UNIVERSITÄT SALZBURG

Second Austrian High Performance Computing Workshop

30. – 31. 5. 2011

hosted by

Paris Lodron Universität Salzburg

Fachbereich Computerwissenschaften

Program – Monday 30.05.

12:00 **Arrival of participants**

13:00 **Opening**
Welcome: Peter Zinterhof (Uni Salzburg)
Statement: Sabine Schindler (Uni Innsbruck)
Statement: Siegfried Reich (Salzburg Research)
Organizatorial: Marian Vajteršic (Uni Salzburg)

13:30 **Invited talk** (chair: Marian Vajteršic)
Esmond G. Ng (Lawrence Berkeley National Laboratory)
Large-scale Matrix Computation in Scientific Applications

14:30 **Session** (chair: Marian Vajteršic)
Dietrich Liko (Öst. Akad. d. Wissenschaft)
Data Intensive Computing at the Federated Tier-2 Center for LHC

14:55 Gregor Mair (Uni Innsbruck)
Grid Computing in Austria for the ATLAS Experiment

15:20 Christian Meisenbichler (Montanuni Leoben)
exciting@web, HPC as a Web Service

15:45 **Coffee break**

16:15 **Session** (chair: Peter Zinterhof)
Rade Kutil (Uni Salzburg)
Minimizing Communication Cost in Succession of Parallel Matrix Operations

16:40 Stefan Wagner (FH Oberösterreich)
HeuristicLab Hive — An Open Source Environment for Parallel and Distributed Execution of Heuristic Optimization Algorithms

17:05 Peter Zinterhof (Uni Salzburg)
Distributed Computation of Feature-Detectors for Medical Image Processing on GPGPU and Cell Processors

17:30 Roland Kwitt (Uni Salzburg)
Estimation of Probabilistic Models in (Semantic) Image Retrieval

18:45 **Dinner**

Program – Tuesday 31.05.

08:30 **Session** (chair: Sabine Schindler)
Manfred Liebmann (KF-Uni Graz)
Large Scale Simulations of the Euler Equations on GPU Clusters

08:55 Aurel Neic (KF-Uni Graz)
Algebraic Multigrid Solver on Clusters of CPUs and GPUs

09:20 Markus Haider (Uni Innsbruck)
Computer Simulations of Galaxy Clusters

09:45 Dominik Steinhauser (Uni Innsbruck)
N-body / Hydrodynamic Simulations of Galaxies in a Cluster Environment

10:10 Andreas Läuchli (Uni Innsbruck)
HPC Challenges in Simulations of Interacting Quantum Many Body Systems

10:35 **Coffee break**

11:05 **Session** (chair: Jens Volkert)
Karl Fürlinger (Uni München)
Consumer Electronic Hardware as a Platform for Parallel Computing

11:30 John Thomson (Uni Innsbruck)
Insieme: A Multi-Objective Optimising Compiler for HPC

11:55 Dirk Draheim (Uni Innsbruck)
The ACSC Supercomputing Infrastructure: News and Trends

12:20 Herbert Störi (TU Wien)
Vienna Scientific Cluster — Status and Outlook

12:45 **Lunch**

14:00 **Session** (chair: Thomas Antretter)
Wolfgang Fenz (RISC Software GmbH)
Simulation of Blood Flow in Intracranical Aneurysms

14:25 Susanne von Grafenstein (Uni Innsbruck)
Dynamics of the Antiviral Target Neuraminidase

14:50 Hannes G. Wallnöfer (Uni Innsbruck)
Molecular Dynamics Simulations of Biological Systems — Specificity in Snake Venom Metalloproteases

15:15 **Coffee break**

15:30 Helmut A. Mayer (Uni Salzburg)
Evolution of Robotic Neurocontrollers with Intrinsic Noise

15:55 Sebastian Hegenbart (Uni Salzburg)
On the Needs of High Performance Computing in Computer Aided Decision Support Systems

