Talk Abstracts

James Hardy Wilkinson

Sven Hammarling
University of Manchester

This talk will be a personal reflection on Wilkinson's life and work; his career, his contributions and his support for people and projects.

Eigenvalue Dating Service

Cleve Moler
MathWorks

Eigenvalues deserve to find their natural mates. Jim Wilkinson knew that a matrix with badly conditioned simple eigenvalues must be close to a matrix with eigenvalues of multiple eigenvalues. How can you find the nearest matrix where some of the eigenvalues have found their mates? One of Jim's favorite examples was the Frank matrix of order 12.

Software for Linear Algebra Targeting Exascale (SLATE)

Jack Dongarra
University of Tennessee, Oak Ridge Laboratory and University of Manchester

We will discuss the SLATE Project. SLATE is to provide fundamental dense linear algebra capabilities from today's high-performance computing (HPC) community. To this end, SLATE will provide basic dense matrix operations (e.g., matrix multiplication, rank-k update, triangular solve), linear systems solvers, least square solvers, singular value and eigenvalue solvers. The ultimate objective of SLATE is to replace the Scalable Linear Algebra PACKage (ScaLAPACK) library, which has become the industry standard for dense linear algebra operations in distributed memory environments. However, after two decades of operation, ScaLAPACK is past the end of its lifecycle and overdue for a replacement, as it can hardly be retrofitted to support hardware accelerators, which are an integral part of today's HPC hardware infrastructure.

Eigenvalue computation for structured problems. A Wilkinson heritage.

Volker Mehrmann
Technische Universität Berlin

Eigenvalue computation is at the heart of many procedures in all areas of science and engineering. In many applications the matrices or matrix pencils have structures that reflect the physics of the underlying problem. Preserving such structures is helpful for computational speed but essential to obtain maximally accurate results in strongly backward stable way. We discuss two classes of (generalized) eigenvalue problems; problems associated with Hamiltonian dynamical systems and problems associated with optimal control problems. We present numerical algorithms for small and large scale problems, error and perturbation analysis and present several numerical examples.

On the Amplification of Rounding Errors

Erin C. Carson
Charles University

In this talk we discuss the importance of distinguishing between the amplification versus the accumulation of rounding errors. The latter may be seen as a relatively harmless, unavoidable consequence of computation in finite precision, whereas the former is a property of the mathematical structure of the method and its algorithmic implementation used to transform the data. In designing algorithms to implement a particular method, it is thus of utmost importance to understand the mechanisms by which accumulated (or even elementary local) rounding errors can be amplified.

Perturbing doubly stochastic matrices

Philip Knight
University of Strathclyde

Many nonnegative matrices can be transformed to doubly stochastic form by the application of diagonal scaling. The transformed matrices have found in a role in an incredibly diverse set of applications from proportional representation to optimal transport. In particular, the spectral properties of doubly stochastic matrices mean that they can offer superior fidelity when uncovering hidden structure in data over other eigenvector-based methods. In this talk we consider the following question: if we take two nearby matrices and apply diagonal scaling to both, how close are the resulting doubly stochastic matrices? By taking into account the action of algorithms to achieve balance we are able to establish a new result. In particular, we show that conditions for the scalings to remain close are related to the fidelity of the block structure in the original matrices.

Squeezing a Matrix Into Half Precision, with an Application to Solving Linear Systems

Srikara Pranesh
University of Manchester

Motivated by the demand in machine learning, modern computer hardware is increasingly supporting reduced precision floating-point arithmetic, which provides advantages in speed, energy, and memory usage over single and double precision. Given the availability of such hardware, mixed precision algorithms that work in single or double precision but carry out part of a computation in half precision are now of great interest for general scientific computing tasks. Because of the limited range of IEEE half precision arithmetic, in which positive numbers lie between 6×10^{-8} and 7×10^{4}, a straightforward rounding of single or double precision data into half precision can lead to overflow, underflow, or subnormal numbers being generated, all of which are undesirable. An algorithm for converting a matrix from single or double precision to half precision will be discussed in this poster. This proposed algorithm will be used for solving a system of linear equations, using Zdenek Strakosthe GMRES-based iterative refinement. In this regard, matrices from actual applications will be used to numerically demonstrate the effectiveness of the proposed algorithm.

