TOWARD HIGH PERFORMANCE TILE DIVIDE AND CONQUER ALGORITHM FOR THE DENSE SYMMETRIC EIGENVALUE PROBLEM *

AZZAM HAIDAR †, HATEM LTAIEF ‡, AND JACK DONGARRA §

Abstract. Classical solvers for the dense symmetric eigenvalue problem suffer from the first step involving a reduction to tridiagonal form that is dominated by the cost of accessing memory during the panel factorization. The solution is to reduce the matrix to a banded form, which then requires the eigenvalues of the banded matrix to be computed. The standard D&C algorithm can be modified for this purpose. The paper combines this insight with tile algorithms that can be scheduled via a dynamic runtime system to multicore architectures. A detailed analysis of performance and accuracy is included. Performance improvements of 14-fold and 4-fold speedups are reported relative to LAPACK and Intel's MKL library.

Key words. divide and conquer, symmetric eigenvalue solver, tile algorithms, dynamic scheduling

AMS subject classifications. 15A18, 65F15, 65F18, 65Y05, 65Y20, 68W10

1. Introduction. The objective of this paper is to introduce a new high performance *Tile Divide and Conquer* (TD&C) eigenvalue solver for dense symmetric matrices on homogeneous multicore architectures. The necessity of calculating eigenvalues emerges from various computational science disciplines e.g., in quantum physics [35], chemistry [38] and mechanics [28] as well as in statistics when computing the principal component analysis of the symmetric covariance matrix. As multicore systems continue to gain ground in the high performance computing community, linear algebra algorithms have to be redesigned or new algorithms have to be developed in order to take advantage of the architectural features brought by these processing units.

In particular, **tile algorithms** have recently shown very promising performance results for solving linear systems of equations on multicore architectures using Cholesky, QR/LQ and LU factorizations available in the PLASMA [3] library and other similar projects like FLAME [46]. The PLASMA concepts consist of splitting the input matrix into square tiles and reorganizing the data within each tile to be contiguous in memory (block data layout) for efficient cache reuse. The whole dataflow execution can then be represented as a directed acyclic graph where nodes are tasks operating on tiles, and edges represent dependencies between them. An efficient and light weighted runtime system environment named QUARK [33] (internally used by the PLASMA library) is exploited to dynamically schedule the different tasks and to ensure that the data dependencies are not violated. As soon as the dependencies are satisfied, QUARK initiates and executes the corresponding tasks on the available computational resources. This engenders an out-of-order execution of tasks which removes unnecessary synchronization points and allows different computational stages to overlap. The authors propose to extend these ideas to the symmetric TD&C eigenvalue solver case using a two-stage approach, in which the symmetric dense matrix is first reduced to band form followed by the TD&C eigenvalue solver directly applied on the symmetric band structure.

*Research reported here was partially supported by the National Science Foundation, Department of Energy and Microsoft Research.

†Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville. (haidar@eecs.utk.edu).

‡KAUST Supercomputing Laboratory, Thuwal, Saudi Arabia. (Hatem.Ltaief@kaust.edu.sa).

§Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville. Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee. School of Mathematics & School of Computer Science, University of Manchester. (dongarra@eecs.utk.edu).

The resulting TD&C symmetric eigenvalue solver algorithm has been extensively evaluated across many matrix types and against similar D&C symmetric eigenvalue solver implementations from state-of-the-art numerical libraries. The performance results show that the proposed TD&C symmetric eigenvalue solver achieves up to 14X speed up compared to the reference LAPACK [5] implementation and up to 4X speed up as compared to the vendor Intel MKL library [1]. Performance results are also reported comparing the TD&C symmetric eigenvalue solver with other standard methods such as the bisection, the QR iteration and MRRR. A study on the accuracy of the computed eigenvalues is provided, which gives a certain confidence on the quality of the overall TD&C symmetric eigenvalue solver framework presented in this paper.

The remainder of this paper is organized as follows: Section 2 recalls some backgrounds on symmetric dense matrix eigenvalue solvers and gives a detailed overview of previous projects. Section 3 provides detailed information on the symmetric D&C eigenvalue solver framework applied on band matrices. Section 4 extends the main concepts behind the standard symmetric band D&C eigenvalue solver to instead use tile algorithms associated with a dynamic scheduler. Section 5 presents the algorithmic complexity of the resulting TD&C symmetric eigenvalue solver. Section 6 presents performance results of the overall algorithm on shared-memory multicore architectures against the corresponding standard D&C routines from LAPACK [5] and the Intel MKL library V10.2 [1] on various matrix types. Also, performance comparisons against other numerical methods are illustrated along with a study on the accuracy of the computed eigenvalues. Finally, Section 7 summarizes the results of this paper and discusses the ongoing work.

2. Background and Related Work. This section recalls the standard algorithm to compute the eigenvalues of a symmetric dense matrix and its inherent bottlenecks. One of the eigenvalue solvers is based on the D&C strategy applied to the final condensed tridiagonal form. This Section also shows how to extend the D&C eigenvalue solver to compute the eigenvalues from a symmetric band structured matrix, which corresponds to one of the main contributions of the paper.

2.1. The Standard Symmetric Eigenvalue Solver Approach. The common way of stating the eigenvalue problem for a symmetric dense matrix is:

$$Ax = \lambda x,\ A \in \mathbb{R}^{n\times n},\ x \in \mathbb{R}^{n},\ \lambda \in \mathbb{R},$$

with A being a symmetric or Hermitian matrix ($A = A^T$ or $A = A^H$), λ – an eigenvalue, and x the corresponding eigenvector. The goal is first to transform the matrix A into a symmetric (or hermitian) tridiagonal matrix S via orthogonal transformations (e.g., Householder reflectors):

$$S = Q \times A \times Q^T,\ A, Q, S \in \mathbb{R}^{n\times n}.$$

Once the tridiagonal form is reached, eigenvalues can then be computed using various methods [27, 43] e.g., the bisection algorithm, the QR method, D&C, MRRR, etc... However, the tridiagonalisation step ($O(n^3)$) is the most time consuming phase in a symmetric eigenvalue solver. It can reach more than 90% of the elapsed time when calculating the eigenvalues ($O(n^2)$) and roughly 50% when both eigenvalues and eigenvectors are to be computed ($O(n^3)$).

The next Section describes the inherent drawbacks of the original aforementioned approach.

2.2. Bottlenecks of the Standard Approach. Because it is the most time consuming, it is very important to identify the bottlenecks of the tridiagonalisation step, as implemented in the state-of-the-art numerical package, namely LAPACK [5]. The block LAPACK algorithms are characterized by two successive computational steps: the panel factorization and

the update of the trailing submatrix. The panel factorization computes the transformations within a specific rectangular region using Level 2 BLAS operations (memory-bound) and accumulates them so that they can be applied into the trailing submatrix using Level 3 BLAS operations (compute-bound). The parallelism in LAPACK resides within the BLAS library, which follows the expensive fork and join model. This produces artifactual synchronization points between the panel factorization and the trailing submatrix update phases.

In particular, the tridiagonal reduction implemented in LAPACK performs successive panel-update sequences. The symmetric matrix is reduced following a one-stage approach where the reduced form is obtained without intermediate steps. The panel factorization requires loading the entire trailing submatrix into memory. As memory is a very scarce resource, this will obviously not scale for large matrices, and thus will generate a tremendous amount of cache and TLB misses. Later, the SBR toolkit [10] introduced a two-stage approach where the matrix is first reduced to a band form during a compute intensive phase and eventually to the final required form through a memory-bound phase (the bulge chasing procedure). The two-stage approach permits the casting of expensive memory operations occurring during the panel factorization into faster compute intensive ones. However, the toolbox still implements block algorithms as in LAPACK, and therefore relies on optimized multithreaded BLAS for parallel performance (based on the fork and join paradigm).

The next Section recalls the principles of the symmetric eigenvalue solver using the D&C strategy, which presents interesting parallel implementation features.

2.3. The Symmetric Eigenvalue Solver using D&C Approach. Introduced by Cuppen [15], the D&C algorithm computes the eigenvalues of the tridiagonal matrix S. Many serial and parallel Cuppen-based eigensolver implementations for shared and distributed memory have been proposed in the past [21, 24, 29, 30, 39, 40, 42]. The D&C approach can then be expressed in three phases: a) partition the tridiagonal matrix T into several subproblems in a recursive manner, b) compute the eigen decompositions of each subproblem, which are defined by a rank-one modification of a diagonal matrix and c) merge the subproblems and proceed to the next level with a bottom-up fashion. In fact, each rank-one modification consists of computing the eigen system of a matrix of the form

$$\tilde{Q}\tilde{D}\tilde{Q}^T = (D + \sigma uu^T), \tag{2.1}$$

where D is a real $n \times n$ diagonal matrix, σ is a nonzero scalar, and u is a real vector of order n and has Euclidean norm equal to 1. We seek a formula or an eigen pair of the system for the matrix on the right hand side of (2.1). The eigen pair can be computed based on solving what we call the "secular equation" which restates results of [11, 25, 41, 45], where we refer the reader for more details on the proof. The overall method amounts to constructing a binary tree with each node representing a rank-one tear and hence, a partition into two subproblems. Each of these problems may be considered as an independent problem without any data dependencies with the other nodes of the same level of the tree. There are two *simple* eigenvalue problems to be solved at each node of the tree. The tree is then traversed in reverse order where at each node the results are merged from the left and the right son calculations. More details on the high-level ideas behind the standard D&C algorithm can be found in [15, 21, 24, 29]. The D&C approach is *sequentially* one of the fastest methods currently available if all eigenvalues and eigenvectors are to be computed [16]. It also has attractive parallelization properties as shown in [42]. Finally, it is noteworthy to mention the *deflation* process, which occurs during the computation of the low rank modifications. It consists of preventing the computation of an eigenpair of a submatrix (matrix of a son in the tree) that is acceptable to be an eigenpair of the larger submatrix (matrix of a father in the tree). Therefore, the greater the amount of deflations, the lesser the number of required operations,

which leads to better performance. The amount of deflations depends on the eigenvalue distribution as well as the structure of the eigenvectors. In practice, most of the application matrices arising from engineering areas provide a reasonable amount of deflations, and so the D&C algorithm runs at less than $\mathcal{O}(n^{2.5})$ instead of $\mathcal{O}(n^3)$.

The next Section presents previous works on how to extend the original symmetric D&C eigenvalue solver to start from a symmetric band matrix.

2.4. Previous Work. The standard D&C eigenvalue solver for symmetric matrices can be extended to start from a symmetric band matrix form. Moreover, band matrices naturally arise in many areas such as the electronic simulations in quantum mechanics, vibrating systems approximated by finite elements or splines, and also in the block Lanczos algorithm where a sequence of increasingly larger band symmetric eigenvalue problems are generated. Besides the overhead of further reducing the symmetric band matrix to the final tridiagonal form, this motivates attempts to compute the eigen decomposition directly from the band form using variants of the standard D&C algorithm.

Arbenz et al. [6, 7, 8] investigated a generalized D&C approach for computing the eigenpairs of symmetric band matrices. The authors provide many important theoretical results concerning the eigen analysis of a low rank modification of a diagonal matrix. Arbenz [6], proposed two methods for computing eigenpairs of a rank-b modification of a diagonal matrix. The first approach consists of computing the rank b modification as a sequence of rank-one modifications. The second approach lies in compressing the rank b modification to a small $b \times b$ eigen problem, solving it and then reconstructing the solution of the original problem [7, 8]. The first approach requires more floating point operations than the second one, which has serial complexity in the $\mathcal{O}(n^3)$ term. The author opted for the second approach. Unfortunately, numerical instabilities in the computation of the eigenvectors have been observed [6], and currently no numerically stable implementation of the second approach exists.

