The QR algorithm: 50 years later its genesis by John Francis and Vera Kublanovskaya and subsequent developments

GENE GOLUB

'February 29, 1932 – November 16, 2007, formerly Fletcher Jones Professor of Computer Science, Stanford University'

AND

FRANK UHLIG†

Department of Mathematics and Statistics, Auburn University, Auburn, AL 36849-5310

Fifty years after the invention of the QR algorithm by John Francis and Vera Kublanovskaya we reconstruct the ideas and the influences that led to its genesis from the originators' own recollections and their sources and give an account of some of its subsequent developments.

Keywords: QR algorithm; numerical algorithm; matrix computation; matrix eigenvalue; matrix eigenvector; Francis Implicit Q Theorem; John Francis; Vera Kublanovskaya; history of math; history of computation.

1. Introduction

On 29 October 1959 John Francis submitted his first theoretical QR algorithm paper. On 6 June 1961 he resubmitted the first part together with a second part on its implementation to *The Computer Journal* where they appeared in October 1961 in volume 4 (see Francis, 1961/1962a,b).

Vera Kublanovskaya submitted her first QR summary (Kublanovskaya, 1961a) on 5 July 1960, followed by two subsequent papers (Kublanovskaya, 1961b, 1962) with details. She is still active today in St Petersburg with research and teaching of numerical analysis and numerical linear algebra.

At the turn of the century and millennium in 2000, Dongarra & Sullivan (2000, p. 22) set out to name the top ten algorithms, sorted chronologically, 'with the greatest influence on the development and practice of science and engineering in the 20th century'; there is a secondary write-up by Cipra (2000). John Francis's *QR Algorithm for Computing Eigenvalues* is number six in that list. A brief sketch of QR by Parlett (2000) appears in the same issue as Dongarra & Sullivan (2000).

Yet, at the same time in 2000, nobody in the numerical analysis community of mathematicians had any recollection of ever having seen John Francis himself, nor was there any knowledge of John Francis—even if and where he lived! In the early 2000s the second author of this article was working on a book of applied numerical analysis for engineers (Elnashaie *et al.*, 2007). In the process he became aware of many an engineering apparatus, constant or equation to which a person's name was attached, but whose originator had become lost in the course of history. One of these was John Francis. So Frank searched for and found personal data on John Francis. Simultaneously and quite independently, the first author, Gene Golub, was engaged in a similar effort to learn about John Francis. Happenstance brought

---

†Email: uhligfd@auburn.edu

the two authors together in the Steklov Institute coffee room during the Second International Conference on Matrix Methods and Operator Equations in Moscow in July 2007. There Frank broached the subject of lost persons with named entities in modern science and engineering. Gene immediately asked, 'What about Francis?', to which Frank replied with John Francis's full name, date of birth, schools attended, places of work, etc., all to Gene's total amazement. Gene was planning to visit John on 7 August 2007 at his home in England and invited Frank to join him, but unfortunately Frank was not able to do so. Both tentatively agreed to try and publish John Francis's biography. Gene drove alone from Oxford to the south coast of England and met John Francis. The meeting was pleasant and Gene sent out a short note about it to his friends. The note, dated 13 August 2007, ended with: 'John Francis did remarkable work and we are all in his debt. Along with the conjugate gradient method, the QR algorithm provided us with one of the basic tools of numerical analysis.' Unfortunately, Gene died unexpectedly on 16 November 2007 at Stanford. Frank instantly vowed to finish the task. In late July 2008 Frank was able to visit with John for a couple of days. Much of the biography that follows is from my daily talks and walks with John Francis then and from subsequent letters, phone calls and e-mails.
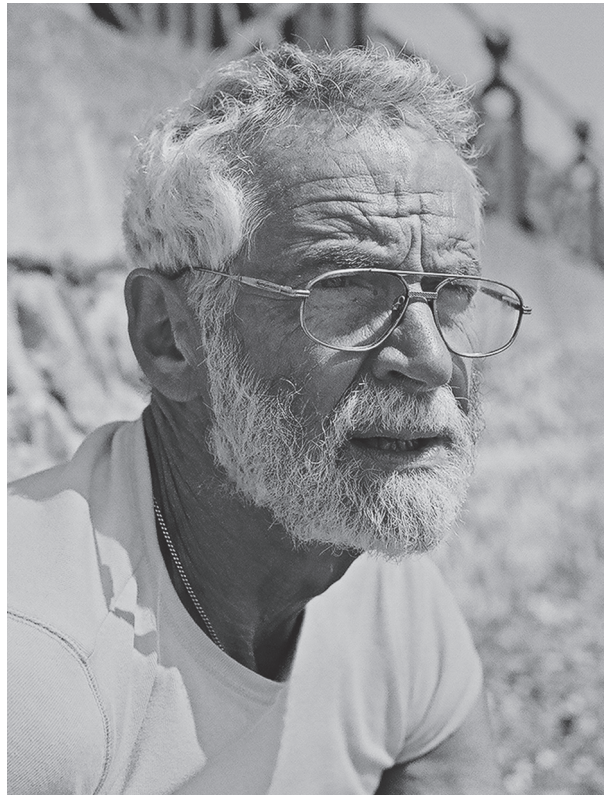
John Francis in July 2008

2. A short biography of John Francis

John Guy Figgis Francis was born on 10 October 1934 into a large family living on the northern outskirts of London, England. John's father was a British Army officer all his working life and in both world wars.

He died in the early 1960s, the same year that John's QR papers appeared in print. John's last given name is his mother's maiden name, the Figgises being an extensive Anglo-Irish family. John grew up in Mill Hill near where the M1 motorway now starts. In his youth he enjoyed the fields and paths into the country that lay to one side of his home. In school he was always quite good at mathematics and he took his A-levels from nearby Harrow School in 1952.

John was then conscripted into the British Army for 2 years of National Service. He spent part of his first Army year in Germany and he was posted for the last 13 months to Korea just after the ceasefire.

The year 1954 was a gap year in which Christopher Strachey employed John to work on the Pegasus computer wiring diagrams. Strachey had been a schoolmaster at Harrow and had taught him physics in 1950/1951. See *Catalogue of the Papers and Correspondence of Christopher Strachey* (1916–1975) for an account of Strachey's unorthodox path into early British computer science via a draughts (or checkers) playing computer program and Editorial (2000), Gordon (2000) and Roger Penrose (2000) for a number of recollections about his contribution to the foundations of computer science. At this time Strachey ran a computer group at the National Research and Development Corporation (NRDC) in Tilney Street in London. In 1955 John entered Christ's College at Cambridge University to study math. After his military experience, however, the theoretical work and the rigour of university math did not agree with him and he returned to NRDC, sharing various flats with friends in north London. The director of NRDC was Lord Halsbury (1908–2000), who saw the development and promotion of computers as a main objective of the Corporation and had been responsible for employing Strachey. Strachey regarded the eigenvalue problem, as exemplified by the need to solve the aircraft flutter equations, as an important task in numerical computation. When John returned to NRDC in 1956 as Strachey's assistant, he implemented several experimental matrix methods on the Pegasus computer, including the Schwarz algorithm for identifying eigenvalues with positive real parts. This work led to the development of the QR algorithm on the side (with Strachey's blessing) (see Section 3 below). Francis stayed at NRDC until 1961, 2 years longer than both Strachey and Lord Halsbury. Then he started work at Ferranti Ltd. By that time John was married and the couple had a baby daughter, the first of three children. Soon after, he dropped all connections with numerical analysis and his young family moved to Berkshire, where John commuted to work at Ferranti.

His work now involved developing a commercial language compiler for the Orion computer. When Ferranti folded in 1963, his division survived under the company names of ICT and later ICL. In 1967, looking for different directions, John joined funded research projects at the University of Sussex, the first one in a team working on Skinner box-type experimental psychology. Here he designed a computer language for use by the experimental psychologists and later joined a group of researchers who worked on artificial intelligence. These associations resulted in a couple of papers that include John's name, but no great breakthroughs. In 1972, with another change of direction, he began to work as a systems engineer involved in designing and building industrial computer systems, installed in several European countries. He continued this work until his retirement in 1991.

John Francis currently lives on the south coast of England, enjoys walking and sailing and is studying for a degree in mathematics and science through the Open University. When Gene first contacted and met him in 2007, John was entirely unaware and quite amazed that there had been many references to and extensions of his early work and that his QR algorithm was considered one of the 10 most important algorithms of the twentieth century. Actually, as far as John recalled in 2008, there had been no reaction, none whatsoever in the early 1960s when his seminal papers appeared, and then he had left the field.

