HPC Activities at University Politehnica of Bucharest

Emil Sluşanschi, Nicolae Ţăpuş, Alexandru Herişanu, Alexandru Olteanu, Răzvan Dobre

Computer Science & Engineering Department

November 29-30 2011, Bucharest, Romania
RO-LCG 2011 Workshop

Agenda – Computing Infrastructure

- Hardware
- NCIT Project Domains
- UPB user domain
- Software

- **HPC** / Cloud / Virtualization / Grid Computing
- The Storage & I/O Department – Connecting Clusters

Current Infrastructure

- Magic word: diversity

  32 dual quad-core Xeon + **20 dual hex-core Opteron**

  + 4 dual PowerXCell 8i + ~~50 P4 HT + 32 dual Xeon~~

  = 642 cores

  (total: 918 cores / HPC 642 – Virtualization 232)
- GbE/Infiniband Interconnect
- Total storage of 36TB

HP-SEE
High-Performance Computing Infrastructure for South East Europe's Research Communities

+ other "friendly clusters"

Our Network

- Full IPv6 stack (routed, not tunneled)
- Infiniband
- Dual Gigabit

Lots of experimenting with network Optimization / monitoring Limbo

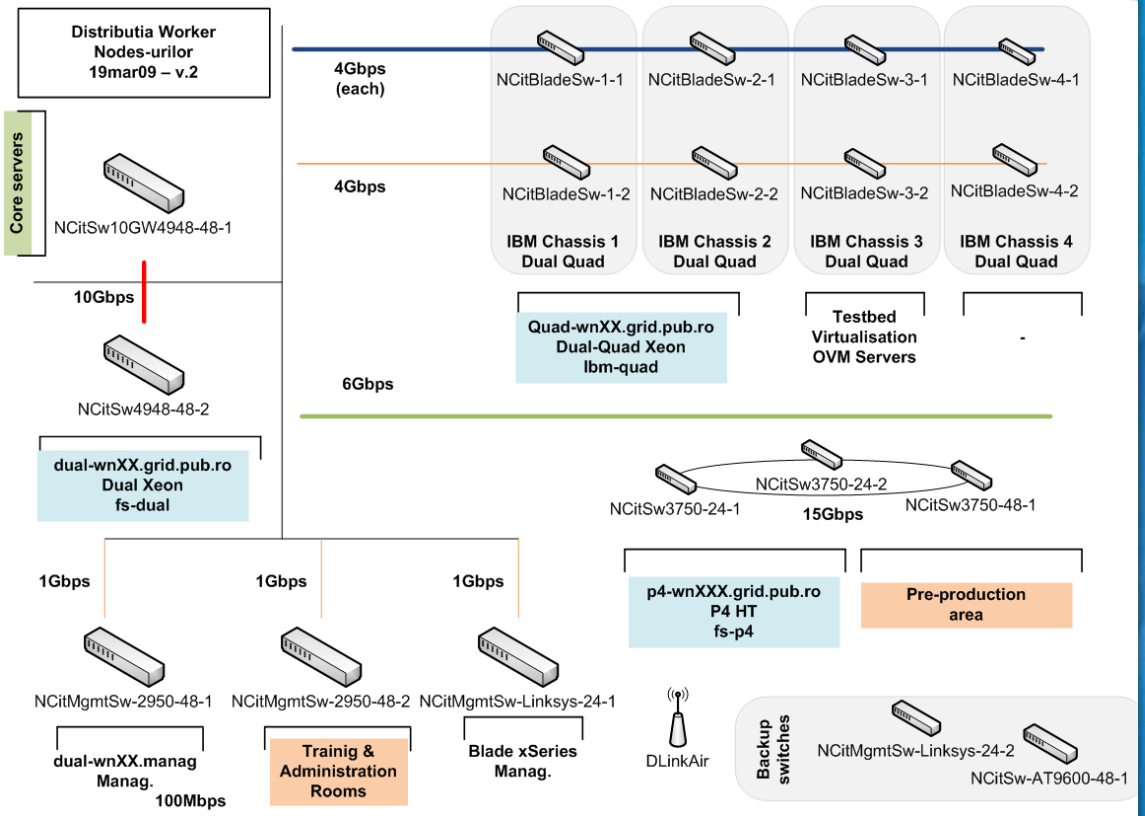

- Jumbo frames
- Network stack configuration, bonding

HPL Performance

- 1TFlops sustained for 230.000 dense linear systems
- Current work:
  - Use IBM XL C/C++ Compiler Suite
  - Use the Intel Cluster Toolkit Compiler Suite
  - Use optimized Intel MKL, MASS libraries
  - Optimize for LS22 Infiniband
  - Use OpenCL version of LINPACK for GP-GPU computing

