



# Computational Science and HPC Software-Development in Europe

**Thomas Lippert / Bernd Mohr**  
Forschungszentrum Jülich, JSC  
and Gauss Centre for Supercomputing e.V.

**1<sup>st</sup> Workshop of the International Exascale  
Software Project (IESP), Santa Fé, April 7-8, 2009**



## Thanks to

- Jean-Yves Berthou (EDF)
  - Michel Marechal (ESF, Lincei Initiative, CSEC)
  - Achim Bachem and all friends from PRACE
  - Catherine Riviere (GENCI, PRACE)
  - Peter Michielse (PRACE WP6)
  - Herbert Huber (PRACE-STRATOS)
  - Wolfgang Nagel (Gauß Alliance)
  - Stefan Heinzel (DEISA)
  - Kimmo Koski (HET, COSI-HPC)
  - Wanda Andreoni et al. (CECAM)
  - Martyn Guest (Editor of HET/HPC-EUR–Scientific Case)
- and many others

# Science Fields & Drivers in Europe

From

- HET Scientific Case
- PRACE Initiative
- Lincei Initiative Report

# HET: Scientific Case White Paper

Area	Application	Science Challenges & Potential Outcomes
<b>Weather, Climatology and Earth Sciences</b>	<b>Climate change</b>	Quantify uncertainties on the degree of warming and the likely impacts by increasing the capability and complexity of 'whole earth system' models that represent the scenarios for our future climate ( <b>IPCC</b> ).
	<b>Oceanography</b>	Build the most efficient modelling and prediction systems to study, understand and predict ocean properties and variations at all scales, and develop <b>economically relevant applications</b> to inform policy
	<b>Meteorology, Hydrology</b>	Predict weather and flood events with <b>high socio-economic and environmental impact</b> within a few days. Understand and predict the quality of air at the earth's surface; development of advanced real-time forecasting systems for early enough warning and practical mitigation in the case of pollution crisis.
	<b>Earth Sciences</b>	Challenges span a range of disciplines and have scientific and social implications, such as the <b>mitigation of seismic hazards</b> , treaty verification for nuclear weapons, and increased discovery of economically recoverable petroleum resources and monitoring of waste disposal. Increased computing capability will make it possible to address the issues of resolution, complexity, duration, confidence and certainty.
<b>Astrophysics, HEP and Plasma Physics</b>	<b>Astrophysics</b>	Deal with systems and structures which span a large range of different length and time scales; almost always non-linear coupled systems differential equations have to be integrated, in 3 spatial dimensions and explicitly in time, with rather complex material functions as input. Grand challenges range from formation of stars and planets to questions concerning the <b>evolution of the Universe</b> as a whole. Evaluate the huge amount of data expected from future space experiments such as the European Planck Surveyor satellite.
	<b>Element. Part. Physics</b>	Quantum field theories like QCD (quantum chromodynamics) are the topic of intense theoretical and experimental research by a large and truly international community involving large European centers like <b>GSI</b> and <b>CERN</b> . This research promises a much deeper understanding of the standard model as well as nuclear forces, but is also to discover yet unknown physics beyond the standard model.
	<b>Plasma physics</b>	The science and technology challenge raised by the construction of the magnetic confinement fusion reactor <b>ITER</b> calls for a major theory and modelling activity. Both the success of the experiment and its safety rely on such simulators. The quest to realize thermonuclear fusion by magnetically confining a high temperature plasma poses computationally most challenging problems of nonlinear physics.
<b>Materials Science, Chemistry and Nanoscience</b>	<b>Understanding Complex Materials</b>	The determination of electronic and transport properties is central to many devices in the electronic industry and hence to progress in the understanding of technologically relevant materials. Simulations of nucleation, growth, self-assembly and polymerization for design and performance of many diverse materials e.g., rubbers, paints, fuels, detergents, functional organic materials, cosmetics and food. Multiscale descriptions of the mechanical properties of materials to determine the relation between process, conditions of use and composition e.g., in <b>nuclear energy production</b> . Such simulations are central to the prediction of the lifetime of high performance materials in energy technology.
	<b>Understanding Complex Chemistry</b>	<b>Catalysis</b> is a major challenge in the chemistry of complex materials, with many applications in industrial chemistry. The knowledge of atmospheric chemistry is crucial for environmental prediction and protection (clean air). Improving the knowledge of chemical processing would improve the durability of chemicals. Supra molecular assemblies open new possibilities for the extraction of heavy elements from spent nuclear fuels. In biochemistry, a vast number of reactions in the human body are not understood in any detail. A key step for <b>clean fuels</b> of the future requires the realistic treatment of supported catalytic nanoparticles.
	<b>Nanoscience</b>	The advance of faster information processing or the development of <b>new generations of processors</b> requires the shrinking of devices, which leads inevitably towards nanoelectronics. Moreover, many new devices, such as nanomotors can be envisioned, which will require simulation of mechanical properties at the nanolevel. Composite high performance materials in the fields e.g. adhesion and coatings will require an atomistic based description of nanorheology, nanofluidics and nanotribology.

# HET: Scientific Case White Paper II

Area	Application	Science Challenges & Potential Outcomes
<b>Life sciences</b>	<b>Systems Biology</b>	The use of increasingly sophisticated models to represent the entire behaviour of cells, tissues, and organs, or to evaluate degradation routes predicting the final excretion product of any drug. <b>In silico cell.</b>
	<b>Chromatine Dynamics</b>	The organization of DNA in nucleosomes largely modifies the accessibility of transcription factors recognition sites playing then a key role in the regulation of gene function. The understanding of nucleosome dynamics will be crucial to understand the mechanism of <b>gene regulation.</b>
	<b>Large Scale Protein Dyn.</b>	The study of large conformational changes in proteins. Major challenges appear in the simulation of protein missfolding, unfolding and refolding (understanding of prion-originated pathologies).
	<b>Protein association and aggregation</b>	One of the greatest challenges is the simulation of crowded <b>“not in the cell” protein environments.</b> To be able to represent “in silico” the formation of the different protein complexes associated with a signalling pathway opens the door to a better understanding of cellular function and to the generation of new drugs.
	<b>Supramolecular Systems</b>	The correct representation of protein machines is still out of range of European groups using current simulation protocols and computers. The challenge will be to analyze systematically how several of these machines work e.g., ribosome, topoisomerases, polymerases.
	<b>Medicine</b>	Genome sequencing, massive genotyping studies are providing massive volumes of information e.g. the simulation of the determinants triggering the development of <b>multigenic-based diseases</b> and the prediction of secondary effects related to bad metabolism of drugs in certain segments of population.
<b>Engineering</b>	<b>Helicopter Simulation</b>	The European helicopter industry has a strong tradition of innovation in technology and design. Computational Fluid Dynamics (CFD) based simulations of aerodynamics, aeroacoustics and coupling with dynamics of rotorcraft play a central role and will have to be improved further in the design loop.
	<b>Biomedical Flows</b>	Biomedical fluid mechanics can improve healthcare in many areas, with intensive research efforts in the field of the human circulatory system, the artificial heart or <b>heart valve prostheses</b> , the respiratory system with nose flow and the upper and lower airways, and the human balance system.
	<b>Gas Turbines &amp; Internal Combustion Engines</b>	Scientific challenges in gas turbines or piston engines are numerous. First, a large range of physical scales should be considered from fast chemical reaction characteristics (reaction zone thicknesses of about tens of millimetres, $10^{-6}$ s), pressure wave propagation up to burner scales (tens of cm, $10^{-2}$ s) or system scales.
	<b>Forest Fires</b>	The development of reliable numerical tools able to model and predict <b>fire evolution</b> is critically important in terms of safety and protection fire fighting and could help in real time disaster management.
	<b>Green Aircraft</b>	ACARE 2020 provides the politically agreed targets for an acceptable maximum impact of air traffic on people and environment, while allowing the constantly increasing amount of air travel. The goals deal with a reduction of exhaust gas and noise. Air traffic will increase by a factor of 3, accidents are expected to go down by 80%. Passenger expense should drop (50%) and flights become largely weather independent. The <b>“Green Aircraft”</b> is the answer of the airframe as well as engine manufacturing industry.
	<b>Virtual Power Plant</b>	Safe production of high quality and <b>cost effective energy</b> is one of the major concerns of Utilities. Several challenges must be faced, amongst which are extending the lifespan of power plants to 60 years, guaranteeing the optimum fuel use and better managing waste.

# PRACE: Support of Science Communities

European Organisations and Research Communities	
<b>EFDA</b>	The European Fusion Development Agreement foresees a huge demand for HPC including tier-0. It is interested in cooperation with PRACE regarding benchmarking and code-scaling and provides the HPC-related requirements for Fusion community.
<b>EMBL-EBI</b>	The Euro Bioinformatics Institute within the European Molecular Biology Laboratory foresees huge demands for HPC resources in the future and is interested in investigating access policies to European tier-0 systems for life scientists.
<b>ENES</b>	The European Network for Earth System Modeling has contributed to the scientific case for HPC in Europe and will continue to promote the involvement of the European climate modelling community in PACE. ENES involvement includes porting of applications on prototype systems of PACE and defining of facility requirements.
<b>ESA</b>	ESA is the European Space Agency. The Space and in particular Earth Observation communities have very demanding HPC applications. ESA is pleased to collaborate with PRACE on specific applications.
<b>ESF</b>	The European Science Foundation is interested to contribute to PRACE, in particular to peer-review process dissemination activities and computer technologies beyond 2010.
<b>MOLSIMU</b>	MOLSIMU, a COST action on Molecular Simulations to Nanoscale Experiments, is offering its support for PRACE by porting their major applications to the prototype systems installed by PACE
<b>Psi-k Network</b>	The Psi-k network is the European Umbrella Network for Electronic Structure Calculations. Several groups within Psi-k are interested to port their ab-initio codes like CPMD, VASP, SIESTA, CASTEP, ABINIT, and Wien 2k on the prototype systems of PRACE.

# PRACE: Support of Research Infrastructures

<b>DEISA</b>	EU- Project	DEISA currently deploys and operates the European Supercomputing Grid infrastructure to enable capability computing across remote computing platforms and data repositories at a continental scale.
<b>HPC-Europa</b>	EU- Project	<p>HPC-Europa is a pan-European Research Infrastructure on HPC providing HPC access and scientific support to researchers in challenging computational activities.</p> <p>HPC-Europa expresses its interest in cooperating in the areas of access technologies and integrated advanced computational services.</p>
<b>OMII-Europe</b>	EU- Project	OMII-Europe is the interoperability project in Europe providing open standards based interoperability components on top of the four major Grid middleware systems in the world.
<b>EGI</b>	EU- Project Prop.	The consortium of EGI aims at establishing a sustainable Grid infrastructure in Europe, coordinating national Grid initiatives.

# Lincei Initiative (2007-2009)



## FORWARD LOOK

### EUROPEAN COMPUTATIONAL SCIENCE: THE “LINCEI INITIATIVE”: FROM COMPUTERS TO SCIENTIFIC EXCELLENCE

Computational sciences and computer simulations in particular, are playing an ever growing role in fundamental and applied sciences. The aim of this Forward Look is to develop a vision on how computational sciences will evolve in the coming 10 to 20 years. Based on a scenario of how this field will evolve and on the needs of the scientific community, a strategy will be presented aimed at structuring software and hardware support and development at the European level.

