PRACE

Europe's Supercomputing Research Infrastructure

Cetraro, 22.6.2010

CAPACITIES

e-infrastructure

PRACE History and first steps

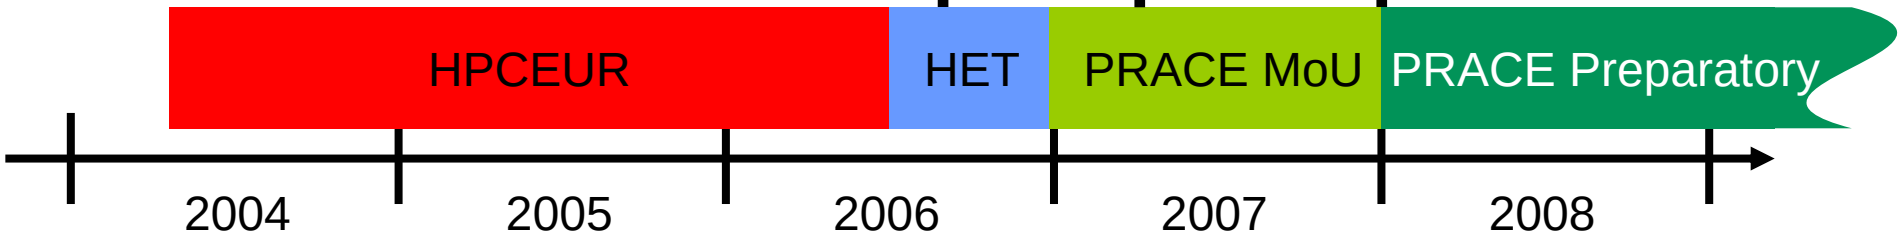

EU-Grant: INFSO-RI-211528, 10 Mio. €

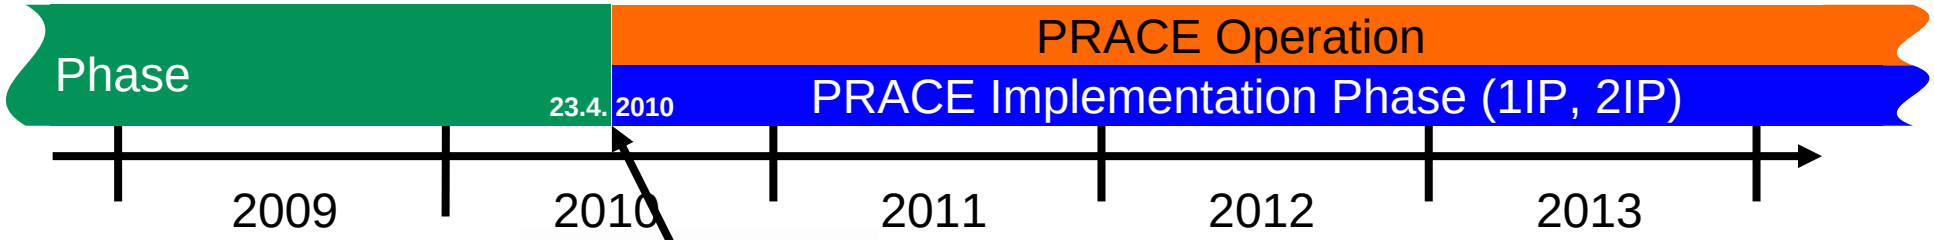

PRACE (AISBL), a legal entity with (current) seat location in Brussels

HPC-Service is item on ESFRI Roadmap

The European Roadmap for Research Infrastructures is the first comprehensive definition at the European level

Research Infrastructures are one of the crucial pillars of the European Research Area

A European HPC service – impact foreseen:

- strategic competitiveness
- attractiveness for researchers
- access based on excellence
- supporting industrial development

The ESFRI Vision for a European HPC service

- European HPC-facilities at the top of an HPC provisioning pyramid
  - Tier-0: 3-6 European Centres for Petaflop
  - Tier-0: ? European Centres for Exaflop
  - Tier-1: National Centres
  - Tier-2: Regional/University Centres

