Booklet of
Participants and Abstracts

Eugen Haug / pixelio.de

Austrian HPC Meeting 2015 – AHPC15

March 15-20, 2015, Obergurgl, Austria

Welcome

We would like to welcome you all to Obergurgl, and we are looking forward to an interesting Austrian HPC Meeting 2015.

As a conference on computational science, as well as a workshop on the successful and productive use of state-of-the-art supercomputers, AHPC15 aims at bringing together scientists and technicians with a background and interest in scientific supercomputing. The meeting is hosted by the Research Focal Point Scientific Computing of the University of Innsbruck and takes place in the beautiful sports and alpine Research Center Obergurgl.

It highlights recent developments and trends for efficient supercomputer applications, as well as strategic issues for the future development of HPC in Austria. Also included is a one-day training on debugging with TotalView, MemoryScape and the ReplayEngine. The meeting intends to foster new insights and collaborations for a sustainable and flourishing HPC landscape in Austria.

We wish you a interesting and sporty time in Obergurgl. If you have any questions, please feel free to contact us.

The organizing team

Christoph Stöckl
Alexander Ostermann
Michael Fink
Sabine Kreidl

Universität Innsbruck
Scientific Computing
Austrian HPC Meeting 2015 - AHPC15

Technikerstraße 13
6020 Innsbruck
AUSTRIA

E-Mail: ahpc15@uibk.ac.at
Web: http://ahpc15.jimdo.com/

Participants

Alexander Michael, Technische Universität Wien, Austria
michael.alexander@tuwien.ac.at

Augustin Christoph, Medizinische Universität Graz, Austria
Christoph.augustin@medunigraz.at

Bilke Andreas, Universität Salzburg, Austria
abilke@cosy.sbg.ac.at

Bischko Jan Christoph, Universität Innsbruck, Austria
jan.bischko@uibk.ac.at

Blaas-Schenner Claudia, Technische Universität Wien, Austria
claudia.blaas-schenner@tuwien.ac.at

Dellago Christoph, Universität Wien, Austria
christoph.dellago@univie.ac.at

Draheim Dirk, Universität Innsbruck, Austria
dirk.draheim@uibk.ac.at

Einkemmer Lukas, Universität Innsbruck, Austria
lukas.einkemmer@uibk.ac.at

Elsässer Robert, Universität Salzburg, Austria
elsa@cosy.sbg.ac.at

Fink Michael, Universität Innsbruck, Austria
michael.fink@uibk.ac.at

Gamnitzer Peter, Universität Innsbruck, Austria
peter.gamnitzer@uibk.ac.at

Gunacker Patrik, Technische Universität Wien, Austria
patrik.gunacker@uibk.ac.at

Hofstetter Günter, Universität Innsbruck, Austria
guenter.hofstetter@uibk.ac.at

Huber Herbert, Leibniz Supercomputing Centre, Germany
Herbert.huber@lrz.de

Kandolf Peter, Universität Innsbruck, Austria
peter.kandolf@uibk.ac.at

Kracher Christian, Universität Wien, Austria
christian.kracher@univie.ac.at

Kutil Rade, Universität Salzburg, Austria
rkutil@cosy.sbg.ac.at

Langer Ulrich, Universität Linz, Austria
ulanger@numa.uni-linz.ac.at

Leitold Christian, Universität Wien, Austria
christian.leitold@univie.ac.at

Liedl Klaus, Universität Innsbruck, Austria
klaus.liedl@uibk.ac.at

Neumüller Martin, Universität Linz, Austria
martin.neumueller@jku.at

Ostermann Alexander, Universität Innsbruck, Austria
alexander.ostermann@uibk.ac.at

Pedit Helmut, Universität Innsbruck, Austria
helmut.pedit@uibk.ac.at

Pöll Martin, Universität Innsbruck, Austria
martin.poell@uibk.ac.at

Prugger Martina, Universität Innsbruck, Austria
martina.prugger@uibk.ac.at

Reichl Irene, Technische Universität Wien, Austria
irene.reichl@tuwien.ac.at

Reif Mathias, Universität Innsbruck, Austria
mathias.reif@uibk.ac.at

Resch Michael, Universität Stuttgart, Germany
resch@hlrs.de

Rotach Mathias, Universität Innsbruck, Austria
mathias.rotach@uibk.ac.at

Schauperl Michael, Universität Innsbruck, Austria
michael.schauperl@uibk.ac.at

Schindler Sabine, Universität Innsbruck, Austria
sabine.schindler@uibk.ac.at

Schreiner Wolfgang, Universität Linz, Austria
wolfgang.schreiner@risc.jku.at

Schulthess Thomas, Swiss National Supercomputing Center (CSCS), Switzerland
schulthess@cscs.ch

