Technische Universität Chemnitz

Sonderforschungsbereich 393

Numerische Simulation auf massiv parallelen Rechnern

David S. Watkins

Infinite Eigenvalues and the QZ Algorithm

Preprint SFB393/99-23

Abstract

The implicitly shifted (bulge chasing) QZ algorithm is the most popular method for solving the generalized eigenvalue problem $Av = \lambda Bv$. This paper explains why the QZ algorithm functions well even in the presence of infinite eigenvalues. The key to rapid convergence of QZ (and QR) algorithms is the effective transmission of shifts during the bulge chase. In this paper the mechanism of transmission of shifts is identified, and it is shown that this mechanism is not disrupted by the presence of infinite eigenvalues.

AMS Subject Classifications. 65F15, 15A18

Preprint-Reihe des Chemnitzer SFB 393

SFB393/99-23 **September 1999**

Contents

Author's addresses:

David S. Watkins
TU Chemnitz
Fakultät für Mathematik
D-09107 Chemnitz

Department of Mathematics
Washington State University
Pullman, WA 99164-3113
USA

1 Introduction

The standard matrix eigenvalue problem for the $n \times n$ matrix A has the form

$$(A - \lambda I)v = 0.$$

Many matrix eigenvalue problems are more naturally presented not in this form but as *generalized eigenvalue problems*

$$(A - \lambda B)v = 0.$$

Generalized eigenvalue problems have some interesting special features; for example, they can have infinite eigenvalues.

The most popular algorithm for solving generalized eigenvalue problems is the QZ algorithm of Moler and Stewart [2]. This is a generalization of the QR algorithm, which solves standard eigenvalue problems. An important feature of the QZ algorithm that was rightly emphasized by its inventors is that it functions perfectly well in the presence of infinite eigenvalues. However, when one looks at explanations of the QZ algorithm, e.g. [1], [2], [5], they always assume from the outset that B is nonsingular, which rules out infinite eigenvalues. As far as I know, all explanations of the QZ algorithm that have been published so far have shared this weakness. Thus there are good explanations of the QZ algorithm in the literature, but none of them holds for the case when infinite eigenvalues are present. The purpose of this paper is simply to demonstrate why the QZ algorithm performs well in the presence of infinite eigenvalues.

In this paper we discuss implicitly shifted (bulge-chasing) QZ algorithms. These are iterative processes that operate on pencils that have been reduced to Hessenberg-triangular form beforehand. Each QZ iteration makes use of shifts to introduce a bulge in the Hessenberg form at the top of the pencil. Then the bulge is chased to the bottom and off of the edge of the pencil, restoring the Hessenberg-triangular form. In the course of the iterations, eigenvalues are deflated one or more at a time. Infinite eigenvalues normally emerge at the top of the pencil, and finite eigenvalues are deflated at the bottom. The key to rapid convergence of the finite eigenvalues is the effective transmission of the shifts from top to bottom of the pencil during the bulge chase. (Of course the shifts must also be good approximations to eigenvalues, but this is not hard to arrange.) In this paper we identify the mechanism by which shifts are transmitted through the pencil during the bulge chase, and we show that the shift-transmission mechanism is not disrupted by the presence of infinite eigenvalues. Thus finite eigenvalues converge rapidly at the bottom of the pencil, regardless of whether or not infinite eigenvalues are present. In either case, every few iterations produces a new finite eigenvalue for deflation. If there are infinite eigenvalues, they rapidly percolate to the top of the pencil for deflation there. Deflation of infinite eigenvalues must not be neglected. If it is, the shift transmission process breaks down, and progress toward convergence comes to a halt.

The implicitly shifted QZ algorithms that we study in this paper are members of the larger family of implicitly shifted GZ algorithms [5]. The ideas presented in here are applicable to the larger family. We will restrict our attention to the QZ case in order to keep the presentation as simple as possible.

2 Basic Facts and Terminology

Given a pair of real or complex $n \times n$ matrices A and B, the matrix polynomial $A - \lambda B$ with indeterminate λ is called a *matrix pencil.* A finite complex number λ is called an *eigenvalue* of the pencil $A - \lambda B$ if there is a nonzero vector v (called an *eigenvector*) such that $(A - \lambda B)v = 0$. The problem of finding the eigenvalues of a matrix pencil is called the *generalized eigenvalue problem.* One easily sees that if the matrix B is nonsingular, the eigenvalues of the matrix pencil $A - \lambda B$ are exactly the eigenvalues of the matrix $B^{-1}A$. There are n of them, and they are (finite) complex numbers.

Regardless of whether or not B is singular, the (finite) eigenvalues of the matrix pencil are exactly the solutions of the *characteristic equation*

$$\det(A - \lambda B) = 0.$$

This is analogous to the standard eigenvalue problem. The difference is that if B is singular, the *characteristic polynomial* $\det(A-\lambda B)$ has degree less than n. In fact, it can even happen that $\det(A - \lambda B)$ is identically zero. For example, this happens when A and B have a common null vector. Then every λ is an eigenvalue. If $\det(A - \lambda B)$ is identically zero, we call $A - \lambda B$ a *singular* pencil. Otherwise it is a *regular* pencil. We will focus on regular pencils.

