

THE PARAMETERIZED SR ALGORITHM FOR SYMPLECTIC (BUTTERFLY) MATRICES

H. FABBENDER*

Abstract. The SR algorithm is a structure-preserving algorithm for computing the spectrum of symplectic matrices. Any symplectic matrix can be reduced to symplectic butterfly form. A symplectic matrix B in butterfly form is uniquely determined by $4n - 1$ parameters. Using these $4n - 1$ parameters, we show how one step of the symplectic SR algorithm for B can be carried out in $\mathcal{O}(n)$ arithmetic operations compared to $\mathcal{O}(n^3)$ arithmetic operations when working on the actual symplectic matrix. Moreover, the symplectic structure, which will be destroyed in the numerical process due to roundoff errors when working with a symplectic (butterfly) matrix, will be forced by working just with the parameters.

Key words. Symplectic Matrix; Eigenvalue Problem; SR Algorithm.

AMS(MOS) subject classifications. 65F15.

1. Introduction. Symplectic (generalized) eigenvalue problems occur in many applications, e.g., in discrete linear quadratic optimal control, discrete Kalman filtering, the solution of discrete algebraic Riccati equations, discrete stability radii and H_∞ -norm computations (see, e.g., [16, 18] and the references therein) and discrete Sturm-Liouville equations (see, e.g., [5]). The solution of the symplectic (generalized) eigenvalue problem has been the topic of numerous publications during the last 30 years. Even so, a numerically sound method, i.e., a strongly backward stable method in the sense of [6], is not yet known. The numerical computation of an invariant (deflating) subspace is usually carried out by an iterative procedure like the QR (QZ) algorithm; see, e.g., [18, 20]. The QR (QZ) algorithm is numerically backward stable but it ignores the symplectic structure. In order to develop fast, efficient, and reliable methods, the symplectic structure of the problem should be preserved and exploited. Then important properties of symplectic matrices like spectral symmetries will be preserved and not destroyed by rounding errors.

Recently there has been renewed interest in constructing structure-preserving methods for the symplectic eigenproblem [1, 2, 3, 4, 11] based on the SR method [10, 17]. This method is a QR -like method based on the SR decomposition. In an initial step, the $2n \times 2n$ symplectic matrix is reduced to a more condensed form, the symplectic butterfly form, which in general contains $8n - 4$ nonzero entries. As in the general framework of GR algorithms [21], the SR iteration preserves the symplectic butterfly form at each step and converges to a form from which eigenvalues and invariant (deflating) subspaces can be read off. The SR algorithm for symplectic butterfly matrices has been fully described and analyzed in [4, 11]. Due to unavoidable roundoff errors the symplectic butterfly structure will be lost in the numerical process. The very compact butterfly form allows one to restore the symplectic structure whenever necessary.

A $2n \times 2n$ symplectic butterfly matrix is determined by $4n - 1$ parameters. As will be shown in this paper, the SR algorithm can be rewritten in a parameterized form that works with $4n - 1$ parameters instead of the $(2n)^2$ matrix elements in each iteration. Thus only $\mathcal{O}(n)$ arithmetic operations per SR step are needed compared to $\mathcal{O}(n^3)$ arithmetic operations when working on the actual symplectic matrix. Moreover, the symplectic structure, which will be destroyed in the numerical process due

*Universität Bremen, Zentrum für Technomathematik, Fachbereich 3 - Mathematik und Informatik, 28334 Bremen, Germany, email: heike@math.uni-bremen.de

to roundoff errors when working with a butterfly matrix, will be forced by working just with the parameters. No additional action have to be taken as in the course of the symplectic butterfly SR algorithm.

The development of the parameterized butterfly SR algorithm has been guided by the unitary case, which the symplectic case resembles to some degree. There has been an earlier attempt by Flaschka, Mehrmann and Zywitz [12] to exploit this resemblance. They proposed a structure-preserving symplectic SR algorithm for symplectic J -Hessenberg matrices. Such matrices (like symplectic butterfly matrices) depend uniquely on $4n - 1$ parameters. A single shift SR step that is purely based on these parameters is derived in [12]. No numerical results are reported, but the authors note: *It forces the symplectic structure, but it has the disadvantage that it needs $4n - 1$ terms to be nonzero in each step, which makes it highly numerically unstable. ... The numerical instability due to extra $2n$ inversions ... seems an unreasonable price to pay compared with the gains in efficiency.* [12, p. 186, last paragraph].

In this paper we will develop a parameterized SR algorithm for computing the eigeninformation of a symplectic matrix based on the initial reduction to a symplectic butterfly matrix. First we will see that, like unitary Hessenberg matrices, any symplectic butterfly matrix B has a unique factorization exhibiting the $4n - 1$ parameters which uniquely determine B . One step of the SR algorithm with shift polynomial q applied to a matrix $B \in \mathbf{R}^{2n \times 2n}$ may be described as follows: Factor $q(B) = SR$ with S symplectic and R J -triangular. Then put $\tilde{B} = S^{-1}BS$. If B is an unreduced symplectic butterfly matrix, then so is \tilde{B} . Hence, B and \tilde{B} can be given in parameterized form. We will derive formulae which given the $4n - 1$ parameters of B compute the $4n - 1$ parameters which determine \tilde{B} without ever forming B , \tilde{B} or S explicitly. If desired, the transformation matrix S can be computed explicitly. But, unfortunately S does not have the same structure as the matrix being transformed. S is symplectic, but not of butterfly form. Therefore, S can not be given in parameterized form.

In Section 2 unreduced butterfly matrices, the reduction of symplectic matrices to butterfly form and the butterfly SR algorithm are reviewed. Like unitary Hessenberg matrices, symplectic butterfly matrices have a unique factorization exhibiting $4n - 1$ parameters which uniquely determine B . Such factorizations are introduced in Section 3. There we also discuss the basic idea of an implicit SR step that makes use of such a factorization. The details of the parameterized butterfly SR algorithm are presented in Section 4. The overall process is discussed in Section 5. In Section 6 numerical examples are presented.

2. Preliminaries. A matrix $M \in \mathbf{R}^{2n \times 2n}$ is called *symplectic* (or *J -orthogonal*) if

$$(2.1) \quad MJM^T = J$$

(or equivalently, $M^TJM = J$) where

$$(2.2) \quad J = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}$$

and I_n is the $n \times n$ identity matrix. Symplectic matrices are nonsingular; their inverses are given by $M^{-1} = JM^TJ^T$. The spectrum of a symplectic matrix is symmetric with respect to the unit circle. Or, in other words, the eigenvalues of symplectic matrices occur in reciprocal pairs: if λ is an eigenvalue of M with right eigenvector x , then λ^{-1} is an eigenvalue of M with left eigenvector $(Jx)^H$. Further, if $\lambda \in \mathbf{C}$ is an eigenvalue of M , then so are $\bar{\lambda}$, λ^{-1} , $\overline{\lambda^{-1}}$.

function q is chosen (such that $q(B) \in \mathbb{R}^{2n \times 2n}$) and the SR decomposition of $q(B)$ is formed, if possible:

$$q(B) = SR.$$

Then the symplectic factor S is used to perform a similarity transformation on B to yield the next iterate, which we will call \widehat{B} :

$$(2.4) \quad \widehat{B} = S^{-1}BS.$$

If $\text{rank}(q(B)) = 2n$ and B is a symplectic butterfly matrix, then so is \widehat{B} in (2.4) [1, 2]. If $\text{rank}(q(B)) = 2n - \nu =: 2k$ and B is an unreduced symplectic butterfly matrix, then \widehat{B} in (2.4) is of the form (see [3, 11] for a proof)

$$(2.5) \quad \widehat{B} = \left[\begin{array}{c|c} \begin{array}{cc} \diagdown & \\ & \square \end{array} & \begin{array}{c} \parallel \\ \parallel \\ \parallel \end{array} \begin{array}{c} \square \\ \square \end{array} \\ \hline \begin{array}{c} \diagdown \\ \square \end{array} & \begin{array}{c} \parallel \\ \parallel \\ \parallel \end{array} \begin{array}{c} \square \\ \square \end{array} \end{array} \right] = \left[\begin{array}{cc|cc} \widehat{B}_{11} & & \widehat{B}_{13} & \\ & \widehat{B}_{22} & & \widehat{B}_{24} \\ \hline \widehat{B}_{31} & & \widehat{B}_{33} & \\ & \widehat{B}_{42} & & \widehat{B}_{44} \end{array} \right] \begin{array}{l} \} k \\ \} n-k \\ \} k \\ \} n-k \end{array}$$

$\underbrace{\hspace{1.5cm}}_k \quad \underbrace{\hspace{1.5cm}}_{n-k} \quad \underbrace{\hspace{1.5cm}}_k \quad \underbrace{\hspace{1.5cm}}_{n-k}$

where

- $\begin{bmatrix} \widehat{B}_{11} & \widehat{B}_{13} \\ \widehat{B}_{31} & \widehat{B}_{33} \end{bmatrix}$ is a symplectic butterfly matrix and
- the eigenvalues of $\begin{bmatrix} \widehat{B}_{22} & \widehat{B}_{24} \\ \widehat{B}_{42} & \widehat{B}_{44} \end{bmatrix}$ are just the ν shifts that are eigenvalues of B .

