

Lawrence Berkeley National Laboratory

Applied Math & Comp Sci

Title

Direction-preserving and Schur-monotonic Semi-separable Approximations of Symmetric Positive Definite Matrices

Permalink

<https://escholarship.org/uc/item/58g7q4qj>

Journal

Lawrence Berkeley National Laboratory

Author

Gu, Ming

Publication Date

2010-10-20

Peer reviewed

Direction-preserving and Schur-monotonic Semi-separable Approximations of Symmetric Positive Definite Matrices

Ming Gu^{*} Xiaoye S. Li[†] Panayot S. Vassilevski[‡]

October 28, 2009

Abstract

For a given symmetric positive definite matrix $A \in \mathbf{R}^{n \times n}$, we develop a fast and backward stable algorithm to approximate A by a symmetric positive-definite semi-separable matrix, accurate to any prescribed tolerance. In addition, this algorithm preserves the product, AZ , for a given matrix $Z \in \mathbf{R}^{n \times d}$, where $d \ll n$. Our algorithm guarantees the positive-definiteness of the semi-separable matrix by embedding an approximation strategy inside a Cholesky factorization procedure to ensure that the Schur complements during the Cholesky factorization all remain positive definite after approximation. It uses a robust direction-preserving approximation scheme to ensure the preservation of AZ . We present numerical experiments and discuss potential implications of our work.

1 Introduction

1.1 Motivation and background

Given any symmetric positive definite (SPD) matrix A and any tolerance $\tau > 0$, in this paper we present a fast backward stable algorithm to construct an SPD semi-separable matrix that approximates A , while preserving the product, AZ , for a given matrix $Z \in \mathbf{R}^{n \times d}$ for $d \ll n$. The idea of preserving the actions of A on certain vectors (directions) goes back to the early pointwise approximate factorization methods by Dupont, Kendall and Rachford [8], Gustafsson [12], and Notay [20].

^{*}861 Evans Hall, Department of Mathematics, University of California, Berkeley, California 94720, U.S.A. mgu@math.berkeley.edu. The work was supported in part by the NSF Career Award CCR-9702866.

[†]Computational Research Division, Lawrence Berkeley National Laboratory, MS 50F-1650, One Cyclotron Road, Berkeley, CA 94720, U.S.A. xsli@lbl.gov. The work was supported in part by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

[‡]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, P.O. Box 808, L-560, Livermore, CA 94551, U.S.A. panayot@llnl.gov. The work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

The motivation there was that by imposing certain row-sum criterion to the incomplete factorization of matrices coming from finite difference approximation of second order elliptic equations, it can lead to improving the condition number of the preconditioned matrix by an order of magnitude better than the one of the original finite difference matrix (i.e., from $O(h^{-2})$ to $O(h^{-1})$). One of our motivations here, is that an approximate factorization of a discretization matrix can lead to Schur complement matrices that can be viewed as coarse discretization matrices, if they preserve the near null-space of the original fine-grid matrix. Our goal is to have a general procedure that can ensure this property for any given number of directions d . For example, in the application of 2D elasticity equations it is important to preserve the so-called rigid body modes in which case we have $d = 3$. For other applications, such as the “adaptive algebraic multigrid” (cf., e.g., [1]) it is important that the coarse space contains several “algebraically smooth” directions. Although in the present paper we do not pursue the application of our direction preserving approximate factorization method to algebraic multigrid (or AMG), this is one of our main motivations to develop and study the proposed approximate factorization technique.

In what follows we adopt the so-called semi-separable matrix structure which in certain applications by using high enough rank in the approximation can lead to virtually exact factorization of the matrix. Thus, by choosing the rank we have a whole spectrum of approximate block-factorization methods that can vary in accuracy from simple preconditioners (comparable to symmetric Gauss-Seidel) to highly accurate (but potentially expensive for large rank) and virtually exact factorizations.

The semi-separable structure is a matrix analog of semi-separable integral kernels as described by Kailath in [14]. This matrix analog was most likely first described by Gohberg, Kailath and Koltracht in [18]. In that paper it is shown that, under further technical restrictions, an LDU factorization is possible with a complexity n^2N where n is the complexity of the semi-separable description and N the dimension of the matrix — in effect an algorithm linear in the size of the matrix, when n is small. In a number of papers Alpay, Dewilde and Dym introduce a new formalism for time-varying systems which provides for a framework closely analogous to the classical time invariant state space description and which allows for the generalization of many time invariant methods to the time-varying case [6, 9]. When applied to matrices, this formalism generalizes the formalism used in [18] and allows for more general types of efficient operations (by ‘efficient’ we mean operations that are linear in the size of the matrix). In the book *Time-varying Systems and Computations* [7], Dewilde and van der Veen describe the various operations that are possible on time-varying systems in great detail, including the efficient application of orthogonal transformations. In particular, they show how a URV type transformation on a general, (possibly infinite dimensional) semi-separable system can be done with an efficient recursive procedure. This procedure is based on the ideas by Dewilde and van der Veen in [22] and by Eidelman and Gohberg in [11]. In the former paper the connection with Kalman filtering as a special case of the procedures is also discussed.

In the literature, several efficient algorithms have been developed [14, 6, 9, 18, 13, 16, 17] for approximating a symmetric matrix A by a symmetric semi-separable matrix, accurate to any given tolerance $\tau > 0$. Fast backward stable algorithms have also been constructed to approximate A with an SPD semi-separable matrix (see [21]).

This current work was also motivated by such work as well as work on construction of mono-

tonic preconditioners for sparse symmetric positive definite matrices. Recent work on superfast direct methods for discretized matrices from elliptic operators uses the semi-separable matrix structure as a basic tool in solving discretized elliptic PDE problems (see [3, 2, 5, 19]). In the process of generalizing these methods to construct robust and effective preconditioners, we are led to the problem of constructing semi-separable SPD matrices to approximate a given dense SPD matrix A for a very large given tolerance $\tau > 0$. Additionally, as mentioned earlier (e.g., in the AMG application), it is often unnecessary (potentially expensive) to maintain high order of approximations along a small number of known directions defined by a given matrix $Z \in \mathbf{R}^{n \times d}$. Values such as 1, 2 or 3 are typical for d in these cases..

