Linear Algebra Appl. ••• (••••) •••–•••

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization

Nicola Mastronardi [a,*,1], Paul Van Dooren [b,2]

[a] *Istituto per le Applicazioni del Calcolo "M. Picone", sede di Bari, Consiglio Nazionale delle Ricerche, Via G. Amendola, 122/D, I-70126 Bari, Italy*
[b] *Catholic University of Louvain, Department of Mathematical Engineering, Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium*

ARTICLE INFO

Article history:
Received 7 February 2015
Accepted 3 September 2015
Available online xxxx
Submitted by V. Mehrmann

MSC:
65F30
65F35
15A23

Keywords:
Indefinite symmetric matrix
Rank revealing
Inertia

ABSTRACT

We present an algorithm for computing a symmetric rank-revealing decomposition of a symmetric $n \times n$ matrix A, as defined in the work of Hansen & Yalamov [9]: we factorize the original matrix into a product $A = QMQ^T$, with Q orthogonal and M symmetric and in block form, with one of the blocks containing the dominant information of A, such as its largest eigenvalues. Moreover, the matrix M is constructed in a form that is easy to update when adding to A a symmetric rank-one matrix or when appending a row and, symmetrically, a column to A: the cost of such an updating is $O(n^2)$ floating point operations.

The proposed algorithm is based on the block anti-triangular form of the original matrix M, as introduced by the authors in [11]. Via successive orthogonal similarity transformations this form is then updated to a new form $A = \hat{Q}\hat{M}\hat{Q}^T$, whereby the first k rows and columns of $\hat{M}$ have elements bounded by a given threshold τ and the remaining bottom

* Corresponding author.
E-mail addresses: n.mastronardi@ba.iac.cnr.it (N. Mastronardi), paul.vandooren@uclouvain.be (P. Van Dooren).

[1] The work of this author was partially supported by the GNCS-INdAM project "Metodi di regolarizzazione per problemi di ottimizzazione e applicazioni".
[2] The work of this author is partly supported by the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office and by the Short Term Mobility program 2015 of Consiglio Nazionale delle Ricerche, Italy.

http://dx.doi.org/10.1016/j.laa.2015.09.037

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

right part of $\hat{M}$ is maintained in block anti-triangular form. The updating transformations are all orthogonal, guaranteeing the backward stability of the algorithm, and the algorithm is very economical when the near rank deficiency is detected in some of the anti diagonal elements of the block anti-triangular form. Numerical results are also given showing the reliability of the proposed algorithm.

1. Introduction

Rank-revealing decompositions of dense matrices are widely used in applications such as signal and image processing, where accurate and reliable computation of the numerical rank, as well as the numerical range and null space, are required [4,12]. In such applications it is crucial to compute in a fast and reliable way an updating of such a factorization when a row or a column is appended/deleted to the initial matrix (updating/downdating) or when the initial matrix is modified by a symmetric rank-one matrix (rank-one modification).

For the unsymmetric case many rank-revealing algorithms have been proposed in the literature based on the QR factorization and URV decomposition [4,3,1,2,6]. The singular value decomposition (SVD) is of course a decomposition that reveals the numerical rank, but in general updatings or rank-one modifications cannot be computed in an efficient way [9,5].

In many applications the underlying matrix is symmetric [9,5] and it is therefore useful to consider rank revealing factorizations exploiting this symmetry. Recently, a new factorization of symmetric indefinite matrices $A = QMQ^T$, with Q orthogonal and M block-anti-triangular (BAT) has been introduced [11]. In particular, given a symmetric indefinite matrix $A \in \mathbb{R}^{n\times n}$ with inertia (n_-, n_0, n_+), the following decomposition can be computed,

$$A = QMQ^T, \quad M = \begin{array}{c} \left[\begin{array}{cccc} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & Y^T \\ \mathbf{0} & \mathbf{0} & X & Z^T \\ \mathbf{0} & Y & Z & W \end{array}\right] \end{array} \begin{array}{l} \}n_0 \\ \}n_1 \\ \}n_2 \\ \}n_1 \end{array} \tag{1}$$

with $Q \in \mathbb{R}^{n\times n}$ orthogonal, $Z \in \mathbb{R}^{n_1\times n_2}$, $W \in \mathbb{R}^{n_1\times n_1}$ symmetric, $Y \in \mathbb{R}^{n_1\times n_1}$ nonsingular lower anti-triangular and $X \in \mathbb{R}^{n_2\times n_2}$ symmetric definite if $n_2 > 0$, i.e., $X = \omega LL^T$ with L lower triangular and

$$\omega = \begin{cases} 1, & \text{if } n_+ > n_- \\ -1, & \text{if } n_+ < n_- \end{cases}.$$

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

When $n_+ = n_-$, ω is not defined. Hence, X is symmetric positive definite if $\omega = 1$ and is symmetric negative definite if $\omega = -1$. This BAT factorization of a symmetric matrix can be efficiently updated/downdated for a symmetric rank-one modification [11]. Moreover, a BAT decomposition modified by appending to it one more row and, symmetrically, one more column, can be updated in $O(n^2)$ floating point operations, where n is the matrix order.