The algorithm is made compact and efficient by using Laurent polynomials, instead of standard polynomials, to drive the iterations. The shifts should be chosen according to the generalized Rayleigh-quotient strategy. The resulting algorithm is typically cubic convergent. For a detailed discussion on the choice of the spectral transformation function q , the choice of the shifts and convergence properties see [4, 11].

An algorithm for computing S and R explicitly is presented in [8]. As with explicit QR steps, the expense of explicit SR steps comes from the fact that $q(B)$ has to be computed explicitly. A preferred alternative is the implicit SR step, an analogue to the Francis QR step [13, 14, 15]. The first implicit transformation S_1 is selected so that the first columns of the implicit and the explicit S are equivalent. That is, a symplectic matrix S_1 is determined such that

$$S_1^{-1}q(B)e_1 = \alpha e_1, \quad \alpha \in \mathbb{R}.$$

Applying this first transformation to the butterfly matrix yields a symplectic matrix $S_1^{-1}BS_1$ with almost butterfly form having a small bulge. The remaining implicit transformations perform a bulge-chasing sweep down the subdiagonals to restore the butterfly form. That is, a symplectic matrix S_2 is determined such that $S_2^{-1}S_1^{-1}BS_1S_2$ is of butterfly form again. Banse presents in [1] an algorithm to reduce an arbitrary symplectic matrix to butterfly form. The algorithm uses the following elementary symplectic transformations:

- symplectic Givens transformation

$$G(k, c, s) = \left[\begin{array}{cc|cc} I_{k-1} & & & \\ & c & & s \\ \hline & & I_{n-k} & \\ -s & & & I_{k-1} \\ \hline & & & c & \\ & & & & I_{n-k} \end{array} \right];$$

- symplectic Householder transformation

$$H(k, v) = \left[\begin{array}{c|c} I_{k-1} & P \\ \hline & I_{k-1} \\ & & P \end{array} \right], \text{ where } P = I_{n-k+1} - 2 \frac{vv^T}{v^T v};$$

- symplectic Gauss transformation

$$L(k, c, d) = \left[\begin{array}{cc|cc} I_{k-2} & & & \\ & c & & d \\ & & c & \\ \hline & & & I_{n-k} \\ & & & & I_{k-2} \\ & & & & & c^{-1} \\ & & & & & & c^{-1} \\ & & & & & & & I_{n-k} \end{array} \right].$$

The symplectic Givens and Householder transformations are orthogonal, while the symplectic Gauss transformations are nonorthogonal. Algorithms to compute the entries of the abovementioned transformations can be found, e.g., in [19] and [9]. The Gaussian transformations can be computed such that among all possible transformations satisfying the same purpose, the one with the minimal condition number is chosen.

Let us briefly describe the algorithm to reduce an arbitrary symplectic matrix M to butterfly form. Zeros in the rows of M will be introduced by applying one of the above mentioned transformations from the right, while zeros in the columns will be introduced by applying the transformations from the left. Of course, in order to perform a similarity transformation, the inverse of each transformation applied from the right/left has to be applied from the left/right as well. The basic idea of the algorithm can be summarized as follows

- for $j = 1$ to n
 - bring the j th column of M into the desired form
 - bring the $(n + j)$ th row of M into the desired form

The remaining rows and columns in M that are not explicitly touched during the process will be in the desired form due to the symplectic structure. The algorithm for reducing an arbitrary symplectic matrix to butterfly form as given in [1] can be summarized as given in Table 2.1 (in MATLAB¹-like notation). For simplicity we will assume here that all symplectic Gauss transformations needed exist; on how to proceed otherwise see [11].

¹MATLAB is a trademark of The MathWorks, Inc.

Algorithm: Reduction to Butterfly Form

Given a $2n \times 2n$ symplectic matrix M compute its reduction to butterfly form. M will be overwritten by its butterfly form.

```

for  $j = 1 : n - 1$ 
  for  $k = n : -1 : j + 1$ 
    compute  $G_k$  such that  $(G_k M)_{k+n,j} = 0$ 
     $M = G_k M G_k^T$ 
  end
  if  $j < n - 1$ 
  then compute  $H_j$  such that  $(H_j M)_{j+2:n,j} = 0$ 
     $M = H_j M H_j^T$ 
  end
  compute  $L_{j+1}$  such that  $(L_{j+1} M)_{j+1,j} = 0$ 
   $M = L_{j+1} M L_{j+1}^{-1}$ 
  for  $k = n : -1 : j + 1$ 
    compute  $G_k$  such that  $(M G_k)_{j,k} = 0$ 
     $M = G_k^T M G_k$ 
  end
  if  $j < n - 1$ 
  then compute  $H_j$  such that  $(M H_j)_{j,j+2+n:2n} = 0$ 
     $M = H_j^T M H_j$ 
  end
end

```

TABLE 2.1

Reduction to Butterfly Form

3. The Basic Idea. The key to the development of a butterfly SR algorithm working only on the parameters is the observation that at any point in the implicit SR step only a certain, limited number of rows and columns of the symplectic butterfly matrix is worked on. In the leading part of the intermediate matrices the butterfly form is already retained and is not changed any longer, while the trailing part has not been changed yet. Hence, from the leading part the first parameters of the resulting butterfly matrix can be read off, while from the trailing part the last parameters of the original butterfly matrix can still be read off. Recall the implicit SR step as described in Section 2. The first implicit transformation S_1 is selected in order to introduce a bulge into the symplectic butterfly matrix B . That is, a symplectic matrix S_1 is determined such that

$$S_1^{-1} q(B) e_1 = \alpha e_1, \quad \alpha \in \mathbf{R},$$

where $q(B)$ is an appropriately chosen spectral transformation function. Applying this first transformation to the butterfly matrix yields a symplectic matrix $S_1^{-1} B S_1$ with almost butterfly form having a small bulge. The remaining implicit transformations perform a bulge-chasing sweep down the subdiagonals to restore the butterfly form. That is, a symplectic matrix S_2 is determined such that $S_2^{-1} S_1^{-1} B S_1 S_2$ is of butterfly form again. If B is an unreduced butterfly matrix and $\text{rank}(q(B)) = 2n$,

then $\tilde{B} = S_2^{-1}S_1^{-1}BS_1S_2$ is also an unreduced butterfly matrix. Hence, there will be parameters $\tilde{a}_1, \dots, \tilde{a}_n, \tilde{b}_1, \dots, \tilde{b}_n, \tilde{c}_1, \dots, \tilde{c}_n, \tilde{d}_2, \dots, \tilde{d}_n$ which determine \tilde{B} . During the bulge-chasing sweep the bulge is successively moved down the subdiagonals, one row and one column at a time. Consider for simplicity a double shift implicit SR step. As discussed in [4, 11] for a double shift the shift polynomial $q_2(B) = (B + B^{-1}) - \beta I$ should be chosen where $\beta = \mu + \mu^{-1}$ if $\mu \in \mathbf{R}$ or $\beta = \mu + \bar{\mu}$ for $\mu \in \mathbf{C}, |\mu| = 1$. The shift μ is chosen corresponding to the generalized Rayleigh-quotient strategy. The bulge is introduced by a transformation of the form

$$(3.1) \quad S_1 = \left[\begin{array}{cc|ccc} \alpha & \beta & & & \\ -\beta & \alpha & & & \\ & & I_{n-2} & & \\ \hline & & & \alpha & \beta \\ & & & -\beta & \alpha \\ & & & & & I_{n-2} \end{array} \right].$$

In a slight abuse of notation, we will call matrices of the form (3.1) symplectic Householder transformations in the following, although they are the direct sum of two Givens transformations. Whenever a transformation of the form (3.1) is used in the following, one can just as well use a symplectic Householder transformation as defined in Section 2.

