arXiv:1611.02435v3 [math.NA] 19 Jul 2018

FAST AND BACKWARD STABLE COMPUTATION OF ROOTS OF POLYNOMIALS, PART II: BACKWARD ERROR ANALYSIS; COMPANION MATRIX AND COMPANION PENCIL*

JARED L. AURENTZ†, THOMAS MACH¶, LEONARDO ROBOL‖, RAF VANDEBRIL‡, AND DAVID S. WATKINS§

Abstract. This work is a continuation of *Fast and backward stable computation of roots of polynomials* by J.L. Aurentz, T. Mach, R. Vandebril, and D.S. Watkins, SIAM Journal on Matrix Analysis and Applications, 36(3): 942–973, 2015. In that paper we introduced a companion QR algorithm that finds the roots of a polynomial by computing the eigenvalues of the companion matrix in $O(n^2)$ time using $O(n)$ memory. We proved that the method is backward stable. Here we introduce, as an alternative, a companion QZ algorithm that solves a generalized eigenvalue problem for a companion pencil. More importantly, we provide an improved backward error analysis that takes advantage of the special structure of the problem. The improvement is also due, in part, to an improvement in the accuracy (in both theory and practice) of the turnover operation, which is the key component of our algorithms. We prove that for the companion QR algorithm, the backward error on the polynomial coefficients varies linearly with the norm of the polynomial's vector of coefficients. Thus the companion QR algorithm has a smaller backward error than the unstructured QR algorithm (used by MATLAB's `roots` command, for example), for which the backward error on the polynomial coefficients grows quadratically with the norm of the coefficient vector. The companion QZ algorithm has the same favorable backward error as companion QR, provided that the polynomial coefficients are properly scaled.

Key words. polynomial, root, companion matrix, companion pencil, eigenvalue, Francis algorithm, QR algorithm, QZ algorithm, core transformation, backward stability

AMS subject classifications. 65F15, 65H17, 15A18, 65H04

1. Introduction.

Consider the problem of computing the n zeros of a complex polynomial

$$p(z) = a_n z^n + a_{n-1} z^{n-1} + \cdots + a_1 z + a_0, \quad a_n \neq 0, \quad a_0 \neq 0,$$

expressed in terms of the monomial basis. One way to do this is to form the companion matrix

$$A = \begin{bmatrix} & & & -a_0/a_n \\ 1 & & & -a_1/a_n \\ & \ddots & & \vdots \\ & & 1 & -a_{n-1}/a_n \end{bmatrix}, \tag{1.1}$$

* The research was partially supported by the Research Council KU Leuven, project C14/16/056 (Inverse-free Rational Krylov Methods: Theory and Applications), by an INdAM/GNCS project, and by the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centers of Excellence in R&D (SEV-2015-0554).

†Instituto de Ciencias Matemáticas, Universidad Autónoma de Madrid, Madrid, Spain (JaredAurentz@gmail.com).

‡Department of Computer Science, University of Leuven, KU Leuven, Leuven, Belgium; (Raf.Vandebril@cs.kuleuven.be).

‖Istituto di Scienza e Tecnologie dell Informazione 'A. Faedo' (ISTI), CNR, Pisa, Italy; (Leonardo.Robol@isti.cnr.it).

§Department of Mathematics, Washington State University, Pullman, WA 99164-3113, USA; (watkins@math.wsu.edu).

¶Department of Mathematics, School of Science and Technology, Nazarbayev University, 010000 Astana, Kazakhstan; (thomas.mach@nu.edu.kz).

and compute its eigenvalues. This is what MATLAB's `roots` command does. But `roots` does not exploit the special structure of the companion matrix, so it requires $O(n^2)$ storage (one matrix) and $O(n^3)$ flops using Francis's implicitly-shifted QR algorithm [17]. It is natural to ask whether we can save space and flops by exploiting the structure. This has, in fact, been done by Bini et al. [7], Boito et al. [11], Chandrasekaran et al. [12], and Aurentz et al. [4]. All of the methods proposed in these papers use the unitary-plus-rank-one structure of the companion matrix to build a special data structure that brings the storage down to $O(n)$. Then Francis's algorithm, operating on the special data structure, has a flop count of $O(n^2)$. Based on the tests in [4], our method appears to be the fastest of the several that have been proposed. Moreover our algorithm is the only one that has been proved to be backward stable. In this paper we will refer to our method as the *companion QR algorithm*.

In cases where the polynomial has a particularly small leading coefficient, one might hesitate to use the companion matrix, since division by a tiny a_n will result in very large entries in (1.1). This might adversely affect accuracy. An alternative to division by a_n is to work with a *companion pencil*

$$A - \lambda B = \begin{bmatrix} 0 & & & & -a_0 \\ 1 & 0 & & & -a_1 \\ & \ddots & \ddots & & \vdots \\ & & 1 & 0 & -a_{n-2} \\ & & & 1 & -a_{n-1} \end{bmatrix} - \lambda \begin{bmatrix} 1 & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & & a_n \end{bmatrix}, \tag{1.2}$$

which also has the roots of p as its eigenvalues. More generally we can consider any pencil of the form

$$V - \lambda W = \begin{bmatrix} 0 & & & -v_1 \\ 1 & & & -v_2 \\ & 1 & & -v_3 \\ & & \ddots & \vdots \\ & & 1 & -v_n \end{bmatrix} - \lambda \begin{bmatrix} 1 & & & w_1 \\ & 1 & & w_2 \\ & & \ddots & \vdots \\ & & 1 & w_{n-1} \\ & & & w_n \end{bmatrix}, \tag{1.3}$$

where

$$\begin{aligned} v_1 &= a_0, \\ v_{i+1} + w_i &= a_i, \text{ for } i = 1, \ldots, n-1, \text{and} \\ w_n &= a_n. \end{aligned} \tag{1.4}$$

One easily checks that all pencils of this form also have the zeros of p as their eigenvalues [5, 22]. This generalized eigenvalue problem is in Hessenberg-triangular form and can be solved by the Moler-Stewart variant [24, 30, 31] of Francis's algorithm, commonly called the QZ algorithm. In this paper we introduce a generalization of the method of [4] to matrix pencils, which we will call the *companion QZ algorithm*. This is a straightforward exercise, and it is not the main point of this publication.

This paper is really about backward error analysis. After the publication of [4] we realized that the analysis in that paper was not quite right. There is a factor that is constant in exact arithmetic and that we treated as a constant. We should have taken into account the backward error in that factor. In this paper we take the opportunity to repair that error. Moreover we will improve the analysis by taking into account

the structure of the backward error that follows from the structure of the companion matrix or pencil.

Our intuition told us that there would be cases where the companion QR method fails but companion QZ succeeds: just take a polynomial with a tiny leading coefficient! Once we had implemented companion QZ, we looked for examples of this kind, but to our surprise we were not able to find any. The companion QR method is much more robust than we had realized!

This discovery led us to take a closer look at the backward error analysis and try to explain why the companion QR algorithm works so well. This paper is the result of that investigation. We prove that the companion matrix method is just as good as the companion pencil method. Both methods have significantly better backward errors than a method like MATLAB's `roots` that computes the eigenvalues of the companion matrix without exploiting the special structure. This is the main message of this paper.

The paper is organized as follows. Section 2 briefly discusses previous work in this area. The memory-efficient $O(n)$ factorization of a companion pencil $V - \lambda W$ of the form (1.3) is introduced in Section 3, together with some necessary background information. In Section 4 we introduce the companion QZ algorithm and demonstrate its $O(n^2)$ performance.

The heart of the paper is the backward error analysis in Section 5. The main results are as follows: Let p be a monic polynomial with coefficient vector $a = \begin{bmatrix} a_0 & \cdots & a_{n-1} & 1 \end{bmatrix}$. Suppose we compute the zeros of p by some method, and let $\hat{p}$, with coefficient vector $\hat{a}$, be the monic polynomial that has the computed roots as its exact zeros. The *absolute normwise backward error on the coefficients* is $\|a - \hat{a}\|$. If our companion QR method is used to do the computation, the backward error satisfies $\|a - \hat{a}\| \lesssim u\|a\|$, where u is the unit roundoff and $\lesssim$ means less than or equal up to a modest multiplicative constant depending on n as a low-degree polynomial.[1] This is an optimal result, and it is better than can be achieved by the unstructured Francis algorithm applied to the companion matrix. The latter gives only $\|a - \hat{a}\| \lesssim u\|a\|^2$. If the companion *pencil* is used, we do not get the optimal result unless we apply the companion QZ algorithm (or any stable algorithm) to a rescaled polynomial $p/\|a\|$. If we do this, we get the optimal result $\|a - \hat{a}\| \lesssim u\|a\|$.

2. Earlier work. There are many ways [23] to compute roots of polynomials. Here we focus on companion matrix and pencil methods. Computing roots of polynomials in the monomial basis via the companion matrix has been the subject of study of several research teams. See [4] for a summary.

There is, to our knowledge, only one article by Boito, Eidelman, and Gemignani [9] that presents a structured QZ algorithm for computing roots of polynomials in the monomial basis. The authors consider a matrix pencil, say (V, W), where both V and W are of unitary-plus-rank-one form, V is Hessenberg and W is upper triangular. Both matrices are represented efficiently by a quasiseparable representation [9]. To counter the effects of roundoff errors, some redundant quasiseparable generators to represent the unitary part are created; a compression algorithm to reduce the number of generators to a minimal set is presented. The computational cost of each structured

[1] This result is valid for the improved turnover that is introduced § 5.1. If the old turnover is used, we get $\|a - \hat{a}\| \lesssim u\|a\|^2$. We reported this in [4], but for a valid proof see [3]. Even with the old turnover we can get a better result as follows. Instead of comparing p with $\hat{p}$, we can compare with $\gamma\hat{p}$, where γ is chosen so that $\|a - \gamma\hat{a}\|$ is minimized. Then we get $\|a - \gamma\hat{a}\| \lesssim u\|a\|$, as we have shown in [3].

QZ iteration is $O(n)$. A double shift version of this algorithm is presented by Boito, Eidelman, and Gemignani in [10].

The companion pencil (1.2) is the most frequently appearing one in the literature, but, there is a wide variety of comparable matrix pencils with the same eigenvalues [6, 14, 22], many of which are highly structured. In this article we will focus on companion pencils of the form (1.3).

3. Core transformations and factoring companion matrices. Core transformations will be used throughout the paper and are the building blocks for a fast algorithm and an efficient representation of the companion pencil.

3.1. Core transformations. A nonsingular matrix G_i identical to the identity matrix except for a 2×2 submatrix in position $(i : i+1, i : i+1)$ is called a *core transformation*. The subscript i refers to the position of the diagonal block $(i : i+1, i : i+1)$ called the *active part* of the core transformation. Core transformations G_i and G_j commute if $|i - j| > 1$.