Schwarz Martin, Universität Innsbruck, Austria
martin.schwarz@uibk.ac.at

Steinhauser Dominik, Universität Innsbruck, Austria
dominik.steinhauser@uibk.ac.at

Stewart Dean, Rogue Wave Software, United Kingdom
dean.stewart@roguewave.com

Stöckl Christoph, Universität Innsbruck, Austria
christoph.stoeckl@uibk.ac.at

Stöhr Markus, Technische Universität Wien, Austria
markus.stoehr@tuwien.ac.at

Thaler Martin, Universität Innsbruck, Austria
martin.thaler@uibk.ac.at

Vajtersic Marian, Universität Salzburg, Austria
marian@cosy.sbg.ac.at

Wall Wolfgang A., Technische Universität München, Germany
wall@lnm.mw.tum.de

Wietek Alexander, Universität Innsbruck, Austria
alexander.wietek@uibk.ac.at

Yang Huidong, Universität Linz, Austria
huidong.yang@ricam.oeaw.ac.at

Zabloudil Jan, Technische Universität Wien, Austria
jan.zabloudil@tuwien.ac.at

Abstracts of Talks

Bridging time, length and energy scales in simulations of condensed matter

Christoph Dellago

The microscopic dynamical behavior of many condensed matter systems occurring in nature and technology is characterized by wide ranges of time, length and energy scales, which are a challenge for atomistic computer simulation. For instance, a supercooled liquid can exist in a metastable state without freezing for very long times before it eventually crystallizes. Typical waiting times vastly exceed the time scales accessible to molecular dynamics simulations. Similarly, the time scales of chemical reactions or of biomolecular reorganizations such as protein folding can be orders of magnitudes longer than those of basic molecular motions. Another major challenge arising in simulations of condensed matter systems is the calculation of accurate energies and forces required for determine the dynamics of the system. While the energies of chemical bonding are very large, the behavior of many systems is determined by small energy differences that are difficult to compute. Here, I will discuss several methods developed to address computational problems arising from wide ranges of time, length and energy scales, and illustrate these approaches with several examples ranging from freezing and cavitation to structural phase transitions.

ACOnet - High Performance Communication Infrastructure for Science and Research in Austria

Christian Kracher

ACOnet (http://www.aco.net/) is the Austrian Academic Computer Network for science, research, education and culture. ACOnet, legally represented by the University of Vienna, operates a national high-performance fibre-optic backbone network. It interconnects all Austrian public universities and some other 200 member institutions via GÉANT (http://www.geant.net/) to international academic networks, as well as to the commodity Internet. ACOnet offers to its participating organizations a cost-efficient data-highway for access to scientific database and academic HPC resources inside Austria, in Europe and abroad.

Does Flexibility Control the Specificity of Protein-Protein and Protein-Drug Interfaces?

Klaus Liedl

Specificity and promiscuity of intermolecular interfaces are typically studied and understood in terms of intermolecular interactions, i.e., enthalpy. However, conformational selection on binding suggests an important contribution originating from flexibility, i.e., entropy. We use proteases as prototypic protein-protein interfaces to define their specificity by analyzing their substrate tolerance. We established a sub-pocket-wise specificity metric named cleavage entropy relying on information theory. Thus, we are able to quantify and map specificity on to the intermolecular interface of proteases. Consequently, we compare this substrate-based and localized cleavage entropy with the local flexibility of the intermolecular interface.

We generate flexibility information by computer simulations and demonstrate that local backbone and side chain flexibility is highly correlated with the cleavage entropy observed by analyzing substrate promiscuity. E.g., a substitution of serine by proline in the binding pocket seems to induce specificity for caspases, even though a potential interaction via a hydrogen bond is lost. Also a correct ordering of the subpockets' specificity by only considering their flexibility is possible in the case of the serine protease thrombin. Other serine proteases of the chymotrypsin fold clan will be discussed as well as an inclusion of enthalpic contributions.

