PDC Newsletter

1997: A Great Year of Scientific Computing

The preparation process of the *Parallelldatorcentrum Progress Report 1997* has now been concluded. Most PDC users will have already received their copy of the report by the time they read this newsletter. The PDC progress report is not mainly a document about the facilities of PDC and the activities at the center; rather, the focus has traditionally been on scientific achievements using the equipment and computational knowledge at PDC.

Reading the progress reports of previous years and comparing them to the current issue reveals a dramatic development. We have progressed from a few articles in the first report to this year's document, which contains 48 scientific articles written by more than 100 different authors.

This growth is only natural, of course;PDC has evolved from being a research project dedicated to parallel computing applied to scientific research, to being selected as the nation's main computational resource. The amount of research performed on the machines at PDC today is also reflected in the fact that the SP system was used by scientific batch jobs for 600.000 node hours.

The report shows the wide range of scientific disciplines that take advantage of supercomputing today. In accordance with worldwide scientific computations,there are many articles in the area of physics and chemistry, but these cover a wide range of applications: from atomic-scale quantum-mechanical computations to macroscopic phenomena, such as paper optics and geophysics.Scientific computing truly covers all possible time and length scales.

The quality of many of the illustrations in the PDC report has also improved over the years. Because the report's audience is mixed,many authors have chosen to illustrate the problems they are addressing – and in many cases with great success. The need of, and progress in, three-dimensional visualization is apparent.

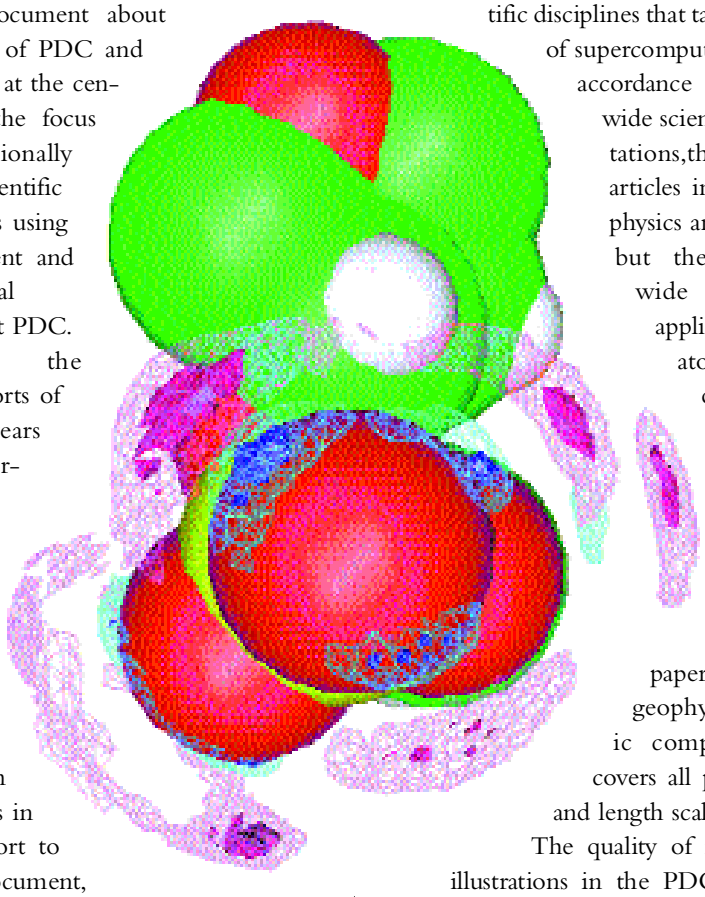

The spatial distribution functions of water atoms around a phosphate group of DNA. Atoms of phosphate group of DNA and neighbouring sugar and ester grours are displayed as spheres, the oxygen and hydrogen spatial distribution functions as correspondingly red and blue surfaces around. The surfaces are drown at intensities 2.4 (solid) and 1.5 (net)

In This Issue

From the editor

The PDC Progress Report for 1997 is now available, and I highly recommend it for reading or browsing. It gives an excellent picture of the diversity of applications of high-performance computing in the present Swedish academic community. The progress report contains close to fifty articles written by the scientists who were using the PDC facilities last year. There are many fascinating examples of scientific progress, resulting from both novel and traditional uses of high-performance computing.

It is obvious from these reports that the application of parallel computing has matured to become an integral part of many important Swedish research projects. This is also clear from the latest round of computer resource applications. Most of the applications were of high scientific level, and the requests for computer resources were well motivated. We apologize that many very good projects could not be fully supported.