16:20 **Conclusion**

Large-scale Matrix Computation in Scientific Applications

Esmond G. Ng
Lawrence Berkeley National Laboratory

The U.S. Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program is a multidisciplinary and diverse R&D program that brings together application scientists, applied mathematicians, and computer scientists to tackle challenging problems in several scientific domains, such as accelerator science, astrophysics, climate, fusion energy, materials science, nuclear physics, and subsurface flow. Most, if not all, of these problems are modeled using differential equations. After discretization, the innermost kernels are often the solution of algebraic systems, including systems of linear equations and eigenvalue problems. In many cases, the level of realism required in the modeling results in simulations that have to be performed at high resolutions. Consequently, the solution of these algebraic systems can be extremely challenging, in terms of the size and/or the conditioning of the systems. In this talk, we will describe some of our collaborations with application scientists. We will discuss some of the large-scale matrix problems that we encounter and our approaches in solving them.

The ACSC Supercomputing Infrastructure: News and Trends

Dirk Draheim
Universität Innsbruck, ZID

In this talk we approach these and similar questions: Which concrete supercomputing infrastructure initiatives are currently planned by ACSC partners? Which initiatives are possible in the future? Which cooperation models will be possible from an infrastructure viewpoint? Which infrastructure services and features could we offer to attract additional partners?

Simulation of Blood Flow in Intracranical Aneurysms

Wolfgang Fenz
RISC Software GmbH, Research Unit Medical Informatics

Cerebral hemorrhages (ruptured brain aneurysms) are among the most frequent causes of death in western countries. Since every patient's physiology forms an individual flow pattern with potential locations of higher risk for ruptures of the blood vessel wall, it is important to find these high risk areas in order to decide on an appropriate treatment. Therefore, we are providing a software tool for the physically correct calculation of hemodynamics in a 3D geometry based on medical image data that is easy to use for physicians and avoids the high costs associated with commercial simulation software. Due to the time-critical nature of the application, we exploit efficient state-of-the-art numerical methods and technologies together with high performance computing (Austrian Grid, CUDA). Our system first builds a three-dimensional virtual replica of the patient's cerebrovascular system from medical image data. The physician can then select a region of interest which is automatically transformed into a tetrahedral mesh. The blood is modeled as an incompressible Newtonian fluid, the surrounding vessel wall as an isotropic linear elastic material. Both the Navier-Stokes equations for the fluid domain and the Navier-Lamé equations for the solid domain are handled with the finite element method (FEM), and the resulting linear equation systems are solved via conjugate gradient (CG) and algebraic multigrid (AMG) methods. At regular intervals during a cardiac pulse cycle, mechanical equilibrium between blood flow and wall elasticity is achieved using an iterative fluid-structure interaction (FSI) algorithm, deforming the fluid mesh according to the wall displacement. We present simulation results calculated on patient-specific geometries, including velocity and pressure fields, oscillatory shear index, as well as vessel deformation and stress distribution.

Consumer Electronic Hardware as a Platform for Parallel Computing

Karl Fürlinger
Ludwig-Maximilians-Universität München, Informatik

Total power consumption and energy efficiency have become important considerations for the design of data centers and high performance computing facilities. At the same time, consumer electronic and handheld devices have always been designed with energy efficiency in mind, with ARM-based CPUs powering the majority of devices in the recent history.

Driven by use cases such as HD video streaming, rich Web 2.0 browser applications, virtual reality, and 3D gaming, there are strong market forces driving the future development of theses mobile devices towards more powerful compute capabilities. For example, dual and even quad core mobile CPU designs with GPU integration have been announced for the next generation of tablet computers and smartphones.

To evaluate the current state of energy efficient parallel and distributed computing on consumer electronic devices we have built a small cluster of second generation Apple TV (ATV2) devices. Like the iPad, the ATV2 contains the Apple A4 processor which combines a ARM Cortex-A8 running at 1 GHz with a PowerVR SGX535 GPU and 256 MB RAM. The device is small, inexpensive, and consumes only about 2-3 Watts. In this talk we present early results of our evaluation of this AppleTV-Cluster.

