PDC Center for
High Performance Computing

no:2 2009

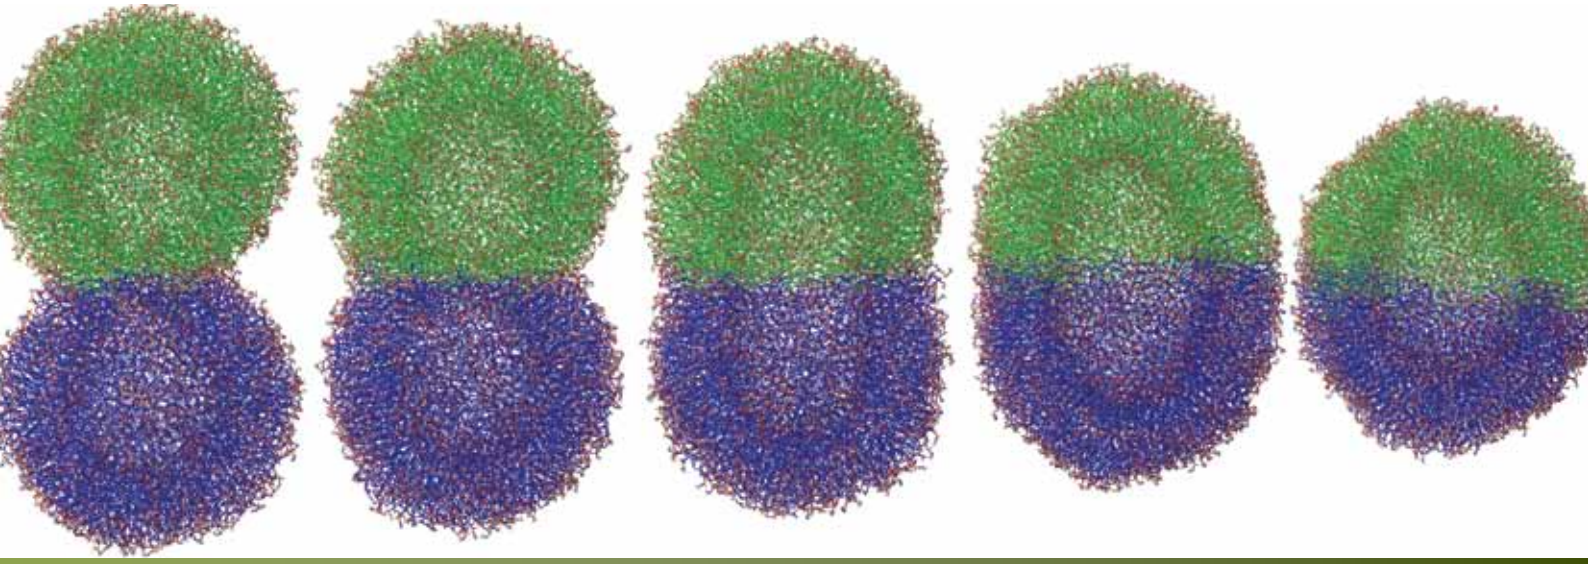

PDC Newsletter

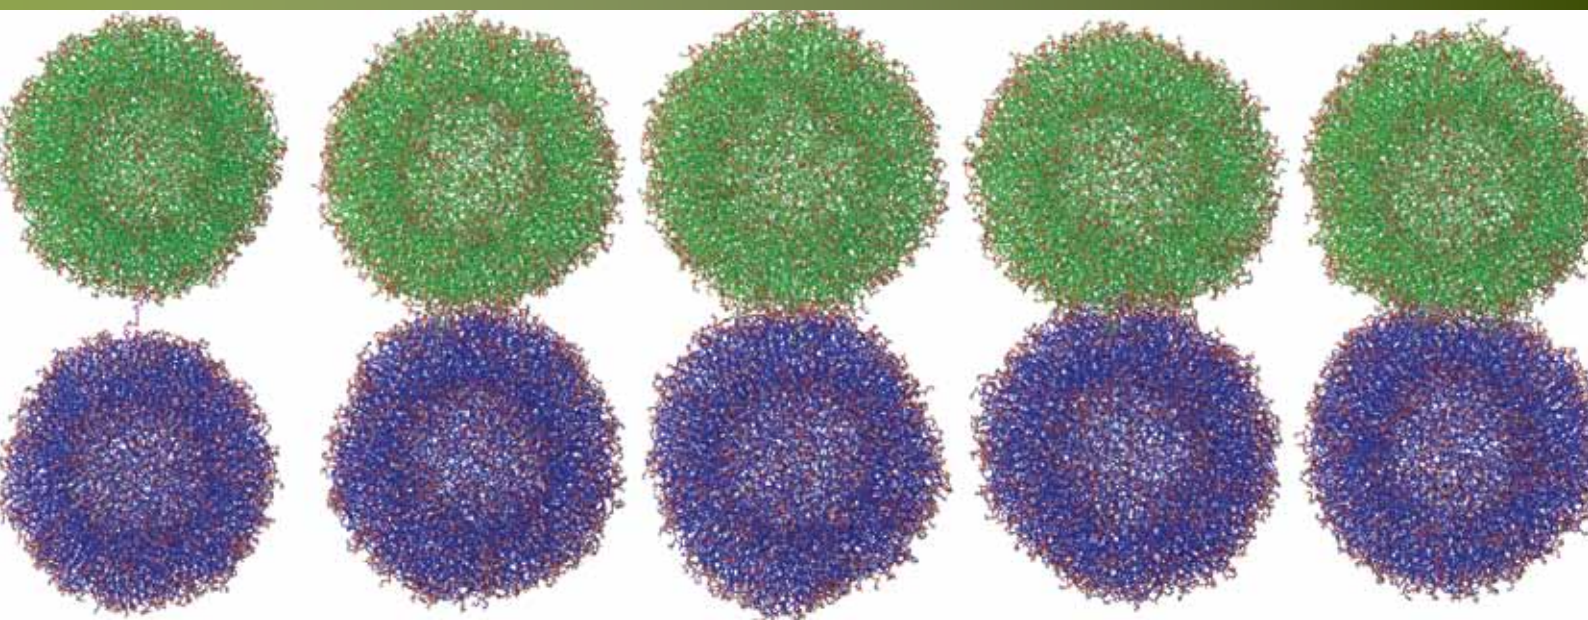

PDC Editorial

Lennart Johnsson,
Director PDC

Erwin Laure,
Director PDC-HPC

Published by PDC at KTH CSC

PDC operates leading-edge, high-performance computers as easily accessible national resources. These resources are primarily available for Swedish academic research and education. PDC receives funding from the Swedish Research Council (VR), and KTH.

Editors:
Lennart Johnsson and Erwin Laure

Editorial staff: Mike Hammill

Layout & Graphic Design:
Maria Malmqvist

E-mail: pdc-newsletter@pdc.kth.se

ISSN 1401-9671

Energy efficiency has been a driving force for the IT industry for several years. With the IT industry's CO_2 emissions already surpassing that of the aviation industry and the cost for electricity and cooling over the lifetime of a computing system being significantly higher than the original investment cost, it is clear that future growth of computing systems require significant changes in energy consumption. PDC has a long history in environmentally friendly computing (sometimes referred to as "green computing"), for example, we re-use the heat produced by our computing systems for district heating and we have initiated a plan to even better exploit our heat directly within KTH. But heat re-use is only part of the solution--we also have to work on more energy-efficient computing system that are capable of delivering increased compute power at reduced energy consumption. The Partnership for Advanced Computing in Europe (PRACE) project is engaging in a number of prototyping activities for energy efficient computing and PDC has recently won a competition for building such a prototype.

In this issue of the PDC newsletter we discuss this prototype system and also continue our series of articles on scientific endeavors supported by PDC with an article titled "Combined Parallel & Distributed Molecular Simulations Unravel Vesicle Fusion" by Peter Kasson & Erik Lindahl.

Commercial "pay-per-use" services, often referred to as "cloud computing" are rapidly gaining interest as viable platform for computing. PDC has engaged in a number of related studies and prototypes and is now leading a Northern European effort to understand what implications cloud technologies may have on electronic infrastructure (e-Infrastructure) for science in Northern Europe.

