Computing matrix-vector products with centrosymmetric and centrohermitian matrices *

Heike Fassbender † Khakim D. Ikramov ‡

Abstract

An algorithm proposed recently by A. Melman reduces the costs of computing the product Ax with a symmetric centrosymmetric matrix A as compared to the case of an arbitrary A. We show that the same result can be achieved by a simpler algorithm, which requires only that A be centrosymmetric. However, if A is hermitian or symmetric, this can be exploited to some extent. Also, we show that similar gains are possible when A is a skew-centrosymmetric or a centrohermitian matrix.

AMS classification: 65F05

Keywords: Centrosymmetric matrix; Skew-centrosymmetric matrix; Centrohermitian matrix; Matrix-vector multiplication

1 Introduction

The conventional algorithm for computing the matrix-vector product Ax, where both A and x are of dimension n, requires n^2 multiplications and $n^2 - n$

*This work was supported by Deutsche Forschung Gemeinschaft grant 436 RUS 17/103/01.

†Zentrum Mathematik, Technische Universität München, D-80290 München, Germany

‡Faculty of Computational Mathematics and Cybernetics, Moscow State University, 119899 Moscow, Russia

additions, or approximately $2n^2$ flops. An algorithm proposed recently by A. Melman [2] computes Ax with a real symmetric centrosymmetric A, using $\frac{1}{2}n^2 + n$ multiplications and $\frac{3}{4}n^2 + n$ additions, that is, altogether $\frac{5}{4}n^2 + O(n)$ flops. Each additional multiplication by the same matrix costs only $n^2 + O(n)$ flops.

Recall that a (generally complex) $n \times n$ matrix A is said to be centrosymmetric if

$$\pi_n A \pi_n = A, \tag{1}$$

where

$$\pi_n = \begin{pmatrix} & & & 1 \\ & & \cdot & \\ & \cdot & & \\ & & & \\ 1 & & & \end{pmatrix}. \tag{2}$$

It is well known (although is not used in [2]) that there exists a very simple similarity transformation which makes every centrosymmetric $n \times n$ matrix a direct sum of two blocks of (roughly, for n odd) half the order. This transformation does not require that A be (real) symmetric or hermitian. However, if these properties are present, they are preserved.

In this short paper, we show that the use of the fact above leads to an algorithm for multiplying a centrosymmetric matrix A by a vector x, which is simpler than Melman's algorithm but ensures the same savings in computational costs. If A is real symmetric or hermitian, this can be exploited for further savings. Moreover, similar gains are possible when A is a skew-centrosymmetric or a centrohermitian matrix.

All the formulas become slightly more complicated when n is odd. For simplicity, we restrict ourselves to the case of even n.

We note that the motivation for the algorithm in [2] was the need to compute products Tx in a method for solving systems $Tz = b$ with a (real) symmetric Toeplitz matrix T. Such a matrix T is, of course, centrosymmetric. However, this seems to be not a very good motivation, because every Toeplitz $n \times n$ matrix (not necessarily symmetric or hermitian) can be multiplied by a vector at the cost of $O(n \log n)$ flops (see, for example, [4, p.194] and the references therein).

2 Preliminaries

For definiteness, matrices throughout the paper are assumed to be complex. When necessary, we make remarks about real matrices.

Suppose that A is a centrosymmetric matrix of order $n = 2m$. Partitioning A into four $m \times m$ blocks,

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}, \tag{3}$$

we observe that A is fully determined by its upper block row because

$$A_{21} = \pi_m A_{12} \pi_m \qquad \text{and} \qquad A_{22} = \pi_m A_{11} \pi_m. \tag{4}$$

Define

$$Q = \frac{1}{\sqrt{2}} \begin{pmatrix} I_m & I_m \\ \pi_m & -\pi_m \end{pmatrix}. \tag{5}$$