Dynamics of the Antiviral Target Neuraminidase

Susanne von Grafenstein,
H. G. Wallnöfer, J. Kirchmair, R. G. Huber, J. E. Fuchs, K. R. Liedl
Universität Innsbruck, Institute of General, Inorganic and Theoretical Chemistry

Influenza neuraminidase (NA) is a tetrameric surface protein of the influenza virus and target for antiviral drugs, e.g., oseltamivir and zanamivir. The conformational diversity of this protein was revealed by protein X-ray analysis [1]. The underlying flexibility opens up new possibilities for drug design. There is an urgent need for new antiviral against influenza as drug resistance against oseltamivir is spreading globally. Dynamics of neuraminidase are investigated by molecular dynamics (MD) simulations [2, 3].

We present a systematic analysis of two neuraminidase subtypes (avian 2005, pandemic 2009) with all-atom, explicit solvent MD simulations applying the Amber forcefield ff99SB. Comparative simulations of monomeric and tetrameric systems reveal that the sampled conformational space for the tetramer is distinct from monomer simulations. These findings point to a stabilization of the active conformations by the protein-protein-interface.

Therefore, we conclude flexibility to be important for protein-protein interactions in tetrameric neuraminidase. In consequence, these interactions are crucial for drug development and drug resistance.

References

[1] R.J. Russell, L.F. Haire, D.J. Stevens, P.J. Collins, Y.P. Lin., G.M. Blackburn, A.J Hay, S.J. Gamblin, J.J. Skehel. Nature, 2006, 443, 45-49.

[2] R.E. Amaro, X.L. Cheng, I. Ivanov, D. Xu, J.A. McCammon, J Am Chem Soc, 2009, 131, 4702-4709.

[3] U. Grienke, M. Schmidtke, J. Kirchmair, K. Pfarr, P. Wutzler, R. Durrwald, G. Wolber, K.R. Liedl, H. Stuppner, J.M. Rollinger, J Med Chem, 2010, 53, 778-786.

Computer Simulations of Galaxy Clusters

Markus Haider
Universität Innsbruck, Astro- und Teilchenphysik

We present simulations of galaxy clusters aimed at studying the enrichment of the intra-cluster-medium with heavy elements. Galaxy clusters are the largest gravitationally bound structures in the universe. Besides galaxies, clusters consist of dark matter and a thin hot gas between the galaxies, the so called intra-cluster-medium. The evolution of the intra-cluster-medium is modeled via a grid-based hydrodynamics code which experiences the gravitational potential of dark matter which is simulated using an N-body code. These simulations allow us to study the evolution of galaxy clusters over a period of billions of years, and we can compare simulated results to observations.

On the Needs of High Performance Computing in Computer Aided Decision Support Systems

Sebastian Hegenbart
Universität Salzburg, FB Computerwissenschaften

Several aspects in medical image classification are subject to high performance computing workloads. The core of computer aided support systems is composed of methods used for extracting discriminative information from images or videos. This process is known as feature extraction. A common approach towards extracting features is a statistical analysis of specific signals in the spatial domain, the frequency domain (Fourier domain) or the wavelet domain. Other approaches are based on stochastic models to describe discriminative information within a signal. All those methods share a high computational effort. Endoscopes equipped with high definition cameras are gaining popularity. The increased amount of data adds to the computational requirements.

The second part essential to a computer aided decision support system is the specific classifier, assigning a class label to each extracted feature vector. Common classification methods such as SVM or KNN are heavily based on numerical computations. In the context of medical image classification, a trained classifier is usually capable of performing in real time. However in case of high definition video streams this might not be true in the future.