In previous work [1, 2] the authors have used non-unitary core transformations, but here we will only use unitary core transformations. Thus, in this paper, the term *core transformation* will mean *unitary* core transformation; the active part could be a rotator or a reflector, for example.

To avoid excessive index usage, and to ease the understanding of the interaction of core transformations, we depict them as $\vec{\gamma}$, where the tiny arrows pinpoint the active part. For example, every unitary upper Hessenberg matrix Q can be factored as the product of $n-1$ core transformations in a *descending* order $Q = G_1 G_2 \cdots G_{n-1}$. Such a *descending sequence* of core transformations is represented pictorially by

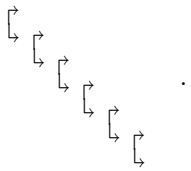

All of the algorithms in this paper are described in terms of core transformations and two operations: the *fusion* and the *turnover*.

Fusion. The product of two unitary core transformations G_i and H_i is again a unitary core transformation. Pictorially we can write this as

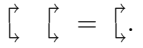

Turnover. The product of three core transformations $F_i G_{i+1} H_i$ is an essentially 3×3 unitary matrix that can be factored also as $F_{i+1} G_i H_{i+1}$, depicted as

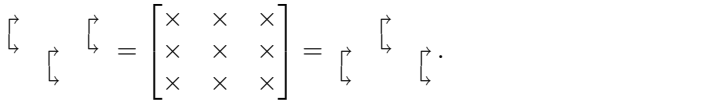

Pictorial action of a turnover and fusion. We will see, when describing the algorithms, that there are typically one or more core transformations not fitting the pattern (called the *misfit(s)*), that need to be moved around by executing turnovers and similarities. To describe clearly the movement of the misfit we use the following pictorial description:

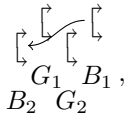

where B_1 is the misfit before the turnover and B_2 after the turnover. The core transformations G_1 and G_2 are involved in the turnover and change once B_2 is created. The picture is mathematically equivalent to $G_1G_2B_1 = B_2\hat{G}_1\hat{G}_2$. Other possible turnovers are

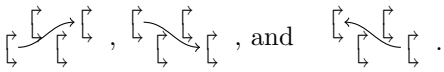

Directly at the start and at the end of each QZ (and QR) iteration, a misfit is fused with another core transformation, so that it vanishes. We will describe this pictorially as

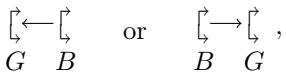

where B is to be fused with G.

3.2. A factorization of the companion pencil. The pencil matrices V and W from (1.3) are both unitary-plus-rank one, V is upper Hessenberg and W is upper triangular. We store V in QR decomposed form: $V = QR$, where Q is unitary and upper Hessenberg, and R is upper triangular and unitary-plus-rank-one. In fact

$$Q = \begin{bmatrix} 0 & & & 1 \\ 1 & & & 0 \\ & 1 & & 0 \\ & & \ddots & \vdots \\ & & 1 & 0 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 1 & & & -v_2 \\ & 1 & & -v_3 \\ & & \ddots & \vdots \\ & & & -v_n \\ & & & -v_1 \end{bmatrix}. \tag{3.1}$$

We need efficient representations of Q, R, and W. Q is easy; it is the product of $n-1$ core transformations: $Q = Q_1 \cdots Q_{n-1}$, with $Q_i(i:i+1, i:i+1) = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$.

Factoring an upper triangular unitary-plus-rank-one matrix. The matrices R (3.1) and W (1.3) have exactly the same structure and can be factored in the same way. This factorization was introduced and studied in detail in [4], so we will just give a brief description here. We focus on R. It turns out that for this factorization we need to add a bit of room by adjoining a row and column. Let

$$\underline{R} = \left[\begin{array}{cccc|c} 1 & & & -v_2 & 0 \\ & 1 & & -v_3 & 0 \\ & & \ddots & \vdots & \vdots \\ & & & -v_n & 0 \\ & & & -v_1 & 1 \\ \hline & & & 0 & 0 \end{array}\right]. \tag{3.2}$$

This is just R with a zero row and nearly zero column added. The 1 in the last column ensures that $\underline{R}$ is unitary-plus-rank-one: $\underline{R} = Y_n + \underline{z}e_n^T$, where

$$Y_n = \left[\begin{array}{cccc|c} 1 & & & & \\ & 1 & & & \\ & & \ddots & & \vdots \\ & & & 1 & \\ & & & 0 & 1 \\ \hline & & & 1 & 0 \end{array}\right] \quad \text{and} \quad \underline{z} = -\left[\begin{array}{c} v_2 \\ v_3 \\ \vdots \\ v_n \\ v_1 \\ \hline 1 \end{array}\right]. \tag{3.3}$$

Let $C_1, \ldots, C_n$ be core transformations such that $C\underline{z} = C_1 \cdots C_n \underline{z} = \alpha e_1$, where $|\alpha| = \|\underline{z}\|_2$. Let B be the unitary Hessenberg matrix $B = CY_n$. Clearly $B = B_1 \cdots B_n$, where $B_i = C_i$ for $i = 1, \ldots, n-1$, and $B_n = C_n Y_n$. This gives us a factorization of $\underline{R}$ as

$$\underline{R} = C_n^* \cdots C_1^* (B_1 \cdots B_n + \alpha\, e_1 \underline{y}^T) = C^*(B + \alpha\, e_1 \underline{y}^T), \tag{3.4}$$

with $\underline{y}^T = e_n^T \in \mathbb{R}^{n+1}$. In the course of our algorithm, the core transformations B_i, C_i, and the vector $\underline{y}$ will be modified repeatedly, but the form (3.4) for $\underline{R}$ will be preserved. $\underline{R}$ remains upper triangular with its last row equal to zero. The theory that supports these claims can be found in [4, § 4]. Notice that the core transformations B_n and C_n both make use of row and column $n+1$. Had we not added a row and column, this factorization would not have been possible.

Multiplying (3.4) on the left by e_{n+1}^T, we find that $0 = e_{n+1}^T \underline{R} = e_{n+1}^T C^* B + \alpha\, e_{n+1}^T C^* e_1 \underline{y}^T$, so [4, Thm. 4.6]

$$\alpha\, \underline{y}^T = -(e_{n+1}^T C^* e_1)^{-1} e_{n+1}^T C^* B. \tag{3.5}$$

This equation demonstrates that the information about $\alpha\, \underline{y}^T$, which determines the rank-one part, is encoded in the core transformations. This means that we will be able to develop an algorithm that does not keep track of $\underline{y}$; the rank-one part can be simply ignored. If at any time we should need $\underline{y}$ or some part of $\underline{y}$, we could recover it from C and B using (3.5). However, as we show in [4, § 4], it turns out that we never need to use (3.5) in practice.

Let $P = I_{(n+1)\times n}$, the $(n+1) \times n$ matrix obtained by deleting the last column from the $(n+1) \times (n+1)$ identity matrix. Then $R = P^T \underline{R} P$, so our factored form of R is

$$R = P^T C_n^* \cdots C_1^* (B_1 \cdots B_n + \alpha\, e_1 \underline{y}^T) P. \tag{3.6}$$

The matrices P and P^T are included just so that the dimensions of R come out right. They play no active role in the algorithm, and we will mostly ignore them. Pictorially, for $n = 8$, $\underline{R}$ (and hence also R) can be represented as

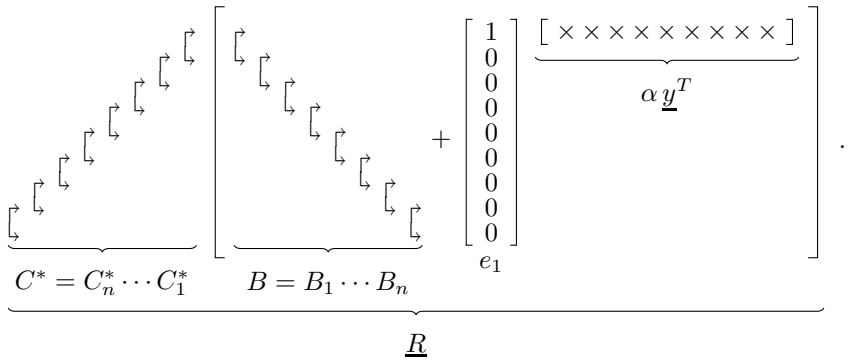

Since we can ignore $\alpha\, \underline{y}^T$, we see that R is represented by two sequences of core transformations, C and B. Hence $A = QR$ is represented by three sequences of core transformations.

The matrix W admits a factorization of the same form as (3.6)

$$W = P^T C_W^* (B_W + \alpha_W e_1 \underline{y}_W^T) P,$$

with $C_W [w_1, \ldots, w_n, 1]^T = \alpha_W e_1$.

Thus W is also represented by two sequences of core transformations. Altogether the pencil (V, W) is represented by five sequences of core transformations.

3.3. Core transformations and upper triangular matrices. In the next section we will show how to compute the eigenvalues of the matrix pencil (V, W) via the QZ algorithm as described by Vandebril and Watkins [27]. An important operation is to refactor the product of an upper triangular matrix times a core transformation RG_i as the product of a core transformation times an upper triangular matrix[2] $\hat{G}_i \hat{R}$. The other way proceeds similarly. Considering dense matrices we get pictorially

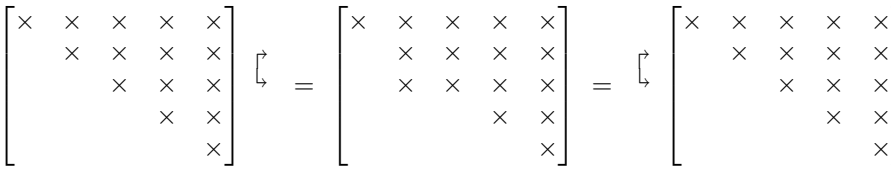

Applying a nontrivial core transformation from the right on the upper triangular matrix creates a non-zero subdiagonal, which can be removed by pulling out a nontrivial core transformation from the left.

In our case the upper triangular matrix is represented in a data-sparse way (3.4). Instead of explicitly creating the upper triangular matrix to interchange the order of the upper triangular matrix and a core transformation we operate on the factored form directly. There are two versions, *passing a core transformation* from right to left and vice versa.