Also, Gansterer et al. [22] developed a D&C algorithm for band symmetric matrices which computes the eigen decomposition. Their approach is based on the separation of the eigenvalue and the eigenvector computations. The eigenvalues are computed recursively by a sequence of rank-one modifications of the standard D&C technique. Once the eigenvalues are known, the corresponding eigenvectors are computed using modified QR factorizations with restricted column pivoting.

Later, Gansterer et al. [23] and Ward et al. [44] proposed another alternative for the generalized D&C algorithm, which computes *approximate* eigenvalues and eigenvectors of symmetric block tridiagonal matrices. Bai and Ward [9] proposed a parallel and distributed version of it. Their algorithm calculates all eigenvalues and eigenvectors of a block-tridiagonal matrix to reduced accuracy by approximating the rank deficient off-diagonal blocks of lower magnitudes with rank-one matrices.

The tile D&C symmetric eigenvalue solver algorithm (TD&C) presented in this paper differs from Arbenz's algorithm [7, 8] in the fact that it uses the first approach (sequence of b rank-one updates) for the solution of the low-rank modifications along with tile algorithms. The first approach, based on a sequence of b rank-one updates, has also been used by Gansterer et al [22]. However, our TD&C algorithm differs from [22] in the representation of the subdiagonal blocks, which are not rank deficient. Moreover, our proposed algorithm has been developed in close analogy to the work of Ward et al. [44], although we do not compute an approximation of the subdiagonal blocks. We consider the subdiagonal blocks as having a full dense matrix structure. The entire algorithm can then be expressed using tiles as a basic block. And once the different computational tasks are defined, a dynamic runtime library is employed to efficiently schedule them on the available resources.

The next Section provides more details on how the three computational phases of the standard symmetric eigenvalue solver based on D&C (See Section 2.3), and originally applied on a tridiagonal matrix, can be extended to the symmetric band case.

3. The *Standard* Band D&C Symmetric Eigenvalue Solver. We assume in this Section that the symmetric matrix is originally in band form or has been reduced to band structure by a pre-processing step. We then describe the D&C methodology applied on a symmetric band matrix A, which is divided into p parts. For simplicity purposes, we define $p = 2$, but it is easy to generalize for any $p < n$, with n being the matrix size (see Section 3.4).

3.1. Partitioning into Subproblems. Similarly to Cuppen's D&C for tridiagonal matrices, the band matrix A can be divided into p parts. If we split A into two parts, A can be written as follows:

$$A = \begin{pmatrix} B_1 & C_1^T \\ C_1 & B_2 \end{pmatrix}, \tag{3.1}$$

where $B_i \in \mathbb{R}^{p_i \times p_i}$, $p_1 + p_2 = n$, and $C_1 \in \mathbb{R}^{b \times b}$ is the upper triangular block of A, with b the bandwidth size. A can be rewritten in the following form:

$$A = \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + R, \quad where \quad \tilde{B}_i \in \mathbb{R}^{p_i \times p_i}, \quad p_1 + p_2 = n. \tag{3.2}$$

There are several ways to define the matrix R. It can be shown that the complexity to compute the eigen decomposition of A in the assembly phase is proportional to the square of rank(R). It is therefore important to keep rank(R) as low as possible. Arbenz [6] shows that it is possible to have the rank of $R = b$ if one defines:

$$R = \begin{pmatrix} O_{p_1-b} & O & O & O \\ O & M & C_1^T & O \\ O & C_1 & C_1 M^{-1} C_1^T & O \\ O & O & O & O_{p_2-b} \end{pmatrix}, \tag{3.3}$$

where $M = \Theta\Delta\Theta^T$, $\Delta = \Delta^T$ and Δ and Θ are any regular matrices. Gansterer and al. [22] have chosen $M = I$. For this choice, R is in most cases unbalanced, meaning that $\| C_1 M^{-1} C_1^T \|$ can be much smaller or larger than $\| M \| = 1$. However, if we choose $M = (C_1^T C_1)^{1/2}$, we get a matrix which is balanced in that sense, and, if $X\Sigma Y^T$ is the singular value decomposition (SVD) of C_1, then $M = Y\Sigma Y^T$ and $C_1^T M^{-1} C_1 = X\Sigma X^T$. Both matrices are symmetric and have the same norm and condition number as C_1. Muller and Ward [9, 23] choose to "*approximate*" the off-diagonal tile C_1 by lower rank matrices using their singular value decompositions.

Going back to the standard Cuppen's algorithm, we plug in the new notations and A can then be rewritten:

$$A \quad = \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + R \quad = \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + Z\Sigma Z^T, \tag{3.4}$$

where $\tilde{B}_1 = B_1 - Y\Sigma Y^T$, $\tilde{B}_2 = B_2 - X\Sigma X^T$ and $Z = \begin{pmatrix} Y \\ X \end{pmatrix}$. The next step is the computation of the eigen decompositions for each updated block $\tilde{A}_i$.

3.2. Solution of the subproblems. The second phase corresponds to the spectral decomposition of each symmetric diagonal tile ($\tilde{B}_i$). This is done through calling the appropriate LAPACK routine for symmetric matrices. As a result, we can write:

$$\tilde{B}_i = Q_{0i} D_{0i} Q_{0i}^T. \tag{3.5}$$

The diagonal matrices D_{0i} contains the eigenvalues of $\tilde{B}_i$ and the matrices Q_{0i} are the corresponding eigenvectors. Thus, the original matrix A can be rewritten as follows:

$$\begin{aligned} A &= \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + Z\Sigma Z^T \\ &= \begin{pmatrix} Q_{01} & 0 \\ 0 & Q_{02} \end{pmatrix} \left\{ \begin{pmatrix} D_{01} & 0 \\ 0 & D_{02} \end{pmatrix} + U\Sigma U^T \right\} \begin{pmatrix} Q_{01} & 0 \\ 0 & Q_{02} \end{pmatrix}^T, \end{aligned} \tag{3.6}$$

where $U = Q_0^T Z$.

The eigenvalues of A are therefore equal to those of $\left\{ \begin{pmatrix} D_{01} & 0 \\ 0 & D_{02} \end{pmatrix} + U\Sigma U^T \right\}$. This is equivalent to $(D + UU^T)y = \lambda y$, where $U \in \mathbb{R}^{n \times b}$ is a matrix of maximal rank $b \leq n$.

Having computed the local spectral decompositions of each tree leaf, the global spectral decomposition can then be calculated by merging the solution subproblem as described in the next subsection.

3.3. Amalgamation of the subproblems. The amalgamation phase consists of traversing the tree in a bottom-up fashion where at each node the results are merged from the left son and the right son calculations. The merging step handles the computation of the spectral decomposition of $(D + UU^T)y = \lambda y$. Here again, it is obvious that all the merging operations, which perform a *rank-b* updating process between two adjacent eigenvalue subproblems, can concurrently run within the corresponding levels.

There are different methods for computing the eigen decomposition of a low rank modification of an Hermitian matrix $(D + UU^T)y = \lambda y$. Two main strategies can be distinguished:
– "$b \times 1$" approach that computes a sequence of b rank-one modifications of a diagonal matrix.
– "$1 \times b$" approach that computes a rank-b modification of a diagonal matrix, Arbenz [6] followed the "$1 \times b$" approach to compute the eigen pairs of a rank-b modification by transforming the problem into a $b \times b$ smaller eigen problem. Unfortunately, numerical instabilities in the computation of the eigenvectors have been observed [6]. The orthogonality of the eigenvectors may be lost if close eigenvalues are present. The resulting algorithm can lose two to four decimal digits of accuracy compared to the LAPACK library.

As previously suggested by Gansterer et al. [22], one can take advantage of the stability of the existing methodology of the standard D&C algorithm to follow the first approach and compute a sequence of b rank-one modifications for calculating the eigenvalues of $(D + UU^T)y = \lambda y$.

The problem to solve now is how to merge the right and the left subproblems. As mentioned above, each merging step consists of a sequence of rank-one updates.

$$A = \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + Z\Sigma Z^T = \begin{pmatrix} \tilde{B}_1 & 0 \\ 0 & \tilde{B}_2 \end{pmatrix} + \sigma_1 z_1 z_1^T + \sigma_2 z_2 z_2^T + \cdots + \sigma_b z_b z_b^T. \tag{3.7}$$

Substituting (3.5) into (3.7) yields:

$$\begin{aligned} A &= \begin{pmatrix} Q_{01} & 0 \\ 0 & Q_{02} \end{pmatrix} \left\{ \begin{pmatrix} D_{01} & 0 \\ 0 & D_{02} \end{pmatrix} \right\} \begin{pmatrix} Q_{01} & 0 \\ 0 & Q_{02} \end{pmatrix}^T + \sigma_1 z_1 z_1^T + \sigma_2 z_2 z_2^T + \cdots + \sigma_b z_b z_b^T \\ &= Q_0 \{D_0\} Q_0^T + \sigma_1 z_1 z_1^T + \sigma_2 z_2 z_2^T + \cdots + \sigma_b z_b z_b^T. \end{aligned} \tag{3.8}$$

Then, the rank-b modification will be successively computed as a sequence of b rank-one updates.

First rank-one update:

$$\begin{aligned} A &= Q_0 \, \{D_0 + \sigma_1 u_1 u_1^T\} \, Q_0^T \; + \sigma_2 z_2 z_2^T + \cdots + \sigma_b z_b z_b^T \\ &= Q_0 \, \{Q_1 \; D_1 \; Q_1^T\} \, Q_0^T \; + \sigma_2 z_2 z_2^T + \cdots + \sigma_b z_b z_b^T, \end{aligned} \tag{3.9}$$

where $u_1 = Q_0^T z_1$, D_1 is the rank-one updated eigenvalues of D_0 and Q_1 contains the eigenvectors resulting from this rank-one update.

Second rank-one update:

$$\begin{aligned} A &= Q_0 \, Q_1 \, \{D_1 + \sigma_2 u_2 u_2^T\} \, Q_1^T \; Q_0^T \; + \cdots + \sigma_b z_b z_b^T \\ &= Q_0 \, Q_1 \, \{Q_2 \; D_2 \; Q_2^T\} \, Q_1^T \; Q_0^T \; + \cdots + \sigma_b z_b z_b^T, \end{aligned} \tag{3.10}$$

where $u_2 = Q_1^T \; Q_0^T z_2$, D_2 is the rank-one updated eigenvalues of D_1 and Q_2 contains the eigenvectors resulting of this rank-one updates. This process of rank-one updates is successively applied to all the z_i rank-one vectors $1 \leq i \leq b$.

bth rank-one update:

$$\begin{aligned} A &= Q_0 \, Q_1 \cdots Q_{b-2} Q_{b-1} \, \{D\} \, Q_{b-1}^T Q_{b-2}^T \cdots Q_1^T \; Q_0^T \; + \sigma_b z_b z_b^T \\ &= Q_0 \, Q_1 \cdots Q_{b-2} Q_{b-1} \left\{ D + \sigma_b u_b u_b^T \right\} Q_{b-1}^T Q_{b-2}^T \cdots Q_1^T \; Q_0^T \\ &= Q_0 \, Q_1 \cdots Q_{b-2} Q_{b-1} Q_b \, \{\Lambda\} \, Q_b^T Q_{b-1}^T Q_{b-2}^T \cdots Q_1^T \; Q_0^T, \end{aligned} \tag{3.11}$$

where $u_b = Q_{b-1}^T \; Q_{b-2}^T \; \cdots \; Q_1^T \; Q_0^T \; z_b$ and $\Lambda = diag\{\lambda_1, \lambda_2, \cdots, \lambda_n\}$ are the eigenvalues of A.

