JÜLICH
FORSCHUNGSZENTRUM

Simulation Laboratories: A Community-Oriented Software Support & Development Strategy for HPC

Paul Gibbon

Jülich Supercomputing Centre

1st JUQUEEN Porting & Scaling Workshop, 4-6 February 2013

Supercomputer Usage by Scientific Discipline

Leadership-Class System

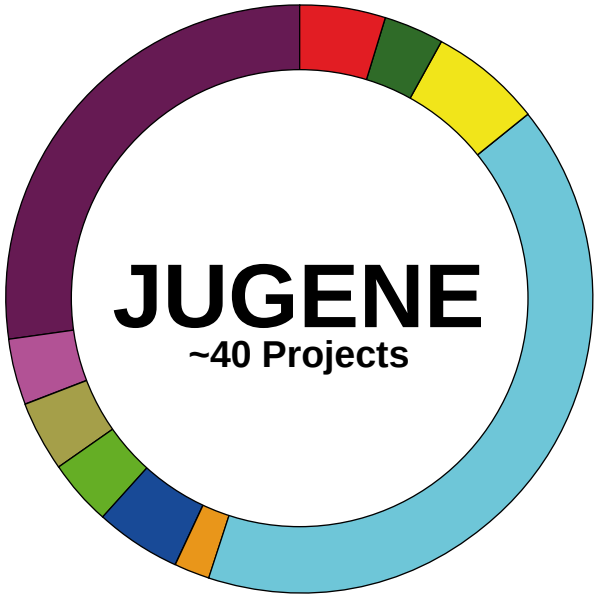

General-Purpose Supercomputer

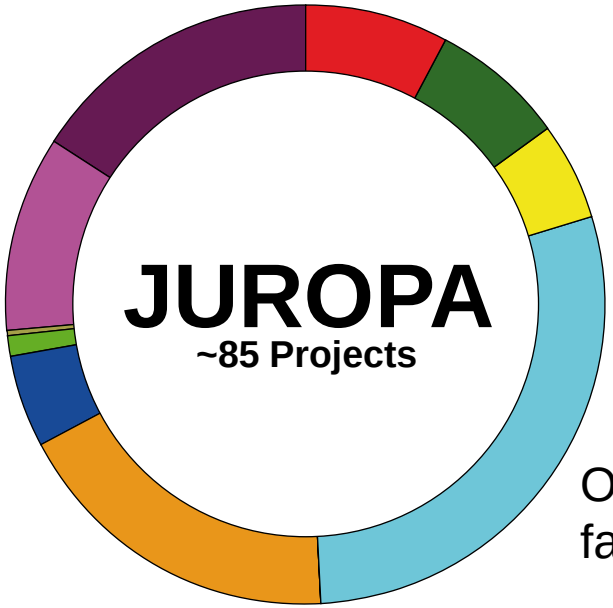

Oversubscription factor > 5

- Astrophysics
- Biophysics
- Chemistry
- Earth & Environment
- Plasma Physics
- Soft Matter
- Fluid Dynamics
- Elementary Particle Physics
- Computer Science
- Condensed Matter
- Material Science