Consider an arbitrary upper triangular unitary-plus-rank-one matrix, say, $R = P^T C^* (B + \alpha\, e_1 \underline{y}^T) P$. Let G_i $(1 \le i \le n-1)$ be the core transformation we want to move from right to left. We have $PG_i = \underline{G}_i P$, where $\underline{G}_i$ is the $(n+1) \times (n+1)$ version of G_i. We have $RG_i = P^T C^* (B + \alpha\, e_1 \underline{y}^T) \underline{G}_i P = P^T C^* (B \underline{G}_i + \alpha\, e_1 \underline{y}^T \underline{G}_i) P$, so $\underline{y}^T$ changes into $\hat{\underline{y}}^T = \underline{y}^T \underline{G}_i$. The next step is to execute a turnover $B_i B_{i+1} \underline{G}_i =$

[2] We assume that eigenvalues at zero or infinity are deflated beforehand, so the involved upper triangular matrices are nonsingular.

$\underline{\tilde{G}}_{i+1}\hat{B}_i\hat{B}_{i+1}$. We now have $RG_i = P^TC^*(\underline{\tilde{G}}_{i+1}\hat{B} + \alpha\, e_1\underline{\hat{y}}_1^T)P = P^TC^*\underline{\tilde{G}}_{i+1}(\hat{B} + \alpha\, e_1\underline{\hat{y}}_1^T)P$. In the last step we have used the fact that $\tilde{G}_{i+1}e_1 = e_1$ because $i+1 > 1$. To complete the procedure we just need to do one more turnover, in which $\underline{\tilde{G}}_{i+1}$ interacts with C_i^* and C_{i+1}^*. Specifically $C_{i+1}^*C_i^*\underline{\tilde{G}}_{i+1} = \underline{\hat{G}}_i\hat{C}_{i+1}^*\hat{C}_i^*$, resulting in $RG_i = P^T\underline{\hat{G}}_i\hat{C}^*(\hat{B} + \alpha\, e_1\underline{\hat{y}}^T)P$. Finally we have $P^T\underline{\hat{G}}_i = \hat{G}_iP^T$, where $\hat{G}_i$ is the $n \times n$ version of $\underline{\hat{G}}_i$. Here it is important that $i \leq n-1$. The final result is $RG_i = \hat{G}_iP^T\hat{C}^*(\hat{B} + \alpha\, e_1\underline{\hat{y}}^T)P = \hat{G}_i\hat{R}$.

The total computational effort required for the task is just two turnovers. The operation $\underline{\hat{y}}^T = \underline{y}^T\underline{G}_i$ is not actually performed, as we do not keep track of $\underline{y}$.

Pictorially for $n = 8$ and $i = 1$, we have

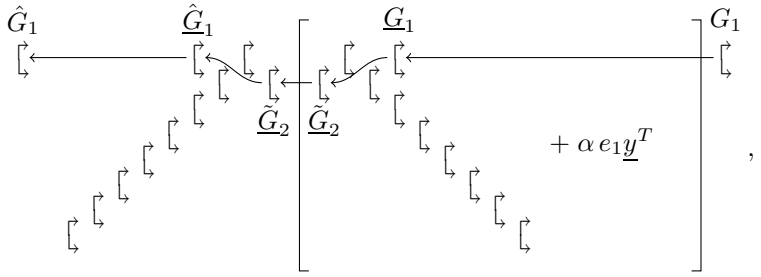

where we have not depicted P or P^T, and we have ignored the action on $\underline{y}^T$.

Notice that when we apply a core transformation G_{n-1}, we temporarily create $\underline{\tilde{G}}_n$, which makes use of row/column $n+1$. Here we see that the existence of an extra row/column is crucial to the functioning of the algorithm.

We can pass a core transformation from left to right through R simply by reversing the above procedure. We omit the details.

It is clear now that one can move a single core transformation through a factored upper triangular matrix in either direction by executing only two turnovers. From now on, to ease the notation, we will depict our Hessenberg-triangular pencil in a simpler format:

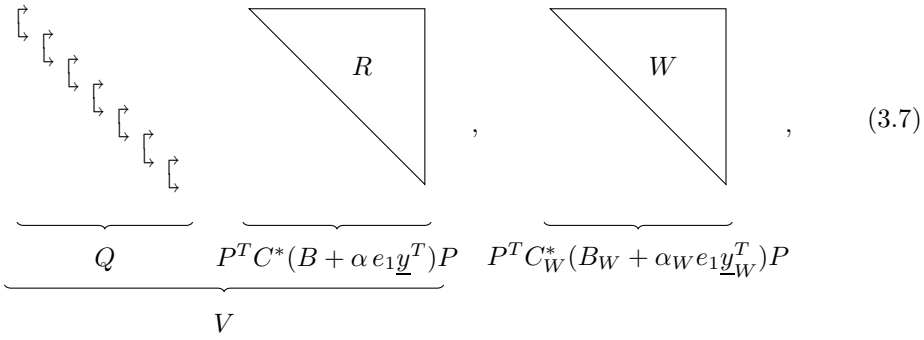

(3.7)

where we have replaced each upper triangular factor by a triangle. With this description, we can immediately apply the algorithms from Vandebril and Watkins [27]. For completeness, however, we will redescribe the flow of a single QZ step.

To facilitate the theoretical description we consider the product $S = RW^{-1}$, which is another upper triangular matrix through which we need to pass core trans-

formations. Moving a core transformation from the right to the left has two stages:

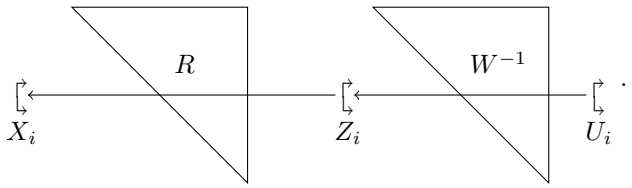

(3.8)

Since we do not wish to invert W, we do not literally execute the operation depicted on the right of (3.8). Instead we do the equivalent operation

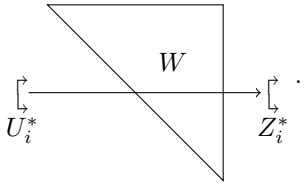

In fact, the latter operation remains valid, even in the case when W is singular. We consider W^{-1} to simplify the description in the next section.

Since both R and W are unitary-plus-rank-one matrices stored in the factored form described in Section 3, each of the two stages costs two turnovers. Thus the computational cost of passing a core transformation through S is four turnovers.

4. The companion QZ algorithm. We have implemented both single-shift and double-shift companion QZ algorithms. For simplicity we will describe only the single-shift case, as the double-shift iteration is a straightforward extension; we refer to Aurentz et al. and Vandebril and Watkins [4, 27]. The companion QZ algorithm is easily described by viewing it as a version of the companion QR algorithm applied to the matrix VW^{-1}. Clearly

$$VW^{-1} = QRW^{-1} = QS,$$

where $S = RW^{-1}$ is upper triangular. As discussed in Section 3.3 the matrix W does not need to be invertible and one could as well work on the pair (R,W) instead which is equivalent to the classical description of QZ algorithms. There are no issues when considering W and we can rely on standard techniques [28–30].

Pictorially

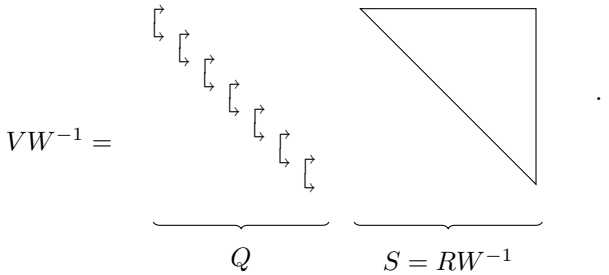

To begin the iteration we select a suitable shift μ and compute $q = (V - \mu W)e_1$. Only the first two entries of q are nonzero, so we can construct a core transformation U_1 such that $U_1^* q = \alpha e_1$ for some α. Our first modification to VW^{-1} is to apply a similarity transformation by U_1:

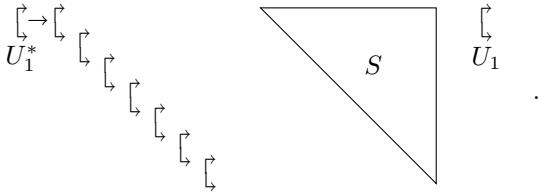

We can immediately fuse U_1^* with Q_1 to make a new Q_1. (To keep the notation under control, we do not give the modified Q_1 a new name; we simply call it Q_1.) We can also pass U_1 through S to obtain

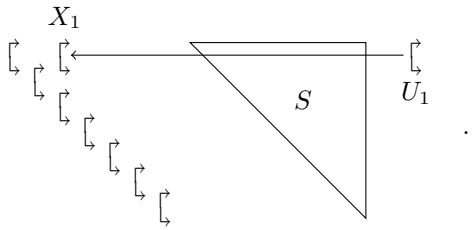

(4.1)

The details of passing a core transformation through S where described in Section 3.3.

If we were to multiply the factors together, we would find that the matrix is no longer upper Hessenberg; there is a bulge in the Hessenberg form caused by a nonzero entry in position $(3, 1)$. The standard Francis algorithm chases the bulge until it disappears off the bottom of the matrix. In our current setting we do not see a bulge. Instead we see an extra core transformation X_1 in (4.1), which is in fact the cause of the bulge. X_1 is the *misfit*. Instead of chasing the bulge, we will chase the misfit through the matrix until it disappears at the bottom. We therefore call this a *core chasing* algorithm.

Proceeding from (4.1), the next step is to do a turnover $Q_1Q_2X_1 = U_2\hat{Q}_1\hat{Q}_2$. Core transformations $\hat{Q}_1$ and $\hat{Q}_2$ become the new Q_1 and Q_2. Pictorially

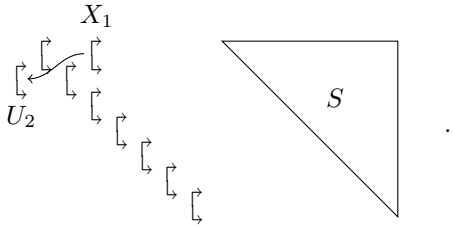

Next we do a similarity transformation, multiplying by U_2^* on the left and U_2 on the right. This has the effect of moving U_2 from the left side to the right side of the matrix. We can also pass U_2 through S to obtain

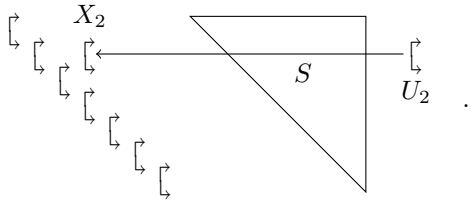

Now we are in the same position as we were at (4.1), except that the misfit has moved downward one position. The process continues as before:

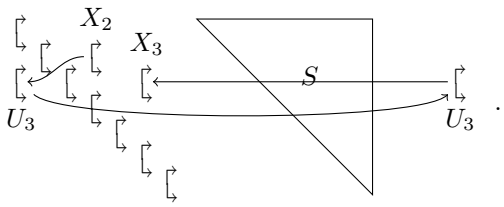

After $n-1$ such steps we arrive at

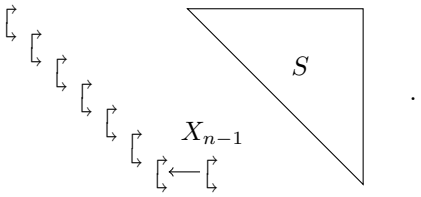

We can now fuse X_{n-1} with Q_{n-1}, completing the iteration.

