A Parallel Structured Divide-and-Conquer Algorithm for Symmetric Tridiagonal Eigenvalue Problems

Xia Liao, Shengguo Li, Yutong Lu, and Jose E. Roman

Abstract—In this article, a parallel structured divide-and-conquer (PSDC) eigensolver is proposed for symmetric tridiagonal matrices based on ScaLAPACK and a parallel structured matrix multiplication algorithm, called PSMMA. Computing the eigenvectors via matrix-matrix multiplications is the most computationally expensive part of the divide-and-conquer algorithm, and one of the matrices involved in such multiplications is a rank-structured Cauchy-like matrix. By exploiting this particular property, PSMMA constructs the local matrices by using generators of Cauchy-like matrices without any communication, and further reduces the computation costs by using a structured low-rank approximation algorithm. Thus, both the communication and computation costs are reduced. Experimental results show that both PSMMA and PSDC are highly scalable and scale to 4096 processes at least. PSDC has better scalability than PHDC that was proposed in [16] and only scaled to 300 processes for the same matrices. Comparing with PDSTEDC in ScaLAPACK, PSDC is always faster and achieves 1.4x–1.6x speedup for some matrices with few deflations. PSDC is also comparable with ELPA, with PSDC being faster than ELPA when using few processes and a little slower when using many processes.

Index Terms—PSMMA, PUMMA algorithm, ScaLAPACK, divide-and-conquer, rank-structured matrix, cauchy-like matrix

1 Introduction

Computing the eigendecomposition of a symmetric tridiagonal matrix is an important linear algebra problem and is widely used in many fields of science and engineering. It is usually solved by divide-and-conquer (DC) [1], [2], QR [3], MRRR [4], [5], and some other methods. The DC algorithm is now the default method in LAPACK [6] and ScaLAPACK [7] when the eigenvectors are required. DC is usually very efficient in practice, requiring only $O(n^{2.3})$ flops for matrices with dimension n on average [8], but its complexity can still be $O(n^3)$ for some matrices with few deflations. The performance of an algorithm not only depends on the number of floating point operations, but also other ingredients such as communication and data movements, etc. Some communication-avoiding algorithms have been developed recently [9], [10]. In this work, we aim to accelerate the parallel DC algorithm on distributed memory machines by reducing both the *computation* complexity and *communication* complexity. Our algorithm can be much faster than the ScaLAPACK routine PDSTEDC for those difficult matrices, and can scale to thousands of processes. Experimental results are included in Section 4.

It is known that some Cauchy-like matrices with off-diagonally low-rank properties appear in the DC algorithm [1], [2], which can be approximated by hierarchically semiseparable (HSS) matrices [11], [12]. Then, the worst case complexity of DC can be reduced from $O(n^3)$ to $O(n^2 r)$, where r is a modest number and is usually much smaller than a large n, see [13]. This technique was extended for the bidiagonal and banded DC algorithms for the SVD problem on a shared memory multicore platform in [14], [15]. For distributed memory machines, a parallel DC algorithm is similarly proposed in [16] by using STRUctured Matrices PACKage (STRUMPACK) [17], which provides some distributed parallel HSS algorithms. The accelerated DC proposed in [16] was called *PHDC*. However, numerical results show that for the simple matrix-matrix multiplication operations, STRUMPACK is not as scalable as PDGEMM, and may become slower than PDGEMM when using 300 or more processes on Tianhe-2 supercomputer. See [16] for details.

Instead of using HSS, this work exploits a much simpler type of rank-structured matrix, called BLR (Block low-rank format [18]). Compared with HSS, BLR abandons the hierarchy but compresses the off-diagonal blocks. It loses the near-linear complexity of other hierarchical matrices such as $\mathcal{H}$-matrix [19], $\mathcal{H}^2$-matrix [20], and HSS matrix. Because of its simple structure, BLR is easy to implement in parallel, and often improves the scalability of corresponding algorithms, see [18], [21] for more details.

In this paper, we propose a parallel structured matrix-matrix multiplication algorithm for Cauchy-like matrices, which will be named *PSMMA*. It exploits the off-diagonal low-rank property of matrices like BLR, and it can further

- *Xia Liao and Shengguo Li are with the College of Computer Science, National University of Defense Technology, Changsha, China. E-mail: liaoxia@nudt.edu.cn, nudtlsg@gmail.com.*
- *Yutong Lu is with the National Supercomputer Center in Guangzhou, and the School of Data and Computer Science, Sun Yatsen University, Guangzhou 510006, China. E-mail: yutong.lu@nscc-gz.cn.*
- *Jose E. Roman is with the D. Sistemes Informàtics i Computació, Universitat Politècnica de ValènciaCamí de Vera s/n46022 València, Spain. E-mail: jroman@dsic.upv.es.*

Manuscript received 1 Feb. 2020; revised 20 July 2020; accepted 19 Aug. 2020. Date of publication 25 Aug. 2020; date of current version 9 Sept. 2020. (Corresponding author: Shengguo Li.) Recommended for acceptance by J. Zhai. Digital Object Identifier no. 10.1109/TPDS.2020.3019471

1045-9219 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

reduce the communication cost by constructing local submatrices using the *generators*, which will be explained in Section 3. Our main contributions include the following:

- We propose a parallel structured matrix multiplication algorithm (PSMMA) for structured matrices including Cauchy-like, Toeplitz, Hankel, Vandermonde, etc. PSMMA can reduce both the communication and computation costs by using low-rank approximations. To the best of the authors' knowledge, none of the matrix multiplication algorithms has been developed to reduce the communication cost by exploiting the structure of matrices.
- PSMMA works for matrices both in the block cyclic data distribution (BCDD) form (like ScaLAPACK) and block data distribution (BDD) form (2D block partitioning). It also works for general process grids.
- Combining PSMMA with the DC algorithm in ScaLAPACK, we propose a parallel structured DC algorithm (PSDC), which can be much faster than `PDSTEDC` in ScaLAPACK. PSDC is also competitive with ELPA [22].

The process of PSMMA is similar to Cannon [23] and Fox [24] algorithms. However, PSMMA works for matrices in BCDD form and works for any rectangular process grids. From this perspective, PSMMA is more like PUMMA [25], a generalized Fox algorithm. PSMMA is more efficient than PUMMA for structured matrices and details are shown in Section 3.1.2. It has three advantages compared with PUMMA. One advantage is that PSMMA constructs the required submatrix locally by using the generators without communication and thus requires less *communication*. Another one is that PSMMA combines with low-rank approximations and therefore the *computation complexity* is also reduced. The third one is that PSMMA requires less workspace and the size of local matrix multiplications is also larger than PUMMA. In this paper, SRRSC [14], [26] is used to compute the low-rank approximations of Cauchy-like matrices in PSMMA, which only requires linear storage. Note that `PDGEMM` implements an algorithm similar to SUMMA [27]. Compared with SUMMA [27], which is based on the outer product form of matrix multiplication, PSMMA can naturally exploit the off-diagonal low-rank property of matrices.

By incorporating PSMMA into the DC algorithm in ScaLAPACK [8], we obtain a highly scalable DC algorithm, which has much better scalability than the previous PHDC algorithm [16]. To distinguish from PHDC, we call the newly proposed algorithm *parallel structured DC* algorithm. Numerical results show that PSDC is always faster than `PDSTEDC` in ScaLAPACK, and scales to 4,096 processes at least. That is because PSDC requires both less computations and communications than `PDSTEDC`. The speedups of PSDC over `PDSTEDC` can be up to 1.4x-1.6x for some matrices with dimension 30,000 on Tianhe-2 supercomputer. Note that PHDC in [16] can only scale to 300 processes for the same matrices.

The remaining sections of this paper are organized as follows. Section 2 introduces the DC algorithm, the SRRSC algorithm for constructing low-rank approximations of Cauchy-like matrices, and some classical parallel matrix multiplication algorithms. Section 3 presents the newly proposed parallel structured matrix multiplication algorithm PSMMA, and describes the implementation details of PSDC, which combines PSMMA with the parallel tridiagonal DC algorithm in ScaLAPACK. All the experimental results are reported in Section 4, and some future works are included in Section 4.2. Conclusions are drawn in Section 5.

2 Preliminaries

Assume T is a symmetric tridiagonal matrix,

$$T = \begin{bmatrix} a_1 & b_1 & & \\ b_2 & \ddots & \ddots & \\ & \ddots & a_{n-1} & b_{n-1} \\ & & b_{n-1} & a_n \end{bmatrix}. \tag{1}$$

We briefly introduce some formulae of Cuppen's divide-and-conquer algorithm [1], [8]. The ScaLAPACK routine also implements this version of DC algorithm [8] based on rank-one update.

