The newsletter of the Centro Svizzero di Calcolo Scientifico / *Swiss Center for Scientific Computing*

March 1997
volume 6
number 1

# CROSSCUTS

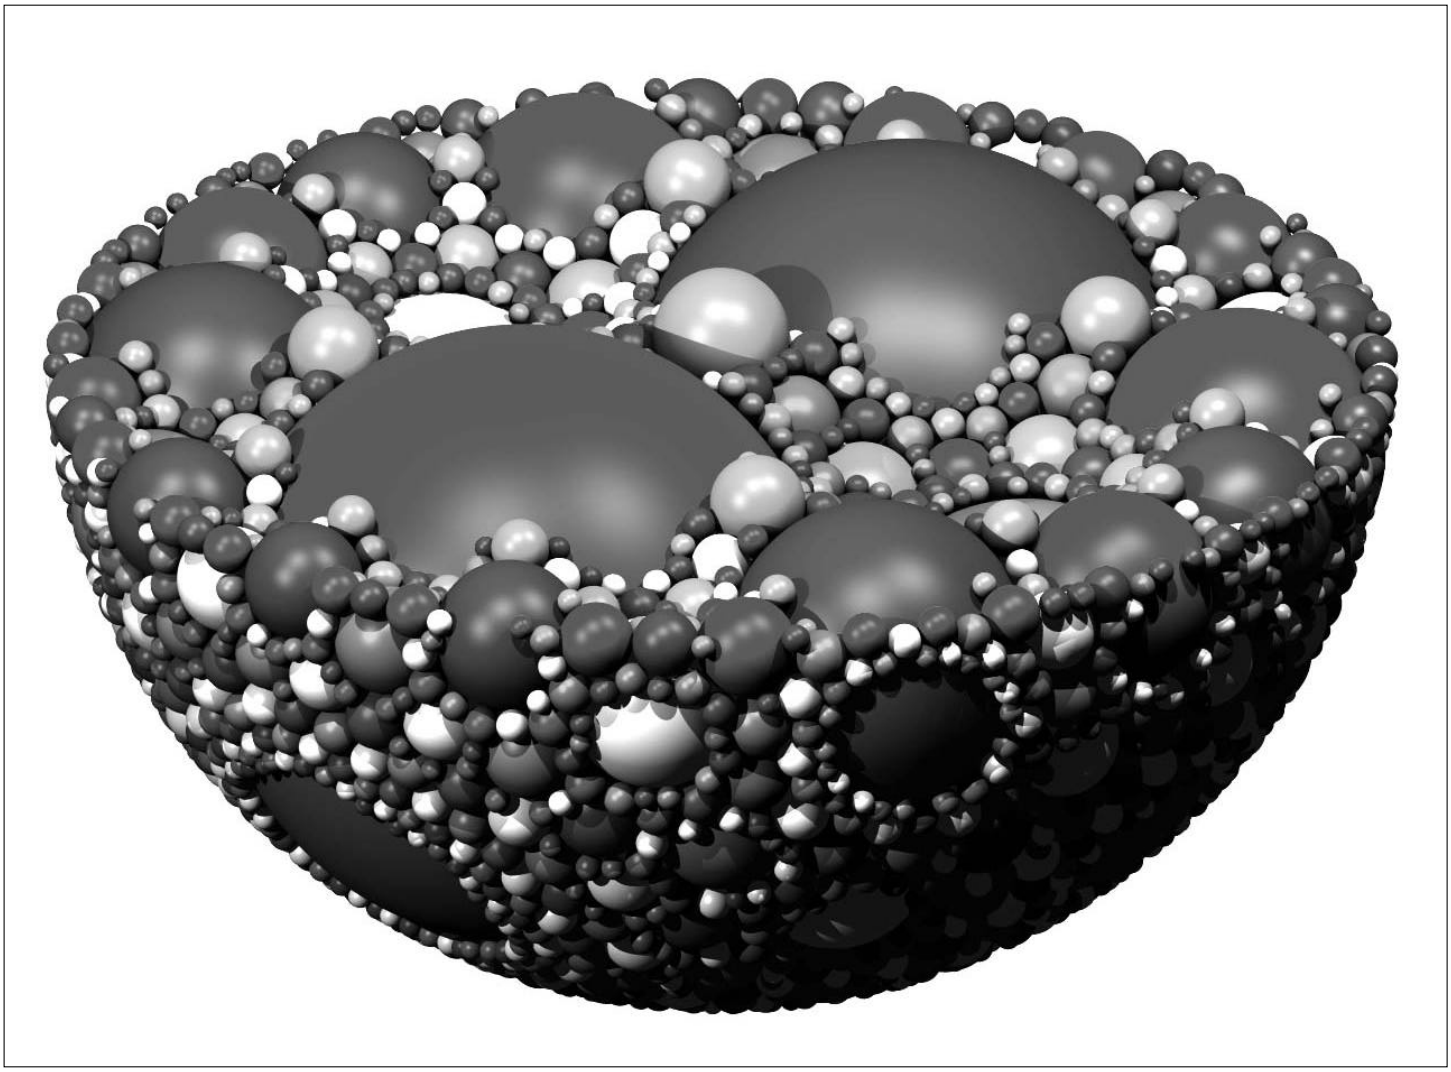

*Size distributions of soil particles. See article on p. 8.*

T A B L E O F C O N T E N T S

*The purpose of scientific computing is insight, not numbers*
Richard Hamming, mathematician and "father" of error correcting codes

With *insight* we often denote *understanding* and *perception* and by the latter we mean the way to grasp the meaning and to become aware of something through one of the senses, especially sight.

No one can deny the persuasive and dominant role of sight in our lives—especially given that an estimated fifty percent of the human brain's neurons are associated with vision and vision processing. Seeing is closely coupled with perception for humans: most natural languages use visual metaphors to communicate perception ("do you *see* what I mean?"). Few developments in computer-related technologies have caused as much as excitement in the general press as Virtual Reality, immersing the user in a synthesized visual environment. Visualization techniques and algorithms for the commercial and entertainment marketplaces deliver more and more "realistic-looking" images.

Science uses computing as a tool for number crunching on a huge scale. Visualization offers us a way to see the unseen amongst all the numbers and can bring enormous leverage to bear on scientific breakthroughs. By generating images from complex multidimensional data sets it enables researchers to visually perceive their simulations and computations, and thereby more quickly comprehend geometric and spatial relationships.

Ronald Peikert, responsible for the CSCS/SCSC visualization team, was awarded the "Best Case Study" award at *Visualization '96* for his video animation "Flow Visualization for Turbomachinery Design". You'll get to know the group and its activities on pages 8–10 of this *CROSSCUTS* issue.

Sight and visual communication are also essential for remote and distributed user support and training activities as well as collaborative development between geographically distant sites. The teleconferencing environment at CSCS/SCSC is described on pages 3–5. In the following pages the interested reader will furthermore find valuable information on numerical libraries with a special focus on the scientific Fortran library package ASL/SX to be installed on the Center's NEC SX-4. An overview on the joint collaboration between CSCS/SCSC and Electricité de France—one of the world's largest electrical utilities and the leading electricity exporter—is given on pages 29–31 of the technology transfer section.

I hope this *CROSSCUTS* issue will update you on ongoing CSCS/SCSC activities, being undertaken both by the Center itself as well as its user community. As always we try to make the newsletter interesting reading. I hope you'll enjoy it! ■

*Claudia Moor*

## More dedication to application portfolio

To better meet users expectations and to leverage its reinforced expertise, the portfolio offering of mathematical libraries and scientific applications is one of the primary 1997 CSCS/SCSC objectives. The driving concept is the reinforcement of the available application portfolio on either the Zurich-based or Manno-based HPC servers. CSCS/SCSC focuses special attention on user requirements for choosing a specific application package: the adequacy of the mathematical model, the numerical merits and the performances of the underlying solver, the direct compatibility with other solvers or pre-/post-processors, and ongoing collaboration with third research partners. Other important constraints are the commercial availability, the numerical effectiveness and/or the sustained performances of the required applications packages on the different HPC architectures currently available to the national community. For each of the required applications or libraries, CSCS/SCSC's responsibility is to identify the most appropriate computing environment and to help the concerned users to leverage it through effective tools and data/graphical interfaces.

However the license and maintenance fees for commercial application packages running on HPC servers can be expensive (even with academic pricing). And since the available investment budget for hardware and software components remains limited, a careful analysis of the needs expressed by groups of users involved in similar or overlapping scientific fields should be undertaken. The duty of CSCS/SCSC is therefore to workout with the user the most appropriate computing solutions and—whenever needed—to negotiate with vendors an appropriate and affordable licensing scheme.

The main focus in 1996 of such analysis and licensing strategy has been on the computational chemistry and molecular modeling. A similar effort targeting the mathematical and numerical libraries environments is currently underway. CSCS/SCSC will continue this effort for other traditional or emerging applications fields such as computational fluid dynamics, structural analysis, process engineering, signal treatment, visualization and animation.

Nevertheless, it is not sufficient to identify and install the required and the most effective application package on the most appropriate computing platform. It is just as crucial to develop and maintain the technical expertise in order to advice and support the users in their use of the available applications, as well as to interface these applications with effective pre- and post-processor software whenever needed. This is probably the biggest challenge for CSCS/SCSC! ■

*Jean-Pierre Therre*

PUBLISHING EDITOR: J.-P. THERRE
MANAGING EDITOR: C. MOOR
TECHNICAL EDITOR: N. VECCHI
PRODUCTION: S. WALTHER
EDITORIAL BOARD: M. GUTKNECHT, D. MARIC, R. MASTROPIETRO, W. P. PETERSEN
DESIGN: A. DELUCCHI

COMMUNICATING WITH CROSSCUTS: ARTICLES AND CORRESPONDENCE SHOULD BE DIRECTED TO: CROSSCUTS, CENTRO SVIZZERO DI CALCOLO SCIENTIFICO, VIA CANTONALE, CH-6928 MANNO. ELECTRONIC CORRESPONDENCE MAY BE SENT TO CROSSCUTS@CSCS.CH

CROSSCUTS IS PUBLISHED BY THE CENTRO SVIZZERO DI CALCOLO SCIENTIFICO/SWISS CENTER FOR SCIENTIFIC COMPUTING (CSCS/SCSC), ETHZ. CROSSCUTS CAN BE OBTAINED FREE OF CHARGE ON REQUEST. 

ARTICLES PUBLISHED IN CROSSCUTS DO NOT NECESSARILY REFLECT THE OPINIONS OR POLICIES OF CSCS/SCSC–ETHZ, NOR DOES PUBLICATION OF AN ARTICLE IMPLY AN ENDORSEMENT BY CSCS/SCSC.

PRINTED IN SWITZERLAND BY SOCIETÀ D'ARTI GRAFICHE GIÀ VELADINI & CO. S.A., LUGANO.

Eidgenössische Technische Hochschule Zürich
Ecole polytechnique fédérale de Zurich
Politecnico federale di Zurigo
Swiss Federal Institute of Technology Zurich

# PROGRESSES IN VIDEO-CONFERENCING AT CSCS/SCSC

*Angelo Mangili, Stefano Klett and Djordje Maric (CSCS/SCSC)*

## Introduction

Remote and distributed user support activities are vital for large national computing centers like CSCS/SCSC. In addition to traditional communication techniques, effective interaction between people can also take place using tele-conferencing techniques to provide support, training [1], collaborative development, and information sharing.

To meet specific users' requests and needs, CSCS/SCSC is in development of a tele-conferencing environment to support such educational activities as courses and tutorials, as well as regularly scheduled events such as High-Performance Computing Seminars. CSCS/SCSC is planning to extend this activity to conferences and workshops for Swiss academia. Another aspect of the tele-conferencing project is direct user support, promoting distributed work in the direction of the so-called Collaboratories [2]. Different technologies for tele-conferencing are currently in use or under evaluation by our team. These include high-performance dedicated ATM networks, standard ISDN connections and workstation-based solutions over conventional Internet.

## Objectives

CSCS/SCSC is currently looking for a solution for each of the following scenarios mainly because they are different, non-overlapping alternatives which reach a wide Swiss academic audience.

## TelePoly Project Integration

The integration of CSCS/SCSC in the TelePoly project [3] offers us a way to participate at ETHZ/EPFL events and vice-versa. These include: CSCS/SCSC High-Performance Computer Seminars (HPCS project), ETHZ/EPFL computing seminars, demonstrations and presentations of new products, supercomputer courses, user updates and information lessons, lectures on selected topics.

The TelePoly project is based on ATM technology. Audio and video signals are converted to ATM cells by a special box (called Cell-Stack module manufactured by K-Net Inc.). ATM cells are switched over the ATM network to the destination side, where another Cell-Stack provides the conversion back from digital ATM cells to audio and video signals.

During a seminar between ETHZ and EPFL, three video signals and one audio signal are sent from one site to the other. This implies the usage of three Cell-Stack modules. A minimal multi-site tele-teaching configuration requires three video signals [Fig. 1]. Some additional ATM ports had to be installed on the CSCS/SCSC ATM switch in order to provide full functionality for multi-point tele-conferencing. This equipment was acquired in September 1996 and is fully operational.

This scenario is very attractive; it provides a perfect video/audio quality (full Pal size), gives CSCS/SCSC a chance to be integrated in the ETHZ/EPFL tele-teaching network, and offers good experience with ATM technology for video/audio.

Some disadvantages are related to this solution; the current equipment is a well-suited, point-to-point solution but not perfect for multi-point tele-conferencing, at least one Cell-Stack module for each new site (future products of K-Net should provide an integrated solution to this problem). From a user support point of view this technology currently covers ETHZ and EPFL "only", and finally the costs are quite high. Additionally, this solution requires some "expert" work to setup each ATM permanent

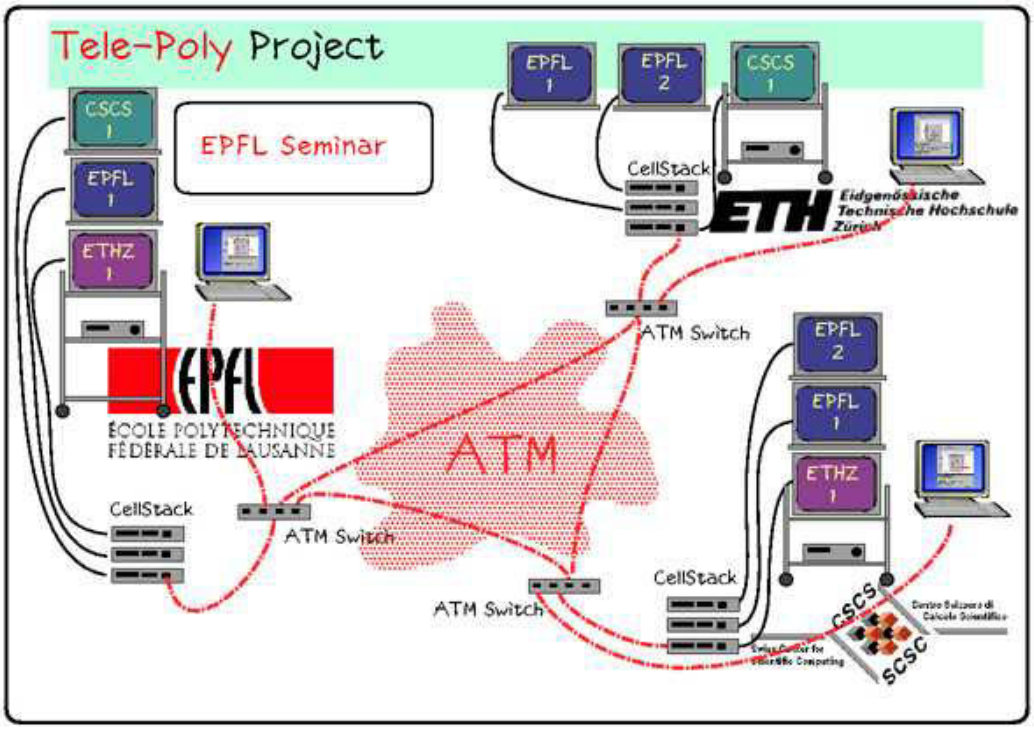


***Figure 1:*** *Example of a multi-site TelePoly session.*

*(Continued on page 4)*

*(Continued from page 3)*

virtual connection (PVC), this is not as easy as setting up phone or Internet connections.

There have been several sessions with TelePoly:

- 8 Oct 96: EPFL-CAST conference: "CAPA inauguration" (three sites EPFL-CSCS-ETHZ)
- 30 Oct 96: CSCS/SCSC presentation: Tera Inc. (two sites CSCS-ETHZ)
- 25 Nov 96: ETHZ seminar: "Can Tiny Computers Thrive?" Charles Thacker, DEC, (two sites ETHZ-CSCS).

## Workstation-Based Solutions

Image quality and transmission costs are all important issues in a true tele-teaching environment. Several tele-conferencing tools are currently available for desktop workstations (MBONE-tools [4], CUSeeMe, SGI InPerson, etc.), however, not all of them are fully multi-platform. These systems have the great advantage of reaching the audience directly on the desktop, but their generally poor video signal (highly dependent on CPU power and network load), is unsuitable for good quality slide transmission. When a good video signal transmission is possible (e.g., ATM), dynamic documents can use the video stream but at a very high cost. Such problems could be solved using additional software tools like whiteboards; unfortunately some of them have limitations on the size and the type of the supported documents.

The CSCS/SCSC Web SlideShow project [5] has been driven by the simple need to provide a wide and remote audience with an electronic and user-friendly access to the usual way of doing presentations: a speaker commenting a set of viewgraphs in front of an audience.

This tool is based on WWW capabilities such as electronic slide transmission with associated hyper-media notes and audio comments. The WWW media has been chosen for its flexibility and multi-platform capabilities. While the tool was at first meant for use in conjunction with standard tele-conferencing systems for live events [Fig. 2], the integration of hyper-media notes and audio comments describing the slides opened the possibility of asynchronous presentations. The first experiences have been very encouraging, and the Web SlideShow will now undergo further testing in the CSCS/SCSC tele-teaching environment. The first prototype is accessible at "http://www.cscs.ch/slide_show".

## ISDN Standard Connections

ISDN has been, since its inception, a world-wide accepted standard for tele-conferencing. Several integrated and highly reliable products exist, unfortunately because of the low bandwidth, the standard 128Kb/s provides a poor video quality especially on moving subjects.

Advantages:

- delivers reasonable video/audio quality using three ISDN channels (384Kb/s),
- offers Plug&Play products (low operational costs),
- integrates easily in an international tele-teaching network, possible candidates are all Swiss academia (universities and technical schools), and universities of Milano and Como,
- enables CSCS/SCSC to participate to any ISDN tele-conferencing event and vice-versa,
- provides well-suited, multi-point solution via a multi-point video-conferencing server (this service could be rented from PTT).

This technology has been evaluated and tested during a workshop on tele-education in Bellinzona organized by CSCS/SCSC in collaboration with Telecom Bellinzona and the Department of Education and Culture of Canton Ticino.