At the end of October PDC decommissioned its two long-serving systems Lenngren and Lucidor. To replace these systems SNIC has awarded PDC with funds for a new capability system whose procurement has started and which should become available to SNIC and PDC users in July 2010.

Other good news is that Swedish eScience will be boosted by two major eScience projects that have been granted as part of the Swedish strategic research program. One of these projects, the "Swedish eScience Research Center - SeRC" is led by KTH with partners from Karolinska Institute, Linköping University, and Stockholm University. More information on SeRC and the new PDC HPC system will be featured in future issues of this newsletter.

Lennart Johnsson, Director PDC Erwin Laure, Director PDC-HPC

Energy Efficient Computing: The SNIC PRACE Prototype

By Lennart Johnsson, PDC

Energy efficiency in computation has become the number one concern for infrastructure providers both for environmental and cost reasons. The CO_2 emissions of the IT industry have surpassed that of the aviation industry in recent years and have a growth rate about twice that of the aviation industry. The IT industry is estimated to contribute 2 – 3% of the global CO_2 emissions today, and the fraction is growing. The cost of computing infrastructure in the form of power and cooling systems and their operation over the lifetime of computing equipment has grown rapidly and in many countries now exceed the cost of the computing equipment. A decade ago these infrastructure costs typically amounted to about 15% of the cost of computing systems. For PDC, we estimate the infrastructure cost for the lifetime of the most recently procured systems to be about 1.5 times the cost of the hardware. The reason for this very significant change is rapidly increasing energy costs and, despite about a million-fold improvement over about 30 years in integrated circuit energy efficiency, an increase in power consumption of High Performance Computing equipment under fairly constant economic budgets.

To address these problems the computing industry started close to a decade ago to focus on the energy consumption at several levels; components, systems and complete data centers. In recent years, user industries, in particular the Internet companies (e.g. Amazon, Google, Microsoft, Yahoo), but also Universities (e.g. UC Berkeley, MIT and Stanford) have also started several initiatives addressing energy efficiency of computing systems.