1.2 The paper outline

In this paper, we present an efficient and backward stable algorithm for solving such problem, for any given tolerance $\tau > 0$. This work will form the basis of our efficient construction of effective preconditioners for sparse matrices arising from discretized PDEs. As in [21], we embed the semi-separable matrix construction scheme of [2] inside the Cholesky factorization procedure for A to ensure that each approximate Schur complement during the Cholesky factorization remains positive definite. In addition, we ensure that the matrix-matrix product AZ remains unchanged throughout the entire procedure, up to rounding errors.

To be more specific, let B be a semi-separable $N \times N$ matrix. Then there exist n positive integers $m_1, \dots, m_n$ with $N = m_1 + \dots + m_n$ to block-partition A as

$$B = (B_{i,j}), \quad \text{where } B_{i,j} \in \mathbf{R}^{m_i \times m_j} \text{ satisfies } B_{i,j} = \begin{cases} D_i, & \text{if } i = j, \\ U_i W_{i+1} \cdots W_{j-1} V_j^T, & \text{if } j > i, \\ P_i R_{i-1} \cdots R_{j+1} Q_j^T, & \text{if } j < i. \end{cases} \quad (1)$$

The sequences $\{U_i\}_{i=1}^{n-1}$, $\{V_i\}_{i=2}^n$, $\{W_i\}_{i=2}^{n-1}$, $\{P_i\}_{i=2}^n$, $\{Q_i\}_{i=1}^{n-1}$, $\{R_i\}_{i=2}^{n-1}$ and $\{D_i\}_{i=1}^n$ are all matrices whose dimensions are defined in Table 1. While any matrix can be represented in this form for large enough k_i 's and l_i 's, our main focus will be on matrices of this special form that have relatively small values for the k_i 's and l_i 's (see Section 3). In the above equation, empty products are defined to be the identity matrix. For $n = 4$, the matrix B has the form

$$B = \begin{pmatrix} D_1 & U_1 V_2^T & U_1 W_2 V_3^T & U_1 W_2 W_3 V_4^T \\ P_2 Q_1^T & D_2 & U_2 V_3^T & U_2 W_3 V_4^T \\ P_3 R_2 Q_1^T & P_3 Q_2^T & D_3 & U_3 V_4^T \\ P_4 R_3 R_2 Q_1^T & P_4 R_3 Q_2^T & P_4 Q_3^T & D_4 \end{pmatrix}$$

Throughout this paper we will assume that the D_i 's are square matrices. It is shown in [7] that this class of matrices is closed under inversion and includes banded matrices, semi-separable matrices as well as their inverses as special cases.

The semi-separable structure of a given matrix B depends on the sequence m_i . Different sequences will lead to different representations.

If D_k is symmetric, $P_k = V_k$, $Q_k = W_k^T$, and $Q_k = U_k$ for all possible values of k , then B is a symmetric matrix. On the other hand, if D_k is upper triangular and $P_k = 0$ for all possible values of k , then B is an upper triangular semi-separable matrix.

Matrix	U_i	V_i	W_i	P_i	Q_i	R_i
Dimensions	$m_i \times k_i$	$m_i \times k_{i-1}$	$k_{i-1} \times k_i$	$m_i \times l_i$	$m_i \times l_{i+1}$	$l_{i+1} \times l_i$

Table 1: Dimensions of matrices in (1). k_i and l_i are column dimensions of U_i and P_i , respectively.

As is well-known, the Cholesky factor of an SPD semi-separable matrix is upper triangular semi-separable. Conversely, let R be a non-singular upper triangular semi-separable matrix. Then $R^T R$ is an SPD semi-separable matrix.

In Sections 2 we present the construction algorithm. In Section 3, we discuss numerical experimental results with this construction algorithm. In Section 4 we discuss potential applications of this work and draw some conclusions.

2 The Construction Algorithm

The main goal of this section is to present our semi-separable matrix construction algorithm. To this end, we need to establish some notation.

2.1 Notation

Let $A \in \mathbf{R}^{N \times N}$ be a symmetric positive definite (SPD) matrix, with block partitioning

$$A = \begin{pmatrix} A_{1,1} & A_{1,2} & \cdots & A_{1,n} \\ A_{1,2}^T & A_{2,2} & \cdots & A_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{1,n}^T & A_{2,n}^T & \cdots & A_{n,n} \end{pmatrix}, \quad (2)$$

where $A_{k,k} \in \mathbf{R}^{m_k \times m_k}$ so that $N = \sum_{k=1}^n m_k$. With a slight abuse of notation, we will denote

$$A_{k,s:t} = (A_{k,s} \quad \cdots \quad A_{k,t}) \quad \text{and} \quad A_{i;j,s:t} = \begin{pmatrix} A_{i,s} & \cdots & A_{i,t} \\ \vdots & \ddots & \vdots \\ A_{j,s} & \cdots & A_{j,t} \end{pmatrix}.$$

For any given matrix H and a given tolerance τ , we consider an orthogonal decomposition of the form

$$H = (U \quad \widehat{U})(V \quad \widehat{V})^T, \quad (3)$$

where the matrix $(U \quad \widehat{U})$ is column-orthonormal, so that $U^T \widehat{U} = 0$. Throughout this paper, we will decompose various matrices in the form of (3) such that U has as few columns as possible and such that $\|\widehat{V}\|_2 = O(\tau)$. Equation (3) will be our main tool for performing low numerical rank approximations.

2.2 Direction-Preserving Approximations

We start by considering direction-preserving low-rank approximations. Let $H \in \mathbf{R}^{m \times n}$, $F \in \mathbf{R}^{n \times d}$, and $G \in \mathbf{R}^{m \times d}$, we seek approximations of the form (3) that further preserve the matrix-matrix products HF and $G^T H$, for $d \ll \min(m, n)$. That is, we would like to preserve the following equalities when H is approximated by UV^T :

$$HF = UV^T F \quad \text{and} \quad G^T H = G^T UV^T. \quad (4)$$

To this end, we first QR-factorize F to get

$$F = Q_F \begin{pmatrix} R_F \\ 0 \end{pmatrix} = (Q_F^1 \quad Q_F^2) \begin{pmatrix} R_F \\ 0 \end{pmatrix}, \quad (5)$$

where $Q_F^1 \in \mathbf{R}^{n \times d}$ and $Q_F^2 \in \mathbf{R}^{n \times (n-d)}$. It is immediate that

$$HF = HQ_F^1 R_F \quad (6)$$

Next, we QR-factorize the $m \times (2d)$ matrix $(G \quad HQ_F^1)$ to get

$$(G \quad HQ_F^1) = Q_G \begin{pmatrix} R_G \\ 0 \end{pmatrix} \equiv Q_G \begin{pmatrix} R_G^1 & R_G^2 \\ 0 & 0 \end{pmatrix}, \quad (7)$$

where $Q_G \in \mathbf{R}^{m \times m}$ and $R_G \in \mathbf{R}^{(2d) \times (2d)}$. Let $R_G \equiv (R_G^1 \quad R_G^2)$, we then have

$$G = Q_G \begin{pmatrix} R_G^1 \\ 0 \end{pmatrix}, \quad (8)$$

where $R_G^1 \in \mathbf{R}^{(2d) \times d}$.

