PRACE

SPRING SCHOOL 2014

15 -17 April 2014
Castle of Hagenberg
Austria

www.prace-ri.eu

Hosted and organised jointly by the Research Institute for Symbolic Computation / Johannes Kepler University Linz (Austria), IT4Innovations / VSB-Technical University of Ostrava (Czech Republic) and PRACE.

SOFTWARE ENGINEERING FOR SUPERCOMPUTERS IN RESEARCH AND INDUSTRY

PRACE Spring School 2014

Software Engineering for Supercomputers in Research and Industry

15 - 17 April 2014

Castle of Hagenberg, Austria

JKU JOHANNES KEPLER UNIVERSITY LINZ

RISC RESEARCH INSTITUTE FOR SYMBOLIC COMPUTATION

RISC Software GmbH

V Š B 1849 OSTRAVA TECHNICKÁ UNIVERZITA

Welcome

Dear participant,

the Research Institute for Symbolic Computation (RISC) of the Johannes Kepler University Linz and the IT4Innovations National Supercomputing Center of the VSB-Technical University of Ostrava are happy to welcome you at the PRACE 2014 Spring School! PRACE ("Partnership for Advanced Computing in Europe", www.prace-project.eu) is a European non-profit association with 25 member countries; its mission is to enable high impact scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society. PRACE also seeks to strengthen the European users of HPC in academy and industry through various initiatives; one of them is the regular organization of "PRACE Seasonal Schools" all over Europe. We are proud that the 2014 Spring School takes place under the general theme "Software Engineering for Supercomputing in Research and Industry" in the Castle of Hagenberg, Austria.

This medieval castle is located in Hagenberg, a small village 20km outside of Linz, the capital of Upper Austria; it was renovated in 1989 and since then houses the RISC Institute (www.risc.jku.at); in 2013 an extension building was opened nearby the castle pond. On initiative of RISC and its industrial spinoff, the RISC Software GmbH, the Softwarepark Hagenberg (www.softwarepark.at) was established, an industrial park with more than 1000 employees working in software development and related subjects. Furthermore, the Upper Austrian University of Applied Sciences (FH Oberösterreich) has established its School of Informatics, Communications and Media (www-en.fh-ooe.at) in Hagenberg. By the FH and various departments of the Johannes Kepler University located in Hagenberg, more than 1500 students are now educated in this village on subjects related to information technology and media. Furthermore, the Software Competence Center Hagenberg (www.scch.at) has been established as one of the largest independent Austrian research centers in the area of software. Despite of all these exciting developments, Hagenberg still has preserved its charme as a small rural village which provides a high quality for living and working. We hope that you will experience some of this flavor while you enjoy the scientific program of the spring school.

On the first day of the program, you will hear a keynote presentation, several reports, and a discussion on the experience with the application of High-Performance Computing in industry. On the other two days, you will take part in one of the tracks on "Software Engineering for Parallel & Emerging Computing Architectures" or "Elmer - Finite Element Software for Multiphysical Problems" with tutorial presentations and hands-on programming sessions. Further keynote talks will complement your experience. On the evening of the second day, we will visit the "Ars Electronica Center" in Linz, a "museum of the future" whose goal is to let people of every age and background get into contact with next-generation technologies; we will then enjoy dinner in the associated Cubus restaurant with a nice view over the Danube and the center of Linz.

We hope that you find this program attractive and interesting and that your participation in the PRACE 2014 Spring School will have a lasting effect on your future carreer.

Wolfgang Schreiner, Michael Krieger, Ondrej Jakl, David Horak, Karoly Bosa, Volker Strumpen
(local organization, programme, respectively admission committee)

April 2014

I. Accommodation and Transport

Accommodation

Hotel Sommerhaus
Julius-Raab-Straße 10
4040 Linz

Tel: 0043 (0)732 / 2457 – 376
Fax: 0043 (0) 732 / 2457 – 39

Mail: hotel@studentenwerk.at

Web: http://www.sommerhaus-hotel.at/en

During the spring school, the shuttle to the spring school location in Hagenberg leaves every day at 8:30 from the parking lot in front of the Hotel Sommerhaus Linz (you can use it, even if you are not accommodated in the Sommerhaus, just be there on time).