Applying a transformation of the form (3.1) to B to introduce a bulge, results in a matrix of the form

$$S_1^{-1}BS_1 = \left[\begin{array}{cccc|cccc} x & + & & & x & x & + & \\ + & x & & & x & x & x & \\ & & x & & + & x & x & x \\ & & & x & & & x & x \\ & & & & & & & \ddots \\ \hline x & + & & & x & x & + & \\ + & x & & & x & x & x & \\ & & x & & + & x & x & x \\ & & & x & & & x & x \\ & & & & & & & \ddots \end{array} \right].$$

Now a symplectic Givens transformation to eliminate the $(n+2, 1)$ element and a symplectic Gauss transformation to eliminate the $(2, 1)$ element are applied, resulting in

$$\left[\begin{array}{cccc|cccc} x & + & & & x & x & + & \\ x & & & & x & x & x & \\ + & x & & & + & x & x & x \\ & & x & & & & x & x \\ & & & \ddots & & & & \ddots \\ \hline x & + & & & x & x & + & \\ x & & & & x & x & x & \\ + & x & & & + & x & x & x \\ & & x & & x & x & x & \\ & & & \ddots & & & & \ddots \end{array} \right].$$

This bulge is chased down the subdiagonals one row and one column at a time. The $(1, 1)$ and the $(n+1, 1)$ element are not altered in any subsequent transformation. Hence, at this point we can already read off \tilde{a}_1 and \tilde{b}_1 . The bulge-chase is done using the algorithm for reducing a symplectic matrix to butterfly form as given in

Similarly, N can be decomposed

$$Y_n Y_{n-1} \cdots Y_1 J^T = N$$

where

$$Y_k = \left[\begin{array}{ccc|ccc} I_{k-1} & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ & & & I_{n-k-1} & & \\ \hline & -c_k & -d_{k+1} & & I_{k-1} & \\ & -d_{k+1} & & & & 1 \\ & & & & & & I_{n-k-1} \end{array} \right],$$

$$Y_n = \left[\begin{array}{c|c} I_{n-1} & \\ \hline & I_{n-1} \\ & -c_n & 1 \end{array} \right].$$

Because of their special structure, most of the X_k , Y_k , the symplectic Givens transformations G_j , the symplectic Householder transformations H_j , and the symplectic Gauss transformations L_j as defined in Section 2 commute:

$$\begin{aligned} X_j X_k &= X_k X_j && \text{for all } j, k, \\ Y_j Y_k &= Y_k Y_j && \text{for all } j, k, \\ X_j Y_k &= Y_k X_j && \text{for } j \neq k, j \neq k-1, \\ G_j X_k &= X_k G_j && \text{for } j \neq k, \\ H_j X_k &= X_k H_j && \text{for } j \neq k, j \neq k+1, \\ L_j X_k &= X_k L_j && \text{for } j \neq k, j \neq k-1, \\ G_j Y_k &= Y_k G_j && \text{for } j \neq k, j \neq k-1, \\ H_j Y_k &= Y_k H_j && \text{for } j \neq k, j \neq k-1, j \neq k+1, \\ L_j Y_k &= Y_k L_j && \text{for } j \neq k, j \neq k-1, j \neq k+1. \end{aligned}$$

Here we assume that

$$H_k = \text{diag}(I_{k-1}, P, I_{n-k-1}, I_{k-1}, P, I_{n-k-1})$$

where $P \in \mathbb{R}^{2 \times 2}$ is a Givens transformation, as all H_k considered in this section are of this special form. Hence, we can write

$$(3.2) \quad B = X_n Y_n X_{n-1} Y_{n-1} \cdots X_2 Y_2 X_1 Y_1 J^T.$$

Now let us take a closer look at a double shift bulge chase. We will start with an unreduced symplectic butterfly matrix B decomposed as in (3.2). The resulting matrix \tilde{B} will have a decomposition of the same form as B ,

$$\tilde{B} = \tilde{X}_n \tilde{Y}_n \tilde{X}_{n-1} \tilde{Y}_{n-1} \cdots \tilde{X}_2 \tilde{Y}_2 \tilde{X}_1 \tilde{Y}_1 J^T,$$

As noted before, the bulge is introduced by the transformation $S_1^{-1}BS_1$ with a matrix S_1 of the form (3.1). This leads to a matrix of the form

$$S_1^{-1}BS_1 = \left[\begin{array}{ccc|ccc} \otimes & \oplus & & \otimes & \otimes & \oplus \\ \oplus & \otimes & & \otimes & \otimes & \otimes \\ & & x & \oplus & \otimes & \otimes & x \\ & & & & & x & x \\ & & & & & & \ddots \\ \hline \otimes & \oplus & & \otimes & \otimes & \oplus \\ \oplus & \otimes & & \otimes & \otimes & \otimes \\ & & x & & \oplus & \otimes & x \\ & & & & & x & x \\ & & & & & & \ddots \end{array} \right],$$

where x denotes desired entries in the butterfly form, $+$ undesired entries, and \otimes and \oplus desired and undesired elements that are changed by the current transformation. As S_1 is a symplectic Householder transformation, S_1 and most of the factors of B commute:

$$S_1^{-1}BS_1 = X_n Y_n \cdots X_3 Y_3 S_1^{-1} X_2 Y_2 X_1 Y_1 J^T S_1.$$

Since S_1 is unitary and symplectic, we have $S_1^{-1} = S_1^T$ and $J^T S_1 = S_1 J^T$. Hence,

$$S_1^T BS_1 = X_n Y_n \cdots X_3 Y_3 S_1^T X_2 Y_2 X_1 Y_1 S_1 J^T.$$

Next a symplectic Givens transformation G_2 is applied to zero the $(n+2, 1)$ element:

$$G_2 S_1^T BS_1 G_2^T = \left[\begin{array}{ccc|ccc} x & \oplus & & x & \otimes & + \\ \oplus & \otimes & & \otimes & \otimes & \otimes \\ & & \oplus & x & + & \otimes & x & x \\ & & & & & & x & x \\ & & & & & & & \ddots \\ \hline x & \oplus & & x & \otimes & + \\ \mathbf{0} & \otimes & & \otimes & \otimes & \otimes \\ & & \oplus & x & + & \otimes & x & x \\ & & & & & & x & x \\ & & & & & & & \ddots \end{array} \right].$$

As G_2 and most of the factors of B commute and as G_2 is unitary and symplectic (hence, $J^T G_2^T = G_2^T J^T$) we obtain

$$G_2 S_1^T BS_1 G_2^T = X_n Y_n \cdots X_3 Y_3 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T J^T.$$

Now a symplectic Gauss transformation L_2 is chosen to eliminate the $(2, 1)$ element

such that

$$B^{(1)} := L_2 G_2 S_1^T B S_1 G_2^T L_2^{-1} = \left[\begin{array}{ccc|ccc} \otimes & \oplus & & \otimes & \otimes & \oplus \\ 0 & \otimes & & \otimes & \otimes & \otimes \\ & \oplus & x & \oplus & \otimes & x & x \\ & & & & & x & x \\ & & & & & & \ddots \\ \hline \otimes & \oplus & & \otimes & \otimes & \oplus \\ & \otimes & & \otimes & \otimes & \otimes \\ & \oplus & x & \oplus & \otimes & x & x \\ & & & & & x & x \\ & & & & & & \ddots \end{array} \right].$$

At this point the actual bulge, which is chased down the subdiagonal, is formed. That is, now a sequence of symplectic Givens, Householder and Gauss transformations is applied to successively chase the bulge of the above form down the subdiagonal.

L_2 is symplectic, but not unitary. Hence, $J^T L_2^{-1} = L_2^T J^T$. Moreover, as L_2 and most of the factors of B commute, we have

$$B^{(1)} := L_2 G_2 S_1^T B S_1 G_2^T L_2^{-1} = X_n Y_n \cdots X_3 Y_3 L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T J^T.$$

The $(1, 1)$ and the $(n + 1, 1)$ elements of $B^{(1)}$ are not altered by any subsequent transformation. Therefore, at this point we can read off \tilde{a}_1 and \tilde{b}_1 of the final \tilde{B} . In other words, we can rewrite

$$L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T J^T$$

in terms of \tilde{X}_1 times an appropriate symplectic matrix Z_1 times J^T . That is,

$$(3.3) \quad L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T J^T = \tilde{X}_1 Z_1 J^T,$$

where Z_1 is symplectic. Moreover, as \tilde{X}_1 commutes with $X_n, \dots, X_3, Y_n, \dots, Y_3$ we obtain

$$B^{(1)} = \tilde{X}_1 X_n Y_n \cdots X_3 Y_3 Z_1 J^T.$$

Now the bulge is chased down the subdiagonals one row and one column at a time. This is done using the algorithm for reducing a symplectic matrix to butterfly form as given in Table 2.1. First a symplectic Givens transformation is applied to eliminate the $(n + 1, 2)$ element. This yields

$$G_2^T B^{(1)} G_2 = \left[\begin{array}{ccc|ccc} x & 0 & & x & \otimes & + \\ & \otimes & & \otimes & \otimes & \otimes \\ & \oplus & x & + & \otimes & x & x \\ & & & & & x & x \\ & & & & & & \ddots \\ \hline x & 0 & & x & \otimes & + \\ & \otimes & & \otimes & \otimes & \otimes \\ & \oplus & x & + & \otimes & x & x \\ & & & & & x & x \\ & & & & & & \ddots \end{array} \right],$$

or in terms of $B^{(1)}$

$$G_2^T B^{(1)} G_2 = \tilde{X}_1 X_n Y_n \cdots X_3 Y_3 G_2^T Z_1 G_2 J^T.$$

Then a symplectic Householder transformation H_2 is used to zero the $(n+1, n+3)$ element:

$$H_2^T G_2^T B^{(1)} G_2 H_2 = \left[\begin{array}{ccc|ccc} x & & & x & \otimes & 0 \\ & \otimes & \oplus & \otimes & \otimes & \otimes & \oplus \\ & \oplus & \otimes & \oplus & \otimes & \otimes & \otimes \\ & & & & \oplus & \otimes & x \\ & & & & & & \ddots \\ x & & & x & \otimes & 0 & \\ & \otimes & \oplus & \otimes & \otimes & \otimes & \oplus \\ & \oplus & \otimes & \oplus & \otimes & \otimes & \otimes \\ & & & & \oplus & \otimes & x \\ & & & & & & \ddots \end{array} \right].$$