As a result, once the sequence of b rank-one updates has been computed, we obtain $\Lambda = diag\{\lambda_1, \lambda_2, \cdots, \lambda_n\}$, which corresponds to the actual eigenvalues of the matrix A.

The whole algorithm for $p > 2$ can be easily described in the following Section.

3.4. General Case. The band matrix A can be divided into p parts.

$$A = \begin{pmatrix} B_1 & C_1^T & & & \\ C_1 & B_2 & C_2^T & & \\ & C_2 & \ddots & \ddots & \\ & & \ddots & B_{p-1} & C_{p-1}^T \\ & & & C_{p-1} & B_p \end{pmatrix} \in \mathbb{R}^{n \times n}, \tag{3.12}$$

A can be rewritten as follows:

$$A = diag(\tilde{B}_i) + \sum_{i=1}^{p-1} Z_i \Sigma_i Z_i^T, \tag{3.13}$$

where $X_i\Sigma_i Y_i^T$ is defined as the SVD of C_i, $\tilde{B}_1 = B_1 - Y_1\Sigma_1 Y_1^T$, $\tilde{B}_i = B_i - Y_{i-1}\Sigma_{i-1}Y_{i-1}^T - X_i\Sigma_i X_i^T$, for $2 \le i \le p-1$, $\tilde{B}_p = B_p - X_{p-1}\Sigma_{p-1}X_{p-1}^T$, and

$$Z_1 = \begin{pmatrix} Y_1 \\ X_1 \\ 0 \\ 0 \end{pmatrix}, Z_i = \begin{pmatrix} 0 \\ Y_i \\ X_i \\ 0 \end{pmatrix}, \quad for \quad 2 \le i \le p-2, \text{ and } Z_{p-1} = \begin{pmatrix} 0 \\ 0 \\ Y_{p-1} \\ X_{p-1} \end{pmatrix}.$$

In this case, the binary tree has a depth $k = log_2(p)$ if p is a power of two; $k = floor(log_2(p)) + 1$ otherwise. The nodes at the bottom of the tree (level=k) are the matrices $\tilde{B}_i$. As mentioned above, the tree is then traversed in a bottom-up fashion, where at each level l, 2^{l-1} independent merging problems are computed. The desired eigenvalues of A are finally computed at the top level $l = 1$ i.e., the root of the tree.

4. The *Tile* D&C Symmetric Eigenvalue Solver. This section highlights the skeleton of the Tile D&C eigenvalue solver for dense symmetric matrices after recalling the main ideas behind tile algorithms. The dense matrix is first reduced to band form with a given bandwidth b. The three stages of the band D&C symmetric eigenvalue solver (see Section 3) are then developed using tile algorithms to calculate all eigenvalues, starting from the symmetric band matrix.

4.1. Concepts of Tile Algorithms. Tile algorithms have shown lots of interests in the recent years, especially for solving linear systems of equations [4, 12, 31, 32, 36, 37]. The main idea consists of overcoming, especially, the fork-join bottleneck of block algorithms, described in Section 2.2. The original matrix stored in column-major format is split in a two-dimensional grid of tiles, where the elements are now stored contiguous in memory within each tile (tile data layout), as in Figure 4.1. This may demand a complete reshaping of the standard numerical algorithm. The parallelism is then no longer hidden inside the BLAS routine calls, but rather it is brought to the fore. The whole computation can then be represented as a *Directed Acyclic Graph* (DAG) where nodes represent tasks operating on tiles and edges represent dependencies among them. A dynamic runtime environment system QUARK [33]

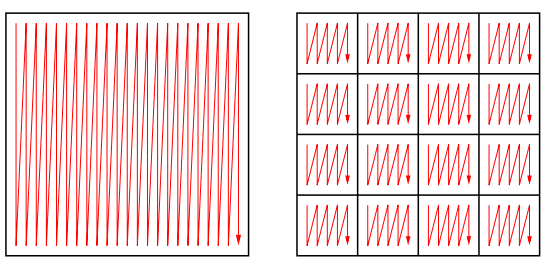

FIG. 4.1. *Translation from LAPACK Layout (column-major) to Tile Data Layout*

is employed to ensure the correct processing of all generated tasks at runtime in a holistic fashion, and is highlighted in the next Section.

4.2. Dynamic Scheduling with QUARK. The main issue when writing a parallel algorithm, and perhaps the most critical, is how to distribute the data and the work across the available processing units in an efficient way. We have integrated the standalone QUARK [33] framework (engine of the PLASMA [3] library) – a dynamic runtime system environment – into our TD&C symmetric eigenvalue solver algorithm. It is worth mentioning that there exist various scheduling frameworks similar to QUARK e.g., SuperMatrix [14], part of the FLAME [46], library or SMPSs [2] as a general scheduling library. Our scheduler is designed

```
int QUARK_core_dgeqrt( Quark *quark, int n,
                       double *A, int lda, double *T, int ldt, int *info )
{
  QUARK_Insert_Task( quark, TASK_core_dgeqrt, 0x00,
    sizeof(int),          &n,          VALUE,
    sizeof(double)*n*n,   A,           INOUT | LOCALITY,
    sizeof(int),          &lda,        VALUE,
    sizeof(double)*n*n,   T,           OUTOUT,
    sizeof(int),          &ldt,        VALUE,
    sizeof(int),          info,        OUTPUT,
    0);
}
void TASK_core_dgeqrt(Quark *quark)
{
  int n; double *A; int lda;  int *T; int ldt; int *info;
  quark_unpack_args_6( quark, n, A, lda, T, ldt, info );
  dgeqrt_( &n, A, &lda, T, &ldt, info );
}
```

FIG. 4.2. *Example of inserting and executing a task in the scheduler. The QUARK_core_dgeqrt routine inserts a task into the scheduler, passing it the sizes and pointers of arguments and their usage (INPUT, OUTPUT, INOUT, VALUE). Later, when the dependencies are satisfied and the task is ready to execute, the TASK_core_dgeqrt routine unpacks the arguments from the scheduler and calls the actual dgeqrt factorization routine.*

to use sequential nested-loop code. This is intended to make it easier for algorithm designers to experiment with algorithms and design new algorithms. Each of the original kernel calls (i.e., tasks) is substituted by a call to a wrapper that decorates the arguments with their sizes and their usage (INPUT, OUTPUT, INOUT, VALUE). As an example, in Figure 4.2 we can see how the DGEQRT call (QR factorization) is decorated for the scheduler. The tasks are inserted into the scheduler, which stores them to be executed when all the dependencies are satisfied. That is, a task is ready to be executed when all parent tasks have completed. The execution of ready tasks is handled by worker threads that simply wait for tasks to become ready and execute them using a combination of default task assignments and work stealing. The thread doing the task insertion is referred to as the master thread. Under certain circumstances, the master thread will also execute computational tasks. Figure 4.3 provides an idealized overview of the architecture of the dynamic scheduler. Moreover, QUARK does

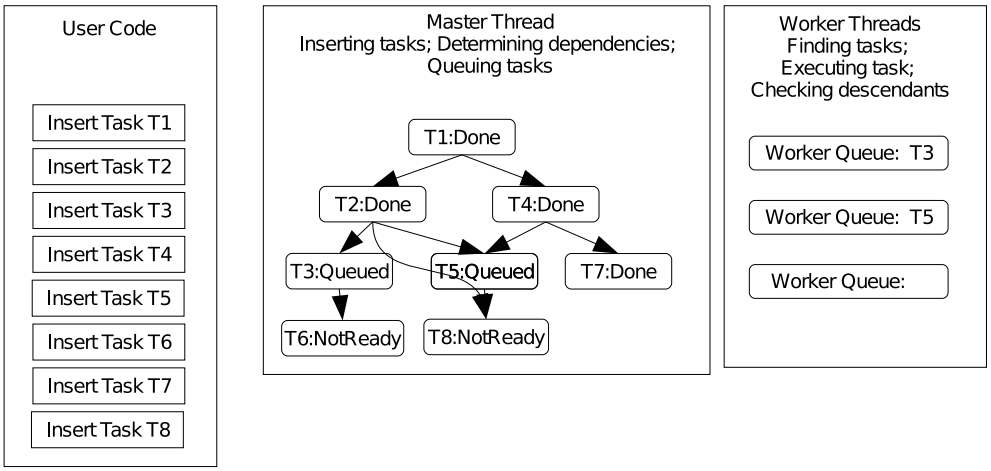

FIG. 4.3. *Idealized architecture diagram for the dynamic scheduler. Inserted tasks go into a (implicit) DAG based on their dependencies. Tasks can be in NotReady, Queued or Done states. Workers execute queued tasks and then determine if any descendants have now become ready and can be queued.*

not explicitly build the DAG prior to the execution for scalability purposes, but rather unrolls it on the fly within a parametrized window of tasks. Therefore, the execution flow is solely driven by the data dependencies. By appropriately hinting QUARK, we are able to resolve two interrelated issues which come into play with regard to scheduling optimality: data locality and the pursuit of the critical path (possibly lookahead opportunities). There are many other details about the internals of the scheduler including its dependency analysis, memory management, and other performance enhancements that are not covered here. However, information about this scheduler can be found in [33]. The authors would like to use the QUARK's highly productive features, along with applying the concepts of tile algorithms to tackle the challenging two-sided reductions, and in particular, the symmetric *tile* D&C eigenvalue solver. To our knowledge, this is the first time a dynamic runtime system coupled with tile algorithms has been employed to solve one of the two-sided reductions on multicore architecture.

The next Section describes the symmetric matrix reduction to band form using tile algorithms, which is the first step toward computing all eigenvalues of a symmetric dense matrix.

4.3. First Step: Symmetric Matrix Reduction to Band Form. This Section describes the numerical kernels involved in the reduction to symmetric band form and highlights the grouping technique used to further enhance data locality during this first step while achieving a small bandwidth size.

