1st ExaHyPE Colloquium

Challenges for Simulation Software and Applications at Exascale

3 April 2017, LRZ Hörsaal 1

Presentations (13.00–17.15)

13.00–13.40 David Keyes (KAUST): *Algorithmic Adaptations to Extreme Scale*

13.40–14.20 Eric Petit, Marie-Christine Sawley (Intel): *Code Modernization for Next Generation Hardware: Models and Runtime Exploration for Efficient and Scalable Algorithms*

14.20–15.00 James Hobro (Schlumberger): *Advances in Scalability for Large-Scale Domain Decomposition in Finite-Difference Simulations*

15.00–15.20 Coffee Break

15.20–15.50 Michael Bader (TUM): *Towards an Exascale Hyperbolic PDE Engine: High Order ADER-DG on Tree-Structured Cartesian Meshes*

15.50–16.20 Roland Haas (NCSA, UIUC): *Simulating Multiphysics Astrophysical Problems Using Current and Future Codes*

16.20–16.50 Dave May (University of Oxford): *Two-phase Flow in Subduction Zones*

David Keyes

David Keyes is the director of the Extreme Computing Research Center at King Abdullah University of Science and Technology, where he was a founding dean in 2009, and an adjoint professor of applied mathematics at Columbia University. Keyes earned his BSE in Aerospace and Mechanical Engineering from Princeton and his PhD in Applied Mathematics from Harvard. He works at the algorithmic interface between parallel computing and the numerical analysis of partial differential equations. He is a Fellow of SIAM and AMS and has received the ACM Gordon Bell Prize and the IEEE Sidney Fernbach Award.

Algorithmic Adaptations to Extreme Scale

Algorithmic adaptations to use next-generation computers closer to their potential are underway. Instead of squeezing out flops – the traditional goal of algorithmic optimality, which once served as a reasonable proxy for all associated costs – algorithms must now squeeze synchronizations, memory, and data transfers, while extra flops on locally cached data represent only small costs in time and energy.

After decades of programming model stability with bulk synchronous processing, new programming models and new algorithmic capabilities must be co-designed with the hardware. We briefly recap the architectural constraints and application opportunities. We then concentrate on two types of tasks each of which occupies a large portion of all scientific computing cycles: large dense symmetric/Hermitian systems and large sparse Poisson/Helmholtz systems. We examine progress in porting solvers for these tasks to the hybrid distributed–shared programming environment, including the GPU and the MIC architectures that make up the cores of the top scientific systems on the floor and on the books.

Eric Petit

Eric Petit is a research engineer at Intel, since 2016. He received his Ph.D. degree in computer science from the University of Rennes and INRIA, in 2009. He spent two years in University of Perpignan as an associate professor and six years at the University of Versailles as a research scientist. Dr. Petit's current research interests are on preparing HPC applications for future exascale many-core systems. He focuses on concurrency and locality characterization and optimization, auto-tuning frameworks, new parallelization strategies based on innovative runtime environment and the floating-point effect on numerical accuracy.

Code Modernization for Next Generation Hardware: Models and Runtime Exploration for Efficient and Scalable Algorithms

With the shift in HPC system paradigm, applications have to be re-designed to efficiently address large numbers of manycores. With the advance in network technology and communication library, new opportunities to explore advanced programming model and load-balancing runtime system in large HPC cluster emerged. However, it requires deep code modification that cannot be practically applied at full-application scale. We propose proto-applications as proxy for code modernization and we will demonstrate two specialized libraries for HPC workload load-balancing based on the GASPI communication library (PGAS) and task based programming. Our first example features unstructured mesh computation using task-based parallization. The second demonstrates load-balancing for combustion simulation. Both proto-applications are open-source and can serve to the development of genuine HPC application.

James Hobro

James Hobro is a Principal Research Scientist at Schlumberger Gould Research in Cambridge. He received his Ph.D. in Earth Sciences from the University of Cambridge in 2000. After a period of post-doctoral research he joined Schlumberger and has since worked in research and product development on a variety of problems related to seismic modeling, imaging and inversion. His work has focused on developing techniques in these areas for application to increasingly complex Earth structures, encompassing ray-based methods and wave equation modelling. His current research areas include waveform inversion and efficient, scalable elastic modelling for complex, anisotropic media.

