EESI – Working Group 3.3

WP3 Task 3 - Fundamental Sciences - Chemistry, Physics

Chair: CECAM-FZJ (Godehard Sutmann)
Vice-Chair: CEA (Jean-Philippe Nominé)

Partner
Logo

EESI Time-Line

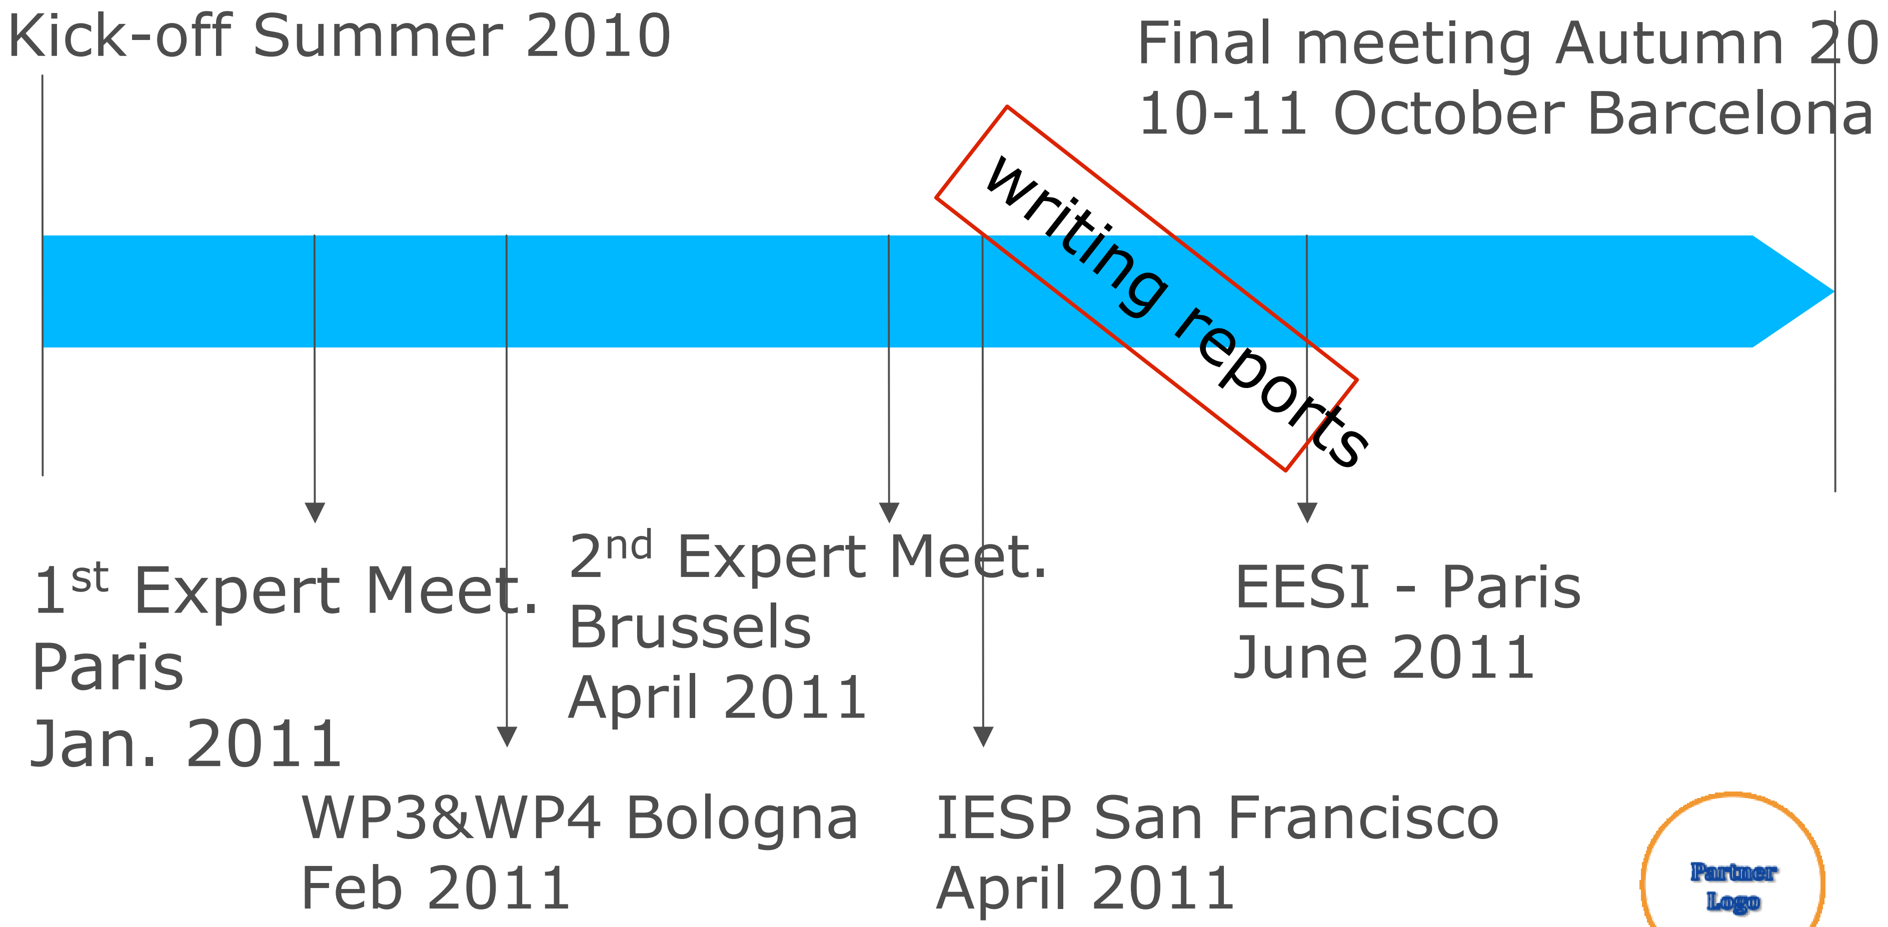

Kick-off Summer 2010

Final meeting Autumn 20
10-11 October Barcelona

writing reports

1st Expert Meet.
Paris
Jan. 2011

WP3&WP4 Bologna