<http://ccp2007.ulb.ac.be/FL-Lincei.pdf>

# Lincei Initiative: Steering Committee

**Doctor Vassilis Pontikis, *Chair*,**

Commissariat à l'Énergie Atomique, Saclay , Gif-sur-Yvette, France

**Professor Carmen N. Afonso, PESC rapporteur,**

Consejo Superior de Investigaciones Cientifica, Instituto de Optica, Madrid, Spain

**Professor Isabel Ambar, LESC rapporteur,**

Directora Instituto de Oceanografia Faculdade de Ciências da Universidade de Lisboa

**Professor Kenneth Badcock,**

Dept. of Engineering, The University of Liverpool, Liverpool, United Kingdom

**Professor Giovanni Ciccotti,**

Dept. of Physics, Università "La Sapienza", Roma, Italy

**Professor Peter H. Dederichs,**

Institut für Festkörperforschung, Jülich Research Centre, Jülich, Germany

**Doctor Paul Durham,**

Daresbury Laboratory, Warrington, United Kingdom

**Professor Franco Antonio Gianturco,**

Dept. of Chemistry, Università "La Sapienza", Roma, Italy

**Professor Volker Heine,**

Cavendish Laboratory (TCM), Cambridge University, Cambridge, United Kingdom

**Professor Ralf Klessen,**

Institute für Astrophysik, Zentrum für Astronomie der Universität Heidelberg, Heidelberg, Germany

**Professor Peter Nielaba,**

Lehrstuhl für Theoretische Physik, Fachbereich Physik, Universität Konstanz, Konstanz, Germany

**Doctor Simone Meloni, Scientific Secretary,**

Consorzio per le Applicazioni del Supercalcolo per Università e Ricerca - CASPUR, Roma, Italy

# Lincei Initiative: Six Fields Addressed

## **Astrophysics**

- Institut für Theoretische Astrophysik, Heidelberg (DE), Dec. 1st-2nd 2006

## **Fluid Dynamics**

- Daresbury Lab., Warrington (UK), Nov. 29th-30th 2006

## **Meteorology and Climatology**

- Swiss Supercomputing Centre, Manno (CH), Jan. 27th 2007

## **Life sciences**

- Chilworth Manor, Southampton (UK), Nov. 19th-21st 2006

## **Material Science and Nanotechnology**

- Jülich Research Centre, Jülich (DE), Nov. 13th-14th 2006

## **Quantum Molecular Sciences**

- Accademia dei Lincei, Rome, Nov. 25th-26th 2006

- State of infrastructure for scientific computing
- Needs in relation to future challenges, in 10-20 year timeframe

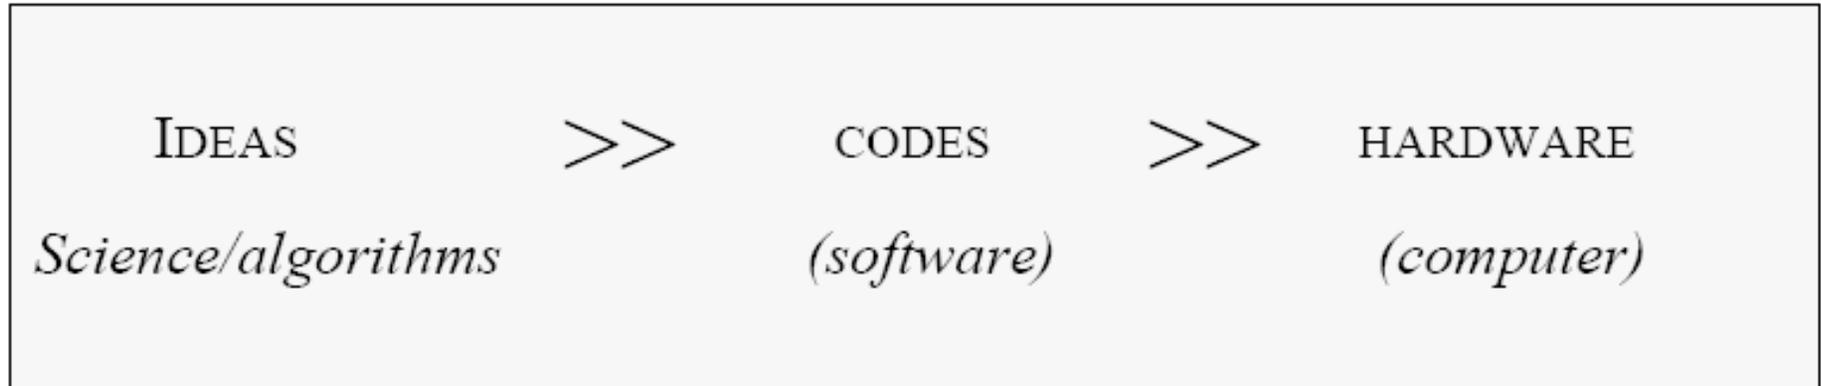
# Some EU Scientific and Engineering Codes

(From Lincei Forward Look Report (for the ESF))



Name	Scientific Area	Brief Description	Licensing	Users
ABINIT	Condensed Matter	DFT+PW+Pseudopotentials	Free	~1000
ESPreSo	Condensed Matter	coarse grained off-lattice	Free	~20 groups
VASP	Condensed Matter	DFT+PW+Pseudopotentials	Licensed	800 li
CP2K	Condensed Matter	DFT-(gaussian+PW)+classical	Free	~100
CPMD	Condensed Matter	DFT+PW+Pseudopotentials	Licensed, free acad.	>1000
Wien2K	Condensed Matter	Full-electrons Augmented PW	Licensed	~1100
Quantum Espresso	Condensed Matter	DFT+PW+Pseudopotentials	Free	~700
Code_Aster	Engineering	Mechanical and thermal analysis	Free	300 (EDF)
			22k downl.	
Code_Saturne	Fluid Dynamics +heat transfer+combustion	Incompressible+expandable	Free	80 (EDF) + 25 groups
OpenFOAM	Fluid Dynamics +Structural Mechanics	Finite volume on unstructured grid	Free, fee for support	~2000
Salome	framework for multiphysics	Used in engineering	Free	50 (EDF) + 21 groups
COSMO-Model	Climatology, Meteorology and scientific research	Operational Weather forecasting	Special agreement	7 Centres 80 groups

# Importance Hierarchy



# Comments from FL-Lincei-Report

- Current [application] software is very complex
- Typical size is 400000 lines of code and 2500 routines/classes
- Large number of variables pass through the code in obscure data flow
- Few strictly object oriented (OpenFOAM, C++, CP2K, FORTRAN95)
- Will be confronted with a software sustainability crisis
- Will be very difficult to adapt most existing complex codes to the coming massively parallel computers
- Structure of many of the codes strongly dependent on the parallel programming paradigm adopted in the early stage of the development
- Current shift from hundreds to tens of thousands of CPUs will require a change in the parallelization scheme
- Very difficult to implement in such very complex community codes

## Findings

- Bottleneck is the support to software, effort mainly focused on Hardware
- Less support is given to the writing, maintenance and dissemination of sc. codes
- Scientific computer programs do not comply with best practices in programming
- Successful efforts in all the technical areas required to support scientific computing: **hardware, system and application software**

## Recommendations

- National science funding agencies in Europe must undertake a coordinated and sustained effort in scientific software development
- Set up a Computational Sciences Expert Committee (CSEC) attached to ESF which would speak for the whole community of computational sciences.
- Its purpose would be to start setting up a durable plan for European cooperation in each of the fields of science using computers

# Example 1: MD for Radiation Hard Materials

Ian J. Bush, Ilian T. Todorov, CCSRC Daresbury, UK

DL-POLY3 classical molecular dynamics

First time on more than 1000 processors

Radiation damage in a fluoritized Zirconium pyrochlore

100 keV recoil of one Uranium atom after alpha decay

15 million particles, supercell very large

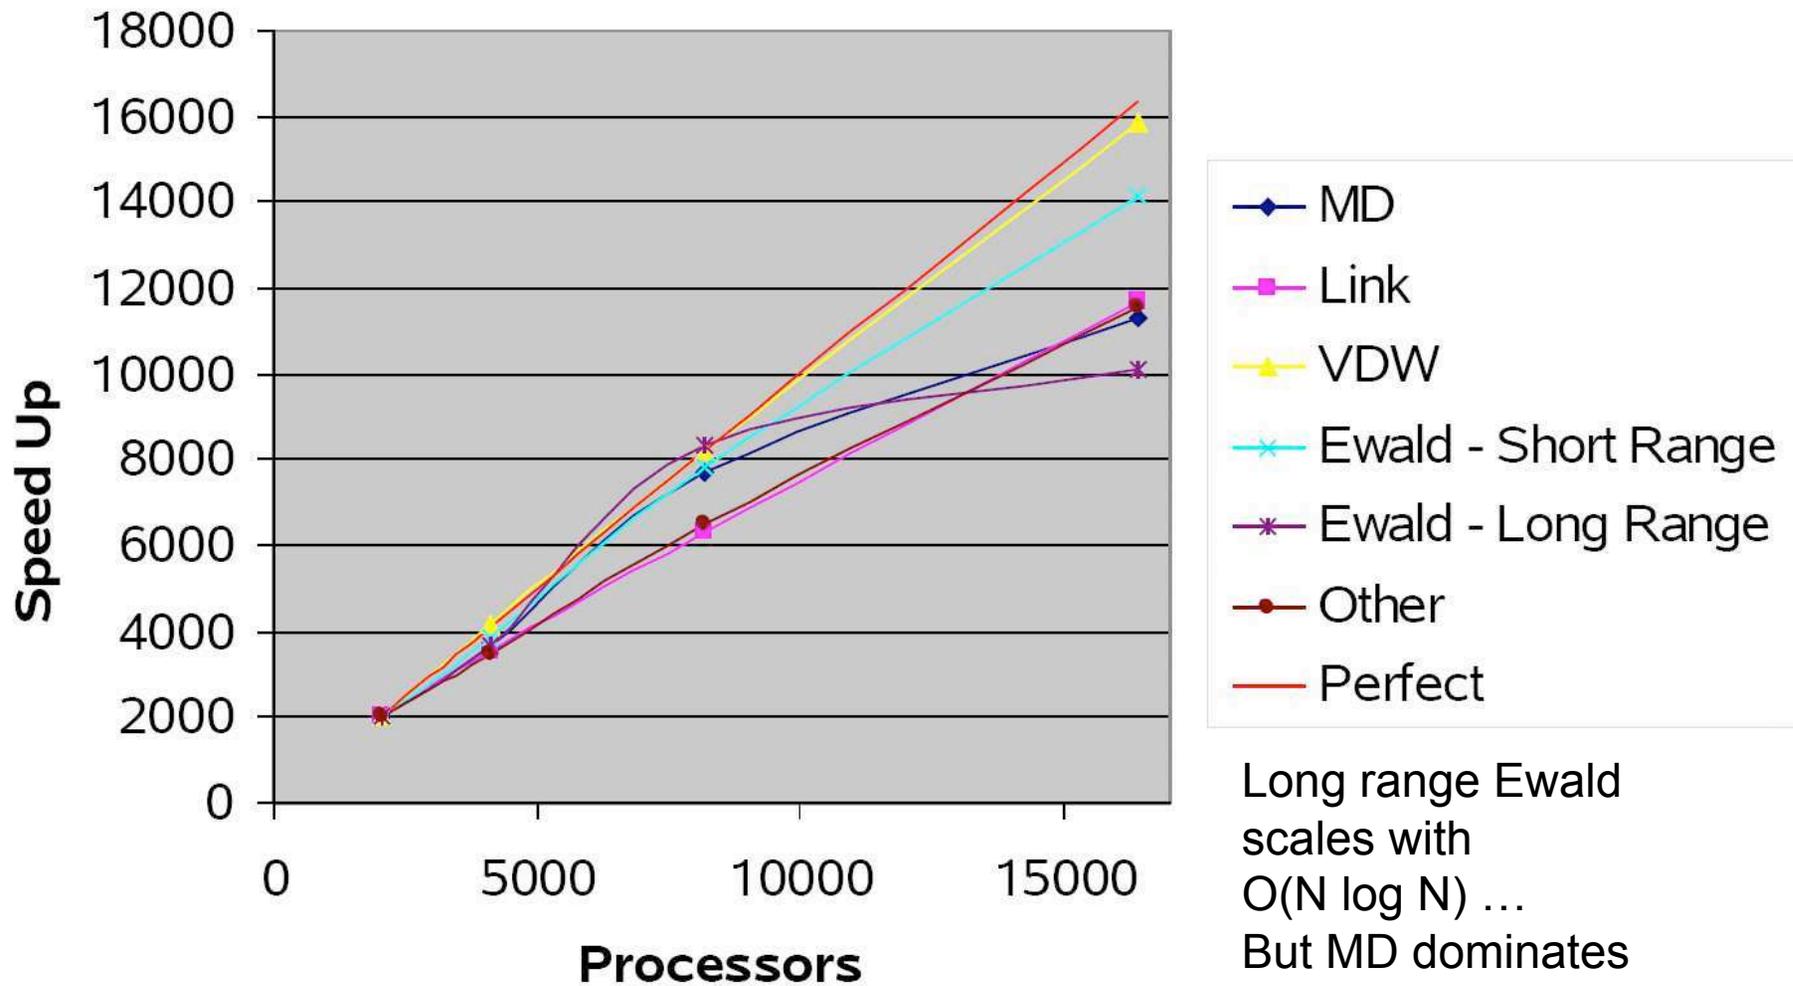
Forces: short range, van der Waals, Coulomb