Traditional application support model at JSC

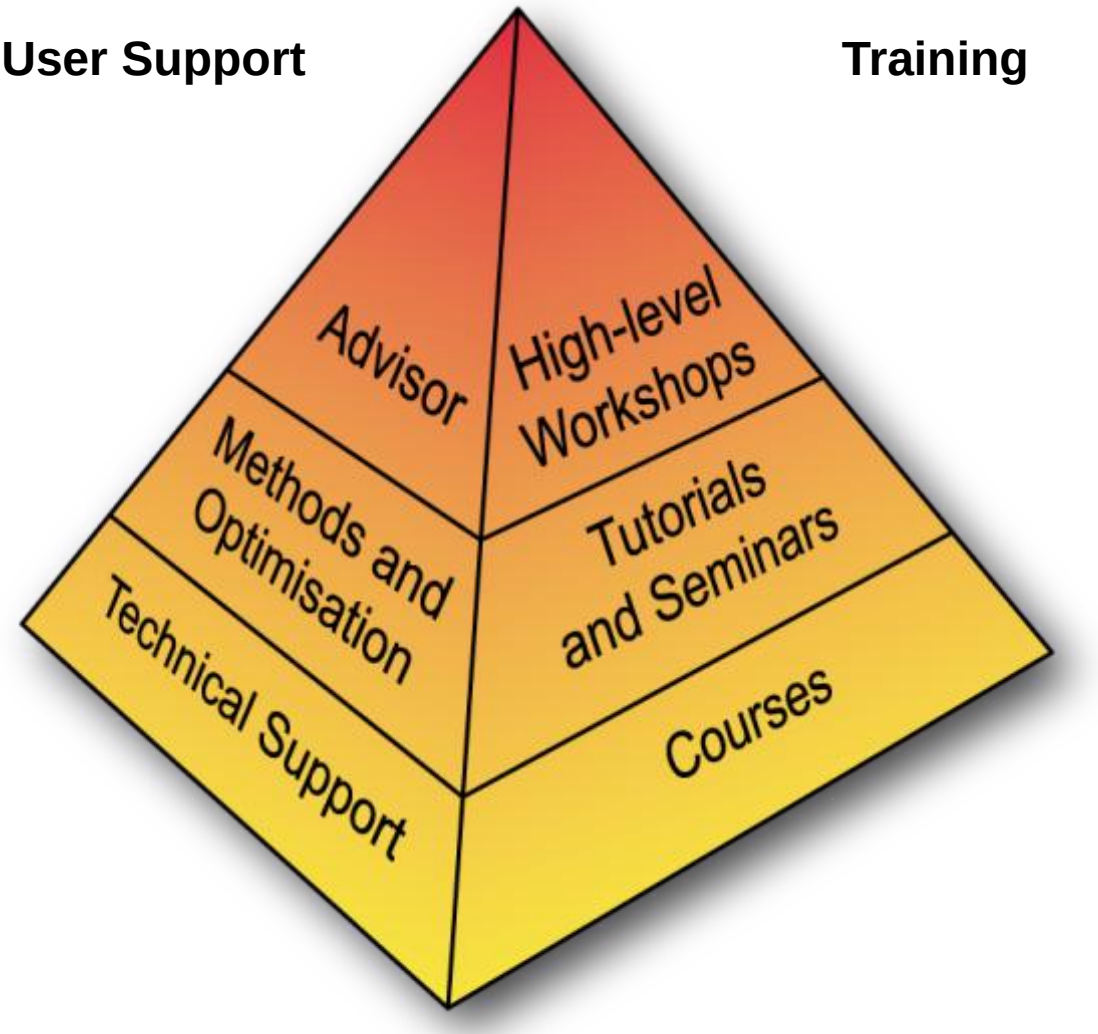

Exascale is placing a strain on software efficiency

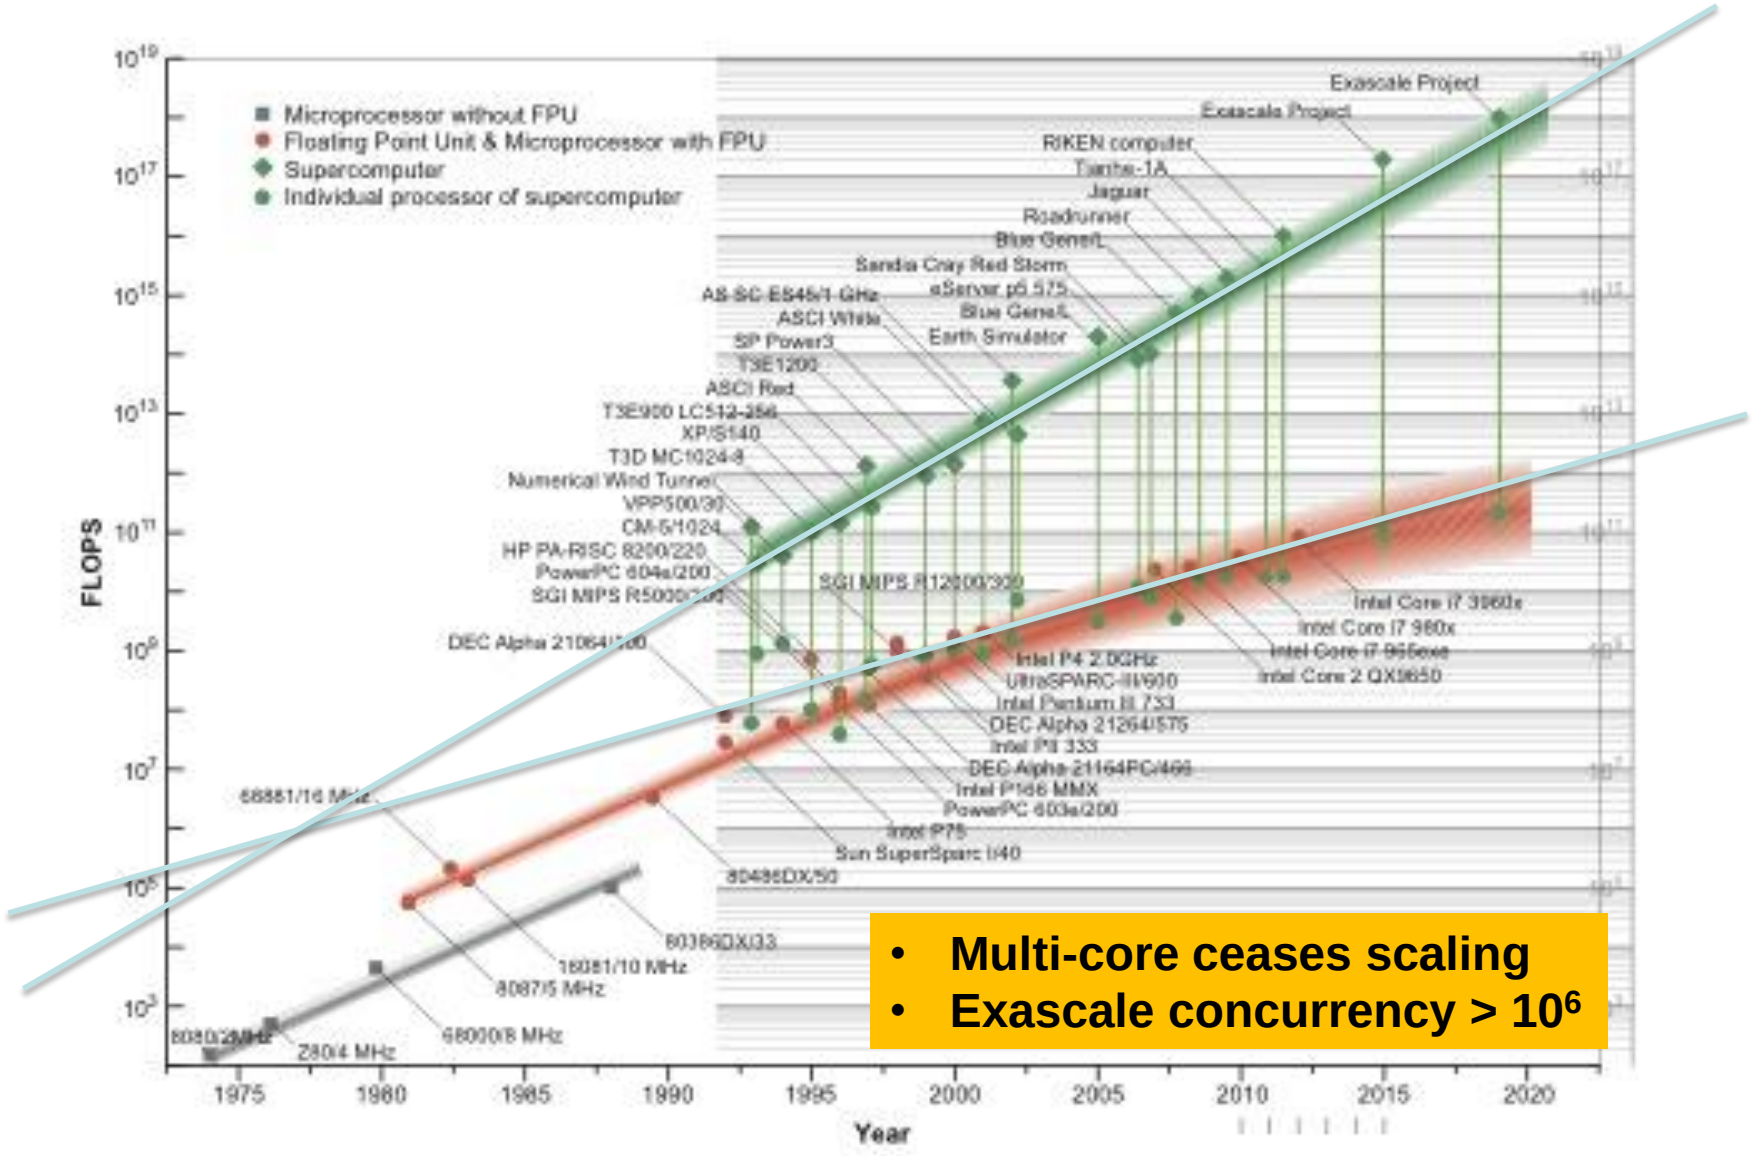

- **Multi-core ceases scaling**
- **Exascale concurrency > 10^6**

~ 2005 AD: first brainstorming @ JSC

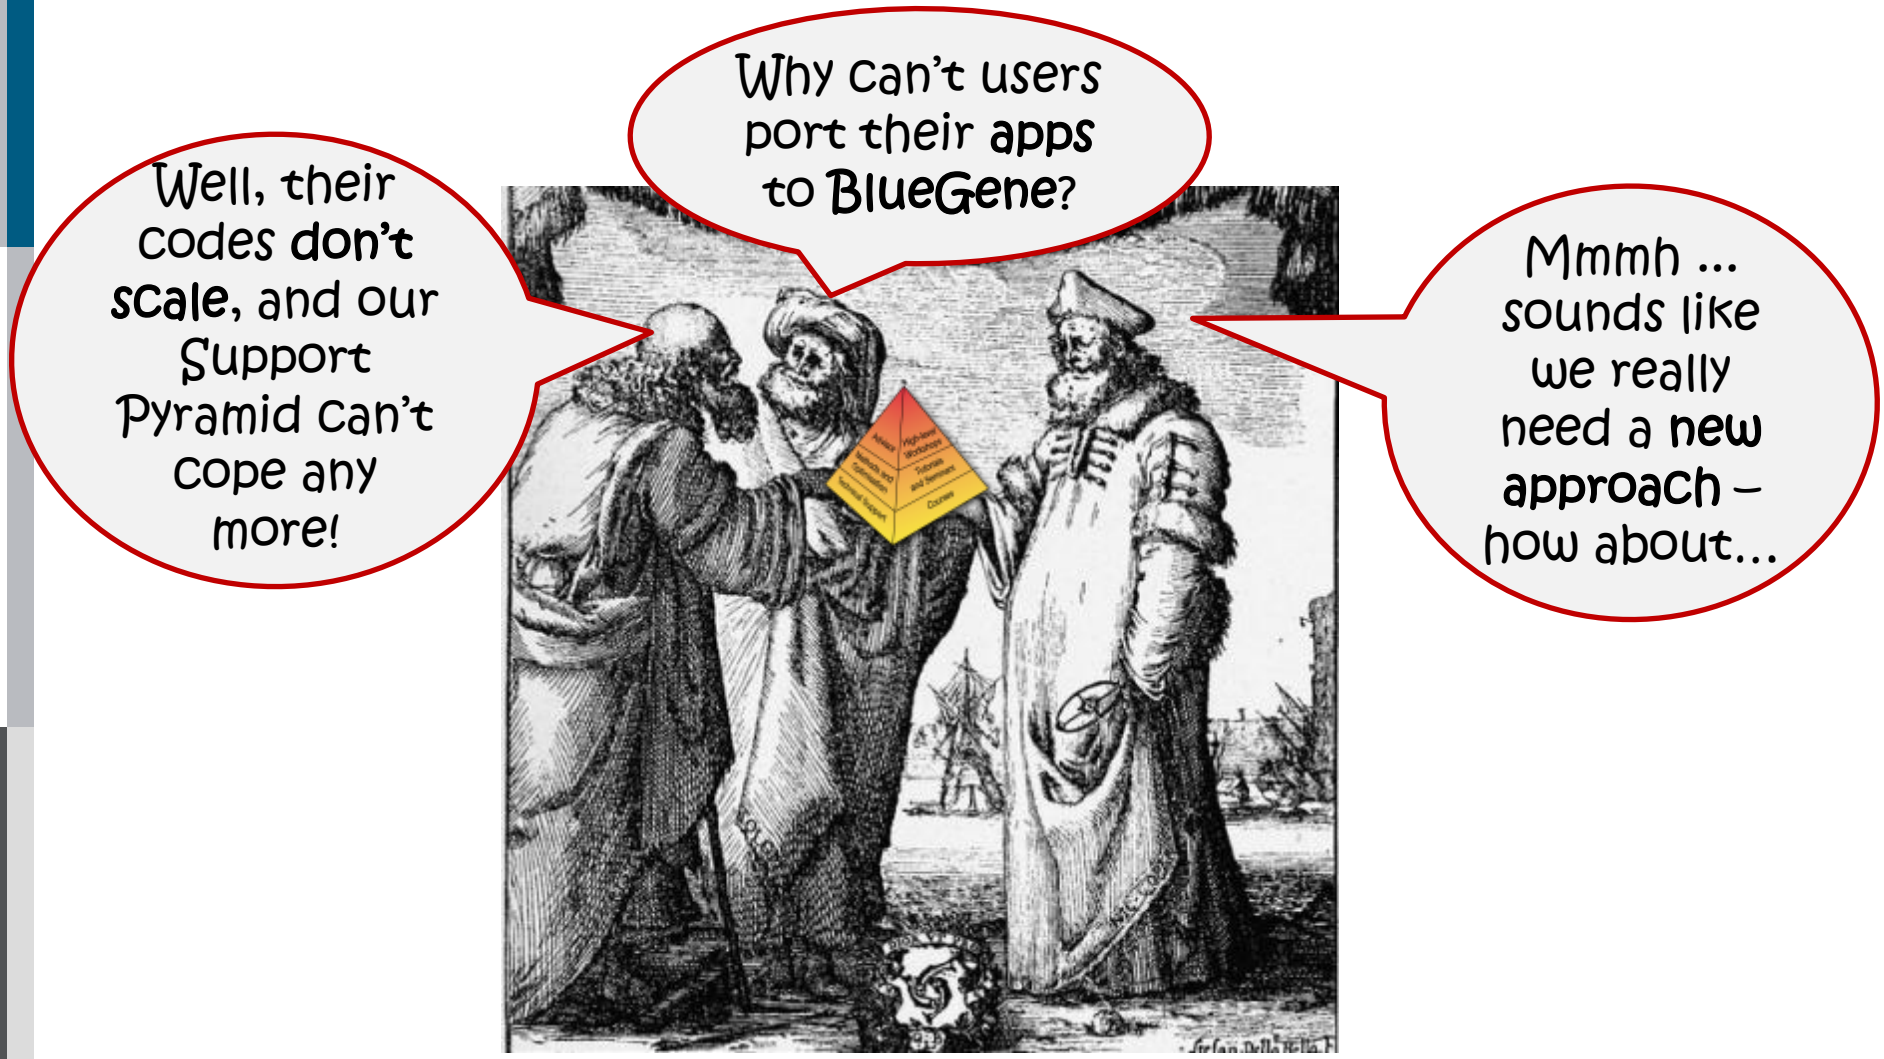

… Simulation Laboratories