3. On the genesis of the QR algorithm

Fifty years later it is appropriate to study John Francis's and Vera Kublanovskaya's QR papers again, to note what influenced their thinking and what their work brought to numerical analysis and matrix computations.

3.1 *The contents of Francis's QR papers*

One source of Francis's two papers (Francis, 1961/1962a,b) is Rutishauser's (1958) LR algorithm paper. As was standard at the time (see Wilkinson, 1959), Francis's papers call what is now universally called the QR factorization of a matrix a 'unitary triangular decomposition' and he refers to the Hessenberg form, named after Karl Adolf Hessenberg (1904–1959), as an 'almost triangular matrix'. On page 266 of Francis (1961/1962a), Francis sets out to find the Schur normal form of a square matrix. The first four pages of Francis (1961/1962a) describe the similarity progression and properties of QR's factor-and-reverse-order-multiply steps for general matrices in analogy to the LR steps proposed by Rutishauser the year before (and actually earlier; see Section 5). In Francis (1961/1962a, pp. 268, 269) there is a proof of the convergence of the lower triangle entries of the unshifted explicit QR iterates $A^{\{k\}}$ to zero provided that A is nonsingular and has eigenvalues of distinct moduli (Theorem 3). The next pages establish the invariance of normalcy, lower bandedness, symmetry and Hermitianness of a matrix during the QR iterations (Theorems 4 and 5). On page 270 Francis mentions the well-known connection between the kth QR iterate and the $2k$th LR iterate for positive-definite symmetric matrices in terms of the Choleski factors. On the same page Francis shows that first reducing a given matrix A to Hessenberg form reduces the cost of the subsequent QR iterations by an order of magnitude to $\mathcal{O}(n^2)$ for each. He proposes elimination methods to achieve Hessenberg form. In Theorem 6 Francis establishes that the Hessenberg form is preserved in the QR iterates. Francis emphasizes the convergence of the QR iterates to A's Schur normal form in Theorem 9. The rate of convergence of the iterates' subdiagonal entries $a_{ij}^{\{k\}}$ is established as $(|\lambda_i/\lambda_j|)^k$ on the final page 271 of Part 1 (Francis, 1961/1962a). Origin shift strategies, as well as deflation at the lower right corner of the iterated matrix, are then introduced for acceleration purposes. Throughout the paper Francis is keenly aware of stability issues connected with the use of unitary transformations and of the backward stability of the QR algorithm itself.

Even today the paper has a thoroughly modern feel to it, except for its discussions of fixed- versus floating-point arithmetic. This modernity is undoubtedly due to John's familiarity with Jim Wilkinson's fundamental numerical work at the National Physical Laboratory at Teddington. Francis visited there several times with Strachey between 1957 and 1959 to meet with James Hardy Wilkinson (1919–1986). As an aside, immediately after obtaining his Ph.D., the first author, Gene Golub, visited Cambridge University on an NSF Fellowship from spring 1959 through to July 1960. At that time John went to Cambridge colloquia regularly and Gene visited Teddington, where Wilkinson was based. So it is possible that both had actually met or been in the same room about 50 years ago.

Francis's second paper (Francis, 1961/1962b) deals with the implementation and applications of QR. The first implementation and test of QR took place at NRDC on the Pegasus computer. Francis's programs were written in assembly language. Algol-like pseudo-codes are given towards the end of Part 2 in Francis (1961/1962b). Francis was mainly interested in finding the eigenvalues of real matrices since the aircraft flutter matrix models that had to be solved at NRDC were real. On pages 332 and 333 he devises a stable way to avoid complex arithmetic for eigenvalue problems with real matrices that have real or complex conjugate eigenvalues. The Francis Implicit Q Theorem is established on page 333 in formula (13) and Theorem 11. This helped modify the QR algorithm from its first explicit QR

factorization—reverse-order multiply version in Part 1 (Francis, 1961/1962a)—to its modern implicit use. Due to the Implicit Q Theorem, Francis (1961/1962b, Section 11, Second Method) notes that two complex conjugate shifts of explicit QR can be performed simultaneously in one implicit double real step. Francis's trick is to create a bulge that protrudes from the Hessenberg form and involves the first column of the real matrix $(A - \lambda)(A - \overline{\lambda}I)$ if $\lambda \in \mathbb{C}\backslash\mathbb{R}$ and then chase this protrusion down the diagonal unitarily until Hessenberg form is re-established. All of this can be implemented over the reals. On page 334 he discusses the near singularity of a shifted QR algorithm iterate when the shift parameter is close to an eigenvalue of the trailing 2×2 diagonal block and of A—this is beneficial in much the same way as it is in inverse iteration.

Inverse iteration was much discussed at NRDC at the time and it inspired John to try shifts in QR. The noted near singularity of $A - \alpha I$ if α is close to an eigenvalue of A is exploited for deflation of the last row of A and subsequent reduction of the order by 1, and likewise for pairs of complex conjugate eigenvalues and a reduction of order 2 through implicit shifts. On pages 335—339 Francis discusses the implementation of Givens rotations and Householder eliminations in detail for their stability properties and the overall backward stability of the algorithm. The inspiration and reference for the latter were Fike's 1959 paper (Fike, 1959) on the preservation of matrix and eigenvalue condition numbers under unitary similarities. On page 339 Francis details a close precursor of what later became known as the Wilkinson shift. His conditional shift is taken as the eigenvalue of the trailing 2×2 diagonal block of an iterate $A^{\{k\}}$ that is closest to its (n, n) entry, with certain restraints on use. On page 338 the vectorial update formula and $\mathcal{O}(n^2)$ implementation of the Householder transformation inside each QR step are explained for Hessenberg matrices. Pages 338 and 339 contain discussions on the merits and use of fixed-point arithmetic as well as floating-point arithmetic in various parts of the algorithm's implementation on Pegasus. Two appendices explore the computed results for a number of real and complex matrices up to dimension 27×27. Appendix B of Francis (1961/1962b) contains tables of computed results, with Tables 8–13 on page 343 annotated in John Francis's hand.

A third paper by Francis (Strachey & Francis, 1961/1962), joint with Strachey, appears in the same volume of the *Computer Journal*. It deals with reducing a general matrix to Hessenberg form first and then to general tridiagonal form by using elementary Gaussian elimination similarities. The paper explains how this process relates to the Lanczos (1950) method and that it shares its difficulties.

When the papers Francis (1961/1962a,b) and Strachey & Francis (1961/1962) appeared in 1961/1962, Jim Wilkinson invited Francis to accompany him to some conferences in the USA. But John could not go due to his commitments to his family and new job with Ferranti Ltd.

Francis's QR papers were reviewed in Mathematical Reviews by Heinz Rutishauser (1918–1970). The referee of Part 1 (Francis, 1961/1962a) was William (Velvel) Kahan at Berkeley who 'delayed publication (for a short time) while trying, without success, to figure out why the QR iteration deserved to work, not that there was anything wrong with Francis's explanation of how it worked if it ever got close enough to an eigenvalue.'

In further private correspondence Kahan has summarized his thoughts on QR as follows: 'I think John's QR iteration is marvellous. Subsequent refinements, mainly in shift strategies, have made relatively little difference compared with the change from before QR to afterwards. And all the other eigenvalue-computing schemes (before) had failure modes that occurred at least as often as not, despite their association with respected names.'

Parlett (2000, p. 38) summarized his assessment of QR in the following words: 'The QR algorithm solves the eigenvalue problem in a very satisfactory way ... What makes experts in matrix computations happy is that this algorithm is a genuinely new contribution to the field of numerical analysis and not just a refinement of ideas given by Newton, Gauss, Hadamard, or Schur.'

And Higham (2003) gives the following assessment: 'The QR algorithm for solving the nonsymmetric eigenvalue problem is one of the jewels in the crown of matrix computations.'

3.2 *An outline of Francis's QR genesis*

John Francis developed the whole complex of the QR algorithm on his own, apart from the obvious influence of Rutishauser's paper (Rutishauser, 1958). John did not collaborate with anybody. The ideas, theorems and implementation for and of QR were all his. Of course, he was influenced by what came before him and went on around him. We shall explain these influences as John remembered them in the interview last July.