Finally, substrate-driven mapping of the degradome by comparison of sequence logos can be used to graft the complete protease tree by substrate specificity independent of evolutionary relationships. Application of this degradome tree to predict off-target effects of drugs is suggested.

„Specificity of a protein–protein interface: Local dynamics direct substrate recognition of effector caspases“, J.E. Fuchs, S.v. Grafenstein, R.G. Huber, H.G. Wallnoefer, K.R. Liedl, Proteins: Structure, Function and Bioinformatics 2014, 82, DOI: 10.1002/prot.24417

„Cleavage Entropy as Quantitative Measure of Protease Specificity“, J.E. Fuchs, S.v. Grafenstein, R.G. Huber, M.A. Margreiter, G.M. Spitzer, H.G. Wallnoefer, K.R. Liedl, PLoS Computational Biology 2013, 9(4) e1003007, DOI: 10.1371/journal.pcbi.1003007

„Substrate Sequences Tell Similar Stories as Binding Cavities: Commentary“, J.E. Fuchs, K.R. Liedl, Journal of Chemical Information and Modeling, 2013, 53(12), 3115-3116, DOI: 10.1021/ci4005783

„Substrate-Driven Mapping of the Degradome by Comparison of Sequence Logos“, J.E. Fuchs, S.v. Grafenstein, R.G. Huber, C. Kramer, K.R. Liedl, PLoS Computational Biology, 2013, 9(11) e1003353, DOI: 10.1371/journal.pcbi.1003353

HPC und Gesellschaft

Michael Resch

Ein Gericht in l'Aquila in Italien verurteilte am 22. Oktober 2012 sechs Wissenschaftler zu hohen Haftstrafen. Begründung war, dass die Experten trotz kleinerer Vorbeben ein größeres Erdbeben ausgeschlossen hätten und damit Mitschuld an der hohen Opferzahl beim Erdbeben vom 6. April 2009 trügen. Hintergrund des Prozesses war eine Diskussion über die wissenschaftliche Prognosefähigkeit bei Erdbeben, die sich im Wesentlichen auf Empirie und Simulation stützt. Tausende italienische Wissenschaftler protestierten mit einer Unterschriftensammlung gegen die Verknüpfung von Wissenschaft und Politik.

Auch wenn das Urteil für die meisten Beschuldigten am 10. November 2014 wieder aufgehoben wurde, lohnt es sich, darüber nachzudenken was hier passiert ist. Zunächst scheint klar, dass ein Sündenbock für das Erdbeben gesucht wurde. Genauere Untersuchungen zeigen jedoch, dass Simulationen auf schnellen Rechnern eine unklare Rolle gespielt haben. Und zunehmend werden derartige Simulationen in einer Reihe von Gebieten als Grundlage für politische und gesellschaftliche Entscheidungen herangezogen. Während das im Bereich des Wetters durchaus meist unproblematisch ist, gibt es eine Reihe von Bereichen in denen Probleme vorhanden sind oder dabei sind zu entstehen. Angesicht der zunehmenden Leistung auch kleinerer Systeme und angesichts der Tendenz, Simulation im politischen Entscheidungsprozess verstärkt einzusetzen, wirft dieser Vortrag einen Blick auf die allgemeinen Vorstellungen die vom HPC in einer Gesellschaft vorhanden sind und diskutiert die Auswirkungen die solche Vorstellungen haben. Eingegangen wird dabei auch auf Möglichkeiten, mit solchen Vorstellungen umzugehen bzw. sie zu beeinflussen.

Reflecting on the “P” in HPC

Thomas Schulthess

Floating point performance has traditionally been the primary concern in scientific computing, which is why benchmarks like High-Performance Linpack (HPL) and the related top500 list are so popular in High-Performance Computing (HPC). However, more recent programs, such as DARP’s High-Productivity Computing Systems projects, are concerned with productivity as well. Furthermore, there has been quite some debate over the metrics used to determine if a computing systems performs well. In the present lecture I will argue why one should go back to generic measures such as time and energy to solution. While time is straightforward to comprehend in terms of complexity analysis, the energy consumption of a simulation is difficult to measure and hard to reason about. Nevertheless, with examples from climate and quantum materials science simulations, I will show how a systematic analysis of energy to solution can be used to optimize computing systems, including software and architectural choices. As for the discussion of productivity, I will argue that this requires a profound rethinking of the way we develop simulations codes in HPC.