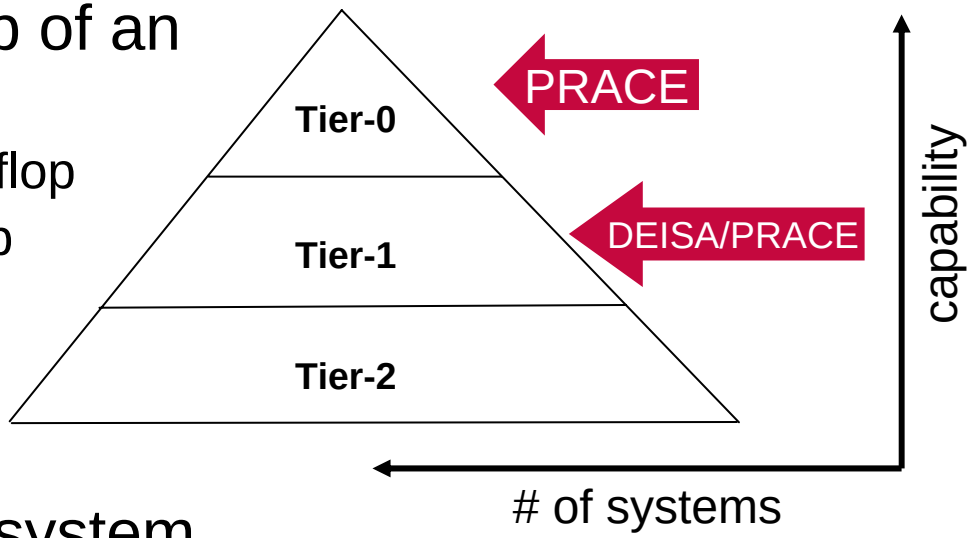

- Creation of a European HPC ecosystem involving all stakeholders
  - HPC service providers on all tiers
  - Grid Infrastructures
  - Scientific and industrial user communities
  - The European HPC hard- and software industry

PRACE – A Partnership with a Clear Objective

- **Provide world-class HPC systems for word-class science**
- **Support Europe in attaining global leadership in public and private research and development**

and a Mission to realize…

- **Create a world-leading persistent HPC infrastructure**
  - Deploy 3 – 6 systems of the highest performance level
  - Ensure a diversity of architectures meeting needs of users
  - Continue DEISA's mission on the Tier-1 level
  - Provide support and training