Two matrix pencils $A-\lambda B$ and $\tilde{A}-\lambda\tilde{B}$ are called *strictly unitarily equivalent* if there are unitary matrices U and V such that $\tilde{A} - \lambda\tilde{B} = U(A - \lambda B)V$. Obviously strictly unitarily equivalent pencils have the same eigenvalues. The Generalized Schur Theorem [1, Theorem 7.7.1] states that every pencil is strictly unitarily equivalent to a pencil $\tilde{A} - \lambda\tilde{B}$ for which $\tilde{A}$ and $\tilde{B}$ are upper triangular. Letting $\alpha_1, \ldots, \alpha_n$ and $\beta_1, \ldots, \beta_n$ denote the main diagonal entries of $\tilde{A}$ and $\tilde{B}$, respectively, we see that the characteristic equation of $\tilde{A} - \lambda\tilde{B}$ is

$$\prod_{i=1}^{n}(\alpha_i - \lambda\beta_i) = 0.$$

If $\alpha_i = \beta_i = 0$ for some i, the pencil is singular. Otherwise it is regular, and each pair (α_i, β_i) for which $\beta_i \neq 0$ gives rise to an eigenvalue α_i/β_i. If the pencil is regular, but the matrix B (and $\tilde{B}$) is singular, there will be at least one pair for which $\beta_i = 0$ (and $\alpha_i \neq 0$). It is reasonable to say that each of these gives rise to an *infinite eigenvalue.* (Each corresponds to a zero eigenvalue of the reciprocal pencil $\mu A - B$.) If we make this convention, then each regular pencil has exactly n eigenvalues, counting the infinite ones.

Since the generalized Schur form tells everything about the eigenvalues of a pencil, one would naturally like to have an algorithm that transforms a pencil to generalized Schur form by a sequence of unitary equivalence transformations. A big step in this direction is to transform the pencil to *Hessenberg-triangular form.* Every pencil is strictly unitarily equivalent to a pencil $\hat{A} - \lambda\hat{B}$ for which $\hat{A}$ is upper Hessenberg ($\hat{a}_{ij} = 0$ if $i > j + 1$) and $\hat{B}$ is upper triangular. The reduction can be carried out by a direct procedure [1, §7.7.4] in $O(n^3)$ flops. This is already quite close to the generalized Schur form, but Galois theory requires that the rest of the reduction be done by an iterative method, in this case the QZ algorithm.

Suppose we have now a Hessenberg-triangular pencil. We will now drop the hats and call it $A-\lambda B$. If any one of the subdiagonal entries $a_{j+1,j}$ is zero, we can split the eigenvalue problem into two (or more) subproblems involving subpencils. Therefore, we can assume without loss of generality that our A is a *proper* upper Hessenberg matrix, i.e. $a_{j+1,j} \neq 0$ for all j.

If B is singular, at least one of the b_{jj} will be zero. There are at least two ways to deal with this situation. One is to chase the zero entry to either the top or the bottom of B by the procedure described in [1, §7.7.5]. Then an infinite (or possibly 0/0) eigenvalue can be deflated from either the top or the bottom of the pencil. The other course of action is simply to leave the zero in place and apply the QZ algorithm to the pencil with a singular B. The objectives of this paper are to demonstrate why there is no harm in doing this and to convince the reader that this is the most efficient course of action.

3 The implicitly-shifted QZ algorithm

The implicitly-shifted QZ algorithm operates on a pencil $A-\lambda B$ that is in proper Hessenberg-triangular form. A QZ iteration begins with an equivalence transformation that disturbs the Hessenberg form of A by introducing a bulge near the upper left hand corner. The rest of the iteration consists of returning A to Hessenberg form by chasing the bulge from one end of the matrix to the other and, finally, off the edge.

Chasing the bulge

We will describe first the procedure by which the bulge is chased through the matrix, deferring until later the description of how the bulge is created. An iteration of degree m produces a bulge that protrudes m diagonals beyond the subdiagonal. For example, a bulge of degree two looks like this:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ & a_{52} & a_{53} & a_{54} & a_{55} & a_{56} \\ & & & & a_{65} & a_{66} \end{bmatrix} - \lambda \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} & b_{15} & b_{16} \\ & b_{22} & b_{23} & b_{24} & b_{25} & b_{26} \\ & & b_{33} & b_{34} & b_{35} & b_{36} \\ & & & b_{44} & b_{45} & b_{46} \\ & & & & b_{55} & b_{56} \\ & & & & & b_{66} \end{bmatrix}.$$

This bulge has been chased one row down and one column to the right from its original position. To chase it further we apply a pair of Givens rotators on the left to annihilate the entries in positions $(5,2)$ and $(4,2)$. We also apply compensating transformations on the right to maintain the triangular form of B. Specifically, the first Givens rotator (or Householder reflector or other unitary transformation matrix) acts on rows 4 and 5 and transforms the $(5,2)$ entry to zero. We must apply the same transformation to B. This recombines rows 4 and 5, creating a new nonzero entry in position $(5,4)$. The newly transformed pencil has the form

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ & & a_{53} & a_{54} & a_{55} & a_{56} \\ & & & & a_{65} & a_{66} \end{bmatrix} - \lambda \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} & b_{15} & b_{16} \\ & b_{22} & b_{23} & b_{24} & b_{25} & b_{26} \\ & & b_{33} & b_{34} & b_{35} & b_{36} \\ & & & b_{44} & b_{45} & b_{46} \\ & & & b_{54} & b_{55} & b_{56} \\ & & & & & b_{66} \end{bmatrix}.$$