The evaluation of the accuracy of methods is an essential part of the development of such a system. Cross validation is used to predict how well methods will generalize on an independent data set. Combined with parameter optimization, cross validation leads to a huge number of single classifications and can lead to evaluation times in the magnitude of weeks.

This talk will be specifically targeted at the need of high performance computing in computer aided support systems used in the field of medical image classification.

Minimizing Communication Cost in Succession of Parallel Matrix Operations

Rade Kutil
Universität Salzburg, FB Computerwissenschaften

Parallel matrix operations are well implemented in a variety of parallel libraries. Applications usually consist of a succession of such operations, possibly within a loop. Between such matrix operations, data has to be redistributed to satisfy the initial data distribution requirements of the subsequent operation. However, there is freedom to choose the data distribution arbitrarily in most parallel matrix operations. This can be used to reduce the cost of data distribution between consecutive operations.

Several things have to be developed to do so. First, one needs a way to generate all possible splits in several dimensions given a certain number of processor elements. These splits have to satisfy the

reasonable condition that no dimension may be split into more divisions without producing more total divisions than the number of processor elements. Second, matrix operations are implemented as classes derived from a base operator class to make automatic optimization possible. Third, each operator has to calculate the computational cost given the operand data size and data divisions. Fourth, the optimization algorithm tests all possible divisions for each operator, adds computational and communication cost, and determines the one with the smallest cost.

A special operator that appears even in simple matrix multiplication, when data is split along the common index of the two operands, is the reduce-operator. It is a combination of computation and communication, capable of producing very arbitrary output data distributions. The integration of this operator together with automatic optimization offers a potential that is usually too complicated to exploit in manual parallel programming.

Overall, the system presents an automatic parallelization tool for a common class of applications at high abstraction level compared to e.g. loop parallelizing compilers. It is still in an early development phase, though. Possible future developments are the incorporation of loops and branches, generalization to other than matrix operations (e.g. signal processing with division border overlaps), a simple programming language to express the succession of matrix operations and connection of operands, and automatic calibration of the performance model by test runs of user-supplied operators.

Estimation of Probabilistic Models in (Semantic) Image Retrieval

Roland Kwitt
Universität Salzburg, FB Computerwissenschaften

In this talk, we discuss estimation issues for generative model based approaches in a (semantic) image retrieval, image annotation or scene categorization, a set of core areas of computer vision research. We particularly focus on the challenges involved in estimating generative feature models for high-dimensional features based on millions of feature vectors. Especially in decision-theoretic formulations of the aforementioned vision problems, generative feature model represent the core element of a system in many cases. Furthermore, computer vision tasks are usually subject to strong runtime constraints and thus require computationally efficient algorithms to perform all necessary computations. Just consider the case of automatically searching a massive digital image database based on some user- formulated (semantic) query. It is intuitive that given the usually huge amount semantic concepts (e.g. car, building, person, etc.) and associated generative models, the search task might be quite time-consuming. However, we have to cope with severe timing constraints, due to the user's limited tolerance to wait for the query result.

HPC Challenges in Simulations of Interacting Quantum Many Body Systems

Andreas Läuchli
Universität Innsbruck, Institut für Theoretische Physik

The quest for an understanding of interacting quantum many body systems is of central importance in a variety of fields of physics ranging from quantum magnetism, strongly correlated electrons, ultracold atomic gases and quantum information. Numerical approaches to tackle these systems are particularly important because of the difficulty to obtain accurate analytical solutions due to the inherent complexity. In this talk we give a brief overview of a few powerful computational methods and the HPC challenges one encounters when parallelizing these algorithms for large scale simulations.

