Research Review 1968-97

My research in numerical analysis started end of 1968 with a Master thesis ("Diplomarbeit") on error bounds for zeros of polynomials computed successively in finite-precision arithmetic using deflation with the Horner scheme. This Master thesis was supervised by Peter Henrici, who in the sequel also became my PhD thesis advisor. It lead to my first two publications[1] [1, 2], and I was proud that the more serious one, finished later with some help of Brian Smith, was accepted by *Numer. Math.*

After I had finished this Master thesis, Peter Henrici gave me an article on digital filters and suggested that I find a PhD topic in this area. Then he went to the US for a sabbatical, and, hence, from then on I worked mainly on my own on my dissertation, in which I worked out an algorithm for complex rational L^∞ approximation, a topic that was algorithmically completely untouched at that time, but also quite hard. I defined my algorithm in a very abstract setting, and for the application to digital filters I had to specify the nonlinear family of approximants as set of stable rational functions and the approximation domain as unit circle. The PhD thesis [3], finished in 1973, got published in 250 copies, but did not give rise to further publications. The fact is that, right after my thesis, Peter Henrici wanted me to work on the posthumous publication of some of Heinz Rutishauser's lectures. The preparation of the resulting two volumes (which became one in the English translation by Walter Gautschi) took me one and a half years.

Meanwhile, Peter Henrici had developed much interest in numerical conformal mapping. A number of Master and PhD students experimented with the various methods known at that time, most of which turned out to be very restricted in their application, that is they converged only for very limited classes of domains sufficiently fast. Theodorsen's method, one that had been widely applied in airfoil design (where a suitable auxiliary mapping improved the convergence crucially, however), also turned out to be very limited. But using underrelaxed nonlinear Jacobi, second order Richardson, and SOR iterations I was able to improve the convergence domain and speed dramatically. This work lead to my Habilitation thesis [8], which was published in three long journal articles [9, 10, 32] and to two shorter papers [7, 11].

At the same time I continued my work in complex approximation theory. On the one hand, I wrote a long, unfinished manuscript on a numerical method improving the one of my thesis. Only a brief sketch of it got published in the proceedings paper [6]. On the other hand, I got involved in theoretical questions. First, I pointed out that the best polynomial approximant needs not be so-called strongly unique [5]. Five years later, Lloyd N. Trefethen and I discovered that, even on the unit disk, rational approximants need not be unique [22, 24] and that, on an interval, complex approximants can be arbitrarily much closer to a real function than real ones [23,

[1] The numbers refer to those on my list of publications.

24], surprising results that were well received by the community. The second one was generalized in [27].

The main topic of my extremely successful collaboration with Nick Trefethen was what we called Carathéodory-Fejér approximation [13, 14, 18, 21, 29, 38], a constructive method of near-best approximation that is closely related to what engineers call H^∞ approximation or approximation in the Hankel norm, and which, as we learned soon after Nick had finished his first, single-authored paper, can be based on a remarkable theory due to Adamian, Arov, and Krein. Very early we applied this technique to digital filter design [12], but it took us a long time (and some help from Julius O. Smith) to adjust this paper to engineering standards and to provide more relevant examples [20]. We also quickly adjusted the method to the polynomial and rational approximation problem on an interval. With Steve W. Ellacott I solved a special Carathéodory-Fejér problem [17] and, using Faber polynomials, generalized the Carathéodory-Fejér approach to other regions than the disk and the interval [16]. Later, with Ed Saff I adapted the de Montessus theorem of Padé approximation to Carathéodory-Fejér approximation [33]. Many of the results on complex approximation and the Carathéodory-Fejér method are summarized in the two long papers [21]. Also touched there is the fact that nontrivial cases of algebraically computable best L^∞ approximations on the disk and the interval are actually Carathéodory-Fejér approximations with no truncation part, that is, with exactly circular error curve in case of the disk, and a projection of such a curve in case of an interval [25]. A long survey paper on this is still unfinished.

Also with Nick Trefethen I looked at the continuity of Chebyshev and Padé approximants of given numerator and denominator degree with respect to various metrics [26]. In [34], we introduced stable Padé approximants and a new construction for Chebyshev-Padé approximants.

The paper [19], dedicated to Peter Henrici's on the occasion of his 60th birthday, shows how to compute efficiently the imaginary part of a rational function on the unit circle from its real part and has several potential applications, including methods for conformal mapping. A general framework and a survey of old and new such methods based on function conjugation is given in [28].