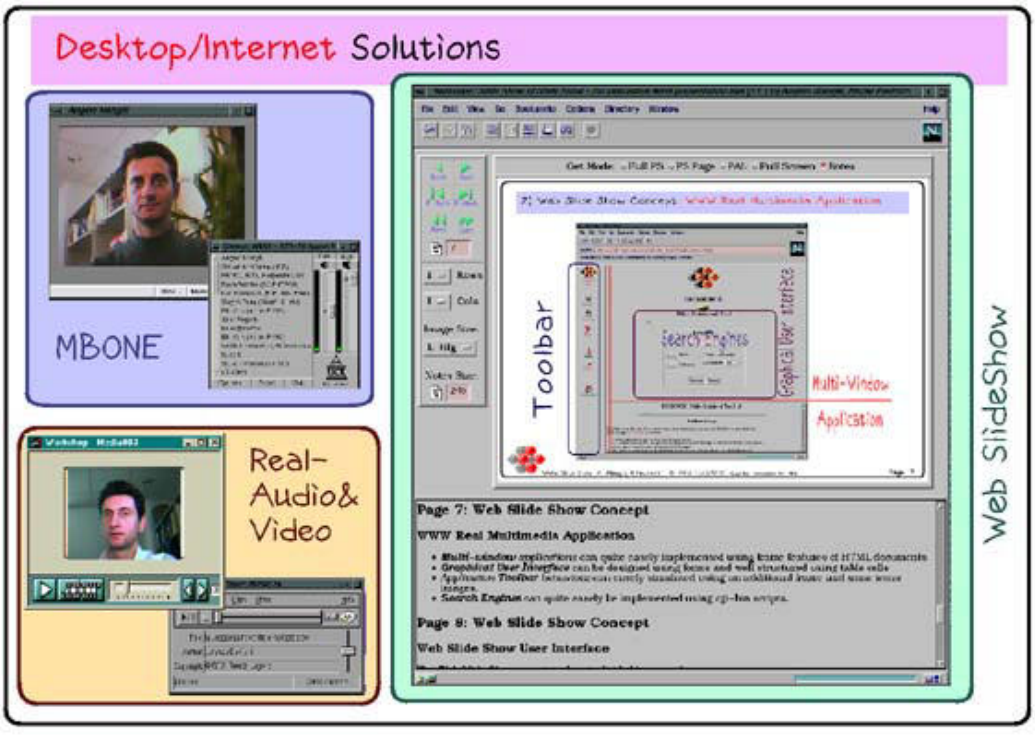


***Figure 2:*** *Possible usage of the Web SlideShow.*

## Future Plans

Promote the usage of the ATM-TelePoly solution which provides a comfortable and powerful environment for tele-teaching for all kind of seminars and workshops that could be of national interest.

Workstation-based solutions

will also be promoted in the future so as to be of benefit for the whole CSCS/SCSC user community. Our current goal is to be able to deliver full tele-seminars and tele-courses directly on the user desktop. This will be possible in a "natural" way because of the current trend to develop multimedia hardware as extensions of the main CPU unit (especially video/audio compression/decompression), the improvement of the network bandwidth, and the development of specific WWW applications (i.e., the Web SlideShow) dedicated to tele-teaching and tele-conference.

Meanwhile tele-conferencing and tele-working software are beginning to approach real professional integrated solutions, that could be helpful in the near future.

Long term plans will be to promote tele-working activities for real-time user consulting, training, and active contribution of CSCS/SCSC staff in user projects.

## References


[1] G. Vacek, A. Hilger, P. Flükiger, W. Klopper, H. P. Lüthi, "A Practical Ab Initio Quantum Chemistry Course Using the WWW as a Teaching Environment", CSCS/SCSC Technical Report (TR-96-24), 1996.

[2] R. T. Kouzes, J. D. Myers, and W. A. Wulf, "Collaboratories: Doing Science on the Internet", IEEE Computer, August 1996, p.40–46.

[3] T. Walter, "TelePoly—A Real Teleteaching Experiment", Institut für Technische Informatik und Kommunikationsnetze (TIK-ETHZ), http://www.tik.ee.ethz.ch/~walter/Telepoly/Telepoly.html

[4] S. Klett, "MBONE: Videoconferencing over Internet", CSCS/CSCS CrossCut, vol. 5, no. 1, April 1996.

[5] A. Mangili, N. Pedrozzi, "Web SlideShow, a Virtual Auditorium for Synchronous and Asynchronous Presentations", CSCS/SCSC Technical Report, to be published.


# Latest Trends in Mass Storage Technologies

*Roberto Mastropietro (CSCS/SCSC)*

The 5th NASA Goddard Conference on Mass Storage Systems and Technologies was held at the University of Maryland during the third week of September. Researchers, developers, suppliers and users of mass storage products got together in this annual conference to share experiences, exchange information and understand what can be expected in the next years in the mass storage products market.

The main topics of the conference were research and development of mass storage devices and hierarchical storage management systems.

## Mass Storage Devices

In the domain of mass storage devices, there is an obviously clear trend toward higher capacity and higher transfer rate media and devices. Tape devices suppliers especially, tend to provide users with a set of tape types and devices with each set targeting a specific application domain, e.g., backup, nearline, large files scientific archive.

Users can compare the characteristics of each kind of device/media (such as cost/MB, tape size, load time, positioning time and transfer rates) and choose the most appropriate combination of device/media to match their application requirements. Special consideration by both users and vendors was given to the load and positioning time for tape devices, primarily for nearline and commercial applications.

Developments in this direction include the recently announced IBM product Magstar MP, DVD disks from Sony, the work being done on optical tapes and other products that will be announced in upcoming months.

*(Continued on page 6)*

# News from the Services Section

Web-based services (such as information about system status/availability, online manuals, an updated list of recently installed and/or upgraded software and an overview of libraries and application software) are now available to users of CSCS/SCSC facilities.

Due to copyright issues, the documentation providers have requested that this site is password protected. To gain access, simply point your browser to: http://www-users.cscs.ch/ and use your frontend username and password.

We encourage you to have a look at it!

*(Continued from page 5)*

Large centers equipped with libraries and HSM are now giving less importance to long duration media; it is evident that due to the rapid technology developments data will not reside on the same tape for more than a few years, but rather be transfered to faster, cheaper, higher capacity and more reliable media before the previous media get close to the expected usage limit.

In terms of interfaces, most of the vendors are quickly moving in the direction of Fibre Channel (FC). On top of FC, different kind of protocols can be run, e.g. SCSI and IP. The current generation products has 100 MB/s bandwidth, while next generation will offer 400 MB/s. Fiber Channel switches are already available. Should FC really become widely adopted, the concept of network attached devices will eventually have a chance to play an important role in production environments.

For marketing and strategic reasons vendors are typically reluctant to reveal what they are working on in their labs and consequently what users might expect from them in the next years. Observing what leading researchers and institutes are doing gives us a clue to what vendors may be doing behind closed doors.

For the most of the researchers an achievable target is a TB per square centimeter. The technology to reach it is typically blue-green laser based, media vary from holographic media to new polymers where information can be stored not only on the surface but in the entire media. Work is also in progress to bring optical devices to similar performance levels as traditional magnetic devices within a few years. Carving silicon or other media with a ion beam is another technology being investigated with the goal of ensuring very long duration to stored information.

Some of the technologies mentioned above could be available today, were it not for the size and the cost of the drives which makes any attempt to manufacture and sell them unsuccessfull.

There is also more exotic research being conducted: for example, DNA-based storage or using a microorganism that can be found in the San Francisco Bay.

## Hierarchical Storage Management (HSM) Systems

In the HSM arena, the number of players is somewhat limited. Most of them are targeting the market of the medium range systems and expanding the supported platforms towards NT-based servers. Among the products that explicitly address the high-performance computing market is High Performance Storage System (HPSS) which is a joint development by IBM and some leading U.S. research centers. The product is now commercially available via IBM. HPSS has a parallel and scalable architecture. The only supported platforms are IBM RISC/6000 and SP2, HPSS is based on the IEEE Mass Storage Reference Model V5 architecture and relies on DCE to implement the distributed architecture. DFS support is expected next year. HPSS has about a dozen installations in U.S. computing centers, with some of them nearing production. The other two well known players in the high-performance computing market, Cray DMF and HP/Convex UniTree+ have gone through a plan revision due to the integration of the two companies respectively within SGI and HP organizations. SGI did not have a product similar to DMF and is now expected to support DMF on the MIPS-based systems. On the contrary, HP already has products that, at least partially, overlap UniTree functionalities. Decisions about HP future strategies for UniTree+ have only recently been taken: unitree+ will be supported on the Exemplar S and X class while support for HP K class system has been dropped.

The other products that have relevant installations are ADSM (by IBM), AMASS (by EMASS) SAM-FS (by Large Storage Configurations) and UniTree (by UniTree Inc.). Most of these products supports RISC based systems and a variety of tape devices and robotics.

A commonly agreed issue for the user community is the possibility to ensure an easy migration from one HSM to another. Currently there is no standard either for database and tapes format or for interchange formats. The same lack is reported in the domain of HSM benchmarking tools. ■

## NEW SOFTWARE INSTALLATION

NAG Fortran 77 Library, Mark 17, is now installed on *gottardo* (NEC SX-4, float0) and on *castor* (Convex Exemplar). Upon request the library could be installed on *cervino* and/or the HP workstations. ■

## NEC SX-4 MONTHLY MAINTENANCE

SX-4 hardware and software maintenance will take place every first Monday of the month. Hardware preventive maintenance is a fixed time (4 hours) from 8am till noon, then possible dedicated time for benchmarking from noon till 2pm, then finally software maintenance (variable time) in the afternoon. After shutdown and reboot the system will be open again to the users. ■

# THE TI-EDU NETWORK

*Mario Gay[1] (Università della Svizzera Italiana)*

TI-EDU is the newest of Canton Ticino's computer networks; the network connects a variety of institutes of higher learning (universities, colleges and *scuole universitarie professionali*), research institutes and libraries/documentation centers.

The 1996 realization of the *Università della Svizzera italiana* (USI) and the transformation of Ticino's technical trade schools into *scuole universitarie professionali* (SUP: university-level professional schools—the equivalent of *Fachhochschule* in German) were the catalysts that created the TI-EDU network. These events were strategic in forming a cantonal approach on the topic of network connectivity for those cantonal institutions not administrative in nature. In fact, although the need for a network such as TI-EDU existed for quite some time, the 1996 events lay the actual foundation for the network.

Prior to 1996, the majority of cantonal educational institutions were without a network connection; those that did have a network connection were connected either through various private service providers or through Canton Ticino's administrative network. Those utilizing the administration network were limited by the fact that the needs of an administrative body are quite different from those of an educational body—security measures just to mention one aspect.

So, the acknowledgment that non-administrative, educational-based institutions in Canton Ticino required network connectivity, and the formation of USI and SUP, brought about a concerted effort to offer a new network. With only eight months before the USI inauguration and limited personnel resources, Canton Ticino contacted CSCS/SCSC to profit from the expertise and know-how available at the center. The result was an agreement to allocate available resources to the TI-EDU network project from January to October 1996.

The short time frame forced a basic technical decision: to use proven technologies in off-the-shelf solutions. The network is based on TCP/IP protocol, relies on leased lines furnished by PTT and uses SWITCH as its Internet service provider.

TI-EDU has been an operational network since October 1996. The institutional users of TI-EDU are the *Università della Svizzera italiana* (Lugano and Mendrisio sites), the future SUP institutions, research institutes, documentation centers and cantonal libraries. Some high schools also take advantage of the TI-EDU network.

The realization of TI-EDU would have been very difficult without the expertise and resources of CSCS/SCSC; as such, the TI-EDU network was the fruit of a successful collaboration between the Canton Ticino and CSCS/SCSC. The future will undoubtedly hold other opportunities for beneficial collaboration. ■

1) At the time of writing, the author was a staff member of CSCS/SCSC.

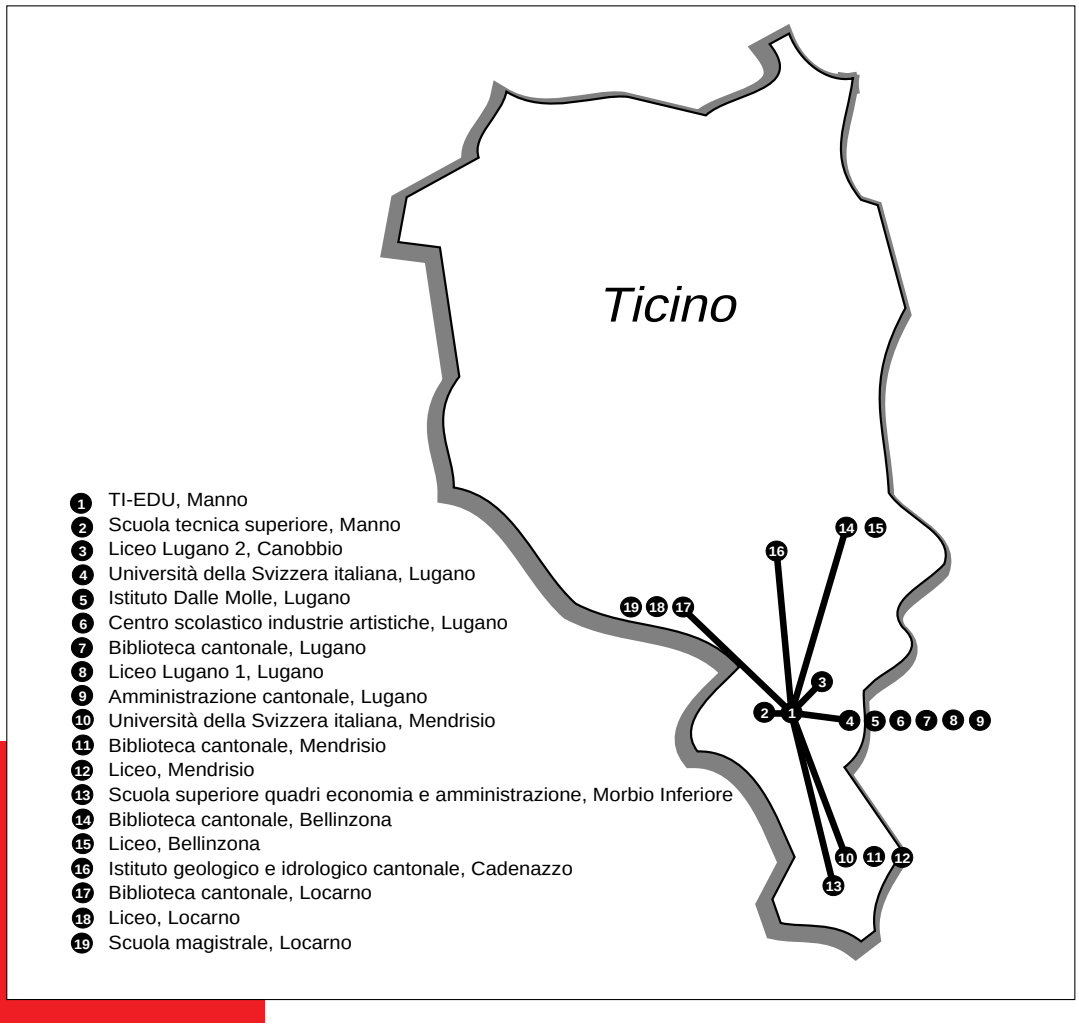


*Configuration of the TI-EDU network for spring 1997 (planned).*

# THE CSCS/SCSC SCIENTIFIC VISUALIZATION GROUP

*Ronald Peikert (CSCS/SCSC)*

## Introduction by Jean-Pierre Therre

*Data post-processing and scientific visualisation/ animation are stated today as critical parts of the overall numerical simulation/ modeling process.*

*In this perspective, CSCS/SCSC, being fully committed to anticipate and support the needs of the HPC community, has introduced a dedicated horizontal task force involving the experts from its scientific visualisation team (see below the article) but also various specialists of its "user support" group.*

*And therefore CSCS/SCSC provides in Manno and in Zurich a reinforced kernel of leading and interdisciplinary competences and a wide scope of complementary equipments in all the related technical aspects. This expertise is used to deal with always difficult issues like the coherent handling of multiple data format (e.g., CAD<->mesh generator interface) up to the video/movie production including the rapid prototyping of WEB-oriented and JAVA based human interfaces for over-the-net interactive visualisation (A. Mangili).*

*To leverage its expertise and dedication to user support, CSCS/SCSC (D. von Büren) has been recently appointed to operate and support the high-end graphical/visualisation environment pertaining to the newly introduced ETHZ's interdepartmental research project "VISDOME".*

## Who we are

The scientific visualization group of CSCS/SCSC, formerly the IPS graphics group, has been in operation since 1988. Our members, in alphabetic order, are:

- Rory Chisholm: responsible for the system administration of the hard- and software installed at the SCSC Zurich site.
- Jean Favre: responsible for visualization activities at Manno.
- Ronald Peikert: group leader.
- Martin Roth: currently writing his doctoral thesis supported by Sulzer Hydro and a KWF grant.
- Daniel von Büren: will operate the new Visdome lab.

Former members of the group were Hanspeter Scherbel, Thomas Koller and Andreas Thurnherr.

## The "philosophy" of our group

Scientific visualization applies computer graphics to scientific data, obtained either by experiments or by numerical simulation, typically on a high-performance computer. Such data sets consist of discrete data, structured or unstructured, and are often three-dimensional and/or time-dependent.

Work done by our group has always concentrated on visualization techniques more than on visualization products. From this, we believe that the most effective way to support ETH or university institutes is by establishing a project-based cooperation. In such a setting, chances are best that our know-how (and in many cases an existant piece of code) can be successfully applied. In an ideal situation, our project partners even have some knowledge of computer graphics.

The opposite of this would be a "help desk" type service. While we are happy to help, we don't try to systematically build up such a type of expertise. In other words, rather than knowing a little bit about all hardware/software products on the computer graphics market, we concentrate on a relatively small selection. Criteria for software packages are: discipline-independence, modularity, extensibility. Since commercial packages, even the most comprehensive ones, usually fall short when it comes to visualizing latest research data, the criterion of extensibility is of particular importance.

A good example of this is AVS, a software which has established a following within the visualization community: it had an excellent

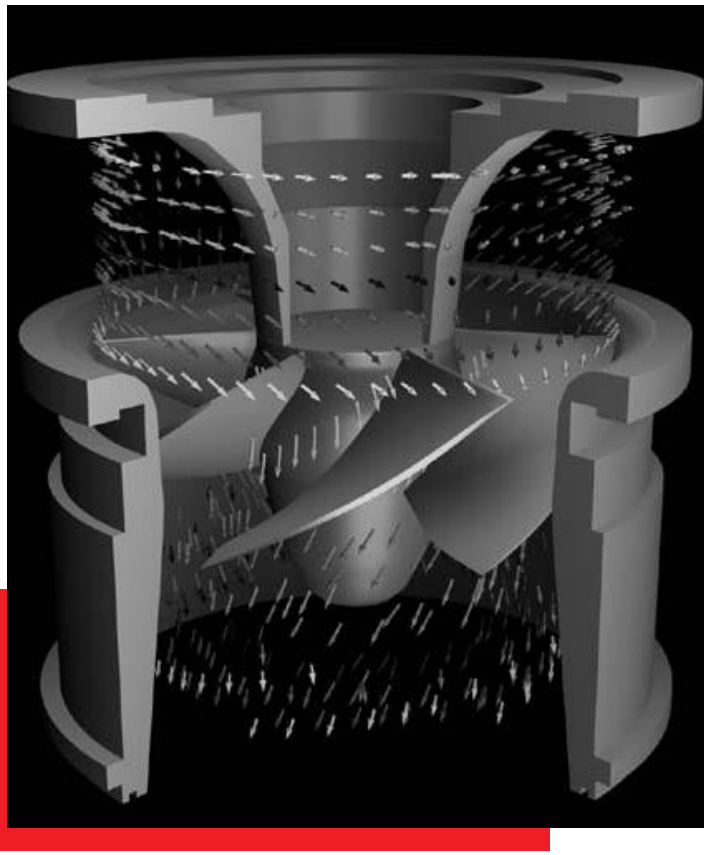

***Figure 1:*** *Simulation model of a Kaplan turbine runner, a type often used in river power plants. This picture is taken from a video we produced in collaboration with Sulzer Hydro. Often, in an initial sequence, the entire flow field is shown from a distant view.*

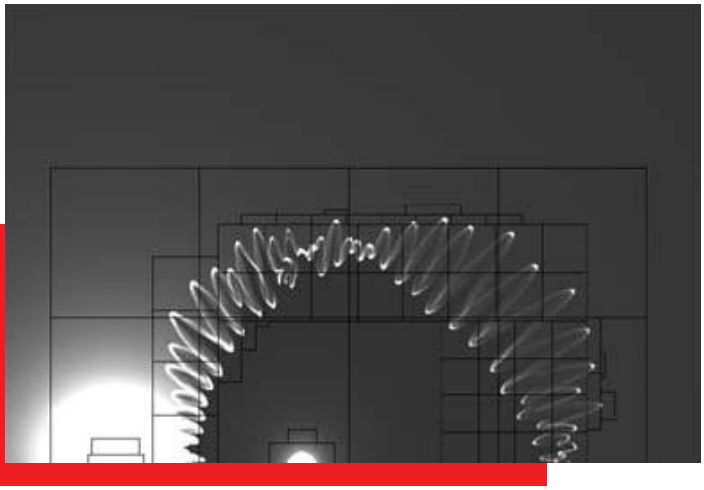

***Figure 2:*** *A visualization of astrophysical data illustrates our collaboration with the Institute for Astronomy. In this density plot of two colliding star winds, the winding structure represents the instability of the shock. The black lines are the block boundaries of the adaptive multi-resolution grid used for the numerical simulation.*

Usenet newsgroup, an active user's group an a large archive of user-supplied modules. For some years, it was the first choice for scientists who wanted to exchange quick implementations of new ideas. Our group handled most projects using AVS. In particular, our videos were produced mainly with either AVS or our own program Ren/Flux. As complementary tools, we made use of a few public-domain image processing libraries, Rayshade, and our own software for video titles and video editing.