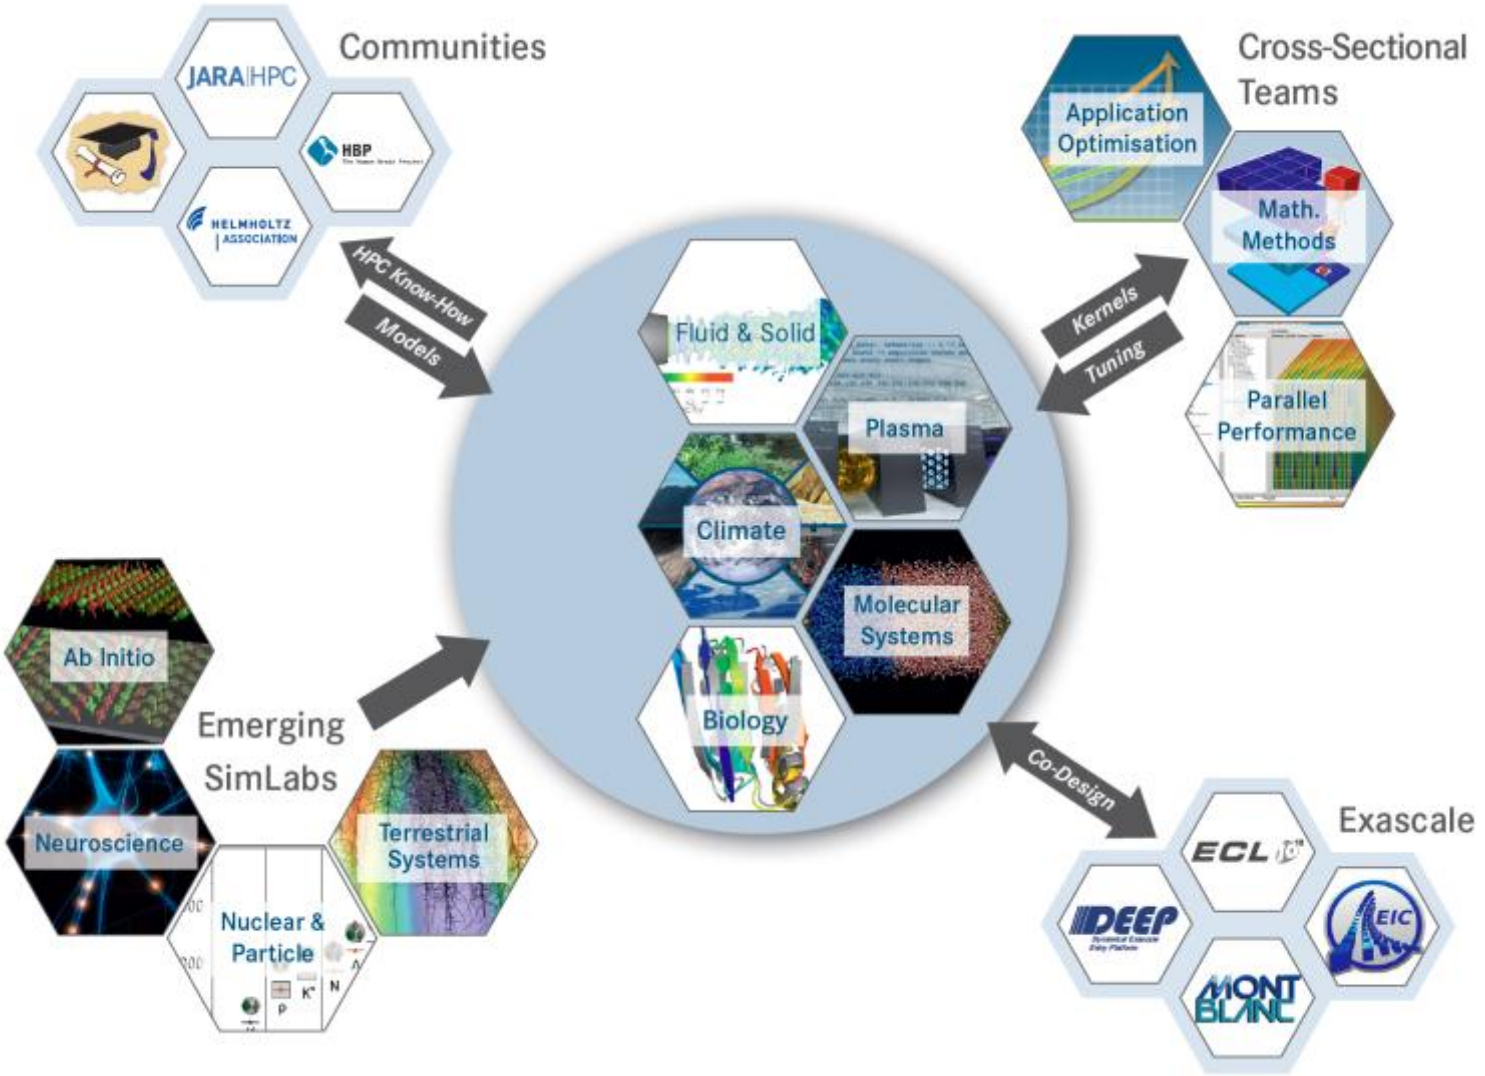

Active Simulation Labs @ JSC

Biology

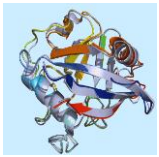



Climate Science

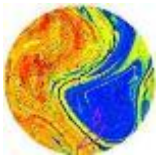


Molecular Systems









Terrestrial Systems

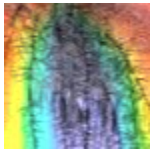


Nuclear & Particle

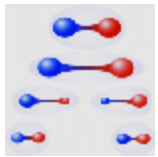


Fluid & Solid Eng.