In this paper we describe an algorithm that uses the BAT factorization as a preprocessing step and computes a factorization of the form $\hat{Q}\hat{M}\hat{Q}^T$, where $\hat{Q}$ is an orthogonal matrix and $\hat{M}$ is a rank-revealing block anti-triangular one,

$$\hat{M} = \left[\begin{array}{c|ccc} \hat{M}_{11} & & \hat{M}_{21}^T & \\ \hline & \mathbf{0} & \mathbf{0} & \hat{Y}^T \\ \hat{M}_{21} & \mathbf{0} & \hat{X} & \hat{Z}^T \\ & \hat{Y} & \hat{Z} & \hat{W} \end{array}\right] \tag{2}$$

with $\|\hat{M}_{11}\|_2 + \|\hat{M}_{12}\|_2 \approx \tau$ and all the entries of the main diagonal of the anti-triangular matrix $\hat{Y}$ greater than τ in absolute value, with τ a fixed tolerance. Such a decomposition can then be exploited in various applications where an approximation of the numerical rank and/or range are needed.

The paper is organized as follows. In Section 2 the notations together with known results used in the manuscript are listed. In Section 3 the main ideas, on which the rank-revealing algorithm is based, are described followed by the section of numerical examples and conclusions.

2. Notation, definitions and known results

In this section we describe the notation, definitions and known results used in the manuscript.

- Matrices are indicated with upper-case letters $A, B, \ldots$, vectors with bold lower-case letters $\mathbf{u}, \mathbf{v}, \ldots$, scalars with lower-case letters $\alpha, \beta, \ldots a, b, \ldots$ Moreover, the element i, j of a matrix A is denoted by a_{ij} and the subvector of elements $i, i+1, \ldots, i+j$ of a vector $\mathbf{b}$ is denoted by $\mathbf{b}_{i:i+j}$.
- The identity matrix of order n is denoted by I_n, while the zero matrix of size (m, n) is denoted by $\mathbf{0}_{mn}$ or simply by $\mathbf{0}$ if there is no ambiguity.
- The eigenvalues of a symmetric matrix $A \in \mathbb{R}^{n\times n}$ are denoted by
$$\lambda_i(A),\ \ i = 1, \ldots, n, \quad \text{with} \quad |\lambda_i(A)| \geq |\lambda_{i+1}(A)|,\ \ i = 1, \ldots, n-1.$$
If there is no ambiguity, the eigenvalues of A are denoted simply by λ_i, $i = 1, \ldots, n$.
- The singular values of a matrix $A \in \mathbb{R}^{m\times n}$; $m \geq n$ are denoted by
$$\sigma_i(A),\ \ i = 1, \ldots, n, \quad \text{with} \quad \sigma_i(A) \geq \sigma_{i+1}(A),\ \ i = 1, \ldots, n-1.$$

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

If there is no ambiguity, the singular values of A are denoted simply by σ_i, $i = 1, \ldots, n$.
- The ith vector of the canonical basis of $\mathbb{R}^n$ is denoted by $\mathbf{e}_i^{(n)}$, or just $\mathbf{e}_i$ if there is no ambiguity.
- Depending on the context, a Givens rotation can be either a 2×2 orthogonal matrix

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix}, \quad c^2 + s^2 = 1$$

or the matrix of order n

$$\begin{bmatrix} I_{i-1} & & & \\ & c & s & \\ & -s & c & \\ & & & I_{n-i-1} \end{bmatrix}.$$

- The numerical rank k of A, with respect to the threshold τ, is defined as the number of singular values greater than or equal to τ, i.e., $\sigma_k(A) \geq \tau \geq \sigma_{k+1}(A)$ [8, Section 3.1].
- The factorization

$$A = QMQ^T, \tag{3}$$

of a symmetric matrix $A \in \mathbb{R}^{n \times n}$, with $Q = \left[Q_0 \,\middle|\, Q_R \right] \in \mathbb{R}^{n \times n}$ orthogonal, $Q_0 \in \mathbb{R}^{n \times (n-k)}$, $Q_R \in \mathbb{R}^{n \times k}$,

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{12}^T & M_{22} \end{bmatrix}, \tag{4}$$

$M_{11} \in \mathbb{R}^{(n-k) \times (n-k)}$, $M_{22} \in \mathbb{R}^{k \times k}$, $M_{12} \in \mathbb{R}^{k \times (n-k)}$, and $0 \leq k \leq n$, is said to be rank-revealing[3] [9] if

$$\operatorname{cond}(M_{22}) \simeq \sigma_1/\sigma_k \ \text{ and } \ \|M_{11}\|_F^2 + \|M_{12}\|_F^2 \simeq \sigma_{k+1}^2 + \cdots + \sigma_n^2.$$

- The inertia of a symmetric matrix $A \in \mathbb{R}^{n \times n}$ is defined as the triple $\operatorname{Inertia}(A) = (n_-, n_0, n_+)$, where n_-, n_0 and n_+ are the number of negative, zero and positive eigenvalues of A, respectively, and $n_- + n_0 + n_+ = n$ [7].
- The *numerical* inertia of a symmetric matrix $A \in \mathbb{R}^{n \times n}$ with respect to $\tau \in \mathbb{R}$, $\tau > 0$, is defined as the quadruple $\operatorname{Inertia}(A, \tau) = (n_\tau, n_-, n_0, n_+)$, where n_0, n_τ, n_-, and n_+ are, respectively, the number of zero eigenvalues of A, the number of non zero eigenvalues of A that, in absolute value, are smaller than τ, the number of negative

[3] The definition of rank-revealing factorization introduced in [9] has the block matrices in (4) permuted.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix} \begin{bmatrix} \times & \times & 0 & \times & \times & \varepsilon & \varepsilon \\ \otimes & \times & 0 & 0 & \varepsilon & 0 & \varepsilon \end{bmatrix} = \begin{bmatrix} \times & \times & 0 & \times & \times & \varepsilon & \varepsilon \\ & \times & 0 & \times & \times & \varepsilon & \varepsilon \end{bmatrix}$$

Fig. 1. Application of a Givens rotation to a matrix. (For interpretation of the references to colour in this figure, the reader is referred to the web version of this article.)

eigenvalues smaller than $-\tau$, and the number of positive eigenvalues of A greater than τ, and $n_\tau + n_- + n_0 + n_+ = n$.

