Journal of Computational and Applied Mathematics 259 (2014) 281–293

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Computing eigenvalues of normal matrices via complex symmetric matrices[☆]

Micol Ferranti *, Raf Vandebril

Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3000 Leuven, Belgium

ARTICLE INFO

Article history:
Received 13 September 2012
Received in revised form 29 August 2013

Keywords:
Normal matrix
Complex symmetric
Takagi factorization
Unitary similarity
Symmetric singular value decomposition
Eigenvalue decomposition

ABSTRACT

Computing all eigenvalues of a modest size matrix typically proceeds in two phases. In the first phase, the matrix is transformed to a suitable condensed matrix format, sharing the eigenvalues, and in the second stage the eigenvalues of this condensed matrix are computed. The main purpose of this intermediate matrix is saving valuable computing time. Important subclasses of normal matrices, such as the Hermitian, skew-Hermitian and unitary matrices admit a condensed matrix represented by only $\mathcal{O}(n)$ parameters, allowing subsequent low-cost algorithms to compute their eigenvalues. Unfortunately, such a condensed format does not exist for a generic normal matrix.

We will show, under modest constraints, that normal matrices also admit a memory cheap intermediate matrix of tridiagonal complex symmetric form. Moreover, we will propose a general approach for computing the eigenvalues of a normal matrix, exploiting thereby the normal complex symmetric structure. An analysis of the computational cost and numerical experiments with respect to the accuracy of the approach are enclosed. In the second part of the manuscript we will investigate the case of nonsimple singular values and propose a theoretical framework for retrieving the eigenvalues. We will, however, also highlight some numerical difficulties inherent to this approach.

1. Introduction

Most of the so-called direct eigenvalue methods are based on a two step approach. First the original matrix is transformed to a suitable shape taking $\mathcal{O}(n^3)$ operations, followed by the core method computing the eigenvalues of this suitable shape, e.g., divide-and-conquer, MRRR, *QR*-methods [1–3]. Consider, e.g., the *QR*-method; starting with an arbitrary unstructured matrix, one first performs a unitary similarity transformation to obtain a Hessenberg matrix in $\mathcal{O}(n^3)$ operations. Next, successive *QR*-steps, which cost $\mathcal{O}(n^2)$ each, are performed until all eigenvalues are revealed.

For some subclasses of normal matrices, e.g., Hermitian, skew-Hermitian, and unitary matrices, the intermediate matrix shapes admit a low storage cost $\mathcal{O}(n)$ and, as such, permit the design of *QR*-algorithms with linear complexity steps [2,4,5]. Unfortunately, for the generic normal matrix class, the intermediate structure is of Hessenberg form, requiring $\mathcal{O}(n^2)$ storage and resulting in a quadratic cost for each *QR*-step. An alternative intermediate condensed form might thus result in significant computational savings. To achieve this goal we propose the use of intermediate complex symmetric

☆ The research was partially supported by the Research Council KU Leuven, projects OT/11/055 (Spectral Properties of Perturbed Normal Matrices and their Applications), CoE PFV/10/002 Optimization in Engineering (OPTEC), by the Fund for Scientific Research—Flanders (Belgium) project G034212N (Reestablishing smoothness for matrix manifold optimization via resolution of singularities) and by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office, Belgian Network DYSCO (Dynamical Systems, Control, and Optimization).

* Corresponding author.
E-mail addresses: micol.ferranti@cs.kuleuven.be (M. Ferranti), raf.vandebril@cs.kuleuven.be (R. Vandebril).
URL: http://people.cs.kuleuven.be/raf.vandebril/ (R. Vandebril).

0377-0427/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cam.2013.08.036

matrices that can be constructed using unitary similarities. The problem of determining whether a square complex matrix is unitarily similar to a complex symmetric one has been intensively studied; see, for instance, [6–8]. Such a similarity always exists for normal matrices [9, Corollary 4.4.4]. One method to perform the unitary transformation of a normal matrix to complex symmetric form was proposed by Ikramov in [10]. It utilizes the Toeplitz decomposition of the normal matrix and symmetries at the same time the two Hermitian terms.

The aim of this article is to provide an initial theoretical basis on which we can continue to build numerical algorithms. Again we rely on the two-step principle: First, the matrix is transformed by a unitary similarity transformation to block matrix form, of which the diagonal blocks are complex symmetric. In the simplest case only one block exists [11], and well-known techniques [12–14] can be used to compute the symmetric singular value decomposition (SSVD), also called Autonne–Takagi factorization [15,9], of this complex symmetric matrix. Based on the SSVD one can retrieve the eigenvalue decomposition. When multiple blocks are present, it is possible to use the same techniques and diagonalize all blocks at once, obtaining a sparse matrix with all blocks diagonal. After that, a specifically designed version of the Jacobi method for normal matrices [16–18] is used, in order to annihilate the last nonzero off-diagonal entries. Numerical experiments illustrate the effectiveness of the proposed method. Whenever the number of block exceeds one, it will be shown, however, that severe numerical difficulties can appear. More precisely, many articles and authors rely on the property that an irreducible Hermitian tridiagonal matrix cannot have coinciding eigenvalues. Though theoretically correct, this statement might fail in a numerical setting, with nonnegligible impact on the accuracy of the proposed methods (see Section 6 or the discussion in [3, Section 5.45]).