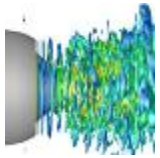


Plasma Physics

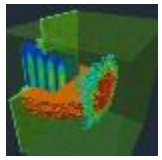





Ab Initio

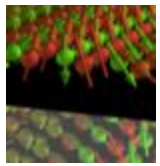


Simulation Labs: Structure

- **Staff**
  - Senior scientist recruited from field
  - 1-2 postdocs
  - 1 technical staff (informatics)
  - Jointly supervised PhD & MSc students
- **Support**
  - Porting/tuning/benchmarking; code hosting
  - Algorithm scaling; code clinics
  - Workshops; schools
- **Research**
  - Common/generic simulation methods
  - Scalable algorithms
  - Project work with SL partners

SimLab Activities

1. **NIC/VSR advisory**
   - Project mentoring; Parateam; technical & scientific reviews
2. **Code Clinics**
   - Short visits to JSC: hands-on performance checks & tuning
3. **Training workshops**:
   - 1st Porting Workshop, June 2010; Heraeus Summer School, Sept 2010
   - CECAM Workshops: September 2011, March & Sept 2012
4. **Advanced application support**
   - http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/2ndSupportCall
   - Source-code tuning, redesign, refactoring, scaling
5. **Research cooperations**
   - Research groups: FZJ, JARA, DE, EU, RoW
   - 3rd party projects; HGF, BMBF, DFG, EU, G8
   - Exascale labs EIC, ECL, Nvidia – ***co-design***

SimLab Activities

1. **NIC/VSR advisory**
   - Project mentoring; Parateam; technical & scientific reviews
2. **Code Clinics**
   - Short visits to JSC: hands-on performance checks & tuning
3. **Training workshops**:
   - 1st Porting Workshop, June 2010; Heraeus Summer School, Sept 2010
   - CECAM Workshops: September 2011, March & Sept 2012
4. **Advanced application support**
   - http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/2ndSupportCall
   - Source-code tuning, redesign, refactoring, scaling
5. **Research cooperations**
   - Research groups: FZJ, JARA, DE, EU, RoW
   - 3rd party projects; HGF, BMBF, DFG, EU, G8
   - Exascale labs EIC, ECL, Nvidia – ***co-design***

Advanced Support Concept

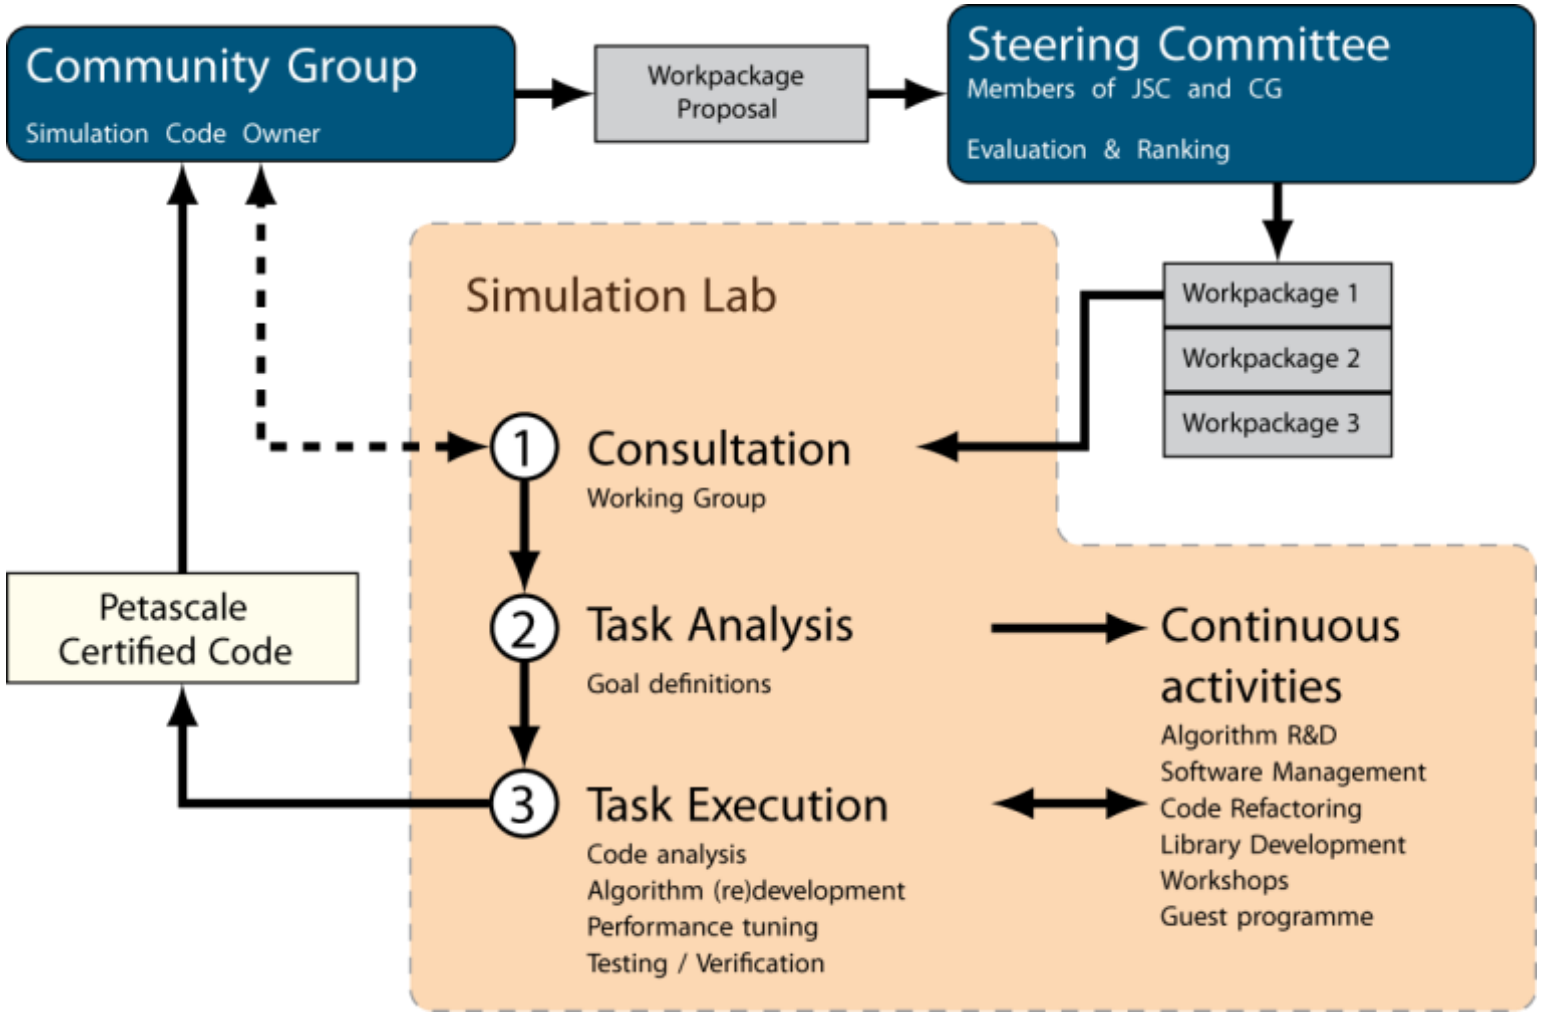

Advanced Support Pilot Project 2010/11

- 1st Call issued October 2010
- SimLabs involved:
  - Biology
  - Molecular Systems
  - Plasma Physics
  - Climate Research
- 23 applications; 18 accepted (~ 4 per SimLab)

Seeking finite-time singularities in CFD with *racoon* (U. Bochum)