His knowledge of Fike's (1959) and Wilkinson's work on the stability of unitary matrix transformations helped John shift the emphasis from the Gaussian elimination-based LR algorithm of Rutishauser to the use of orthogonal Householder and Givens eliminations in QR. John was aware of the (slow) convergence of Jacobi's method that annihilates the largest subdiagonal entry in a symmetric matrix A via similarity until the matrix becomes diagonal and displays A's eigenvalues. This helped him gain an intuitive grasp of the convergence of Rutishauser's LR algorithm. He and his colleagues at NRDC had become aware of the power of inverse iteration for finding eigenvalues accurately from approximations. Inverse iteration can be traced back to the Aerodynamische Versuchsanstalt, Göttingen paper of 1944 by Helmut Wielandt (1910–2001) (Wielandt, 1944). The method was apparently brought to the group by Roger Penrose (private correspondence with the second author, 2009) when he was a consultant at NRDC. It inspired John to apply shifts in QR. The abundance of real matrix models for real-world problems further guided him to develop and implement the Implicit Q Theorem to obtain complex conjugate eigenvalue pairs of real matrices in real arithmetic. The resulting QR algorithm with conditional Wilkinson shifts, as proposed by Francis (1961/1962a,b), computes the eigenvalues of real and complex matrices quickly and accurately. His tests involved flutter matrix models of dimensions up to 27.

During World War II and the 1940s, Olga Taussky (1906–1995) had been working on flutter matrix models for the British Ministry of Aircraft Production at the National Physical Laboratory in Teddington. Their stability had to be established before building and flying the planes. Olga learned of Gershgorin's Theorem and applied it with good results to these models of dimensions below 10. Around 1960 John was excited to be able to extend this dimensional limit threefold.
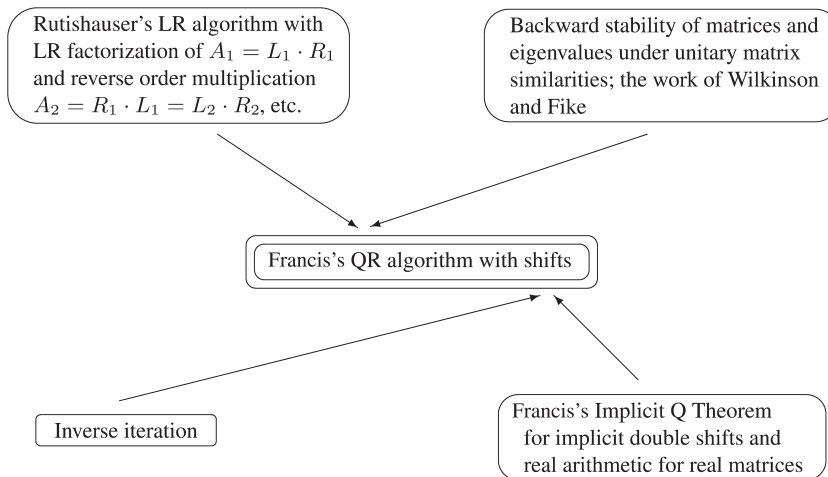
In our diagram the influences that appear in the top row make the QR algorithm accurate and backward stable, while the two items in the bottom row ensure the fast convergence of the method. Relying on all four, Francis's QR algorithm succeeds.

3.3 *The atmosphere for John Francis's work*

John Francis was a member of Christopher Strachey's team at the NRDC in the 1950s when he developed QR. The team was involved in the design and configuration of various computers that were built by Ferranti Ltd. The NRDC was also charged with solving the aircraft flutter problem numerically and needed to find the eigenvalues of real matrices that were deemed too large to handle by hand or by numerical methods of the time.

The atmosphere at NRDC was congenial and exhilarating. Strachey's team included excellent people, both on long-time appointments and as short-term visitors, among them Roger Penrose (2000) who recalls 'long and penetrating conversations ... with Strachey ... and later, John Francis's. All the important mathematical journals of the time were available at NRDC and the Bureau of Standards Math Series issue with Rutishauser's 1958 paper (Rutishauser, 1958) landed on John's desk one day. He never learnt where it came from or who brought it. Strachey's presence always seemed to have a catalytic effect on the people around him. His published work is small, but his influence on computing and computer science is immeasurable. While at NRDC, Francis and his colleagues would often take Thursday or Friday afternoons off to attend varying colloquia in Cambridge. The ritual was to go by train and learn about and discuss science, engineering, medicine, etc. Talks by Jim Wilkinson and others at the Wilkes' Colloquia in Cambridge's Mathematical Laboratory enriched John's understanding of modern numerical practices. In John's view, these were glorious times for computing and science. Strachey knew about Francis's endeavours with QR and gave him sufficient time and computer resources to finish the eigenvalue algorithm on the job.

3.4 *A short biography of Vera Kublanovskaya and an outline of her ideas for QR*

Vera Nikolaevna Kublanovskaya was born on 21 November 1920, one of nine siblings of a farming and fishing family in Krokhono in Vologda Oblast, Russia. Krokhono was a small village on the bank of Lake Beloye where the river Sheksna had its source. When the Volga–Baltic Waterway was rebuilt for larger vessels in the 1960s, Krokhono was swallowed by the lake as its level was raised by a dam. Vera went to normal school in Belozersk, 18 km from Krokhono. This city dates back to 862 and lies about 430 km east of Leningrad and 470 km north of Moscow. In 1939 she began her studies at the Gertzen Pedagogical Institute in Leningrad to become a teacher. There she met and was encouraged to pursue a career in mathematics by Dmitrii Konstantinovich Faddeev (1907–1989). In 1942 her mother's serious illness required her to go back home. After the siege of Leningrad had ended she wrote to D. K. Faddeev in 1945. Upon his recommendation she was immediately accepted to study mathematics at Leningrad State University without the obligatory entrance test. She graduated in 1948 and then joined the Steklov Mathematical Institute of the USSR Academy of Sciences (Leningrad branch, or LOMI), where she is still engaged today in research and publishing. At first she worked in a secret nuclear engineering project from which she retired in 1955. During that time she bore two sons and finished her candidates' thesis for a Ph.D. In 1955 Leonid Vitalievich Kantorovich (1912–1986) organized a group in Leningrad to develop a universal computer language, called Prorab, for BESM, the first electronic computer of the USSR that was completed in 1954 in Moscow. The easternmost corner tower of Lomonosov Moscow State University is dedicated to Professor Kantorovich, who lived in that part of the building with his

family until his death. At first, Kantorovich's group developed analytic computational tools in Prorab for algebraic and trigonometric polynomials, for integer arithmetic and series, etc. When Vera joined, her task was to select and classify matrix operations that are useful in numerical linear algebra. Linear algebra subroutines were included in Prorab much later. This experience brought Vera close to numerical algebra and computation. In 1972 she obtained her secondary doctorate (Habilitation). More extended biographies of Vera's life, achievements and long and productive career are available in Golub *et al.* (1990) and Konkova *et al.* (2003), for example.

Vera Kublanovskaya in August 2008

In the Russian literature the QR algorithm was initially called the *method of one-sided rotations*. Vera Kublanovskaya started to develop her version of the QR algorithm in 1958 after reading Rutishauser's LR algorithm paper (Rutishauser, 1958). She represented a matrix $A = L \cdot Q$ as the product of a lower triangular matrix L and an orthogonal matrix Q that she suggested to compute as the product of elementary rotations or reflections. Her eigenvalue algorithm factors $A = A_1 = L_1 \cdot Q_1$, reverse order multiplies $A_2 = Q_1 \cdot L_1$ and then factors A_2 again as $A_2 = L_2 \cdot Q_2$, etc. In 1959 she performed numerical experiments with her LQ decomposition-based algorithm on an electromechanic 'Mercedes' calculator in low dimensions by hand since linear algebra computer codes had not been developed yet in Prorab. Vera's hand computations indicated that her explicit factorization—reverse-order multiply algorithm— is convergent. However, D. K. Faddeev feared that the obtained results may only have been accidental. Vera's QR-like algorithm was briefly mentioned in 1960 in the first edition of a monograph

written by D. K. Faddeev jointly with his wife Vera Nikolaevna Faddeeva (1906–1983). Its second edition (Faddeev & Faddeeva, 1963, Chapter 8) discusses the QR-type eigenvalue algorithms of both Kublanovskaya and Francis. Vera's first short paper (Kublanovskaya, 1961a) on this subject states the basics of her LQ factorization and reverse-order multiply algorithm for nonsingular matrices A with real and distinct modulus eigenvalues, as well as the linear convergence of the diagonal entries of the L_k to the eigenvalues of A. Her second paper (Kublanovskaya, 1961b) proves at least linear convergence of the diagonal entries of L_k to the eigenvalues of such matrices and introduces simple shifts for improved convergence. Moreover, Kublanovskaya (1961b) indicates how the LQ method can be adapted to find the eigenvalues of $A^T \cdot A$ or $A \cdot A^T$, i.e., the singular values of A without forming the matrix product $A^T A$ or $A A^T$ explicitly. For real matrices, now with arbitrary eigenvalues, Vera's third LQ paper (Kublanovskaya, 1962) proves the convergence of the L_k to a lower block triangular form with 1×1 or 2×2 real diagonal blocks by using determinants. Thus she put Faddeev's earlier fears to rest.