One solution to exploit increased capability without increased power consumption is multi-core CPUs, i.e., processor chips with more than one complete processor. Technology demonstration systems based on dual-core AMD CPUs and PowerPC CPUs appeared in 2004. Today, six-core CPUs are available from both AMD and Intel and CPUs with up to 12 cores are scheduled to appear around the end of the year. Specialized CPUs, such as Graphics Processing Units (GPUs), today typically have hundreds of cores with, as examples, the nVidia Tesla C1060 CPU having 240 stream processor cores and the AMD FireStream 9270 having 800 stream processor cores.

A good example of the industry's efforts at energy efficiency at the computer systems level is IBM's Blue Gene/L (BG/L) system introduced in 2004. It is based on the dual-core PowerPC CPU. The BG/L system not only set a record in terms of performance, but also in terms of energy efficiency. (PDC is operating a BG/L system for the Stockholm Brain Institute (SBI)). The current version of the Blue Gene series of systems, the Blue Gene/P, holds five of the top ten positions in terms of energy efficiency among systems on the Top500 list, a list of the 500 most powerful computer systems. The other five of the ten most energy efficient systems use some form of accelerator technology, at the expense of programmability.

In This Issue:

On the Cover:
"Evolution of 15-nm vesicle fusion over 100 ns. By using a combination of parallel and distributed molecular simulations we have been able to complete over 10 microseconds of simulations of this system."

Another design targeting energy efficiency was SiCortex's MIPs based system (the company closed this spring), that used relatively slow cores, like the BG/P (850MHz at ~7W/core), the Lawrence Berkley National Laboratories Green Flash proposed architecture based on Tensilica CPUs (650 MHz at ~0.7W/core) and SGI's concept computer, the Molecule, based on Intel's ATOM processor (~1.8GHz at ~1W/core).

For the Partnership for Advanced Computing in Europe (PRACE) that is in the initial steps towards establishing an infrastructure of Leadership Class computing systems in Europe, energy efficiency is of critical importance for both environmental and cost reasons. Without a significant change in technology and possibly architecture, leadership class systems towards the end of the next decade have been estimated to consume 70-130 MW, something that may very well prevent such systems from being built or deployed in Europe. PRACE has therefore undertaken a number of technology prototype efforts that seeks to gain better understanding of critical issues such as programmability, scalability of systems and applications, and energy efficiency of different architectural approaches. The impact of architectural changes on programmability and application codes is the number one concern in offering useful resources and user support for scientists.

PDC with the support of the Committee for Research Infrastructures (KFI) of the Swedish Research Council, SNIC, KTH, the Center for Biomembrane Research (CBR) at Stockholm University, and the PRACE project is building a prototype system that seeks to assess energy efficiencies achievable with standard platform technologies with as small an impact as possible on programming models and application codes. This prototype is built out of standard components for cost reasons and without acceleration for programmability and application porting reasons.

The SNIC PRACE prototype is developed jointly with Supermicro, AMD and Southpole. This choice of collaborating vendors was made for several reasons, but in particular for the vendor's willingness to work with the PRACE project to advance the energy efficiency of HPC server designs, for the energy efficiency and power management features of AMD CPUs and Supermicro platforms in existence at the time of procurement, and expected component availability at the time of delivery. The SNIC PRACE prototype is expected to have an energy efficiency as measured for the Green500 list close to that of the BG/P, but with more memory/core. For some applications, such as the Gromacs molecular dynamics code, it is expected to have considerably better energy efficiency than the BG/P.

The prototype consists of 180 4-socket blade servers designed for PRACE by Supermicro with 10 blades housed in a new 7U chassis. With more than 5.7 CPU sockets per U the prototype is significantly denser than the currently popular 1U twin-server and 2U quad-server solutions with a density of 4 CPU sockets/U. (U is commonly used as the unit height measure for rack based computer systems: 1U=4.445cm.) The SNIC PRACE blade design can house CPUs with a max power consumption of about 100W each. Blade based systems typically have an energy efficiency advantage over 1U or 2U server solutions in that chassis housing blades can use larger fans which are typically more efficient. The prototype will also have new power supplies with slightly enhanced efficiency. It is estimated that only close to 15% of the energy will be consumed by power supply losses and chassis fans at full load, in about equal portions.

Each prototype blade is equipped with AMD six-core Istanbul 2.1 GHz High Efficiency (HE) CPUs. This CPU choice was predicted to offer the best energy efficiency of the CPU choices expected to be available at the time of system delivery. Each CPU can accept four memory modules, with the prototype having two 4 GB memory modules for each CPU (32 GB per 4-socket node). The new blades make use of AMD's new chipset that support the Hyper-Transport-3 (HT3) technology and the PCI Express Gen-2 I/O bus. Each blade chassis has a peak performance of 2TF and each standard 19" rack a peak performance of 12 TF. The peak performance for a rack using 2.6 GHz CPUs, the current highest frequency of AMD's latest generation CPUs with a max power consumption of about 100W would be 15TF, but the energy efficiency of those CPUs is expected to be lower than the selected ones. Each chassis is equipped with a 36-port Quad Data Rate (QDR) Infiniband (IB) switch. In total, the prototype has

180 nodes that are interconnected as a full bi-section width fat tree with two levels of 36-port IB switches. The theoretical peak performance of the SNIC Prototype cluster is 36.3 TF.