4.3.1. Description of the Numerical Kernels. The matrix reduction to symmetric band form follows the tile algorithm strategy and relies on highly optimized compute-intensive kernels to achieve high performance. It is composed of eight kernels total. Four kernels come directly from the one-sided QR factorization [13], and the four others have been recently implemented to handle the symmetry property of the matrix when updating the trailing submatrix around the diagonal. Figure 4.4 highlights the execution breakdown at the second step of the reduction to band form. Since the matrix is symmetric, only the lower part is referenced, and the upper part (gray color) stays untouched. We recall that xGEQRT computes a QR factorization of a single sub-diagonal tile, as presented in Figure 4.4(a). The left and right applications of a Householder reflector block on a symmetric diagonal tile is done by the new xSYRFB kernel, as shown in Figure 4.4(b). The right applications on the single tiles then proceed along the same tile column using the xORMQR kernel, as depicted in Figure 4.4(c). Figure 4.4(d) shows how xTSQRT computes a QR factorization of a matrix composed by the sub-diagonal tile, and a square tile located below it, on the same tile column. Once the Householder reflectors have been calculated, they need to be applied to the trailing submatrix. We developed a new kernel to appropriately handle the symmetric property of the trailing matrix. Indeed, the xTSMQRLR kernel in Figure 4.4(e) loads three tiles together – two of them are symmetric diagonal tiles – and carefully applies left and right orthogonal transformations on them. Once the special treatment for the symmetric structure has completed, we then apply the same transformations to the left side of the matrix as well as to the right side (down to the bottom of the matrix) using the xTSMQR kernel with the corresponding side variants, as displayed in Figure 4.4(f) and 4.4(g), respectively. It is necessary to further explain the left variant, as it should require the light gray tile located in position (3,4) from Figure 4.4(f). In fact, this tile is not referenced since only the lower part of the matrix is being operated. Therefore, by taking the transpose of the tile located in position (4,3) from the same Figure, we are able to compute the right operations. Finally, Figure 4.4(h) represents the matrix structure at the end of the second step of the reduction. The symmetric band structure starts to appear at the top left corner of the matrix (the dashed area). Since the entire algorithm operates on tiles, the tile size b naturally corresponds to the bandwidth size. This matching between tile and bandwidth sizes is in fact more than just a choice for productivity purposes. Indeed, it is

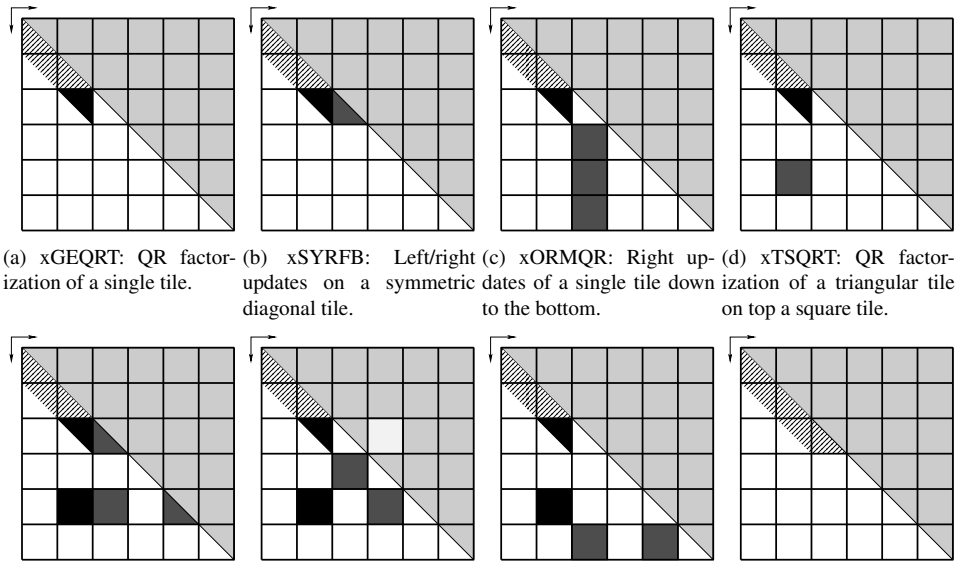

(a) xGEQRT: QR factorization of a single tile.

(b) xSYRFB: Left/right symmetric updates on a diagonal tile.

(c) xORMQR: Right updates of a single tile down to the bottom.

(d) xTSQRT: QR factorization of a triangular tile on top a square tile.

(e) xTSMQRLR: Left-/right applications on the diagonal region.

(f) xTSMQR: Left application using the transpose of the light gray upper tile.

(g) xTSMQR: Right application down to the bottom.

(h) Matrix structure at the end of the reduction step 2.

FIG. 4.4. *Kernel execution breakdown of the TD&C algorithm during the first stage.*

an algorithmic restriction to avoid destroying the obtained zeroed structure while computing the symmetric band form. For instance, in the lower symmetric case, the left updates have to be shifted one tile down (i.e., the actual bandwidth size) from the diagonal tile such that the right ones do not produce fill-in elements in the already annihilated part of the matrix.

Also, the parameter b is critical to be tuned, as explained in the next Section.

4.3.2. Enhancing Data Locality by Grouping Tasks. A small bandwidth b will dramatically affect the efficiency of the reduction phase and increase the elapsed time due to (1) the huge bus traffic required to load all the kernel data into memory, (2) the inefficiency of the kernels on small b, and (3) the runtime system overhead of scheduling very fine granularity tasks. At the same time, as analyzed later in Section 5, the number of floating point operations of the TD&C symmetric eigenvalue solver strongly depends on the band size b. The smaller the b, the less the number of operations. This trade-off between achieving high performance for the reduction step and diminishing the number of flops of the TD&C symmetric eigenvalue solver has then to be addressed. To overcome those three limitations, we have implemented a *task grouping technique*, which consists of aggregating different data tiles as well as the computational kernels operating on them, in order to build a *super* tile. The idea is that the kernels operate on a *super* tile overcoming limitations (1) and (2), also increasing data reuse where we are forced to apply all the possible operations that could be applied to the *super* tile. Moreover, such a technique allows us to decrease the number of tasks, as all the operations occurring on a *super* tile are considered as only one task. This later helps us to sweep over the limitation (3) by removing the scheduler overhead. Figure 4.5 shows the super tiles of size 2. The kernels operating on those super tiles are actually a combination of the kernels used without the grouping technique. For instance, in Figure 4.5, to enhance the data locality, the single call to the update kernel on the dark gray super tiles will be internally decoupled by multiple calls to the corresponding kernel variants (i.e., left update, followed by

a left transposed and right updates). For instance, the elapsed time for reducing a symmetric dense matrix ($n > 4000$) to band form with $b = 20$ on an Intel Xeon system using 16 cores can be reduced up to one order of magnitude by enabling the task grouping technique, and the number of tasks can decrease up to two orders of magnitude compared to the original version i.e., without the task grouping technique. Therefore, applying the task grouping technique permits us to substantially improve the reduction phase, and at the same time, to obtain a small bandwidth size required for the second step to be effective.

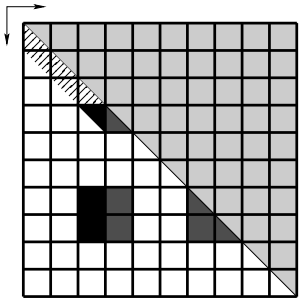

FIG. 4.5. *A schematic description of the task grouping technique using super tiles of size* 2.

Once the band form is obtained, the goal is to compute the eigenvalues directly from the band form using the TD&C symmetric eigenvalue solver, which is the main topic of next Section.

4.4. Second Step: Computing All Eigenvalues from the Band Form using TD&C. Applying the D&C algorithm on the symmetric band form using tile algorithms turns out to be straightforward, thanks to the high level of productivity provided by the QUARK framework. The partitioning into subproblems consists of computing the SVDs of each off-diagonal tile, which are independent from each other and can therefore run concurrently. Once the SVDs are computed, the symmetric diagonal tiles are then updated accordingly, and the solutions of each subproblem can be triggered by calculating the spectral decompositions. Here again, the spectral decompositions are completely independent and can run in parallel. The computational routines used for the two phases are DGESVD and DSYEVD from LAPACK, respectively. The last phase consists of the amalgamation of the subproblems in which the contributions of a pair of subproblems are merged together (b rank-one modifications) following a reverse tree traversal. A customized kernel has been implemented to merge the different subproblems, which are independent from each other when located on the same tree level. The three phases can thus be easily mapped into the tile algorithms. It is also noteworthy to mention that a dynamic-type scheduler is preferred over a static one especially, because the work of the TD&C algorithm depends strongly on the amount of deflations, which can produce load balancing issues. Therefore, there is no need to wait the end of the merging procedure of all subproblems within a tree level. The merge of the subsequent subproblems at higher level of the tree can proceed as soon as the previous subproblem amalgamations terminate. Moreover, the fine granularity of tile algorithms, together with the dynamic scheduler, permit the removal of the artifactual synchronization points seen in the standard D&C parallel implementation located (1) inside the fork/join paradigm and (2) between the different phases of the algorithm. Tile algorithms alleviate (1) by bringing the parallelism to the fore and relying on sequential high performance kernels with a smaller granularity. Since the execution flow of the TD&C symmetric eigenvalue solver through QUARK is solely driven by the data dependencies, a task immediately proceeds after the dependencies coming from

its parent are satisfied, regardless of the application global status. Therefore, the unnecessary synchronization points between the different phases of the algorithm (2) are completely removed. Tasks from the reduction phase and from the band D&C phase can substantially overlap. The band D&C phase can in fact start even before the dense matrix has completely achieved the band form.

The next Section details the algorithmic complexity of the overall algorithmic framework.

5. Algorithmic Complexity. In this section, we calculate the number of floating-point operations (flops) required by the tile DC (TD&C) symmetric eigenvalue solver algorithm. We recall that b corresponds to the tile and the bandwidth sizes. The matrix is divided into p blocks/tiles i.e., $p = n/b$ with n being the size of the matrix.

The algorithmic complexity of the full tridiagonal reduction is $\mathcal{O}(4/3n^3)$. The calculation of the number of flops for the band reduction step is then straightforward. It requires $4/3n \times (n-b) \times (n-b)$ flops and gets therefore closer to $\mathcal{O}(4/3n^3)$ for small bandwidth size b. The flop calculation for the band D&C algorithm to get the actual eigenvalues is more complicated. The flop count for each phase is reported below:

- *Partitioning phase:* this phase is characterized by the singular value decomposition of the off-diagonal blocks (estimation bound for small square matrices of size k is $6.67k^3$) and the construction of the new diagonal blocks $\tilde{A}_i$ as defined by Equation (3.4). The number of flops is then at most:
  $(p-1) \times (6.67)b^3 + (p-1) \times 2 \times 2b^3$ flops.
- *Subproblems solution phase:* this phase computes the spectral decompositions of each symmetric $\tilde{B}_i$ as described by Equation (3.5). Its cost depends on the algorithm used and the convergence of the eigenvalues (an upper bound for small square matrices of size k is $9k^3$ [26]), it requires at most:
  $p \times 9b^3$ flops.
- *Amalgamation phase:* each merging step of two-by-two nodes involves the application of a sequence of b rank-one modifications. Each rank-one modification necessitates the computation of a rank-one vector $u_i = Q_i^T z_i$, and an eigen solution of $D + \sigma_i u_i u_i^T$. The computation of the rank-one vector u_i involves a set of matrix-vector products with the computed Q_i^T from the current level (cost detailed in α) and also with all the Q_i^T from the previous deeper levels of the tree (cost detailed in β). The cost of the eigensolution will be reported in γ.
  - $\alpha-$ For the current level, for each rank-one modification "j", as it have been described by (3.11), we will have to apply a matrix-vector product with all the previously computed Q_k^T, (where $k = 1 \ldots j-1$) of the current level. Thus, for the b rank-one modifications, this gives rise to $\frac{b}{2}(b-1)$ matrix-vector products of a matrix of size $(\frac{n}{2^i})$. Note that, for each level, there are 2^i leaves to be amalgamated, giving rise to $2^i \times \frac{b}{2}(b-1) \times 2(\frac{n}{2^i})^2$. Thus, for all the levels, this is expressed by the first component of the formula below.
  - $\beta-$ Similarly, for each rank-one modification, the cost of the matrix-vector product with the Q_i^T from the previous deeper levels of the tree can be expressed as $i \times b$ matrix-vector of a matrix of size $\times 2(\frac{n}{2^i})^2$. Thus for the b rank-one modifications, its $b \times i \times b \times 2(\frac{n}{2^i})^2$. For each level, there are 2^i leaves to be amalgamated giving rise to the second component of the formula below.
  - $\gamma-$ The eigen decomposition of $D + \sigma uu^T \in \mathbb{R}^{k \times k}$ demands two distinct computation steps; the zero finding algorithm, which requires an average of $5k^2$ operations and the calculation of k eigenvectors, which costs $2k^2$ operations. The number of flops involved in the eigen solution of a rank-one modification is

$\mathcal{O}(7n^2)$. In practice, this number is the worst case scenario as the remarkable phenomenon of deflation takes place and may dramatically reduce this cost depending on the number of non-deflated vectors. Finally, for b rank-one modifications, for all the 2^i leaves and for all the levels, the cost is reported in the third component of the formula below.