- If A and E are two symmetric matrices of order n then

$$\sum_{i=1}^{n} (\lambda_i(A+E) - \lambda_i(A))^2 \leq \|E\|_F^2, \tag{5}$$

$$|\lambda_i(A+E) - \lambda_i(A)| \leq \|E\|_2, \quad i = 1, \dots, n, \tag{6}$$

$$\lambda_i(A) + \lambda_n(E) \leq \lambda_i(A+E) \leq \lambda_i(A) + \lambda_1(E), \quad i = 1, \dots, n. \tag{7}$$

The latter result is known as the Wielandt–Hoffman theorem [7, p. 442].

- Let $A = \begin{bmatrix} \mathbf{a}_1 | \cdots | \mathbf{a}_n \end{bmatrix} \in \mathbb{R}^{m \times n}$ be a column partitioning with $m \geq n$. If $A_r = \begin{bmatrix} \mathbf{a}_1 | \cdots | \mathbf{a}_r \end{bmatrix}$, $r = 1, \dots, n-1$, then

$$\sigma_1(A_{r+1}) \geq \sigma_1(A_r) \geq \sigma_2(A_{r+1}) \geq \cdots \geq \sigma_r(A_{r+1}) \geq \sigma_r(A_r) \geq \sigma_{r+1}(A_{r+1}). \tag{8}$$

- If $A \in \mathbb{R}^{m \times n}$, $m \geq n$ and $\omega = \pm 1$, then

$$|\lambda_i(\omega A^T A)| = \sigma_i^2(A), \quad i = 1, \dots, n. \tag{9}$$

Before considering the rank revealing algorithm, we analyze, in the same fashion described in [12], the effects of applying a Givens rotation to a matrix $A \in \mathbb{R}^{2 \times 7}$ with entries of different size. The Givens rotation has been chosen to introduce a zero in position $(2, 1)$. In general, in the figures of the manuscript we denote by $\otimes$ the entry of the matrix to be annihilated by the multiplication of the Givens rotation while the entries modified by the multiplication are in red. Moreover, $\times$'s, 0's, and ε's, respectively represent non zero entries, zero entries, and tiny entries (below a fixed threshold τ) of the matrix. Fig. 1 shows the modification of the matrix multiplied by the Givens matrix. In particular,

R.1: a pair of $\times$'s remains a pair of $\times$'s;
R.2: a $\times$ and a 0 are replaced by a pair of $\times$'s;
R.3: a $\times$ and an ε are replaced by a pair of $\times$'s;
R.4: a pair of ε's remains a pair of ε's;
R.5: an ε and a 0 are replaced by a pair of ε's.

We observe that two tiny entries remain tiny after the multiplication due to the orthogonality of the Givens matrix.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

3. Rank-revealing block anti-triangular factorization

The aim of this section is to describe an algorithm that computes a rank-revealing factorization of type (3), after having computed the BAT factorization (1) of a symmetric indefinite matrix $A = QMQ^T$ by the BAT algorithm described in [11].

Without loss of generality, we assume the matrix M nonsingular and the central block X positive definite. In fact, the singularities of A are already detected by the reduction of the original matrix into BAT form by the BAT algorithm.

When we thus consider the symmetric indefinite matrix

$$M = \begin{bmatrix} \mathbf{0} & \mathbf{0} & Y^T \\ \mathbf{0} & X & Z^T \\ Y & Z & W \end{bmatrix} \begin{matrix} \}n_1 \\ \}n_2 \\ \}n_1 \end{matrix} \tag{10}$$

where Z, W are symmetric, $Y \in \mathbb{R}^{n_1 \times n_1}$ is lower anti-triangular and $X \in \mathbb{R}^{n_2 \times n_2}$ symmetric definite, if one of the entries of the main anti-diagonal of Y is zero, then the whole matrix M is singular [11]. The latter entry can be removed from the main anti-diagonal of Y, transforming M into the BAT form (1).

In Subsection 3.1 we show that if an entry of the main anti-diagonal of Y is, in absolute value, below a fixed threshold τ, then the matrix M can be modified via orthogonal similarity transformations into the matrix (2).

Let

$$\hat{M}^{(1)} = \left[\begin{array}{c|ccc} 0 & & 0 & \\ \hline & \mathbf{0} & \mathbf{0} & {Y^{(1)}}^T \\ 0 & \mathbf{0} & X^{(1)} & {Z^{(1)}}^T \\ & Y^{(1)} & Z^{(1)} & W^{(1)} \end{array}\right], \quad E = \left[\begin{array}{c|c} M_{11}^{(1)} & {M_{21}^{(1)}}^T \\ \hline M_{21}^{(1)} & \\ & \end{array}\right].$$

Then, by (7),

$$\lambda_i(\hat{M}^{(1)}) \leq \lambda_i(M) \leq \lambda_i(\hat{M}^{(1)}) + \tau, \; i = 1, \ldots, n.$$

Hence, if an entry of the main anti-diagonal of Y in (10) is, in absolute value, less than τ, then $\lambda_n(M) \leq \tau$, since $\lambda_n(\hat{M}^{(1)}) = 0$. Moreover, we have transformed the matrix (10) into a matrix of the form (4).