The most important mission of a national center is to supply the services requested by the user community in the form of computing cycles, storage, help-desk assistance, courses, etc. It is also important, however, to introduce new technology – in both hardware and software, for the future. PDC played an important role in this process by introducing parallel computing in the early 90:s as a research tool for the academic and industrial communities.

PDC plans to continue its leadership role by exposing our users to the latest computing developments. This can be seen by the addition of the new SMP nodes to the IBM SP system. We hope these processors can introduce the next level of hierarchical parallel processing to a wider audience and at the same time be useful for the codes already running on Strindberg. We realize the importance of continuity. The investment in time by the user community in algorithm and code development must be protected. The visualization capability at PDC has also recently been enhanced by an upgrade of the Onyx2, Boye. This is the first essential step in a future improvement of the visualization equipment.

On the software and infrastructure side, we benefit greatly from our international collaboration. PDC is an associated member of NPACI, a partner of the GLOBUS project in the USA, and a node in the European TTN network (see articles in this and previous newsletters). In particular, we have learned a lot from the recent progress in software development for having distributed computers and centers working together.

We believe the time is ripe for a tighter cooperation between the high-performance centers in Sweden. One component is the technological development with a faster network and supporting software of the type discussed above. Another is the expressed interest from the centers. Recently a joint letter, signed by seven centers, was sent to the investigation comittee for research initiated by the Swedish government (Forskning 2000, http://www.hsv.se/F2000/). The letter discussed the conditions for Swedish high-performance computing in the future. Finally, and always of fundamental importance, this cooperation is what the funding agencies in Sweden expect today.

On September 24-25 a workshop at PDC will address both the technical and organizational aspects of the potential in the cooperation between the centers and their users. One interesting model is the NSF-sponsored partnerships in the USA, which focus the resources for high cost-efficiency but also involve several centers and research groups each. Some of the hardware and most of the research and development projects are distributed.

Björn Engquist
Director

The letter mentioned in the editorial is available at **http://www.pdc.kth.se/info/newsletter/**

Fujitsu VX/3

The Fujitsu VX machine, selma, has been upgraded to 3 CPUs. The machine now has a peak performance of 6.6 GFlop/s. The new CPU has 2GB of internal memory just like the older ones. Both PVM and MPI are available for parallel programming using either Fortran or C.

IBM SP SMP Nodes

As previously announced, six new symmetric multiprocessor (SMP) nodes have been introduced into the IBM SP system, strindberg. These nodes are equipped with four Power PC 604e CPUs each, and they are running at a frequency of 332MHz. The aggregated peak floating point performance of each SMP node is thus 2656 MFlop/s. The IBM SP system now has an aggregated performance of more than 100GFlop/s and an aggregated memory bandwidth of 441 GFlop/s.

The nodes have been acquired as part of a research agreement with IBM Sweden to investigate the applicability of SMP in general, and parallel SMP systems on some scientific codes in particular.

Users interested in investigating the applicability of SMP processing on their codes – and who are willing to take the time and effort to tune their codes against the SMP architecture of the M-nodes are invited to contact **sp2-staff@pdc.kth.se** for more information.

Clustered Computing Workshop at PDC

PDC is planning to hold a workshop around the subject of cluster computing September 24–25, 1998. This is one of the fastest growing areas of high-performance computing and has vast potential for delivering cost-effective and yet very powerful computing resources. There are many different subjects that could be covered, and we have made the following preliminary selection:

Clustered Computing Workshop

- Meta-Computing Overview: possibilities and obstacles
- Hardware: computer architectures and interconnection networks
- Software: operating systems, security, and libraries
- Clusters and Throughput Computing: scheduling and CONDOR
- Centralized Services:long-term storage, backup, and distrusted file systems

The second day of the meeting will be dedicated to distributed computing. Topics addressed will include co-scheduling computer resources between computer sites, the GLOBUS project, and other aspects of resource sharing.

We would like to receive comments and ideas of other possible subjects and also invite presentations of the trials and tribulations of practical cluster projects. Please send suggestions to **ccw98@pdc.kth.se.**

Future Chip Technologies

The technical progress made during the fifty years since the invention of the transistor makes the semiconductor industry the envy of many other industries. We have become accustomed to continued rapid progress (see the article on Moore's Law in the sping issue of the PDC newsletter). The major force behind the improvements we have seen are related to the decreasing size of the transistor. Information processing is about the movement of bits of information, and a smaller bit usually also means faster movement.

