IST AUSTRIA
Institute of Science and Technology

Mathematics and CS Seminar

Sparse Linear Algebra in the Exascale Era

Erin Carson (NYU)

Host: Chris Wojtan

Sparse linear algebra problems, typically solved using iterative methods, are ubiquitous throughout scientific and data analysis applications and are often the most expensive computations in large-scale codes due to the high cost of data movement. Approaches to improving the performance of iterative methods typically involve modifying or restructuring the algorithm to reduce or hide this cost. Such modifications can, however, result in drastically different behavior in terms of convergence rate and accuracy. A clear, thorough understanding of how inexact computations, due to either finite precision error or intentional approximation, affect numerical behavior is thus imperative in balancing the tradeoffs between accuracy, convergence rate, and performance in practical settings.In this talk, we focus on two general classes of iterative methods for solving linear systems: Krylov subspace methods and iterative refinement. We present bounds on the attainable accuracy and convergence rate in finite precision s-step and pipelined Krylov subspace methods, two popular variants designed for high performance. For s-step methods, we demonstrate that our bounds on attainable accuracy can lead to adaptive approaches that both achieve efficient parallel performance and ensure that a user-specified accuracy is attained. We present two such adaptive approaches, including a residual replacement scheme and a variable s-step technique in which the parameter s is chosen dynamically throughout the iterations. Motivated by the recent trend of multiprecision capabilities in hardware, we present new forward and backward error bounds for a general iterative refinement scheme using three precisions. The analysis suggests that on architectures where half precision is implemented efficiently, it is possible to solve certain linear systems up to twice as fast and to greater accuracy.As we push toward exascale level computing and beyond, designing efficient, accurate algorithms for emerging architectures and applications is of utmost importance. We discuss extensions to machine learning and data analysis applications, the development of numerical autotuning tools, and the broader challenge of understanding what increasingly large problem sizes will mean for finite precision computation both in theory and practice.

Wednesday, April 4, 2018 10:00am - 11:00am

IST Austria Campus Mondi Seminar Room 2, Central Building

This invitation is valid as a ticket for the IST Shuttle from and to Heiligenstadt Station. Please find a schedule of the IST Shuttle on our webpage: https://ist.ac.at/en/campus/how-to-get-here/ The IST Shuttle bus is marked IST Shuttle (#142) and has the Institute Logo printed on the side.