SIAM J. MATRIX ANAL. APPL.
Vol. 15, No. 3, pp. 943-967, July 1994

016

THEORY OF DECOMPOSITION AND BULGE-CHASING ALGORITHMS FOR THE GENERALIZED EIGENVALUE PROBLEM*

DAVID WATKINS† AND LUDWIG ELSNER‡

Abstract. A generic GZ algorithm for the generalized eigenvalue problem $Ax = \lambda Bx$ is presented. This is actually a large class of algorithms that includes multiple-step QZ and LZ algorithms, as well as QZ-LZ hybrids, as special cases. First the convergence properties of the GZ algorithm are discussed, then a study of implementations is undertaken. The notion of an elimination rule is introduced as a device for studying the QZ, LZ and other algorithms simultaneously. To each elimination rule there corresponds an explicit GZ algorithm. Through a careful study of the steps involved in executing the explicit algorithm, it is discovered how to implement the algorithm implicitly by bulge chasing. The approach taken here was introduced by Miminis and Paige in the context of the QR algorithm for the ordinary eigenvalue problem. It is more involved than the standard approach, but it yields a much clearer picture of the relationship between the implicit and explicit versions of the algorithm. Furthermore, it is more general than the standard approach, as it does not require the use of a theorem of "Implicit-Q" type. Finally a generalization of the implicit GZ algorithm, the generic bulge-chasing algorithm, is introduced. It is proved that the generic bulge-chasing algorithm implicitly performs iterations of the generic GZ algorithm. Thus the convergence theorems that are proved for the generic GZ algorithm hold for the generic bulge-chasing algorithm as well.

Key words. generalized eigenvalue problem, QZ algorithm, GZ algorithm, chasing the bulge

AMS subject classifications. 65F15, 15A18

1. Introduction.

The standard algorithm for finding the eigenvalues of a dense, indefinite, matrix pencil $A - \lambda B$ with B nonsingular is the QZ algorithm of Moler and Stewart [11]. Related methods are the LZ algorithm of Kaufman [8] and the combination-shift QZ algorithm of Ward [14]. In this paper we introduce and study a generic GZ algorithm, which is actually a large class of algorithms that contains these and many other algorithms as special cases. For example, QZ-LZ hybrids are also included. Our coverage is not restricted to single- or double-step algorithms; we allow multiple steps of arbitrary multiplicity.

The QZ algorithm is an extension of the QR algorithm, which is one of the most widely used algorithms for the standard eigenvalue problem. The QR algorithm has both explicit and implicit versions. The explicit version is useful for introducing the algorithm and discussing theoretical aspects such as convergence theory, but it is usually the implicit version that is actually implemented. The standard approach to the QZ algorithm, as presented in contemporary textbooks [6], [12], mentions only an implicit version, which is interpreted as a way of applying the QR algorithm to the matrix AB^{-1} without actually forming AB^{-1} or even B^{-1}. Earlier approaches [11], [8] started from an explicit version and derived the implicit version therefrom. In every instance the focus was on the matrix AB^{-1}. Our approach also starts with an explicit version, but our explicit QZ algorithm differs from earlier formulations in that it effectively applies the QR algorithm to both AB^{-1} and $B^{-1}A$. The advantage of this approach is that it reveals symmetries in the algorithm that are obscured by the usual approaches. In particular, it puts the "Q" and "Z" transformations on an

* Received by the editors December 16, 1991; accepted for publication (in revised form) February 12, 1993.

† Department of Pure and Applied Mathematics, Washington State University, Pullman, Washington 99164-3113 (na.watkins@na-net.ornl.gov).

‡ Fakultät für Mathematik, Universität Bielefeld, Postfach 8640, W-4800 Bielefeld 1, Germany (elsner@mathi.mathematik.uni-bielefeld.de).

equal footing. Of course our discussion is couched in more general terms. We consider a generic GZ algorithm that amounts to the generic GR algorithm [17] applied to the matrices $p(AB^{-1})$ and $p(B^{-1}A)$ simultaneously, where p is a polynomial whose degree is the multiplicity of the step. Our generic GZ algorithm is quite similar to the FGZ algorithm of [16].

The explicit algorithm is not a practical algorithm, because it would be too costly to implement and quite likely unstable as well. However, it is a useful vehicle for both the study of convergence and the introduction of implicit versions of the algorithm. We introduce the generic GZ algorithm in §2. In §3 we study the convergence properties of the algorithm and in §4 we consider questions of implementation. Sections 3 and 4 can be read independently of one another.

The convergence theorem that we prove in §3 is a generalization of a theorem on the convergence of the GR algorithm that we proved in [17]. The theorem says roughly that if the eigenvalues can be separated, and the shifts converge, and the condition numbers of the accumulated transforming matrices remain bounded, then the algorithm converges. We also introduce the generalized Rayleigh quotient shift strategy and discuss its asymptotic convergence properties without proof. Usually the convergence rate is quadratic.

In §4 we consider how to implement the GZ algorithm. Our approach is inspired by Miminis and Paige [10]. They showed that by taking a detailed look at how one would carry out an iteration of the QR algorithm in its explicit form, one can discover how it can be done implicitly. As Miminis and Paige pointed out, this approach is more involved than the usual approach, which invokes the Implicit-Q Theorem [6, p. 367], but it gives a much clearer picture of the relationship between the explicit and implicit versions of the algorithm. Miminis and Paige also stated that their approach is quite general. Our vehicle for introducing the desired generality is the idea of an elimination rule, which allows us to discuss the QZ, LZ, and all related algorithms simultaneously. Each elimination rule gives rise to a specific implementation of the GZ algorithm. Following Miminis and Paige, we take a close look at the steps involved in implementing the GZ algorithm explicitly. By studying the form of the intermediate matrices so produced, we discover how the algorithm can be implemented implicitly, that is, without forming or operating on the matrices $p(AB^{-1})$ or $p(B^{-1}A)$.

Once we have derived the implicit GZ algorithm, we introduce a generalization called the generic bulge-chasing algorithm and prove that each iteration of the generic bulge-chasing algorithm amounts to an iteration of the generic GZ algorithm. The purpose of making this last generalization is to allow additional flexibility in implementing the algorithm. This flexibility can be exploited to build more efficient and stable algorithms. In particular it allows the introduction of variants that do not break down when B happens to be singular. (For the originators of the QZ and LZ algorithms this was an important point.) When it comes to implementing the algorithm in practice, these are the variants that should be used.

In [18] we introduced a generic bulge-chasing algorithm for the standard eigenvalue problem. Our aim there was to lay common foundations for implicit versions of GR algorithms of all types (e.g., QR, LR with or without pivoting, SR, hybrids, etc.). To achieve the desired level of generality, we devised an approach that, like the Miminis–Paige approach, avoids using a theorem of the Implicit-Q type. However, unlike Miminis and Paige, we did not establish a close correspondence between the operations in the implicit and explicit versions of the algorithms. The results of this paper generalize those of [18].

2. The generic GZ algorithm. We consider the generalized eigenvalue problem

$$(A - \lambda B)x = 0,$$

where A and B are square matrices whose entries are complex numbers. Recall that the pencil $A - \lambda B$ is said to be *singular* if its determinant is zero for all λ and *regular* otherwise. We focus here on the regular case. If the given pencil is singular (or not known a priori to be regular), the staircase algorithm of Van Dooren [13] can be used to remove the singular part. (See also Demmel and Kågström [3], [4].) This algorithm also removes the infinite eigenvalue and its associated structure (which may be present if B is singular) and the zero eigenvalue and its associated structure (which may be present if A is singular). What is left is a regular pencil for which both A and B are nonsingular. We assume throughout (with few exceptions, when we explicitly state otherwise) that our pencil has a nonsingular B; we do not need to assume that A is nonsingular.[1]

Recall that the pencils $A - \lambda B$ and $\hat{A} - \lambda \hat{B}$ are said to be *strictly equivalent* if there exist nonsingular matrices G and Z such that

$$\hat{A} = G^{-1}AZ \quad \text{and} \quad \hat{B} = G^{-1}BZ.$$

Strictly equivalent pencils have the same eigenvalues, and the eigenvectors are related in a simple way through the transforming matrices G and Z. The generic GZ algorithm generates a sequence of strictly equivalent pencils $(A_i - \lambda B_i)$ that converges (we hope) to upper triangular or block triangular form, thus exposing the eigenvalues of the pencil. The eigenvectors can be found by a back-substitution process that utilizes the final upper-triangular matrices and the accumulated transforming matrices.

We assume that before we start our iterations of the GZ algorithm, we transform the pencil to some initial form

$$A_0 = G_0^{-1}AZ_0, \qquad B_0 = G_0^{-1}BZ_0.$$

For example, it is possible to make A_0 upper Hessenberg and B_0 upper triangular, as described in [6] and elsewhere. Later on we assume that A_0 and B_0 have this form, but for now we allow them to have any form; for example, we could take $G_0 = Z_0 = I$.

The ith iteration of the GZ algorithm transforms $A_{i-1} - \lambda B_{i-1}$ to $A_i - \lambda B_i$ by transformations obtained from GR decompositions. By a GR decomposition of a square matrix M, we mean any decomposition

$$M = GR$$

in which G is nonsingular and R is upper triangular. Every matrix has many different GR decompositions. To obtain A_i and B_i we first take GR decompositions of $p_i(A_{i-1}B_{i-1}^{-1})$ and $p_i(B_{i-1}^{-1}A_{i-1})$, where p_i is a polynomial. Thus we find nonsingular G_i and Z_i and upper triangular R_i and S_i such that

$$p_i(A_{i-1}B_{i-1}^{-1}) = G_iR_i \quad \text{and} \quad p_i(B_{i-1}^{-1}A_{i-1}) = Z_iS_i.$$

[1] However, if A is known to be nonsingular, one has the possibility of reversing the roles of A and B and considering the pencil $B - \mu A$, where $\mu = 1/\lambda$. If neither A nor B is known to be nonsingular, a prudent course of action is to run the staircase algorithm to determine the fine structure of the pencil.

We then let

$$A_i = G_i^{-1} A_{i-1} Z_i \quad \text{and} \quad B_i = G_i^{-1} B_{i-1} Z_i.$$

In the special case $B_{i-1} = I$, $Z_i = G_i$, $S_i = R_i$, this algorithm reduces to the generic GR algorithm for the standard eigenvalue problem.

The GZ algorithm is really a large class of algorithms. Specific instances are obtained by specifying the exact form of each GR decomposition and how the p_i are to be chosen. For example, variants of the QZ and LZ algorithms are obtained by specifying that each decomposition be a QR or LR decomposition, respectively. The p_i are chosen so that their roots, which we call the *shifts* for the ith iteration, are estimates of eigenvalues. The degree of p_i is called the *multiplicity* of the iteration.

We will see that it is possible to carry out the GZ iterations implicitly without even calculating matrices of the form AB^{-1} or $B^{-1}A$, much less $p(AB^{-1})$ or $p(B^{-1}A)$. Were this not the case, there would be no point in discussing these algorithms at all. First we look at convergence.

3. Convergence of GZ algorithms. An easy computation shows that

$$A_i B_i^{-1} = G_i^{-1} (A_{i-1} B_{i-1}^{-1}) G_i.$$

Since G_i was obtained from the decomposition $p_i(A_{i-1}B_{i-1}^{-1}) = G_i R_i$, we see that the transformation $A_{i-1}B_{i-1}^{-1} \to A_i B_i^{-1}$ is an iteration of the GR algorithm [17]. At the same time we have

$$B_i^{-1} A_i = Z_i^{-1} (B_{i-1}^{-1} A_{i-1}) Z_i,$$

where

$$p_i(B_{i-1}^{-1} A_{i-1}) = Z_i S_i,$$

so the transformation $B_{i-1}^{-1}A_{i-1} \to B_i^{-1}A_i$ is also a GR iteration. It follows from the theorems in [17] that both of the sequences $(A_i B_i^{-1})$ and $(B_i^{-1} A_i)$ converge to (block) upper triangular form, provided that the condition numbers of the accumulated transforming matrices $\hat{G}_i = G_1 \cdots G_i$ and $\hat{Z}_i = Z_1 \cdots Z_i$ remain bounded and the shifts converge, as $i \to \infty$. Preferably the shifts should converge to eigenvalues of the pencil, in which case the convergence of $(A_i B_i^{-1})$ and $(B_i^{-1} A_i)$ is superlinear.