One general principle of this algorithm is that whenever the triangular form of B is disturbed, we restore it immediately. Thus we apply a transformation on the right that acts on columns 4 and 5 and annihilates the entry b_{54}. When we apply this same transformation to A, it recombines columns 4 and 5, creating a nonzero entry in position $(6,4)$. The pencil now has the form

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ & & a_{53} & a_{54} & a_{55} & a_{56} \\ & & & a_{64} & a_{65} & a_{66} \end{bmatrix} - \lambda \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} & b_{15} & b_{16} \\ & b_{22} & b_{23} & b_{24} & b_{25} & b_{26} \\ & & b_{33} & b_{34} & b_{35} & b_{36} \\ & & & b_{44} & b_{45} & b_{46} \\ & & & & b_{55} & b_{56} \\ & & & & & b_{66} \end{bmatrix}.$$

The next transformation from the left acts on rows 3 and 4 and annihilates the entry a_{42}. When we apply this same transformation to B, it creates a new nonzero entry in position $(4,3)$. We then apply a transformation on the right to columns 3 and 4 to restore B to triangular form. When we apply this transformation to A, we recombine columns 3 and 4, creating a new nonzero entry in position $(6,3)$. The pencil now has the form

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ & & a_{43} & a_{44} & a_{45} & a_{46} \\ & & a_{53} & a_{54} & a_{55} & a_{56} \\ & & a_{63} & a_{64} & a_{65} & a_{66} \end{bmatrix} - \lambda \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} & b_{15} & b_{16} \\ & b_{22} & b_{23} & b_{24} & b_{25} & b_{26} \\ & & b_{33} & b_{34} & b_{35} & b_{36} \\ & & & b_{44} & b_{45} & b_{46} \\ & & & & b_{55} & b_{56} \\ & & & & & b_{66} \end{bmatrix}.$$

The bulge has now been pushed down and to the right one position. The cost was two rotations on the left, accompanied by two compensatory rotations on the right.

Chasing a bulge of arbitrary degree m follows the same principle. To push it forward one notch, we need to apply m rotations on the left, accompanied by m compensatory rotations on the right.

In our 6×6 example, the bulge is now near the bottom of A. The next step will push it part way off the bottom, and the final step (requiring only half as many transformations) finishes the job.

This is a very small example. The reader should envision a much bigger pencil with, say, $n = 100$ or $n = 1000$. Then the bulge chase, which pushes the bulge all the way from top to bottom of the matrix, is a much longer process.

Typical implementations of the QZ algorithm have degree $m = 1$ or $m = 2$. There is no theoretical reason why m cannot be higher, say 4 or 6.[1]

Initiating the QZ iteration

Now that we know how to chase the bulge, we consider the question of getting the bulge started. We will assume initially that B is nonsingular. The QZ iteration is the same as an implicit QR iteration on AB^{-1}. If we are going to do an iteration of degree m, we need *m shifts*, $\mu_1, \ldots, \mu_m$. Choice of shifts will be discussed later. Given m shifts, we calculate a vector

$$\begin{aligned} x &= \alpha(AB^{-1} - \mu_1 I) \cdots (AB^{-1} - \mu_m I)e_1 \\ &= \alpha(A - \mu_1 B)B^{-1} \cdots (A - \mu_m B)B^{-1}e_1, \end{aligned} \tag{1}$$

where α is any convenient nonzero scale factor. Because A is upper Hessenberg and B is upper triangular, only the first $m + 1$ components of x are nonzero.

The next step is to build a unitary transformation Q whose first column is proportional to x, i.e. $Qe_1 = \beta x$ or $Q^{-1}x = \beta^{-1}e_1$. If we think of Q^{-1} as a matrix that introduces zeros into the vector x, we see that one way to build Q^{-1} is as a product of m Givens rotations. The first, Q_m^{-1}, acts on rows m and $m + 1$ and transforms x_{m+1} to zero. The next, Q_{m-1}, acts on rows $m - 1$ and m and transforms the mth component of (the transformed) x to zero. The next transformation annihilates the $(m - 1)$st component, and so on. In the end, we can take

$$Q^{-1} = Q_1^{-1}Q_2^{-1} \cdots Q_m^{-1}, \tag{2}$$

and we have $Q^{-1}x = \gamma e_1$.

Next we transform the pencil by multiplying it by Q^{-1} on the left. (This has the same effect as the similarity transformation $AB^{-1} \to Q^{-1}AB^{-1}Q$.) If we apply the transformations Q_i^{-1} one at a time, we can maintain the upper triangularity of B by following each left transformation immediately by a compensating right transformation, just as in the bulge chase. For example, when we apply Q_m^{-1}, it recombines rows m and $m + 1$ of B, creating a new nonzero entry in position $(m + 1, m)$. We then immediately apply a right transformation Z_m, acting on columns m and $m + 1$, that annihilates that new entry. In the end we will have made a transformation $A - \lambda B \to Q^{-1}(A - \lambda B)Z$ (which also corresponds to a similarity transformation $AB^{-1} \to Q^{-1}AB^{-1}Q$). The transformed B is upper triangular, but the transformed A is no longer upper Hessenberg. It is easy to check that it has a bulge that protrudes m diagonals beyond the superdiagonal. For example, in