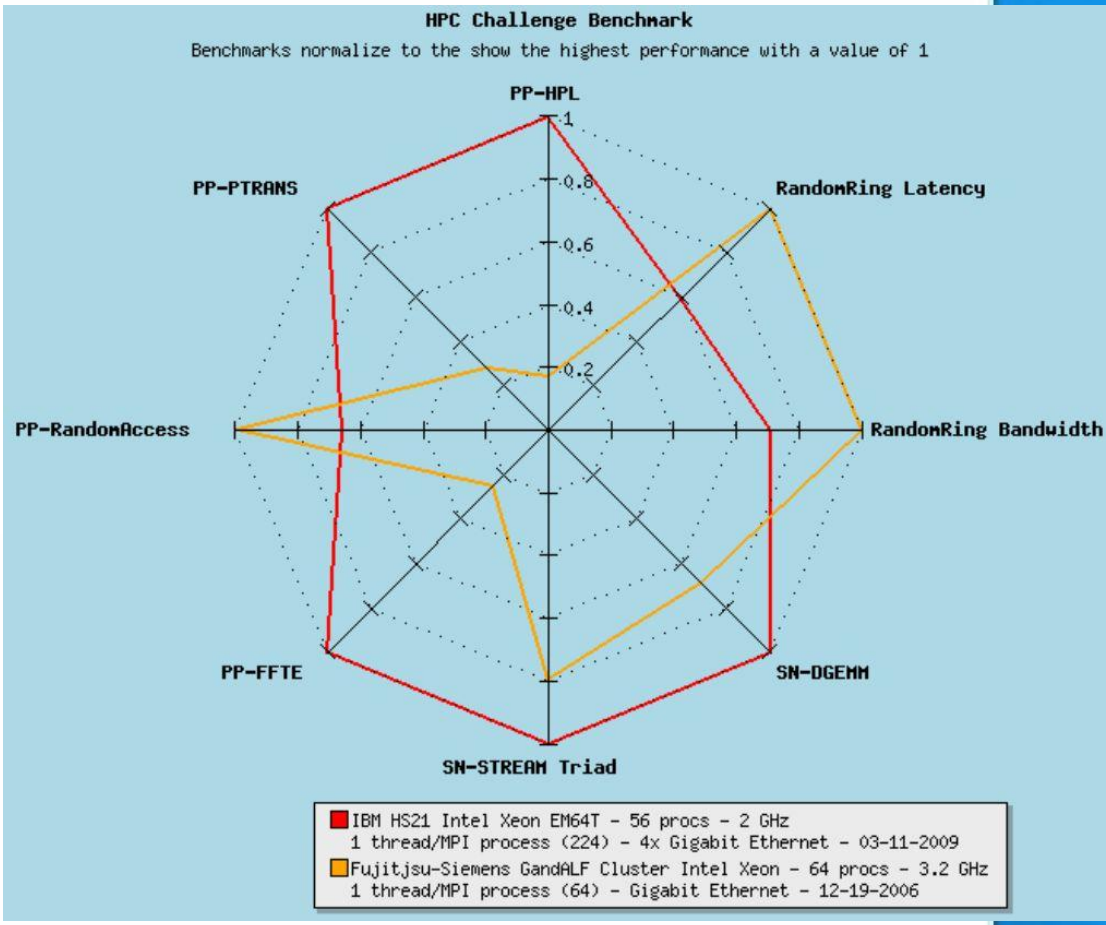

NCIT Domains of Interest

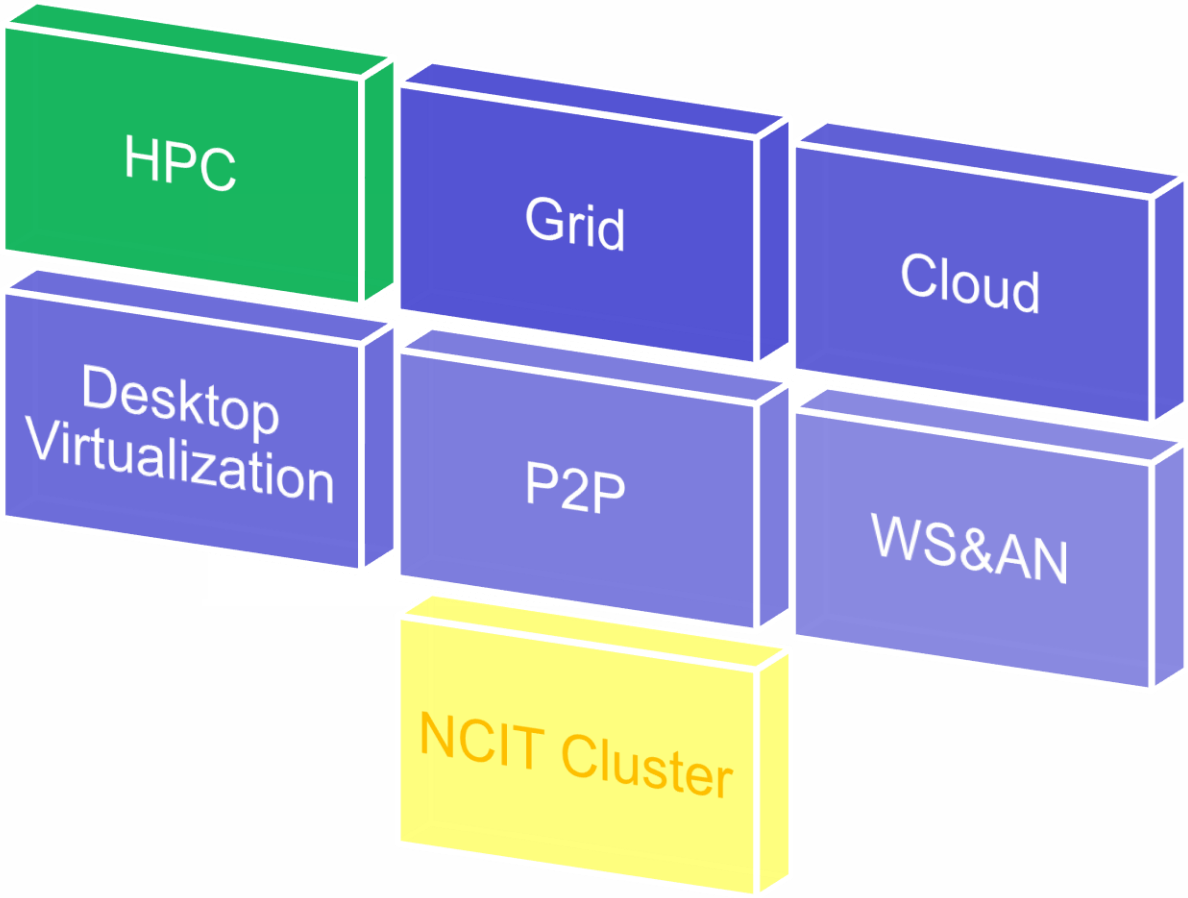

Our user domain

- Hardware needs people to function

MPI OpenMP

Cell Computing

HPC Summer School

Profiling / Tuning

Diploma

Master

first sem.

second sem.

Σ

Grid / HPC

3rd year of study

4th year of study

2 years

Research Projects / PhD

Software

- Sun Grid Engine 6.2u6 / cfEngine software provisioning
- Compilers
  - Ibm XL, Intel C/Fortran, PGI, SunStudio, gcc
- Debuggers
  - TotalView
- Profilers
  - VTune, Sun Studio Analyser
- Libraries
  - Intel MKL, NAG, Blas
- Tools
  - Code Saturne, Charm++, Gaussian09, OpenFoam, Paraview etc.

http://cluster.grid.pub.ro

The Storage & I/O Department

- How do you store your data?
  - Disk, NFS, Lustre
- Transport? TCP/IP, Infiniband

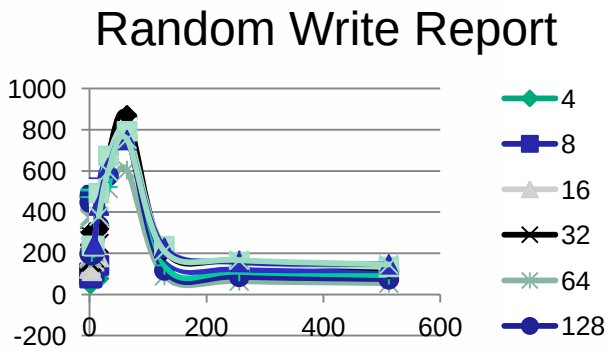

LNET Storage Routing

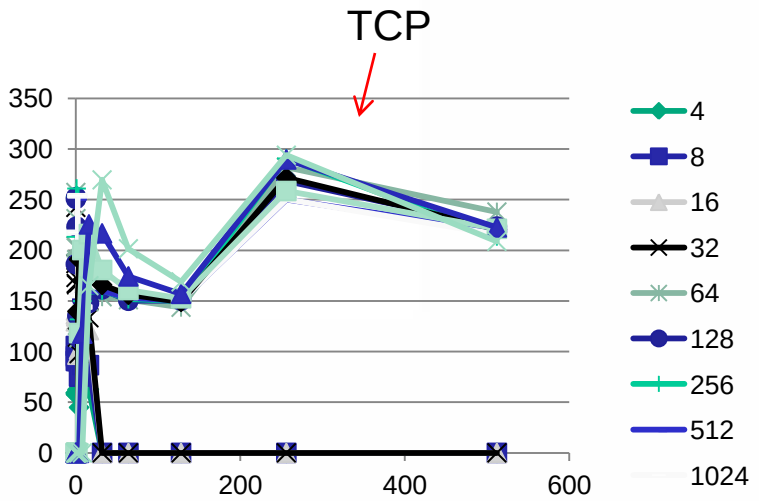

- Large storage requires operational knowledge

Connecting Clusters

- Hardware is everywhere
- How do you connect multiple clusters together?

users, storage, cpu, software

NCIT Cluster – ICF-HPC (Currently)

- We have to write the Cookbook for that

QS22 chassis, Myrinet, Infiniband,
HS22, LS21, BlueGene/P (UVT)?

HP-SEE
High-Performance Computing Infrastructure for South East Europe's Research Communities

Infrastructure Outlook

- GPGPU Programming in curricula probably in 1-2 years
- At least 4 NVidia Fermi Engines in IBM Blade Servers
- RDMA over Infiniband GPU Memory to GPU Memory
- OpenCL & CUDA Programming

- Everything is a ~~Object~~ VM

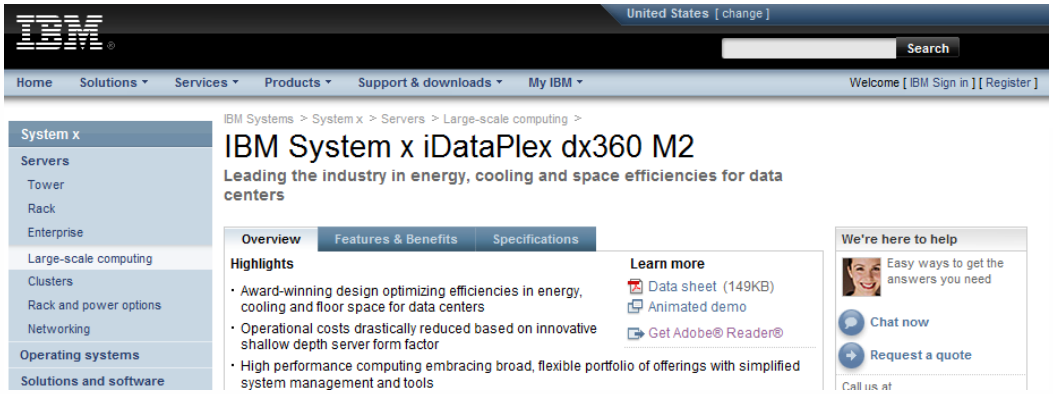

Agenda – HPC Applications

- 3D Romania – online high-resolution model; element and feature extraction – EagleEye (Image processing)
- Modeling and Simulation of Aerodynamic Phenomena: INCAS (Aerospace Research)
- Weather Prediction Models: COSMO, HRM, WRF – ANM (Meteorology)
- Atomic Scale Simulation in Material Science – ICF (Molecular Dynamics and Physical Chemistry)
- N-Body & SPH Simulations (Astrophysics)
- Numerical Simulation of Earthquakes in the Vrancea Region – INFP/NIPNE (Earth Sciences)
- HPC /Multicore Training @ cs.pub.ro