We would like to be able to say something about the convergence of the sequences (A_i) and (B_i) separately, since these are the matrices with which we actually work. To do this we recall some nomenclature. Let $\mathcal{T}_d$ and $\mathcal{T}_r$ be subspaces of $\mathbb{C}^n$ of equal dimension. The pair $(\mathcal{T}_d, \mathcal{T}_r)$ is called a *deflating pair* for the regular pencil $A - \lambda B$ if and only if

$$A\mathcal{T}_d \subseteq \mathcal{T}_r \quad \text{and} \quad B\mathcal{T}_d \subseteq \mathcal{T}_r.$$

The subscripts d and r are mnemonics for domain and range, respectively. Since we are assuming that B is nonsingular, the condition $B\mathcal{T}_d \subseteq \mathcal{T}_r$ implies $B\mathcal{T}_d = \mathcal{T}_r$. Clearly $(\mathcal{T}_d, \mathcal{T}_r)$ is a deflating pair for $A - \lambda B$ if and only if $\mathcal{T}_d$ is invariant under $B^{-1}A$, $\mathcal{T}_r$ is invariant under AB^{-1}, and $B\mathcal{T}_d = \mathcal{T}_r$. The following lemma generalizes Lemma 6.1 of [17]. Here $d(\mathcal{S}, \mathcal{T})$ denotes the usual distance (or gap) between two subspaces and $\kappa_2(G)$ denotes the condition number of G with respect to the spectral norm.

LEMMA 3.1. *Let $A, B \in \mathbb{C}^{n\times n}$ and let $(\mathcal{T}_d, \mathcal{T}_r)$ be a deflating pair of k-dimensional subspaces for the pencil $A-\lambda B$. Let Z, $G \in \mathbb{C}^{n\times n}$ be nonsingular matrices, and let $\mathcal{S}_d$ and $\mathcal{S}_r$ be the spaces spanned by the first k columns of Z and G, respectively. (Think of $\mathcal{S}_d$ and $\mathcal{S}_r$ as approximations to $\mathcal{T}_d$ and $\mathcal{T}_r$, respectively.) Let C denote either A or B, and let $\hat{C} = G^{-1}CZ$. Consider the partition*

$$\hat{C} = \begin{bmatrix} \hat{C}_{11} & \hat{C}_{12} \\ \hat{C}_{21} & \hat{C}_{22} \end{bmatrix},$$

where $\hat{C}_{11} \in \mathbb{C}^{k\times k}$. Then

$$\frac{\|\hat{C}_{21}\|_2}{\|\hat{C}\|_2} \le \sqrt{2}\kappa_2(G)\kappa_2(Z)\left[d(\mathcal{S}_d, \mathcal{T}_d) + d(\mathcal{S}_r, \mathcal{T}_r)\right].$$

Proof. Let $Z = PU$, $G = QR$ be the QR decompositions of Z and G, respectively. Thus P and Q are unitary, and U and R are upper triangular. Partition these decompositions as

$$\begin{bmatrix} Z_1 & Z_2 \end{bmatrix} = \begin{bmatrix} P_1 & P_2 \end{bmatrix} \begin{bmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{bmatrix},$$

where Z_1 and P_1 are $n \times k$, and similarly for the decomposition $G = QR$. Since $\hat{C} = G^{-1}CZ = R^{-1}Q^*CPU$, we have $\hat{C}_{21} = R_{22}^{-1}Q_2^*CP_1U_{11}$, from which

$$\|\hat{C}_{21}\|_2 \le \|R_{22}^{-1}\|_2 \|Q_2^*CP_1\|_2 \|U_{11}\|_2.$$

Since $\|R_{22}^{-1}\|_2 \le \|R^{-1}\|_2 = \|G^{-1}\|_2$, $\|U_{11}\|_2 \le \|U\|_2 = \|Z\|_2$, and $\|C\|_2 \le \|G\|_2\|\hat{C}\|_2\|Z^{-1}\|_2$, we see that

$$\frac{\|\hat{C}_{21}\|_2}{\|\hat{C}\|_2} \le \kappa_2(G)\kappa_2(Z)\frac{\|Q_2^*CP_1\|_2}{\|C\|_2}. \tag{1}$$

Since $Z_1 = P_1U_{11}$ and $G_1 = Q_1R_{11}$, we have $\mathcal{S}_d = \mathcal{R}(Z_1) = \mathcal{R}(P_1)$ and $\mathcal{S}_r^{\perp} = \mathcal{R}(G_1)^{\perp} = \mathcal{R}(Q_1)^{\perp} = \mathcal{R}(Q_2)$. Therefore, by Lemma 4.1 of [17], there exist $T_1 \in \mathbb{C}^{n\times k}$ and $T_2 \in \mathbb{C}^{n\times n-k}$ with orthonormal columns, such that $\mathcal{T}_d = \mathcal{R}(T_1)$, $\mathcal{T}_r^{\perp} = \mathcal{R}(T_2)$,

$$\|P_1 - T_1\|_2 \le \sqrt{2}d(\mathcal{S}_d, \mathcal{T}_d)$$

and

$$\|Q_2 - T_2\|_2 \le \sqrt{2}d(\mathcal{S}_r, \mathcal{T}_r).$$

We use here the fact that $d(\mathcal{S}_r, \mathcal{T}_r) = d(\mathcal{S}_r^{\perp}, \mathcal{T}_r^{\perp})$. Now

$$\|Q_2^*CP_1\|_2 \le \|(Q_2 - T_2)^*CP_1\|_2 + \|T_2^*C(P_1 - T_1)\|_2 + \|T_2^*CT_1\|_2.$$

Since $\mathcal{R}(T_1) = \mathcal{T}_d$, $C\mathcal{T}_d \subseteq \mathcal{T}_r$, and $\mathcal{R}(T_2) = \mathcal{T}_r^{\perp}$, the product $T_2^*CT_1$ is zero. Thus

$$\begin{aligned}\|Q_2^*CP_1\|_2 &\le \|Q_2 - T_2\|_2\|C\|_2\|P_1\|_2 + \|T_2\|_2\|C\|_2\|P_1 - T_1\|_2 \\ &\le \sqrt{2}\,\|C\|_2\,[d(\mathcal{S}_d, \mathcal{T}_d) + d(\mathcal{S}_r, \mathcal{T}_r).]\end{aligned}$$

Combining this inequality with (1), we obtain the desired result. ∎

Define the cumulative transforming matrices by

$$\hat{G}_i = G_1 \cdots G_i, \qquad \hat{R}_i = R_i \cdots R_1,$$

$$\hat{Z}_i = Z_1 \cdots Z_i, \qquad \hat{S}_i = S_i \cdots S_1.$$

Then

$$C_i = \hat{G}_i^{-1} C_0 \hat{Z}_i,$$

where C can stand for either A or B. In the following theorem we prove the convergence of the GZ algorithm by applying Lemma 3.1 with the roles of C, G, Z, and $\hat{C}$ played by C_0, $\hat{G}_i$, $\hat{Z}_i$, and C_i, respectively. The symbol $\langle e_1, \ldots, e_k \rangle$ denotes the space spanned by the vectors $e_1, \ldots, e_k$.

THEOREM 3.2. *Let A_0, $B_0 \in \mathbb{C}^{n \times n}$ with B_0 nonsingular, and let p be a polynomial of degree $\leq n$. Let $\lambda_1, \ldots, \lambda_n$ denote the eigenvalues of the pencil $A_0 - \lambda B_0$, ordered so that $|p(\lambda_1)| \geq |p(\lambda_2)| \geq \cdots \geq |p(\lambda_n)|$. Suppose k is a positive integer less than n such that $|p(\lambda_k)| > |p(\lambda_{k+1})|$, and let $\rho = |p(\lambda_{k+1})|/|p(\lambda_k)|$. Let (p_i) be a sequence of polynomials of degree $\leq n$ such that $\lim_{i \to \infty} p_i = p$ and $p_i(\lambda_j) \neq 0$ for $j = 1, \ldots, k$ and all i. Let $(\mathcal{T}_d, \mathcal{T}_r)$ and $(\mathcal{U}_d, \mathcal{U}_r)$ be the deflating subspaces of $A_0 - \lambda B_0$ associated with $\lambda_1, \ldots, \lambda_k$ and $\lambda_{k+1}, \ldots, \lambda_n$, respectively. Suppose $\langle e_1, \ldots, e_k \rangle \cap \mathcal{U}_d = \{0\}$ and $\langle e_1, \ldots, e_k \rangle \cap \mathcal{U}_r = \{0\}$. Let $(A_i - \lambda B_i)$ be the sequence of iterates of the GZ algorithm using the given (p_i), starting from $A_0 - \lambda B_0$. If there is a constant $\hat{\kappa}$ such that the cumulative transforming matrices $\hat{G}_i$ and $\hat{Z}_i$ all satisfy $\kappa_2(\hat{G}_i) \leq \hat{\kappa}$ and $\kappa_2(\hat{Z}_i) \leq \hat{\kappa}$, then $(A_i - \lambda B_i)$ tends to block triangular form in the following sense. Let C_i denote either A_i or B_i, and partition C_i as*

$$C_i = \begin{bmatrix} C_{11}^{(i)} & C_{12}^{(i)} \\ C_{21}^{(i)} & C_{22}^{(i)} \end{bmatrix},$$

where $C_{21}^{(i)} \in \mathbb{C}^{(n-k) \times k}$. Then for every $\hat{\rho}$ satisfying $\rho < \hat{\rho} < 1$ there exists a constant M such that

$$\frac{\| C_{21}^{(i)} \|_2}{\| C_i \|_2} \leq M \hat{\rho}^i \quad \text{for all } i. \tag{2}$$

Remark 1. If A_0 is upper Hessenberg with no zeros on the subdiagonal and B_0 is upper triangular, then the subspace conditions $\langle e_1, \ldots, e_k \rangle \cap \mathcal{U}_d = \{0\}$ and $\langle e_1, \ldots, e_k \rangle \cap \mathcal{U}_r = \{0\}$ are satisfied for all k, as is explained in [17] for the standard eigenvalue problem. The reasoning is no different for the generalized eigenvalue problem.

Remark 2. The conditions $p_i(\lambda_j) \neq 0$ for $j = 1, \ldots, k$ may occasionally be violated, but this is not undesirable. If $p_i(\lambda_j) = 0$, then $p_i(A_i B_i^{-1})$ is singular. Theorem 4.3 shows that in this case the eigenvalue λ_j can be deflated from the problem after the ith iteration.

Remark 3. The conclusion of the theorem implies that the eigenvalues of $A_{11}^{(i)} - \lambda B_{11}^{(i)}$ and $A_{22}^{(i)} - \lambda B_{22}^{(i)}$ converge to $\lambda_1, \ldots, \lambda_k$ and $\lambda_{k+1}, \ldots, \lambda_n$, respectively, as can be shown by standard techniques.

Remark 4. If p has $\lambda_{k+1}, \ldots, \lambda_n$ among its roots, then $\rho = 0$, so (2) holds for all $\hat{\rho} > 0$. Thus the convergence is superlinear.

Remark 5. The hypotheses of the theorem usually hold for many values of k simultaneously, thereby giving a limiting form that is block triangular with many small blocks on the main diagonal. If the conditions hold for all k $(1 \le k \le n-1)$, the limiting form is upper triangular.