Algorithm 1. DC(T, Q, Λ) Algorithm for Computing the Eigendecomposition of a Symmetric Tridiagonal Matrix

Input: $T \in \mathbb{R}^{n\times n}$
Output: eigenvalues Λ, eigenvectors Q
if *the size of T is small enough* **then**
  apply the QR algorithm and compute $T = Q\Lambda Q^T$;
  return Q and Λ;
else
  form $T = \begin{bmatrix} T_1 & \\ & T_2 \end{bmatrix} + b_k vv^T$;
  call DC(T_1, Q_1, Λ_1);
  call DC(T_2, Q_2, Λ_2);
  form $M = D + b_k uu^T$ from Q_i, Λ_i and v $(i = 1, 2)$, where $D = \mathrm{diag}(\Lambda_1, \Lambda_2)$;
  find eigenvalues Λ and eigenvectors $\widehat{Q}$ of M;
  compute the eigenvectors of T as $Q = \begin{bmatrix} Q_1 & \\ & Q_2 \end{bmatrix}\widehat{Q}$;
  return Q and Λ;
end

First, T is decomposed into the sum of two matrices,

$$T = \begin{bmatrix} T_1 & \\ & T_2 \end{bmatrix} + b_k vv^T, \tag{2}$$

where $T_1 \in \mathbb{R}^{k\times k}$, b_k is the off-diagonal element at the kth row of T and $v = [0, \ldots, 1, 1, \ldots, 0]^T$ with ones at the kth and $(k+1)$th entries. If $T_1 = Q_1\Lambda_1 Q_1^T$ and $T_2 = Q_2\Lambda_2 Q_2^T$, then T can be written as

$$T = \begin{bmatrix} Q_1 & \\ & Q_2 \end{bmatrix}\left(\begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix} + b_k uu^T\right)\begin{bmatrix} Q_1^T & \\ & Q_2^T \end{bmatrix}, \tag{3}$$

where $u = \begin{bmatrix} Q_1^T & \\ & Q_2^T \end{bmatrix} v = \begin{bmatrix} \text{last col. of } Q_1^T \\ \text{first col. of } Q_2^T \end{bmatrix}$.

Since Q_1 and Q_2 are orthogonal matrices, the problem is reduced to computing the spectral decomposition of a diagonal plus rank-one matrix,

$$M \equiv D + b_k uu^T = \widehat{Q}\Lambda\widehat{Q}^T, \tag{4}$$

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

where $D = \mathrm{diag}(\Lambda_1, \Lambda_2)$ is a diagonal matrix, Λ is a diagonal matrix whose diagonal elements are the eigenvalues of matrix M, and $\widehat{Q}$ is the eigenvector matrix of M. Then, the eigenvector matrix of T is computed as

$$\begin{bmatrix} Q_1 & \\ & Q_2 \end{bmatrix} \widehat{Q}. \tag{5}$$

The eigenvalues λ_i of $D + b_k u u^T$ are the roots of the secular equation

$$f(\lambda) = 1 + b_k \frac{u_k^2}{d_k - \lambda} = 0, \tag{6}$$

where d_k is the kth diagonal entry of D, u_k is the kth component of u. Then, the eigenvector is computed as

$$\hat{q}_i = (D - \lambda_i I)^{-1} u. \tag{7}$$

The main observation of works [13], [16] is that

$$\widehat{Q} = \left(\frac{u_i v_j}{d_i - \lambda_j} \right)_{i,j}, \tag{8}$$

where $v_j = 1/\sqrt{\sum_{k=1}^{n} \frac{u_k^2}{(d_k - \lambda_j)^2}}$, is a Cauchy-like matrix, and the vectors $u, v \in \mathbb{R}^n$ and $d = (d_1, \ldots, d_n)^T, \lambda = (\lambda_1, \ldots, \lambda_n)^T$ are called *generators*.

The whole classical DC algorithm is shown in Algorithm 1. The main computational task of DC lies in computing the eigenvectors via matrix-matrix multiplications (MMM) (5), which costs $O(n^3)$ flops. Since $\widehat{Q}$ is a Cauchy-like matrix and off-diagonally low-rank, MMM (5) can be accelerated by using HSS matrix algorithms, and the computational complexity can be reduced significantly, see [13], [16] for more details. The aim of this work is not only to reduce the computation cost of MMM (5) but also its communication cost in the distributed memory environment. For simplicity, we do not consider *deflation* in (5). About the *deflation* process, we refer the interested readers to [1], [28] and Section 3.2.

2.1 SRRSC Low-Rank Approximation

A novel low-rank approximation method for Cauchy-like matrix is proposed in [14], [26], which only requires linear storage. For completeness, this method is introduced briefly in this section, which only works on the generators.

Assume that A is an $n \times n$ Cauchy-like matrix, $A = (\frac{u_i v_j}{d_i - w_j})_{i,j}$. The following factorization is called the kth *Schur complement factorization* of A,

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & \\ A_{21} & A^{(k)} \end{bmatrix} \begin{bmatrix} I & Z^{(k)} \\ & I \end{bmatrix}, \tag{9}$$

where $A_{11} \in \mathbb{R}^{k \times k}$ and $A^{(k)}$ is called the kth Schur complement.

One good property of a Cauchy-like matrix is that its kth Schur complement $A^{(k)}$ is also Cauchy-like [29], [30], and its generators can be computed recursively as follows.

Theorem 1. *The kth Schur complement $A^{(k)}$ satisfies*

$$D_k A^{(k)} - A^{(k)} W_k = u^{(k)}(k+1:n) \cdot v^{(k)T}(k+1:n),$$

with $D_{k+1} = \mathrm{diag}(d_{k+1}, \ldots, d_n)$ and $W_{k+1} = \mathrm{diag}(w_{k+1}, \ldots, w_n)$. Then

$$u^{(k)}(k+\ell) = u^{(k-1)}(k+\ell) \cdot \frac{d_{k+\ell} - d_k}{d_{k+\ell} - w_k}, \tag{10}$$

$$v^{(k)}(k+\ell) = v^{(k-1)}(k+\ell) \cdot \frac{w_{k+\ell} - w_k}{w_{k+\ell} - d_k}, \tag{11}$$

with $1 \le \ell \le n-k, k \ge 1$, $A^{(0)} = A$, $u^{(0)} = u$ and $v^{(0)} = v$.

Corollary 1. *For $1 \le \ell \le n-k, k \ge 1$, the generators of $A^{(k)}$ are*

$$u^{(k)}(k+\ell) = \prod_{j=1}^{k} \frac{d_{k+\ell} - d_j}{d_{k+\ell} - w_j} \cdot u^{(0)}(k+\ell), \tag{12}$$

$$v^{(k)}(k+\ell) = \prod_{j=1}^{k} \frac{w_{k+\ell} - w_j}{w_{k+\ell} - d_j} \cdot v^{(0)}(k+\ell). \quad \square \tag{13}$$

It is easy to prove that $Z^{(k)}$ is also Cauchy-like, and its generators can also be computed recursively.

Theorem 2. *The generators of $Z^{(k)}$ satisfy the displacement equation*

$$W_1^{(k)} Z^{(k)} - Z^{(k)} W_2^{(k)} = y^{(k)}(1:k) \cdot v^{(k)}(k+1:n)^T,$$

where $W_1^{(k)} = \mathrm{diag}(w_1, \ldots, w_k)$ and $W_2^{(k)} = \mathrm{diag}(w_{k+1}, \ldots, w_n)$, and

$$y^{(k)}(i) = \prod_{1 \le j \le k, j \ne i} \frac{d_j - w_i}{w_j - w_i} \cdot \frac{d_i - w_i}{v^{(0)}(i)}, \tag{14}$$

and $v^{(k)}(k+1:n)$ are computed recursively by Equation (13). Furthermore,

$$y^{(k)}(i) = \begin{cases} y^{(k-1)}(i) \cdot \frac{d_k - w_i}{w_k - w_i}, & \text{if } 1 \le i \le k-1, \\ \prod_{j=1}^{k-1} \frac{w_k - d_j}{w_k - w_j} \cdot \frac{d_k - w_k}{v^{(0)}(k)}, & \text{if } i = k. \end{cases} \quad \square$$

Since permutation does not destroy the structure of Cauchy-like matrices, we can permute generators $u^{(k)}, v^{(k)}, d$ and w to make the first entry of $A^{(k)}$, $A^{(k)}(1,1) = \frac{u^{(k)}(k+1) v^{(k)}(k+1)}{d_{k+1} - w_{k+1}}$, large. A complete pivoting strategy could be used. Some efficient pivoting strategies have been proposed in [30] and [31]. We used the pivoting strategy proposed in [14].

If the entries of $A^{(k)}$ are negligible, and then by ignoring $A^{(k)}$, we get a low-rank approximation to A,

$$AP \approx A(1:n, \mathscr{T}) \begin{bmatrix} I & Z^{(k)} \end{bmatrix}, \tag{15}$$

where $\mathscr{T} = \{i_1, i_2, \ldots, i_k\} \subset \{1, 2, \ldots, n\}$ and P is a permutation matrix which records the column permutations during the pivotings. To be more specific, $A(1:n, \mathscr{T})$ consists of a subset of columns of A, and from Theorem 2,

$$Z_{ij}^{(k)} = \frac{u^{(k)}(k+i) y^{(k)}(j)}{w(i) - w(k+j)}, \tag{16}$$

where w is the array after permutation.

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

2.2 Parallel Matrix-Matrix Multiplications

Matrix-matrix multiplication is a very important computational kernel of many scientific applications. In this subsection, we briefly introduce some parallel matrix multiplication algorithms, which compute $C = A \times B$.

Cannon algorithm [23] was the first efficient algorithm for parallel matrix multiplication providing theoretically optimal communication cost. However it requires the process grid to be square, which limits its practical usage. Fox algorithm proposed in [24] has the same problem. The PUMMA algorithm [25] is a generalized Fox algorithm, and it works for a general $P \times Q$ processor grid. PUMMA was designed for ScaLAPACK and used the BCDD form.

