IT4Innovations
national01$#&0
supercomputing
center@#01%101

HPCSE 2015

High Performance Computing in Science and Engineering

Hotel Soláň, May 25-28, 2015

organized by the IT4Innovations National Supercomputing Center,

VSB-Technical University of Ostrava

Collection of abstracts

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

This project is supported by the ESF and the Government of the Czech Republic.

New creative teams in priorities of scientific research, CZ.1.07/2.3.00/30.0055.

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

Preface

The aim of the conference High Performance Computing in Science and Engineering 2015, organized by the IT4Innovations National Supercomputing Center, VSB-Technical University of Ostrava, is to bring together foreign specialists on applied mathematics, numerical methods, and parallel computing, exchange the common experience, and initiate new research cooperation. We are glad that our invitation was accepted by distinguished experts from leading world research institutions. We hope that the conference will motivate new students and researchers to join the growing community focused on supercomputing.

The event takes place in the heart of the Beskydy Mountains, in one of the most beautiful parts of Wallachia. The hill of Soláň offers a tranquil atmosphere and a spectacular view of the Javorníky Mountains. For many years it was a favorite location of Czech artists.

The conference is supported by the IT4Innovations National Supercomputing Center and the project OP EC "New creative teams in priorities of scientific research" no. CZ.1.07/2.3.00/30.0055. This project is supported by Operational Program Education for competitiveness funded by Structural Funds of the European Union and state budget of the Czech Republic. It is our pleasure to acknowledge this support.

We would like to wish HPCSE 2015 to be a fruitful event, providing interesting lectures, showing new ideas, beauty of applied mathematics, numerical linear algebra, optimization methods, and high performance computing and starting or strengthening collaboration and friendship.

On behalf of the Programme and Organizing Committee of HPCSE 2015,
Radim Blaheta, Zdeněk Dostál, Tomáš Kozubek, Miroslav Tůma, Jakub Šístek,
Zdeněk Strakoš, Vít Vondrák

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

About IT4Innovations

IT4Innovations National Supercomputing Center at VSB-Technical University of Ostrava (http://www.it4i.eu) is the strategic research infrastructure in the Czech Republic. In the framework of the national supercomputing center and its mission, the most powerful supercomputing facilities in the Czech Republic are being built and excellent research in HPC technologies and embedded systems is performed. The center is providing state-of-the-art technology and expertise in high performance computing and embedded systems and makes it available for Czech and international research teams from academia and industry. It aspires to improve the quality of life, to increase the competitiveness of industrial sector and to promote the cross-fertilization of high-performance computing, embedded systems and other scientific and technical disciplines. The center conducts research in the areas of earth science simulations for disaster management, traffic management, numerical modelling for engineering, physics and chemistry, development of libraries for parallel computing, modelling for nanotechnologies, soft-computing methods, recognition and presentation of multimedia data and safe and reliable architectures, networks and protocols. The center has also strong link to industrial sector. Currently, the largest Intel® Xeon Phi™ co-processor-based cluster in Europe with the peak performance 2 PFLOP/s is being installed at IT4Innovations National Supercomputing Center premises.

List of speakers

Invited speakers

Owe Axelsson	Institute of Geonics AS CR, Ostrava, Czech Republic
Cevdet Aykanat	Bilkent University, Ankara, Turkey
Marc Baboulin	INRIA, Paris, France
Santiago Badia	CIMNE, Barcelona, Spain
Jed Brown	Argonne National Laboratory, Illinois, USA
Fehmi Cirak	University of Cambridge, United Kingdom
Jacek Gondzio	University of Edinburgh, United Kingdom
Akhtar A. Khan	Rochester Institute of Technology, Rochester, USA
Johannes Kraus	University of Essen, Germany
Jaroslav Kruis	Czech Technical University in Prague, Czech Republic
Julien Langou	University of Colorado, Denver, USA
Dan Negrut	University of Wisconsin-Madison, Madison, USA
Ulrich Rüde	University of Erlangen, Germany
Valeria Simoncini	University of Bologna, Italy
Wim Vanroose	University of Antwerpen, Belgium
Barbara Wohlmuth	Technical University Munich, Germany
Walter Zulehner	Johannes Kepler University, Linz, Austria

Special keynote speakers

Roman Wyrzykowski	Czestochowa University of Technology, Poland
Jan Mandel	University of Colorado, Denver, USA

Other speakers

Simplice Donfack	Universita della Svizzera italiana, Lugano, Switzerland
Václav Hapla	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Jiří Jaroš	Brno University of Technology, Brno, Czech Republic
Lubomír Říha	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Patrick Sanan	Universita della Svizzera italiana, Lugano, Switzerland
Zdeněk Strakoš	Charles University in Prague, Czech Republic
Stanislav Sysala	Institute of Geonics AS CR, Ostrava, Czech Republic
Jakub Šístek	Institute of Mathematics AS CR, Prague, Czech Republic
Dalibor Štys	University of South Bohemia in České Budějovice, Czech Republic
Miroslav Tůma	Institute of Computer Science AS CR, Prague, Czech Republic
Vít Vondrák	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Jan Zapletal	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

Scientific and organizing committees

Scientific committee

Radim Blaheta	Institute of Geonics AS CR, Ostrava, Czech Republic
Zdeněk Dostál	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Tomáš Kozubek	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Miroslav Tůma	Institute of Computer Science AS CR, Prague, Czech Republic
Jakub Šístek	Institute of Mathematics AS CR, Prague, Czech Republic
Zdeněk Strakoš	Charles University in Prague, Czech Republic
Vít Vondrák	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic

Organizing committee

Tomáš Kozubek	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Martin Čermák	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Lucie Bestová	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic
Ivana Bieliková	IT4Innovations, VSB-Technical University of Ostrava, Czech Republic