3.1. Removal of tiny entries of the main anti-diagonal of Y

As previously discussed, if an entry of the main anti-diagonal of Y in a BAT matrix M is, in absolute value, below a fixed threshold τ, then so is $\lambda_n(M)$. In this subsection we show how to transform such a BAT matrix into form (2). For the sake of brevity, we describe the effects of algorithm only on the lower anti-triangular matrix Y. The extension of the algorithm to the whole matrix M is straightforward.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

$$\underset{(a)}{\begin{bmatrix} & & & & \times \\ & & & \varepsilon & \times \\ & & \times & \times & \times \\ & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(b)}{\begin{bmatrix} & & & & \times \\ & & & \varepsilon & \times \\ & & \otimes & \times & \times \\ & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(c)}{\begin{bmatrix} & & & & \times \\ & & \varepsilon & \varepsilon & \times \\ & & & \times & \times \\ & \otimes & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(d)}{\begin{bmatrix} & & & & \times \\ & \varepsilon & \varepsilon & \varepsilon & \times \\ & & & \times & \times \\ & & \times & \times & \times \\ \otimes & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(e)}{\begin{bmatrix} & & & & \times \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \times \\ & & & \times & \times \\ & & \times & \times & \times \\ & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(f)}{\begin{bmatrix} & & & & \times \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \times \\ & & & \times & \times \\ & & \times & \times & \times \\ & \times & \times & \times & \times \end{bmatrix}} \Rightarrow$$

$$\underset{(g)}{\begin{bmatrix} & & & & \times \\ & & & \times & \times \\ & & \times & \times & \times \\ & \otimes & \times & \times & \times \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \times \end{bmatrix}} \Rightarrow \underset{(h)}{\begin{bmatrix} & & & & \times \\ & & & \times & \times \\ & & \otimes & \times & \times \\ \varepsilon & & \times & \times & \times \\ \varepsilon & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(i)}{\begin{bmatrix} & & & & \times \\ & & & \otimes & \times \\ \varepsilon & & & \times & \times \\ \varepsilon & & \times & \times & \times \\ \varepsilon & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(j)}{\begin{bmatrix} & & & & \otimes \\ \varepsilon & & & & \times \\ \varepsilon & & & \times & \times \\ \varepsilon & & \times & \times & \times \\ \varepsilon & \times & \times & \times & \times \end{bmatrix}} \Rightarrow \underset{(k)}{\begin{bmatrix} \varepsilon & & & & \\ \varepsilon & & & & \times \\ \varepsilon & & & \times & \times \\ \varepsilon & & \times & \times & \times \\ \varepsilon & \times & \times & \times & \times \end{bmatrix}}$$

Fig. 2. Removal of tiny entries of the main anti-diagonal of Y by multiplications of Givens rotations and a permutation matrix.

Let us suppose the element in position $(n_1 - k + 1, k)$ of the main anti-diagonal of $Y \in \mathbb{R}^{n_1 \times n_1}$ (and, hence, the element $(n_0 + 2n_1 + n_2 - k + 1, n_0 + k)$ of M) is less than τ in absolute value.

A sequence of Givens rotations $\hat{G}_i^{(a)} \in \mathbb{R}^{n_1 \times n_1}$, $i = 1, \ldots, k-1$, are applied to the right of Y. Each $\hat{G}_i^{(a)}$ acts on the columns $k - i$ and $k - i + 1$ of $Y\hat{G}_1^{(a)^T} \cdots \hat{G}_{i-1}^{(a)^T}$ annihilating the element in position $(n_1 - k + i + 1, k - i)$. Let

$$Y := Y\hat{G}_1^{(a)^T} \cdots \hat{G}_{k-1}^{(a)^T}. \tag{11}$$

This step is graphically depicted in Fig. 2, $(a) \Rightarrow (e)$, for $n_1 = 5$ and $k = 4$. After this step, we observe that, by R.4, all the entries in the $(n_1 - k + 1)$th row of Y, in columns $1, \ldots, k - 1$, are less than τ in absolute value. Thus, the row $n_1 - k + 1$ is moved after the last row of Y by means of the multiplication by a permutation matrix P (Fig. 2, $(f) \Rightarrow (g)$). Let

$$Y := PY. \tag{12}$$

A second sequence of Givens rotations $\tilde{G}_i^{(b)} \in \mathbb{R}^{n_1 \times n_1}$, $i = 1, \ldots, n_1 - 1$, is now applied to the left of Y. Each $\tilde{G}_i^{(b)}$ acts on the rows $n_1 - i$ and $n_1 - i + 1$ annihilating the entry in position $(n_1 - i, i + 1)$ of $\tilde{G}_{i-1}^{(b)} \cdots \tilde{G}_1^{(b)} Y$. Let $Y := \tilde{G}_{n_1-1}^{(b)} \cdots \tilde{G}_1^{(b)} Y$. This step is graphically depicted in Fig. 2, $(g) \Rightarrow (k)$. We observe that, because of R.4, all the entries in the first column of the new Y are less than τ in absolute value. Moreover $Y(2 : n_1, 2 : n_1)$ is in lower anti-triangular form. Let

$$Q^{(1)} = G_{n_1-1}^{(b)} \cdots G_1^{(b)} P G_{k-1}^{(a)} \cdots G_1^{(a)}$$

and

$$Q := QQ^{(1)^T}, \tag{13}$$

where

$$G_i^{(a)} = \begin{bmatrix} \hat{G}_i^{(a)} & \\ & I_{n_1+n_2} \end{bmatrix}, \ i = 1, \ldots, k-1, \quad G_i^{(b)} = \begin{bmatrix} \tilde{G}_i^{(b)} & \\ & I_{n_1+n_2} \end{bmatrix}, \ i = 1, \ldots, n_1 - 1.$$

Then $M^{(1)} = Q^{(1)} M Q^{(1)^T}$ and $A = QM^{(1)}Q^T$, with $M^{(1)}$ is in form (4) and the submatrix $M^{(1)}(2:n, 2:n)$ in a BAT form.