Smooth particle-mesh Ewald algorithm → FFT

Implementation on BGL

# Scaling DL\_Poly3

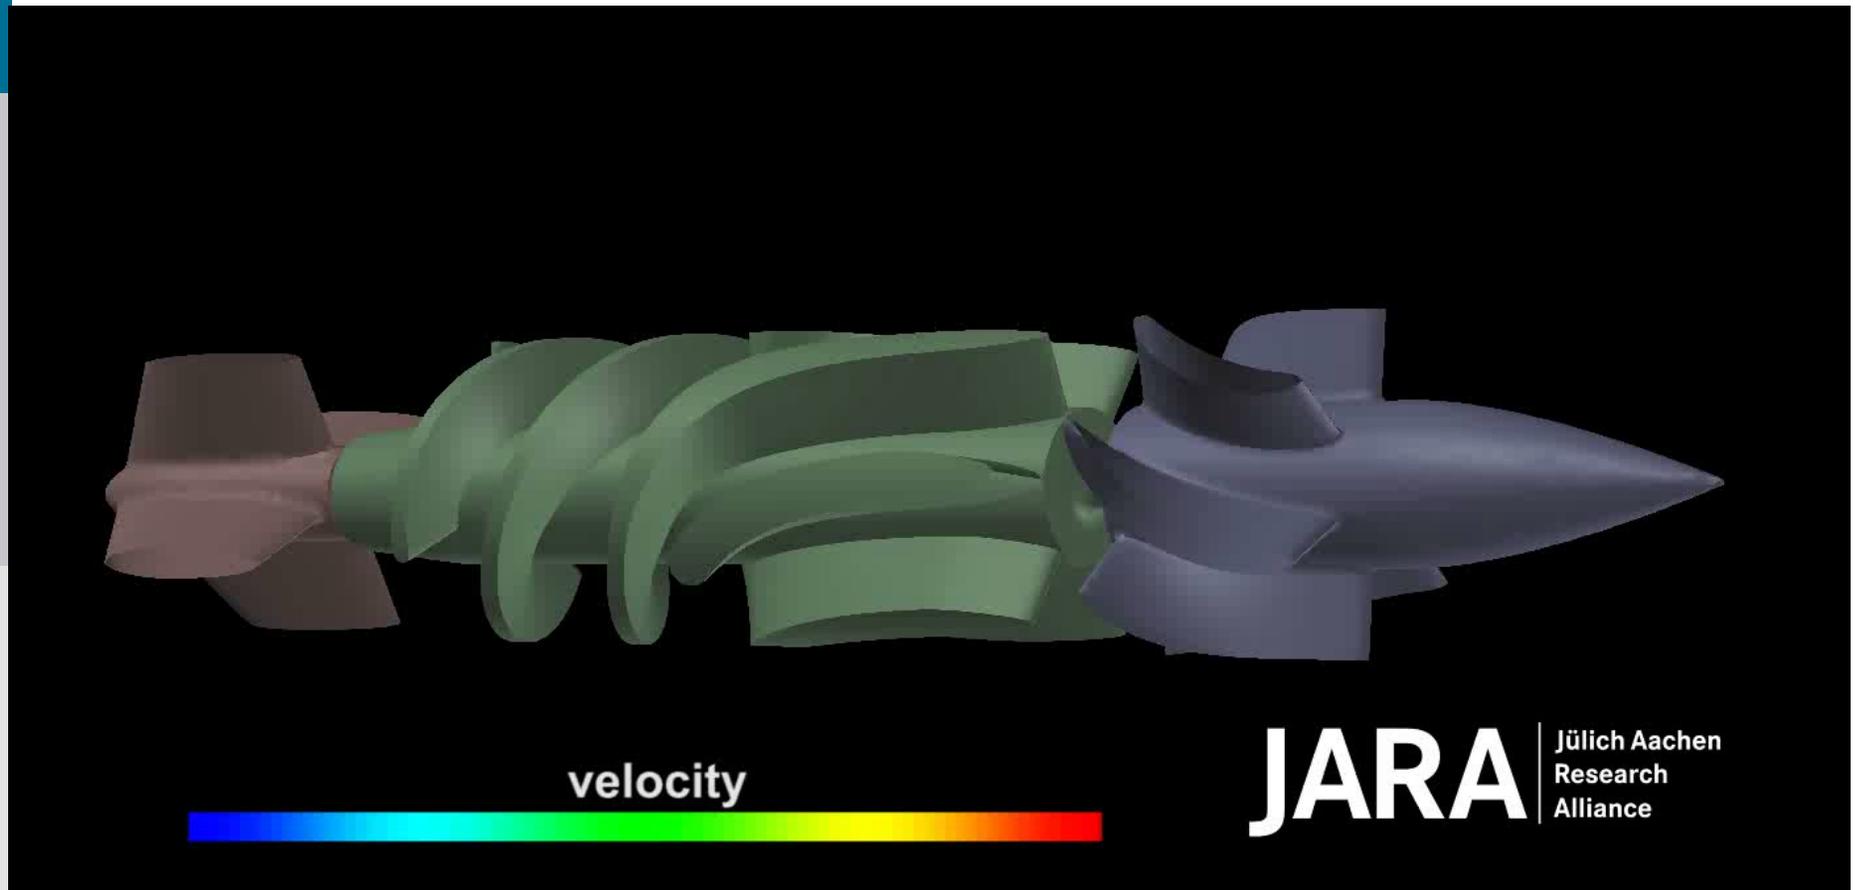
## Substantial improvements by performance analysis tool Scalasca



Long range Ewald scales with  $O(N \log N)$  ...  
But MD dominates

## Example 2: Engineering – Biomedical Flows

### Simulation of Blood Flow in a Ventricular Assist Device Marek Behr, RWTH Aachen



# Code + Analysis tools → large Improvements

## XNS CFD solver

- 3D space-time simulation of MicroMed DeBaKey axial blood pump
- 4 million elements
- Partitioning by Metis graph partitioning package
- Incompressible Navier-Stokes Eq.
- FEM, GMRES, 3 time steps, 4 Newton-Raphson iterations

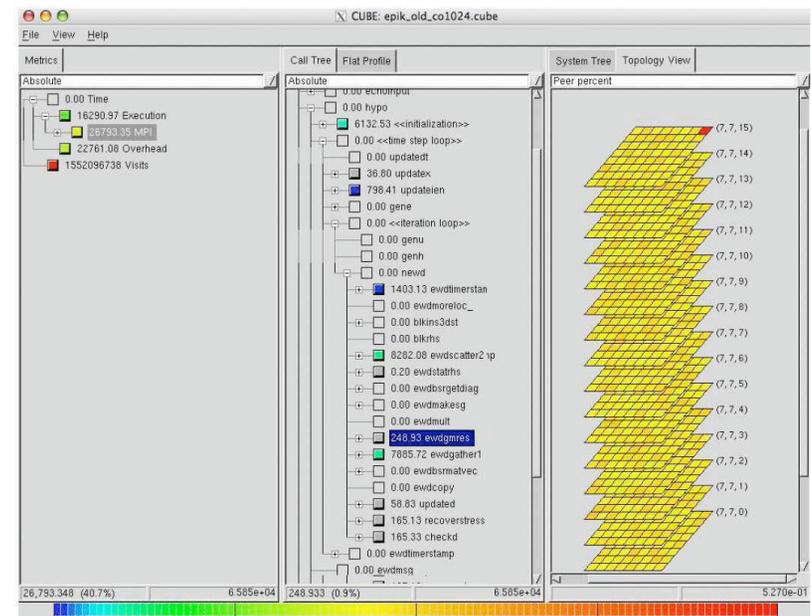
## Analysis by SCALASCA package

(Bernd Mohr, Felix Wolf)

- Too many MPI\_Sendrecv with zero-byte transfers
- Good speedup to 4096 processors

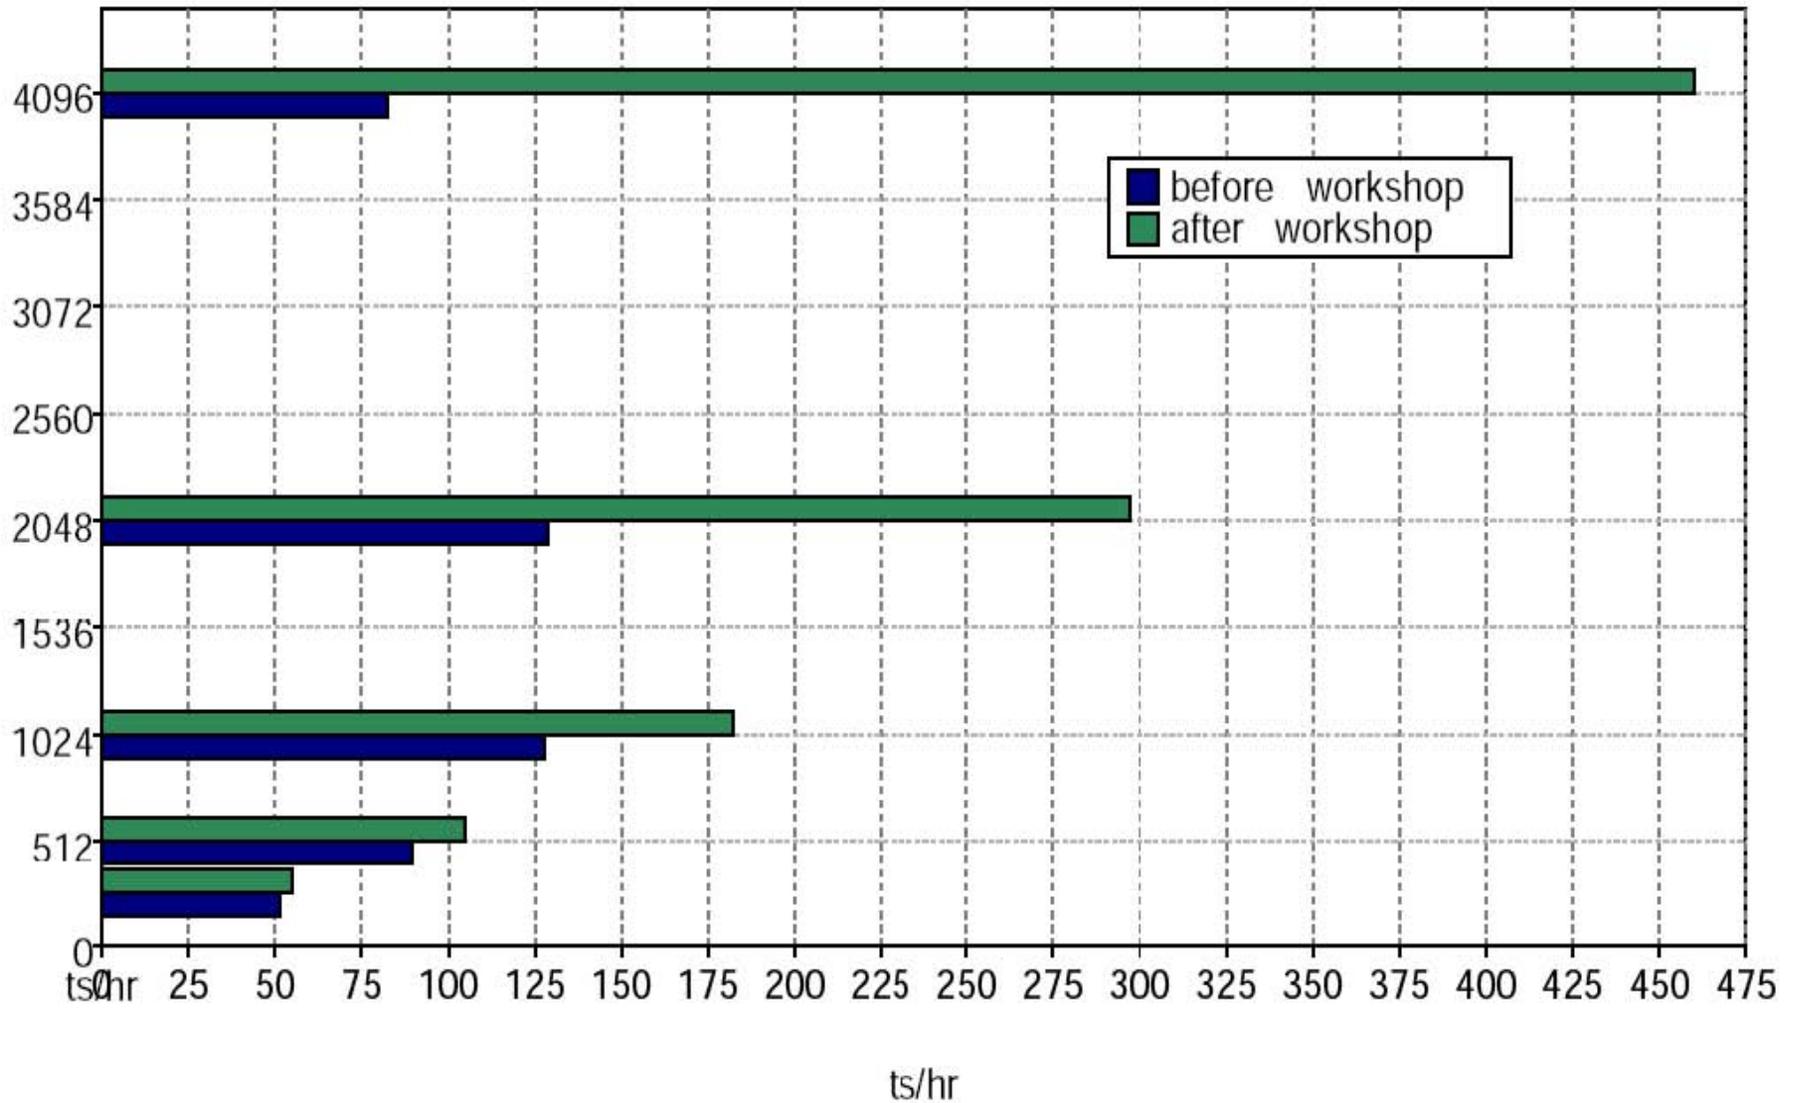
## Still: strong load imbalance in GMRES