When the SP User Group gathered at PDC in April, Dan Edelstein from IBM Research told the participants that the pursuit of smaller semiconductors continues to create many challenges. For example, the smaller size of features has required the development of near-ultraviolet lithography. As transistors get smaller, the role of interconnects also becomes increasingly important. A narrow conductor has to sustain higher current densities than a broad one. Because of this, the semiconductor industry has settled on using aluminum for interconnects within a chip, although it is well known that copper is a better conductor.

IBM has been doing research in the interconnect area for years, and during the fall of 1997 IBM announced that it had managed to solve the related problems. Dan Edelstein and his team have received recognition for this breakthrough, and the copper chip technology has made headlines in most newspapers. IBM can now build chips using copper instead of aluminum for interconnects, which will mean smaller chips and faster speeds. Because the new chip fabrication processes also require fewer manufacturing steps, it is clear that the IBM researchers and chip makers within IBM Micro-electronics are excited about this new development.

IBM Micro-electronics is the part within IBM that produces chips for IBM and increasingly for other vendors. Copper chips will be produced in 1998 and are expected to appear in new products within six or seven months. So don't be surprised if you see a sticker on your new computer that says "Cu inside."

Carl G. Tengwall, IBM

Parallel MD Simulations of DNA Systems

Alexander Lyubartsev and Aatto Laaksonen, Division of Physical Chemistry, Arrhenius Laboratory, Stockholm University

The full version of this article is available in the *PDC Progress Report 1997*
http://www.pdc.kth.se/info/reports/ Introduction

Computer simulation methods such as Molecular Dynamics (MD) and Monte Carlo (MC) have now become important techniques to study fluids and solids. These methods provide a link between theory and experiment, and they are also the only way to study complex many-body systems when both experimental techniques and analytical theories are unavailable.

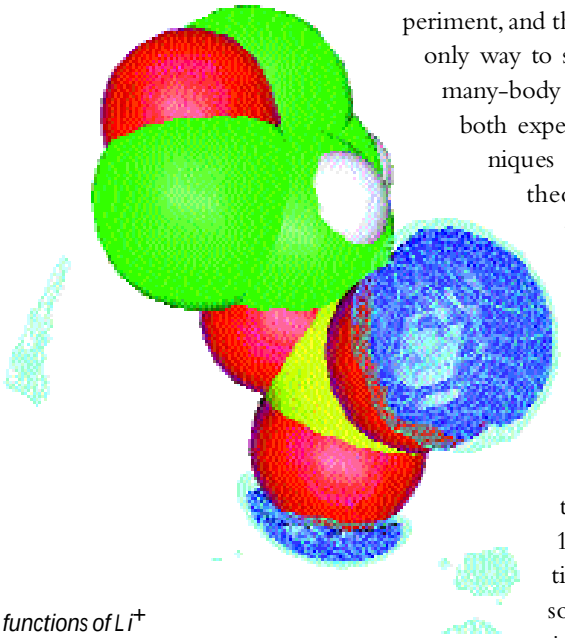

e spatial distribution functions of Li^+ n around a phosphate group of DNA lue surfaces). The surfaces are drawn at tensities 10 (solid) and 3 (net)

Now it is a standard routine procedure to simulate molecular systems consisting of the order of 100–1000 particles, which in some cases (*eg* simple liquids) is sufficient to give a good description of the corresponding macrosystem. For other systems, a larger number of particles is needed in order to describe them in a realistic way. Complex bio- and organic molecules (*eg* proteins, nucleic acids, membranes, carbohydrates, *etc*) immersed into a solvent increase the number of involved atoms with one to two or more orders of magnitude. Also, the larger the molecular systems grow, the longer simulations are needed to follow low-amplitude motions and slow conformational transitions. It is clear that the rate of the progress towards more complex molecular models is set, to a large extent, by advances in microprocessor technology and computer architecture as well as by development of appropriate software.

Parallel algorithms

Computer simulations of many-particle systems are well suited for parallel computer systems. The basic reason for this is that the forces acting on each particle can be calculated independently in different processors. However, the most optimal parallel scheme for a particular problem depends both on the hardware in hand and on the system under investigation (size, type of interaction, *etc*). Electrostatic interactions, fast intramolecular motions due to explicit modeling of hydrogens, angle and torsional angle forces of macromolecules – all these kinds of forces require a special treatment to create an effective parallel code.