Finally, we compute the matrix

$$\widehat{H} = Q_G^T H Q_F \equiv (Q_G^T H Q_F^1 \quad Q_G^T H Q_F^2) \stackrel{(7)}{\equiv} \begin{pmatrix} R_G^2 \\ 0 \end{pmatrix} Q_G^T H Q_F^2.$$

Our goal is to approximate H by approximating $\widehat{H}$ instead. Note that $H = Q_G \widehat{H} Q_F^T$.

By construction, it is sufficient to preserve the first d columns and rows of $\widehat{H}$ in order to preserve HF and $G^T H$. Furthermore, our choices of the QR factorizations result in the lower left corner of $\widehat{H}$ being 0, as below:

$$\widehat{H} = \begin{pmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & \widehat{H}_{2,2} \end{pmatrix} \quad (9)$$

with $\widehat{H}_{1,1} \equiv R_G^2 \in \mathbf{R}^{(2d) \times d}$.

We now compute an orthogonal decomposition in the style of (3) for $\widehat{H}_{2,2}$:

$$\widehat{H}_{2,2} = (U_1 \quad U_2) \begin{pmatrix} V_1^T \\ V_2^T \end{pmatrix},$$

with columns of $(U_1 \ U_2)$ orthonormal and $\|V_2\|_2 = O(\tau)$. This leads to an orthogonal decomposition of the form (3) for $\widehat{H}$ with

$$\widehat{H} = \begin{bmatrix} I & \\ & U_1 \end{bmatrix} \begin{bmatrix} 0 \\ U_2 \end{bmatrix} \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & V_1^T \\ 0 & V_2^T \end{bmatrix}, \quad (10)$$

Since $H = Q_G \widehat{H} Q_F^T$, we can define

$$U = Q_G \begin{bmatrix} I & \\ & U_1 \end{bmatrix}, \quad \widehat{U} = Q_G \begin{bmatrix} 0 \\ U_2 \end{bmatrix},$$

and

$$V = Q_F \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & V_1^T \end{bmatrix}^T, \quad \text{and} \quad \widehat{V} = Q_F \begin{bmatrix} 0 & V_2^T \end{bmatrix}^T,$$

which leads to an orthogonal decomposition of the form (3) for H with

$$H = (U \ \widehat{U}) (V \ \widehat{V})^T. \quad (11)$$

We now show that (4) is true under this approximation. To verify the first part, we have

$$\begin{aligned} UV^T F &= Q_G \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & U_1 V_1^T \end{bmatrix} Q_F^T F = Q_G \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & U_1 V_1^T \end{bmatrix} \begin{bmatrix} R_F \\ 0 \end{bmatrix} \\ &= Q_G \begin{bmatrix} R_G^2 \\ 0 \end{bmatrix} R_F \stackrel{(7)}{=} H Q_F^1 R_F \stackrel{(6)}{=} H F. \end{aligned}$$

To verify the second part, we have

$$G^T H \stackrel{(8)}{=} \begin{bmatrix} R_G^{1^T} & 0 \end{bmatrix} Q_G^T H = \begin{bmatrix} R_G^{1^T} & 0 \end{bmatrix} \widehat{H} Q_F^T = R_G^{1^T} \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \end{bmatrix} Q_F^T,$$

and

$$G^T UV^T \stackrel{(8)}{=} \begin{bmatrix} R_G^{1^T} & 0 \end{bmatrix} Q_G^T Q_G \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \\ 0 & U_1 V_1^T \end{bmatrix} Q_F^T = R_G^{1^T} \begin{bmatrix} \widehat{H}_{1,1} & \widehat{H}_{1,2} \end{bmatrix} Q_F^T.$$

Thus, the above two quantities are equal.

It costs $O((m+n)d^2)$ flops to compute both QR factorizations; it costs $O(mnd)$ flops to compute $\widehat{H}$; and it costs $O(mnr)$ flops to compute an orthogonal decomposition for $\widehat{H}_{2,2}$, where r is the column dimension of V_1 ; and it costs about $O((m+n)(d+r)^2)$ flops to compute the representation (11). So the total cost of this compression scheme is about $O(mn(r+d))$ flops.

2.3 Construction of Approximate Cholesky Factorization

To begin our procedure, we first recall the following standard block Cholesky factorization procedure:

for $k = 1, 2, \dots, n$:

Cholesky factorize $R_{k,k}^T R_{k,k} := A_{k,k}$;
 Compute $R_{k,k+1:n} := R_{k,k}^{-T} A_{k,k+1:n}$;
 Schur complement $\mathcal{A}_k := A_{k+1:n,k+1:n} - R_{k,k+1:n}^T \cdot R_{k,k+1:n}$;

end for

For each k , the first step in this procedure computes the Cholesky factorization of the k -th diagonal block; the second step computes the rest of k -th block row; and the last step computes the Schur complement of the k -th block. The output of this procedure is the upper triangular matrix

$$R = \begin{pmatrix} R_{1,1} & R_{1,2} & \cdots & R_{1,n} \\ & R_{2,2} & \cdots & R_{2,n} \\ & & \ddots & \vdots \\ & & & R_{n,n} \end{pmatrix} \quad \text{such that} \quad A = R^T R.$$

In the following, we will modify the above procedure to find an approximate Cholesky factorization satisfying

$$S^T S = A + O\left(\sqrt{\|A\|_2} \tau\right) \quad \text{and} \quad S^T S Z = AZ, \quad (12)$$

where

$$Z = \begin{pmatrix} Z_1 \\ \vdots \\ Z_n \end{pmatrix},$$

and where S is an upper triangular semi-separable matrix of the form (cf. (1))

$$S = \begin{pmatrix} D_1 & S_{1,2} & \cdots & S_{1,n} \\ & D_2 & \cdots & S_{2,n} \\ & & \ddots & \vdots \\ & & & D_n \end{pmatrix}, \quad (13)$$

with the D_k 's being upper triangular, and $S_{k,t} = U_k W_{k+1} \cdots W_{t-1} V_t^T$.