Techniques I had learned and applied in my early work on conformal mapping [8, 9, 10] were the starting point to a series of papers on iterative k-step and (k, l)-step methods for linear and nonlinear systems of equations: [30](with André Kaiser), [31](with Wilhelm Niethammer and Richard S. Varga), [35](which is connected to Pick-Nevanlinna interpolation), [40], [42], [49]; the most important of these is [40].

For a conference in Princeton, 1987, where many old pioneers of numerical analysis gathered, I wrote a paper on how Eduard Stiefel founded 1948 the Institute of Applied Mathematics at ETH Zurich, how he acquired the Zuse Z4 computer and kept it operational till 1955, and how he developed and built with Heinz Rutishauser and Ambros Speiser a Swiss electronic computer called ERMETH [37]. This paper was supposed to be my first joint paper with Peter Henrici, but he was too ill to contribute and died before the conference.

At this time I got more and more interested into the construction of rational interpolants. First, I considered the characterization of these interpolants, considering them as meromorphic functions defined on the extended plane, with ∞ as a potential interpolation point [39]. In [43], I extended the continuity results of [26] to rational interpolation. Then, I improved an algorithm due to Helmut Werner, making it fully recursive and introducing corresponding generalizations of Thiele continued fractions [41]. This paper inseminated much of my subsequent work on rational interpolation and the Lanczos algorithm. The recursive construction I proposed overcame breakdowns of other recursive algorithms for rational interpolation; however, it did not yet overcome near-breakdowns. The paper provided also the background for my joint work with Gene H. Golub on constructing formal orthogonal polynomials from modified moments with respect to an indefinite weight function [44]. In this work, I learned from Gene about the connection between moments and the Lanczos algorithm, two of his most favorite topics. On the other hand, I knew that my generalized Thiele fractions of [41] were special cases of so-called P-fractions (a favorite topic of Bill Gragg), and, studying once again Rutishauser's booklet on the QD algorithm, I realized that all fits together: an extension of the Lanczos algorithm that overcomes exact breakdowns could be defined in analogy to the construction of P-fractions (actually, Bill Gragg had found this some 15 years earlier, but never written down the details), formal orthogonal polynomials (FOPs), and an extension of Rutishauser's progressive QD algorithm that also overcomes exact breakdowns.

In parallel, I wrote up this theory for the generic case in which no breakdowns occur [47] and for the nongeneric case with exact breakdowns [45, 46], a task that took me more than two years (since, at the same time I was just building up the Interdisciplinary Project Center for Supercomputing (IPS)). In these papers, a variety of Lanczos-type algorithms are put into relation with each other and with the underlying theory of FOPs, Padé approximants, and Hankel matrices. As an invited paper for the first Copper Mountain Conference on Iterative Methods, 1990, [47] was published in the Preliminary Proceedings, but nowhere else. An extra section on various forms of the Conjugate Gradient Squared (CGS) method was distributed separately at the conference. When I had [45] nearly finished, I learned that Dan Boley had also just found a Lanczos algorithm that overcame exact breakdowns and that Beresford Parlett's student Derek R. Taylor had described one in his thesis of 1982. Dan's proposal, along with other material, got included into my paper with Daniel L. Boley, Sylvan Elhay, and Gene H. Golub [48]. Taylor's approach, although he did not address near-breakdowns explicitly, made it easier to include those. For the resulting algorithm that can cope both with exact and near-breakdowns, and for variations of it, Taylor's denomination "look-ahead Lanczos" is therefore used. Such algorithms were described in a new paper by Parlett, in Sections 9 and 10 of [46], and in my joint paper with Roland W. Freund and Noël M. Nachtigal [50], who provided a sophisticated FORTRAN implementation.

When Henk van der Vorst proposed BiCGStab in 1990, I realized quickly that eigenvalues with large imaginary part had no chance of being approximated well. My extension, BiCGStab2, fixed this problem in a simple way [51].