It is easy to see that Q is a (real) orthogonal matrix. Moreover, the similarity transformation

$$A \to B = Q^{-1} A Q \tag{6}$$

takes A to the block diagonal matrix

$$B = B_1 \oplus B_2, \tag{7}$$

where

$$B_1 = A_{11} + A_{12} \pi_m \tag{8}$$

and

$$B_2 = A_{11} - A_{12} \pi_m. \tag{9}$$

Note that Q is independent of a specific matrix A. The first m columns in Q are (the simplest possible) symmetric vectors; that is, vectors x such that

$$x = \pi_n x.$$

The last m columns in Q are skew-symmetric vectors; i.e., they satisfy the equation

$$x = -\pi_n x.$$

Combined, relations (5)–(7) are just a reformulation of the well-known fact: the linear subspace of symmetric vectors and that of skew-symmetric vectors are invariant w.r.t. any centrosymmetric matrix A.

Since Q is orthogonal, the matrix B is real symmetric or hermitian if A is. As a consequence, the blocks B_1 and B_2 in (7)–(9) are also symmetric or hermitian.

Definition. A $n \times n$ matrix A is said to be skew-centrosymmetric if

$$\pi_n A \pi_n = -A. \tag{10}$$

For a skew-centrosymmetric matrix A with partition (3), we have

$$A_{21} = -\pi_m A_{12} \pi_m \qquad \text{and} \qquad A_{22} = -\pi_m A_{11} \pi_m, \tag{11}$$

rather than relations (4). As a consequence, the similarity transformation (6) applied to a skew-centrosymmetric A produces a block matrix B of the form

$$B = \begin{pmatrix} 0 & B_2 \\ B_1 & 0 \end{pmatrix} \tag{12}$$

with B_1 and B_2 given, respectively, by (8) and (9).

Observe that, if A is real symmetric or hermitian, the matrices (8) and (9), in general, are not. However, they do satisfy the relation

$$B_2 = B_1^T$$

or

$$B_2 = B_1^*,$$

making B a symmetric or a hermitian matrix.

Definition. A $n \times n$ matrix A is said to be centrohermitian if

$$\pi_n A \pi_n = \bar{A}. \tag{13}$$

For a centrohermitian matrix A with partition (3), we have

$$A_{21} = \pi_m \bar{A}_{12} \pi_m \qquad \text{and} \qquad A_{22} = \pi_m \bar{A}_{11} \pi_m. \tag{14}$$

Define

$$Q = \frac{1}{\sqrt{2}} \begin{pmatrix} I_m & iI_m \\ \pi_m & -i\pi_m \end{pmatrix}. \tag{15}$$

It is easy to see that Q is a unitary matrix. The similarity transformation in (6) takes a centrohermitian A to the *real* matrix

$$B = \begin{pmatrix} \Re(A_{11} + \bar{A}_{12}\pi_m) & -\Im(A_{11} + \bar{A}_{12}\pi_m) \\ \Im(A_{11} - \bar{A}_{12}\pi_m) & \Re(A_{11} - \bar{A}_{12}\pi_m) \end{pmatrix}. \tag{16}$$

This matrix is symmetric if A is hermitian.

Remark. The most widely known example of centrohermitian matrices are hermitian Toeplitz matrices. It was shown in [5] that every matrix of the latter type is unitarily similar to a real Toeplitz-plus-Hankel matrix. To our knowledge, the general term "centrohermitian matrices" has first appeared in [1].

3 Algorithms

In this section, we state and discuss algorithms for calculating the product Ax, where A is, respectively, a centrosymmetric or a centrohermitian matrix.