Lessons Taught by James Wilkinson

Margaret Wright
Courant Institute of Mathematical Sciences

Jim Wilkinson was not a university faculty member during most of his career, but his published works nonetheless constitute a matchless body of fundamental research. Beyond his publications, Jim's legacy includes classroom notes and homework problems for courses that he taught during 1977-82 as a visiting professor in Stanford University's Computer Science Department. Copies of his detailed notes, prepared before TeX (!), have been obtained from former Stanford students and visitors. This talk will highlight some of the treasures in Jim's notes, including his thoughts about good ways to present the ideas to students.

Wilkinson, Numerical, and Me

Nick Trefethen
University of Oxford

The two courses I took from Wilkinson as a graduate student at Stanford influenced me greatly. Along with some reminiscences of those days, this talk will touch upon backward error analysis, Gaussian elimination, and Evariste Galois.

Probabilistic Versus Worst-Case Rounding Error Analysis

Nicholas J. Higham
University of Manchester

Wilkinson developed the theory and practice of backward error analysis and carried out analyses of many fundamental algorithms. His bounds were worst-case bounds, but he recognized that "In general, the statistical distribution of the rounding errors will reduce considerably the function of n occurring in the relative errors. We might expect in each case that this function should be replaced by something which is no bigger than its square root." We will give, for the first time, a rigorous foundation for Wilkinson's expectation. using a new probabilistic rounding error analysis that can be applied to a wide range of algorithms. Our results have three key features: they are backward error bounds; they are exact, not first order; and they are valid for any n, unlike results obtained by applying the central limit theorem, which apply only as $n \to \infty$. Our numerical experiments show that for both random and real-life matrices the bounds can be much smaller than the standard deterministic bounds and can have the correct asymptotic growth with n. This is joint work with Theo Mary.

Redundant yet rapid approximation of functions

Marcus Webb
Katholieke Universiteit Leuven

Redundancy can be used to approximate functions when a well-conditioned basis with rapidly convergent approximations is not practically available. For example, when a function lives on an irregularly-shaped domain or if it experiences singularities. Extremely ill-conditioned linear systems (with condition numbers much greater than 1-over-machine-epsilon), arise from these redundant approximation problems, but they can nevertheless be solved in a sense which yields rapidly convergent approximations to the given function. In this talk we discuss recent progress in fast, direct, randomised solvers for such problems.

The perfidious condition number

Zdeněk Strakoš
Charles University

The condition number has been introduced as a single-number characterization of the sensitivity of a computed result to small changes in the data. With the concept of ill-conditioning came the idea of preconditioning linear algebraic systems. Later on, bounds based on condition numbers became widely used in the context of the convergence of iterative computations. With Krylov subspace methods, however, much precaution is needed, as strongly nonlinear phenomena can be meaningfully linearized only locally.

This contribution will present a case based on elliptic PDEs. It demonstrates that the naive condition number bound can provide a completely wrong assessment of convergence of the conjugate gradient method. On the other hand, a careful look at the PDE problem and at its discretization provides a-priori information on the whole matrix spectrum and allows a precise quantitative convergence description based on the reduced condition numbers. The result is based on a new theorem which uses a perturbation argument together with the classical Hall's theorem on bipartite graphs. It would not be attainable without incorporating numerical stability analysis of the conjugate gradient method.

This is a joint work with Tomáš Gergelits, Kent-André Mardal and Bjørn Nielsen.