Feb 2011

2nd Expert Meet.
Brussels
April 2011

IESP San Francisco
April 2011

EESI - Paris
June 2011

Scientific Domains in WG 3.3

- Fundamental sciences:
  Physics, Chemistry, Material Sciences, Astrophysics

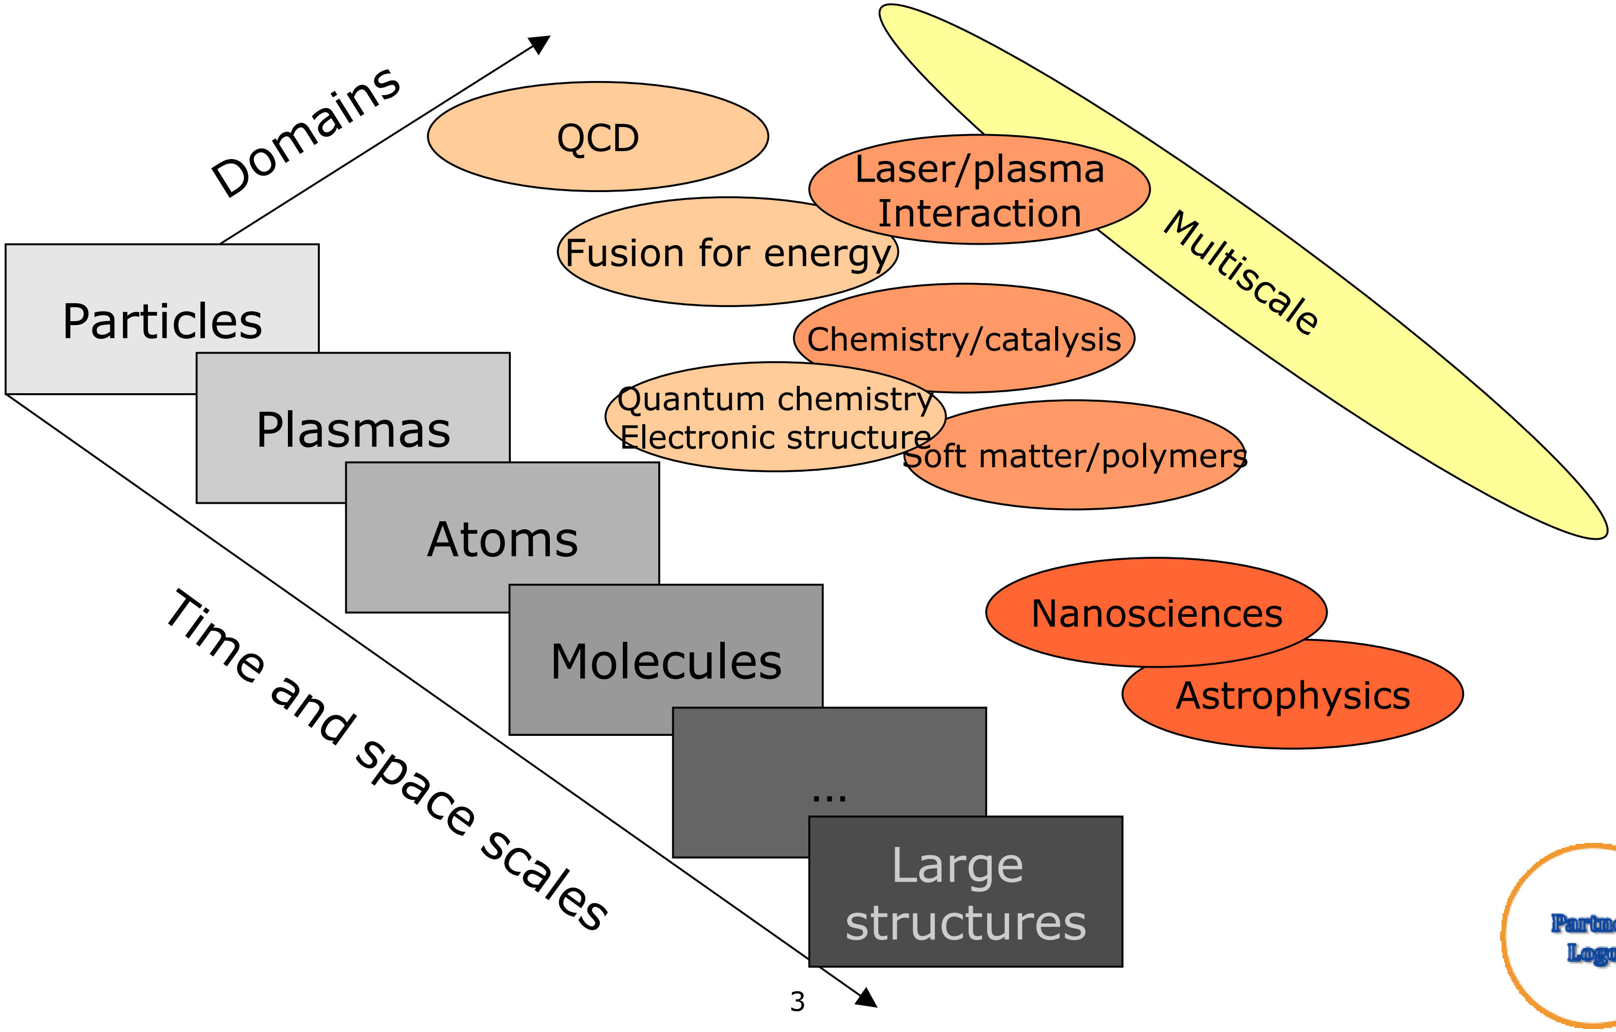

Composition of WG 3.3

- Description of the scientific and technical perimeter of the WG
  - Address science drivers and grand challenge problems in the fields of physics and chemistry
  - How are communities/science organizations going to prepare future software issues?
  - Workgroup composition:
    - Astrophysics
    - Laser- / plasma-physics
    - Fusion
    - Material Sciences
    - Quantum Chemistry
    - Soft Matter Physics
    - Software Engineering and Algorithms