Proof. Let $\hat{p}_i = p_i \cdots p_1$, let $\mathcal{S} = \langle e_1, \ldots, e_k \rangle$, $\mathcal{S}_{ri} = \hat{p}_i(A_0 B_0^{-1})\mathcal{S}$, and $\mathcal{S}_{di} = \hat{p}_i(B_0^{-1}A_0)\mathcal{S}$. All of the hypotheses of Theorem 5.4 of [17] are satisfied, with the role of A in that theorem played by either $A_0 B_0^{-1}$ or $B_0^{-1}A_0$. Consequently there exists $\hat{M}$ such that $d(\mathcal{S}_{ri}, \mathcal{T}_r) \le \hat{M}\hat{\rho}^i$ and $d(\mathcal{S}_{di}, \mathcal{T}_d) \le \hat{M}\hat{\rho}^i$. Recall that (as shown in [17] and elsewhere) $\hat{p}_i(A_0 B_0^{-1}) = \hat{G}_i \hat{R}_i$ and $\hat{p}_i(B_0^{-1}A_0) = \hat{Z}_i \hat{S}_i$. Consider the partition $\hat{G}_i = [\hat{G}_1^{(i)} \ \hat{G}_2^{(i)}]$, $\hat{Z}_i = [\hat{Z}_1^{(i)} \ \hat{Z}_2^{(i)}]$, where $\hat{G}_1^{(i)}$, $\hat{Z}_1^{(i)} \in \mathbb{C}^{n \times k}$. Since $\hat{R}_i$ and $\hat{S}_i$ are upper triangular, $\mathcal{S}_{ri} = \mathcal{R}(\hat{G}_1^{(i)})$ and $\mathcal{S}_{di} = \mathcal{R}(\hat{Z}_1^{(i)})$. This is true even if $\hat{p}_i(A_0 B_0^{-1})$ and $\hat{p}_i(B_0^{-1}A_0)$ are singular, as the assumptions $\mathcal{S} \cap \mathcal{U}_d = \{0\}$, $\mathcal{S} \cap \mathcal{U}_r = \{0\}$, and $|p_i(\lambda_j)| > 0$, $j = 1, \ldots, k$ guarantee that $\mathcal{S}$ contains no nontrivial null vectors of $\hat{p}_i(A_0 B_0^{-1})$ or $\hat{p}_i(B_0^{-1}A_0)$. Therefore the spaces $\mathcal{S}_{di}$ and $\mathcal{S}_{ri}$ have dimension k for all i. Applying Lemma 3.1 with the roles of C, G, Z, and $\hat{C}$ played by C_0, $\hat{G}_i$, $\hat{Z}_i$, and C_i, respectively, we conclude that

$$\frac{\|C_{21}^{(i)}\|_2}{\|C_i\|_2} \le 2\sqrt{2}\kappa_2(\hat{G}_i)\kappa_2(\hat{Z}_i)\hat{M}\hat{\rho}^i \le M\hat{\rho}^i,$$

where $M = 2\sqrt{2}\hat{\kappa}^2\hat{M}$. □

3.1. The generalized Rayleigh quotient shift. Suppose we plan to perform GZ iterations of multiplicity m, where $m \ll n$. A natural way of choosing the shift polynomials is to let p_i be the characteristic polynomial of the $m \times m$ lower right-hand corner pencil $A_{22}^{(i)} - \lambda B_{22}^{(i)}$. We call this the *generalized Rayleigh quotient shift* strategy.

In [17] we proved that for the standard eigenvalue problem, the asymptotic convergence rate of the GR algorithm with the generalized Rayleigh quotient shift strategy is quadratic, provided that the eigenvalues of the given matrix are simple. This result also holds for the generalized eigenvalue problem. Specifically, if the GZ algorithm converges under the conditions of Theorem 3.2, and generalized Rayleigh quotient shifts with $m = n - k$ are used, the asymptotic convergence rate will be quadratic, provided $A_0 B_0^{-1}$ is simple. We omit the proof. The details are more tedious than they are for the standard problem, but the ideas are the same.

4. Implementation of GZ algorithms. We assume from now on that the initial transformation

$$A_0 = G_0^{-1} A Z_0, \qquad B_0 = G_0^{-1} B Z_0$$

makes A_0 upper Hessenberg and B_0 upper triangular. We even assume that A_0 is a *proper* upper Hessenberg matrix; that is, all of its subdiagonal entries $a_{i+1,i}^{(0)}$ are nonzero. This implies no loss of generality, for if some of the entries $a_{i+1,i}^{(0)}$ are zero, we can reduce the problem to two or more subproblems, each of which has a properly upper Hessenberg coefficient matrix. Since we are now concerned with the problem of implementing one iteration of the GZ algorithm, we drop the subscripts and consider the single iteration

$$\hat{A} = \hat{G}^{-1} A \hat{Z}, \qquad \hat{B} = \hat{G}^{-1} B \hat{Z}, \tag{3}$$

where

$$p(AB^{-1}) = \hat{G}\hat{R}, \qquad p(B^{-1}A) = \hat{Z}\hat{S}. \tag{4}$$

Here all matrices are in $\mathbb{C}^{n\times n}$. The degree of the polynomial p is m, which is assumed to be less than n. Normally $m \ll n$. Dropping the subscripts allows us to reintroduce subscripts later for a different purpose.

An important relationship that follows directly from (3) and (4) is given in the following lemma, which plays a key role in determining the structure of $\hat{A}$, $\hat{B}$ and intermediate matrices that arise during the execution of a GZ iteration. (It is a generalization of [10, (3.5)].) Although the lemma is used to study structure, it is not itself dependent on any special structure of the matrices involved, except that it is crucial that $\hat{G}$ and $\hat{Z}$ be nonsingular.

LEMMA 4.1. *Suppose A, $\hat{A}$, B, $\hat{B}$, $\hat{R}$, and $\hat{S}$ are any $n \times n$ matrices related by (3) and (4), where $\hat{G}$ and $\hat{Z}$ are nonsingular matrices, and p is a polynomial. Then*

$$\hat{A}\hat{S} = \hat{R}A \quad \textit{and} \quad \hat{B}\hat{S} = \hat{R}B.$$

Proof. $\hat{A}\hat{S} = \hat{A}\hat{Z}^{-1}p(B^{-1}A) = \hat{G}^{-1}Ap(B^{-1}A) = \hat{G}^{-1}p(AB^{-1})A = \hat{R}A$. The same argument shows that $\hat{B}\hat{S} = \hat{R}B$, since the equation $Bp(B^{-1}A) = p(AB^{-1})B$ also holds. □

As a first application of Lemma 4.1, consider a GZ iteration in which the matrices $p(AB^{-1})$ and $p(B^{-1}A)$ are nonsingular, as is usually the case. Then $\hat{R}$ and $\hat{S}$ are also nonsingular, so the equations in Lemma 4.1 can be rewritten in the form

$$\hat{A} = \hat{R}A\hat{S}^{-1}, \qquad \hat{B} = \hat{R}B\hat{S}^{-1}.$$

Since A is properly upper Hessenberg and B, $\hat{R}$, and $\hat{S}^{-1}$ are all upper triangular, we see immediately that $\hat{A}$ is properly upper Hessenberg and $\hat{B}$ is upper triangular. Thus the special form is preserved from one iteration to the next.

4.1. The singular case. When $p(AB^{-1})$ and $p(B^{-1}A)$ are singular, the upper Hessenberg-triangular form is not preserved, but something even better happens. A small subpencil at the lower right-hand corner of the matrix can be deflated from the pencil after the iteration. The part of the pencil that remains after deflation remains in Hessenberg-triangular form. We consider this case in detail.

The matrices AB^{-1} and $B^{-1}A$ are both properly upper Hessenberg. The proper upper Hessenberg matrix W that appears in the following lemma is taken to be AB^{-1} or $B^{-1}A$ in the application.

If $p(W)$ is singular, then at least one of the shifts (roots of p) is an eigenvalue of W, and conversely. Any shift that is an eigenvalue is called a *perfect* shift.

LEMMA 4.2. *Let W be a proper upper Hessenberg matrix, and let p be a polynomial that has ν roots that are perfect shifts for W. Then*

$$\operatorname{rank}(p(W)) = n - \nu.$$

Furthermore, the leading $n - \nu$ columns of $p(W)$ are linearly independent.

Remark. When we count perfect shifts, we allow repeated shifts, but we count a repeated shift no more times than it appears as a root of the characteristic polynomial of W.

This result is also proved in [10] as part of Theorem 4.1.

Proof. The statement about the rank is just Lemma 4.4 of [18]. To get the other assertion, let $x = p(W)e_1$. Let $K(W,x)$ denote the Krylov matrix

$$K(W,x) = [x, Wx, W^2x, \ldots, W^{n-1}x].$$

Then

$$K(W,x) = p(W)T, \tag{5}$$

where $T = K(W,e_1)$. Since W is properly upper Hessenberg, T is upper triangular and nonsingular. Thus, for any k, the span of the first k columns of $K(W,x)$ is the same as the span of the first k columns of $p(W)$. In particular, $\operatorname{rank}(K(W,x)) = n-\nu$. The form of a Krylov matrix implies that if a given column is a linear combination of previous columns, all subsequent columns will also be linear combinations of the previous columns. Thus the first $n-\nu$ columns of $K(W,x)$, and hence also of $p(W)$, must be linearly independent. □

THEOREM 4.3. *Consider the GZ iteration* (3), (4), *in which* ν *of the shifts are eigenvalues of* $A - \lambda B$. *Then*

$$\hat{R} = \begin{bmatrix} \hat{R}_{11} & \hat{R}_{12} \\ 0 & 0 \end{bmatrix}, \qquad \hat{S} = \begin{bmatrix} \hat{S}_{11} & \hat{S}_{12} \\ 0 & 0 \end{bmatrix},$$

$$\hat{A} = \begin{bmatrix} \hat{A}_{11} & \hat{A}_{12} \\ 0 & \hat{A}_{22} \end{bmatrix}, \quad \textit{and} \quad \hat{B} = \begin{bmatrix} \hat{B}_{11} & \hat{B}_{12} \\ 0 & \hat{B}_{22} \end{bmatrix},$$

where $\hat{R}_{11}, \hat{S}_{11}, \hat{A}_{11}, \hat{B}_{11} \in \mathbb{C}^{(n-\nu)\times(n-\nu)}$, $\hat{R}_{11}$ *and* $\hat{S}_{11}$ *are nonsingular,* $\hat{A}_{11}$ *is properly upper Hessenberg, and* $\hat{B}_{11}$ *is upper triangular. The eigenvalues of the subpencil* $\hat{A}_{22} - \lambda\hat{B}_{22}$ *are exactly the* ν *perfect shifts.*

Proof. In light of (4) and Lemma 4.2, the upper triangular matrices $\hat{R}$ and $\hat{S}$ both have rank $n-\nu$ and their first $n-\nu$ columns are linearly independent. This proves that they have the stated form.

Writing the equations of Lemma 4.1 in partitioned form, we have

$$\begin{bmatrix} \hat{C}_{11} & \hat{C}_{12} \\ \hat{C}_{21} & \hat{C}_{22} \end{bmatrix} \begin{bmatrix} \hat{S}_{11} & \hat{S}_{12} \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \hat{R}_{11} & \hat{R}_{12} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix},$$

where C can denote either A or B. Equating the $(2,1)$ blocks of the partitioned equation, we find that $\hat{C}_{21}\hat{S}_{11} = 0$. Since $\hat{S}_{11}$ is nonsingular, we have $\hat{C}_{21} = 0$. Equating the $(1,1)$ blocks and multiplying on the right by $\hat{S}_{11}^{-1}$, we get

$$\hat{C}_{11} = \hat{R}_{11}C_{11}\hat{S}_{11}^{-1} + \hat{R}_{12}C_{21}\hat{S}_{11}^{-1}.$$

In the case $C = B$, we have $B_{21} = 0$, so $\hat{B}_{11} = \hat{R}_{11}B_{11}\hat{S}_{11}^{-1}$, which shows that $\hat{B}_{11}$ is upper triangular. Now consider the case $C = A$. Since A is upper Hessenberg, $A_{21} = \alpha e_1 e_{n-\nu}^T$, where $\alpha = a_{n-\nu+1,n-\nu}$. Since $\hat{S}_{11}^{-1}$ is upper triangular, we have $\alpha e_{n-\nu}^T\hat{S}_{11}^{-1} = \beta e_{n-\nu}^T$ for some β. Let $x = \hat{R}_{12}e_1$. Then

$$\hat{A}_{11} = \hat{R}_{11}A_{11}\hat{S}_{11}^{-1} + \beta x e_{n-\nu}^T.$$

The first term on the right-hand side is properly upper Hessenberg and the second term has nonzero entries only in the last column. Thus $\hat{A}_{11}$ is a proper upper Hessenberg matrix.