PRACE History and first steps

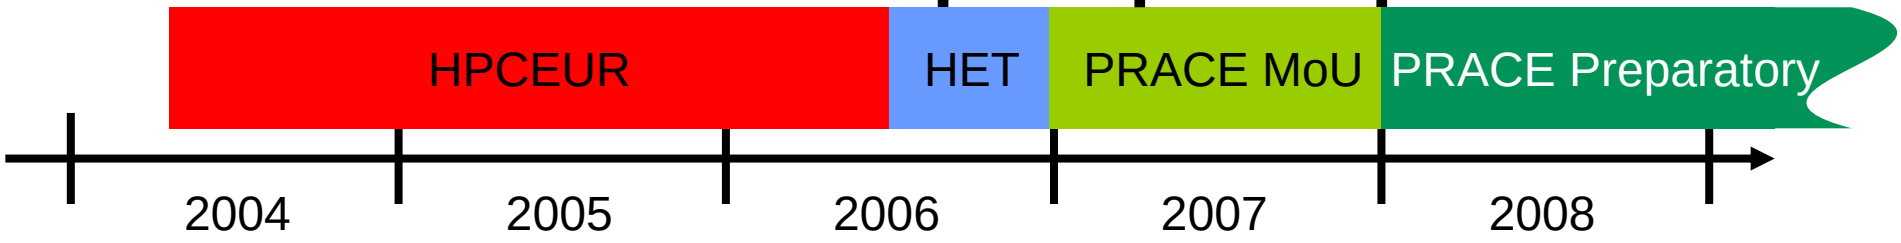

EU-Grant: INFSO-RI-211528, 10 Mio. €

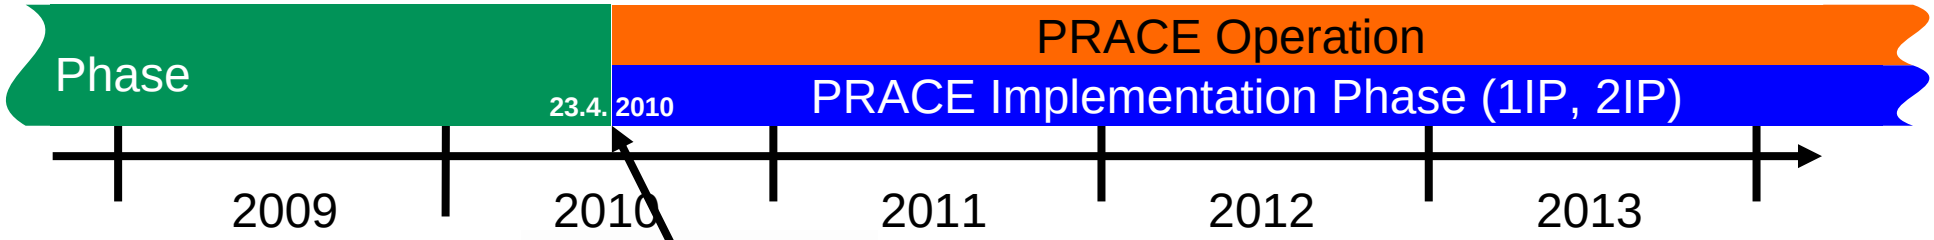

PRACE (AISBL), a legal entity with (current) seat location in Brussels

April, 23rd 2010 creation of the legal entity (AISBL) PRACE with seat location in Brussels, Belgium

16 founding members (today: 19 European members, +1 observer, +1 applicant)
400 Mio. € committments of 4 hosting members for the infrastructure

COMMITMENT

Provision of Capacity and Access

- Binding commitments by Germany, France, Italy, Spain
  - 100 Mio € over 5 years in terms of cycles
  - Contribution accounted as TCO
- Access strictly by peer review at a European level
  - Early access call: May 2010
  - Start of provision: 1.8.2010
    1. Test / evaluation access
    2. Project access – for a specific project, grant period ~ 1 year
    3. Programme access – resources managed by a community
  - Free-of-charge for European scientific communities

Provision of Capability and Support

- PRACE goes for a set of machines with successively increasing capability
  - 1 PF (2010) + 1.5 PF (2011) + 2 PF (2012) + 3 PF (2012) + 5 PF (2013)

    … and will add upgrade steps
  - Accumulated capability of more than10 PF in 2013
  - PRACE will Include Tier-1 sites
- PRACE will provide support competence over several sites
  - PRACE will create full ecosystem (Tier-0/1/2 …)

GOVERNANCE AND ORGANIZATION

Governance of the Association

- Modelled after successful examples of existing RIs
  - Council as main decision making body
  - Director with strong managing mandate
  - Scientific Steering Committee and Access Committee to give scientific advice and to steer the Peer Review process
  - Further committees will be instantiated by the Council as needed

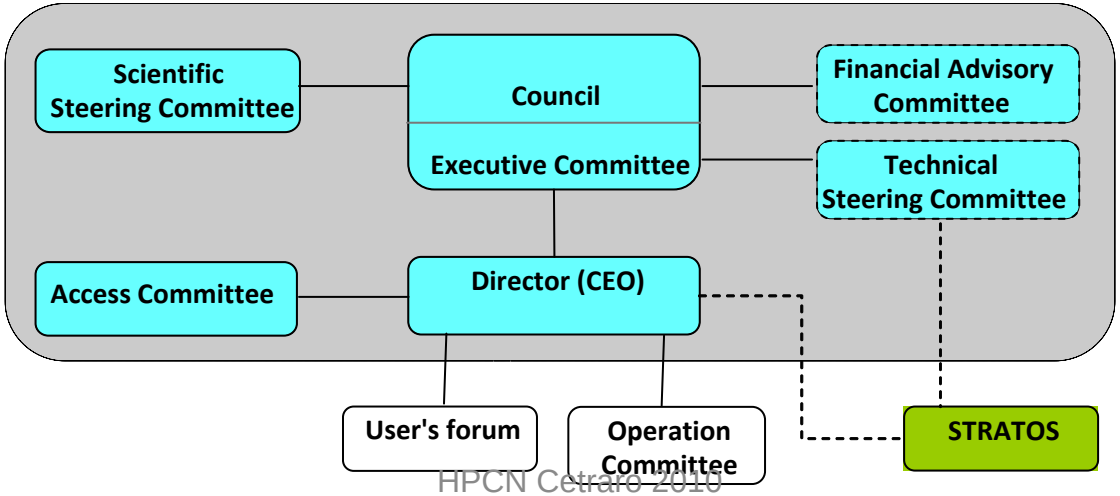

PRACE (interims) Board of Directors

General Partners:

Lennart Johnsson (Dissemination, Financial Work)

- **BSC**:

  Sergi Girona (Business Plan, Chair)

- **GENCI**:

  Jean-Philipe Nominé (Financial Work, Dissemination)

- **CINECA** :

  Sergio Bernardi (Legal Work, other tasks)

- **NCF**:

  Peter Michielse (Peer Review, Brussels-local)

Scientific Steering Committee (Article 22-24)

- Scope and basic rules defined in the Statutes of the AISBL
  - The SSC is responsible for giving opinions on all matters of a scientific and technical nature
  - Maximum of 21 members
  - Members appointed by Council based on a list of candidates prepared by the SSC
  - Two year term (renewable twice)
  - Propose the members of the Access Committee
  - Resolutions by simple majority

SSC “bootstrap” group (approved on 20.11.2009 by Management Board)

- Richard Kenway (UK, chair)
- E.J.Baerends (Netherland)
- Kurt Binder (Germany)
- Miquel Coll (Spain, biology)
- Filippo Giorgi (Italy)
- Olivier Pironneau (France)

Access Committee (Article 25-26)

- Responsible for giving opinions on the scientific use of Tier-0 Infrastructure, and providing recommendations on the allocation of Association computational resources based on the Peer Review process
  - Proposed by the SSC based on their personal experience in the areas of science
  - Appointed by the Council
  - Minimum of 5 members
  - Two years term (renewable once)
  - Half of the members shall be replaced every year
  - The Access Committee shall define its internal working rules

PRACE Project Status

- The preparatory phase project will end in June 2010 after 2 ½ years of successful work
  - Contracts for legal entity prepared
  - 400 Mio € funding for the next 5 years secured from France, Germany, Italy, Spain
  - Decisions about additional 100 Mio € from The Netherlands expected at the end of 2010
  - Architectures for Tier-0 systems identified
  - IBM BG/P in Gauss@Jülich selected as first Tier-0 system
- Contract Negotiations with the EC about an Implementation Phase Project to start in July 2010 are finished
  - Budget: 28.5 Mio € (20 Mio € EC funding)

Preparation Phase Project: Example of Activity

EVALUATION OF PETASCALE SYSTEMS

PRACE Preparatory Phase

- WP1 Management
- WP2 Organizational concept → Statutes
- WP3 Dissemination, outreach and training
- WP4 Distributed computing
- WP5 Deployment of prototype systems
- WP6 Software enabling for prototype systems
- WP7 Petaflop/s systems for 2009/2010
- WP8 Future Peta to Exaflop/s technologies

Categorisation of Applications (2009)

- Surveys of PRACE partners' HPC systems and major applications
- 24 systems and 69 applications
- Quantitative basis for selecting representative applications

