NSC News
National Supercomputer Centre at Linköpings universitet
2007:3

LCSC

LINUX CLUSTERS FOR SUPERCOMPUTING

For the eight year in a row, NSC invites you to the Linux Clusters for Supercomputing workshop.

The LCSC workshop will take place on October 16–18 at Linköping University's Campus Valla.

Technical programme • Tutorials • Inauguration of the new 60 Tflops cluster

For more information and to register for the event, please visit http://www.nsc.liu.se/lcsc

NSC
National Supercomputer Centre

Multiphysics Modelling

With "multiphysics modelling" methods can be combined that have different physical content in order to obtain a more complete view of an applied problem. In materials science this can involve methods that bridge length and time scales; in length from atomistic to macroscopic levels through "electrons, atoms, grains and grids"; in time from femtosecond electron dynamics to slow processes perceived by humans.

Read further on page 3

Neolith—the new NSC computer

NSC's new supercomputer cluster will allow users of Swedish high performance computer systems to tackle problems on a new scale.

Read further on page 5

Big bills

Sven Stafström
NSC Director
stafstrom@nsc.liu.se

Cost for electricity and cooling of supercomputers is becoming more and more of an issue. This is indeed understandable considering the amount of electrical energy converted into heat in modern supercomputers. It is of course not only a matter of economy but also an environmental problem, which we share with other energy intensive activities initiated by humans.

Monolith has a FLOPS (peak)/W ratio of about 50 MFLOPS/W, while for our most recent HPC resource, Neolith (with a peak performance of 60 TFLOPS), the ratio is 200 MFLOPS/W. Taking the leap towards petascale computing with today's ratio we end up with a power of 5 MW and more. The good thing is that performance increases faster than power but on the other hand the cost of energy will probably increase (which in fact also is a positive trend since it will stimulate energy saving). We can expect to reach 1MW for an affordable (considering the SNIC budget for hardware) system before the year 2010, but can we afford the energy costs? A rough estimate of the electricity (and cooling) cost for such a system is around 10 Mkr/year. A few more years ahead the costs for electricity will be higher than the hardware investments. Clearly, if Sweden is interested in providing state-of-the-art HPC facilities there has to be a discussion between universities, which today pay the electricity bill, and SNIC and/or KFI how these costs should be covered. My personal opinion is that the investments and the running costs should be considered together in order to find the most efficient solution from the point of view of total costs.

The environmental issue should also be discussed. Even though LiU has the policy to use "green electricity" there is still an environmental cost for using electrical power. The specific problem for NSC is of course that we want big machines and big machines need MWatts of power. Can we still be environmentally friendly? In order to take a very explicit step in the right direction I propose that LiU, together with NSC, should invest in a wind power turbine. Today, multi-MW turbines are available at a cost which, on a long term basis, can reduce the electricity bill. Why not put the turbine, with the LiU and NSC logo clearly visible along the E4 motorway to show everybody our intentions!

NSC storage capacity

NSC's new storage system consists of a combination of SATA and FC-disks. This combination is designed to fulfil requirements that combine storage capacity and access capability. The mass storage system is a Nexsan SATAbeast with a total capacity of 105 TiB. In addition we acquired an IBM DS4700 system with a total of 7 TiB of FC-disks. The storage system is delivered by Pulsen AB. This is a first step in a major upgrade of NSC's storage capacity, which in addition to the above described general purpose storage system also includes storage systems for the CERN/LHC project WLCG, the neutrino detection project IceCube, and for Climate Research. In total we expect to have a disk storage capacity exceeding 500 TiB some time next year and tape storage with about the same volume.

Reminder of visualization service at NSC

NSC offers a visualization service to our SNAC users. This service is predominantly concerned with the visualization of molecular and crystal structures and electron densities. NSC provides access to our supercomputers from a common and dedicated visualization server named Inspector that collects graphical software for the interaction with our set of electronic structure programs. It is our sincere hope that you, by use of the visualization server, find it even easier to do computational research at NSC.