Advances in Scalability for Large-Scale Domain Decomposition in Finite-Difference Simulations

Finite-difference modelling has seen widespread commercial use in seismic applications since the mid-2000s, most notably in reverse time migration and full waveform inversion. These applications require modelled wavefields to be computed independently for large numbers of shots and are therefore loosely-coupled and scale readily, with large-scale data management providing the main challenges. However, the move towards elastic modelling at higher frequencies, developments in computing hardware, the ready availability of cloud computing resources, and the demands of interactive applications are all leading to an increasing focus on large-scale domain decomposition in seismic modelling. This is a disruptive phase for modelling technology with many promising developments being explored. In this talk I will focus on some of the improvements that can be achieved in the scalability of finite-difference modelling by moving towards a highly-asynchronous actor-based design. This minimises the loss of performance due to the artificial synchronisation present in conventional formulations and also streamlines network communication. I will also explore the application of theoretical scaling models in investigating and characterising real scaling behaviour in HPC systems for this type of algorithm.

Michael Bader

Michael Bader is associate professor at the Technical University of Munich and leads the research group for hardware-aware algorithms and software for HPC. Since 2015 he coordinates the Horizon 2020 project "ExaHyPE – An Exascale Hyperbolic PDE Engine".

Towards an Exascale Hyperbolic PDE Engine

The ExaHyPE project develops an engine for the solution of hyperbolic PDEs at exascale. Our approach is based on a space-tree discretization of the computational domain, high-order discontinuous Galerkin space-time discretization and a-posteriori subcell limiting. This setup shall combine highest numerical accuracy with the exploitation of massive parallelism and an intense level of adaptive mesh refinement. Designated applications of the ExaHyPE engine will be grand challenge simulations in astrophysics and seismology. In this talk, we present the current status of the ExaHyPE engine, which has been released as open source software in March 2017.

Roland Haas

Roland Haas is a senior research programmer at the Blue Waters project at the National Center for Supercomputing Applications. He received his Ph.D. in theoretical physics from the University of Guelph, Canada in 2008. His tour as a postdoc included positions at Georgia Tech, Caltech and the Albert Einstein Institute. Dr. Haas' current research interest focus on the simulation of gravitational waves generated by binary neutron star mergers. He is one of the core developers of the Einstein Toolkit with particular interest in scalability and user friendliness of the toolkit.

Simulating Multiphysics Astrophysical Problems Using Current and Future Codes

The birth and death of neutron stars requires truly multiphysics simulations to handle general relativity, relativistic magneto-hydrodynamics and neutrino radiation transport resolving multiple length scales from the 100m range to hundreds of km and time-scales from microseconds to minutes or even days. These simulations continue to require the largest available computing resources and provide challenges both in how to use these resources efficiently and also on how to orchestrate simulations and data. I present an overview of current simulations using the Einstein Toolkit as well as glimpses for future developments to overcome limitations of current codes.

Dave May

Dave May is a senior research fellow in computational geosciences at the University of Oxford. He received his PhD in mathematics in 2009 from Monash University, Australia. From 2009–2010 we has a postdoc in the Institute of Geophysics at ETH Zurich, and from 2011–2016 he was an ober-assistant in the same institute. His research interests include the development of discretisations, massively parallel preconditioners and software for simulating highly viscous flows.

Dr. May has been using and developing the Portable Extensible Toolkit for Scientific computing (PETSc) since 2004.

Two-phase Flow in Subduction Zones

Subduction defines the process in which the oceanian lithospheric plunges underneath a strong, more buoyant over-riding continental lithosphere. As the oceanic plate descends to greater depths, the plate increases in temperature and is subjected to higher confining pressures. This gives rise to the release of fluids along the upper surface of the plate and results in partial melting of the mantle. We are interested in understanding the long-time scale evolution of flow in the mantle wedge (the region between the subducting plate and the over-riding plate). To that end we consider the deformation of the mantle rock, together with fluids and melts, as a coupled two-phase system.

In this presentation I will provide an overview of the governing equations and highlight the challenges these non-linear PDEs pose with respect to the spatial-temporal discretisation, and the linear/non-linear solvers applied to the resulting discrete system of equations. I will also discuss recent methodological developments underway within the magma group at Oxford, which aim to simulate the evolution of volatiles, such as carbon, in two-phase systems. The objectives of this work are to develop a first order understanding of how volatiles influence the thermal structure within the subduction wedge.