Exploiting the representation of the factors W and V we get as a total cost for passing a core transformation through $VW^{-1} = QS$ five turnovers. The corresponding cost for companion QR [4] is three turnovers, so we expect the companion QZ code to be slower than the companion QR code by a factor of 5/3. During a QR or QZ iteration the misfit gets passed through the matrix about n times, so the total number of turnovers is $5n$ for a QZ step and $3n$ for a QR step. Either way the flop count is $O(n)$. Reckoning $O(n)$ total iterations, we get a flop count of $O(n^2)$ for both companion QR and QZ, with companion QZ expected to be slower by a factor of 5/3.

In Figure 4.1 we show execution times for our companion QR and QZ codes on polynomials of degree from 4 up to about 16000. We also make a comparison with the code from Boito, Eidelman, and Gemignani (BEGQZ) [9] and the LAPACK QR and QZ codes ZHSEQR and ZHGEQZ. Straight lines indicating $O(n^2)$ and $O(n^3)$ performance are included for comparison purposes. Our codes are the fastest. The lower plot in the figure corroborates our expectation that companion QZ will be slower than companion QR by a factor of about 5/3. For polynomials of low degree the LAPACK QR code DHSEQR is roughly as fast as our codes, but from degree 100 or so we are much faster than all other methods. The execution time curves for our companion QR and QZ codes are about parallel to the $O(n^2)$ line, indicating $O(n^2)$ execution time. The same is true of the BEGQZ method.

Table 4.1 shows the execution times for a few selected high degrees.

For this experiment we used a single core of an Intel® Xeon® CPU E5-2697 v3 running at 2.60GHz with 35 MB shared cache and 128 GB RAM. We used the GNU Compiler Collection gcc version 4.8.5 on an Ubuntu 14.04.1. For the comparison we used LAPACK version 3.7.1.

5. Backward Error Analysis. The norm symbol $\|\cdot\|$ will denote the 2-norm, i.e. the Euclidean norm for vectors and the spectral norm for matrices. These are the norms that we use in our backward error analysis. In our numerical tests we use the Frobenius matrix norm $\|\cdot\|_F$ instead of the spectral norm because it is easier to

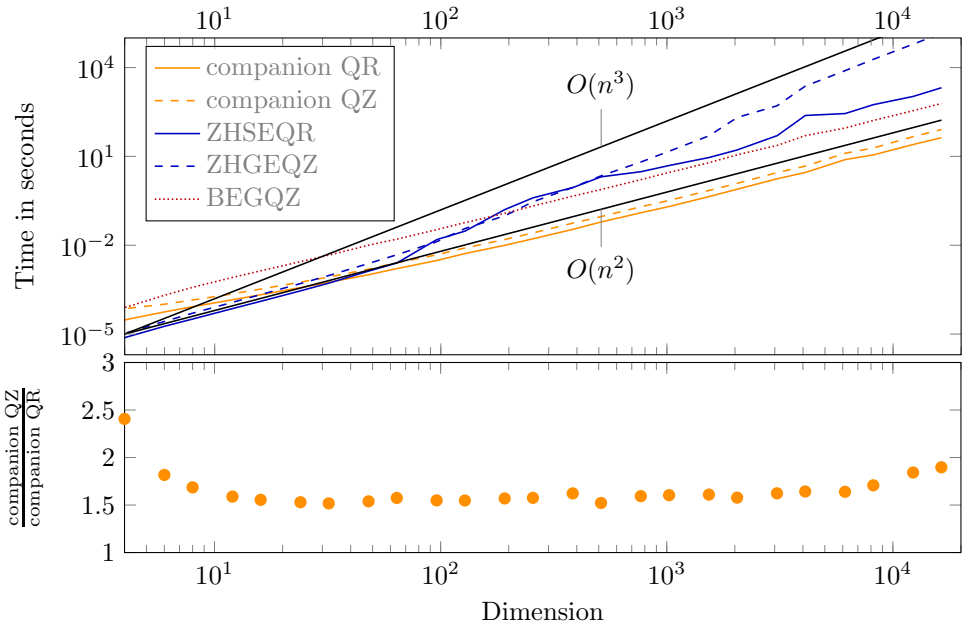

FIG. 4.1. *Execution times for several methods on polynomials of degree from 4 up to 16384. $O(n^2)$ and $O(n^3)$ lines are included for comparison.*

TABLE 4.1
Execution times in seconds for selected high degrees

degree	3072	6144	12288
comp. QR	2	8	25
comp. QZ	3	13	47
BEGQZ	23	91	350
ZHSEQR	50	280	1062
ZHGEQZ	514	7885	61856

compute. In fact the choice of norms is not important to the analysis; we could use other common norms such as $\|\cdot\|_1$ or $\|\cdot\|_\infty$ with no change in the results.

In addition we use the following conventions: $\doteq$ denotes an equality where second and higher order terms are dropped, $\lesssim$ stands for less than or equal up to a modest multiplicative constant typically depending on n as a low-degree polynomial, $\approx$ denotes equal up to a modest multiplicative constant. The symbol u denotes the *unit roundoff*, which is about 10^{-16} for IEEE binary64 arithmetic.

Van Dooren and Dewilde [26] were the first to investigate the backward stability of polynomial root finding via companion matrices and pencils. Edelman and Murakami [15] revisited this analysis, focusing on scalar polynomials. Jónsson and Vavasis [21] presented a clear summary of these results.

There are two important measures of backward accuracy when dealing with companions: the backward error (i) on the companion matrix or pencil and (ii) on the coefficients of the original polynomial. Let $a = \begin{bmatrix} a_0 & \cdots & a_n \end{bmatrix}^T$, the coefficient vector of p. Edelman and Murakami [15] showed that pushing an unstructured error

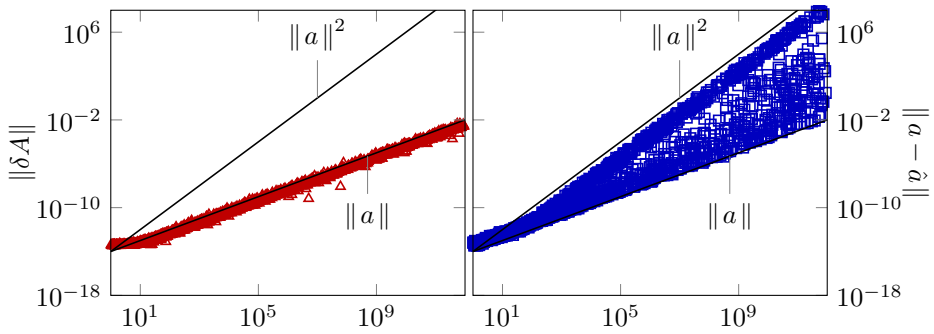

FIG. 5.1. *Backward error on the companion matrix (left) and the coefficient vector (right) as a function of* $\|a\|$ *when roots are computed by unstructured LAPACK code.*

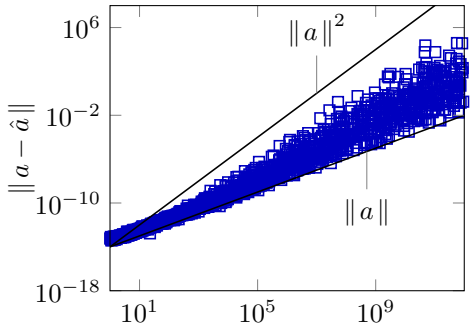

FIG. 5.2. *Same experiment as in Figure 5.1, except that the matrix is balanced before the eigenvalue computation. Only the backward error on the polynomial coefficient vector is shown.*

further back from the pencil (or matrix) to the polynomial coefficients introduces an additional factor $\|a\|$ in the backward error. We will show that we can do better since the backward error produced by our companion QR method is highly structured.

To illustrate what happens in the unstructured case we show in Figure 5.1 the absolute backward error as a function of $\|a\|$ when the companion matrix eigenvalue problem is solved by the QR algorithm from LAPACK. Twelve hundred random polynomials with varying norm between 1 and 10^{12} are produced by a method described below. For each sample we plot the backward error against $\|a\|$ (a single point). Black lines with slopes corresponding to $\|a\|$ and $\|a\|^2$ performance are also provided.

The graph on the left is the backward error on the companion matrix A. We see that this grows linearly as a function of $\|a\|$. This is consistent with the backward stability of the QR algorithm, which guarantees that the computed eigenvalues are the exact eigenvalues of a slightly perturbed matrix $A+\delta A$ with $\|\delta A\| \lesssim u\|A\| \approx u\|a\|$.

The graph on the right shows the backward error on the coefficients of the polynomial as a function of $\|a\|$. Note that the growth is quadratic in $\|a\|$ in the worst cases, consistent with the analysis of Edelman and Murakami [15].

In this and all subsequent experiments we used polynomials of varying norm such that the coefficients within each polynomial have widely varying magnitude. Specifically, we produced polynomials with increasing coefficient norm, parametrized by an integer $\rho = 1, \ldots, 12$. For each ρ we ran 100 experiments. We chose polynomials of degree 50, but the results for degree 6, 20, and 200 were very similar. For each of

the 51 coefficients of each polynomial we choose three random numbers ν, μ, and η, uniformly distributed in $[0,1]$. The coefficient is a complex number with argument $2\pi\nu$ and absolute value $(2\mu-1)10^{\rho(2\eta-1)}$. Polynomials of similar type have been used by De Terán, Dopico, and Pérez in [13].

In this and other experiments for which monic polynomials were needed, we made them monic by dividing through by the leading coefficient a_n.

Balancing [25] is often touted as an important step in solving the companion eigenvalue problem. In Figure 5.2 we have repeated the same experiment as in Figure 5.1, except that a balancing step is added. We see that balancing reduces the backward error on the polynomial coefficients. Later on we will show that our companion QR algorithm (with no balancing) has a much better backward error.[3] As a step in this direction we begin with some basic facts about the turnover.

5.1. Backward Stability of the Turnover. Anyone who has ever tried to program the turnover operation has discovered that there are many ways to do it, and some are much more accurate than others. As we have noticed recently, even among the "good" implementations, some are better than others.