As member of an active SNAC project you are welcome to request an account on Inspector and start using this service (send an e-mail to support@nsc.liu.se). You may learn more about our new service from our homepage: http://www.nsc.liu.se/software/visualization/

Multiphysics Modelling

With "multiphysics modelling" we combine methods that have different physical content in order to obtain a more complete view of an applied problem. In materials science this can involve methods that bridge length and time scales; in length from atomistic to macroscopic levels through "electrons, atoms, grains and grids"; in time from femtosecond electron dynamics to slow processes perceived by humans, see Fig 1. The combination can be obtained in terms of integrated algorithms or simply by piping datasets from one model to the other. At our department we have applied the multiphysics concept to some contemporary biochemistry problems and to the area of non-linear optical materials. We have, for instance, studied optical power limiting through the combined use of quantum mechanics (QM) and classical electrodynamics [1]; modelled the design of electro-optical switching materials through a combination of QM and classical dynamics [2]; studied multiphoton induced light emission in quantum nanoparticles with QM coupled to statistical mechanics [3]; estimated solvent effects on so-called optical Kerr effects with polarizable continuum models [4], i.e. combining QM with dielectric theory, to mention a few examples.

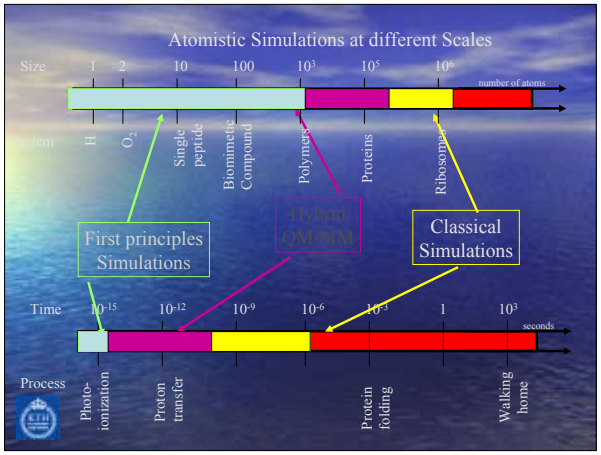

Fig. 1 Time and length scales.

Self-activating materials

The project concerned with finding optical power limiting materials is part of a FOI supported effort within their Nanotechnology program to derive materials that protect sensors and eyes from strong laser light. My group at KTH, together with the group of Patrick Norman at LiU, serve as the theoretical modelling unit in this effort, which also contains Swedish groups involved in chemical synthesis, measurements and device fabrication. The project has presented challenges of basic academic nature as well as very applied problems; the goal is a product with specified performance. The key point of the material is that it should be self-activating, that is, while it should be transparent for light of normal intensity, it should self-activate itself so that it becomes opaque when it is hit by a strong laser pulse.

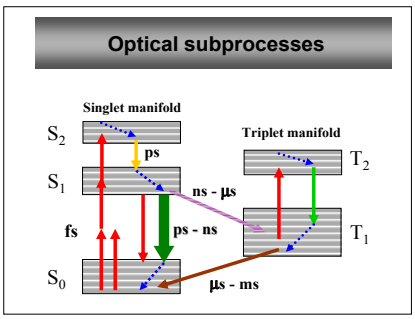

Fig. 2 Jablonski diagram.

The idea is to find a material which is two-photon active as the two-photon cross section depends quadratically on the light intensity. However, nature is not so easily controlled; at high intensities various saturation phenomena occur, and many photons from one and the same laser pulse will be involved in side processes that are crucial for the final outcome of the power limiting capability. In principle we need to master the full so-called Jablonski diagram, see Fig. 2, a quite formidable challenge. Luckily, our basic bread-and-butter program DALTON has a battery of non-linear properties at its disposal, and

New staff member: Pär Andersson

I am a Computer Science student at Linköping University and since this spring I am doing my master thesis project at NSC. My project is to modify a quantum chemistry program from MPI parallelization to a hybrid model using both MPI and OpenMP which will be more memory efficient on multi-core processors.

Since June I am also part time employed at NSC. I mainly work as a system administrator on Neolith, installing and configuring a working software environment on this new resource. This of course also involves support for our test pilot users as Neolith is now in a testing phase.

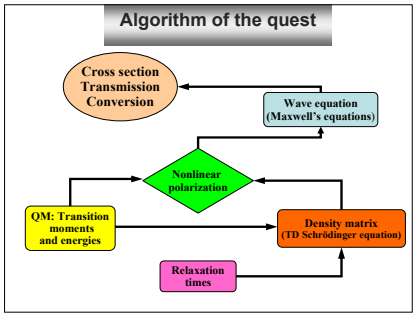

Fig. 3 Multiphysics modelling.

can be used for modelling the electronic interactions in the diagram.