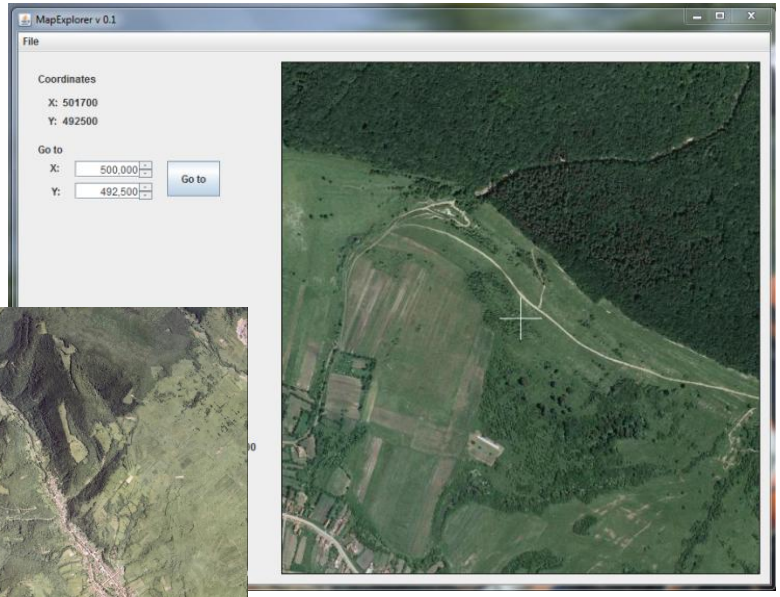
MapExplorer v 0.1
File
Coordinates
X: 501700
Y: 492500
Go to
X:
500,000
Y:
492,500
Go to

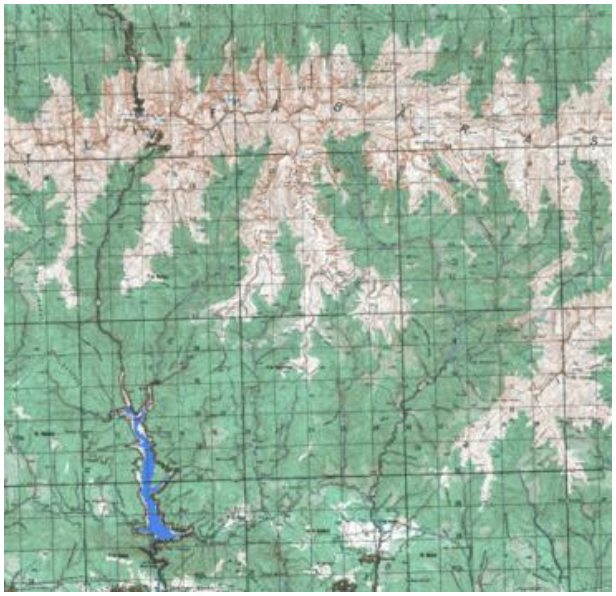

3D Romania

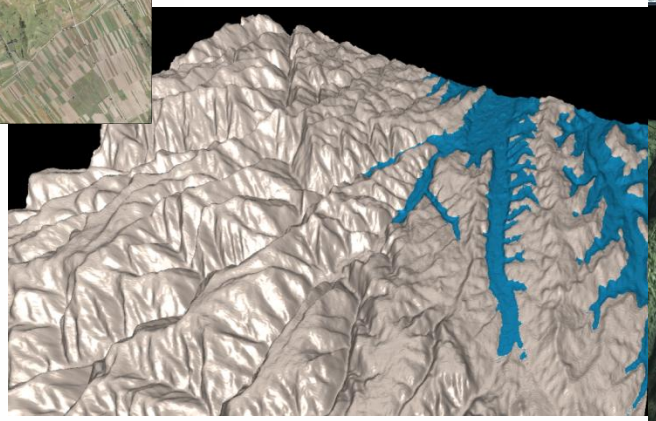

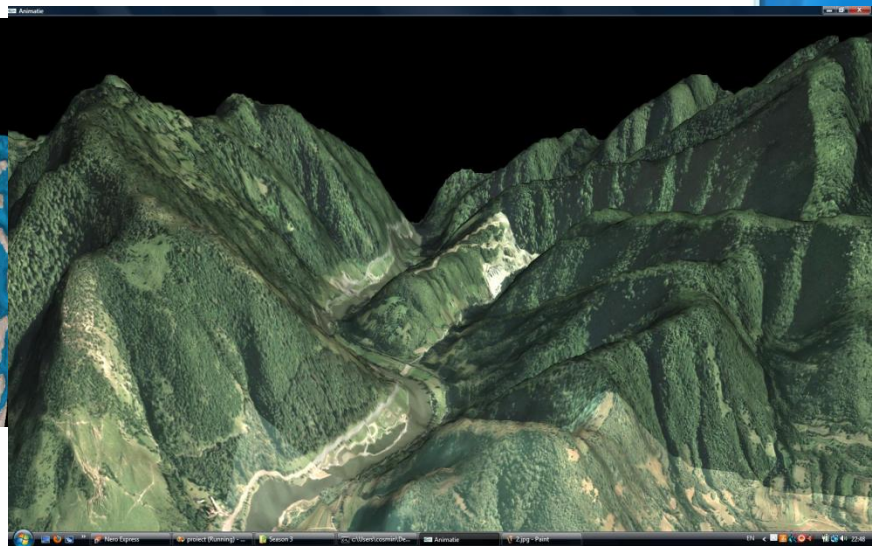

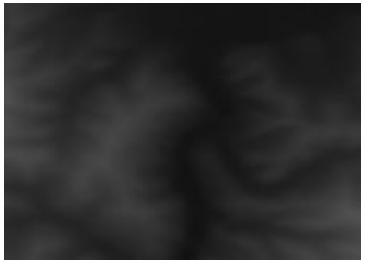

Element Detection & Feature Extraction

Median Filter

Grayscale

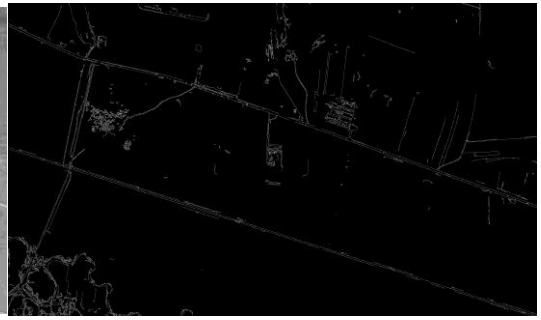

Canny Edge Detector

From Images to Features

Original

Grayscale

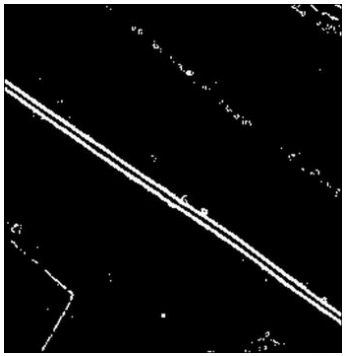

Image Detection (Sobel)

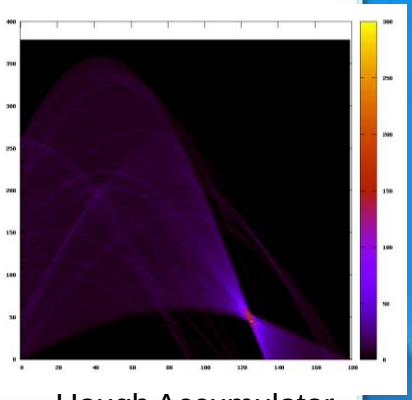

Hough Accumulator

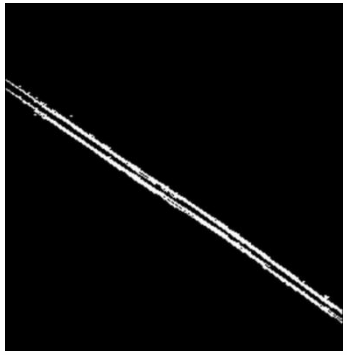

Hough Peaks over image edges

Mark road segment edges

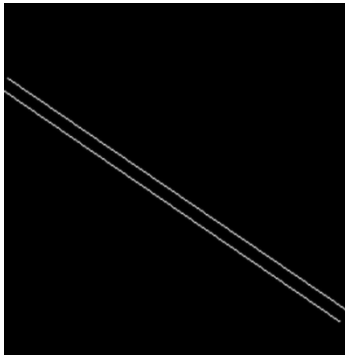

Final identified feature (road)