In light of the block Cholesky factorization procedure above, we begin by computing

$$D_1^T D_1 = A_{11} \quad \text{and} \quad H_1 = D_1^{-T} A_{1,2:n}.$$

Our next step is to compute a low-rank approximation to H_1 without changing AZ . Note that

$$AZ = \begin{pmatrix} D_1^T D_1 Z_1 + D_1^T H_1 Z_{2:n} \\ H_1^T G_1 + A_{2:n,2:n} Z_{2:n} \end{pmatrix},$$

where $G_1 = D_1 Z_1$. To preserve AZ , we only need to find a low-rank approximation to H_1 while preserving both $H_1 Z_{2:n}$ and $G_1^T H_1$. Here we compute an orthogonal decomposition of H_1 in the style of equation (11) as follows:

$$H_1 = (U_1 \quad \widehat{U}_1) (Q_1 \quad \widehat{Q}_1)^T,$$

where the matrix $\begin{pmatrix} U_1 & \widehat{U}_1 \end{pmatrix}$ is column orthogonal and $\|\widehat{Q}_1\|_2 \leq \tau$. It follows that

$$H_1^T H_1 = Q_1 Q_1^T + \widehat{Q}_1 \widehat{Q}_1^T.$$

According to the block Cholesky factorization procedure, $(D_1 \ H_1)$ is actually the first block row of R . Hence the Schur complement of the first block becomes

$$\mathcal{A}_1 = A_{2:n,2:n} - H_1^T H_1 = A_{2:n,2:n} - Q_1 Q_1^T - \widehat{Q}_1 \widehat{Q}_1^T.$$

We now approximate $\mathcal{A}_1$ by

$$\widetilde{\mathcal{A}}_1 = A_{2:n,2:n} - Q_1 Q_1^T = \mathcal{A}_1 + \widehat{Q}_1 \widehat{Q}_1^T = \mathcal{A}_1 + O(\tau^2).$$

Since A is symmetric positive definite, both the Schur complement $\mathcal{A}_1$ and its approximation $\widetilde{\mathcal{A}}_1$ must also be symmetric positive definite. We note that this approximation amounts to adding a symmetric positive semi-definite matrix of norm at most τ^2 to the original matrix.

We further approximate $H_1 = R_{1,2:n}$ by $U_1 Q_1^T$. Since this approximation is done on the Cholesky factor of A , the amount of perturbation to A is only $O(\|D_1\|_2 \tau) = O(\sqrt{\|A\|_2} \tau)$.

After these two approximations, we obtain the first block row in the Cholesky factor as

$$\begin{pmatrix} D_1 & U_1 Q_1^T \end{pmatrix}$$

and the Schur complement is now $\widetilde{\mathcal{A}}_1$. We will only store the current $A_{2:n,2:n}$ and Q_1 instead of computing $\widetilde{\mathcal{A}}_1$ explicitly.

To continue, partition

$$Q_1^T = \begin{pmatrix} V_2^T & \widehat{H}_1 \end{pmatrix}.$$

The Schur complement becomes

$$\widetilde{\mathcal{A}}_1 = \begin{pmatrix} A_{2,2} - V_2 V_2^T & A_{2,3:n} - V_2 \widehat{H}_1 \\ (A_{2,3:n} - V_2^T \widehat{H}_1)^T & A_{3:n,3:n} - \widehat{H}_1^T \widehat{H}_1 \end{pmatrix}.$$

For approximations on the second block, we first compute

$$A_{2,2} := A_{2,2} - V_2 V_2^T \quad \text{and} \quad A_{2,3:n} := A_{2,3:n} - V_2 \widehat{H}_1.$$

We then Cholesky factorize $A_{2,2} := D_2^T D_2$, compute $H_2 := D_2^{-T} A_{2,3:n}$, and define

$$\mathcal{D}_2 = \begin{pmatrix} D_1 & U_1 V_2^T \\ & D_2 \end{pmatrix} \quad \text{and} \quad \mathcal{H}_2 = \begin{pmatrix} U_1 & \\ & I \end{pmatrix}.$$

With this notation and the approximation to H_1 , we can rewrite the matrix A as

$$A \approx \begin{pmatrix} \mathcal{D}_2^T \mathcal{D}_2 & \mathcal{D}_2^T \mathcal{H}_2 \begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix} \\ \begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix}^T \mathcal{H}_2^T \mathcal{D}_2 & A_{3:n,3:n} \end{pmatrix}. \quad (14)$$

The product AZ was preserved in approximation to H_1 . In approximating $\begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix}$, we only need to preserve the product of matrix on the right hand side of (14) and Z . But this can be done by preserving $\begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix} Z_{3:n}$ and $G_2^T \begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix}$, where $G_2 = \mathcal{H}_2^T \mathcal{D}_2 Z_{1:2}$.

Preserving these directions, we compute an orthogonal decomposition in the style of (11) as follows:

$$\begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix} = (\mathcal{U}_2 \quad \widehat{\mathcal{U}}_2) (\mathcal{Q}_2 \quad \widehat{\mathcal{Q}}_2)^T,$$

where the matrix $(\mathcal{U}_2 \quad \widehat{\mathcal{U}}_2)$ is column orthogonal and $\|\widehat{\mathcal{Q}}_2\|_2 \leq \tau$. As before, approximating $\begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix}$ by $\mathcal{U}_2 \mathcal{Q}_2^T$ will not change the original matrix-matrix product AZ .

We write the Schur complement of $A_{2,2}$ as

$$\begin{aligned} \mathcal{A}_2 &= A_{3:n,3:n} - \widehat{H}_1^T \widehat{H}_1 - H_2^T H_2 = A_{3:n,3:n} - \begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix}^T \begin{pmatrix} \widehat{H}_1 \\ H_2 \end{pmatrix} \\ &= A_{3:n,3:n} - \mathcal{Q}_2 \mathcal{Q}_2^T - \widehat{\mathcal{Q}}_2 \widehat{\mathcal{Q}}_2^T. \end{aligned}$$

We now approximate $\mathcal{A}_2$ by

$$\widetilde{\mathcal{A}}_2 = A_{3:n,3:n} - \mathcal{Q}_2 \mathcal{Q}_2^T$$

and the first two blocks of the Cholesky factor by

$$\begin{pmatrix} D_1 & U_1 V_2^T & U_1 W_2 \mathcal{Q}_2^T \\ & D_2 & U_2 \mathcal{Q}_2^T \\ & & \ddots \end{pmatrix},$$

where we have used the partition

$$\mathcal{U}_2 = \begin{pmatrix} W_2 \\ U_2 \end{pmatrix}.$$

Again this approximation ensures that the matrix-matrix product AZ remains unchanged, and the Schur complement $\widetilde{\mathcal{A}}_2$ remains SPD.