Using again the commuting properties and the fact that H_2 is unitary and symplectic, we obtain

$$H_2^T G_2^T B^{(1)} G_2 H_2 = \tilde{X}_1 X_n Y_n \cdots X_4 Y_4 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 J^T.$$

A symplectic Givens transformation G_3 annihilates the $(n+3, 2)$ element. This yields

$$G_3 H_2^T G_2^T B^{(1)} G_2 H_2 G_3^T = \left[\begin{array}{ccc|ccc} x & & & x & x & & \\ & x & \oplus & x & x & \otimes & + \\ & \oplus & \otimes & \oplus & \otimes & \otimes & \otimes \\ & & \oplus & x & & + & \otimes & x \\ & & & & & & & \ddots \\ x & & & x & x & & & \\ & x & \oplus & x & x & \otimes & + \\ & 0 & \otimes & 0 & \otimes & \otimes & \otimes \\ & & \oplus & x & + & \otimes & x \\ & & & & & & \ddots \end{array} \right],$$

and

$$G_3 H_2^T G_2^T B^{(1)} G_2 H_2 G_3^T = \tilde{X}_1 X_n Y_n \cdots X_4 Y_4 G_3 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 G_3^T J^T.$$

Finally, a symplectic Gauss transformation L_3 to eliminate the $(3, 2)$ element completes the bulge chase: $B^{(2)} := L_3 G_3 H_2^T G_2^T B^{(1)} G_2 H_2 G_3^T L_3^{-1}$ is of the form

$$B^{(2)} = \left[\begin{array}{ccc|ccc} x & & & x & \otimes & & \\ & \otimes & \oplus & \otimes & \otimes & \otimes & \oplus \\ & 0 & \otimes & 0 & \otimes & \otimes & \otimes \\ & & \oplus & x & \oplus & \otimes & x \\ & & & & & & \ddots \\ x & & & x & \otimes & & \\ & \otimes & \oplus & \otimes & \otimes & \otimes & \oplus \\ & & \otimes & \oplus & \otimes & \otimes & \otimes \\ & & \oplus & x & \oplus & \otimes & x \\ & & & & & & \ddots \end{array} \right].$$

The bulge has been chased exactly one row and one column down the subdiagonal in each block. The form of $B^{(2)}$ is the same as the form of $B^{(1)}$, just the bulge can be found one row and one column further down in each block. The same sequence of symplectic Givens, Householder and Gauss transformation as in the last four steps can be used to chase the bulge one more row and column down in each block.

Furthermore, due to the commuting properties and the symplecticity of L_3 we have

$$\begin{aligned} B^{(2)} &:= L_3 G_3 H_2^T G_2^T B^{(1)} G_2 H_2 G_3^T L_3^{-1} \\ &= \tilde{X}_1 X_n Y_n \cdots X_4 Y_4 L_3 G_3 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 G_3^T L_3^T J^T. \end{aligned}$$

In subsequent transformations the elements of $B^{(2)}$ in the positions $(2, 2)$, $(n+2, 2)$, $(1, n+1)$, $(1, n+2)$, $(2, n+1)$, $(n+1, n+1)$, $(n+1, n+2)$ and $(n+2, n+1)$ are not altered. Hence, at this point we can read off \tilde{a}_2 , \tilde{b}_2 , \tilde{c}_1 , and \tilde{d}_2 of the final \tilde{B} . Note that \tilde{X}_2 and \tilde{Y}_1 do not commute. In other words, we can rewrite

$$L_3 G_3 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 G_3^T L_3^T J^T$$

in terms of $\tilde{X}_2 \tilde{Y}_1$ times an appropriate symplectic matrix Z_2 times J^T . That is,

$$L_3 G_3 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 G_3^T L_3^T J^T = \tilde{X}_2 \tilde{Y}_1 Z_2 J^T.$$

As \tilde{X}_2 and \tilde{Y}_1 commute with most of the factors of $B^{(2)}$ we obtain

$$B^{(2)} = \tilde{X}_1 \tilde{X}_2 \tilde{Y}_1 X_n Y_n \cdots X_4 Y_4 Z_2 J^T.$$

Continuing in this fashion, we obtain for $j = 2, \dots, n-1$

$$B^{(j)} := L_{j+1} G_{j+1} H_j^T G_j^T B^{(j-1)} G_j H_j G_{j+1}^T L_{j+1}^{-1},$$

and

$$\begin{aligned} B^{(j)} &:= L_{j+1} G_{j+1} H_j^T G_j^T B^{(j-1)} G_j H_j G_{j+1}^T L_{j+1}^{-1} \\ &= \tilde{X}_1 \cdots \tilde{X}_{j-1} \tilde{Y}_1 \cdots \tilde{Y}_{j-2} X_n Y_n \cdots X_{j+2} Y_{j+2} \cdot \\ &\quad \cdot L_{j+1} G_{j+1} H_j^T X_{j+1} Y_{j+1} G_j^T Z_{j-1} G_j H_j G_{j+1}^T L_{j+1}^T J^T \\ &= \tilde{X}_1 \cdots \tilde{X}_{j-1} \tilde{X}_j \tilde{Y}_1 \cdots \tilde{Y}_{j-2} \tilde{Y}_{j-1} X_n Y_n \cdots X_{j+2} Y_{j+2} Z_j J^T, \end{aligned}$$

where $X_{n+1} = Y_{n+1} = I$. Thus,

$$B^{(n-1)} := L_n G_n H_{n-1}^T G_{n-1}^T B^{(n-2)} G_{n-1} H_{n-1} G_n^T L_n^{-1},$$

and

$$B^{(n-1)} = \tilde{X}_1 \cdots \tilde{X}_{n-1} \tilde{Y}_1 \cdots \tilde{Y}_{n-2} Z_{n-1} J^T.$$

One last symplectic Givens transformation has to be applied to $B^{(n-1)}$ to obtain the new butterfly matrix \tilde{B}

$$G_n^T B^{(n-1)} G_n = \tilde{B}.$$

Hence,

$$G_n^T Z_{n-1} G_n = \tilde{X}_n \tilde{Y}_{n-1} \tilde{Y}_n.$$

and

$$\tilde{B} = \tilde{X}_1 \cdots \tilde{X}_n \tilde{Y}_1 \cdots \tilde{Y}_n J^T.$$

4. The Details. How can the above observations be used to derive an algorithm which works solely on the parameters that determine B without forming B , \tilde{B} or any of the intermediate matrices? Let us start with (3.3),

$$L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T = \tilde{X}_1 Z_1.$$

X_1, X_2, Y_1 , and Y_2 are known. S_1 is determined by the choice of the spectral transformation function which drives the current SR step. As discussed in [4, 11] for a double shift the shift polynomial $q_2(B) = (B + B^{-1}) - \beta I$ should be chosen where $\beta = \mu + \mu^{-1}$ if $\mu \in \mathbb{R}$ or $\beta = \mu + \bar{\mu}$ for $\mu \in \mathbb{C}, |\mu| = 1$. Here the shift μ is chosen corresponding to the generalized Rayleigh-quotient strategy. This implies

$$q_2(B)e_1 = (b_1 + a_1 c_1 - b_n - a_n c_n)e_1 + a_1 d_2 e_2.$$

Hence, for S_1 as in (3.1), α and β have to be determined such that

$$\begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix} \begin{bmatrix} b_1 + a_1 c_1 - b_n - a_n c_n \\ a_1 d_2 \end{bmatrix} = \begin{bmatrix} \star \\ 0 \end{bmatrix}.$$

Next a symplectic Givens transformation G_2 has to be determined such that

$$(G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{n+2, n+1} = 0.$$

This implies that $G_2 = G(2, \alpha_2, \beta_2)$ has to be chosen such that

$$\begin{bmatrix} \alpha_2 & \beta_2 \\ -\beta_2 & \alpha_2 \end{bmatrix} \begin{bmatrix} (S_1^T X_2 Y_2 X_1 Y_1 S_1)_{2, n+1} \\ (S_1^T X_2 Y_2 X_1 Y_1 S_1)_{n+2, n+1} \end{bmatrix} = \begin{bmatrix} \star \\ 0 \end{bmatrix}$$

where

$$\begin{aligned} (S_1^T X_2 Y_2 X_1 Y_1 S_1)_{2, n+1} &= \beta \alpha (b_1 - b_2), \\ (S_1^T X_2 Y_2 X_1 Y_1 S_1)_{n+2, n+1} &= \beta \alpha (a_1 - a_2). \end{aligned}$$

