



HAL
open science

Asynchronous Task-Based Polar Decomposition on Single Node Manycore Architectures

Dalal Sukkari, Hatem Ltaief, Mathieu Faverge, David Keyes

► **To cite this version:**

Dalal Sukkari, Hatem Ltaief, Mathieu Faverge, David Keyes. Asynchronous Task-Based Polar Decomposition on Single Node Manycore Architectures. *IEEE Transactions on Parallel and Distributed Systems*, 2017, XX, 10.1109/TPDS.2017.2755655 . hal-01585079

HAL Id: hal-01585079

<https://inria.hal.science/hal-01585079v1>

Submitted on 11 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Asynchronous Task-Based Polar Decomposition on Single Node Manycore Architectures

Dalal Sukkari, Hatem Ltaief, Mathieu Faverge, and David Keyes

Abstract—This paper introduces the first asynchronous, task-based formulation of the polar decomposition and its corresponding implementation on manycore architectures. Based on a new formulation of the iterative QR dynamically-weighted Halley algorithm (QDWH) for the calculation of the polar decomposition, the proposed implementation replaces the original and hostile LU factorization for the condition number estimator by the more adequate QR factorization to enable software portability across various architectures. Relying on fine-grained computations, the novel task-based implementation is also capable of taking advantage of the identity structure of the matrix involved during the QDWH iterations, which decreases the overall algorithmic complexity. Furthermore, the artificial synchronization points have been weakened compared to previous implementations, unveiling look-ahead opportunities for better hardware occupancy. The overall QDWH-based polar decomposition can then be represented as a directed acyclic graph (DAG), where nodes represent computational tasks and edges define the inter-task data dependencies. The StarPU dynamic runtime system is employed to traverse the DAG, to track the various data dependencies and to asynchronously schedule the computational tasks on the underlying hardware resources, resulting in an out-of-order task scheduling. Benchmarking experiments show significant improvements against existing state-of-the-art high performance implementations (i.e., Intel MKL and Elemental) for the polar decomposition on latest shared-memory vendors' systems (i.e., Intel Haswell/Broadwell/Knights Landing, NVIDIA K80/P100 GPUs and IBM Power8), while maintaining high numerical accuracy.

Index Terms—Polar decomposition; Asynchronous execution; Dynamic runtime system; Fine-grained execution; Directed acyclic graph; High performance computing



1 INTRODUCTION

Today's most powerful supercomputers are composed of fat computational nodes over-provisioned by floating-point units [1], which may distort the balance of characteristics systems with respect to other hardware resources, such as memory per core, aggregated bandwidth, I/O nodes, interconnect, etc. Although scientific applications are often memory-bound with low arithmetic intensity kernels, and therefore limited by the bus bandwidth, we revisit the polar decomposition, an important dense linear algebra (DLA) algorithm, which can make an effective use of the predominant floating-point units provided by the current state-of-the-art hardware commercial chips (for instance, Intel Knights Landing and NVIDIA Pascal P100). The polar decomposition consists in decomposing a dense matrix $A = U_p H$, where U_p is the orthogonal polar factor and H is the positive semi-definite Hermitian polar factor. One way to achieve this decomposition is the QR -based dynamically weighted Halley (QDWH) iteration introduced by Nakatsukasa et. al [2]. The polar decomposition is a key algorithm for various scientific applications, e.g.; in continuum mechanics to decompose stress tensors and to simulate the deformation of an object; in aerospace computations [3] during strapdown inertial navigation; and other aerospace systems

to describe the rotation of one coordinate system relative to a reference coordinate system; and in chemistry [4] to help the understanding of properties in terms of electron pair (chemical bond) transferability, etc. Further applications are also reported by Higham in [5].

This paper describes the first asynchronous, task-based formulation of the QDWH-based polar decomposition and its corresponding implementation on manycore architectures. The standard algorithm requires up to six iterations to converge in double precision and to calculate the polar factor, depending on the condition number of the input matrix, involving $O(n^3)$ matrix operations at each operation. Its algorithmic complexity may, therefore, be prohibitive. Nevertheless, this challenge can be compensated for the high level of concurrency exposed at each iteration [6], [7].

This paper proposes to considerably improve previous works [2], [6], [7] from two distinct algorithmic and implementation perspectives. The former consists in replacing the *hostile* LU-based matrix condition number estimation by an *adequate* QR -based implementation for broader code portability across vendor architectures. This enables to remove the dependence on the partial pivoting, which may be poorly supported, for example, on high-end GPU-based systems, without increasing the overall algorithmic complexity. The latter algorithmic improvement has twofold aspects: (1) it permits to take advantage and exploit the structure of the identity matrix involved at each QR -based QDWH iteration, which significantly reduces the algorithmic complexity, thanks to fine-grained computations associated with a dynamic asynchronous execution; and (2) the artificial synchronization points are weakened, unveiling look-ahead opportunities for better hardware occupancy. Although this

- D. Sukkari, H. Ltaief and D. Keyes are with the King Abdullah University of Science and Technology, Thuwal, KSA, in the Division of Computer, Electrical, and Mathematical Sciences and Engineering, at the Extreme Computing Research Center.
E-mail: Dalal.Sukkari@kaust.edu.sa
- M. Faverge is with Bordeaux INP, CNRS, INRIA et Université de Bordeaux, Talence, France.

Manuscript received M XX, YYYY; revised M XX, YYYY.

paper studies the performance impact of weakening synchronization points through fine-grained computations in the context of dense linear algebra, the authors in [8], [9] have also demonstrated it for the Conjugate Gradient solver. The overall QDWH-based polar decomposition can then be represented as a directed acyclic graph (DAG), where nodes represent computational tasks and edges define the inter-task data dependencies. We employ the StarPU dynamic runtime system to unroll the DAG. StarPU tracks the various data dependencies and asynchronously schedules the computational tasks on the underlying hardware resources. The algorithm may then be executed with an out-of-order task scheduling.

StarPU increases user-productivity by establishing a separation of concerns consisting in hiding the hardware complexity from library developers. This enables end-users to target various hardware architectures with a single source code. Extensive benchmarking experiments show significant improvements against existing state-of-the-art high performance implementations (i.e., MKL and Elemental) for the polar decomposition on latest shared-memory systems (i.e., Intel Haswell/Broadwell/Knights Landing, NVIDIA K80/P100 GPUs and IBM Power8), while maintaining high numerical accuracy for well and ill conditioned matrices.

The remainder of the paper is organized as follows. Section 2 presents related work. Section 3 highlights our research contributions. Section 4 briefly recalls the polar decomposition and its main computational phases. Section 5 describes the algorithmic paradigm shift that current state-of-the-art DLA software libraries have witnessed, as implemented in LAPACK [10], MAGMA [11] and more recently in PLASMA [12] and Chameleon [13] libraries. The StarPU dynamic runtime system is illustrated in Section 6 as a scheduling engine for task-based programming model. The implementation details of the high performance task-based asynchronous QDWH are given in Section 7. Section 8 provides new upper-bound for the QDWH algorithmic complexity. Numerical accuracy, implementations assessments and performance comparisons with existing state-of-the-art DLA software are given in Section 9 and we conclude in Section 10.

2 RELATED WORK

The polar decomposition algorithm has been well studied in the last three decades in terms of complexity and numerical robustness/accuracy [14], [15], [16], [17], [18], [19], [20]. Initially designed with Newton's method based on an explicit matrix inversion calculation, numerical instability was reported, especially in presence of ill-conditioned matrices. An algorithm based on Halley's iteration was then introduced with asymptotically cubic rate of convergence in obtaining the final polar factor. To solve the numerical accuracy issues due to the matrix inversion computation, an inverse-free QR -based dynamically weighted Halley (QDWH) has finally been proposed by Nakatsukasa et. al [2]. However, the polar decomposition has not been implemented in a high performance computing environment, most likely due to its excessive algorithmic complexity, which does not reflect a practical assessment of the method. More recently, Nakatsukasa and Higham [21] have shown that QDWH

can be used as a building block for the dense symmetric eigensolver and singular value decomposition [22], [23], which has brought to the fore further research directions. Indeed, previous works from the authors have implemented QDWH-based singular value decomposition on hardware accelerators [6] and distributed-memory systems [7], where the calculation of the polar factor is the most-time consuming phase. The aforementioned implementations have somewhat demonstrated limited performance scalability on multiple GPUs and large clusters. This is mostly due to the low hardware resource occupancy achieved by the inherent bulk synchronous programming model (BSP), which both implementations rely on for parallel performance. By the same token, it is also noteworthy to mention that the high performance software library Elemental [24] provides a QDWH implementation for distributed-memory systems.

Last but not least, the polar decomposition can alternatively be computed through an SVD as follows: $A = U\Sigma V^T = UV^T V\Sigma V^T = U_p V\Sigma V^T = U_p H$. This strategy has shown some performance scalability issues, due to the slow convergence of the QR algorithm on the condensed bidiagonal form [7].

3 CONTRIBUTIONS

This section enumerates our contributions, given the previous related work section, which represent the crux of the paper: (1) improve the standard QDWH algorithm by replacing the LU-based condition estimator with QR , without increasing the overall algorithmic complexity, while enabling software portability across hardware architectures; (2) develop the first task-based QDWH implementation based on fine-grained computations, which enables to exploit the identity data structure during the QDWH iterations, reducing up to 20% the overall complexity; (3) rely on a dynamic runtime system (i.e., StarPU) to asynchronously schedule the computational tasks among available processing units in order to improve hardware occupancy; and (4) provide a comprehensive performance assessment and comparisons on a myriad of high-end architectures.