[1] However, shift blurring [3] limits the effectiveness of the QZ algorithm when m is large (e.g. $m = 20$).

the case $m = 2$, the transformed pencil looks like

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ & & & a_{54} & a_{55} & a_{56} \\ & & & & a_{65} & a_{66} \end{bmatrix} - \lambda \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} & b_{15} & b_{16} \\ & b_{22} & b_{23} & b_{24} & b_{25} & b_{26} \\ & & b_{33} & b_{34} & b_{35} & b_{36} \\ & & & b_{44} & b_{45} & b_{46} \\ & & & & b_{55} & b_{56} \\ & & & & & b_{66} \end{bmatrix}.$$

Now the preliminary transformation is complete. The rest of the QZ iteration consists of chasing the bulge through the matrix, as described above. The resulting transformation of $A - \lambda B$ is the same as that of an iteration of the QR algorithm of degree m on AB^{-1} [2], [1], [5].

4 The objective of the QZ iteration

The iteration begins with the choice of m shifts. A common strategy is to take them to be the eigenvalues of the lower right hand $m \times m$ subpencil. That is, if

$$A - \lambda B = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} - \lambda \begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix},$$

where $A_{22} - \lambda B_{22}$ is $m \times m$, then we take the shifts to be the eigenvalues of $A_{22} - \lambda B_{22}$. The hope is that these shifts will be good approximations to eigenvalues of the pencil. Notice that A_{21} has only one nonzero entry, $a_{n-m+1,n-m}$. If this entry is small, all of the shifts will be excellent approximations to eigenvalues of the pencil, except in ill-conditioned situations.

The progress that one makes on a QZ iteration is, roughly speaking, determined by ratios of eigenvalues of the matrix

$$p(AB^{-1}) = (AB^{-1} - \mu_1 I)(AB^{-1} - \mu_2 I) \cdots (AB^{-1} - \mu_m I). \tag{3}$$

If $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of the pencil $A - \lambda B$, then $p(\lambda_1), \ldots, p(\lambda_n)$ are the eigenvalues of $p(AB^{-1})$. Suppose we number the eigenvalues so that $|\,p(\lambda_1)\,| \geq |\,p(\lambda_2)\,| \geq \cdots \geq |\,p(\lambda_n)\,|$. Then the ratios

$$|\,p(\lambda_{k+1})/p(\lambda_k)\,| \qquad k = 1, \ldots, n-1,$$

determine the progress one makes. If any one of these ratios is much less than 1, the step will make good progress toward convergence.[2] If $|\,p(\lambda_{k+1})/p(\lambda_k)\,| \ll 1$, we will typically have $|\,\hat{a}_{k+1,k}\,| \ll |\,a_{k+1,k}\,|$, where $\hat{a}_{k+1,k}$ denotes the entry after the iteration. Indeed (typically)

$$|\,\hat{a}_{k+1,k}\,| \approx |\,p(\lambda_{k+1})/p(\lambda_k)\,|\,|\,a_{k+1,k}\,|,$$

[2]Technically the ratios determine asymptotic convergence rates that would be achieved if the same shifts were used over and over again. Strictly speaking, they should not be applied to a single step. However, experience indicates they usually do give a good indication of what will happen in a single step as well.

Since $\hat{a}_{k+1,k}$ is much smaller than $a_{k+1,k}$, we have made good progress toward block triangular form and a deflation.

Our shift strategy aims to choose shifts that are very close to eigenvalues. If $\mu_1, \ldots, \mu_m$ are very close to $\lambda_{n-m+1}, \ldots, \lambda_n$, then $p(\lambda_{n-m+1}), \ldots, p(\lambda_n)$ will all be small numbers, typically much smaller than the next smallest eigenvalue $p(\lambda_{n-m})$. Thus the ratio $|p(\lambda_{n-m+1})/p(\lambda_{n-m})|$ will be small, so $\hat{a}_{n-m+1,n-m}$ will be much smaller than $a_{n-m+1,n-m}$. Our next iteration will use the eigenvalues of the subpencil $\hat{A}_{22} - \lambda \hat{B}_{22}$ as shifts. Since $\hat{a}_{n-m+1,n-m}$ is much smaller than $a_{n-m+1,n-m}$, the new shifts will be much better estimates of eigenvalues than the previous ones were, and the new ratio $|\hat{p}(\lambda_{n-m+1})/\hat{p}(\lambda_{n-m})|$ will be much better than the previous one. Thus the next step will make much better progress than the previous one did. The positive feedback cycle of improving shifts and improving convergence results in a quadratic local convergence rate [4], [5]. In practice only a few iterations are needed before $a_{n-m+1,n-m}$ becomes small enough that we can set it to zero and deflate an $m \times m$ chunk from the problem. We then go to work on the remaining submatrix. A few iterations later, we deflate off another $m \times m$ chunk, and so on. In practice, the deflations sometimes occur in bigger or smaller chunks. Keep in mind that $m \ll n$; small chunks are being deflated from the bottom of the pencil.

5 Transmission of shifts

The whole point of the previous section was to demonstrate that shifts are important. Good shifts lead to rapid convergence. If one reviews what happens in a QZ iteration, one sees that the shifts are used only at the very beginning, in the determination of the vector $x = p(AB^{-1})e_1$, which is used to build the initial bulge at the top of the pencil. The shifts are then forgotten in the ensuing bulge chase. Yet the good shifts are what causes the rapid convergence, which takes place at or around $a_{n-m+1,n-m}$, i.e. near the bottom of the pencil.