Now a symplectic Gauss transformation $L_2 = L_2(\tau_1, \psi_1)$ is used such that

$$(L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T)_{2, n+1} = 0.$$

Hence, we have to compute τ_1 and ψ_1 such that

$$\left[\begin{array}{c|c} \tau_1 & \psi_1 \\ \hline \tau_1 & \psi_1 \\ \tau_1^{-1} & \\ \hline & \tau_1^{-1} \end{array} \right] \begin{bmatrix} (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{1, n+1} \\ (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{2, n+1} \\ (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{n+1, n+1} \\ 0 \end{bmatrix} = \begin{bmatrix} \star \\ 0 \\ \star \\ 0 \end{bmatrix},$$

where

$$\begin{aligned} (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{1, n+1} &= \alpha^2 b_1 + \beta^2 b_2, \\ (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{2, n+1} &= \alpha_2 (\beta \alpha (b_1 - b_2)) + \beta_2 (\beta \alpha (a_1 - a_2)), \\ (G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T)_{n+1, n+1} &= \alpha^2 a_1 + \beta^2 a_2. \end{aligned}$$

Now we can read off \tilde{a}_1 and \tilde{b}_1

$$\tilde{a}_1 = (\alpha^2 a_1 + \beta^2 a_2) / \tau_1^2, \quad \tilde{b}_1 = \alpha^2 b_1 + \beta^2 b_2.$$

Moreover, $L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T$ is a matrix of the form

$$\left[\begin{array}{ccc|cc} x & x & x & x & x \\ x & x & x & & x \\ & & x & & \\ & & & \ddots & \\ \hline x & x & x & x & x \\ x & x & x & & x \\ x & x & & & x & x \\ & & & & & \ddots \end{array} \right].$$

Now we form

$$\tilde{X}_1 = \left[\begin{array}{c|c} \tilde{a}_1^{-1} & \tilde{b}_1 \\ \hline I_{n-1} & \tilde{a}_1 \\ & I_{n-1} \end{array} \right],$$

and build $Z_1 = \tilde{X}_1^{-1} L_2 G_2 S_1^T X_2 Y_2 X_1 Y_1 S_1 G_2^T L_2^T$. This is a matrix of the form

$$\left[\begin{array}{ccc|cc} 1 & & & & \\ \delta_{21} & \delta_{22} & \delta_{23} & & \varepsilon_{22} \\ & & 1 & & \\ & & & \ddots & \\ \hline \mu_{11} & \mu_{12} & \mu_{13} & 1 & \zeta_{12} \\ \mu_{21} & \mu_{22} & \mu_{23} & & \zeta_{22} \\ \mu_{31} & \mu_{32} & & & \zeta_{32} & 1 \\ & & & & & \ddots \end{array} \right],$$

where the entries that will be used in the subsequent transformations are given by

$$\begin{aligned} \mu_{11} &= (\tau_1^2 g_1 + \psi_1 \tau_1 h_6) / (\alpha^2 a_1 + \beta^2 b_1), & \delta_{22} &= \tau_1^2 (\alpha_2 h_1 + \beta_2 h_2) + \tau_1 \psi_1 h_5, \\ \mu_{12} &= \psi_1 \tau_1 + \tau_1^2 h_5 / (\alpha^2 a_1 + \beta^2 b_1), & \delta_{23} &= -\tau_1 \alpha_2 \alpha b_2 d_3 - \tau_1 \beta_2 \alpha d_3 a_2, \\ \mu_{13} &= -\tau_1 \beta d_3 a_2 / (\alpha^2 a_1 + \beta^2 b_1), & \varepsilon_{22} &= -\beta_2 h_1 + \alpha_2 h_2 + \psi_1 h_6 / \tau_1, \\ \mu_{22} &= \alpha_2 h_3 + \beta_2 h_4, & \zeta_{12} &= h_6 / (\alpha^2 a_1 + \beta^2 b_1), \\ \mu_{23} &= (\beta_2 \alpha b_2 d_3 - \alpha_2 \alpha d_3 a_2) / \tau_1, & \zeta_{22} &= (-\beta_2 h_3 + \alpha_2 h_4) / \tau_1^2, \\ \mu_{32} &= -\tau_1 \alpha_2 \alpha d_3, & \zeta_{32} &= \beta_2 \alpha d_3 / \tau_1, \end{aligned}$$

where

$$\begin{aligned} h_1 &= \alpha_2 g_2 + \beta_2 g_3, & h_2 &= \alpha_2 (\beta^2 b_1 + \alpha^2 b_2) + \beta_2 (\beta^2 a_1 + \alpha^2 a_2), \\ h_3 &= \alpha_2 g_3 - \beta_2 g_2, & h_4 &= \alpha_2 (\beta^2 a_1 + \alpha^2 a_2) - \beta_2 (\beta^2 b_1 + \alpha^2 b_2), \\ h_5 &= \alpha_2 g_4 + \beta_2 \alpha \beta (a_1 - a_2), & h_6 &= \alpha_2 \alpha \beta (a_1 - a_2) - \beta_2 g_4, \end{aligned}$$

and

$$\begin{aligned} g_1 &= -\alpha^2 a_1 c_1 + \alpha \beta d_2 (a_1 - a_2) - \beta^2 c_2 a_2, \\ g_2 &= \beta^2 (a_1^{-1} - b_1 c_1) - \alpha \beta d_2 (b_1 + b_2) + \alpha^2 (a_2^{-1} - b_2 c_2), \\ g_3 &= -\beta^2 a_1 c_1 - \alpha \beta d_2 (a_1 + a_2) - \alpha^2 a_2 c_2, \\ g_4 &= -\alpha^2 a_1 d_2 + \alpha \beta (c_2 a_2 - a_1 c_1) + \beta^2 a_2 d_2. \end{aligned}$$

Next we have to consider

$$L_3 G_3 H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 G_3^T L_3^T J^T.$$

First a symplectic Givens transformation G_2 eliminates the $(n+1, n+2)$ element of Z_1 . This implies that $G_2 = G_2(\alpha_3, \beta_3)$ has to be chosen such that

$$(4.1) \quad [\mu_{12} \quad \zeta_{12}] \begin{bmatrix} \alpha_3 & \beta_3 \\ -\beta_3 & \alpha_3 \end{bmatrix} = [\star \quad 0].$$

The resulting transformed matrix is given by

$$G_2^T Z_1 G_2 = \left[\begin{array}{cccc|ccc} 1 & & & & & & \\ \delta_{21}^{(1)} & \delta_{22}^{(1)} & \delta_{23}^{(1)} & & & \varepsilon_{22}^{(1)} & \\ & & 1 & & & & \\ & & & \ddots & & & \\ & & & & & & \\ \hline \mu_{11}^{(1)} & \mu_{12}^{(1)} & \mu_{13}^{(1)} & & 1 & & \\ \mu_{21}^{(1)} & \mu_{22}^{(1)} & \mu_{23}^{(1)} & & & \zeta_{22}^{(1)} & \\ \mu_{31}^{(1)} & \mu_{32}^{(1)} & & & & \zeta_{32}^{(1)} & 1 \\ & & & & & & \ddots \end{array} \right],$$

where the relevant entries are

$$(4.2) \quad \begin{aligned} \mu_{12}^{(1)} &= \alpha_3 \mu_{12} - \beta_3 \zeta_{12}, & \mu_{23}^{(1)} &= \beta_3 \delta_{23} + \alpha_3 \mu_{23}, \\ \mu_{22}^{(1)} &= \alpha_3^2 \mu_{22} + \beta_3 \alpha_3 (\delta_{22} - \zeta_{22}) - \beta_3^2 \varepsilon_{22}, & \mu_{32}^{(1)} &= \alpha_3 \mu_{32} - \beta_3 \zeta_{32}, \\ \delta_{22}^{(1)} &= \alpha_3^2 \delta_{22} - \alpha_3 \beta_3 (\varepsilon_{22} + \mu_{22}) + \beta_3^2 \zeta_{22}, & \mu_{32}^{(1)} &= \alpha_3 \mu_{32} - \beta_3 \zeta_{32}, \\ \varepsilon_{22}^{(1)} &= \alpha_3^2 \varepsilon_{22} + \beta_3 \alpha_3 (\delta_{22} - \zeta_{22}) - \beta_3^2 \mu_{22}, & \delta_{23}^{(1)} &= \alpha_3 \delta_{23} - \beta_3 \mu_{23}, \\ \zeta_{22}^{(1)} &= \alpha_3^2 \zeta_{22} + \alpha_3 \beta_3 (\mu_{22} + \varepsilon_{22}) + \beta_3^2 \delta_{22}, & \zeta_{32}^{(1)} &= \beta_3 \mu_{32} + \alpha_3 \zeta_{32}. \end{aligned}$$

The $(1, 1)$ and the $(1, 3)$ entry are not altered by this transformation: $\mu_{11}^{(1)} = \mu_{11}$ and $\mu_{13}^{(1)} = \mu_{13}$.