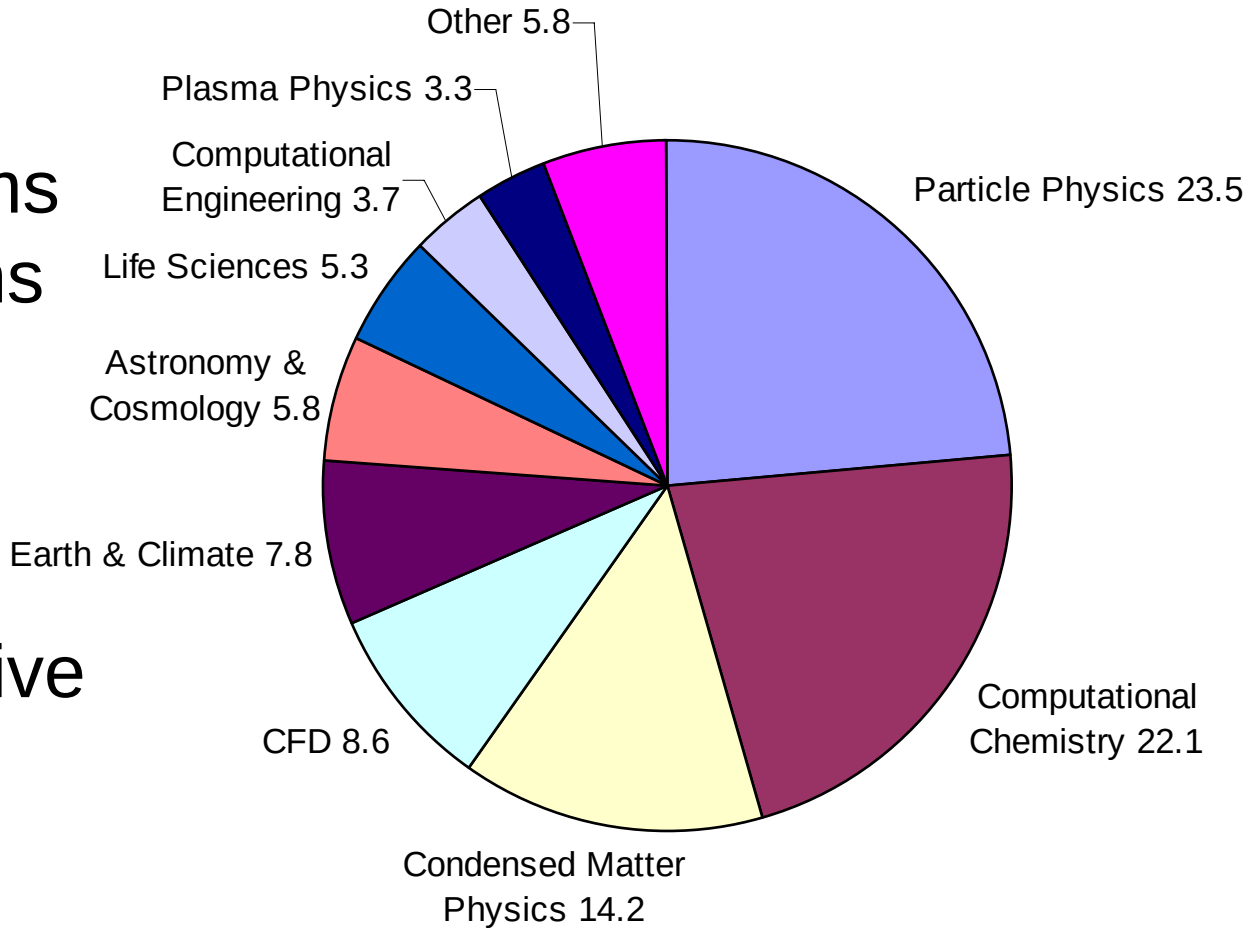

→ PRACE Benchmark Suite

- 12 core applications, plus 8 additional applications
  - NAMD, VASP, QCD, CPMD, GADGET, Code Saturne, TORB, ECHAM5, NEMO, CP2K, GROMACS, N3D
  - *Additional:* AVBP, HELIUM, TRIPOLI_4, PEPC, GPAW, ALYA, SIESTA, BSIT
  - Synthetic benchmarks for architecture evaluation
- Integrated into JuBE (**Ju**elich **B**enchmark **E**nvironment)

Prototypes for Petaflop/s systems

IBM BlueGene/P (FZJ)
01-2008 / 06-2009

IBM Power6 (SARA)
07-2008

Cray XT5 (CSC)
11-2008

IBM Cell/Power (BSC)
12-2008

NEC SX9, vector part (HLRS)
02-2009

Intel Nehalem/Xeon (CEA/FZJ)
06-2009

Mapping Applications to Architectures

Code	**MPP** (i.e. BlueGene L/P or CRAY XT4/5)	**Thin node clusters** (i.e. Bull INCA or SGI ICE)	**Fat node clusters** (i.e. Bull MESCA, SGI UltraViolet or IBM Power6)	**Vector systems** (NEC SX8-9, Cray X2	**Accelerated systems** (i.e. scalar or vector + GPU, FPGA or Clearspeed).	**Accelerated systems – Cell based** (i.e. Roadrunner, Maricell)
NAMD				E	E	E
CPMD						E
VASP						E
QCD				E	E	E
GADGET						
Code Saturne				E	E	E
TORB						E
NEMO						
ECHAM5						
CP2K	E					
GROMACS					E	E
N3D		E	E			
AVBP	E					
HELIUM						
TRIPOLI 4	E		E			
GPAW						
ALYA						E
SIESTA						E
BSIT						E
PEPC				E	E	E