At this time, it is less clear which of the candidates AVS/Express, NAG Explorer, Khoros (with new 3D enhancements), IBM Data Explorer or VTK will reach the same level of acceptance as AVS had. Therefore, our group will continue to support AVS and AVS/Express and in a limited fashion also Explorer and VTK.

## Projects with ETH and university institutes

We work together with a large number of institutes, but the most intense collaboration so far occurred with the ETH institutes of applied mathematics, astronomy, energy systems, and fluid dynamics. The advanced type of problems arising in these projects gives us the chance to explore more thoroughly the capabilities both of software packages and of visualization techniques.

A current example is multi-resolution grids used by Rolf Walder in astronomy. Jean Favre extended the array-of-fields feature of AVS/Express by several libraries of modules to make it suitable for practical use with multi-resolution grids.

An easier problem we are currently working on is component-labelled isosurfaces for an archeology project with Jürg Leckebusch.

## Industry projects

A long-term industry cooperation exists with Sulzer Hydro (formerly Escher Wyss). Resulting from this are two IrisGL-based programs: Ren and its successor Flux. Both programs, written primarily by Martin Roth, are used on a regular basis by Sulzer's engineers. During the design of turbines and pumps, the geometric shape of individual machine parts is repeatedly optimized, each time followed by a numerical simulation. At this point, our visualization programs help the engineers to better understand the effects of such optimization steps.

Besides the development of interactive visualization programs, we have produced about eight videos for Sulzer Hydro. Each of these videos was created for one of the company's world-wide power plant projects. The purpose of such videos is to demonstrate Sulzer's efforts in gaining efficiency by optimizing the design of relevant machine parts. The videos allow customers to understand how the flow is affected by the new designs and can thus help Sulzer to win competitions.

Another industry partner is Technopark-based company Bitplane. This company develops software for confocal microscopy and medical applications. We were able to contribute a module for isosurface generation written at CSCS/SCSC in a different context. A module for triangle decimation is currently under development by Lukas Vonblon, a computer science student.

## Research activities

After being a purely support-oriented group for many years, we became more present in research during the past two years. This is due to two events. First, ETH gained a professor for computer graphics: Markus Gross. With him

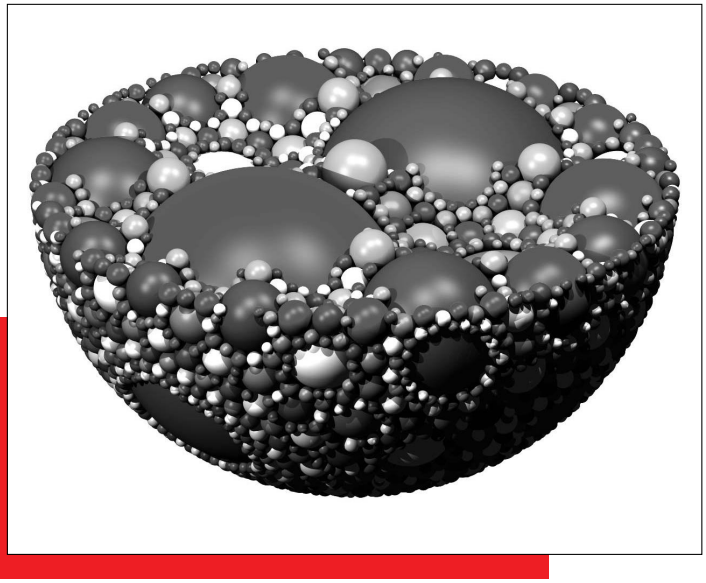

***Figure 3:*** *The problem of computing the fractal dimension of the Apollonian sphere packing came up in a project on size distributions of soil particles. In collaboration with the Institute of Terrestrial Ecology we obtained a result of 2.4739465… improving the previously known estimate of 2.4. Visualization has been helpful in checking the sphere generating program.*

*(Continued on page 10)*

*(Continued from page 9)*

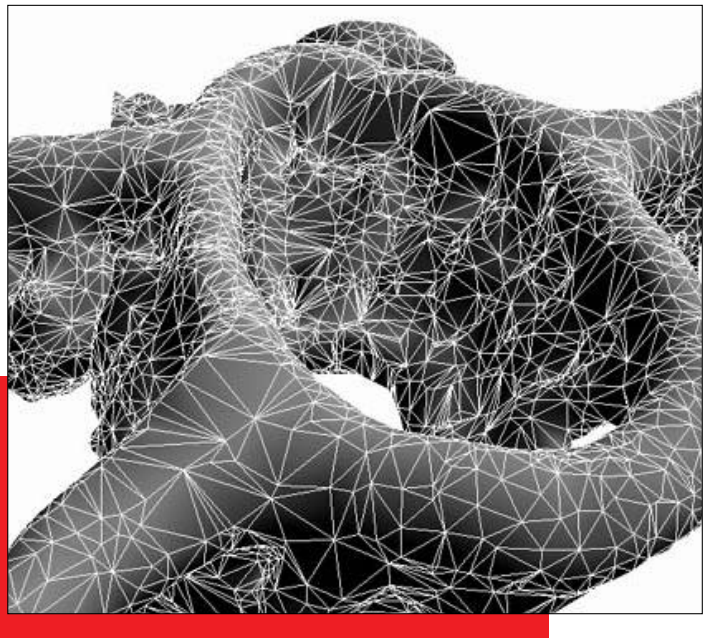

***Figure 4:*** *From our collaboration with Bitplane AG, this picture shows a pollen of Cabaea Scandens. An isosurface from a filtered 3D microscopy image has been computed and then decimated to the triangle mesh shown in white. The resolution of the original image is about 0.1 micrometers.*

and his group we have had many fruitful cooperations. In particular, Prof. Gross acts as the thesis advisor of Martin Roth. Then, when Jean Favre joined our group, we gained a computer graphics expert who has been active in research for many years.

We see several reasons why a group like this must participate in international research. Besides making jobs in the visualization group more attractive, a benefit arises also for our users by getting fast access to newest visualization techniques. For example, the following trends have been observed at recent conferences and already applied to our practical problems:

- in general: multi-resolution methods
- unstructured data: decimation, alpha-shapes
- volume visualization: acceleration techniques, texture-based techniques, component labelling
- flow visualization: vector field topology, line integral convolution, vortex recognition

## The "Visdome"

As some readers may be aware, a leading-edge visualization lab is under construction at ETH Zurich. The location in the dome of the main building inspired its name, the "Visdome".

The lab will be used as a shared resource by the members of AGVIS (Arbeitsgruppe Visualisierung), an organization comprising those ETH institutes with heavy requirements in computer graphics.

AGVIS defined four "pillars" for the lab, namely

- High-performance graphics: The Onyx2, configured with 3 graphics pipes and 8 CPUs, certainly meets this specification.
- Digital video environment: Real-time and 8-bit accurate recording from hard disk to video will be possible.
- Advanced I/O: Three video beamers will allow a 160-degree projection to a screen of 10 meters in diameter. Virtual reality I/O devices including force-feedback will also be added.
- High-speed networking: Fast-Ethernet and ATM will deliver the bandwidth, e.g. for SGI's new GLR render server concept.

CSCS/SCSC, and in particular our group, play an active role within AGVIS. While our director J.-P. Therre coordinated the GATT procedure leading to the purchase of the Onyx 2, three members of our group were involved in the technical evaluation. Daniel von Büren planned the audio and most of the video configuration. When the lab is operational, he will be responsible for the system administration, for help with video productions and with computer graphics-related problems.

## Other activities of the group

For historical reasons, the scientific visualization group is responsible for the system administration of the CSCS/SCSC servers and workstations located in Zurich. We are currently replacing our old SPARCservers 4/490 and 4/690MP by a Sun Enterprise 3000. Our twelve Sun workstations are now all upgraded to Solaris 2.5, whereas the twelve SGI workstations run mostly Irix 5.3. Recently, we completed the transition from data-less (not diskless!) clients to stand-alone machines.

On the software side, installed application packages include AVS, AVS/Express, Matlab, Maple, Mathematica, UniChem and an extensive list of public-domain packages. Our hosts run AFS for file-sharing with high-performance computers (Cray and Paragon). On AFS, our former member Hanspeter Scherbel installed TeX for various platforms. On the other hand, we still profit from the dozens of public domain programs installed in the Irix version by Wolfram Willuhn (Kommunikationstechnik). The idea of ETH-wide installation of public-domain programs can make life easier for many system administrators, and we would be happy to see more of it in the future. Two more things we don't have to do ourselves: For data backup and for high-quality printing we rely on the excellent services Net-Backup and VPP provided by the computing services of ETH.

Relatively often we are asked for video editing work, starting from given sequences of digital images. When our resources are not used otherwise, we can provide such a service in analog Betacam quality (both in Zurich and in Manno). A next-generation digital video system (Panasonic D5) will soon be part of the Visdome environment.

Further information on our group can be found under http://www.scsc.ethz.ch/SV/index.html. ■

# BUILDING ON SOLID GROUND

*Armin Friedli (CSCS/SCSC)*

Numerical libraries can provide high-quality building blocks for technical and scientific programming. Users should make use of them in order to benefit from effective numerical algorithms and to avoid duplication of efforts. A survey of the current offerings at CSCS/SCSC regarding numerical libraries and user support in numerical analysis is given. Also, a CSCS /SCSC T-shirt or CD can be won.

## The problem (a true story)...

A scientist was recently faced with having to extend an old program (the author of which naturally had long since left the university) in order to compute larger cases. She observed that the program behaved rather slowly at a point where the inverse of a matrix is being computed. Wisely, she called on the *Advanced Numerical Algorithms Group* of CSCS/SCSC's Support Section. Already during the first telephone conversation with the user, the adviser doubted that the inverse of the matrix really had to be determined. The following unfolded: For a given matrix A and a given vector y, a vector x is sought, given by $x = A^{-1}y$.

## ... and the wrong solution

Inspection of the code revealed that the original program author had copied a routine to invert an n by n matrix A out of a textbook. Unfortunately, the method used not only was based on the exchange algorithm (using $3n^3$ floating point operations), but also multiplied together A and the computed inverse (to see how accurate the computation was?), unnecessarily adding another $2n^3$ operations. The best algorithm for inverting A, based on Gaussian elimination, could have done the job in basically $4/3n^3$ operations—almost four times faster.

The second issue here, of course, is the fact that the inverse of A was not needed at all. Solving Ax= y for vector x is a standard problem in numerical linear algebra and can be done in basically $2/3n^3$ operations.

All in all, the original piece of code can be sped up by a factor of nearly eight by simply calling a routine out of one of those proven numerical libraries like NAG or LAPACK!

## Lessons to learn

- Consult a specialist in numerical analysis. Your problem could be reformulated to lend itself to an optimal numerical algorithm.
- Try to put your numerical computation into a standard form for which robust, optimized and ready made routines in numerical libraries are available.
- Don't blindly code up formulas and algorithms out of papers and textbooks. Even copying code out of a textbook may be a trap, as illustrated above.

## User support in numerical analysis

Help regarding numerical libraries can be sought by sending an email to help@cscs.ch and specifying the key word *library* in the subject field.

Another way is to get into direct contact with one of the members of the Advanced Numerical Algorithms Group in the Scientific Projects and Support Section of CSCS/SCSC. The members are: Andrea Bernasconi (aberna@cscs.ch), Eric de Sturler (sturler@cscs.ch), Armin Friedli (friedli@cscs.ch), Martin Gutknecht (mhg@cscs.ch), Wesley Petersen (wpp@cscs.ch).

Whereas all these members have general knowledge in numerical analysis, special expertise exists in the following areas: Numerical software, numerical libraries, ordinary differential equations, stochastic differential equations, fast Fourier transforms, general problems in linear algebra, iterative methods for linear equations, High Performance Fortran, concepts for vector and parallel algorithms, approximation theory, simulated annealing, neural nets. In those cases where a particular problem area is not sufficiently covered by the team, there exists a considerable reservoir of expertise at ETH that can be tapped, especially so in the Seminar für Angewandte Mathematik and the Institut für Wissenschaftliches Rechnen.

## NAG (Numerical Algorithms Group)

The NAG Fortran Library is a comprehensive collection of Fortran 77 routines for the solution of numerical and statistical problems. An effort is being made that all central computing facilities at CSCS/SCSC and ETH Zurich are equipped with the latest version, i.e., Mark 17, of the NAG Fortran Library. At CSCS/SCSC in particular, the SX-4 and the HP/Convex Exemplar are now furnished accordingly; the HP cluster and Convex 3820 can be served if user demand exists. The routines are not parallelized.

*(Continued on page 12)*

*(Continued from page 11)*

## ASL/SX (NEC's Advanced Scientific Library)

The ASL/SX is NEC's proprietary Advanced Scientific Library for the SX-4, currently at release 13.0. Using routines out of this Library will therefore make applications less portable. However, we decided to provide ASL/SX in order to benefit from the excellent performance some of the routines exhibit on the SX-4. And the most important reason: Some routines (in the area matrix times matrix, solving linear equations, fast Fourier transforms) are parallelized and show really superior performance.

Editor's note: See also pages 13–16

## LAPACK

LAPACK is a library of Fortran 77 subroutines for solving the most commonly occurring problems in numerical linear algebra, i.e., systems of linear equations, linear least squares problems, eigenvalue problems and singular value problems. LAPACK uses direct methods; iterative solvers for sparse matrices are not included. It has been designed to be efficient on a wide range of modern high-performance computers. The routines of LAPACK installed on the SX-4 are not parallelized at the moment. Being in the public domain, LAPACK is the most portable of the libraries described here. LAPACK supersedes LINPACK (systems of linear equations and linear least squares problems) and EISPACK (eigenvalue and singular value problems). Routines out of LINPACK and EISPACK should be replaced by routines from LAPACK for reasons of improved performance (speed, accuracy, stability).

## BLAS

The Basic Linear Algebra Subroutines are the most basic building blocks for numerical linear algebra and cover vector operations (BLAS-1), matrix times vector type operations (BLAS-2) and matrix times matrix type operations (BLAS-3). The BLAS are used extensively in LAPACK and many other packages and libraries. The BLAS exist on most computer systems at CSCS/SCSC in highly optimized versions.

## Other libraries

There are other numerical libraries on CSCS/SCSC computers and many more available from vendors and from public domain sources like *netlib*. Space doesn't permit us to list them all here. Nor would it make much sense to do so: We suggest to sit together with one of our advisors, discuss the problem to be solved and then decide how to proceed further.

## Documentation

For electronically accessible documentation for the NAG, ASL/SX, LAPACK and BLAS routines as well as an overview of numerical libraries actually installed on CSCS/SCSC machines start at http://www-users.cscs.ch/ (see box on page 5).

NAG: There is an Introductory Guide to the Fortran Library, Mark 17, available from NAG Ltd. for GBP 25.- This Guide provides an excellent overview of the capabilities provided by the NAG Fortran Library. The detailed description of all routines is available in the 12 volume Fortran Library Mark 17 Manual and costs GBP 185.- A complete set of NAG documentation is located at the following places: CSCS Library, Manno; SCSC ETHZ, RZ F10; ID ETHZ RZ G11. For ordering your own copies, the email contact to NAG is: infodesk@nag.co.uk

LAPACK: Printed documentation is available from SIAM, Philadelphia: LAPACK Users' Guide, second edition, 1995, US$<20.- Copies can be viewed at CSCS in Manno and SCSC at ETH Zurich, RZ F. The book is also available over the Internet: http://www.netlib.org/lapack/lug/lapack_lug.html.

## Floating point formats on the NEC SX-4

Although there are three different floating point formats available on the SX-4, we definitely recommend that only 'float0', i.e., the IEEE standard for binary floating point arithmetic (ANSI/IEEE Std 754-1985), be used. It is our policy to restrict the installation of numerical libraries to the IEEE format as much as possible. The ASL/SX library is available in the float0 format; the float2 (Cray) format will be supported only until the end of 1997. The NAG library is supported only in the float0 format. We ask users who think that they need libraries other than in float0 format (in the long run) to get in contact with us. ■

## WIN A CSCS/SCSC T-SHIRT OR CD

As mentioned, the elements of the inverse of a matrix are rarely needed. However, the need does exist sometimes. If in one of your applications you are computing some or all inverse entries, let us know. Describe the problem you are solving on one or two pages. If we can't prove to you that you could do without explicitly inverting the matrix, you win one of those famous CSCS/SCSC T-shirts or classical music CD. ■

# ASL/SX Advanced Scientific Vectorized Serial/Parallel Library

*Tatsunobu Kokubo (NEC ASL Development Group), Charles Henriet, René Hausammann (NEC European Supercomputer Systems)*

The ASL/SX is a comprehensive scientific Fortran library package that provides powerful support on the SX-4 for numerical calculations, such as basic matrix algebra, simultaneous linear equations (direct and iterative methods), eigenvalues and eigenvectors, FFT, ordinary differential and integral equations, numerical integration, special functions, random numbers, among others. Developed by NEC, the library incorporates the latest numerical computational techniques and takes advantage of the superior performance of the SX-4 series. The ASL/SX provides the scientific programmer with easy access to the powerful parallel-vector architecture of the SX-4 series supercomputers, without special knowledge of vector and parallel tuning and programming techniques.

First benchmarks at CSCS/SCSC with major user codes have demonstrated the substantial optimization work ASL/SX has undergone for the SX-4 series, which results in optimum performance in serial mode and very good scalability for parallel processing. The ASL/SX is built upon basic functions that are optimally vectorized, extended functions adapted to the SX-4 architecture at the algorithmic level, and parallelized functions that fully exploit the parallel architecture of the SX-4 series.

The ASL/SX Library will be installed on CSCS/SCSC's SX-4 in IEEE floating point data format, as well as Cray format until such time when all users have migrated to the IEEE format. This article presents a first introduction to the ASL/SX.

## 1. Functionality

### 1.1 Basic and Extended Functions

These functions work in serial mode, based on optimally designed algorithms, as indicated in the following.

Basic Matrix Algebra: The basic matrix algebra operations are included. For some operations, performance is enhanced by data type conversion functions. For example, a 2-dimensional array of a real symmetric matrix is converted to a 1-dimensional, column-oriented, compressed data array.

For real matrix multiplication, the Strassen algorithm is used. For a real *M*N* matrix multiplied by a real *N*L* matrix, about *2*M*N*L* floating point operations are required. If the Strassen algorithm is used, about 12% less operations are needed which results in reduced computation times.

