Computational Challenges and Needs for Academic and Industrial Applications Communities

Breakout Session 2

IESP, Saclay, June 28 & 29, 2009

Participants: th.lippert@fz-juelich.de, david.keyes@columbia.edu, andre@cerfacs.fr, lichnewsky@genci.fr, aoyagi@cc.kyushu-u.ac.jp, jean-philippe.nomine@cea.fr, guillaume.colin-de-verdiere@cea.fr, kale@illinois.edu, mike.ashworth@stfc.ac.uk, voevodin@parallel.ru, claude.puech@inria.fr, jy.berthou@edf.fr, stefan.heinzel@rzg.mpg.de, wolfgang.nagel@tu-dresden.de, giovanni.aloisio@unile.it

Objectives

- Establish roadmap to exascale for scientific domains
- Document software issues (type of issues, time frame)

Methodology

- Identify the application domains (HPC EUR, DoE)
- Identify the scientific and technical questions
- Identify the experts and « interview » them

The application domains and experts

Excel file

Computational Challenges and Needs for Academic and Industrial Applications Communities

Scientific and computational challenges

Brief overview of the underlying scientific and computational challenges and potential impact

Software issues – 2012, 2015, 2020

Expected scientific and technical hurdles

Software issues - 2009

Brief overview of identified software issues for addressing state of the art machines

Issues for experts

(from the last machine change, expectations for the next step)
(asking these questions may lead to a better understanding of the applications requirements, of the maturity of the domain wrt exascale computing)

- Load imbalance
- Communication scheme: synchronization, point-to-point and collective communications
- Scaling: weak, strong (time per step)
- I/O issues
- Parallel file systems
- Multilevel of parallelism: core and heterogeneity at core- and chip-level
- Memory issues: size per node, bandwidth, latency
- Measure of efficiency: per core, per node, global
- Solvers
- Pre- & post-processing: data management visualization, …
- ….
- Do not know

CS techniques

- Parallel code optimizations
- Parallel programming abstractions and support
- Adaptive optimizations (load balancing, etc.)
- Post-processing of large data sets, pre-processing
- Coupling multiple modules: computational schemes, control and data aspects, parallel composition of modules
- Data distribution, replication, integration, integrity, security
- Metadata, ontologies management
- Storage management and parallel I/O
- Integrated framework: workflow, common data models, components interoperability

Math techniques

- Partial differential equations
- Solution adaptive methods
- Molecular dynamics modeling and algorithmics
- Particle methods (distinct from classical MD in terms of inhomogeneous distributions, e.g., cosmology)
- DFT/electronic structure
- (hybrid) Monte Carlo
- Stochastic optimization
- Scalable implicit solvers (linear and nonlinear)
- Eigenanalysis techniques
- FFTs and other fast transforms
- Smoothed particle hydrodynamics
- Agent-based methods
- Coupling multiple models (stability aspects)
- Data assimilation/fusion of observations and models
- Data mining, clustering, classification, statistical analysis and representation
- Uncertainty quantification

Sample roadmapping exercise

- Nuclear reactor vessel (Code Saturne)
- Aerodynamics and aeronautics (Airbus France)
- Climate

Code_Saturne – HPC Roadmap: *application examples*

2003

Consecutive thermal fatigue event

Computations enable to better understand the wall thermal loading in an injection.

Knowing the root causes of the event ⇒ define a new design to avoid this problem.

2006

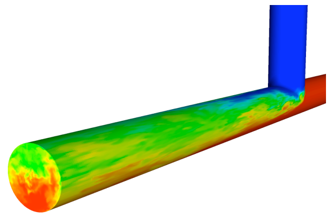

Computation with an L.E.S. approach for turbulent modelling

Refined mesh near the wall.

2007

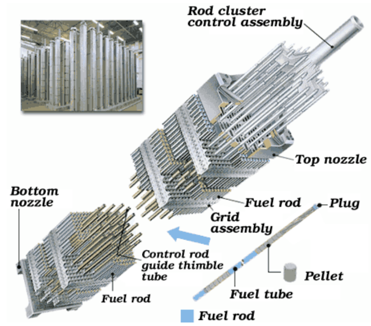

Part of a fuel assembly
3 grid assemblies

2010

9 fuel assemblies

No experimental approach up to now

Will enable the study of side effects implied by the flow around neighbour fuel assemblies.

Better understanding of vibration phenomena and wear-out of the rods.