As many photons in the pulse are active at the same time in the Jablonski diagram, the OPL outcome will be highly dependent on the characteristics of the pulse; its strength, rate, duration, its temporal and spatial distribution. The "multiphysics" character of the problem appears as having to carry out three kinds of modelling, see chart in Fig 3, quantum chemistry (response theory) to obtain the basic cross sections; solving differential equations to get the

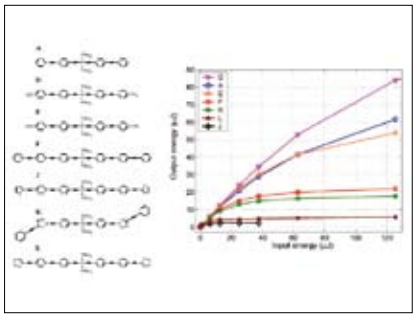

Fig. 4 Optical power limiting of some Platinum complexes.

population rates and total polarization; and solving the wave equation to get the pulse propagation in the material. Each of these contains its own theoretical considerations, approximations and numerical challenges. The project, which thus has had strong modelling support throughout, has been successful, a material with the specified performance has been derived and fabricated. The remaining efforts concern incorporation of the material into a device, and field tests, see Figure 4 for some results in this project.

Combining QM and MD modelling

The second example of multiphysics modelling combines QM and molecular dynamics (MD). The goal of the project, which is EU based, is to derive an all-organic electro-optic switching device, Fig. 5. It relies on the action of electric field induced second harmonic generation (EFISH), which is a nonlinear optical property. The quantum part is to find organic chromophores with an optimum product of dipole moment and second harmonic generation (SHG) coefficient. With a dipole moment the chromophore can be poled by an external field and can so be aligned to cooperate in generating a macroscopic SHG effect. However, in this lies a problem in that dipolar molecules tend

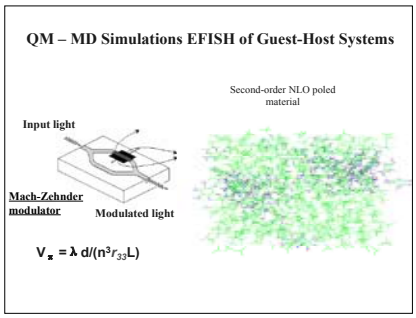

Fig. 5 Electro-optic switching device.

to aggregate and thus cancel the possibility to become poled.

A solution to this problem is to put the chromophores in a host, a polymer, the property and action of which must be mastered at widely different temperatures. Poling occurs above the glass transition temperature, the system should then be cooled, followed by relaxation without electric field. Moreover, a problem, here as well as in general, is to know how the external laser field transforms into the local region, as this cannot be measured. These issues have all been attacked by a combined QM-MD formalism, see Fig 6 (see page 5), which has been used to screen several guest-host materials for synthesis.

Computational challenges

It is clear that the use of multiphysics approaches, like those briefly reviewed above, put extra demands on the capability and flexibility of the computer platform on which they are

New staff member: Daniel Johansson

Hello, my name is Daniel Johansson and I began working for Nordic Data Grid Facility (NDGF) here at the National Supercomputer Centre in June.

I have studied Computer Science at Linköping University and I have worked part-time both as a system-administrator at the Institution for Computer Science and as a programmer here at the National Supercomputer Centre. I will work mainly on the ARC grid middleware, a software that is used for job-management within grids such as for example, NorduGrid.

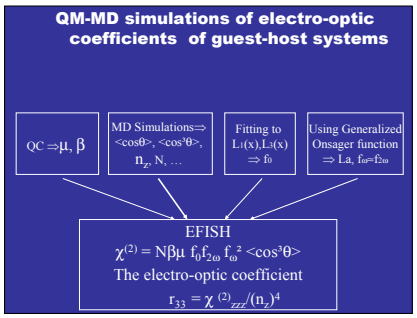

Fig. 6 Simulation of electro-optic coeffcients.

implemented. For instance, quantum chemistry and molecular dynamics have widely different requirements and hardware limitations. Quantum chemical calculations perform time- and memory-consuming steps, mostly involving linear algebra operations, either diagonalization or spectral transformations consisting of a series of matrix-matrix multiplications. Larger calculations require distributed sparse matrix representations and effective matrix rebalancing algorithms. Linear sets of equations need to be solved in parallel as well. For very large calculations, fast multipole expansion methods are used which involve a considerable communication between nodes. Thus future breakthrough calculations call for large node memory with fast, low-latency interconnect for a satisfactory scaling with the number of nodes. This should be compared with Molecular Dynamics which is normally limited by the hardware floating-point performance, but to enable simulations of large systems it is also crucial to scale to hundreds or thousands of nodes, which puts severe restraints on the balance between floating point and internode communication bandwidth performance.