Large Scale Simulations of the Euler Equations on GPU Clusters

Manfred Liebmann
Karl-Franzens-Universität Graz, Inst. für Mathematik und wissenschaftliches Rechnen

We investigate the scalability of a parallel Euler solver, using the Vijayasundaram method, on CPU and GPU clusters. The aim of this research is to enable large scale fluid dynamics simulations with up to one billion elements. We investigate communication protocols especially for the GPU cluster to compensate for the slow Gigabit Ethernet network between the GPU compute nodes and to maintain overall efficiency. A diesel engine intake-port and a nozzle, meshed in different resolutions, give good real world examples for the scalability tests on the CPU and GPU cluster configurations.

Data Intensive Computing at the Federated Tier-2 Center for LHC

Dietrich Liko
Öst. Akad. d. Wissenschaft, Inst. für Hochenergiephysik

The Institute for High Energy Physics of the Austrian Academy of Sciences and the Institute for Astro and Particle Physics of the University of Innsbruck are operating a federated Tier-2 computing center for the LHC. The ATLAS and the CMS experiment at Large Hadron Collider (LHC) at CERN in Geneva are recoding data in the range of Petabyte by year. The computing and storage capacity of the Tier-2 computing center is the basis for the successful analysis of this data in Austria. The emphasis of the analysis in Vienna is in the search for new physics in the area of super-symmetry and for standard physics in the area of QCD precision measurements. Since the start of data-taking of the CMS experiment, data is continuously flowing to Vienna. A rate of about 2 TB per day has been reached. An important step is the further data-reduction using the local resources, to allow the physicist to analyze a large number of events. This work is carried out in a local group of computing centers specialized on super-symmetry. To perform their analysis the physicists profit from the use of global and local resources in parallel. In absence of an official Austrian Grid Initiative, a modus operandi with EGI has been found, where the Tier-2 center is using some EGI services. The optimal use of the possibilities offered by EGI would require a full membership. Within Austria the Tier-2 centers sees itself close to the HPC initiatives and the ACSC. For the future a further integration to the Austrian HPC landscape is under consideration by hosting the center near the Vienna Scientific Cluster (VSC) of the Technical University.

Grid Computing in Austria for the ATLAS Experiment

Gregor Mair
Universität Innsbruck, Astro- und Teilchenphysik

In the recent decades the landscape for distributed computing changed drastically due to the increasing hardware requirements of scientific and industrial applications. With requirements of several thousand CPU hours or several PiB of storage per month, user groups are confronted with a variety of different approaches, such as local clusters, grids, clouds, supercomputers or desktop grids, with different strengths and weaknesses.

In particle physics the experiments at the LHC in Geneva faced tremendous requirements on their computing infrastructures very early. In order to process, to store and to make the data available to researchers worldwide, the Worldwide LHC Computing Grid (WLCG) was founded. Also Austria participates within the WLCG: the Institute for High Energy Physics of the Austrian Academy of Sciences and the Institute of Astro- & Particle Physics of the University of Innsbruck operate the Austrian Federated WLCG Tier-2.

Already in 2003 a first WLCG grid site was installed in Innsbruck and the institute participated in several projects and tests of the grid technology. Since the start of operation of the LHC in 2010

Innsbruck operates as an ATLAS associated Tier-2 within the German WLCG cloud.

In this talk the different approaches for distributed computing as seen by particle physics are characterized and compared. Then, experiences with the grid in Innsbruck are discussed and some statistics and results are shown. Finally, ongoing research in Innsbruck and future plans for the site in Innsbruck are presented.

Evolution of Robotic Neurocontrollers with Intrinsic Noise

Helmut A. Mayer
Universität Salzburg, FB Computerwissenschaften

We report on experiments with robotic neurocontrollers with intrinsic noise evolved for a peg pushing task. The specific controller of the simulated robot is a feed–forward network with noisy weights, i.e, the weight values are perturbed by additive, normal noise. Artificial evolution is employed to generate and gradually improve the neurocontroller (the robotic "brain") according to a fitness function measuring the success of the robot in solving the task. Due to the non-deterministic nature of the evolutionary process a sufficient number of evolutionary runs has to be performed in order to gather statistically reliable results. Naturally, these runs can be performed in parallel on a high performance computing system. Considering that a single evolutionary run may take hours or even days on a standard computer, the benefit of parallel computation becomes evident.