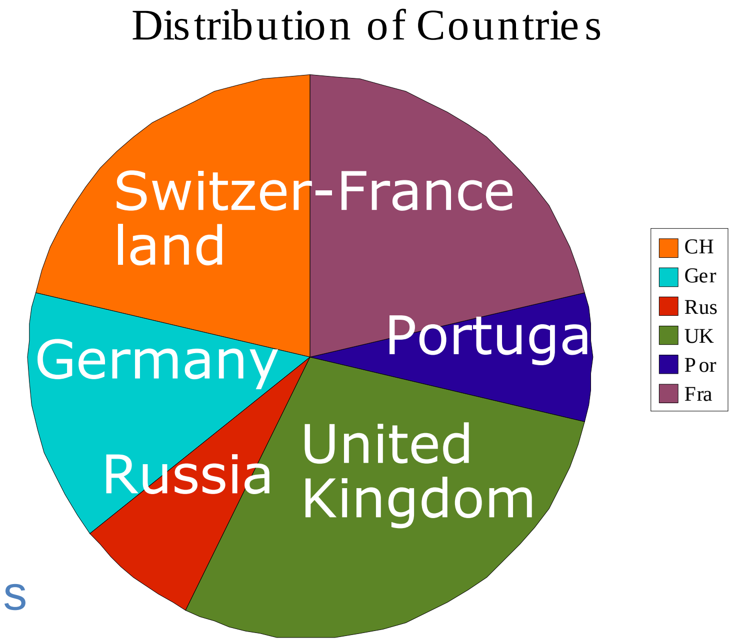

EESI

WG3.3 List of Experts

Name	Organization	Country	Area of Expertise
Volker Springel	Garching, MPI Astrophysik	Ger	Astrophysics
Romain Teyssier	ETH Zürich	Sui	Astrophysics
Maurizio Ottaviani	CEA	Fra	Fusion
Luis Silva	Universidade Tecnica de Lisboa	Por	Laser Plasma Interaction
Alexei Removich Kokhlov	Moscow State University	Rus	Soft Matter
Alessandro Curioni	IBM Research - Zurich	Sui	Materials Sciences
Gilles Zerah	CECAM - CEA	Fra	Materials Sciences
Nicola Marzari	University of Oxford	UK	Materials Sciences
Adrian Wander	STFC Daresbury	UK	Materials Sciences
Mike Payne	University of Cambridge	UK	Quantum Chemistry
Thierry Deutsch	CEA	Fra	Quantum Chemistry
Mike Ashworth	STFC Daresbury	UK	Methods and algorithms
Thomas Schultess	CSCS	Sui	Materials Sciences
Pieter in t'Veld	BASF	Ger	Soft matter

Ger: 2 Sui: 3 Fra: 3 Por: 1 Rus: 1 UK: 4 Total: 14

Partner Logo

Example for codes in WG 3.3

- Quantum Chemistry / Material Science
  - AbInit, BigDFT, CASTEP, ONETEP, CP2K, CPMD, Quantum-Espresso, Wannier90, Octopus, GPAW, Crystal, Dalton, Turbomole, Columbus
- Molecular Dynamics (Soft Matter)
  - DL_POLY, Gromacs, Espresso, LAMMPS, NAMD
- Laser / Plasma
  - TORB, ORB5, Euteurpe, ELMFIRE, GYSELA
- Astrophysics
  - Gadget, AREPA, PKDGRAV, Pluto, RAMSES

List not complete ..

Partner Logo

WG3.3

- Scientific activities and software issues
  - scientific activities
    - astrophysics: large scale structure of the universe
    - fusion: ITER facility
    - plasma: cold plasmas, magnetic plasmas
    - material sciences: catalysis, cracks, magnetic properties
    - soft matter: polymers, membranes for fuel cells, self-aggregation
    - algorithms: fault tolerance, energy efficiency, locality, optimal order algorithms

  - potential need for exascale performance proved

    - Material Science and Quantum Chemistry have potential for sustained PetaFlop applications at present
    - > 1PFlop/s sustained performance on Jaguar