- process with highest rank overloaded



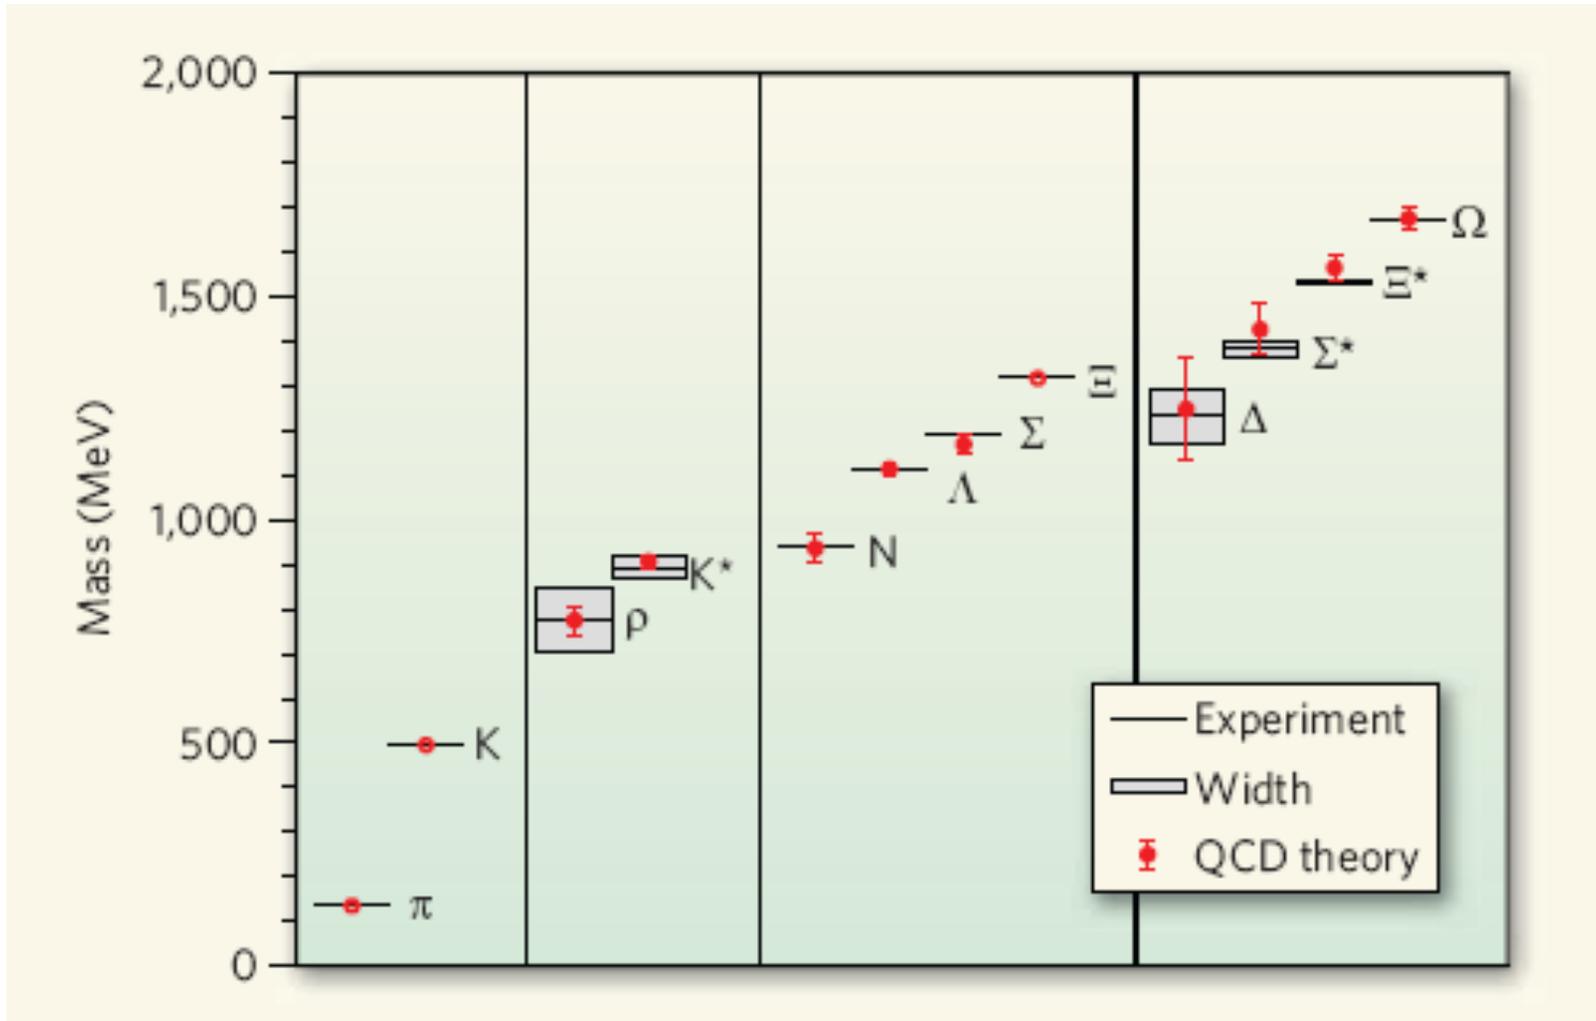
# After Improvements

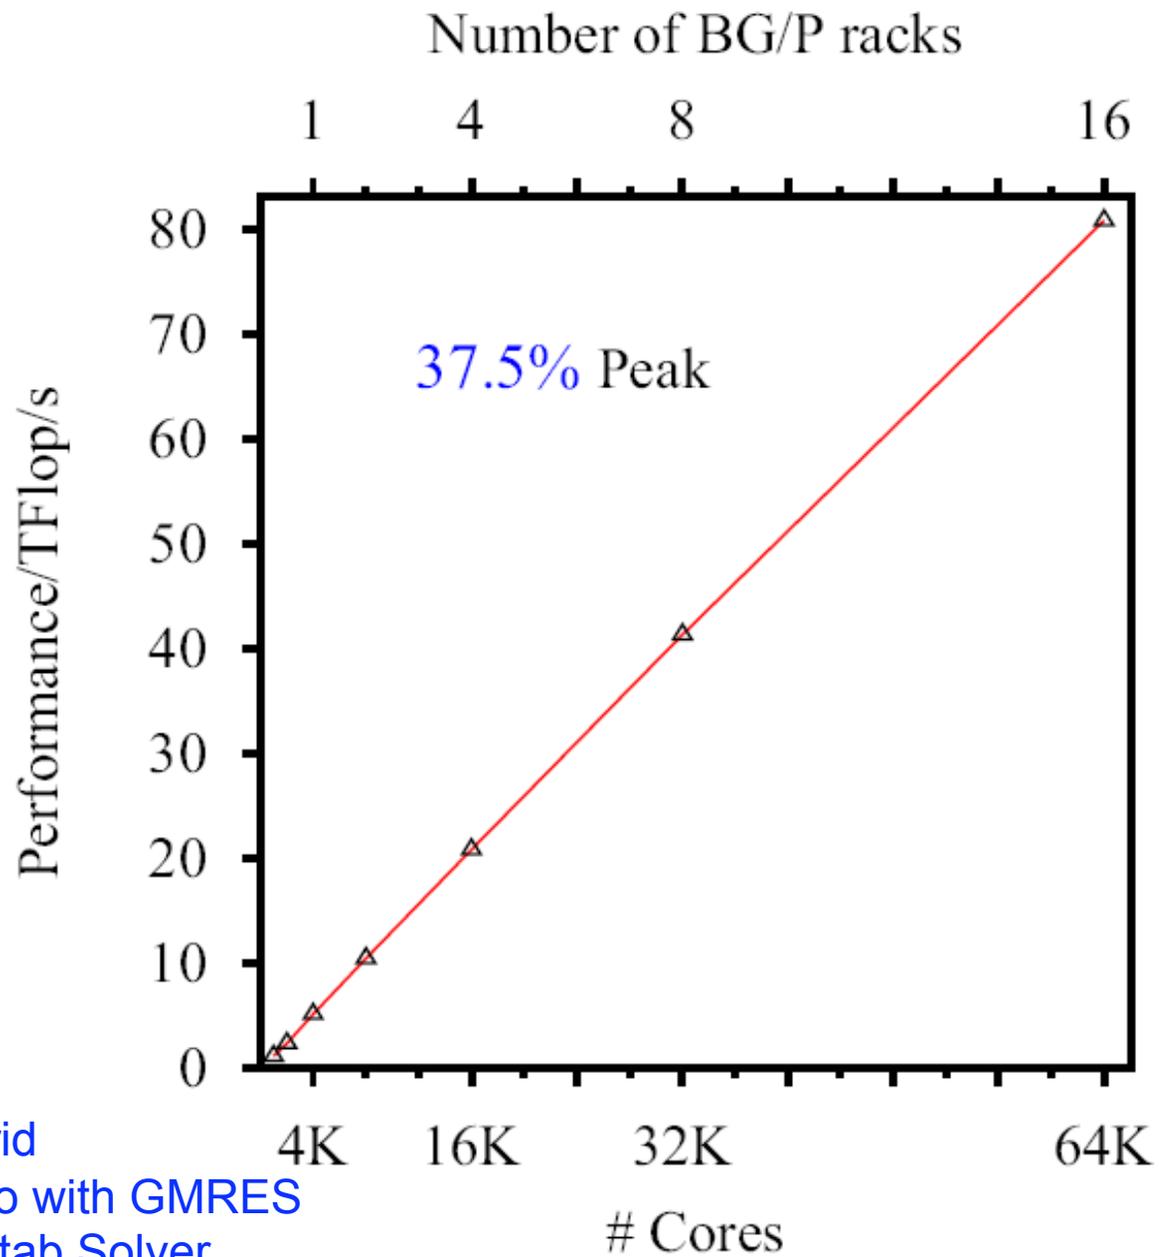
Scaling Workshop Blue Gene Juelich



# Example 3: Theoretical Particle Physics

Fodor et al. 2008: Validation of Quantum Chromodynamics  
Among 10 SCIENCE-breakthroughs of 2008