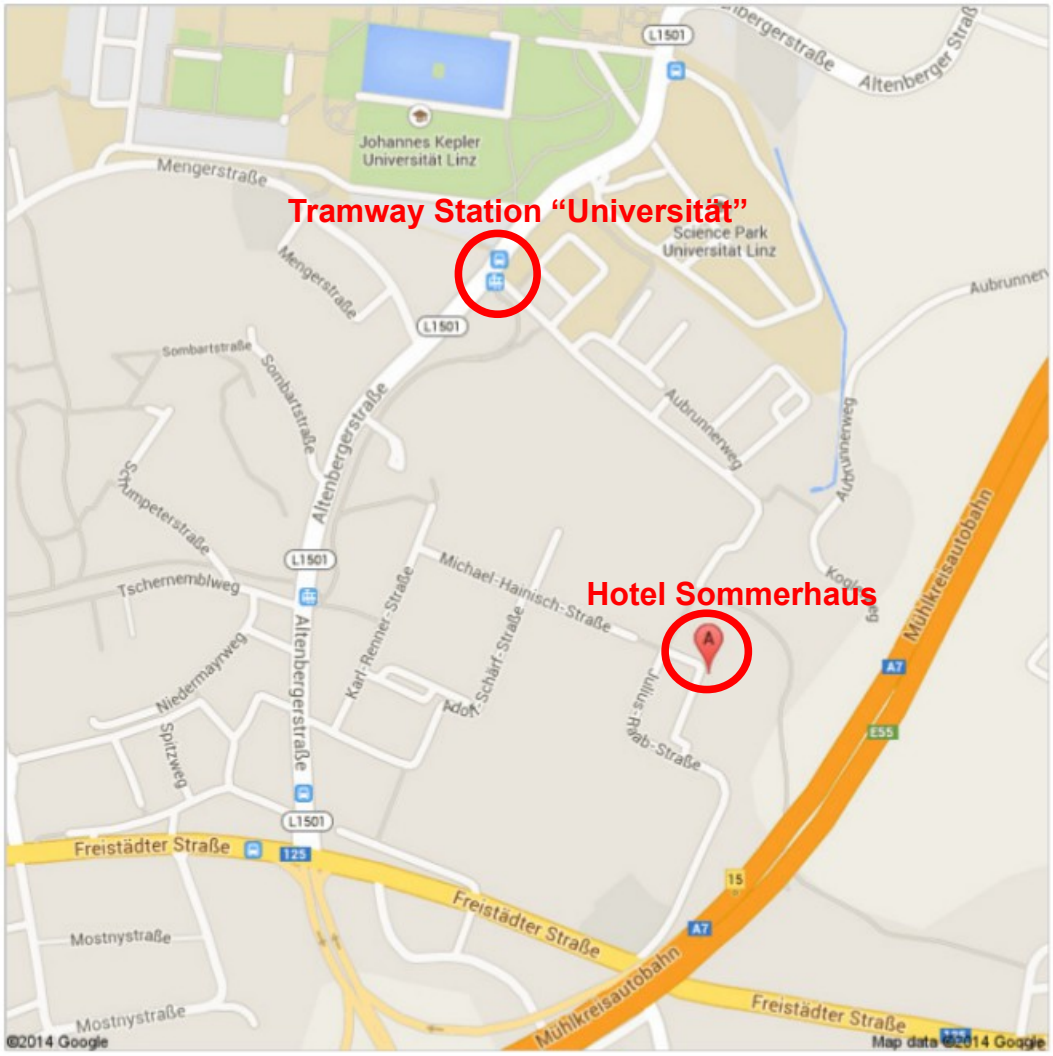

Public Transport in Linz

Take the tramway line number 1 or 2 to the station "Universität" to get to the Hotel Sommerhaus Linz. Along this line, there are also the main railroad station (station "Hauptbhf."), the Ars Electronica Center (station "Rudolfstr."), where the social programme takes place on Wednesday, and the "Hauptplatz", where you can find dinner.