The fact that the eigenvalues of $\hat{A}_{22} - \lambda\hat{B}_{22}$ are just the ν perfect shifts can be deduced in the same way as in the standard eigenvalue problem by considering the form of $\hat{A}\hat{B}^{-1}$. See Theorem 4.5 of [18]. □

Remarks. We have opted for a brief algebraic proof using Lemma 4.1. Alternatively one could prove the form of $\hat{A}$ and $\hat{B}$ geometrically, using the Hessenberg form of A, the relationships between the underlying subspaces (spanned by the leading columns of $\hat{G}$ and $\hat{Z}$), and the fact that the $n-\nu$ dimensional spaces $\mathcal{T}_d = \mathcal{R}(p(B^{-1}A))$ and $\mathcal{T}_r = \mathcal{R}(p(AB^{-1}))$ form a deflating pair for $A - \lambda B$. Such a proof would be lengthier but perhaps more revealing.

Theorem 4.3 generalizes Theorem 4.5 of [18] and some aspects of Theorem 4.1 of Miminis and Paige [10]. However, the Miminis–Paige result addresses certain details that we have chosen to ignore.

Theorem 4.3 shows that if ν of the shifts are perfect, then a $\nu \times \nu$ subpencil can be deflated from the problem after the iteration. The pencil $\hat{A}_{22} - \lambda\hat{B}_{22}$ may not have Hessenberg-triangular form, but it is normally small enough that it can easily be returned to that form and its eigenvalues found. Subsequent iterations can focus on the pencil $\hat{A}_{11} - \lambda\hat{B}_{11}$, which does have Hessenberg-triangular form. Of course this is only a theoretical result. In a GZ step with roundoff errors, $\hat{A}_{21}$ will be not quite zero. Usually it will be far enough from zero to prevent deflation. In that case a subsequent GZ step with the same p will often produce the deflation.

4.2. *GR* decompositions and elimination rules. To introduce specific versions of the GZ algorithm, we need to consider how GR decompositions are carried out in practice. The standard way to perform a GR decomposition of any type is to "reduce the matrix to triangular form" by introducing zeros into the matrix one column at a time. Each column of zeros is obtained by multiplying on the left by a nonsingular matrix of a specified form. Algorithms of this type have the following general structure. A matrix $M \in \mathbb{C}^{n\times n}$ is reduced to upper triangular form in $n-1$ steps. After $i-1$ steps, M has been transformed to a matrix $\hat{R}_{i-1}$ whose first $i-1$ columns have been reduced to upper triangular form. That is,

$$\hat{R}_{i-1} = \begin{bmatrix} T & E \\ 0 & F \end{bmatrix},$$

where $T \in \mathbb{C}^{i-1\times i-1}$ is upper triangular. The ith step transforms $\hat{R}_{i-1}$ to $\hat{R}_i = G_i^{-1}\hat{R}_{i-1}$, where G_i has the form

$$G_i = \begin{bmatrix} I & 0 \\ 0 & \tilde{G}_i \end{bmatrix},$$

and $\tilde{G}_i \in \mathbb{C}^{(n-i+1)\times(n-i+1)}$ is chosen so that $\tilde{G}_i^{-1}x = \alpha e_1$, where x is the first column of F, and α is a scalar. After $n-1$ such steps, M will have been transformed to the upper triangular matrix $\hat{R} = \hat{R}_{n-1}$. Clearly $\hat{R} = G_{n-1}^{-1}\cdots G_1^{-1}M$, or

$$M = \hat{G}\hat{R},$$

where $\hat{G} = G_1\cdots G_{n-1}$.

Given any vector $x \in \mathbb{C}^n$ (for any $m \geq 2$) we say that a matrix $\tilde{G} \in \mathbb{C}^{m\times m}$ is an *elimination matrix* for x if $\tilde{G}$ is nonsingular and $\tilde{G}^{-1}x = \alpha e_1$ for some scalar α. If $\tilde{G}$ is an elimination matrix for x, then $\tilde{G}$ is an elimination matrix for all nonzero multiples of x.

Usually an elimination matrix $\tilde{G}$ is embedded in a larger matrix G. We also refer to the larger matrix as an elimination matrix.

An *elimination rule* is a map $x \mapsto \tilde{G}$ having the following properties. (i) The domain of the map is a subset of $\bigcup_{i=2}^{\infty} \mathbb{C}^i$. (ii) Each vector x in the domain is mapped to a matrix $\tilde{G}$ that is an elimination matrix for x. (iii) The map is homogeneous, that is, if x is in the domain, so are all nonzero multiples of x, and they are all mapped to the same elimination matrix. (iv) Zero vectors are in the domain, and each is mapped to the identity matrix of the same size. (v) If $x = \begin{bmatrix} y \\ 0 \end{bmatrix} \in \mathbb{C}^n$, where $y \in \mathbb{C}^k$ with $k < n$, then x is in the domain if and only if y is, and x is mapped to the matrix $\tilde{G} = \operatorname{diag}\{\tilde{H}, I_{n-k}\}$, where $\tilde{H}$ is the elimination matrix assigned to y by the map.

A *complete* elimination rule is one whose domain is all of $\bigcup_{i=2}^{\infty} \mathbb{C}^i$. A *partial* elimination rule is one whose domain is a proper subset of $\bigcup_{i=2}^{\infty} \mathbb{C}^i$.

Probably the simplest elimination rule is *Gaussian elimination without pivoting.* It is a partial elimination rule, as it is undefined on those nonzero x that satisfy $x_1 = 0$. An example of a complete elimination rule is Gaussian elimination with pivoting, which interchanges x_1 with the entry in x of largest magnitude before performing the elimination. Another complete elimination rule is *elimination by reflector* (*Householder transformation*). All of these types of elimination are discussed in [6] and [15], for example. One can also build hybrid elimination rules from other rules. For example, one can pick a tolerance τ satisfying $0 < \tau < 1$ and specify that x should be eliminated by Gaussian elimination (without pivoting) if $\max_{2\leq i\leq k} |x_i| \leq \tau |x_1|$ and by a reflector otherwise. This type of strategy has been used successfully in some of the algorithms in [7]. There are also more exotic types of elimination rules. For example, a symplectic (partial) elimination rule, which gives rise to the SR algorithm, is described in [1].

Every elimination rule induces a rule for carrying out GR decompositions; namely, carry out the "reduction to triangular form" described above using the specified elimination rule. Hence each elimination rule, together with a mechanism for choosing p, induces a GZ algorithm. If the elimination rule is not complete, the algorithm will break down (fail) if at some point it needs to perform an elimination on a vector that is not in the domain of the rule.

4.3. The explicit GZ algorithm. We now assume that we have chosen an elimination rule and will perform all of our eliminations with that rule.[2] Let us examine closely the steps involved in performing a GZ iteration explicitly. First the matrices $p(AB^{-1})$ and $p(B^{-1}A)$ are calculated. Then GR decompositions of both matrices are performed, using our chosen elimination rule. As above, we assume that ν of the shifts are perfect. If $\nu > 0$, then by Theorem 4.3, the resulting upper triangular matrices $\hat{R}$ and $\hat{S}$ have rank $n - \nu$, and their bottom ν rows zero. This

[2] However, everything we do could be cast in greater generality. For example, we could allow a different rule to be used at each step of the decomposition. Another possibility is to prescribe different rules for the two different GR decompositions on which the GZ iteration is based. Such an algorithm was once proposed by Kaufman [9]. In this algorithm all of the "G" transformations are unitary and all of the "Z" transformations are stabilized elementary (i.e., Gaussian elimination) transformations. This is a GZ algorithm in which the decomposition of $p(AB^{-1})$ is a QR decomposition and that of $p(B^{-1}A)$ is an LR decomposition with partial pivoting.

implies that the reductions to triangular form will be completed after $n - \nu$ steps. Thus, letting $\rho = \min\{n - \nu, n - 1\}$, the reductions have the form

$$\hat{R} = G_\rho^{-1} \cdots G_2^{-1} G_1^{-1} p(AB^{-1}),$$

$$\hat{S} = Z_\rho^{-1} \cdots Z_2^{-1} Z_1^{-1} p(B^{-1}A).$$

Therefore

$$p(AB^{-1}) = \hat{G}\hat{R} \quad \text{and} \quad p(B^{-1}A) = \hat{Z}\hat{S},$$

where

$$\hat{G} = G_1 \cdots G_\rho \quad \text{and} \quad \hat{Z} = Z_1 \cdots Z_\rho.$$

We then complete the iteration by performing the equivalence transformations

$$\hat{A} = \hat{G}^{-1} A \hat{Z} \quad \text{and} \quad \hat{B} = \hat{G}^{-1} B \hat{Z}. \tag{6}$$

Remark. In the case of the standard eigenvalue problem ($B = I$), we have $p(AB^{-1}) = p(B^{-1}A) = p(A)$, so $G_i = Z_i$, $i = 1, \ldots, \rho$, $\hat{G} = \hat{Z}$, $\hat{R} = \hat{S}$, $\hat{A} = \hat{G}^{-1} A \hat{G}$, and $\hat{B} = I$. This is one iteration of a GR algorithm [17].

To determine how to do these operations implicitly, we break the transformations (6) down into small steps and study the intermediate results. Let C denote A or B, as before; define $\hat{C}_0 = C$, and

$$\left.\begin{aligned} \hat{C}_{i-1/2} &= G_i^{-1} \hat{C}_{i-1} \\ \hat{C}_i &= \hat{C}_{i-1/2} Z_i \end{aligned}\right\} \quad i = 1, \ldots, \rho. \tag{7}$$

Then $\hat{C} = \hat{C}_\rho$.

We also give names to the intermediate matrices in the GR decompositions. Let $\hat{R}_0 = p(AB^{-1})$, $\hat{S}_0 = p(B^{-1}A)$, and

$$\left.\begin{aligned} \hat{R}_i &= G_i^{-1} \hat{R}_{i-1} \\ \hat{S}_i &= Z_i^{-1} \hat{S}_{i-1} \end{aligned}\right\} \quad i = 1, \ldots, \rho.$$

Then $\hat{R} = \hat{R}_\rho$ and $\hat{S} = \hat{S}_\rho$.

Since AB^{-1} and $B^{-1}A$ are proper upper Hessenberg matrices, $\hat{R}_0$ and $\hat{S}_0$ satisfy

$$\hat{r}_{j+m,j}^{(0)} \neq 0, \quad \hat{s}_{j+m,j}^{(0)} \neq 0, \quad j = 1, \ldots, n - m$$

and

$$\hat{r}_{ij}^{(0)} = \hat{s}_{ij}^{(0)} = 0 \quad \text{when } i > j + m,$$

where m is the degree of p. Thus they have the form

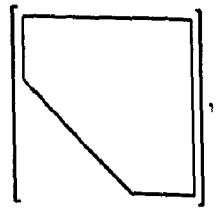

where the entries outside of the outlined area are all zero. Since G_1 is the elimination matrix for the first column of $\hat{R}_0$, $G_1 = \text{diag}\{\tilde{G}_1, I\}$, where $\tilde{G}_1 \in \mathbb{C}^{(m+1)\times(m+1)}$ is the elimination matrix of $[\hat{r}_{11}^{(0)}, \ldots, \hat{r}_{m+1,1}^{(0)}]^T$. The form of Z_1 is similar. In general $\hat{R}_{i-1}$ and $\hat{S}_{i-1}$ have the form

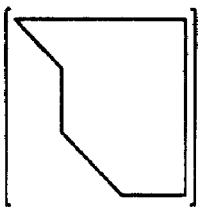

for $i = 2, \ldots, n-m$. The first $i-1$ columns are in upper triangular form. The ith column has m nonzeros below the main diagonal. Thus $G_i = \text{diag}\{I_{i-1}, \tilde{G}_i, I\}$ and $Z_i = \text{diag}\{I_{i-1}, \tilde{Z}_i, I\}$, where $\tilde{G}_i$ and $\tilde{Z}_i$ are the elimination matrices of

$$[\hat{r}_{ii}^{(i-1)}, \ldots, \hat{r}_{m+i,i}^{(i-1)}]^T \quad \text{and} \quad [\hat{s}_{ii}^{(i-1)}, \ldots, \hat{s}_{m+i,i}^{(i-1)}]^T,$$

respectively. For $i > n-m$ the transformations have the same form, except that the vectors to be eliminated are shorter because we have reached the bottom of the matrix. We then have $G_i = \text{diag}\{I_{i-1}, \tilde{G}_i\}$, where $\tilde{G}_i$ is the elimination matrix of $[\hat{r}_{ii}^{(i-1)}, \ldots, \hat{r}_{n,i}^{(i-1)}]^T$, and similarly for Z_i.

