High Performance Computing Infrastructure of Babeş-Bolyai University

Virginia Niculescu Darius Bufnea Adrian Sterca

Faculty of Mathematics and Computer Science
Babes-Bolyai University

Outline

- Need for HPC infrastructure at BBU
- MADECIP Project
- HPC Cluster
- MOS – Research Center of Modeling, Optimization and Simulation
- Collaborations
- Technical view of the infrastructure
  - Hardware architecture
  - Software
  - Usage

Babes-Bolyai University, Cluj-Napoca

Faculty of Mathematics and Computer Science

Good positions in different Academic Rankings

- Mathematics
  - *Shanghai Academic Ranking of World Universities in Mathematics*
  - *US News & World Report*
- Computer Science & Mathematics
  - *QS World University Rankings by Subjects*
  - *Times Higher Education World University Rankings*

  - An important interest and orientation on:
    - Model of computations
    - Parallel and distributed computing
    - Cloud and cluster computing;
    - GPGPU programming
  - …
- New courses on these directions

- New Master Program:

High Performance Computing and Big Data Analytics

Master Program – HPC & BDA

- *High Performance Computing and Big Data Analytics*
  - *Master Program at Faculty of Mathematics and Computer Science started in 2014*
  - *http://www.cs.ubbcluj.ro/education/academic-programmes/masters-programmes/high-performance-computing-and-big-data-analytics-programme-profile/*
- *Internationally advertised by*
  - Keystone Academic Solutions - higher education web marketing (http://keystoneacademic.com/)
- Core disciplines:
  - Programming Paradigms
  - Parallel and Distributed Operating Systems
  - Formal Modelling of Concurrency
  - Advanced Methods in Data Analysis
  - Functional parallel programming for big data analytics
  - Models in parallel programming
  - General Purpose GPU Programming
  - Workflow Systems
  - Resource-aware computing
  - Data Mining
  - Grid, Cluster and Cloud Computing
  - Knowledge Discovery in Wide Area Networks

High Performance Computing Needs

- before 2015
  - small HPC clusters exist at Faculty of Physics, Faculty of Chemistry, Faculty of Mathematics and Computer Science, FSEGA
- --->2013
  - a complex analysis of the computation requirements
  - analysis done at the university level
  - initiated by the *Department of Computer Science* (FMCS)
- The demand for an infrastructure of High Performance Computing comes from many research centers of UBB
  - Environment Science
  - Physics
  - Chemistry
  - Mathematics and Computer Science
  - Biology
  - Geography

MADECIP – "*Dezvoltarea infrastructurii de cercetare pentru managementul dezastrelor bazat pe calcul de înaltă performanţă*"
(Project POSCEE no. ID 1862)

UNIUNEA EUROPEANĂ | GUVERNUL ROMÂNIEI | UNIVERSITATEA BABEŞ-BOLYAI CLUJ-NAPOCA | Instrumente Structurale 2007-2013

GOALS:

- R&D Infrastructure in the field of **Disaster Management**

- UBB general research infrastructure

=> modernization

  - innovative character
  - multi-disciplinary research

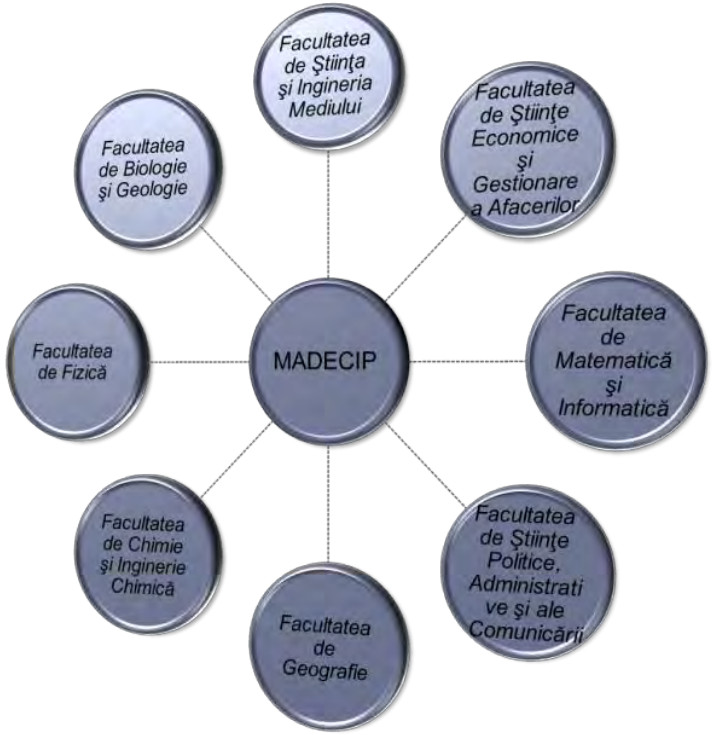

MADECIP – Effective Objectives

Developing

UBB Research Infrastructure

- Research centers and laboratories

Computational Infrastructure

- High Performance Computing Center

MADECIP - Usability

- Disaster management from micro to macro levels
- Interdisciplinary Research:
  - Mathematics
  - Computer science
  - Physics
  - Chemistry
  - Biology
  - Geography
  - Geology
  - Meteorology
  - Economics and Political sciences
- Potential usage: all research areas in UBB.

High Performance Computing Infrastructure

http://hpc.cs.ubbcluj.ro/

- The system can be used for different jobs types:
  - *computation intensive*
  - *data intensive*

=> The solution is based on hybrid architecture:
  - *High Performance Computing system*
  - *Private cloud system*

- Two interaction platforms for easy satisfying:
  - *computation jobs*
  - *storage requirements*
  - *interaction*

High Performance Computing system - Cluster structure

- IBM Solution
- Computation power:
  - Performance (Linpack Benchmark):
    - 62Tflops (Rpeak) and
    - 40Tflops (sustained performance).
- **HPC (Nextscale)**
  - Classical cluster architecture
  - Classical parallel programming
- **Private cloud system (Flex System)**
  - Machine virtualization
  - Cloud computing
- **Common storage** component: 72TB raw
- *Tape* System

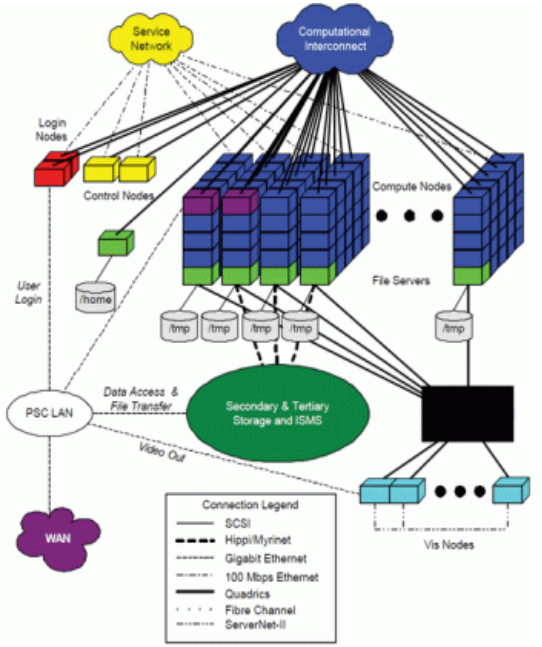

Typical Architecture for a cluster system
source: www.enginsoft.net/activities/hpc2.html

System Software

- Integrated solution for the management of the HPC system and the Cloud System: ***IBM Cluster Platform Manager***
  - cluster management,
  - jobs management,
  - monitoring and reports,
  - MPI compilers and libraries,
  - web interface – easy access,
  - support for GPU and Intel Phi.

- Integrated management solution for the cloud system: ***OpenStack***
  - support for virtualization,
  - resource allocation control

- IBM GPFS (General Parallel File System)

Parallel Programming Tools

- **Intel Parallel Studio, cluster edition**:
  - Compilers for the C, C++ and Fortran, and Python interpreter
  - MPI implementations
  - Debugging and profiling tools
  - Intel Math Library
  - Data Analytics and Machine Learning Library
  - Optimized Building Blocks for Image, Signal, and Data Applications
  - Intel Threading Building Blocks
  - Intel VTune Amplifier
- **Rogue Wave - TotalView**:
  - Debugging and profiling parallel applications

Domain Specific Software

- Matlab
- Mathematica

- Ansys
- CFD Comsol

- Multyphysics Gaussian
- Lumerical FDTD

- WRF-Chem

- …

MOS – Modeling, Optimization and Simulation

Research Center

http://www.cs.ubbcluj.ro/mos

Why MOS?

- What should be the main driving force in future parallel computing developments?
  - Hardware

  or
  - Software
    - important recent development!