First of all, as has been known for many years, it is crucial to maintain the unitary property of the core transformations. Each time a new core transformation is produced by computation of a c and an s, the unitary property $|c|^2 + |s|^2 = 1$ must be enforced by a rescaling operation. Assuming that c and s have been produced by valid formulas, the magnitude of the correction will be extremely tiny, on the order of the unit roundoff. This seemingly trivial correction is crucial. If it is neglected, the core transformations will gradually drift away from unitarity over the course of many operations, and the algorithm will fail. If it is not neglected, the programmer has a good chance of producing an accurate turnover.

It is worth mentioning as an aside that each of these rescaling operations requires a square root, but the number whose square root is being taken differs from 1 by a very small amount on the order of the unit roundoff. Therefore we can use the Taylor expansion $\sqrt{1+w} = 1 + .5\,w + O(w^2)$ to compute the square root to full precision cheaply or, more conveniently, we can use the formula $1/\sqrt{1+w} = 1 - .5\,w + O(w^2)$ to compute the reciprocal of the square root. Since the cost of square root operations contributes significantly to the cost of doing a turnover, and both the companion QR and companion QZ algorithms are dominated by turnovers, this shortcut has a significant impact.

In the data structure for our companion QR code, the core transformations are arranged in descending and ascending sequences $Q = Q_1 \cdots Q_{n-1}$, $B = B_1 \cdots B_n$, and $C^* = C_n^* \cdots C_1^*$ (3.4). Obviously the last of these is equivalent to the descending sequence $C = C_1 \cdots C_n$. In the course of the algorithm these sequences are modified repeatedly by turnover operations and (in Q only) occasional fusions. The following theorem shows that the backward errors associated with these modifications are tiny, on the order of the unit roundoff u.

THEOREM 5.1. *Suppose that in the course of the companion QR algorithm the descending sequence $G = G_1 \cdots G_n$ is transformed to $\hat{G}$ by multiplication on the right by U and on the left by V^*, where each of U and V is a product of core transformations.*

[3] However, a smaller backward error does not necessarily imply more accurate roots. It can happen that the balancing operation improves the condition number of the eigenvalue problem enough to more than offset the disadvantage in backward error. This is a matter for further study.

Then, in floating-point arithmetic

$$\hat{G} = V^*(G + \delta G)U,$$

where the backward error δG *satisfies* $\|\delta G\| \lesssim u$.

Here we can view U as the product of all core transformations that are absorbed by G on the right via turnovers (or fusions), and V is the product of all the core transformations that were ejected from G on the left. This is the right view if we are passing core transformations from right to left, but one can equally well think of passing core transformations from left to right. Then V^* is the product of all the transformations that are absorbed by G on the left, and U^* is the product of the core transformations that are ejected on the right. Either way Theorem 5.1 is valid.

Proof. Each fusion is just a matrix multiplication operation, and this is backward stable. Each turnover begins with a matrix multiplication operation that forms an essentially 3×3 unitary matrix. That matrix is then refactored by a process that is in fact a QR decomposition operation (on a unitary matrix!). These operations are all normwise backward stable [19]. Combining the small backward errors, and considering that all transformations are unitary and therefore do not amplify the errors, we conclude that the combined backward error δG satisfies $\|\delta G\| \lesssim u$. □

With Theorem 5.1 in hand we already have enough to build a satisfactory error analysis, but we can make the analysis stronger (and simpler) by taking into account a certain quantity that is conserved by the turnover operation. In $G = G_1 \cdots G_n$ each core transformation G_i has active part

$$\begin{bmatrix} c_i & -s_i \\ s_i & \overline{c_i} \end{bmatrix},$$

where s_i is real.[4] The interesting quantity that is preserved is the product $s_1 \cdots s_n$, as the following theorem shows. Exact arithmetic is assumed.

THEOREM 5.2. *Suppose* $G = G_1 \cdots G_n$ *is transformed to* $\hat{G} = \hat{G}_1 \cdots \hat{G}_n$ *by turnovers only (no fusions). Then*

$$\hat{s}_1 \hat{s}_2 \cdots \hat{s}_n = s_1 s_2 \cdots s_n.$$

Proof. It suffices to consider a single turnover. For the sake of argument we will think of passing a core transformation from right to left through G, but we could equally well go from left to right. Each turnover alters exactly two adjacent core transformations in the descending sequence. Before the turnover we have three adjacent cores

$$G_j G_{j+1} M_j = \begin{bmatrix} c_j & -s_j & \\ s_j & \overline{c}_j & \\ & & 1 \end{bmatrix} \begin{bmatrix} 1 & & \\ & c_{j+1} & -s_{j+1} \\ & s_{j+1} & \overline{c}_{j+1} \end{bmatrix} \begin{bmatrix} \gamma_j & -\sigma_j & \\ \sigma_j & \overline{\gamma}_j & \\ & & 1 \end{bmatrix}. \tag{5.1}$$

G_j and G_{j+1} belong to the descending sequence, and M_j is the misfit coming in from the right. We form the 3×3 product, which we then factor into three new core transformations

$$M_{j+1} \hat{G}_j \hat{G}_{j+1} = \begin{bmatrix} 1 & & \\ & \gamma_{j+1} & -\sigma_{j+1} \\ & \sigma_{j+1} & \overline{\gamma}_{j+1} \end{bmatrix} \begin{bmatrix} \hat{c}_j & -\hat{s}_j & \\ \hat{s}_j & \overline{\hat{c}}_j & \\ & & 1 \end{bmatrix} \begin{bmatrix} 1 & & \\ & \hat{c}_{j+1} & -\hat{s}_{j+1} \\ & \hat{s}_{j+1} & \overline{\hat{c}}_{j+1} \end{bmatrix}. \tag{5.2}$$

[4] It is always possible to make the s_i real, and that is what we have done in our codes. It allows for some simplification, but it is not crucial to the theory.

M_{j+1} will be the new misfit, which is going to be ejected to the left, while $\hat{G}_j$ and $\hat{G}_{j+1}$ will replace G_j and G_{j+1} in the descending sequence.

It is not hard to work out the products (5.1) and (5.2) explicitly, but for our present purposes we just need to compute the $(1,3)$ entry of each. These are easily seen to be $s_j s_{j+1}$ and $\hat{s}_j \hat{s}_{j+1}$, respectively. Since (5.1) and (5.2) are equal, we conclude that $\hat{s}_j \hat{s}_{j+1} = s_j s_{j+1}$. This proves the theorem. ∎

We remark that Theorem 5.2 is not applicable to the Q sequence, as Q is subjected to fusions. Moreover, Q is the sequence in which we search for deflations: a zero s_i is a lucky event signaling a deflation. The B and C sequences satisfy the conditions of Theorem 5.2 ensuring thereby that all turnovers are well-defined implying correctness of the algorithm [4].

Next we consider how Theorem 5.2 holds up in floating-point arithmetic. It turns out that this depends upon how the turnover is implemented. Using notation from the proof of the theorem, suppose we have computed $\hat{s}_j$ by some means. Then we can compute $\hat{s}_{j+1}$ by

$$\hat{s}_{j+1} = \frac{s_j s_{j+1}}{\hat{s}_j}. \tag{5.3}$$

If the turnover uses this formula, the product will be preserved to high relative accuracy. The multiplication and the division will each have a tiny relative error not exceeding u, and the computed $\hat{s}_j$ and $\hat{s}_{j+1}$ will satisfy $\hat{s}_j \hat{s}_{j+1} = s_j s_{j+1}(1+\delta)$, where $|\delta| \lesssim u$. This proves the following theorem.

THEOREM 5.3. *Suppose $G = G_1 \cdots G_n$ is transformed to $\hat{G} = \hat{G}_1 \cdots \hat{G}_n$ by turnovers only (no fusions) using floating point arithmetic. Suppose that the turnover uses the formula (5.3). Then*

$$\hat{s}_1 \hat{s}_2 \cdots \hat{s}_n = s_1 s_2 \cdots s_n (1+\delta),$$

where $|\delta| \lesssim u$.

In words, the product $s_1 \cdots s_n$ is preserved to high relative accuracy. The results presented in this paper were obtained using a (new!) turnover that uses the formula (5.3), so we will be able to use Theorem 5.3.

The good results in our publication [4] were obtained using a turnover (the *old* turnover) that did not use (5.3) and violated Theorem 5.3. That turnover preserves $s_1 \cdots s_n$ to high absolute accuracy but not to high relative accuracy. A satisfactory backward error analysis is still possible, and we did one. We have preserved that analysis in the technical report [3]. We were about to submit it for publication when we realized that a small change in the turnover yields more accurate results and a stronger and simpler backward error analysis. That is what we are presenting here.

5.2. Backward error on the companion matrix. We consider the companion QR algorithm (with the improved turnover) first, leaving companion QZ for later. We start with the backward error on the companion matrix. We will fix the error in the analysis of [4] and make other substantial improvements. We will take p to be monic with coefficient vector $a = \begin{bmatrix} a_0 & \cdots & a_{n-1} & 1 \end{bmatrix}^T$. The companion matrix is

$$A = \begin{bmatrix} & & & -a_0 \\ 1 & & & -a_1 \\ & \ddots & & \vdots \\ & & 1 & -a_{n-1} \end{bmatrix}. \tag{5.4}$$

Clearly $1 \le \|a\| \approx \|A\|$.

When we run the companion QR algorithm on A, we transform it to U^*AU, where U is the product of all of the core transformations that took part in similarity transformations. At the same time Q and R are transformed to U^*QX and X^*RU, where X is the product of all core transformations that were ejected from R and absorbed by Q. (Similarly we can view U as the product of all core transformations that were ejected from Q and absorbed by R.)

In floating-point arithmetic we have

$$\hat{A} = U^*(A + \delta A)U.$$

where δA is the backward error. The roots that we compute are exactly the eigenvalues of $\hat{A}$ and of $A + \delta A$. We would like to get a bound on $\|\delta A\|$.

We begin by looking at the backward error on R, and to this end we consider first the larger matrix $\underline{R}$ (3.2), which we can write in the factored form (3.4).

$$\hat{\underline{R}} = \underline{X}^*(\underline{R} + \underline{\delta R})\underline{U},$$

where $\underline{U} = \mathrm{diag}\{U, 1\}$ and $\underline{X} = \mathrm{diag}\{X, 1\}$. Consider the unitary and rank-one parts separately: $\underline{R} = \underline{R}_u + \underline{R}_o$, where $\underline{R}_u = C^*B$ and $\underline{R}_o = \alpha\, C^* e_1 \underline{y}^T = \alpha\, \underline{x}\, \underline{y}^T$. Here we have introduced a new symbol $\underline{x} = C^* e_1 = \alpha^{-1}\underline{z}$. We note that $\|\underline{x}\| = \|\underline{y}\| = 1$, and $|\alpha| = \|\underline{z}\| = \|a\|$ in this (monic) case. We will determine backward errors associated with these two parts: $\underline{\delta R} = \underline{\delta R}_u + \underline{\delta R}_o$.