The first one hundred days of VSC-3

Markus Stöhr

Presentation of the latest VSC cluster - the VSC-3. The key features of the new installation in Vienna are the very efficient oil submersion cooling and the dual rail infiniband network. We will present our experiences and problems with these technologies and give an overview of other capabilities/features of VSC-3.

HPC based on high performance modeling in Engineering Sciences and beyond

Wolfgang A. Wall

High Performance Computing (HPC) nowadays plays a crucial role in Science and Technology and tremendous achievements are continuously realized in HPC. But sometimes there is a tendency to see HPC topics a bit isolated or to see HPC as the goal itself. But looking at it from the perspective of Engineering Sciences or the Applied Sciences, FLOPS and Peak Performance are very important, but just to the point that they carry over to the real goal. And as Richard Hamming nicely stated in the 60's, "The purpose of computing is insight, not numbers". Hence, the key interest is in the "insight factor" (IF), i.e. "insight" over "real time". It is obvious that the best parallelization strategies, cache-aware implementations, etc. are all crucial ingredients to make the IF as high as possible. But it is also important to note that sometimes more complex discretization schemes or algorithms, that even might not fit the processor well, are providing more insight per FLOP and hence one is willing to pay the prize of a lower peak performance. And besides all of this, one should never forget the "biggest hammer of all", i.e. the role of an adequate model.

In this presentation I will try to give some coupled multifield and/or multiscale examples from different fields of the Engineering Sciences but also from some other fields and specifically will try to elaborate on all three aspects, i.e. the models, the methods/algorithms and the implementations - all of them as high performance as possible in order to increase the IF.

The VSC project

Jan Zabloudil

The presentation will describe how the VSC project has originated and developed and give an overview of its organizational structure. A brief technical description of the first two HPC Clusters

VSC-1 and VSC-2 - compute nodes, network, and cooling concepts - will be given and experiences with these systems in recent years will be discussed. In addition I will explain the necessary prerequisites for getting access to the VSC and explain the management of projects.

Contributed Talks

Strongly scalable algorithms for the modeling of human whole heart electromechanics

Christoph Augustin, Aurel Neic, Anton Prassl and Gernot Plank

Developing efficient frameworks for multiscale-multiphysics models of cardiac electromechanics is a challenging task due to the vast computational costs and the numerical complexities involved in coupling different physics.

Due to preferential orientations of fibers, such as collagen and myocytes, the passive mechanical behavior of myocardial tissue is anisotropic and highly nonlinear. Moreover, an active contraction is driven by electrical processes, which is modeled by an additional active stress contribution. This additive term is computed using a particular cell model and a cardiac electrophysiology model, based on the bidomain equations. To specify prescribed displacements or tractions Dirichlet and Neumann

boundary conditions are incorporated. Windkessel models of the circulatory system are fitted to match pressure-volume relations during ejection and diastolic filling phases.

In order to obtain a numerical solution we apply variational and finite element techniques and the nonlinear system is linearized using Newton's method.

However, such detailed multiphysics simulations are computationally vastly demanding, since the resulting discretized systems of equations are large, with up to 10^9 degrees of freedom.

While current trends in high performance computing (HPC) hardware promise to alleviate this problem, exploiting the potential of such architectures remains challenging for various reasons. On one hand, strongly scalable algorithms are necessitated to achieve a sufficient reduction in execution time by engaging a large number of cores, and, on the other hand, alternative acceleration technologies such as graphics processing units (GPUs) are playing an increasingly important role which imposes further constraints on design and implementation of solver codes.

We present a domain decomposition algebraic multigrid (AMG) preconditioner for an iterative Krylov solver (CG, GMRES) which is custom-tailored for the specific problem. Detailed benchmarking studies were performed on the HPC facility SuperMUC (within the PRACE project CAMEL). Strong scalability results up to 8192 cores for electromechanical simulations will be shown and we discuss advantages and limitations of the particular models and numerical methods.

References:

CM Augustin, GA Holzapfel, and O Steinbach. “Classical and all-floating FETI methods for the simulation of arterial tissues”. In: Int. J. Numer. Meth. Engrg. 99.4 (2014), pp. 290–312.