Our initial motivation to investigate the effects of noise affecting the internal computations in an Artificial Neural Network (ANN) stems from Biological Neural Networks (BNNs) facing the stability–plasticity problem. Noise is ubiquitous in BNNs as the physical and chemical states of cells change constantly even without external excitation. Then, the question arises, how the remarkable stability of biological systems can be achieved, e.g., the precise movements of a human hand. Actually, there is evidence that intrinsic noise may contribute positively to stability in BNNs.

Consequently, we investigate the effects of intrinsic noise in a robot's neurocontroller, which may not only serve as a model for biological phenomena, but may also improve the capability of a robot to deal with external noise or, more generally, with imprecise information. In order to analyze the impact of intrinsic noise the neurocontrollers are evolved in a noise–free environment, and the best–performing networks are then tested in noisy environments, where peg movement and sensor signals are afflicted by noise. We find that the internal (robotic brain) noise is beneficial in coping with external noise, especially, in the case of noisy sensors.

All experiments have been conducted using the high performance computing facilities at the Department of Computer Sciences at the University of Salzburg.

exciting@web, HPC as a Web Service

Christian Meisenbichler
Montanuniversität Leoben, Atomistic Modelling and Design of Materials

We will report on exciting@web, a soft-ware development that allows for setting up calculations through a graphical user interface (GUI) and, at the same time, serves as a database. Our tool has the ability to connect to an HPC service via an HTTP API to perform the calculations. These are carried out by the *exciting* package, which is a full-potential all-electron code based on density-functional theory. This approach reflects the idea to use the software as a service (SAS), that is present in todays internet and cloud computing. In this talk, we will demonstrate the functionality of the software and give an outlook on the opportunities which are provided by the concepts behind.

Algebraic Multigrid Solver on Clusters of CPUs and GPUs

Aurel Neic
Karl-Franzens-Universität Graz, Inst. für Mathematik und wissenschaftliches Rechnen

Elliptic solvers are used in a wide area of scientific problems. We will present a highly parallel implementation of a conjugate gradient solver with an algebraic multigrid preconditioner in a package called Parallel Toolbox. In addition also a current application using our toolbox will be investigated. Benchmark results of computations on clusters of CPUs and GPUs will be presented. They will show that a linear equation system with 25 million unknowns can be solved in about 1 second.

N-body / Hydrodynamic Simulations of Galaxies in a Cluster Environment

Dominik Steinhauser
Universität Innsbruck, Astro- und Teilchenphysik

Clusters of galaxies are the largest, gravitationally bound structures in the universe. About twenty percent of the mass of such a cluster belongs to the so called intra-cluster medium (ICM), a very hot and thin gas. The galaxies themselves only account for about 5 % of the mass, the rest is dark matter. The galaxies, moving in the cluster's gravitational potential, experience a force acting on the gaseous part of the galaxy, the inter-stellar medium (ISM), caused by the ICM. This ram pressure induces star formation and enforces the galaxy to lose part of its gas. This process is called ram-pressure stripping.

For a better understanding of this process and to better interpret observations of galaxies undergoing ram pressure, we perform combined N-Body and Hydrodynamic simulations of such scenarios. For this purpose we are using the Tree-SPH code GADGET-2, which is a highly parallelised code which uses the message passing interface (MPI) for communication between processing nodes and is hence well suited for our distributed memory machines. For the simulations of ram pressure stripping, up to 10^8 particles are used and the computation is done using up to 256 computing nodes.

In our simulations, radiative cooling and a recipe for star formation is included. To simulate the ram pressure acting on the model galaxies, a special wind tunnel setup was developed. Using this setup, the influence of a bulge on the morphology and the star formation of a galaxy undergoing ram pressure was investigated. Model galaxies with different bulge sizes were used to accomplish several ram-pressure stripping simulations with varying ICM densities and different inclination angles.