In this article, the following notation is used: A^T refers to the transpose of A, $\overline{A}$ to the conjugate of A and $A^H = \overline{A}^T$ denotes the Hermitian conjugate. With $A(i : j, \ell : k)$ the submatrix of a matrix A consisting of rows i up to and including j and columns ℓ up to and including k is depicted. With a_i we refer to the i-th column of A. A matrix is said to be symmetric if $A = A^T$ and Hermitian if $A = A^H$. A matrix is real orthogonal if $AA^T = A^TA = I$ and A is real, and unitary if $AA^H = A^HA = I$. We might use the expressions real and complex symmetric to stress that the symmetric matrix is real or possibly complex. The elements of a matrix A are denoted by a_{ij}, when taking subblocks out of a partitioned matrix, we refer to them as $A_{k\ell}$. The square root of -1 is denoted by $\imath$.

The article is organized in two main parts. One part of the article discusses the easy setting in which the intermediate matrix is of complex symmetric form. The second part of the article presents a theoretical approach to deal with the block form, and discusses possible numerical issues. Section 2 recapitulates some known results on normal matrices, the singular value decomposition and results from [11]. In Section 3, under some constraints, the theoretical setting for eigenvalue computations of normal matrices whose distinct eigenvalues have distinct absolute values is considered. The unitary similarity transformation as well as the link with the SSVD is presented to reveal the eigenvalues. Section 4 supports the theoretical discussion by numerical experiments. In Section 5 the generic nonsingular case, is investigated. The similarity transformation will now result in a block structured matrix, which can be diagonalized efficiently. The eigenvalues of this latter sparse block matrix are then computed via a Jacobi-like diagonalization procedure. In Section 6 some numerical experiments and observations with respect to the latter structure are presented. We also compare the performance of our method with that of [10] in relation to different distributions of eigenvalues and singular values: we show that both methods can suffer from discrepancies between their theoretical and practical behavior.

2. Preliminaries

This section highlights some essential properties of normal matrices, the singular value decomposition, and some other results required in the remainder of the text.

A singular value decomposition of A is a factorization of the form $A = U\Sigma V^H$, where U, V are unitary matrices, and Σ is a diagonal matrix with nonnegative real entries $\sigma_1, \ldots, \sigma_n$, we write $\Sigma = \text{diag}(\sigma_1, \ldots, \sigma_n)$. The diagonal elements of Σ are called the singular values of A and the columns of U and V are called the left and right singular vectors respectively. A singular value σ_i is said to be a multiple singular value if it appears more than once on the diagonal of Σ. A standard choice consists of ordering the singular values such that $\sigma_1 \geq \cdots \geq \sigma_n$ [1]. We will implicitly assume that every singular value decomposition in this article has this conventional form, except when stated otherwise, and we will stress this by naming it an unordered singular value decomposition. It is well-known that the singular value decomposition for a matrix with n distinct singular values is essentially unique [1], which signifies unique up to unimodular scaling. The unordered version is also unique up to permutations of the diagonal element as long as the singular values are unique.

Suppose that the matrix has singular values of multiplicities exceeding one, so that uniqueness is lost. One then still has uniqueness of the subspaces associated with equal singular values, as given by the following theorem.

Lemma 1 (*Autonne's Uniqueness Theorem, Theorem 2.6.5 in [9]*). *Let $A \in \mathbb{C}^{n\times n}$ and let $A = U\Sigma V^H = W\Sigma Z^H$ be two, possibly distinct, singular value decompositions. Then there exist unitary block diagonal matrices $B = \text{diag}(B_1, B_2, \ldots, B_d)$ and $\tilde{B} = \text{diag}(\tilde{B}_1, \tilde{B}_2, \ldots, \tilde{B}_d)$, such that $U = WB$, $V = Z\tilde{B}$ and $B_i = \tilde{B}_i$ whenever the associated singular value differs from zero.*

In general, given a matrix $A \in \mathbb{C}^{n\times n}$ and a singular value decomposition $A = U\Sigma V^H$, we have that $AA^H = U\Sigma^2U^H$ and $A^HA = V\Sigma^2V^H$ are eigenvalue decompositions of AA^H and A^HA respectively, having orthonormal eigenvectors. If $A \in \mathbb{C}^{n\times n}$ is normal, then $AA^H = A^HA$, so the columns of U and V both form a basis of $\mathbb{C}^n$, made out of eigenvectors of the same matrix.