Let $\hat{G}_i = G_1 \cdots G_i$ and $\hat{Z}_i = Z_1 \cdots Z_i$, $i = 1, \ldots, \rho$. Then for $i < n-m$, $\hat{G}_i$ has the form

$$\hat{G}_i = \begin{bmatrix} \hat{G}_{11}^{(i)} & 0 \\ 0 & I \end{bmatrix},$$

where $\hat{G}_{11}^{(i)}$ is $(m+i) \times (m+i)$. This is clear from the form of the factors. The form of $\hat{Z}_i$ is the same.

The initial pencil $A - \lambda B = \hat{A}_0 - \lambda \hat{B}_0$ is in Hessenberg-triangular form, and so is the final pencil $\hat{A} - \lambda \hat{B} = \hat{A}_\rho - \lambda \hat{B}_\rho$, except possibly for a small subpencil that can be removed by deflation. The intermediate pencils $\hat{A}_{i-1/2} - \lambda \hat{B}_{i-1/2}$ and $\hat{A}_i - \lambda \hat{B}_i$ are not in Hessenberg-triangular form, but, as we shall see, they do not deviate from it by too much. First note that

$$\hat{A}_i = \hat{G}_i^{-1} A \hat{Z}_i \quad \text{and} \quad \hat{B}_i = \hat{G}_i^{-1} B \hat{Z}_i, \tag{8}$$

for $i = 1, \ldots, \rho$. Since also

$$p(AB^{-1}) = \hat{G}_i \hat{R}_i \quad \text{and} \quad p(B^{-1}A) = \hat{Z}_i \hat{S}_i, \tag{9}$$

we see that the pencil $\hat{A}_i - \lambda \hat{B}_i$ is the result of a partial GZ iteration driven by the partial GR decompositions (9). Applying Lemma 4.1 to (8) and (9), we find that also

$$\hat{A}_i \hat{S}_i = \hat{R}_i A \quad \text{and} \quad \hat{B}_i \hat{S}_i = \hat{R}_i B. \tag{10}$$

We use these two equations in Lemmas 4.8 and 4.5, respectively, to help determine the shape of $\hat{A}_i$ and $\hat{B}_i$.

Similarly, we have

$$\hat{A}_{i-1/2} = \hat{G}_i^{-1} A \hat{Z}_{i-1} \quad \text{and} \quad \hat{B}_{i-1/2} = \hat{G}_i^{-1} B \hat{Z}_{i-1}, \tag{11}$$

so the pencil $\hat{A}_{i-1/2} - \lambda\hat{B}_{i-1/2}$ can be viewed as the result of a partial GZ iteration driven by the partial GR decompositions

$$p(AB^{-1}) = \hat{G}_i\hat{R}_i \quad \text{and} \quad p(B^{-1}A) = \hat{Z}_{i-1}\hat{S}_{i-1}. \tag{12}$$

Applying Lemma 4.1 to (11) and (12), we obtain

$$\hat{A}_{i-1/2}\hat{S}_{i-1} = \hat{R}_iA \quad \text{and} \quad \hat{B}_{i-1/2}\hat{S}_{i-1} = \hat{R}_iB. \tag{13}$$

These equations are used to help to determine the shape of $\hat{A}_{i-1/2}$ and $\hat{B}_{i-1/2}$.

We study first the shape of the "B" matrices.

LEMMA 4.4. *For $i = 1, \ldots, n-m-1$, the last $n-m-i$ rows of $\hat{B}_{i-1/2}$ and $\hat{B}_i$ are in upper triangular form. That is,*

$$\hat{b}_{jk}^{(i-1/2)} = \hat{b}_{jk}^{(i)} = 0 \quad \text{if } j > k \quad \text{and} \quad j > i+m.$$

Pictorially, $\hat{B}_{i-1/2}$ and $\hat{B}_i$ have the form

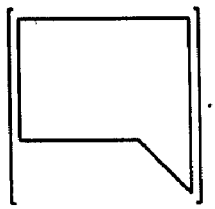

Proof. Writing the transformation $\hat{B}_i = \hat{G}_i{}^{-1}B\hat{Z}_i$ in partitioned form, we have

$$\begin{bmatrix} \hat{B}_{11}^{(i)} & \hat{B}_{12}^{(i)} \\ \hat{B}_{21}^{(i)} & \hat{B}_{22}^{(i)} \end{bmatrix} = \begin{bmatrix} X & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix} \begin{bmatrix} Y & 0 \\ 0 & I \end{bmatrix},$$

where the $(1,1)$ blocks all have dimension $(m+i) \times (m+i)$. Clearly $\hat{B}_{21}^{(i)} = 0$, and $\hat{B}_{22}^{(i)} = B_{22}$, which is upper triangular. Thus $\hat{B}_i$ has the stated form. To prove that $\hat{B}_{i-1/2}$ also has this form, apply the same partition to the equation $\hat{B}_{i-1/2} = \hat{G}_i^{-1}B\hat{Z}_{i-1}$. □

Lemma 4.4 suggests that $\hat{B}_{n-m-1/2}$ and $\hat{B}_{n-m}$ should be completely filled in. Fortunately the transformations do not only destroy zeros, they create zeros as well, as we see in Lemma 4.5. For the purpose of avoiding distracting complications in the statement of this lemma, we define $\hat{B}_{\rho+1/2} = \hat{B}$.

LEMMA 4.5. *For $i = 1, \ldots, \rho$, the first i columns of $\hat{B}_i$ and $\hat{B}_{i+1/2}$ are in upper triangular form. That is,*

$$\hat{b}_{jk}^{(i)} = \hat{b}_{jk}^{(i+1/2)} = 0 \quad \text{if } j > k \quad \text{and} \quad k \le i.$$

Pictorially, both $\hat{B}_i$ and $\hat{B}_{i+1/2}$ have the form

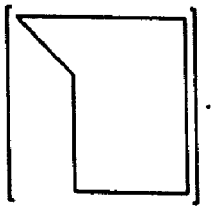

Proof. By (10) $\hat{B}_i\hat{S}_i = \hat{R}_iB$. We write this equation in partitioned form as

$$\begin{bmatrix} \hat{B}_{11}^{(i)} & \hat{B}_{12}^{(i)} \\ \hat{B}_{21}^{(i)} & \hat{B}_{22}^{(i)} \end{bmatrix} \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix},$$

where the $(1,1)$ blocks are all $i \times i$. The matrices S_{11}, R_{11}, and B_{11} are upper triangular. We know that the first $n-\nu$ columns of $\hat{S}_i$ are linearly independent, and since $i \le \rho \le n-\nu$, S_{11} must be nonsingular. Therefore

$$\begin{bmatrix} \hat{B}_{11}^{(i)} \\ \hat{B}_{21}^{(i)} \end{bmatrix} = \begin{bmatrix} R_{11} \\ 0 \end{bmatrix} B_{11}S_{11}^{-1}.$$

Consequently, $\hat{B}_{11}^{(i)} = R_{11}B_{11}S_{11}^{-1}$ is upper triangular and $\hat{B}_{21}^{(i)} = 0$. This proves that $\hat{B}_i$ has the stated form. To obtain the same result for $\hat{B}_{i+1/2}$, partition the equation $\hat{B}_{i+1/2}\hat{S}_i = \hat{R}_{i+1}B$ (from (13)) exactly as above. □

Remark. In the nonsingular case, or even in the case $\nu = 1$, Lemma 4.5 shows that $\hat{B} = \hat{B}_{n-1}$ is upper triangular. If $\nu \ge 2$, we can conclude that $\hat{B} = \hat{B}_\rho$ has only its first $\rho = n-\nu$ columns upper triangular; it is not guaranteed that the final ν columns get reduced to upper triangular form. But we already know from Theorem 4.3 that this portion of the matrix will be removed by deflation at the end of the iteration.

Combining Lemmas 4.4 and 4.5, we have the following result.

THEOREM 4.6. *For $i = 1, \ldots, \rho$,*

$$\hat{b}_{jk}^{(i-1/2)} = 0 \quad \textit{if } j > k \textit{ and either } j > i+m \textit{ or } k \le i-1,$$

$$\hat{b}_{jk}^{(i)} = 0 \quad \textit{if } j > k \textit{ and either } j > i+m \textit{ or } k \le i.$$

Thus we see that when $i < n-m$, $\hat{B}_{i-1/2}$ is upper triangular in its first $i-1$ columns and its last $n-i-m$ rows. Its nonzero pattern is

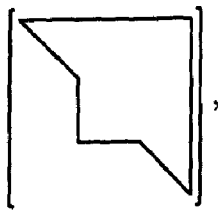

which would be upper triangular, except that it has a bulge. The tip of the bulge is at the $(i+m, i)$ position. We call this an *m-bulge* or a *bulge of order m* because it protrudes m diagonals below the upper triangular part of the matrix. The form of $\hat{B}_i$ is similar, but the tip of its bulge is at the $(i+m, i+1)$ position. This is a bulge of order $m-1$. We see thus that the transformation $\hat{B}_{i-1/2} \to \hat{B}_i$ shrinks the bulge by deleting one column from the left side. On the other hand, the transformation $\hat{B}_i \to \hat{B}_{i+1/2}$ enlarges the bulge by adding one row to the bottom. Thus the bulge is chased downward and to the right as the GZ iteration proceeds. When $i = n-m$, the bulge has reached the bottom of the matrix and begins to be pushed off the edge. If $\nu \le 1$, the bulge is eventually eliminated completely. If $\nu > 1$, the iteration ends with the last ν columns uncleared.

We now turn our attention to the "A" matrices.

LEMMA 4.7. *For $i = 0, \ldots, n-m-2$, the last $n-i-m-1$ rows of $\hat{A}_i$ and $\hat{A}_{i+1/2}$ have upper Hessenberg form. That is,*

$$\hat{a}_{jk}^{(i)} = \hat{a}_{jk}^{(i+1/2)} = 0 \quad \text{if } j > k+1 \quad \text{and} \quad j > i+m+1.$$

Proof. Write the transformation $\hat{A}_{i+1/2} = \hat{G}_{i+1}^{-1} A \hat{Z}_i$ in the unsymmetric partitioned form

$$\begin{bmatrix} \hat{A}_{11}^{(i+1/2)} & \hat{A}_{12}^{(i+1/2)} \\ \hat{A}_{21}^{(i+1/2)} & \hat{A}_{22}^{(i+1/2)} \end{bmatrix} = \begin{bmatrix} X & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix} \begin{bmatrix} Y & 0 \\ 0 & I \end{bmatrix},$$

where $\hat{A}_{11}^{(i+1/2)}$, $A_{11} \in \mathbb{C}^{(i+m+1)\times(i+m)}$, $X \in \mathbb{C}^{(i+m+1)\times(i+m+1)}$, $Y \in \mathbb{C}^{(i+m)\times(i+m)}$. Clearly $\hat{A}_{21}^{(i+1/2)} = 0$, and $\hat{A}_{22}^{(i+1/2)} = A_{22}$, which is the bottom right-hand corner of an upper Hessenberg matrix. Thus $\hat{A}_{i+1/2}$ has the stated form. To prove that $\hat{A}_i$ also has this form, apply the same partition to the equation $\hat{A}_i = \hat{G}_i^{-1} A \hat{Z}_i$. □

Thus it appears that $\hat{A}_{n-m-1}$ is completely filled in. But again it turns out that the transformations are not only destroying zeros, they are creating zeros as well.