CM Augustin, A Neic, M Liebmann, AJ Prassl, SA Niederer, G Haase, and G Plank. “Anatomically accurate high resolution modeling of cardiac electromechanics: a strongly scalable algebraic multigrid solver method for non-linear deformation”. In: Biomechanics and modeling in mechanobiology (submitted).

CM Augustin and G Plank. “Simulating the mechanics of myoardial tissue using strongly scalable parallel algorithms”. In: Biomed Tech (Berl). Vol. 35. 12. 2013.

From Community Detection to Graph Models in Social Networks

Andreas Bilke, Robert Elsässer

In social network science there is growing interest in the analysis of large real-world social graphs. Since crawling existing well-known social networks like Facebook or Google+ is not allowed due to privacy issues, there is a need for artificial networks which have similar structural properties, e.g. degree distribution or community structure. A social network is usually modeled by a graph, where the nodes represent individuals and there is an edge between two nodes if the corresponding individuals are connected in the network.

We first analyze the (overlapping) community structure of two different real-world social networks. Then, we study the resulting communities using the eigenvalues of the corresponding graphs. Finally, we develop a randomized model for constructing graphs, which are able to capture several structural and algorithmic properties of real-world social networks. To obtain the communities as well as to analyze their structural properties, we used an MPI implementation which run on our university cluster doppler.

Efficient numerical solution of the Vlasov equation

Lukas Einkemmer

The numerical solution of the Vlasov equation poses a number of challenges both from an algorithmic as well as from an HPC standpoint. Advances in both areas are needed in order to simulate many interesting physical phenomena (which are posed in an up to 6 dimensional phase space).

The recently developed semi-Lagrangian discontinuous Galerkin methods are considered an attractive alternative to more traditional approaches. These schemes are conservative, preserve a local approximation, and limit numerical diffusion.

In this talk we will focus on the parallel implementation of these semi-Lagrangian discontinuous Galerkin methods on distributed memory clusters. In this context, scaling studies on the Vienna Scientific Cluster 2 (VSC-2) will be presented. The framework discussed here provides a dimension independent implementation (based on C++ templates) that can be used for constructing scientific computer codes.

Development of a parallel Simulation Environment for multiphase problems

Peter Gamnitzer

Structures of engineering interest are often made of porous materials like sand, clay, or concrete. They can be successfully described using finite element approaches based on a multiphase description of the underlying continuum. These approaches result in the iterative solution of a coupled nonlinear system of balance, kinematic, and constitutive equations. The efficient simulation of such problems thus poses several multiphase specific demands on the software and the models used. This contribution aims at the discussion of the impact of these requirements on i) the specific code structure implemented in our parallel multiphase FEM software, highlighting the benefits of the new code framework in comparison to our implementation of the multiphase models in OpenSees (a third-party finite element open-source software we used previously), ii) the robust formulation of complex material laws for the use in parallel simulations, and iii) the structure of the linear system and future prospects for an improved handling of the linear sub-problem.

The Leja method in the context of HPC

Peter Kandolf

The Leja method is a well established scheme for computing the action of the matrix exponential, see [1]. The approximation of this action is a computationally involved tasked heavily depending on the matrix-vector multiplication. This task can be highly parallelized on graphical processing units.

The aim of this project is to develop a state-of-the-art implementation on GPUs.

References

[1] Caliari, M., Kandolf, P., Ostermann, A., Rainer, S., 2014. Comparison of software for computing the action of the matrix exponential. BIT Numerical Mathematics, 54 (1), 113-128.

An Object-Oriented Programming Framework for MPI Matrix Algorithms

Rade Kutil

An approach previously reported at AHPC is extended from shared memory to MPI. Its automated parallelization efficiency is evaluated. Parallelizing matrix algorithms can be an arduous task because keeping track of the data distribution as it propagates through the operations is not easy. Together with memory management and communication operations, programs can become confusingly complex. This work aims at freeing the programmer from these obligations so that they can concentrate on on their algorithms. Computations are performed in parallel without requiring the user to handle threading and MPI. The idea is to decompose basic matrix operations into communication-free sub-operations. For instance, the multiplication of a m x k and a k x n matrix is not free of communication if the data is split in the k dimension. By moving the summation over divisions of the k-dimension into a separate operation, the communication is transferred to data redistribution steps between sub-operations, which can be handled by the framework.

