EuroGPU2009
at

ParCo 2009

PROGRAM & ABSTRACTS

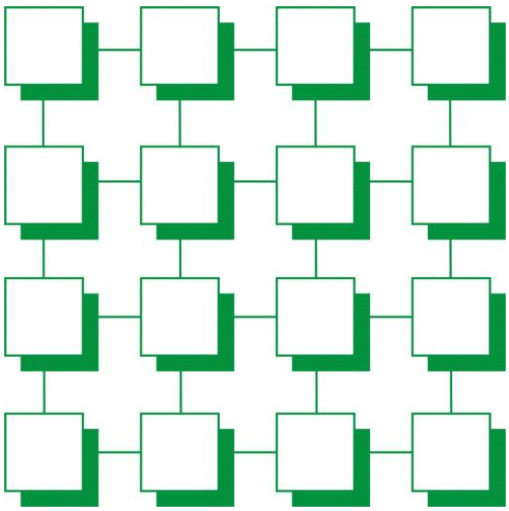

Parallel Computing 2009
1 – 4 September 2009

Ecole Normale Superieure de Lyon (ENS),
Lyon, France

EuroGPU'09: Parallel Computing with GPU's.

Organisers

Anne C. Elster (Norwegian University of Science and Technology (NTNU), Norway) and Stephane Requena (Genci, Paris, France) with Guillaume Colin de Verdière, CEA0

Summary

The success of the gaming industry is now pushing processor technology like we have never seen before. Since recent graphics processors (GPU's) have been improving both their programmability as well as have been adding more and more floating point processing, it makes them very appealing as accelerators for general-purpose computing.

This minisymposium gives an overview of some of these advancements by bringing together experts working on the development of techniques and tools that improve the programmability of GPU's as well as the experts interested in utilizing the computational power of GPU's for scientific applications.

The minisymposium will start with a short history and overview as seen by the organizer. Abstractions of programming models accross multi-core and GPU architectures will then be presented following a talk describing latency issues with GPU's. The first set of talks then are rounded off with a talk on asynchronous algorithms for solving PDEs on a cluster of GPU's.

In the second part of the minisymposium, we will focus on topics related to GPU's and visualization by starting off with describing visualization of snow using the power of modern GPU's to help simulate the interaction of several million snow particles in a wind field and building up on a given terrain. This talk will be followed by a presentation on how to use GPU's to accellerate a wall-sized high-resolution tiled display wall. Finally, further talks will be included, for example, on voxalization algorithms on GPU's.

TUESDAY Sept. 1, 2009 Euro GPU 2009 – DAY 1

Time	*Title/Speaker*
11:30-12:00	**Intro. – GPU Computing**, *Anne C. Elster, Norwegian University of Science and Technology (NTNU), Trondheim, Norway*
12:00-12:30	**Throughput Computing on Future GPUs** *Rune Johan HOVLAND and Anne C. ELSTER, NTNU, Norway*
12:30-14:00	LUNCH
14:00-14:30	I1: **OpenCL a new standard for GPU programming** *Francois BODIN -- Caps Entreprise, Rennes, France*
14:30-15:00	I2: **Heterogeneous Multicore Parallel Programming** *Stéphane BIHAN -- Caps Entreprise, Rennes, France*
15:00-15:30	I3: **Cosmological reionisation powered by multi-GPUs** *Dominique AUBERT [a], Romain TEYSSIER [b]* *[a] Université de Strasbourg, France [b] CEA, France*
15:30-16:00	Coffee/Tea Break & Exhibitions
16:00-16:30	I3: **Efficient use of hybrid computing clusters for nanosciences** *Lugi GENOVESE [a], Matthieu OSPICI [b],* *Jean Francois MEHAUT [c], Thierry DEUTSCH [d]* *[a] ESFR, Grenoble, [b] BULL, UJF/LIG, CEA, Grenoble* *[c] UJF/INRIA, Grenoble, [d] CEA, Grenoble, France*
16:30-17:00	I4*:* **Accelerating depth imaging seismic application on GPUs, status and perspectives,** *Henri CALANDRA, TOTAL, Pau*
17:00-17:30	I5: **Debugging for GPUs with DDT** *David LECOMBER, Allinea Ltd, Bristol, UK*
19:00-20:00	Reception: City Hall