2015

The whole vessel reactor

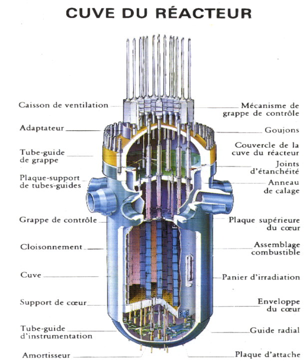

Computations with smaller and smaller scales in larger and larger geometries
⇒ a better understanding of physical phenomena ⇒ a more effective help for decision making
⇒ A better optimisation of the production (margin benefits)

2003	2006	2007	2010	2015
10^6 cells **3.10^{13} operations**	10^7 cells **6.10^{14} operations**	10^8 cells **10^{16} operations**	10^9 cells **3.10^{17} operations**	10^{10} cells **5.10^{18} operations**
Fujistu VPP 5000 1 of 4 vector processors 2 month length computation	Cluster, IBM Power5 400 processors 9 days	IBM Blue Gene/L « Frontier » 20 Tflops during 1 month	600 Tflops during 1 month	10 Pflops during 1 month
# 1 Gb of storage 2 Gb of memory	# 15 Gb of storage 25 Gb of memory	# 200 Gb of storage 250 Gb of memory	# 1 Tb of storage 2,5 Tb of memory	# 10 Tb of storage 25 Tb of memory
Power of the computer	Pre-processing not parallelized	Pre-processing not parallelized Mesh generation	… ibid. … … ibid. … Scalability / Solver	… ibid. … … ibid. … … ibid. … Visualisation

Weather, Climate and Earth Sciences: Roadmap

- Current simulations of climate change at **30-Km grid resolution** require a **1 Petaflop** sustained performance computer → i.e the Oak Ridge Nat.Lab's 149,504 cores Jaguar sytem (1 Petaflop sustained performance)

- More accurate, comprehensive and integrated Earth System Models (ESMs) are required

- Important processes in the climate system have to be simulated at their native spatial and temporal scales → critical organized features in the atmosphere and acean (including clouds and ocean eddies) have characteristic sizes of **1** to **10 Km**

- A global cloud-resolving modeling capability is needed within 5 years to predict "regional" changes in water, ice and clouds

Weather, Climate and Earth Sciences: Roadmap

- **Climate Coupled Model (Ocean-Atm-Med)**

  - Resolution: Atm: 80 Km T159L31 – Med: 1/16°

  - FLOPS required: about 3 *10 exp 14

  - Memory: about 110 GB

  - Storage: about 8TB

  - NEC-SX9 48 vector processors: about 40 days (century simulation)

  - bottleneck

    - inefficient memory access management (bank conflict)

    - I/O bandwidth

    - Load not balanced

- **Echam5 stand-alone**
  - Resolution: 50 Km - T255L31
  - - Cluster IBM Power6 450 cores : 3h (month simulation

Weather, Climate and Earth Sciences: Roadmap

2009

Resolution : 80 Km

$FLOPS \approx 3 * 10^{14}$

Memory: ≈110 GB

Storage: ≈ 8 TB

NEC-SX9 48 vector procs: ≈ 40 days run

2015

Resolution : 20 Km

$FLOPS \approx 1 * 10^{16}$

MemSory: ≈ 3,5 TB

Storage: ≈ 180 TB

High resolution model with complete carbon cycle model

Challenges: data viz and post-processing, data discovery, archiving

2020

Resolution : 1 Km

$FLOPS \approx 1 * 10^{19}$

Memory: ≈ 4 PB

Storage: ≈ 150 PB

Higher resolution with global cloud resolving model

Challenges: data sharing, transfer memory management, I/O management

Weather, Climate and Earth Sciences: Data Challenges

- Data distribution & Sharing
- Data Replication/Caching
- Metadata Agreement (Schema)
- Metadata Management (Automatic extraction, ingestion, validation, etc.)
- Metadata Harvesting
- Storage management, access, transport protocols, etc.
- Filesystem & Parallel I/O
- Scientific Data Gataway (unified access, integrated environment, transparent and secure access)
- Distributed Aggregation, Subsetting (slicing and dicing of data, etc,)
- Security (authorization, authentification, confidentiality, integrity, etc.)

Time schedule

Refine group notes with off-line research by panelists

- Consult as needed with experts, by mid-September
- Complete before 18-20 October workshop in Japan