Sub-operations operate on “slabs", i.e. data slices of arbitrary dimensions. The user can use predefined basic operators, or they can also easily add new operators by inheriting the operator class and overloading the constructor and the operate method. The constructor specifies input and output operands and how their dimensions correspond. The operate method performs the actual operation on slabs delivered by the framework. As the sub-operations are communication-free, specifying the operation on slabs is sufficient.

Additionally, optimal divisions of matrices can be found so that the computational and communication cost is minimized. Cost functions estimating the performance of the sub-operations are provided by operator classes.

To enable all this, the framework has to capture matrix calculations in symbolic form. The symbolic form of a matrix calculation is constructed by performing the desired operations, e.g. multiplication and addition, not on data itself but on place-holder objects whose job it is to instantiate the corresponding operator objects and connect their input and output operands. Finally, a calculation is a sequence of connected operators which are executed successively on distributed slabs of operands. Thus, the user only has to program the sequence of matrix operations, provide matrix dimensions and initial data, and start the parallel calculation.

A string reaction coordinate for the folding of a polymer chain

Christian Leitold, Wolfgang Lechner, Christoph Dellago

Using extensive molecular-level simulations on the Vienna Scientific Cluster, we investigate the crystallization mechanism of a single, flexible homopolymer chain with short range at tractions. For a sufficiently narrow attractive well, the system undergoes a first-order like freezing transition from an expanded disordered coil to a compact crystalline state. Based on a maximum likelihood analysis of committor values computed for configurations obtained by Wang–Landau sampling, we construct a non-linear string reaction coordinate for the coil-to-crystal transition. In contrast to a linear reaction coordinate, the string reaction coordinate captures the effect of different degrees of freedom controlling different stages of the transition. Our analysis indicates that a combination of the energy and the global crystallinity parameter Q6 provide the most accurate measure for the progress of the transition. While the crystallinity parameter Q6 is most relevant in the initial stages of the crystallization, the later stages are dominated by a decrease in the potential energy.

Parallel space-time multigrid methods

Martin Neumüller

For evolution equations we present a space-time method based on Discontinuous Galerkin finite elements. Space-time methods have advantages if we have to deal with moving domains and if we need to do local refinement in the space-time domain. For this method we present a multigrid approach based on space-time slabs. This method allows the use of parallel solution algorithms. In particular it is possible to solve parallel in time and space. Furthermore this multigrid approach leads to a robust method with respect to the polynomial degree which is used for the DG time stepping scheme. Numerical examples for some optimal control problem and for the Heat and Stokes equations will be given, which show the parallel performance of this space-time multigrid approach.

Carbon Dioxide Exchange in complex topography

Matthias Reif, Mathias Rotach, Georg Wohlfahrt, Alexander Gohm

On a global scale the budget of carbon dioxide (CO2) bears a quite substantial uncertainty, which is commonly understood to be mainly due to land-surface exchange processes. In this project we investigate to what extent complex topography can amplify these land-surface exchange processes. The hypothesis is that, on the meso-scale, topography adds additional atmospheric mechanisms that drive the exchange of CO2 at the surface.

This sensitivity model study investigates an idealized sine shaped valley with the atmospheric numerical model Weather Research and Forecasting (WRF) coupled to the community land model (CLM) to study the effect of complex topography on the CO2 budget compared to flat terrain. The experiment is designed to estimate the effect of the topography during maximum ecosystem exchange in summer using meteorological and ecosystem conditions at solstice, the 21. of June. Systematic variation of meteorological initial conditions, plant functional types and the topography creates an ensemble that unveils the fundamental factors that dominate the differences of CO2 between simulations with topography compared to plain surfaces in the model.

The sign and magnitude of the difference between the CO2 exchange over topography and over a plain simulation are strongly dependent on the CLM plant functional type, the initial temperature, the initial relative humidity, the latitude and the area height distribution of the topography. However, in this model experiment the topography is, in the mean, a sink to the CO2 budget in the order of 5% per day.

Needs for high-performance computing in applications of atmospheric science

Mathias W Rotach, Alexander Gohm

With the early ‚numerical weather prediction' atmospheric sciences have certainly been standing at the cradle of the development of HPC. Even if other applications may in the meantime have become even more (computing) time consuming atmospheric applications (e.g., weather forecast and climate scenario modeling) are still among the major players for HPC resources and their products largely limited by computational efficiency and availability of resources. We illustrate the main fields of application in atmospheric sciences, the latest developments as well as the role technical improvements (as opposed to physical developments) may play in the future. Also, some of our own research directions and results addressing atmospheric flows in complex, mountainous terrain will be presented.