**Improvements  
through low-level  
programming**

Code: Hybrid  
Monte Carlo with GMRES  
and BiCGStab Solver

# HPC Software & Benchmark Codes

Disclaimer:

List of software is a selection and not comprehensive

# HPC Software Challenges

- **Extreme scalability**
  - Exascale: number of cores beyond any reasonable, manageable limit
- **Extreme complexity**
  - Machine architecture gets more complicated instead of becoming simpler (KISS!)
- **Little to not existing fault-tolerance in existing base software**
  - e.g. MPI, OpenMP, schedulers, ...
- **Rapid grows in system size /change in HW architecture**
  - SW developers cannot keep pace

## MPI

- Open MPI European partners HLRS, INRIA, Univ. Jena, Univ. Chemnitz, TU-Dresden, BULL
- MPICH-V fault tolerance, MPI Madeleine (INRIA)
- HLRS, Bull, NEC, (ZIH, JSC) participating in MPI-3

## OpenMP

- EU ARB members: EPCC, RWTH, (BSC?)
- BSC Mercurium compiler framework

## Pragma-based task parallelism

- SuperScalar (BSC)
  - Subject in future EU proposals (EU ITEA2 H4H, FP7 FET EXACT)
- HPMM (INRIA/CAPS)

## Parallel Object-oriented

- Kaapi (C++) + PROACTIVE (Java) (INRIA), PM2 (LaBRI, INRIA)

## Numerical Middleware

- superLU (INRIA), MUMPS (ENSEEIH)
- Scilab (Digiteo)

## Benchmarks

- DEISA
  - 14 full applications
  - HPCC
- PRACE
  - 20 full applications
  - Various low-level
- EPCC micro benchmarks

# Application Software Benchmarks: DEISA



**Astrophysics:**

[GADGET](#), [RAMSES](#)

**CFD and combustion:**

[FenflOSS](#)

**Earth sciences and climate research:**

[ECHAM5](#), [IFS](#), [NEMO](#)

**Life sciences and informatics:**

[NAMD](#), [IQCS](#)

**Materials science:**

[CPMD](#),

[QuantumESPRESSO](#)

**Plasma physics:**

[GENE](#), [PEPC](#)

**Quantum chromodynamics:**

[BQCD](#), [SU3\\_AHiggs](#)

**DEISA benchmark represents major EU HPC applications**

# Application Software Benchmarks: PRACE

(see White Paper by Peter Michielse)

<b>Application</b>	<b>Application area</b>	<b>Application</b>	<b>Application area (to be considered)</b>
QCD	Particle physics	AVBP	Computational fluid dynamics
VASP	Computational chemistry, condensed matter physics	CP2K	Computational chemistry, condensed matter physics
NAMD	Computational chemistry	GROMACS	Computational chemistry
	life sciences	HELIUM	Computational physics
CPMD	Computational chemistry, condensed matter physics	SMMP	Life sciences
		TRIPOLI4	Computational engineering
Code_Saturne	Computational fluid dynamics	PEPC	Plasma physics
		RAMSES	Astronomy and cosmology
GADGET	Astronomy and cosmology	CACTUS	Astronomy and cosmology
TORB	Plasma physics	NS3D	Computational fluid dynamics
ECHAM5	Atmospheric modelling		
NEMO	Ocean modelling		

# Porting Codes

Application	MPP-BG	MPP-Cray	SMP-TN-x86	SMP-FN-pwr6	SMP-FN+Cell	SMP-TN+vector
QCD	Done	Done		Done		
VASP	Done			Done	Stopped	Yet to start
NAMD	Done	Done		Done	Yet to start	
CPMD	Done			Done	Done	Yet to start
Code_Saturne	Done	Done		Done	Stopped	Done
GADGET	Done		Done	Done		
TORB	Done			Done	Yet to start	
ECHAM5	Stopped	Done	In progress	Done		Yet to start
NEMO	Done	Done		Done		In progress
CP2K	Done	Done		Done		
GROMACS	Done	Done		Done		
NS3D		Yet to start	Done	Yet to start		Done
AVBP	Yet to start		Done	Done		
HELIUM	In progress	Done		Done		
TRIPOLI_4	Yet to start		Done			
PEPC	Done	Done		Done		
GPAW	Done	Done		Done		
ALYA					Done	
SIESTA					Done	
BSIT					Done	

**Table 4: Summary on porting efforts for benchmark codes and prototype architectures.**

## System / cluster tools

- Benchmarking: JuBe (JSC)
- Resource allocation: OAR (INRIA)
- System monitoring: LLview (JSC)
- Cluster middleware: ParaStation (ParaStation-Consortium: ParTec, JSC, Karlsruhe, Heidelberg, Wuppertal)

## Grid Middleware

- UNICORE (UNICORE forum, JSC, ...)
- GLite (CERN, LHC)
- dCache (DESY)
- DIET: grid RPC system (INRIA, CNRS, LIP/ENS Lyon, ...)

## Programming tools

- Debugging: DDT (Allinea)
- MPI debugging: Marmot (ZIH –TU-Dresden / HLRS)

## Performance

- OPT (Allinea)
- Paraver/Dimemas (BSC)
- KOJAK/Scalasca (JSC)
- Vampir (ZIH-TU-Dresden)
- Periscope (TU Munich)
- SlowSpotter/ThreadSpotter (Acumem)

## European tool integration projects

- EU ITEA2 ParMA project (17 partners, FR, DE, ES, UK)
- German BMBF SILC

# Existing Working Collaborations

- **MPI standardization and Open MPI project**
- **OpenMP standardization**
- **Global Grid Community**
- **Example: Performance tools community**
  - Voluntary US participation in EU APART WG (1998-2004)
  - Common Dagstuhl seminars (2002, 2005, 2007, 2010)
  - CScADS workshops (2007, 2008, 2009)
  - Collaborating collaboration projects
    - POINT (UO, ICL, NCSA, PSC)
    - VI-HPS (RWTH, ZIH, JSC, ICL)
  - New: DOE ASCR funding for non-U.S. partners!
    - PRIMA (UO, JSC): 2009-2012
    - “PTP++” (IBM, LANL, ORNL, JSC, Monarch): 2009-2012

# Collaboration and Funding

## Lessons learned

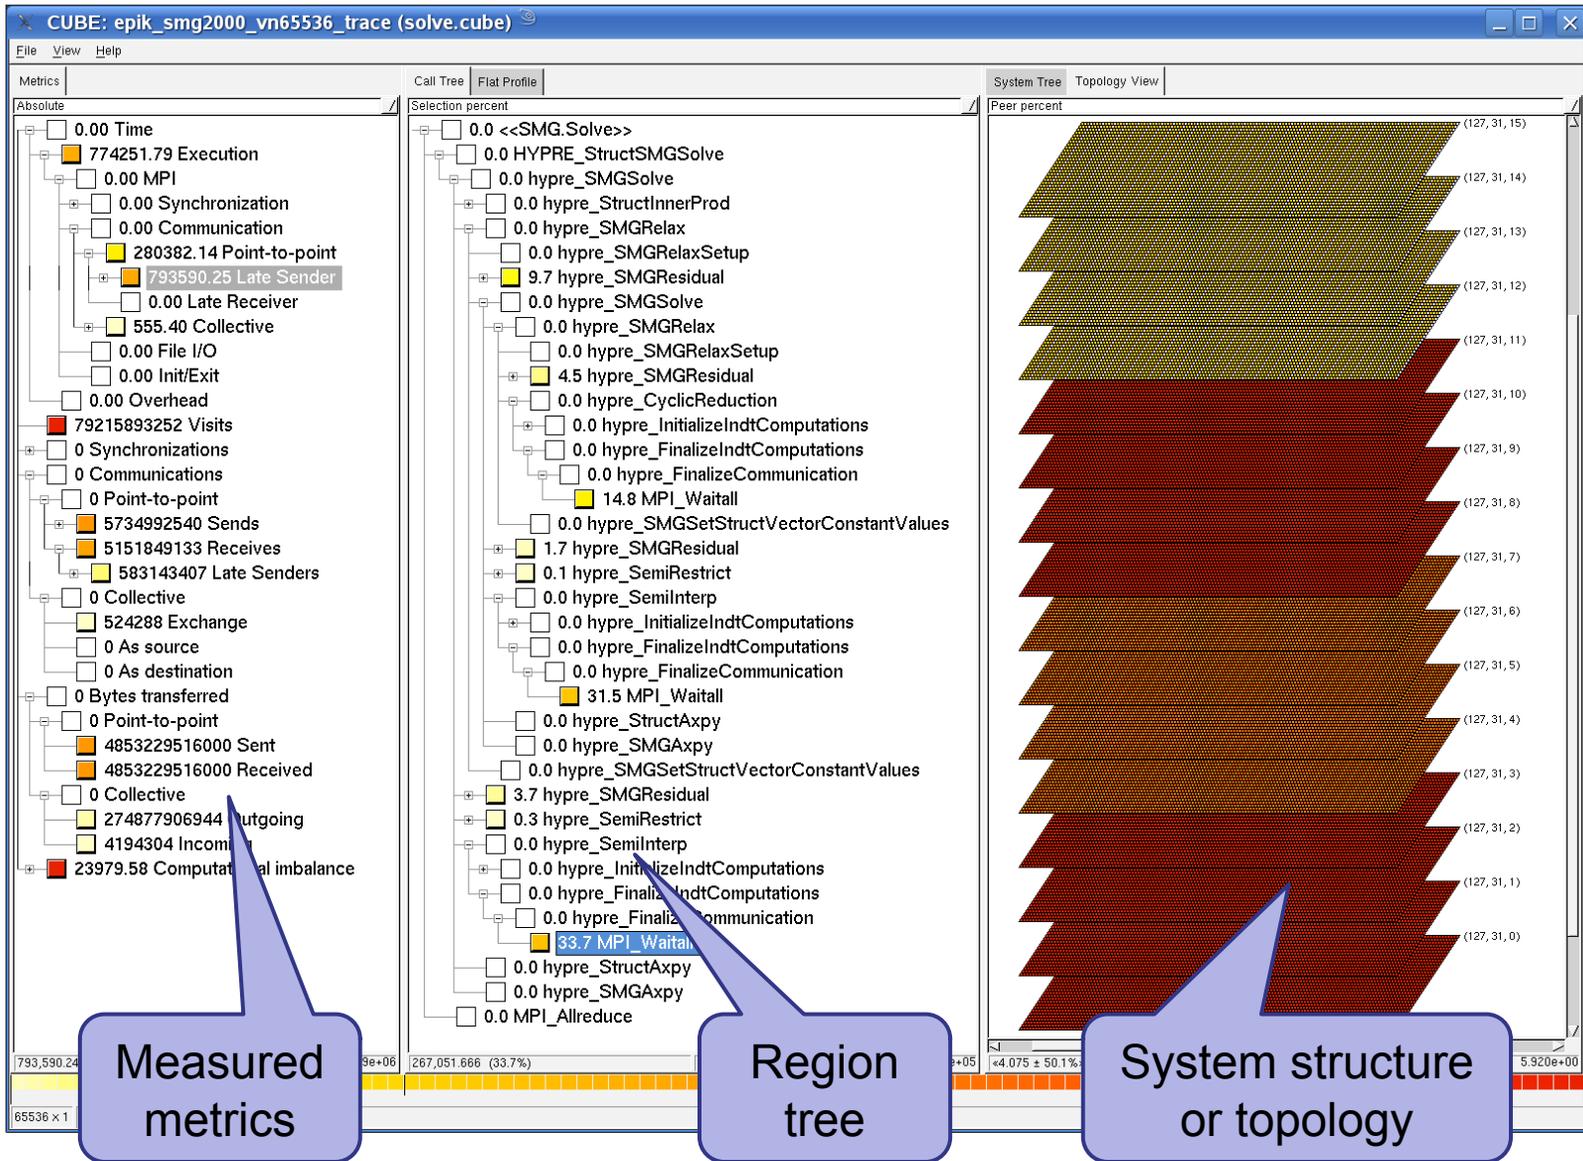
- Collaboration projects need
  - Strong leadership + Funding
  - Examples of failures: PTOOLS, OSPAT, ....
- Bottom-up, technology-driven, friendship approaches work much better than top-down, politically-driven, mandated ones
- Top-down provides funding
- **Need combined approach: bottom-up meets top-down and long-term commitments of funding agencies**