Advances in high accuracy matrix computations

Zlatko Drmac
University of Zagreb

Modern theoretical developments and exciting applications in applied sciences and engineering demand efficient and numerically sound algorithms for matrix computations. Our aim is to illustrate how some recent developments in numerical linear algebra (accurate algorithms for numerical computation of eigenvalues and singular values, and corresponding theory) improve numerical computations in various applications. In particular, we stress the importance of error and perturbation analysis that identifies relevant condition numbers and guides computation with noisy data. Proper identification of a condition number that governs the forward error under a class of perturbations is the key for the design of a robust matrix algorithm.

Probabilistic roundoff error analysis for fundamental matrix computations

Ilse Ipsen
North Carolina State University

The objective is a probabilistic and realistic model for analysing the accumulation of floating point errors in fundamental matrix computations. We derive perturbation and forward roundoff error bounds, with new condition numbers, based on Azuma's inequality and Martingales.

Starting with the basic problem of computing an inner product between two real vectors, we represent the perturbations as independent, discrete, bounded random variables, and derive the bound from Azuma's inequality for independent random variables. In contrast, the roundoff errors, also represented as discrete, bounded random variables, are not required to be independent. We derive forward error bounds by constructing a Martingale that mirrors the particular order of computation, followed by an appropriate Azuma inequality. Time permitting, this will be extended to matrix vector products and matrix multiplication.

Numerical experiments confirm that our bounds are more realistic (often by several orders of magnitude) than existing deterministic worst-case bounds—even for small vector dimensions and stringent success probabilities.

Insights into block rational Krylov methods

Stefan Güttel
University of Manchester

Block (rational) Krylov methods are used in model order reduction, matrix function approximation (e.g. in multisource electromagnetic modeling), and for the solution of matrix equations, eigenvalue problems, and linear systems with multiple right-hand sides. I will give an overview of some of these applications and discuss theoretical aspects of the methods.

Wilkinson's Improvement Algorithm and Mixed Precision Eigensolvers

Françoise Tisseur
University of Manchester

Thanks to recent work by Carson and Higham, we now have a clear and thorough understanding of how to use mixed precision arithmetic to accelerate the solution of linear systems using iterative refinement. Can the same said about iterative refinement for eigenvalue problems? We will discuss this question, including Wilkinson's improvement algorithm and Newton's method in floating point arithmetic.

Preconditioned iterative methods for Toeplitz systems

Jennifer Pestana
University of Strathclyde

Linear systems with real, nonsingular Toeplitz and block Toeplitz matrices arise in many applications, including the discretisation of partial differential equations and fractional differential equations. Although the Toeplitz matrices may be dense, these linear systems are amenable to solution by iterative methods, e.g., Krylov subspace methods.

To keep the number of iterations low preconditioning is typically required, but for nonsymmetric Toeplitz matrices these preconditioners are usually heuristically chosen. However a simple flip to a Hankel matrix generates an equivalent problem with a symmetric coefficient matrix. In this talk we discuss the implications of flipping for the spectrum of (block) Toeplitz matrices, and present new preconditioners that are guaranteed to be effective.

Numerical Linear Algebra Routines for Emerging Computer Architectures

Mawussi Zounon
The Numerical Algorithms Group

NAG's numerical library provides state of the art mathematical algorithms for solving many scientific and engineering problems efficiently and reliably. NAG's core software runs predominantly on Intel architectures and it exploits the Intel Math Kernel Library (MKL) for solving numerical linear algebra problems at the heart of its algorithms. But recently, new and emerging architectures from vendors including AMD, Fujitsu and ARM are starting to completely change the hardware landscape. These vendors provide few or no optimized low-level numerical linear algebra routines. Therefore, NAG is investing in the design of a new generation of numerical linear algebra routines to exploit these emerging computer architectures efficiently. In this talk we will present a promising alternative to MKL BLAS that will underpin the NAG library in the future. As an alternative to MKL LAPACK routines, we have redesigned the Parallel Linear Algebra Software for Multicore Using OpenMP library (PLASMA) by reducing its memory overhead from 100% to zero, while significantly improving the performance up to 30%. We will present new results using the refactored PLASMA library on the very latest architectures and summarize our experience.