During this time and later, neither Vera nor her co-workers implemented her method of one-sided rotations or QR for electronic computers. Her papers include no implementations or codes and her version of QR was never used to solve any applied eigenvalue problems. Instead, Vera and members of her group solved low-dimensional matrix eigenproblems with $n \leqslant 5$ by using her method and hand computations with 'Mercedes' calculators imported from East Germany.

The papers Kublanovskaya (1961b, 1962) were reviewed for Mathematical Reviews by Alston Scott Householder (1904–1993) of the Oak Ridge National Laboratory. Alston's work is also listed among 'The 10 Top Algorithms' (Dongarra & Sullivan, 2000) for his 1951 contribution of the 'Decompositional Approach to Matrix Computations', i.e., matrix factorizations in numerical linear algebra.

Later on, Vera was not involved in further investigations of the QR algorithm itself. However, some of her studies can be considered as extensions of the ideas underlying QR. The ideas of QR were exploited by Vera and her students and colleagues in designing methods for solving spectral problems in different situations such as for matrix pencils, for polynomial and rational matrices and, in particular, for two-parameter and multiparameter problems.

4. Subsequent uses and extensions of John Francis's and Vera Kublanovskaya's QR work

As of March 2009, according to the Mathematical Reviews, there are 361 papers that include QR in their title and 1660 papers whose review mentions QR. For the Zentralblatt, these numbers are 397 and 1802, respectively. When searching for 'QR algorithm', Google returns 363,000 items, while Google Scholar finds 108,000.

In this section we list a limited number of extensions to Francis's and Kublanovskaya's QR algorithm and we limit our references to selected author names only. Mathematical Reviews or Zentralblatt searches are more appropriate to find the full extent of applications and extensions of the original QR algorithm over the last 50 years.

4.1 *The Implicit Q Theorem of Francis, applications, generalizations, as well as QZ*

The reduction to Hessenberg form of a square matrix A via an orthonormal basis change is not unique. However, if the first unit vector of the basis change matrix is specified then the Hessenberg matrix becomes essentially unique. This was first observed by Francis in Francis (1961/1962a, p. 333 and Theorem 11).

> FRANCIS'S IMPLICIT Q THEOREM. Let U and V be two orthogonal $n \times n$ matrices with respective columns u_i and v_i in $\mathbb{C}^n$ that transform a matrix $A \in \mathbb{C}^{n,n}$ to the upper Hessenberg matrices $U^T A U = H_u$ and $V^T A V = H_v$, respectively.

Let j be the minimal index with $H_{u(j+1,j)} = 0$ in the case when H_u reduces and $j = n$ otherwise. If $u_1 = v_1$ then $u_k = e^{i\theta_k} v_k$ for $\theta_k \in \mathbb{R}$ and $|H_{u(k,k-1)}| = |H_{v(k,k-1)}|$ for all $k \leqslant j$. If $j < n$ then $H_{v(j+1,j)} = 0$ as well.

This theorem has many applications. Francis (1961/1962a,b) originally used it to replace two complex conjugate shifts for a real matrix A by an implicit double shift over the reals. The QR multishift techniques of Braman, Byers and Mathias from 2002 (Braman *et al.*, 2002a,b) do so for speed-up of QR with large sized matrices with n up to 10,000, using chains of bulges to avoid numerical instabilities when incorporating a large number of simultaneous shifts. These shifts are determined—as is customary—from a bottom right submatrix of appropriate size from an iterate. These authors also described advanced deflation techniques for multishifted QR.

Implicit QR steps usually create one or more protrusions in the QR iterates' Hessenberg structure, called a bulge, that is subsequently chased down the diagonal cheaply via local orthogonal similarities to achieve Hessenberg form again. Once that form is regained, all the shifts have been performed simultaneously, cheaply and implicitly.

Implicit shifts performed on an orthogonal bidiagonalization of a matrix A allow Golub and Kahan's SVD algorithm of 1965 (Golub & Kahan, 1965) to work exclusively with two vectors and chase bulges accordingly, rather than compute the eigenvalues of the worse conditioned matrices $A^T A$ or $A A^T$ as theory would have us do to find the singular values of A.

The Implicit Q Theorem generalizes to an Implicit DQ Theorem that applies the same principle to and allows the same applications for generalized D-orthogonal matrices A defined as follows. For $D = \text{diag}(\pm 1)$ a matrix $A \in \mathbb{R}^{n,n}$ is called D-orthogonal if $A^T D A = D$. If $D = I_n$ then D-orthogonality is ordinary orthogonality. Note that D-orthogonality is sometimes called J-orthogonality in the literature. Recent applications of D-orthogonality and the Implicit DQ Theorem in eigenvalue computations are the DQR algorithm of Uhlig (1997), its application to polynomial root finding in $\mathcal{O}(n^2)$ time with good results for moderately often repeated roots by Uhlig (1999), as well as recent work by Gemignani & Uhlig (2009) with implicit double DQR shifts for structured eigenvalue computations and for the secular equation.

Another generalization of QR involves the generalized eigenvalue problem $Ax = \lambda Bx$ for two $n \times n$ matrices A and B. Here the QZ algorithm of Moler & Stewart (1973) is an extension of QR. In QZ the given matrix pair A and B is first reduced to upper Hessenberg $\overline{A}$ and upper triangular form $\overline{B}$, respectively, through simultaneous multiplication of both A and B on the left by unitary Q and on the right by unitary Z. Thus the name. Then a QR style algorithm of creating bulges and chasing is performed implicitly on the matrix $\overline{A}\overline{B}^{-1}$ to find the generalized eigenvalues for A and B once A has been reduced to quasi-triangular form.

4.2 *How to understand Francis's QR algorithm in a numerical context*

Right from its invention around 1960, the QR algorithm has given numerical analysts a difficult task in understanding, explaining and proving. While both inventors of QR had supplied simple proofs of convergence for the nonshifted QR algorithm, namely Francis (1961/1962a, pp. 269, 270) by connecting QR to the matrix powers A^k and their effect on a basis and Kublanovskaya (1962) by using a complicated determinant argument, the shifted QR algorithm seems to require special proofs for each new shifting strategy. The QR algorithm with Wilkinson shifts, for example, was established and recognized very quickly as very good, i.e., reliable, fast and accurate. But a complete explanation was lacking for a considerable time (see below). To illustrate, even today nobody has found a way to guarantee the

convergence of QR with shifts for every non-Hermitian matrix. All working QR codes execute special *ad hoc* shifts when the algorithm does not converge as quickly as anticipated. These shifts are still derived from the entries in the left top or bottom right corner of the current iterate as Francis suggested. Very few problem matrices are known for which this is necessary, however.

After Rutishauser, Francis and Kublanovskaya, our understanding of the convergence behaviour of QR progressed very slowly. Householder's (1964) book gives a first glimpse in Section 7.9, followed by Wilkinson's (1965) book with 84 pages on LR and QR in Chapter 8 that connect QR with simultaneous iteration. The thesis of Buurema (1970) gives a geometric proof of QR's convergence with and without shifts using simultaneous iteration and Krylov subspace ideas. Stewart's (1973) book links QR to the power and inverse power method and Rayleigh quotient iteration in Section 7.2, with the whole of Chapter 7 devoted to QR and its extension to the SVD and generalized matrix eigenproblems. A detailed, reasoned analysis of why and how QR converges has been given by David Watkins's in two *SIAM Review* articles in 1982 and 2008 (Watkins, 1982, 2008). The insights of Watkins on QR also appear in Lectures 28 and 29 of Trefethen and Bau's (1997) textbook with SIAM. Watkins's point is that QR is a successful implementation of simultaneous iteration on $A_{n,n}$ combined with a change of basis in each iteration. Simultaneous iteration extends the power method that led to Francis's first proof. In practical terms, QR or simultaneous iteration is executed in terms of shifts of A for $(A - \alpha_1 I_n) \cdots (A - \alpha_k I_n)$. This creates a bulge of the intermediate associated Hessenberg forms of A of small size $k \times k$ in $\mathcal{O}(k^3)$ time. This small bulge is then chased down and to the right orthogonally in $\mathcal{O}(n^2)$ time to form another Hessenberg matrix as the next QR iterate in the implicitly shifted or multishifted version of QR. Thus implicitly shifted QR is a subspace iteration that finds invariant subspaces under A, thereby reducing the $n \times n$ eigenvalue problem for A to smaller ones until all eigenvalues are found.

