NIC Workshop: Hybrid Particle-Continuum Methods in Computational Materials Physics

From March 4 to 7, 2013, the NIC Workshop "Hybrid Particle-Continuum Methods in Computational Materials Physics" was held at Jülich Supercomputing Centre (JSC). It was organized jointly by the NIC research group "Computational Materials Physics" and the Institute of Advanced Simulation. The goal of the workshop was to foster the exchange of ideas between the communities working on complex fluids and on complex solids. 65 participants could learn about the new developments in hybrid particle-continuum methods from 15 invited and 10 contributed talks as well as from 30 posters.

During the workshop, particular emphasis was placed on continuum-mediated interactions between particles as well as on the adaptive and non-adaptive coupling between particle-based and continuum-based descriptions of materials. The subjects covered included the modeling of hydrodynamic interactions between particles in complex fluids or environments, through coarse-grained descriptions of biological systems, to the coupling of atomically represented regions with various continuum-based theories for fluids and solids. Special aspects were long time scale properties of systems with slow collective dynamics, the development of efficient adaptive resolution algorithms, and the coupling of quantum-mechanically treated regions with continuum descriptions.

The proceedings of the workshop were edited by Martin Müser, Godehard Sutmann and Roland Winkler (Forschungszentrum Jülich). The volume is available either as hard copy or as PDF file on the web http://juser.fz-juelich.de/record/132949/files/

• Martin Müser

Jülich Supercomputing Centre

NVIDIA Application Lab at Jülich

Accelerating scientific HPC applications using GPUs has become popular. For many applications it has indeed proved to be a very successful approach. Nevertheless, there is still a lot to learn about algorithms and methodologies for porting applications. Enabling more scientific applications for GPU-based architectures is a core goal of the NVIDIA Application Lab at Jülich. The lab is jointly operated by Jülich Supercomputing Centre (JSC) and NVIDIA since July 2012. For JSC it is yet another step to further strengthen links to vendors that are actively pursuing an exascale strategy.

During its first half year of operation the main focus of the lab was to establish a broad application portfolio encompassing computational neuroscience, high-energy physics, radio astronomy, data analytics and others. In common with each other, the applications have a high level of parallelism, ideally with few dependencies between tasks or task groups. One example is an application from the JuBrain project developed at the Jülich Institute for Neuroscience and Medicine INM-1 [1]. The project will result in an accurate, highly detailed computer model of the human brain. This atlas is created by reconstructing fibre tracks from pictures of thousands of slices (see Fig. 1). The process of mapping the pictures, called registration, requires repeated computation of a metric that measures how pixels of two pictures map to each other, a computationally expensive process that maps well to the GPU.

In applications from experimental physics, a natural data decomposition may be according to how the data leaves the detector. For instance, in high-energy physics experiments data is generated at such extremely high data rates that data for different time slices has to be distributed to different processing devices. Another example is search for pulsars in radio astronomy, where data sets from different beams and different measurements have to be repeatedly processed resulting in thousands or even millions of Fast Fourier Transforms.

Large quantities of data which need further processing are created not only by experiments and observatories, but increasingly from "computational experiments" such as Monte Carlo simulations of protein folding. Here clustering, a standard method of data analytics, is applied to identify regions of similar objects in multi-dimensional data sets. At the lab we have shown that sub-space clustering algorithms can be very efficiently implemented on GPUs opening the path to analysis of high-dimensional data sets in other areas as well [2].

For many applications, using a single GPU is not sufficient, either because more computing power is required, or because the problem size is too large to fit into the memory of a single device. This forces application developers to not only consider parallelization at device level, but also to manage an additional level of parallelism. Depending on the application this may be challeng-

ing since on most architectures device-to-device communication through the network is implemented as split transactions via the host processor. Software solutions like CUDA-aware MPI implementations can help mitigate this problem, but ultimately better hardware support is needed to interconnect GPU and network devices.

The goal of the lab is not solely to provide service to application developers and achieve performance improvements for their applications. To improve future architectures and their usability for scientific applications it is necessary to better understand how well these applications map onto such architectures. The introduction of a new architecture, Kepler, or more specifically, the GK110 GPU, is a good opportunity to learn about the effects of architectural changes by means of comparison with the previous architecture. How can application kernels cope with a significant increase in compute performance when bandwidth to device memory becomes only moderately larger? Do other changes in the memory hierarchy allow compensating for the increased flops-per-byte ratio? First experience shows that a careful analysis of the utilization of all levels of the memory hierarchy helps to make the application use the new architecture as efficiently as its predecessor. A better understanding at this level helps not only application developers but also processor and systems architects to improve GPU-based architectures for scientific computations.

