From *SIAM News*, Volume 37, Number 9, November 2004

Accurate Eigenvalues, Anyone?

The fifth workshop on accurate solutions to eigenproblems, IWASEP5, was held in Hagen, Germany, at the Arcadeon Conference Centre, from June 28 to July 1, 2004. Many of the more than sixty participants were regulars who have attended at least four of these lively biannual gatherings.

The Arcadeon is an attractive meeting place with delicious cuisine (but no fitness centre). We met in a light, airy space that is more like a classroom than a hotel conference room. There was enough room at the back for the posters to be available throughout the workshop, and they elicited small discussions before and after each lecture session. The posters get better with each meeting. This time, three jurors chose the poster by Daniel Kressner, "Structured Condition Numbers for Invariant Subspaces," as the best one.

These workshops have something of the feeling of the early Householder reunions (Gatlinburg was the name used then) that the family, or clan, has come together. This time the make-up of the group, by nation, was Germany 24, USA 10, Croatia 6, Spain 5, and 3 or fewer from other countries. As you might guess, the new faces came mainly from the host country. We were fortunate to have with us, for the second time, Academician Sergei Godunov from Russia, who spoke about his recent work in a talk titled "Spectral Portraits of Symplectic Matrices."

These workshops began in 1996 as a forum for the few people interested in obtaining small eigenvalues to high relative accuracy, when that is warranted. There is no mention of this topic in Wilkinson's *The Algebraic Eigenvalue Problem*, nor in Golub and Van Loan's *Matrix Computations*, nor even in Parlett's *The Symmetric Eigenvalue Problem*. Only with the publication in 1990 of Demmel and Kahan's seminal paper "Accurate Singular Values of Bidiagonal Matrices" (*SIAM Journal on Matrix Analysis and Applications*, Vol. 11, pages 873–912) did the notion that such accuracy can sometimes be attained begin to seep into our collective consciousness. It follows that for a positive definite tridiagonal matrix T with Cholesky factor L, $T = LL'$, L bidiagonal, it is possible to compute each eigenvalue of LL', however small it may be, to high relative accuracy.

Other results soon followed—by Eisenstat/Ipsen and by Ren-Cang Li, and by Veselic and his students on "relative" perturbation theory. Demmel and co-authors worked on methods that could achieve increased accuracy when warranted by the matrix. Dhillon and Parlett showed that, for tridiagonals, small eigenvalues of LDL' can be determined, and computed, to high relative accuracy even when the larger ones are not relatively robust. What a contrast with the general symmetric case, where it is only the large eigenvalues that are determined to high relative accuracy! The next realization was that the cost of this increased accuracy is sometimes low. The new methods are actually faster than the backward stable methods that, for symmetric problems, produce eigenvalues with an error proportional to the norm of the input matrix.

The range of topics has broadened in recent meetings of the clan, but the old preoccupations have not gone away. At this year's meeting, we began with a fascinating talk by Volker Mehrmann on computing the eigenvalues of German trains (see article by Ilse Ipsen on page 1). Zlatko Drmac showed that his carefully crafted one-sided Jacobi method, when applied to a general rectangular matrix, can compute the singular values (possibly to good relative accuracy) even before conventional schemes have completed the reduction to bidiagonal form. This algorithm makes use of techniques devised by Veselic and by Hari. Jesse Barlow presented a different approach, a new one-sided reduction to bidiagonal form (completing work described by Ralha at the third workshop in the series), that might stimulate Drmac to make further refinements to his one-sided Jacobi method.

Describing recent work by himself and by Knyazev, Klaus Neymeyr discussed the task of extending the extensive theory and successful implementation of preconditioned iterative elliptic solvers to the positive definite eig-enproblem. Hubert Schwetlick presented theory for new eigensolvers of the shift-and-invert type that can solve nonlinear eigenproblems and have asymptotic properties superior to those of Jacobi–Davidson methods for linear non-normal problems.

Most, but not all, of the speakers addressed eigenvalue or singular value computations. One exception was Nick Higham, who discussed the computation of $\exp(A)$ for small (MATLAB-sized) matrices A by the venerable "scaling and squaring" (S&S) technique. MATLAB first scales the given A by 2^s, so that $\|A/2^s\| \le 1/2$, then approximates $\exp(A/2^s)$ by $r_6\,(A/2^s)$, where r_m is the $[m/m]$ Padé approximant, and finally squares the approximant s times. Higham asks: Why scale down to 1/2? Why take $m = 6$? The answer to the first question lies in the concern that Padé approximations are good only near the origin.

As to the second question, $m = 6$ is the smallest value that makes an elegant (well-known) error bound for r_m less than 4 roundoff units in MATLAB. Higham bypasses the well-known error bound and considers instead the (known) series expansion of the error in r_m in the form $\rho\,(x) = \exp(-x)\, r_m\,(x) - 1 = O(x^{(2m+1)})$, and the related bound obtained by taking the absolute values of the coefficients. Because these coefficients are quite complicated, Higham does a one-shot symbolic computation, using 250-digit arithmetic (aren't computers wonderful?), to determine, for $m = 1{:}20$, the largest value ω_m such that the bound on the error in $r_m(A/2^s)$ does not exceed MATLAB's roundoff unit. These ω values increase steadily, from 0.085 for $m = 4$ to 13.0 for $m = 20$. Of course, larger values of m require more matrix multiplications (matmuls, for short) to evaluate r_m but need fewer scalings (smaller s), so there is a trade-off. Lo and behold, the normalized cost in matmuls with $m = 6$ is 4.9, while the optimal cost is 3.6 with $m = 13$ and $\omega_{13} = 5.4$. Until $\|A\|$ exceeds 5.4, no scaling is needed!

Higham also tightens the 1977 bounds of R.C.Ward to show that the rounding error in evaluating r_m is dominated by the truncation error in r_m. The resulting algorithm requires two to three fewer matmuls than MATLAB's exp$m$. Higham also discussed other recent methods and ways to evaluate performance, but the point to be emphasized here is that this research provides another example of a more accurate method that turns out to be faster than its rivals.

Christopher Beattie also discussed approximating $f(A)b$ but in a much more general setting.

The one sad note at this cheerful meeting was the news that this was the last workshop to be organized by Kresimir Veselic, who plans to retire from Fern-universität, the principal host of IWASEP5, in 2005. All who have become addicted to these stimulating informal gatherings should be grateful to him for establishing them. Participants appreciated the successful efforts of organizers Barlow, Ivan Slapnicar, and Veselic to produce another rewarding reunion.—*Beresford Parlett, University of California, Berkeley.*