WEDNESDAY Sept. 2, 2009 Euro GPU 2009 – DAY 2

Time	*Title/Speaker*
10:00-10:30	**Porus Rock Simulations and Lattice Boltzmann on GPUs** *Erik Ola AKSNES and Anne C. ELSTER, NTNU, Norway*
10:30-11:00	**An efficient multi-algorithms sparse linear solver for GPUs** *Stéphane VIALLE; Thomas JOST and Sylvain CONSTASSOT-VIVIER, Supélec Campus de Metz, France*
11:00-11:30	Coffee/Tea Break & Exhibitions
11:30-12:00	**Abstraction of Programming Models Across Multi-Core and GPGPU Architectures,** *Ian GRIMSTEAD and David R. WALKER, Cardiff University, UK*
12:00-12:30	**Modeling Communication on Modern GPU Systems**, Anne C. ELSTER, Thorvald NATVIG, and Daniele G. SPAMPINATO, NTNU, Norway
12:30-14:00	LUNCH
14:00-15:00	PANEL DISCUSSION – ParCo 2009
15:00-15:30	Coffee/Tea Break & Exhibitions
15:30-16:00	**An Efficient GPGPU Implementation of the Revised Simplex Method** *Jakob BIELING, Patrick PRESCHLOW, Peter MARTINI, University of Bonn, Germany*
16:00-|17:00	**PANEL DISCUSSION on GPU Computing**
18:30-22:00	Excursion & Conference Dinner

EuroGPU'09 ABSTRACTS

Session C-EGPU1	EuroGPU	Tuesday, 1.9.2009	11:30–12:30

Intro. – GPU Computing

Anne C. ELSTER,
Norwegian University of Science and Technology (NTNU), Trondheim, Norway

This talk will serve as an intro to EuroGPU 2009 and include the histrory of GPUs and GPGPUs. Dr. Elster's research la bat NTNU has had a focus on GOU computing the last several years and will also highlight some of her and her students recent achievements. This talk will then run directly into the following talk followed by highlights of the rest of EuroGPU 2009.

Throughput Computing on Future GPUs

Rune Johan HOVLAND and Anne C. ELSTER,
Norwegian University of Science and Technology (NTNU), Trondheim, Norway

The process of moving data to and from the GPU is inuenced by two important metrics, latency and bandwidth. For small transfers, the latency can severely impact the performance, while for larger transfers the bandwidth comes more and more into play. These metrics are therefore the metrics that will be used throughout this report to judge the important of various properties of the host system. These properties include processor clock frequencies, chipsets, memory frequencies and architecture, as well as the PCI Express bus.

Our measurements performed shows how the PCI Express bus is a major bottleneck for the transfers to and from the GPU, making overclocking this bus an action worth considering. The CPU clock frequency which one would assume to have great influence, proved not to affect the bandwidth at all, and affected the latency only to a small extent. The architecture of the CPU, however provedto be a crucial aspect

Session C-EGPU2	EuroGPU	Tuesday, 1.9.2009	14:00-14:30

OpenCL a new standard for GPU programming

Francois BODIN [a]
[a] *Caps Entreprise, Rennes*

In this presentation we give an overview of OpenCL for programming Graphics Processing Units (GPUs). OpenCL is an initiative launched by Apple to ensure application portability across various types of GPUs. It aims at being an open standard (royalty free and vendor neutral) developed by the Khronos OpenCL working group (http://www.khronos.org). OpenCL which is based on the ISO C99, shares many features with CUDA and exposes data and task parallelism.

Session C-EGPU2	EuroGPU	Tuesday, 1.9.2009	14:30-15:00

Heterogeneous Multicore Parallel Programming

Stéphane BIHAN
Caps Entreprise, Rennes

Hybrid parallel multicore architectures based on Graphics Processing Units (GPUs) can provide tremendous computing power. Current NVIDIA and AMD Graphics Product Group hardware display a peak performance of hundreds of gigaflops. However, exploiting GPUs from existing applications is a difficult task that requires non-portable rewriting of the code. In this talk, we present HMPP, a Heterogeneous Multicore Parallel Programming workbench with compilers, developed by CAPS entreprise, that allows the integration of heterogeneous hardware accelerators in a non-intrusive manner while preserving legacy codes.