- Scaling enhancement: 2 racks ⇨ 32 racks: increase in *resolution*

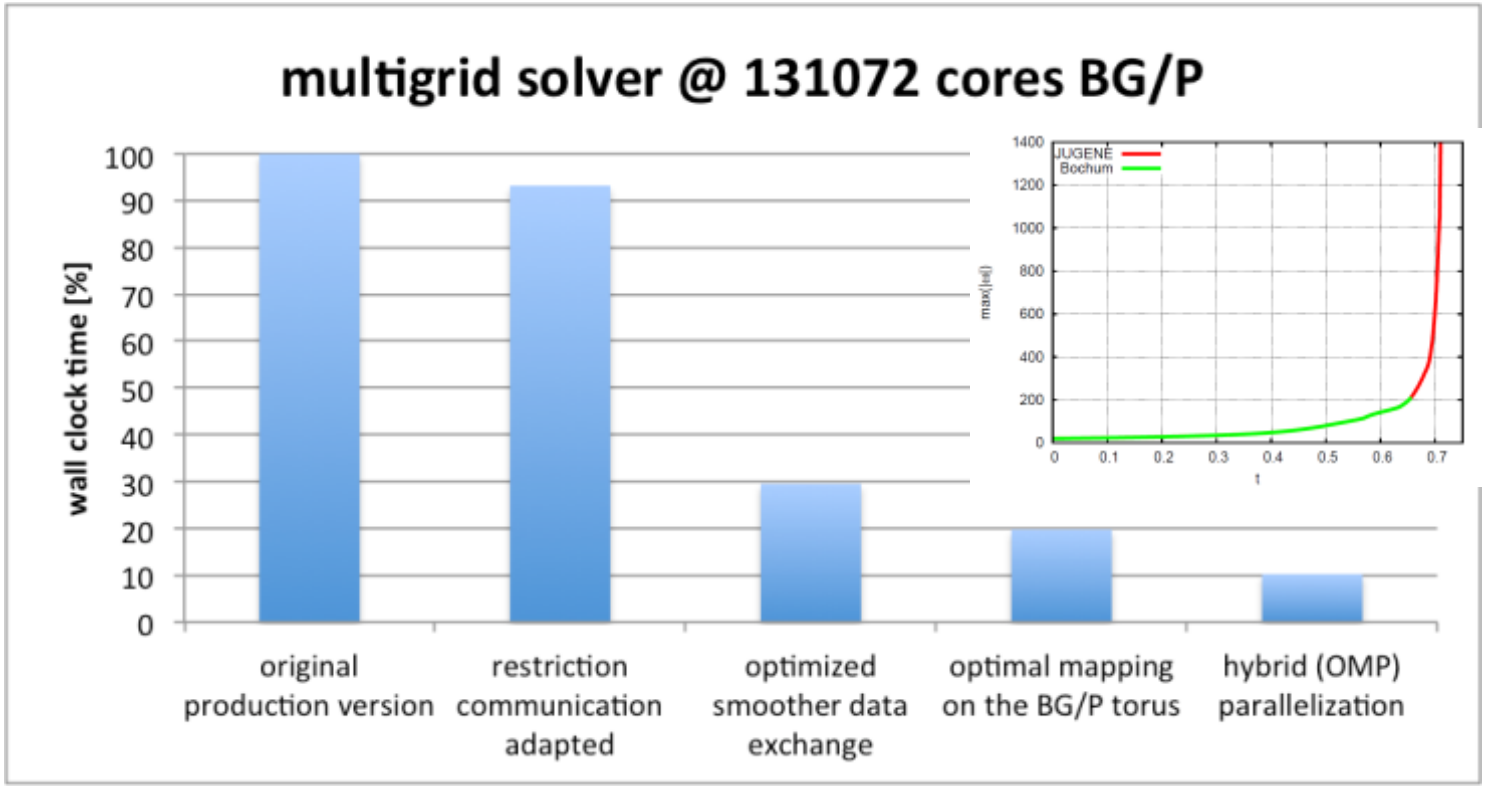

Advanced Support Initiative: lessons learned

Outcomes:

- 1-2 PMs effort (from JSC) per proposal; variable in-kind
- Enhanced scaling in 50% of apps
- Follow-ups: i) proposals for computing time ii) 3rd-party projects iii) joint publications
- Report available here:

http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/2ndSupportCall

Conclusions:

- Better project vetting to sift out non-starters
- Closer connection to established NIC/VSR projects
- Selected ‚priority' projects (6PMs) to encourage joint dissemination

SimLab Activities

1. **NIC/VSR advisory**
   - Project mentoring; Parateam; technical & scientific reviews
2. **Code Clinics**
   - Short visits to JSC: hands-on performance checks & tuning
3. **Training workshops**:
   - 1st Porting Workshop, June 2010; Heraeus Summer School, Sept 2010
   - CECAM Workshops: September 2011, March & Sept 2012
4. **Advanced application support**
   - http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/2ndSupportCall
   - Source-code tuning, redesign, refactoring, scaling
5. **Research cooperations**
   - Research groups: FZJ, JARA, DE, EU, RoW
   - 3rd party projects; HGF, BMBF, DFG, EU, G8
   - Exascale labs EIC, ECL, Nvidia – ***co-design***

ScaFaCoS: BMBF project 2009-2012

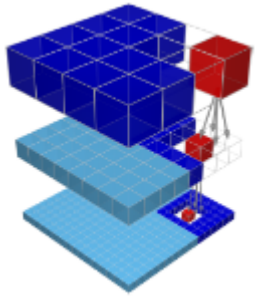

- Fast Electrostatics Library
  - Unified parallel library for various methods of long range interactions
  - Multiple boundary conditions: open, 1d-,2d-,3d-periodic
  - Error control
  - OpenSource distribution under LGPL license

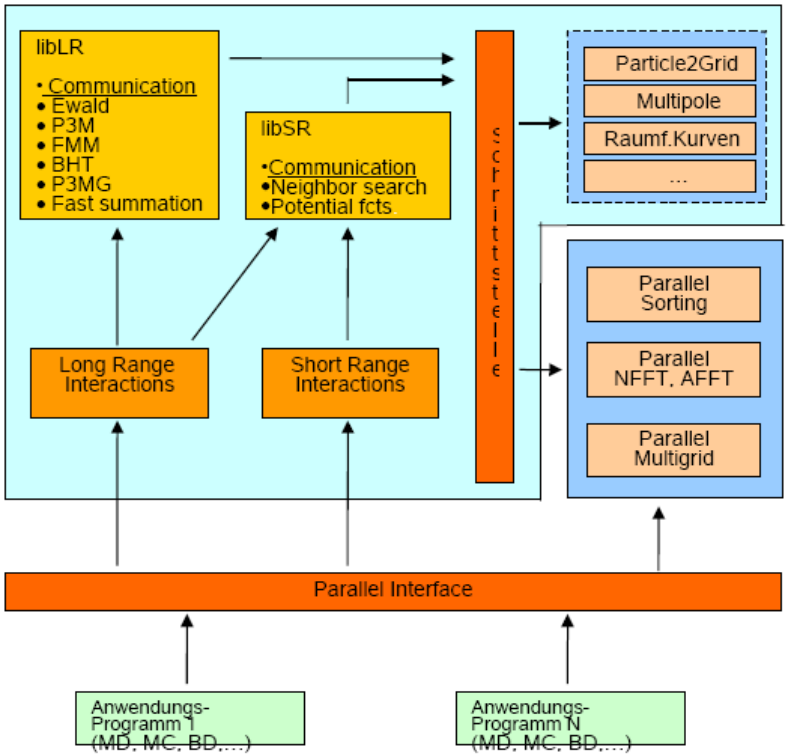

Exascale Lab support

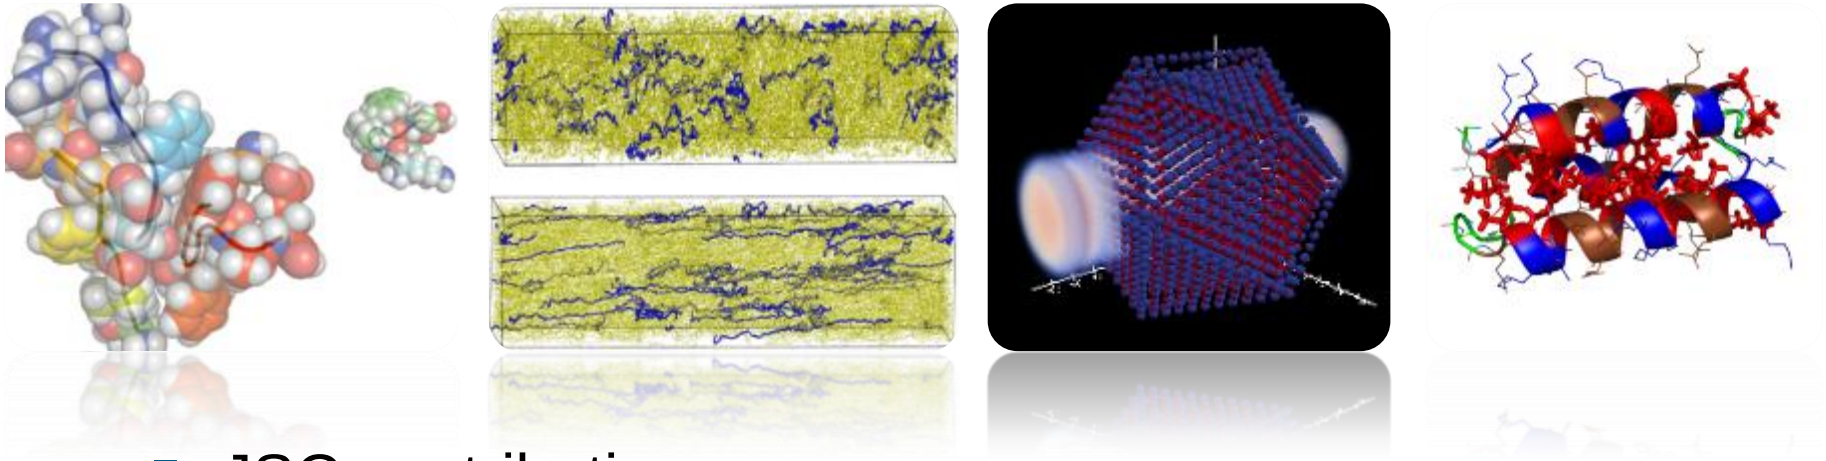

- JSC contributions:
  - Tools: Scalasca
  - Applications: SMMP, MP2C, PEPC, ProFASI
- JSC (SimLab) participants:
  - Dirk Broemmel, Thomas Fieseler, Paul Gibbon, Rene Halver, Daniel Lorenz, Jan Meinke, Sandipan Mohanty, Bernd Mohr, Annika Schiller, Mark Schlütter, Godehard Sutmann, Olav Zimmermann