Three Faces of Computing
(Pervasive Parallelism Laboratory – Stanford Univ.; K. Olukotun,et. all)

- **Predicting the future**
  - Modeling and simulation (weather, materials, products)
  - Decide what to build and experiment or instead of build and experiment ⇒ third pillar of science
- **Coping with the present (real time)**
  - Embedded systems control (cars, planes, communication)
  - Virtual worlds (second life, facebook)
  - Electronic trading (airline reservation, stock market)
  - Robotics (manufacturing, cars, household)
- **Understanding the past**
  - Big data set analysis (commerce, web, census, simulation)
  - Discover trends and develop insight

MOS – Modeling, Optimization and Simulation

Main Research Domains:

- Mathematical modeling for various phenomena and processes
- Numerical and statistical simulations
- Simulation of natural phenomena
- Visualization and image processing
- Models for parallel and distributed computing
- Domain Specific Languages
- Optimization $\rightarrow$
  - Mathematical models
  - Software implementation
    - performance
    - productivity

Applicative Research – (not restrictive view)

- Modelling and Simulation for:
  - torrents, analysis, floods, dangerous substances overflowing, dam breakdown, dangerous substances dispersion in fluid or poriferous environments etc.
- Big Data Analytics – for specific data needed in disaster managements:
  - web interrogation
  - big databases management
- GIS maps
- Satellite image processing.
- Decision Support Tools – DDST
- Tools for communications, informing and alarming in disaster management domain.
- Simulation /Visualization of different scenarios
  - different disaster types.
- Frameworks and libraries development based on high performance computation.

Models

- Mathematical Models
- Computing Models
  - with higher degree of abstraction
  - General Goal: to assure high level of performance and robustness of the developed software, by assisting the process of the development in order to reduce the time and the work-load.

Domain specific languages/frameworks

...through Models to Performance...

a simplified view...