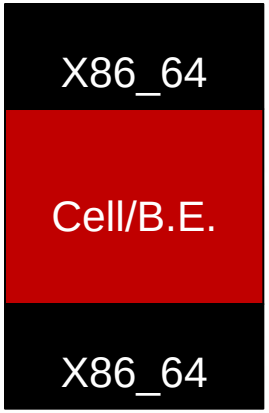

Modeling and Simulation of Aerodynamic Phenomena

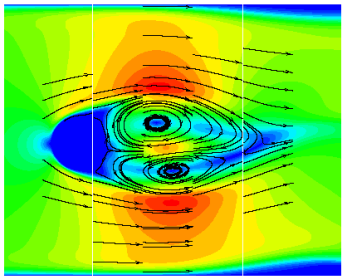

- Preprocessing & Grid generation – Metis/ParaMetis

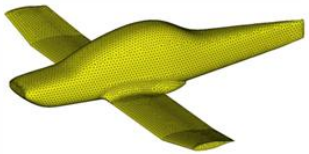

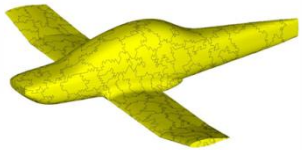

- Open-Source complex Navier-Stokes CFD codes
  - Cluster porting of existing CFD Codes: Code Saturne & OpenFOAM
  - Tuning & improving serial performance
  - Improve MPI/OpenMP/GPGPU scaling
- Postprocessing – ParaView & Salome