Next a symplectic Householder transformation H_2 is used to zero the $(n+1, 3)$ element of $G_2^T Z_1 G_2$. H_2 is a matrix of the form (3.1); we denote its entries by α_4 and β_4 . The scalars α_4 and β_4 have to be chosen such that

$$(4.3) \quad [\mu_{12}^{(1)} \quad \mu_{13}^{(1)}] \begin{bmatrix} \alpha_4 & \beta_4 \\ -\beta_4 & \alpha_4 \end{bmatrix} = [\star \quad 0].$$

This results in

$$H_2^T X_3 Y_3 G_2^T Z_1 G_2 H_2 = \left[\begin{array}{cccc|cc} 1 & & & & & \\ \delta_{21}^{(2)} & \delta_{22}^{(2)} & \delta_{23}^{(2)} & \delta_{24}^{(2)} & \varepsilon_{22}^{(2)} & \varepsilon_{23}^{(2)} \\ \delta_{31}^{(2)} & \delta_{32}^{(2)} & \delta_{33}^{(2)} & \delta_{34}^{(2)} & \varepsilon_{32}^{(2)} & \varepsilon_{33}^{(2)} \\ & & 1 & & & \\ & & & \ddots & & \\ \hline \mu_{11}^{(2)} & \mu_{12}^{(2)} & & & 1 & \\ \mu_{21}^{(2)} & \mu_{22}^{(2)} & \mu_{23}^{(2)} & \mu_{24}^{(2)} & \zeta_{22}^{(2)} & \zeta_{23}^{(2)} \\ \mu_{31}^{(2)} & \mu_{32}^{(2)} & \mu_{33}^{(2)} & \mu_{34}^{(2)} & \zeta_{32}^{(2)} & \zeta_{33}^{(2)} \\ & \mu_{42}^{(2)} & \mu_{43}^{(2)} & & & 1 \\ & & & & & \ddots \end{array} \right],$$

Forming $\tilde{X}_2\tilde{Y}_1$ we see that $Z_2 = \tilde{Y}_1^{-1}\tilde{X}_2^{-1}L_3G_3H_2^T X_3Y_3G_2^T Z_1G_2H_2G_3^T L_3^T$ is given by

$$Z_2 = \left[\begin{array}{ccc|ccc} 1 & & & & & \\ & 1 & & & & \\ & \delta_{32}^{(5)} & \delta_{33}^{(5)} & \delta_{34}^{(5)} & & \varepsilon_{33}^{(5)} \\ & & & 1 & & \\ & & & & \ddots & \\ & & & & & 1 \\ \hline & \mu_{22}^{(5)} & \mu_{23}^{(5)} & \mu_{24}^{(5)} & 1 & \zeta_{23}^{(5)} \\ & \mu_{32}^{(5)} & \mu_{33}^{(5)} & \mu_{34}^{(5)} & & \zeta_{33}^{(5)} \\ & \mu_{42}^{(5)} & \mu_{43}^{(5)} & & & \zeta_{43}^{(5)} \\ & & & & & 1 \\ & & & & & \ddots \end{array} \right],$$

where only the elements

$$(4.10) \quad \begin{aligned} \zeta_{23}^{(5)} &= \zeta_{23}^{(4)} / \zeta_{22}^{(4)}, \\ \mu_{2k}^{(5)} &= \mu_{2k}^{(4)} / \zeta_{22}^{(4)}, \end{aligned} \quad \text{for } k = 2, 3, 4.$$

changed.

Comparing Z_1 and Z_2 , the bulge has been chased down exactly one row and column in each block. The same sequence of symplectic Givens, Householder and Gauss transformations as in the last four steps can be used to chase the bulge one more row and column down in each block. Therefore, renaming

$$\begin{aligned} \delta_{22} &= \delta_{33}^{(5)}, & \delta_{23} &= \delta_{34}^{(5)}, \\ \varepsilon_{22} &= \varepsilon_{33}^{(5)}, & \mu_{11} &= \mu_{22}^{(5)}, \\ \mu_{12} &= \mu_{23}^{(5)}, & \mu_{13} &= \mu_{24}^{(5)}, \\ \mu_{22} &= \mu_{33}^{(5)}, & \mu_{23} &= \mu_{34}^{(5)}, \\ \mu_{32} &= \mu_{43}^{(5)}, & \zeta_{12} &= \zeta_{23}^{(5)}, \\ \zeta_{22} &= \zeta_{33}^{(5)}, & \zeta_{32} &= \zeta_{43}^{(5)}, \end{aligned}$$

and repeating the computations (4.1) – (4.10) we obtain

$$\begin{aligned} \tilde{a}_3 &= \zeta_{22}^{(4)}, & \tilde{b}_3 &= \varepsilon_{22}^{(4)}, \\ \tilde{c}_2 &= -\mu_{11}^{(4)}, & \tilde{d}_3 &= -\mu_{12}^{(4)}. \end{aligned}$$

Iterating like this, the parameters $\tilde{a}_1, \dots, \tilde{a}_{n-1}, \tilde{b}_1, \dots, \tilde{b}_{n-1}, \tilde{c}_1, \dots, \tilde{c}_{n-2}$, and $\tilde{d}_2, \dots, \tilde{d}_{n-1}$ can be computed.

For the final step of the algorithm, let us consider the matrix Z_{n-1} . It has the form

$$\left[\begin{array}{cc|cc} I_{n-2} & & & \\ & 1 & & \\ & \delta & \delta & \varepsilon \\ \hline & \mu & \mu & I_{n-2} \\ & \mu & \mu & 1 & \zeta \\ & & & & \zeta \end{array} \right].$$

A symplectic Givens transformation G_n has to be applied to zero the $(2n - 1, 2n)$ entry of Z_{n-1} . The transformation $G_n^T Z_{n-1} G_n$ does not cause any fill-in. Hence, the remaining parameters $\tilde{a}_n, \tilde{b}_n, \tilde{c}_{n-1}, \tilde{c}_n$, and \tilde{d}_n can be read off, as

$$\tilde{X}_n \tilde{Y}_{n-1} \tilde{Y}_n = \left[\begin{array}{cc|c} I_{n-2} & & \\ & 1 & \\ & -\tilde{b}_n \tilde{d}_n & \tilde{a}_n^{-1} - \tilde{b}_n \tilde{c}_n \\ \hline & & \tilde{b}_n \\ & -\tilde{c}_{n-1} & -\tilde{d}_n \\ & -\tilde{a}_n \tilde{d}_n & -\tilde{a}_n \tilde{c}_n \\ \hline & & I_{n-2} \\ & & 1 \\ & & \tilde{a}_n \end{array} \right],$$

and $G_n^T Z_{n-1} G_n$ are symplectic butterfly matrices of the same form.

Using the same renaming convention as above, this implies that for the Givens transformation G_n , the scalars α_6 and β_6 have to be determined such that

$$[\mu_{12} \quad \zeta_{12}] \begin{bmatrix} \alpha_6 & \beta_6 \\ -\beta_6 & \alpha_6 \end{bmatrix} = [\star \quad 0].$$

Applying the transformation, the following matrix entries change:

$$\begin{aligned} \mu_{12} &= \alpha_6 \mu_{12} - \beta_6 \zeta_{12}, \\ \mu_{22} &= \beta_6 (\alpha_6 \delta_{22} - \beta_6 \varepsilon_{22}) + \alpha_6 (\alpha_6 \mu_{22} - \beta_6 \zeta_{22}), \\ \zeta_{22} &= \alpha_6 (\beta_6 \mu_{22} + \alpha_6 \zeta_{22}) + \beta_6 (\beta_6 \delta_{22} + \alpha_6 \varepsilon_{22}), \\ \varepsilon_{22} &= \alpha_6 (\beta_6 \delta_{22} + \alpha_6 \varepsilon_{22}) - \beta_6 (\beta_6 \mu_{22} + \alpha_6 \zeta_{22}), \\ \delta_{22} &= \alpha_6^2 \delta_{22} - \alpha_6 \beta_6 (\varepsilon_{22} + \mu_{22}) + \beta_6^2 \zeta_{22}. \end{aligned}$$

The parameters $\tilde{a}_n, \tilde{b}_n, \tilde{c}_{n-1}, \tilde{c}_n$, and \tilde{d}_n are given by

$$\begin{aligned} \tilde{a}_n &= \zeta_{22}, & \tilde{b}_n &= \varepsilon_{22}, & \tilde{c}_{n-1} &= -\mu_{11}, \\ \tilde{c}_n &= -\mu_{22}/\alpha_n, & \tilde{d}_n &= -\mu_{12}. \end{aligned}$$

REMARK 4.1.