Session C-EGPU2	EuroGPU	Tuesday, 1.9.2009	15:00-15:30

Cosmological reionisation powered by multi-GPUs

Dominique AUBERT[a], *Romain TEYSSIER*[b]
[a] *Université de Strasbourg* [b]*CEA*

By definition, cosmology cannot rely on lab experiments to reproduce the phenomenons observed in the sky and test its theories. For this very reason, the use of numerical simulations is widely spread within this community in order to understand the formation of the astrophysical objects and to put constrains on the physical ingredients that lead to the Universe as it is currently observed.
One of the key question in cosmology is how was the matter organised in the early Universe and how can it be observed? In particular, the first stars (around 500 million years after the Big-Bang) are expected to emit a strong ultra-violet radiation that ionize the gas which fills the Universe. This process (called reionisation) can be observed by looking at radio wavelengths very far in the Universe using ground based installation such as LOFAR or the forthcoming SKA radio antennas. These experiments will probe the gas distribution at the earliest epochs. We developped a code (CUDATON) which take the simulated distribution of gas and stars in the early Universe and it models the propagation of ionising radiation and its effect on the gas. This modelling will help to understand the radio observations in a near future and the impact of this first stellar light on the formation of galaxies.
Basically, the code explicitely solves a set of conservative equations on a fixed grid in a similar manner to hydrodynamics and it follows the evolution of a fluid made of photons. However due to typical velocities close to the speed of light, the stringent CFL condition implies that a very large number of timesteps must be computed, making the code intrinsically slow. However, we ported it on GPU architecture using CUDA and it removed this intrinsic limitation by accelerating the code by a factor close to 80. Furthermore, using a MPI layer, I expanded it to a multi-GPU version and CUDATON is currently running on 128 GPUs installed on the new CCRT calculator of the French atomic agency (CEA). The code is able to perform 60 000 timesteps on a 10243 grid in ~2.5 hours (elapsed). For comparison, the largest calculation made so far on the same topic involved a 4003 grid and required 11 000 cores to be run.
Such a boost in the performance demonstrates the relevance of multigpu calculations for computational cosmology. It also opens bright perspectives for a systematic exploration of the impact of the physical ingredients on high resolution simulations since the these calculations are extremely fast to complete.

Session C-EGPU3	EuroGPU	Tuesday, 1.9.2009	16:00-16:30

Efficient use of hybrid computing clusters for nanosciences

Lugi GENOVESE [a], Matthieu OSPICI [b],
Jean Francois MEHAUT [c], Thierry DEUTSCH [d]
[a] ESFR, Grenoble
[b] BULL, UJF/LIG, CEA, Grenoble
[c] UJF/INRIA, Grenoble
[d] CEA, Grenoble

In this talk, we will study the programming and the utilization of hybrid clusters in the field of computational physics. These massively parallel computers are composed of a fast network (Infiniband) connecting classical nodes with multicore Intel processors and accelerators. In our case, the accelerators used are GPUs from NVIDIA.
First, we will analyze some ways to use with efficiency CPUs cores and GPUs together in a code (BiGDFT, http://inac.cea.fr/L_Sim/BigDFT) without hotspot routines. Starting from this analyze, we have designed a new library: S_GPU, used to share GPUs between the CPU cores of a node. The implementation and the usage of S_GPU will be described in this talk.

Then, we will evaluate and compare performances between S_GPU and others approaches to share GPUs with CPUs. This performance evaluation will be based on BigDFT, an ab-initio simulation software designed to take advantage of massively hybrid parallel clusters as the Titane cluster (CCRT). Our experiments will be performed, on a first time, on one hybrid node, and on a second time, on a large number of nodes of the hybrid cluster.

Session C-EGPU3	EuroGPU	Tuesday, 1.9.2009	16:30-17:00

Accelerating depth imaging seismic application on GPUs, status and perspectives

Henri CALANDRA, TOTAL, Pau

The extraordinary challenge that the oil and gas industry must face for hydrocarbon exploration requires the development of leading edge technologies to recover an accurate representation of the subsurface. Seismic modeling and Reverse Time Migration (RTM) based on the full wave equation discretization, are tools of major importance since they give an accurate representation of complex wave propagation areas. Unfortunately, they are highly compute intensive. Advances in High Performance Computing technologies resulted in renewed attention from the seismic community to these techniques.
The recent development in GPU technologies with unified architecture and general-purpose languages coupled with the high and rapidly increasing performance throughput of these components made General Purpose Processing on Graphics Processing Units (GPGPU) an attractive solution to speed up diverse applications.
In this presentation we will first present the challenges in O&G and the need in terms of computing power for solving seismic depth imaging problems. We then show how GPUs can be part of the solution and the the solutions developped in TOTAL.