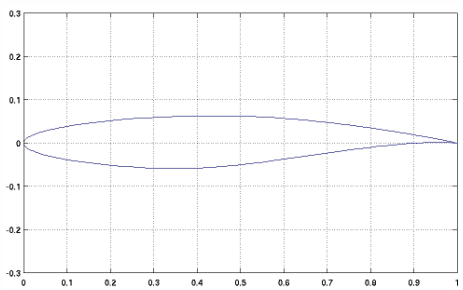

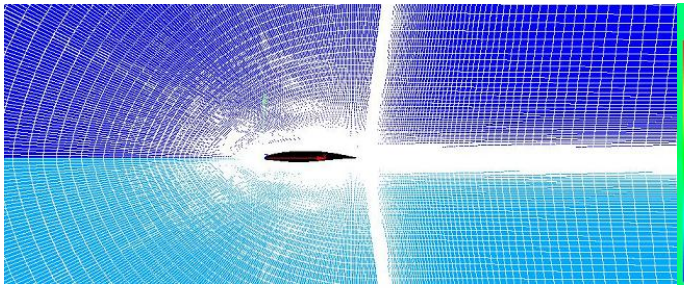

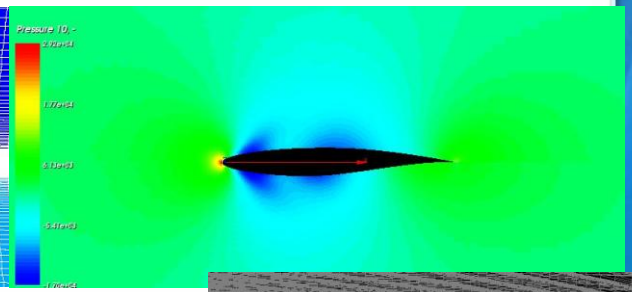

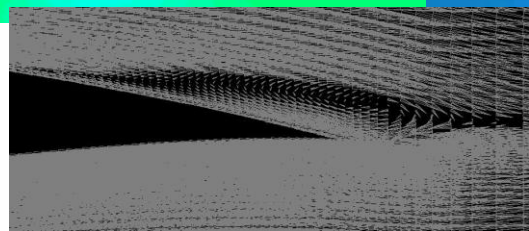

Weather Prediction Models

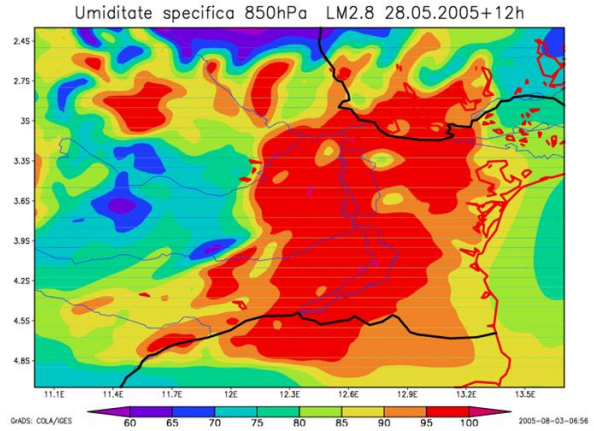

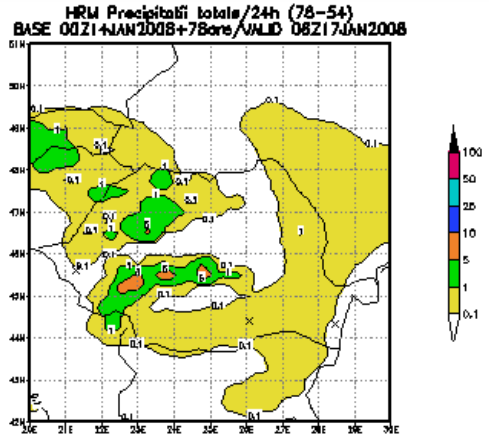

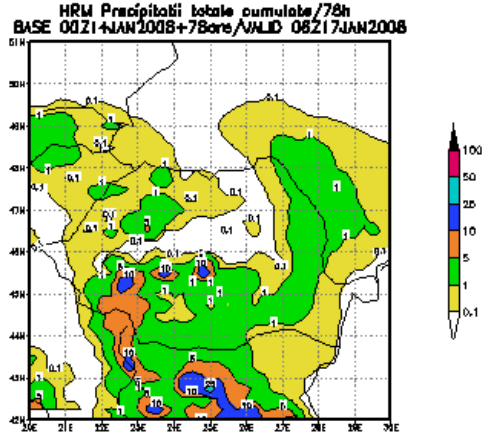

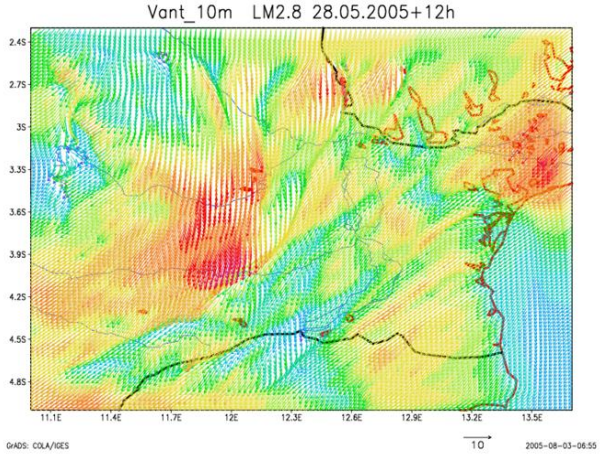

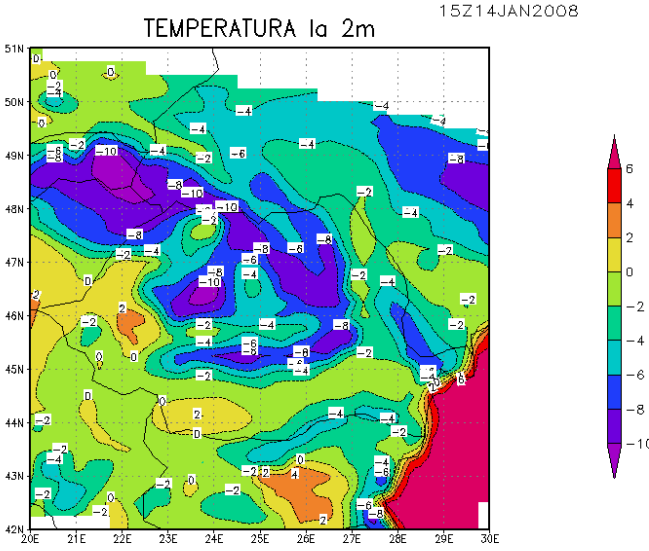

Prognoza temperaturii la 2m la 15h HRM Made by Ion Victor PESCARU

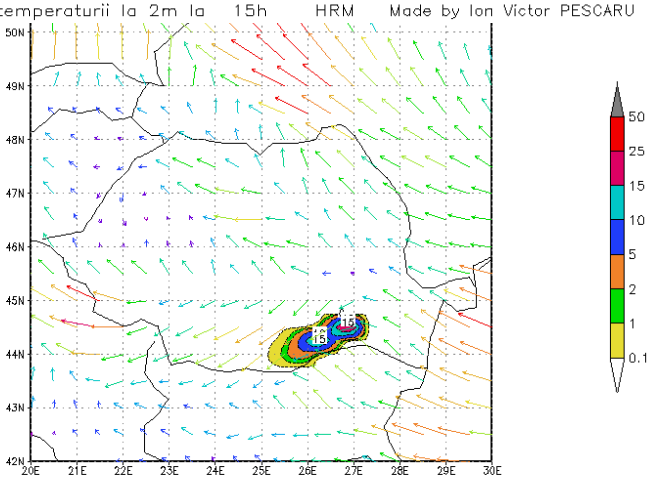

Prognoza norului de poluant la 24h Administratia Nationala de Meteorologie

HRM & COSMO Scaling

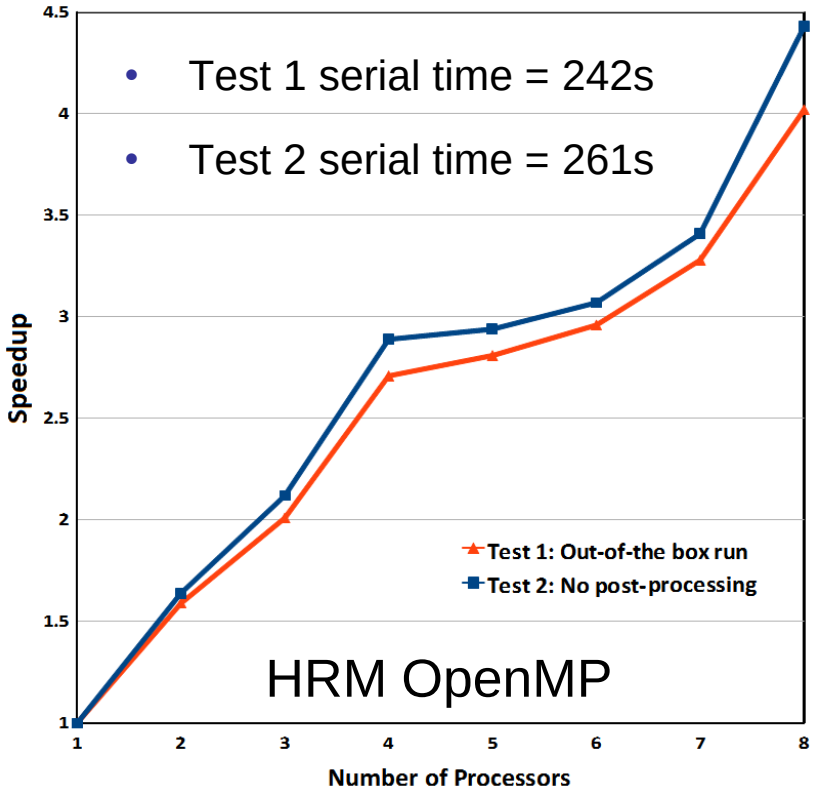

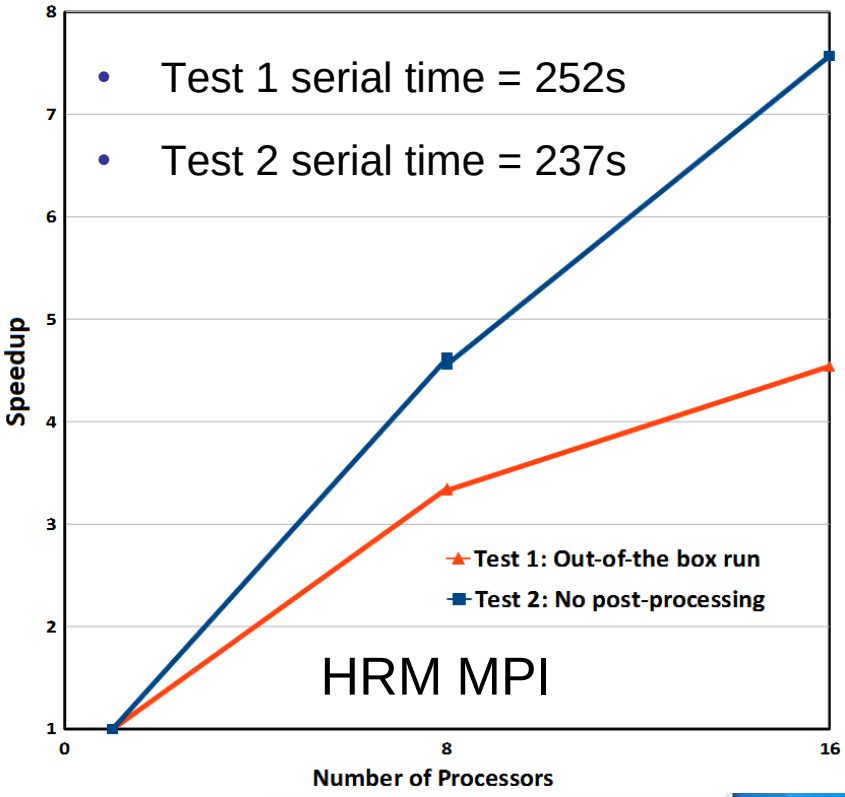

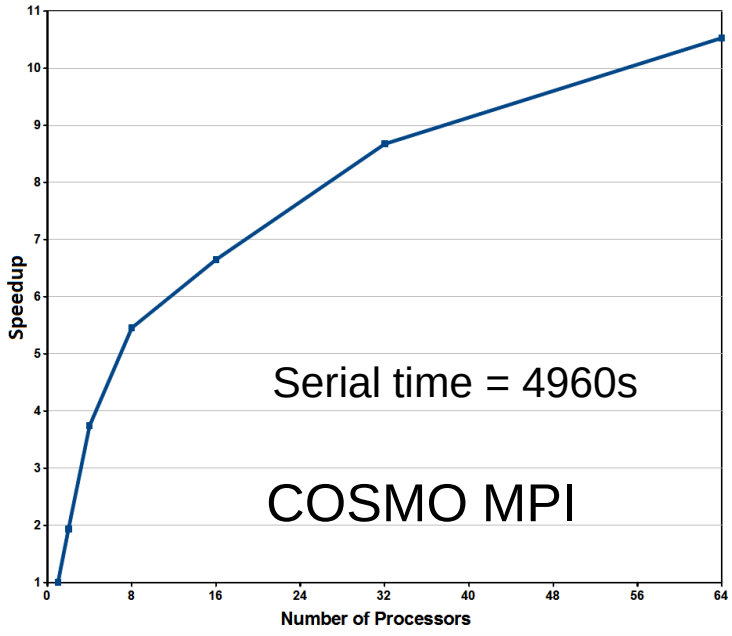

Processor Architecture Scaling

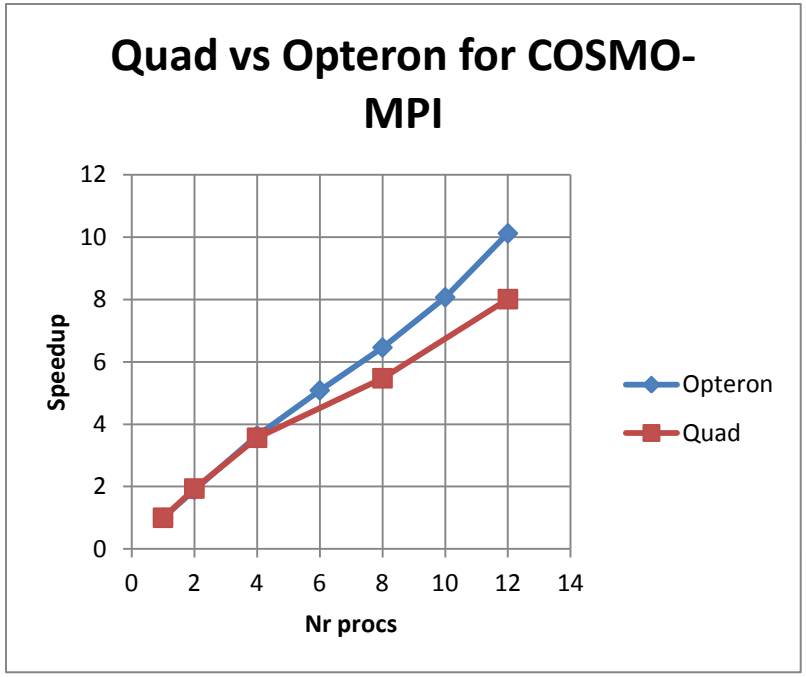

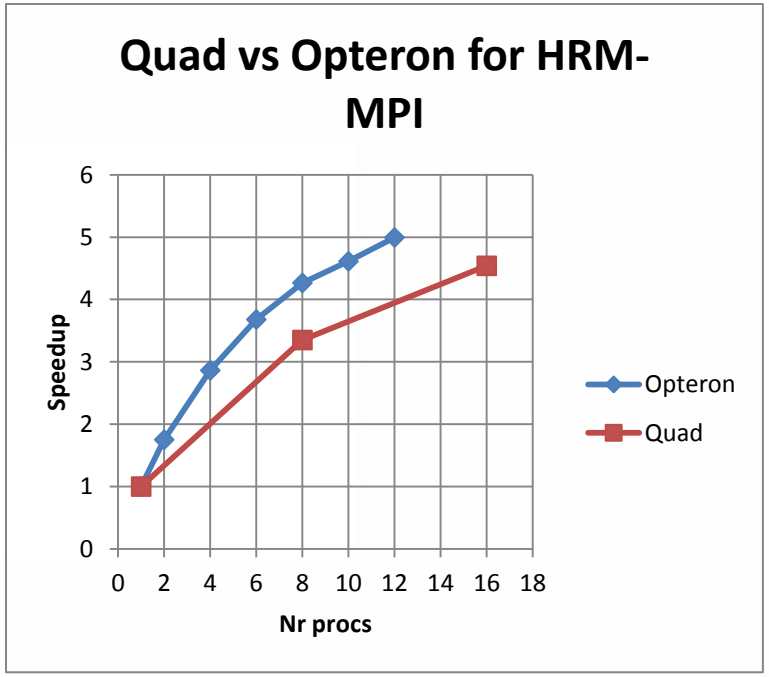

Intel Xeon Quad-Core vs AMD Opteron Six-Core

Ongoing tests on Intel Nehalems & planned tests on Intel Sandy Bridge

Atomic Scale Simulation in Material Science

- Paramagnetic materials simulation
  - OpenMP CS parallelization: Speedup 8.3x on 8 procs – **superlinear** due to improved cache performance
- GAMESS
  - MPI original program – Speedup 6.4x on 8 cores
- MOPAC
  - Propose a parallelization scheme for modern multi-core architectures
- NAMD / GROMACS / Gaussian / CPMD
  - Scale to production systems