To summarize, we can say that the shifts go in at the top and cause rapid convergence at the bottom. Thus it makes sense to ask how the shifts are transmitted through the matrix during the bulge chase. This question was answered for the standard eigenvalue problem in [3]. For the generalized problem, the answer turns out to be about the same. We briefly sketch it here, as it is the key to our understanding of the effects of infinite eigenvalues.

The bulge pencil

Consider a pencil that has a bulge somewhere in the middle. Say the initial bulge has been created, and it has been chased $j-1$ positions down and to the right. The current pencil $A_j - \lambda B_j$ has a bulge (in A_j) starting in column j. If the degree of the QZ iteration is m, the tip of the bulge is at $a_{j+m+1,j}$. We define the *bulge pencil* $C_j - \lambda F_j$ to be the $(m+1) \times (m+1)$, non-principal subpencil of $A_j - \lambda B_j$ consisting of rows $j+1$ through

$j+m+1$ and columns j through m. Thus

$$C_j = \begin{bmatrix} a_{j+1,j} & a_{j+1,j+1} & a_{j+1,j+2} & \cdots & a_{j+1,j+m} \\ a_{j+2,j} & a_{j+2,j+1} & a_{j+2,j+2} & \cdots & a_{j+2,j+m} \\ \vdots & \vdots & & & \vdots \\ a_{j+m,j} & a_{j+m,j+1} & a_{j+m,j+2} & \cdots & a_{j+m,j+m} \\ a_{j+m+1,j} & a_{j+m+1,j+1} & a_{j+m+1,j+2} & \cdots & a_{j+m+1,j+m} \end{bmatrix} \tag{4}$$

and

$$F_j = \begin{bmatrix} 0 & b_{j+1,j+1} & b_{j+1,j+2} & \cdots & b_{j+1,j+m} \\ 0 & 0 & b_{j+2,j+2} & \cdots & b_{j+2,j+m} \\ \vdots & \vdots & & & \vdots \\ 0 & 0 & 0 & \cdots & b_{j+m,j+m} \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}. \tag{5}$$

The bulge pencil is centered on the subdiagonal of the big pencil, and it is just big enough to accommodate the bulge. One can show by induction on j that the entry $a_{j+m+1,j}$ cannot be zero (The original A is properly upper Hessenberg). F_j is strictly upper triangular. If all of the superdiagonal entries of F_j are nonzero (as is the case when B is nonsingular), then the degree of the characteristic polynomial $\det(C_j - \lambda F_j)$ is exactly m. In this case the bulge pencil has m finite eigenvalues and one infinite eigenvalue.

The main theorem is that the m finite eigenvalues of the bulge pencil are the shifts μ_1, $\ldots$, μ_m. Thus the shifts are transmitted from top to bottom of the matrix as eigenvalues of the bulge pencil. In order to prove this, we need to introduce a "zeroth" bulge pencil.

The zeroth bulge

Define the *zeroth bulge pencil* $C_0 - \lambda F_0$ by

$$C_0 = \begin{bmatrix} x_1 & a_{1,1} & \cdots & a_{1,m-1} & a_{1,m} \\ x_2 & a_{2,1} & \cdots & a_{2,m-1} & a_{2,m} \\ \vdots & \vdots & & \vdots & \vdots \\ x_m & 0 & \cdots & a_{m,m-1} & a_{m,m} \\ x_{m+1} & 0 & \cdots & 0 & a_{m+1,m} \end{bmatrix}$$

and

$$F_0 = \begin{bmatrix} 0 & b_{1,1} & \cdots & b_{1,m-1} & b_{1,m} \\ 0 & 0 & \cdots & b_{2,m-1} & b_{2,m} \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & b_{m,m} \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

The entries x_1, x_2, $\ldots$, x_{m+1} are the nonzero entries of the vector x defined by (1), and the entries a_{ij} and b_{ij} are from the pencil $A - \lambda B$ before the beginning of the iteration.

The "bulge" in this pencil is caused by x. We can view $C_0 - \lambda F_0$ as a subpencil of the augmented pencil obtained by adjoining a "zeroth" column $x - \lambda 0$ to the pencil $A - \lambda B$. If we take this view, then $C_0 - \lambda F_0$ is not so different from the other bulge pencils $C_j - \lambda F_j$.

Because A is Hessenberg and B is triangular, the computation of x (cf. (1)) uses only the upper left-hand corner entries of A and B. One easily checks that the entries that participate in the computation are exactly those that are contained in $C_0 - \lambda F_0$. This is the only part of the computation that uses the shifts. We therefore expect that it should be possible to recover the shifts from $C_0 - \lambda F_0$.

Notice also that we can compute x even if B is singular; all that is needed is that the upper left-hand $m \times m$ submatrix of B is nonsingular, i.e. $b_{kk} \neq 0$ for $k = 1, \ldots, m$.