4 THE QDWH-BASED POLAR DECOMPOSITION

This section focuses on the inverse-free QDWH-based iterative procedure to calculate the polar decomposition [2], [21] of a matrix $A \in \mathbb{R}^{m \times n}$ ($m \geq n$), such that $A = U_p H$. To ensure the paper is self-contained, we briefly recall the convergent sequence as follows, with A the initial matrix:

$$U_0 = A/\alpha, \quad \alpha = \|A\|_2 \quad (1)$$

$$U_{k+1} = \frac{b_k}{c_k} U_k + \frac{1}{\sqrt{c_k}} \left(a_k - \frac{b_k}{c_k} \right) Q_1 Q_2^T, \quad k \geq 0 \quad (2)$$

where $U_p = \lim_{k \rightarrow \infty} (U_k)$, and $\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R = \begin{bmatrix} \sqrt{c_k} U_k \\ I \end{bmatrix} \cdot H$ can then be found with the two steps formula:

$$H' = U_p^T * A, \quad H = \frac{1}{2} (H' + H'^T) \quad (3)$$

The main goal consists in calculating the optimal parameters (a_k, b_k, c_k) so that cubical convergence can be attained during the QDWH iteration. The expression of the

parameters (a_k, b_k, c_k) can be written as follows:

$$\begin{aligned} a_k &= h(l_k), \quad b_k = (a_k - 1)^2/4, \quad c_k = a_k + b_k - 1, \\ l_0 &= \frac{\beta}{\alpha}, \quad l_k = \frac{l_{k-1}(a_{k-1} + b_{k-1}l_{k-1}^2)}{1 + c_{k-1}l_{k-1}^2}, \quad k \geq 1, \\ h(l) &= \sqrt{1+d} + \frac{1}{2} \sqrt{8-4d + \frac{8(2-l^2)}{l^2\sqrt{1+d}}}, \quad d = \sqrt[3]{\frac{4(1-l^2)}{l^4}}, \end{aligned} \quad (4)$$

with $\beta = 1/\|A^{-1}\|_2$. For ill-conditioned matrices, the number of iterations k can vary up to six. We refer to [2] for further details on the theoretical proof. When U_k becomes well-conditioned, it is possible to replace Eq. (2) with a Cholesky-based implementation as follows:

$$\begin{aligned} U_{k+1} &= \frac{b_k}{c_k} U_k + \left(a_k - \frac{b_k}{c_k} \right) (U_k W_k^{-1}) W_k^{-\top}, \\ W_k &= \text{chol}(Z_k), \quad Z_k = I + c_k U_k^\top U_k. \end{aligned} \quad (5)$$

This algorithmic switch at runtime allows to further speed up the overall computation, thanks to a lower algorithmic complexity, while still maintaining numerical stability. In practice, this transition is monitored by setting a threshold for c_k . Once convergence is reached, the polar factor is $U_p = U_k$ and the positive semi-definite Hermitian polar factor corresponds to $H = U_p^\top A$.

All in all, the number of floating-point operations depends on the number of iterations required to converge, which is dictated by the condition number of the original matrix problem. Typically, for ill-conditioned matrices, QDWH will perform three QR -based QDWH iterations (Eq. (2)), followed by three Cholesky-based QDWH iterations (Eq. (5)), besides executing other compute-intensive Level 3 BLAS operations, i.e., triangular solves, applications of Householder reflectors, matrix-matrix multiplications, etc.

In fact, although the QDWH-based polar decomposition is a challenging and complex algorithm, it relies on conventional dense linear algebra operations, e.g., QR /Cholesky-based linear solvers. These building block operations are well-supported by several open-source and vendor-optimized state-of-the-art numerical libraries.

Thanks to its highly-parallel and compute-bound kernels, the QDWH-based polar decomposition may take advantage of the current manycore era and a foreseen period of convergence, where hardware/software co-design plays now a major role in designing future systems and numerical libraries for exascale.

5 ALGORITHMIC PARADIGM SHIFT

High performance dense linear algebra software libraries have witnessed an algorithmic paradigm shift in response to hardware evolution, moving from block to tile algorithms.

5.1 Block Algorithms

Block algorithms rely on successive panel and update sequences to perform matrix computations. The panel phase is memory-bound and does not benefit from thread parallelism, while the phase of the trailing submatrix update is highly parallel, in which computations are applied by means of multithreaded Level 3 BLAS kernel

executions. These sequences are characteristic of the fork-join paradigm, alternating sequential and parallel computational phases, and therefore, suffer from performance losses due to low hardware occupancy engendered by unnecessary in-between synchronization points. In fact, this BSP model corresponds to the backbone of many open-source and vendors' state-of-the-art numerical libraries such as LAPACK [10], MAGMA [11] and ScaLAPACK [25] for shared-memory, accelerator-based and distributed-memory systems, respectively.

As highlighted in the exascale software roadmap [26], which summarizes the HPC community consensus on an urgent call for sustainable software development for extreme scale, the BSP model may need to be reconsidered, especially in presence of millions of cores, which already constitute today's supercomputers [1].

5.2 Tile Algorithms

To answer this call for action and provide a solution for the challenge brought by the manycore era, the DLA community has initiated a decade ago a profound redesign of matrix computation algorithms in order to benefit from the high level of concurrency. This translated into breaking down the dense matrix data structure into tiles following a tile data layout as opposed to the standard column-major format, which is the standard for block algorithms. The various matrix operations can then be represented as a directed acyclic graph (DAG), where nodes represent sequential computational tasks and edges define the inter-task data dependencies. The resulting fine-grained computations permit to weaken the artificial synchronization points by bringing to the fore opportunities for look-ahead, where subsequent tasks may already have their data dependencies satisfied and be ready for execution. In return, this can be exploited by dynamic runtime systems in keeping threads in a busy state throughout the entire execution. The performance gain of block versus tile algorithms has been thoroughly addressed in the literature [27], [28], [29], in the context of PLASMA [30] and FLAME [31] numerical software libraries.

5.3 The Chameleon Library

More recently, in a community effort to enhance user productivity by abstracting the hardware complexity, the Chameleon library [13] has been developed to target multiple hardware architectures with a single source code. This is achieved by standardizing existing dynamic runtime system APIs (e.g., OpenMP [32], OmpSs [33], [34], [35], QUARK [36], StarPU [37], [38], PaRSEC [39], SuperMatrix [40]) through a thin layer of abstraction, making the user developer experience oblivious to the underneath runtime system and its corresponding hardware deployment. For instance, this hardware/runtime-oblivious software infrastructure has been already used with StarPU [41], and more recently with OmpSs [42], in the context of computational astronomy applications. And this is in the Chameleon software environment that we deploy our QDWH implementation.

6 THE STARPU DYNAMIC RUNTIME SYSTEM

Dynamic runtime systems are critical scheduling engines in supporting task-based programming models at large scale [43]. In particular, StarPU [37] is the *de facto* dynamic runtime system for Chameleon. StarPU deals with the execution of generic task graphs, given through the sequential task flow (STF) programming model where tasks are inserted to the runtime in a sequential manner with additional hints on the data directions (i.e., read, write, read-write). StarPU is then in charge of dynamically scheduling the tasks while enforcing those dependencies. Although Chameleon supports other runtimes (e.g., ParSEC [39], QUARK [36]), we decided to solely rely on the StarPU runtime system to implement this algorithm, since it is probably one of the most mature runtime systems when it comes to supporting various hardware architectures. Comparing the performance of various runtime systems on a given hardware is interesting in the context of QDWH, but beyond the scope of this paper.

One of the main advantages of using the task-based implementation is to become oblivious of the targeted architectures. This improves the user productivity, and it is even more realistic for runtimes such as StarPU, which are able to transparently handle single heterogeneous nodes, and eventually multiple heterogeneous nodes in case the StarPU-MPI [38] extension is used. To enable such portability, StarPU tasks are associated to codelets which groups under the same name multiple implementations of the same task: CPU, CUDA, OpenCL, OpenMP, etc. At runtime, StarPU will automatically decide which implementation of the task is better suited to achieve the highest performance based on cost models. These cost models are automatically generated by StarPU when executing the application and kept for subsequent executions. These models are especially important to the Heterogeneous First Time [44] (HeFT) scheduling strategy used by StarPU, when accelerators are involved in the computations.

Further benefits to using such programming models are the capabilities offered to the programmer to submit simultaneously independent steps of an application. This permits to raise the resources occupancy, and adds a single synchronization point when all steps are performed. The `MORSE_xxxx_Tile_Async` interface of the Chameleon library offers this capability to interleave multiple dense linear algebra operations when it is possible. Conversely, the synchronous interface, `MORSE_xxxx_Tile`, enforces a synchronization call at the end of the function to wait for the end of all submitted tasks.

7 HIGH PERFORMANCE IMPLEMENTATION

In this section, we describe the task-based implementation of the QDWH algorithm and the novel optimizations introduced to increase hardware occupancy and overall performance, in the context of the Chameleon library [13].