## Proposal

- Local (US, EU, Asian) funding programs need to allow to fund additional global partners
- New global funding for networking (coordination, dissemination, synchronization efforts)

- **Scalable Analysis of Large Scale Applications**
- **Follow-up project to well-known KOJAK project**
- **Installed on many leadership class systems (EU, US)**
- **Successfully used on 65536 cores**
- **Integration with TAU and Vampir toolsets**
- **Approach**
  - **Instrument** C, C++, and Fortran parallel applications
    - Based on MPI, OpenMP, SHMEM, or hybrid
  - **Collect** event traces (or callpath profiles)
  - **Search** trace for event patterns representing inefficiencies **in parallel**
  - **Categorize and rank** inefficiencies found
- **<http://www.scalasca.org/>**

# Trace analysis SMG2000@64k



## ParaStation Cluster Middleware

### ParaStation V5:

- Multi-core aware cluster operating and management software
- Open source → GPL licensed
- ParaStation Consortium: ParTec, Forschungszentrum Jülich, Universities of Karlsruhe, Heidelberg, Wuppertal
- Daemon based
- MPI-2
- Grid Monitor (full awareness of complete cluster status)
- IB, Ethernet, Myrinet, just everything

# ParaStation Research

(Projects funded by Federal Ministry of Education & Research)



## ISAR project (2008 – 2011)

Integrated system and application analysis for massive parallel computer

Members:

Uni Munich, Leibniz Compute Center (LRZ),  
Compute Center Garching (Max-Planck), ParTec, IBM

## D-Grid 2, (2007-2010)

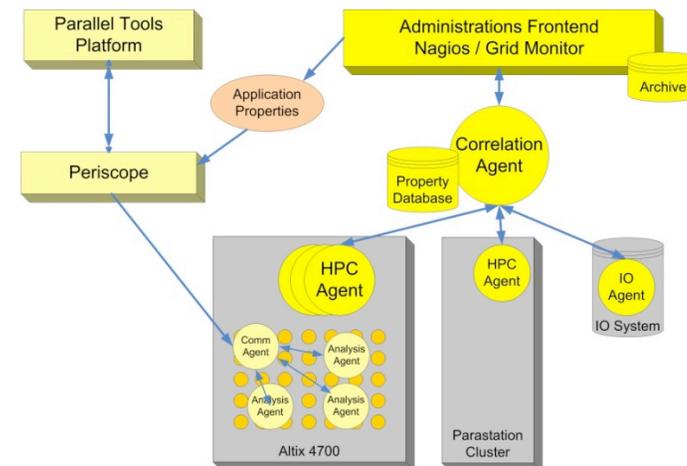
German Grid initiative  
(funded by BMBF)

## ee-Clust project (2008 – 2011)

Energy efficient cluster computing

Members:

Uni Heidelberg, TU Dresden,  
Research Centre Julich, ParTec



- Goals
- Scalable cluster OS
- Fighting OS-jitter

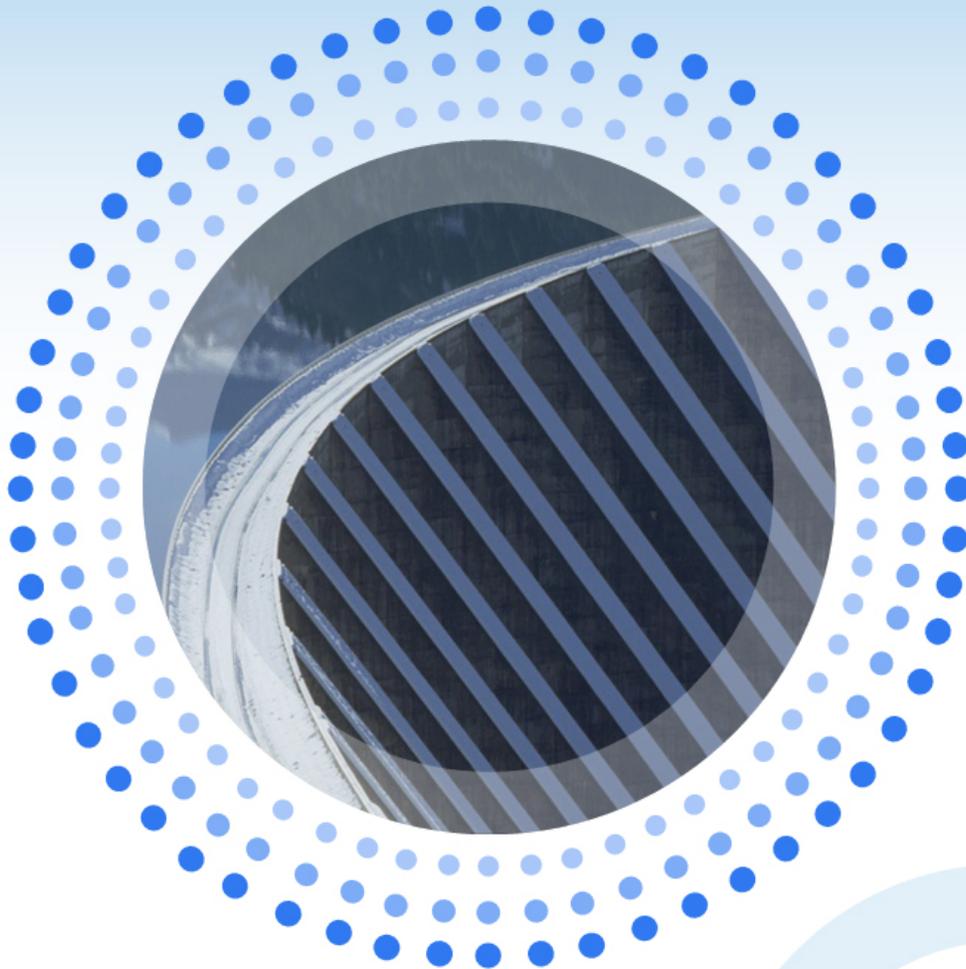
# Plans for Exascale Activities and Initiatives in Europe

1. EESI (International HPC Software Coordination and Development)
2. COSI-HPC Proposal (HPC-Software - Coordination)
3. Lincei Initiative (Comp. Science)
4. CECAM (Comp. Science)
5. PRACE (ESFRI-Infrastructure)



# Establishing the European Exascale Software Initiative

Contribution by  
Jean-Yves Berthou, EDF R&D





## Context: International Exascale Software Project

SC'08 (nov. 2008) : DOE/NSF/DOD launched the International Exascale Software Project (IESP)

*Plan to build an international partnership that joins together **industry, the HPC community (CS and Apps), and production HPC facilities** in a collective effort to design, coordinate, and integrate software for leadership-class machines.*

*Specifically, engagement in the following activities should be started:*

- Build international collaborations in the areas of high-performance computing software and applications.*
- Development of open source systems software, I/O, data management, visualization, and libraries of all forms targeting tera/peta/exascale computing platforms,*
- Research and development of new programming models and tools addressing extreme scale, multicore, heterogeneity and performance,*
- Cooperation in large-scale systems deployments for attacking global challenges,*
- Joint programs in education and training for the next generation of computational scientists.*
- Vendor engagement to coordinate on how to deal with anticipated scale.”*



# European Exascale Software Initiative (EESI)

## Main goals

- Building and promoting European position inside the IESP initiative
- Identifying Grand Challenge applications, from academia and industry, with a strong economical, societal and/or environmental impact that will benefit of Petaflop capacities in 2010 and Exaflops in 2020
- Identify critical software issues for Peta-ExaScale systems
- Building a European/US/Japan program in education and training for the next generation of computational scientists
- **Output : Proposition of a strategic research action agenda for Peta-Exascale Software and Grand Challenge applications at the European level coordinated with US and Japanese agendas**



## European Exascale Software Initiative (EESI) Preparatory phase Project Proposal – 12 months

Establish a European position inside the IESP initiative

- o Promote and represent the European position
- o Influence on decisions and actions
- o Synchronize European agenda with other international agenda

Contribute to the International dialogs between US and Europe and Japan and Europe and be a bridge between some EU organizations including the European commission and IESP

Identify main HPC European actors both at end users level and at academic level

Define and implement the organization and governance rules of EESI

Identify main European HPC existing or planned projects

Built a first European and international vision of the on-coming HPC challenges and work to achieve



# European Exascale Software Initiative (EESI)

## Preparatory phase Project Proposal – 12 months

Submitted to ICT 2009.9.1 International cooperation a) Support to Information Society policy dialogues and strengthening of international cooperation

Partners	Country	Contact		Title
EDF (leader)	France	Jean-Yves	Berthou	EDF R&D Information Technology Program Manager
		Jean-François	Hamelin	EDF R&D Information System Director
GENCI, FR	France	Catherine	Rivière	Chairman and CEO of GENCI
		Virgine	Mahdi	
INRIA	France	Frank	Cappello	Director of the joint INRIA/NCSA laboratory
EPSRC	UK	Jane	Nicholson	High End Computing & E-Science Program Manager
Forschungszentrum Jülich GmbH	Germany	Thomas	Lippert	Director of Institute for Advanced Simulation, Head of Jülich Supercomputing Centre
		Bernard	Mohr	
BSC	Spain	Mateo	Valero	Director of BSC
		Sergi	Girona	Operations Director BSC
NCF	Netherlands	Patrick	Aerts	Director of NCF
		Peter	Michielse	Deputy Director of NCF
Arrtic	France	Thierry	<u>Bidot</u>	



# European Exascale Software Initiative (EESI) Preparatory phase Project Proposal – 12 months

## Supporting partners

### US

IESP, Executive Director, J. Dongarra

U. Urbana-Champaign, Deputy Director for Research, B. Gropp

U. Urbana-Champaign, Professor, M. Snir

### Japan

Tokyo Institute of Technology, Professor & Director Research Infrastructures Division

GSIC, Satoshi Matsuoka

### Europe

PRACE, Current Chairman of the Initiative Management Board, Jane Nicholson

European Science Fondation, the Physics and Engineering Sciences Unit, Science Officer, Dr Thibaut Lery

European Network for Earth System modelling, Chairman of the Scientific Board, S. Jousaume