Theorem 5.1 *Suppose $b_{kk} \neq 0$ for $k = 1, \ldots, m$. Then the eigenvalues of the zeroth bulge pencil $C_0 - \lambda F_0$ are the shifts $\mu_1, \ldots, \mu_m$ and ∞.*

Proof. Since A is properly upper Hessenberg, we deduce easily that $x_{m+1} \neq 0$. This condition and the conditions $b_{kk} \neq 0$, $k = 1, \ldots, m$, together imply that the characteristic polynomial $\det(C_0 - \lambda F_0)$ has degree exactly m. Thus $C_0 - \lambda F_0$ has one infinite eigenvalue and m finite eigenvalues. To see that each shift μ_i is an eigenvalue of $C_0 - \lambda F_0$, write $p(AB^{-1})$ (cf. 3) in partially factored form: $p(AB^{-1}) = (A - \mu_i B)B^{-1}q(AB^{-1})$, where q has degree $m - 1$. Then $x = (A - \mu_i B)y$, where $y = B^{-1}q(AB^{-1})$. Let $\hat{y}$ be the subvector of y consisting of the first m entries. Notice that the rest of y is zero. Then the equation $x = (A - \mu_i B)y$ can be recast as

$$(C_0 - \mu_i F_0) \begin{bmatrix} 1 \\ -\hat{y} \end{bmatrix} = 0.$$

Thus μ_i is an eigenvalue of $C_0 - \lambda F_0$.

This argument holds even if B^{-1} does not exist; all that is needed is that the upper left-hand corner of B is invertible.

If $\mu_1, \ldots, \mu_m$ are distinct, then there can be no other finite eigenvalues. If $\mu_1, \ldots, \mu_m$ are not distinct, we draw the same conclusion by a continuity argument: Perturb the shifts slightly so that they are distinct. This implies a small perturbation of x. The m perturbed shifts are the m finite eigenvalues of the slightly perturbed bulge pencil. Now move the shifts continuously back to their original values and invoke continuity of eigenvalues of a pencil. □

Theorem 5.2 *Suppose $b_{kk} \neq 0$ for $k = 1, \ldots, m$. Then all of the bulge pencils $C_j - \lambda F_j$, $j = 0, 1, 2, \ldots$ have $\mu_1, \ldots, \mu_m$ and ∞ as their eigenvalues.*

Proof. The proof is by induction. We just need to show that $C_{j+1} - \lambda F_{j+1}$ has the same eigenvalues as $C_j - \lambda F_j$.

Suppose we have pushed the bulge forward to the point where we have reached the pencil $A_j - \lambda B_j$. The bulge begins in column j. In preparation for pushing the bulge further, consider the $(m+2)\times(m+2)$ subpencil of $A_j - \lambda B_j$ that consists of $C_j - \lambda F_j$ plus one additional column on the right and one additional row on the bottom. This augmented bulge pencil, which we will call $\hat{C}_j - \lambda\hat{F}_j$ has the same eigenvalues as $C_j - \lambda F_j$, except for one additional infinite eigenvalue. The transformation that moves the bulge one row down and one column to the right transforms $\hat{C}_j - \lambda\hat{F}_j$ to a new pencil $\check{C}_j - \lambda\check{F}_j$, which has the same eigenvalues, because the transformation is a strict equivalence. If we now delete the first row and column from $\check{C}_j - \lambda\check{F}_j$, we obtain the new bulge pencil $C_{j+1} - \lambda F_{j+1}$. The effect of the deletion is just to remove an infinite eigenvalue. Thus $C_{j+1} - \lambda F_{j+1}$ has exactly the same eigenvalues as $C_j - \lambda F_j$.

This argument is applicable even in the case $j = 0$. The transformation Q^{-1} (2) that is used to set up the initial bulge is exactly the transformation one would use to chase the "bulge" x from $C_0 - \lambda F_0$. Thus $C_1 - \lambda F_1$ is produced from $C_0 - \lambda F_0$ in exactly the same way as each subsequent bulge pencil is produced from its predecessor. This completes the proof. $\square$

6 The singular case

If B is singular, at least one of the entries b_{kk} is zero. A simple procedure for removing the zero and deflating an infinite eigenvalue is given in [1, §7.7.5]. We have a choice of removing all such zeros by this deflation procedure or going ahead with the QZ algorithm without first removing the zeros.

The cost of deflation is not great. The deflation procedure chases the zero to one end of the matrix or the other. The infinite eigenvalue is then deflated at either the top or the bottom. Thus zeros near either the top or the bottom of B can be removed at negligible cost. Zeros near the middle are the most expensive to remove; the cost is about half that of a single QZ iteration.

Now let us consider what happens if we simply perform QZ iterations without first removing the zeros. How do the QZ iterations affect the zeros? What effect do the zeros have on the QZ iteration? We consider the first question first.

Suppose $b_{kk} = 0$. We consider the effect of a QZ iteration of degree one on this zero. The bulge consists of a single entry that is chased from position a_{31} to a_{42} to a_{53}, and so on. Eventually it arrives at position $a_{k,k-2}$. At this point the zero entry at b_{kk} has not been touched. Annihilation of the bulge at $a_{k,k-2}$ is done by a rotator on rows $k-1$ and k. When this rotator is applied to B, it touches b_{kk}. Focus on the 2×2 submatrix of B consisting of rows and columns $k-1$ and k. Before the rotator, it has the form

$$\begin{bmatrix} b_{k-1,k-1} & b_{k-1,k} \\ 0 & 0 \end{bmatrix}.$$

Its rank is obviously one. When the rotator is applied, it disturbs both of the zeros. The submatrix now looks like

$$\begin{bmatrix} \tilde{b}_{k-1,k-1} & \tilde{b}_{k-1,k} \\ \tilde{b}_{k,k-1} & \tilde{b}_{k,k} \end{bmatrix},$$

but its rank is still one. The next step in the QZ iteration is to apply a rotator to columns $k-1$ and k to annihilate $\tilde{b}_{k,k-1}$. Application of this rotator transforms the submatrix to

$$\begin{bmatrix} 0 & \hat{b}_{k-1,k} \\ 0 & \hat{b}_{k,k} \end{bmatrix},$$

since the rank is still one. The zero has been moved from position b_{kk} to position $b_{k-1,k-1}$. The rest of the QZ iteration leaves that entry untouched. Thus the effect of a single QZ iteration is to move the zero up by one position.