LEMMA 4.8. *For $2 \le i \le \rho$, the first $i-1$ columns of $\hat{A}_{i-1/2}$ and $\hat{A}_i$ are in upper Hessenberg form. That is,*

$$\hat{a}_{jk}^{(i-1/2)} = \hat{a}_{jk}^{(i)} = 0 \quad \text{if } j > k+1 \quad \text{and} \quad k \le i-1.$$

Proof. By (13) we know that $\hat{A}_{i-1/2}\hat{S}_{i-1} = \hat{R}_i A$. Consider the unsymmetric partition

$$\begin{bmatrix} \hat{A}_{11}^{(i-1/2)} & \hat{A}_{12}^{(i-1/2)} \\ \hat{A}_{21}^{(i-1/2)} & \hat{A}_{22}^{(i-1/2)} \end{bmatrix} \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix},$$

where $\hat{A}_{11}^{(i-1/2)}$, $A_{11} \in \mathbb{C}^{i\times(i-1)}$, $S_{11} \in \mathbb{C}^{(i-1)\times(i-1)}$, and $R_{11} \in \mathbb{C}^{i\times i}$. Both S_{11} and R_{11} are upper triangular and nonsingular. Thus

$$\begin{bmatrix} \hat{A}_{11}^{(i-1/2)} \\ \hat{A}_{21}^{(i-1/2)} \end{bmatrix} = \begin{bmatrix} R_{11} \\ 0 \end{bmatrix} A_{11} S_{11}^{-1}.$$

Therefore $\hat{A}_{21}^{(i-1/2)} = 0$ and $\hat{A}_{11}^{(i-1/2)} = R_{11} A_{11} S_{11}^{-1}$. Since R_{11} and S_{11}^{-1} are upper triangular, and A_{11} satisfies $a_{jk} = 0$ if $j > k+1$, $\hat{A}_{11}^{(i-1/2)}$ must also have this zero pattern. This proves that $\hat{A}_{i-1/2}$ has the stated form. We obtain the same result for $\hat{A}_i$ by partitioning the equation $\hat{A}_i \hat{S}_i = \hat{R}_i A$ in exactly the same way. □

Remark. As long as $\nu \le 1$, Lemma 4.8 shows that both $\hat{A}_{\rho-1/2}$ and $\hat{A}_\rho$ are upper Hessenberg. If $\nu > 1$, Lemma 4.8 states that the first $\rho - 1$ columns of $\hat{A} = \hat{A}_\rho$ are upper Hessenberg. In fact the situation is better than that. From Theorem 4.3 we know that $\hat{A}$ is block triangular; all of the entries in column ρ below the main diagonal are automatically zero. Of course this is a theoretical result that is valid only in the absence of roundoff error.

Combining the results of Lemmas 4.7 and 4.8, we have the following theorem.

THEOREM 4.9. *For* $1 \le i \le \rho$,

$$\hat{a}_{jk}^{(i-1/2)} = 0 \quad \text{if } j > k+1 \quad \text{and either } j > i+m \text{ or } k \le i-1,$$

$$\hat{a}_{jk}^{(i)} = 0 \quad \text{if } j > k+1 \text{ and either } j > i+m+1 \text{ or } k \le i-1.$$

For $i = 1, \ldots, n-m-1$, $\hat{A}_{i-1/2}$ has its first $i-1$ columns and its last $n-m-i$ rows in upper Hessenberg form. Thus it has the form

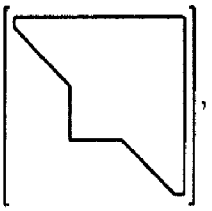

It has a bulge that has its tip at position $(i+m, i)$. We call this a bulge of *order* $m-1$ because it protrudes $m-1$ diagonals below the Hessenberg part of the matrix. $\hat{A}_i$ has its first $i-1$ columns and last $n-m-i-1$ rows upper Hessenberg. That is, it has a bulge whose tip is at $(i+m+1, i)$. This is an m-bulge.

The transformation $\hat{A}_{i-1/2} \to \hat{A}_i$ enlarges the bulge by adding one row to the bottom. On the other hand, the transformation $\hat{A}_i \to \hat{A}_{i+1/2}$ shrinks the bulge by removing one column from the left side. Thus the bulge in $\hat{A}_i$ moves downward and to the right, just as it does in $\hat{B}_i$. When $i = n-m-1$, the bulge has reached the bottom of the matrix and begins to be pushed off the edge.

Now let us consider the effects of the transformations on the A and B matrices together. The transformation $\hat{C}_{i-1/2} \to \hat{C}_i$ enlarges the "A"-bulge while shrinking the "B"-bulge. On the other hand, the transformation $\hat{C}_i \to \hat{C}_{i+1/2}$ shrinks the "A"-bulge while expanding the "B"-bulge. The relative positions of the bulges in $\hat{A}_i$ and $\hat{B}_i$ can be seen by superimposing them on one array.

$$\begin{bmatrix}
b & b & b & b & b & b & b & b \\
a & b & & & & & & b \\
 & a & b & & & & & b \\
 & a & b & & & & & b \\
 & a & b & b & b & & & b \\
 & a & a & a & a & b & & b \\
 & & & & & a & b & b \\
 & & & & & & a & b
\end{bmatrix}.$$

Here we have pictured the case $n = 8$, $m = 3$, $i = 2$. The nonzero part of $\hat{B}_2$ is the area outlined by the letter b. The nonzero part of $\hat{A}_2$ includes the nonzero part of $\hat{B}_2$ and in addition the entries marked with the letter a. This is an m-bulge for $\hat{A}_2$ and an $(m-1)$-bulge for $\hat{B}_2$. After the transformation $\hat{C}_2 \to \hat{C}_{5/2}$, the situation is

as follows.

$$\begin{bmatrix} b & b & b & b & b & b & b & b \\ a & b & & & & & & b \\ & a & b & & & & & b \\ & 0 & b & & & & & b \\ & 0 & b & b & b & & & b \\ & 0 & b & b & b & b & & b \\ & & & & & a & b & b \\ & & & & & & a & b \end{bmatrix}.$$

The "A"-bulge has been shrunk by the elimination of one column from the left (marked by zeros), and the "B"-bulge has been enlarged by the addition of one row at the bottom. Now the bulges coincide perfectly. This is an $m-1$ bulge in $\hat{A}_{5/2}$ and an m bulge in $\hat{B}_{5/2}$.

4.4. The implicit GZ algorithm. Now we are ready to use the information amassed in the previous section to see how to carry out an iteration of the GZ algorithm without actually forming the matrices $\hat{R}_0 = p(AB^{-1})$, $\hat{S}_0 = p(B^{-1}A)$, or any of the derived matrices $\hat{R}_i$, $\hat{S}_i$.

The first step is to find G_1, which is the elimination matrix for the first column of $\hat{R}_0 = p(AB^{-1})$. Thus we need to find $x = p(AB^{-1})e_1$. This can be computed relatively inexpensively if $m \ll n$. Indeed, p is given in the factored form $p(\lambda) = (\lambda - \sigma_1)\cdots(\lambda - \sigma_m)$. Thus x can be calculated by the recursion

$$x^{(j)} = (AB^{-1} - \sigma_j I)x^{(j-1)}, \qquad j = 1, \ldots, m, \tag{14}$$

with $x^{(0)} = e_1$. Then $x = x^{(m)}$. Since it does not matter whether we get x or a multiple of x, in practice we would also rescale at each step to avoid over/underflow. This is inexpensive. Since AB^{-1} is upper Hessenberg, only the first j components of $x^{(j-1)}$ are nonzero. Now consider the jth step. If we let $y^{(j)} = AB^{-1}x^{(j-1)}$, then $x^{(j)} = y^{(j)} - \sigma_j x^{(j-1)}$. We can find $y^{(j)}$ by solving $Bz^{(j)} = x^{(j-1)}$ for $z^{(j)}$ and then calculating $y^{(j)} = Az^{(j)}$. Because only the first j components of $x^{(j-1)}$ are nonzero and B is upper triangular, the system $Bz^{(j)} = x^{(j-1)}$ is in fact only a $j \times j$ upper triangular system, whose solution requires only $O(j^2)$ operations. Only the first j entries of $z^{(j)}$ are nonzero. Thus the product $Az^{(j)}$ involves only the first j columns of A. The nonzero entries in these columns are confined to the first $j+1$ rows, so $y^{(j)}$ can be obtained in $O(j^2)$ operations. The computation $x^{(j)} = y^{(j)} - \sigma_j x^{(j-1)}$ requires only $O(j)$ operations. Thus the total operation count for the jth step is $O(j^2)$. This must be done for $j = 1, \ldots, m$, so the total cost of computing x is $O(m^3)$, which is small if $m \ll n$.[3]

Once we have x, we can let $G_1 = \mathrm{diag}\{\tilde{G}_1, I\}$ be the elimination matrix for x given by the chosen elimination rule and calculate

$$\hat{A}_{1/2} = G_1^{-1}A \quad \text{and} \quad \hat{B}_{1/2} = G_1^{-1}B.$$

Logically the next question would be how to find Z_1. However, we postpone that and ask instead how one finds G_i, $i = 2, 3, \ldots, \rho$, to make the transformations

$$\hat{A}_{i-1/2} = G_i^{-1}\hat{A}_{i-1} \quad \text{and} \quad \hat{B}_{i-1/2} = G_i^{-1}\hat{B}_{i-1},$$

[3] This procedure can be modified in various ways. For example, a more accurate formula for the case $m = 2$ is given in [11]. Also, although we are assuming throughout this paper that B is nonsingular, one might reasonably ask whether this process can be salvaged in case one of $b_{11}, \ldots, b_{mm}$ happens to be zero. Deflation strategies for singular B are discussed in [11], [9], and [6, p. 400].

(see (7)) given that $\hat{A}_{i-1}$ and $\hat{B}_{i-1}$ are available. To deal with the two cases $i \leq n-m$ and $i > n-m$ simultaneously, let $k = \min\{i+m, n\}$. Then G_i has the form $\operatorname{diag}\{I_{i-1}, \tilde{G}_i, I_{n-k}\}$, where $\tilde{G}_i \in \mathbb{C}^{(k-i+1)\times(k-i+1)}$ is the elimination matrix for $[\hat{r}_{i,i}^{(i-1)}, \ldots, \hat{r}_{k,i}^{(i-1)}]^T$. Call this (column) vector y. We need to find y or a multiple of y and we need to do it without knowledge of $\hat{R}_{i-1}$. We know that $\tilde{G}^{-1}y = \alpha e_1$ for some scalar α. The fact that the first ρ columns of $\hat{R}_i$ are linearly independent guarantees that $\alpha \neq 0$.

We know from Theorem 4.9 that the operation

$$G_i^{-1}\hat{A}_{i-1} = \hat{A}_{i-1/2}$$

shrinks the bulge in $\hat{A}_{i-1}$ by removing column $i-1$ from the bulge. All this means is that the entries in positions $(i+1, i-1), \ldots, (k, i-1)$ get set to zero. Let us focus on this column. In transforming $\hat{A}_{i-1}$ to $\hat{A}_{i-1/2}$, the submatrix $\tilde{G}_i^{-1}$ acts only on rows i through k. As far as column $i-1$ is concerned, the action is

$$\tilde{G}_i^{-1}\begin{bmatrix} \hat{a}_{i,i-1}^{(i-1)} \\ \hat{a}_{i+1,i-1}^{(i-1)} \\ \vdots \\ \hat{a}_{k,i-1}^{(i-1)} \end{bmatrix} = \begin{bmatrix} \hat{a}_{i,i-1}^{(i-1/2)} \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

In other words, defining $x \in \mathbb{C}^{k-i+1}$ by $x = [\hat{a}_{i,i-1}^{(i-1)}, \hat{a}_{i+1,i-1}^{(i-1)}, \ldots, \hat{a}_{k,i-1}^{(i-1)}]^T$, we have $\tilde{G}_i^{-1}x = \beta e_1$, where $\beta = \hat{a}_{i,i-1}^{(i-1/2)}$. One easily checks that $\beta \neq 0$. Indeed, subsequent transformations do not alter the entry in position $(i, i-1)$. Thus $\beta = \hat{a}_{i,i-1}$, and this is nonzero by Theorem 4.3. The relationship $\tilde{G}_i^{-1}x = \beta e_1$ says that $\tilde{G}_i$ is an elimination matrix for x. Since it is also an elimination matrix for y, x must be a multiple of y. Indeed, $x = \beta\tilde{G}_i e_1 = \beta\alpha^{-1}y$.