TERATEC, Chairman, C. Saguez

ORAP, Chairman of the Scientific Council, JC André

Daresbury Lab., Acting Director CS & E dpt., R. Blake

CERFACS, Director, JC André

### Industry/Editor

TOTAL, Scientific Director, JF Minster

SNECMA, Vice President Engineering & Technology, P. Thouraud

NAG, Chief Tech Officer/Vice President HPC Business, M. Dewar/A. Jones



## European Exascale Software Initiative (EESI) Implementation phase (draft)

### Building a research agenda and directions for future

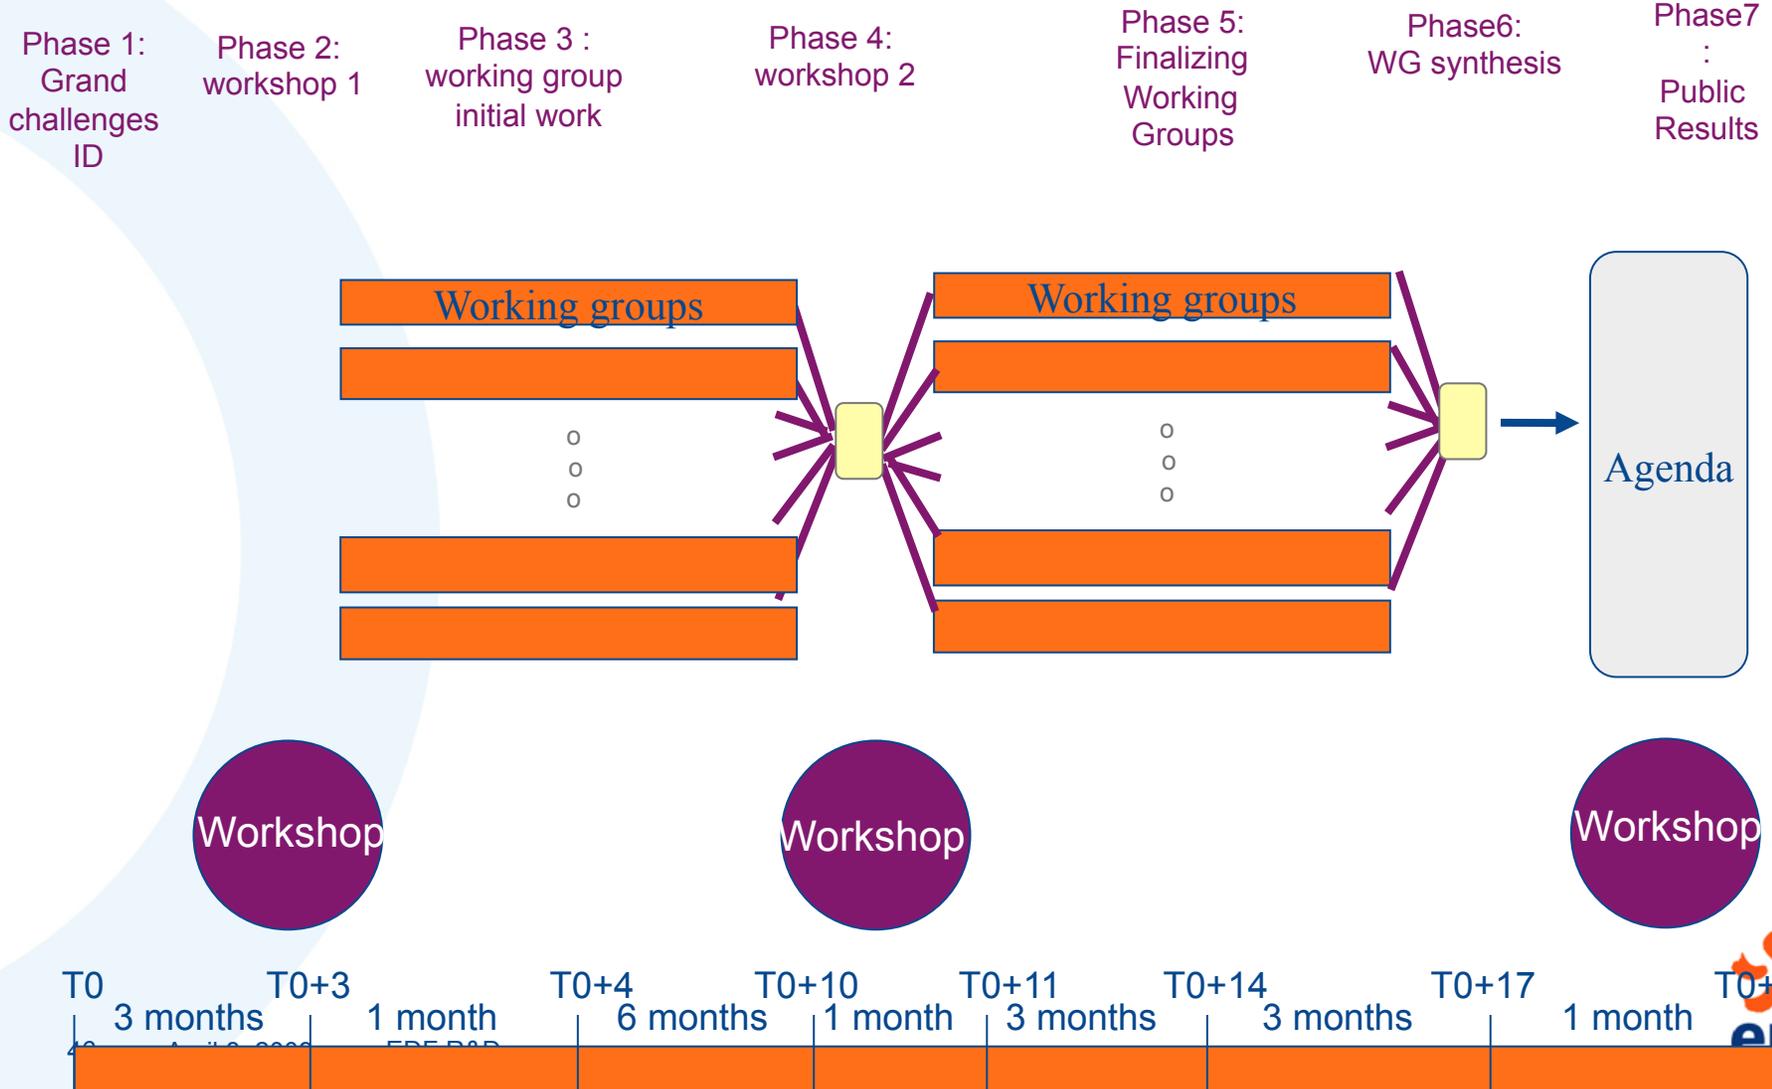
- Identifying Grand Challenge applications, from academia and industry, with a strong economical, societal and/or environmental impact that will benefit of Petaflop capacities in 2010 and Exaflop around 2020
- Identify critical software issues for Peta-ExaScale systems
- Building a EU/US/Japan program in education and training for the next generation of computational scientists
- Proposition of a strategic research action agenda for Peta-Exascale Software and Grand Challenge applications at the European level coordinated with US and Japan agendas



# European Exascale Software Initiative (EESI)

## Implementation phase – 18 months (draft working program)

Input from EESI Preparatory Phase : identification of keyplayers (End user communities, techno. providers, ...)





# COSI-HPC

(proposal, lead by CSC-Finland)

- The Coordination for Software Initiatives in HPC (COSI-HPC) project is designed to promote key elements in an innovation and service ecosystem around the future European Petascale computing research infrastructure (RI).
- Set of actions aimed at coordinating activities in the area of software engineering and software services for large-scale computing, targeting the planned European Petascale facilities as well as future Exascale systems.
- Coordination of existing and future research and industry initiatives such as PRACE, DEISA, PROSPECT, and STRATOS
  - Analysis of HPC software activities in Europe
  - Building up a software community for HPC
  - Address future software challenges

# 3

## Lincei Initiative

Contribution by Michel Marechal

**EUROPEAN  
SCIENCE  
FOUNDATION**  
SETTING SCIENCE AGENDAS FOR EUROPE

### FORWARD LOOK

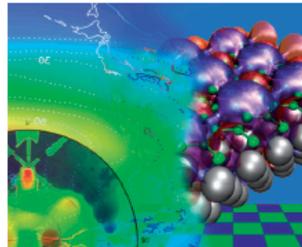
EUROPEAN COMPUTATIONAL  
SCIENCE: THE "LINCEI INITIATIVE":  
FROM COMPUTERS TO  
SCIENTIFIC EXCELLENCE

Computational sciences and computer simulations in particular, are playing an ever growing role in fundamental and applied sciences. The aim of this Forward Look is to develop a vision on how computational sciences will evolve in the coming 10 to 20 years. Based on a scenario of how this field will evolve and on the needs of the scientific community, a strategy will be presented aimed at structuring software and hardware support and development at the European level.

Computational approaches are becoming an increasingly important tool in modern science. Computational sciences have reached such a level of importance as to now be considered the third pillar of science, after experimental and theoretical approaches.

The complexity of the modern codes has caused a transition. A few years ago, each computational group had its own "home-brewed" software. At present, an increasing number of groups rely on the availability of such codes. In the field of computational sciences Europe is playing a leading role, also thanks to the availability of these codes. To consolidate its position, a scientific software must be developed at transnational level, playing the same role as large European facilities (i.e. neutron, X-ray) do for experimental research.

The aim of this Forward Look is to develop a vision



The Organising Committee of this Forward Look adopted web based tools to publish preliminary reports and to organise the day-by-day life of the initiative.

**ESF:** European Science Foundation does coordinate National Research funding organizations in Europe

80 members in 30 countries  
<http://www.esf.org>

A **Forward Look** has been set up by ESF Panel of 12 high level computational scientists has produced a report

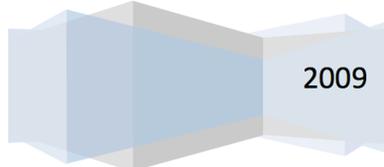
<http://ccp2007.ulb.ac.be/FL-Lincei.pdf>

## Recommendations (I)

- National science funding agencies in Europe undertake a coordinated and sustained effort in scientific software development, including documentation, updating, maintenance and dissemination.
- This necessarily implies the means for training and cooperation.
- **Restructure and federate**, within an European-scale infrastructure, existing and expanded activities on scientific software and other forms of cooperation and dissemination in Europe through **European Computational Collaborations specific to each scientific area**.
- This would be guided by active research scientists and deliver the infrastructural services to the working scientists.

European Science Foundation

European Computational Science  
Forum: The "Linco Initiative":  
from computers to scientific  
excellence



One such example:

- CECAM upgrade
- (multinode, multi-disciplines

CECAM is organizing code developers in condensed matter

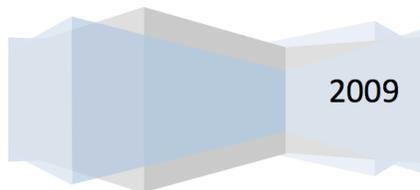
<http://www.cecarn.org/>

- **Recommendations (II)**

- To achieve those goals, it is proposed to set up a **Computational Sciences Expert Committee (CSEC) attached to ESF** which would **speak for the whole community of computational sciences**.
- Its purpose would be to start setting up a durable plan for European cooperation in each of the fields of science using computers.
- It would address the policy issues involved, and work with national and European organisations to optimize the development of scientific computing in Europe.

European Science Foundation

European Computational Science  
Forum: The "Lincei Initiative":  
from computers to scientific  
excellence



ESF is now  
considering  
establishing CSEC



## Scientific software development: a new CECAM initiative

**On March 30-31, 09, the director (Wanda Andreoni) and vice-president (Paul Durham) of CECAM convened a meeting at CECAM Headquarters in Lausanne of a group of scientists with the aim of reflecting upon the possible role CECAM could play in enhancing European scientific software development and support**

- [Alessandro Curioni](#) (*IBM Research Zurich*)
- [Stefano de Gironcoli](#) (*SISSA, Trieste*)
- [Mauro Ferrario](#) (*University of Modena*)
- [Xavier Gonze](#) (*University of Louvain*)
- [Christian Holm](#) (*University of Stuttgart*)
- [Wim Klopper](#) (*University of Karlsruhe*)
- [Mike Payne](#) (*University of Cambridge*)
- [Bill Smith](#) (*Daresbury Laboratory*)
- [Godehard Sutmann](#) (*Research Centre Jülich*)
- [Doros Theodorou](#) (*University of Athens*)
- Other scientists will be invited to join the group.

CECAM (Centre Européen de Calcul Atomique et Moléculaire) is a European organization devoted to the promotion of fundamental research on advanced computational methods and to their application to important problems in frontier areas of science and technology

5.