In assessing the efficiency of the prototype we plan to make use of the various control features available at the CPU and platform level that affect energy consumption and performance. The CoolCore technology turns off unused circuits and entire cores, The Smart Fetch technology copies data held in local cache to a higher level cache so that an idle core does not have to be woken up to access data that otherwise would only have been available in a local cache. CoolCore and SmartFetch features are under hardware control, and not accessible to the user. The CPU also allows for separate control of the voltage setting for the memory controller and jointly for the cores through the Dual Dynamic Power Management feature. The frequencies for the cores can be controlled individually through the PowerNow feature by choosing among a set of predefined sequences, from 2.1 GHz for the selected CPU down to 0.8 GHz. It also allows for control of the voltage and frequency settings of the HT3 channels. The settings have a small impact on energy consumption, but possibly a more significant impact on performance depending on the application.

AMD's Advanced Platform Management Link (APML) that is supported on the Supermicro blades design will allow for detailed observation of the state and power consumption of many components of the platform. Control of many features is possible through IPMI (Intelligent Platform Management Interface). In addition to the CPUs we will also seek to control fans and load distribution among the power supplies through these features.

In order to assess the characteristics of the various PRACE technology prototypes it has been decided to use a set of "simple" standard benchmarks, and some applications: High Performance Linpack (HPL) for power consumption and some application workloads, Stream for memory bandwidth, GUPS for memory latency, and functions from EuroBen for application kernels. For internode communication we will make use of the Intel MPI Benchmark. As one "true" application code we have chosen Gromacs not only because it is a very important code for the SNIC PRACE prototype partner CBR with Erik Lindahl's group being one of the lead Gromacs development teams, but also because it is a widely used molecular dynamics code. Gromacs is one of the application codes used for architecture assessment by the PRACE software effort. The Lindahl group has ported and optimized Gromacs for a wide range of platforms, including platforms making use of GPGPUs enabling comparisons of the prototype with several other modern platforms.

We hope to make the SNIC PRACE prototype available to SNIC users and the KTH and PRACE communities in December. We welcome any and all scientists interesting in studying approaches to enhanced computational energy efficiency and scalability of their codes to use the SNIC PRACE prototype and share their findings with the SNIC and PRACE user communities and partners. The prototype is intended for research in energy efficient computing and scalability, not for production, at least in the near future.

How to Contact PDC

Visiting address:
Teknikringen 14, plan 5, KTH, Stockholm

Mail: PDC, KTH,
SE-10044, Stockholm, Sweden

E-mail: support@pdc.kth.se

www: http://www.pdc.kth.se

Phone: +468790780o

Fax: +468247784

Combined Parallel & Distributed Molecular Simulations Unravel Vesicle Fusion

Peter Kasson & Erik Lindahl, Center for Biomembrane Research,
Department of Biochemistry & Biophysics, Stockholm University

Figure 1

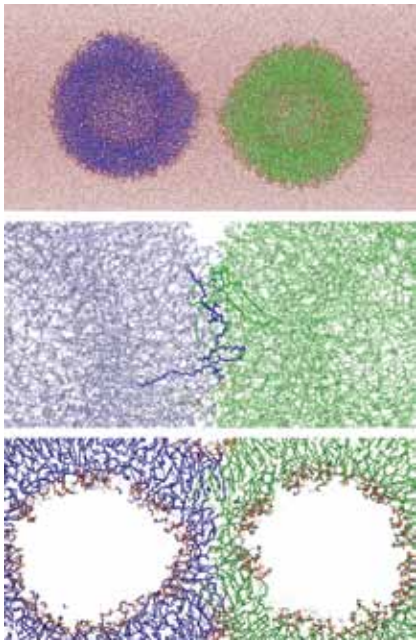

Figure 1) Solvated POPC/POPE vesicles in a hexagonal box; including the solvent water, the system comprises 1.1 million atoms. (Middle) The transition state with 50/50 probability of stalk formation occurs when the first lipid chains interaction over the water layer. (Bottom) Successful stalk formation rapidly leads to a metastable hemifused state.

Fusion of membranes is an extremely important process in cells. Not only is it critical for synaptic transmission in our nerve cells, but viruses such as influenza and HIV infect a host by fusing with the cellular membrane. Exactly how this happens on the molecular level is not fully understood, but we do know that viruses need special "fusion peptide" proteins to catalyze this process - simply bringing the two membranes together is not sufficient [1,2]. Since membrane and membrane protein systems are hard to study experimentally on the atomic level, molecular simulation techniques have evolved as a highly useful complement to laboratory experiments. Simulations have not only been able to explain important experimental results, but have also been used to make predictions e.g., about drug binding and how voltage-gated ion channels function in our cells when potentials are applied or removed [3,4].

Still, systems of this size are very challenging even with modern computers; even a medium-size protein in a membrane reaches >100,000 atoms, and to describe a pair of vesicles in water with atomic detail requires more than 1,000,000 atoms. In addition, since the point of these studies is to understand dynamics rather than structure, we need to evaluate energies and forces repeatedly - a microsecond simulation needs 500,000,000 steps. For this reason we typically employ classical mechanics with semi-empirical force fields parameterized to reproduce experimental properties, rather than extrapolating quantum chemistry results. Even with classical mechanics these are far from trivial problems; fusion is expected to take hundreds of nanoseconds (months of computer time), not to mention that we also need good statistics.