To continue this procedure by induction, we assume that at the k -th step for $k < n - 1$, the first k blocks of the approximate Cholesky factor has the form

$$\begin{pmatrix} D_1 & U_1 V_2^T & \cdots & U_1 W_2 \cdots W_{k-1} V_k^T & U_1 W_2 \cdots W_k \mathcal{Q}_k^T \\ & D_2 & \cdots & U_2 W_3 \cdots W_{k-1} V_k^T & U_2 W_3 \cdots W_k \mathcal{Q}_k^T \\ & & \ddots & \vdots & \vdots \\ & & & D_k & U_k \mathcal{Q}_k^T \end{pmatrix},$$

and the approximate Schur complement has the form

$$\widetilde{\mathcal{A}}_k = A_{k+1:n,k+1:n} - \mathcal{Q}_k \mathcal{Q}_k^T.$$

As before, partition

$$\mathcal{Q}_k^T = \begin{pmatrix} V_{k+1}^T & \widehat{H}_k \end{pmatrix}$$

so that

$$\widetilde{\mathcal{A}}_k = \begin{pmatrix} A_{k+1,k+1} - V_{k+1} V_{k+1}^T & A_{k+1,k+2:n} - V_{k+1} \widehat{H}_k \\ (A_{k+1,k+2:n} - V_{k+1} \widehat{H}_k)^T & A_{k+2:n,k+2:n} - \widehat{H}_k^T \widehat{H}_k \end{pmatrix}.$$

We explicitly compute

$$A_{k+1,k+1} := A_{k+1,k+1} - V_{k+1} V_{k+1}^T \quad \text{and} \quad A_{k+1,k+2:n} := A_{k+1,k+2:n} - V_{k+1} \widehat{H}_k.$$

We then Cholesky factorize $A_{k+1,k+1} := D_{k+1}^T D_{k+1}$ and compute

$$H_{k+1} := D_{k+1}^{-T} A_{k+1,k+2:n}.$$

Define

$$\mathcal{D}_{k+1} = \begin{pmatrix} D_1 & U_1 V_2^T & \cdots & U_1 W_2 \cdots W_{k-1} V_k^T & U_1 W_2 \cdots W_k V_{k+1}^T \\ & D_2 & \cdots & U_2 W_3 \cdots W_{k-1} V_k^T & U_2 W_3 \cdots W_k V_{k+1}^T \\ & & \ddots & \vdots & \vdots \\ & & & D_k & U_k V_{k+1}^T \\ & & & & D_{k+1} \end{pmatrix}$$

and

$$\mathcal{H}_{k+1} = \begin{pmatrix} U_1 W_2 \cdots W_k \\ U_2 W_3 \cdots W_k \\ \vdots \\ U_k \\ I \end{pmatrix}.$$

We can write the matrix approximation as

$$A \approx \begin{pmatrix} \mathcal{D}_{k+1}^T \mathcal{D}_{k+1} & \mathcal{D}_{k+1}^T \mathcal{H}_{k+1} \begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix} \\ \left(\begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix} \right)^T \mathcal{H}_{k+1}^T \mathcal{D}_{k+1} & A_{k+2:n,k+2:n} \end{pmatrix}.$$

In order to keep the matrix-matrix product AZ unchanged, we only need to preserve $\begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix} Z_{k+2:n}$ and $G_{k+1}^T \begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix}$ for $G_{k+1} = \mathcal{H}_{k+1}^T \mathcal{D}_{k+1} Z_{1:k+1}$. As before, we compute an approximation of $\begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix}$ in the style of equation (11) as follows:

$$\begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix} = (\mathcal{U}_{k+1} \quad \widehat{\mathcal{U}}_{k+1}) (\mathcal{Q}_{k+1} \quad \widehat{\mathcal{Q}}_{k+1})^T, \quad (15)$$

where the matrix $(\mathcal{U}_{k+1} \quad \widehat{\mathcal{U}}_{k+1})$ is column orthogonal and $\|\widehat{\mathcal{Q}}_{k+1}\|_2 \leq \tau$.

It follows that the Schur complement for block $k+1$ is

$$\mathcal{A}_{k+1} = A_{k+2:n,k+2:n} - \widehat{H}_k \widehat{H}_k^T - H_{k+1}^T H_{k+1} = A_{k+2:n,k+2:n} - \begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix} \begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix}^T.$$

Again, this allows us to write

$$\mathcal{A}_{k+1} = A_{k+2:n, k+2:n} - Q_{k+1} Q_{k+1}^T - \widehat{Q}_{k+1} \widehat{Q}_{k+1}^T,$$

which is then approximated by

$$\widetilde{\mathcal{A}}_{k+1} = A_{k+2:n, k+2:n} - Q_{k+1} Q_{k+1}^T.$$

Since the difference between $\widetilde{\mathcal{A}}_{k+1}$ and $\mathcal{A}_{k+1}$ is a symmetric positive semi-definite matrix, $\widetilde{\mathcal{A}}_{k+1}$ must itself be a symmetric positive definite matrix.