3.1 Centrosymmetric matrices

Assume that a centrosymmetric matrix A and a vector x of an even dimension $n = 2m$ are given. Partition x in accordance with (3):

$$x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}. \tag{17}$$

To find the vector $u = Ax$, do the following:

Preparatory step Find the matrices

$$B_1 = A_{11} + A_{12}\pi_m \qquad \text{and} \qquad B_2 = A_{11} - A_{12}\pi_m.$$

Step 1 Find the vector (see (5))

$$y = \sqrt{2}Q^{-1}x = \sqrt{2}Q^T x =$$

$$\begin{pmatrix} I_m & \pi_m \\ I_m & -\pi_m \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 + \pi_m x_2 \\ x_1 - \pi_m x_2 \end{pmatrix} \equiv \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}.$$

Step 2 Find the vector

$$z = By = \begin{pmatrix} B_1 y_1 \\ B_2 y_2 \end{pmatrix} \equiv \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}.$$

Step 3 Find the vector

$$u = \frac{1}{\sqrt{2}} Qz =$$

$$\frac{1}{2} \begin{pmatrix} I_m & I_m \\ \pi_m & -\pi_m \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} z_1 + z_2 \\ \pi_m(z_1 - z_2) \end{pmatrix}.$$

Now we calculate the number of operations in this algorithm. Each entry in B_1 and B_2 is a sum or a difference of two entries in A. It follows that the preparatory step takes $n^2/2$ additions/ subtractions. Step 1 amounts to n additions/subtractions. The same amount of additive operations plus n divisions by 2 are required on Step 3. Step 2 is the calculation of two matrix-vector products of dimension $n/2$. By the conventional algorithm, it takes $n^2/2$ multiplications and about the same number of additions. On the whole, we need $n^2/2 + O(n)$ multiplicative operations and $n^2 + O(n)$ additive ones to compute $u = Ax$. This seems to be worse than $\frac{5}{4}n^2 + O(n)$ flops required in Melman's algorithm. Note, however, that, unlike the latter, our algorithm does not assume A to be symmetric. If A *is* real symmetric, then we need to compute only $\approx n^2/4$ entries on the preparatory step, which reduces the total number of operations to $\frac{5}{4}n^2 + O(n)$, the same amount as in Melman's algorithm. The same number of *complex* operations are required in the case of a hermitian A.

Suppose that B_1 and B_2 computed on the preparatory step are stored. If, later, we need to calculate the product $\tilde{u} = A\tilde{x}$ with the same matrix A, then only Steps 1 to 3 in the algorithm must be repeated for the new vector x. This means that each additional product with the matrix A takes only $n^2/2+O(n)$ multiplications and $n^2+O(n)$ additions to compute. Again, this is the same amount as in Melman's algorithm.

Remark. In the introduction of [2], the author, after stating that n^2 multiplications and $\approx n^2$ additions are usually required for a general matrix-vector product of dimension n, then adds: "This remains true even for a symmetric matrix." However, suppose that one needs to compute many products with the same symmetric or Hermitian matrix A. Then, as shown by Mou [3], after some preparatory work amounting to performing $\approx n^2/2$

subtractions, roughly half the multiplications in each product $u = Ax$ can be replaced by the same number of additions. If the times for multiplication and addition differ considerably for the computer used, this may lead to significant gains in the overall performance.

For the benefit of the reader, we explain the basic idea of Mou's algorithm, using a fourth-order example with a symmetric 4×4 matrix A. First, we replace A by the matrix $\hat{A}$ with the same off-diagonal entries as in A and the diagonal entries

$$\hat{a}_{11} = a_{11} - a_{12} - a_{13} - a_{14},$$
$$\hat{a}_{22} = a_{22} - a_{12} - a_{23} - a_{24},$$
$$\hat{a}_{33} = a_{33} - a_{13} - a_{23} - a_{34},$$
$$\hat{a}_{44} = a_{44} - a_{14} - a_{24} - a_{34}.$$