Example: parallel tree code PEPC

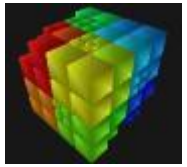

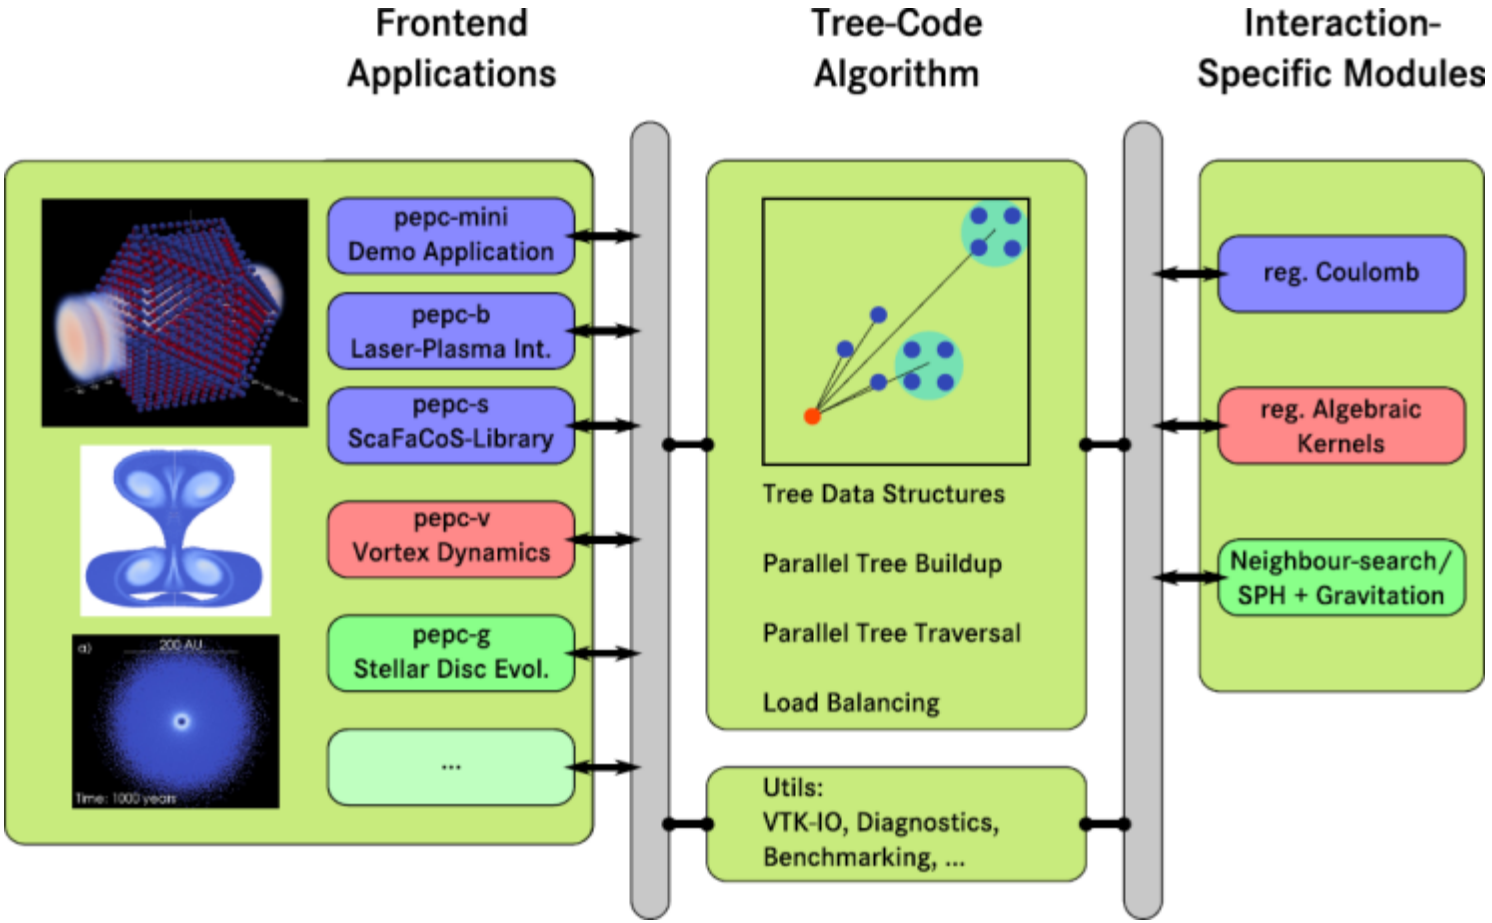

https://trac.version.fz-juelich.de/pepc

Scaling of pure MPI scheme: breakdown

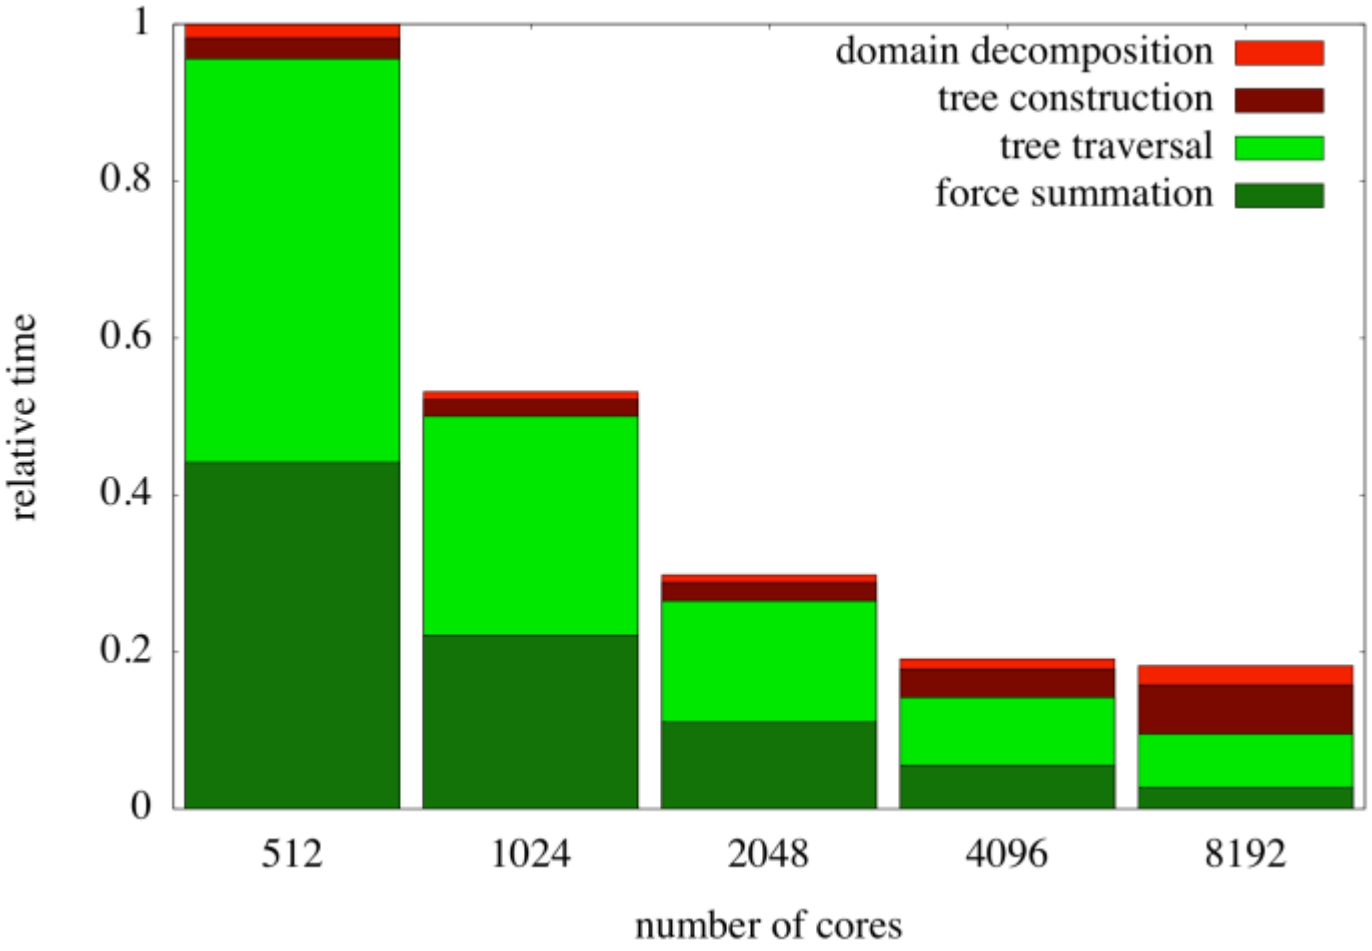

R. Speck, L. Arnold, PG, J. Comp. Sci. 2, 138 (2011)

Hybrid scheme: basic idea

Pure MPI