Simultaneous Linear Equations Both the direct method and the iterative method are provided:

- The direct method is applied to real matrices, complex matrices, positive symmetric matrices and Hermitian matrices. For real matrix LU-decomposition, the Block Gauss method is used. The block algorithm takes advantage of the SX-4 vector type architecture.
- The iterative method is used for real symmetric and asymmetric sparse matrices. For symmetric sparse matrices, preconditioned conjugate gradient (PCG) method is applied.
  For asymmetric sparse matrices, preconditioned bi-conjugate gradient (PBCG) method and preconditioned conjugate gradient square (PCGS) method are used.
  The following algorithms are provided as preconditioning algorithms to accelerate convergence: the incomplete-LU (ILU) decomposition and the tridiagonal factorization (TF) for regular sparse matrices, and the ILU decomposition and the symmetric SOR matrix splitting (SSOR) for random sparse matrices. In addition, Gustafsson acceleration is implemented in the ILU and TF.

Eigenvalues and Eigenvectors The eigenvalue and eigenvector solver corresponds to the standard and the generalized eigenvalue problem.

The standard eigenvalue problem is available to real matrices, complex matrices, symmetric matrices and Hermitian matrices. In particular, for the transformation of symmetric real matrices to tridiagonal matrices, the Block Householder transformation algorithm is applied. In addition, QR method, root-free QR method, bisection method and inverse iteration method are used for getting eigenvalues and eigenvectors. Again, the block algorithm takes advantage of the vector type architecture.

1D, 2D, 3D FFT FFTs are available for 1-, 2-, and 3-dimensional complex and real Fourier transforms and 1-dimensional cosine, sine, cosine-q, and sine-q transforms. The FFTs use the Temperton algorithm (real and complex) and the Gentleman-Sande algorithm (cosine, sine, cosine-q, and sine-q). In the coding technique, needless floating-point calculation is omitted.

Spline Functions Interpolation, smoothing, numerical derivatives and numerical integration use cubic and bicubic spline functions.

Roots of Equations Roots to algebraic equations, nonlinear equations and simultaneous nonlinear equations are provided.

*(Continued on page 14)*

*(Continued from page 13)*

For real coefficient algebraic equations, the Cardano method, the Ferrari method and the Hirano method are used.

Special Functions Bessel-related functions, Gamma-related functions, elliptic functions, indefinite integrals of elementary functions, associated Legendre functions, orthogonal polynomials, and other functions are available.

Random Numbers Uniform random numbers, distribution random numbers, and their random number tests are provided. For uniform random number generation, the linear congruence method (short period) and the maximum-length linearly recurring sequence method (long period) are used.

### 1.2 Parallel Processing Functions

Parallel versions of routines are available for the following functions:

- Basic Matrix Algebra
- Simultaneous Linear Equations (Direct Method)
- 2D, 3D real and complex FFT

The parallel functions are based on almost the same algorithms as the serial ones. Optimal balance in granularity was of prime importance to their parallel implementation.

## 2. Scalability of Important Parallel Functions

SX-4 performance charts, as a function of number of processors, are given below for the examples of matrix multiply, complex three-dimensional Fast Fourier Transform with complex arguments, and LU decomposition of real and complex matrices:

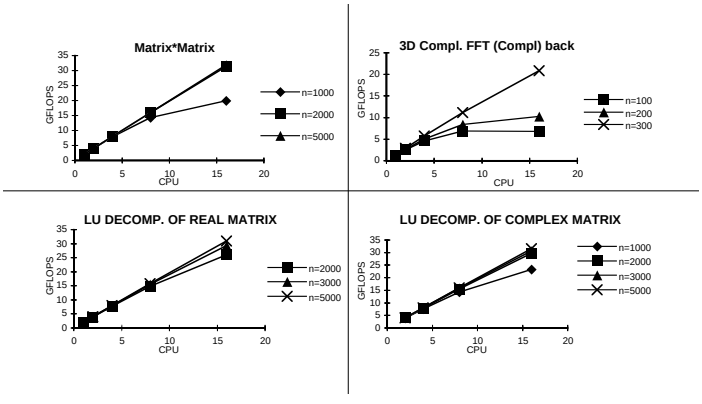


## 3. Practical Example: 3-D Complex Fourier Transform for 3-D Complex Data with Arbitrary Radix

### 3.1. Introduction

Transforms corresponding to the following equations are performed when your input data includes three-dimensional complex arrays:

Forward transform

$$d(j,t,v)= \sum_{k=1}^{nx} \sum_{s=1}^{ny} \sum_{u=1}^{nz} c(k,s,u) * \exp\{-2*\pi*(j-1)(k-1)/nx\}* \exp\{-2*\pi*(t-1)(s-1)/ny\}* \exp\{-2*\pi*(v-1)(u-1)/nz\}$$

$$j=1,...,nx \quad t=1,...,ny \quad v=1,....,nz$$

Backward transform

$$d(j,t,v)= \sum_{k=1}^{nx} \sum_{s=1}^{ny} \sum_{u=1}^{nz} c(k,s,u) * \exp\{2*\pi*(j-1)(k-1)/nx\}* \exp\{2*\pi*(t-1)(s-1)/ny\}* \exp\{2*\pi*(v-1)(u-1)/nz\}$$

$$j=1,...,nx \quad t=1,...,ny \quad v=1,....,nz$$

where *i* represents the imaginary unit

The above equations are obtained by dividing the interval $[0,2\pi]$ for each dimension into *nx,ny* and *nz* parts, respectively.

### 3.2. Serial and Parallel Processing

ASL/SX provides a way to compute the transforms corresponding to the previous expressions in serial and parallel mode by using different sets of subroutines:

- for serial processing: the cfc3fb/cfc3bf subroutines are used
- for parallel processing: the gfc3fb/gfc3bf subroutines are used

### 3.3. Normalization

If these subroutines are used to compute a backward transform after performing a forward transform, the data obtained will have values that are several times as large as those of the original data.

Therefore, the results of either the forward or the backward transform must be normalized.

The normalization has to be added explicitly by the user.

### 3.4. Processing Switch

The parameter which controls the transform direction (forward or backward) is defined as:

forward transform: *isw* = +1

backward transform: *isw* = - 1

In some other references, the definition of forward and backward transforms may be reversed from the ones used within ASL/SX.

### 3.5. Initialization

The processing switch parameter *isw* can also be used for doing the initialization step. In this case *isw* is equal to zero. This is very useful if a large number of transforms have to be processed and the number of input data values *nx,ny* and *nz* do not change.

This way of working enables processing to be performed more efficiently since calculations such as factorization or the creation of trigonometric function tables are performed only once.

In this case the work areas for factorization and trigonometric function tables must be retained in memory so they can be used later as input data.

### 3.6. Dimensions of the Arrays

Using ASL/SX for performing three-dimensional complex Fourier transforms on complex input data, one has to distinguish between adjustable dimensions of the arrays (*lx,ly,lz*) and the number of input data in each direction (*nx,ny,nz*).

Odd numbers should be set for the adjustable dimensions (*lx,ly* and *lz*) of the array C(*lx,ly,lz*). If nx is even then *lx=nx*+1is the best value. The same rule is valid for *ly* and *lz*.

However, if *nx<lx* (or *ny<ly* or *nz<lz*), array values of elements other than those in which input/output data is stored in array *C* will be undefined after the ASL/SX subroutine is used.

In fact, ASL/SX uses rotation for speed enhancements because the rotation gives maximum vector lengths: *ly*lz* (*x* direction), *lz*lx* (*y* direction) and *lx*ly* (*z* direction).

For the direction along the *x* direction: *C*(*lx,ly,lz*) is rotated to *C*(*ly,lz,lx*) and a multi-FFT in the *x* direction with length *ly*lz* is performed. For the direction along the *y* direction: *C*(*ly,lz,lx*) is rotated to *C*(*lz,lx,ly*) and a multi-FFT in the *y* direction with length *lz*lx* is performed. For the direction along the *z* direction: *C*(*lz,lx,ly*) is rotated to *C*(*lx,ly,lz*) and a multi-FFT in the *z* direction with length *lx*ly* is performed.

According to these rotations the array *C* is not retained after the transform: for the first dimension the array *C*(*nx+1* through *lx*,*1* through *ly*,*1* through *lz*) will be undefined and some additional vector and parallel operations must be implemented if one wants to save the structure of the array for later use after performing the FFT transforms.

If we want to perform a 3D complex FFT on an array *c*(*lx=nx+1*, *ly=ny+1,lz=nz+1*) in using *nx,ny* and *nz* complex data values and we want to save the structure of array *c* for later use, we have to define three arrays *cx*(*lx,ly*), *cy*(*lx,lz*) and *cz*(*ly,lz*) and compute the following:

a) saving the original *c* structure

```
do i=1,lx
  do j=1,ly
    cx(i,j)=c(i,j,lz)
  enddo
enddo
do i=1,lx
  do j=1,lz
    cy(i,j)=c(i,ly,j)
  enddo
enddo
do i=1,ly
  do j=1,lz
    cz(i,j)=c(lx,i,j)
  enddo
enddo
```

b) computing the transform by calling ASL/SX on the array *c*(*lx,ly,lz*) and using *nx,ny* and *nz* data values

c) restoring the original structure of the *c* array

```
do i=1,lx
  do j=1,ly
    c(i,j,lz)=cx(i,j)
  enddo
enddo
do i=1,lx
  do j=1,lz
    c(i,ly,j)=cy(i,j)
  enddo
enddo
do i=1,ly
  do j=1,lz
    c(lx,i,j)=cz(i,j)
  enddo
enddo
```

### 3.7. Moving from Previous Libraries to ASL/SX

In order to help users to understand how they can modify their existing programs for using ASL/SX and to get better performance, the web page http://www-users.cscs.ch/all/sx4/docs/ASL.appendix.html (only for NEC users) contains the source of four small programs computing a three-dimensional complex Fourier forward and backward transforms using as input data an array *c* defined by

c(i,j,k)=cmplx(i+j+k) for i=1,..5,j=1,..4,k=1,..3

and the following libraries: libfft, scilib, ASL/SX (serial) and ASL/SX (parallel).

The following table summarizes the needed changes and can be used to understand the short source files provided. It only applies to float2 data format.

| program used | libfft.f | scilib.f | aslsxs.f | aslsxp.f |
|---|---|---|---|---|
| library used | libfft2.a | libscilib2.a | libasl2.a | libasl2.a |
| initialization of fft | cftfax | xcftfax | cfc3fb | cfc3fb |
| computation of fft | cfftmlt | xcfftmlt | cfc3fb<br>gfc3fb | cfc3bf<br>gfc3bf |
| forward isw | -1 | +1 | +1 | +1 |
| backward isw | +1 | -1 | -1 | -1 |
| normalization | included | needed | needed | needed |
| f77sx options | -float2 | -float2 | -float2 | -multi-float2 |
| initialization of data | init.f | init.f | init.f | init.f |

The program and initialization sources can be seen: http://www-users.cscs.ch/all/sx4/docs/ASL.appendix.html (only for NEC users)

## 4. Application to the NAS Parallel Benchmarks

To show how these ASL/SX FFT routines perform and scale versus the number of CPUs, one can compare the figures obtained from the well-known FT NAS Parallel Benchmark.

In this benchmark a 3D partial differential equation is solved using FFTs. This kernel performs the essence of many spectral methods. The rules of the NPB specify that assembly-coded library routines may be used to perform matrix multiplication and 1-dimensional, 2-dimensional, or 3-dimensional FFTs. Thus this benchmark is somewhat unique in that computational library routines may be legally employed. Results are shown below.

Results of the 3D FFT PDE benchmark NPB FT:

| #CPU | SX-4 (seconds) | ratio C90/1CPU | ratio T916/1CPU |
|---|---|---|---|
| 1 | 53.90 | 2.05 | 1.20 |
| 2 | 26.97 | 4.10 | 2.40 |
| 4 | 13.49 | 8.20 | 4.80 |
| 8 | 6.81 | 16.24 | 9.52 |
| 16 | 3.43 | 32.24 | 18.90 |
| 32 | 1.85 | 59.78 | 35.03 |

*(Continued on page 16)*

*(Continued from page 15)*

## 5. Application to a State-of-the-Art Molecular Dynamics User Code

If too small a number of input data is used, the parallel processing overhead significantly affects the calculation costs, and performance may decrease (see graph in paragraph 3). However, due to the fact that ASL/SX uses rotation for speed enhancements, quite good speedup can be obtained even if the number of data processed in each direction is small. Below we show how the performance of two major sections s1 and s2 of a real program can be improved by using the parallel ASL/SX 3D Complex FFT subroutines. We compare program execution times between serial and parallel ASL 3D complex FFT routines, as a function of CPUs (the number of input data processed is *nx*=54 *ny*=54 *nz*=120):

| # CPU | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|---|---|---|---|---|---|---|---|---|
| s1 parall | 952.68 | 529.32 | 377.31 | 322.03 | 260.50 | 245.86 | 237.14 | 223.95 |
| s1 serial | 926.43 | 926.43 | 926.43 | 926.43 | 926.43 | 926.43 | 926.43 | 926.43 |
| s2 parall | 445.13 | 249.91 | 178.96 | 154.39 | 125.29 | 118.84 | 115.68 | 110.05 |
| s2 serial | 437.66 | 437.66 | 437.66 | 437.66 | 437.66 | 437.66 | 437.66 | 437.66 |

In this case by simply switching from ASL/SX serial FFT to ASL/SX parallel FFT, one can improve the global performance of the whole application dramatically as demonstrated by the elapsed times and the MFLOPS obtained from the entire application program, in which parts other than FFT are also parallelized:

| #CPU | fft serial | | fft parall | |
|---|---|---|---|---|
| | elapsed | MFLOPS | elapsed | MFLOPS |
| 1 | 8527.90 | 1801 | 8557.89 | 1801 |
| 2 | 5003.86 | 3053 | 4434.69 | 3466 |
| 3 | 3857.80 | 3961 | 3047.03 | 5047 |
| 4 | 3302.68 | 4628 | 2379.47 | 6487 |
| 5 | 2903.30 | 5262 | 1940.20 | 7958 |
| 6 | 2683.05 | 5693 | 1697.94 | 9105 |
| 7 | 2537.50 | 6018 | 1527.20 | 10128 |
| 8 | 2408.19 | 6341 | 1372.50 | 11266 |

## 6. Practical Example: Eigenvalues and Eigenvectors of a Real Symmetric Matrix

ASL/SX can also be used to compute efficiently frequently encountered problems like the determination of the eigenvalues and the corresponding eigenvectors of matrices:

For the example of a real symmetric matrix, one can use ASL/SX to compute:

- the complete set of eigenvalues and eigenvectors
- only the set of eigenvalues
- some set of eigenvalues or eigenvector
- some set of eigenvalues only

This functionality is implemented in serial processing mode.

We present here performance figures obtained using the so called Frank matrix as an application example. The elements of the Frank matrix are defined as $a(i,j)=n+1-max(i,j)$, where $n$ is the dimension of the problem.

Real Symmetric Matrix

All Eigenvalues And Eigenvectors

| N | TIME(SEC) | MFLOPS |
|---|---|---|
| 100 | 0.01617 | 407.450 |
| 200 | 0.07084 | 717.992 |
| 300 | 0.19466 | 873.873 |
| 400 | 0.39411 | 1027.497 |
| 500 | 0.69711 | 1127.450 |
| 600 | 1.14769 | 1181.081 |
| 700 | 1.72132 | 1246.695 |
| 800 | 2.51599 | 1269.494 |
| 900 | 3.44520 | 1315.374 |
| 1000 | 4.57243 | 1356.293 |

Real Symmetric Matrix

All Eigenvalues

| N | TIME(SEC) | MFLOPS |
|---|---|---|
| 100 | 0.00665 | 317.591 |
| 200 | 0.02855 | 585.984 |
| 300 | 0.07376 | 758.218 |
| 400 | 0.14792 | 916.724 |
| 500 | 0.25303 | 1039.213 |
| 600 | 0.40267 | 1125.987 |
| 700 | 0.59130 | 1210.077 |
| 800 | 0.83590 | 1273.928 |
| 900 | 1.15468 | 1309.008 |
| 1000 | 1.50947 | 1368.623 |

Real Symmetric Matrix

One Eigenvalue And Eigenvector

| N | TIME(SEC) | MFLOPS |
|---|---|---|
| 100 | 0.00466 | 443.544 |
| 200 | 0.01881 | 879.703 |
| 300 | 0.05077 | 1093.494 |
| 400 | 0.10634 | 1268.440 |
| 500 | 0.18766 | 1395.238 |
| 600 | 0.30807 | 1466.510 |
| 700 | 0.46259 | 1542.087 |
| 800 | 0.66820 | 1589.417 |
| 900 | 0.94279 | 1599.432 |
| 1000 | 1.24872 | 1650.901 |

Real Symmetric Matrix

One Eigenvalue

| N | TIME(SEC) | MFLOPS |
|---|---|---|
| 100 | 0.00412 | 494.765 |
| 200 | 0.01766 | 930.556 |
| 300 | 0.04894 | 1129.494 |
| 400 | 0.10378 | 1295.878 |
| 500 | 0.18422 | 1417.973 |
| 600 | 0.30375 | 1484.560 |
| 700 | 0.45730 | 1557.388 |
| 800 | 0.66144 | 1603.423 |
| 900 | 0.93515 | 1610.534 |
| 1000 | 1.23910 | 1661.911 |

## 7. On-line Documentation

The ASL/SX Advanced Scientific Library comes with a complete on-line manual. A major effort has been spent in providing optimal reference and handy help to the programmer. To this end, source code of tested and working program examples are included. These program templates can be copied into the application program in order to exemplify the correct calling sequence for each ASL/SX subroutine. ■

# PARALLEL PERFORMANCE OF ADF ON THE NEW CONVEX SPP-1600

*Norbert Paschedag[1] (Fritz-Haber-Institute, Berlin)*

ADF is the Amsterdam Density Functional Program, an application for calculations on polyatomic systems, including problems in solid state and organometallic chemistry, material science, catalysis etc. It is based on the Kohn-Sham approach to the density functional theory (DFT) and has a number of features that have made it a popular tool in chemistry and physics.

Beginning in 1995, several groups of users installed their own versions of ADF on the CSCS/SCSC platforms; since autumn 1996 CSCS/SCSC, provides a common installation of the code on the scalar machines, including a parallel version for the Convex SPP [1]. The current approach to parallelization in ADF follows the Single Program, Multiple Data (SPMD) paradigm, i.e., several identical copies of the program are started, each being assigned part of the work [2]. This scheme has the advantage of making the maintenance of both serial and parallel versions easier—at the cost, however, of running several full–sized executables at the same time. Also, the static work assignment to each process simplifies the program structure, but may turn out to be non–optimal for environments without sets of dedicated processors.

At CSCS/SCSC, the primary platform for running ADF is clearly the Convex SPP. For this platform an MPI–based parallel version was installed, as opposed to one using PVM, mostly due to its simpler startup procedure. First experiences have shown that ADF scales very well for problem sizes where the overhead that is introduced for data distribution and communication among the processes becomes negligible compared to the time spent by each process in the numerical integration and diagonalization sections. For this problem size the speedup actually is quite good and it can be estimated that ADF would perform well on up to 16 processors (see figure). Comparing the execution times on the old SPP–1000 and the newer SPP–1600, that was installed in December, it can be seen that ADF generally runs about 30% faster, a value that was predicted due to the minor architectural changes.

