

SPPEXA @ ISC 2013

09:00 – 09:15	Welcome Addresses	
09:15 – 10:00	Invited	Pete Beckman (<i>Argonne National Laboratory</i>): New Directions for Extreme Scale System Software
10:00 – 10:30	SPPEXA	Christian Lengauer (<i>University of Passau</i>): Modern Software Technology for Exascale Computing
10:30 – 11:30	Coffee Break	
11:30 – 12:00	SPPEXA	Felix Wolf (<i>German Research School of Simulation Sciences</i>): Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes
12:00 – 12:30	SPPEXA	Oliver Rheinbach (<i>TU Bergakademie Freiberg</i>): Local Flops are Free – New Approaches to Nonlinear Domain Decomposition
12:30 – 13:00	SPPEXA	Frank Jenko (<i>Max Planck Institute of Plasma Physics</i>): Simulations of Turbulence in Fusion and Astrophysical Plasmas on Peta- to Exascale Platforms
13:00 – 14:15	Lunch Break	
14:15 – 15:00	Invited	Chuck Hansen (<i>University of Utah</i>): Big Data: A Scientific Visualization Perspective
15:00 – 16:15	<p>The potentials and limits of simulation and HPC A round-table discussion with</p> <p>Omar Ghattas (<i>University of Texas at Austin</i>), Gabriele Gramelsberger (<i>Freie Universität Berlin</i>), and Friedel Hoßfeld (<i>Jülich Supercomputing Centre</i>).</p> <p>Moderated by Patrick Regan (<i>Technische Universität München</i>)</p>	

Pete Beckman is Director of the Argonne Leadership Computing Facility and co-founder of the International Exascale Software Project, which has built an international software roadmap for exascale software and co-design.

Omar Ghattas is the John A. and Katherine G. Jackson Chair in Computational Geosciences, Professor of Geological Sciences and Mechanical Engineering, and Director of the Center for Computational Geosciences in the Institute for Computational Engineering and Sciences at the University of Texas at Austin.

Gabriele Gramelsberger works in the Institute of Philosophy at Freie Universität Berlin on the influence of computation on science and society.

Chuck Hansen is an IEEE Fellow and a Professor of Computer Science in the School of Computing and an Associate Director of the Scientific Computing and Imaging Institute at the University of Utah.

Friedel Hoßfeld was Director of the Institute for Advanced Simulation at Jülich Supercomputing Centre from 1973 to 2002.

Frank Jenko leads the Plasma Turbulence Group at the Max Planck Institute of Plasma Physics in Garching.

Christian Lengauer holds the Chair for Programming in the Department of Informatics and Mathematics at the University of Passau.

Patrick Regan is responsible for international public relations and video production at TUM-IAS. He came to TUM from PBS and NPR affiliate New Jersey Public Television and Radio, where he was the senior correspondent for science and technology.

Oliver Rheinbach is professor in the Institut of Numerical Analysis and Optimization at TU Bergakademie Freiburg.

Felix Wolf is head of the Parallel Programming laboratory at the German Research School for Simulation Sciences.

New Directions for Extreme Scale System Software

Pete Beckman

Argonne National Laboratory, USA

HPC systems are quickly evolving and moving in new directions. Features such as deep memory hierarchies, non-volatile memory, and dynamic power management are new frontiers for extreme scale system software. We have many challenges before us. How will the operating system and run-time software dynamically manage power and storage? How will the message layers, which have been very slow to change, adapt to the new system designs? This presentation will explore the new layers and capabilities that extreme scale system will adapt to for our next-generation platforms.

Modern Software Technology for Exascale Computing

Christian Lengauer

University of Passau, Germany

The practical and economical development and maintenance of reliable exascale applications requires a radical shift of the software technology of high-performance computing – from comparatively machine-oriented, manual coding in the general-purpose languages Fortran or C to a flexible, multi-step refinement, employing various, domain-specific languages and optimization techniques and supported by highly automatic development tools.

The talk will present an approach pursued in the project "Advanced Stencil-Code Engineering" as part of the DFG Priority Programme 1648 "Software for Exascale Computing".

Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes

Alexandru Calotoiu¹, Torsten Höfler², Marius Poke¹, Felix Wolf¹

¹*German Research School of Simulation Sciences, Germany*

²*ETH Zurich, Switzerland*

Many parallel applications suffer from latent performance limitations that may prevent them from scaling to larger machine sizes. Often, such scalability bugs manifest themselves only when an attempt to scale the code is actually being made - a point where remediation can be difficult. However, creating analytical performance models that would allow such issues to be pinpointed earlier is so laborious that application developers attempt it at most for a few selected kernels, running the risk of missing harmful bottlenecks. In this paper, we show how both coverage and speed of this scalability analysis can be substantially improved. Generating an empirical performance model automatically for each part of a parallel program, we can easily identify those parts that will reduce performance at larger core counts. Using a climate simulation as an example, we demonstrate that scalability bugs are not confined to those routines usually chosen as kernels.

Local Flops are Free – New Approaches to Nonlinear Domain Decomposition

Oliver Rheinbach

TU Bergakademie Freiberg, Germany

As a consequence of exploding concurrency, algorithms and software developed for today's supercomputers will not automatically be able to profit from the future generation of supercomputers. It is therefore a challenge to bring current implicit solvers for nonlinear problems, e.g., in nonlinear structural mechanics, to the scale of future massively parallel computers. In this talk, new nonlinear, nonoverlapping domain decomposition schemes will be presented that may help to overcome limitation of current solver schemes.

Simulations of Turbulence in Fusion and Astrophysical Plasmas on Peta- to Exascale Platforms

Frank Jenko

Max Planck Institute of Plasma Physics, Germany

Plasma turbulence is a ubiquitous phenomenon, influencing the dynamics in most of the visible universe and playing a crucial role in countless laboratory experiments of plasma science including, in particular, fusion research. Yet, various fundamental aspects of this prototypical nonlinear process are only poorly understood at present and our predictive capability is limited. In order to make progress on these fronts, one has to resort to extreme computing on the peta- to exascale.

For this purpose, cutting-edge numerical tools like the GENE code have been developed and are run on some of the largest available supercomputers. In this presentation, I will introduce the GENE code and discuss its use to predict and optimize the energy confinement characteristics of the upcoming international flagship project ITER. In this context, I will also describe the targeted GENE-related developments within the SPPEXA-EXAHD project.

Big Data: A Scientific Visualization Perspective

Chuck Hansen

University of Utah, USA

Modern high performance computers have speeds measured in petaflops and handle data set sizes measured in terabytes and petabytes. In the next decade (or less) we will see exascale computational resources. Although these machines offer enormous potential for solving very large-scale realistic computational problems, their effectiveness will hinge upon the ability of human experts to interact with their simulation results and extract useful information. One of the greatest scientific challenges of the 21st century is to effectively understand and make use of the vast amount of information being produced. Visual data analysis will be among our most important tools in helping to understand such large-scale information.

Our research at the Scientific Computing and Imaging (SCI) Institute at the University of Utah has focused on innovative, scalable techniques for large-scale 3D visual data analysis. In this talk, I will present state-of-the-art visualization techniques, including scalable visualization algorithms and software, cluster-based visualization methods and innovative visualization techniques applied to problems in computational science, engineering, and medicine. There is a bright future for large scale visual analysis and many challenges to overcome.