Now, if we sum up all eigen decomposition computations occurring at all tree levels with a depth $k = \lfloor log_2(p) \rfloor$, the cost of the amalgamation step is equal to:

$$\sum_{i=0}^{k-1} 2^i \times \frac{b}{2}(b-1) \times 2(\frac{n}{2^i})^2 + \sum_{i=1}^{k-1} 2^i \times b \times i \times b \times 2(\frac{n}{2^i})^2 + \sum_{i=0}^{k-1} 2^i \times b \times 7(\frac{n}{2^i})^2$$

$$= \sum_{i=0}^{k-1} \frac{b}{2}(b-1) \times 2n^2(\frac{1}{2})^i + \sum_{i=1}^{k-1} b \times i \times b \times 2n^2(\frac{1}{2})^i + \sum_{i=0}^{k-1} b \times 7n^2(\frac{1}{2})^i$$

$$= \frac{b^2}{2} \times 2n^2 \sum_{i=0}^{k-1}(\frac{1}{2})^i + 2b^2n^2 \sum_{i=1}^{k-1} i(\frac{1}{2})^i + 7bn^2 \sum_{i=0}^{k-1}(\frac{1}{2})^i$$

we know that, $\sum_{i=0}^{\infty} ar^i = \frac{a}{1-r}$ and $\sum_{i=0}^{\infty} ir^i = \frac{r}{(1-r)^2}$ for $|r| < 1$, thus,
$= 2b^2n^2 + 4b^2n^2 + 14bn^2 \approxeq \mathcal{O}(6\, b^2n^2)$ flops.

The total arithmetic complexity of the band D&C algorithm will be dominated by the amalgamation phase and will require at most $\mathcal{O}(6\, b^2n^2)$ flops.

We note that the cost of the TD&C for computing all eigenvalues of a dense symmetric matrix is $\mathcal{O}(4/3n^3 + 6\, b^2n^2)$. It is straightforward to see that the size of the bandwidth b will have a huge impact on the overall algorithm complexity. The TD&C symmetric eigenvalue solver algorithm will be cost-effective if and only if $b \lll n$. Also, if all corresponding eigenvectors are to be calculated, the TD&C symmetric eigenvalue solver will not be suitable as the cost will be dominated by the accumulation of the intermediate eigenvector matrices for each single rank-one modification. This would involve matrix-matrix multiplications and thus, the overall cost will substantially increase and become dominated by $\frac{8}{3}bn^3$ flops. Although these Level 3 BLAS operations are well-suited for hardware accelerators such as GPUs, other methods, such as MRRR [19, 20], may be more competitive and need to be explored if the eigenvectors are required.

The next section presents the parallel performance of the TD&C symmetric eigenvalue solver against the state-of-the-art open-source and commercial numerical libraries.

6. Performance Results and Accuracy Analysis. This section summarizes our main performance results of the TD&C symmetric eigenvalue solver algorithm.

6.1. Machine Description. All of our experiments have been run on a shared-memory multicore architecture composed of a quad-socket quad-core Intel Xeon EMT64 E7340 processor operating at 2.39 GHz. The theoretical peak is equal to 9.6 Gflop/s/ per core or 153.2 Gflop/s for the whole node, composed of 16 cores. The practical peak achieved with DGEMM on a single core is equal to 8.5 Gflop/s or 136 Gflop/s for the 16 cores. The level-1 cache, local to the core, is divided into 32 kB of instruction cache and 32 kB of data cache. Each quad-core processor is actually composed of two dual-core Core2 architectures and the level-2 cache has 2×4 MB per socket (each dual-core shares 4 MB). The machine provides Intel Compilers 11.0 together with the Intel MKL V10.2 vendor library.

6.2. Tuning the Tile/Band Size b. From Sections 4 and 5, it is clear that the tile/band size is the paramount parameter. Figure 6.1 highlights the effect of this parameter on the reduction phase (blue curve), on the band D&C phase (green curve) and on the overall TD&C eigenvalue solver (red curve). For a very large b, the band D&C phase is the dominant part

of the general algorithm and reciprocally, for a very small b, the reduction phase governs the whole application. From these experiments, a tile size $b = 20$ looks to be the best compromise for a matrix of order less than 20000, while $30 < b < 40$ appears to be the best choice for a matrix of larger order.

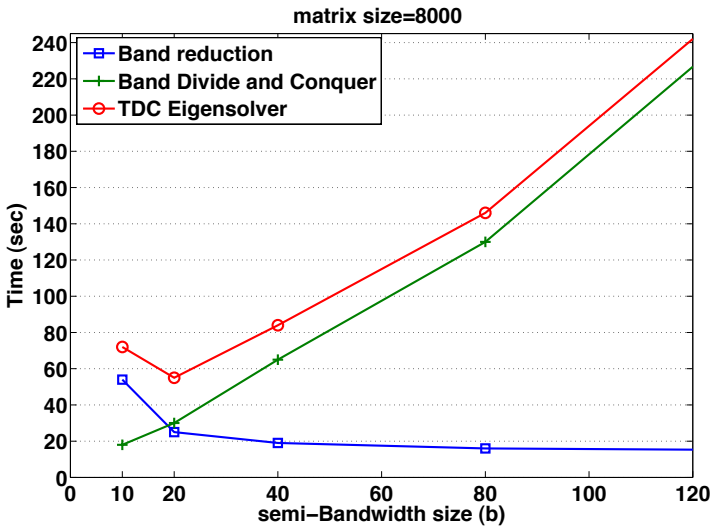

FIG. 6.1. *Effect of the tile/band size.*

6.3. Performance Comparisons of the first stage reduction from dense to band. Figure 6.2(a) highlights the performance of the first stage (PLASMA-DSYRDB) when varying the semi bandwidth size. We compare our first stage procedure against the equivalent function from the SBR toolbox using two different tile sizes i.e., b= 200 and b= 20. Our implementation of the first stage performs very well, thanks to the increased parallelism degree brought by tile algorithms. Although this stage is very compute-intensive, we could easily note that the performance of the first stage is very sensitive to the semi bandwidth size. The perfor-

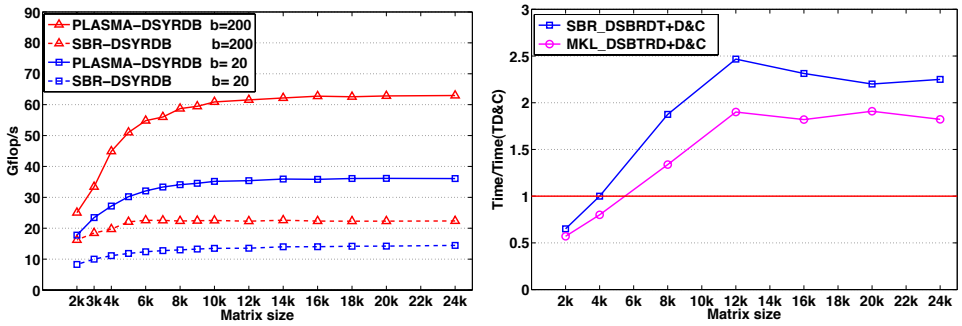

(a) Performance comparisons of the first stage when varying the semi bandwidth size.

(b) Performance comparisons of finding the eigenvalues of band matrices with a bandwidth equal to 20.

FIG. 6.2. *Kernel execution breakdown of the TD&C algorithm during the first stage.*

mance decreases as the tile size gets smaller, however, it remains sustainably acceptable. The equivalent SBR function to reduce the symmetric dense to band form does not perform sim-

ilarly, surprisingly, mainly due to the overhead of the nested-parallelism within the fork-join paradigm.

6.4. Performance Comparisons of the second stage finding the eigenvalues of band matrices. Figure 6.2(b) presents timing comparisons in seconds of the second stage; finding the eigenvalue of band matrices either using our band divide and conquer or using the standard techniques of reducing the band matrices to tridiagonal then finding the eigenvalue of the tridiagonal matrices. The elapsed time of the standard techniques is divided by the elapsed time of our band divide and conquer. We compare against the MKL-DSBTRD and also against the SBR-DSBRDT which reduce a band matrix to tridiagonal. The bandwidth size is equal to 20 which corresponds to our choice of optimal bandwidth.

6.5. Performance Comparisons with other Symmetric D&C Implementations. This section shows the performance results of our TD&C symmetric eigenvalue solver and compares them against the similar routine (DSTEDC) available in the state-of-the-art open source and vendor numerical libraries i.e., multi-threaded reference LAPACK V3.2 with optimized Intel MKL BLAS and Intel MKL V10.2, respectively. Since DSTEDC operates on a tridiagonal matrix, the dense symmetric matrix first needs to be reduced to tridiagonal form. Intel MKL actually provides two interfaces to reduce the symmetric dense matrix to tridiagonal form. The first corresponds to the optimized version of DSYTRD from LAPACK (MKL-LAPACK), and the second one integrates the optimized version of the corresponding routine named DSYRDD from the Successive Band Reduction toolbox (SBR) [10] (MKL-SBR). Figure 6.3 summarizes the experiments in Gflop/s. The number of flops used as a reference is $4/3n^3$, which is roughly the total number of flops for the tridiagonal reduction, since the D&C eigenvalue solver of a tridiagonal matrix is of order n^2. All experiments have been performed on 16 cores, with random matrices, in double precision arithmetic.

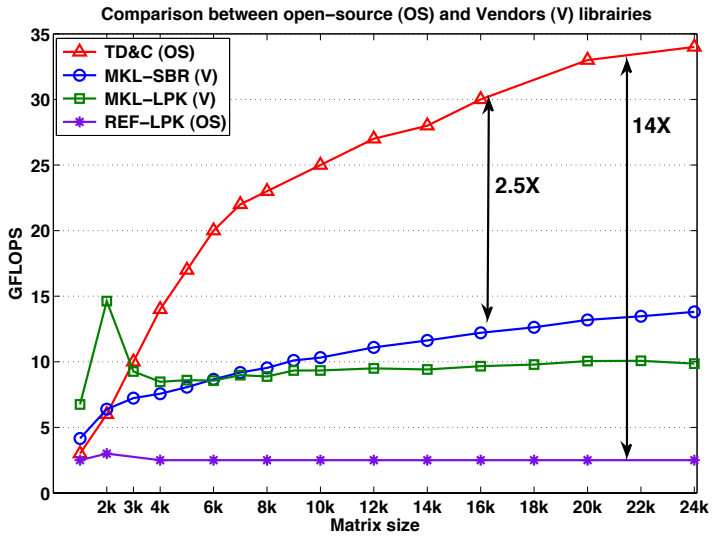

FIG. 6.3. *Performance comparisons of the TD&C symmetric eigenvalue solver against open-source (OS) and vendor (V) libraries.*

The performance results show that the proposed TD&C symmetric eigenvalue solver achieves up to 14X speed up compared to the reference LAPACK implementation, and up to 2.5X speed up as compared to the vendor Intel MKL library on random matrices.