Programme

Monday, May 25, 2015

10:00 – 12:00	*Registration of participants, discussion*
12:00 – 13:50	*Lunch*
13:50 – 14:00	*Opening*
14:00 – 14:40	Johannes Kraus: **Stable discretization and robust iterative solution methods for the Brinkman problem**
14:40 – 15:20	Dan Negrut: **From Granular Dynamics to Fluid-Solid Interaction, and from Handling Large Optimization Problems to Solving Sparse Linear Systems**
15:20 – 16:00	Jed Brown: **HPGMG: Relevant benchmarking for scientific computing**
16:00 – 16:30	*Coffee break*
16:30 – 18:00	Vít Vondrák: **IT4Innovations National Supercomputing Center**
	Václav Hapla: **PERMON toolbox**
	Lubomír Říha: **ExaScale PaRallel FETI SOlver (ESPRESO)**
	Jan Zapletal: **BEM4I - A Parallel Library for Boundary Element Discretization of Engineering Problems**
	Simplice Donfack: **Improving the applicability of the highly efficient compilers to a wider class of problems**
	Zdeněk Strakoš: **On the Vorobyev method of moments in the context of Krylov subspace methods**
19:00	*Welcome Grill Party*

Tuesday, May 26, 2015

07:00 – 09:00	*Breakfast*
09:00 – 09:40	Ulrich Rüde: **Coupled Models for Extreme-Scale Computing**
09:40 – 10:20	Jacek Gondzio: **Preconditioners for higher order methods in big data optimization**
10:20 – 10:50	*Coffee break*
10:50 – 11:30	Wim Vanroose: **Rounding errors in pipelined Krylov Methods**
11:30 – 12:10	Jan Mandel: **Data Assimilation of Satellite Active Fire Detection in Coupled Atmosphere-Fire Simulation**
12:10 – 12:50	Roman Wyrzykowski: **Using GPU and Intel Xeon Phi coprocessors to accelerate stencil algorithms**
12:50 – 14:00	*Lunch*
14:00 – 19:00	*Networking time (trip to Wallachian Open Air Museum, hiking, or wellness)*
19:00 – 20:00	*Dinner*
20:00	*Poster session with short presentation per poster*

Wednesday, May 27, 2015

07:00 – 09:00	*Breakfast*
09:00 – 09:40	Owe Axelsson: **Title TBC saddle point systems Brinkman flow**
09:40 – 10:20	Cevdet Aykanat: **Optimizing right communication cost metrics to scale sparse kernels on distributed memory systems**
10:20 – 11:00	*Coffee break*
11:00 – 11:40	Akhtar A. Khan: **Computational Methods for Elastography Inverse Problem**
11:40 – 12:20	Barbara Wohlmuth: **From optimal algorithms to fast solvers**
12:30 – 14:00	*Lunch*
14:00 – 14:40	Fehmi Cirak: **Large-scale computational fluid-membrane interaction with application to supersonic decelerators**
14:40 – 15:20	Marc Baboulin: **The story of the butterflies**
15:20 – 16:00	Santiago Badia: **Multilevel domain decomposition at extreme scales**
16:00 – 16:30	*Coffee break*
16:30 – 18:00	Jakub Šístek: **Parallel performance of iterative solvers for pressure-correction methods for incompressible flows**
	Dalibor Štys: **Simulation of the state space trajectory: are concrete implementations by necessity multilevel cellular automata?**
	Stanislav Sysala: **How to Simplify Return-Mapping Algorithms in Computational Plasticity**
	Miroslav Tůma: **Preconditioning of symmetric indefinite systems**
	Jiří Jaroš: **Large scale ultrasound simulations**
	Patrick Sanan: **Aggressive Accelerator-enabled Local Smoothing via Incomplete Factorization, with Applications**
19:00	*Conference dinner*

Thursday, May 28, 2015

07:00 – 09:00	*Breakfast*
09:00 – 09:40	Walter Zulehner: **Hilbert space interpolation as a tool for analyzing and preconditioning mixed variational problems**
09:40 – 10:20	Julien Langou: **Improving the communication lower bounds for matrix-matrix multiplication**
10:20 – 11:00	*Coffee break*
11:00 – 11:40	Valeria Simoncini: **Multiterm linear matrix equations and the numerical solution of (S)PDEs**
11:40 – 12:20	Jaroslav Kruis: **High performance computing in multi-scale problems**
12:30 – 14:00	*Lunch*
14:00 – 17:00	*Discussion, closing*

Abstracts - Invited speakers

Owe Axelsson, Radim Blaheta, and Martin Hasal – A comparison of three preconditioning methods for solving saddle point problems with an application for modelling fluid flow around porous media

The choice of a preconditioner for a matrix is mainly influenced by the two factors, ease of implementation of the preconditioner and efficiency of the convergence of Krylov subspace methods for the resulting spectrum of the preconditioned operator.
During the years many preconditinoning methods for saddle point problems have been presented. In this paper three methods which either give an easy implementation but somewhat less efficient spectrum or give a very efficient spectrum but with more involved implementation are presented.
The preconditioning methods are:
(i) The HSS method, a matrix factorization based on a Hermitian–skewsymmetric splitting.
(ii) Factorization of a modified saddle point matrix that results in a block triangular preconditioned matrix, with a controllable upper triangular off-diagonal block matrix.
(iii) A special block augmented matrix triangular preconditioner.
Simple proofs are given to show that method (i) gives a clustering of eigenvalues near 0 and 1, method (ii) gives a somewhat more accurate block triangular preconditioned matrix than the standard block triangular preconditioner and method (iii) gives a highly efficient clustered spectrum near unity for large values of the augmented matrix coefficient.
Hence methods (i) and (ii) require many more outer iterations than method (iii). However, normally method (i) and (ii) are easier to implement.
It is also shown how the methods (ii) and (iii) can be improved by use of explicit approximations of the inverse of the given pivot matrix and how method (iii) can be

implemented by use of a Cholesky incomplete factorization, by element by element assembly or by use of a projection matrix framework.