Let $Y^{(1)} = Y(2:n_1, 2:n_1)$ and $\mathbf{v}_\varepsilon = Y(1:n_1, 1)$, i.e., $\mathbf{v}_\varepsilon$ is made by the tiny entries of Y moved in its first column. After this step, the matrix $M^{(1)}$ has the following structure:

$$M^{(1)} = \left[\begin{array}{c|c|cc|c} & & & \varepsilon & \mathbf{v}_\varepsilon^T \\ \hline & & & & Y^{(1)^T} \\ \hline & & X^{(1)} & \mathbf{v} & \\ \varepsilon & & \mathbf{v}^T & \gamma & Z^{(1)^T} \\ \hline \mathbf{v}_\varepsilon & Y^{(1)} & Z^{(1)} & & W^{(1)} \end{array}\right].$$

To reduce $M^{(1)}(2:n, 2:n)$ in BAT form, we need to analyze if the submatrix

$$\hat{X}^{(1)} = \begin{bmatrix} X^{(1)} & \mathbf{v} \\ \mathbf{v}^T & \gamma \end{bmatrix} \tag{14}$$

is definite, being $X^{(1)}$ definite.

If $\hat{X}^{(1)}$ is definite, we only need to extend the Cholesky factorization of $X^{(1)} = \omega L^{(1)} L^{(1)^T}$, $L^{(1)} \in \mathbb{R}^{n_2 \times n_2}$, to $\hat{X}^{(1)} = \omega \hat{L}^{(1)} \hat{L}^{(1)^T}$, $\hat{L}^{(1)} \in \mathbb{R}^{(n_2+1)\times(n_2+1)}$. It turns out that

$$\begin{aligned} \hat{L}^{(1)}_{1:n_2,1:n_2} &= L^{(1)}, \\ \hat{L}^{(1)}_{n_2+1,1:n_2} &= \omega (L^{(1)} \backslash \mathbf{v})^T, \\ \hat{L}^{(1)}_{n_2+1,n_2+1} &= \sqrt{\omega\gamma - \hat{L}^{(1)}_{n_2+1,1:n_2} \hat{L}^{(1)^T}_{n_2+1,1:n_2}}. \end{aligned}$$

Hence, we set $Inertia(M, \tau) = (n_\tau + 1, n_- - 1, n_0, n_+)$.

If $\hat{X}^{(1)}$ is indefinite, we need to transform it into the following form (see [11, § 2.1])

$$\begin{bmatrix} & & \beta \\ & \tilde{X}^{(1)} & \tilde{\mathbf{v}} \\ \beta & \tilde{\mathbf{v}}^T & \tilde{\gamma} \end{bmatrix},$$

with $\tilde{X}^{(1)} \mathbb{R}^{(n_2-1)\times(n_2-1)}$. Hence, we set $Inertia(M, \tau) = (n_\tau + 1, n_-, n_0, n_+ - 1)$.

In both cases, the new matrix $M^{(1)}$ can be computed with $O(n^2)$ floating point operations.

3.1.1. Computational complexity

The multiplication of M by the first sequence of Givens rotations $G_i^{(a)}$, $i = 1, \ldots, k-1$, requires $O(k^2)$ floating point operations, while the multiplication of the new M by the second sequence of Givens rotations $G_i^{(b)}$ requires $O(n_1(n_1+n_2))$ floating point operations. The updating of the orthogonal matrix Q in (13) requires $O(n_1(n_1+n_2))$ floating point operations.

3.2. Completing the rank-revealing BAT factorization

As already emphasized in [4] for the unsymmetric rank-revealing QR factorization, even though M has $n-k$ singular values smaller than τ, it is very well possible that less than $n-k$ entries in main anti-diagonal of Y are smaller than τ.

Let $j \le n-k$ be the numerical rank of (2) detected after having removed the entries of Y less than τ in absolute value, i.e., j is the number is of columns of $M_{11}^{(1)}$ and $M_{21}^{(1)}$. We now show how the full rank-revealing BAT factorization can be retrieved.

The smallest eigenvalue $\lambda_{n-j}(M_{22}^{(1)})$ and the corresponding eigenvector $\mathbf{u}_{n-j}(M_{22}^{(1)})$ of $M_{22}^{(1)} \in \mathbb{R}^{(n-j)\times(n-j)}$, can be computed via inverse iteration, requiring $O(n_1^2+n_2^2)$ floating point operations [10]. In the same paper an algorithm is described for transforming $M_{22}^{(1)}$ into $M_{22}^{(2)}$, where

$$M_{22}^{(2)} = \hat{U}^{(2)^T} M_{22}^{(1)} \hat{U}^{(2)} = \left[\begin{array}{c|ccc} \lambda_{n-j}(M_{22}^{(1)}) & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{0} & Y^{(2)^T} \\ \mathbf{0} & \mathbf{0} & X^{(2)} & Z^{(2)^T} \\ \mathbf{0} & Y^{(2)} & Z^{(2)} & W^{(2)} \end{array}\right], \tag{15}$$

with $\hat{U}^{(2)}$ orthogonal such that $\hat{U}^{(2)}\mathbf{u}_{n-j}(M_{22}^{(1)}) = \pm\mathbf{e}_{n-j}^{(1)}$. This can be accomplished with $O(n_1^2+n_2^2)$ floating point operations [10].