The next sections describe a collection of different matrix types and compare our TD&C symmetric eigenvalue solver against D&C (DSTEDC) as well as other methods to compute the eigenvalues of a symmetric dense matrix e.g., the bisection BI (DSTEVX), QR (DSTEV) and MR (DSTEMR). For simplicity, we compare in the experiments shown in the next sections, our implementation with those four routines from the Intel MKL library, and skip the comparison with the reference LAPACK library.

6.6. Description of the Matrix Collections. In this section, we present the different matrix collections used during the extensive testing to get performance results (Section 6.7) as well as the numerical accuracy analysis (Section 6.8). There are three matrix collections:
– The matrices from the first set represented in Table 6.1, are "*synthetic testing matrices*" chosen to have extreme distributions of eigenvalues as well as other specific properties that exhibit the strengths and weaknesses of a particular algorithm. More information about these matrix collections can be found in [17, 18, 34].

Type	Description
Type 1	$\lambda_1 = 1,\ \ \lambda_i = \frac{1}{k},\ \ i = 2,3,\cdots n$
Type 2	$\lambda_i = 1,\ \ i = 2,3,\cdots n-1,\ \ \lambda_n = \frac{1}{k}$
Type 3	$\lambda_i = k^{-(\frac{i-1}{n-1})},\ \ i = 2,3,\cdots n$
Type 4	$\lambda_i = 1-(\frac{i-1}{n-1})(1-\frac{1}{k}),\ \ i = 2,3,\cdots n$
Type 5	n random numbers in the range $(\frac{1}{k},k)$, their logarithms are uniformly distributed
Type 6	n random numbers from a specified distribution
Type 7	$\lambda_i = ulp \times i,\ \ i = 2,3,\cdots n-1,\ \ \lambda_n = 1$
Type 8	$\lambda_i = ulp,\ \ \lambda_i = 1 + i \times \sqrt{ulp},\ \ i = 2,3,\cdots n-1,\ \ \lambda_n = 2$
Type 9	$\lambda_1 = 1,\ \ \lambda_i = \lambda_{i-1} + 100 \times ulp,\ \ i = 2,3,\cdots n$

TABLE 6.1

Synthetic testing matrices from LAPACK *testing (Type 1-6) and from [18] (Type 7-9). Note that for distributions of Type 1-5, the parameter k has been chosen to be equal to the machine precision ulp in double precision arithmetic. We use the* LAPACK *routines* DLATMS *or* DLAGSY *to generate $A = QDQ^T$. Given the eigenvalues λ, the dense matrix A is generated by multiplying $D = diag(\lambda)$ by an orthogonal matrix Q generated from random entries.*

Type	Description
Type 10	(1,2,1) tridiagonal matrix
Type 11	Wilkinson-type tridiagonal matrix
Type 12	Clement-type tridiagonal matrix
Type 13	Legendre-type tridiagonal matrix
Type 14	Laguerre-type tridiagonal matrix
Type 15	Hermite-type tridiagonal matrix

TABLE 6.2

Matrices with interesting properties.

Type	Description
Type 16	Matrices from an application on quantum chemistry and electronic structure
Type 17	The bcsstruc1 set in the Harwell-Boeing collection
Type 18	Matrices from Alemdar, NASA, and cannizzo, etc. sets at the University of Florida

TABLE 6.3

Matrices from real life applications.

– The second set of matrices are "*matrices with interesting properties*", which are represented in Table 6.2. Those matrices will only be used for the accuracy analysis in Section 6.8

since they are already in tridiagonal form.
– The third set are "*practical matrices*" which are based on a variety of practical applications, and thus are relevant to a large group of users. Some of these matrices are described in Table 6.3.

6.7. Performance Comparisons with other Symmetric Eigenvalue Solvers. This Section compares the TD&C approach against other state-of-the-art symmetric eigenvalue solvers.

6.7.1. Understanding the Various Graphs. All experiments have been performed on 16 cores with matrix types from Table 6.1 and Table 6.3. The matrix sizes vary from 1000 to 24000, and the computation is done in double precision arithmetic. Our TD&C symmetric eigenvalue solver is compared against the bisection BI (DSTEVX), QR (DSTEV), MRRR or MR (DSTEMR), as well as the D&C (DSTEDC) symmetric eigenvalue solvers. Similarly to DSTEDC, the other eigenvalue solver implementations necessitate the matrix being reduced into tridiagonal form using either the MKL LAPACK-implementation routine DSYTRD, or the MKL SBR-implementation routine DSYRDD. There are two types of graphs: (1) *speed up graphs* obtained by computing the ratio between the elapsed time of these four methods over the TD&C total execution time – in other words, a "*value t above* 1" means that our TD&C is "*t times faster*" than the corresponding algorithm, and a "*value t below* 1" means that our TD&C is "$1/t$ *times slower*" than the corresponding algorithm, and (2) *performance graphs* in Gflop/s, which show how much of the peak performance the different symmetric eigenvalue solvers are able to reach. Also, we note that the number of flops used as a **reference is 4/3**n^3. The color and symbol codes used for all plots are as follows: BI curves are green, using "x", QR curves are blue, using "*", MR curves are magenta using "$\diamond$", D&C curves are black, using "o", and TD&C curves are red and are either considered as a reference (speed up graphs) or use "o" to mark data points (performance graphs).

6.7.2. Synthetic testing matrices (Table 6.1). Figure 6.4 shows respectively, the speed up and the performance in Gflop/s obtained by the TD&C symmetric eigenvalue solver as compared to BI, QR, MR and D&C when the MKL LAPACK-implementation routine DSYTRD is used to reduce the first stage.

Figure 6.5 shows respectively, the speed up and the performance in Gflop/s obtained by the TD&C symmetric eigenvalue solver as compared to BI, QR, MR and D&C when the MKL SBR-implementation routine DSYRDD is used to reduce the first stage.

The results show that the TD&C symmetric eigenvalue solver algorithm performs better than all other represented algorithms and for all synthetic testing matrices for $n \geq 3000$. The results also indicate that the TD&C speed up ranges from 2X faster for matrices of size $3000 \leq n \leq 10000$ (27 Gflop/s) to 4X faster for matrices of size $n \geq 10000$ (35 Gflop/s). For small matrix sizes with $n \leq 3000$, TD&C runs slower than the other eigenvalue solvers. It is clear, from the algorithmic complexity study in Section 5, that our TD&C algorithm requires $4/3n^3 + 6b^2n^2$ flops, which for example, for a $b = 20$ and $n = 3000$ means $4/3n^3 + 2400n^2 \approxeq 2n^3$ flops, while the other algorithms roughly perform $4/3n^3$ flops.

On the other hand, BI seems to be the slowest algorithm for computing eigenvalues although its efficiency mostly depends on the distribution of the eigenvalues. Indeed, for matrices with strongly clustered eigenvalues (e.g., type 2), BI performance is comparable with the other symmetric eigenvalue solver methods. Furthermore, the TD&C symmetric eigenvalue solver is able to take full advantage of a situation where a significant amount of deflations occur for special matrices of type 1 and type 2. By avoiding unnecessary calculations, especially in the amalgamation phase thanks to deflations, the TD&C runs even at higher performance. This is also emphasized when comparing matrices from type 2 and type 6. While a type 6 matrix represents the worst case scenario for the TD&C symmetric

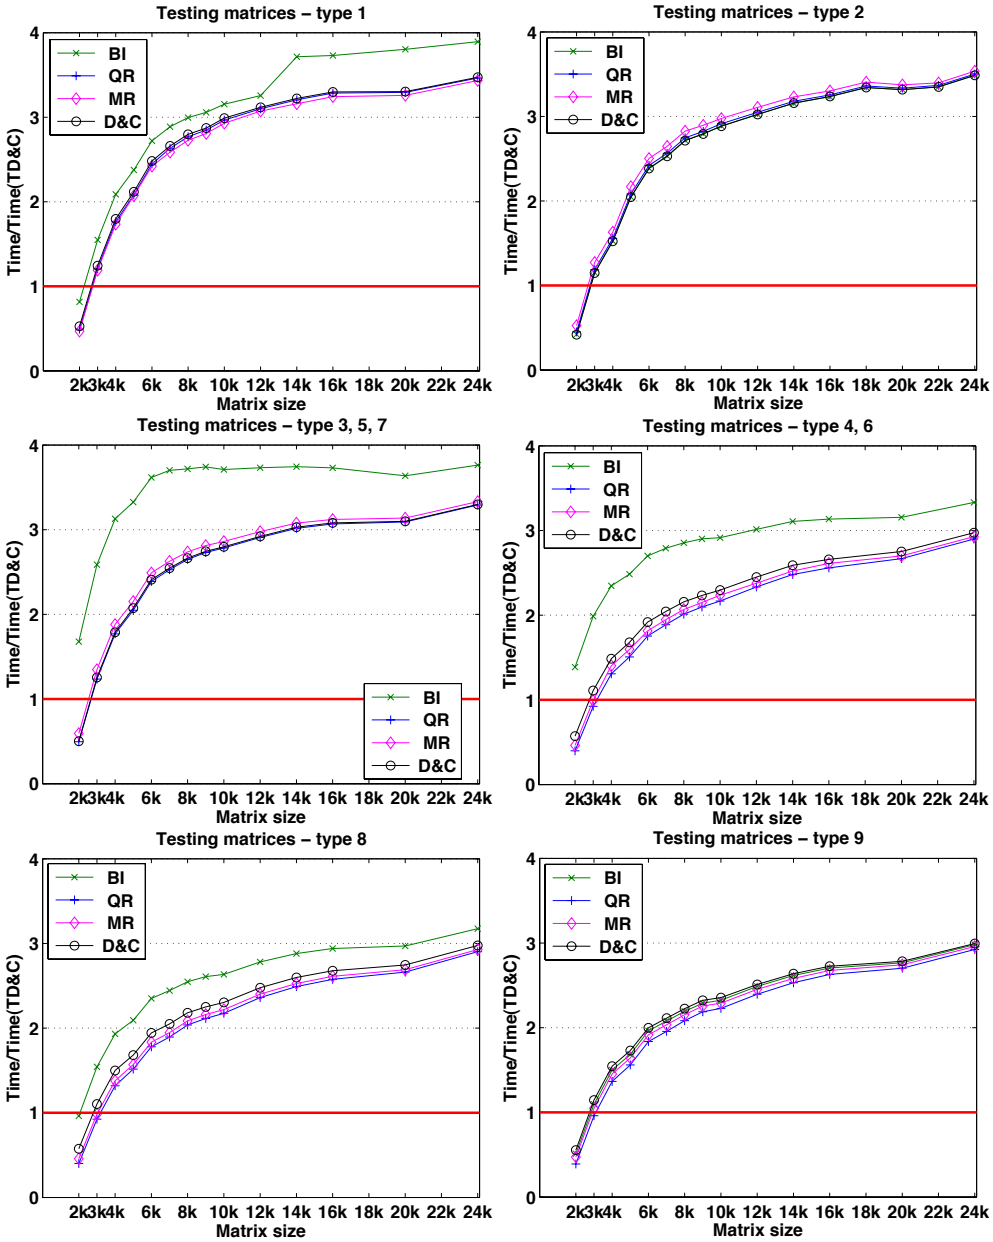

FIG. 6.4. *TD&C Speed up as compared to BI, QR, MR, and D&C for all synthetic testing matrices from Table 6.1. For BI, QR, MR, and D&C, the MKL LAPACK-implementation routine DSYTRD is used to transform the dense symmetric matrix A to tridiagonal.*

eigenvalue solver in which less than 2% of deflations happen, a type 2 matrix corresponds to a matrix where a considerable amount of deflation occurs.