Since $\underline{R}_u = C^*B$, we have $\underline{R}_u + \underline{\delta R}_u = (C + \delta C)^*(B + \delta B)$, where $\|\delta B\| \lesssim u$ and $\|\delta C\| \lesssim u$ by Theorem 5.1. We deduce that $\underline{\delta R}_u \doteq \delta C^* B + C^* \delta B$, and $\|\underline{\delta R}_u\| \lesssim u$.

For the rank-one part recall from (3.5) that

$$\alpha \underline{y}^T = -\rho^{-1} e_{n+1}^T \underline{R}_u, \quad \text{where} \quad \rho = e_{n+1}^T C^* e_1.$$

It is not hard to show that the quantity ρ remains invariant under QR iterations in exact arithmetic. Our mistake in the backward error analysis in [4] was to treat it as a constant, when in fact we should have taken the error into account. The simplest way to do this is to note that

$$\hat{\rho} = \rho + \delta\rho = e_{n+1}^T (C + \delta C)^* e_1,$$

so $\delta\rho = e_{n+1}^T \delta C^* e_1$, and $|\delta\rho| \lesssim u$. This computation, which shows that $\delta\rho$ is small in an absolute sense but not necessarily small relative to ρ, is valid for both the old and the new turnover.

With our new turnover we can get a better result: Let

$$\begin{bmatrix} c_i & -s_i \\ s_i & \overline{c}_i \end{bmatrix} \quad i = 1, \ldots, n$$

denote the active parts of the core transformations in the descending sequence $C = C_1 \cdots C_n$. A straightforward computation shows that $\rho = e_{n+1} C^* e_1 = (-1)^n s_1 \cdots s_n$. Therefore we can invoke Theorem 5.3 to deduce that the error in ρ is tiny relative to ρ:

$$\hat{\rho} = \rho(1 + \delta\rho_r) \quad \text{where } |\delta\rho_r| \lesssim u. \tag{5.5}$$

This result enables a stronger and easier error analysis.

Another straightforward computation shows that $-\rho^{-1} = \alpha$. Indeed,

$$\rho = e_{n+1}^T C^* e_1 = e_{n+1}^T \underline{x} = x_{n+1} = \alpha^{-1} z_{n+1} = -\alpha^{-1},$$

with $\underline{z}$ defined in (3.3). Thus $\underline{y}^T = e_{n+1}^T \underline{R}_u$ and

$$\underline{R}_o = \alpha\, \underline{x}\, \underline{y}^T = \alpha\, C^* e_1 e_{n+1}^T \underline{R}_u.$$

The backward error in $\underline{R}_o$ is given by

$$\underline{R}_o + \underline{\delta R}_o = \alpha(1+\delta\rho_r)^{-1}(C+\delta C)^* e_1 e_{n+1}^T (\underline{R}_u + \underline{\delta R}_u)$$

Making the approximation $(1+\delta\rho_r)^{-1} \doteq 1 - \delta\rho_r$, we get

$$\underline{R}_o + \underline{\delta R}_o = \alpha(\underline{x} + \underline{\delta x})(\underline{y} + \underline{\delta y})^T, \tag{5.6}$$

where $\underline{\delta x} \doteq (\delta C^* - \delta\rho_r C^*)e_1$ and $\underline{\delta y}^T = e_{n+1}^T \underline{\delta R}_u$. Note that $\|\underline{\delta x}\| \lesssim u$ and $\|\underline{\delta y}\| \lesssim u$.

Recall that $|\alpha| = \|a\| = \|\underline{z}\| = \|\underline{R}_o\| \approx \|R\| = \|A\|$, and the approximation is excellent when $\|A\|$ is large. Thus we will use the approximation $|\alpha| = \|a\| \approx \|A\|$ without further comment.

THEOREM 5.4. *If the companion eigenvalue problem is solved by our companion QR algorithm, then*

(a) $\hat{A} = U^*(A + \delta A)U$, *where* $\|\delta A\| \lesssim u\|A\|$.
(b) $\hat{Q} = U^*(Q + \delta Q)X$, *where* $\|\delta Q\| \lesssim u$.
(c) $\hat{R} = X^*(R + \delta R)U$, *where* $\|\delta R\| \lesssim u\|A\|$.
(d) More precisely,

$$\delta R \doteq \delta R_u + \alpha(\delta x\, y^T + x\, \delta y^T),$$

where $\|\delta R_u\| \lesssim u$, $\|\delta x\| \lesssim u$, *and* $\|\delta y\| \lesssim u$.

Proof. The work is mostly done; let's start with part (d). A first-order expansion of (5.6) gives $\underline{\delta R}_o \doteq \alpha(\underline{\delta x}\underline{y}^T + \underline{x}\underline{\delta y}^T)$, so

$$\delta R_o = P^T \underline{\delta R}_o P = \alpha(\delta x\, y^T + x\, \delta y^T),$$

where $\delta x = P^T \underline{\delta x}$ and so on. This establishes (d).

The factor α in the expression for δR_o implies that $\|\delta R\| \lesssim u\,\|a\| \approx u\,\|A\|$, so (c) holds. Part (b) is true by Theorem 5.1. Part (a) follows from (b) and (c) because $\delta A \doteq \delta Q\, R + Q\, \delta R$. ∎

Example 5.5. As a numerical test of Theorem 5.4 we computed the backward error on the companion matrix as follows. During the companion QR iterations we accumulated the transforming matrix U. We then computed $\check{A} = U\hat{A}U^*$. The backward error on A is $\delta A = \check{A} - A$. Figure 5.3 shows that the backward error on A grows in proportion to $\|A\|$ as claimed.

5.3. Backward error on the polynomial coefficients. We continue to study the backward error of the companion QR algorithm, leaving the companion QZ case for later. The monic polynomial

$$p(\lambda) = a_0 + a_1\lambda + a_2\lambda^2 + \cdots + a_{n-1}\lambda^{n-1} + \lambda^n$$

is associated with the coefficient vector

$$a = \begin{bmatrix} a_0 & \cdots & a_{n-1} & 1 \end{bmatrix}^T \in \mathbb{C}^{n+1}.$$

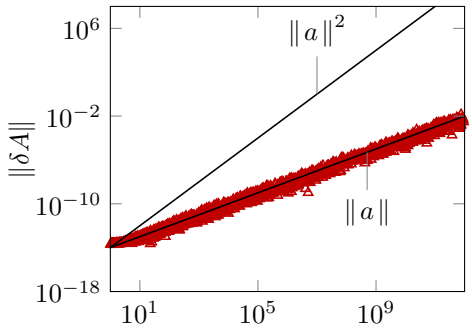

FIG. 5.3. *Companion QR method; norm of the backward error on the companion matrix A plotted against $\|a\|$*

When we compute the zeros of p, we don't get the exact zeros, but we hope that they are the zeros of a "nearby" polynomial. Let $\lambda_1, \ldots, \lambda_n$ denote the computed zeros and

$$\tilde{p}(\lambda) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_n).$$

This is the monic polynomial that has the computed roots as its zeros. We can compute the coefficients of $\tilde{p}$:

$$\tilde{p}(\lambda) = \tilde{a}_0 + \tilde{a}_1\lambda + \tilde{a}_2\lambda^2 + \cdots + \tilde{a}_{n-1}\lambda^{n-1} + \lambda^n$$

and the corresponding coefficient vector

$$\tilde{a} = \begin{bmatrix} \tilde{a}_0 & \cdots & \tilde{a}_{n-1} & 1 \end{bmatrix}^T \in \mathbb{C}^{n+1}.$$

This computation is done in multiple precision arithmetic using the multi precision engine of MPSolve 3.1.5 [8]. The quantity $\delta a = \tilde{a} - a$ is the backward error on the coefficients. We would like to show that $\|\delta a\|$ is tiny.

Theorem 5.4 shows that the norm of the backward error on the companion matrix A is directly proportional to the norm of the matrix. According to Edelman and Murakami [15], the backward error on the polynomial coefficients should then grow like the square of the norm, that is, $\|\delta a\| \lesssim u\,\|a\|^2$. In this section we show that we can do better: by exploiting the structure of the problem, we can make an argument that shows that the backward error on the polynomial depends linearly on the norm: $\|\delta a\| \lesssim u\,\|a\|$. This is an optimal result and is better than what is achieved by LAPACK's unstructured QR algorithm or Matlab's `roots` command.

Now let's get started on the analysis. In the factorization $A = QR$, the triangular factor can be written as $R = I + (\alpha\, x - e_n)y^T$, using notation established earlier in this section. Thus

$$\lambda I - A = (\lambda I - Q) - Q(\alpha\, x - e_n)y^T. \tag{5.7}$$

To make use of this equation we need the following known result [20, p. 26].

LEMMA 5.6. *Let $K \in \mathbb{C}^{n\times n}$, and let w, $v \in \mathbb{C}^n$. Then*

$$\det(K + w\,v^T) = \det(K) + v^T \operatorname{adj}(K)w,$$

where $\operatorname{adj}(K)$ *denotes the adjugate matrix of* K*.*

COROLLARY 5.7. *If* $A = QR = Q(I + (\alpha\, x - e_n)y^T)$*, then*

$$\det(\lambda I - A) = \det(\lambda I - Q) - y^T \operatorname{adj}(\lambda I - Q)Q(\alpha\, x - e_n).$$

The entries of the adjugate matrix are determinants of order $n-1$, so $\operatorname{adj}(\lambda I - Q)$ is a matrix polynomial of degree $n-1$:

$$\operatorname{adj}(\lambda I - Q) = \sum_{k=0}^{n} G_k \lambda^k, \tag{5.8}$$

with $\|\, G \,\|_k \approx 1$ for $k = 1, \ldots, n-1$, and $G_n = 0$.

The characteristic polynomial of Q is $\det(\lambda I - Q) = \lambda^n - 1$, which we will also write as

$$\det(\lambda I - Q) = \sum_{k=0}^{n} q_k \lambda^k,$$

with $q_n = 1$, $q_0 = -1$, and $q_k = 0$ otherwise. Using Corollary 5.7 and (5.8) we can write the characteristic polynomial of A as

$$p(\lambda) = \det(\lambda I - A) = \sum_{k=0}^{n} \left[q_k - y^T G_k Q(\alpha\, x - e_n) \right] \lambda^k. \tag{5.9}$$

The roots that we actually compute are the zeros of a perturbed polynomial

$$\tilde{p}(\lambda) = \det(\lambda I - (A + \delta A)) = \sum_{k=0}^{n} (a_k + \delta a_k)\lambda^k.$$

Our plan now is to use (5.9) to determine the effect of the perturbation δA on the coefficients of the characteristic polynomial. That is, we want bounds on $|\, \delta a_k \,|$.