If one studies a QZ iteration of degree m in the same way, one easily deduces that its effect is to move the zero up by m positions. This is to be expected, since an iteration of degree m is equivalent to m iterations of degree 1.

After $k-1$ iterations of degree 1, the zero has been moved to the top. An infinite eigenvalue can now be deflated.[3]

Effect of zeros on the performance of QZ iterations

We have seen that QZ iterations automatically move zeros upward and into position for deflation. We now ask whether or not the presence of those zeros has any adverse effect on the QZ iterations.

The key to effective QZ iterations is effective transmission of shifts. Theorem 5.2 tells how this is done, in theory at least (and perhaps also in practice). The key observation is that Theorem 5.2 holds even if B is singular. All that is needed is that $b_{kk} \neq 0$ for $k = 1, \ldots, m$. If there are some zero entries b_{kk} for $k > m$, then, according to the theorem, these zeros should not adversely affect the transmission of the shifts. As the zeros gradually float upward through the matrix, accurate shifts are transmitted in the bulge, through these zeros, to the bottom of the pencil, resulting in rapid convergence and deflation of eigenvalues at the bottom.

The hypothesis

$$b_{kk} \neq 0 \quad \text{for } k = 1, \ldots, m$$

is important; we cannot even compute (nor define) the vector x nor the zeroth bulge pencil unless the upper left-hand $m \times m$ submatrix of B is nonsingular. This means that if we are doing QZ iterations of degree m, as soon as a zero floats up into the first m positions of B, it must be removed by the deflation procedure outlined in [1, §7.7.5]. The cost of this is negligible since m is small.

We call Theorem 5.2 a theoretical result because it ignores the effects of roundoff errors, which can be significant. In [3] it was shown that if m is large (e.g. $m = 20$), roundoff

[3]The finite emergence of infinite eigenvalues at the top of the pencil is consistent with the convergence theory of the QZ algorithm [5]. The largest (in magnitude) eigenvalues should emerge at the top. Suppose the pencil has j infinite eigenvalues. The rate at which they should converge (assuming, for simplicity, we are using a fixed spectral transformation polynomial p) should be determined by the ratio $|p(\lambda_{j+1})|/|p(\lambda_j)|$, which is zero since $\lambda_j = \infty$. Thus the convergence should be superlinear. Finite emergence is certainly superlinear.

errors often prevent the effective transmission of shifts. This is cause for caution. We cannot say for sure that the shift transmission mechanism will work well unless we try it on the computer.

Here we are not concerned about large m. We want to ascertain whether the shift mechanism works well for small values of m in the presence of infinite eigenvalues. We conducted numerous experiments with pencils of various sizes and types and found that it does.

We will report on just a couple of examples. Consider first a random, complex 20×20 Hessenberg-triangular pencil with no infinite eigenvalues, to which we apply the QZ algorithm with $m = 1$. Although we are discussing a single pencil, the results reported here are typical of many examples that we looked at. We find that at each point in the bulge chase, the single finite eigenvalue of the bulge pencil agrees with the intended shift to fifteen or more decimal places. Thus the shift is transmitted effectively. Checking the convergence pattern, we observe that the first four eigenvalues are deflated after 9, 12, 16, and 23 iterations, respectively. Each eigenvalue converges quadratically, as evidenced by the rate at which the bottom subdiagonal entry tends to zero.

Now suppose we alter the pencil by setting $b_{5,5}$ and $b_{15,15}$ to zero. Since the zeros are pushed up by one position per iteration, we expect to deflate infinite eigenvalues at the top after 4 and 13 iterations, and we do. At the same time we hope to have normal convergence behavior at the bottom of the pencil. Checking the single finite eigenvalue of the bulge pencil, we find that at each stage of the bulge chase it agrees with the shift to fifteen or more decimal places. Thus the shift is transmitted effectively. We have no reason to believe that the altered pencil will have the same convergence pattern as the original pencil, but we hope that it will be comparable. Indeed it is; the first four eigenvalues are deflated after iterations 8, 11, 15, and 18, respectively. Quadratic convergence is observed. For example, Table 1

iteration	$\lvert a_{20,19} \rvert$	shift transmission error
3	2.1×10^{-2}	1.2×10^{-15}
4	4.3×10^{-3}	1.1×10^{-15}
5	1.8×10^{-4}	3.3×10^{-16}
6	1.1×10^{-6}	1.0×10^{-15}
7	3.9×10^{-11}	4.4×10^{-16}
8	5.2×10^{-20}	7.0×10^{-16}

Table 1: Quadratic convergence of eigenvalue

shows the values of $\lvert a_{20,19} \rvert$ in iterations three through eight. The approximate doubling of the exponent of $a_{20,19}$ from one iteration to the next indicates quadratic convergence to zero. Thus $a_{20,20}$ converges quadratically to an eigenvalue. The presence of infinite eigenvalues does not in any way impede convergence. The *shift transmission error* given in Table 1 is the difference between the intended shift and the finite eigenvalue of the bulge pencil when the bulge pencil has reached the bottom of the matrix. We see that these errors are always tiny.