Let

$$U^{(2)} = \begin{bmatrix} I_j & \\ & \hat{U}^{(2)} \end{bmatrix}.$$

Then $A = Q^{(2)} M^{(2)} Q^{(2)^T}$, where $Q^{(2)} = Q\hat{U}^{(2)}$ and

$$M^{(2)} = \left[\begin{array}{c|ccc} M_{11}^{(2)} & & M_{21}^{(2)^T} & \\ \hline & \mathbf{0} & \mathbf{0} & Y^{(2)^T} \\ M_{21}^{(2)} & \mathbf{0} & X^{(2)} & Z^{(2)^T} \\ & Y^{(2)} & Z^{(2)} & W^{(2)} \end{array}\right] \begin{array}{l} \}j+1 \\ \left.\begin{array}{l} \\ \\ \\ \end{array}\right\} n-j-1. \end{array} \tag{16}$$

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

We observe that the anti-triangular matrix $Y^{(2)}$ has been modified by the orthogonal transformation $U^{(2)}$. Hence, it can happen that, if $j+1 < n-k$, one of the entries in its main anti-diagonal is less than τ in absolute value. Hence we reapply over and over the procedure to remove the entries on the main anti-diagonal of $Y^{(2)}$ below τ in absolute value followed by the computation of the smallest eigenvalue and the corresponding eigenvector of the BAT submatrix $M_{22}^{(2)}$ and its displacement to the left-up corner of the submatrix until the computed eigenvalue is grater than τ in absolute value.

The described method can be summarized in the following algorithm.

Algorithm Rank-Revealing Block Anti-Triangular

% `input` : $A \in \mathbb{R}^{n\times n}$ symmetric indefinite,
% τ, the tolerance for the numerical rank
% `output` : $M \in \mathbb{R}^{n\times n}$, rank revealing BAT matrix
% $Q \in \mathbb{R}^{n\times n}$ orthogonal such that $A = QMQ^T$
% k_τ, the numerical rank
$[M, Q, \mathit{Inertia}] = \mathtt{BAT}(A)$;
$k_\tau = 0$;
$\mathrm{Inertia}(\mathrm{M}, \tau) = [0, \mathrm{Inertia}]$;
$\mathit{flag} = 0$;
`while` flag $== 0$,
  $[M, Q, \mathit{Inertia}(M, \tau), k_\tau] = \mathtt{move}\varepsilon\mathtt{\_out\_}Y(M, \tau, \mathit{Inertia}(M, \tau), k_\tau)$;
  $[\lambda, \mathbf{u}] = \mathtt{inv\_it}(M(k_\tau + 1 : n, k_\tau + 1 : n))$;
  `if` $|\lambda| < \tau$,
    $[M, Q, \mathit{Inertia}(M, \tau), k_\tau] = \mathtt{move}\lambda\mathtt{\_out}(M, \tau, \mathit{Inertia}(M, \tau), \lambda, \mathbf{u}, k_\tau)$;
  `else`
    $\mathit{flag} = 1$;
  `end`
`end`

The function `BAT`, given as input the symmetric indefinite matrix $A \in \mathbb{R}^{n\times n}$, yields as output $M, Q \in \mathbb{R}^{n\times n}$, and *Inertia*, with M in BAT form, Q orthogonal and Inertia the inertia of A.

The function $\mathtt{move}\varepsilon\mathtt{\_out\_}Y$ having as input M, τ, $\mathit{Inertia}(M, \tau)$, k_τ, with

$$M = \left[\begin{array}{c|ccc} M_{11} & & M_{21}^T & \\ \hline & \mathbf{0} & \mathbf{0} & Y^T \\ M_{21} & \mathbf{0} & X & Z^T \\ & Y & Z & W \end{array}\right] \begin{array}{l} \}k_\tau \\ \left.\begin{array}{l} \\ \\ \\ \end{array}\right\} n - k_\tau \end{array} \tag{17}$$

"moves" all the entries of the anti-diagonal matrix Y smaller than τ in absolute value to M_{21} delivering as output the matrix

$$\hat{M} = \left[\begin{array}{c|ccc} \hat{M}_{11} & & \hat{M}_{21}^T & \\ \hline & \mathbf{0} & \mathbf{0} & \hat{Y}^T \\ \hat{M}_{21} & \mathbf{0} & \hat{X} & \hat{Z}^T \\ & \hat{Y} & \hat{Z} & \hat{W} \end{array}\right] \begin{array}{l} \}\hat{k}_\tau \\ \left.\begin{array}{l} \\ \\ \\ \end{array}\right\} n - \hat{k}_\tau, \end{array}$$

where $k = \hat{k}_\tau - k_\tau$ are the number of entries of the main anti-diagonal of Y detected, the updated orthogonal matrix Q, the updated numerical inertia and rank $\hat{k}_\tau$.