We find that the star formation is enhanced with increasing ICM density. However, a bulge suppresses the star-formation rate when the same ram pressure is applied. Furthermore, the denser the surrounding gas, the more of the inter-stellar medium (ISM) is stripped, but again a bulge prevents the stripping of the ISM. Thereby, less stars are formed in the wake if the galaxy contains a bulge. The dependence of the star-formation rate on the disc tilt angle is not very pronounced. Just a slight trend of decreasing star formation with increasing inclination angle can be determined. Also, with increasing disc tilt angle, less gas is stripped and therefore less stars are formed in the wake.

Vienna Scientific Cluster – Status and Outlook

Herbert Störi,
P. Berger, E. Haunschmid, B. Hermann, D. Kvasnicka, J. Zabloudil
Technische Universität Wien, Angewandte Physik

The first Vienna Scientific Cluster, now termed VSC-1, was installed in summer 2009 and went into operation in late autumn. VSC-1 is a joint project by the University of Vienna, the University of Natural Resources and Life Sciences and the Vienna University of Technology, managing the system.

The System contains roughly 4000 2.66 GHz Intel i7 cores and a brief technical overview of the system will be given. The fact that full load was reached within a month points to the fact, that high performance computing resources were urgently needed at the time. Presently a lively community of roughly 70 active projects is using VSC-1. Access and operation policies will be shortly reviewed and a very brief overview over the areas of application will be given. After some upgrades, the LINPACK performance is in the order of 40 TFlops.

The second system, VSC-2 is presently being installed and the status at the time of the seminar will be reported. VSC-2 is based on low power 2.2 GHz 8-core AMD processors and contains a total of 21000 cores. This system is mainly intended for highly parallel jobs. While presently the same universities as for VSC-1 are part of the project, the project is open for participation by other Austrian universities.

Finally the economics, especially aspects of energy consumption, of the systems will be reviewed and an outlook to VSC-3 and the more distant future will be given.

Insieme: A Multi-Objective Optimising Compiler for HPC

John Thomson
Universität Innsbruck, Institute of Computer Science

Optimisation in the HPC world has, in the past, focussed purely on performance. As power design constraints and energy costs have risen in prominence, so too has the focus on optimisation changed to account for these two optimisation objectives - performance and energy.

Insieme is an optimising compiler, developed at the University of Innsbruck, which aims to not only optimise performance and power, but provider a generic multi-objective optimiser which can take into account n optimisation objectives, such as computing costs, security and reliability. Important to note is that these objectives interact with each other - for instance, increasing the performance may also increase the computing cost and energy consumed. Insieme aims to find the best possible optimisation, taking into account multiple optimisation objectives.

HeuristicLab Hive – An Open Source Environment for Parallel and Distributed Execution of Heuristic Optimization Algorithms

Stefan Wagner
FH Oberösterreich, Studiengang Software Engineering

HeuristicLab is an open source environment for solving hard optimization problems with heuristic and evolutionary algorithms. It is developed by the members of the Heuristic and Evolutionary Algorithms Laboratory (HEAL) of the Upper Austria University of Applied Sciences and is used in education, research and industry projects.

Evolutionary algorithms are often applied on hard optimization problems and usually require much CPU time in order to find satisfying solutions. Additionally, they provide various parameters which have a significant impact on the algorithm's behavior and which therefore have to be configured by the user. Finding good parameter settings is not a trivial task and requires to run an evolutionary algorithm several times which results in an even longer runtime.