During the last few years, we have all benefited from a significant increase in available high performance computing resources both at PDC, NSC, and other sites in Sweden, and this trend is set to continue. 5,000-core clusters are standard today, and sooner than we think this will have increased at least tenfold. This is great news, but it leads to entirely new challenges about how to parallelize code and utilize these resources efficiently. About a year ago, we completed a very efficient domain decomposition parallelization in our code Gromacs (largely due to heroic efforts by our colleague Berk Hess) [5], which has enabled us to run simulations of these million-atom systems up to roughly 1,000 cores. In fact, after removing all global communication we have been able to use 150,000 cores for systems that are large enough; the scaling limit is rather that it is hard for us to efficiently go below ~250 atoms per core - at that point we spent more time communicating than calculating.

While this new generation of software has transformed the way we do simulations, we are now faced with a shortage of capability systems, and more serious: we have realized that no matter what we do, simulations of normal systems will not scale to future supercomputers with up to a billion cores that might be reality in ten years. We will almost certainly improve scaling to 100 atoms per core, possibly 10 if lucky, but definitely not under 1. Instead, we have started to employ techniques based on ensemble simulations [6], that turn out to be an amazingly efficient way of combining capability and capacity systems at PDC and NSC, and that

Figure 2

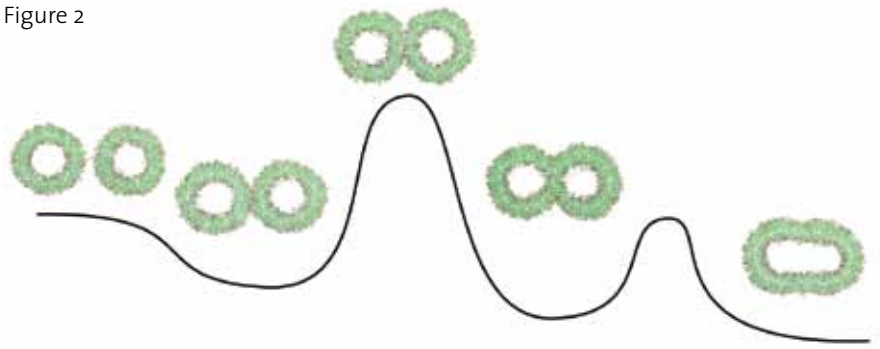

Figure 2) Schematic free energy variation for the fusion process. Starting from a conformation with a chemical linker, the vesicles first form an metastable contact patch, and eventually pass a transition state where the first side chains from different vesicles interact. This leads to formation of a stalk, and subsequently fusion of the outer leaflets into a metastable hemifused state. Finally the inner leaflets merge and the contents mix. This process is driven by perturbations similar to those caused by fusion peptides, which will help us understand the role of those proteins e.g. in influenza infection.

might be of interest to a wider range of users.

The approach we have used for our present studies is to first simulate fusion of real vesicles connected by a small chemical linker, as shown in figure 1. This system mimics the smallest vesicles found in cells (roughly 900 lipids each) and contains just under 1.1 million atoms when solvated in water. This is quite demanding to simulate with atomic detail, but critical to accurately model hydrogen bonds, dewetting and similar effects.

Since the composition of membranes is critical for fusion, we simulated a total of nine systems with different lipid mixtures. As a result, seven pairs of vesicles with 75-100% POPE lipids (the remaining lipids are POPC, with the difference that PE headgroups are much more polar) fused in a few 100 ns, but systems with 50/50 remained unfused up to half a microsecond. This part of the project was carried out on a capability (Infiniband) system at NSC, and used almost 2 million core hours in total. These results are fascinating per se, but they suffer from the same shortcoming as so many other simulations; it is merely a handful of semi-random observations in a statistical mechanics system - what can we predict quantitatively from it?

To address this, we have use ensembles of simulations and committer analysis [7], with the initial successful fusion pathways as seeds. The technique is quite straightforward: we simply extract the conformations every 5 ns from a 100 ns fusion trajectory, and from each such point we start 20 copies of 20 ns simulations with different random seeds. This means 400 simulations in total (close to 10 microseconds), which required an order of magnitude more computational resources than the initial project. Normally we use the Folding@Home distributed computing project for this type of work, but the vesicles presented a new problem - most normal home computers simply cannot simulate a million-atom system since it would take years on a single CPU and require gigabytes of memory. Instead, we chose to use limited parallelization (8-16 cores per system), and use nodes on capacity clusters, primarily Ferlin at PDC (but also some high-end Folding@Home users!), with our Folding@Home servers at PDC as meta-job-managers. This gave us an order of magnitude more data in just a few months - in total the ensemble part of the project comprised roughly 10 million core hours.