The function `inv_it`, having as input the matrix

$$M_{22} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & Y^T \\ \mathbf{0} & X & Z^T \\ Y & Z & W \end{bmatrix},$$

yields its smallest eigenvalue and the corresponding eigenvector. The function `move`λ`_out`, has as input the matrix (17), $\lambda_{n-\tilde{k}_\tau}(M_{22})$, $\mathbf{u}_{n-\tilde{k}_\tau}(M_{22})$. It has, as output, the matrix

$$\tilde{M} = \left[\begin{array}{c|cccc} \tilde{M}_{11} & & \tilde{M}_{21}^T & & \\ \hline & \lambda_{n-\tilde{k}_\tau}(M_{22}) & & & \\ \tilde{M}_{21} & & \mathbf{0} & \mathbf{0} & \tilde{M}^T \\ & & \mathbf{0} & \tilde{X} & \tilde{Z}^T \\ & & \tilde{Y} & \tilde{Z} & \tilde{W} \end{array}\right] \begin{array}{l} \}\tilde{k}_\tau \\ \left.\begin{array}{l} \\ \\ \\ \\ \end{array}\right\} n - \tilde{k}_\tau, \end{array}$$

the updated orthogonal matrix Q, the updated numerical inertia and rank $\hat{k}_\tau$.

4. Numerical results

Example 1. The purpose of this section is to illustrate the behavior of the described algorithm. All the experiments were carried out in `matlab`.

We considered 1000 test matrices of size $n = 100$, constructed as the product $A = QDQ^T$, where Q is a random orthogonal matrix generated by the function `gallery('qmult',`n`)`, D is a diagonal matrix with diagonal elements $d_{ii} = \pm\sigma_i$, $i = 1, \ldots, n$, where the sign is chosen randomly. Furthermore, σ_i are geometrically distributed between σ_1 and σ_k and also between σ_{k+1} and σ_n, with $k = 80$, $\sigma_1 = 1$, $\sigma_k = 10^{-5}$, $\sigma_{k+1} = 10^{-7}$, $\sigma_n = 10^{-10}$. Hence, the numerical rank ρ_τ with respect $\tau = 10^{-6}$ is 80. Moreover, $\sigma_k/\sigma_{k+1} = 100$, $\kappa_2(M_{n-k+1:n,n-k+1:n}) = 10^5$ and $\sqrt{\sum_{i=k+1}^n \sigma_i^2} = 1.39 \times 10^{-7}$. The distribution of the singular values of the considered matrices is displayed in Fig. 3 (on the right in logarithmic scale).

After computing the rank-revealing BAT matrix

$$M = \begin{bmatrix} M_{11} & M_{21}^T \\ M_{21} & M_{22} \end{bmatrix}$$

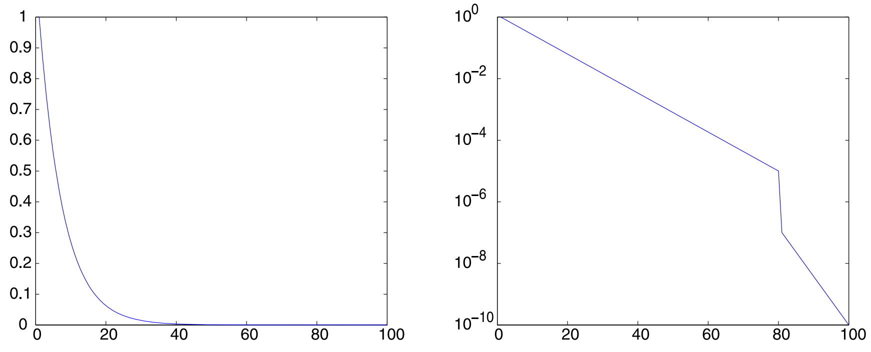

Fig. 3. Distribution of the singular values of the matrices considered in Example 1 (log scale, right).

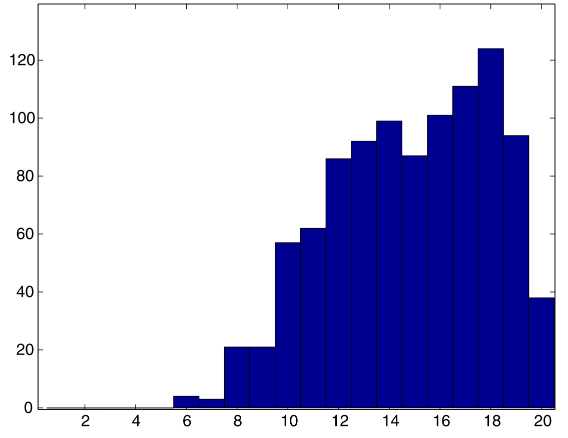

Fig. 4. Number of times (y axis), over 1000 runs, ρ_{τ_1} columns of M_{21} (x axis) are detected by the function `move`ε`_out`. (For interpretation of the references to colour in this figure, the reader is referred to the web version of this article.)

with $M_{11} \in \mathbb{R}^{(n-\rho_\tau)\times(n-\rho_\tau)}$, $M_{21} \in \mathbb{R}^{\rho_\tau\times(n-\rho_\tau)}$, $M_{22} \in \mathbb{R}^{\rho_\tau\times\rho_\tau}$, where ρ_τ is the numerical rank wit respect to τ delivered by the algorithm, we construct the matrix

$$\bar{M}_{22} = \begin{bmatrix} M_{11} & M_{21}^T \\ M_{21} & \mathbf{0} \end{bmatrix}.$$

We denote by ρ_{τ_1} the number of columns of M_{21} in (2) "detected" by the function `move`ε`_out` and by ρ_{τ_2} the number of columns of M_{21} "detected" by the function `move`λ`_out` such that $\rho_\tau = \rho_{\tau_1} + \rho_{\tau_2}$.

