Talk Abstracts

Jack

Iain Duff

STFC-RAL and Cerfacs

After a somewhat irreverent and cursory look at the future for NLAand HPC, I will dedicate my talk to the person whom this conference is hono(u)ring, Jack Dongarra.

Through mainly the medium of photographs (pictures with a story), I will reflect on aspects of Jack's life that maybe even he is not aware of! This talk is meant to be fun, as much as it can be with an online only presentation, so those with a nerd index (to be defined) of greater than 7 may want to join the meeting later!

Happy Birthday, Jack!

New communication-avoiding algorithms, and fixing old bugs in the BLAS and LAPACK

James Demmel

University of California, Berkeley

(1) Algorithms have two costs: arithmetic and communication, i.e. moving data between levels of a memory hierarchy or processors over a network. Communication costs (measured in time or energy per operation) greatly exceed arithmetic costs, so our goal is to design algorithms that minimize communication. We survey known algorithms that communicate asymptotically less than their classical counterparts, for a variety of linear algebra and machine learning problems, and present some new ones for mixed precision computation, and sketching.
(2) As more aspects of society become automated, eg self-driving cars, health monitors, and cyber-physical systems more generally, it is becoming increasingly important to design software that is resilient to exceptions, and that responds to them in a consistent way. We present some inconsistencies in the BLAS, show how they propagate into LAPACK, and present our overall plan for making numerical software more resilient against exceptions.

Numerical methods and benchmarking across scales, precisions, and hardware platforms

Piotr Luszczek

University of Tennessee

For many decades now, HPC served scientific community in a variety of ways that advanced the computational simulation capability based on models that not so long ago remained theoretical. Measuring the progress made by the high-end hardware has become its own field with a vibrant community of performance engineers exchanging the practical experiences and even building theory of performance models that allow the users to gain a deeper comprehension of the multi-million dollar machines. With transition to data science and great demand for deep models for statistical inference, the need for accurate measurement of hardware capability has remained strong. With comparative performance studies that inform the decision makers ranging from purchasing to software implementation and deployment, this talk will offer glimpse of how scientifically important numerical methods can serve to give insightful information into the potential and the drawbacks of the modern hardware platforms. The talk will feature description of algorithms, historical perspective on the HPC hardware, and results from a variety of systems from relevant benchmarks.

Computing in the Old Days

Cleve Moler

MathWorks

I want to reminisce about how we did our computing in the days before everybody had one in their lap. 205, 7090, 4081, VAX, Sun, Ardent, FPS, JCL, and especially BOF.

25+ years of scheduling at ICL

Yves Robert

Ecole Normale Suprieure de Lyon

This talk reviews a few results in parallel linear algebra and resilience techniques that were achieved at ICL under Jack's guidance. The story started in 1996. It continues today.

Mixed Precision Tall and Thin QR Factorization with Applications

Françoise Tisseur

The University of Manchester

We present an algorithm for the reduced QR factorization of tall and skinny matrices. Low (half and single) precision is used to precondition the matrix into one that has condition number of order 1 at working precision. As a result the Cholesky-QR algorithm applied to the preconditioned matrix yields a QR factorization with Q orthonormal to working precision and a relative residual of the order of the working precision. This algorithm does not require a priori knowledge of the matrix condition number. It is rich in BLAS 3 operations, is communication-avoiding, and has the potential to be significantly faster than TSQR.

This is joint work with Srikara Pranesh

Adaptive Nonlinear Preconditioning for PDEs with Error Bounds on Output Functionals

David Keyes

King Abdullah University of Science and Technology

Nonlinear preconditioning refers to transforming a nonlinear algebraic system into a form for which Newton-type algorithms have improved success through quicker advance to the domain of quadratic convergence. We place these methods, which go back at least as far as the Additive Schwarz Preconditioned Inexact Newton (ASPIN, 2002) in the context of a proliferation distinguished by being left- or right-sided, multiplicative or additive, and partitioned by field, subdomain, or other criteria. We present the Nonlinear Elimination Preconditioned Inexact Newton (NEPIN, 2021) based on a heuristic "bad/good" heuristic splitting of equations and corresponding degrees of freedom. We augment nonlinear preconditioning with two features of practical interest: an adaptive switchover to ordinary Newton as the domain of convergence is approached and error bounds on output functionals of the solution, and we illustrate on various nonlinearly stiff discretized PDEs. Joint work with Lulu Liu, Li Luo, Xiao-Chuan Cai, and others.

Many eigenpair computation via Hotelling's deflation

Zhaojun Bai
University of California, Davis

There are continual and compelling needs for computing many eigenpairs of large Hermitian matrices in physical simulation and data analysis. Although Krylov eigensolvers are effective for computing eigenvalues on the periphery of the spectrum of Hermitian matrices, they are expensive for computing eigenvalues located deep inside the spectrum. In this talk, we discuss an algorithm that combines the Lanczos method and the classical Hotelling's deflation for many eigenpair computation. We will present a recent study on the backward stability of Hotelling's deflation and dynamical selection of shifts for numerical backward stability. In addition, we will also discuss a preliminary implementation of a communication-avoid variant of the algorithm. This is a joint work with Jack Dongarra, Chao-Ping Lin, Ding Lu and Ichitaro Yamazaki.

A few observations about summation algorithms

Ilse Ipsen
North Carolina State University

Motivated by the fast summation algorithms of Blanchard, Higham and Mary (2020) and the probabilistic approach of Higham and Mary (2019, 2020), we consider different algorithms for the summation of real numbers: binary, blocked, centered/shifted, and compensated. We present deterministic and probabilistic bounds numerical experiments, and comparisons. This is joint work with Eric Hallman.

TOP500 and Accidental Benchmarking

Erich Strohmaier
TOP500

The TOP500 list (www.top500.org) has served as the defining yardstick for supercomputing performance since 1993. Published twice a year, it compiles the worlds 500 largest installations and some of their main characteristics. Systems are ranked according to their performance of the Linpack benchmark, which solves a dense system of linear equations and was invented by Jack J. Dongarra. Over time, the data collected for the list has enabled the early identification and quantification of many important technological and architectural trends related to high-performance computing. We briefly describe the projects history and our main observations. We also reflect on the reasons for the continued success of the list over almost three decades and what we see as it's main challenges and limitations.

Solving Dense Linear Systems: A Brief History and Future Directions

Nick Higham

The University of Manchester

We take a brief look at some landmarks in the history of numerical methods for solving dense linear systems $Ax = b$, with a particular focus on Jack Dongarra's contributions. We also discuss current challenges and future research directions.

Still having Fun After 50 Years

Jack Dongarra

University of Tennessee, Oak Ridge Laboratory and The University of Manchester

In this talk, we will look back at some of the projects and software development efforts in which I have been fortunate to be involved

---