Verkehrslinienplan

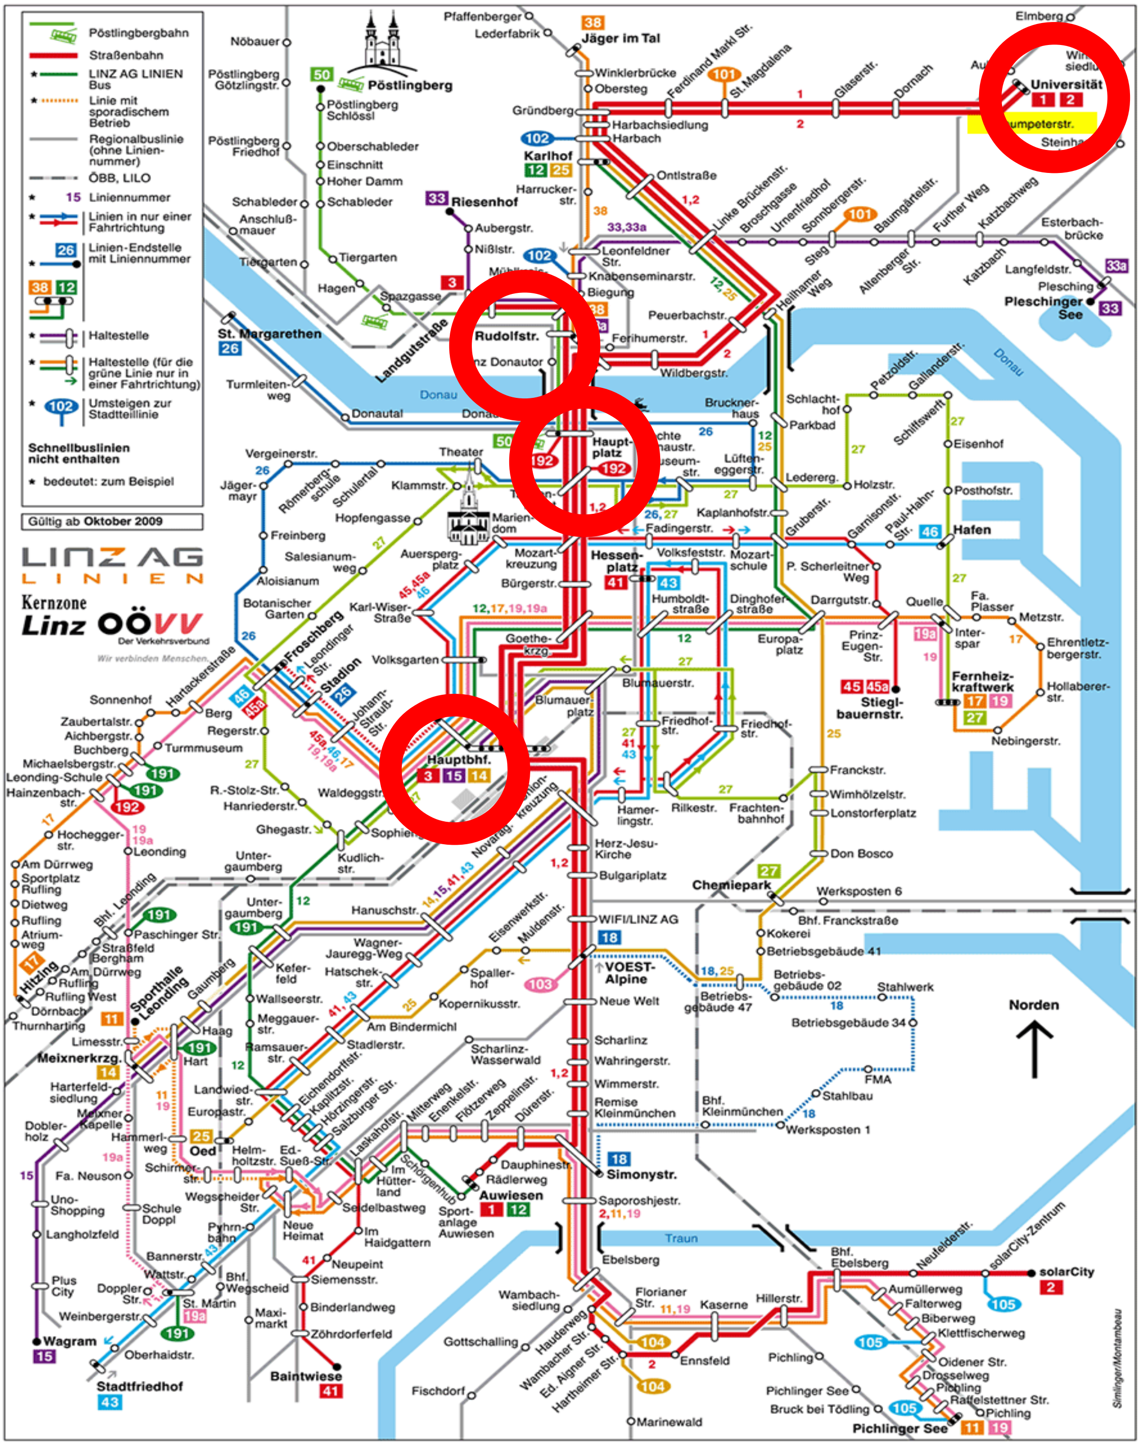

Dinner in Linz

If you want to enjoy dinner in Linz, go with the tramway to the central square ("Hauptplatz") south of the Danube (before you cross the Danube, you may notice to the east the "Ars Electronica Center", where we the social programme takes place on Wednesday). To the west of the Hauptplatz there is the old part of the city ("Altstadt"). Here and along to main road "Landstraße" to the south you will find many restaurants and bars.