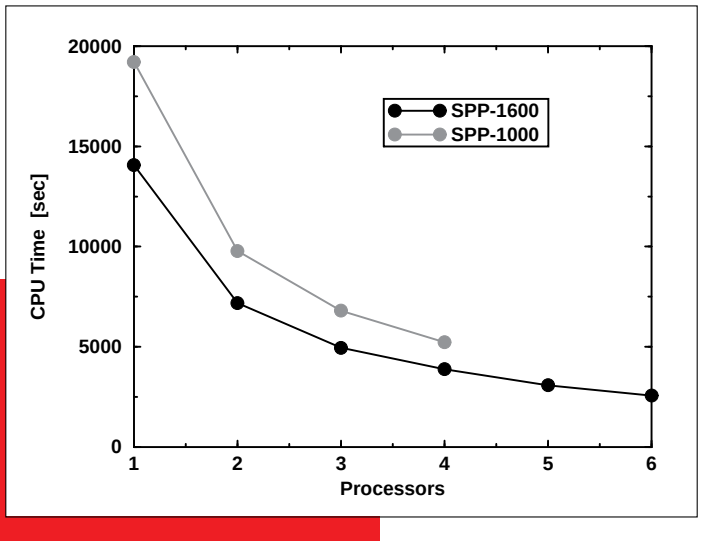


***Figure 1:*** *CPU times for a ferrocene geometry optimization using ADF on the SPP-1000 and SPP-1600.*

## References

[1] http://www.cscs.ch/Official/Services/AppSoft_CCM/adf/

[2] C. Fonseca Guerra, et al., "Parallelisation of the Amsterdam Density Functional Program", Methods and Techniques in Computational Chemistry, ed. E. Clementi and G. Corongiu. Vol. METECC-95. 1995, STEF: Cagliari.

1) At the time of writing, the author was a staff member of CSCS/SCSC.

## NEW APPLICATION SPECIALIST AT CSCS/SCSC

Stefan Portmann is the new CSCS/SCSC specialist for chemistry applications and succeeds Norbert Paschedag. After completion of his studies at the Department of Chemistry of ETH Zurich in 1993, Stefan Portmann began graduate study in the area of macromolecular crystallography with Prof. M. Dobler and Dr. M. Egli. In 1995, Stefan began work at the Northwestern University Medical School (Department of Molecular Pharmacology and Biological Chemistry) in Chicago when Dr. Egli accepted an assistant professorship at the university. In 1996 he submitted his thesis entitled "Kristallstrukturanalysen von RNA-Sekundaerstrukturen und Antisense-Oligonukleotiden" (Structural Analysis of Crystals in RNA-secundary Structures and Antisense Oligonucleotids) to obtain a doctoral degree from ETH Zurich.

With Stefan Portmann, the CSCS/SCSC chemistry group gains a specialist with a background in structural- and biochemistry, an area of research where demand for high-performance computing services has grown rapidly over the past few years.

You can contact Stefan Portmann at portmann@cscs.ch.

# COMPUTING AND COMMUNICATIONS CAMP C³

*Andrea Danani (CSCS/SCSC)*

## Introduction

In 1996, the second annual Computing and Communication Camp (C³) was organized by CSCS/SCSC, together with *Ingenieure für die Schweiz von morgen (INGCH)* and the *Dipartimento Istruzione e Cultura (DIC) del Cantone Ticino*. These camps target Swiss high school students and provide students with the possibility to get a practical approach to computer applications by using them in self-defined projects[1]. After the approval of the jury, two camps each with a duration of two weeks were organized with a total of twenty-three students. After a one-week introduction course on Unix, the Internet and computer applications, the students work an additional week on their projects. At the end of the camps the students write a short report about their work. Financially it was possible to cover the costs of the preparation and realization of the camps with the help of private and public sponsors.

This report summarizes the activities during the camps. For more information about C³ see also http://www.cscs.ch/C3

## Announcement and Competition of C³

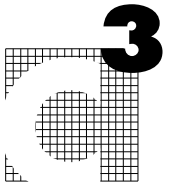

The C³ program was announced in February in all high schools of Switzerland. Along with the announcement a detailed description of the project competition was included which served as selection criteria for the camps. For the competition, the students have to describe a science project they would like to carry through during the camp. The deadline for the competition was mid-May. All these projects were then judged by members of the jury and a decision for acceptance was made mid-June. In general, the quality of the project proposals were quite good and only a few projects were rejected.

## The two C³ sessions of 1996

In summer 1996, two sessions of the C³ Camp took place at CSCS/SCSC. The first (July 22–August 3), had twelve students: five from Swiss-German cantons, six from Ticino and one from the Canton of Geneva. For the second camp (August 5–August 17), the participants totaled eleven: seven from Swiss-German cantons and four from Swiss-French cantons. Because of the big load on supervisors in 1995, we introduced an advisor (or group of advisor) from CSCS/SCSC for every project; these advisors followed the students during the second week. Since there were two projects concerning Artificial Intelligence (AI), researchers of the *Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA)* at Lugano were contacted to give a lecture and to support the students of the AI project.

As previously mentioned, the first week of the camp serves to introduce students to the working environment through a series of courses.

Apart from these courses they also had the chance to tour the CSCS/SCSC infrastructure including the machine room with the NEC SX-3 supercomputer and the technical infrastructure with power supply, battery backup system and water cooling circuit.

After the first few introduction courses, the advisors individually discussed the projects with each student to achieve a project plan for the second week. The seventeen projects covered different subjects like physics, chemistry, biology and informatics. Here is the list of the projects together with the students and their advisors:

Projects: Camp 1 (22/7–3/8)

| Name | Title |
|---|---|
| Manuel Altermatt, Manuel Studer (Kantonsschule Olten (SO)) Advisors : A. Bernasconi, A. Danani (CSCS/SCSC) | Abenteuer Achterbahn |
| Rudi Belotti (Liceo Cantonale Bellinzona (TI)) Advisors : N. Paschedag (CSCS/SCSC) | Rappresentazione grafica delle strutture terziarie e quaternarie della melittina |
| Sébastien Caille (Collège de Saussure (GE)) Advisors : G. Gobbi, E. Gerteisen (CSCS/SCSC) Luca Cereghetti Christian Schuster (Liceo Cantonale Mendrisio (TI)) Advisors: A. Bernasconi, A. Danani (CSCS/SCSC) | Rappresentazione tridimensionale di urti |
| Davide Cescato (Liceo Cantonale Mendrisio (TI)) Advisor: C. Domain (CSCS/SCSC) | Evoluzione di caratteristiche genetiche |
| Tibor Deér (Kantonsschule St Gallen (SG)) Advisor: J. Favre (CSCS/SCSC) | Darstellung dreidimensionaler Umgebung |
| Papik Meli, Alessio Villa (Liceo Cantonale Mendrisio (TI)) Advisor: J. Favre (CSCS/SCSC) | Jump |
| Christian Meyer (Gymnasium Muttenz (BL)) Advisors : N. Paschedag (CSCS/SCSC) | Restriktion eines Plasmids durch ein Enzym |
| Gregor Riepl (Realgymnasium Basel (BS)) Advisor : J. Favre (CSCS/SCSC) | Menschliche Bewegungen in 3D |

Projects: Camp 2 (5/8 - 17/8)

| Name | Title |
|---|---|
| Amaury Boutry, Andres Petralli (Gymnasium Muenchenstein (BL)) Advisors: G. Gobbi , E. Gerteisen (CSCS/SCSC) | Simulation eines Staudammbruches |
| Thomas Bugnon, Claude Fahrni (Lycée cantonal de Porrentruy (JU)) Advisors: M. Hohenadel, P. Pagny (CSCS/SCSC) | La voiture du futur |
| Thomas Burg, Boris Lutz (Realgymnasium Rämibühl (ZH)) Advisors: A. Bernasconi, A. Danani (CSCS/SCSC) | Simulation der Kinematik von Mehrkörpersystemen |
| Yvan Charbonnier (Collège de Saussure (GE)) Advisors: A. Bernasconi, A. Danani (CSCS/SCSC) | Conservation de quantité de mouvement |
| Tiziano Mengotti (Bündner Kantonsschule Chur (GR)) Advisor: S. Pedrazzini (IDSIA) | Experiment künstlicher Intelligenz |
| Monika Meyer, Susanne Reffert (Realgymnasium Rämibühl (ZH)) Advisors: M. Johnson, P. Pagny (CSCS/SCSC) | Fraktale auf dem WWW |
| Yann Thoma (Collège de Saussure (GE)) Advisor: C. Lepori (IDSIA) | Programmes évolutifs |

The second week was fully devoted to work on the projects. Although there were always one or two advisors in the classroom to assist with general environment problems, the students tried to find solutions themselves or helped each other. Some of the students also managed to get solutions, advice or software through the Internet. Experts at CSCS/SCSC were consulted about the project as a last resort.

Besides working in the computer laboratory, DIC's office "Youth and Sport" organized some excursions. A half-day excursion to Alpe Pairolo took place during the first camp on Sunday but was rained out during the second session. On holidays, students worked in the morning but were able to enjoy the afternoon with an introduction to windsurfing on Lugano lake. They enjoyed this initiative very much, especially with the beautiful weather.

Accommodation, breakfast and dinner were organized in the students' hostel of the DIC of the Canton Ticino. A shuttle bus brought the students every day from Lugano to CSCS/SCSC.

## Achievements

The capabilities of the students differed only slightly. Although some of them already had good knowledge in programming, their capabilities concerning the systems at CSCS/SCSC were almost at the same level. This forced them to help each other to solve a problem with the computer or the software and created a good working atmosphere.

Some students were never completely comfortable with the Unix and X-Windows environment, especially those who prepared work at home with some other systems. The mathematical environments Maple and Matlab proved to be very useful for a lot of projects. On the communications side, the connection to the Internet caught most of the attention of the students. They rapidly managed how to search for information, programs, tools and images and although not all the material gathered from the Internet was useful for the project, it helped to acquire experience with the media and its possibilities.

A big accomplishment of the advisors was to lower the level of the projects, which were too ambitious in certain cases, and to simplify the initial objectives the students had set. During the realization of their projects, students were able to appreciate the computing power of the workstations, the

*(Continued on page 20)*

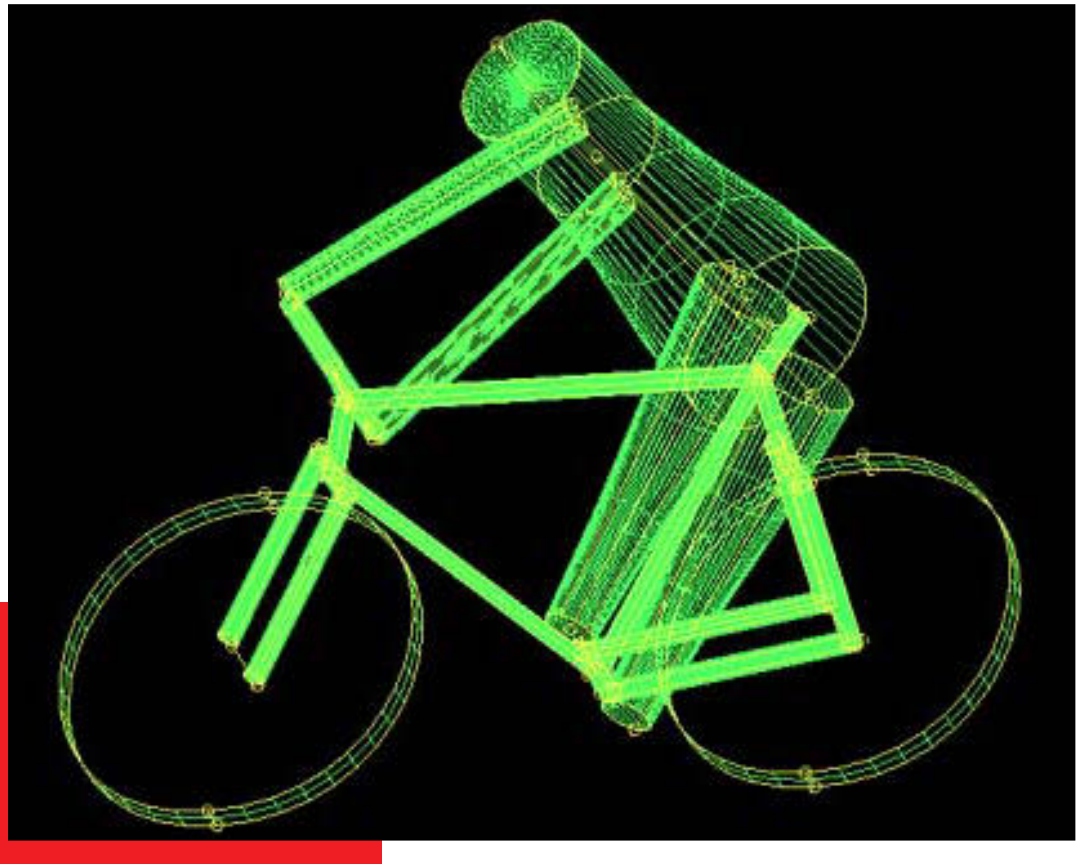

*Aerodynamics of a bicycle, L. Palli, C3 '95*

*(Continued from page 19)*

tools and the huge amount of memory and disk space they had available (compared to a PC) and because of this fact most of them assumed that they would accomplish more in the projects. They nevertheless soon realized what kind of problems can emerge during a numerical simulation and they got a fairly realistic impression of what is still needed to solve "real" problems on a computer.

This aspect of the camps can probably be improved by providing the students with more differentiated and adapted tools, especially for what concerns the visualization, because powerful tools (such as AVS) have a long learning curve. As already said, the introduction time of computers and software should also be shortened in favor of more time for the projects since it does not seem possible to conceive sessions with a duration of three weeks.

## Future Plans

We intend to continue the C³ initiative in the years to come. Funding is a problem and for this reason, we approached the departments of the public education of each canton for subsidies of 600.- CHF for each of their students accepted to the C³. For 1996, sixteen cantons agreed to fund their students. Nevertheless, in the future, we hope to find other financial supports which could be as stable as possible, at least for some years.

While the feedback from the students was very positive, there are some things which can be improved in the future. In particular, a more personalized introductory part could be foreseen, even if this could require more organizatorial efforts. Indeed, many students agreed on the fact that the work on the project could already begin in the first week, because the second week seems to be too restricted to allow a complete development of the project. We could therefore think to organize parallel sessions of lectures during the first three days, where the students follow the lectures with the topics tightly related to their projects.

## Acknowledgments

The C³ would not have been such a success without the voluntary assistance of many people. Thanks go to CSCS/SCSC and IDSIA staff personnel who were willing to give courses during the camps, to Marina de Senarclens (INGCH) , Giambattista Ravano (DIC) and Bernardo Zumthor (CSCS/SCSC), who gave a valuable help in the organization and public relations activities, and to all other people at CSCS/SCSC who helped the students in their stay. Special thanks go to the sponsors ABB Asea Brown Boveri, Cantone Ticino, DEC Digital Equipment, ETH Zürich, KGF (Ciba-Geigy, Hoffmann-La Roche, Lonza and Sandoz), Schweizerische Mobiliar, StorageTek, as well as to the Swiss Cantons who accepted to support their students who participated in the C³ Camp.

1) Student selection is based on project evaluation by a jury.

# THE PRSS STUDENTS MOBILITY PROGRAM: 1997 SESSION

With the 'Project-Related Student Stages' (PRSS) mobility program, already carried out in 1995 and 1996, CSCS/SCSC has an internship program which gives young researchers and students from Swiss Academia the opportunity to relocate to Manno (Ticino) and work for up to three months at the Center. As these internships should offer the possibility to closely interact with the Center's specialists and take advantage of the scientific teams, presentations given by CSCS/SCSC staff members and invited speakers constitute the basis of the program.

These presentations cover file and archive serving, queue structures (NQS, LSF), the programming environments of scalar scalable shared-distributed memory (HP Cluster and Convex Exemplar) and parallel-vector shared memory (NEC SX-4) compute platforms as well as the post-processing and visualization of computed results. Personalized programs for each student are made on the basis of these topics and will help the intern in gaining the necessary skills in the field of high-performance computing. Furthermore, the students will be supported throughout their internship by a scientific consultant providing the appropriate scientific support.

If you are interested in participating, please send your application (signed by your professor) detailing the goals of your project to:

Dr. D. Maric
CSCS/SCSC
Via Cantonale
6928 Manno

# Professional and Technical Courses: Courses for the Scuola Tecnica Superiore (STS)

*Andrea Bernasconi (CSCS/SCSC)*

The *Scuola Tecnica Superiore (STS)* in Manno consists of two sections: one for computer science and another for electronics. In 1997 the STS will become a *Scuola Universitaria Professionale (SUP)* with an explicit research mandate.

The professional and technical courses offered by CSCS/SCSC to the STS are highly appropriate in the context of the CSCS/SCSC goals. Briefly, the first CSCS/SCSC goal is to provide a high level support in the field of supercomputer applications to the swiss Academy and if possible to industry as well. The second goal is the creation of synergies with the Italian part of Switzerland and eventually with the north of Italy.

A course on vector and parallel programming originally was requested by Dr. Carlo Spinedi, vice-director of the STS. The main reason for this request was a particular interest in the unique possibility offered by CSCS/SCSC to the STS students in the field of high performance computing.

## Description

The audience is represented by the third year computer science students of the STS. From the CSCS/SCSC side the following personnel are involved: Andrea Bernasconi (responsible), Mauro Ballabio and Matteo Boverat.

To some extent the course is based on the *Cycle d'études postgrades le parallélisme* given at the EPFL in the years 1990–1991. Basically the course presents an introduction to the elements of vector and parallel programming and in particular to scientific applications, supercomputer architectures, and programming methods. It is divided into six modules plus four exercises sessions as presented by the Table 1. All course material and instruction are in Italian.

The course editions have been: November 1994, November-December 1995 and September–November 1996. In 1996 an examination was performed. The notes were integrated into a corresponding STS course. Concerning requested semester and diploma works, the firsts proposals have been submitted.

On the other hand, part of the course was presented to the students that followed the projects PRSS (1995) and $C^3$ (1995, 1996) and also at the Physics Department of the Milano University (1996).

The start-up of the course was time consuming. For 1997 we foresee four man months of work including the completition of the documents corresponding to the six modules mentioned before. The required time for the full course preparation will decrease in 1998.

## Contents

The table Modules gives more details on the course contents. The first module explains why the main scientific fields need huge simulation programs pushing the development of supercomputers. The basic principles of vector computers are presented, in particular the pipelining techniques. These principles are illustrated by the SX architectures.

Modules 2 and 3 give a general overview of the Fortran 77 language and of vector programming techniques. The differences between scalar and vector instructions are explained as well as instructions dependencies, optimization using global and local directives, and tools for analyzing program executions.

Module 4 presents the general characteristics of parallel computers and the networks used to interconnect the processing elements.

Module 5 is mainly based on the problems concerning communication between processors. After a general introduction to multiprogramming theory, types and modes of communications are presented. Message routing and communications costs are also extensively analyzed. At the end, concepts and examples of macro-communications are given, as well as an overview of the existing communications and parallel linear algebra libraries.

The scope of Module 6 is an introduction of some programming models in the following way. After the description of a possible work method in the parallelism field, the three goals (go faster, divide and conquer, and create new concepts) and the three sources of parallelism (data, control, and flow) are presented. This presentation of parallelism probably is in contradiction to others views. Finally, some aspects of the parallelism are illustrated using short theoretical explanations and practical examples.

## Exercises

The exercises on vector programming gradually familiarize the programmer with key points such as memory access, data dependences, and optimal usage of vector registers and pipelines. These topics are discussed and analyzed by coding in Fortran77 basic algorithms like sum, dot product, and matrix operations.

The exercises on parallel programming consist of the development of simple algorithms for solving problems such as the one-dimensional diffusion equation. The algorithms are implemented in Fortran or C and optimized for vector processing. The parallelization is done first automatically, then using the Message Passing

*(Continued on page 22)*

*(Continued from page 21)*

Interface (MPI) and the results are compared. When MPI is used, MIMD-SM and MIMD-DM implementations are presented.

## Conclusion

We hope the course will continue to be of interest to the SUP and that CSCS/SCSC will reinforce the professional and technical courses activity in the canton Ticino. Clearly we are ready to modify and expand the course according to the SUP and CSCS/SCSC requests.