7.1 Task-Based QDWH Pseudo-Code

Algorithm 1 presents the pseudo-code of the task-based QDWH implementation on top of the Chameleon library. It is decomposed in three main code sections. The first one

Algorithm 1 QDWH pseudo-code on top of Chameleon

```

1: /* Estimate the condition number */
2: dlapcy_Async( A, U )                                ▷ U = A
3: dlapcy_Async( A, B )                                ▷ B = A
4: Anorm = dlange_Async( A )                            ▷ ||A||1
5: dgenm2( A, α )                                       ▷ α ≈ ||A||2
6: RUNTIME_sequence_wait()
7: /* Compute U0 and l0 */
8: dlascl_Async( U, 1, /α )                             ▷ U0 = A/α
9: if lu then
10:  dgetrf_Async( B );                                  ▷ A = LU
11:  l0 = dgecon( B, Anorm )                            ▷ l0 ≈ 1/(||A-1||1 ||A||1)
12: else
13:  dgeqrf_Async( B )                                   ▷ A = QR
14:  dtrtri_Async( B )                                   ▷ Compute R-1
15:  Ainvnorm = dlantr_Async( B )                        ▷ ≈ ||A-1||1
16:  RUNTIME_sequence_wait()
17:  l0 = 1. / (Ainvnorm * Anorm)
18: end if
19: l0 = (α/1.1) * l0
20:
21: /* Compute the polar decomposition A = UpH using QDWH */
22: k = 1, Li = l0, conv = 100
23: while (conv ≥ √[3]{5eps} || |Li - 1| ≥ 5eps) do
24:  L2 = Li2, dd = √[3]{(4(1 - L2)/L22)}
25:  sqd = √(1 + dd)
26:  a1 = sqd + √(8 - 4 × dd + 8(2 - L2)/(L2 × sqd))/2
27:  a = real(a1); b = (a - 1)2/4; c = a + b - 1
28:  Li = Li(a + b × L2)/(1 + cL2)
29:  dlapcy_Async( U, U1 )                               ▷ Backup Uk-1
30:
31: /* Compute Uk from Uk-1 */
32: if c > 100 then
33:  C = [ C1 ] = [ √c Uk-1 ]
34:  dgeqrf_Async( C )                                   ▷ C = QR = [ Q1 ] R
35:  dorgqr_Async( C )                                   ▷ C = Q = [ Q1 ]
36:  dgemm_Async( Q1, Q2⊤, U ) ▷ Uk =  $\frac{1}{\sqrt{c}} (a - \frac{b}{c}) Q_1 Q_2^{\top} + \frac{b}{c} U_{k-1}$ 
37: else
38:  dlaset_Async( Z, 0, 1. )                             ▷ Z = I
39:  dgemm_Async( U⊤, U, Z )                             ▷ Zk = I - cUk-1⊤Uk-1
40:  dgeadd_Async( U, B )                                   ▷ B = Uk-1⊤
41:  dposv_Async( Z, B )                                   ▷ Solve Zkx = Uk-1⊤
42:  dgeadd_Async( B, U )                                   ▷
43:  Uk =  $\frac{b}{c} U_{k-1} + (a - \frac{b}{c}) (U_{k-1} W_{k-1}^{-1}) W_{k-1}$ 
44: end if
45:  dgeadd_Async( U, U1 )                               ▷ Uk - Uk-1
46:  dlange_Async( U1, conv )                             ▷ conv = ||Uk - Uk-1||F
47:  RUNTIME_sequence_wait()
48:  k = k + 1
49: end while
50: /* Compute H */
51: dgemm_Async( Uk, A, H )                               ▷ H = Up⊤A
52: dlapcy_Async( H, B )                                   ▷ B = H
53: dgeadd_Async( B, H )                                   ▷ H =  $\frac{1}{2}(H + H^{\top})$ 
54: RUNTIME_sequence_wait()

```

from row 1 to 6 evaluates the two-norm of the input matrix A , as in Eq. (1) from Section 4, that is required to start the iterative process. The two-norm estimator corresponds to the largest singular value of the matrix. It relies on the power iteration, which involves repeated multiplication by the matrix A and A^{\top} . The power iteration converges when the difference between two successive estimates fall within the specified relative tolerance. We have introduced `genm2` in the Chameleon library through an iterative procedure, in which we minimized the number of synchronizations, thanks to fine-grained computations and look-ahead techniques. The second section of the algorithm computes the initial condition number l_0 from row 7 to 19, as in Eq. (4) from Section 4. The classical way consists in computing an LU factorization of the matrix A , and its one-norm. Then, it is possible to compute an estimator of the condi-

tion number with d_{gecon} by means of those two results. The main challenge here resides in the LU factorization with partial pivoting, which is difficult to implement using task-based programming models. Indeed, searches for pivot candidates and row swapping generate many global synchronization points within the panel factorization and its resulting updates. Some solutions have been proposed on shared-memory systems [45] but there are no existing solutions that are oblivious of heterogeneous architectures. We thus propose a QR -based solution which consists in estimating the norm of A^{-1} by computing the norm of R^{-1} with $A = QR$. This solution, which turns out to be less costly, alleviates the pivoting issue all together, uses only regular tile algorithms and allows code portability across various architectures, thanks to the underlying runtime system. The third section of the algorithm, rows 21 to 48, is the main loop of the algorithm, which iterates on U_k and converges to the polar factors. This section of the algorithm is straightforward and follows the mathematical description of the problem using either a QR or a Cholesky factorization to calculate the next U , as in Eq. (2) or Eq. (5) from Section 4, respectively. Finally, the last section, rows 49 to 53, computes the Hermitian polar factor H from the polar factor computed out of the main loop.

7.2 Code Optimizations

The Chameleon library provides two APIs to perform dense matrix computations. The first one, `MORSE_xxxx_Tile`, is a synchronous implementation of a linear algebra operation. This means that all the tasks required for the computations are submitted to the runtime, and then the library internally waits for the completion of all tasks before returning the control to the programmer. This is the first version we implemented in the Algorithm 1. To highlight the benefit of using a task-based programming model (through tile algorithms) as opposed to the fork-join paradigm, as implemented in the LAPACK library, we have manually integrated synchronization points within the QR /Cholesky factorization kernel calls, at the end of each panel and update computations, to better emphasize on the performance discrepancy between both approaches. We refer to this reference implementation as *Sync*.

The second optimization is the possible acceleration of the QR -based Halley iterations. This optimization consists in exploiting the identity matrix structure of the C_2 matrix in the QR factorization (line 34 in Algorithm 1) and the corresponding Q generation (line 35 in Algorithm 1). Indeed, thanks to tile algorithms, it is possible to design a specific QR factorization algorithm in order to factorize a dense matrix on top of an identity matrix. This new QR factorization takes into account the identity matrix structure so that only non-zero tiles are operated on during the factorization. By the same token, during the Q generation step, only the non-zero tiles containing the Householder reflectors will be accessed. This optimization is important as it reduces the number of FLOPs as well as data movement. We refer to this implementation as *OptId*.

The last optimization, `MORSE_xxxx_Tile_Async`, ensures that all the tasks of an algorithm are submitted to the runtime, but their completion is not ensured when the

function call returns. Thus, it is possible to simultaneously submit tasks of multiple operations. This may unveil look-ahead opportunities at runtime, once the data dependencies are satisfied, and may engender out-of-order task execution. Indeed, the runtime is in charge of keeping the data coherency of tasks, generated from different kernel calls, since these tasks may operate on the same data. Operations that are asynchronously submitted to the runtime are indicated in bold font in Algorithm 1. At some point of the algorithm, synchronization points are however required to guarantee the consistency of the results. This is made through a call to `RUNTIME_sequence_wait()`, which waits for the completion of all tasks. It is then possible to release synchronization in the three steps of the algorithm to ensure a better occupancy of the resources, especially on small to medium test cases, as presented in Section 9. We refer to this implementation as *Async*. It is also noteworthy to mention that it is possible to estimate *offline* the minimal number of iterations performed in the main loop. In that case, the synchronization in line 45 can be safely removed for the first iterations and introduced only for the last iteration as a sanity check on the value *conv* against the convergence threshold.

These three code optimizations (i.e., *Sync*, *OptId* and *Async*) can be combined for further performance. While *Sync* and *Async* have a direct impact on task scheduling, *OptId* actually changes the algorithm and reduces the algorithmic complexity.

8 ARITHMETIC COMPLEXITY

In this section, we present the algorithmic complexity (FLOPs) of the polar decomposition using two variants based on the Halley iteration (QDWH) and the SVD. For simplicity purposes, we consider only square dense matrices, but QDWH works also for rectangular matrices [21].

8.1 The QDWH-based Polar Decomposition

The condition number estimation l_0 can be calculated using the LU factorization, which requires $\frac{2}{3}n^3$ operations, followed by two triangular solvers $LX = Id$ and $UA^{-1} = X$, adding $2n^3$ FLOPs. Alternatively, l_0 can be calculated using the QR factorization, $A = QR$ which needs $\frac{4}{3}n^3$ operations, followed by inverting the upper triangular matrix R with $\frac{1}{3}n^3$ operations. Calculating l_0 using the QR factorization needs less FLOPs overall. Moreover, the resulting QR factors can be reused during the first iteration of QDWH, thanks to fine-grained computations.