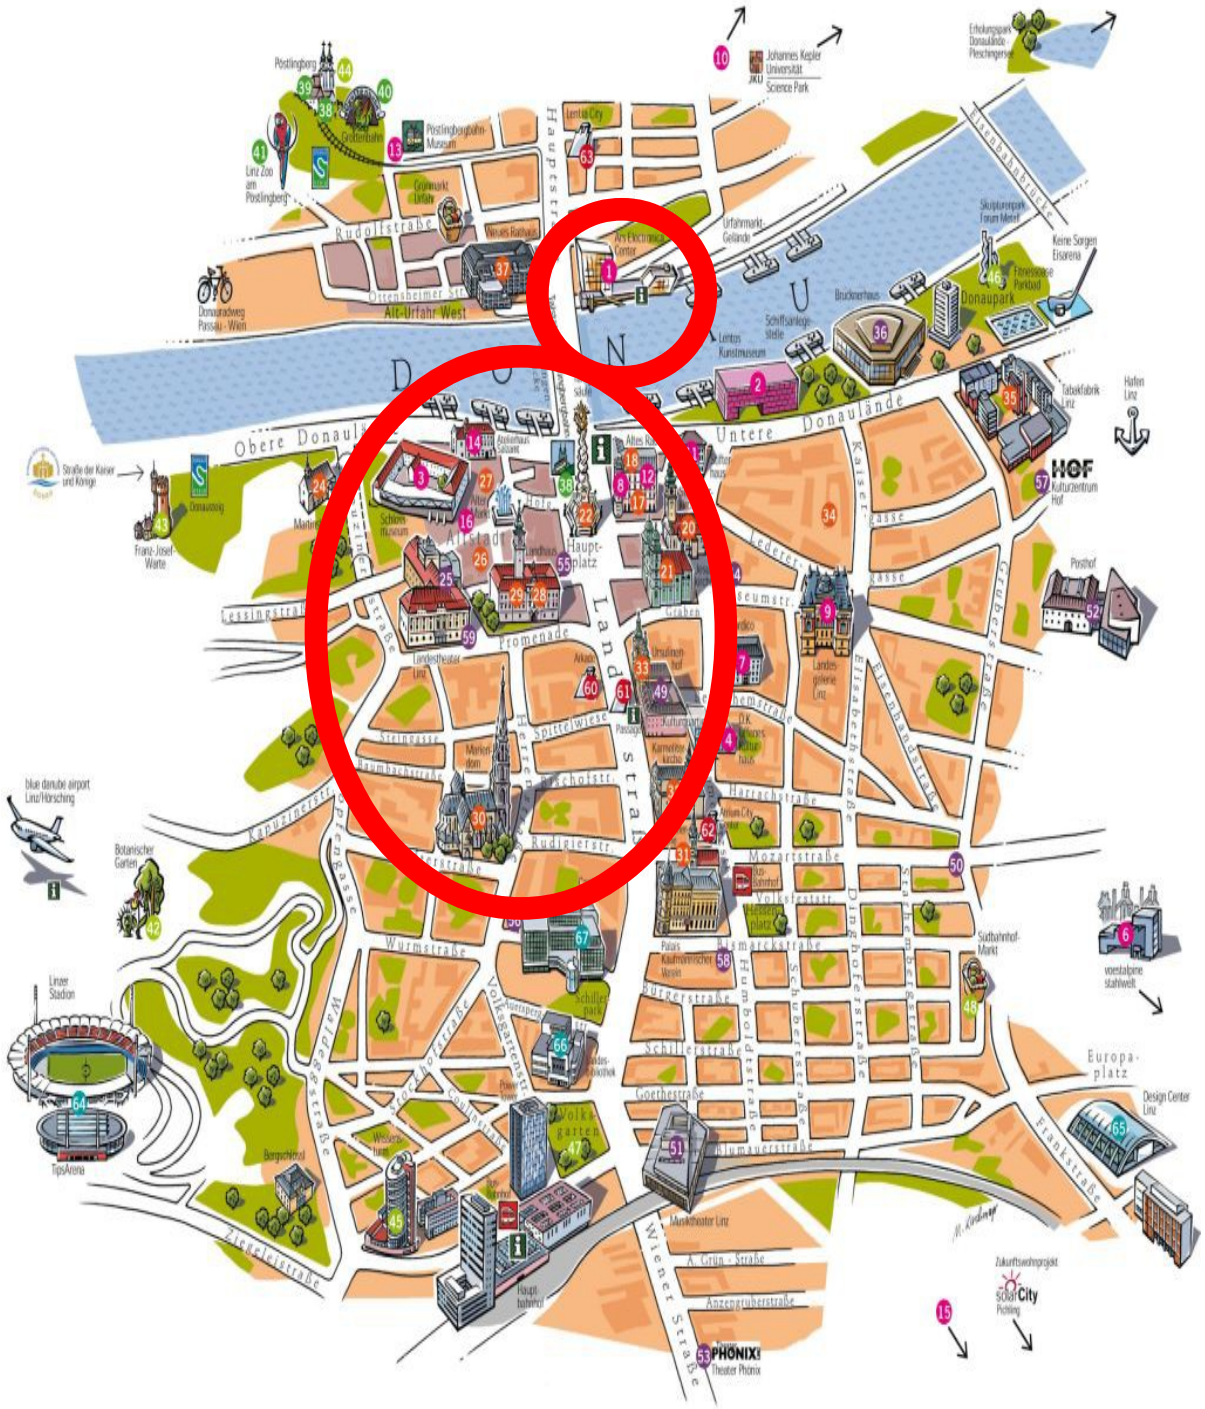

Spring School Location

Schloss Hagenberg, Kirchenplatz 5b, 4232 Hagenberg im Mühlkreis, Austria.

The shuttle bus will bring you at 8:30 from the Hotel Sommerhaus in Linz to the spring school location and back in the evening. See the directions on http://www.risc.jku.at/about/map/, if you want to get to the spring school location on your own.