The current version of ScaLAPACK implements SUMMA [27], which was proposed in the mid-1990s and designed for a general processor grid. It also uses the block-cyclic data distribution form. It implements the outer product form of matrix multiplication, and allows to pipeline them. PUMMA implements the inner product form, and requires the largest possible matrices for computations and communications. Some more efficient matrix multiplication algorithms have been proposed recently, like 2.5D algorithm [32], CARMA [33], CTF [34], and COSMA [35], and many others. Since they are not quite related to this current work, we do not attempt to give a complete literature review.

PUMMA is more appropriate for the rank-structured matrices than SUMMA, since the large off-diagonal blocks can be compressed by low-rank approximations. So are Cannon and Fox algorithms. It is not obvious for SUMMA to exploit the rank-structured property of the input matrices.

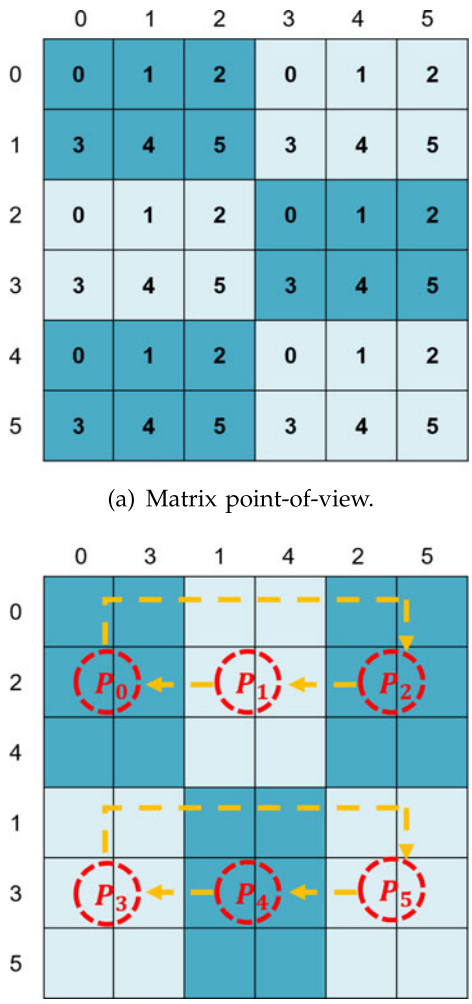
In this paper, we propose a parallel structured matrix multiplication algorithm, which is called PSMMA for short. We assume at least one of the two matrices is a structured matrix. By 'structured matrices' we mean those matrices which can be expressed using $O(n)$ parameters where n is the dimension of matrix, e.g., Cauchy-like, Toeplitz, Hankel, and Vandermonde matrices [36]. PSMMA can be based on the BCDD structure like PUMMA and ScaLAPACK, and it can also be based on the BDD structure like Cannon or Fox algorithms. As shown later, the second approach is more efficient to exploit the off-diagonal low-rank structure. Its drawback is that it requires to redistribute the matrix from the BCDD form to BDD form, since the matrices are initially stored in BCDD form in ScaLAPACK routines. The PSMMA in BCDD form fits well with ScaLAPACK routines and does not need any data redistribution.

By exploiting the special structure of matrix, PSMMA can reduce both the computation and communication costs. To the best of the authors' knowledge, this work is the first one proposing a reduction of the communication cost of matrix multiplication algorithms by exploiting the structures of matrices.

3 Algorithm Proposed

The main contributions of this paper consist of two parts. First, we design a parallel structured matrix multiplication algorithm for structured matrices in Section 3.1, which is called PSMMA and can reduce both the computation and communication costs.

Second, Section 3.2 shows how to combine PSMMA with the parallel tridiagonal DC algorithm in ScaLAPACK. It illustrates how to modify several routines in ScaLAPACK in order to use PSMMA. The parallel structured DC algorithm is called PSDC, and its whole procedure is summarized in Algorithm 3. A cartoon is included in Fig. 4 to show the whole workflow of PSDC.



(a) Matrix point-of-view.

(b) Process point-of-view.

Fig. 1. A matrix with 6×6 blocks is distributed over a 2×3 process grid.

3.1 PSMMA for Structured Matrices

In this subsection, we introduce PSMMA to compute $C = A \times B$, where $A \in \mathbb{R}^{m \times k}$ is a general matrix, $B \in \mathbb{R}^{k \times n}$ is a structured matrix, and its entries can be represented by $O(n + k)$ parameters. The case that A is a structured matrix and B is a general matrix is similar. When both A and B are structured matrices, all operations can be performed locally without any communication, which will be explained later. In the following sections, we assume A is stored in BCDD form and B is represented by its generators.

Suppose the matrix A has M block rows and K block columns, and the matrix B has K block rows and N block columns. Block (I, J) of C is then computed by

$$C(I, J) = \sum_{\ell=0}^{K-1} A(I, \ell) \cdot B(\ell, J), \tag{17}$$

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

where $I = 0, 1, \ldots, M-1, J = 0, 1, \ldots, N-1$. Cannon and Fox algorithms initially considered only the case of matrices in which each processor contains a single row or a single column of blocks [23], [24]. PUMMA considered the matrix multiplication algorithms with BCDD form [27]. Fig. 1 shows a 6×6 block matrix stored in a 2×3 process grid. It is easy to see that the matrix in the block cyclic form is obtained from the original matrix by performing row and column block permutations.

Since the matrix B can be represented by its generators, any submatrices of B can be formed easily. This fact enables us to treat the submatrices of A in each process as a whole continuous block, and we only need to construct the proper submatrices of B correctly. This is our *main observation*. After discovering this fact, the proposal of the PSMMA algorithm becomes very natural and simple, just following the Equation (17). The whole procedure is shown in Algorithm 2.

Algorithm 2. PSMMA for a Structured Matrix B

Input: $A \in \mathbb{R}^{m \times k}, B \in \mathbb{R}^{k \times n}$, where A is distributed over a $p \times q$ process grid, B is a structured matrix and all processes have a copy of its generators;

Output: $C = A \times B$.

1) Each process constructs the column indexes (*CIndex*) of matrix B based on its process column in the process grid;
2) Set $C = 0$.
3) **do** $\ell = 0, q-1$
   (a) For each process (i, j), construct the column indexes (*RIndex*) of matrix A based on the process column $\bmod (j + \ell, q)$;
   (b) Calculate the required B subblock $B(RIndex, CIndex)$, and construct its low-rank approximation (if needed) by using its generators, $B \approx U_B V_B$;
   (c) Multiply the copied A subblock with the currently residing B subblock: $C = C + (A \cdot U_B) \cdot V_B$;
   (d) Shift matrix A leftward cyclically along each process row;
   end do

To illustrate the algorithm from the process point of view, we show how the submatrices of C stored at process P_0 (located at position (0,0) of the process grid) are computed for the matrix shown in Fig. 1b. This consists of three steps, and the process is depicted in Fig. 2. The column indexes of B are fixed, and its row indexes are determined by the column indexes of A. After each step, matrix A would be shifted leftward, and the local matrix A on process P_0 is updated by the matrix on process P_1. As shown in Fig. 1b, the local matrix A of process P_0 at step 1 is updated by that of process P_1, which is located at the right of process P_0. After process P_0 has received the matrix A from all other processes, the algorithm stops.

Algorithm 2 works both for block cyclic data distribution and block data distribution. It only depends on the distribution of A to determine the row indexes of local matrix B. In step 3(b) of Algorithm 2, we construct a low rank approximation only when the local matrix B is probably low rank. For the tridiagonal DC algorithm in Section 2, matrix B is a Cauchy-like matrix and we can check whether the intersection of *CIndex* and *RIndex* is empty or not. If the intersection is empty, the B submatrix is probably numerically low-rank. Otherwise, the B submatrix is probably full rank. For Cauchy-like matrices, we use SRRSC discussed in Section 2.1 to construct a low-rank approximation to matrix B.

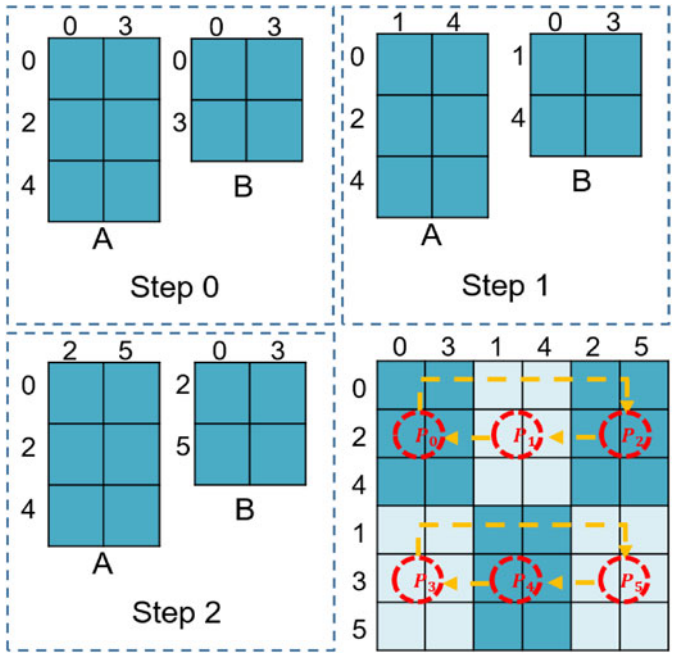


Fig. 2. The process for computing the submatrices of C located at process (0,0).

REMARK 1. The PSMMA algorithm introduced in this section is also suitable for other structured matrices, such as Toeplitz, Vandermonde, and Discrete Fourier Transform (DFT) matrices. Some results will be shown in our future works.