The TD&C considerably performs less flops in the latter case, and thus, it runs more than twice as fast as when no deflations take place. For example, for a matrix of size $n = 16000$, the reduction to band form requires 145 seconds. The band D&C phase with lots of deflations requires 30 seconds instead of 74 seconds when no deflations occur, giving a total execution time of 175 seconds and 219 seconds , respectively (1.3X overall speed up). From a global point of view, the elapsed time and the performance results of the other methods (BI, QR, MR and D&C) are very similar to one another. In addition to that, in the presence of deflation, the improvement is easily noticed for the TD&C algorithm, but it is rather moderate for the standard D&C algorithm. Therefore, once again, it is important to mention that when

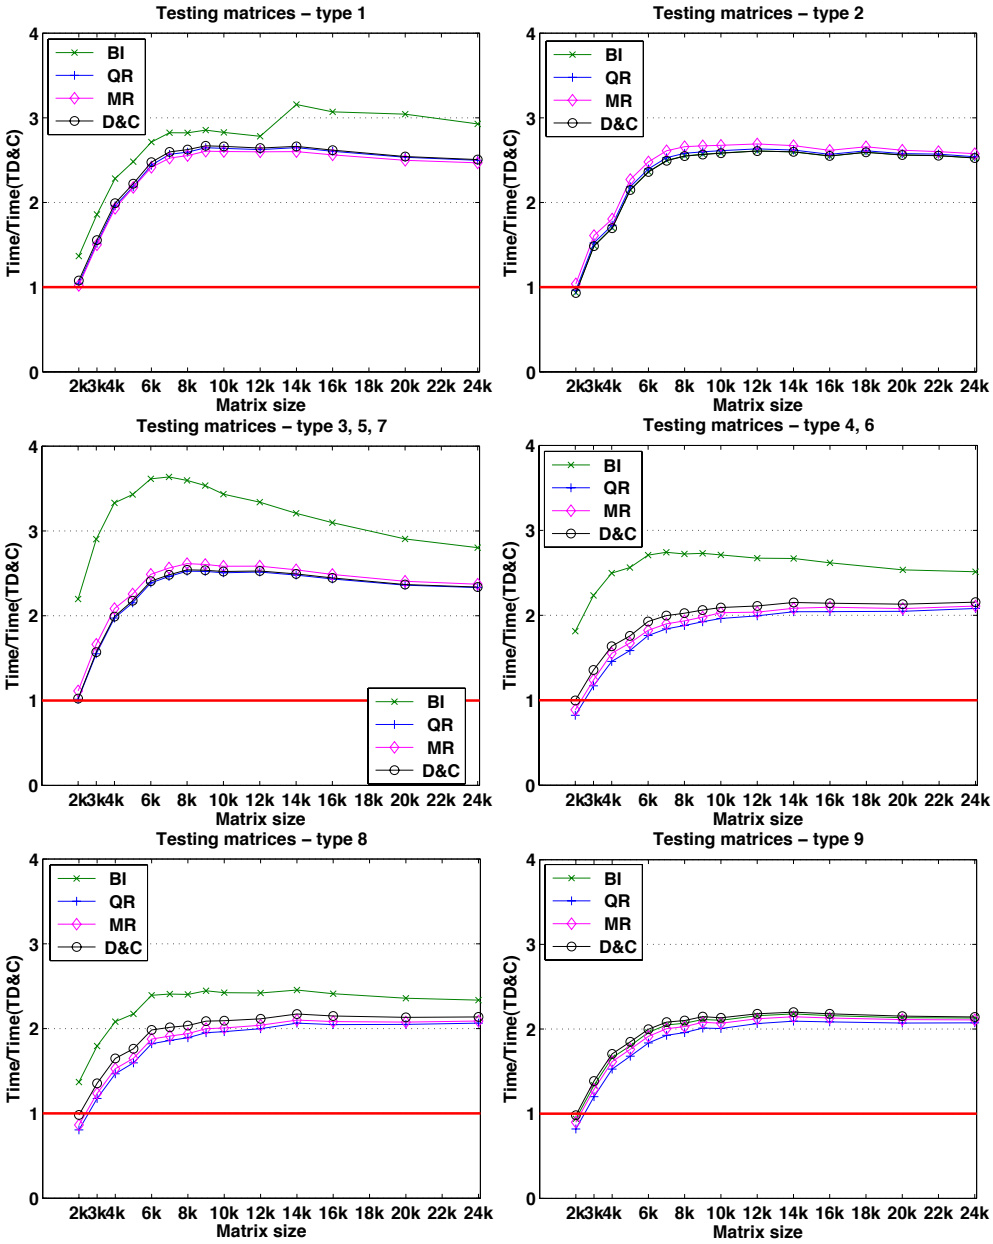

FIG. 6.5. *TD&C Speed up as compared to BI, QR, MR, and D&C for all synthetic testing matrices from Table 6.1. For BI, QR, MR, and D&C, the MKL SBR-implementation routine DSYRDD is used to transform the dense symmetric matrix A to tridiagonal.*

only eigenvalues are computed, BI, QR, MR, and D&C count only for $O(n^2)$ on the overall algorithm complexity and the leading term corresponding to the tridiagonal reduction phase is about $4/3n^3$. The time spent in this reduction phase can be as high as 90% of the total execution time when only eigenvalues are needed but more than 50% when eigenvectors are to be computed. Thus, no matter how D&C or the other methods get improved, the impact on the overall algorithm performance is rather negligible for large matrix sizes.

6.7.3. Practical matrices (Table 6.3). The most important testing matrices are perhaps those which come from real applications. Figure 6.6 shows the speed up and the performance in Gflop/s obtained by the TD&C symmetric eigenvalue solver as compared to BI, QR, MR and D&C for a set of more than 100 matrices arising in different scientific and engineering

areas. Similar to synthetic testing matrices, the TD&C symmetric eigenvalue solver algorithm out-performs all the other symmetric eigenvalue solvers, between two to three times faster, and achieves up to 36 Gflop/s.

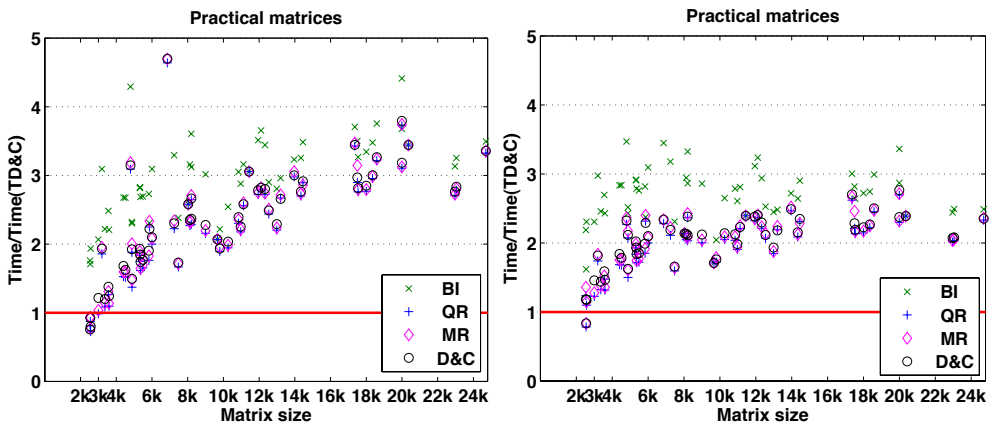

FIG. 6.6. *Timings comparison for practical matrices. The tridiagonalisation has been done either by the MKL-LAPACK-DSYTRD (left) or by the MKL-SBR-DSYRDD (right).*

6.8. Accuracy Analysis. This section is dedicated to the analysis of the TD&C symmetric eigenvalue solver accuracy as compared to the other four symmetric eigenvalue solvers: QR, BI, MR and D&C.

6.8.1. Metric Definitions. For a given symmetric matrix B, computed eigenvectors $Q = [q_1, q_2, \cdots, q_n]$ and their corresponding eigenvalues $\Lambda = diag(\lambda_1, \lambda_2, \cdots, \lambda_n)$, we use the following accuracy tests by using the LAPACK testing routines (DLANSY, DSTT21 for tridiagonal and DSYT21 for dense symmetric):

$$\frac{\|I - QQ^T\|}{n \times ulp} \tag{6.1}$$

which measures the orthogonality of the computed eigenvectors, and

$$\frac{\|B - Q\Lambda Q^T\|}{\|B\| n \times ulp} \tag{6.2}$$

which measures the accuracy of the computed eigenpairs, and

$$\frac{\|\lambda_i - \delta_i\|}{\|\lambda_i\| \times ulp} \tag{6.3}$$

which measures the accuracy of the computed eigenvalues compared to the reference eigenvalue δ, which are the exact eigenvalues (analytically known), or the ones computed by the QR iteration routine using LAPACK (DSTEV). The value ulp represents the machine precision computed by the LAPACK subroutine DLAMCH. Its value on the Intel Xeon, where the accuracy experiments were conducted, is equal to 2.220446049259313E-16.

For most of the tests presented here, the original matrix is either generated from a multiplication of a given diagonal eigenvalue matrix by an orthogonal matrix or given directly as a symmetric matrix. Then, for the standard algorithms (QR, BI, MR and D&C), the symmetric dense matrix is first reduced to tridiagonal form using the DSYTRD routine and then

solved using one of the algorithms cited above. While for our TD&C symmetric eigenvalue solver, the original matrix is reduced to band form, and then solved using the band D&C method according to the approach described in this paper. Therefore, the matrix B considered in the metrics (6.1), (6.2), (6.3), is either tridiagonal obtained from DSYTRD routine or band tridiagonal obtained from our TD&C symmetric eigenvalue solver.

6.8.2. Accuracy Evaluations. We present the results obtained using the accuracy metrics defined above, for all sets of the matrices described in the subsection 6.6. Our TD&C symmetric eigenvalue solver provides the three metrics with the same order of magnitude as compared to the other eigenvalue solvers.

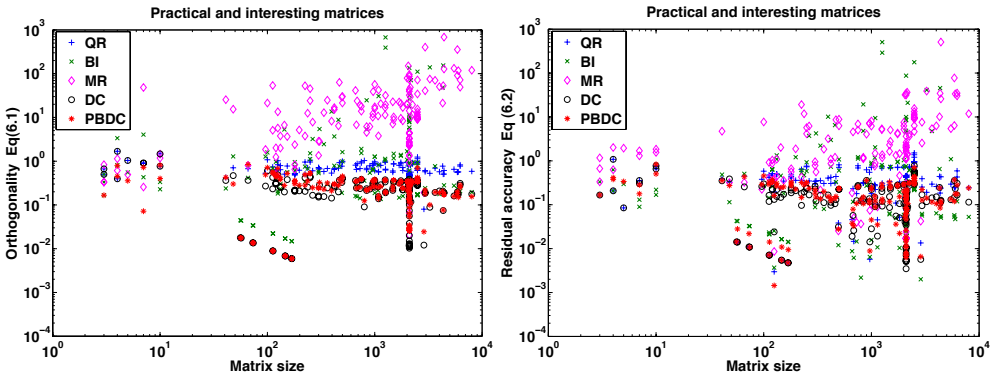

FIG. 6.7. *Summary of accuracy results observed on matrices from Table 6.2 and Table 6.3.*

Figure 6.7 and Figure 6.8 depict the losses of orthogonality (6.1) and the maximal residual norm (6.2) for practical and interesting matrix categories and synthetic testing matrix type, respectively. These figures allow for the study of the errors with respect to the matrix size n. The error of our TD&C symmetric eigenvalue solver decreases when n increases, similarly to QR and D&C eigenvalue solvers. The observed accuracy is on of order of $\mathcal{O}(\sqrt{n\varepsilon})$. However, it is clear that MR and BI do not achieve the same level of accuracy as QR, D&C and TD&C. Finally, after measuring the eigenvector orthogonality (6.1), the residual norms (6.2) and the eigenvalue accuracy (6.3), our TD&C framework is one of the most accurate symmetric eigenvalue solver algorithms, which gives us a certain confidence on the quality of the overall TD&C symmetric eigenvalue solver presented in this paper.