The idea behind the committer analysis technique is to collect statistics about the likelihood of a conformation leading to fusion/folding (or whatever the target state of the trajectory is), when we restart the trajectory from each point. For many applications this provides a very efficient solution to the "statistical" prefix problem of statistical mechanics, and it is obviously not limited to biological macromolecules. The result of the analysis is effectively a free energy profile of the process (figure 2); for the vesicle fusion this means we can show that there is a first metastable state when the head groups of the vesicles first make an extended contact patch, and we can identify and characterize the transition barrier state for the entire process - the point with 50% success probability in the committer analysis. Interestingly, this turns out to be quite specific, and happens when one lipid tail from each vesicle make contact in the thin hydrophilic layer between the vesicles. With equal probability, the tails will either separate again, or other tails will join them and grow into a larger hydrophobic 'stalk' and then fusion of the outer leaflets of the two vesicles (which is

Staff Focus

Photo: Harald Barth

Ali Gholami

Ali Gholami received a BSc in Computer Science at the Polytechnic of Tehran. After that, he travelled to Stockholm and began studying at the Royal Institute of Technology where he received a Master's degree in Software Engineering in Distributed Systems.

Ali worked as researcher and software developer at the Open Middleware Infrastructure Institute for Europe (OMII-Europe). This was at PDC. Later, he joined Telenor and Ericsson as a consultant. Just now, Ali has resumed working at PDC, this time collaborating on the Nordic Data Grid Facility (NDGF), specifically with the Data Management issues of the Advanced Resource Connector (ARC) middleware.

Photo: Harald Barth

the second metastable state before full fusion). We have then examined the fluctuations of lipids to better understand which events lead to this transition state, and been able to show that these effects are quite similar to membrane perturbations caused by the fusion peptides active during influenza viral infection. These are still only initial results, and we have great hopes to be able to scale it up on future clusters to include larger systems including real peptide-mediated fusion as well as systematic studies of effects from mutations.

We are quite enthusiastic about the prospects of using techniques like these that combine different levels of parallelization. For the vesicle case, we had to start from normal parallel simulations since we had little information about the intermediate states, but when a pathway is known the seed conformations can often be constructed either directly or by faster non-equilibrium simulations. However, we have found committer analysis to be at least equally useful in combination with large-scale parallel simulations. They can easily augment the amount of collected data 1-2 orders of magnitude beyond the largest allocations on capacity systems. They provide a way to use hundreds of thousands of nodes in parallel without having full intersectional bandwidth in the entire system. The concept can be extended to entire networks of transitions where it is possible to employ Markov models to map out the entire free energy landscape and even decide dynamically which simulations to run next.

In summary, we think this approach represents a great marriage of capability and capacity systems for future supercomputing; objectively the capability systems will always be 'better', but since capacity systems such as Ferlin are significantly more cost-efficient it is quite attractive to be able to use them for a wider set of problems!

1. Bissonnette ML, Donald JE, DeGrado WF, Jardetzky TS, & Lamb RA, J Mol Biol 386(1):14-36 (2009)
2. Chernomordik LV & Kozlov MM (2008) Mechanics of membrane fusion. (Translated from eng) Nat Struct Mol Biol 15(7):675-683 (in eng)
3. Lindahl, E., Sansom, M. Membrane Proteins: Molecular Dynamics Simulations, Curr. Opin. Struct. Biol. 18(4), 425-431 (2008)
4. Bjelkmar, P., Niemelä, P., Vattulainen, I., Lindahl, E., Conformational Changes and Slow Dynamics through Microsecond Polarized Atomistic Molecular Simulation of an Integral Kv1.2 Ion Channel, PLoS Comp. Biol 5 (2): e1000289 (2009)
5. Hess B, Kutzner C, van der Spoel D, & Lindahl E (2008) GROMACS 4: Algorithms for highly efficient, load-balanced, and scalable molecular simulation. Journal of Chemical Theory and Computation 4:435-447.
6. Kasson PM, et al. (2006) Ensemble molecular dynamics yields sub-millisecond kinetics and intermediates of membrane fusion. . Proc Natl Acad Sci U S A 103(32):11916-11921.
7. Du R, Pande VS, Grosberg AY, Tanaka T, & Shakhnovich ES (1998) On the transition coordinate for protein folding. (Translated from English) Journal of Chemical Physics 108(1):334-350 (in English).

Cloud Computing at PDC: A Northern Europe Matter!

By Åke Edlund, PDC

NEON - the Northern Europe Cloud Initiative <http://www.necloud.org/> started this fall.

The overall objectives of the NEON project are to build competence in cloud computing technologies for scientific applications, involving both computation and data, and to understand what implications cloud technologies may have on electronic infrastructure (e-Infrastructure) for everyday science in Northern Europe.

The collaboration will include the study and use of open-source software (for private clouds) and commercial products. It will build competence in the end-user interfaces provided by cloud technologies, and it will also identify their added value to today's grid and standard HPC interfaces. Additionally, the collaboration will try to provide understanding for the economic implications and cost-efficiency of e-Infrastructure, by possibly adopting interfaces to external (commercial) providers. NEON participants will collaborate and share experiences with a number of other European cloud initiatives, including those in the Netherlands, the U.K., Spain, and Greece.

The participating countries are Sweden, Norway, Denmark, Finland, Iceland, Estonia, Latvia, Lithuania, Belarus, and Poland. The project is led by Åke Edlund, PDC, and coordinated by a steering group that initially contains representatives from Sweden, Norway, Denmark, and the Nordic Datagrid Facility (NDGF). The project duration is one year.