Oguz Selvitopi, Seher Acer, and Cevdet Aykanat – Optimizing right communication cost metrics to scale sparse kernels on distributed memory systems

Sparse kernels are characterized with low computational density, making them difficult to scale on large-scale parallel systems. Partitioning of matrices with the aim of reducing communication costs in these kernels is a common strategy employed for efficient parallelization. Most approaches aim at minimizing total volume while maintaining computational balance. However, different sparse kernels have different communication requirements, requiring other cost metrics also to be taken into account. To this end, we discuss combinatorial models and techniques to meet the specific communication requirements of different sparse kernels. For sparse methods that involve sparse matrix-vector multiplication (SpMV), we consider minimizing latency overhead, which especially becomes critical in the case of strong scaling. The first approach is tailored to iterative solvers for symmetric linear systems and is applicable to a wide range of Krylov methods (such as CG, basic GMRES, etc.). It uses the observation that each SpMV is followed by an inner product that meets a certain dependency and exploits it for a computation and communication rearrangement scheme. Via these rearrangement schemes, maximum number of messages communicated per processor becomes lgP for a system with P processors, granting a bound on the latency overhead. The second approach is tailored to iterative solvers for nonsymmetric linear systems (such as CGNE, CGNR, etc.). These solvers accommodate more potential for addressing latency-based cost metrics due to the flexibility of adopting different partitions on the input and output vectors of SpMV. Here, we investigate hypergraph-partitioning-based combinatorial models which are applied on top of any readily available 1D and 2D partitioning methods. These models try to address both latency and bandwidth costs by aiming to minimize total number of messages communicated and maintain a balance on communication loads of processors. A wide range of experiments up to 8K processors using the widely adopted PETSc toolkit are discussed by comparing eight partitioning models. Another important sparse kernel we consider is the sparse matrix dense matrix multiplication (SpMM), a common operation used in block iterative methods. This kernel on distributed systems is characterized with its high communication volume requirements, making bandwidth costs more important than latency costs. For this kernel, we discuss a generic framework that can minimize multiple and different volume-based communication costs metrics, such as total volume, maximum send volume, maximum receive volume, etc. This framework consists of combinatorial graph and hypergraph models utilized within the context of recursive bipartitioning paradigm and can handle multiple cost metrics simultaneously in any state-of-the-art partitioner. The models are quite lightweight, introducing only an additional complexity proportional to the number of nonzeros in the matrix to each bipartitioning level.

Marc Baboulin – The story of the butterflies

We present an overview of how random butterfly transformations (RBT) can accelerate the solution of linear systems by preventing the communication overhead due to pivoting. We explain how it has been successfully applied to dense general and symmetric indefinite linear systems, resulting in efficient RBT solvers for current parallel architectures, including multicore, GPU and Intel Xeon Phi. We also present experiments on direct sparse factorizations (serial and parallel). In particular we describe how RBT can be combined with sparsity-preserving strategies.

Santiago Badia, Alberto F. Martín, and Javier Principe – Multilevel domain decomposition at extreme scales

We present a fully-distributed, communicator-aware, recursive, and interlevel-overlapped message-passing implementation of the multilevel balancing domain decomposition by constraints (MLBDDC) preconditioner. The implementations highly rely on subcommunicators in order to achieve the desired effect of coarse-grain overlapping of computation and communication, and communication and communication among levels in the hierarchy (namely inter-level overlapping). Essentially, the main communicator is split into as many non-overlapping subsets of MPI tasks (i.e., MPI subcommunicators) as levels in the hierarchy. Provided that specialized nodes resources and memory are devoted for each level, a careful re-scheduling and mapping of all the computations and communications in the algorithm lets a high degree of overlapping to be exploited among levels. All subroutines and associated data structures are expressed recursively, and therefore MLBDDC preconditioners with an arbitrary number of levels can be built while re-using significant and recurrent parts of the codes. This approach leads to excellent perfect scalability results as soon as level-1 tasks can mask coarser-levels duties, which can be attained properly choosing coarsening rations and number of levels. We have provided a model to indicate how to choose the number of levels and coarsening ratios between consecutive levels and determine qualitatively the scalability limits for a given choice. We have carried out a comprehensive weak scalability analysis of the proposed implementation for the Laplacian and linear elasticity problems. Excellent weak scalability results have been obtained up to 458,752 processors and 1.8 million MPI tasks.

Jed Brown – HPGMG: Relevant benchmarking for scientific computing

High-performance LINPACK (HPL) is currently the dominant benchmark for comparing supercomputers, but is recognized as not being representative of many real applications. Premier supercomputing centers have extensive procurement suites intended to represent the needs of their intended workload, but it requires a great deal of effort to compare machines using these suites. Meanwhile, many applications will face performance degradation on newer "more powerful" machines. Scientists and engineers need a way to evaluate whether a machine is merely "difficult to program" or in fact ill-suited to their target application, lest they invest undue effort in development that does not improve their production requirements. In this talk, I will introduce our new HPGMG benchmark, its algorithmic characteristics, and interpretation of the performance spectrum from the perspective of common scientific and engineering workloads.

Fehmi Cirak – Large-scale computational fluid-membrane interaction with application to supersonic decelerators

Supersonic membrane decelerators, like parachutes or inflatable aero-shells, are presently used for planetary reentry in space exploration missions. It is known from flight tests and experiments that such membrane decelerators are susceptible to aeroelastic instabilities in certain flight regimes. To study this, we developed a simulation technique by combining several novel developments in membranes, fluid-membrane interaction, fluid dynamics and block-structured adaptive mesh refinement. The compressible, inviscid fluid flow is discretised with a shock-capturing Eulerian finite volume technique and an adaptive block-structured Cartesian grid. The membrane is discretised with a Lagrangian subdivision finite element technique and an unstructured surface mesh. Both discretizations are advanced in time with an explicit integration scheme and their coupling is achieved with a staggered approach using a variant of the ghost-fluid method. The fluid and membrane solvers are parallelised for distributed memory machines using two independent domain decompositions. For the exchange of fluid-membrane interface information a non-blocking point-to-point communication library is used. The developed simulation approach is applied to the study of the breathing oscillations of disk-gap-band parachutes and the progressive buckling of supersonic tension-cone decelerators. In both cases good qualitative and quantitative agreement between experiments and simulations have been achieved.

Jacek Gondzio – Preconditioners for higher order methods in big data optimization

We address efficient preconditioning techniques for the second-order methods applied to solve various sparse approximation problems arising in big data optimization. The preconditioners cleverly exploit special features of such problems and cluster the spectrum of eigenvalues around one. The inexact Newton Conjugate Gradient method excels in these conditions. Numerical results of solving L1-regularization problems of unprecedented sizes will be presented.

Akhtar A. Khan – Computational Methods for Elastography Inverse Problem

This talk will focus on various optimization formulations for solving the elastography inverse problem of identifying cancerous tumors. Computational efficiency of the various approaches for the identification of smooth as well as nonsmooth coefficient will be compared. Stability as well as the the effect of the noise in the measured data will be discussed. Detailed computational results will be presented.