Session C-EGPU3	EuroGPU	Tuesday, 1.9.2009	17:00-17:30

Debugging for GPUs with DDT

David LECOMBER, Allinea Ltd, Bristol, UK

GPGPUs have become a hot topic in high-performance computing for their remarkable power efficiency and peak performance. Developers are experimenting with CUDA, OpenCL and others to port (or rewrite) their code to take advantage of this technology, but are discovering there is more to programming than writing code.
Finding bugs and optimizing performance are essential tasks - particularly so with a new complex model of program execution. In this talk I will review the state of play - exploring what is possible now, and what is being done by Allinea and others to improve the lot of GPU developers who need to debug or optimize their codes.

WEDNESDAY

Session C-EGPU3	EuroGPU	Wednesday, 2.9.2009	10:00-10:30

Porus Rock Simulations and Lattice Boltzmann on GPUs

Erik Ola AKSNES and Anne C. ELSTER, NTNU, Norway

In this paper, we show how the Lattice Boltzmann method can be used to calculate the porous Rocks' ability to transport fluids (permeability), an important problem for oil industry, among others. Our work is done in collaborations with Numerical Rocks AS and the Dept. of Petroleum Engineering at the Norwegian Univ. of Science and Technology.

To better evaluate our GPU implementation, a sequential CPU implementation is first prepared. We then develop our GPU implementation and test both implementation using three porous data sets with known permeabilities provided by Numerical Rocks AS. Our simulations of fluid flow get high performance on modern GPUs showing that it is possible to calculate the permeability of porous rocks of simulations sizes up to 3683, which fit into the 4GB memory of the NVIDIA Quadro FX 5800 card. The performancesof the CPU and GPU implementations are measured in MLUPS (million lattice node updates per second). Both implementations achieve their highest performances using single floating-point precision, resulting in the maximum performance equal to 1.59 MLUPS and 184.30 MLUPS. Techniques for reducing round-off errors are also discussed and implemented

Session C-EGPU3	EuroGPU	Wednesday, 2.9.2009	10:00-10:30

An efficient multi-algorithms sparse linear solver for GPUs

*Stéphane **VIALLE**; Thomas JOST and Sylvain CONSTASSOT-VIVIER*
Supélec Campus de Metz, France

We present a new sparse linear solver for GPUs. It is designed to work With structured sparse matrices where all the non-zeros are on a few diagonals. Several iterative algorithms are implemented, both on CPU and GPU. The GPU code is designed to be fast yet simple to read and understand. It aims to be as accurate as possible, even on chips that do not support double-precision floating-point arithmetic. Several benchmarks show that GPU algorithms are much faster than their CPU counterpart while their accuracy is satisfying.

Session C-EGPU3	EuroGPU	Wednesday, 2.9.2009	10:00-10:30

Abstraction of Programming Models Across Multi-Core and GPGPU Architectures

*Ian **GRIMSTEAD** and David R. WALKER*
Cardiff University, Wales, United Kingdom

Keywords: programming models, application acceleration, GPGPU, ClearSpeed
Work in the field of application acceleration devices is showing great promise, but still remains a tool largely for computer scientists with domain knowledge, given the complexity of porting existing algorithms to new architectures or environments. Such porting is hindered by the lack of abstraction available.

We present our latest work in the development of a novel solution to this abstraction problem; an intelligent semi-automatic porting system. This allows a higher level of abstraction where the user does not have to intervene or annotate their source code, while maintaining reasonable levels of performance. We present comparisons between manual and automatic code ports on two different platforms (NVIDIA CUDA and ClearSpeed C^n), showing the versatility of this approach.

Session C-EGPU3	EuroGPU	Wednesday, 2.9.2009	10:00-10:30

An Efficient GPGPU Implementation of the Revised Simplex Method

Jakob BIELING, Patrick PRESCHLOW, Peter MARTINI
University of Bonn, Germany

General purpose computing on a GPU (GPGPU) describes the utilization of graphics cards to solve problems not exclusively related to graphics. As such, a GPU is used much like a co-processor similar to a math co-processor to supplement the CPU. While the clock rate is usually below that of modern CPUs, the high number of processors a GPU provides (typically around 50 to 500 as of today) makes it suitable for problems that are highly parallelizable. A good example are matrix multiplications, as each resultant element can be calculated independently from all others. A particular promising application is the revised simplex method, which, due to a heavy use of various matrix operations, promises to greatly benefit from GPGPU.

Different approaches exist in order to make use of a GPU for general purpose computing. One is to use a vendor-specific programming language for direct access to the GPU. The other is to use the traditional graphical approach and treat matrices as textures (i.e., images), mapping each element of the matrix to one texture element at the respective position. Shaders, initially provided to allow more sophisticated shading of geometry, are then used to perform calculations using these textures. In our work, we choose the second approach to implement the revised simplex algorithm.