Pictures from PDC's Machine Room, page 8 PDC's Ekman page 9: PDC Tapes and Storage

Staff Focus

Photo: Harald Barth

Jonathan Vincent

Jonathan Vincent obtained his Ph.D. in semiconductor physics from the University of Exeter in the UK. He has since worked on a wide variety of compuational science problems including quantum and classical molecular dynamics and weather models. This work was carried out at several different locations around Europe, including Reading in the UK, Leiden in the Netherlands, and Gothenburg and Uppsala in Sweden. At PDC he will be working as an application expert reinforcing the PDC support team. He will particularly focus on helping users to efficiently use PDC resources and also work directly with selected applications.

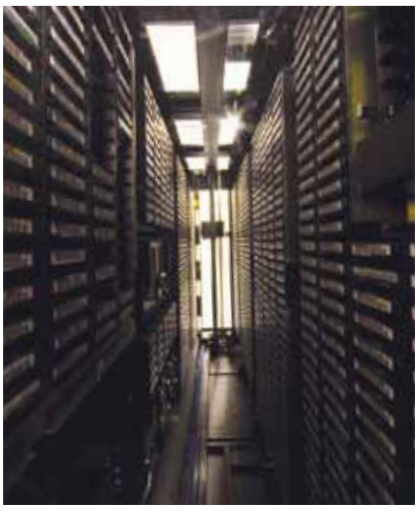

Photo: Harald Barth

Staff Focus

PDC's Financal Administration Is Handled by Two Staff Members: Anette Arling and Elisabeth Hegrad

Anette Arling has worked at the School of Computer Science and Communication (CSC, <http://www.kth.se/csc/>) at KTH for several years. Before she came to PDC, she worked at the CSC's Human-Computer Interaction (MDI, <http://hci.csc.kth.se/>) department. At the beginning of 2009, she started working at PDC, where her duties include budgeting and financial planning. In addition, she is responsible for doing the accounting for all European Union projects in-which PDC is involved. Anette has taken a variety of different courses at Stockholm University <http://www.su.se/>. Before she was hired at KTH, she worked at Stockholm School of Economics (Händelshögskolan). Anette is married with two children, a daughter and a son. She rides horses in her free time.

Read about Elisabeth at page 11...

PDC's financal administration Anette Arling and Elisabeth Hegrad

More Students Dive into HPC at PDC

By Mike Hammill, PDC

More students than ever took the "Introduction to High-Performance Computing" course offered at PDC this summer. With 78 students, participation was up 20% from the previous year and up 42% from the average. In particular, the number of Ph.D. students showed a marked increase, up 53% from the year before and up 58% from average.

Being at the leading edge of high-performance computing is what the class is all about. The students get both a strong conceptual foundation in state-of-the-art HPC as well as experience with practical aspects through lab work and projects.

To help the students make the most of the two-week course, a seasoned cadre of professors and professionals from around Sweden are joined by world-famous practitioners to cover various topics. These include parallel programming (OpenMP, MPI), modern computer architectures, parallel algorithms, efficient programming, and case studies. Lecturers include such well-known leaders in the field as Björn Engquist (KTH/Univ. of Texas at Austin), Lennart Johnsson (KTH/Univ. of Houston), Scott Baden (Univ. of Cal., San Diego), Thomas Ericsson (Chalmers), Erik Hagersten and Sverker Holmgren (Uppsala Univ.), Martin Ingvar (Karolinka Inst.), Mats Nilsson (Acumem AB), and Nils Smeds (IBM Sweden).

Each student is expected to complete a project, which is often related to his or her research area. This year's projects include "Parallel Short DNA Sequence Mapping," "Parallel Attack on MD5," and "Parallel Computation of Fractional Quantum Hall Effect."

To help the students complete this work, a tutor is assigned. In addition to helping answer students' questions, the tutor reviews the work in conjunction with the course examiner, Jesper Oppelstrup, a numerical analysis professor at KTH.

For PDC, the course is a win-win situation. Not only do the students gain valuable knowledge that will help them with their research, but PDC also benefits by encouraging more efficient use of its resources by the next generation of its users.

Photo: Maria Malmqvist

Photo: Harald Barth

What do the students think? They gave the overall course a 3.6 score out of 4.0 with 4 being "very good" and 1 being "poor." Here are some quotes from this year's course evaluation.

"In general, very good lecturers and relevant topics. Good mix between lectures and lab work. Very nice summing up dinner. The lab assistants were great!"

"Overall the course was excellent and enlightened the students with modern high-performance computing practices. Both the software and hardware aspects were covered, and personally I was satisfied with the content and quality of the course."

"Overall an awesome course, I learned one hell of a lot."

The course, which is given the last two weeks in August, is within the National Graduate School in Scientific Computing (NGSSC), and the KTH Computational Science and Engineering Centre (KCSE), both of which receive funding by the Swedish Research Council (Vetenskapsrådet). It is also open to KTH masters students as well as Ph.D. students worldwide.

This is the 14th year the course has been given since the first summer school in 1996.

For more information, see the course Web page at <http://www.pdc.kth.se/education/summer-school/>.