To Hagenberg

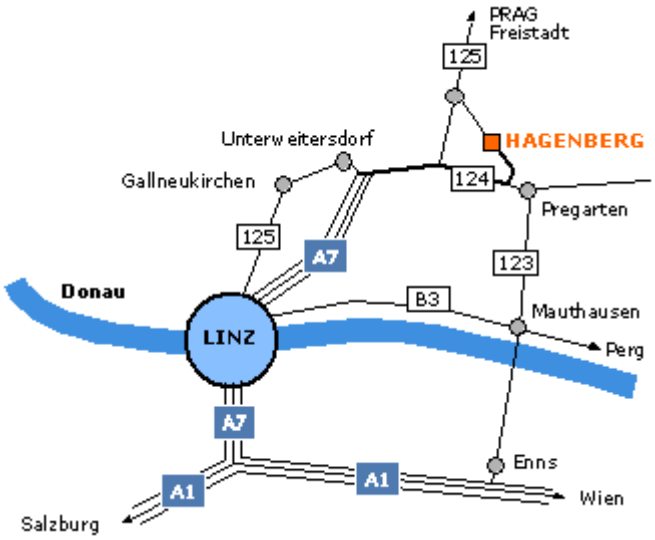

In Hagenberg

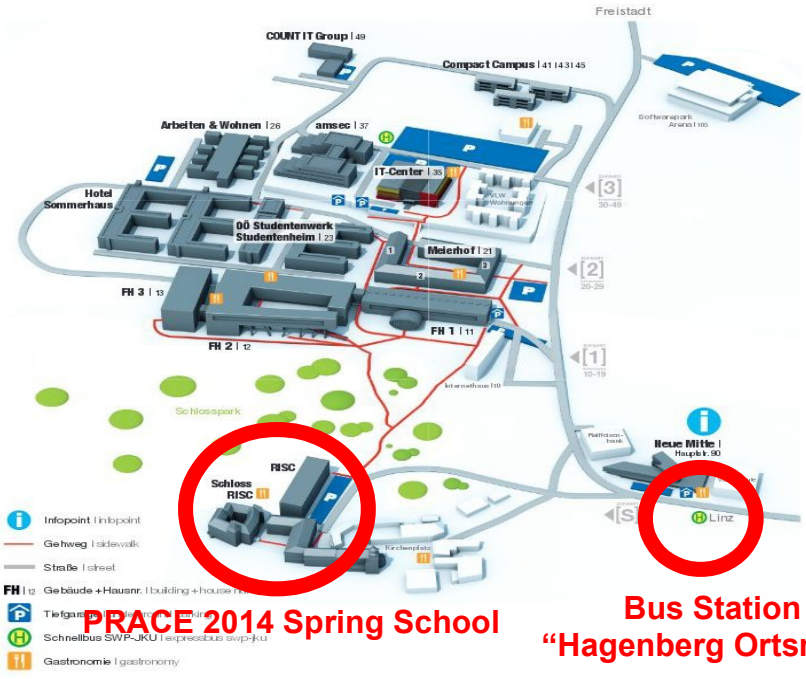

Castle and Pond Building Hagenberg

The spring school takes place in the seminar rooms in the castle and in the extension building at the nearby pond.


Schlossrestaurant Hagenberg

Lunches take place in the Schlossrestaurant next to the castle.


II. Programme

List of Presenters

(in alphabetical order)

Dr. Wolfgang FENZ
Research Department Medical Informatics of the RISC Software Company

Prof. Andrew GRIMSHAW
University of Virginia

Bernhard GRUBER, BSc.
RISC Software GmbH

Michael HAVA, MSc.
RISC Software GmbH

Prof. Dr. Dieter KRANZLMÜLLER
Leibniz Supercomputing Centre (LRZ)

Dipl.-Ing. (FH) Alexander LEUTGEB
RISC Software GmbH

Dr. Oskar MENCER
Maxeler Technologies

Wolfgang PLANER
Catalysts

Thomas PONWEISER
RISC Software GmbH

Dr. Peter RABACK
CSC - Scientific Computing Ltd.

Prof. Oswaldo TRELLES
University of Malaga

Dipl.Inf. Torsten WELSCH
RISC Software GmbH

Tuesday, April 15

8:30 Hotel Sommerhaus Linz

Bus Transfer to Hagenberg

9:15-9:30 Hagenberg Castle Seminar Room

Welcome

9:30-10:30 Hagenberg Castle Seminar Room

Keynote: Prof. Andrew GRIMSHAW, University of Virginia

The XSEDE Global Federated File System (GFFS/GFFS.eu) - Breaking Down Barriers to Secure Resource Sharing

10:30-10:45 Hagenberg Castle 1st Floor

Coffee Break

10:45-11:45 Hagenberg Castle Seminar Room

Dr. Wolfgang FENZ, Research Department Medical Informatics of the RISC Software Company

Simulating Brain Aneurysms with HPC

11:45-13:00 Schlossrestaurant Hagenberg

Lunch

13:00-14:00 Hagenberg Castle 1st Floor

Poster Presentations and Coffee, Spring School Photograph

14:00-15:30 Hagenberg Castle Seminar Room

Wolfgang PLANER, Catalysts

HPC for Processing Satellite Aerosol Observations

Prof. Dr. Dieter KRANZLMÜLLER, Leibniz Supercomputing Centre (LRZ)

Extreme Scale Computing at the Leibniz Supercomputing Centre (LRZ).