In the paper, we give a detailed description of the algorithm implementation as well as an extensive comparison with a commonly used CPU implementation, the GNU Linear Programming Kit (GLPK). Using several optimization techniques, we are able to achieve a speedup of up to 26 over the GLPK when solving problems with 1300 decision variables and 1300 constraints. The results show a strong tendency towards an even greater speedup for larger problem

Program Committee of ParCo 2009

(with Anne Elster's afficilation corrected)

Last Name	First Name	Institution	Country
Abramson	David	Monash University	Australia
Alam	Sadaf	Oak Ridge National Laboratory	USA
Allen	Gabrielle	Louisiana State University	USA
Badia	Rosa	BSC	Spain
Bartz	Dirk	Univ. Tübingen	Germany
Berthou	Jean-Yves	EDF R&D	France
Bjørstad	Petter	Univ. Bergen	Norway
Bode	Arndt	TU Munich	Germany
Bubak	Marian	AGH Univ. Krakow	Poland
Bull	Mark	EPCC	UK
Calandra	H.	Total	France
Chauhan	Arun	Indiana University	USA
Clematis	Andrea	CNR	Italy
D'Ambra	Pasqua	ICAR-CNR	Italy
D'Amore	Luisa	Univ. Naples	Italy
D'Hollander	Erik H.	Univ. Gent	Belgium
Dayde	Michel	ENSSEIHT	France
Dehne	Frank	Carleton Univ.	Canada
DeRose	Luiz	Cray Inc.	USA
Elster	Anne	Norwegian University of Tech. (NTNU)	Norway
Gabriel	Edgar	University of Houston	USA
Gallopoulos	Efstratios	Univ. Patras	Greece
Huang	Lei	University of Houston	USA
Huck	Kevin	University of Oregon	USA
Jeannot	Emmanuel	INRIA	France
Jin	Hai	Huazhong University	China
Kessler	Christoph	Univ. Linköping	Sweden
Kranzlmüller	Dieter	Univ. Linz	Austria
Kuchen	Herbert	Univ. Muenster	Germany
Lastovetsky	Alexey	Univ. College Dublin	Ireland
Leca	Pierre	CEA DAM	France
Ludwig	Thomas	Univ. Heidelberg	Germany
Luque	Emilio	Univ. Autonomica Barcelona	Spain
Massaioli	Federico	CASPUR	Italy
Matsuoka	Satoshi	Tokyo Institute of Technology	Japan
Nagel	Wolfgang	TU Dresden	Germany
Nakajima	Kengo	Univ. Tokyo	Japan
Namyst	Raymond	LABRI / University of Bordeaux I	France

Pérez	Christian	INRIA/ENS Lyon	France
Petiton	Serge	Ecole Universitaire des Ingenieurs de Lille	France
Petkov	Nicolai	Univ. Groningen	Netherlands
Plata	Oscar	University of Malaga	Spain
Roman	Jean	LABRI	France
Rüde	Ullrich	Univ. Erlangen-Nürnberg	Germany
Rünger	Gudula	TU Chemnitz	Germany
Sawley	Marie-Christine	ETH Zurich	Switzerland
Schulz	Martin	Lawrence Livermore National Laboratory	USA
Sørevik	Tor	Univ. Bergen	Norway
Strohmaier	Erich	Lawrence Berkeley National Lab.	USA
Suter	Frédéric	CNRS/IN2P3	France
Talia	Domenico	Univ. Calabria	Italy
Tirado	Paco	Univ. Computense Madrid	Spain
Trystram	Denis	IMAG	France
Vanneschi	Marco	Univ. Pisa	Italy
Xu	Zhiwei	ICT	China
Zomaya	Albert	Univ. Sidney	Australia

Advisory Committee of ParCo 2009

Jean-Yves Berthou	Electricité de France - Research and Development, France
Jack Dongarra	University of Tennessee, USA
Alain Lichnewsky	Université de Paris Sud - GENCI, France
Thomas Lippert	Forschungszentrum Jülich, Germany

ParCo 2009 Locations

From the subway to the conference site :