After these computational steps, the approximate Cholesky factor becomes

$$\begin{pmatrix} D_1 & U_1 V_2^T & \cdots & U_1 W_2 \cdots W_{k-1} V_k^T & U_1 W_2 \cdots W_k V_{k+1}^T & U_1 W_2 \cdots W_k \widehat{H}_k \\ & D_2 & \cdots & U_2 W_3 \cdots W_{k-1} V_k^T & U_2 W_3 \cdots W_k V_{k+1}^T & U_2 W_3 \cdots W_k \widehat{H}_k \\ & & \ddots & \vdots & \vdots & \vdots \\ & & & D_k & U_k V_{k+1} & U_k \widehat{H}_k \\ & & & & D_{k+1} & H_{k+1} \end{pmatrix}.$$

Partitioning

$$\mathcal{U}_{k+1} = \begin{pmatrix} W_{k+1} \\ U_{k+1} \end{pmatrix},$$

in the numerical low rank approximation of $\begin{pmatrix} \widehat{H}_k \\ H_{k+1} \end{pmatrix}$ (see (15)) leading to $\widehat{H}_k \approx W_{k+1} Q_{k+1}^T$ and $H_{k+1} \approx U_{k+1} Q_{k+1}^T$, thus ending up with a new approximate Cholesky factor of the form

$$\begin{pmatrix} D_1 & U_1 V_2^T & \cdots & U_1 W_2 \cdots W_{k-1} V_k^T & U_1 W_2 \cdots W_k V_{k+1}^T & U_1 W_2 \cdots W_k W_{k+1} Q_{k+1}^T \\ & D_2 & \cdots & U_2 W_3 \cdots W_{k-1} V_k^T & U_2 W_3 \cdots W_k V_{k+1}^T & U_2 W_3 \cdots W_k W_{k+1} Q_{k+1}^T \\ & & \ddots & \vdots & \vdots & \vdots \\ & & & D_k & U_k V_{k+1} & U_k W_{k+1} Q_{k+1}^T \\ & & & & D_{k+1} & U_{k+1} Q_{k+1}^T \end{pmatrix}.$$

Throughout the steps, the matrix-matrix product AZ have always been kept unchanged.

This completes the induction for $k < n - 1$. For $k = n - 1$, the new approximate Cholesky factor still has the form similar to above, without the last column. This is exactly the form of the matrix S defined in (13). This ends the proof. $\diamond$

It can easily be seen from the algorithm description that every approximate Schur complement during the Cholesky factorization is obtained by adding symmetric positive semi-definite matrices of norm at most τ^2 to the true one. We also perform an approximation of the order $O(\sqrt{\|A\|_2} \tau)$ for low-rank approximation at every step of the algorithm. Hence the total truncation error $O(\sqrt{\|A\|_2} \tau)$ in equation (12) could be $O(n)$ times larger than $\sqrt{\|A\|_2} \tau$.

Assume that each diagonal block in A has roughly the same number of columns. Let p be the maximum dimension in all the diagonal blocks, and assume that p is bigger than the column dimension of every matrix U_k . Then the cost for each step is $O(Np^2)$ flops, leading to a total cost of $O(n^2 p^3) = O(N^2 p)$ flops for the whole construction algorithm.

As is shown in [2], the column dimensions of U_k in S turns out to be precisely the rank of $S_{1:k,k+1:n}$, for $k = 1, \dots, n-1$. If $A_{1:k,k+1:n}$ has small numerical rank for the given tolerance for $k = 1, \dots, n-1$, the matrix S constructed above will also have small rank in each of its upper off-diagonal blocks. Otherwise some U_k 's would need to have large number of columns.

We have presented our construction algorithm using SVDs. However, any rank-reveal decomposition satisfying equation (3) will also work. Good examples are rank-revealing QR factorizations and rank-revealing modified Gram-Schmidt procedures. It is likely that this will lead to considerable speed-up for a small loss in compression.

For the purpose of computing a preconditioner, we can further require that the number of columns in U_k not to exceed a certain pre-set number, like *MaxRank*. This is equivalent to restricting the number of columns in U in equation (3) never to exceed a certain pre-set number, like *MaxRank*. In our numerical experiments, we simply set the submatrix $\widehat{H}_{2,2} = 0$ in equation (9). This simple strategy has still led to very effective preconditioners, see Section 3.

3 Numerical Results

We have written a C code implementing our construction algorithm. In the following we report some numerical experimental results with this code. Here we concentrate on demonstrating the effectiveness of our semi-separable matrix approximations as preconditioners.

First, we consider finite-element discretizations on uniform triangular mesh of size h , with piecewise linear functions of the following model problem defined on the unit square $\Omega = [0, 1] \times [0, 1]$:

$$-\operatorname{div}(k(x, y)\nabla u) = f(x, y), \quad (16)$$

where the coefficient $k(x, y)$ is a two-by-two matrix of the form $\epsilon I + \mathbf{b}\mathbf{b}^T$ for a given $\epsilon > 0$ and variable direction vector $\mathbf{b} = \begin{bmatrix} \cos \alpha(1 - x \cos \alpha) \\ \sin \alpha(1 - y \sin \alpha) \end{bmatrix}$. In the test we chose $\epsilon = 0.01$ and $\alpha = \frac{\pi}{3}$. We assume a mixture of Dirichlet and Neumann boundary conditions.

We use standard lexicographic ordering of the unknowns (or mesh-points). The block-structure of the matrix is obtained by putting together every p consecutive nodes in a block. In the test we varied the block size p , the number of direction vectors, d , between zero and three, and the maximum rank r . Note that, our algorithm requires $r \geq 2d$ and the block size p to be not smaller than the rank r . We present results of two settings of block size and max rank, the smaller one with $p = 8$ and $r = 2d + 2$, and the larger one with $p = 20$ and $r = 2d + 10$. The direction vectors correspond to the constant vector for $d = 1$, and $d = 2$ and $d = 3$ correspond additionally to the vectors coming from the linear functions x and y evaluated at the nodes of the mesh. We use the thus constructed block-factorization matrix as a preconditioner in the preconditioned conjugate gradient (or PCG) method. We list in Table 2 the number of iterations m for which the respective residuals satisfy $\sqrt{\mathbf{r}_m^T \mathbf{r}_m} \leq 10^{-6} \sqrt{\mathbf{r}_0^T \mathbf{r}_0}$. We do not use the preconditioned residual norm since we want to compare the different preconditioners corresponding to different d (the number of directions) using fixed norm. We also include the time to construct the approximate factor preconditioner. The tests were run on an 1.9 GHz IBM Power5 machine at the National Energy Research Scientific Computing Center.

The results in Table 2 show some (non-linear) improvement of the number of iterations for increased number of directions used. It is clear that the preconditioner for larger d is more expensive to construct and apply. Also as expected, larger rank results in better approximate factorization. It is good that the extra construction cost is acceptable—with more than doubling the block size and rank, the construction time is not more than doubled, and the increase is smaller as the problem size increases.