Stable and efficient QR factorization and least-squares solver based on CholeskyQR

Yuji Nakatsukasa
University of Oxford

CholeskyQR is an efficient algorithm for computing the QR factorization that is not backward stable and has been rarely used. However, by regarding a CholeskyQR-like step as a means to improve the conditioning, we can derive stable and efficient algorithms applicable to ill-conditioned matrices with $\kappa_2(A) = O(u^{-1})$. Such algorithms are attractive particularly when working with a non-standard inner product (e.g. B-orthogonal, quasimatrices). I will also discuss the use of CholeskyQR for solving (moderately ill-conditioned) least-squares problems, which results in a 5x speedup compared with a standard solver based on Householder QR, with comparable accuracy.

Computation of Matrix Gamma Function

Joao R. Cardoso
Coimbra Polytechnic-ISEC

An important matrix function, which has connections with certain matrix differential equations and other special matrix functions, is the matrix gamma function. This presentation is focused on its numerical computation. Well-known techniques for the scalar gamma function, such as Lanczos and Stirling approximations, are extended to the matrix case. This extension raises many challenging issues and several strategies used in the computation of matrix functions, like Schur decomposition and block Parlett recurrences, need to be incorporated to make the methods more effective. We also propose a third technique based on the reciprocal gamma function that is shown to be competitive with the other three methods in terms of accuracy, with the advantage of being rich in matrix multiplications. Strengths and weaknesses of the proposed methods are illustrated with examples.

Probabilistic Linear Solvers

Jonathan Cockayne
University of Warwick

A fundamental task in numerical computation is the solution of large linear systems. The conjugate gradient method is an iterative method which offers rapid convergence to the solution, particularly when an effective preconditioner is employed. However, for more challenging systems a substantial error can be present even after many iterations have been performed. The estimates obtained in this case are of little value unless further information can be provided about the numerical error. We propose a novel statistical model for this numerical error set in a Bayesian framework. Our approach is a strict generalisation of the conjugate gradient method, which is recovered as the posterior mean for a particular choice of prior. The estimates obtained are analysed with Krylov subspace methods and a contraction result for the posterior is presented.

Creating an Adjoint Library from the NAG Primal Code Base

Lawrence Mulholland
The Numerical Algorithms Group

Sensitivities of outputs with respect to inputs for a large complicated algorithmic code are increasingly required. An adjoint method can greatly reduce the computational complexity of financial derivative calculations traditionally performed by a finite-difference approach known as "bumping". If a code contains calls to a third party library and adjoints are required then an adjoint of the library call must be provided. For a large library code base, algorithmic differentiation (AD) tools must be used to automatically transform the code into a version from which adjoints can be computed; such a tool is dco/c++. However, linear algebra operations are often at the core of these calculations, and many of these operations have adjoints that are simple to represent symbolically and can just be variant calls to standard linear algebra routines. Hence a hybrid approach of top-down transformed code bottom-up symbolic replacement might be appropriate. Here we discuss how dco/c++, designed primarily for C++ code, can be used to transform a large Fortran primal code base into a library of adjoint routines callable from C++ and Fortran.

Wilkinson's bus: weak condition numbers, with applications

Vanni Noferini
Aalto University

We propose a new approach to the theory of conditioning for numerical analysis problems where both the classical (worst-case) and stochastic perturbation theory fail to capture the observed perturbation behaviour. To motivate the new ideas, we also present and discuss a class of relevant examples coming from polynomial eigenvalue problems. We also discuss how to compute, in practice, the weak condition number for these problems.