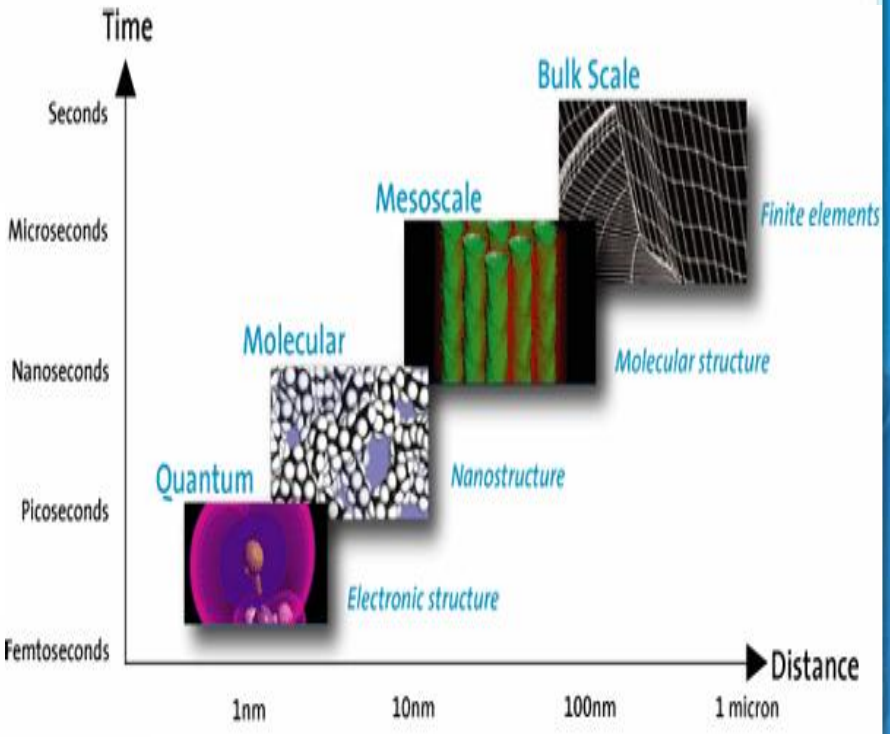

Atoms/ Molecules		Nanoscale Particles		Condensed Matter	
1		125	70,000	6×10^6	∞N^oAtoms
	1		10	100	∞ Diameter (nm)
Quantum Chemistry		?		Solid State Physics	

Nanotubes and systems of nanotubes

- Hysteresis phenomenon 1 nanotube
  - Serial run: 2165s
  - Optimized serial run: 9s
  - Optimized parallel* run: 3.8s
  - Total Speedup: **569x**
- Systems of nanotubes
  - 100x100 tubes
  - Serial run: 350s
  - Optimized serial run: 17s
  - Optimized parallel* run: 10s
  - Total Speedup: **35x**
- Intel Parallel Studio analysis proved essential!

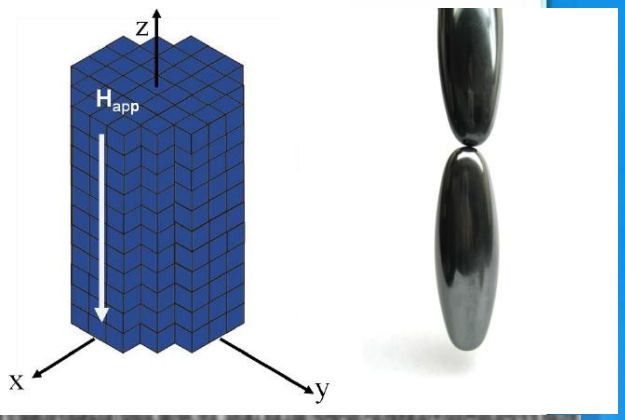

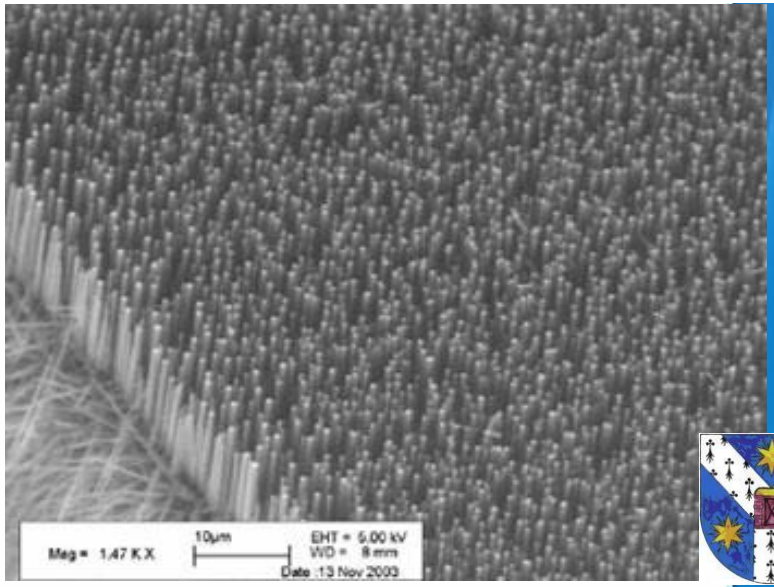

*Dual Penryn quad-core → 8 cores

NAMD & VMD Using Charm++

NAMD = **N**ot**A**nother **M**olecular **D**ynamics
VMD = **V**isual **M**olecular **B**ynamics