Partner Logo

WG3.3 Performance in Material Science

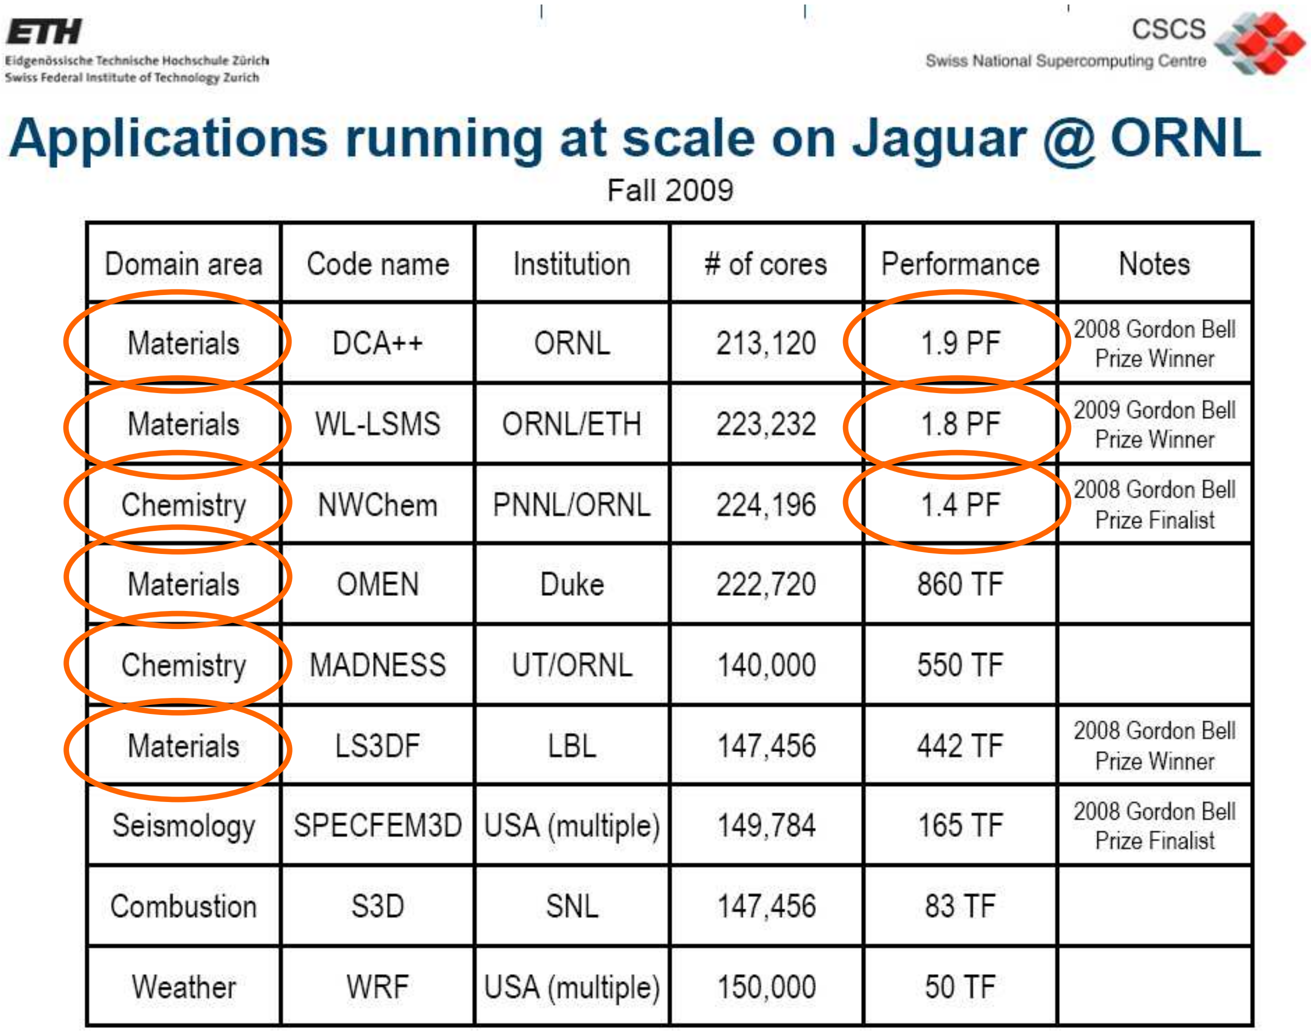

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

Applications running at scale on Jaguar @ ORNL

Fall 2009

Domain area	Code name	Institution	# of cores	Performance	Notes
Materials	DCA++	ORNL	213,120	1.9 PF	2008 Gordon Bell Prize Winner
Materials	WL-LSMS	ORNL/ETH	223,232	1.8 PF	2009 Gordon Bell Prize Winner
Chemistry	NWChem	PNNL/ORNL	224,196	1.4 PF	2008 Gordon Bell Prize Finalist
Materials	OMEN	Duke	222,720	860 TF	
Chemistry	MADNESS	UT/ORNL	140,000	550 TF	
Materials	LS3DF	LBL	147,456	442 TF	2008 Gordon Bell Prize Winner
Seismology	SPECFEM3D	USA (multiple)	149,784	165 TF	2008 Gordon Bell Prize Finalist
Combustion	S3D	SNL	147,456	83 TF	
Weather	WRF	USA (multiple)	150,000	50 TF	