Poster Abstracts

Improvements to the 2DRMP code for the simulation of electron scattering from H-like atoms

Maksims Abalenkovs
University of Manchester

The R-matrix approach for solving the close-coupling equations arising in electron and photon collisions uses a matrix representation to describe the interaction between two particles, by discretising each particle's wave-function using an orthonormal set of energy-independent basis functions. The R-matrix propagation method breaks down the large and densely populated global R-matrix into a series of sub-matrices, propagating the influence of each sub-matrix throughout the domain in order to obtain a global solution for the particle interaction.

This work seeks to improve the linear algebra component of an existing code for 2D R-Matrix propagation, known as 2DRMP. Originally developed at Queen's University Belfast, the code has been modified to improve the portability and user accessibility of the underlying linear algebra routines, and to improve performance on non-Intel architectures. By introducing the PLASMA linear algebra library we are able to access emerging architectures, such as ARM, and take advantage of task-based parallelism and a tile-based approach to solving dense linear algebra in shared memory machines.

The poster will demonstrate some of the modifications made to the 2DRMP code, including performance on non-Intel architectures, automatic optimisation of numerical linear algebra parameters, and streamlining of existing LAPACK routines.

Fast and Accurate Algorithms for Large-scale Low-precision Computations

Pierre Blanchard
University of Manchester

Modern scientific computing faces a double challenge with the rise of large-scale, low-precision computations. On the one hand, larger and larger problems are being solved routinely; on the other hand, the use of low precision arithmetics, such as half precision (fp16), is becoming increasingly attractive. Conventional algorithms, for example in numerical linear algebra, are unfortunately not often able to exploit the speed benefits of using lower precisions while still delivering low floating-point errors for large problems. We present an overview of recent advances in developing new algorithms that are both fast and accurate. We illustrate the superiority of these new techniques for numerical linear algebra algorithms through their rounding error analysis as well as with numerical experiments on a wide range of problems and architectures, including NVIDIA GPUs equipped with tensor cores. This is joint work with Nick Higham and Theo Mary.

Investigation of Numerical Stability of Algebraic Linearizations

Eunice Chan
Western University

A matrix polynomial is defined by $P(\lambda) = \sum_{i=0}^{n} A_i \lambda^i$, where A_i is an $n \times n$ matrix, and it is commonly solved using linearization. Chan et al. (2019) introduced a new construction of linearizations to solve matrix polynomials of the form $za(z)b(z) + c$, where $a(z)$ and $b(z)$ are matrix polynomials, and c is a matrix of the same size. In this poster, we will be giving the preliminary results of the backward stability of the solutions given by this new linearization. We will be using one of the techniques given by Higham, Li, and Tisseur (2007) to investigate backward stability. This technique perturbs the coeffients of the matrix polynomials $a(z)$ and $b(z)$ and see how much the solutions change. We hope to find the upper bound of the backward error of this new linearization.

A Comparison of 2x2 and 3x3 Block Saddle Point Formulations of Weak Constraint 4D-Var

Ieva Dauzickaite
University of Reading

Data assimilation is used to estimate the state of a dynamical system by combining observations of it with numerical model. The estimate can be achieved by minimising a weak constraint 4D-Var cost function. The minimisation may be performed by using an iterative method to solve a sequence of saddle point systems. The solutions of these systems provide an increment to the state vector. Some blocks in the previously proposed 3x3 block saddle point matrix depend on the available observations. We present a novel examination of how introducing new observations influences the convergence of the iterative method. The convergence of linear iterative solvers depends on the spectrum of the coefficient matrix, hence we derive bounds for the eigenvalues of the 3x3 block matrix. Also, inspired by the practice in solving saddle point systems that arise from interior point methods, we reduce the 3x3 block system to a 2x2 block formulation and derive new eigenvalue bounds for this system.

On the Structure of Linearizations for Rational Approximations of Nonlinear Eigenvalue Problems

Maria del Carmen Quintana Ponce
Universidad Carlos III de Madrid

In this poster we present the notion of local linearizations of rational matrices. The main property of a local linearization of a rational matrix R(z) is that it preserves the zeros and poles, together with its partial multiplicities, of R(z) locally, that is, in subsets of the complex plane and/or at infinity. By using this new notion of linearization, we study the structure of linearizations constructed in the literature for solving nonlinear eigenvalue problems by using rational approximation on a target set.