As shown in 2, the QDWH FLOPs using QR -based iteration includes the QR decomposition of $2n \times n$ matrix for a cost of $(3 + \frac{1}{3})n^3$ FLOPs. Then, forming $\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}$ explicitly, needs $(3 + \frac{1}{3})n^3$ FLOPs. The product $Q_1 Q_2^T$ requires $2n^3$ FLOPs. Therefore, the arithmetic cost of each QR -based iteration is $(8 + \frac{2}{3})n^3$ FLOPs. For the Cholesky-based iteration in 5, matrix-matrix multiplication involves $2n^3$ operations, the Cholesky factorization needs $\frac{1}{3}n^3$, and solving two linear systems requires $2n^3$ operations. Therefore, the arithmetic cost of Cholesky-based iteration is $(4 + \frac{1}{3})n^3$ operations. Computing the Hermitian polar factor $H = U_p^T A$

requires $2n^3$ operations. Hence, the overall cost of QDWH is $(8 + \frac{2}{3})n^3 \times \#it_{QR} + (4 + \frac{1}{3})n^3 \times \#it_{Chol} + 2n^3$, where $\#it_{QR}$ and $\#it_{Chol}$ correspond to the number of QR -based and Cholesky-based iterations, respectively. As discussed in [21], the flop count of QDWH depends on l_0 , which involves during the QDWH iteration. The total flop count of QDWH for dense matrices ranges then from $(10 + \frac{2}{3})n^3$ (for $l_0 \approx 1$ with $\#it_{Chol} = 2$) to $41n^3$ (for $l_0 \gg 1$, with typically $\#it_{QR} = 3$ and $\#it_{Chol} = 3$). Furthermore, taking advantage of the trailing identity matrix structure in the QR factorization (OptId) reduces the flop count of the iteration in 2. Forming the upper triangular matrix R by applying the Householder reflectors with $n + 1$ nonzero elements $\prod_{k=1}^{n-1} H_k A = R$ to k vectors requires $4(n + 1)k$ FLOPs, therefore forming R needs $\sum_{k=1}^{n-1} 4(n - k)(n + 1) = 2n^3$ FLOPs. Accumulating the Householder reflectors to form $\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} = \prod_{k=1}^{n-1} H_{n-k} \begin{bmatrix} I_n \\ 0 \end{bmatrix}$ requires $\sum_{k=1}^{n-1} 4nk = 2n^3$, as explained in [21]. Table 1 summarizes the total flop count of QDWH (including condition number estimation and Halley iteration) (1) when using LU to estimate l_0 (original implementation), (2) when using QR to estimate l_0 and reusing the QR factors in the first iteration of QDWH and (3) when additionally taking advantage of the identity matrix structure in QR -based iterations (2).

TABLE 1
Algorithmic complexity of the QDWH-based polar decomposition.

	well	ill
(1) QDWH+LU	$(13 + \frac{1}{3})n^3 \leq \dots$	$\leq (43 + \frac{2}{3})n^3$
(2) QDWH+QR	$(12 + \frac{1}{3})n^3 \leq \dots$	$\leq (41 + \frac{1}{3})n^3$
(3) QDWH+QR+OptId	$(12 + \frac{1}{3})n^3 \leq \dots$	$\leq (33 + \frac{1}{3})n^3$

8.2 The SVD-based Polar Decomposition

The polar decomposition can be calculated via SVD as follows, $A = U\Sigma V^T$, then, $U_p = UV$, $H = V^T \Sigma V$. Therefore, the flop count of this approach includes the cost of an SVD decomposition, a matrix-matrix multiplication to compute the orthogonal polar factor U_p and a matrix-matrix multiplication to calculate the Hermitian polar factor H .

The standard approach to compute the SVD of a dense matrix is to first reduce it to bidiagonal form $A = U_1 B V_1^T$. The subsequent left and right singular vectors from the bidiagonal solver are then accumulated during the back transformation phase, i.e., $U = U_1 U_2$ and $V = V_2 V_1$, to calculate the singular vectors of the original matrix A . The final estimated flop count to calculate the SVD is $22n^3$, as implemented in the divide-and-conquer DGESDD [46]. Then, we need to add $2n^3$ operations to compute $U_p = UV$, and n^3 to compute $H = V^T \Sigma V$ (symmetric rank- k update operation). The final estimated cost of the polar decomposition using SVD is, therefore, $25n^3$.

Compared to the QDWH-based polar decomposition (3) in Table 1, this is 30% less than in case of ill-conditioned

matrices and almost twice the FLOPs in case of well-conditioned matrices. In theory, it seems there is a clear advantage to use SVD-based for the polar decomposition in presence of ill-conditioned matrices. However, the SVD algorithm inherently suffers from lack of parallelism, due to a very expensive panel factorization phase and may not be as competitive as QDWH-based approaches.

9 PERFORMANCE RESULTS AND ANALYSIS

This section provides a comprehensive performance analysis of the task-based QDWH implementation in the context of the Chameleon library with the dynamic runtime system StarPU on various architectures.

9.1 Environment Settings

We have considered three different single node systems, which are representative of the current manycore-based hardware trends. The first system is composed of dual-socket 16-core Intel Haswell Xeon CPU E5-2698 v3 running at 2.30GHz equipped with 8 K80 dual-boards with 16 effective GPUs. In the following, we call this system Haswell when no GPU are used, and we add the suffix '+8xK80' whenever both CPUs and GPUs are exploited. The second system hosts the latest Intel commodity chip with dual-socket 14-core Intel Broadwell Xeon E5-2680 v4 running at 2.4GHz. The third system has the latest Intel manycore Knights Landing (KNL) 7210 chips with 64 cores. For simplicity purposes, each system is named after its chip codename.

All numerical accuracy and parallel performance (time in seconds and GFLOP/s) graphs report experiments performed on the whole system, unless mentioned otherwise. The GFLOP/s graphs are shown in linear scale while the accuracy and time to solution graphs are shown in logarithmic scale.

Our QDWH implementation has been compiled with Intel compiler 16 and linked against the Chameleon library v0.9.0 with hwloc v1.11.4, StarPU v1.2.0 and Intel MKL v11.3.1. Each dense synthetic matrix $A = QDQ^T$ is generated by initially setting a diagonal matrix $D = \text{diag}(\Sigma)$ containing the singular values, with a specific condition number and from an orthogonal matrix Q generated by calculating the QR factorization of a random entries.

We have considered well and ill-conditioned randomly generated matrices, with the latter representing the worse case scenario, where QDWH performs a maximum of six iterations. In particular, in the subsequent experiments, our QDWH implementation switches Equations from (2) to (5) from Section 4 if c_k is smaller than 100 (see Algorithm 1), which generates QR -based iterations for the first three followed by three Cholesky-based iterations.

The QDWH performance graphs in GFlop/s are estimated as the ratio of the algorithmic complexity (see Section 8.1) divided by the execution time. All performance runs have been repeated three times and only the average is reported, since consistent timing results are obtained.

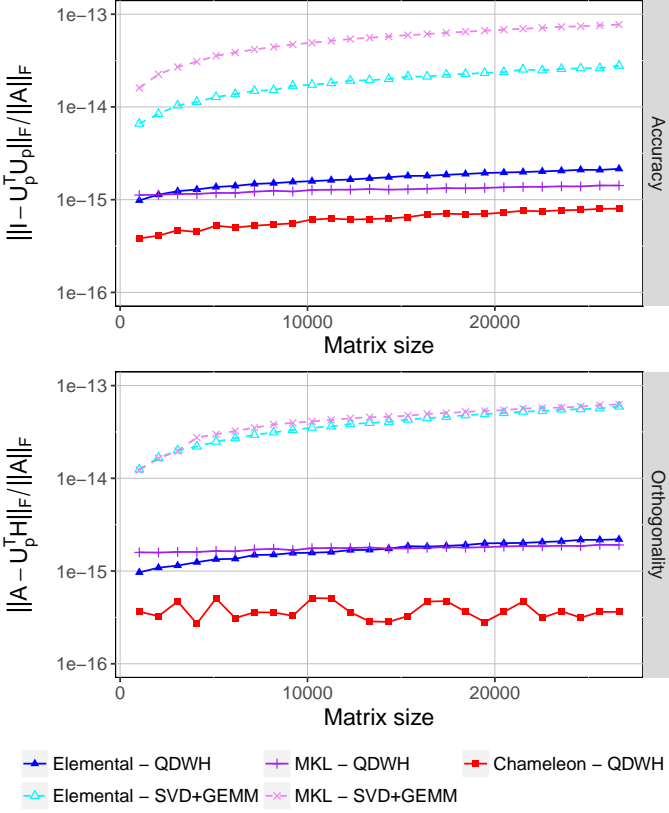


Fig. 1. Assessing the numerical accuracy/robustness of the task-based QDWH.

9.2 Numerical Accuracy

We recall the polar decomposition of a given general matrix $A \in \mathbb{R}^{n \times n}$: $A = U_p H$. The norm $\|\cdot\|_F$ denotes the Frobenius norm. To highlight the numerical robustness of the method, we use the following two accuracy metrics: $\frac{\|I - U_p^T U_p\|_F}{\|A\|_F}$ for the orthogonality of the polar factor U_p and $\frac{\|A - U_p H\|_F}{\|A\|_F}$ for the accuracy of the overall computed polar decomposition. Fig. 1 presents the orthogonality of U_p and the accuracy of the polar decomposition $A = U_p H$ for ill-conditioned matrix on the KNL system (very similar numerical results on other systems). We can distinguish two clusters, i.e., QDWH-based and SVD-base polar decomposition, with up to two digits difference in the orthogonality and accuracy magnitudes. Although both mostly employ orthogonal transformations, the SVD variant of the polar decomposition necessitates the DC algorithm, which may show some convergence limitations with ill-conditioned matrices, as shown later in Section 9.6.

9.3 Incremental Optimizations

Fig. 2 highlights the performance impact of various incremental optimizations on the task-based QDWH, as described in Section 7.2. Taking advantage of the identity matrix structure (OptId), by only operating on non-zero tiles, engenders up to 20% performance improvements compared to the oblivious approach on all studied systems.