It is interesting to see how the algorithm behaves if we neglect to deflate infinite eigenvalues as they emerge. Table 2 shows the same information as Table 1, except that in this run we do not deflate out the infinite eigenvalue that emerges after the fourth iteration. In principle the algorithm should crash because of a division by zero (b_{11}). In practice it does not, because roundoff errors prevent b_{11} from being exactly zero. Instead we have a breakdown of the shift-transmission process, as evidenced by the large shift transmission errors. As a consequence, the convergence process stalls.

iteration	$\lvert a_{20,19} \rvert$	shift transmission error
3	2.1×10^{-2}	1.2×10^{-15}
4	4.3×10^{-3}	1.1×10^{-15}
5	2.2×10^{-3}	1.7×10^{-1}
6	2.2×10^{-3}	7.1×10^{0}
7	2.2×10^{-3}	9.7×10^{0}

Table 2: Breakdown caused by failure to deflate infinite eigenvalue

Our second example is also a random Hessenberg-triangular pencil, but this one is 30×30 and has real entries. We apply the double-shift QZ algorithm ($m = 2$). The first four pairs of eigenvalues are deflated after 9, 12, 17, and 19 iterations, respectively. Quadratic convergence of $a_{n-1,n-2}$ to zero is observed.

Calculating the two finite eigenvalues of the bulge pencil near the end of the bulge chase, we observe that they never differ from the intended shifts by more then 2×10^{-14}. Thus the shift transmission error is a bit larger than in the case $m = 1$ but still tiny.

If we modify the pencil by setting $b_{13,13}$ and $b_{17,17}$ to zero, we get comparable results. Since the zeros float up by two positions per iteration, we deflate an infinite eigenvalue at the top after six iterations, and again after seven. While this is happening, finite eigenvalues are emerging at the bottom. The first four pairs of finite eigenvalues are deflated from the bottom after 7, 14, 20, and 26 iterations, respectively.[4] Quadratic convergence is observed.

Comparing the intended shifts with the eigenvalues of the bulge pencil near the end of the bulge chase, we find that the shift transmission error never exceeds 1.3×10^{-13}. Thus the shifts are transmitted effectively.

These good results depend upon deflation of the infinite eigenvalues as they emerge. If we fail to do this, the shift transmission mechanism breaks down, and progress toward convergence comes to a halt.

[4] Some of the eigenvalues are real, and they don't always emerge in pairs. For example, the "pair" of eigenvalues that emerged after 14 iterations was really two real eigenvalues that were deflated after 12 and 14 iterations, respectively.

7 Avoidance of infinite shifts

The infinite eigenvalues move naturally to the top of the pencil. If we wish to deflate finite eigenvalues rapidly at the bottom, then we should always use finite shifts. This means that if we choose our shifts in the usual way, i.e. as eigenvalues of the lower right-hand $m \times m$ subpencil $A_{22} - \lambda B_{22}$, we should make sure that none of the main diagonal entries of B_{22} is zero. If one of them is zero, we can chase the zero to the bottom at trivial cost and deflate an infinite eigenvalue there [1, §7.7.5]. Alternatively we can just ignore the infinite shift, replacing it by some finite shift chosen in any way. The subsequent QZ iteration will move the zero up and out of B_{22}.

Finally we note that if the Hessenberg-triangular form was obtained by the standard reduction algorithm [1, §7.7.4], then zero diagonal entries will not be found in B_{22}; they will be at the the top of B. This is so because the reduction algorithm has the same effect of moving zeros upward as the QZ iteration has. The reduction algorithm begins by transforming B to upper triangular form by a QR decomposition. Then it transforms A to Hessenberg form one entry at a time. Each entry is eliminated by a rotator on the left. The triangular form of B is then restored by a compensating rotator on the right. If these rotators act in the $(j-1, j)$ plane, and if there is a zero entry at b_{jj}, the combined effect of these two rotators will be to move the zero entry from the b_{jj} position up to $b_{j-1,j-1}$. The upward sweeping effect of the reduction algorithm is much stronger than that of the QZ iteration, because there are many more entries to eliminate from the full matrix A, and hence many more rotators are applied. In fact, the QZ iteration is just a sparse special case of the reduction algorithm.

References

[1] G. H. Golub and C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, third ed., 1996.

[2] C. B. Moler and G. W. Stewart, *An algorithm for generalized matrix eigenvalue problems*, SIAM J. Numer. Anal., 10 (1973), pp. 241–256.

[3] D. S. Watkins, *The transmission of shifts and shift blurring in the QR algorithm*, Linear Algebra Appl., 241–243 (1996), pp. 877–896.

[4] D. S. Watkins and L. Elsner, *Convergence of algorithms of decomposition type for the eigenvalue problem*, Linear Algebra Appl., 143 (1991), pp. 19–47.

[5] ——, *Theory of decomposition and bulge-chasing algorithms for the generalized eigenvalue problem*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 943–967.