15:30-15:45 Hagenberg Castle 1st Floor

Coffee Break

15:45-16:45 Hagenberg Castle Seminar Room

Panel Discussion: *HPC for Industry: Status and Challenges*

17:00 Hagenberg Castle

Bus Transfer to Linz

Wednesday, April 16

8:30 Hotel Sommerhaus Linz

Bus Transfer to Hagenberg

9:30-10:30 Hagenberg Pond Building Seminar Room

Keynote: Dr. Oskar MENCER, Maxeler Technologies

OpenSPL - The Open Spatial Programming Language

10:30-10:45 Hagenberg Castle Ground Floor

Coffee Break

10:45-12:15 Castle Seminar Room (Session A) and Pond Building Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Dipl.-Ing. (FH) Alexander LEUTGEB: *Modern CPU Architectures*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK: *Introduction to Elmer Finite Element Software*

12:15-13:30 Schlossrestaurant Hagenberg

Lunch

13:30-15:00 Castle Seminar Room (Session A) and Poind Building Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Dipl.-Ing. (FH) Alexander LEUTGEB: *Introduction to Vectorization*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK: *Hands-on session using ElmerGUI*

15:00-15:15 Hagenberg Castle Ground Floor

Coffee Break

15:15-16:45 Castle Seminar Room (Session A) and Pond Building Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Michael HAVA, M.Sc: *Programming for the Intel Xeon Phi*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK.: *OpenLab*

17:00 Hagenberg Castle

Bus Transfer to Ars Electronica Center (AEC) in Linz

Wednesday, April 16

Social Programme

17:30-21:30 Ars Electronica Center (AEC) Linz

17:30 Reception

18:00 Guided Tour through the AEC

19:00 Dinner in the Cubus Restaurant

21:30 (or later) Tramway to Hotel Sommerhaus Linz

Thursday, April 17

8:30 Hotel Sommerhaus Linz

Bus Transfer to Hagenberg

9:30-10:30 Hagenberg Pond Building Seminar Room

Keynote: Prof. Oswaldo TRELLES, University of Malaga

About the Importance of HPC for Life Sciences

10:30-10:45 Hagenberg Castle Ground Floor

Coffee Break

10:45-12:15 Castle Seminar Room (Session A) and Pond Building Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Thomas PONWEISER: *Profiling Techniques for Parallel Applications*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK: *Advanced Use of Elmer*

12:15-13:30 Schlossrestaurant Hagenberg

Lunch

13:30-14:30 Castle Seminar Room (Session A) and Pond Building Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Bernhard GRUBER: *Introduction to OpenCL Programming in Heterogeneous Architectures*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK.: *Hands-on Session with Advanced Features and Parallel Computation*

14:30-14:45 Hagenberg Castle Ground Floor

Coffee Break

14:45-15:45 Pond Building Seminar Room (Session A) and Castle Seminar Room (Session B)

Session A: Software Engineering for Parallel & Emerging Computing Architectures

Dipl.Inf. Torsten WELSCH: *Introduction to GPGPU Programming with CUDA*

Session B: Elmer - Finite Element Software for Multiphysical Problems

Dr. Peter RABACK.: *OpenLab*

16:00-16:15 Hagenberg Castle Seminar Room

Closing

16:30 Hagenberg Castle

Bus Transfer to Hotel Sommerhaus Linz

Abstracts

The XSEDE Global Federated File System (GFFS/GFFS.eu) - Breaking Down Barriers to Secure Resource Sharing

Prof. Andrew GRIMSHAW, University of Virginia

The GFFS offers scientists a simplified means through which they can interact with and share resources. Currently, many scientists struggle to exploit distributed infrastructures because they are complex, unreliable,and require the use of unfamiliar tools. For many scientists, such obstacles interfere with their research; for others, these obstacles render their research impossible. It is therefore essential to lower the barriers to using distributed infrastructures. The first principle of the GFFS is simplicity. Every researcher is familiar with the directory-based paradigm of interaction; the GFFS exploits this familiarity by providing a global shared namespace. The namespace appears to the user as files and directories so that the scientist can easily organize and interact with a variety of resources types. Resources can include compute clusters, running jobs, directory trees in local file systems, groups, as well as storage resources at geographically dispersed locations. Once mapped into the shared namespace, resources can be accessed by existing applications in a location-transparent fashion, i.e., as if they were local.

In this talk I will present the GFFS, its functionality, its motivation, as well as typical use cases. I will demonstrate many of its capabilities, including: 1) how to securely share data, storage, and computational resources with collaborators; 2) how to access data at the centers from campus and vice versa; 3) how to create shared compute queues with collaborators; and 4) how to create jobs and interact with them once started.

See http://genesis2.virginia.edu/wiki/Main/GFFSeu for information on the European instance.