	$p = 8, \quad r = 2d + 2$						$p = 20, \quad r = 2d + 10$					
h^{-1}	$d = 0$	time	$d = 1$	$d = 2$	$d = 3$	time	$d = 0$	time	$d = 1$	$d = 2$	$d = 3$	time
12	28	0.00	24	21	20	0.01	7	0.00	1	1	1	0.00
24	61	0.05	55	51	51	0.07	28	0.05	24	23	20	0.13
48	115	0.57	113	121	110	0.91	77	1.00	65	65	53	1.14
96	233	8.52	221	216	210	13.74	158	15.48	139	185	118	18.49

Table 2: Number of PCG iterations for anisotropic diffusion equation: $\epsilon = 0.01$, $\alpha = \frac{\pi}{3}$. The times (in seconds) for constructing the preconditioner are shown for $d = 0$ and $d = 3$.

The purpose of the second test that we performed is to achieve high tolerance in the approximation, when we factor a dense s.p.d. matrix. We consider the model anisotropic diffusion problem (16), for a set of diffusion direction vectors $\mathbf{b}$. The dense matrix under consideration is obtained as follows. We order the nodes using the nested-dissection ordering [10, 15]. In this ordering, the last $n \times n$ ($n = 1/h$) dimensional Schur complement, S , is a dense symmetric positive definite matrix, costing traditional direct solvers $O(n^3)$ operations to factorize. We approximate this matrix by $R^T R$, where R is an upper-triangular semi-separable matrix with maximum off-diagonal rank at most 2. We require that a single direction $Z = (1, \dots, 1)^T$ be preserved under our compression scheme. Z in this case is a well-known rigid-body mode of our model problem under our discretization. This implies that we must set $\widehat{H}_{2,2} = 0$ in equation (11) at every step of compression, even though the matrix S in consideration can be very ill-conditioned. Let $\widehat{S} = R^{-T} S R^{-1}$. Obviously, $R^T R$ is a good preconditioner if $\kappa(\widehat{S})$, the condition number of $\widehat{S}$, is much smaller than $\kappa(S)$, the condition number of S . Table 3 summarizes our results for this problem. We observe that $\kappa(\widehat{S})$ always hovers around 1, indicating highly effectiveness of $R^T R$ as a preconditioner for S . In other words, the last $n \times n$ dimensional dense Schur complement in the traditional Cholesky factorization can be well-represented by a semi-separable representation with off-diagonal rank 2.

Finally, we considered the two dimensional linear elasticity equation

$$-(\mu \nabla \vec{u} + \lambda \nabla \nabla \bullet \vec{u}) = \vec{f} \quad \text{in } \Omega = (0, 1) \times (0, 1), \quad (17)$$

$$\vec{u} = \vec{0} \quad \text{on } \partial\Omega, \quad (18)$$

here $\vec{u} \in R^2$ is the displacement vector field, λ and μ are the Lamé constants. This PDE is very ill-conditioned when the ratio λ/μ is very large; this limit is known as the incompressible limit and is associated with the mechanical behavior of elastomeric materials and plastic flow in metals, for example. Iterative methods including standard geometric multigrid converge very slowly or even diverge for very large λ/μ . However, such situations are important as they are ubiquitous in nature;

	$n = 200, \mathbf{b}$ is unit random				$n = 400, \mathbf{b}$ is unit random			
ϵ	1	10^{-4}	10^{-8}	10^{-12}	1	10^{-4}	10^{-8}	10^{-12}
$\kappa(S)$	4.7×10^2	5.1×10^3	1.3×10^2	5.6×10^2	4.9×10^2	2.9×10^2	6.4×10^5	4.7×10^2
$\kappa(\widehat{S})$	2.9	1.3	1.5	1.9	3.2	1.7	2.4×10^1	2.0
	$n = 200, \mathbf{b} = (1, 0)^T$				$n = 400, \mathbf{b} = (1, 0)^T$			
ϵ	1	10^{-4}	10^{-8}	10^{-12}	1	10^{-4}	10^{-8}	10^{-12}
$\kappa(S)$	2.8×10^2	2.0×10^5	2.0×10^9	2.0×10^{13}	5.7×10^2	4.0×10^5	4.0×10^9	4.2×10^{13}
$\kappa(\widehat{S})$	2.8	1.6	1.5	1.0	3.2	2.2	1.0	1.0

Table 3: Approximation on the Schur complements for model problem (16)

one of our chosen example problems in fact possess this behavior in its linearized form. The two direction vectors correspond to the two well-known rigid-body modes. Let $u = (u_1 \ u_2)$. One of the rigid-body modes is such that all the discretized u_1 nodes are 1 and all the discretized u_2 nodes are 0; and the other is such that all the discretized u_1 nodes are 0 and all the discretized u_2 nodes are 1. Table 4 shows the PCG convergence history and the condition number of $\widehat{A} = R^{-T}AR^{-1}$, where R is the approximate semi-separable Cholesky factor. It is clear that with higher ratio λ/μ , the system is much more ill-conditioned, and requires much many more PCG iterations. When $\lambda/\mu = 1$, preserving directions and increasing block/rank size are beneficial. When $\lambda/\mu = 10^4$, preserving directions is generally beneficial, but larger block/rank size is not helpful.

(λ, μ)	h^{-1}	$p = 8, \ r = 2d + 2$				$p = 20, \ r = 2d + 10$			
		$d = 0$	$\kappa(\widehat{A})$	$d = 2$	$\kappa(\widehat{A})$	$d = 0$	$\kappa(\widehat{A})$	$d = 2$	$\kappa(\widehat{A})$
(1.0, 1.0)	8	32	1.5×10^1	25	9.7×10^1	16	2.9×10^1	11	1.9×10^1
	16	62	6.4×10^2	48	4.7×10^2	64	8.6×10^2	31	2.0×10^2
	32	123	2.5×10^3	92	1.7×10^3	83	3.0×10^3	62	1.2×10^3
(1.0, 10^{-4})	8	243	3.1×10^5	236	3.5×10^5	12	1.3×10^1	9	1.3×10^1
	16	549	1.1×10^6	440	9.7×10^5	1230	1.7×10^6	1203	2.0×10^6
	32	1216	4.5×10^6	1258	4.3×10^6	1867	7.0×10^6	1996	8.6×10^6

Table 4: Number of PCG iterations and the condition number $\kappa(\widehat{A} = R^{-T}AR^{-1})$ for the elasticity equations.

For the elasticity problem, we also examined the last Schur complement matrix arising from direct Cholesky factorization with nested dissection ordering. This time, we allow the maximum off-diagonal rank to be at most 4 in the semi-separable representation, and still require our compression scheme to preserve the two rigid-body modes. The results are shown in Table 5. This time, $\kappa(\widehat{S})$ hovers around 1, even when S is ill-conditioned.