Thanks to fine-grained computations, look-ahead opportunities are unveiled, which directly translate into asynchronous out-of-order task execution at runtime. In fact,

this has been previously reported in the literature [40], [42], [47] for many of the dense linear algebra matrix operations, which compose our task-based QDWH implementation (e.g., Cholesky, QR , etc). Therefore, running additionally in asynchronous mode (Async) further reduces time to solution (up to 2.8x). This is true, especially for medium range of matrix sizes, where processing units run out of work and look-ahead techniques jump right in to fill the performance gap. For asymptotic matrix sizes, although work is abundant, the asynchronous mode still provides additional performance. In particular, on KNL and Haswell+8xK80 systems, data movement engendered by NUMA and PCIe channels is expensive and can be overlapped by computations, thanks to the Async optimization.

9.4 Execution Traces

Fig. 3 shows the execution traces when running in synchronous (*Sync* API) and asynchronous (*Async* API) modes. We have added additional synchronization points within the *Sync* API, after each panel/update computation, so that we can better capture the performance gain against coarse-grained computations engendered by block algorithms, as described in Section 5. These traces have been obtained on the KNL system for a matrix size of 10K. Since the matrix is ill-conditioned, the task-based QDWH performs six iterations (three QR -based and three Cholesky-based). The green, yellow and blue blocks correspond to QR , Cholesky/Level 3 BLAS and Level 1/2 BLAS, respectively. We can clearly notice the idle time during the first three QR -based iterations when running with a synchronous mode (Fig. 3(a)). The performance impact of synchronous execution for the next three Cholesky-based iterations is not as severe as QR -based iterations because the Cholesky panel factorization involves only the diagonal block (Fig. 3(b)). For the subsequent graphs, the performance curves of the task-based QDWH correspond to performance when all optimizations are enabled (i.e., Async and OptId).

9.5 Performance Scalability

Fig. 4 demonstrates the performance scalability of the task-based QDWH implementation. The scalability is almost linear for the commodity CPU systems. For instance, 7.5/3.7-fold speedups are achieved using 32/28 threads from the reference points of 4/7 threads on the Haswell/Broadwell systems, respectively. On the KNL system, the task-based QDWH implementation obtains a 1.9-fold speedup on 64 threads, compared to 32 threads.

On the densely GPU populated Haswell+8xK80 system, with a total of 16 GPUs, moving data between host and device memory turns out to be challenging. The performance bottleneck has been reported in a recent study [48]. By using the StarPU framework, the data movement overhead of moving data through the thin PCIe bus is partially hidden. StarPU is able to cope with some of these communication overheads by mitigating and adjusting to the memory congestion, thanks to its asynchronous mode of execution. The obtained speedup is 1.4-fold on 16 GPUs, compared to 8 GPUs. Further optimizations are possible by integrating into StarPU some of the performance models discussed in [48].

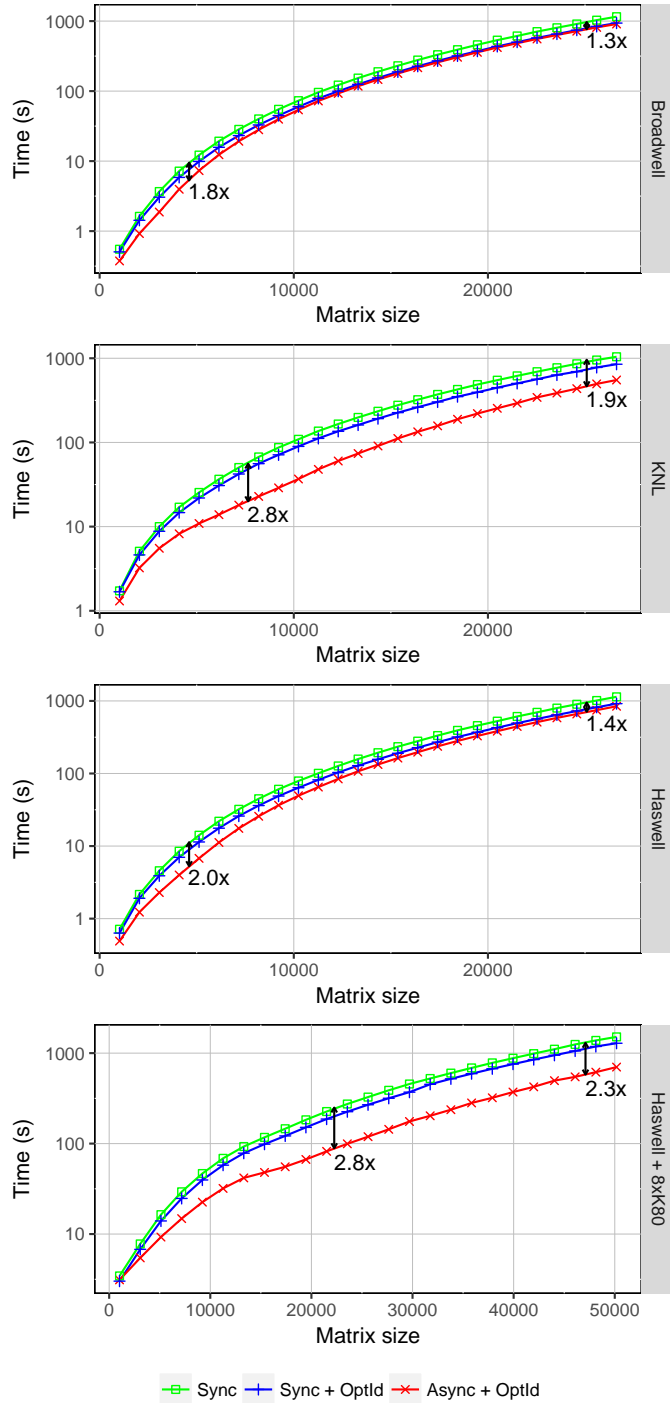
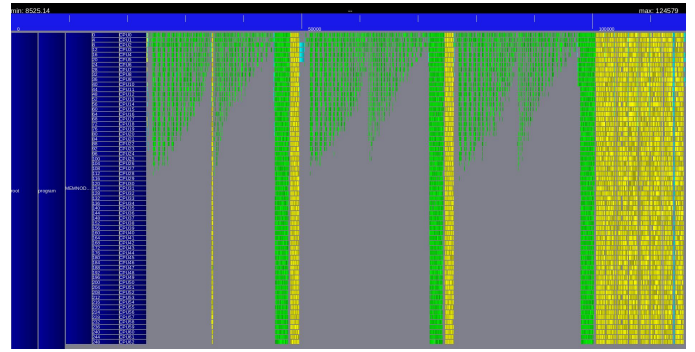


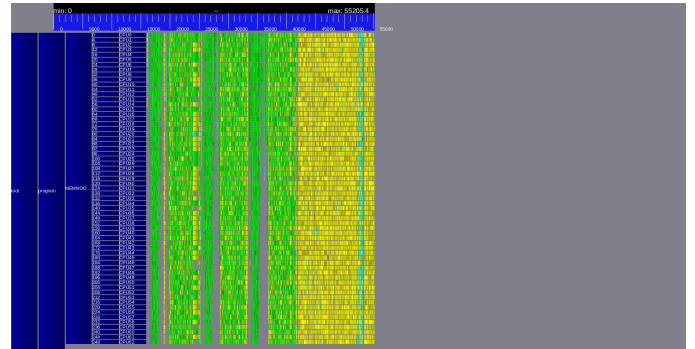
Fig. 2. Assessing the performance of various incremental optimizations.

9.6 Performance Comparisons of QDWH Variants

Fig. 5 reports task-based QDWH performance against other various existing QDWH implementations and SVD-based polar decomposition on ill (left) and well (right) conditioned matrices, across the three systems. The missing data points for the polar decomposition variant based on the SVD from MKL correspond to runs, which did not achieve the proper accuracy, as defined in Section 9.2, probably due to the convergence failure encountered by the algorithm in the SVD. The corresponding variant with Elemental does not face this problem because it uses a different imple-



(a) Synchronous task-based QDWH.



(b) Asynchronous task-based QDWH.

Fig. 3. Comparison of synchronous vs asynchronous execution traces of task-based QDWH on the KNL system with a matrix size of 10240.

mentation of the QR algorithm in the SVD instead. For well-conditioned matrices, time to solution is much more shortened for the QDWH implementation variants, thanks to less iterations for convergence. The SVD variants of the polar decomposition do not seem to take advantage of such matrices since the bidiagonal reduction and the matrix-matrix multiplication have still to be performed in the same manner, regardless of the matrix condition number.

All in all, the task-based QDWH achieves gains up to [6%, 8%] on Haswell and [39%, 17%] on Broadwell, [85%, 82%] on Haswell+8xK80, and [63%, 67%] on KNL against the best (non task-based) implementation for [ill, well]-conditioned matrices, respectively. Highest performance are achieved on systems where data movement are most expensive (e.g., NUMA for KNL and PCIe for Haswell+8xK80) since the asynchronous mode may still mitigate the overhead of data transfers by overlapping communications with task computations. But recent work [48] may further enhance the reported work.

It is noteworthy to mention that our task-based QDWH implementation using Chameleon on the 16 GPUs does not perform well on small matrix sizes. At that matrix scale, the workloads are too small to saturate all the devices' floating-point units, and therefore, performance is ultimately limited by the overhead of off-loading data back and forth.

Also, compared to MAGMA_QDWH [6], the task-based QDWH achieves gains up to [71%, 22%] on Haswell+4xK80 for [ill, well]-conditioned matrices, respectively.