Johannes Kraus – Stable discretization and robust iterative solution methods for the Brinkman problem

We consider discontinuous Galerkin $H(div,\Omega)$-conforming discretizations of elasticity type and Brinkman equations. We analyze their uniform stability and describe a simple Uzawa iteration for the solution of the Brinkman problem, which is based on solving a nearly incompressible linear elasticity type problem with mass term on every iteration. Based on a special subspace decomposition of $H(div,\Omega)$, we analyze variable V-cycle and W-cycle multigrid methods with nonnested bilinear forms. We prove that these algorithms are robust and their convergence rates are independent of the material parameters such as Lamé parameters and Poisson ratio and of the mesh size. Numerical results that confirm the theoretical analysis are presented.

This is a joint work with Qingguo Hong (RICAM, Austrian Academy of Sciences).

Jaroslav Kruis – High performance computing in multi-scale problems

Behaviour of structures made from highly heterogeneous materials, like masonry, is often described by models with two or more scales. As an example can serve coupled heat and moisture transfer in historical masonry bridges. Masonry composition is too complicated and therefore enormous number of finite elements is needed for description of the whole bridge. Apart from the very high number of degrees of freedom (DOF), randomness has very strong influence on the analysis. Therefore two levels are usually used. The macro-scale level describes the bridge geometry with homogenized material parameters while meso-scale level describes the masonry composition with the help of representative unit cell. The typical number of DOFs on the macro-scale level is hundred thousand but only hundreds or thousands DOFs are used in unit cells. The material parameters of the macro-scale level finite elements are determined from solution on micro-scale level. It means, a problem with thousands of DOFs is solved hundred thousand times. Moreover, a time integration method has to be used on the macro-scale level because the heat and moisture transfer is nonstationary. The methodology described is very demanding and it is suitable for parallel computing. In contrast to solution of systems of algebraic equations, the amount of data sent among processors is significantly smaller and therefore very good speedup can be achieved.

Julien Langou – Improving the communication lower bounds for matrix-matrix multiplication

We consider communication lower bounds for matrix-matrix multiplication in the sequential case (number of messages and total volume of messages). Our new proof technique improves the known lower bound for the number of reads and writes in the sequential memory case. This lower bound can be adapted to hierarchical memory and parallel distributed memory cases. We will show how extend our technique to other applications than the matrix-matrix multiplication kernel.

Jan Mandel, Adam Kochanski, Martin Vejmelka, and Sher Schranz – Data Assimilation of Satellite Active Fire Detection in Coupled Atmosphere-Fire Simulation

Active fire detection products from the VIIRS and MODIS instruments on polar-orbiting NASA and NOAA satellites provide planet-wide monitoring of fire activity several times daily as detection squares or polygons at a resolution of a fraction of a km up to 1 km. Thus, they present an ideal automated data source for fire behavior simulations. However, false negatives are common, there is no detection under cloud cover, geolocation errors can be significant, and active fire detection data are at a much lower resolution than the resolution of fire behavior models. Thus, fire detection seems to be better suited to improve fire behavior simulations in a statistical sense rather than as a direct input.

Data assimilation allows steering the simulation periodically and avoiding the accumulation of modeling errors. Data assimilation cycles at every satellite overpass also help to account for uncertainties caused by fuel information and fire fighting efforts. We propose a new data assimilation method for satellite fire detection, inspired by techniques for contour detection used in computer vision, particularly the Microsoft Kinect. The method takes advantage of encoding the state of the fire propagation in the fire arrival time, which can be modified by an additive correction.

We employ a Bayesian approach and obtain a Maximum-a-Posteriori (MAP) estimate as a solution of a least squares problem to maximize simultaneously a log likelihood function of fire detection, and minimize the change in the spatial gradient of the fire arrival time (which is proportional to the change in the reciprocal of the spread rate). The method is implemented effectively as preconditioned steepest descend using Fast Fourier Transform (FFT), and its computational cost is negligible compared to the coupled atmosphere-fire simulation itself.

Another challenge in data assimilation in coupled fire-atmosphere models is how to change in the state of the atmosphere when the state of the fire model changes in response to data. Heat flux from the fire changes the atmospheric circulation over time, and when the state of the fire model changes, the correspondence between the state of fire and the state of the atmosphere is lost. Encoding the fire model state as the fire arrival time allows to go back in time, blend the new fire arrival time with the original one, and spin up the atmospheric model with heat fluxes generated by replaying the modified fire arrival time instead of running the fire propagation model itself. Then, at the fire detection time (the satellite overpass time), the fire spread model takes over, and the simulation continues.

The method is illustrated on the case of the 2012 Barker fire simulated by WRF-SFIRE, with the assimilation of MODIS and VIIRS Active Fire detection data.

Supported by NSF grant DMS-1216481, NASA grant NNX13AH59G, and Czech Science Foundation grant 13-34856.

<u>Dan Negrut</u> – From Granular Dynamics to Fluid-Solid Interaction, and from Handling Large Optimization Problems to Solving Sparse Linear Systems

The talk outlines the key research thrusts in the Simulation-Based Engineering Lab at the University of Wisconsin-Madison, with the hope that some of the ideas presented will resonate with colleagues who might face similar challenges despite working in vastly different applications areas. The presentation will highlight how parallel computing is used to gauge the dynamics of granular material, to qualitatively characterize the dynamics of mechanical systems such as wheeled or tracked vehicles, and to describe the fluid-solid dynamic coupling in applications with engineering relevance. We touch on several orthogonal research directions, e.g., the solution of large optimization problems, sparse linear algebra, and parallel computing benchmarking, which turned out to play critical roles in our efforts to use modeling and simulation to solve Mechanical Engineering problems.

<u>Ulrich Rüde</u> – Coupled Models for Extreme-Scale Computing

Exploiting heterogeneous extreme scale computers to their full capability requires innovation on many levels. New algorithmic paradigms must address unprecedented levels of concurrency and a new performance oriented design methodology must be created. We will report on recent work on simulating complex flows based on the lattice Boltzmann method (LBM) and on using parallel multigrid methods. A special focus will be on coupled problems as they arise in the direct numerical simulation of suspensions, where the fluid-structure-interaction between the hydrodynamics and large ensembles of geometrically resolved particles must be modelled and simulated. Validation, scalability and performance results will be presented.

