

Distinguished Lecture Series

The Limits of Matrix Computations at Extreme Scale and Low Precisions



Professor Nicholas J. Higham

Department of Mathematics, University of Manchester

Biography:

Nick Higham is Royal Society Research Professor and Richardson Professor of Applied Mathematics in the School of Mathematics, University of Manchester. He is well known for his research on the accuracy and stability of numerical algorithms.

He was awarded a Royal Society Research Professorship in 2018. He is a Fellow of the Society for Industrial and Applied Mathematics (SIAM), a Fellow of the ACM, and a Member of Academia Europaea. He held a Royal Society-Wolfson Research Merit Award (2003–2008). He was President of SIAM (2017–2018), and was Vice President at Large (2010–2013) and served for over ten years on the SIAM Board of Trustees and the SIAM Council. He has also served on the Board of Directors of the International Linear Algebra Society, and as Chair of the SIAM Activity Group on Linear Algebra.

Honours include the Alston S. Householder Award VI, 1987 (for the best Ph.D. thesis in numerical algebra 1984–1987), the 1988 Leslie Fox Prize in Numerical Analysis, a 1999 Junior Whitehead Prize from the London Mathematical Society, the 2008 Fröhlich Prize of the London Mathematical Society, the 2019 Naylor Prize and Lectureship by the London Mathematical Society, the 2020 IMA Gold Medal of the Institute of Mathematics and its Applications, and the 2021 George Pólya Prize for Mathematical Exposition from SIAM.

Date: 27 April 2022 (Wednesday)

Time: 4:00-5:00 p.m. GMT+8 (Hong Kong Time)

Venue: Online via Zoom (Meeting ID: 921 4525 8273)

Abstract

As computer architectures evolve and the exascale era approaches, we are solving larger and larger problems. At the same time, much modern hardware provides floating-point arithmetic in half, single, and double precision formats, and to make the most of the hardware we need to exploit the different precisions. How large can we take the dimension n in matrix computations and still obtain solutions of acceptable accuracy? Standard rounding error bounds are proportional to $p(n)u$, with p growing at least linearly with n . We are at the stage where these rounding error bounds are not able to guarantee any accuracy or stability in the computed results for extreme-scale or low-accuracy computations. We explain how rounding error bounds with much smaller constants can be obtained. The key ideas are to exploit the use of blocked algorithms, which break the data into blocks of size b and lead to a reduction in the error constants by a factor b or more; to take account of architectural features such as extended precision registers and fused multiply-add operations; and to carry out probabilistic rounding error analysis, which provides error constants that are the square roots of those of the worst-case bounds. Combining these different considerations provides new understanding of the limits of what we can compute at extreme scale and low precision in numerical linear algebra.

Sponsored by: Centre for Mathematical Imaging and Vision

 CMIV

Joint Research Institute for Applied Mathematics



✦ ✦ ✦ All are welcome ✦ ✦ ✦

<https://www.math.hkbu.edu.hk/>