There exist two main strategies in parallelization of a molecular dynamics code: replicated data and domain decomposition. The first scheme implies that all the nodes know the positions and velocities of all the atoms, and calculation of different forces goes in parallel. In the domain decomposition approach the atoms are distributed over the nodes, and each node calculates forces only for a certain set of atoms. Both schemes have pros and cons. We applied the replicated data method, which is more suitable for simulations of macromolecules surrounded by a solvent. The effectiveness of this method does not depend critically on the system structure and the force field. The weak point of this approach is that it requires more communications between the processors. Still, the prevailing point of view is that for several or several dozen processors, the replicated data method is preferable.

The MDynaMix program

During 1994–1996 we have developed a general purpose molecular dynamics code

Execution time (in minutes) for 1 ps simulation of 3 different systems on the IBM SP2,depending on number of processors. Box sizes and cutoff distances R_{cut} are given in Ångström.

System	#atoms	Box sizes	R_{cut}	Nodes 1	4	8	16	32
DNA	3825	33x33x34	13	120	32	17	9	5.5
Ion	5920	39x39x39	13	220	57	28	17	9.5
Lipid	7616	44x44x64	14	300	77	45	23	14.0

(MDynaMix) for simulations of arbitrary mixtures of rigid or flexible molecules employing the most modern simulation techniques: double time step algorithm for fast and slow modes, optimized Ewald method for electrostatic interactions, constant temperature-constant pressure algorithm. The program can be used for simulation of mixtures of molecules interacting with AMBER- or CHARMM-like force fields which includes the following terms:

1) Atom-Atom short-range interactions (Lennard-Jones potential),
2) Atom-Atom electrostatic interactions,
3) Intramolecular interactions: covalent bonds, covalent angles and torsional angles,
4) Some optional terms (hydrogen bonds, *etc*).

The code is highly universal and well suited for simulation of both simple molecules and complex biological macromolecules. It can easily be made to run on a single-processor computer. In the latest version (4.2) additional features were included: separate pressure control in different directions (for simulation of anisotropic systems), generalized reaction field method for electrostatic interactions, truncated octahedron or hexagonal simulation cell, parallel SHAKE algorithm for constrained dynamics, different types of torsion angle potential, and a few other options.

The replicated data method implies that each node knows the positions of all the atoms. This easily allows one to distribute calculation of different forces over the nodes, providing equal load to all the nodes. After completing the calculations on each node, all the forces acting on a given atom are transfered to the "home" node corresponding to this atom and summed up. This is done by an MPI_Reduce_scatter call from the MPI library. Then the atoms are moved according to a chosen integration scheme. After completing an MD step, new positions of atoms are broadcast to all nodes. The two global communication operations, MPI_Reduce _scatter and MPI_Broadcast, cause some computational overhead. Other parts of the program are almost 100% scalable.

Benchmarks

The performance of the program was tested on several molecular systems qualitatively differing by the composition:

1) A periodic fragment of DNA in ionic aqueous solution: DNA (635 atoms), 1050 waters, and 40 Na^+ and Cl^- ions, 3825 atoms in total.
2) An NaCl ion solution: 20 NaCl ion pairs and 1960 water molecules, 5920 atoms in total.
3) A lipid bilayer: 64 DPPC lipid molecules (50 atoms each) and 1472 water molecules, 7616 atoms in total.

All the simulations were carried out with flexible molecular models employing the double time step algorithm with 0.2 fs small time step and 2 fs long time step. The Ewald method was applied for calculations of electrostatic contributions. The table shows total execution time on the IBM SP2 for 1 ps simulations depending on the number of processors.

The color pictures of the spatial distribution functions were produced using the gOpenMol package by Leif Laaksonen,CSC, Finland

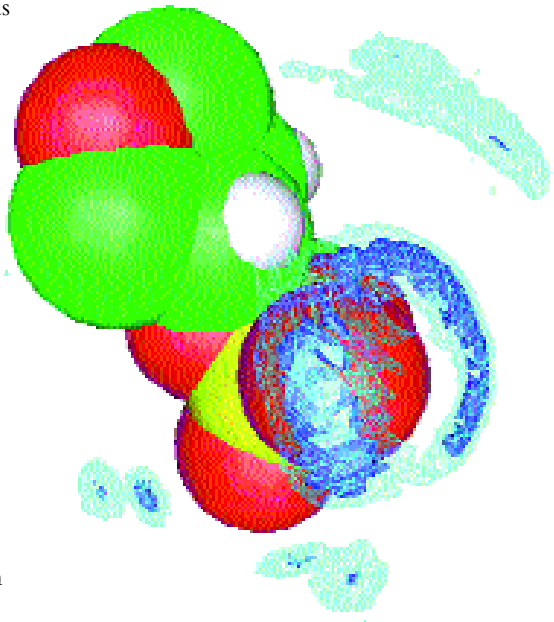

The spatial distribution functions of Na^+ ion around a phosphate group of DNA (blue surfaces). The surfaces are drawn at intensities 4 (solid) and 2 (net)