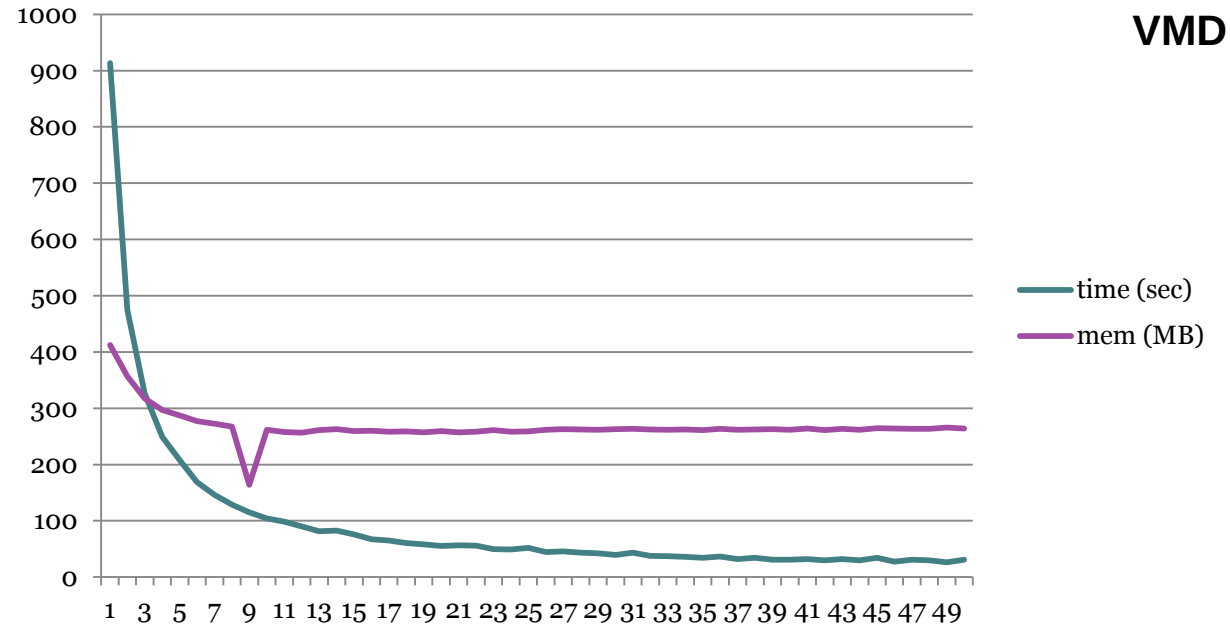

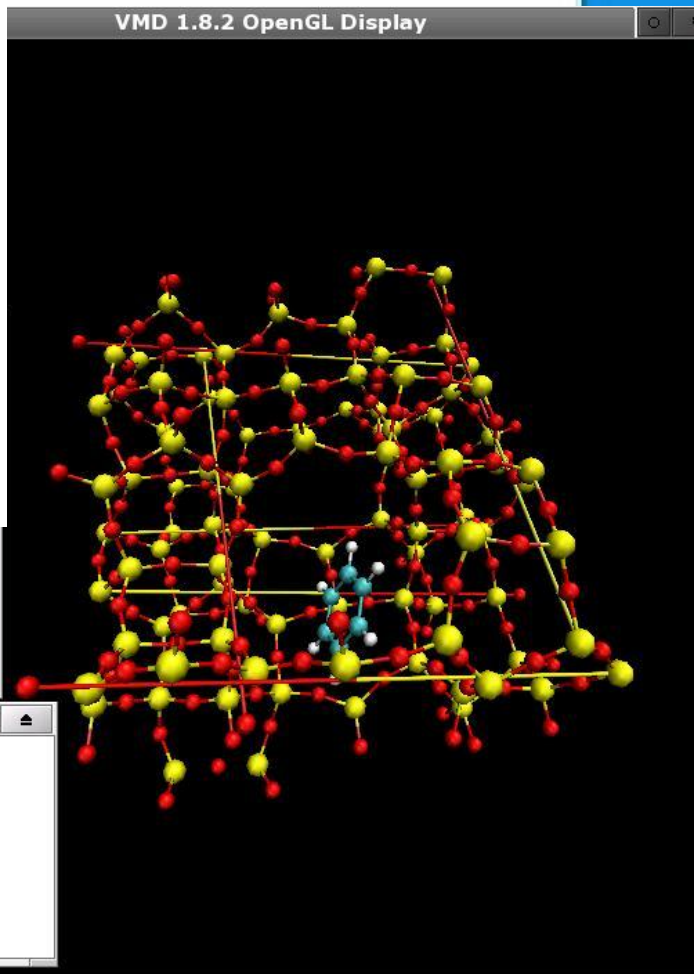

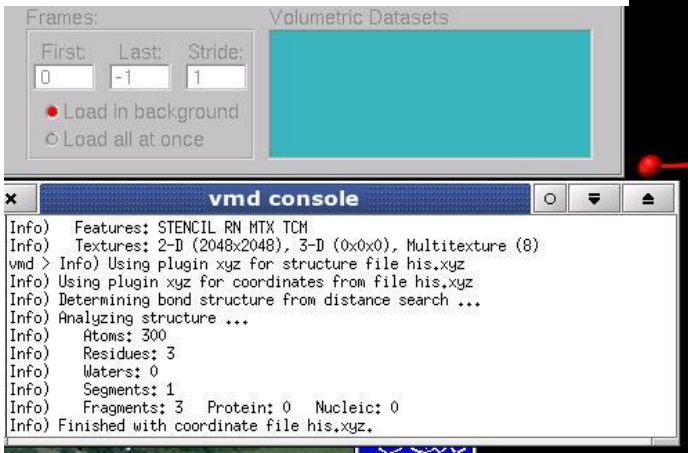

Total Speedup:
37.5x / 4x cores

Astrophysics

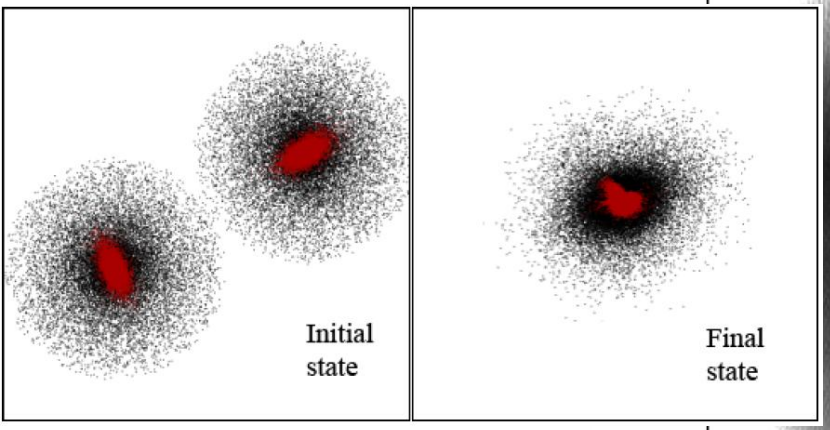

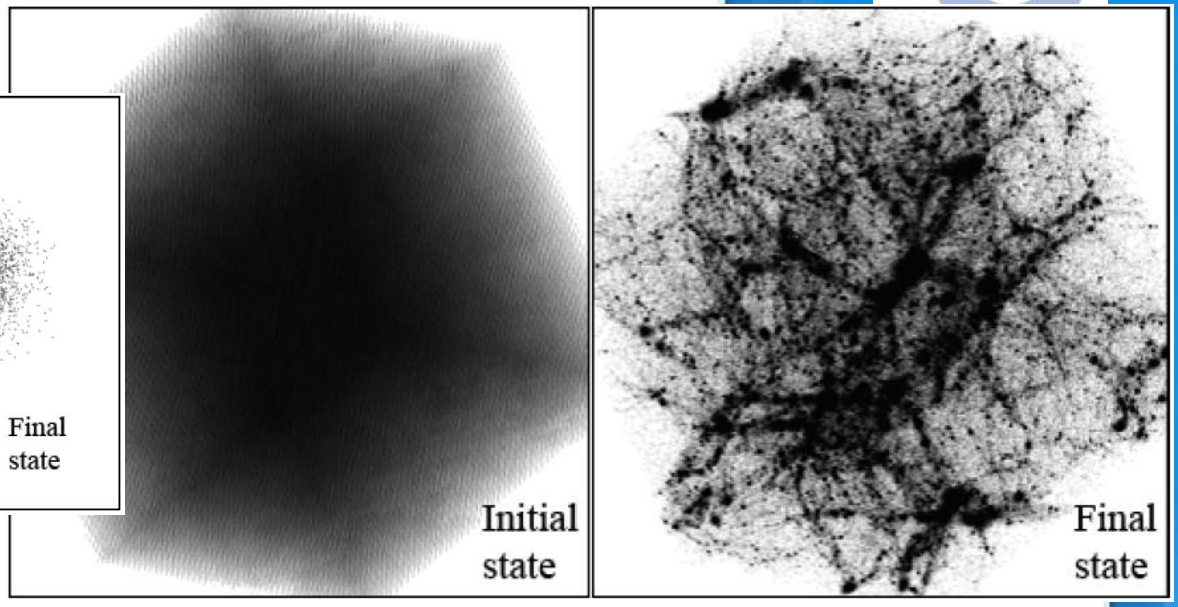

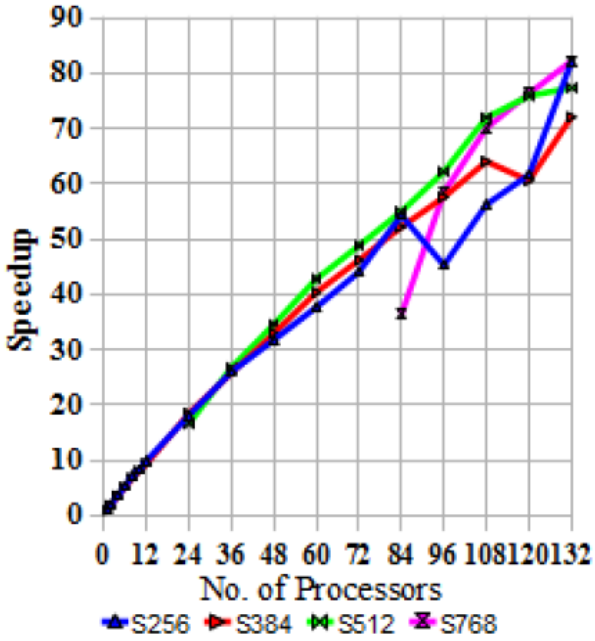

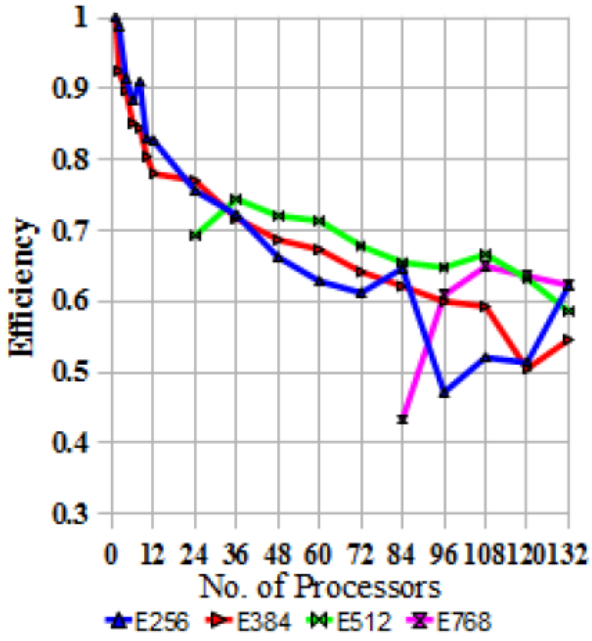

Numerical Simulations of Earthquakes in the Vrancea Region

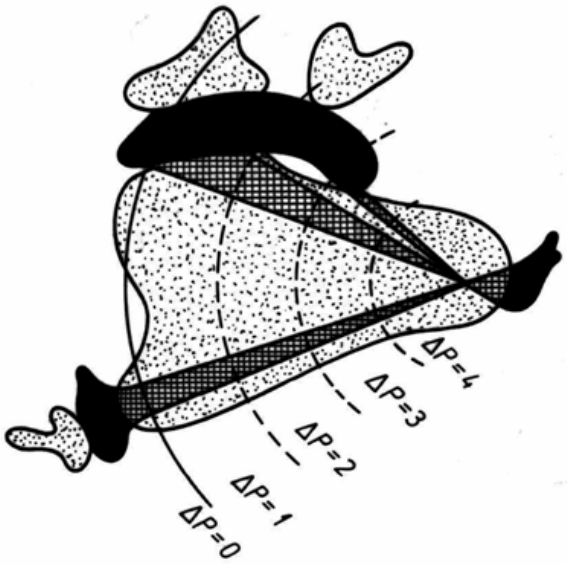

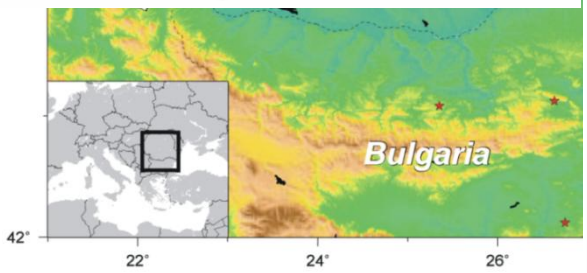

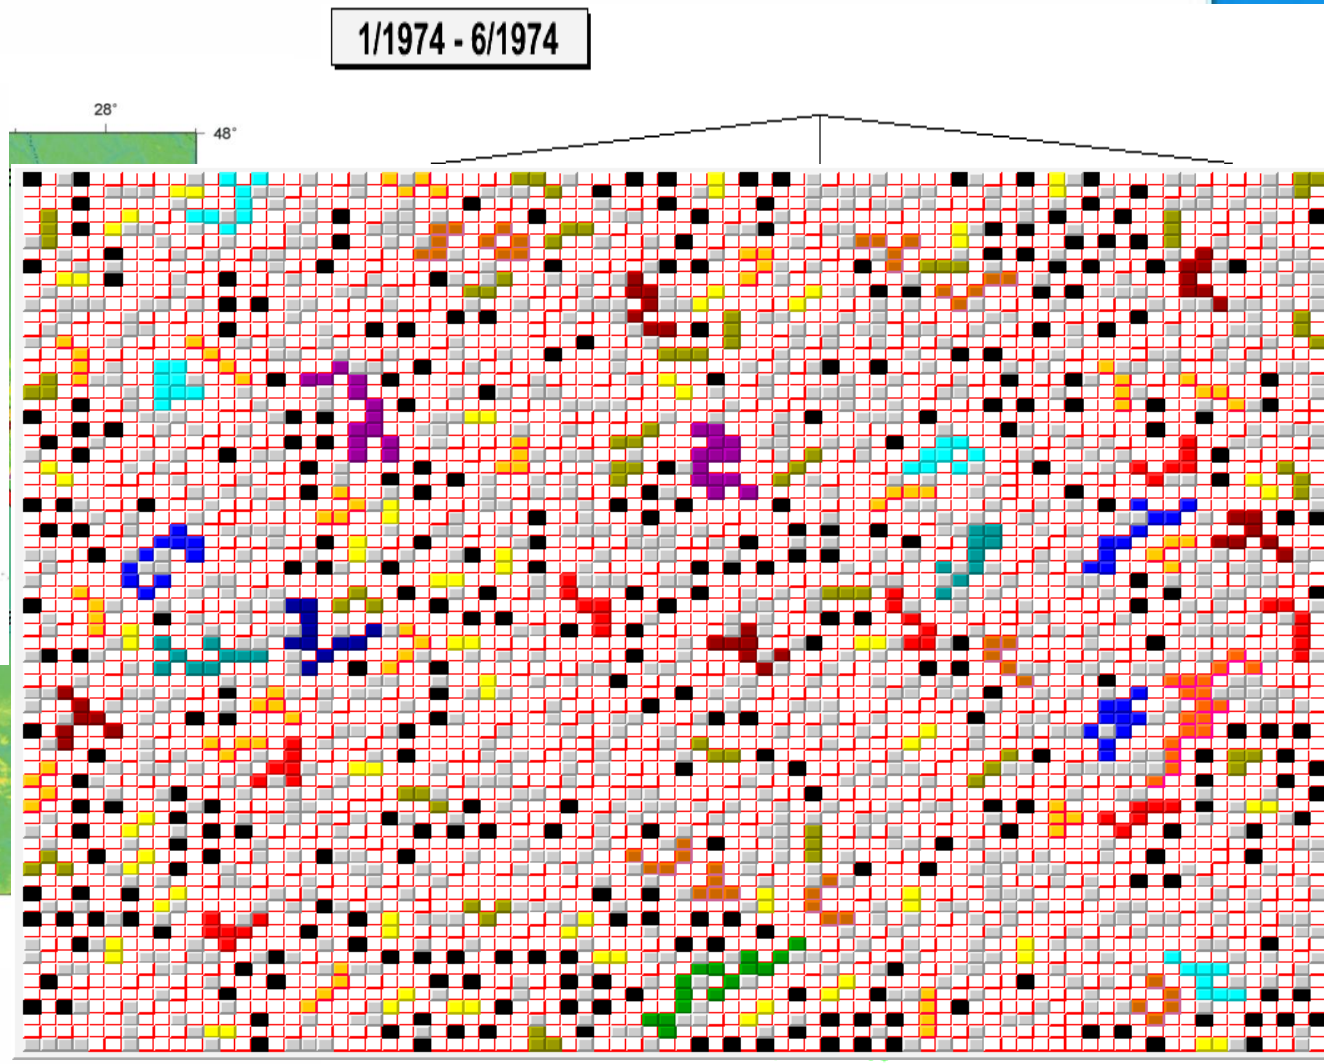

HPC Related Lectures & Training @ CS

- Grid/HPC Initiative Summer school:
  - First GridInit was in **2004**
  - Usually debated grid middleware tasks
  - Fom 2008 the main focus is on developing **HPC Applications** using architectures with **multicore** processors
- Undergraduate Lectures:
  - Parallel Computing Algorithms and Data Structures, (Parallel) Computer Systems Architecture, Distributed Programming Languages