4.3 *Fast and structured QR-like eigensolvers*

Originally, the QR algorithm was conceived for dense and unstructured matrices. For a Hermitian matrix A, however, the initial $\mathcal{O}(n^3)$ QR reduction step to Hessenberg form leads to a tridiagonal Hermitian matrix for which QR computes the eigenvalues in $\mathcal{O}(n^2)$ time instead of $\mathcal{O}(n^3)$ required for dense Hessenbergs. Recent efforts have achieved similar speed-up for certain classes of rank-structured matrices such as quasi-separable Hermitian matrices under low-rank perturbations, rank-one perturbations of unitary matrices, tridiagonal plus arrow matrices and symmetric arrow matrices. Fast eigensolvers that are based on QR or QR-like methods for these matrices have been studied by Marc van Barel, Dario Bini, Shivkumar Chandrasekaran, Yuli Eidelman, Luca Gemignani, Ming Gu, Israel Gohberg and others (Gemignani, 1999; Bini *et al.*, 2003, 2005, 2007; Bindel *et al.*, 2005; Eidelman *et al.*, 2005, 2008; Vandebril *et al.*, 2005; Chandresekaran *et al.*, 2008; Delvaux & van Barel, 2008).

At first, three independent research groups in Pisa, California and Leuven were simultaneously interested in designing fast and effective eigensolvers for suitable classes of structured matrices. The Pisa group wanted to find roots of polynomials and other algebraic equations efficiently. At Berkeley and the University of California, Santa Barbara, eigenvalue problems for discretized integral equations were the driving force, while in Leuven the interest lay in adapting linear algebraic techniques to deal with computational problems in the theory of orthogonal rational functions. Now the separate interests have converged and cooperation has led to growth.

A different type of structured matrices are the Hamiltonian and symplectic matrices. They occur in problems of particle physics and quantum mechanics, etc. For the skew-symmetric matrix J defined as

$$J = \begin{pmatrix} O_n & I_n \\ -I_n & O_n \end{pmatrix},$$

a $2n \times 2n$ real matrix A is called Hamiltonian if JA is symmetric. A is skew-Hamiltonian if $JA = -(JA)^{\mathrm{T}}$ and symplectic if $A^{\mathrm{T}}JA = J$. Symplectic spaces generalize the indefinite inner product spaces used for DQR (see Section 4.1) further to spaces with a skew-symmetric metric. QR-like eigensolvers for such matrices with names such as MQR, HR and SR are associated with Peter Benner, Angelika Bunse-Gerstner, Ralph Byers, Heike Faßbender, Daniel Kressner, Volker Mehrmann, Charlie van Loan, David Watkins and others (Bunse-Gerstner, 1981, 1986; van Loan, 1984; Bunse-Gerstner & Mehrmann, 1986; Byers, 1986; Bunse-Gerstner *et al.*, 1989, 1992; Benner *et al.*, 1998, 1999; Fassbender, 2000, 2006, 2007; Kressner, 2005).

Moreover, speed-up has also been achieved by implementing QR-like eigenvalue algorithms for specific sparse and structured matrices in terms of their individual entries, much as the qd forerunner of QR (see Section 5) did for Hermitian matrices.

Separately, the QR algorithm has been used in a backwards way to achieve desirable pole or eigenvalue locations for a system under feedback control, first by Miminis & Paige (1988) and refined later in a fast algorithm by Daniel Kressner, Nicola Mastronadi, Vasile Sima, Paul van Dooren, Sabine van Huffel and others (see, for example, Mastronardi *et al.*, 2001).

4.4 *The computation of Jordan structures and Jordan and Kronecker canonical forms*

Vera Kublanovskaya's work and understanding of QR prepared her to use orthogonal transformations such as Householder reflections and Givens rotations to solve the Jordan normal form and general matrix pencil problems numerically. Her first paper extending QR is Kublanovskaya (1968). It describes an orthogonal matrix algorithm to compute the Jordan normal form of a matrix with defective repeated eigenvalues. Further extensions of QR by Vera and co-authors are to compute the Kronecker normal form and Jordan and Kronecker indices of matrix pencils, both linear and polynomial, as well as numerically solve generalized eigenvalue problems for regular and singular pencils, λ-matrices and spectral problems for matrices with polynomial parameters in general. Vera's work on spectral problems of algebra was done throughout her career in Leningrad or St Petersburg, often with students and co-authors. A review of these results appeared in 1988/1989 in a series of three articles by Vera, Vladimir Khazanov and Vladimir Bely on 'Spectral problems for matrix pencils, I, II, III' (Kublanovskaya & Khazanov, 1988a,b; Kublanovskaya *et al.*, 1989). The ideas and algorithms of Vera and her co-workers have found further extensions in Sweden through the work of Axel Ruhe, Bo Kågström, their students and co-authors, as well as by many others. We just mention 'An algorithm for the numerical computation of the Jordan normal form of a complex matrix' of Kågström & Ruhe (1980). These extensions of QR are far reaching and regularly used in control theory and its applications.

4.5 *The nature and structure of generalized orthogonalities*

For a given square matrix $A \in \mathbb{C}^{n,n}$, the QR algorithm finds an orthonormal basis and an associated orthogonal matrix U for A so that U^*AU is in numerical (block) triangular Schur normal form. Inside QR, the matrix A is reduced to its Schur normal form via a sequence of Householder reflections $I_n - 2v \cdot v^*$ for $v^* \cdot v = 1$ and Givens planar rotations, named after Wallace Givens (1910–1993). In algebraic group theory, Householder reflections are called symmetries. Their role in generating the group of unitary matrices had been studied for a long time before numerical analysis made use of them for their stability properties.

The fundamental theorem of algebraic groups, for example, states that, over any field $\mathbb{F}$ with char $(\mathbb{F}) \neq 2$, if $S = S^{\mathrm{T}} \in \mathbb{F}^{n,n}$ is nonsingular and symmetric then, for every S-orthogonal matrix $U \in \mathbb{F}^{n,n}$ (with $U^{\mathrm{T}}SU = S$), the minimal number of factors to express U as the product of generalized

S-Householder matrices $H_v = I_n - 2vv^\mathrm{T}S$ for $v^\mathrm{T}Sv = 1$ is $n - \dim(\ker(U - I_n))$, unless $S(U - I_n)$ is skew-symmetric. If $S(U - I_n)$ is skew-symmetric then the minimal number of S-Householder factors of U is $n - \dim(\ker(U - I_n)) + 2$.

This fundamental result dates back to Élie Joseph Cartan (1869–1951) in 1938 (Cartan, 1938) for $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$ with the bound of at most n factors. Jean Alexandre Eugène Dieudonné (1906–1992) extended Cartan's result in 1948 to hold for arbitrary fields $\mathbb{F}$ of characteristic not two (Dieudonné, 1948). And finally, Peter Scherk (1910–1985) established the minimal number of factors in 1950 using an existence proof (Scherk, 1950).

One intriguing question in this realm is to learn how to factor a unitary matrix $U \in \mathbb{C}^{n,n}$ numerically. What can the Householder vectors v_i that define H_{v_i} tell us about $U = \prod H_{v_i}$? This open problem of finding a constructive proof for the above Cartan–Dieudonné–Scherk theorem was posed as one of the Challenge Problems in *Linear Algebra and its Applications* (Uhlig, 2001) together with partial results.

4.6 *High-accuracy symmetric QR and SVD algorithms and their implementation*

Over the last decades, much progress has been made to find the eigenvalues and/or singular values of certain matrices accurately in both the high absolute and high relative sense. In the 1960s Kahan observed that a bidiagonal matrix determines all its singular values to high relative accuracy, however small they may be. No notice was taken of this result and it was not until 1980 that he and Jim Demmel published a paper that included Kahan's early result and the observation that, with zero shift, the bidiagonal QR algorithm preserves high relative accuracy (Demmel & Kahan, 1990). Their suggestion was to use a zero shift until all the small singular values were found and then revert to the standard shift. No mention was made of Rutishauser's qd algorithm. Here q and d refer to 'quotient' and 'difference' and the lower case letters indicate that q and d are vectors, not matrices. Fernando saw the intimate connection with qd first. A new variation of Rutishauser's qd algorithm, called dqds, was needed to preserve high relative accuracy, and in dqds shifts may be used provided that they fall below the smallest singular value. Fernando and Parlett published their results in 1994 (Fernando & Parlett, 1994). This is now the preferred way to compute singular values of bidiagonals. The dqds algorithm permits a parallel implementation quite naturally. Parlett and his student Inderjit Dhillon used some of these ideas in their algorithm to compute eigenvectors of symmetric tridiagonals with high accuracy (Dhillon, 1997; Dhillon *et al.*, 1997). Demmel and his co-authors, on the other hand, have worked mainly with accurate QR-type implementations for eigenvalues of symmetric tridiagonals and for bidiagonals, i.e., singular values (Deift *et al.*, 1991; Demmel, 1992). Their results invoke rank revealing decompositions and certain displacement rank classes of matrices. Both schools' success depends on the interpretation and numerical quality of implementation for Francis's QR and Rutishauser's qd algorithm, respectively.

