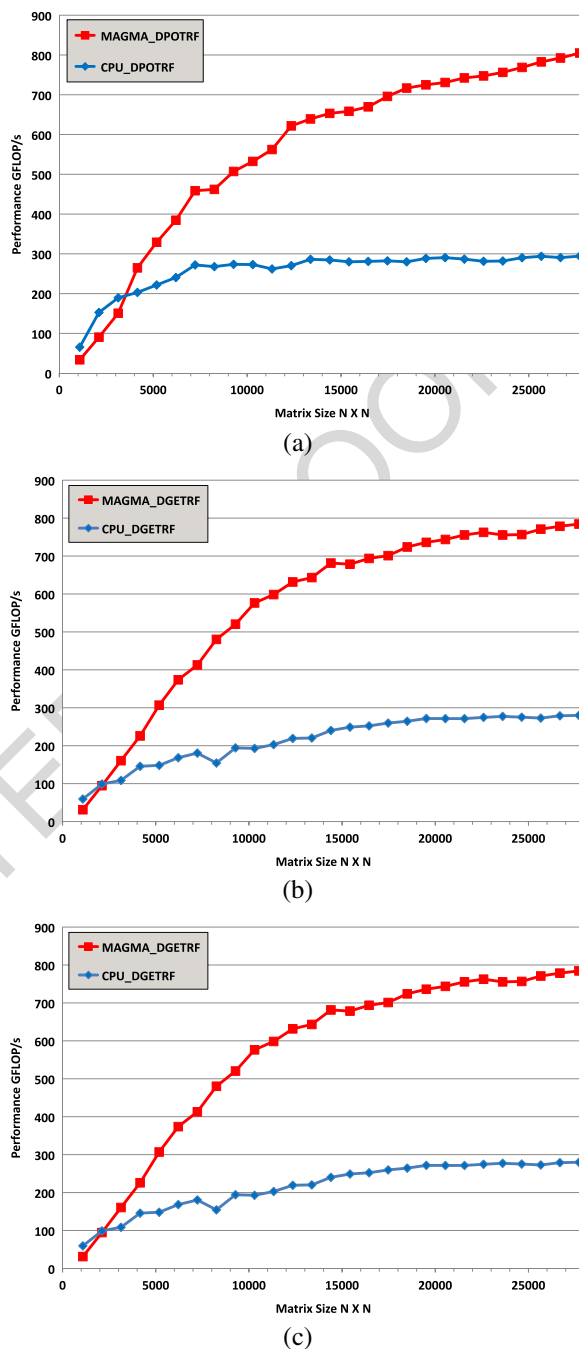


Fig. 4. Comparison of the performance versus the optimized CPU version of the MKL libraries for the three one-sided factorizations. (a) Cholesky factorization (magma_zpotrf_mic). (b) LU factorization (magma_zgetrf_mic). (c) QR factorization (magma_zgeqrf_mic). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140404>.)

obvious from the graphs that the combination of a CPU and a Phi coprocessor with a tuned implementation

provides substantial performance benefits as opposed to a CPU-only implementation. The figures show that the MIC MAGMA hybrid algorithms are capable of completing any of the three factorization algorithms as much as twice as fast as the CPU optimized version for a matrix of size larger than 10,000; and more than three times faster when the matrix size is large enough (larger than 20,000). The actual curves of Fig. 4 illustrates the efficiency of our hybrid techniques where we note that the performance obtained by our implementation, achieves a very close level to the practical peak of the Intel Xeon Phi coprocessor computed by running the GEMM routine (which is around 850 Gflop/s). This gain is mostly obtained by two improvements. First, the nature of the operations involved on the Phi side which are mostly BLAS Level 3 operations were redesigned and implemented as a combination of vendor's DGEMM calls. For more details we denote below the routines executed on the Xeon Phi coprocessor:

- The DSYRK operations for the Cholesky factorization where the DSYRK has been redesigned as a combination of DGEMM's routines;
- The DGEMM for the LU factorization;
- The DLARFB for the QR decomposition where also its has been redesigned as a combination of DGEMM's.

Second, all of the Level 2 BLAS routines that are memory bound and that represent a limit for the performance (i.e., DPOTF2, DGETF2 and DGEQR2 for Cholesky, LU and QR factorization, respectively) are executed on the CPU side while being overlapped with the Phi coprocessor execution as described in Section 6.4.