This relationship can also be inferred directly from the equation $\hat{A}_{i-1}\hat{S}_{i-1} = \hat{R}_{i-1}A$ by comparing column $i-1$ of the left-hand side product with the same column of the right-hand side, which is analogous to the approach used in [10]. Focusing on rows i through k, and taking into account the zero structure of the various matrices, we find that

$$\begin{bmatrix} \hat{a}_{i,i-1}^{(i-1)} \\ \vdots \\ \hat{a}_{k,i-1}^{(i-1)} \end{bmatrix} \hat{s}_{i-1,i-1}^{(i-1)} = \begin{bmatrix} \hat{r}_{i,i}^{(i-1)} \\ \vdots \\ \hat{r}_{k,i}^{(i-1)} \end{bmatrix} a_{i,i-1}. \tag{15}$$

We know that $a_{i,i-1} \neq 0$ for all i, and $\hat{s}_{i-1,i-1}^{(i-1)} \neq 0$ as long as $i \leq \rho + 1$. Thus x is a multiple of y.

This solves the problem of how to find G_i. We see that we do not need $\hat{R}_{i-1}$, as the required information for building G_i is also present in $\hat{A}_{i-1}$. We summarize our findings as a theorem.

THEOREM 4.10. *Let $2 \leq i \leq \rho$, and let $k = \min\{i+m, n\}$. Then $G_i = \operatorname{diag}\{I_{i-1}, \tilde{G}_i, I_{n-k}\}$, where $\tilde{G}_i$ is the elimination matrix of*

$$[\hat{a}_{i,i-1}^{(i-1)}, \hat{a}_{i+1,i-1}^{(i-1)}, \ldots, \hat{a}_{k,i-1}^{(i-1)}]^T.$$

We now turn to the question of how to compute the transformations Z_i. Given $1 \leq i \leq \rho$, let $k = \max\{i+m, n\}$. We know that $Z_i = \operatorname{diag}\{I_{i-1}, \tilde{Z}_i, I_{n-k}\}$, where $\tilde{Z}_i$

is the elimination matrix for the vector

$$x = [\hat{s}_{i,i}^{(i-1)}, \hat{s}_{i+1,i}^{(i-1)}, \ldots, \hat{s}_{k,i}^{(i-1)}]^T,$$

which is part of the ith column of $\hat{S}_i$. We wish to determine this vector (or a multiple thereof) without computing $\hat{S}_i$. We have $\tilde{Z}_i^{-1}x = \alpha e_1$, where $\alpha \neq 0$.

We know from Theorem 4.6 that the transformation $\hat{B}_{i-1/2}Z_i = \hat{B}_i$ (from (7)) shrinks the bulge in $\hat{B}_{i-1/2}$ by setting the entries in positions $(i+1,i),\ldots,(k,i)$ to zero. To see how Z_i^{-1} acts, we consider the inverse equation $Z_i^{-1}\hat{B}_{i-1/2}^{-1} = \hat{B}_i^{-1}$. We know that $\hat{B}_{i-1/2}$ has the form

$$\hat{B}_{i-1/2} = \begin{bmatrix} B_{11} & B_{12} & B_{13} \\ 0 & B_{22} & B_{23} \\ 0 & 0 & B_{33} \end{bmatrix},$$

where $B_{11} \in \mathbb{C}^{(i-1)\times(i-1)}$, $B_{22} \in \mathbb{C}^{(k-i+1)\times(k-i+1)}$, and $B_{33} \in \mathbb{C}^{(n-k)\times(n-k)}$. B_{11} and B_{33} are upper triangular, but B_{22} is full, as it holds the bulge. Clearly $\hat{B}_{i-1/2}^{-1}$ has exactly the same zero structure; it is upper triangular with a bulge of order $k-i$. $\hat{B}_i$ has the same form as $\hat{B}_{i-1/2}$, except that the first column of B_{22} has been set to upper triangular form. It is upper triangular with a bulge of order $k-i-1$. Its inverse has exactly the same structure. Thus the transformation $Z_i^{-1}\hat{B}_{i-1/2}^{-1} = \hat{B}_i^{-1}$ shrinks the bulge in $\hat{B}_{i-1/2}^{-1}$ by setting the entries in positions $(i+1,i),\ldots,(k,i)$ to zero. The submatrix $\tilde{Z}_i^{-1}$ acts on rows i through k. Within these rows, our interest focuses on column i, the first column of B_{22}^{-1}, as this is where the elimination takes place. Call this column y; that is, $y = B_{22}^{-1}e_1$. Then $\tilde{Z}_i^{-1}y = \beta e_1$, where $\beta \neq 0$ because $\hat{B}_i^{-1}$ is nonsingular. Thus $\tilde{Z}_i$ is an elimination matrix of y. Since it is also an elimination matrix of x, y must be a multiple of x. Indeed $y = \beta\tilde{Z}_i e_1 = \beta\alpha^{-1}x$. Since $y = B_{22}^{-1}e_1$, we can obtain it by solving the small system $B_{22}y = e_1$. We have now proved the following theorem.

THEOREM 4.11. *For $1 \leq i \leq \rho$ let $k = \min\{i+m, n\}$. Let B_{22} denote the principal submatrix of $\hat{B}_{i-1/2}$ consisting of rows and columns i through k. Let y be the unique solution of $B_{22}y = e_1$. Then $Z_i = \operatorname{diag}\{I_{i-1}, \tilde{Z}_i, I_{n-k}\}$, where $\tilde{Z}_i$ is the elimination matrix of y.*

Theorem 4.11 can also be inferred from the equation

$$\hat{B}_{i-1/2}\hat{S}_{i-1} = \hat{R}_i B, \tag{16}$$

which holds by (13). The vector x lies in the ith column of $\hat{S}_{i-1}$, so consider the ith column of (16), partitioned as

$$\begin{bmatrix} B_{11} & B_{12} & B_{13} \\ 0 & B_{22} & B_{23} \\ 0 & 0 & B_{33} \end{bmatrix} \begin{bmatrix} z \\ x \\ 0 \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \\ 0 & R_{32} \end{bmatrix} \begin{bmatrix} b \\ 0 \end{bmatrix},$$

where $\hat{B}_{i-1/2}$ is partitioned as before, $z \in \mathbb{C}^{i-1}$, $x \in \mathbb{C}^{k-i+1}$, $R_{11} \in \mathbb{C}^{(i-1)\times i}$, $R_{21} \in \mathbb{C}^{(k-i+1)\times i}$, and $b \in \mathbb{C}^i$. Since the first i columns of $\hat{R}_i$ are upper triangular, R_{21} has only one nonzero entry, $\hat{r}_{ii}^{(i)}$, which lies in the upper right-hand corner. Thus $R_{21} = \gamma e_1 e_i^T$, where $\gamma = \hat{r}_{ii}^{(i)} \neq 0$. We seek x. But clearly $B_{22}x = R_{21}b = \gamma e_1 e_i^T b = \delta e_1$, where $\delta = \gamma e_i^T b = \gamma b_{ii} \neq 0$. Thus we can find $y = \delta^{-1}x$ by solving $B_{22}y = e_1$.

Theorems 4.10 and 4.11 justify the implicit GZ algorithm, which is summarized in (17). Notice that the algorithm takes $n-1$ steps, rather than stopping after ρ steps. In practice ρ is usually unknown because ν is unknown. Even if ν is known in principle, it is not well determined in the presence of roundoff error. When it happens that $\nu > 1$, the implicit algorithm differs from the explicit algorithm only in one way. In this case a small subpencil can be deflated from the bottom of the matrix. The implicit algorithm operates on this subpencil, reducing it to Hessenberg-triangular form, whereas the explicit algorithm does not. This further reduction is useful because we need to calculate the eigenvalues of the subpencil anyway.

IMPLICIT GZ ALGORITHM

(17)

```
for i = 1, ..., n - 1
    k ← min{i + m, n}
    if (i = 1) then
        x ← p(AB^{-1})e_1        (See the discussion following (14).)
        y ← first k entries of x
    else
        y ← [a_{i,i-1}, ..., a_{k,i-1}]^T
    end if
    G̃ ← elimination matrix of y                (*)
    G ← diag{I_{i-1}, G̃, I_{n-k}}
    A ← G^{-1}A,   B ← G^{-1}B
    B_22 ← rows, columns i through k of B
    z ← B_22^{-1} e_1
    Z̃ ← elimination matrix of z                (**)
    Z ← diag{I_{i-1}, Z̃, I_{n-k}}
    A ← AZ,   B ← BZ
```

Remarks. For standard eigenvalue problems ($B = I$), we have already noted that $Z_i = G_i$ for $i = 1, \ldots, n-1$. Thus we have $Z = G$ at each step of (17). This special case is the implicit GR algorithm.

In [2] Bunse-Gerstner and Elsner developed a new version of the QZ algorithm for unitary pencils, which are of interest in certain signal-processing applications. Instead of using the Hessenberg-triangular form, they introduced a more condensed block-diagonal form. Since our development is built upon the Hessenberg-triangular form, it does not encompass the algorithm of [2] as a special case. However, the methodology used here can be adapted to that situation and used to derive that algorithm and its generalizations to higher multiplicity. In particular, if A and B are unitary, the equations in Lemma 4.1 are joined by the two related equations

$$\hat{A}^* \hat{R} = \hat{S} A^* \quad \text{and} \quad \hat{B}^* \hat{R} = \hat{S} B^*.$$

Whereas the equations of Lemma 4.1 yield information about zeros below the main diagonal (e.g., upper Hessenberg form is preserved from one iteration to the next), these new equations give information about zeros above the main diagonal. In particular, block diagonal forms are preserved from one iteration to the next. Of course, we must now insist that the G and Z transformations be unitary, so that the pencil stays unitary from one iteration to the next.

4.5. The generic bulge-chasing algorithm. If one compares the implicit GZ algorithm (17) with the standard double-step QZ algorithm, as presented in, for

example, EISPACK [5], one notices an important difference. In (17) the transformation Z_i is designed to clear out one column of the bulge in B. In contrast, the corresponding transformation in the standard code eliminates the entire B bulge before proceeding to the computation of G_{i+1}. The bulge is annihilated row by row, using one elimination matrix for each row. In the case $m = 2$ (as in the EISPACK code) the added cost of doing this is small. However, as m is made larger, the cost difference becomes significant. On the other hand, the standard procedure may be substantially more stable. The reason for this is that the Z_i calculated in (17) is designed so that the transformation $Z_i^{-1}\hat{B}_{i-1/2}^{-1} = \hat{B}_i^{-1}$ will clear out one column in the bulge in the inverse matrix. As a consequence, the transformation $\hat{B}_{i-1/2}Z_i = \hat{B}_i$ will also, in principle, remove one column from the bulge in $\hat{B}_{i-1/2}$. However, these zeros are introduced only incidentally; they are not enforced by the transformation. Therefore, in the presence of roundoff errors these numbers will not be exactly zero. They may sometimes be far enough from zero that they cannot be set to zero without compromising the stability of the algorithm. After all, if the submatrix B_{22} should be ill conditioned at some point, then the solution of $B_{22}z = e_1$ and the resulting Z_i may not be well determined. In contrast, the standard procedure introduces the desired zeros explicitly through the mechanism of eliminating the entire bulge from B. Therefore there is no question of having to set to zero some numbers that should be, but are not quite, zero. We note finally that in the extreme case of singular B, the implicit GZ algorithm (17) breaks down, whereas the standard procedure does not.