4.7 *Parallelizing the QR and QZ algorithms*

Parallel implementations of the QR factorization of a matrix via Householder reflections or Givens rotations have been studied since the late 1970s. Their efficiency naturally depends on the computer architecture. To parallelize the QR factorizations is very desirable on parallel computers for speed-up inside QR or QZ matrix eigenvalue codes, especially today as multithreaded CPUs with two to eight cores are becoming standard fare on laptops, Macs and PCs. Various approaches involve blocking of the given matrix, look-ahead strategies, load-balancing schemes, grain aggregation, pipelining of iterations or dimensional analysis. Associated with these efforts at QR parallelization are, for example, Greg Henry, Ilan Bar-On, Robert A. van de Greijn, Daniel Boley, Erik Elmroth, Bo Kågström, Fred G. Gustavson, Robert W. Numrich, Daniel Kressner and their respective co-authors (Boley *et al.*, 1989;

Bar-On & Codenotti, 1995; Henry & van de Geijn, 1996; Elmroth & Gustavson, 2000, 2001; Henry *et al.*, 2002; Gunter & van de Geijn, 2005; Adlerborn *et al.*, 2008; Numrich, 2008).

5. A short historical sketch of Rutishauser's LR method

Both John Francis and Vera Kublanovskaya started independently to develop QR around 1958/1959. Their efforts were admittedly based on Rutishauser's Bureau of Standards LR algorithm paper (Rutishauser, 1958) of 1958.

However, the LR method did not originate in 1958, but well before in 1954 or 1955 (see Rutishauser, 1955; Rutishauser & Bauer, 1955).

Eduard Stiefel (1909–1978) and Heinz Rutishauser (1918–1970) had developed the qd algorithm, or the 'Quotienten-Differenzen Algorithmus' in German, to find the zeros of polynomials described in the power basis (see Stiefel, 1953; Rutishauser, 1954a). Their aim was to compute the eigenvalues of a matrix A from the Schwarz constants $s_k = x_0^T \cdot A^k \cdot y_0$ of A (see Schwarz, 1885) by using qd and thereby avoid forming the characteristic polynomial of A. Inside qd the qd constants are manipulated in a quotient difference scheme that is reminiscent of the Neville–Aitken *et al.* polynomial interpolation schemes of old. The qd method is related to earlier work by Cornelius Lanczos (1893–1974) (Lanczos, 1950) and to the conjugate gradient algorithm of Magnus Hestenes (1906–1991) and Stiefel of 1952 (Hestenes & Stiefel, 1952). It works with the data stored in a vector $W = (q_1, e_1, q_2, e_2, \ldots, q_{n-1}, e_{n-1}, q_n)$ and concerns continued fractions. It was further developed by Rutishauser (1954b). He interpreted qd in Rutishauser (1955) for the first time via matrices, and he noted that its rhombus iteration rules $\hat{e}_k + \hat{q}_{k+1} = q_{k+1} + e_{k+1}$ and $\hat{e}_k \hat{q}_k = q_{k+1} e_k$ can be translated into lower–upper matrix factorizations and reverse-order multiplications applied to the bidiagonal matrices

$$
L = \begin{pmatrix} 1 & & & \\ e_1 & 1 & & \\ & \ddots & \ddots & \\ & & e_{n-1} & 1 \end{pmatrix} \quad \text{and} \quad U = \begin{pmatrix} q_1 & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & q_n \end{pmatrix}
$$

formed from W. The results were the preliminary reports (Rutishauser, 1955; Rutishauser & Bauer, 1955) with the earliest LR-type algorithms of factoring and reverse-order matrix multiplication to obtain matrix eigenvalues. These 1955 papers developed a precursor to LR for positive-definite real symmetric matrices. If $A = A^T \in \mathbb{R}^{n,n}$ is positive definite, Rutishauser (1955) showed that the matrices A_k in the Choleski factorization and reverse-order multiplication sequence $A = A_1 = L_1 \cdot L_1^T$, $A_2 = L_1^T \cdot L_1 = L_2 \cdot L_2^T$, etc. are symmetric positive definite and similar to A. Moreover, the A_k converge to a diagonal matrix that contains the eigenvalues of A. An improved version by Rutishauser and Friedrich Ludwig Bauer that uses the square-root-free Choleski factorization $A = L \cdot D \cdot L^T$ instead with a unit lower triangular L is described in Rutishauser & Bauer (1955). In 1956 Rutishauser completed a 51-page internal ETH report in English that contains most of the theoretical and implementation results of Rutishauser (1958). He was well aware of the convergence to two-block upper triangular form of the LR algorithm for real matrices with complex eigenvalues (see Rutishauser, 1958, pp. 55, 56 and Theorem 6), as well as of the invariance of banded matrix structures inside LR. But Rutishauser was unaware of preliminary reductions to Hessenberg form. In fact, he acknowledged in Rutishauser (1958, p. 48) that 'the method seems very time-consuming on first sight', which is true when applied to dense matrices. First ideas about shifts occurred in Rutishauser (1958, pp. 71/73 and Section 11). Many low-dimensional classical matrices, as well as two symmetric banded matrices that result from discretizations of DEs were tested in Rutishauser

(1958). The latter have dimensions $n = 50$ and $n = 89$ and bandwidth seven or three, respectively. The testing and implementation of LR took place on the ERMETH electronic computer at the ETH in Zürich. The lowest five eigenvalues of tridiagonal symmetric 20×20 matrices, for example, were computed in 30 min via LR on ERMETH. Shifts for convergence acceleration were studied further by Rutishauser (1960) for the LR algorithm.

Interestingly enough, during my visit with John Francis in July 2008, John tried to reconstruct how the Schwarz constants (Schwarz, 1885), named after Hermann Amadeus Schwarz (1843–1921), relate to his work with QR because he had worked with them just before that time and had implemented a Schwarz constants-based eigenvalue algorithm on Pegasus. But he could not recall the details. Had he been aware of qd and the roots of Rutishauser's LR papers in some (sub-)conscious way?

Acknowledgements

References

(1916–1975) *Catalogue of the Papers and Correspondence of Christopher Strachey*. Reference NCUACS 71.1.80 (Administrative history). Bodleian Library, Special Collections, Oxford University. Available at http://www.nationalarchives.gov.uk/a2a/records.aspx?cat =161-csac71180.

Editorial. (2000). *Higher-Order Symbolic Comput.*, **13**, 5–6.

BJÖRN ADLERBORN, BO KÅGSTRÖM & DANIEL KRESSNER (2008) Parallel variants of the multishift QZ algorithm with advanced deflation techniques. *Applied Parallel Computing. State of the Art in Scientific Computing*, vol. 4699. Lecture Notes in Computer Science. Springer, pp. 117–126.

ILAN BAR-ON & BRUNO CODENOTTI (1995) A fast and stable parallel QR algorithm for symmetric tridiagonal matrices. *Linear Algebra Appl.*, **220**, 63–95.

PETER BENNER, HEIKE FASSBENDER & DAVID S. WATKINS (1998) Two connections between the SR and HR eigenvalue algorithms. *Linear Algebra Appl.*, **272**, 17–32.

PETER BENNER, HEIKE FASSBENDER & DAVID S. WATKINS (1999) SR and SZ algorithms for the symplectic (butterfly) eigenproblem. *Linear Algebra Appl.*, **287**, 41–76.

DAVID BINDEL, SHIV CHANDRASEKARAN, JAMES DEMMEL, DAVID GARMIRE & MING GU (2005) A fast and stable nonsymmetric eigensolver for certain structured matrices. *Technical Report*. Berkeley, CA: University of California.

DARIO BINI, YULI EIDELMAN, LUCA GEMIGNANI & ISRAEL GOHBERG (2007) Fast QR eigenvalue algorithms for Hessenberg matrices which are rank-one perturbations of unitary matrices. *SIAM J. Matrix Anal. Appl.*, **29**, 566–585.

DARIO BINI, LUCA GEMIGNANI & VICTOR PAN (2003) QR-like algorithms for generalized semiseparable matrices. *Technical Report 1470*. Pisa: Department of Mathematics, University of Pisa.

DARIO BINI, LUCA GEMIGNANI & VICTOR PAN (2005) Fast and stable QR eigenvalue algorithms for generalized companion matrices and secular equations. *Numer. Math.*, **100**, 373–408.