Table 4 : application mapping to Petaflop/s systems archite

JUGENE – FIRST PRACE SYSTEM

JUGENE@Jülich: #5 worldwide, #1 in Europe

IBM Blue Gene/P

1st PRACE system

Jülich Blue Gene/P Configuration

- **72 Racks Blue Gene/P**
  - 73728 Compute Nodes (4 processor cores, 2 GB memory)
  - 294912 CPU cores, 144 TB memory
  - 1 PFlop/s peak performance
  - 825.5 TFlop/s Linpack
  - 600 I/O nodes (10GigE) → >60 GB/s I/O
  - 2.2 MW power consumption
- **2 Frontend Nodes + 2 Service Nodes**
  - IBM p6 550, 8 x Power6 (4.2 GHz), 128 GB memory
  - local storage device DS4700 (28 TB)

ACCESS

Scheme of Calls for JUGENE

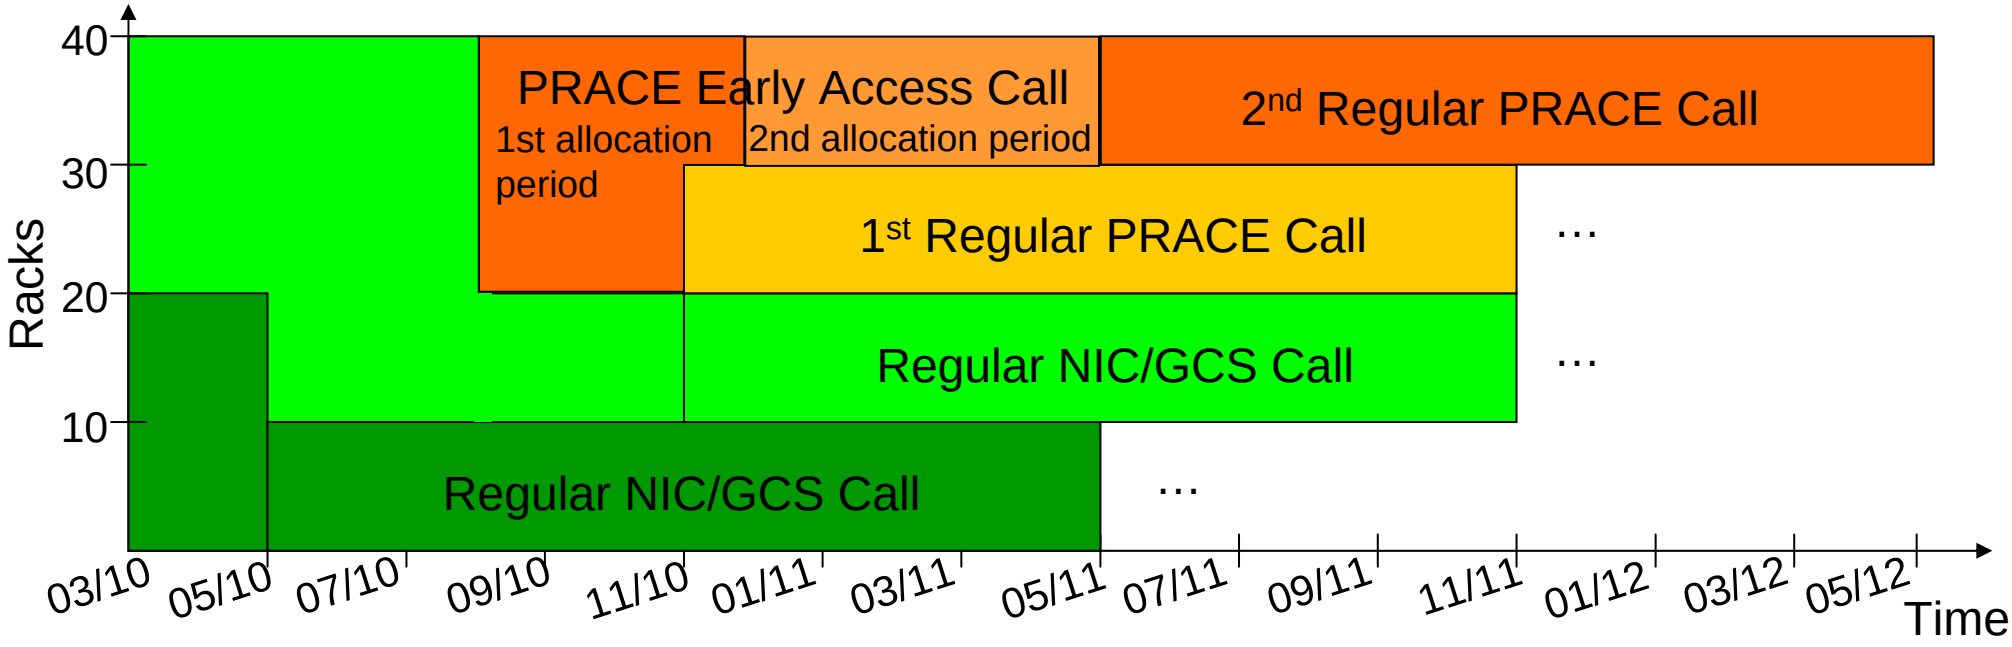

35% of Capacity provided by PRACE corresponding to 40 Racks
Early Access Call assisted by the SSC Establishment Group

Proposal Statistics

- 65 proposals
  - 47 for project access
  - 11 for the 1st of December
  - 36 for the 1st of August
  - 18 for combined access
- Project leaders from 16 countries
  - Belgium – 1
  - Bulgaria - 2
  - Denmark – 1
  - Finland - 2
  - France – 5
  - Germany - 11
  - Greece – 2
  - Ireland - 2
  - Italy – 4
  - Netherlands – 2
  - Poland – 2
  - Portugal – 4
  - Spain – 11
  - Sweden – 1
  - Switzerland – 3
  - UK - 12

Proposal Statistics II

- 1,471,509,270 core hours
- Oversubscription: factor 5
  - Astrophysics - 7
  - Chemistry and Materials - 11
  - Earth Sciences and Environment - 4
  - Engineering and Energy - 15
  - Fundamental Physics - 17
  - Mathematics and Computing - 3
  - Medicine and Life Sciences - 8

Ana Bela Dias (NCF) and Susie Douglas (EPSRC)

PEER REVIEW IMPLEMENTATION

PRINCIPLES

- Transparency
- Expert assessment
- Confidentiality
- Right to reply
  - reviewers comments and final decision appeal
- Prioritisation
- Managing interests
- No parallel assessment
- Ensure fairness to the science proposed

Types

- Preparatory access
  - only technical peer review
- Project
  - both technical and scientific peer review
- Programme
  - both technical and scientific peer review