It is therefore desirable to broaden our class of algorithms to include procedures of this type. To this end we introduce a *generic bulge-chasing algorithm*, which is exactly algorithm (17), except that at the two steps labelled $(*)$ and $(**)$, where $\tilde{G}$ and $\tilde{Z}$ are chosen, we do not require that they be determined by a specific elimination rule. At $(*)$ we allow any nonsingular $\tilde{G}$ for which $\tilde{G}^{-1}y = \alpha e_1$ for some α, and likewise for $\tilde{Z}$ at $(**)$. That $\tilde{G}^{-1}y = \alpha e_1$ means exactly that (at the ith step) the premultiplication by G^{-1} causes entries $(i+1, i-1), \ldots, (i+k, i-1)$ of the A matrix to be transformed to zero. Similarly, that $\tilde{Z}^{-1}z = \beta e_1$, where $z = B_{22}^{-1}e_1$, means neither more nor less than that (at the ith step) the postmultiplication by Z causes entries $(i+1, i), \ldots, (i+k, i)$ of the B matrix to be transformed to zero. We do not require that the vector z be calculated explicitly, only that the transformations produce the desired zeros. (Indeed this can be done even if B_{22} is singular, in which case z is not well defined.) Thus algorithms that annihilate the entire bulge in B at each step fit into this structure.

It is worth mentioning a class of algorithms that avoids the high cost of eliminating the entire bulge in B at each step by never allowing that bulge to build up in the first place [9, p. 75]. Suppose that after $i-1$ steps of a generic bulge-chasing algorithm we have $\hat{A}_{i-1} - \lambda\hat{B}_{i-1}$, where $\hat{B}_{i-1}$ has no bulge. The next step is to build an elimination matrix G_i such that G_i^{-1} annihilates the entries in positions $(i+1, i-1), \ldots, (i+k, i-1)$ of $\hat{A}_{i-1}$. One can build such a G_i^{-1} as a product of simpler matrices $G_{i,i+1}^{-1} \cdots G_{i,i+k}^{-1}$, where each $G_{i,j}^{-1}$ acts only on rows $j-1$ and j (that is, its nontrivial part is a 2×2 matrix) and annihilates the $(j, i-1)$ entry. One easily checks that the matrix $\hat{B}_{i-1/2} = G_i^{-1}\hat{B}_{i-1}$ fails to be upper triangular only in that the entries in positions $(i+1, i), (i+2, i+1), \ldots, (i+k, i+k-1)$ are nonzero. These can be annihilated by a transformation $\hat{B}_i = \hat{B}_{i-1/2}Z_i$, where $Z_i = Z_{i,i+k} \cdots Z_{i,i+1}$, and each $Z_{i,j}$ acts only on columns $j-1$ and j and annihilates the entry in position $(j, j-1)$. Thus $\hat{B}_i$ is upper triangular. Algorithms of this type conform to the framework of the generic bulge-chasing algorithm.

Our analysis of the generic bulge-chasing algorithm will show that each iteration amounts to an iteration of a generic GZ algorithm. Thus the generic bulge-chasing algorithm lies within the class of algorithms whose convergence properties we studied in §3. To this end we introduce some notation. In fact the notation is identical to notation that we used earlier, but the symbols now carry slightly different meanings. Let G_i and Z_i denote the transformations produced at the ith step of the generic chasing algorithm (as we have already done in the previous paragraph), let $\hat{G}_i = G_1 \cdots G_i$, $\hat{Z}_i = Z_1 \cdots Z_i$, $\hat{A}_i = \hat{G}_i^{-1} A \hat{Z}_i$, $\hat{B}_i = \hat{G}_i^{-1} B \hat{Z}_i$, $\hat{A}_{i-1/2} = \hat{G}_i^{-1} A \hat{Z}_{i-1}$, and $\hat{B}_{i-1/2} = \hat{G}_i^{-1} B \hat{Z}_{i-1}$, $i = 1, \ldots, n-1$. These matrices may be different from those featured in Theorems 4.6 and 4.9, but they have the same bulge structure; it is exactly the function of the generic bulge-chasing algorithm to enforce this structure. Let $\hat{A} = \hat{A}_{n-1}$ and $\hat{B} = \hat{B}_{n-1}$. These are the final products of (one iteration of) the generic bulge-chasing algorithm. $\hat{A}$ is upper Hessenberg, and $\hat{B}$ is upper triangular. Assuming once again that B is nonsingular, let $\hat{R}_0 = p(AB^{-1})$ and $\hat{S}_0 = p(B^{-1}A)$, as before, and let $\hat{R}_i = G_i^{-1} \hat{R}_{i-1}$ and $\hat{S}_i = Z_i^{-1} \hat{S}_{i-1}$ for $i = 1, \ldots, n-1$. Now we are using the matrices G_i and Z_i to define $\hat{R}_i$ and $\hat{S}_i$, whereas in the development of the explicit GZ algorithm we used $\hat{R}_i$ and $\hat{S}_i$ to define G_i and Z_i. It is an immediate consequence of the new definitions that

$$p(AB^{-1}) = \hat{G}_i \hat{R}_i \quad \text{and} \quad p(B^{-1}A) = \hat{Z}_i \hat{S}_i$$

for $i = 1, \ldots, n-1$. The matrices $\hat{R}_i$ and $\hat{S}_i$ defined in connection with the explicit algorithm were partially upper triangular. Whether or not the new $\hat{R}_i$ and $\hat{S}_i$ have that property is not immediately clear from the definition. In fact they do, as the following theorem shows.

THEOREM 4.12. *For $i = 1, \ldots, n-1$, the matrices $\hat{R}_i$ and $\hat{S}_i$ defined in the previous paragraph both have the form*

$$\begin{bmatrix} X_{11} & X_{12} \\ 0 & X_{22} \end{bmatrix}, \tag{18}$$

where $X_{11} \in \mathbb{C}^{i \times i}$ is upper triangular.

Proof. The proof is by induction on i. First let $i = 1$. The transformation G_1 is designed to annihilate $p(AB^{-1})e_1$, the first column of $p(AB^{-1}) = \hat{R}_0$. Since $\hat{R}_1 = G_1^{-1} \hat{R}_0$, $\hat{R}_1$ must have its first column in upper triangular form, as claimed.

Now we show that, for $i = 1, \ldots, n-1$, if $\hat{R}_i$ has the stated form, then so does $\hat{S}_i$. Since $\hat{B}_i = \hat{G}_i^{-1} B \hat{Z}_i$, $p(AB^{-1}) = \hat{G}_i \hat{R}_i$ and $p(B^{-1}A) = \hat{Z}_i \hat{S}_i$, we can apply Lemma 4.1 to obtain $\hat{B}_i \hat{S}_i = \hat{R}_i B$ or, equivalently,

$$\hat{S}_i = \hat{B}_i^{-1} \hat{R}_i B.$$

Each of the three matrices on the right-hand side has the form (18), where $X_{11} \in \mathbb{C}^{i \times i}$, so $\hat{S}_i$ must also have this form.

We complete the induction by showing that for $i = 2, \ldots, n-1$, if $\hat{R}_{i-1}$ and $\hat{S}_{i-1}$ have the stated form, then so does $\hat{R}_i$. Certainly the first $i-1$ columns of $\hat{R}_i$ are in upper triangular form, for this is true of $\hat{R}_{i-1}$, and the transformation $\hat{R}_i = G_i^{-1} \hat{R}_{i-1}$ does not alter these columns, as one easily checks. Thus we can focus on the ith column of $\hat{R}_i$. Since $\hat{A}_{i-1/2} = \hat{G}_i^{-1} A \hat{Z}_{i-1}$, $p(AB^{-1}) = \hat{G}_i \hat{R}_i$ and $p(B^{-1}A) = \hat{Z}_{i-1} \hat{S}_{i-1}$, we can apply Lemma 4.1 to obtain

$$\hat{A}_{i-1/2} \hat{S}_{i-1} = \hat{R}_i A. \tag{19}$$

We wish to pick out the ith column of $\hat{R}_i$. Noting that $a_{i,i-1} \neq 0$, we examine column $i-1$ of (19), partitioned unsymmetrically as

$$\begin{bmatrix} \hat{A}_{11}^{(i-1/2)} & \hat{A}_{12}^{(i-1/2)} \\ 0 & \hat{A}_{22}^{(i-1/2)} \end{bmatrix} \begin{bmatrix} v \\ 0 \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} w \\ 0 \end{bmatrix}, \tag{20}$$

where $\hat{A}_{11}^{(i-1/2)} \in \mathbb{C}^{i \times i-1}$, $R_{11} \in \mathbb{C}^{i \times i}$, $v \in \mathbb{C}^{i-1}$, and $w \in \mathbb{C}^{i}$. The last entry in w is $a_{i,i-1}$. We already know that the first $i-1$ columns of $\hat{R}_i$ are upper triangular. This implies that R_{11} is upper triangular, and all but the last column of R_{12} is zero. Thus $R_{12} = xe_i^T$ for some x. If we can show that $x = 0$, we will be finished. Equating second components of (20), we have $0 = R_{21}w = xe_i^T w = xa_{i,i-1}$. Since $a_{i,i-1} \neq 0$, we have $x = 0$. □

Let $\hat{G} = \hat{G}_{n-1}$, $\hat{Z} = \hat{Z}_{n-1}$, $\hat{R} = \hat{R}_{n-1}$, and $\hat{S} = \hat{S}_{n-1}$. Then $\hat{R}$ and $\hat{S}$ are upper triangular by Theorem 4.12 with $i = n-1$, and

$$\hat{A} = \hat{G}^{-1} A \hat{Z} \quad \text{and} \quad \hat{B} = \hat{G}^{-1} B \hat{Z}, \tag{21}$$

where

$$p(AB^{-1}) = \hat{G}\hat{R} \quad \text{and} \quad p(B^{-1}A) = \hat{Z}\hat{S}. \tag{22}$$

We conclude that one iteration of the generic bulge-chasing algorithm amounts to one iteration of the generic GZ algorithm.

REFERENCES

[1] W. Bunse and A. Bunse-Gerstner, *Numerische lineare Algebra*, Teubner, Stuttgart, 1985.

[2] A. Bunse-Gerstner and L. Elsner, *Schur parameter pencils for the solution of the unitary eigenvalue problem*, Linear Algebra Appl., 154–156 (1991), pp. 741–778.

[3] J. Demmel and B. Kågström, *Computing stable eigendecompositions of matrix pencils*, Linear Algebra Appl. 88–89 (1987), pp. 139–186.

[4] J. Demmel and B. Kågström, *GUPTRI*, NETLIB, 1991.

[5] B. S. Garbow, J. M. Doyle, J. J. Dongarra, and C. B. Moler, *Matrix Eigensystem Routines–EISPACK Guide Extension*, Springer-Verlag, New York, 1977.

[6] G. Golub and C. Van Loan, *Matrix Computations*, Second Edition, The Johns Hopkins University Press, Baltimore, MD, 1989.

[7] J. B. Haag and D. S. Watkins, *QR-like algorithms for the nonsymmetric eigenvalue problem*, ACM Trans. Math. Software, 19 (1993), pp. 407–418.

[8] L. Kaufman, *The LZ algorithm to solve the generalized eigenvalue problem*, SIAM J. Numer. Anal., 11 (1974), pp. 997–1024.

[9] ———, *Some thoughts on the QZ algorithm for solving the generalized eigenvalue problem*, ACM Trans. Math. Software, 3 (1977), pp. 65–75.

[10] G. S. Miminis and C. C. Paige, *Implicit shifting in the QR algorithm and related algorithms*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 385–400.

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

[12] G. W. Stewart, *Introduction to Matrix Computations*, Academic Press, New York, 1973.

[13] P. Van Dooren, *The computation of Kronecker's canonical form of a singular pencil*, Linear Algebra Appl., 27 (1979), pp. 103–140.

[14] R. C. Ward, *The combination shift QZ algorithm*, SIAM J. Numer. Anal., 12 (1975), pp. 835–853.

[15] D. S. WATKINS, *Fundamentals of Matrix Computations*, John Wiley and Sons, New York, 1991.
[16] D. S. WATKINS AND L. ELSNER, *Self-equivalent flows associated with the generalized eigenvalue problem*, Linear Algebra Appl., 118 (1989), pp. 107–127.
[17] ———, *Convergence of algorithms of decomposition type for the eigenvalue problem*, Linear Algebra Appl., 143 (1991), pp. 19–47.
[18] ———, *Chasing algorithms for the eigenvalue problem*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 374–384.