Calendar of Events

August

9–11 5th International Symposium on Solving Irregularly Structured Problems in Parallel, Berkeley, California
http://www.nersc.gov/Irregular98

17–28 NGSSC Summer School "Introduction to High-Performance Computing," PDC, Stockholm
http://www.pdc.kth.se/training/

19–21 20th World Conference on the Boundary Element Method BEM 20, Orlando, Florida
http://www.wessex.ac.uk/conferences/bem20/

19–23 4th International Conference on Numerical Methods and Applications, Sofia, Bulgaria
http://www.math.acad.bg/~nma98/

24–27 Numerical Methods and Computational Mechanics 98, Miskolc, Hungary
http://www.uni-miskolc.hu:8080/home/nmcm98/conf98.html

September

1–4 Euro-Par'98, Southampton, United Kingdom
http://www.europar98.ecs.soton.ac.uk/

2–5 1998 Conference on Computational Physics (CCP 1998), Granada, Spain
http://dalila.ugr.es/~ccp1998/

7–9 EuroPVM/MPI98, Liverpool, United Kingdom
http://www.csc.liv.ac.uk/~pvmmpi98/

14–16 SciTools'98: Modern Software Tools for Scientific Computing, Oslo, Norway
http://www.oslo.sintef.no/SciTools98/

22–25 3rd High-Performance Computing Asia Conference & Exhibition , Singapore, Malaysia
http://www.nsrc.nus.edu.sg/HPCAsia98

24-25 Clustered Computing Workshop & Distributed Computing Meeting
PDC, Stockholm
http://www.pdc.kth.se/

October

5–8 10th GAMM-Workshop on Multigrid Methods, Bonn, Germany
http://wwwwissrech.iam.uni-bonn.de/mg10

18–20 4th IMACS International Symposium on Iterative Methods in Scientific Computation, Austin, Texas
http://www.ticam.utexas.edu/dmy98

28–31 Workshop on High-Performance Scientific and Engineering Computing , Las Vegas, Nevada
http://www.cps.udayton.edu/~pan/pdcs98

November

7–13 Supercomputing '98, Orlando, Florida
http://www.supercomp.org/

15–20 Parallel Computing in Mechanical Engineering, Anaheim, California
http://www.asme.org/

December

16 PSCI/PDC Workshop on Computational Electromagnetics
KTH, Stockholm, Sweden

17–18 PDC Annual Conference: High-Performance Computing and Visualization,KTH, Stockholm, Sweden
http://www.pdc.kth.se/news/events

README

The standing section of short tips and tricks for users

- **Multiprocessor nodes on Strindberg.** The newly installed SMP nodes are available for batch jobs as well as interactive access through the EASY scheduler. The nodes are called M-nodes. The corresponding compiler commands are **xlf_r**, **xlf90_r**, and **mpxlf_r**. Compiler flags to use are **-qsmp** and **-qreport=smplist**.

- **NQS(qsub) batch files on Fujitsu** When transferring NQS batch files from other systems, option selection lines in the batch file need to be edited. #QSUB of many other NQS systems should be edited to read **#␣␣@$**<nqs option>. See chapter 5.5 of the NQS manual, available at **http://www.pdc.kth.se/doc/fujitsu/manuals/C/nqs/index.htm**

- **Batch job statistics on Fujitsu.** Jobs running on the Fujitsu may ask for job accounting statistics by issuing the command **eoj `jid`** as the last command in the batch file. The output of this command gives information about vectorization, memory usage, and disk I/O, for example.

- **CAC command.** To investigate the members of the charge allocation cxxx (CAC), use the command cac. **cac members**<username> lists the CACs that the user name is a member of. **cac members**<cacname> lists the members of that CAC. More information is available by the command **cac help**.