Peer review process

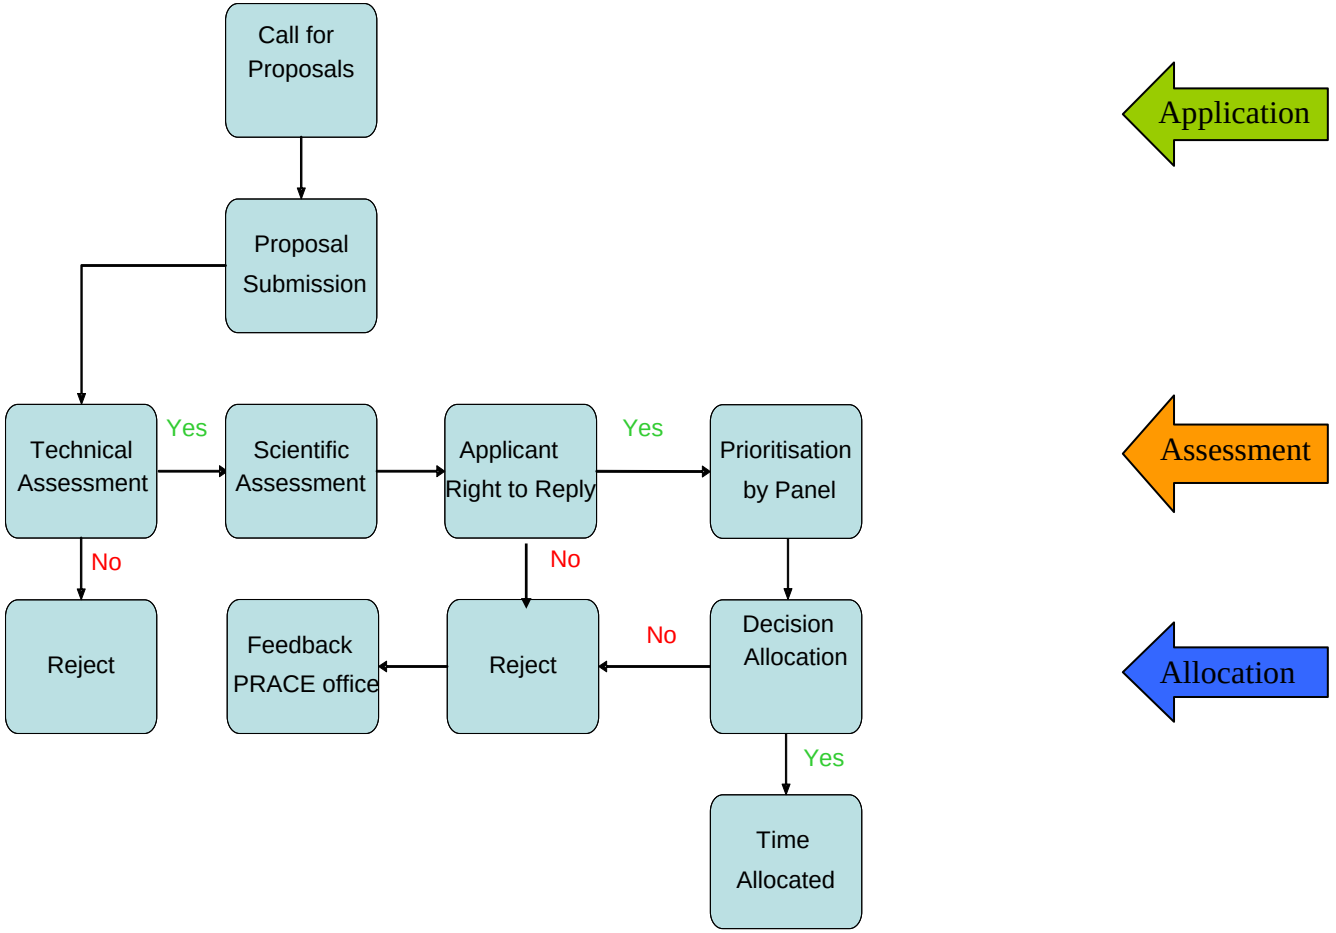

Project and programme proposals:

- Technical peer review (system and code suitability) by hosting centre representatives
- Scientific peer review by 3 external reviewers
- Applicants have the right to comment on the reviewers remarks; these remarks are sent together with the reviewers comments to the (prioritisation) panel
- Formally the Council ratifies the prioritisation list
- The Director informs (in writing) the applicants of the computing grants, on behalf of the Council
- Applicants have the right to appeal to a decision of the Council; the Council will have a formal procedure to handle such appeals

(MACHINE) DEVELOPMENT

Machine Ramp-Up Until 2012

Country		2010	2011	2012
Germany /GCS		FZJ 1 PF peak	HLRS	LRZ
France			GENCI	
Italy			CIN	ECA
Spain				BSC

PRACE Preparatory Phase

- WP1 Management
- WP2 Organizational concept → Statutes
- WP3 Dissemination, outreach and training
- WP4 Distributed computing
- WP5 Deployment of prototype systems
- WP6 Software enabling for prototype systems
- WP7 Petaflop/s systems for 2009/2010
- **WP8 Future Peta to Exaflop/s technologies**

- STRATOS: Evaluation and Development of Next-Generation Technology Technology
- 1 IP First PRACE Implementation Project
- EU: ICT call 7 “Objective ICT-2011.9.13 Exascale computing
- 2 IP Second PRACE Implementation Project
- 3 IP?