LEMMA 5.8. *If* $\|\, \delta Q \,\| \lesssim u$ *then*

$$\operatorname{adj}(\lambda I - (Q + \delta Q)) = \sum_{k=0}^{n} (G_k + \delta G_k)\lambda^k \tag{5.10}$$

and

$$\det(\lambda I - (Q + \delta Q)) = \sum_{k=0}^{n} (q_k + \delta q_k)\lambda^k, \tag{5.11}$$

with $\|\, \delta G_k \,\| \lesssim u$ *and* $|\, \delta q_k \,| \lesssim u$*,* $k = 0, \ldots, n-1$*.*

Proof. For the bounds $|\, \delta q_k \,| \lesssim u$ we rely on Edelman and Murakami [15]. We have $\|\, \delta q \,\| \lesssim u \,\|\, q \,\|^2 = 2u$.

The adjugate and the determinant are related by the fundamental equation $B \operatorname{adj}(B) = \det(B) I$. Applying this with $B = \lambda I - Q$, we get

$$(\lambda I - Q) \sum_{k=0}^{n} G_k \lambda^k = \sum_{k=0}^{n} q_k I \lambda^k.$$

Expanding the left-hand side and equating like powers of λ, we obtain the recurrence

$$G_k = QG_{k+1} + q_{k+1}I, \quad k = n-1, \ldots, 0. \tag{5.12}$$

This is one half of the Faddeev-Leverrier method [16, p. 260], [18, p. 87]. Starting from $G_n = 0$, and knowing the coefficients q_k, we can use (5.12) to obtain all of the coefficients of $\mathrm{adj}(\lambda I - Q)$. The recurrence holds equally well with Q replaced by $Q + \delta Q$. We have

$$G_k + \delta G_k = (Q + \delta Q)(G_{k+1} + \delta G_{k+1}) + (q_{k+1} + \delta q_{k+1})I,$$

so

$$\delta G_k \doteq \delta Q\, G_{k+1} + Q\, \delta G_{k+1} + \delta q_{k+1} I. \tag{5.13}$$

If $\|\,\delta G_{k+1}\,\| \lesssim u$, we can deduce from (5.13) that $\|\,\delta G_k\,\| \lesssim u$. Since we have $\delta G_n = 0$ to begin with, we get by induction that $\|\,\delta G_k\,\| \lesssim u$ for all k. ∎

LEMMA 5.9. *If the companion QR algorithm is applied to the companion matrix $A = QR$, where $R = I + (\alpha\, x - e_n)y^T$, then the backward error δR satisfies*

$$R + \delta R = (I + \delta I) + (\alpha(x + \delta x) - e_n)(y + \delta y)^T,$$

where $\|\,\delta I\,\| \lesssim u$, $\|\,\delta x\,\| \lesssim u$, and $\|\,\delta y\,\| \lesssim u$.

Proof. From (5.6) it follows that

$$\underline{R} + \underline{\delta R} = \underline{R}_u + \underline{\delta R}_u + \alpha\,(\underline{x} + \underline{\delta x})(\underline{y} + \underline{\delta y})^T.$$

Projecting this down to $n \times n$ matrices we get

$$R + \delta R = R_u + \delta R_u + \alpha\,(x + \delta x)(y + \delta y)^T, \tag{5.14}$$

with $\|\,\delta R_u\,\| \lesssim u$, $\|\,\delta x\,\| \lesssim u$, and $\|\,\delta y\,\| \lesssim u$.

Recalling the form of $\underline{R}_u$, we see that $R_u = I - e_n y^T$, so

$$\begin{aligned} R_u + \delta R_u &= I - e_n y^T + \delta R_u \\ &= I - e_n(y + \delta y)^T + e_n\,\delta y^T + \delta R_u \\ &= I + \delta I - e_n(y + \delta y)^T, \end{aligned} \tag{5.15}$$

where $\delta I = e_n \delta y^T + \delta R_u$, and $\|\,\delta I\,\| \lesssim u$. Substituting (5.15) into (5.14), we get

$$R + \delta R = I + \delta I + (\alpha(x + \delta x) - e_n)(y + \delta y)^T,$$

completing the proof. ∎

THEOREM 5.10. *Suppose we apply the companion QR algorithm to the monic polynomial p with coefficient vector a. Let $\tilde{p}$, with coefficient vector $\tilde{a} = a + \delta a$, denote the monic polynomial that has the computed roots as its exact zeros. Then*

$$\|\,\delta a\,\| \lesssim u\,\|\,a\,\|.$$

Proof. From Theorem 5.4 we know that $\tilde{p}$ is the characteristic polynomial of a matrix $A + \delta A = (Q + \delta Q)(R + \delta R)$, with $\|\,\delta Q\,\| \lesssim u$. The form of $R + \delta R$ is given by Lemma 5.9.

$$\begin{aligned} A + \delta A &= (Q + \delta Q)(R + \delta R) \\ &= (Q + \delta Q)[(I + \delta I) + (\alpha(x + \delta x) - e_n)(y + \delta y)^T] \\ &= (Q + \widetilde{\delta Q}) + (Q + \delta Q)(\alpha(x + \delta x) - e_n)(y + \delta y)^T, \end{aligned}$$

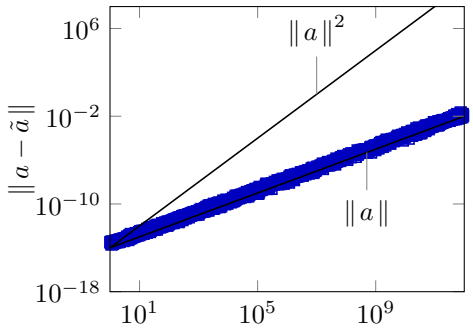

FIG. 5.4. *Backward error of companion QR on the polynomial coefficients plotted against* $\|a\|$

where $\widetilde{\delta Q} \doteq Q\,\delta I + \delta Q$ and $\|\widetilde{\delta Q}\| \lesssim u$. Now, using (5.9) with A replaced by $A + \delta A$, we get

$$\begin{aligned}\tilde{p}(\lambda) &= \det(\lambda I - (A + \delta A)), \qquad (5.16)\\ &= \sum_{k=0}^{n} \left[(q_k + \delta q_k) + (y + \delta y)^T (G_k + \delta G_k)(Q + \delta Q)(\alpha\,(x + \delta x) - e_n)\right] \lambda^k.\end{aligned}$$

where $\|\delta G_k\| \lesssim u$ and $|\delta q_k| \lesssim u$ by Lemma 5.8. Expanding (5.16) and ignoring higher order terms, we obtain

$$\begin{aligned}\delta a_k \doteq\ & \delta q_k + \delta y^T\, G_k\, Q(\alpha\, x - e_n) + y^T \delta G_k Q(\alpha\, x - e_n)\\ &+ y^T G_k\, \delta Q(\alpha\, x - e_n) + y^T G_k Q(\alpha\, \delta x)\end{aligned}$$

for $k = 0, \ldots, n-1$. Each term on the right-hand side has one factor that is $\lesssim u$. All terms except the first contain exactly one factor α and other factors that are ≈ 1. Thus $|\delta a_k| \lesssim u\,|\alpha| = u\,\|a\|$, and therefore $\|\delta a\| \lesssim u\,\|a\|$. ∎

Figure 5.4 verifies that the backward error grows linearly in $\|a\|$. If we compare this with Figures 5.1 and 5.2, we see that the companion QR algorithm (with no balancing step) has a significantly smaller backward error than the unstructured QR algorithm, with or without balancing.

So far throughout this section we have assumed for convenience that we are dealing with a monic polynomial. In practice we will often have a non-monic p, which we make monic by rescaling it. The following theorem covers this case.

THEOREM 5.11. *Suppose we compute the roots of a non-monic polynomial p with coefficient vector a by applying the companion QR algorithm to the monic polynomial p/a_n with coefficient vector a/a_n. Let $\tilde{p}$ denote the monic polynomial that has the computed roots as its exact zeros, let $\hat{p} = a_n\tilde{p}$, and let $a + \delta a$ denote the coefficient vector of $\hat{p}$. Then*

$$\|\delta a\| \lesssim u\,\|a\|.$$

Proof. Apply Theorem 5.10 to p/a_n, then rescale by multiplying by a_n. ∎

5.4. Backward error of the companion QZ algorithm. We now consider the backward error of the companion QZ algorithm, which finds the zeros of a non-monic polynomial

$$p(z) = a_0 + a_1 z + \cdots + a_{n-1} z^{n-1} + a_n z^n$$

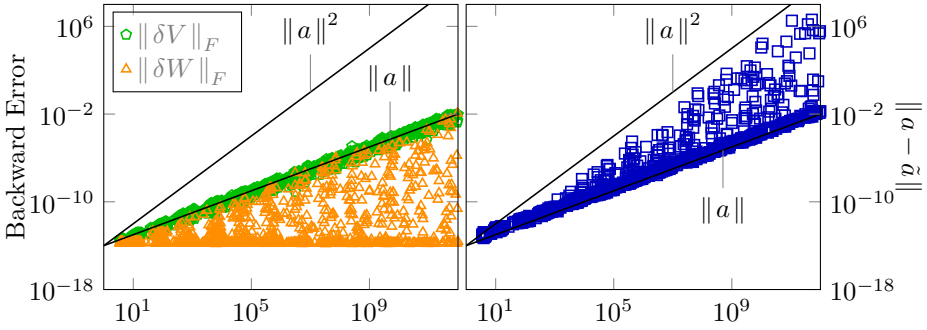

FIG. 5.5. *Backward errors* $\|\delta V\|_F$ *and* $\|\delta W\|_F$ *(both left) and* $\|a-\hat{a}\|$ *(right) of companion QZ plotted against* $\|a\|$

by computing the eigenvalues of a pencil $V-\lambda W$ of the form (1.3) with vectors v and w satisfying (1.4). We will assume a reasonable choice of v and w so that $\max\{\|v\|,\|w\|\}\approx\|a\|$, where a is the coefficient vector of p as before. (In fact, in all of our numerical experiments we have made the simplest choice, namely the one given by (1.2).)

Notice that in this setting we have the freedom to rescale the coefficients of the polynomial by an arbitrary factor. Thus we can always arrange to have $\|a\|\approx 1$, for example. This is the advantage of this approach, and this is what allows us to get an optimal backward error bound in this case.