Simulating Brain Aneurysms with HPC

Dr. Wolfgang FENZ, Research Department Medical Informatics of the RISC Software Company

We present two projects dealing with the numerical simulation of brain aneurysms. Aneurysms are balloon-like dilations of blood vessels that can rupture and cause an internal bleeding (brain hemorrhage) or stroke. Treatment nowadays mostly involves a minimally invasive procedure (insertion of a stent and coil), although in some cases brain surgery (clamping with a metal clip) is necessary.

Recent studies have shown that patient-specific blood flow patterns and flow-related quantities such as wall shear stress and oscillatory shear index have a great influence on growth and rupture of aneurysms. In order to facilitate a more accurate rupture risk assessment for physicians, the project MEDVIS 3D, tries to introduce blood flow analysis via computational fluid dynamics (CFD) into clinical practice. Using modern numerical methods, parallelization and GPU computing, detailed simulations of blood flow coupled with wall elasticity over several cardiac cycles can be performed on a desktop PC within a reasonable time frame. Moreover, the system also provides three-dimensional reconstruction of medical image data, volume segmentation, visualization, and efficient mesh generation for the finite element calculations.

The second project, currently in development, is a training simulator allowing neurosurgeons to perform virtual clipping procedures using haptic input devices. For this, we have to accelerate the CSD (computational structural dynamics) part of the flow simulation for real-time application. In addition, collisions between instrument and vessel meshes have to be detected and transformed into corresponding boundary conditions for the simulation, while appropriate feedback forces have to be sent back to the user.

Catalysts - HPC for Processing Satellite Aerosol Observations

Wolfgang PLANER, Catalysts

Aerosols are small particles in the atmosphere. They have great influence to world climate due to their high absorption of sun light. Various satellites are continuously scanning the electromagnetic radiation. Using this data the distribution of aerosols in the atmospheric layers can be computing by various algorithms. Until now only simple algorithms could be used for performance reasons. Using accelerators with a high number of concurrent processors enables the scientists to use accurate physical models in their algorithms to produce detailed aerosol distribution models in nearby real time speed. The amount of date to be processed in one day is in the range of some terabytes, while the amount of data for reprocessing the accumulated data of some years is in the petabyte range.

The lecture gives an overview of the techniques used to accelerate these algorithms.

Extreme Scale Computing at the Leibniz Supercomputing Centre (LRZ)

Prof. Dr. Dieter KRANZLMÜLLER, Leibniz Supercomputing Centre (LRZ)

One of the most powerful computing systems of the world, SuperMUC, is hosted at the Leibniz Supercomputing Centre (LRZ) in Garching near Munich, Germany. Its performance of currently up to 3 Petaflop/s (soon doubled to >6 Petaflop/s) is provided by its extreme scale, in concrete by more than 155.000 cores. These cores are connected using a complex structure of thin and fat islands and the high speed Infiniband interconnect. While the x86-compatible cores allow portability with smaller systems, the full performance is only achievable with scalable software, where power consumption and energy restrictions apply. This talk presents the LRZ and its environment in Europe, especially its contribution to the Partnership for Advanced Computing in Europe (PRACE).

OpenSPL - The Open Spatial Programming Language

Dr. Oskar MENCER, Maxeler Technologies

The Open Spatial Programming Language (OpenSPL) consortium was founded by CME Group, Juniper, Chevron and Maxeler Technologies with initial academic members from Stanford University, Imperial College London, Tsinghua University and the University of Tokyo. OpenSPL is a model for describing computations in 2D space, suitable for a wide range of chip substrates. For example, Maxeler's Multiscale Dataflow Engines (DFEs) are a currently available instantiation. Most recently, the UK Science and Technology Facilities Council invested into a significant DFE based machine with MPC-X nodes capable of an equivalent 8.52TFLOPs per 1U and 8.97 GFLOPs/Watt.

Modern CPU Architectures

Dipl.-Ing. (FH) Alexander LEUTGEB, RISC Software GmbH

This session gives at first the motivation for parallelization in central processing units (CPUs). The different levels of parallelism implemented in hardware are presented in the case study of the Intel x86 Sandy Bridge architecture. These are namely task level, control level, data level, and instruction level parallelism. For each level the implementation in hardware is illustrated. We analyse the relevance of each level from a programmers point of view. Because parallel algorithms are not only bound by the computing power (peak floating point performance) of the CPU, but also limited by the memory bandwidth between CPU and the main memory, memory hierarchies are presented in more detail. We show the motivation for Caching in Hardware and what kind of problems arise from Caching in a parallel context. Finally the Roofline Model is presented. It allows for any parallel algorithm with a certain operational density to calculate the attainable floating point performance considering the peak floating point performance and the peak memory bandwidth of the target machine.