Travel Report : EHUG and SuperHIPPI

The European High-Performance Networking User Group (EHUG) held its spring meeting at the Swiss Center for Scientific Computing (CSCS) in Lugano, Switzerland, on March 9, 1998.

The aim of EHUG is to bring users of HIPPI and Serial-HIPPI together to exchange information and to discuss and solve common problems.

The meeting was attended by more than twenty people from all over Europe and was also transmitted over the Internet via MBONE. Several interesting presentations were held with a focus on the new Gigabyte System Network (GSN) standard. Slides are available under the EHUG home page.

GSN is a new full duplex network technology for switched networks with a transfer rate of 6.4 Gbit/s in each direction. It has also been known as HIPPI-6400 and as SuperHIPPI. This networking standard is more than sixty times faster than Fast Ethernet, which work stations and PCs use today. Furthermore it is already a standard, and public domain drivers are even being developed for LINUX.

The EHUG home page can be found at **http://www.cern.ch/HSI/hippi/hug/ehug.htm**

For an interesting popular article on GSN, see *PC Week*, January 23, 1998, **http://www.zdnet.com/zdnn/content/pcwk/1504/277384.html**

Nordic SP User Group Meeting

On April 20-21 the first meeting of the Nordic SP User Group (NSPUG) was held at Parallelldatorcentrum at KTH; it was organized and cosponsored by PDC and IBM Sweden. The meeting was intended to gather owners of IBM SP systems from the Nordic and Baltic countries to create a forum for exchanging information and concerns about running IBM SP systems in production. NSPUG encourages participation from academia, industry, and commerce; and it is open to any organization that owns an IBM SP system.

This first meeting of NSPUG lasted a day and a half. The purpose of the first day was to let people from different sites become acquainted. Consequently most of the time was spent doing site presentations, but there was also time for discussions. Site presentations were done by people from PDC, Tele Denmark, ICA, UNI*C, Phillips Petroleum Company Norway, and HPC2N. The seminars, "SP Trends and Directions" and "Copper Technology in Chip Manufacturing," were led by IBM staff members Clive Harris and Dan Edelstein, respectively.

On the second day the necessary decisions were made to form the NSPUG organizations. The tentative name Nordic SP User Group was accepted. The organization's mission can be briefly described: help ourselves in running IBM SP systems by sharing information. There will be two meetings per year, one in the fall and one in the spring. The NSPUG'98 fall meeting will be held in conjunction with the SP World meeting, September 8-10 in Dusseldorf.

VIDEOGRAPH, one of the projects sponsored by the HPCN TTN programme, is a pan-Nordic film post-production project. In the television and movie industries,there is a high demand for more sophisticated computer effects produced in much shorter time than is currently possible. As performance requirements of today's applications are increasing, the need for high-performance computing grows as well. For example, rendering is one of the most popular effects used for many purposes. A 10-second rendering effect, at a resolution of 4096x3072 pixels, can take up to 1050 hours – or 44,5 days – of computing time on either a high-end Macintosh or PC platform. Using a 16-node CYCORE reduces this production time to 80 hours, or 3,5 days.

The entertainment industry's software choice for creating visual effects is a package called Adobe After Effects, with corresponding plug-in effects called Final Effects and Studio Effects; this runs on both Windows and Macintosh platforms.

Three partners are involved in the project. Cycore Computers AB from Uppsala, Sweden, is the owner of the source code for Final Effects and Studio Effects; FilmEffekt AS in Norway is a post-production company for both film and video; and PELAB is a research group at the Department of Computer and Information Science at Linköping University, Sweden.

The project started in May of 1997 and is expected to finish in September of 1998.

More information can be found at **http://www.pdc.kth.se/pdcttn/**

B

Sverige
Porto
Betalt

PDC Newsletter

Published quarterly by Parallelldatorcentrum at KTH.

The center 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 Council for Planning and Coordination of research, FRN and KTH.

Editor: Björn Engquist

Editorial staff: Nils Jönsson, Faith Short

Email: pdc-quarterly@pdc.kth.se

ISSN 1401-9671

How to Contact PDC

Visiting address: Lindstedtsvägen 3, Plan 5, KTH, Stockholm

Mail: PDC,KTH,S-100 44 STOCKHOLM

Email: pdc-staff@pdc.kth.se

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

Phone: +46 8 790 79 07

Fax: +46 8 24 77 84

PDCTTN

Contact: Marina Backer Skaar

Email: pdcttn@pdc.kth.se

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

Phone: +46 8 790 69 23