WG3.3 Science Drivers

- **Materials Science**: first principles design of materials
  - e.g. energy conversion
    - chemical energy to electricity (fuel cells)
    - sunlight to electricity (photovoltaic cells)
    - nuclear energy to electricitiy (nuclear fission or fusion)
  - e.g. energy storage
    - batteries
    - supercapacitors
    - chemical storage in high energy density fuels (hydrogen, etanol, methanol)
  - e.g. energy use
    - solid state lighting, smart windows, low power computing, leightweight materials for transporting

WG3.3 Science Drivers

- **Materials Science**: first principles design of materials
  - High throughput design

    “Once a material property can be calculated accurately from quantum simulations, it becomes straightforward to systematically explore hundreds or thousands of compunds for improved performance.”

    N.Marzari
  - Acceleration of invention and discovery in science and technology
    - reduce costs
    - reduce interval of time-of-discovery to time-to-commercialisation

Partner Logo

WG3.3 Science Drivers

- **Soft-Matter Research**: overcome time- and length scales for device simulations

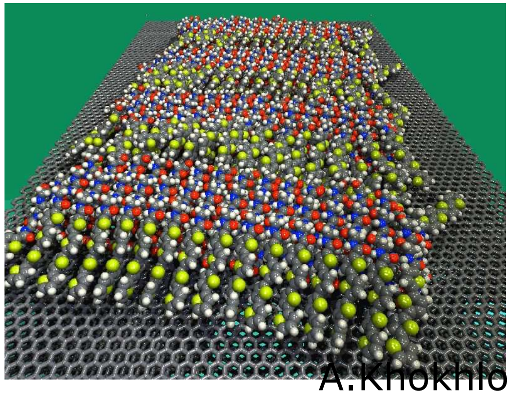

A.Khokhlov

  - Catalysis
    - temperature effects
    - non-equilibrium
    - chemical reactions
  - Self-organization and self-assembly of nano-structures
    - length- and time-scale wall (e.g. for polymers timescales ~N^2)
  - Coupled device modeling
    - Fuel cells: device modeling with explicit system description -> multiscale modeling for gas-liquid-solid + chemical reactions

WG3.3 Science Drivers

- **Plasma-Fusion Research**: ITER - magnetic fusion for energy production

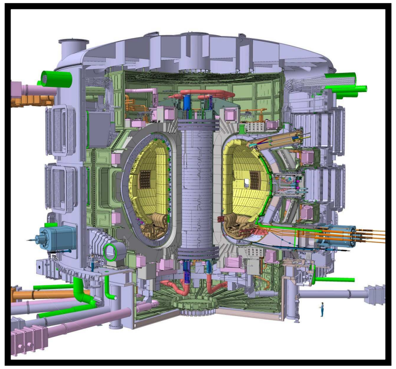

  - Characterisation of plasma dynamics
  - Understanding energy and particle confinement
  - Predict instabilities and their consequences
  - Challenges:
    - Spatial domains: electron scales (❍m), micro-turb. (mm), machine scal. (m)
    - Time-scale separation: energy confinement time (s), microturbulence (❍s)
    - Fully electromagnetic simulations, ions + electrons
    - time scale simulations up to ~1s (1ms at present)
    - spatial resolutions down to mesh sizes 10^5 -> electronic scales
    - memory limitations at present (need for fat nodes)