- a) No 'optimality' is claimed for the form of the algorithm as discussed above either with regard to operation counts or numerical stability. Variants are certainly possible.
- b) A careful flop count shows that one parameterized SR step as described above requires $219n - 233$ flops² (assuming that the parameters of a symplectic Givens transformation are computed using 6 flops, while those of a symplectic Gauss are computed using 7 flops). This can be seen as follows: The initial step requires 166 flops, the final one 39 flops. The computation of (4.2), (4.4), (4.5), (4.7), (4.9), (4.10) require 47, 90, 48, resp. 2 flops. These computations have to be repeated $n - 2$ times, resulting in $219(n - 2)$ flops. These flops counts assume that quantities like $\alpha_j^2, \beta_j^2, \alpha_j \beta_j$ which are used more than once are computed only once in order to save computational time. If the transformation matrix S is required, then $64n^2 - 128n$ flops have to be added as $2n - 4$ symplectic Givens transformations, $n - 2$ symplectic Gauss transformations, and $n - 2$ symplectic Householder transformations with $v \in \mathbf{R}^2$ are used.

²Following [14], we define each floating point arithmetic operation together with the associated integer indexing as a flop.

- c) The development of a parameterized quadruple shift SR step is possible.
d) The presented parameterized double shift SR algorithm can not be used to mimic a quadruple shift. For a quadruple shift the spectral transformation function

$$q_4(\lambda) = (\lambda + \lambda^{-1})^2 - (\mu + \mu^{-1} + \bar{\mu} + \bar{\mu}^{-1})(\lambda + \lambda^{-1}) + (\mu + \mu^{-1})(\bar{\mu} + \bar{\mu}^{-1}) - 2$$

should be used. The shift μ should be chosen according to the generalized Rayleigh-quotient strategy as explained in [4, 11]. That is, for a quadruple shift, the eigenvalues of the 4×4 symplectic matrix

$$G = \left[\begin{array}{cc|cc} b_{n-1} & & b_{n-1}c_{n-1} - a_{n-1}^{-1} & b_{n-1}d_n \\ & b_n & b_nd_n & b_nc_n - a_n^{-1} \\ \hline a_{n-1} & & a_{n-1}c_{n-1} & a_{n-1}d_n \\ & a_n & a_nd_n & a_nc_n \end{array} \right]$$

are chosen. We can not work with a double shift step in the case that the matrix G has eigenvalues $\mu, \bar{\mu}, \mu^{-1}, \bar{\mu}^{-1} \in \mathbb{C}$, $|\mu| \neq 1$. One might have the idea to first apply a double SR step with the driving polynomial

$$q_2^{(1)} = (B - \mu I)(B - \bar{\mu} I)B^{-1}$$

followed by a double shift SR step with the driving polynomial

$$q_2^{(2)} = (B - \mu^{-1} I)(B - \bar{\mu}^{-1} I)B^{-1}$$

as this is equivalent to applying a quadruple SR step. The vectors $q_2^{(1)}e_1$ and $q_2^{(2)}e_1$ are of the form

$$\xi_1 e_1 + \xi_2 e_2 + \xi_3 e_{n+1}.$$

But the parameterized double shift SR step relies on the fact that for the driving polynomial q_2 we have

$$q_2(B)e_1 = \varsigma_1 e_1 + \varsigma_2 e_2.$$

5. The Overall Process. By applying a sequence of parameterized double shift SR steps to a symplectic butterfly matrix B , it is possible to reduce the tridiagonal blocks in B to diagonal form if B has only real eigenvalues or eigenvalues on the unit circle. The eigenproblem decouples into simple symplectic 2×2 eigenproblems. Decoupling occurs if $d_j = 0$ for some j . Therefore it is necessary to monitor the parameters d_j in order to bring about decoupling whenever possible. We proceed with the process of applying double shift SR steps until the problem has completely split into subproblems of dimension 2. That is, until all parameters d_j are equal to zero. The complete process is given in Table 5.1. In a final step we then have to solve the small subproblems.

In case the $(2, 1)$ entry is zero, the problem is already in the desired form; but we might have to reorder the eigenvalues on the diagonal such the smaller one is in the $(1, 1)$ position. Assume we have

$$\begin{bmatrix} \alpha & t \\ 0 & \alpha^{-1} \end{bmatrix},$$

where $|\alpha| \geq 1$. The reordering can be done as described in [14, Section 7.6.2] using a Givens rotation Q_D such that the second component of $Q_D^T \begin{bmatrix} t \\ \alpha - \alpha^{-1} \end{bmatrix}$ is zero.

Otherwise, the subproblems are of the form

$$M = \begin{bmatrix} b_j & b_j c_j - a_j^{-1} \\ a_j & a_j c_j \end{bmatrix}.$$

The eigenvalues are given by $\lambda_{\pm} = (a_j c_j + b_j)/2 \pm \sqrt{(a_j c_j + b_j)^2/4 - 1}$. If these eigenvalues are real, choose the one that is inside the unit circle and denote it by λ . The corresponding eigenvector is given by

$$\begin{bmatrix} \lambda - a_j c_j \\ a_j \end{bmatrix}.$$

Then the orthogonal symplectic matrix

$$Q = \frac{1}{\sqrt{(\lambda - a_j c_j)^2 + a_j^2}} \begin{bmatrix} \lambda - a_j c_j & -a_j \\ a_j & \lambda - a_j c_j \end{bmatrix}$$

transforms M into upper triangular form

$$Q^T M Q = \begin{bmatrix} \lambda & \star \\ 0 & \lambda^{-1} \end{bmatrix}.$$

In case $|\lambda_{\pm}| = 1$, we leave M as it is. Embedding Q into a $2n \times 2n$ symplectic Givens transformation in the usual way, we can update the $2n \times 2n$ problem. The above described process computes the real Schur form of M using a (symplectic) Givens transformation. In our implementation we use the MATLAB routine 'schur' for this purpose instead of the above, explicit approach. In this case we might have to order the eigenvalues on the diagonal as there is no guarantee that 'schur' puts the eigenvalue inside the unit circle into the (1, 1) position.

REMARK 5.1. *The parameterized double shift SR algorithm for butterfly matrices as given in Table 5.1 requires about $146n^2$ flops. If the transformation matrix S is accumulated, an additional $28n^3$ flops have to be added. This estimate is based on the observation that $\frac{2}{3}$ SR iterations per eigenvalue are necessary.*

6. Numerical Examples. The parameterized butterfly SR algorithm for computing the eigenvalues of symplectic matrices was implemented in MATLAB Version 5.1. Numerical experiments comparing the algorithm presented here with the double shift SR algorithm were performed on a SPARC Ultra 1 creator workstation.

For the tests reported here $n \times n$ diagonal matrices D were generated using MATLAB's 'rand' function. Then a symplectic matrix S was constructed such that $S = M^T \text{diag}(D, D^{-1})M$ where $M \in \mathbf{R}^{2n \times 2n}$ are randomly generated symplectic orthogonal matrices. This guarantees that all test matrices have only real-valued pairs of eigenvalues $\{\mu, \mu^{-1}\}$, $\mu \in \mathbf{R}$. Hence, using only double shift Laurent polynomials to drive the SR step, the corresponding butterfly matrices can be reduced to butterfly matrices such that the (1, 2) and the (2, 2) block is diagonal (that is, all parameters d_j are zero).

In order to detect deflation in the parameterized SR algorithm, parameters d_j were declared to be zero during the iteration when

$$d_j \leq 10 \cdot n \cdot \text{eps}$$

Algorithm: Parameterized Double Shift SR Algorithm for Butterfly Matrices

Given the parameters $a_1, \dots, a_n, b_1, \dots, b_n, c_1, \dots, c_n, d_1, \dots, d_n$ of a symplectic butterfly matrix B , the following algorithm computes the parameters $\tilde{a}_1, \dots, \tilde{a}_n, \tilde{b}_1, \dots, \tilde{b}_n, \tilde{c}_1, \dots, \tilde{c}_n, \tilde{d}_1, \dots, \tilde{d}_n$ of a symplectic butterfly matrix \tilde{B} that is similar to B . All \tilde{d}_j are zero. Thus the eigenproblem for \tilde{B} decouples into 2×2 symplectic eigenproblems. B is assumed to have only real eigenvalues or eigenvalues on the unit circle.

```

let  $d_1 = 0$ 
 $q = n + 1, p = 1$ 
repeat until  $q = p$ 
  set all  $d_j$  to zero that satisfy  $d_j \leq \epsilon$ 
  find the largest nonnegative  $q$  and the smallest nonnegative  $p$  such that

       $d_1 = \dots = d_p = 0 \neq d_{p+1} \quad d_{q-1} \neq d_q = \dots = d_n = 0$ 

  if  $q \neq p$ 
    perform a parameterized double shift  $SR$  step on
       $a_{p+1}, \dots, a_{q-1}, b_{p+1}, \dots, b_{q-1}, c_{p+1}, \dots, c_{q-1}, d_{p+1}, \dots, d_{q-1}$ 
  end
end
solve the  $2 \times 2$  subproblems as described in the text

```

TABLE 5.1

Parameterized Double Shift SR Algorithm for Butterfly Matrices

was fulfilled, where the dimension of the problem is $2n \times 2n$ and $\text{eps} \approx 2.2204 \cdot 10^{-16}$ is MATLAB's floating point relative accuracy. Deflation in the double shift SR algorithm was determined by a condition of the form

$$(6.1) \quad |h_{p+1,p}| \leq 10 \cdot n \cdot \text{eps} (|h_{pp}| + |h_{p+1,p+1}|).$$

While symplecticity is forced by the parameterized SR algorithm, it has to be enforced after each double shift SR step. Otherwise symplecticity is lost in the double shift SR algorithm.