Computing Matrix Functions in Arbitrary Precision Arithmetic

Massimiliano Fasi
University of Manchester

In some applications, results computed in IEEE double precision arithmetic are not sufficiently accurate and higher precision is needed, often due to ill conditioned or badly scaled problems. Yet in other situations, single or even half precision is sufficient. The number of programming languages supporting arbitrary precision arithmetic, either natively or via dedicated libraries, is increasing, and includes C++, Julia, MATLAB and Python. Here, we present algorithms for computing the matrix logarithm and the matrix exponential in arithmetic of arbitrary precision. The algorithms adapt their underlying approximations to the precision through the use of rigorous error bounds. We explain, in particular, the algorithmic issues that arise in working at precisions higher than double.

Wilkinson Test Matrices

Carla Ferreira
University of Minho

In the 1950's J. H. Wilkinson introduced two families of symmetric tridiagonal integer matrices. Most of the eigenvalues are close to diagonal entries. We develop the structure of their eigenvectors in a natural way which reveals that the envelopes of these eigenvectors all look the same to the naked eye. The shape is a badly dented bell curve. We also analyze the eigenvectors of the remaining non-integer eigenvalues.

Nearest Correlation Algorithms in the NAG Library

Craig Lucas
The Numerical Algorithms Group

NAG has implemented many different algorithms for computing nearest correlation matrices in its libraries. The starting point is a matrix that is not mathematically a true correlation matrix, in particularly not positive semi-definite, and the result is one that is. Some algorithms use weighting and fixing of elements in the input matrix, some don't use a norm as the measure of nearness. We compare the functionality and performance of our suite of routines on a range of large problems and present the results.

Using random projections to accelerate numerical linear algebra with the NAG library

Philip Maybank
The Numerical Algorithms Group

The NAG library is an extensive collection of numerical and statistical algorithms, the largest such commercial library in the world. The NAG Library is widely used and trusted in industry because of its unrivalled quality, reliability and portability.

Historically NAG was a key contributor to the design and implementation of the widely used LAPACK software for numerical linear algebra. An implementation of LAPACK was subsequently incorporated into the NAG library.

Recently we have been exploring the use of random projection based algorithms for the solution of large-scale numerical linear algebra problems, that can be more efficient than current LAPACK algorithms. This work is motivated by findings in the nascent field of Randomized Numerical Linear Algebra (RNLA), particularly the work of Halko et al. on matrix decompositions and Avron et al. on preconditioned iterative least-squares solvers.

Our work to date has mainly focused on the case of dense matrices that fit into core memory, and on evaluating whether RNLA algorithms outperform LAPACK in terms of computational efficiency. This work has led to the incorporation of new routines for RNLA into the next release of the NAG library. In this talk we outline the process of translating research findings into code that can be used in production, and we present benchmarking results comparing our RNLA implementation with LAPACK.

References: Avron, H., Maymounkov, P., & Toledo, S. (2010). Blendenpik: Supercharging LAPACK's Least-Squares Solver, SIAM Journal on Scientific Computing. Halko, N., Martinsson, P.-G., & Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, SIAM Review.

On deflation process and solving the quadratic eigenvalue problems

Ivana Šain Glibić
University of Zagreb, Faculty of Science, Department of Mathematics

We analyze the numerical properties of the deflation process from the quadeig algorithm for the solution of the quadratic eigenvalue problems developed by Hammarling, Munro and Tisseur. The deflation process attempts to remove zero and infinite eigenvalues from the corresponding linearization before calling the QZ algorithm. We propose several enhancements that lead to more robust procedure. In addition to the theoretical considerations, we provide numerical examples to illustrate the improved accuracy of the new approach.

Author Index