Valeria Simoncini – Multiterm linear matrix equations and the numerical solution of (S)PDEs

When finite differences or bilinear finite elements are employed, the numerical discretization of (deterministic) linear partial differential equations on rectangular or parallelepipedal domains leads to the solution of the following algebraic linear matrix equation

$$A_1 \boldsymbol{X} B_1 + A_2 \boldsymbol{X} B_2 + \cdots + A_m \boldsymbol{X} B_m = C$$

where $A_i \in \mathbb{R}^{n\times n}$, $B_i \in \mathbb{R}^{p\times p}$ and $C \in \mathbb{R}^{n\times p}$, to be solved for the unknown $n\times p$ matrix $\boldsymbol{X}$. When using stochastic Galerkin methods, the numerical solution of partial differential equations with random data also leads to linear matrix equations of the type above. Note that in both cases, it may hold that $n = N^d$, where d is related to the physical space dimension of the considered application. This linear algebra problem is typically solved by resorting to its Kronecker formulation,

$$\mathcal{A}\boldsymbol{x} = c, \quad \mathcal{A} = B_1 \otimes A_1 + \cdots + B_m \otimes A_m, \qquad \boldsymbol{x} = \mathrm{vec}(\boldsymbol{X}), \qquad c = \mathrm{vec}(C),$$

which involves vectors with np components. If both n and p are large, the system dimensions readily explode, making the numerical solution very demanding, both in terms of memory and computational resource allocations. Instead, we present a class of iterative methods for the low-rank approximation of $\boldsymbol{X}$ as $\boldsymbol{X} = \boldsymbol{X}_1 \boldsymbol{X}_2^T$ with $\boldsymbol{X}_1$, $\boldsymbol{X}_2$ tall, that takes full advantage of the possible low rank structure of the right-hand side matrix C. Numerical experiments with application problems will be included in the presentation to illustrate the performance of the new approach, compared to methods available in the literature. *This is joint work with Davide Palitta, Universita di Bologna (I), and with Catherine Powell and David Silvester, University of Manchester (UK).*

Wim Vanroose – Rounding errors in pipelined Krylov Methods

Pipelined Krylov solvers offer better strong scaling behavior because they combine various global reductions and hide it behind other useful computations. However, these methods use additional vectors and recurrence relations what results in a different accumulation of rounding errors. In this lecture we analyse these rounding errors and show how countermeasures can improve the accuracy of the final solution.

Barbara Wohlmuth – From optimal algorithms to fast solvers

In this talk we consider two main ingredients for large scale simulations in coupled flow problems. Firstly we discuss a local a posteriori flux correction technique which can be applied to all standard Stokes discretizations and guarantee strong mass conservation. The influence on the accuracy of the temperature is illustrated by some geophysically motivated benchmark examples. Secondly we focus on different multigrid schemes for the solution of the indefinit saddle point problem and provide scalability results on currect peta-scale systems. Here we also show test cases with large viscosity contrast and illustrate the influence of the PDE operator in a matrix free, stencil based formulation. This work is part of the DFG initiative SPPEXA and joint work with the group of U. Rüde.

Roman Wyrzykowski, Krzysztof Rojek, and Lukasz Szustak – Using GPU and Intel Xeon Phi coprocessors to accelerate stencil algorithms

EULAG (Eulerian/semi-Lagrangian fluid solver) is an established computational model for simulating thermo-fluid flows across a wide range of scales and physical scenarios. Among the most time-consuming components of EULAG, there is the multidimensional positive definite advection transport algorithm (MPDATA), which represents a sequence of stencil codes.

In the first part of this study, we focus on adapting the 3D MPDATA computations to clusters with graphics processors. Our approach is based on a hierarchical decomposition including the level of cluster, as well as an optimized distribution of computations between GPU resources within each node. We present performance results for the 3D MPDATA code running on the NVIDIA GeForce GTX TITAN graphics card, as well as on the Piz Daint cluster equipped with NVIDIA Tesla K20x GPUs.

In the second part, we discuss possible approaches to parallelizing the 3D MPDATA algorithm on platforms based on the Intel Xeon Phi coprocessors, and present some preliminary performance results.

Walter Zulehner – Hilbert space interpolation as a tool for analyzing and preconditioning mixed variational problems

We consider mixed variational problems, such as mixed methods for fourth-order elliptic boundary value problems and optimality systems of PDE-constrained optimization problems. Typical examples are plate bending problems and stationary distributed optimal control problems. It will be shown how to derive well-posed mixed variational formulations for such problems with the help of interpolation theory of Hilbert spaces. An analogous analysis for associated mixed finite element methods lead to efficient preconditioners for the linear systems of equations resulting from the discretization.

Abstracts – Other speakers

Simplice Donfack, Olaf Schenk, Bram Reps, Wim Vanroose, Matthias Christen – Improving the applicability of the highly efficient compilers to a wider class of problems

We propose an approach that allows the current solvers to adapt to future architectures and continue to scale at exascale, and this by removing the bottleneck introduced in the communication process. Our objective is to increase the arithmetic intensity, that is the number of floating-point operations performed per bytes fetched in the memory during the execution of the solvers, indeed reduce the number and the volume of the data exchanged among the processors and the memory. We use the help of efficient stencil compilers such as PLUTO and PATUS to increase the arithmetic intensity of these solvers. Our experiments on various stencil problems show that we can achieve a speedup up to 5X compared to the classic approach.

David Horák, Zdeněk Dostál, Václav Hapla, Lukáš Pospíšil, Alexandros Markopoulos, Martin Čermák, Alena Vašatová, Radim Sojka – PERMON toolbox

We shall present our new software called PERMON (Parallel, Efficient, Robust, Modular, Object-oriented, Numerical) toolbox [1]. It is based on PETSc and it combines domain decomposition methods especially of FETI type (Total-FETI) and quadratic programming algorithms (such as MPRGP - Modified Proportioning with Reduced Gradient Projections or SMALBE - SemiMonotonic Augmented Lagrangian algorithm for Bound and Equality constraints). This combination enjoys both numerical and parallel scalability for the solution of the contact problems (frictionless, Tresca friction, transient) of elasticity (frictionless, Tresca friction).