DANIEL BOLEY, ROBERT MAIER & JOUNG KIM (1989) A parallel QR algorithm for the nonsymmetric eigenvalue problem. *Comput. Phys. Commun.*, **53**, 61–70.

KAREN BRAMAN, RALPH BYERS & ROY MATHIAS (2002a) The multishift QR algorithm. Part I: maintaining well-focused shifts and level 3 performance. *SIAM J. Matrix Anal. Appl.*, **23**, 929–947.

KAREN BRAMAN, RALPH BYERS & ROY MATHIAS (2002b) The multishift QR algorithm. Part II: aggressive early deflation. *SIAM J. Matrix Anal. Appl.*, **23**, 948–973.

ANGELIKA BUNSE-GERSTNER (1981) An analysis of the HR algorithm for computing the eigenvalues of a matrix. *Linear Algebra Appl.*, **35**, 155–173.

ANGELIKA BUNSE-GERSTNER (1986) Symplectic QR-like methods. *Habilitationsschrift*, University of Bielefeld, Germany.

ANGELIKA BUNSE-GERSTNER, RALPH BYERS & VOLKER MEHRMANN (1992) A chart of numerical methods for structured eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, **13**, 419–453.

ANGELIKA BUNSE-GERSTNER & VOLKER MEHRMANN (1986) A symplectic QR like algorithm for the solution of the real algebraic Ricatti equation. *IEEE Trans. Autom. Control*, **31**, 1104–1113.

ANGELIKA BUNSE-GERSTNER, VOLKER MEHRMANN & DAVID S. WATKINS (1989) An SR algorithm for Hamiltonian matrices based on Gaussian elimination. *Methods Oper. Res.*, **58**, 339–358.

HENDRIK JAN BUUREMA (1970) A geometric proof of convergence for the QR method. *Ph.D. Thesis*, Rijksuniversiteit te Groningen, The Netherlands.

RALPH BYERS (1986) A Hamiltonian QR algorithm. *SIAM J. Sci. Stat. Comput.*, **7**, 212–229.

ÉLIE JOSEPH CARTAN (1938) *La Théorie des Spineurs I, II.* Actualités Scientifiques et Industrielles, nos 643 and 701. Paris: Hermann.

SHIV CHANDRESEKARAN, MING GU, JIANLIN XIA & JIANG ZHU (2008) A fast QR algorithm for companion matrices. *Recent Advances in Matrix and Operator Theory*. Operator Theory: Advances and Applications, vol. 179. Basel: Birkhäuser, pp. 111–143.

BARRY A. CIPRA (2000) The best of the 20th century: editors name top 10 algorithms. *SIAM News*, **33**. Available at http://www.siam.org/news/news.php?id=637.

PERCY DEIFT, JAMES DEMMEL, CHAU LI & CARLOS TOMEI (1991) The bidiagonal singular value decomposition and Hamiltonian mechanics. *SIAM J. Numer. Anal.*, **28**, 1463–1516.

STEVEN DELVAUX & MARC VAN BAREL (2008) A QR-based solver for rank structured matrices. *SIAM J. Matrix Anal. Appl.*, **30**, 464–490.

JAMES DEMMEL (1992) *The Inherent Inaccuracy of Implicit Tridiagonal QR*, vol. 963. Minneapolis, MN: IMA Preprint Series, Institute for Mathematics and its Applications, University of Minnesota.

JAMES DEMMEL & WILLIAM KAHAN (1990) Accurate singular values of bidiagonal matrices. *SIAM J. Sci. Stat. Comput.*, **11**, 873–912.

INDERJIT S. DHILLON (1997) A new $O(n^2)$ algorithm for the symmetric tridiagonal eigenvalue/eigenvector problem. *Ph.D. Thesis*, University of California, Berkeley, CA.

INDERJIT S. DHILLON, GEORGE FANN & BERESFORD N. PARLETT (1997) Application of a new algorithm for the symmetric eigenproblem to computational quantum chemistry. *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing.* Minneapolis.

JEAN DIEUDONNÉ (1948) *Sur les Groupes Classiques.* Actualités Scientifiques et Industrielles, no. 1040. Paris: Hermann.

JACK DONGARRA & FRANCIS SULLIVAN (2000) The top 10 algorithms. *Comput. Sci. Eng.*, **2**, 22–23.

YULI EIDELMAN, LUCA GEMIGNANI & ISRAEL GOHBERG (2008) Efficient eigenvalue computation for quasiseparable Hermitian matrices under low rank perturbations. *Numer. Algorithms*, **47**, 253–273.

YULI EIDELMAN, ISRAEL GOHBERG & VADIM OLSHEVSKY (2005) The QR iteration method for Hermitian quasiseparable matrices of an arbitrary order. *Linear Algebra Appl.*, **404**, 305–324.

ERIK ELMROTH & FRED GUSTAVSON (2000) Applying recursion to serial and parallel QR factorization leads to better performance. *IBM J. Res. Dev.*, **44**, 605–624.

ERIK ELMROTH & FRED GUSTAVSON (2001) High-performance library software for QR factorization. *Applied Parallel Computing. New Paradigms for HPC in Industry and Academia*. Lecture Notes in Computer Science, vol. 1947. Springer, pp. 53–63.

SAID ELNASHAIE, FRANK UHLIG & CHADIA AFFANE (2007) *Numerical Techniques and MATLAB Codes for Chemical and Biological Engineers, A Simple Bifurcation Approach*. Springer.

DIMITRI K. FADDEEV & VERA NIKOLAEVNA FADDEEVA (1963) *Computational Methods of Linear Algebra*, vol. 28 (in Russian) pp. 734. [MR 061454, #666897].

HEIKE FASSBENDER (2000) *Symplectic Methods for the Symplectic Eigenproblem*. New York: Kluwer.

HEIKE FASSBENDER (2006) A detailed derivation of the parameterized SR algorithm and the symplectic Lanczos method for Hamiltonian matrices. *Technical Report*. Germany: Institut Computational Mathematics, TU Braunschweig, Germany.

HEIKE FASSBENDER (2007) The parameterized SR algorithm for Hamiltonian matrices. *ETNA*, **26**, 121–145.

K. VINCE FERNANDO & BERESFORD N. PARLETT (1994) Accurate singular values and differential qd algorithms. *Numer. Math.*, **67**, 191–229.

CHARLES T. FIKE (1959) Note on the practical computation of proper values. *J. Assoc. Comput. Mach.*, **6**, 360–362.

JOHN G. F. FRANCIS (1961/1962a) The QR transformation, a unitary analogue to the LR transformation—part 1. *Comput. J.*, **4**, 265–271. [MR 0130111, vol. 23 #B3143], [Zbl 0104.34304].

JOHN G. F. FRANCIS (1961/1962b) The QR transformation—part 2. *Comput. J.*, **4**, 332–345. [MR 0137289, vol. 25 #744], [Zbl 0104.34304].

LUCA GEMIGNANI (1999) Fast QR factorization of low-rank changes of Vandermonde-like matrices. *CALCOLO*, **36**, 1–15.

LUCA GEMIGNANI & FRANK UHLIG (2009) Numerical behavior of the DQR method for rank-structured matrices (submitted).

GENE GOLUB, BO KÅGSTRÖM, AXEL RUHE & PAUL VAN DOREN (1990) Dedication to Vera N. Kublanovskaya on her 70th birthday. *SIAM J. Matrix Anal. Appl.*, **11**, vii–ix.

GENE GOLUB & WILLIAM KAHAN (1965) Calculating the singular values and pseudo-inverse of a matrix. *J. SIAM Numer. Anal. Ser. B*, **2**, 205–224.

MIKE GORDON (2000) Christopher Strachey: recollection of his influence. *Higher-Order Symbolic Comput.*, **13**, 65–67.

BRIAN GUNTER & ROBERT VAN DE GEIJN (2005) Parallel out-of-core computation and updating of the QR factorization. *ACM Trans Math. Softw.*, **31**, 60–78.

GREG HENRY & ROBERT VAN DE GEIJN (1996) Parallelizing the QR algorithm for the unsymmetric algebraic eigenvalue problem: myths and reality. *SIAM J. Sci. Comput.*, **17**, 870–883.

GREG HENRY, DAVID WATKINS & JACK DONGARRA (2002) A parallel implementation of the nonsymmetric QR algorithm for distributed memory architectures. *SIAM J. Sci. Comput.*, **24**, 284–311.

MAGNUS R. HESTENES & EDUARD STIEFEL (1952) Methods of conjugate gradients for solving linear systems. *J. Res. Nat. Bur. Stand.*, **49**, 409–436.

NICHOLAS J. HIGHAM (2003) SIAG/LA prize winners speed up QR algorithm. *SIAM News*, **36**.