REMARK 2. Only step *3(d)* of Algorithm 2 requires point-to-point communications. It can be overlapped with other computations if implemented carefully. During our numerical experiments, we did not implement this technique.

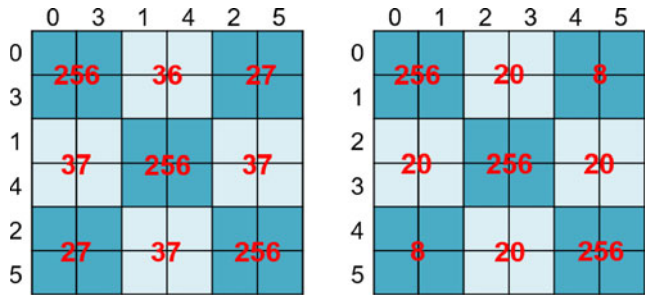
3.1.1 Storage Form Affects the Off-Diagonally Low-Rank Property

For our problem, the eigenvector matrix $\widehat{Q}$ in Equation (4) is a Cauchy-like matrix, see Equation (8). It is further off-diagonally low rank, since $\{d_i\}_{i=1}^n$ and $\{\lambda_i\}_{i=1}^n$ are interlacing. The storage form of matrix A affects the low-rank property of matrix B in Fig. 2. We use the following example to show the main points.

Example 0. Assume that matrix B is defined as $B_{ij} = \frac{u_i v_j}{d_i - w_j}$, where u_i and v_j are random numbers, $d_i = i \cdot \frac{b-a}{n}$, $w_j = d_j + \frac{b-a}{2*n}$, $a = 1.0, b = 9.0$, for $i, j = 1, 2, \ldots, n$. It is known that B is a rank-structured Cauchy-like matrix, see [14]. Let $n = 768$ and assume B is distributed over a 3×3 process grid. Suppose that $N_B = 128$ (the parameter of block size for distribution), we get the BCDD form of matrix B, as shown in Fig. 3a. By choosing $N_B = 256$, we get the BDD form of B in Fig. 3b.

Fig. 3 shows the block partitions of matrix B when choosing different N_B. The numbers in Fig. 3 are the ranks of the corresponding blocks. Fig. 3a shows the ranks of blocks of B when B is in the BCDD form. From it we can see the off-diagonal ranks in Fig. 3a are larger than those when B is in the BDD form, which are shown in Fig. 3b. Many flops can be saved when the ranks are small. When B is initially

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.



(a) Block cyclic distribution of matrix B when $N_B = 128$. (b) Block distribution of B when $N_B = 256$.

Fig. 3. The ranks of B blocks when B is in the BCDD and BDD forms, respectively.

stored in the BCDD form with small N_B, we can transform it to the BDD form by using ScaLAPACK routine PDGEMR2D. We compare these two cases in Example 1 in Section 4.

3.1.2 Comparison With Other Algorithms

Algorithm 2 works on the whole block of local matrices, like the block partition based algorithms such as Cannon and Fox algorithms. The advantages of PSMMA over Cannon and Fox algorithms are that Algorithm 2 also works for matrices of BCDD form and also works for any rectangular process grids.

Comparing with PUMMA, which works for block cyclic data distribution and rectangular process grids, Algorithm 2 requires less workspace and the size of local matrix multiplications is larger than that in PUMMA. Based on the *LCM* concept, PUMMA needs extra workspace to permute the block columns of A and block rows of B together, which can be multiplied simultaneously. Therefore, PUMMA requires roughly as much extra workspace to store another copy of the local matrices A and B, which makes it impractical in real applications, see [25], [37]. Furthermore, only partial block columns (rows) can be merged together in PUMMA, while PSMMA always multiplies the whole local matrix A with B and it does not need to permute the block columns of A or the block rows of B.

Comparing with SUMMA, which is based on the outer product form of matrix multiplication, PSMMA can further exploit the low-rank structure of matrix B. It is not easy for SUMMA to exploit this property.

3.2 Parallel Structured DC Algorithm

The excellent performance of the DC algorithm is partially due to *deflation* [1], [38], which happens in two cases. If the entry z_i of z is negligible or zero, the corresponding $(\lambda_i, \hat{q}_i)$ is already an eigenpair of T. Similarly, if two eigenvalues in D are identical then one entry of z can be transformed to zero by applying a sequence of plane rotations. All the deflated eigenvalues are moved to the end of D by a permutation matrix, and so are the corresponding eigenvectors. Then, after deflation, (3) reduces to

$$T = Q(GP)\begin{pmatrix} \bar{D} + b_k \bar{z}\bar{z}^T & \\ & \bar{D}_d \end{pmatrix}(GP)^T Q^T, \tag{18}$$

where G is the product of all rotations, P is a permutation matrix, and $\bar{D}_d$ are the deflated eigenvalues.

According to (4), the eigenvectors of T are computed as

$$U = \left[\begin{pmatrix} Q_1 & \\ & Q_2 \end{pmatrix} GP\right]\begin{pmatrix} \widehat{Q} & \\ & I_d \end{pmatrix}. \tag{19}$$

To improve efficiency, Gu [39] suggested a permutation strategy for reorganizing the data structure of the orthogonal matrices, which has been used in ScaLAPACK. The matrix in square brackets is permuted as $\begin{pmatrix} Q_{11} & Q_{12} & 0 & Q_{14} \\ 0 & Q_{22} & Q_{23} & Q_{24} \end{pmatrix}$, where the first and third block columns contain the eigenvectors that have not been affected by deflation, the fourth block column contains the deflated eigenvectors, and the second block column contains the remaining columns. Then, the computation of U can be done by two parallel matrix-matrix products (calling PDGEMM) involving parts of $\widehat{Q}$ and the matrices $(Q_{11} \quad Q_{12})$, $(Q_{22} \quad Q_{23})$. Another factor that contributes to the excellent performance of DC is that most operations can take advantage of highly optimized matrix-matrix products.

When there are few deflations, the size of matrix $\widehat{Q}$ in (19) will be large, and most of the time spent by DC would correspond to the matrix-matrix multiplication in (19). Furthermore, it is well-known that matrix $\widehat{Q}$ defined as in (4) is a Cauchy-like matrix with off-diagonally low rank property, see [2], [13]. Therefore, we simply use the parallel structured matrix-matrix multiplication algorithm to compute the eigenvector matrix U in (19). Since PSMMA requires much fewer floating point operations and communications than the plain matrix-matrix multiplication, PDGEMM, this approach makes the DC algorithm in ScaLAPACK much faster.

Algorithm 3. PSDC(T, Q, Λ) Algorithm for Computing the Eigendecomposition of a Symmetric Tridiagonal Matrix

Input: $T \in \mathbb{R}^{n\times n}$
Output: eigenvalues Λ, eigenvectors Q
if *the size of T is small enough* **then**
  apply the QR algorithm and compute $T = Q\Lambda Q^T$;
  return Q and Λ;
else
  form $T = \begin{bmatrix} T_1 & \\ & T_2 \end{bmatrix} + b_k v v^T$;
  call PSDC(T_1, Q_1, Λ_1);
  call PSDC(T_2, Q_2, Λ_2);
  form $M = \bar{D} + b_k \bar{z}\bar{z}^T$ from $Q_1, Q_2, \Lambda_1, \Lambda_2$, and v, $\bar{D} = \mathrm{diag}(\Lambda_1, \Lambda_2)$ after deflations;
  if *the size of matrix M is small* **then**
    find eigenvalues Λ and eigenvectors $\widehat{Q}$ of M;
    use Gu's idea to calculate $Q = \begin{bmatrix} Q_1 & \\ & Q_2 \end{bmatrix}\widehat{Q}$;
  else
    find eigenvalues Λ and eigenvectors $\widehat{Q}$ of M;
    use PSMMA (via SRRSC) to calculate $Q = \begin{bmatrix} Q_1 & \\ & Q_2 \end{bmatrix}\widehat{Q}$;
  end
  return Q and Λ;
end

As mentioned before, the central idea is to replace PDGEMM by PSMMA. The eigenvectors are updated in the ScaLAPACK routine PDLAED1, and therefore we modify it and call PSMMA in it instead of PDGEMM. The whole procedure of PSDC accelerated by PSMMA is summarized in Algorithm 3. Comparing

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

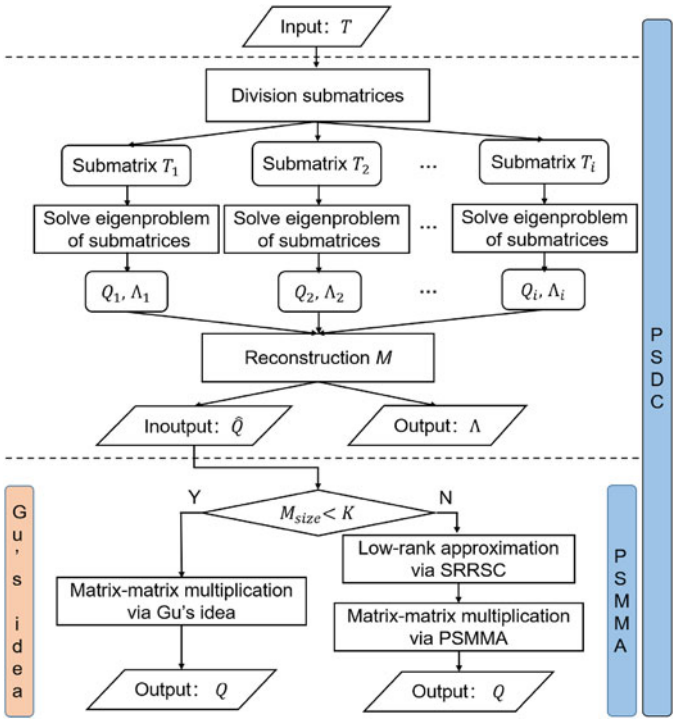


Fig. 4. The stages of the PSDC method for solving the symmetric tridiagonal eigenvalue problem.

with the classical DC algorithm (Algorithm 1), the only difference is that PSMMA is used when the size of matrix M is large. In Fig. 4 the stages of PSDC are graphically represented.