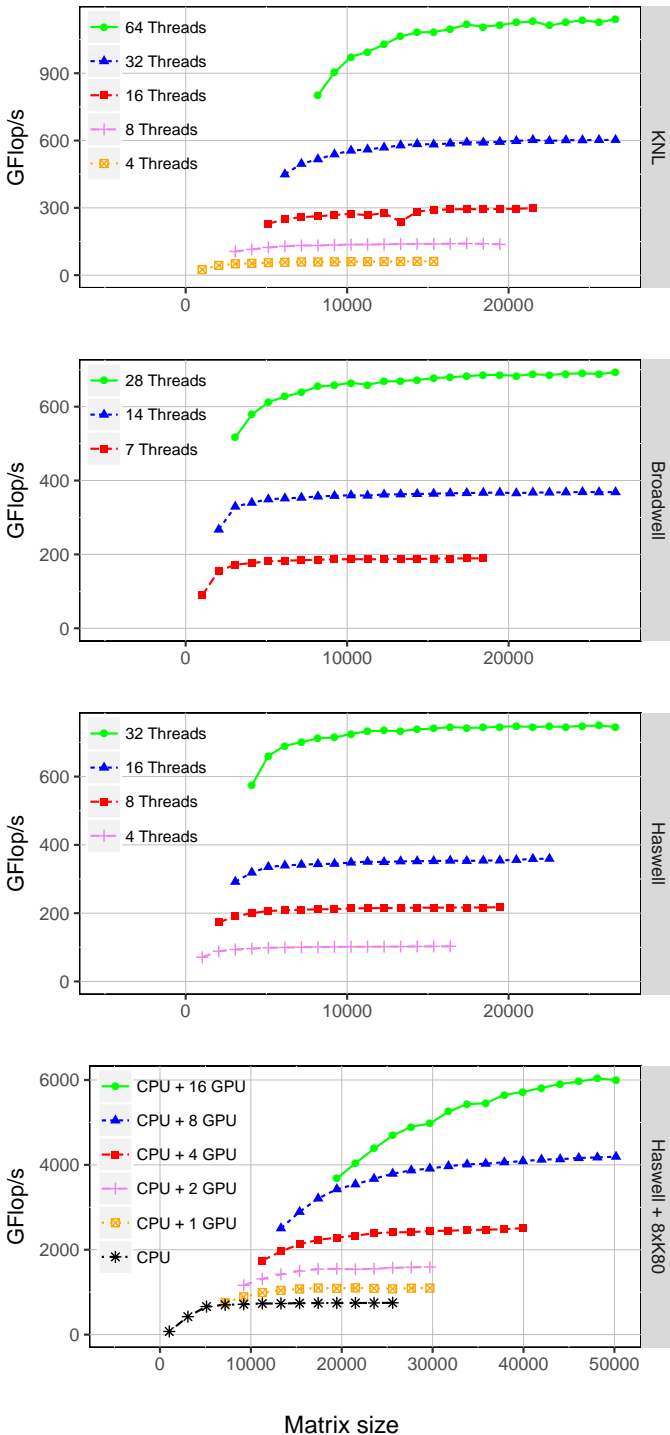


Fig. 4. Assessing the task-based QDWH scalability.

9.7 Performance Comparisons Across Architectures

We have additionally considered two more recent architectures, i.e., a dual-socket 10-cores IBM Power8 (3.69GHz) and a dual-socket 16-cores Intel Haswell equipped with four NVIDIA Pascal P100 GPUs. Fig. 6 presents the performance of the task-based QDWH across all systems investigated in the paper. The main idea is not to cross-compare the performance delivered by each system but rather to show that the task-based QDWH can support not only various architectures but also can achieve decent sustained peak (up

to 90% and up to 60% of the sustained Chameleon DGEMM peak for CPU only systems and for KNL/GPUs platforms, respectively).

10 CONCLUSION AND FUTURE WORK

We have presented a comprehensive performance analysis of a novel asynchronous task-based QDWH algorithm for the polar decomposition of a dense matrix. Thanks to fine-grained computations, we have reduced by 20% the overall complexity by taking advantage of the identity structure of the matrix during the iterations, while exposing look-ahead opportunities to increase hardware occupancy. Furthermore, the Chameleon library and its dynamic runtime system StarPU abstracts the hardware complexity from end-users and is capable of asynchronously scheduling computational tasks on the underlying processing units. Thanks to its wide hardware range support, we demonstrated that StarPU can port a single sequential source code to a myriad of hardware systems. Experimental results of the asynchronous task-based QDWH show significant performance improvement (up to an order of magnitude) against state-of-the-art implementations on ill and well-conditioned matrices across various hardware technologies, which are paving the road to future petascale/exascale systems.

Future work includes using the task-based QDWH implementation as a building block for the dense symmetric eigensolver and SVD on shared and distributed-memory systems. In particular, the extension to distributed-memory systems will necessitate the redesign of the existing QR factorization in the Chameleon library, which is currently based on a flat tree [27]. This will engender excessive communications in distributed-memory systems, although it does not prevent our task-based QDWH implementation from getting high performance for the matrix sizes and systems' configurations studied in this paper. The idea will be to adapt a hierarchical reduction tree [49] into our customized QR factorization to privilege intra-node communications.

Last but not least, we would like to investigate other QDWH algorithmic variants, which may require more FLOPs but entails an even higher level of concurrency [50].

ACKNOWLEDGMENTS

The authors would like to thank Samuel Thibault from Inria for his support with StarPU, Jack Poulson from Google Inc. for his help in tuning Elemental and the vendors Cray/IBM/Intel/NVIDIA for their hardware donations and/or systems' remote accesses in the context of the Cray Center of Excellence, the Intel Parallel Computing Center and the NVIDIA GPU Research Center awarded to the Extreme Computing Research Center at KAUST.

REFERENCES

- [1] "The Top500 List," <http://www.top500.org/>.
- [2] Y. Nakatsukasa, Z. Bai, and F. Gygi, "Optimizing Halley's Iteration for Computing the Matrix Polar Decomposition," *SIAM Journal on Matrix Analysis and Applications*, pp. 2700–2720, 2010.
- [3] J. Meyer and I. Y. Bar-itzhack, "Practical Comparison of Iterative Matrix Orthogonalization Algorithms," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-13, no. 3, pp. 230–235, May 1977.

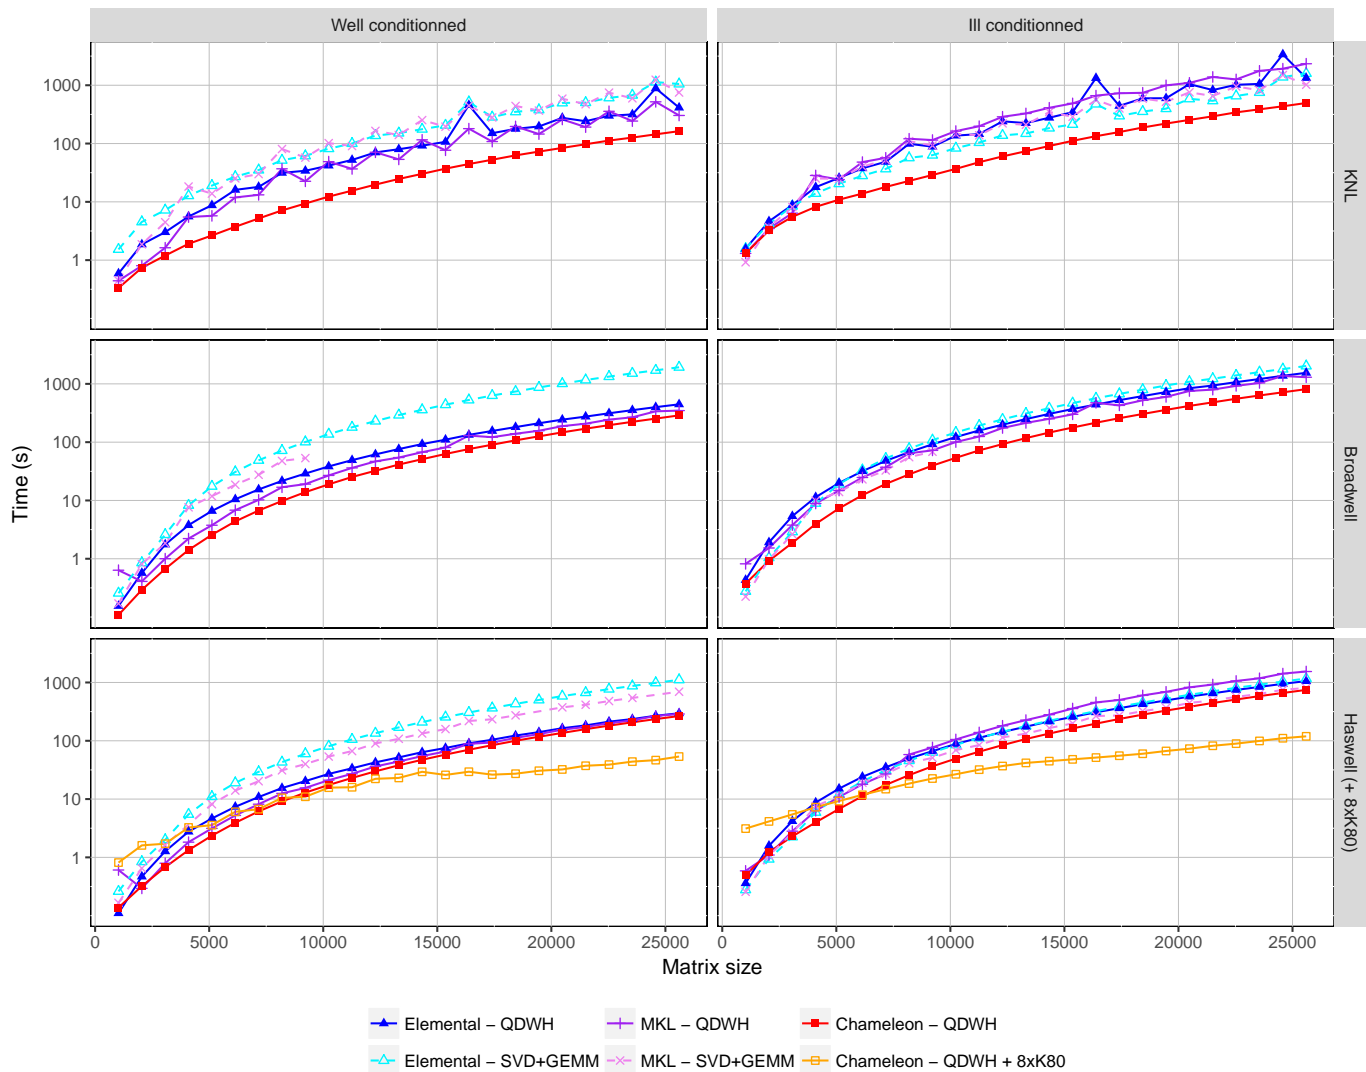


Fig. 5. Assessing task-based QDWH performance against other QDWH variant implementations on ill (left) and well (right) conditioned matrices.