ALSTON S. HOUSEHOLDER (1964) *The Theory of Matrices in Numerical Analysis*. New York: Blaisdell.

BO KÅGSTRÖM & AXEL RUHE (1980) An algorithm for the numerical computation of the Jordan normal form of a complex matrix. *ACM Trans. Math. Softw.*, **6**, 398–419.

T. YA. KONKOVA, V. N. SIMONOVA & VLADIMIR B. KHAZANOV (2003) Vera Nikolaevna Kublanovskaya. Short biography. *J. Math. Sci.*, **114**, 1755–1756.

DANIEL KRESSNER (2005) *Numerical Methods for General and Structured Eigenvalue Problems*. Springer.

VERA N. KUBLANOVSKAYA (1961a) Certain algorithms for the solution of the complete eigenvalue problem. *Dokl. Akad. Nauk. SSSR*, **136**, 26–28 (in Russian). English translation in *Soviet Math. Dokl.*, **2** (1961), 17–19. [MR 0142195, vol. 25, #5588], [Zbl 0104.10102].

VERA N. KUBLANOVSKAYA (1961b) Some algorithms for the solution of the complete eigenvalue problem. *Z. Vycisl. Mat. i Mat. Fiz.*, **1**, 555–570 (in Russian). [MR0136058, vol. 24, #2097], [Zbl 0128.11702].

VERA N. KUBLANOVSKAYA (1962) The solution of the eigenvalue problem for an arbitrary matrix. *Trudy Mat. Inst. Steklov.*, **66**, 113–135 (in Russian). [MR0150940, vol. 27, #B929], [Zbl 0131.14205].

VERA N. KUBLANOVSKAYA (1968) On a method for solving the complete eigenvalue problem for a degenerate matrix. *Z. Vycisl. Mat. i Mat. Fiz.*, **6**, 611–620 (in Russian). English translation in *USSR Comput. Math. Math. Phys.*, **6** (1966), 1–14. [Zbl 0171.36003].

VERA N. KUBLANOVSKAYA & VLADAMIR B. KHAZANOV (1988a) Spectral problems for matrix pencils. Methods and algorithms. I. *Sov. J. Numer. Anal. Math. Model.*, **3**, 337–371 (in Russian).

VERA N. KUBLANOVSKAYA & VLADAMIR B. KHAZANOV (1988b) Spectral problems for matrix pencils. Methods and algorithms. II. *Sov. J. Numer. Anal. Math. Model.*, **3**, 467–485 (in Russian).

VERA N. KUBLANOVSKAYA, VLADAMIR B. KHAZANOV & VLADAMIR A. BELIY (1989) Spectral problems for matrix pencils. Methods and algorithms. III. *Sov. J. Numer. Anal. Math. Model.*, **4**, 19–51 (in Russian).

CORNELIUS LANCZOS (1950) An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Stand.*, **45**, 255–282.

NICOLA MASTRONARDI, DANIEL KRESSNER, VASILE SIMA, PAUL VAN DOOREN & SABINE VAN HUFFEL (2001) A fast algorithm for subspace state-space system identification via exploitation of the displacement structure. *J. Comput. Appl. Math.*, **132**, 71–81.

GEORGE S. MIMINIS & CHRIS C. PAIGE (1988) A direct algorithm for pole assignment of time invariant multi-input linear systems using state feedback. *Automatica*, **24**, 343–356.

CLEVE B. MOLER & G. W. STEWART (1973) An algorithm for generalized matrix eigenvalue problems. *SIAM J. Numer. Anal.*, **10**, 241–256.

ROBERT W. NUMRICH (2008) Dimensional analysis applied to a parallel QR algorithm. *Parallel Processing and Applied Mathematics*. Lecture Notes in Computer Science, vol. 4967. Springer, pp. 148–155.

BERESFORD PARLETT (2000) The QR algorithm. *Comput. Sci. Eng.*, **2**, 38–42.

ROGER PENROSE (2000) Reminiscences of Christopher Strachey. *Higher-Order And Symbolic Computation*, **13**, 83–84.

HEINZ RUTISHAUSER (1954a) Der Quotienten-Differenzen-Algorithmus. *Z. Angew. Math. Phys.*, **5**, 233–251 (in German). [MR0063763 (16,176c)], [Zbl 0055.34702].

HEINZ RUTISHAUSER (1954b) Anwendungen des Quotienten-Differenzen-Algorithmus. *Z. Angew. Math. Phys.*, **5**, 496–508 (in German). [MR0068314 (16,863c)], [Zbl 0056.35001].

HEINZ RUTISHAUSER (1955) Une méthode pour la détermination des valeurs propres d'une matrice. *C. R. Acad. Sci.*, **240**, 34–36 (in French). [MR0067844 (16,785c)], [Zbl 0056.35002].

HEINZ RUTISHAUSER (1958) Solution of eigenvalue problems with the LR-transformation. *Nat. Bur. Stand. Appl. Math. Ser.*, **49**, 47–81. [MR0090118 (19,770c)], [Zbl 0123.11303].

HEINZ RUTISHAUSER (1960) Über eine kubisch konvergente Variante der LR-Transformation. *Z. Angew. Math. Mech.*, **40**, 49–54 (in German). [MR0114295 (22 #5119)], [Zbl 0128.11701].

HEINZ RUTISHAUSER & FRIEDRICH L. BAUER (1955) Détermination des vecteurs propres d'une matrice par une méthode itérative avec convergence quadratique. *C. R. Acad. Sci.*, **240**, 1680–1681 (in French). [MR0070269 (16,1157a)].

PETER SCHERK (1950) On the decomposition of orthogonalities into symmetries. *Proc. Am. Math. Soc.*, **1**, 481–491.

HERMANN AMADEUS SCHWARZ (1885) Über ein die Flächen kleinsten Flächeninhalts betreffendes Problem der Variationsrechnung. *Acta Soc. Sci. Fennicae*, **15**, 315–362.

G. W. STEWART (1973) *Introduction to Matrix Computations*. New York: Academic Press.

EDUARD STIEFEL (1953) Zur Interpolation von tabellierten Funktionen durch Exponentialsummen und zur Berechnung von Eigenwerten aus den Schwarzschen Konstanten. *Z. Angew. Math. Mech.*, **33**, 260–262 (in German). [Zbl 0051.09604].

CHRISTOPHER STRACHEY & JOHN G. F. FRANCIS (1961/1962) The reduction of a matrix to codiagonal form by eliminations. *Comput. J.*, **4**, 168–176. [MR0129125 (23 #B2162)], [Zbl 0103.10005].

LLOYD N. TREFETHEN & DAVID BAU III (1997) *Numerical Linear Algebra*. Philadelphia, PA: SIAM.

FRANK UHLIG (1997) The DQR algorithm, basic theory, convergence and conditional stability. *Numer. Math.*, **76**, 515–553.

FRANK UHLIG (1999) General polynomial roots and their multiplicities in $O(n)$ memory and $O(n^2)$ time. *Linear Multilinear Algebra*, **46**, 327–359.

FRANK UHLIG (2001) Constructive ways for generating (generalized) real orthogonal matrices as products of (generalized) symmetries. *Linear Algebra Appl.*, **332–334**, 459–467.

CHARLES VAN LOAN (1984) A symplectic method for computing all the eigenvalues of a Hamiltonian matrix. *J. Linear Algebra Appl.*, **61**, 233–251.

RAF VANDEBRIL, MARC VAN BAREL & NICOLA MASTRONARDI (2005) An Implicit QR algorithm for semiseparable matrices. *Numer. Linear Algebra Appl.*, **12**, 625–658.

DAVID S. WATKINS (1982) Understanding the QR algorithm. *SIAM Rev.*, **24**, 427–440.

DAVID S. WATKINS (2008) The QR algorithm revisited. *SIAM Rev.*, **50**, 133–145.

HELMUT WIELANDT (1944) Beiträge zur mathematischen Behandlung komplexer Eigenwertprobleme, Teil V: Bestimmung höherer Eigenwerte durch gebrochene Iteration. *Berichte Aerodynamische Versuchsanstalt Göttingen*, B 44/J/37.

JAMES H. WILKINSON (1959) Stability of the reduction of a matrix to almost triangular and triangular forms by elementary similarity transformations. *J. Assoc. Comput. Mach.*, **6**, 336–359.

JAMES H. WILKINSON (1965) *The Algebraic Eigenvalue Problem*. Oxford: Clarendon Press.

In Memoriam
Gene Howard Golub 1932–2007

In Moscow, July 2007

Dedicated to the authors' mothers who both were born in Latvia and grew up there nearly a century ago.