Then, instead of the conventional formulas

$$u_i = \sum_{j=1}^{4} a_{ij} x_j, \qquad i = 1, 2, 3, 4$$

we compute $u_1, \ldots, u_4$ from the formulas

$$u_1 = \hat{a}_{11} x_1 + a_{12}(x_1 + x_2) + a_{13}(x_1 + x_3) + a_{14}(x_1 + x_4),$$
$$u_2 = a_{12}(x_1 + x_2) + \hat{a}_{22} x_2 + a_{23}(x_2 + x_3) + a_{24}(x_2 + x_4),$$
$$u_3 = a_{13}(x_1 + x_3) + a_{23}(x_2 + x_3) + \hat{a}_{33} x_3 + a_{34}(x_3 + x_4),$$
$$u_4 = a_{14}(x_1 + x_4) + a_{24}(x_2 + x_4) + a_{34}(x_3 + x_4) + \hat{a}_{44} x_4.$$

Here, once a partial product of the form $a_{ij}(x_i + x_j)$ has been computed, it is immediately added to the intermediate sums for the two components u_i and u_j.

Remark. The case when A is a skew-centrosymmetric matrix differs from what was said above by obvious and insignificant details.

3.2 Centrohermitian matrices

Let A be a centrohermitian matrix of order $n = 2m$, and x be a vector partitioned as in (17). To find the vector $u = Ax$, do the following:

Preparatory step Find the matrices

$$B_1 = A_{11} + \bar{A}_{12}\pi_m \qquad \text{and} \qquad B_2 = A_{11} - \bar{A}_{12}\pi_m.$$

Form the (real) matrix

$$B = \begin{pmatrix} \Re B_1 & -\Im B_1 \\ \Im B_2 & \Re B_2 \end{pmatrix}.$$

Step 1 Find the vector (see (15))

$$y = \sqrt{2}Q^{-1}x = \sqrt{2}Q^{*}x =$$

$$\begin{pmatrix} I_m & \pi_m \\ -iI_m & i\pi_m \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 + \pi_m x_2 \\ i(\pi_m x_2 - x_1) \end{pmatrix} \equiv y_R + iy_I.$$

Step 2 Find the real matrix-vector products

$$z_R = By_R \qquad \text{and} \qquad z_I = By_I$$

and define

$$z = z_R + iz_I = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}.$$

Step 3 Find the vector

$$u = \frac{1}{\sqrt{2}}Qz =$$

$$\frac{1}{2}\begin{pmatrix} I_m & iI_m \\ \pi_m & -i\pi_m \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \frac{1}{2}\begin{pmatrix} z_1 + iz_2 \\ \pi_m(z_1 - iz_2) \end{pmatrix}.$$

The preparatory step amounts to performing $n^2/2$ complex additions, which is equivalent to n^2 real additions. If A is hermitian, then B is symmetric, and only roughly half the entries in the symmetric blocks $\Re B_1$ and $\Re B_2$ need to be computed. Also, the blocks $\Im B_2$ and $-\Im B_1$ are the transposed versions of each other; thus, only, say, $\Im B_1$ should be computed.

The costs of Steps 1 and 3 are negligible. The two real matrix-vector products on Step 2 take about half of the work required for a single complex matrix-vector product of the same dimension. This says that the algorithm above ensures the same relative savings as the algoritm in the preceding subsection.

References

[1] A. Lee, Centrohermitian and skew-centrohermitian matrices, Linear Algebra Appl. 29 (1980) 205–210.

[2] A. Melman, Symmetric centrosymmetric matrix-vector multiplication, Linear Algebra Appl. 320 (2000) 193–198.

[3] Z.J. Mou, Fast algorithms for computing symmetric/Hermitian matrix-vector products, Electronics Letters 27 (1991) 1272–1274.

[4] E.E. Tyrtyshnikov, A Brief Introduction to Matrix Analysis, Birkhaeuser Boston, Boston, MA, 1997.

[5] D.M. Wilkes, S.D. Morgera, F. Noor, M.H. Hayes, A Hermitian Toeplitz matrix is unitarily similar to a real Toeplitz-plus-Hankel matrix, IEEE Trans. Signal Processing 39 (1991) 2146–2148.