Partner Logo

WG3.3 Science Drivers

Astrophysics

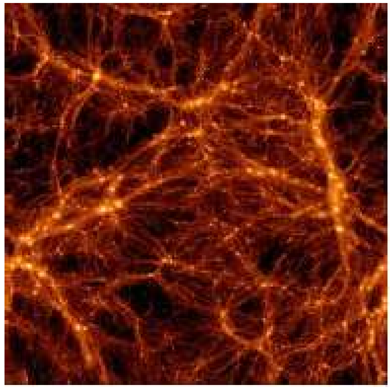

- Dark energy, large scale structure and cosmolgy
  - From galaxies to Hubble volume
  - Large n-body simulations (16384^3 part. for resolution) current limit 8192^3 particles – memory limits
- Physics of clusters and galaxy formation
  - very inhomogenous systems, different time scales – load balancing
  - need for lots of simulations to explore variance
- Physics of planetary formation
  - from molecular clouds to planetesimals
  - time- and length scales – need for multiscale modeling
  - new physics with MHD and dust-gas coupling
- Stellar interiors and model of the Sun
  - Supernovae mechanisms, compact stars
  - goal: global model of stellar structure & dynamics
  - MHD & radiative transport

Partner Logo

WG3.3

- Main methods used in application fields:
  - Particle methods
    - e.g. molecular simulations, particle based hydrodynamics
    - PIC, MD, Brownian Dynamics, Monte Carlo
    - Long range- / non-local interactions
  - Mesh-based methods
    - e.g. Navier-Stokes, MHD
    - Adaptive mesh-refinement
    - Multigrid
    - FFT
  - Ab inito / electronic structure calculations
    - Linear algebra (e.g. Eigenwert solver, Cholesky, matrix-vector)
    - Wavelets

WG3.3

- Goals to be achieved
  - reduce time-to-solution / time-to-market
  - reduce / optimize energy-to-solution
    - new metric for job-accounting foreseen: (time-to-solution) x (energy)
  - Algorithmic targets:
    - seek for optimal complexity (O(N))
      * Fast Multipole Methods
      * Decomposition methods
      * Multigrid
      * H-matrices
    - seek for locality (reduce data movement and therefore energy)
      * communication-friendly or -avoiding algorithms
      * time-scale splitting schemes
      * real space methods
      * wavelets

WG3.3

- Preparing for the next steps: **Short term perspectives**
  - Modularisation of codes
    - share components of codes between different groups
  - Most groups start thinking in terms of
    - extending codes to hybrid:
      MPI + OpenMP / P-Threads
    - writing codes or parts of codes for:
      GPU
    - planning extensions for codes in multi-stage parallelism
      MPI + OpenMP + accelerator (GPU, FPGA,...)

WG3.3

- Preparing for the next steps: **Long term perspectives**
  - The exa-scale challenge – the threefold way
    - Strong- / weak scaling
    - Multiscale (horizontal / vertical)
    - Ensemble simulations

WG3.3: strong- / weak-scaling

- Strong- / weak scaling
  - Common opinion of experts:

    - do not build up on existing codes
    - rewrite legacy codes
    - adjust / choose algorithms for exa-scale hardware
    - address hardware specific features and design special algorithms for specific hardware features

    - implies several man-years for redesign of functional units in programs

    New design and optimal implementation of codes

WG3.3: Multiscale Simulations (I)

- Multiscale Simulations

  - partial solution to escape from the dilemma of hyper-scaling
    - * strong scaling not possible for a lot of codes
      (e.g. for some commercial codes or quantum chemistry)

  - solution to weak scaling (WS) problem, since WS not always
    - * desired (e.g. in Quantum Chem. not everything is “worth” to
      be calculated in full precision)
    - * or possible (e.g. non-linear increase in memory consumption)

  - combine codes and run concurrently

Horizontal Multiscale

WG3.3: Multiscale Simulations (II)

- Multiscale Simulations

  - not all experts need / want full exascale performance in a single program application

    - solution is in running different codes coupled simultaneously
    - run codes concurrently on smaller number of nodes

    - **chance for a survival of legacy codes**

    - codes should be coupled through *standard interface*
    (this has to be developed and agreed for in the community)

    - develop multiscale simulation codes which can be used as
    plug & play

Partner
Logo

WG3.3: Ensemble Simulations

- Ensemble Simulations
  - “perfect parallel scaling”

    - used to increase statistics
    - verification of methods
    - method towards fault tolerant computing

WG3.3 Summary

- Software issues for Fundamental Sciences
  - Fault tolerant and energy efficient algorithms
  - Software support to measure or estimate energy
  - Data locality
  - Optimal order algorithms
  - Mesh generation
  - Algorithms with low communications
  - Standard interfaces for multiscale simulations
  - Support of several executables in one job
  - Parallel I/O
  - Load-balancing

Partner Logo