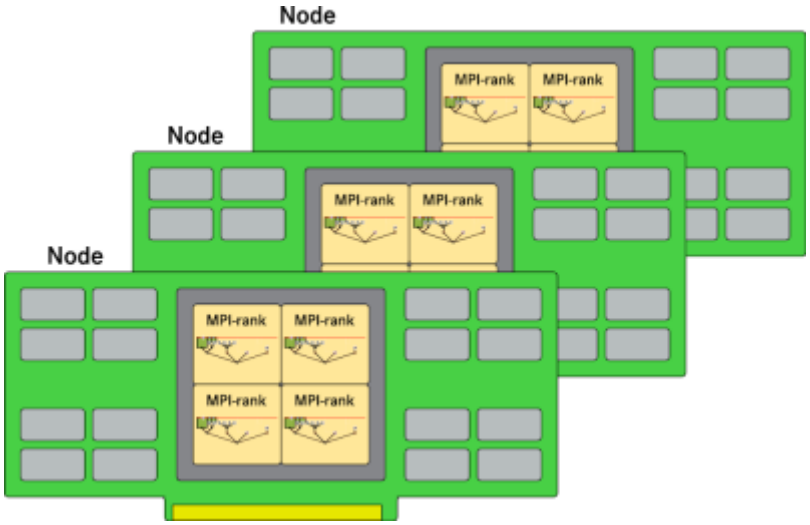

Hybrid

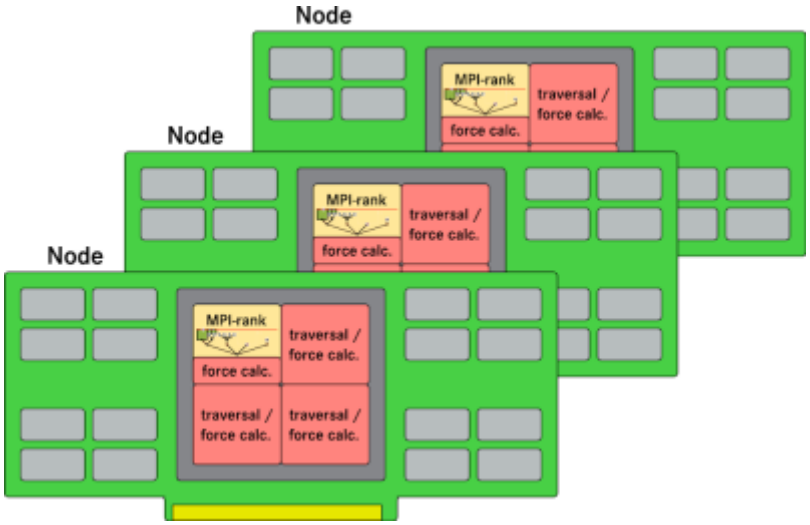

- Fewer multipole copies = less MPI comm. and lower storage
- Reduced reply latency for remote tree node requests

Hybrid algorithm: P-threads + MPI

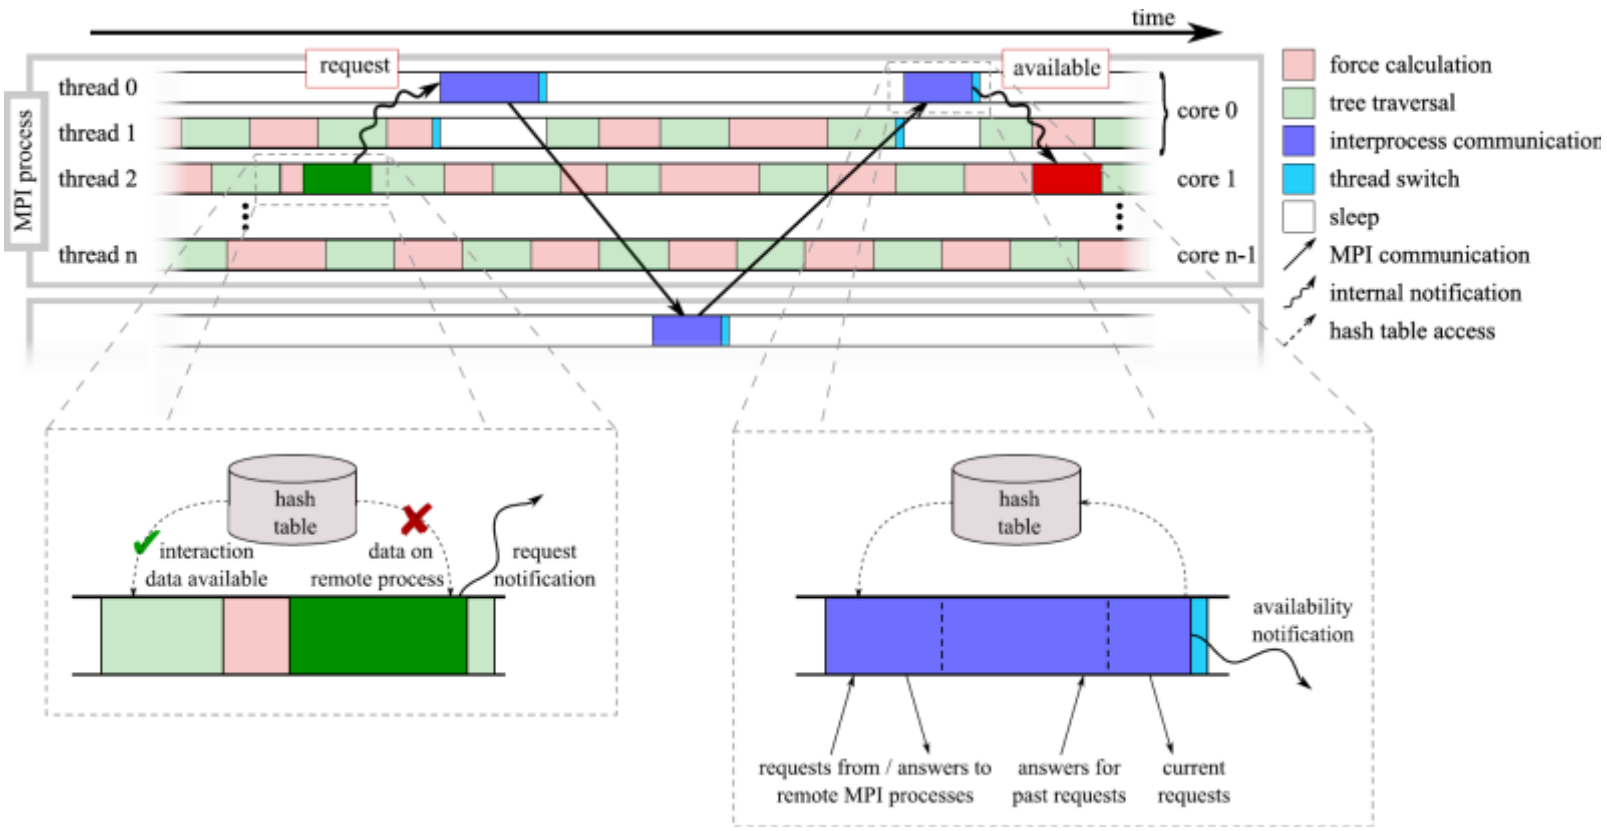

Scaling things up: tree code world record (2010)