Elmer - Finite Element Software for Large Scale Multiphysical Problems

Dr. Peter RABACK, CSC - Scientific Computing Ltd

Elmer is an open source finite element software suite particularly well suited for the study of multiphysical problems. In Elmer each physical model is a dynamically linked solver and there is no upper limit to the number of different solvers. Elmer includes at least basic models for many branches of computational engineering, e.g. fluid mechanics, solid mechanics, heat transfer, electromagnetics, acoustics. Elmer has been parallized with MPI and has demonstrated good scalability up to thousands of cores.

After the tutorial the participants should be able to solve basic engineering problems using the graphical user interface of Elmer and also run these problems in parallel. The tutorial gives a good starting point for more complex simulations using Elmer software.

Peter Råback holds a D.Sc.(Tech.) from Helsinki University of Technology. He has been involved in the development of Elmer since 2001 contributing to many aspects of the code. His main specialities lie in the solution of coupled problems. Currently Råback is managing the development of Elmer at CSC - IT Center for Science, Finland.

Introduction to Vectorization

Dipl.-Ing. (FH) Alexander LEUTGEB, RISC Software GmbH

This session gives an introduction into vetorization on Intel x86 CPUs from a programmers point of view. We will present the potential of vector units. The different levels of the programmers control over the vectorization are shown. Special focus lies in the auto vectorization support of the Intel C/C++ compiler. The inhibitors of auto vectorization and techniques how to get rid of them will be presented. Finally we will talk about portability issues, if we use the latest vector unit features of Intel CPUs.

Programming for the Intel Xeon Phi

Michael HAVA, MSc., RISC Software GmbH

This session starts with an introduction to the Intel Xeon Phi co-processor, giving a short overview on the hardware architecture and unique design characteristics. A brief overview on the different development and execution models supported by the Xeon Phi will be given. The new OpenMP 4.0 standard (published in July 2013), especially its support for dedicated accelerator hardware, will be a central aspect of this session. In addition to the standard OpenMP functionality selected Intel-specific extensions - not necessarily limited to the Xeon Phi - will be presented.

About the Importance of HPC for Life Sciences

Prof. Oswaldo TRELLES, University of Malaga

Large scale genomics projects exploiting high throughput leading technology have produced and continue to produce massive data sets with exponential growing rates. So far, only a small part of this data can be abstracted, managed and processed, giving an incomplete understanding of the biological process being observed. The lack of processing power is a bottle neck in acquiring results.

A promising approach to address the processing of such massive data sets is the creation of new computer software that makes effective use of distributed, grid, parallel and cloud computing.

Comparative genomics is a good example since it includes all the ingredients: huge and ever growing datasets, complex applications that demands large computational resources and new mathematical and statistical models for analysing and synthetizing genomic information.

This talk will provide an overview of the main bioinformatics demands for data processing and the ways to provide solutions in the framework of high performance computing.

Profiling Techniques for Parallel Applications

Thomas PONWEISER, RISC Software GmbH

This session gives an introduction to effective strategies for analyzing program performance. The focus will lie on profiling tools utilizing event-based sampling and hardware counters (such as HPCToolkit, Amplifier, etc.). After a short overview over useful metrics for different problems, practical examples will provide attendants with knowledge for improving their applications.

Introduction to OpenCL Programming in Heterogeneous Architectures

Bernhard GRUBER, BSc., RISC Software GmbH

In this session, I will give an introduction into key concepts that influenced OpenCL's design, followed by a detailed example of how to implement a simple algorithm (e.g. matrix matrix multiplication) for the GPU.

Introduction to GPGPU Programming with CUDA

Dipl.Inf. Torsten WELSCH, RISC Software GmbH

In this session, we give an introduction to the basic concepts behind general-purpose computing on graphics processing units (GPGPU). Our one-hour presentation focuses on the following items: A brief explanation of the general differences between classical central processing units (CPUs) and GPUs, an overview of today’s most important vendors and their current products, a detailed description of NVIDIA’s Kepler architecture, an introduction to NVIDIA’s CUDA programming language, including the programming model, the memory hierarchy and performance guidelines, and finally, the analysis of a real world application – see the RISC enlight project – containing the description of profiling results generated by the NVIDIA Nsight profiler.

III. Further Information

Wireless LAN

- **Castle**: you may connect to the WLAN hotspot

  `risc_public_hotspot`

  without authentication.

- **Pond Building:** you may connect to the WLAN hotspot

  `eduroam`

  with the eduroam account from your home academic institution. In case, you do not have such an account, please use the temporary account that you find on a separate sheet in your conference bag.

Addresses and Contacts

Spring School Location:

Research Institute for Symbolic Computation (RISC)
Schloss Hagenberg (Castle of Hagenberg)
Kirchenplatz 5b
A-4232 Hagenberg im Mühlkreis, Austria

Phone: +43 732 2468 9921
Fax: +43 732 2468 9930
Email: secretary@risc.jku.at
Web: http://www.risc.jku.at

Emergency Contacts:

Secretary:	+43 732 2468 9921	(during office hours)
Cornelia Staub:	+43 699 1300 9022	(outside office hours)
Wolfgang Schreiner:	+43 699 11441228	(outside office hours)
Michael Krieger:	+43 699 17331744	(outside office hours)