

Blue Gene Extreme Scaling Workshop 2011

From February 14 to 16, JSC organized its 2011 Blue Gene Extreme Scaling Workshop. Similar to the previous workshops in 2009 and 2010, the main focus was on application codes able to be scaled up during the workshop to the full Blue Gene/P system JUGENE, which consists of 72 racks with a total of 294,912 cores – still the highest number of cores available worldwide in a single system.

Interested application teams had to submit short proposals which were evaluated. Selection criteria were the required extreme scaling, the application-related constraints which had to be fulfilled by the JUGENE software infrastructure, and the scientific impact that the codes could produce. The process was very competitive: 8 out of the submitted 15 high-quality proposals

were selected. Participating teams came from the Princeton Plasma Physics Laboratory in the US, the Royal Institute of Technology (KTH) in Sweden, the King Abdullah University of Science and Technology (KAUST) in Saudia Arabia, the Mickiewicz University in Poland, University College London in UK, the University of the Basque Country in Spain, RZG MPI, and the University of Heidelberg.

During the workshop, the teams were supported by JSC parallel application experts, the JUGENE system administrators and two IBM MPI and compiler experts. In addition, the participants shared their expertise and knowledge. The workshop was extremely successful: the 8 teams succeeded in submitting one or more successful full 72-rack jobs for 11(!) different applications

during the course of the workshop as some teams experimented with more than one code. A total of 308 jobs were launched using 122 rack days of the total 157 rack days reserved for the workshop.

Many interesting results were achieved: The team of the University of the Basque Country uses the time-dependent density functional theory code Octopus

to perform first principle simulations of the excited states of chlorophyll molecules combined with other chromophores and proteins which form the different complexes that take part in photosynthetic processes. This type of simulations has a direct impact on our understanding of photo-induced processes in biological systems. The large size of these molecules, of the order of thousands or tens of thousands of atoms, makes them very challenging to model computationally. During the workshop, they performed runs simulating the absorption of light by these large molecular systems. For the first time, they were able to simulate molecules with 2,676 and 5,879 atoms (the whole chlorophyll molecule) using all 72 racks of the machine.

Researchers from the Royal Institute of Technology (KTH) in Sweden investigated the scaling of large-scale neural simulations using an experimental neural simulator (ANSCore) and an experimental code simulating a model of the neocortical network of the brain (BrainCore). Both codes successfully scaled up to the full system scale when applied to large-scale neural simulations. Scaling of core components for building a variety of neural models was featured in ANSCore. There, communication was handled by MPI collective calls while the BrainCore model featured a large homogeneous single network and a straight-forward application in terms of associative memory using point-to-point communication. The results achieved open the path for simulations of neural models of sizes comparable to real mammalian nervous systems with a much higher complexity than so far attempted. They will allow to handle spiking communication in models of

the neocortical network as these scale up to sizes of real mammalian brains. Furthermore, their study paves the way for the use of extremely scalable brain network models for information processing of data obtained with e.g., large-scale sensor arrays. The knowledge gained can also be used to investigate the design of dedicated hardware.

Finally, the group of KAUST investigated the scaling of their Billions-Body Molecular Dynamics (BBMD) code which is a highly optimized, parallel C++ MD code for Lennard-Jones particle systems. The code is used to study for the first time the behaviour of large-size structured glasses characterized by tens of billions of particles. This will permit to answer long-standing questions in the field of complex systems concerning the existence and the dynamics of specific wave vibrations of structurally disordered systems like glasses. The BBMD code was successfully run on all 72 racks of the Blue Gene/P system for experiments with 10 billion particles. The scaling from 32,786 cores (8 racks of BG/P) to 294,912 cores (72 racks of BG/P) shows an efficiency of 89%. The amount of communication compared to elapsed time can be seen to be less than 2% for 8 racks and only grows to 7% for 72 racks. This shows that utilizing mainly nearest neighbour communication, with a very limited amount of global communication, is very advantageous for extreme scaling.

All experience gathered during the workshop will be summarized in a technical report. For more information on the workshop as well as on the reports of the previous workshops see:

<http://www2.fz-juelich.de/jsc/bg-ws11>

Activities

Activities



• Bernd Mohr

Jülich
Supercomputing
Centre