The core solver layer consists of several modules: PermonQP for unconstrained and equality-constrained QP, its PermonFLLOP extension for FETI, and PermonIneq extension providing algorithms for inequality-constrained QP. Other modules include the PermonCube benchmark generator include and application-specific modules like PermonPlasticity for plasticity, PermonImage for the image registration, PermonMultiBody for particle dynamics, and others.

REFERENCES

[1] PERMON web page: http://industry.it4i.cz/en/products/permon/

Jiří Jaroš, Bradley E. Treeby – Large scale ultrasound simulations

The simulation of ultrasound wave propagation through biological tissue has a wide range of practical applications including planning therapeutic ultrasound treatments of various brain disorders such as brain tumours, essential tremor, and Parkinson's disease. The major challenge is to ensure the ultrasound focus is accurately placed at the desired target within the brain because the skull can significantly distort it. Performing accurate ultrasound simulations, however, requires the simulation code to be able to exploit several thousands of processor cores and work with datasets on the order of tens of TB.

We have recently developed an efficient full-wave ultrasound model based on the pseudospectral method using pure-MPI with 1D slab domain decomposition that allows simulations to be performed using up to 1024 compute cores. However, the slab decomposition limits the number of compute cores to be less or equal to the size of the longest dimension, which is usually below 1024.

This talk will present an improved implementation that exploits 2D hybrid OpenMP/MPI decomposition. The 3D grid is first decomposed by MPI processes into slabs. The slabs are further partitioned into pencils assigned to threads on demand. This allows 8 to 16 times more compute cores to be employed compared to the pure-

MPI code, while also reducing the amount of communication among processes due to the efficient use of shared memory within compute nodes.

The hybrid code was tested on the Anselm with up to 2048 compute cores and the SuperMUC with up to 8192 compute cores. The simulation domain sizes ranged from 256^3 to 1024^3 grid points. The experimental results show that the hybrid decomposition can significantly outperform the pure-MPI one for large simulation domains and high core counts, where the efficiency remains slightly below 50%. For a domain size of 1024^3, the hybrid code using 8192 cores enables the simulations to be accelerated by a factor of 4 compared to the pure-MPI code. Deployment of the hybrid code has the potential to eventually bring the simulation times within clinically meaningful timespans, and allow detailed patient specific treatment plans to be created.

Lubomír Říha – ExaScale PaRallel FETI SOlver (ESPRESO)

ESPRESO is a massively parallel implementation of the Hybrid Total Finite Element Tearing and Interconnecting (FETI) method, which is designed to solve extremely large problems using multilevel decomposition. In Hybrid FETI a small number of neighboring subdomains is aggregated into clusters and each cluster is processed by single compute node. When compared to a two level FETI method, it significantly reduces its main bottleneck caused by solving the coarse problem. As of now the solver is able to solve problems larger than 6 billions of unknowns using 2000 compute nodes with 32 GB of RAM.

In order to test and fine tune the ESPRESO solver a massively parallel problem generator PermonCUBE, with nearly linear scaling, is in parallel developed by our team. Its main capability is generating the mesh with billion of unknowns in parallel over the cubical domain, and also preparing and assembling linear algebra objects, mostly for the problems related to mechanically deformable bodies.

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

We will show the scalability tests of the PermonCUBE and ESPRESO for both strong and weak scalability as well as on extremely large problems. In addition we will discuss the problems, that arise at this scale.

Patrick Sanan, Karl Rupp, Olaf Schenk, Dave May – Aggressive Accelerator-enabled Local Smoothing via Incomplete Factorization, with Applications

The development of hybrid supercomputers such as Piz Daint at the Swiss National Supercomputing Center (CSCS) has provided an opportunity to solve extremely large physics problems. However, there are limitations of current hardware, in particular the latency and bandwidth restrictions involved with launching accelerator kernels and moving data to and from global accelerator memory. As part of the GeoPC PASC project, we investigate the application of ``heavy smoothing'' within a multigrid preconditioner, using more aggressive coarsening paired with more aggressive smoothers which damp the wider range of frequencies required. This can be shown to maintain optimal scaling and allow for a tradeoff between local work and non-local communication, well-known to be the bottleneck on tightly-coupled physics problems solved on today's largest supercomputers. Smoothers are defined as overlapping local smoothers which can leverage accelerators. In particular, we examine the computation and application of the recently-developed Chow-Patel fine-grained ILU decomposition, which is amenable to computation on an accelerator and also allows for `warm-starting' between iterations in a nonlinear solve or time-stepped simulation. Implementation is done within the ViennaCL library, and its interface to PETSc, allowing its use with the pTatin3d library for lithospheric and mantle dynamics. We present software contributions to provide the required capabilities in a portable way, and results of applying the smoothers to precondition highly ill-conditioned linear systems arising from mantle convection simulation with highly heterogeneous viscosity structure.

Zdeněk Strakoš – On the Vorobyev method of moments in the context of Krylov subspace methods

In 1958 Yu. V. Vorobyev published (in Russian) a book called (in English translation that appeared in 1965) *Method of Moments in Applied Mathematics.* As mentioned in the annotation, ``*This book presents the theory behind the moment method for finding the eigenvalues of a linear operator approximately and for solving linear problems."* The book remained within the mathematical community almost unnoticed. Its importance has been pointed out by Claude Brezinski in 1996 in relation to the Lanczos method. Influenced by Gene Golub (through his interest in moments), by Volker Mehrmann (through the discussions on model reduction in control and in PDEs) and others, the author of this contribution would like to recall some ideas of Vorobyev in relation to the context of Krylov subspace methods.

Stanislav Sysala – How to Simplify Return-Mapping Algorithms in Computational Plasticity

The presentation is devoted to numerical realization of elastoplastic problems. The main goal is to improve implementation of the related constitutive problems. This can be done if plastic flow rules are defined by subdifferentials of plastic potentials. Then just one plastic multiplier is used even if the plastic potentials are nondifferentiable for unknown stress tensors. Further, the implicit Euler time discretization scheme is considered and the standard elastic predictor - plastic corrector method is used to find the discretized constitutive solution. Due to the presence of the one multiplier, it is possible to construct a unique system of nonlinear equations within the plastic correction regardless the unknown stress tensor lies on the smooth portion of the yield surface or not.