***Table 1**: 1996 Course Agenda*

Vector programming

| Speaker | Subject |
|---|---|
| Andrea Bernasconi (AB) | Scientific applications and vector architectures |
| Matteo Boverat (MBo) | Introduction to Fortran and vectorization |
| Mauro Ballabio (MBa) | Elements of vector compilation and optimization |
| AB, MBo e MBa | Exercises on NEC SX-4 |
| AB, MBo e MBa | Exercises on NEC SX-4 |

Parallel programming

| Speaker | Subject |
|---|---|
| Matteo Boverat (MBo) | Parallel architectures and interconnection networks |
| Mauro Ballabio (MBa) | The communications |
| Andrea Bernasconi (AB) | Models of parallel programming |
| AB, MBo e MBa | Models of parallel programming. Exercises on NEC SX-4 |
| AAB, MBo e MBa | Exercises on NEC SX-4 |

***Table 2:** Modules**

*The table of contents are not translated in english because all course material is in italian.

M1. Scientific application and vector architectures

1 Applicazioni scientifiche
- 1.1 Esempio illustrativo della propagazione del calore
- 1.2 Programmi di simulazione in climatologia

2 Architetture vettoriali
- 2.1 Il principio del pipelining
- 2.2 Esempio della pipeline aritmetica d'addizione
- 2.3 Esempio della pipeline per le istruzioni
- 2.4 Il principio dell'unità vettoriale
- 2.5 La tecnica del Gather/Scatter
- 2.6 La memoria interlacciata
- 2.7 L'architettura del superordinatore NEC SX-3/24R
- 2.8 Osservazione sulla struttura gerarchica della memoria

M2. Introduction to Fortran and vectorization

1 Introduzione al Fortran77
2 La vettorizzazione
- 2.1 Istruzioni scalari e vettoriali
- 2.2 Barriere alla vettorizzazione
- 2.3 Il compilatore vettoriale

M3. Elements of vector compilation and optimization

1 Elementi di compilazione vettoriale
- 1.1 Vector Mask
- 1.2 Le macrooperazioni
- 1.3 Gather e scatter di vettori

2 L' ottimizzazione
- 2.1 Direttive globali/locali
- 2.2 L' ottimizzazione manuale
- 2.3 Gli analizzatori

M4. Parallel architectures and interconnection networks

1 Computers paralleli
- 1.1 Introduzione
- 1.2 Passato, presente e futuro
- 1.3 Classificazione di Flynn
- 1.4 Organizzazione della memoria
- 1.5 Le strade del parallelismo
- 1.6 Speedup ed efficienza
- 1.7 Esempi d'architetture parallele

2 Reticoli d'interconnessione
- 2.1 Evoluzione
- 2.2 Reticoli SM
- 2.3 Reticoli DM
- 2.4 Comunicazione nei reticoli DM
- 2.5 Differenze tra SM e DM

M5. The communications

1 Le comunicazioni
- 1.1 Alcune nozioni di multiprogrammazione
- 1.2 Le comunicazioni punto a punto
- 1.3 L' invio e la ricezione
- 1.4 Communication link e assenza di tempo globale
- 1.5 I modi sincrono e asincrono
- 1.6 I modi bloccanti e non-bloccanti
- 1.7 Comunicazioni tramite interruzioni

2 Il routing
3 Un modello di costo per comunicazioni
4 Le macro-comunicazioni
- 4.1 Le primitive
- 4.2 La sincronizzazione
- 4.3 Il trasferimento e la diffusione nell' anello
- 4.4 Il trasferimento e la diffusione nel toro
- 4.5 Il total-exchange nell' ipercubo

5 Alcune librerie di comunicazione
- 5.1 librerie di algebra lineare
- 5.2 librerie di comunicazione

M6. Models of parallel programming

Premessa
Introduzione

1 Il metodo di lavoro IDEAO
2 I tre scopi del parallelismo
- 2.1 Andare più veloce
  - 2.1.1 Accelerazione sovralineare
- 2.2 Dividere per regnare
  - 2.2.1 Il problema dello zaino
- 2.3 Creare nuovi concetti
  - 2.3.1 Il modello Gamma

3 Le tre sorgenti del parallelismo
- 3.1 Parallelismo dei dati
- 3.2 Parallelismo di flusso
- 3.3 Parallelismo di controllo
- 3.4 Le sorgenti del parallelismo mediante un esempio
- 3.5 Una classificazione delle architetture parallele

4 Aspetti del parallelismo
- 4.1 Diffusione del calore col metodo delle differenze finite
- 4.2 Moltiplicazione di matrici col metodo sistolico
- 4.3 Risoluzione del TSP col metodo B&B
- 4.4 Risoluzione del TSP col metodo neuronale

http://www.cscs.ch/Official/Services/STSCourses.html

# A User Project: Phase Stability and Thermal Properties of FeSi-systems

*E.G. Moroni (Technical University of Vienna, Austria) and T. Jarlborg (DPMC, University of Geneva)*

The phase stability and structural properties of FeSi systems are studied from first principle density functional calculations. By investigating the electronic and structural properties of different types of crystal structures we explain the competition between pseudomorphic and bulk stable phases. The electronic structure of the ground state phase as a function of temperature T and applied magnetic field H is also discussed. As thermal disorder is introduced, the small band gap (6 mRy) is filled and the Stoner enhancement is activated. The combined effect of Fermi-Dirac smearing and thermal disorder explains the large magnetic susceptibility at large T.

Electronic structure calculations for solids consist of solving the Schrödinger equation for periodic structures with several atoms per unit cell. The quantum mechanical results of band energies and wavefunctions are related to physical properties by subsequent calculations. This type of calculation is a useful tool in material science as a mean of understanding physical mechanisms, and sometimes it can be used for predictions or guidance for obtaining improved properties. At present it is mainly used in basic research, but in the future, with the rapid development of methods and computers, it might be used routinely in applied material research. In this work we present calculations for some FeSi-compounds. Since these structures are fairly simple with few atoms per cell, and our methods are rather efficient, it has been possible to do most calculations on work stations. But thanks to the possibilities offered by CSCS it is more convenient to do interface structures with the pseudo potential on a supercomputer.

The FeSi system can be made in several phases and each one has quite different electronic properties. The stable bulk phase of FeSi is the cubic *B*20 crystal structure (ε-FeSi), while novel metastable, pseudomorphic structures (like FeSi *B*2 and $FeSi_2$ B*2*-defect structures) are found in epitaxial films. The fact that some of these structures have a small lattice mismatch with Si, makes the FeSi systems interesting for application as metal-semiconductor interfaces. In contrast to ε-FeSi and orthorhombic *β*-FeSi2 which are semiconducting, FeSi films epitaxially stabilised on Si(111) substrate exhibit metallic conductivity as the vast majority of transition metal silicides.

Resistivity, optical properties as well as band theory, show that bulk ε-FeSi has a very narrow band gap in the middle of the Fe-*d* band of 4–8 mRy [1]-[4]. Non-stoichiometric compositions are typical, making it easy to have the material doped. At very low temperature most properties are well understood from the band results, for instance the thermopower [5] is very well reproduced by Bolzmann theory [4]. A long standing problem has been the magnetic susceptibility χ(T), which raises very quickly with temperature. Jaccarino et al [1] concluded that unreasonably high peaks in the density-of-states (DOS) on both sides of the gap are required to explain the amplitude of χ. If optical data agree well with LDA bands and shows the gap at T=0, they show also that the gap disappears already at about 200 K [6]. On the other hand it was also proposed that spin-fluctuations are responsible for the large magnetic response at large T [7].

In this work our objective is to understand the electronic properties and some features of the phase diagram of FeSi systems, and we use *ab-initio* methods. For the band calculations and formation energy of FeSi we use the self-consistent, semi-relativistic, all-electron (AE), Linear Muffin-Tin Orbital (LMTO) method [8], in which T-dependent electronic excitations are included via the Fermi-Dirac function [9], as well as plane wave methods [9] using ultrasoft pseudopotential (US PP) method [10]. By using the plane wave method, relaxed geometries are optimized by computing forces acting on atoms using the Hellmann-Feynmann theorem. Both the local density approximation (LDA) and the generalised gradient approximation (GGA) [12] are used for the electronic exchange-correlation energy.

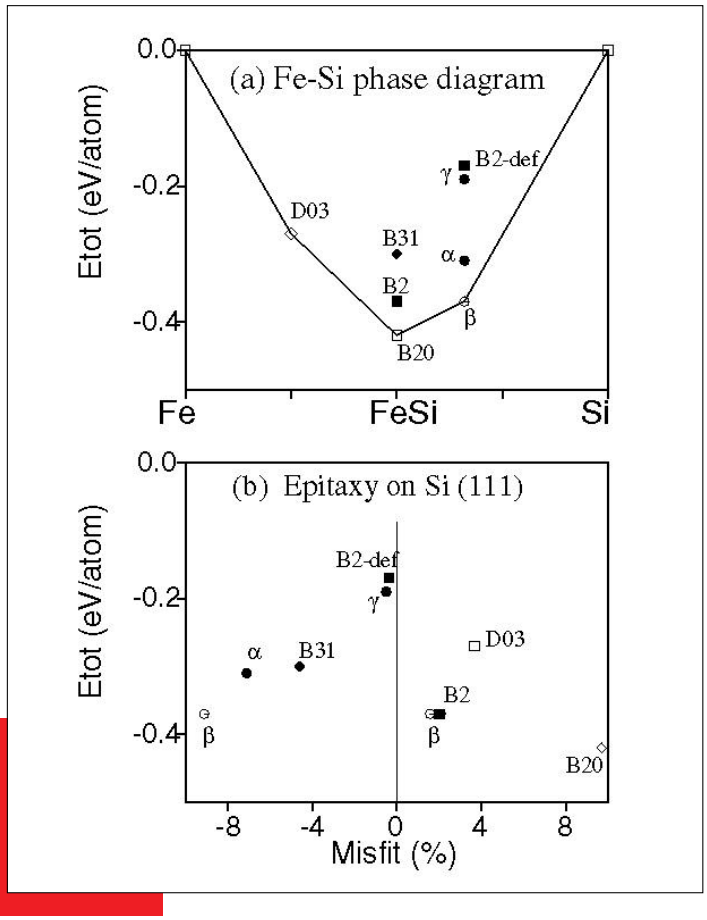


***Figure 1:*** *a) Formation energy (eV/atom) of several Fe-Si intermetallic compounds. b) Formation energy of Fe-Si systems versus lattice misfit with respect to Si interatomic distance.*

*(Continued on page 24)*

*(Continued from page 23)*

The LMTO calculation for ε-FeSi at T=0 gives a band structure and a gap which is very similar to what have been published earlier by use of different methods [2–4]. The LDA LMTO calculated lattice constant, $a_0$=4.43 Å is found at the minimum of the total energy $E_{tot}$. This value is shifted to $a_0$=4.47 Å, when using US PP and GGA potential. This compares better with the experimental value of 4.49 Å. The computed bulk modulus (B) at $a_0$ is calculated to be around 2.2 Mbar. The gap $E_g$ is indirect, ranging from 6.1 mRy at a lattice constant of 0.98 $a_0$ to 5.7 mRy at 1.02 $a_0$. In the *B*20 structure the Fe and Si atoms have the positions ($u$,$u$,$u$), (0.5+$u$,0.5-$u$,-$u$), (-$u$,0.5+$u$,0.5-$u$) and (0.5-$u$,-$u$,0.5+$u$). When u(Fe)= 0.150 and u(Si)= 0.850 each Si (Fe) atoms has seven Fe (Si) nearest neighbors. The calculated relaxed atom-positions are u(Fe)=0.1358 and u(Si)=0.8408 and they compare very well with the experimental values of u(Fe)=0.1358 and u(Si)=0.844.

For FeSi the stability of the *B*20 structure (FeSi type) has been studied with respect to several other close-packed structure types as for instance the *B*2 (CsCl type), the *B*31 (MnP type), the *B*35 (NiAs type) and *B*1 (NaCl type) structures. The *B*20 structure can be seen as a distortion of the *B*31, *B*35 or *B*1 structure, when one considers the plane perpendicular to the trigonal axis, and all show nearly similar density of states (DOS) with characteristic pseudogaps. The *B*20 structure is found to be the most stable, however the energy difference with respect to the *B*31 and the *B*2 is quite small, below 0.5 eV per formula unit, so that the stability of the *B*20 structure might be fragile at large pressure and temperature. The energy minima of the *B*35 and *B*1 are 1–2 eV above the ground-state. The *B*31 becomes more stable than the B20 structure for increased *3d*-valence electron composition systems like NiSi compound. FeSi can be stabilized in the *B*2 structure by applying high pressure (epitaxial FeSi films) or by substitution of Fe atoms with Ru.

The difference in local coordination of the *B*20 and the *B*2 structure is strongly affecting the DOS but the total energies are still close. The *B*2 structure has a DOS that is more bcc-like, dominated by a 2 eV narrow peak of the metal d bands while the DOS of FeSi *B*20 is more fcc-like and without a dominating peak. Using US PP and GGA approximation the difference in total energy between the *B*2 and the *B*20 structure ($\Delta E_{B2\text{-}B20}$) is of ≈0.1 eV. This value shifts to ≈0.6 eV when using LMTO with LDA. The highest coordination (8 nearest-neighbors atoms of the other kind) of the *B*2 structure with respect to the *B*20 and *B*31 explains the findings that the *B*2 structure displays a much lower volume per atom with respect to the other two structures, and we found that $V_{B2} < V_{B20} < V_{B31}$. The present study predicts that the phase transformation of FeSi from the *B*20 to the *B*2 structure is accompanied by an isotropic compression, from an equilibrium volume per f.u. of $V_{B20}$= 22.25 Å$^3$ to $V_{B2}$ = 21.26 Å$^3$. This transformation is understood in terms of competitions between covalent and metallic bonding. The predicted transition pressure is quite imprecise (in the range of 200 to 900 kbar) because of the uncertainty of the relative energy difference.

As function of T, one expects that metallic phases should become relatively more stable compared to structures with low DOS at $E_F$, as has been demonstrated for several transition metals [13]. This is because of electronic entropy which is nearly proportional to the DOS and gives a negative contribution to the free energy. Thus, the difference between the *B*2 and *B*20 energy minima should diminish with T. However, as is discussed below, thermal disorder tend to make even the *B*20 phase metallic from about 200K and above.

Several pseudomorphic structures appearing at the interface to silicon can be interpreted, from a geometrical point of view, as CsCl (*B*2) based lattices with different arrangements of about 50% vacancies on the iron sublattice. One of these defect structures is the cubic fluorite structure (γ-phase) which is unstable in bulk $FeSi_2$ at all temperatures. This structure shows larger volume per f.u. and smaller bulk modulus than both the ground-state that is the $\beta$-$FeSi_2$ phase (orthorhombic centered cell with 24 atoms, a= 9.87, b=7.80, c=7.84 Å) and α-FeSi2 phase (tetragonal cell with 3 atoms, a= 2.70 Å and c= 5.15 Å). In fact, the *B*2-based defect structures with an ordered arrangement of vacancies corresponding to the γ- and α-phase ($c/a$=2.0, $z$=0.25) that have been studied exhibit lattice parameters of 5.40 Å and 5.41 Å, respectively, close to the experimental value of 5.43 Å for silicon while FeSi in the B2 structure has a lattice constant $a_0$=2.77 Å, which exceeds half the Si lattice constant by ~2%. The computed formation energy of several Fe-Si compounds are shown in Fig. 1 versus composition but also as a function of the lattice misfit $f$=($b$-$a$)/$a$ with respect to Si interatomic distance, when one considers the Fe-Si/Si interface. In agreement with experiments [14], these results indicate that the *B*2-defect structures and the γ-$FeSi_2$ have a very small lattice mismatch with Si and are therefore suitable for epitaxial growth on Si(111). On the other hand, *B*2 FeSi or the stable bulk phases ε-FeSi and $\beta$-$FeSi_2$ may stabilize on Si only under different compressive strain.

Since the mismatch of the lattices at the FeSi-Si interface is small, it is sufficient to stabilize *B*2-based defect structures that are not stable in the bulk. With increasing distance from the silicon substrate, where the lattice alignment no longer imposes large strains, these cubic phases subsequently transform to energetically more favorable structures either by disordering the vacancies towards the defected-$FeSi_2$ or by relaxing the $c/a$ and internal coordinates $z$ towards the α-phase. Finally, metal to semiconductor phase transformations may occur to recover ε-FeSi or the $\beta$-$FeSi_2$ structure in very thick films.

Spin-polarized calculations including the applied magnetic field, $H$, give the Stoner enhancement of the susceptibility. When thermal disorder and Fermi-Dirac occupations are taken into account in band calculations, they will effectively smear out the sharpness of the DOS around the gap in ε-FeSi, and the narrow-gap semi-conductor becomes gradually metallic as T is increased. The effective

DOS around $E_F$ leads to a Stoner enhancement that will substitute the role of the 'huge' DOS peaks around the gap. This process was partly demonstrated by including only the Fermi-Dirac function and with a parametrized DOS broadening [4] in self-consistent band calculations. Here we use essentially a parameter free band description of $\chi(T)$. The exchange (or Stoner) enhancement S(T,H) is defined as the ratio between the calculated exchange splitting of the (mostly Fe-d) electrons at the Fermi energy ($E_F$), and the applied field energy (1mRyd corresponds to 230 Tesla).

Thermal disorder is introduced by using a Gaussian distribution function as for harmonic vibrations. The mean displacement $u$ is related to the force constant $K$ ($K=M\ \omega^2$ where $M$ is the atomic mass and $\omega$ the phonon frequency) and the temperature T. The force constant is estimated from the calculated bulk modulus to about 10 eV/Å$^2$. Each atomic position is given a randomized Gaussian displacement $u$, the averaged $u$ is calculated from all positions in the normal 8 atom unit cell, and T is defined from the equation $u^2 = 9\ k_B\ T / K$. In addition to the thermal disorder, the self-consistent calculations are done with the electronic excitations as given by the Fermi-Dirac distribution. Figure 2 shows the paramagnetic DOS with and without a certain disorder corresponding to about 440 K. Several disordered configurations were used in independent band calculations. In all cases one observes a closing of the gap as the disorder (or T) increases. Lower degree of disorder, corresponding to T smaller than about 250 K, makes the gap narrow. This is consistent with the observation from infrared and optical measurements that the gap disappears for T larger than 200 K [6]. The metallization is therefore rather quick and when a magnetic field is added in spin polarized calculations, it leads to a Stoner exchange enhancement S(T).

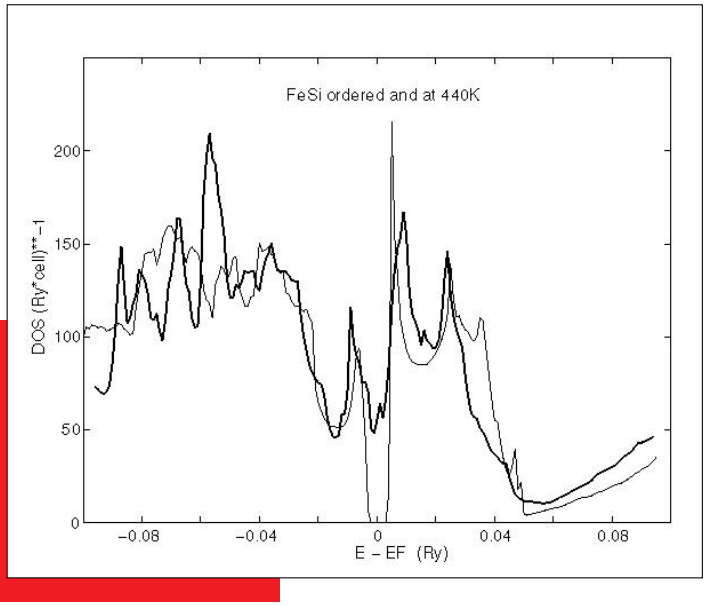


***Figure 2:*** *Paramagnetic density-of-states of e-FeSi in the B20 structure without disorder (thin line) and with disorder corresponding to a temperature of about 440K (heavy line). The Fermi energy is at zero energy. Units in states per cell and Rydbergs.*

The magnetic susceptibility $\chi$, written as $\mu_b^2\ N(E_F,T)\ S(T)$, increases from zero at 0 K to about 4.5 $\mu_B^2 / eV$ at 400 K, when only electronic excitations are taken into account [4]. The Stoner enhancement is about 1.8 at this temperature. But when the band disorder is considered, and S(T) is an average (~4) over the enhancements of several Fe sites, $\chi$ increases further, to about 25 $\mu_B^2 / eV$ at the same T. This is close to estimations from experiment [1, 3]. Hovever, the local S(T) values are much spread among different Fe sites. This can be understood from the spread of local disorder: At some sites the disorder is strong and leads to an attractive Madelung shift of the potential. The Fermi level tends to enter into the high DOS above the (normal) gap in such cases, while on other sites $E_F$ will be at or below the dip in the DOS. From this one can understand why the Stoner enhancement, which is sensitive to the DOS, show local variations from site to site. This is in agreement with the picture of spin-fluctuations [7].