Note that after applying permutations to Q in (19), matrix $\widehat{Q}$ should also be permuted accordingly. From the results in [2], [13], [14], we know that $\widehat{Q}$ is a Cauchy-like matrix and off-diagonally low-rank, the numerical rank is usually around 50-100. When combining with PSMMA, we would not use Gu's idea [39] since permutation may destroy the off-diagonally low-rank structure of $\widehat{Q}$ in (19). We need to modify the ScaLAPACK routine PDLAED2, and only when the size of deflated matrix $\bar{D}$ in (18) is large enough, PSMMA would be used, otherwise use Gu's idea. In Section 4, we denote the size of $\bar{D}$ by K, whose value depends on the matrix as well as the architecture of the particular parallel computer used, and may be different for different computers. In Example 2, PSMMA is used when $K \geq 20,000$.

REMARK 3. To keep the orthogonality of $\widehat{\widehat{Q}}$ (see Equation (8)), $d_i - \lambda_j$ must be computed by

$$d_i - \lambda_j = \begin{cases} (d_i - d_j) - \gamma_j & \text{if } i \leq j \\ (d_i - d_{j+1}) + \mu_j & \text{if } i > j \end{cases}, \tag{20}$$

where $\gamma_i = \lambda_i - d_i$ (the distance between λ_i and d_i), and $\mu_i = d_{i+1} - \lambda_i$ (the distance between λ_i and d_{i+1}), which can be returned by calling the LAPACK routine DLAED4. In our implementation, $\widehat{Q}$ is represented by using *five* generators, $\{d_i\}, \{\gamma_i\}, \{\mu_i\}, \{u_i\}$ and $\{v_i\}$.

4 Experimental Results

All the experimental results are obtained on the Tianhe-2 supercomputer [40], [41], located in Guangzhou, China. Each compute node is equipped with two 12-cores Intel Xeon E5-2692 v2 CPUs and our experiments only use CPU cores. The details of the test platform and environment of compute nodes are shown in Table 1. For all these numerical experiments, we only used plain MPI, run 24 MPI processes per node in principle, and one process per core. For example, we used 171 compute nodes for testing 4096 processes.

TABLE 1
The Test Platform and Environment of One Node

Items	Values
2*CPU	Intel Xeon CPU E5-2692 v2@2.2GHz
Memory size	64GB (DDR3)
Operating System	Linux 3.10.0
Complier	Intel ifort 2013_sp1.2.144
Optimization	-O3 -mavx

Example 1. Assume that $A \in \mathbb{R}^{n\times n}$ is a random matrix and B is defined as $B_{ij} = \frac{u_i v_j}{d_i - w_j}$, where u_i and v_j are random numbers, $d_i = i \cdot \frac{b-a}{n}$, $w_j = d_j + \frac{b-a}{2*n}$, $a = 1.0, b = 9.0$, for $i, j = 1, \ldots, n$. It is known that B is a rank-structured Cauchy-like matrix, see [13], [14]. We compute $C = A \times B$. Let $n = 8192$, 16384 and 32768, respectively, and choose different number of processes, $N_P = 16$, 64, 256, 1024 and 4096, to compare PDGEMM with PSMMA. To avoid performance variance during multiple executions, we evaluated the performance of PDGEMM and PSMMA twice in the same program and called that program three times, and chose the best results among these six executions. Our codes will be released on Github (available at https://github.com/shengguolsg/PSMMA).

It is well-known that the performance of PDGEMM depends on the block size N_B. We tested the performances of PDGEMM by choosing $N_B = 64$, 128 and 256, and we found that their differences are very small. But they are better than choosing $N_B \leq 32$. Therefore, we chose $N_B = 64$ and $n/\sqrt{N_P}$, which corresponds to the BCDD form and BDD form, respectively. Note that a large N_B is better for PSMMA since the ranks of off-diagonal blocks after permutations may be smaller, see Table 2.

As shown in Section 3.1.1, we can redistribute matrix A from BCDD to BDD to exploit the off-diagonal low-rank property of matrix B. In this example, we tested four versions of PSMMA, which are

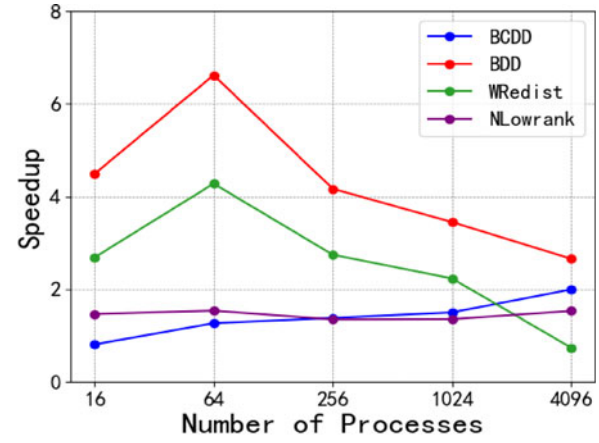
- PSMMA_BCDD: $N_B = 64$ and with low-rank approximation;
- PSMMA_BDD: $N_B = n/\sqrt{N_P}$ and with low-rank approximation;

TABLE 2
The Ranks of Off-Diagonal Blocks of B in the First Block Column When Stored on 4×4 Processes in the BCDD Form

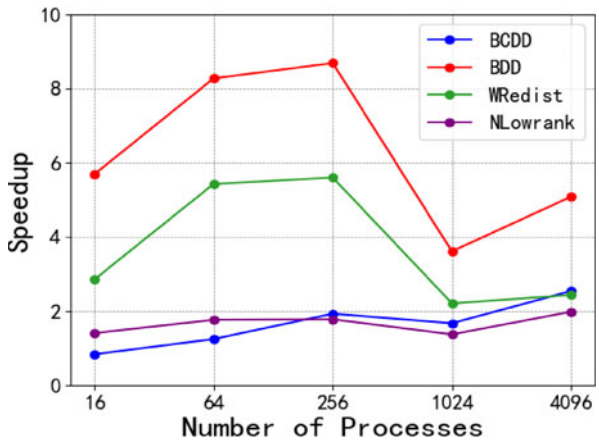
N_B	$B(2,1)$	$B(3,1)$	$B(4,1)$
64	1260	892	1252
128	711	445	699
256	401	221	390
4096	34	11	9

Each block is a 4096×4096 *submatrix.*

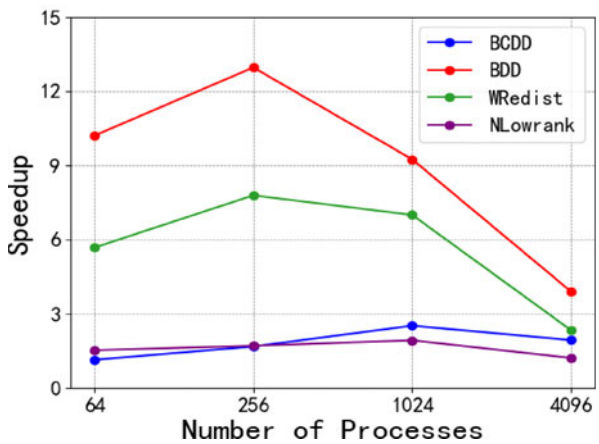
Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.



(a) Dimension $n = 8192$.



(b) Dimension $n = 16384$.



(c) Dimension $n = 32768$.

Fig. 5. The speedup of PSMMA over PDGEMM.

- PSMMA_WRedist: $N_B = 64$ and with data redistribution and low-rank approximation (matrix A is transformed from BCDD to BDD with $N_B = n/\sqrt{N_P}$ and then back);
- PSMMA_NLowrank: $N_B = 64$ and without low-rank approximation;

The speedups of PSMMA over PDGEMM are shown in Figs. 5a, 5b and 5c with dimensions $n = 8192$, 16384, 32768, respectively. From the results, we can see that PSMMA_BCDD is always faster than PDGEMM except for $N_P = 16$. It is because

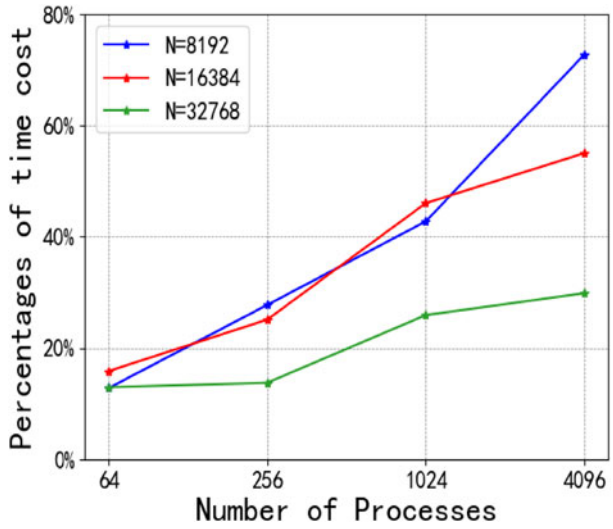


Fig. 6. The percentages of time required by data redistribution.

the ranks of B blocks are very large when using the BCDD form, and most of the time is spent in computing the low-rank approximations. Table 2 shows the ranks of the off-diagonal blocks in the first block column, see Fig. 3 for the structure of matrix B. Without using low-rank approximation, PSMMA can be faster than PDGEMM. It is interesting to see that PSMMA_NLowrank is always faster than PDGEMM for these three matrices. Note that PSMMA_NLowrank requires more floating point operations than PDGEMM since it needs to construct the local B submatrices. However, PSMMA_NLowrank requires fewer communications than PDGEMM.