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

Jakub Šístek and Fehmi Cirak – Parallel performance of iterative solvers for pressure-correction methods for incompressible flows

We deal with a pressure-correction method for solving unsteady incompressible flows. In this approach, five subsequent equations are solved within each time step. These correspond to three scalar convection-diffusion problems, one for each component of velocity, a pure Neumann problem for the correction of pressure, and a problem of the L2 projection for pressure update. We present a comparative study of several parallel preconditioners and Krylov subspace methods from the PETSc library and investigate their suitability for solving the arising linear systems after discretizing by the finite element method. The target application are large-scale simulations of flows around wings of insects.

Dalibor Štys, Anna Zhyrova, Tomáš Náhlík and Renata Rychtáriková – Reaction-diffusion processes, cellular automata, interpretation of chemical self-organisation

The Belousov-Zhabotinsky reaction is a well known process in which from a simple mixture of chemicals spontaneously arise a system of nice dynamic structures. The identity of chemical reactions found in the system is far from understood. Rather good simulations of the process may be achieved by cyclic multilevel cellular automata. We shall show a modification of Dewndey´s hodgepodge machine model including mainly high level of noise and demonstrate its remarkable stability. We shall discuss following aspects: (a) Are simulations of reaction - diffusion processes in silico true simulations of differential equations or more multilevel cellular automata? (b) How far the simulations should be allowed to determine our interpretation of experimental processes? (c) How does the exact implementation of the process in silico affect the simulation of process dynamics and of its final ergodic state?

Jennifer Scott, Miroslav Tůma – Preconditioning of symmetric indefinite systems

Sparse symmetric indefinite linear systems of equations arise in numerous practical applications. In many situations, an iterative method is the method of choice but a preconditioner is normally required for this to be effective. Here the focus is on the development of incomplete factorization algorithms that can be used to compute high quality preconditioners for general indefinite systems, including saddle-point problems. A limited memory approach is used that generalizes recent work on incomplete Cholesky factorization preconditioners. A number of new ideas are discussed with the goal of improving the stability, robustness and efficiency of the resulting preconditioner. These include the monitoring of stability as the factorization proceeds and the use of pivot modifications when potential instability is observed. Numerical experiments involving test problems arising from a range of real-world applications are used to demonstrate the effectiveness of our approach and comparisons are made with a state-of-the-art sparse direct solver.

Vít Vondrák – IT4Innovations National Supercomputing Center

IT4Innovations National Supercomputing Center at VSB-Technical University of Ostrava (http://www.it4i.eu) is the strategic research infrastructure in the Czech Republic. In the framework of the national supercomputing center and its mission, the most powerful supercomputing facilities in the Czech Republic are being built and excellent research in HPC technologies and embedded systems is performed. The center is providing state-of-the-art technology and expertise in high performance computing and embedded systems and makes it available for Czech and international research teams from academia and industry. It aspires to improve the quality of life, to increase the competitiveness of industrial sector and to promote the cross-fertilization of high-performance computing, embedded systems and other scientific and technical disciplines. The center conducts research in the areas of earth science simulations for disaster management, traffic management, numerical modelling for engineering, physics and chemistry, development of libraries for parallel computing, modelling for nanotechnologies, soft-computing methods, recognition and presentation of multimedia data and safe and reliable architectures, networks and protocols. The center has also strong link to industrial sector. Currently, the largest Intel® Xeon Phi™ co-processor-based cluster in Europe with the peak performance 2 PFLOP/s is being installed at IT4Innovations National Supercomputing Center premises.

In the presentation, the IT4I supercomputing infrastructure will be presented as well as research activities related to this infrastructure.

Jan Zapletal, Michal Merta – BEM4I - A Parallel Library for Boundary Element Discretization of Engineering Problems

The BEM4I library under development at the IT4Innovations National Supercomputing Center implements boundary element solvers for the 3D Laplace, Helmholtz, Lamé, and time-dependent wave equations. The library aims at HPC architectures and utilizes modern parallelization strategies including OpenMP and MPI parallelization in shared and distributed memory, respectively. Special attention is paid to in-core vectorization for the evaluation of the singular surface integrals. To overcome the quadratic complexity of the classic BEM, fast boundary element methods are supported.

In the talk we present an overview of the library, scalability tests performed on the Anselm cluster in shared and distributed memory as well as preliminary results on the Intel Xeon Phi coprocessors.

One of the recent applications of BEM4I is the computation of homogenized coefficients for periodic structures with a highly oscillating material function. The usual volume approach is reformulated as a boundary integral equation involving the Steklov-Poincaré operators. The resulting boundary element method only discretizes the boundary of the periodic cell and the surface of the inclusion. The numerical experiments performed with BEM4I on the Anselm cluster validate the approach and agree with the theoretical order of convergence for the homogenized coefficients.

List of posters

M. Hanek, J. Šístek, P. Burda	Application of the BDDC method
R.Blaheta, J. Haslinger, A.A. Khan, R.Hrtus	Identification of material parameters
D. Horák, Z. Dostál, V. Hapla, L. Pospíšil, A. Markopoulos, M. Čermák, A. Vašatová, R. Sojka	PERMON Toolbox
K. Georgiev, I. Georgiev	AComIn: a Project for Advanced Computing in Innovations in IICT-BAS, Bulgaria
I. Georgiev	Advanced methods for structure and properties characterisation of fiber-reinforced silicate composite
M. Golasowski	Floreon+ - Disaster management system
M. Hanek, J. Šístek, P. Burda	Application of the BDDC method for incompressible Navier-Stokes equations
J. Hron, P. Minakowski	Plastic Deformation Treated as Material Flow Through Adjustable Crystal Lattice
P. Kůs, S. Badia, J. Principe	Domain Decomposition for Adaptively Refined Meshes
D. Lukáš, P. Kovář, T. Kovářová, M. Merta, M. Kravčenko	Parallel Fast Boundary Element Methods
C. Matonoha, J. Jablonský, Š. Papáček	Parameter Identification Problem Based on FRAP Images: From Data Processing to Design of Photobleach
J. Martinovič	viaRODOS: Monitoring and Analysis of Current Traffic Situation
D. Štys, A. Zhyrova, R. Rychtáriková, Š. Papáček	Simulation of the state space trajectory of the BZ reaction: Lower and upper flat noise limits