When we run the companion QZ algorithm on (V,W), we obtain

$$\hat{V}=U^*(V+\delta V)Z,\quad \hat{W}=U^*(W+\delta W)Z,$$

where δV and δW are the backward errors. We begin with an analogue of Theorem 5.4.

THEOREM 5.12. *If the companion pencil eigenvalue problem is solved by the companion QZ algorithm, then*

(a) $\hat{V}=U^*(V+\delta V)Z$, *where* $\|\delta V\|\lesssim u\|a\|$,

(b) $\hat{W}=U^*(W+\delta W)Z$, *where* $\|\delta W\|\lesssim u\|a\|$.

Proof. The proof for $V=QR$ is identical to the proof of Theorem 5.4. The proof for W is even simpler because W is already upper triangular; there is no unitary Q factor to take into account. ∎

The left panel of Figure 5.5 gives numerical confirmation of Theorem 5.12. We see that the growth is linear in $\|a\|$ as claimed.

Now let us consider the backward error of the companion QZ algorithm on the polynomial coefficient vector a. We will compute an optimally scaled backward error as follows. Given the computed roots, we build the monic polynomial $\tilde{p}$ (with coefficient vector $\tilde{a}$) that has these as its exact roots. We then let $\hat{p}=\gamma\tilde{p}$ (with coefficient vector $\hat{a}$), where γ is chosen so that $\|a-\gamma\tilde{a}\|$ is minimized. We hope to get a backward error $\|a-\hat{a}\|$ that is linear in $\|a\|$, but the right panel of Figure 5.5 shows what we actually get. The backward error seems to grow quadratically in $\|a\|$, which is a disappointment.

Combining Theorem 5.12 with the analysis of Edelman and Murakami [15], we get the following result.

THEOREM 5.13. *The backward error of the companion QZ algorithm on the polynomial coefficient vector satisfies*

$$\|a-\hat{a}\|\lesssim u\|a\|^2.$$

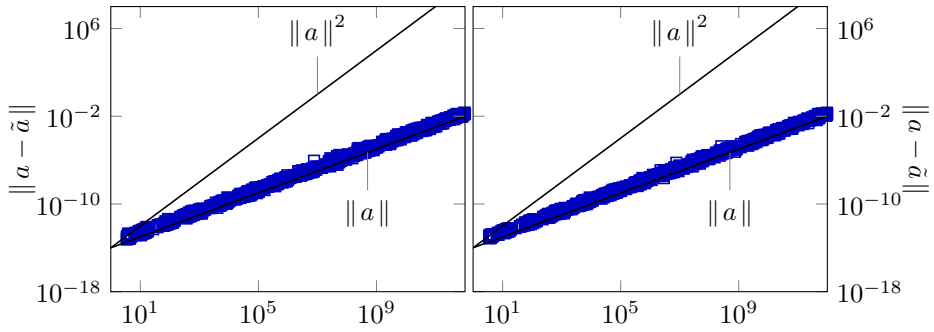

FIG. 5.6. *Backward error of scaled companion QZ algorithm (left) and unstructured QZ algorithm (right)*

So far it looks like companion QR is (surprisingly) more accurate than companion QZ, but we have not yet taken into account the freedom to rescale that we have in the companion QZ case.

THEOREM 5.14. *Suppose we compute the zeros of the polynomial p with coefficient vector a by applying the companion QZ algorithm to polynomial $p/\|a\|$ with coefficient vector $a/\|a\|$. Then the backward error satisfies*

$$\|a - \hat{a}\| \lesssim u\,\|a\|.$$

Proof. By Theorem 5.13 the backward error on $b = a/\|a\|$ satisfies $\|b - \hat{b}\| \lesssim u\,\|b\|^2 = u$. Therefore the backward error on a, which is $a - \hat{a} = \|a\|(b - \hat{b})$ satisfies $\|a - \hat{a}\| = \|a\|\,\|b - \hat{b}\| \lesssim u\,\|a\|$. ∎

In the interest of full disclosure we must point out that this argument applies equally well to any stable method for computing the eigenvalues of the pencil. If we use, for example, the unstructured QZ algorithm on the rescaled polynomial $p/\|a\|$, we will get the same result. Figure 5.6 provides numerical confirmation of Theorem 5.14. In the left panel we have the backward error for companion QZ, and in the right panel we have the backward error of the unstructured QZ code from LAPACK.

6. Conclusions. The companion QR algorithm is not only faster than the unstructured QR algorithm, it also has a smaller backward error on the polynomial coefficients: $\|\tilde{a} - a\| \lesssim u\,\|a\|$. In contrast the unstructured QR algorithm only satisfies $\|\tilde{a} - a\| \lesssim u\,\|a\|^2$.

As an alternative to the companion QR algorithm, we introduced a companion QZ algorithm that acts on a companion pencil. Like the companion QR algorithm it runs in $O(n^2)$ time and uses $O(n)$ storage, but it is slower by a factor of about 5/3. The backward error of the companion QZ algorithm also satisfies $\|\tilde{a} - a\| \lesssim u\,\|a\|$, provided that the polynomial is appropriately scaled before applying that algorithm.

When we began this project we fully expected to find classes of problems for which the companion QZ algorithm succeeds but companion QR fails. So far we have not found any; the companion QR algorithm is much more robust than we had believed. Since companion QR is faster, our recommendation at this time is to use companion QR and not companion QZ. We do not exclude the possibility that classes of problems for which companion QZ has superior performance will be found in the future.

7. Acknowledgments. We thank the referees and the associate editor for comments that improved the paper. In particular, one referee pointed out that we can simplify the error analysis significantly by making use of the Faddeev-Leverrier recurrence.

REFERENCES

[1] J. Aurentz, R. Vandebril, and D. S. Watkins, *Fast computation of the zeros of a polynomial via factorization of the companion matrix*, SIAM Journal on Scientific Computing, 35 (2013), pp. A255–A269.

[2] ———, *Fast computation of eigenvalues of companion, comrade, and related matrices*, BIT Numerical Mathematics, 54 (2014), pp. 7–30.

[3] J. L. Aurentz, T. Mach, L. Robol, R. Vandebril, and D. S. Watkins, *Fast and backward stable computation of roots of polynomials, Part IIA: general backward error analysis*, Tech. Rep. TW683, Dept of Computer Science, University of Leuven, KU Leuven, Belgium, 2017.

[4] J. L. Aurentz, T. Mach, R. Vandebril, and D. S. Watkins, *Fast and backward stable computation of roots of polynomials*, SIAM Journal on Matrix Analysis and Applications, 36 (2015), pp. 942–973.

[5] ———, *Fast and stable unitary QR algorithm*, Electronic Transactions on Numerical Analysis, 44 (2015), pp. 327–341.

[6] J. L. Aurentz, T. Mach, R. Vandebril, and D. S. Watkins, *A note on companion pencils*, Contemporary Mathematics, 658 (2016), pp. 91–101.

[7] D. A. Bini, P. Boito, Y. Eidelman, L. Gemignani, and I. Gohberg, *A fast implicit QR eigenvalue algorithm for companion matrices*, Linear Algebra and its Applications, 432 (2010), pp. 2006–2031.

[8] D. A. Bini and L. Robol, *Solving secular and polynomial equations: A multiprecision algorithm*, Journal of Computational and Applied Mathematics, 272 (2014), pp. 276–292.

[9] P. Boito, Y. Eidelman, and L. Gemignani, *Implicit QR for companion-like pencils*, Mathematics of Computation, 85 (2016), pp. 1753–1774.

[10] P. Boito, Y. Eidelman, and L. Gemignani, *A real QZ algorithm for structured companion pencils*. https://arxiv.org/abs/1608.05395, 2016.

[11] P. Boito, Y. Eidelman, L. Gemignani, and I. Gohberg, *Implicit QR with compression*, Indagationes Mathematicae, 23 (2012), pp. 733–761.

[12] S. Chandrasekaran, M. Gu, J. Xia, and J. Zhu, *A fast QR algorithm for companion matrices*, Operator Theory: Advances and Applications, 179 (2007), pp. 111–143.

[13] F. De Terán, F. M. Dopico, and J. Pérez, *Backward stability of polynomial root-finding using Fiedler companion matrices*, IMA Journal of Numerical Analysis, 36 (2016), pp. 133–173.

[14] B. Eastman, I.-J. Kim, B. Shader, and K. Vander Meulen, *Companion matrix patterns*, Linear Algebra and its Applications, 463 (2014), pp. 255–272.

[15] A. Edelman and H. Murakami, *Polynomial roots from companion matrix eigenvalues*, Mathematics of Computation, 64 (1995), pp. 763–776.

[16] D. K. Faddeev and V. N. Faddeeva, *Computational Methods of Linear Algebra*, W. H. Freeman, 1963.

[17] J. G. F. Francis, *The QR transformation, part II*, Computer Journal, 4 (1961), pp. 332–345.

[18] F. R. Gantmacher, *The Theory of Matrices, Vol. 1*, Chelsea, 1959.

[19] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, Second ed., 2002.

[20] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Second ed., 2013.

[21] G. F. Jónsson and S. Vavasis, *Solving polynomials with small leading coefficients*, SIAM Journal on Matrix Analysis and Applications, 26 (2004), pp. 400–414.

[22] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann, *Vector spaces of linearizations for matrix polynomials*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 971–1004.

[23] J. M. McNamee, *A bibliography on roots of polynomials*, Journal of Computational and Applied Mathematics, 47 (1993), pp. 391–394.

[24] C. B. Moler and G. W. Stewart, *An algorithm for generalized matrix eigenvalue problems*, SIAM Journal on Numerical Analysis, 10 (1973), pp. 241–256.

[25] B. N. Parlett and C. Reinsch, *Balancing a matrix for calculation of eigenvalues and eigenvectors*, Numerische Mathematik, 13 (1969), pp. 293–304.
[26] P. Van Dooren and P. Dewilde, *The eigenstructure of an arbitrary polynomial matrix: Computational aspects*, Linear Algebra and its Applications, 50 (1983), pp. 545–579.
[27] R. Vandebril and D. S. Watkins, *An extension of the QZ algorithm beyond the Hessenberg-upper triangular pencil*, Electronic Transactions on Numerical Analysis, 40 (2012), pp. 17–35.
[28] R. C. Ward, *The combination shift QZ algorithm*, SIAM Journal on Scientific and Statistical Computation, 12 (1975), pp. 835–853.
[29] D. S. Watkins, *Performance of the qz algorithm in the presence of infinite eigenvalues*, SIAM Journal on Matrix Analysis and Applications, 22 (2000), pp. 364–375.
[30] D. S. Watkins, *The Matrix Eigenvalue Problem: GR and Krylov Subspace Methods*, SIAM, Philadelphia, USA, 2007.
[31] D. S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, Inc., New York, Third ed., 2010.