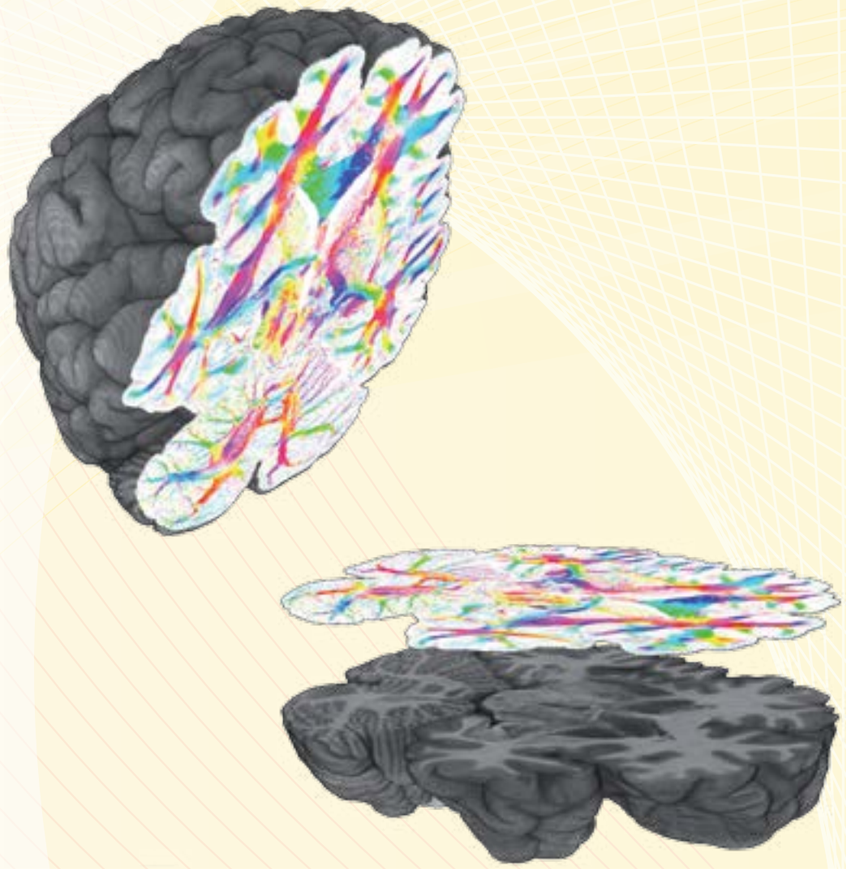

Figure 1: Many cuts of a human brain are required to build an accurate, highly detailed computer model of the human brain [1].

Acknowledgements

We would like to thank all application developers that have been involved in the lab so far, in particular: Markus Axer, Andreas Herten, Marcel Huysegoms, Jan Meinke (all Forschungszentrum Jülich) and David Champion (MPI for Radioastronomy).

References

[1] http://www.fz-juelich.de/JuBrain/EN

[2] **A. Adinetz, J. Kraus, J. Meinke, D. Pleiter,** „GPUMAFIA: Efficient Subspace Clustering with MAFIA on GPUs," submitted to Euro-Par 2013.

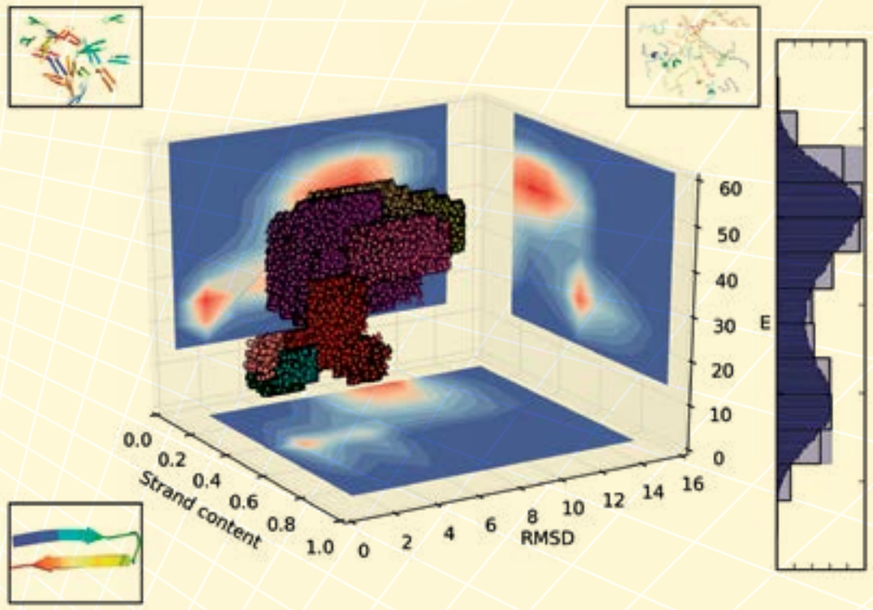

Figure 2: Clusters identified in data obtained by Monte Carlo simulations of protein folding using MAFIA [2].

• Andrew Adinetz
• Jiri Kraus
• Dirk Pleiter

Jülich Supercomputing Centre