Staff Focus

Elisabeth Hegrad has worked at CSC since 2001. Before she came to PDC, she worked at the KTH Network Operation Centre (KTHNOC). Elisabeth has been a financial administrator at PDC since spring 2006. Currently, she works with accounting, billing, and a wide variety of other administrative tasks at PDC. Looking back, Elisabeth has been employed at KTH since 1989, during which time she worked for several departments, including Central Administration and External Relations. Before she was employed at CSC, she was in charge of student housing at KTH, a coordinator position with broad responsibilities. Her education includes studying bookkeeping at Företagsekonomiska Institutet (FEI,<http://www.fei.se/>). She has a grown-up daughter. Elisabeth is originally from Hungary.

PDC Related Events (Sponsored/Associated)

PDC is holding a code-porting workshop on stream computing.

GPU Workshop

December 7-10, 2009, Stockholm, Sweden.

http://www.pdc.kth.se/events/event-repository/stream-computing-workshop/

This user forum is hosted by SNIC, UPPMAX and PDC.

EGEE User Forum in Uppsala

April 12-16, 2010, Uppsala, Sweden. Info: http://egee-uf5.eu-egee.org/

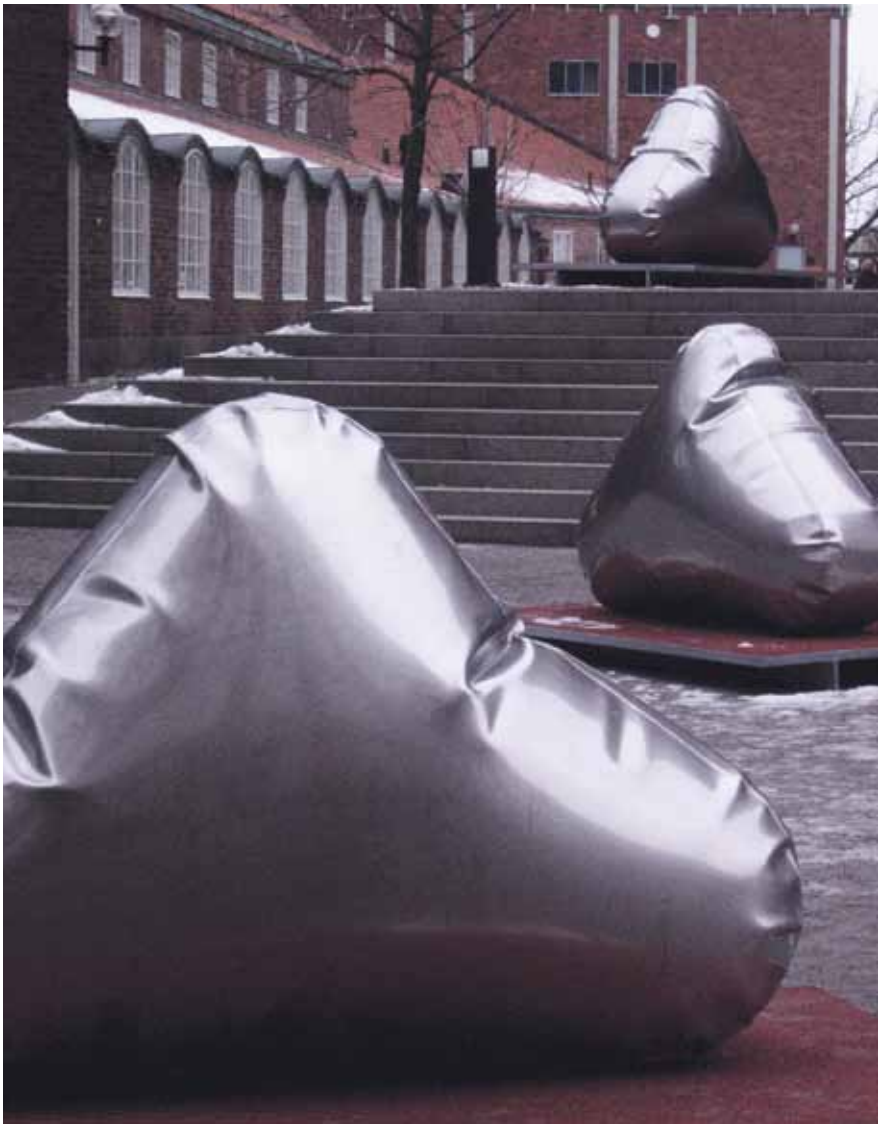

Photo: Maria Malmqvist

Sources

We can recommend the following sources for other interesting HPC opportunities and events:

PRACE

http://www.prace-project.eu/prototype-access

http://www.prace-project.eu/hpc-training-events

DEISA

http://www.deisa.eu/science/deci

http://www.deisa.eu/news_press/symposium/Amsterdam2009/deisa-symposium-amsterdam-may-11-13-2009

TeraGrid

http://www.teragrid.org/tg09/

US Department of Energy

http://hpc.science.doe.gov/

HPC University

http://www.hpcuniv.org/events/current/

HPCwire

http://www.hpcwire.com/events/

Netlib

http://www.netlib.org/confdb/

CERN Courier

http://cerncourier.com/cws/events

EGEE

http://events.eu-egee.org/

Linux Journal

http://www.linuxjournal.com/events