The obtained results suggest that the unusual susceptibility of ε-FeSi has its origin in a rather simple mechanism, that of thermal disorder. This is in contrast to various approaches which propose that unusually strong electron-electron correlation is important in FeSi [2, 3]. Thermal disorder is present in all systems, but FeSi is an ideal case for seeing its effect because of the sharp DOS features very near to $E_F$.

This work has been supported by the Swiss National Science Foundation under grant No: 8220-042840 (E.G.M). E.G.M. acknowledges various discussions with Walter Wolf, Raimund Podloucky, Georg Kresse and Jurgen Hafner.

[1] V. Jaccarino, G. K. Wertheim, J. H. Wernick, L. R. Walker and S. Arajs Phys. Rev. 160, 476 (1967)

[2] C. Fu, M. P. Krijn, and S. Doniach Phys. Rev. B49, 2219 (1994)

[3] V .I. Anisimov, S. Yu Ezhov, I. S. Elfimov, I. V. Solovyev and T. M. Rice, Phys. Rev. Lett. 76, 1735 (1996).

[4] T. Jarlborg, Phys. Rev. B51, 11106 (1995); Phys. Rev. Lett. 77, 3693 (1996).

[5] B. C. Sales, E. C. Jones, B. C. Chakoumakos, J. A. Fernandez-Baca, H. E. Harmon, J. W. Sharp and E. H. Volckmann, Phys. Rev. B50, 8207 (1994).

[6] Z. Schlesinger, Z. Fisk, H-T Zhang, M. B. Maple, J. F. DiTusa and G. Aeppli, Phys. Rev. Lett. 71, 1748 (1993)

[7] Y. Takahashi and T. Moriya, J. Phys. Soc. Jap. 46, 1451 (1979); S. N. Evangelou and D.M. Edvards, J. Phys. C16, 2121 (1983)

[8] O. K. Andersen, Phys. Rev. 12, 3060 (1975); T. Jarlborg and G. Arbman, J. Phys. F7, 1635 (1977)

[9] T. Jarlborg and P. Lerch, Phys. Lett. A176, 130 (1993)

[10] D. Vanderbilt, Phys. Rev. B 41, 7892 (1990); G. Kresse and J. Hafner, J. Phys. Condensed Matter 6, 8245 (1994).

[11] G. Kresse and J. Furthmüller, Comput. Mat. Science 6, 15 (1996); G. Kresse and J. Hafner, Phys. Rev. B 47, 558 (1993).

[12] J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, C. Fiolhais, Phys .Rev. B 46, 6671 (1992)

[13] E. G. Moroni, G. Grimvall and T. Jarlborg, Phys. Rev. Lett. 76, 2758 (1996)

[14] H. von Känel et al, Phys. Rev. 50, 3570 (1994).

# A User Project: System Architecture for Parallel Particle Simulation in Real-Time: Model Calculation and Visualization in Molecular Dynamics

*Erich Stoll (Physics Institute, University of Zurich-Irchel), Christian Stern and Peter Stucki (Computer Science Department, University of Zurich)*

Particle-Simulation is used for the dynamic representation of a wide range of natural phenomena. The aim of the project is to specify and implement a system architecture for parallel Particle-Motion Simulation in real-time. This new simulation technique is used for model calculation and visualization of selected models in Molecular Dynamics. A large number of objects (> 1000) can only be simulated using special algorithms. Here, this is done with Parallel Distributed Computing on a cluster of workstations and by interactive visualization of the results in real-time.

## Introduction

Molecular Dynamics (MD) is considered an important instrument for a wide range of research areas in science. Major potential for such a system exists in research areas such as chemistry and physics. A user-friendly man-machine interface and the object-oriented software technology permits reuse and expansion of individual system-components.

## A Simulation Framework for Molecular Dynamics

In order to use modern software engineering principles, the system architecture of the Interactive Molecular Dynamics System (IMDS) application has been designed using the paradigm of object-oriented programming (see Figure 1). To avoid performance bottlenecks caused by interactive visualization, the Silicon Graphics IRIS-4D-Superworkstation-family was chosen as hardware platform. Computationally-intensive tasks may be performed using Parallel Distributed Computing (PDC). The particles and their interactions are visualized using different rendering techniques (lighted, wireframe, pixels). The graphical representation of the simulation process can be rotated and scaled interactively.

## Object-oriented Implementation of the Lennard-Jones Model in C++

The application is divided in three main modules: model calculation, visualization and user interaction. The inter-particle forces for real systems (e.g., rare gas fluids), are derived from the Lennard-Jones (LJ)-potential [1, 2]

$$V(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right].$$

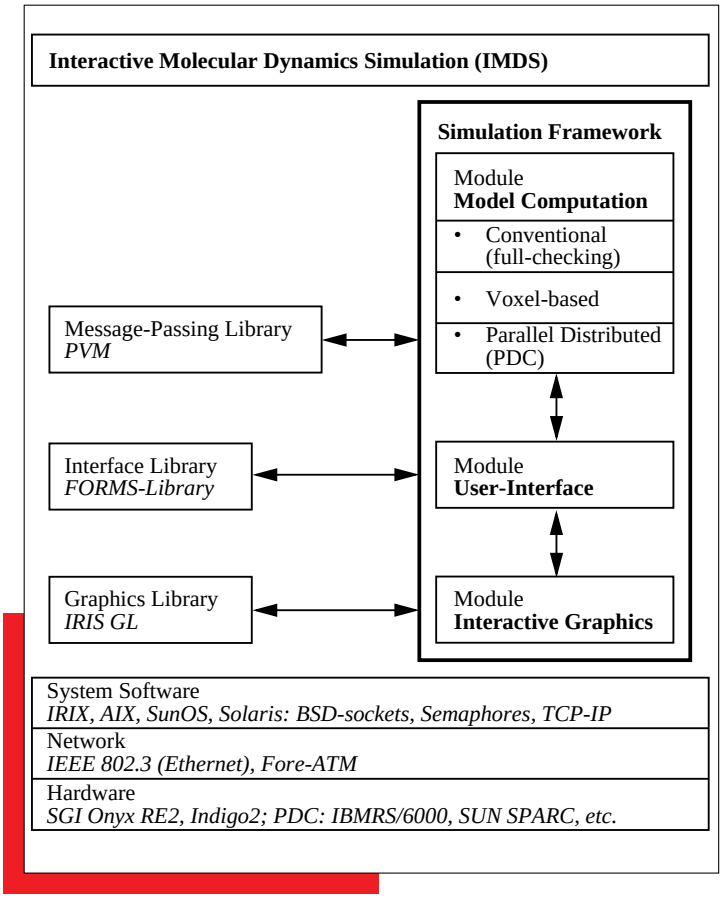


***Figure 1:*** *System architecture of the IMDS application.*

## Acceleration and Speedup Techniques

Due to the fact that MD represents an N-Body-problem, special algorithms and PDC-techniques have been considered. Performance improvements are achieved with a spatial domain decomposition algorithm [1]. The simulation box is divided spatially into cubes (voxels) of the same size. Because the interactions of a certain particle $i$ with the other particles are restricted to a given range $R$ not larger than the voxel extension, the precision of the calculation is not hampered. This algorithm is ideal for parallelization, because the spatial domain-decomposition implies data-locality. For the parallelized version of the algorithm, one processor only needs to know the particle-parameters (positions in space $[x,y,z]$) of a local region, causing a reduced process-to-process data communication overhead.

The runtime complexity of the voxel-algorithm is linear, if the number of particles per voxel remains constant. The parallel version of the IMDS application is based on top of PVM [3]. At CSCS, the computation-intensive parts of the IMDS

application were ported to various Massively Parallel Processor (MPP) systems.

## Results

### Performance

The interactive visualization was performed on an SGI Crimson VGX, R4000 connected to the compute servers via a Fiber Data Distributed Interface (FDDI) network infrastructure. When simulating 1000 atoms in a box of $(6\sigma)^3$, the voxel-algorithm results in a speedup of 1.5 to 2 (64 voxels).

Using eight CPUs on the CONVEX-EXEMPLAR, a speedup factor of two was achieved (25% efficiency). The efficiency is as low, because there were other computation-intensive jobs running on these CPUs. In addition, a redesign of the IMDS-application would be necessary in order to use the EXEMPLAR's architecture (shared memory) optimally. The speedup achieved on the HP-cluster was 4 on 8 CPUs (50% efficiency). Higher speedups could be achieved when using the native C++ compiler (which is not installed at CSCS) instead of the GNU C++ compiler (g++).

### Applications

We have generated isotropic random 3D fractal systems, existing in Silica-Aerogels [4], using the technique of simulated annealing [5] and frozen LJ fluids (see Figure 2). From all nearest neighbor bonds shorter than 1.15σ, a subset is selected randomly with the probability $p \approx p_c$ [6] to avoid an anomalous multifractal spectrum as it was reported recently [7]. The box-counting method and the static-structure factor allowed us to determine the 3D-fractal dimension *D* of 2.4 [6].


### Acknowledgement

We are grateful to the *Schweizerischer Nationalfonds zur Förderung der wissenschaftlichen Forschung, Schwerpunktprogramm Informatik* for financial support.


## References


[1] M. P. Allen, D. J. Tildesley, "Computer Simulation of Liquids", Oxford Science Publications, 1987.

[2] L. Verlet, Phys. Rev. bf 159, 98–103 (1967).

[3] A. Geist et al., "PVM: Parallel Virtual Machine–A User's Guide and Tutorial for Networked Parallel Computing", The MIT Press, Cambridge Massachusetts (USA), 1994.

[4] J. Fricke, Phys. Bl. 51, 935–939 (1996).

[5] Christian Stern, Peter Stucki, and Erich Stoll, "Priority Programme Information Research", Proc. Information Conf. Module 3 Mass. parall. syst., eds. K. Bauknecht and J. M. Grossenbacher, (Swiss Nat. Sci. Fond., Bern, 1994), pg. 29.

[6] E. P. Stoll Ch. Stern P. Stucki, Physica A, 230, 11–18 (1996).

[7] J. H. J. van Opheusden, M. T. A. Bos, and G. van der Kaaden, Physica A 22, 183–196 (1996).


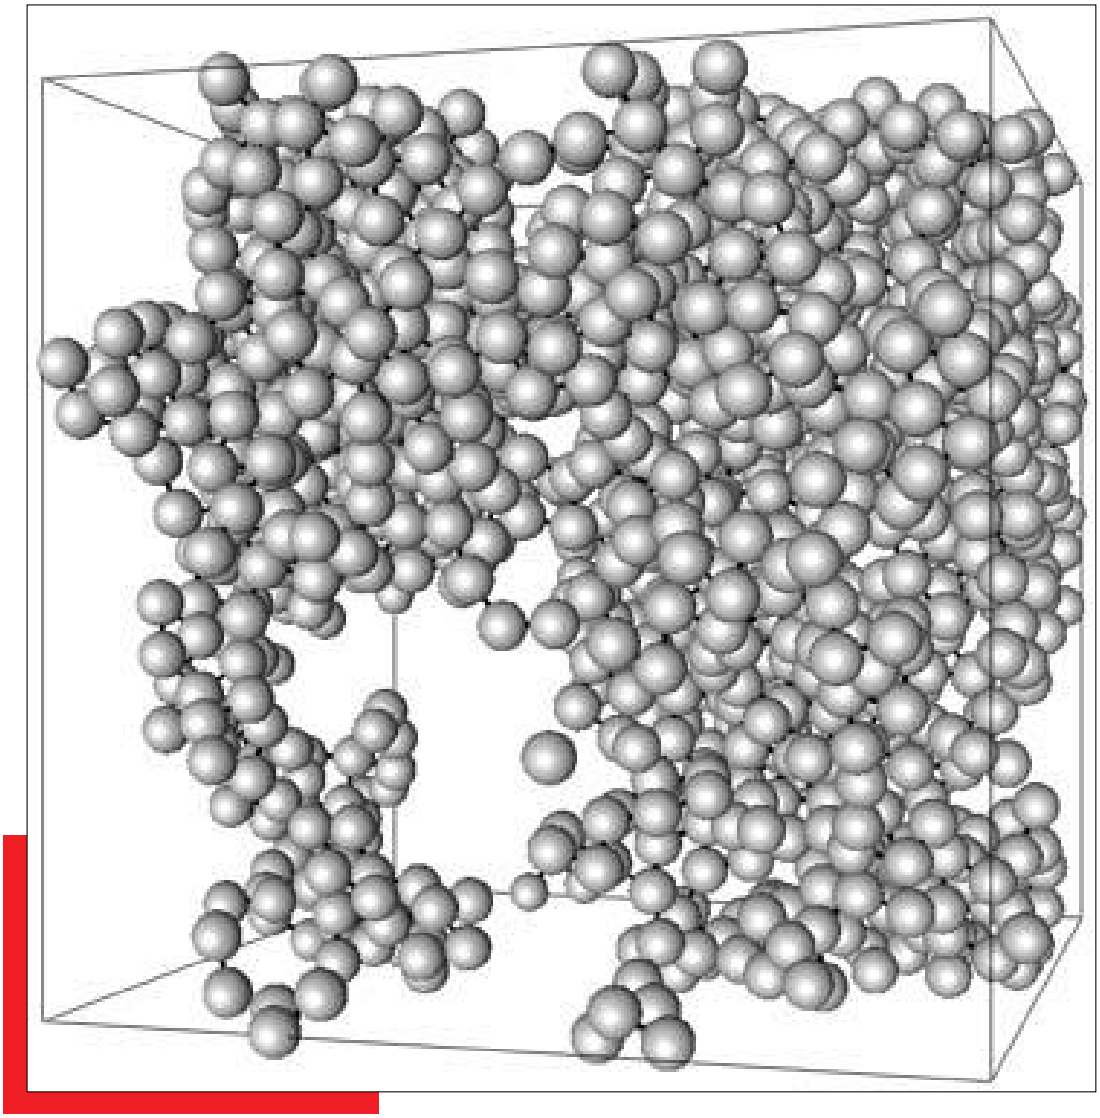

***Figure 2:*** *Cluster of an isotropic fractal system (1369 particles). Bonds are marked for those particles separated less than 1.15 s.*

# INTERDISCIPLINARY AND INDUSTRIAL APPLICATIONS SUPPORT AT CSCS/SCSC

*K.M. Decker (CSCS/SCSC)*

## 1. Introduction

The Swiss Center for Scientific Computing (CSCS/SCSC) has a long tradition in supporting interdisciplinary and industrial applications in a user-oriented and application-driven way. With the recent integration of the *Software Technology Group* (STG) into the *Interdisciplinary and Industrial Applications* (I²A) Group at CSCS/SCSC, this strategy was further reinforced. The new I²A Group offers extended multidisciplinary expertise in various application fields software technologies, both mandatory requirements to achieve leading-edge High-Performance Computing *and* Networking (HPCN) solutions.

It is widely accepted that state-of-the-art HPCN, software technologies, and multi-disciplinarity are critical for developing competitive solutions, both in the private as well as in the public sector. Therefore, the I²A Group is targeted towards collaboration with third party partners to complement and further leverage the competencies of the service and scientific sections of CSCS/SCSC, SSU and SPS, respectively. The goal which has been defined for the group is to exploit multidisciplinary expertise to develop problem-oriented software solutions. To achieve this, computational and software technologies should be applied in an interdisciplinary fashion.

To achieve this goal, the following objectives were defined:

- Technology transfer to third parties,
- Transformation of advanced computational methods into quality software by means of software engineering, and
- Exploitation of research results.

To fulfill these objectives, the I²A Group conducts collaborative applied R&D projects driven by problem- and user-requirements and disseminates expertise by education, training, consulting, and promotional activities both within CSCS/SCSC and third parties.

The common denominator of all group activities is the focus on interdisciplinary aspects such as modeling methods and software technologies, complementary to the in-depth application-specific activities conducted by the application-oriented groups of CSCS/SCSC's scientific projects and support section and any third party partner. In addition, the latest technology developments will be followed to ensure that leading competence in the fields of expertise is maintained, but the group will not push them.

The work of the I²A Group is further characterized by its emphasis on engineering problems for third parties and the focus on development of integrated software solutions. General characteristics comprise the use of professional project management methods, the application of high-performance computing and networking where required, and large flexibility and responsiveness to project requirements.

## 2. Expertise

I²A offers both *practical expertise* and *fundamental and emerging expertise*. Practical expertise has been successfully applied in collaboration projects, used by end-users for some of their real-life applications, and demonstrated for a range of problems and applications. Fundamental and emerging expertise is related to the professional educational background of the group members, and/or areas where basic understanding of theory and concepts exists. This expertise is currently under further development, but not yet mature and competitive, though demonstrators may exist.

Within the group, four clusters of expertise can be identified:

Applied Mechanical Engineering *Practical expertise:* modeling and simulation of complex flow phenomena, discrete models for highly complex geometries, advanced mesh generation techniques, modeling and concurrent engineering of complex multidisciplinary phenomena.

*Fundamental and emerging expertise:* chemical process engineering, structural analysis, linking numerical analysis methods to CAD standard formats, direct interfacing of CAD systems to mesh generators, computational geometry, conception cycle methodologies, gradient free optimization methods.

Engineering for Signal Processing *Practical expertise:* advanced image, signal and data compression algorithms; information retrieval on signals; image segmentation; time, scale and frequency analysis; wavelet analysis.
*Fundamental and emerging expertise:* distributed processing environments dedicated to image compression and satellite images, kinetic measurement on images, audio signal de-noising.

Software Technology *Practical expertise:* parallel programming and application engineering, parallel application engineering and system software tools, distributed computing environments and integrated applications, human-machine interfaces, expert systems, (concurrent) software engineering, problem-solving environments.
*Fundamental and emerging expertise:* information systems, knowledge discovery in data bases (data mining).

Project Management *Practical expertise:* project planning and control; team creation, direction and organization; handling of legal matters and Intellectual Property Rights; project audits.

## 3. Projects

The I²A Group has a diverse spectrum of successfully concluded and running projects, varying with respect to project characteristics (e.g., size, duration), deliverables, and project management requirements.

Successfully concluded projects include

- Interactive Image Processing and Synthesis on Innovative Computer Architectures (EC Human Capital & Mobility Program).
- Joint CSCS/NEC Collaboration in Parallel Processing.
- OPENWAVE: development of an advanced image compression scheme (bilateral cooperation between JRC Ispra and CSCS/SCSC).
- Prototyping problem-solving environments for parallel and distributed programming (Swiss Priority Program Informatics).

Running projects include

- COMPRESS: Development of a distributed software environment applied to image compression schemes assessment (OPENPRESS; EC-DGXIII).
- Interdisciplinary Environment for ESP (Electrostatic Precipitator) simulation (collaboration with ABB).
- Joint CSCS/SCSC-EDF Collaboration in parallel processing (collaboration with Electricité de France).