- [4] J. A. Goldstein and M. Levy, "Linear algebra and quantum chemistry," *Am. Math. Monthly*, vol. 98, no. 10, pp. 710–718, Oct. 1991. [Online]. Available: <http://dx.doi.org/10.2307/2324422>
- [5] N. J. Higham, "Computing the Polar Decomposition with Applications," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 4, pp. 1160–1174, 1986. [Online]. Available: <http://dx.doi.org/10.1137/0907079>
- [6] D. Sukkari, H. Ltaief, and D. Keyes, "A High Performance QDWH-SVD Solver Using Hardware Accelerators," *ACM Trans. Math. Softw.*, vol. 43, no. 1, pp. 6:1–6:25, Aug 2016. [Online]. Available: <http://doi.acm.org/10.1145/2894747>
- [7] —, "High Performance Polar Decomposition on Distributed Memory Systems," in *Best Papers, Euro-Par 2016: Parallel Processing: 22nd International Conference on Parallel and Distributed Computing, Grenoble, France, August 24-26, 2016, Proceedings*, P.-F. Dutoit and D. Trystram, Eds. Cham: Springer International Publishing, 2016, pp. 605–616. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-43659-3_44
- [8] S. Zhuang and M. Casas, "Iteration-Fusing Conjugate Gradient," in *Proceedings of the International Conference on Supercomputing*, ser. ICS '17. New York, NY, USA: ACM, 2017, pp. 21:1–21:10. [Online]. Available: <http://doi.acm.org/10.1145/3079079.3079091>
- [9] E. Agullo, L. Giraud, A. Guermouche, S. Nakov, and J. Roman, *Task-Based Conjugate Gradient: From Multi-GPU Towards Heterogeneous Architectures*. Cham: Springer International Publishing, 2017, pp. 69–82. [Online]. Available: https://doi.org/10.1007/978-3-319-58943-5_6
- [10] E. Anderson, Z. Bai, C. H. Bischof, L. S. Blackford, J. W. Demmel, J. J. Dongarra, J. J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. C. Sorensen, *LAPACK User's Guide*, 3rd ed. Philadelphia: Society for Industrial and Applied Mathematics, 1999.
- [11] MAGMA, "Matrix Algebra on GPU and Multicore Architectures. Innovative Computing Laboratory, University of Tennessee. Available at <http://icl.cs.utk.edu/magma/>," 2009.
- [12] 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," in *Journal of Physics: Conference Series*, vol. 180, 2009.
- [13] "The Chameleon project," January 2016, <https://project.inria.fr/chameleon/>.
- [14] W. Gander, "On Halley's iteration method," *American Mathematical Monthly*, vol. 92, no. 2, pp. 131–134, 1985.
- [15] —, "Algorithms for the polar decomposition," *SIAM J. Scientific Computing*, vol. 11, no. 6, pp. 1102–1115, 1990. [Online]. Available: <http://dx.doi.org/10.1137/0911062>
- [16] C. S. Kenney and A. J. Laub, "On scaling Newton's method for polar decomposition and the matrix sign function," *SIAM J. Matr. Anal. Appl.*, vol. 13, pp. 688–706, 1992, cited in a personal communication by Alan Laub.
- [17] N. J. Higham and P. Papadimitriou, "A parallel algorithm for computing the polar decomposition," *Parallel Computing*, vol. 20, no. 8, pp. 1161–1173, Aug. 1994.
- [18] A. Kielbasinski and K. Zietak, "Numerical behaviour of

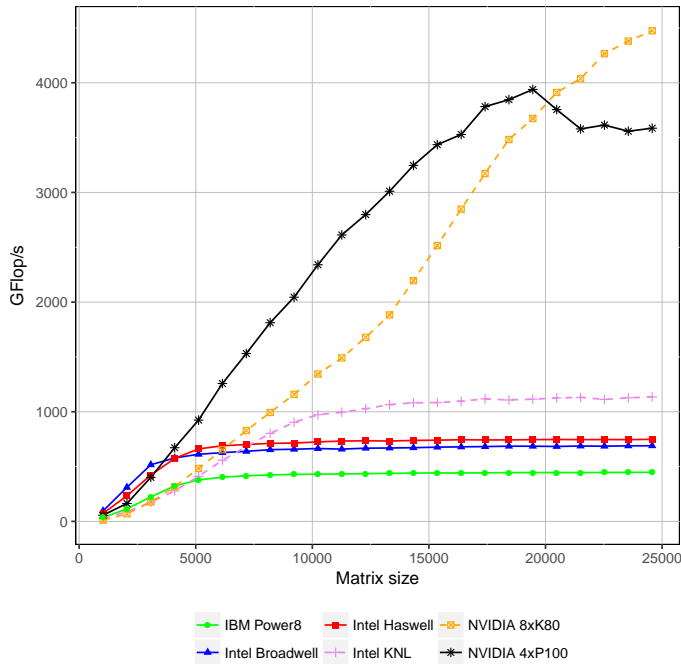


Fig. 6. Task-based QDWH performance across various architectures.

higham's scaled method for polar decomposition," *Numerical Algorithms*, vol. 32, no. 2-4, pp. 105-140, 2003. [Online]. Available: <http://dx.doi.org/10.1023/A:1024098014869>