References

1. A quantum mechanical-electrodynamical approach to nonlinear properties: A. Baev, P. Welinder, R. Erlandsson, J. Henriksson, P. Norman and H. Ågren, J. of Nonlinear Optical Physics & Materials, 16, 157 (2007).

2. Electric field poled nonlinear optical systems, Y. Tu, Q. Zhang, and H. Ågren J. Phys. Chem. B 111, 3591 (2007)

3.Dynamics of multiple-photon optical processes in semiconductor quantum, Y. Fu, T.-T. Han, Y. Luo, and H. Ågren, J. Phys. Condens. Matter, 18, 9071 (2006).

4. Solvation of azide at the water surface: the Polarizable Continuum Model approach, L. Bondesson, L. Frediani, H. Ågren, and B. Mennucci, J. Phys. Chem., 110, 11361 (2006).

Hans Ågren
Department of Theoretical Chemistry, School of Biotechnology
Royal Institute of Technology
SE-10691, Stockholm, Sweden
e-mail: agren@theochem.kth.se

Hans Ågren received his PhD 1979 in Experimental Physics at the former Uppsala University. He is currently Professor and Head of the Department of Theoretical Chemistry at the School of Biotechnology, KTH.

Neolith, new capability system at NSC

NSC's new supercomputer cluster will allow users of Swedish high performance computer systems to tackle problems on a new scale. The cluster has been named Neolith after the first tools that humans started to use in the beginning of the Neolithic, or new stone age, time period 11,000 BC. A small part of the system arrived at NSC in May. It consists of 36 nodes and was installed and subject to system test during June and July. A limited number of test pilots were given access to the systems in early August. Their tests, which include applications such as Gaussian03, Fluent, Star-CD, VASP, Gromacs etc., show very good single node performance and scaling behaviour in accordance with the design goals.

The second and main part of the system will be installed and fused with the first part during September. Neolith will then have 805 computer nodes, each equipped with two Quad-Core Intel® Xeon® E5345 processors, i.e. a total of 6440 cores. The processors are connected with a high-performance Infiniband interconnect with switches delivered by Cisco®.

Our hope is that Neolith will be the tool that allows our users to tackle new scientific problems and challenges. This may involve modifications of existing codes to ensure maximal scalability. The NSC staff is eager to assist in such a development and in any other matter that relates to an efficient use of Neolith. We are also looking forward to all the scientific results that will emerge from the use of Neolith in the years to come.

UPCOMING EVENTS

LCSC 2007; The 8th Annual Workshop on Linux Clusters for Super Computing

October 16-18, 2007, Linköping, Sweden.

http://www.nsc.liu.se/lcsc

HiPC 2007; 14th IEEE International Conference on High Performance Computing

December 18-21, 2007, Goa, India.

http://www.hipc.org

SIAM Conference on Parallel Processing for Scientific Computing

March 12-14, 2008, Atlanta, Georgia, USA.

http://www.siam.org/meetings/pp08/

HiCOMB 2008; 7th IEEE International Workshop on High Performance Computational Biology

April 14-18, 2008, Miami, Florida, USA.

http://www.hicomb.org

IPDPS 2008; 22nd IEEE International Parallel & Distributed Processing Symposium

April 14-18, 2008, Miami, Florida, USA.

http://www.ipdps.org

JSSPP 2008; 14th Workshop on Job Scheduling Strategies for Parallel Processing, in conjuction with IPDGF 2008

April 18, 2008, Miami, Florida, USA.

http://www.cs.huji.ac.il/~feit/parsched/jsspp08/

I-SPAN 2008; The 9th International Symposium on Parallel Architechtures, Algorithms and Networks

May 7-9, 2008, Sydney, Australia.

http://www.cs.usyd.edu.au/~ispan08/

CCGrid'08; 8th IEEE International Symposium on Cluster Computing and the Grid

May 18-22, 2008, Lyon, France.

http://ccgrid2008.ens-lyon.fr

Linköpings universitet

National Supercomputer Centre, Linköpings universitet, 581 83 Linköping, Sweden

Tel: +46 (0) 13-28 26 18, fax: +46 (0) 13-28 25 35, e-mail: info@nsc.liu.se

www.nsc.liu.se