In the histogram in Fig. 4 the height of each blue column denotes how many times, over 1000 runs, the same number of ρ_{τ_1} columns of M_{21} are detected by the function `move`ε`_out`. This means that each time ρ_{τ_1} columns are detected by `move`ε`_out`, $\rho_{\tau_2} = n - \rho_\tau - \rho_{\tau_1}$ are computed by the function `move`λ`_out`. We observe that, on average, the number of columns detected by the function `move`ε`_out` is much larger than the number of columns detected by the function `move`λ`_out`.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

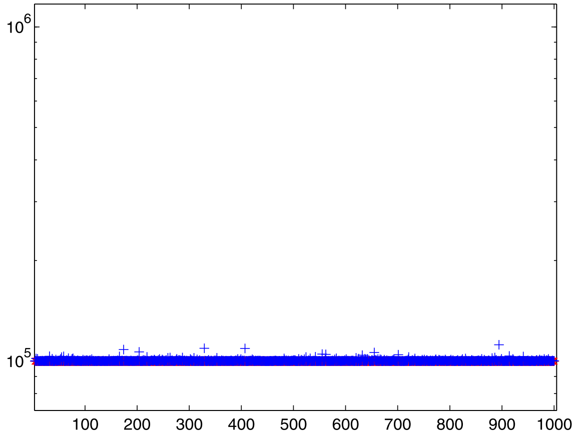

Fig. 5. Condition numbers of the submatrices M_{22} (blue) and ratio σ_1/σ_k (red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

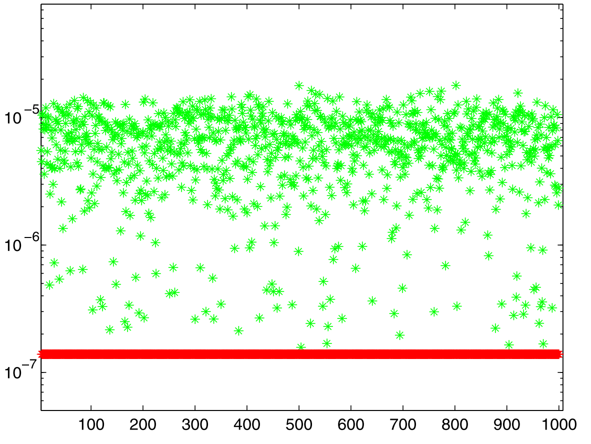

Fig. 6. Frobrenius norms of the computed submatrices $\bar{M}_{22}$ (green) and $\sqrt{\sum_{i=n-k+1}^{n} \sigma_i^2}$ (red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In Fig. 5 the condition numbers of the computed submatrices M_{22} are depicted in blue, while the constant line in red represents the ratio σ_1/σ_k.

For all the considered matrices, the computed numerical rank is $\rho_\tau = 80$.

In Fig. 6 the Frobrenius norms of the computed submatrices $\bar{M}_{22}$ are depicted in green, while the constant line in red denotes $\sqrt{\sum_{i=n-k+1}^{n} \sigma_i^2}$.

5. Conclusions

A rank-revealing algorithm, based on the block anti-triangular factorization of indefinite symmetric matrices, has been proposed. The computed factorization can be easily updated when added a rank-one matrix or one more row and column is appended. The numerical results show the reliability of the proposed algorithm.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037

References

[1] C.H. Bischof, P.C. Hansen, Structure preserving and rank-revealing QR factorizations, SIAM J. Sci. Statist. Comput. 12 (1991) 1332–1350.

[2] C.H. Bischof, P.C. Hansen, A block algorithm for computing rank-revealing QR factorizations, Numer. Algorithms 2 (1992) 371–392.

[3] C.H. Bischof, G. Quintana-Ortí, Computing rank-revealing QR factorizations of dense matrices, ACM Trans. Math. Software 24 (2) (1998) 226–253.

[4] T.F. Chan, Rank revealing QR factorizations, Linear Algebra Appl. 88–89 (1987) 67–82.

[5] R. Fierro, P. Hansen, Truncated VSV solutions to symmetric rank-deficient problems, BIT 42 (3) (2002) 531–540.

[6] R.D. Fierro, P.C. Hansen, Accuracy of TSVD solutions computed from rank-revealing decompositions, Numer. Math. 70 (4) (1995) 453–471.

[7] G.H. Golub, C.F. Van Loan, Matrix Computations, 4th ed., Johns Hopkins University Press, Baltimore, 2013.

[8] P.C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion, SIAM, Philadelphia, 1998.

[9] P.C. Hansen, P.Y. Yalamov, Computing symmetric rank-revealing decompositions via triangular factorization, SIAM J. Matrix Anal. Appl. 23 (2001) 443–458.

[10] N. Mastronardi, P. Van Dooren, Recursive approximation of the dominant eigenspace of an indefinite matrix, J. Comput. Appl. Math. 236 (2012) 4090–4104.

[11] N. Mastronardi, P. Van Dooren, The antitriangular factorization of symmetric matrices, SIAM J. Matrix Anal. Appl. 34 (2013) 173–196.

[12] G. Stewart, An updating algorithm for subspace tracking, IEEE Trans. Signal Process. 40 (1992) 1535–1541.

Please cite this article in press as: N. Mastronardi, P. Van Dooren, Rank-revealing decomposition of symmetric indefinite matrices via block anti-triangular factorization, Linear Algebra Appl. (2015), http://dx.doi.org/10.1016/j.laa.2015.09.037