P. Rozehnalová, A. Kučerová, P. Štěpánek	Processes in concrete during fire
Jan Stebel	Flow123d - simulator of transport processes in porous media flow
P. Strakoš, M. Jaroš, T. Karásek, M. Jarošová, A. Vašatová, L. Říha, T. Kozubek	Advanced image processing methods for automatic liver segmentation
Václav Štumbauer	A step towards cad of photobioreactors for microalgae production
D. Štys, A. Zhyrová, T. Náhlík, R. Rychtáriková	Reaction-diffusion processes, cellular automata, interpretation of chemical self-organization

european social fund in the czech republic
EUROPEAN UNION
MINISTRY OF EDUCATION, YOUTH AND SPORTS
OP Education for Competitiveness
INVESTMENTS IN EDUCATION DEVELOPMENT

List of participants

Owe Axelsson	owe.axelsson@it.uu.se
Cevdet Aykanat	aykanat@cs.bilkent.edu.tr
Marc Baboulin	baboulin@lri.fr
Santiago Badia	sbadia@cimne.upc.edu
Petr Beremlijski	petr.beremlijski@vsb.cz
Radim Blaheta	blaheta@ugn.cas.cz
Jan Březina	jan.brezina@tul.cz
Jed Brown	jed-it4i@jedbrown.org
Tomáš Brzobohatý	tomas.brzobohaty@vsb.cz
Robert Cimrman	cimrman3@ntc.zcu.cz
Fehmi Cirak	fc286@eng.cam.ac.uk
Martin Čermák	martin.cermak@vsb.cz
Simplice Donfack	simplice.donfack@usi.ch
Jan Franců	francu@fme.vutbr.cz
Krassimir Georgiev	krtgeorgiev@gmail.com
Ivan Georgiev	ivan.georgiev@parallel.bas.bg
Tomáš Gergelits	gergelits@karlin.mff.cuni.cz
Martin Golasowski	martin.golasowski@vsb.cz
Jacek Gondzio	j.gondzio@ed.ac.uk
Martin Hanek	martin.hanek@fs.cvut.cz
Václav Hapla	vaclav.hapla@vsb.cz
Martin Hasal	martin.hasal@vsb.cz
Jaroslav Haslinger	hasling@karlin.mff.cuni.cz
David Horák	david.horak@vsb.cz
Radek Hrbáček	ihrbacek@fit.vutbr.cz
Jaroslav Hron	hron@karlin.mff.cuni.cz
Rostislav Hrtus	hrtus@ugn.cas.cz
Ondřej Jakl	ondrej.jakl@ugn.cas.cz
Adam Janečka	janecka@karlin.mff.cuni.cz
Branislav Jansík	branislav.jansik@vsb.cz
Jiří Jaroš	jarosjir@fit.vutbr.cz
Pavla Jirůtková	pavla.jirutkova@vsb.cz
Akhtar A. Khan	aaksma@rit.edu
Roman Kohut	kohut@ugn.cas.cz
Tomáš Kozubek	tomas.kozubek@vsb.cz

INVESTMENTS IN EDUCATION DEVELOPMENT

Johannes Kraus	johannes.kraus@uni-due.de
Michal Kravčenko	michal.kravcenko@vsb.cz
Jaroslav Kruis	jk@cml.fsv.cvut.cz
Jiří Kružík	ext.jiri.kruzik@skoda-auto.cz
Radek Kučera	radek.kucera@vsb.cz
Pavel Kůs	kus@math.cas.cz
Julien Langou	julien.langou@ucdenver.edu
Tomáš Luber	tomas.luber@ugn.cas.cz
Dalibor Lukáš	dalibor.lukas@vsb.cz
Lukáš Malý	lukas.maly@vsb.cz
Jan Mandel	jan.mandel@gmail.com
Alexandros Markopoulos	alexandros.markopoulos@vsb.cz
Pavel Maršálek	pavel.marsalek@vsb.cz
Jan Martinovič	jan.martinovic@vsb.cz
Ctirad Matonoha	matonoha@cs.cas.cz
Dan Negrut	negrut@wisc.edu
Horymír Netuka	horymir.netuka@upol.cz
Jana Nowaková	jana.nowakova@vsb.cz
Martin Palkovič	martin.palkovic@vsb.cz
Štěpán Papáček	spapacek@frov.jcu.cz
Jan Papež	jan@papez.org
Petr Pařík	parik@it.cas.cz
Marek Pecha	marek.pecha@vsb.cz
Lukáš Pospíšil	lukas.pospisil@vsb.cz
Lubomír Říha	lubomir.riha@vsb.cz
Filip Roskovec	roskovec@gmail.com
Petra Rozehnalová	rozehnalova.p@fce.vutbr.cz
Miroslav Rozložník	miro@cs.cas.cz
Ulrich Rüde	ulrich.ruede@fau.de
Martin Řehoř	rehor@karlin.mff.cuni.cz
Patrick Sanan	sananp@usi.ch
Václav Šátek	vaclav.satek@vsb.cz
Valeria Simoncini	valeria.simoncini@unibo.it
Kateřina Slaninová	katerina.slaninova@vsb.cz
Radim Sojka	soj0018@vsb.cz
Martin Stachoň	martin.stachon@vsb.cz
Jiří Starý	jiri.stary@ugn.cas.cz
Jan Stebel	jan.stebel@tul.cz

Zdeněk Strakoš	strakos@karlin.mff.cuni.cz
Petr Strakoš	petr.strakos@vsb.cz
Stanislav Sysala	stanislav.sysala@ugn.cas.cz
Jakub Šístek	sistek@math.cas.cz
Josef Šlapal	slapal@fme.vutbr.cz
Václav Štumbauer	stumbav@gmail.com
Dalibor Štys	stys@jcu.cz
Helena Švihlová	svihlova@karlin.mff.cuni.cz
Miroslav Tůma	tuma@cs.cas.cz
Wim Vanroose	wim.vanroose@uantwerpen.be
Alena Vašatová	vas266@vsb.cz
Oldřich Vlach	vla04@vsb.cz
Vít Vondrák	vit.vondrak@vsb.cz
Barbara Wohlmuth	wohlmuth@ma.tum.de
Roman Wyrzykowski	roman@icis.pcz.pl
Jan Zapletal	jan.zapletal@vsb.cz
Walter Zulehner	zulehner@numa.uni-linz.ac.at

Yours sincerely
Scientific and Organizing Committees