An important remark has to be made here for the Cholesky factorization: the *left-looking* algorithm as implemented in LAPACK is considered as well optimized for memory reuse but at the price of less parallelism and thus is not suitable for massively multi-core machines. This variant delivers poor performance when compared to the *right-looking* variant that allows more parallelism and thus run at higher speed.

8. Conclusions and future work

In this article, we have shown how to extend our hybridization methodology from existing systems to a new hardware platform. The challenge of the porting effort stemmed from the fact that the new coprocessor from Intel, the Xeon Phi, featured programming

models and relative execution overheads, that were markedly different from what we have been targeting on GPU-based accelerators. Nevertheless, we believe that the techniques used in this paper adequately adapt our hybrid algorithm to best take advantage of the new heterogeneous hardware. We have derived an implementation schema of the dense linear algebra kernels that also can be applied to either the two-sided factorization used for solving the eigenproblem and the SVD or to the sparse linear algebra algorithms. We plan to further study the implementation of multi-Xeon Phi algorithms in a distributed computing environment. We think that the techniques presented will become more popular and will be integrated into dynamic runtime system technologies. The ultimate goal is that this integration will help to tremendously decrease development time while retaining high-performance.

In addition, we see an opportunity in fully automating the tuning process of various algorithmic parameters of our implementation including the blocking factor nb , the number of threads used in various computational kernel, etc. This will become even more important as the number of linear algebra operations included in MIC MAGMA grows.

Acknowledgements

Work was funded in part by the Ministry of Education and Science of the Russian Federation, Agreement N 14.607.21.0006 (unique identifier RFMEFI57714X0020).

The authors would like to thank the National Science Foundation for supporting this work under Grant No. ACI-1339822, the Department of Energy and ISTC for Big Data for supporting this research effort.

References

- [1] E. Agullo, J. Demmel, J. Dongarra, B. Hadri, J. Kurzak, J. Langou, H. Ltaief, P. Luszczek and S. Tomov, Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects, *J. Phys.: Conf. Ser.* **180**(1) (2009).
- [2] AMD, AMD Core Math Library (ACML), available at: <http://developer.amd.com/tools/>.
- [3] AMD, clMath libraries: clBLAS 2.0, 13 August 2013, available at: <https://github.com/clMathLibraries>.
- [4] E. Anderson, Z. Bai, C. Bischof, S.L. Blackford, J.W. Demmel, J.J. Dongarra, J. Du Croz, A. Greenbaum, S.J. Hammarling, A. McKenney and D.C. Sorensen, *LAPACK User's Guide*, 3rd edn, SIAM, Philadelphia, 1999.