PSMMA_WRedist is only slower than PDGEMM when the size of the matrix is small $(n = 8192)$ and the number of processes is large. From the last row of Table 2, we can see that the ranks of off-diagonal submatrices in BDD form are much smaller than the sizes (4096) of submatrices. PSMMA_WRedist is generally faster than both PSMMA_BCDD and PSMMA_NLowrank. The disadvantage of PSMMA_WRedist is that it requires to perform data redistribution (communication). When the number of processes is large, the data redistribution represents a large portion of the overall time. Fig. 6 shows the percentages of transforming matrix A from BCDD to BDD and transforming it back. It takes over 70 percent of total time when n is small and N_P is large.

PSMMA_BDD is always the best and can be more than ten times faster than PDGEMM. This is because it assumes that matrix A is initially stored in BDD form and the low-rank property of matrix B is not destroyed either.

It is difficult to know exactly why the speedups of PSMMA_WRedist and PSMMA_BDD over PDGEMM first increase and then decrease. It depends on the properties of these two algorithms. Comparing with PDGEMM, PSMMA_BDD and PSMMA_WRedist save both computations and communications. When the number of processes increases, these two contributions make the speedups first increase. As the number of processes increases, the size of the local submatrix located on each process decreases, and therefore the percentage of floating point operations saved from using low-rank approximations decreases. Since PSMMA_BDD does not save many floating point operations compared to PDGEMM, the speedups decrease when using

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

more processes. For PSMMA_WRedist, the cost of data redistribution also increases as the number of processes increases. Meanwhile, as the number of processes grows, the percentage of saved communication by PSMMA_BDD and PSMMA_WRedist increases. Therefore, PSMMA_BDD can be always faster than PDGEMM since it not only reduces computations but also communications. It is better to use PSMMA_NLowrank instead of PSMMA_WRedist when the number of processes is very large.

REMARK 4. We can adaptively choose a particular PSMMA algorithm based on the size of matrix and the number of processes. It is better to use PSMMA_WRedist when the number of processes is small, and use PSMMA_BCDD or PSMMA_NLowrank when the number of processes is large and/or the size of matrix is small. It is always best to use PSMMA_BDD when it is available and matrix B is off-diagonally low rank.

Example 2. We use some 'difficult' matrices [42] for the DC algorithm, for which few or no eigenvalues are deflated. Examples include the Clement-type, Hermite-type and Toeplitz-type matrices, which are defined as follows.

The Clement-type matrix [42] is given by

$$T = \text{tridiag}\begin{pmatrix} \sqrt{n} & \sqrt{2(n-1)} & & \sqrt{n\cdot 1} \\ 0 & 0 & \cdots & 0 \\ \sqrt{n} & \sqrt{2(n-1)} & & \sqrt{n\cdot 1} \end{pmatrix},$$

where the off-diagonal entries are $\sqrt{i(n+1-i)}, i = 1, \ldots, n$.

The Hermite-type matrix is given as [42],

$$T = \text{tridiag}\begin{pmatrix} \sqrt{1} & \sqrt{2} & & \sqrt{n-1} \\ 0 & 0 & \cdots & 0 \\ \sqrt{1} & \sqrt{2} & & \sqrt{n-1} \end{pmatrix}.$$

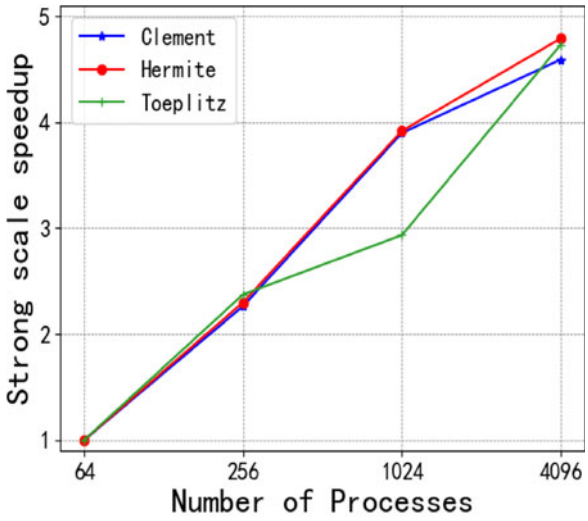
The Toeplitz-type matrix is defined as [42],

$$T = \text{tridiag}\begin{pmatrix} 1 & 1 & & 1 & 1 \\ 2 & 2 & \cdots & 2 & 2 \\ 1 & 1 & & 1 & 1 \end{pmatrix}.$$

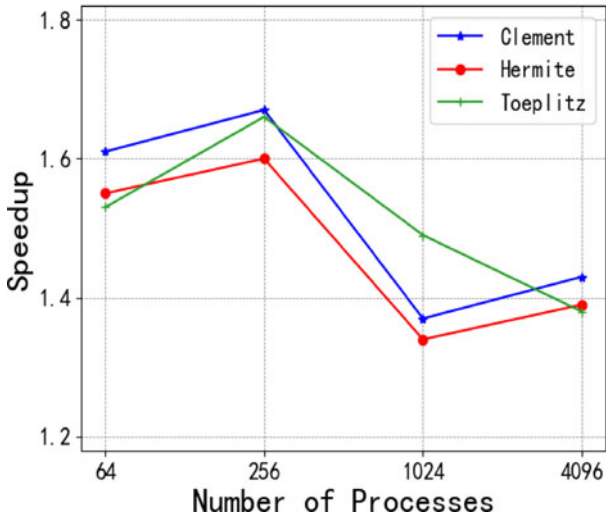
For the results of strong scaling, we let the dimension n be 30,000, and use rank-structured techniques only when the size of the secular equation is larger than $K = 20,000$. We used PSMMA_WRedist and chose $N_B = 64$. The results for strong scaling of PSDC are shown in Fig. 7a. The speed-ups of PSDC over ScaLAPACK are reported in Fig. 7b. We can see that PSDC is about 1.4x–1.6x times faster than PDSTEDC in ScaLAPACK for all cases. Because of deflations, the performances of these three matrices can be different even though they have the same dimensions.

The orthogonality of the eigenvectors computed by PSDC is in the same order as those by ScaLAPACK, as shown in Table 3. The orthogonality of matrix Q is defined as $\|I - QQ^T\|_{\max}$, where $\|\cdot\|_{\max}$ is the maximum absolute value of entries of $(\cdot)$. We confirm that the residuals of eigenpairs computed by PSDC are in the same order as those computed by ScaLAPACK though the results are not included here.

Furthermore, we compare PSDC with PHDC which was introduced in [16] and used STRUMPACK to accelerate the matrix-matrix multiplications. The results are shown in Fig. 8. For these three matrices, PSDC is much faster than PHDC when using many processes. It is better to use STRUMPACK when using few processes since HSS-based multiplications can save more floating point operations than BLR-based multiplications.



(a) The strong scaling of PSDC.



(b) The speedup of PSDC over ScaLAPACK.

Fig. 7. The results for the matrices of Example 2.

4.1 Results From Real Applications

In this subsection, we use three matrices that come from real applications to test PSDC. One comes from the spherical harmonic transform (SHT) [43], which has been used in [16]. One symmetric tridiagonal matrix is defined as follows, which will be denoted by SHT,

$$A_{jk} = \begin{cases} c_{m+2j-2}, & k = j-1 \\ d_{m+2j}, & k = j \\ c_{m+2j}, & k = j+1 \\ 0, & \text{otherwise}, \end{cases} \tag{21}$$

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

TABLE 3
The Orthogonality of the Computed Eigenvectors by PSDC

Matrix	Number of Processes			
	64	256	1024	4096
Clement	3.02e-14	3.73e-14	3.80e-14	3.65e-14
Hermite	2.49e-14	2.75e-14	2.96e-14	3.01e-14
Toeplitz	2.88e-14	3.01e-14	3.03e-14	2.97e-14

for $j, k = 0, 1, \ldots, n-1$, where $\xi_l = l - m$,

$$c_l = \sqrt{\frac{(\xi_l + 1)(\xi_l + 2)(l + m + 1)(l + m + 2)}{(2l + 1)(2l + 3)^2(2l + 5)}},$$

$$d_l = \frac{2l(l + 1) - 2m^2 - 1}{(2l - 1)(2l + 3)},$$

for $l = m, m + 1, m + 2, \ldots$. We assume the dimension of this matrix is $n = 30,000$ and $m = n$.

The other two are sparse matrices obtained from the SuiteSparse matrix collection [44], called SiO and Si5H12. We first reduce each matrix into its tridiagonal form by calling ScaLAPACK routines and then call PSDC to compute its eigendecomposition. It is also the general process for computing the eigenvalue decomposition of any symmetric (sparse) matrices. These matrices are real and symmetric and their dimensions are $n = 33,401$ and 19,896, respectively.