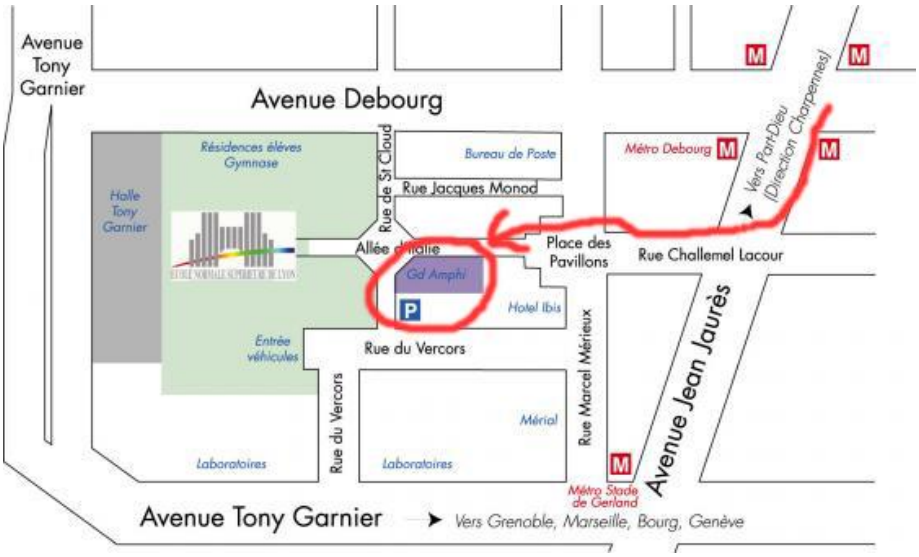

ParCo Rooms Location :

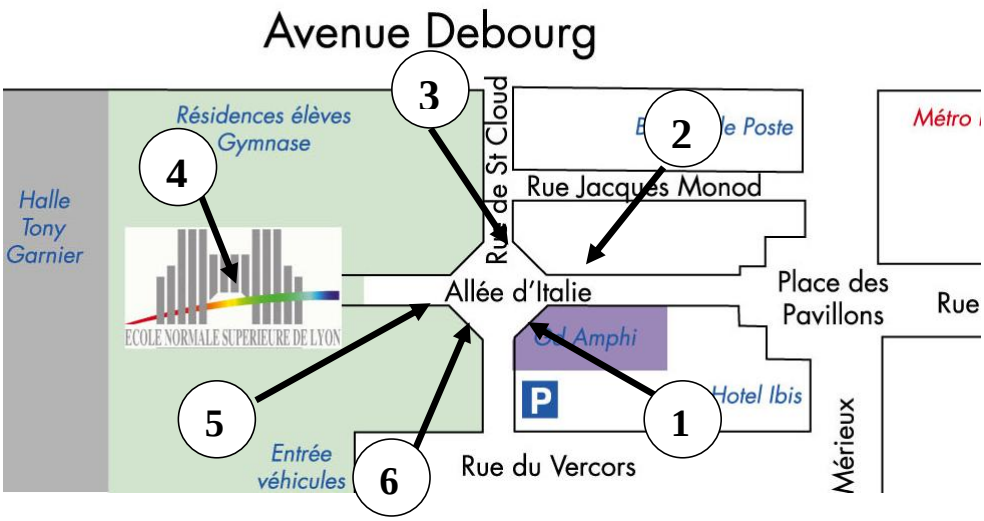

1 : Mérieux Theater
2 : UNESCO Room (Ground floor)
3 : UNESCO Room (1st floor)
4 : I, J, K, L Theaters (ENS)
5 : Tutorial Break Room
6 : ENS Restaurant

Social Events

Tuesday, 1st September 2009 : Welcome reception and cocktail in Lyon City Hall

A bus will take you from the conference site at 19:00 and bring you back around 21:30 (additional bus-stop at Bellecour for those whose hotel is in the center of the town)

Wednesday, 2nd September 2009 : Conference Banquet at the Abbaye de Collonges

Three miles (5km) outside the city of Lyon, on the banks of the river Saône, the "Abbaye of Collonges" stands in all its splendor, offering you a unique welcome.
The sound of the playing of its grand fairground organs, so lovingly restored, will bring thrills and joy to the heart, the moment you enter this domain created expressly for good times.
The prestigious reputation of the Abbaye de Collonges is not its only striking asset. The enthusiasm of its staff always ready to meet your every requirement, adds much to the fabulous attraction of this house, generating that age old sense of festivity, la fête.

A bus will take you from the conference site to the "Abbaye de Collonges" at 19:00 and bring you back to the ENS around 22:30 (additional bus-stop at Bellecour for those whose hotel is in the center of the town).

Comment [unknown1]: This page already on p. 2

ParCo 2009 Sponsors