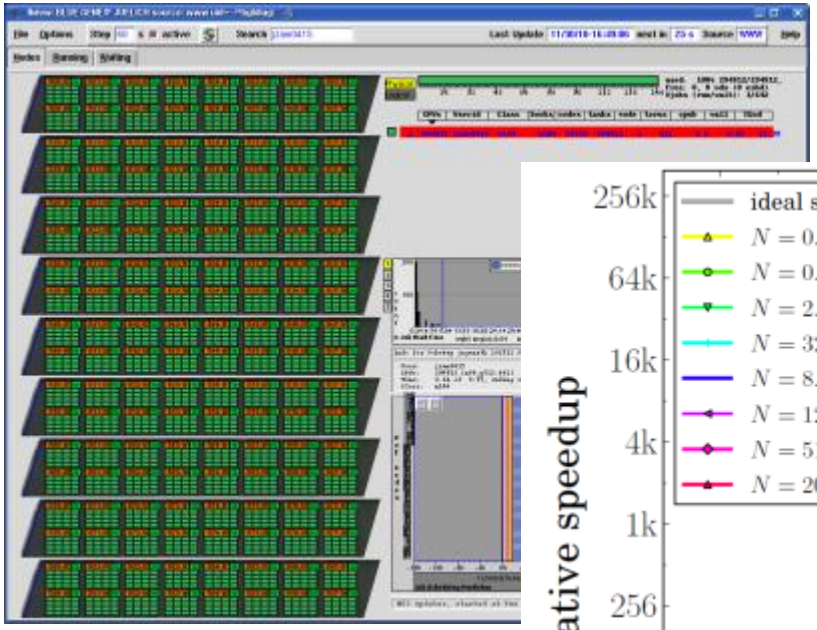

2×10^9 particles on 294 912 tasks of BlueGene/P, load-balanced

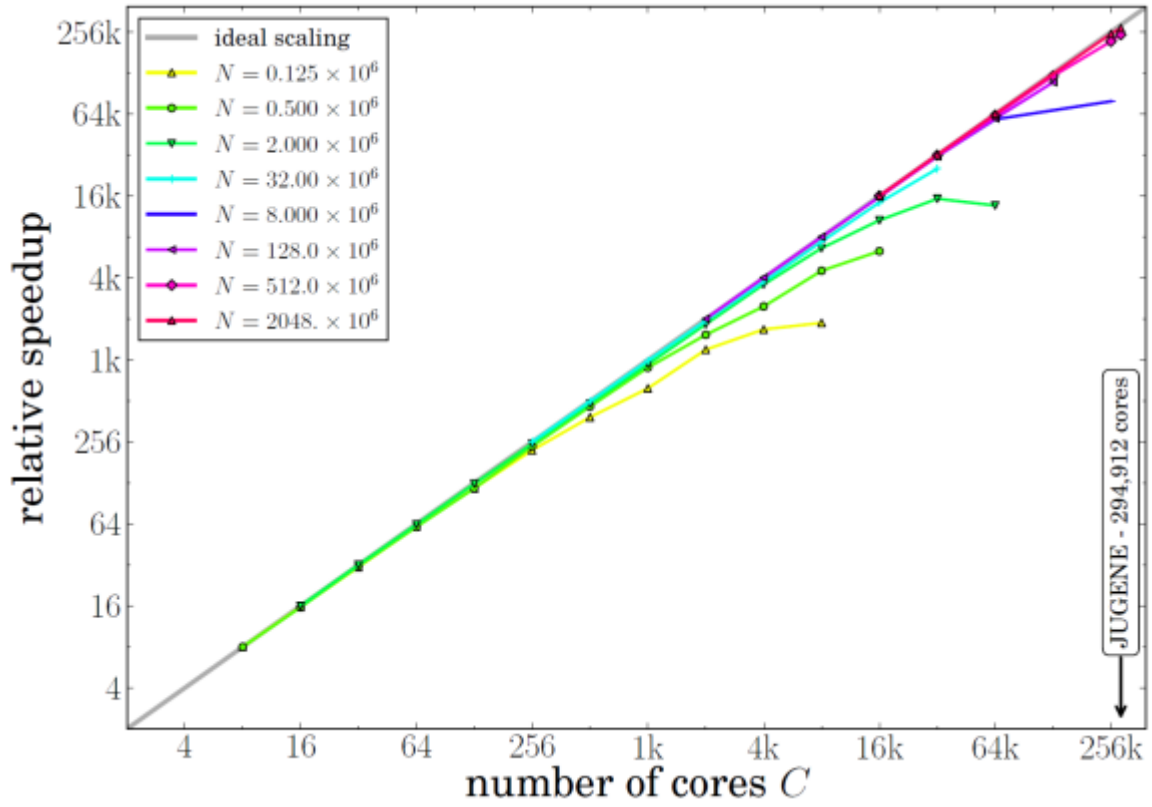

M. Winkel et al., Comp. Phys. Commun. **187**, 880 (2012)

BG/P vs BG/Q

	BG/P (Jugene)	BG/Q (JuQueen)
Processor	PowerPC 850MHz	PPC A2 1.6Gz
Cores/Node	4	16
Memory/Core	512 MB	1 GB
Nodes/Rack	1024	1024
Cores/Rack	4096	16386
# racks	72	28
# cores	294 912	458 752
Compute power	1 Pflop/s	5.9 Pflop/s

Sweet spot determination

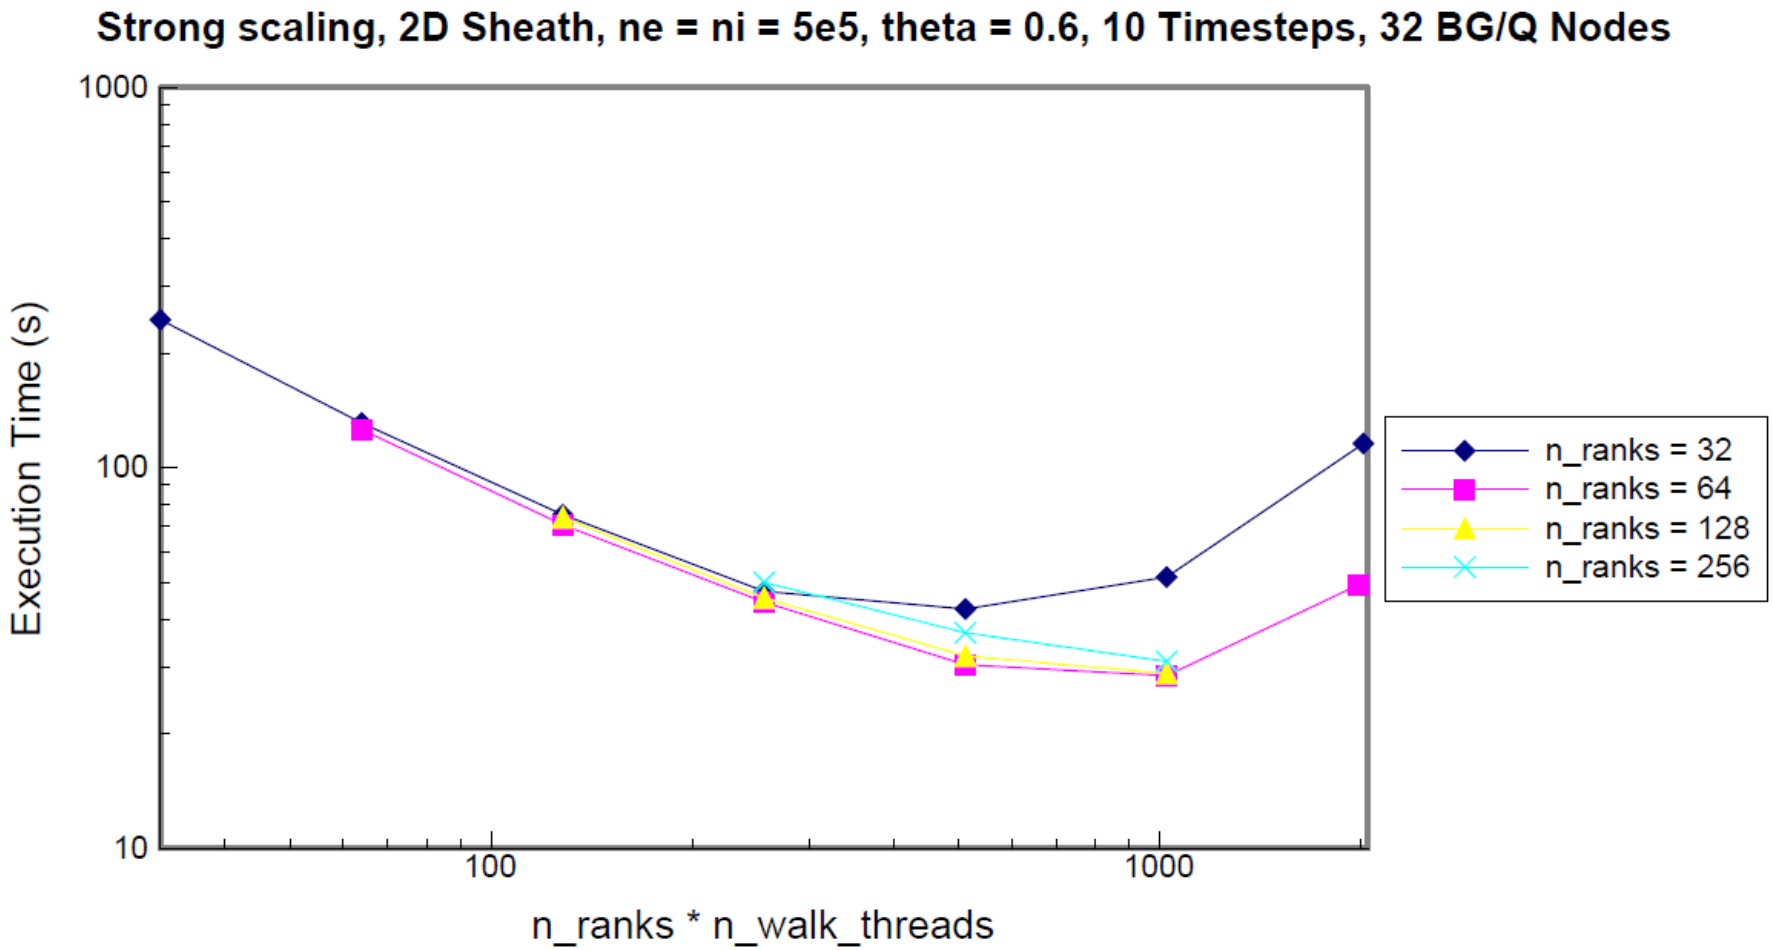

Simulation Labs: Action items 2012/13

- Sept. 2012: Open call for advanced support:
http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/2ndSupportCall

- Feb 2013: BG/Q porting workshop 2012/13

- Sept. 2013: CECAM Tutorial on Fast Coulomb Algorithms

- Engagement with Exascale projects

Exporting the SimLab Blueprint

- **JARA-HPC**:
  - Fluid & Solid Engineering
  - Ab Initio Methods in Physics & Chemistry
  - Terrestrial Systems (Geoverbund ABC/J)

- **Helmholtz Association**:
  - Climate Science/Earth & Environment + KIT
  - Neuroscience Portfolio/Bernstein Center
  - HSL: white paper on 16 potential SimLabs in HGF ⇨ PoF III

- **EU**:
  - Nuclear & Particle Physics (Cyprus Institute)