In addition, a range of new collaborative projects are currently in the planning or realization phase.

Anyone who would be interested in discussing our activities and/or potential future collaboration is invited to contact us. For up-to-date information on CSCS/SCSC's I²A Group, the interested reader is referred to the WWW at URL http://www.cscs.ch/Official/I2A/. ■

# JOINT EDF-CSCS/SCSC COLLABORATION IN MATERIAL SCIENCE

*Christophe Domain (CSCS/SCSC)*

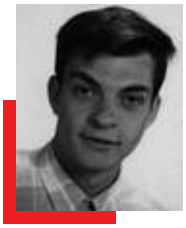

## 1. Introduction

Recognizing that CSCS/SCSC has leading expertise in the area of parallel computing, Electricité de France (EDF) started to discuss possible forms of cooperation in 1993. For the first phase, it was agreed to vectorize and parallelize the Esclangon benchmark code which solves the two-dimensional heat equation with mixed Dirichlet and von Neumann boundary conditions on a rectangular plate using a finite element method [1] for the NEC SX-3, the Cenju-2 and later for the Cenju-3 available at CSCS/SCSC.

After successful completion of phase one, EDF decided to enter into full collaborative research and development effort with CSCS/SCSC on October 1, 1995. The purpose of this collaboration was to parallelize a classical molecular dynamics code for using and evaluating different parallel programming paradigms, e.g., parallelization using the standardized message-passing interface MPI [2] and HPF and extensions prototyped in the context of the *Joint CSCS/NEC Collaboration in Parallel Processing* [3].

Although EDF has a strong tradition in numerical simulations, material science is a relatively new topic. To obtain data which have practical relevance, large systems with several million of atoms must be simulated. Thus vector and parallel machines are mandatory for such simulations.

The molecular dynamics (MD) and Monte Carlo (MC) code DYMOKA studied and further developed is based on the original sequential Fortran code CDCMD from J. Rifkin [4] of the University of Connecticut. The code was parallelized with explicit message-passing for the NEC Cenju-3 distributed memory parallel processor system and and the NEC SX-4 shared memory parallel vector machine. The results are already in practical used by EDF to study the effect of neutron irradiation on steels used in the reactor pressure vessel of nuclear power plants.

## 2. Physical Model

In MD, the classical equation of motions are solved for all atoms where the force between atoms is determined through an empirical interatomic potential. In MC the energy of two successive states is computed, compared and the new state is accepted according to the Metropolis algorithm [5]. The energy calculation is, as in the MD case, based on an interatomic potential.

The metals considered are iron and iron-copper alloy. Since pair potentials are not able to describe the force to the required precision, many-body potentials were developed for better accuracy. In our simulation, the Embedded-Atom Method (EAM) potential [6] was adopted. For this short-range potential, a local "volume force" is included in the calculations as an "embedded energy" which is a function of the local density of electrons.

## 3. Implementation

The interatomic potential is tabulated and interpolated by a 5th order Lagrange polynomial. The equation of motions are integrated with the Gear 5th order predictor-corrector algorithm. In order to have a code fully linear with the number of atoms (and as our potential type is short-range) the neighbor list of interacting atoms is built with the link-cell method [7]. The simulation box is divided into cubes of size larger than the cut-off and for each atom in a given cube, the interaction is only considered for atoms in the

*(Continued on page 30)*

*(Continued from page 29)*

26 nearest- and next-to-nearest cubes of the three-dimensional physical domain, respectively. Taking Newton's third law into consideration, the number of neighboring cubes which needs to be considered can be reduced to 13. Furthermore, the list of interacting atoms is built using a critical distance between atoms larger than the cut-off. This permits that the list must not be rebuild in each MD or MC step (Verlet list method) [7].

## 4. Parallelization

As the potential is short-range, the domain decomposition method [8] was adopted. The three-dimensional simulation box is decomposed into sub-domains and each of them is allocated to one processor element (PE). The link cell method combined with the Verlet list method is applied on each sub-domain like in the sequential code. However, in the parallel version, information about atoms from the cells which are on the neighboring sub-domains need to be communicated. The communication scheme was designed such that the required data from the 13 neighboring sub-domains are communicated into 5 phases by taking care of the evolution of the data locality after each communication phase.

During the simulation of non-equilibrium processes, for instance, crack propagation or irradiation, some atoms can move significantly and thus change of sub-domain. This implies that before rebuilding the neighbor list, atoms are if necessary reallocated to their new sub-domain.

## 5. Performance Results

DYMOKA was developed on the NEC Cenju-3 with the help of the *Annai* [9] integrated application engineering environment and later for the NEC SX-4. The application was tested successfully on a large variety of vector and parallel platforms. Vector machines used include a Cray C90 at EDF and NEC SX-3 and NEC SX-4 systems at CSCS/SCSC. Parallel machines utilized comprise the NEC Cenju-3, Sun SPARCserver 1000 and NEC SX-4 at CSCS/SCSC and a Cray T3D at CEA Grenoble.

The standard Message Passing Interface (MPI) [2] was used in order to achieve portability. The debugging, performance analysis and tuning of the code took advantage of the Annai environment [9] developed in the Joint CSCS/NEC Collaboration in Parallel Processing [3]. The resulting code demonstrated good scalability on all the platforms tested.

The comparison of the performance on different platforms is summarized in Table 1. The cost per atom per MD step is the figure used to compare the performance. For our system, each atom interacts with its 56 neighboring atoms at equilibrium with the EAM potential used [10].

| Machine | #PE | Cost /atom/step ($\mu$s) |
|---|---|---|
| Convex Exemplar (SPP 1000) | 1 | 230 |
| SGI (R10000 195MHz) | 1 | 117 |
| CRAY C90 | 1 | 90 |
| NEC SX-3 | 1 | 30 |
| NEC SX-4 | 1 | 30 |
| | 8 | 4.7 |
| Sun SPARCserver 1000 | 1 | 450 |
| | 2 | 230 |
| | 4 | 160 |
| NEC Cenju-3 | 1 | 270 |
| | 16 | 20 |
| | 128 | 2.5 |
| Cray T3D | 1 | 560 |
| | 64 | 8.8 |
| | 128 | 4.9 |

*Cost per atom per step—Comparison on vector machines (Cray C90 and NEC SX-3) and parallel machines (Sun SPARCserver 1000, NEC Cenju-3, NEC SX-4 and Cray T3D).*

In order to run very large simulation, memory must be saved. The suppression of the neighbor list allows to go from 7 millions to 17.5 millions on the 128 PEs of the Cray T3D; in contrast, the cost per atom and per step is doubled. The replacement of the Gear-5 algorithm by the Verlet integration algorithm allows to reach 21 millions of atoms. Another option is passing from double precision to simple precision, but that incurs a loss of accuracy.

## 6. Simulation Results

The validation of the parallel code was achieved with different physical simulations: the MC component was used to relax a system from a non-equilibrium state and the MD was used for crack propagation and neutron irradiation applications. MD is used to understand the mechanisms of crack propagation and it needs large systems. A crack, initiated using continuum mechanics, propagates (or not) according to the stress applied onto the sample by MD simulation [11]. An example of a propagating crack is illustrated in Figure 1.

Under irradiation, the mechanical properties of the steel of the vessel are changed: hardening, embrittlement, creep and growth are observed. Experimental study does not explain all the mechanism and microstructure of irradiated steels, and MD simulation can be an accurate tool to analyze the irradiation process where the cascade phase in the crystal is very quick (10 ps). Irradiation with energy up to 20 keV were performed. A 20 keV irradiation on the NEC Cenju-3 with a 14 million atoms into iron-2%copper alloy shows the influence of displacement cascades on copper clustering in accordance with experimental results [12]. In the area of the irradiation

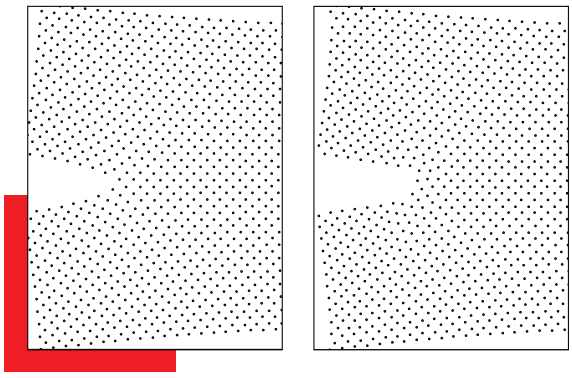

***Figure 1:*** *Crack propagation: initial crack and system after 11 ps.*

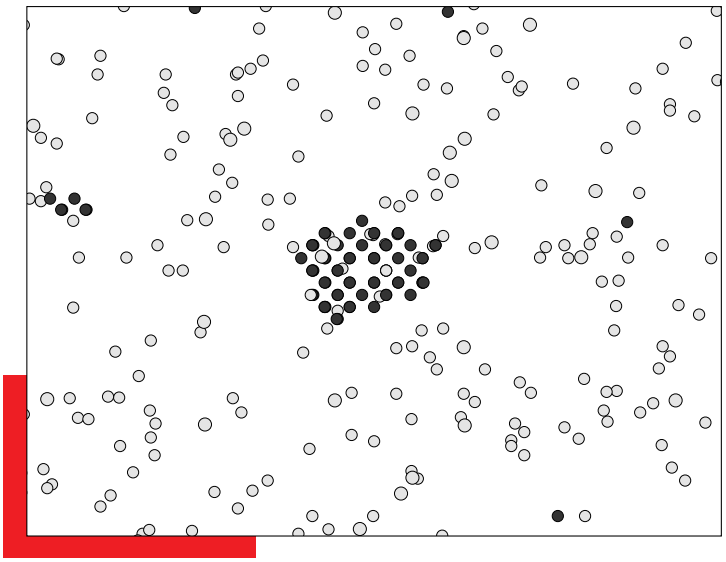

***Figure 2:*** *Cluster of vacancies (black) and copper (white) at the end of a 20 keV irradiation in Fe-2%Cu (NEC Cenju-3)*

a cluster of copper embedded by vacancies was formed into the alloy (Figure 2). Actually, production runs are made on the Cray T3D AC 128-8 (CEA) by EDF in order to investigate in details these new results.

## 7. Interactive Use of DYMOKA

To gain detailed insight into the simulation process itself and to test new simulation scenarios, a simple batch-queuing environment is often not sufficient. To study and evaluate the opportunities of interactive simulations, the DYMOKA application was therefore chosen as a demonstrator application for the project *Interactive Scientific Computing over Networks* (ISCN) currently in progress at CSCS/SCSC in collaboration with L. Lin from the National Yunlin Institute of Technology, Taiwan.

The first ISCN prototype offers simple interactive selection of different target systems, configurations, and batch queues, and interactive access to the running application, even when submitted to a batch queue. The interactive simulation control, step-by-step execution of the application, the immediate visualization of simulation results, and the interactive mechanism to manipulate output data visualization (e.g., stop/go, step-by-step execution, forward/backward replay of critical sequences of data) were proven to be a major help to understand the simulation. An example of an interactive DYMOKA session with ISCN is depicted in Figure 3. For more details on this project and the technologies exploited, the interested reader is referred to [13].

## 8. Perspectives

The new code developed has already demonstrated its added value in material science at EDF with some new results for the neutron irradiation application. DYMOKA is running successfully on different platforms, the performance and the scalability is good. The utilization of the code at EDF will be reinforced.

## Acknowledgments

*We would like to thank all staff at CSCS/SCSC for their help. Provision of computational resources is gratefully acknowledged.*

## References

[1] P. Esclangon. Support Theorique d'un Benchmark Supercalculateur. Technical Report 93NJ00009, Service Informatique et Mathématiques Appliqué, Electricité de France, Apr. 1993. ISSN 1161-059X.

[2] Message Passing Interface Forum. MPI: A Message-Passing Interface Standard. *International Journal of Supercomputer Applications*, 8(3&4):157–416, 1994.

[3] C. Clémençon, K. M. Decker, A. Endo, J. Fritscher, G. Jost, N. Masuda, A. Müller, R. Rühl and W. Sawyer, E. de Sturler, B. J. N. Wylie and F. Zimmermann. Application-Driven Development of an Integrated Tool Environment for Distributed Memory Parallel Processors. In *Proceedings of the First International Workshop on Parallel Processing, Bangalore, India, December 27–30, 1994,* 110–116. Tata McGraw-Hill Publishing, New Delhi, India, Dec.1994. ISBN 0-07-462332-X.

[4] J. Rifkin. *Computer Simulation of Martensitic Transformations.* PhD Thesis, University of Connecticut, 1987.

[5] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller. Equation of State Calculations by Fast Computing Machines. *Jour. Chem. Phys.*, 21(6):1087, 1953

[6] M. Daw and M. Baskes. *Phys. Rev. Lett.*, 50: 1285, 1983.

[7] M. Allen and D. Tildesley, *Computer Simulations of Liquid*. Clarendon Press Oxford, 1987

[8] S. Plimpton. Fast parallel algorithms for short-range Molecular Dynamics. *Journal of Computational Physics*, 117:1, 1995.

[9] Christian Clémençon, Akiyoshi Endo, Josef Fritscher, Andreas Müller, Roland Rühl and B. J. N. Wylie. *Annai*: An Integrated Parallel Programming Environment for Multicomputers. In A. Zaky and T. Lewis, editors, *Tools and Environments for Parallel and Distributed Systems*, chapter 2, pages 33–59. Kluwer Academic Publishers, Feb. 1996. [ISBN:~0-7923-9675-8].

[10] C.S. Becquart, J. Rifkin, J.C. Turbatte, J. Ruste, J.C. Van Duysen. An interatomic potential for FeCu using the embedded-atom method (In preparation).

[11] C.S. Becquart. *Molecular dynamics simulation of crack propagation in ruthenium aluminum*. PhD Thesis, University of Connecticut, 1992.

[12] C.S. Becquart, C. Domain, J. Ruste, Y. Souffez, J.-C. Turbatte and J.C. Van Duysen. Computer simulations study of iron-copper alloy. In *Proceedings of the 3rd International Conference on Computer Simulation of Radiation Effects in Solids (COSIRES, Guildford, United Kingdom, July 22–26, 1996)*, July 1996. To appear in *J. Radiation Effects and Defects in Solids*.

[13] L. Lin, K. M. Decker, M. J. Johnson, C. Domain and Y. Souffez. ISCN: Towards a Distributed Scientific Computing Environment. In *Proceedings of the HPC Asia '97 Conference and Exhibition, Seoul, Korea, April 28–May 2, 1997*, 1997. Accepted for publication.

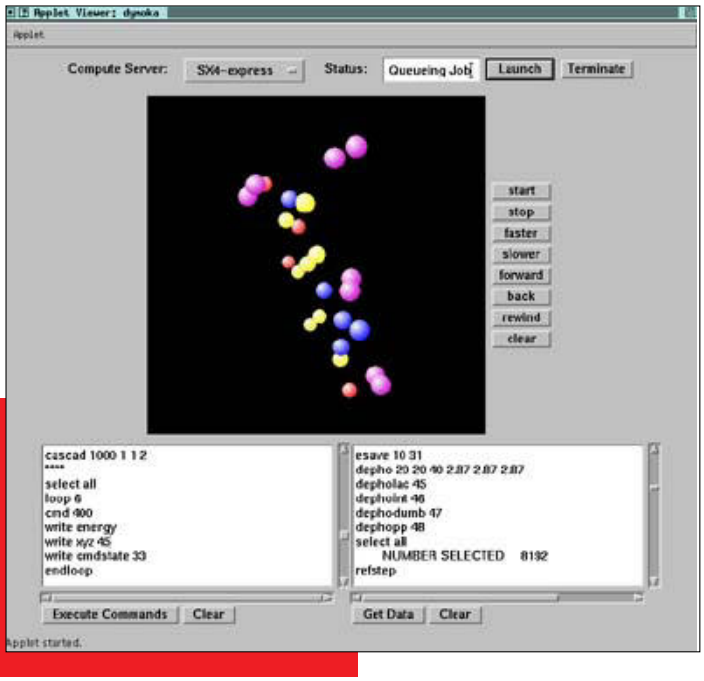


***Figure 3:*** *Interactive DYMOKA session with ISCN.*

## Free Crosscuts Subscription

☐ new ☐ cancel

Mr. ☐ Mrs. ☐ Ms. ☐

Title

Family name | Name

Institute

Organization

Street address

City code | City | Country

E-Mail | Telephone

Return to: CROSSCUTS, Centro Svizzero di Calcolo Scientifico
Via Cantonale, 6928 Manno, Switzerland
cr.subscr@cscs.ch


## Calendar of High-Performance Computing and Related Events

### April 1997

9–11 ISADS 97–Third International Symposium on Autonomous Decentralized Systems
Location: Berlin, Germany
Contact: http://www.fokus.gmd.de/ws/isads97

### May 1997

12–16 ERCOFTAC Summer School: High Performance Parallel Computing for Fluid Dynamic Applications
Location: Daresbury Laboratory, England
Contact:
http://www.dl.ac.uk/TCSC/CompEng/MEETINGS/ERCOFTAC/SummerSchool/

17–18 PDSE97–2nd International Workshop on Software Engineering for Parallel and Distributed Systems. At the 19th International Conference on Software Engineering (ICSE-19)
Location: Boston, Massachusetts, USA
http://www.dis.unina.it/~russo/PDSE97/ or:
www-osl.cs.uiuc.edu/PDSE97/ (mirror site)
and for further information of the ICSE-19 conference:
http://www.ics.uci.edu/icse97 or:
http://www.elet.polimi.it/~fuggetta/icse97/ (mirror site)
Dr Stefano Russo, Dipartimento di Informatica e Sistemistica,
Università di Napoli, Via Claudio 21, 80125 Napoli, Italy
EMAIL: russo@nadis.dis.unina.it
FAX: +39 81 7683186
TEL: +39 81 7682895 (direct line)

19–21 Parallel CFD '97: Implementations and Results Using Parallel Computers
Location: Manchester, England,
Contact: http://www.dl.ac.uk/TCSC/CompEng/MEETINGS/CFD97/

28–30 IWPC '97–5th International Workshop on Program Comprehension
Location: Dearborn, Michigan, USA
Contact:
http://www.dis.unina.it/~iwpc97 and http://www.cacs.usl.edu/~iwpc97

### June 1997

1–13 IRREGULAR'97, 4th International Symposium on Solving Irregularly Structured Problems in Parallel
Location: University of Paderborn, Germany
Contact: http://cuiwww.unige.ch/~rolim/irregular

9–11 Euro PDS'97 Parallel and Distributed Systems
Location: Barcelona, Spain
For more information or to be placed on the mailing list, please contact:
Iasted Secretariat, Euro-PDS'97
#80, 4500-16 Avenue, N.W.
Calgary, AB Canada T3B 0M6
Email: iasted@cadvision.com
Contact: http://www.iasted.com

10–13 PADS'97–11th Workshop on Parallel and Distributed Simulation
Location: Castle Lockenhaus, near Vienna, Austria
Contact: http://www.it.kth.se/labs/sim/pads97

### July 1997

4–5 AIN'97, the 2nd International Workshop on Advanced Intelligent Networks
Location: Cesena, Italy
Contact: http://brahms.fmi.uni-passau.de/bs/organization/

14–16 ASAP 97–11th International Conference on Application-specific Systems, Architectures and Processors
Location: Federal Institute of Technology Zurich, Switzerland
Contact: http://www.tik.ee.ethz.ch/ASAP, mail to:asap@tik.ee.ethz.ch

21–23 FMOODS'97–Second IFIP International conference on Formal Methods for Open Object-based Distributed Systems
Location: Canterbury, United Kingdom
Contact: http://alethea.ukc.ac.uk/Dept/Computing/Research/NDS/FMOODS/July

27–31 SIGIR '97–20th International ACM SIGIR Conference on Research and Development in Information Retrieval
Location: Philadelphia, PA, USA