Antifreezing Protein Dynamics with Polarizable Multipole Force Fields

Michael Schauperl, C. Kramer, K. R. Liedl

In the last two decades a new class of force fields, polarizable multipole force fields, such as AMOEBA (Atomic Multipole Optimized Energetics for Biomolecular Applications) has become increasingly more important. In order to enhance the performance of classical (fixed point charge) force fields, a new treatment of the electrostatics is crucial. To gain a better representation of the real electrostatic behavior of molecules this new force fields include fixed multipoles on each atom providing a higher basis set for the description of the electron density. In this work the AMOEBA force field was used, which covers polarization effects through an induced dipole on every atom [1]. Polarization effects, which are not covered by conventional force fields, are of major interests as atoms are able to react upon their environment [2] resulting in a force field which can in an ideal case be used to describe molecules in vapor and solution. The precise calculation of the electrostatics is very demanding regarding computational time, which makes high performance computing systems necessary for these simulations.

We show how polarizable multipole force fields improve the dynamic description of proteins in water compared to conventional force fields. To validate that, AMOEBA can enhance the results of molecular dynamics simulations by treating the electrostatics in a more accurate way; the simulations of freezing processes and water ordering effects of antifreezing proteins [3] are analyzed.

[1] J. W. Ponder et al., J. Phys. Chem. B, 2010, 114, 2549-2564.

[2] A. Warshel, M. Kato, A. V. Pisliakov, J. Chem. Theory. Comput., 2007, 3, 2034-2045.

[3] Sicher F, Yang DSC. Nature. 1995; 375:427-431

Modeling the evolution of galaxies: N/body and hydrodynamic simulations on a moving mesh

Dominik Steinhauser

We present results of the numerical simulation group of the Institute for Astro- and Particle Physics at the University of Innsbruck.
Especially, we focus on galaxy scale simulations, presenting our new set-up for galaxies undergoing ram-pressure stripping in a realistic cluster environment, using the moving mesh code AREPO.
Besides of this new approach for solving hydrodynamics, we also present our new implementation of a N-body tree code, running entirely on GPUs.

Exact Diagonalization for Strongly Correlated Electron Systems on HPC Clusters

Alexander Wietek

Strongly correlated electron systems may exhibit fascinating novel physics like high temperature superconductivity or quantum spin liquids. In order to theoretically investigate these systems numerical approaches like the Exact Diagonalization method are commonly applied. This method involves the calculation of eigenvalues and eigenvectors of matrices on spaces with up to several hundreds of billion dimensions.
We propose methods how to efficiently deal with the challenges that arise in these computations, especially on distributed memory machines. Furthermore we present a new Code developed by us incorporating these ideas and benchmarks on VSC3.

Robust and Efficient Monolitic Fluid-Structure-Interaction Solvers and their parallelization stratigies

Huidong Yang, Ulrich Langer

In this talk, we construct robust and efficient preconditioner Krylov subspace solvers for the monolithic linear system of algebraic equations arising from the finite element discretitzation and Newton's liniarization of the fully coupled fluid-structure interaction system of Partial Differential Equations in the Arbitrary Lagrangian-Eulerian formulation. We admit nonlinear hyperelastic material in the solid model and cover a large range of flows, e.g,, water, blood, and air with highly varying density. The robust preconditioners are constructed in form of $\hat{L}\hat{D}\hat{U}$, where $\hat{L}, \hat{D}$ and $\hat{U}$ are proper approximations to the matrices L, D and U in the LDU factorization of the fully coupled system matrix, respectively. The inverse of the corresponding Schur complement is approximated by applying one cycle of a special class of algebraic multigrid methods to the perturbed fluid sub-problem, that is obtained by modifying corresponding entries in the original fluid matrix with an explicitly constructed approximation of the exact perturbation coming from the sparse matrix-matrix multiplications. The numerical studies presented impressively demonstrate the robustness and the efficiency of the preconditioners proposed in the talk.

In addition, we will also discuss possible strategies on how to implement the parallel solvers for FSI problems, mainly using Domain Decomposition methods.