To improve performance, a distributed computing infrastructure called HeuristicLab Hive has been developed. Hive is fully integrated into HeuristicLab and speeds up evolutionary algorithms by distributing their execution on multiple computers. HeuristicLab can thereby be used to upload work packages to the Hive server. The server queues all uploaded jobs and, if resources are available, distributes them to its calculation slaves. Theses slaves execute the jobs and store the results on the server again from where they can be obtained. Additionally, Hive supports that running jobs can be paused, changed and then again be redistributed to the clients. As HeuristicLab provides a generic model for representing parallel algorithms, algorithms do not have to be adapted for Hive, making

it very easy to use Hive without any need to write parallelization and distribution specific code.

Hive slaves currently run on a cluster and will be deployed on PCs in the computer labs of the Upper Austria University of Applied Sciences in the near future. These computers can then be used to execute jobs during the time when the labs are not occupied, making it possible to utilize the CPU power of idle computers.

Molecular Dynamics Simulations of Biological Systems – Specificity in Snake Venom Metalloproteases

Hannes G. Wallnöfer,
T. Lingott, J. M. Gutiérrez, I. Merfort, K. R. Liedl
Universität Innsbruck, Institute of General, Inorganic and Theoretical Chemistry

Molecular Dynamics Computer Simulations are a viable tool to study the evolvement of biological systems during time. Consequently, insights into function as well as thermodynamic properties are accessible.

Frequently occurring biological paradigms are protein-protein interfaces having crucial functions in many biological processes [1]. The large interaction areas of such interfaces show complex interaction motifs. Even more challenging is the understanding of (multi-)specificity in protein-protein binding. Many proteins can bind several partners to mediate their function [2].

A perfect case to study such multi-specific protein-protein interfaces are snake venom metalloproteases (SVMPs) [3]. Inherently, they bind to a variety of basement membrane proteins of capillaries, hydrolyze them, and induce profuse bleeding.

Our results indicate that the activity to induce hemorrhage, and thus the capability to bind the potential reaction partners, is related to the backbone flexibility in a certain surface region [4]. A subtle interplay between flexibility and rigidity of two loops seems to be the prerequisite for the proteins to carry out their damaging function. Presumably, a significant alteration in the backbone dynamics makes the difference between SVMPs that induce hemorrhage and the inactive ones.

References

[1] D.J. Mandell, T. Kortemme. *Computer-aided design of functional protein interactions.* Nat Chem Biol 2009, 5:797-807.

[2] J.D. Han, T. Hao, D.S. Goldberg, G.F. Berriz, L.V. Zhang, D. Dupuy, A.J.M. Walhout, M.E. Cusick, F.P. Roth, M. Vidal. *Evidence for dynamically organized modularity in the yeast protein-protein interaction network.* Nature 2004, 430:88-93.

[3] J.B. Bjarnason, J.W. Fox. *Hemorrhagic metalloproteinases from snake venoms.* Pharmac Ther 1994, 62:325-372.

[4] H.G. Wallnöfer, T. Lingott, J.M. Gutiérrez, I. Merfort, K.R. Liedl. *Backbone Flexibility Controls the Activity and Specificity of a Protein-Protein Interface: Specificity in Snake Venom Metalloproteases.* J Am Chem Soc 2010, 132:10330-10337.

Distributed Computation of Feature-Detectors for Medical Image Processing on GPGPU and Cell Processors

Peter Zinterhof
Universität Salzburg, FB Computerwissenschaften

Automated classification of medical (computed tomography) images may ultimately lead to faster and improved diagnosis, benefiting both patients and clinicians. We describe a software system, that

can be trained for classification purposes in the area of medical image processing. The underlying algorithm is based on a set of perceptron-like feature detectors, which are combined to short feature vectors. Those are used to form self-organized Kohonen maps, which will be used for the classification of new image data. The exact description of the feature detectors is derived from a large set of sample images by way of an evolutionary strategy. This leads to a computationally demanding process of iterated image decomposition, Kohonen map training and quality assessment. To make our method feasible, we rely on clusters of rather cheap commodity hardware, namely general purposes graphics processing units (GPGPU) and the STI Cell Broadband Engine Architecture (Cell), as it comes with the PS3 gaming console.