Example 3. We use matrices SHT, SiO and Si5H12 to compare PSDC with PDSTEDC. In this example, we use the rank-structured techniques, i.e., calling PSMMA, whenever the size of the secular equation is larger than $K = 15,000$, since the largest K for matrix Si5H12 is 15,489 when $N_B = 64$. The speedups of PSDC over PDSTEDC are reported in Table 4. The backward errors of the computed eigenpairs are also included in the third column, which are computed as

$$Residual = \frac{\|A - Q\Lambda Q^*\|_c}{\|A\|_2}, \tag{22}$$

where $Q \in \mathbb{R}^{n\times n}$ is orthogonal, $\Lambda \in \mathbb{R}^{n\times n}$ is diagonal with the eigenvalues as its diagonal elements, $\|X\|_c$ denotes the maximum Frobenius norm of each column of X and $\|X\|_2$ denotes the 2-norm of X, its largest singular value.

TABLE 4
The Speedups of PSDC Over PDSTEDC for Matrices From Real Applications

Matrix	K	Residual	Number of Processes			
			64	256	1024	4096
SHT	27,136	1.10e-14	1.20	1.48	1.42	1.41
Si5H12	15,489	3.54e-15	1.27	1.16	1.15	1.05
SiO	24,981	1.55e-14	1.46	1.31	1.10	1.15

Example 4. It is known that ELPA (Eigenvalue soLver for Petascale Applications [22], [45]) has better scalability and is faster than the MKL version of ScaLAPACK. As opposed to ScaLAPACK, ELPA routines do not rely on BLACS and PBLAS, and it can overlap the computations with communications, and the computation is also optimized by using OpenMP and even GPU. For the tridiagonal eigensolver, ELPA rewrites the DC algorithm and implements its own matrix-matrix multiplications and does not use the PBLAS routine PDGEMM. In contrast, PSDC follows the main procedure of PDSTEDC, and only uses PSMMA to accelerate the expensive matrix-matrix multiplication part.

We compare PSDC with ELPA and find that it competes with ELPA (with version 2018.11.001). We use the Clement matrix to do experiments. Fig. 9 shows the execution times of PSDC and ELPA when using different processes. It shows that PSDC is faster than ELPA when using few processes, but PSDC becomes slower when using more than 1024 processes. It is because matrix multiplication is no longer the dominant factor when using many processes. It is shown in [16] that the percentage of time dedicated to the matrix multiplications can be less than 10 percent of the total time. The gains of ELPA are obtained from other optimization techniques.

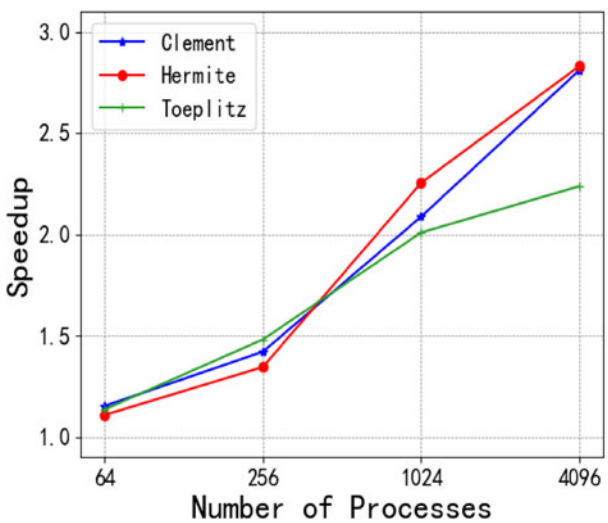


Fig. 8. The speedup of PSDC over PHDC.

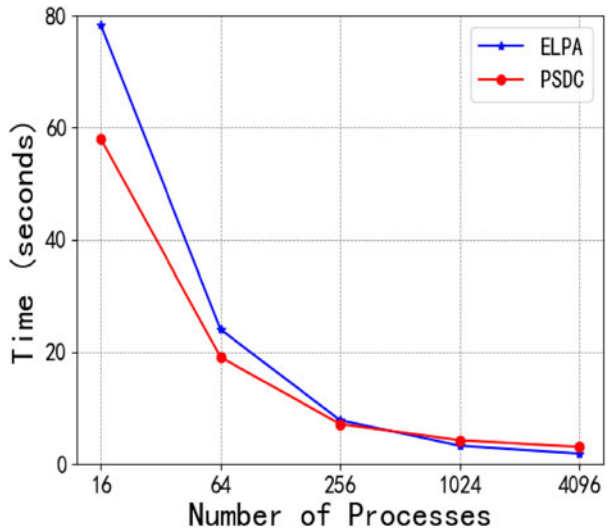


Fig. 9. The comparison of PSDC with ELPA.

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

4.2 Future Works

We discuss some bottlenecks of current implementation and future works in this subsection. A restriction of our current codes is that PSMMA was only used at the top level of the DC tree. It should be used at any level as long as the size of the matrix is large. This will be modified in the near future. We did not use OpenMP or vectorization to optimize our routines. The routines for constructing the local submatrices from generators can be optimized by using vectorization and OpenMP.

Following our current work, there are some interesting research projects to do in the future. First, PSMMA can be used to extend the banded DC algorithms proposed in [15], [46] to distributed memory platforms in a similar manner. Second, the structured matrix-matrix multiplication techniques can be used in heterogeneous architectures, which can reduce the data movements from CPU to the accelerators such as GPU.[1] We only need to transform the generators to GPUs once instead of many submatrices. Last but not least, PSMMA can be adapted for Toeplitz, Hankel, Discrete Fourier Transform and other structured matrices [47]. Some results will be included in our future works.

5 Conclusion

The starting point of this paper is trying to accelerate the tridiagonal DC algorithm implemented in ScaLAPACK [8] for some difficult matrices. It is known that the main task lies in multiplying a general matrix with a rank-structured Cauchy-like matrix, see [1], [2], [13]. The main contribution of this paper is that a highly scalable parallel matrix multiplication algorithm is proposed for rank-structured Cauchy-like matrices, which fits well for the parallel tridiagonal DC algorithm.

The matrix multiplication problem is known to be very compute intensive. However, as HPC moves towards exascale computing, the development of communication-avoiding or communication-decreasing algorithms becomes more and more important. By taking advantage of the particular structures of Cauchy-like matrices, we proposed a parallel structured matrix multiplication algorithm PSMMA, which can reduce both computation and communication costs. The workflow of PSMMA is similar to PUMMA [27] and further exploits the rank-structured property of input matrices. Experimental results show that PSMMA can be much faster than `PDGEMM` for rank-structured Cauchy-like matrices, and the speedups over `PDGEMM` can be up to 12.96.

By combing PSMMA with the parallel tridiagonal DC algorithm, we propose a parallel structured DC algorithm. For these difficult matrices for which DC deflates very few eigenvalues, PSDC is always much faster than the classical DC algorithm implemented in ScaLAPACK. Unlike PHDC which is proposed in [16], PSDC does not have scalability problem and it can scale to 4,096 processes at least.

Acknowledgments

The authors would like to thank the referees for their valuable comments which greatly improve the presentation of this article. This work was supported by National Natural Science Foundation of China (No. NNW2019ZT6-B20, NNW2019ZT6-B21, NNW2019ZT5-A10, U1611261, 61872392, and U1811461), National Key RD Program of China (2018YFB0204303), NSF of Hunan (No. 2019JJ40339), NSF of NUDT (No. ZK18-03-01), Guangdong Natural Science Foundation (2018B030312002), and the Program for Guangdong Introducing Innovative and Entrepreneurial Teams under Grant 2016ZT06D211. The work of Jose E. Roman was supported by the Spanish Agencia Estatal de Investigación (AEI) under project SLEPc-DA (PID2019-107379RB-I00).