To summarize, our results show that for both diffusion and elasticity problems, our direction-preserving factorization method is very efficient and achieves very good approximation for the Schur complement matrices corresponding to the top level separator. Our future main goal is

	$n = 200$				$n = 400$			
λ/μ	1	10^4	10^8	10^{12}	1	10^4	10^8	10^{12}
$\kappa(S)$	1.7×10^2	2.2×10^5	2.2×10^9	1.5×10^{13}	3.3×10^2	4.5×10^5	4.5×10^9	1.8×10^{13}
$\kappa(\tilde{S})$	1.6	2.1	2.1	2.0	2.0	2.4	2.4	2.2

Table 5: Preconditioner effectiveness on the Schur complements for the elasticity equations.

to use this factorization algorithm to construct reduced (Schur complement) matrices that have prescribed actions on certain direction vectors and not as much as stand-alone preconditioners (as explained in the beginning of the introduction of this paper).

4 Conclusions

We presented an efficient and backward stable algorithm for constructing SPD semi-separable matrices that approximate a given dense SPD matrix A with a guaranteed a priori given tolerance $\tau > 0$. In the literature, there are several different classes of semi-separable matrices that have similar low-rank structures [5, 4, 13, 16, 17]. Work has begun to extend our algorithm to such low-rank structures. Ultimately, such algorithms will be used to form the basis of efficient algorithms to construct effective preconditioners for sparse matrices arising from discretized PDEs.

Alternatively, giving-up on the guaranteed tolerance property, the proposed algorithm provides a SPD factorized matrix that has the same actions as the original SPD matrix on a given set of direction vectors. More generally, the proposed algorithm has the property that it provides approximate Schur complement (reduced) matrices that have the same actions as the corresponding exact Schur complements on a given set of direction vectors. The latter property offers the potential to construct coarse matrices for algebraic multigrid methods which is a topic of future research.

References

- [1] M. Brezina, R. Falgout, S. MacLachlan, T. Manteuffel, S. McCormick, and J. Ruge. Adaptive algebraic multigrid. *SIAM J. Sci. Comput.*, 27:1261–1286, 2006.
- [2] S. Chandrasekaran, P. Dewilde, M. Gu, T. Pals, X. Sun, A.-J. van der Veen, and D. White. Some fast algorithms for sequentially semi-separable representations. *SIAM J. Mat. Anal. Appl.*, 27(2):341–364, 2005.
- [3] S. Chandrasekaran and M. Gu. A fast and stable solver for recursively semi-separable systems of equations. In Vadim Olshevsky, editor, *Structured matrices in mathematics, computer science and engineering II*, Contemporary Mathematics. AMS, 2001.
- [4] S. Chandrasekaran, M. Gu, and W. Lyons. A fast and stable adaptive solver for hierarchically semi-separable representations. To appear in *Calcolo*, 2006.

- [5] S. Chandrasekaran, M. Gu, and T. Pals. A fast *ULV* decomposition solver for hierarchically semi-separable representations. *SIAM J. Mat. Anal. Appl.*, 28:603–622, 2006.
- [6] P. Dewilde and D. Alpay. Time-varying signal approximation and estimation. In *Proc. Int. Symposium MTNS-89, Vol. III*, pages 1–22. Birkhäuser Verlag, 1990.
- [7] P. Dewilde and A. van der Veen. *Time-varying systems and computations*. Kluwer Academic Publishers, New York, 1998.
- [8] T. Dupont, R. P. Kendall, and Jr. H. H. Rachford. An approximate factorization procedure for solving self-adjoint elliptic difference equations. *SIAM Journal on Numerical Analysis*, 5:559–573, 1968.
- [9] H. Dym, D. Alpay, and P. Dewilde. Lossless inverse scattering and reproducing kernels for upper triangular matrices. In *Operator Theory Advances and Applications*, pages 61–135. Birkhäuser Verlag, 1990.
- [10] A. George. Nested dissection of a regular finite element mesh. *SIAM J. Num. Anal.*, 10:345–363, 1973.
- [11] I. Gohberg and Y. Eidelman. A modification of the Dewilde van der Veen method for inversion of finite structured matrices. *Linear Algebra and its Applications*, 343:419–450, 2001.
- [12] I. Gustafsson. A class of first order factorization methods. *BIT*, 18:142–156, 1978.
- [13] W. Hackbusch. A sparse arithmetic based on $\langle$ -matrices. part-I: Introduction to $\langle$ -matrices. *Computing*, 62:89–108, 1999.
- [14] T. Kailath. Fredholm resolvents, wiener-hopf equations, and riccati differential equations. *IEEE Trans. on Information Theory*, IT-15(6), 1969.
- [15] M. S. Khairi, G. L. Miller, and T. J. Sheffler. Nested dissection: a survey and comparison of various nested dissection algorithms. Technical Report Technical Report CMU-CS-106R, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1992.
- [16] B. N. Khoromskij and W. Hackbusch. A sparse arithmetic based on $\langle$ -matrices. part-II: application to multi-dimensional problems. *Computing*, 64:21–47, 2000.
- [17] B. N. Khoromskij and W. Hackbusch. A sparse $\langle$ -matrix arithmetic: general complexity estimates. *Journal of Computational and Applied Mathematics*, 125:79–501, 2000.
- [18] I. Koltracht, I. Gohberg, and T. Kailath. Linear complexity algorithms for semiseparable matrices. *Integral Equations and Operator Theory*, 8:780–804, 1985.
- [19] W. Lyons, H. D. Cecineros, S. Chandrasekaran, and M. Gu. Fast algorithms for spectral collocation with non-periodic boundary conditions. *Journal of Computational Physics*, 207(1):173–191, 2005.

- [20] Y. Notay. DRIC: a dynamic version of the RIC method. *Numerical Linear Algebra with Applications*, 1(6):511–532, 1994.
- [21] A.J. van der Veen. Approximate inversion of a large semiseparable positive matrix. In *Proc. 17th Int. Symp. on Mathematical Theory of Networks and Systems (MTNS-04), Brussels (BE)*, July 2004.
- [22] A.J. van der Veen and P. Dewilde. Inner-outer factorization and the inversion of locally finite systems of equations. *Linear Algebra and its Applications*, 313:53–100, 2000.