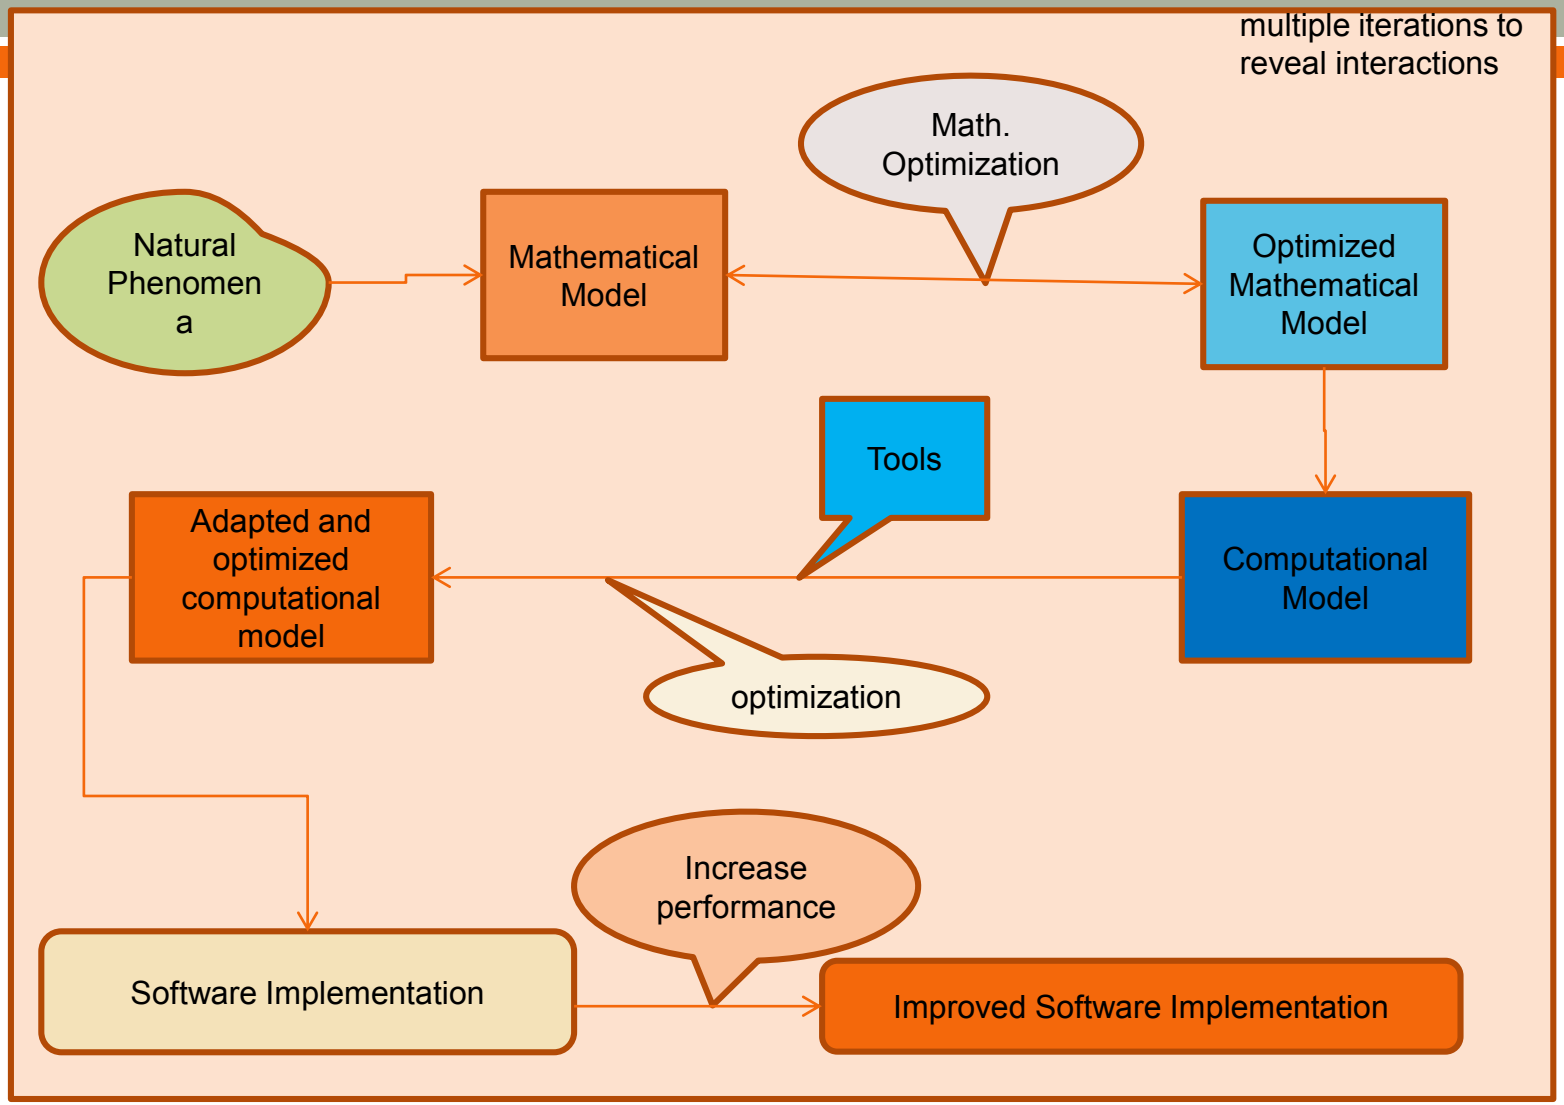

Collaborations

- *Faculty of Environment Science – ISUMADECIP Institute*
- *Faculty of Economics*
- *Faculty of Physics*
- *Faculty of Chemistry*
- *Faculty of Biology*
- *Faculty of Geography*

MOS Goals

Mathematics and Computer science – as a glue –

- **Projects**
  - Expertise in different research areas
  - Open to involvement in complex, national and international projects
- **Collaborations**
  - Other similar national and international centers
  - Foster international scientific collaborations
- **Development**
  - Increase the number of researchers
  - Increase the number of involved master and PhD students
  - Improvement of the infrastructure

MOS – interconnection point between different disciplines.

HPC Infrastructure

– Technical View –

Infrastructure - Hardware Architecture

Managed by **Faculty of Mathematics and Computer Science & Faculty of Economic Studies**