1. The authors would like to thank the referee for pointing out this research direction.

References

[1] J. J. M. Cuppen, "A divide and conquer method for the symmetric tridiagonal eigenproblem," *Numerische Mathematik*, vol. 36, pp. 177–195, 1981.
[2] M. Gu and S. C. Eisenstat, "A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem," *SIAM J. Matrix Anal. Appl.*, vol. 16, pp. 172–191, 1995.
[3] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: Johns Hopkins Univ. Press, 1996.
[4] I. S. Dhillon and B. N. Parlett, "Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices," *Linear Algebra Appl.*, vol. 387, pp. 1–28, 2004.
[5] M. Petschow, E. Peise, and P. Bientinesi, "High-performance solvers for dense Hermitian eigenproblems," *SIAM J. Numer. Anal.*, vol. 35, pp. C1–C22, 2013.
[6] E. Anderson *et al.*, *LAPACK Users' Guide*, 3rd ed. Philadelphia, PA, USA: SIAM, 1999.
[7] J. Choi *et al.*, "ScaLAPACK: A portable linear algebra library for distributed memory computers-design issues and performance," *Comput. Phys. Commun.*, vol. 97, pp. 1–15, 1996.
[8] F. Tisseur and J. Dongarra, "A parallel divide and conquer algorithm for the symmetric eigenvalue problem on distributed memory architectures," *SIAM J. Sci. Comput.*, vol. 20, no. 6, pp. 2223–2236, 1999.
[9] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz, "Minimizing communication in numerical linear algebra," *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 2, pp. 562–580, 2011.
[10] J. Demmel, L. Grigori, M. Hoemmen, and J. Langou, "Communication-optimal parallel and sequential QR and LU factorizations," *SIAM J. Sci. Comput.*, vol. 34, pp. A206–A239, 2012.
[11] S. Chandrasekaran, T. Pals, and M. Gu, "Fast matrix algorithms for hierarchically semi-separable representations," in *Proc. Recent Adv. Comput. Sci. Eng.*, 2002, pp. 365–368.
[12] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, and T. Pals, "A fast solver for HSS representations via sparse matrices," *SIAM J. Matrix Anal. Appl.*, vol. 29, pp. 67–81, 2006.
[13] S. Li, X. Liao, J. Liu, and H. Jiang, "New fast divide-and-conquer algorithm for the symmetric tridiagonal eigenvalue problem," *Numer. Linear Algebra Appl.*, vol. 23, pp. 656–673, 2016.
[14] S. Li, M. Gu, L. Cheng, X. Chi, and M. Sun, "An accelerated divide-and-conquer algorithm for the bidiagonal SVD problem," *SIAM J. Matrix Anal. Appl.*, vol. 35, no. 3, pp. 1038–1057, 2014.
[15] X. Liao, S. Li, L. Cheng, and M. Gu, "An improved divide-and-conquer algorithm for the banded matrices with narrow bandwidths," *Comput. Math. Appl.*, vol. 71, pp. 1933–1943, 2016.
[16] S. Li, F.-H. Rouet, J. Liu, C. Huang, X. Gao, and X. Chi, "An efficient hybrid tridiagonal divide-and-conquer algorithm on distributed memory architectures," *J. Comput. Appl. Math.*, vol. 344, pp. 512–520, 2018.
[17] F. Rouet, X. Li, P. Ghysels, and A. Napov, "A distributed-memory package for dense hierarchically semi-separable matrix computations using randomization," *ACM Trans. Math. Softw.*, vol. 42, no. 4, pp. 27:1–35, 2016.
[18] P. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, and C. Weisbecker, "Improving multifrontal methods by means of block low-rank representations," *SIAM J. Sci. Comput.*, vol. 37, no. 3, pp. A1451–A1474, 2015.
[19] W. Hackbusch, "A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices," *Computing*, vol. 62, pp. 89–108, 1999.
[20] W. Hackbusch, B. Khoromskij, and S. Sauter, "On $\mathcal{H}^2$-matrices," in *Proc. Lecture Appl. Math.*, 2000, pp. 9–29.

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.

[21] I. Yamazaki, A. Ida, R. Yokota, and J. Dongarra, "Distributed-memory lattice $\mathcal{H}$-matrix factorization," *Int. J. High Perform. Comput. Appl.*, vol. 33, no. 5, pp. 1046–1063, 2019.
[22] T. Auckenthaler *et al.*, "Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations," *Parallel Comput.*, vol. 37, no. 12, pp. 783–794, 2011.
[23] L. E. Cannon, "A cellular computer to implement the kalman filter algorithm," Ph.D. dissertation, College Eng., Montana State Univ., Bozeman, MT, 1969.
[24] G. C. Fox, S. W. Otto, and A. J. G. Hey, "Matrix algorithms on a hypercube I: Matrix multiplication," *Parallel Comput.*, vol. 4, no. 1, pp. 17–31, 1987.
[25] J. Choi, D. W. Walker, and J. J. Dongarra, "Pumma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers," *Concurrency Comput., Pract. Experience*, vol. 6, no. 7, pp. 543–570, 1994.
[26] M. Gu and J. Xia, "A multi-structured superfast Toeplitz solver," preprint, 2009.
[27] R. A. V. De Geijn and J. Watts, "SUMMA: Scalable universal matrix mltiplication algorithm," *Concurrency Comput., Pract. Experience*, vol. 9, no. 4, pp. 255–274, 1997.
[28] M. Gu and S. C. Eisenstat, "A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem," *SIAM J. Matrix Anal. Appl.*, vol. 15, pp. 1266–1276, 1994.
[29] I. Gohberg, T. Kailath, and V. Olshevsky, "Fast Gaussian elimination with partial pivoting for matrices with displacement structure," *Math. Comput.*, vol. 64, no. 212, pp. 1557–1576, 1995.
[30] M. Gu, "Stable and efficient algorithms for structured systems of linear equations," *SIAM J. Matrix Anal. Appl.*, vol. 19, no. 2, pp. 279–306, 1998.
[31] C.-T. Pan, "On the existence and computation of rank revealing LU factorizations," *Linear Algebra Appl.*, vol. 316, pp. 199–222, 2000.
[32] E. Solomonik and J. Demmel, "Communication-optimal parallel 2.5D matrix multiplication and LU factorization algorithms," in *Proc. 17th Int. Conf. Parallel Process.*, 2011, pp. 90–109.
[33] J. Demmel *et al.*, "Communication-optimal parallel recursive rectangular matrix multiplication," in *Proc. IEEE 27th Int. Symp. Parallel Distrib. Process.*, 2013, pp. 261–272.
[34] E. Solomonik, D. Matthews, J. Hammond, and J. Demmel, "Cyclops tensor framework: Reducing communication and eliminating load imbalance in massively parallel contractions," in *Proc. IEEE 27th Int. Symp. Parallel Distrib. Process.*, 2013, pp. 813–824.
[35] G. Kwasniewski, M. Kabić, M. Besta, J. VandeVondele, R. Solcà, and T. Hoefler, "Red-blue pebbling revisited: Near optimal parallel matrix-matrix multiplication," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2019, pp. 1–22.
[36] V. Y. Pan, *Structured Matrices and Polynomials, Unified Superfast Algorithms*. New York, NY, USA: Birkhäuser, Springer, 2001.
[37] J. Choi, "A new parallel matrix multiplication algorithm on distributed-memory concurrent computers," *Concurrency: Pract. Experience*, vol. 10, no. 8, pp. 655–670, 1998.
[38] J. R. Bunch, C. P. Nielsen, and D. C. Sorensen, "Rank one modification of the symmetric eigenproblem," *Numerische Mathematik*, vol. 31, pp. 31–48, 1978.
[39] M. Gu, "Studies in numerical linear algebra," Ph.D. dissertation, Dept. Comput. Sci., Yale University, New Haven, CT, 1993.
[40] X. Liao, L. Xiao, C. Yang, and Y. Lu, "Milkyway-2 supercomputer: System and application," *Front. Comput. Sci.*, vol. 8, no. 3, pp. 345–356, 2014.
[41] Y. Liu *et al.*, "623 Tflop/s HPCG run on Tianhe-2: Leveraging millions of hybrid cores," *Int. J. High Perform. Comput. Appl.*, vol. 30, no. 1, pp. 39–54, 2016.
[42] O. A. Marques, C. Voemel, J. W. Demmel, and B. N. Parlett, "Algorithm 880: A testing infrastructure for symmetric tridiagonal eigensolvers," *ACM Trans. Math. Softw.*, vol. 35, pp. 8:1–13, 2008.
[43] M. Tygert, "Fast algorithms for spherical harmonic expansions, II," *J. Comput. Phys.*, vol. 227, pp. 4260–4279, 2008.
[44] T. Davis and Y. Hu, "The University of Florida sparse matrix collection," *ACM Trans. Math. Softw.*, vol. 38, no. 1, pp. 1:1–1:25, 2011.
[45] A. Marek *et al.*, "The ELPA library: Scalable parallel eigenvalue solutions for electronic structure theory and computational science," *J. Phys., Condens. Matter*, vol. 26, pp. 1–15, 2014.
[46] P. Arbenz, "Divide-and-conquer algorithms for the bandsymmetric eigenvalue problem," *Parallel Comput.*, vol. 18, pp. 1105–1128, 1992.
[47] D. Bini and V. Pan, *Polynomial and Matrix Computations, Volume I Fundamental Agorithms*. Basel, Switzerland: Birkhäuser, 1994.



Xia Liao received the MS degree from the College of Computer Science, National University of Defense Technology (NUDT), Changsha, China. She is currently working toward the PhD degree with the College of Computer Science, National University of Defense Technology, China. Her research interests include high performance computing, big data analysis and processing, and parallel computing.



Shengguo Li received the BS, MS, and PhD degrees from the National University and Defense Technology, Changsha, China, in computational mathematics. He is currently an assistant professor with the College of Computer Science, National University and Defense Technology, China. His research interests include numerical linear algebra, high performance computing, and machine learning.



Yutong Lu received the MSc and PhD degrees in computer science from the National University of Defense Technology (NUDT), Changsha, China respectively. She is currently a professor with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China. She is also the director of National Supercomputer Center in Guangzhou. Her research interests include high performance computing, parallel system management, high-speed communication, distributed file systems, and advanced programming environments with the MPI.



Jose E. Roman received the MSc and PhD degrees in computer science from the Universitat Politecnica de Valencia (UPV), Spain. He is currently a professor with the School of Computer Science, Universitat Politecnica de Valencia, Spain. His research interests focus mainly on software engineering for large-scale scientific computing, especially in the field of numerical solution of large-scale eigenvalue problems. He is the lead developer of SLEPc, a parallel software library for eigenvalue computations. Other topics of interest are numerical linear algebra, matrix functions, matrix equations, and high performance computing.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/csdl.

Authorized licensed use limited to: UNIVERSIDAD POLITECNICA DE VALENCIA. Downloaded on December 01,2020 at 08:11:47 UTC from IEEE Xplore. Restrictions apply.