- [5] C. Augonnet, S. Thibault, R. Namyst and P.-A. Wacrenier, StarPU: A unified platform for task scheduling on heterogeneous multicore architectures, *Concurrency and Computation: Practice and Experience* **23**(2) (2011), 187–198.
- [6] Barcelona Supercomputing Center, SMP Superscalar (SMPSs) User’s Manual, Version 2.0, 2008, available at: <http://www.bsc.es/media/1002.pdf>.
- [7] R. Barik, Z. Budimlic, V. Cavè, S. Chatterjee, Y. Guo, D. Peixotto, R. Raman, J. Shirako, S. Taşlılar, Y. Yan, Y. Zhao and V. Sarkar, The habanero multicore software research project, in: *Proceedings of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications, OOPSLA’09*, ACM, New York, NY, USA, 2009, pp. 735–736.
- [8] S.L. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I.S. 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, available at: <http://www.netlib.org/scalapack/slug/>.
- [9] R.D. Blumofe, C.F. Joerg, B.C. Kuszmaul, C.E. Leiserson, K.H. Randall and Y. Zhou, Cilk: An efficient multithreaded runtime system, *ACM SIGPLAN Notices* **30** (1995), 207–216.
- [10] J. Dongarra, J. Bunch, C. Moler and G.W. Stewart, *LINPACK Users’ Guide*, SIAM, Philadelphia, PA, 1979.
- [11] P. Du, R. Weber, P. Luszczek, S. Tomov, G. Peterson and J. Dongarra, From CUDA to OpenCL: Towards a performance-portable solution for multi-platform GPU programming, *Parallel Computing* **38**(8) (2012), 391–407.
- [12] K. Fatahalian, D.R. Horn, T.J. Knight, L. Leem, M. Houston, J.Y. Park, M. Erez, M. Ren, A. Aiken, W.J. Dally and P. Hanrahan, Sequoia: Programming the memory hierarchy, in: *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC’06*, ACM, New York, NY, USA, 2006.
- [13] C.H. González and B.B. Fraguola, A framework for argument-based task synchronization with automatic detection of dependencies, *Parallel Computing* **39**(9) (2013), 475–489.
- [14] Intel, Intel® Xeon Phi™ coprocessor system software developers guide, available at: <http://software.intel.com/en-us/articles/>.
- [15] Intel, Math Kernel Library, available at: <http://software.intel.com/en-us/articles/intel-mkl/>.
- [16] J. Jeffers and J. Reinders, *Intel® Xeon Phi™ Coprocessor High-Performance Programming*, Morgan Kaufmann, 2013.
- [17] Khronos OpenCL Working Group, The OpenCL specification, Version 1.0, document revision: 48, 2009.
- [18] J. Kurzak, P. Luszczek, A. YarKhan, M. Faverge, J. Langou, H. Bouwmeester and J. Dongarra, Multithreading in the PLASMA library, in: *Handbook of Multi and Many-Core Processing: Architecture, Algorithms, Programming, and Applications*, Computer and Information Science Series, Vol. 26, Chapman & Hall/CRC, 2013.
- [19] NVIDIA Corporation, NVIDIA CUDA C Programming Guide, 13 February 2014, retrieved 1 May 2014, available at: <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>.
- [20] J.M. Pérez, R.M. Badia and J. Labarta, A dependency-aware task-based programming environment for multi-core architectures, in: *Proceedings of the 2008 IEEE International Conference on Cluster Computing*, Tsukuba, Japan, 29 September–1 October 2008, IEEE, 2008, pp. 142–151.
- [21] Proposed additions for OpenACC 2.0, OpenACC™ application programming interface, November 2012.
- [22] M.C. Rinard, D.J. Scales and M.S. Lam, Jade: A high-level, machine-independent language for parallel programming, *Computer* **26**(6) (1993), 28–38.
- [23] Software distribution of MAGMA MIC, Version 1.0, 3 May 2013, available at: <http://icl.cs.utk.edu/magma/software/>.
- [24] P.E. Strazdins, Lookahead and algorithmic blocking techniques compared for parallel matrix factorization, in: *10th International Conference on Parallel and Distributed Computing and Systems, IASTED*, Las Vegas, USA, 1998.
- [25] The OpenACC™ application programming interface, Version 1.0, November 2011.
- [26] S. Tomov and J. Dongarra, Dense linear algebra for hybrid GPU-based systems, in: *Scientific Computing with Multicore and Accelerators*, J. Kurzak, D.A. Bader and J. Dongarra, eds, Chapman & Hall/CRC, 2010.
- [27] L.G. Valiant, Bulk-synchronous parallel computers, in: *Parallel Processing and Artificial Intelligence*, M. Reeve, ed., Wiley, 1989, pp. 15–22.
- [28] S. Williams, D.D. Kalamkar, A. Singh, A.M. Deshpande, B. Van Straalen, M. Smelyanskiy, A. Almgren, P. Dubey, J. Shalf and L. Oliker, Optimization of geometric multigrid for emerging multi- and manycore processors, in: *SC’12, Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, IEEE Computer Society Press, Los Alamitos, CA, USA, 2012, pp. 96:1–96:11, available at: <http://dl.acm.org/citation.cfm?id=2388996.2389126>.
- [29] M.M. Wolf, M.A. Heroux and E.G. Boman, Factors impacting performance of multithreaded sparse triangular solve, in: *VECPAR’10, Proceedings of the 9th International Conference on High Performance Computing for Computational Science*, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 32–44, available at: <http://dl.acm.org/citation.cfm?id=1964238.1964246>.
- [30] A. YarKhan, Dynamic task execution on shared and distributed memory architectures, PhD thesis, University of Tennessee, 2012.