IBM Intelligent Cluster - Hybrid Architecture

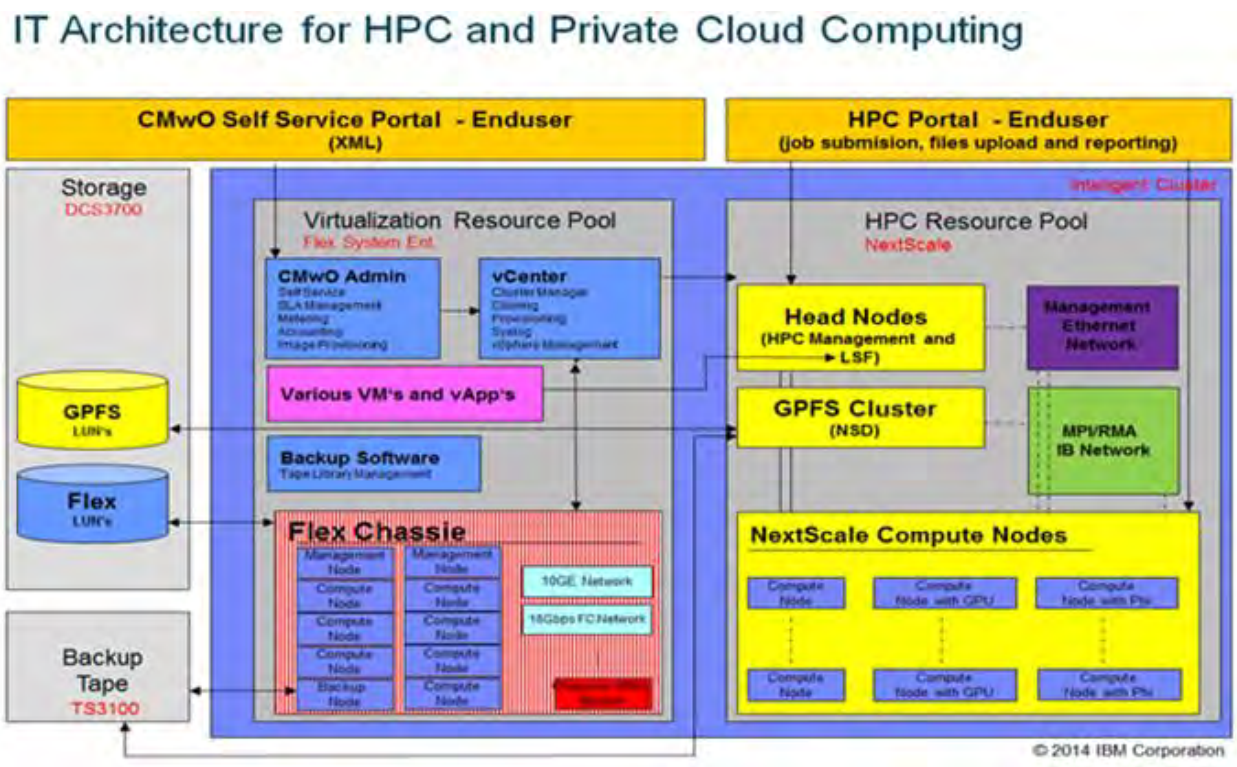

IBM Intelligent Cluster (2016)

- HPC Cluster + Cloud System
- Built by IBM (costed approx. 1 mil. EUR, without VAT)
- Consists of:
  - 4x 42U computing racks,
  - 4x cooling rack(+exterior unit),
  - 2 UPS enclosures
- Performance: Rpeak 62 Tflops (theoretical), Rmax 40 Tflops

IBM NextScale (HPC)

- 68 Nx360 M5 computing nodes
  - 50 nodes: 2x Intel Xeon E5-2660 v3 CPU, 10 cores; 128GB RAM; 2 HDD SATA 500 GB
  - 12 nodes with 2 Nvidia K40X GPU
  - 6 nodes with Intel Phi
- Fast networking: 56 Gb/s (Infiniband Mellanox FDR switch SX6512 with 216 ports, 1:1 subscription rate)
- Storage: IBM GPFS (General Parallel File System) NetApp E5660 Total: 72TB
- Backup: IBM TS3100 Tape library
- Mgmt. Software: IBM Platform HPC 4.2+xCAT and RedHat Linux Enterprise 6 for comp. nodes
- Others: 2 management nodes, 2 NSD, Fast Ethernet switches

IBM Flex System

- 10 virtualization servers Flex System x240
  - 128 Gb RAM / server
  - CPU 2 x Intel Xeon E5-2640 v2 / server
  - HDD 2 x SSD SATA 240 Gb / server
- 1 management server
- Cloud Software: IBM cloud manager with OpenStack 4.2
- Management and Monitor Software: IBM Flex System Manager software stack
- Virtualization Software: Vmware vSphere Enterprise 5.1

Development tools (Programming)

- Compilers: Gcc, Java, Intel C/C++ and Fortran compiler, Python interpreter
- 4 MPI libraries: OpenMPI, IntelMPI, MPICH, IBM/Platform MPI
- Job submission:
  - using CLI interface (bsub of LSF)
  - using the web interface
- Scientific libraries: Intel Mathematical Library

System facilities for development (Programming)

- Shared file systems:
  - /bigdata - 16TB redundant, working files partition
  - /home – user files and applications
  - /shared – applications and library
- SSH public key authentication
- Running commands on a set of nodes: ssh, pdsh

Application types

- **Single host, very intensive, non-distributed applications**
  - the recommendation is to use a cloud resource
- **MPI-based applications**
  - 4 MPI libraries: OpenMPI, IntelMPI, MPICH, IBM/Platform MPI
  - Executed through the LSF job scheduler
- **Non-MPI parallelized applications in a cluster environment**
  - hadoop, etc.

Thank you!