- Graduate Lectures:
  - Distributed Systems, Cluster & Grid Computing, High Performance Computing – Numerical Methods and Programming Techniques, Distributed Algorithms
- HPC Industry Training @cs.pub.ro:
  - Intel Multi-core Programming for Academia – 2007
  - IBM Basic and Advanced Cell Programming – 2008
  - IBM BlueGene Programming – 2009
  - Intel Parallelism Faculty – 2009
  - NVidia Cuda Programming – 2012

Acknowledgements

- ICF: Viorel Chihaia
- INCAS: Victor Pricop, Marius Cojocaru, Claudiu Vadean
- IFIN: Octavian Carbunar
- INFP: Mircea Radulian, Constatin Ionescu
- UBUC: Marian Ivan
- ANM: Victor Pescaru, Rodica Dumitrache, Cosmin Barbu
- AIRA: Marian Suran, Dumitru Pricopi
- CS@UPB: Ovidiu Hupca, Iulian Milas, Andrei Pasatoiu, Maria, Nadejde, Cosmin Constantin, Marius Poke, Andreea Sandu, Vlad Spoiala, Diana Gudu, Aurora Mirea, Cristina Ilie, Victor Spiridon

Thank you for your attention

Q & A

cluster.grid.pub.ro
cs.pub.ro