All tests showed that the parameterized SR algorithm and the double shift SR algorithm (with symplecticity enforced after each SR step) compute the eigenvalues to about the same accuracy. But the parameterized SR algorithm converged slightly faster than the double shift SR algorithm, exhibiting the same cubic convergence behavior (see [4, 11] for a discussion and numerical examples). Figure 6.1 shows the average number of iterations needed for convergence using the parameterized SR algorithm and the double shift SR algorithm. The average number of iterations needed for convergence of a pair of eigenvalues tends to be around $\frac{4}{3}$ iterations.

In order to compute the average number of iterations needed for convergence, 100 symplectic matrices S for each of the dimensions $2n \times 2n$ for $n = 4 : 40$ were constructed as described above. It was observed that the parameterized SR algorithm converges typically slightly faster than the double shift SR algorithm. For most of the test examples, the parameterized SR algorithm was as fast or faster than the double shift SR algorithm. Just for very few examples, the parameterized SR algorithm

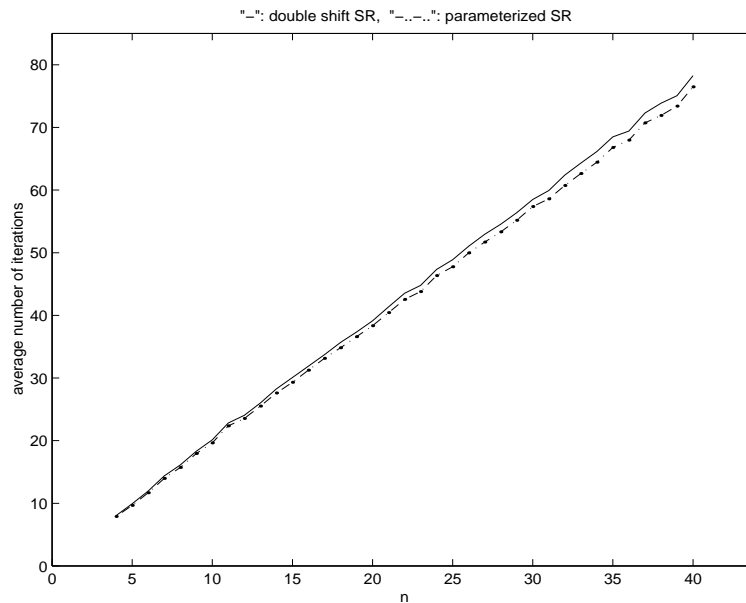


FIG. 6.1. average number of iterations, 100 examples for each dimension

needed more iteration than the double shift SR algorithm; and than only up to 3 iterations more. Mostly this was due to the fact that the deflation criterion for the parameterized SR algorithm is somewhat more strict than the one for the double shift SR algorithm. Similar results were obtained for test matrices $S = M^T \begin{bmatrix} D & F \\ 0 & D^{-1} \end{bmatrix} M$, where D, F are random diagonal $n \times n$ matrices and M is as before.

7. Conclusions. In this paper we have derived a parameterized version of the butterfly SR algorithm that works only on the $4n - 1$ parameters which uniquely determine a butterfly matrix. Symplecticity is forced in every step of the algorithm. The parameterized butterfly SR algorithm is an efficient structure-preserving algorithm for computing the eigenvalues of symplectic matrices. Using Laurent polynomials as shift polynomials cubic convergence can be observed. The parameterized butterfly SR algorithm converges slightly faster than the SR algorithm. The eigenvalues are computed to about the same accuracy.

Acknowledgments My thanks go to Volker Mehrmann and Hongguo Xu for their hospitality in Chemnitz. Sharing his office with me, Hongguo endured my grinding out the details of the parameterized SR .

REFERENCES

- [1] G. BANSE, *Symplektische Eigenwertverfahren zur Lösung zeitdiskreter optimaler Steuerungsprobleme*, PhD thesis, Fachbereich 3 - Mathematik und Informatik, Universität Bremen, Bremen, Germany, 1995.
- [2] G. BANSE AND A. BUNSE-GERSTNER, *A condensed form for the solution of the symplectic eigenvalue problem*, in *Systems and Networks: Mathematical Theory and Applications*, U. Helmke, R. Menniken, and J. Sauer, eds., Akademie Verlag, 1994, pp. 613–616.
- [3] P. BENNER AND H. FASSBENDER, *The symplectic eigenvalue problem, the butterfly form, the SR algorithm, and the Lanczos method*, *Linear Algebra Appl.*, 275–276 (1998), pp. 19–47.

- [4] P. BENNER, H. FASSBENDER, AND D. WATKINS, *SR and SZ algorithms for the symplectic (butterfly) eigenproblem*, Linear Algebra Appl., 287 (1999), pp. 41–76.
- [5] M. BOHNER, *Linear Hamiltonian difference systems: Disconjugacy and Jacobi-type conditions*, J. Math. Anal. Appl., 199 (1996), pp. 804–826.
- [6] J. BUNCH, *The weak and strong stability of algorithms in numerical algebra*, Linear Algebra Appl., 88 (1987), pp. 49–66.
- [7] A. BUNSE-GERSTNER, *Matrix factorizations for symplectic QR-like methods*, Linear Algebra Appl., 83 (1986), pp. 49–77.
- [8] A. BUNSE-GERSTNER AND V. MEHRMANN, *A symplectic QR-like algorithm for the solution of the real algebraic Riccati equation*, IEEE Trans. Automat. Control, AC-31 (1986), pp. 1104–1113.
- [9] A. BUNSE-GERSTNER, V. MEHRMANN, AND D. WATKINS, *An SR algorithm for Hamiltonian matrices based on Gaussian elimination*, Methods of Operations Research, 58 (1989), pp. 339–356.
- [10] J. DELLA-DORA, *Numerical linear algorithms and group theory*, Linear Algebra Appl., 10 (1975), pp. 267–283.
- [11] H. FASSBENDER, *Symplectic Methods for Symplectic Eigenproblems*, Habilitationsschrift, Universität Bremen, Fachbereich 3 – Mathematik und Informatik, 28334 Bremen, Germany, 1998.
- [12] U. FLASCHKA, V. MEHRMANN, AND D. ZYWIETZ, *An analysis of structure preserving methods for symplectic eigenvalue problems*, RAIRO Automatique Productique Informatique Industrielle, 25 (1991), pp. 165–190.
- [13] J. FRANCIS, *The QR transformation, Part I and Part II*, Comput. J., 4 (1961), pp. 265–271 and 332–345.
- [14] G. GOLUB AND C. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, 3rd ed., 1996.
- [15] V. KUBLANOSKAJA, *On some algorithms for the solution of the complete eigenvalue problem*, U.S.S.R. Comput. Math. and Math. Phys., 3 (1961), pp. 637–657.
- [16] P. LANCASTER AND L. RODMAN, *The Algebraic Riccati Equation*, Oxford University Press, Oxford, 1995.
- [17] V. MEHRMANN, *Der SR-Algorithmus zur Berechnung der Eigenwerte einer Matrix*, Diplomarbeit, Universität Bielefeld, Bielefeld, FRG, 1979.
- [18] ———, *The Autonomous Linear Quadratic Control Problem, Theory and Numerical Solution*, no. 163 in Lecture Notes in Control and Information Sciences, Springer-Verlag, Heidelberg, 1991.
- [19] C. PAIGE AND C. VAN LOAN, *A Schur decomposition for Hamiltonian matrices*, Linear Algebra Appl., 14 (1981), pp. 11–32.
- [20] T. PAPPAS, A. LAUB, AND N. SANDELL, *On the numerical solution of the discrete-time algebraic Riccati equation*, IEEE Trans. Automat. Control, AC-25 (1980), pp. 631–641.
- [21] D. WATKINS AND L. ELSNER, *Convergence of algorithms of decomposition type for the eigenvalue problem*, Linear Algebra Appl., 143 (1991), pp. 19–47.