- [19] R. Byers and H. Xu, "A new scaling for Newton's iteration for the polar decomposition and its backward stability," *SIAM J. Matrix Analysis Applications*, vol. 30, no. 2, pp. 822-843, 2008. [Online]. Available: <http://dx.doi.org/10.1137/070699895>
- [20] B. Laszkiewicz and K. Zietak, "Approximation of matrices and a family of gander methods for polar decomposition," *BIT Numerical Mathematics*, vol. 46, no. 2, pp. 345-366, 2006. [Online]. Available: <http://dx.doi.org/10.1007/s10543-006-0053-4>
- [21] Y. Nakatsukasa and N. J. Higham, "Stable and Efficient Spectral Divide and Conquer Algorithms for the Symmetric Eigenvalue Decomposition and the SVD," *SIAM Journal on Scientific Computing*, vol. 35, no. 3, pp. A1325-A1349, 2013. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/120876605>
- [22] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., ser. John Hopkins Studies in the Mathematical Sciences. Baltimore, Maryland: Johns Hopkins University Press, 1996.
- [23] L. N. Trefethen and D. Bau, *Numerical Linear Algebra*. Philadelphia, PA: SIAM, 1997. [Online]. Available: <http://www.siam.org/books/OT50/Index.htm>
- [24] J. Poulson, B. Marker, R. A. van de Geijn, J. R. Hammond, and N. A. Romero, "Elemental: A new framework for distributed memory dense matrix computations," *ACM Trans. Math. Softw.*, vol. 39, no. 2, p. 13, 2013. [Online]. Available: <http://doi.acm.org/10.1145/2427023.2427030>
- [25] L. S. Blackford, J. Choi, A. Cleary, E. F. D'Azevedo, J. W. Demmel, I. S. Dhillon, J. J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. W. Walker, and R. C. Whaley, *ScaLAPACK Users' Guide*. Philadelphia: Society for Industrial and Applied Mathematics, 1997.
- [26] J. Dongarra, P. Beckman, T. Moore, P. Aerts, G. Aloisio, J.-C. Andre, D. Barkai, J.-Y. Berthou, T. Boku, B. Braunschweig, F. Cappello, B. Chapman, X. Chi, A. Choudhary, S. Dosanjh, T. Dunning, S. Fiore, A. Geist, B. Gropp, R. Harrison, M. Hereld, M. Heroux, A. Hoisie, G. Hotta, Z. Jin, Y. Ishikawa, F. Johnson, S. Kale, R. Kenway, D. Keyes, B. Kramer, J. Labarta, A. Lichnewsky, T. Lippert, B. Lucas, B. Maccabe, S. Matsuoka, P. Messina, P. Michielse, B. Mohr, M. S. Mueller, W. E. Nagel, H. Nakashima, M. E. Papka, D. Reed, M. Sato, E. Seidel, J. Shalf, D. Skinner, M. Snir, T. Sterling, R. Stevens, F. Streitz, B. Sugar, S. Sumimoto, W. Tang, J. Taylor, R. Thakur, A. Trefethen, M. Valero, A. Van Der Steen, J. Vetter, P. Williams, R. Wisniewski, and K. Yelick, "The International Exascale Software Project Roadmap," *Int. J. High Perform. Comput. Appl.*, vol. 25, no. 1, pp. 3-60, Feb. 2011. [Online]. Available: <http://dx.doi.org/10.1177/1094342010391989>
- [27] A. Buttari, J. Langou, J. Kurzak, and J. J. Dongarra, "A class of parallel tiled linear algebra algorithms for multicore architectures," *Parallel Comput. Syst. Appl.*, vol. 35, pp. 38-53, 2009. <http://dx.doi.org/10.1016/j.parco.2008.10.002>
- [28] G. Quintana-Orti, E. S. Quintana-Orti, R. A. V. D. Geijn, F. G. V. Zee, and E. Chan, "Programming matrix algorithms-by-blocks for thread-level parallelism," *ACM Trans. Math. Softw.*, vol. 36, pp. 14:1-14:26, July 2009. [Online]. Available: <http://doi.acm.org/10.1145/1527286.1527288>
- [29] E. Agullo, B. Hadri, H. Ltaief, and J. Dongarra, "Comparative study of one-sided factorizations with multiple software packages on multi-core hardware," in *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. New York, NY, USA: ACM, 2009, pp. 1-12.
- [30] *PLASMA Users' Guide, Parallel Linear Algebra Software for Multicore Architectures, Version 2.3*, University of Tennessee, November 2010.
- [31] "The FLAME project," April 2010, <http://zcs.utexas.edu/wiki/flame/wiki/FrontPage>.
- [32] A. OpenMP, "OpenMP application program interface version 3.0," 2008. [Online]. Available: <http://www.openmp.org/wp-content/uploads/spec30.pdf>
- [33] A. Duran, R. Ferrer, E. Ayguadé, R. M. Badia, and J. Labarta, "A Proposal to Extend the OpenMP Tasking Model with Dependent Tasks," *International Journal of Parallel Programming*, vol. 37, no. 3, pp. 292-305, 2009. [Online]. Available: <http://dx.doi.org/10.1007/s10766-009-0101-1>
- [34] J. Bueno, J. Planas, A. Duran, R. M. Badia, X. Martorell, E. Ayguad, and J. Labarta, "Productive Programming of GPU Clusters with OmpSs," in *2012 IEEE 26th International Parallel and Distributed Processing Symposium*, May 2012, pp. 557-568.
- [35] J. Planas, R. M. Badia, E. Ayguadé, and J. Labarta, "AMA: Asynchronous Management of Accelerators for Task-based Programming Models," *Procedia Computer Science*, vol. 51, pp. 130-139, 2015, international Conference On Computational Science, ICCS 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1877050915010200>
- [36] A. YarKhan, J. Kurzak, and J. Dongarra, "QUARK Users' Guide: QUEuing And Runtime for Kernels," *University of Tennessee Innovative Computing Laboratory Technical Report ICL-UT-11-02*, 2011.
- [37] 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*, vol. 23, no. 2, pp. 187-198, 2011. [Online]. Available: <https://hal.inria.fr/inria-00550877>
- [38] E. Agullo, O. Aumage, M. Faverge, N. Furmento, F. Pruvost, M. Sergent, and S. Thibault, "Achieving High Performance on Supercomputers with a Sequential Task-based Programming Model," Inria Bordeaux Sud-Ouest ; Bordeaux INP ; CNRS ; Université de Bordeaux ; CEA, Research Report RR-8927, Jun. 2016. [Online]. Available: <https://hal.inria.fr/hal-01332774>
- [39] G. Bosilca, A. Bouteiller, A. Danalis, T. Herault, P. Lemarinier, and J. Dongarra, "DAGuE: A generic distributed DAG engine for High Performance Computing," *Parallel Computing*, vol. 38, no. 1-2, pp. 37-51, 2012, extensions for Next-Generation Parallel Programming Models. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167819111001347>
- [40] E. Chan, E. S. Quintana-Orti, G. Quintana-Orti, and R. van de Geijn, "Supermatrix Out-of-Order Scheduling of Matrix Operations for SMP and Multi-Core Architectures," in *SPAA '07: Proceedings of the nineteenth annual ACM symposium on Parallel algorithms and architectures*. New York, NY, USA: ACM, 2007, pp. 116-125.
- [41] A. Charara, H. Ltaief, D. Gratadour, D. E. Keyes, A. Sevin, A. Abdelfattah, E. Gendron, C. Morel, and F. Vidal, "Pipelining Computational Stages of the Tomographic Reconstructor for Multi-Object Adaptive Optics on a Multi-GPU System," in *SC'14. IEEE*, 2014, pp. 262-273. [Online]. Available: <http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=7012142>
- [42] H. Ltaief, D. Gratadour, A. Charara, and E. Gendron, "Adaptive Optics Simulation for the World's Largest Telescope on Multicore Architectures with Multiple GPUs," in *Proceedings of the Platform for Advanced Scientific Computing Conference*, ser. PASC '16. New York, NY, USA: ACM, 2016, pp. 9:1-9:12. [Online]. Available: <http://doi.acm.org/10.1145/2929908.2929920>
- [43] J. Bennett, R. Clay, G. Baker, M. Gamell, D. Hollman, S. Knight, H. Kolla, G. Sjaardema, N. Slattengren, K. Teranishi, J. Wilke,

M. Bettencourt, S. Bova, K. Franko, P. Lin, R. Grant, S. Hammond, S. Olivier, L. Kale, N. Jain, E. Mikida, A. Aiken, M. Bauer, W. Lee, E. Slaughter, S. Treichler, M. Berzins, T. Harman, A. Humphrey, J. Schmidt, D. Sunderland, P. McCormick, S. Gutierrez, M. Schulz, A. Bhatele, D. Boehme, P.-T. Bremer, and T. Gamblin, "Asynchronous Many-Task Runtime System Analysis and Assessment for Next Generation Platforms," *SANDIA Report*, no. SAND2015-8312, 2015.

- [44] H. Topcuoglu, S. Hariri, and M.-Y. Wu, "Performance-effective and low-complexity task scheduling for heterogeneous computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 13, no. 3, pp. 260–274, 2002.
- [45] J. Dongarra, M. Faverge, H. Ltaief, and P. Luszczek, "Exploiting fine-grain parallelism in recursive LU factorization." in *PARCO*, 2011, pp. 429–436.
- [46] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, ser. Mathematical Modeling and Computation. Philadelphia: Society for Industrial and Applied Mathematics, 1998. [Online]. Available: http://books.google.com.sa/books?id=A5XWVG_PFFdcC
- [47] A. Haidar, H. Ltaief, A. YarKhan, and J. Dongarra, "Analysis of Dynamically Scheduled Tile Algorithms for Dense Linear Algebra on Multicore Architectures," *Concurr. Comput. : Pract. Exper.*, vol. 24, no. 3, pp. 305–321, Mar. 2011. [Online]. Available: <http://dx.doi.org/10.1002/cpe.1829>
- [48] M. Martinasso, G. Kwasniewski, S. R. Alam, T. C. Schulthess, and T. Hoefler, "A PCIe Congestion-aware Performance Model for Densely Populated Accelerator Servers," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '16. Piscataway, NJ, USA: IEEE Press, 2016, pp. 63:1–63:11. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3014904.3014989>
- [49] J. Dongarra, M. Faverge, T. Héroult, M. Jacquelin, J. Langou, and Y. Robert, "Hierarchical QR factorization algorithms for multi-core clusters," *Parallel Computing*, vol. 39, no. 4-5, pp. 212–232, 2013. [Online]. Available: <https://hal.inria.fr/hal-00809770>
- [50] Y. Nakatsukasa and R. W. Freund, "Computing Fundamental Matrix Decompositions Accurately via the Matrix Sign Function in Two Iterations: The Power of Zolotarev's Functions," *SIAM Review*, vol. 58, no. 3, pp. 461–493, 2016. [Online]. Available: <http://dx.doi.org/10.1137/140990334>



bra algorithms for those systems.

Mathieu Faverge Dr Mathieu Faverge holds an Assistant Professor position at Bordeaux INP and is member of the Inria HiePACS team since 2012. His main research interests are numerical linear algebra algorithms for sparse and dense problems on massively parallel architectures, and especially DAG algorithms relying on dynamic schedulers. He has experience with hierarchical shared memory, heterogeneous and distributed systems, and his contributions to the scientific community include efficient linear algebra



(IEEE).

David Keyes David Keyes is the director of the Extreme Computing Research Center at King Abdullah University of Science and Technology and an adjunct professor of applied mathematics at Columbia University. Keyes earned his BSE from Princeton and his PhD from Harvard. He works at the algorithmic interface between parallel computing and the numerical analysis of partial differential equations. He is a Fellow of SIAM and AMS and has won the Gordon Bell Prize (ACM) and the Sidney Fernbach Award



Dalal Sukkari Dalal Sukkari is a PhD student in Applied Mathematics at KAUST. She has worked on a new high performance implementation of the QR -based Dynamically Weighted Halley Singular Value Decomposition (QDWH-SVD) solver on multicore architecture enhanced with multiple GPUs. She has introduced a high performance QDWH implementation on distributed memory based on the state-of-the-art vendor-optimized numerical library ScaLAPACK. Dalal received her MSc degree in Applied Mathematics

from KAUST in July 2013 with Dean's Academic Excellence Award, and her BS in Mathematics from Hashemite University of Jordan.



Hatem Ltaief Hatem Ltaief is a Senior Research Scientist in the Extreme Computing Research Center at KAUST. His research interests include parallel numerical algorithms, parallel programming models, and performance optimizations for multicore architectures and hardware accelerators. His current research collaborators include INRIA, Observatoire de Paris, Cray, NVIDIA, and Intel.