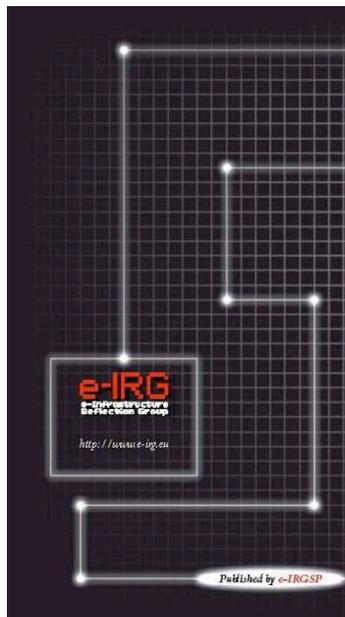


# PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

## Towards the High-End HPC Service for European Science

Thomas Lippert, PRACE Project Coordination@FZ-Jülich

# Computational science infrastructure in Europe



**The European Roadmap for Research Infrastructures is the first comprehensive definition at the European level**

**Research Infrastructures are one of the crucial pillars of the European Research Area**

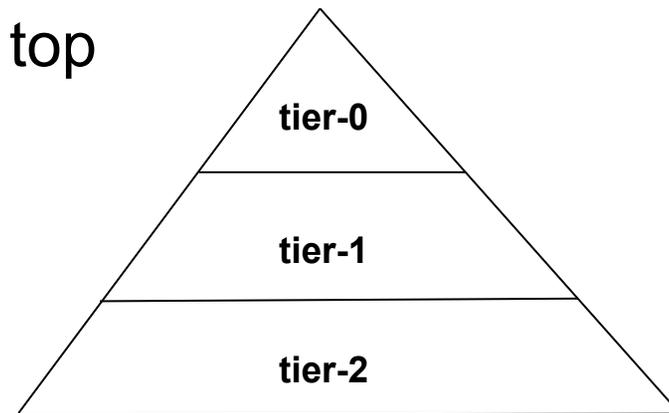
**A European HPC service:**

- **Horizontal**
- **attractive for research communities**
- **supporting industrial development**

## ESFRI Vision for a European HPC service

- Need European HPC-facilities at top of an HPC provisioning pyramid

- Tier-0: 3-5 European Centres
- Tier-1: National Centres
- Tier-2: Regional/University Centres



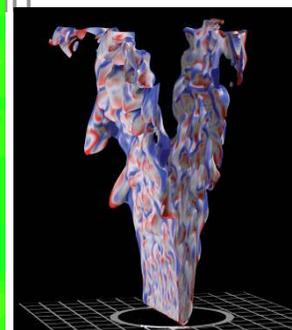
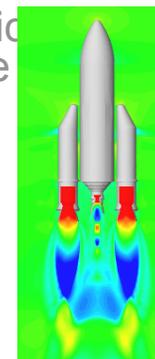
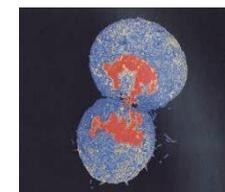
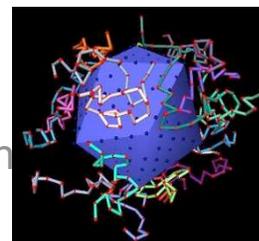
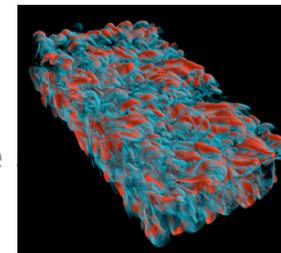
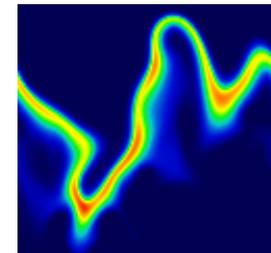
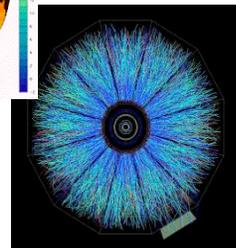
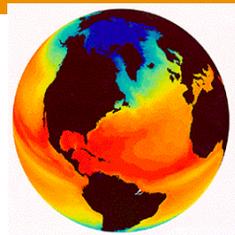
- Part of the Creation of a European HPC ecosystem

- HPC service providers on all tiers
- Grid Infrastructures
- Scientific and industrial communities
- The European HPC industry

**Renewal every 2-3 years**  
**Construction cost 200 – 400 Mio. €**  
**Annual running cost 100 – 200 Mio.€**

# HET: The Scientific Case

- Weather, Climatology, Earth Science
  - degree of warming, scenarios for our future climate.
  - understand and predict ocean properties and variations
  - weather and flood events
- Astrophysics, Elementary particle physics, Plasma physics
  - systems, structures which span a large range of different length and time
  - quantum field theories like QCD → LHC, FAIR
  - ITER
- Material Science, Chemistry, Nanoscience
  - understanding complex materials, complex chemistry, nan
  - the determination of electronic and transport properties
- Life Science
  - system biology, chromatin dynamics, large scale protein dynamic association and aggregation, supramolecular systems, medicine
- Engineering
  - complex helicopter simulation, biomedical flows, gas turbines and internal combustion engines, forest fires, green aircraft





# PRACE – Initiative



## New Partners - since May 2008



PARTNERSHIP  
FOR ADVANCED COMPUTING  
IN EUROPE

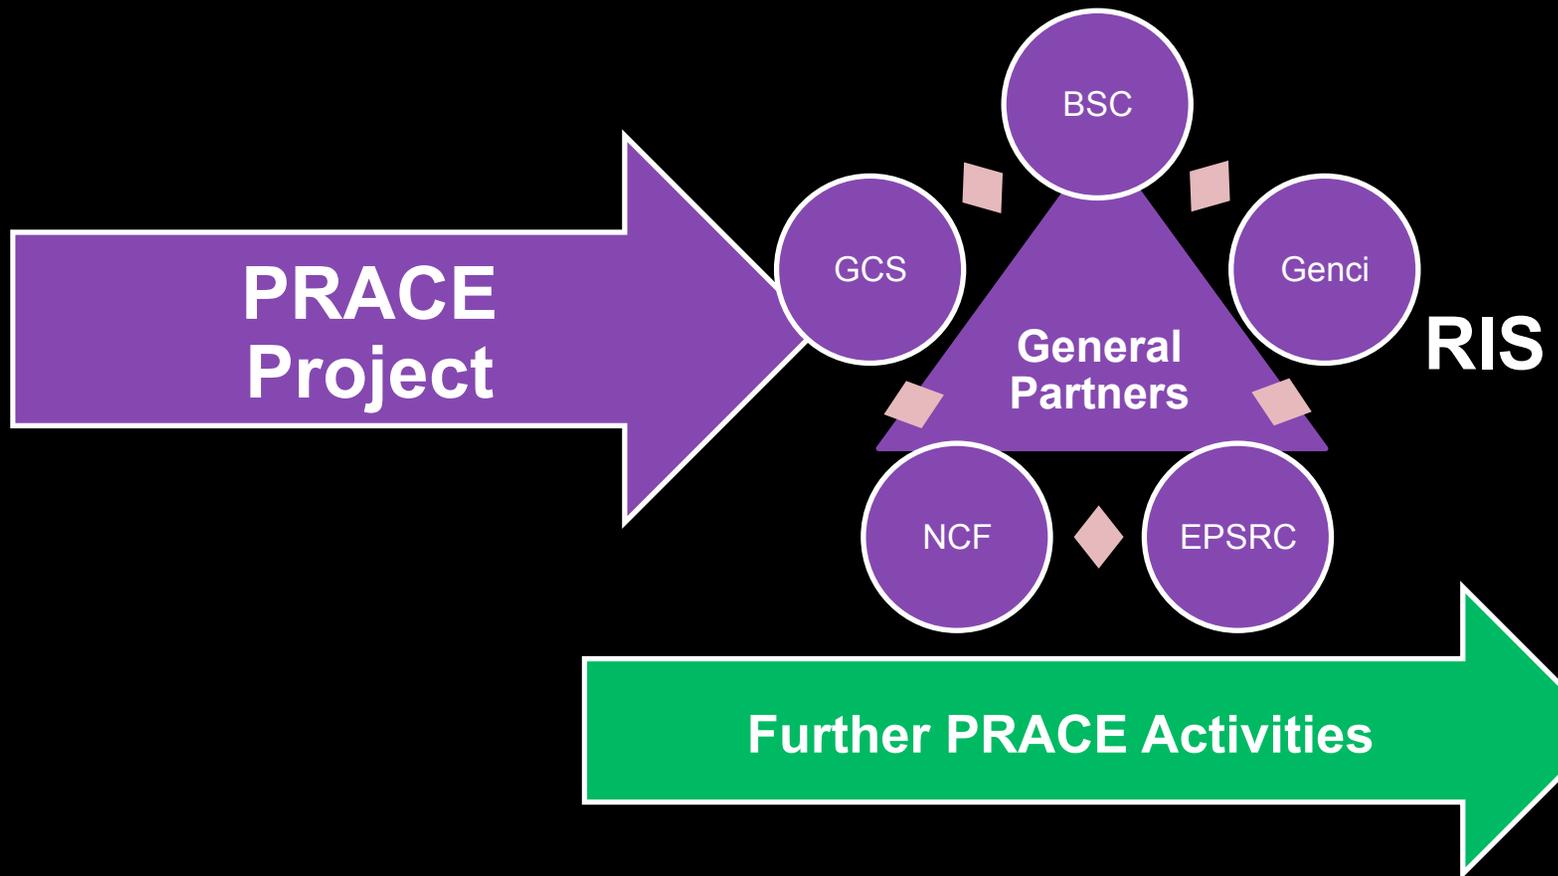


# First Industry Seminar attendees



15 April 2009

# PRACE Initiative





## PRACE Project

- PRACE is **horizontal** ESFRI project
  - **Mission to serve the scientific communities at large**
  - Need to cooperate with communities
- **Software for the Multi-Petaflop/s age**
  - Only few of today's applications are scalable to hundred-thousand CPU-cores
  - PRACE seeks to gain knowledge in Petascaling to educate and support its future users
  - **An additional European effort is needed – international cooperation should be sought for Exascale challenges**
- **Exascale data services for scientific communities**
  - Support efforts to agree on community standards for storing, annotating and retrieving their data, provide reliable data services



## PRACE Project

- Prepare the contracts to establish the PRACE permanent Research Infrastructure as a **single Legal Entity** in 2010 including governance, funding, procurement, and usage strategies.
- Perform the technical work to prepare operation of the **Tier-0 systems in 2009/2010** including deployment and benchmarking of prototypes for Petaflop/s systems and porting, optimising, Peta-scaling of applications



<b>WP6</b>	Software enabling for Petaflop/s systems (RTD)	<b>Prepare key applications to use the future Petaflop/s systems efficiently; capture requirements for WP7 and WP8 and create a benchmark suite.</b>	EPSRC
<b>WP7</b>	Petaflop/s Systems for 2009/2010 (RTD)	Identify potential Petaflop/s systems for PACE that can be installed in 2009/10 with prototypes deployed by WP5. Prepare the procurement process including acceptance criteria.	GENCI
<b>WP8</b>	Future Petaflop/s computer technologies beyond 2010 (RTD)	<b>Start a permanent process to identify technologies for future multi-Petaflop/s systems of the RI and work with hardware and software vendors to influence the direction they are taking. Establish PRACE as a leader in HPC technology.</b>	Gauss Centre

## WP6: Software Enabling for Petaflop/s Systems

- Create an application benchmark suite
- Capture application requirements for Petascale systems
- Port, optimise and scale selected applications
- Evaluate application development environments of the prototypes

## PRACE WP8: STRATOS

- STRATOS is a deliverable of PRACE WP8:  
Create sustained platform for technology watch  
and development for PRACE
- Hardware
  - Identifying and developing components of future multi-Petaflop/s hardware
- Software
  - Plans for Exascale software development within STRATOS

# Areas of Contribution to IESP

## European Science and Engineering Communities

- Coordination with science drivers
- Identify application codes and enabling HPC software

## Performance Tools

## Programming Tools

## Benchmark Codes

## MPI, OpenMPI standardization

## Scalable Cluster OS

## Grid/Cloud-Integration-Middleware