Writing [47] I discovered nice matrix relations between the BiORes, BiOMin, and BiODir versions of the biconjugate gradient (BiCG) method, which also hold for the ORes, OMin, and ODir versions of the conjugate gradient (CG) method. For the latter, these connections must have been known to Rutishauser and Stiefel, since they are exactly reflected in Rutishauser's QD algorithm, but I have not been able to find this in their work. In obscured form, some of these relations can be found in papers of C.C. Paige and M.A. Saunders (1975) and R. Fletcher (1976). With Steven F. Ashby I first reviewed these relations for the versions of CG and CR [55] and then worked out the corresponding facts for ORTHORES, ORTHOMIN, and ORTHODIR, three versions of the generalized conjugate gradient (GCG) and generalized conjugate residual (GCR) methods for nonsymmetric systems of equations; however, the second paper was never completely finished. In concise style, its most important results are also given in [52], along with other material on changing the underlying formal inner product, including a discussion of the residual history smoothing process proposed by Schönauer and Weiss.

At the same time I restarted to investigate recursive algorithms for rational and Padé approximation, also in conjunction with the recursive solution of structured systems of linear equations, in particular Toeplitz systems. This work was on the one hand in the line of [41] and on the other hand influenced by recent work of Stan Cabay and Ron Meleshko, who were the first to come up with a (forward) stable recursive procedure for computing Padé approximants on a (co)diagonal of the Padé table and for solving a real indefinite or complex Hankel system in $\mathcal{O}(N \log^2 N)$ operations. Their algorithm exhibits a close analogy to look-ahead Lanczos algorithms, and therefore, the denomination "look-ahead" is often applied to it and to similar procedures too.

In the invited paper [53] I first discussed the block structure of the Newton-Padé table(a special rational interpolation table) and of a newly introduced similar table, called multipoint Padé table, that included interpolants with $z = \infty$ as a possibly multiple interpolation point; then, I gave general recurrences for the Newton-Padé table, which could be the basis of stable algorithms and which, incidentally, involved entries of the other table. However, establishing the stability of such algorithms is difficult, as is pointed out in a short note with Stan Cabay and Ron Meleshko [57]. In [59], another invited paper, I provided analogous general recurrences for my multipoint Padé table.

However, before, when working on [53], I wrote a paper on the simpler confluent case, that is, on general recurrences in the Padé table instead of the Newton-Padé table. These were analogous to those of Cabay and Meleshko, but not restricted to (co)diagonals. Included were in particular recurrences along rows or adjacent pairs of rows, and these were exactly those needed to generalize the Levinson and Schur algorithms to nonsymmetric Toeplitz matrices (with possibly singular leading principal submatrices). Although I did not prove stability in detail, I am convinced that the look-ahead algorithms presented in this paper were the first forward stable fast algorithms for the indefinite case. Except for cases with more than $\mathcal{O}(1)$ near-singular submatrices, these look-ahead Levinson-like and Schur-like algorithms have still $\mathcal{O}(N^2)$ cost, and, moreover, I also proposed a related superfast look-ahead Schur-like algorithm with $\mathcal{O}(N \log^2 N)$ cost only.

In a series of papers with Marlis Hochbruck [56, 58, 62, 63] these algorithms were modified and further discussed and enhanced. First, in [56], we came up with versions that were true generalizations of the classical Levinson and Schur algorithms, and we described them in a linear algebra setting, not as recurrences for the Padé table. The latter were later given in [62]. In [58] we discussed the stability of the Toeplitz inversion formulas that we applied in some of our algorithms, and in [63] we introduced a version with reduced overhead. In her Habilitation thesis, Marlis was then able to reduce the overhead even further; see her paper in *BIT*.

In my paper with William B. Gragg [60] we reviewed the Cabay-Meleshko algorithm and put it into perspective with related material, in particular the extended Euclidean algorithm and P-fractions; we also pointed out that the Chebyshev algorithm is a special case of the Euclidean algorithm.

For my invited paper at the Lanczos Centenary Conference [61] I reviewed the connections between the Lanczos process and the Padé table. Most of the material was recycled from [47]. For the solicited *Acta Numerica* article [66], I also started from [47], but revised the text completely and added many new sections, in order to provide a complete review of the most important Lanczos-type algorithms for solving linear systems, including some very recent ones. Not included is work I had just finished with Klaus J. Ressel, a look-ahead procedure for Lanczos-type product methods (LTPMs) [63], and a general framework for applying the quasi-minimal residual (QMR) approach to LTPMs [65].

When compiling the material for the survey article [66], I noticed that many interesting questions still remained unsolved, for example, why the gap between recursive and true residuals opens up much earlier in algorithms based on three-term recurrences than in those using coupled two-term recurrences. For the latter type Anne Greenbaum had just analyzed this gap. With Zdeněk Strakoš I was able to analyze the former class and to explain the difference [67].