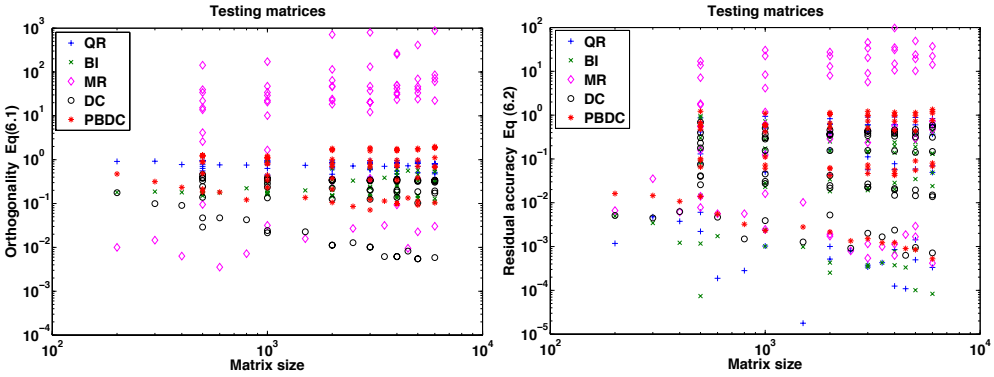

FIG. 6.8. *Summary of accuracy results observed on synthetic testing matrices from Table 6.1.*

7. Summary and Future Work. The Tile Divide and Conquer symmetric eigenvalue solver (TD&C) presented in this paper shows very promising results in terms of performance on multicore architecture as well as in terms of numerical accuracy. TD&C has been extensively tested using different matrix types against other well-known symmetric eigenvalue solvers such as the QR iteration (QR), the bisection (BI), the standard divide and conquer (D&C), and the multiple relatively robust representations (MRRR). The performance results obtained for large matrix sizes and certain matrix types are very encouraging. The proposed TD&C symmetric eigenvalue solver reaches up to 14X speed up compared to the state-of-the-art numerical open source library LAPACK V3.2, and up to 4X speed up against the commercial numerical library Intel MKL V10.2. Our TD&C symmetric eigenvalue solver also proves to be one of the most accurate symmetric eigenvalue solvers available. The authors plan to eventually integrate this symmetric eigenvalue solver within the PLASMA library [3]. The authors are also currently looking at the eigenvector calculations. Indeed, the TD&C symmetric eigenvalue solver is not really appropriate for that purpose due to the large amount of extra flops required to accumulate them. The authors are investigating the possibility of applying a single rank-b modification (instead of b rank-one modifications), which would remove the cost of accumulating the subsequent orthogonal matrices to generate the eigenvectors.

8. Acknowledgment. The authors would like to thank the two anonymous reviewers for their insightful comments, which greatly helped to improve the quality of this article.

REFERENCES

[1] `http://www.intel.com/cd/software/products/asmo-na/eng/307757.htm`.

[2] SMP Superscalar. `http://www.bsc.es/` → Computer Sciences → Programming Models → SMP Superscalar.

[3] *PLASMA Users' Guide, Parallel Linear Algebra Software for Multicore Architectures, Version 2.4.5, University of Tennessee Knoxville*, November 2011.

[4] E. Agullo, B. Hadri, H. Ltaief, and J. Dongarra. Comparative study of one-sided factorizations with multiple software packages on multi-core hardware. *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, pages 1–12, 2009.

[5] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, third edition, 1999.

[6] P. Arbenz. Divide and conquer algorithms for the bandsymmetric eigenvalue problem. *Parallel Computing*, 18(10):1105–1128, 1992.

[7] P. Arbenz, W. Gander, and G. H. Golub. Restricted rank modification of the symmetric eigenvalue problem: Theoretical considerations. *Linear Algebra Appl.*, 104:75–95, 1988.

[8] P. Arbenz and G. H. Golub. On the spectral decomposition of hermitian matrices modified by low rank perturbations. *SIAM J. Matrix Anal. Appl.*, 9(1):40–58, 1988.

[9] Y. Bai and R. C. Ward. A parallel symmetric block-tridiagonal divide-and-conquer algorithm. *ACM Trans. Math. Softw.*, 33(4):25, 2007.

[10] C. H. Bischof, B. Lang, and X. Sun. The sbr toolbox - software for successive band reduction, 1996.

[11] J. R. Bunch, C. P. Nielsen, and D. C. Sorensen. Rank-one modification of the symmetric eigenproblem. *Numer. Math.*, 31:31–48, 1978.

[12] A. Buttari, J. Langou, J. Kurzak, and J. Dongarra. A class of parallel tiled linear algebra algorithms for multicore architectures. *Parallel Computing*, 35(1):38–53, 2009.

[13] A. Buttari, J. Langou, J. Kurzak, and J. J. Dongarra. A class of parallel tiled linear algebra algorithms for multicore architectures. *Parellel Comput. Syst. Appl.*, 35:38–53, 2009.

[14] E. Chan, F. G. V. Zee, P. Bientinesi, E. S. Quintana-Orti, G. Quintana-Orti, and R. van de Geijn. Supermatrix: A multithreaded runtime scheduling system for algorithms-by-blocks. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (23th PPOPP'2008)*, pages 123–132, Salt Lake City, UT, Feb. 2008. ACM SIGPLAN.

[15] J. J. M. Cuppen. A divide and conquer method for the symmetric eigenproblem. *Numer. Math.*, 36:177–195, 1981.

[16] J. W. Demmel. *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997.

[17] J. W. Demmel, O. A. Marques, B. N. Parlett, and C. Vmel. Performance and accuracy of LAPACK's symmetric tridiagonal eigensolvers. Technical Report 183, LAPACK Working Note, Apr. 2007.

[18] I. S. Dhillon. *A new O (N(2)) algorithm for the symmetric tridiagonal eigenvalue/eigenvector problem*. PhD thesis, Berkeley, CA, USA, 1998. UMI Order No. GAX98-03176.

[19] I. S. Dhillon and B. N. Parlett. Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices. *Linear Algebra and its Applications*, 387:1 – 28, 2004.

[20] I. S. Dhillon, B. N. Parlett, and C. Vömel. The design and implementation of the mrrr algorithm. *ACM Trans. Math. Softw.*, 32:533–560, December 2006.

[21] J. J. Dongarra and D. C. Sorensen. A fully parallel algorithm for the symmetric eigenvalue problem. *SIAM J. Sci. Statist. Comput.*, 8:s139–s154, 1987.

[22] W. N. Gansterer, J. Schneid, and C. Ueberhuber. A low-complexity divide-and-conquer method for computing eigenvalues and eigenvectors of symmetric band matrices. *BIT Numerical Mathematics*, 41:967–976, 2001. 10.1023/A:1021933127041.

[23] W. N. Gansterer, R. C. Ward, and R. P. Muller. An extension of the divide-and-conquer method for a class of symmetric block-tridiagonal eigenproblems. *ACM Trans. Math. Softw.*, 28(1):45–58, 2002.

[24] K. Gates and P. Arbenz. Parallel divide and conquer algorithms for the symmetric tridiagonal eigenproblem, 1994.

[25] G. H. Golub. Some modified matrix eigenvalue problems. *SIAM Review*, 15(2):pp. 318–334, 1973.

[26] G. H. Golub and C. F. V. Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore and London, third edition, 1996.

[27] G. H. Golub and C. F. Van Loan. *Matrix Computation*. John Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, Maryland, third edition, 1996.

[28] R. Grimes, H. Krakauer, J. Lewis, H. Simon, and S.-H. Wei. The solution of large dense generalized eigenvalue problems on the cray x-mp/24 with ssd. *J. Comput. Phys.*, 69:471–481, April 1987.

[29] M. Gu and S. C. Eisenstat. A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. *SIAM J. Matrix Anal. Appl.*, 16(1):172–191, 1995.

[30] L. C. F. Ipsen and E. R. Jessup. Solving the symmetric tridiagonal eigenvalues problem on the hypercube. *SIAM J. Sci. Stat. Comput.*, 11:203–229, March 1990.

[31] J. Kurzak, A. Buttari, and J. J. Dongarra. Solving systems of linear equations on the CELL processor using Cholesky factorization. *IEEE Transactions on Parallel and Distributed Systems*, 19(9):1–11, Sept. 2008.

[32] J. Kurzak and J. Dongarra. QR Factorization for the CELL Processor. LAPACK Working Note 201, May 2008.

[33] J. Kurzak and J. Dongarra. Fully dynamic scheduler for numerical scheduling on multicore processors. Technical Report LAWN (LAPACK Working Note) 220, UT-CS-09-643, Innovative Computing Lab, University of Tennessee, 2009.

[34] O. A. Marques, C. Vömel, J. W. Demmel, and B. N. Parlett. Algorithm 880: A testing infrastructure for symmetric tridiagonal eigensolvers. *ACM Trans. Math. Softw.*, 35:8:1–8:13, July 2008.

[35] R. M. Martin. Electronic structure: Basic theory and practical methods. *Cambridge University Press*, 2008.

[36] E. S. Quintana-Ortí and R. A. van de Geijn. Updating an LU factorization with pivoting. *ACM Transactions on Mathematical Software*, 35(2), July 2008.

[37] G. Quintana-Ortí, E. S. Quintana-Ortí, E. Chan, R. A. van de Geijn, and F. G. Van Zee. Scheduling of QR factorization algorithms on SMP and multi-core architectures. In *PDP*, pages 301–310. IEEE Computer Society, 2008.

[38] N. Rösch, S. Krüger, V. Nasluzov, and A. Matveev. ParaGauss: The density functional program paragauss for complex systems in chemistry. pages 285–296, 2005. DOI: 10.1007/3-540-28555-5-25.

[39] J. Rutter and J. D. Rutter. A serial implementation of cuppen's divide and conquer algorithm for the symmetric eigenvalue problem, 1994.

[40] D. C. Sorensen and P. T. P. Tang. On the orthogonality of eigenvectors computed by divide-and-conquer techniques. *SIAM J. Numer. Anal.*, 28(6):1752–1775, 1991.

[41] R. C. Thompson. The behavior of eigenvalues and singular values under perturbations of restricted rank. *Linear Algebra and its Applications*, 13(1-2):69 – 78, 1976.

[42] F. Tisseur and J. Dongarra. Parallelizing the divide and conquer algorithm for the symmetric tridiagonal eigenvalue problem on distributed memory architectures. *SIAM J. SCI. COMPUT*, 20:2223–2236, 1998.

[43] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.

[44] R. C. Ward and R. P. Muller. *Unpublished manuscript 1999*.

[45] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Claredon Press, Oxford, UK, 1965.

[46] F. G. V. Zee, E. Chan, and R. A. van de Geijn. libflame. In D. A. Padua, editor, *Encyclopedia of Parallel Computing*, pages 1010–1014. Springer, 2011.