Bridging Global Computing with Grid (BIGG) Programming model section

Marco Danelutto
Dept. Computer Science - Univ. Pisa &
CoreGRID Programming model Institute

Sophia Antipolis, November 29, 2006

GRID definition

A fully distributed, dynamically reconfigurable, scalable and autonomous infrastructure to provide location independent, pervasive, reliable, secure and efficient access to a coordinated set of services encapsulating and virtualising resources (computing power, storage, instruments, data, etc.) in order to generate knowledge.

Programmer perspective (today)

- *collection of heterogeneous resources, dynamically available in time, subject to faults & misfunctioning, searched, recruited, secured, deployed and used to perform complex tasks*

- concurrent activity set up, mapping and scheduling; resource management (recruiting, monitoring, repairing); communication & synchronization management; fault tolerance and security management; heterogeneity management (cross compiling, architectural neutral data formats), ...

Programmer perspective (today) (2)

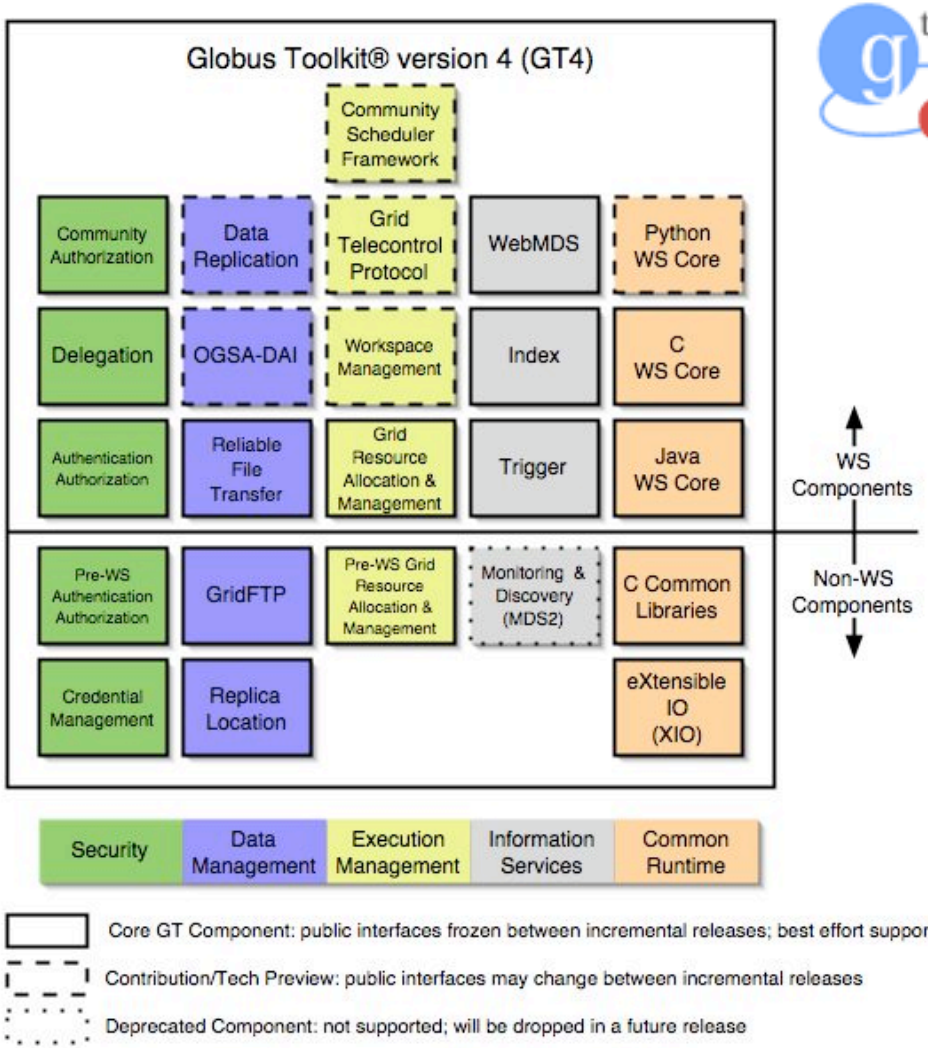

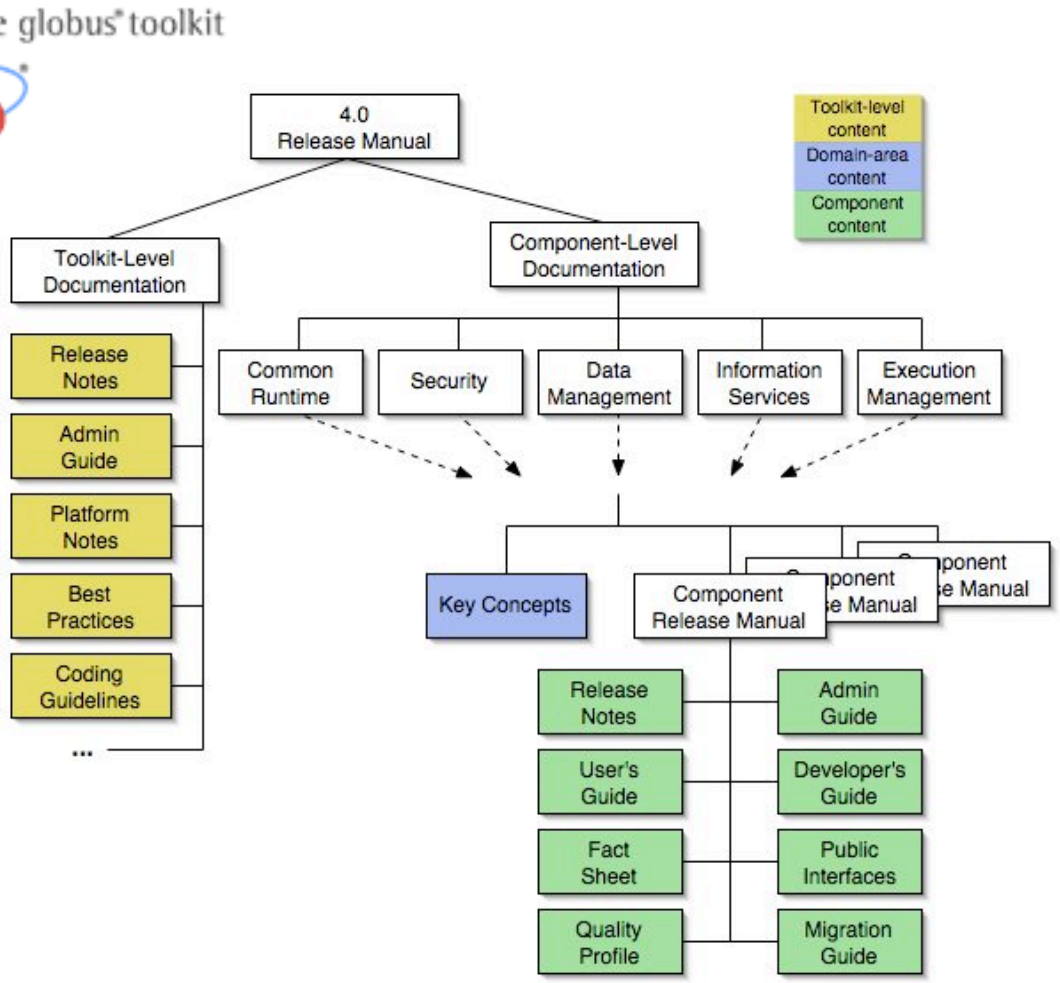

Programmer perspective (today) (3)

Functional core

Support code

NGG recommendations

In the first half of 2003 a group of high level experts was convened by the European Commission, Unit INFSO/F2, in order to produce a report titled "Next Generation Grids, European Grid Research 2005-2010". In this report, known as the NGG report, the experts have pioneered the vision of the 'Invisible Grid', i.e. whereby the complexity of the Grid is fully hidden to users and developers through the complete virtualisation of resources, and have sketched the research priorities underpinning the realisation of the Next Generation Grids.

6.4 Raising the Level of Abstraction

Substantial research efforts, greater than in the past, have to be invested to raise the level of abstraction of future generation grid systems at all the levels. In particular, this is necessary to raise the level of abstraction in such a way that the users/programmers are provided with higher level programming models and tools, as well as with better management abstractions. Such programming models, tools and abstractions must actually be able to relieve the programmers from most of (possibly all) the burden involved in the direct management of the specific, demanding and error prone grid related issues. These research efforts should be specifically finalised

How to implement NGG recommendation ?

- Exploit *layered* design of grid middleware + programming environments

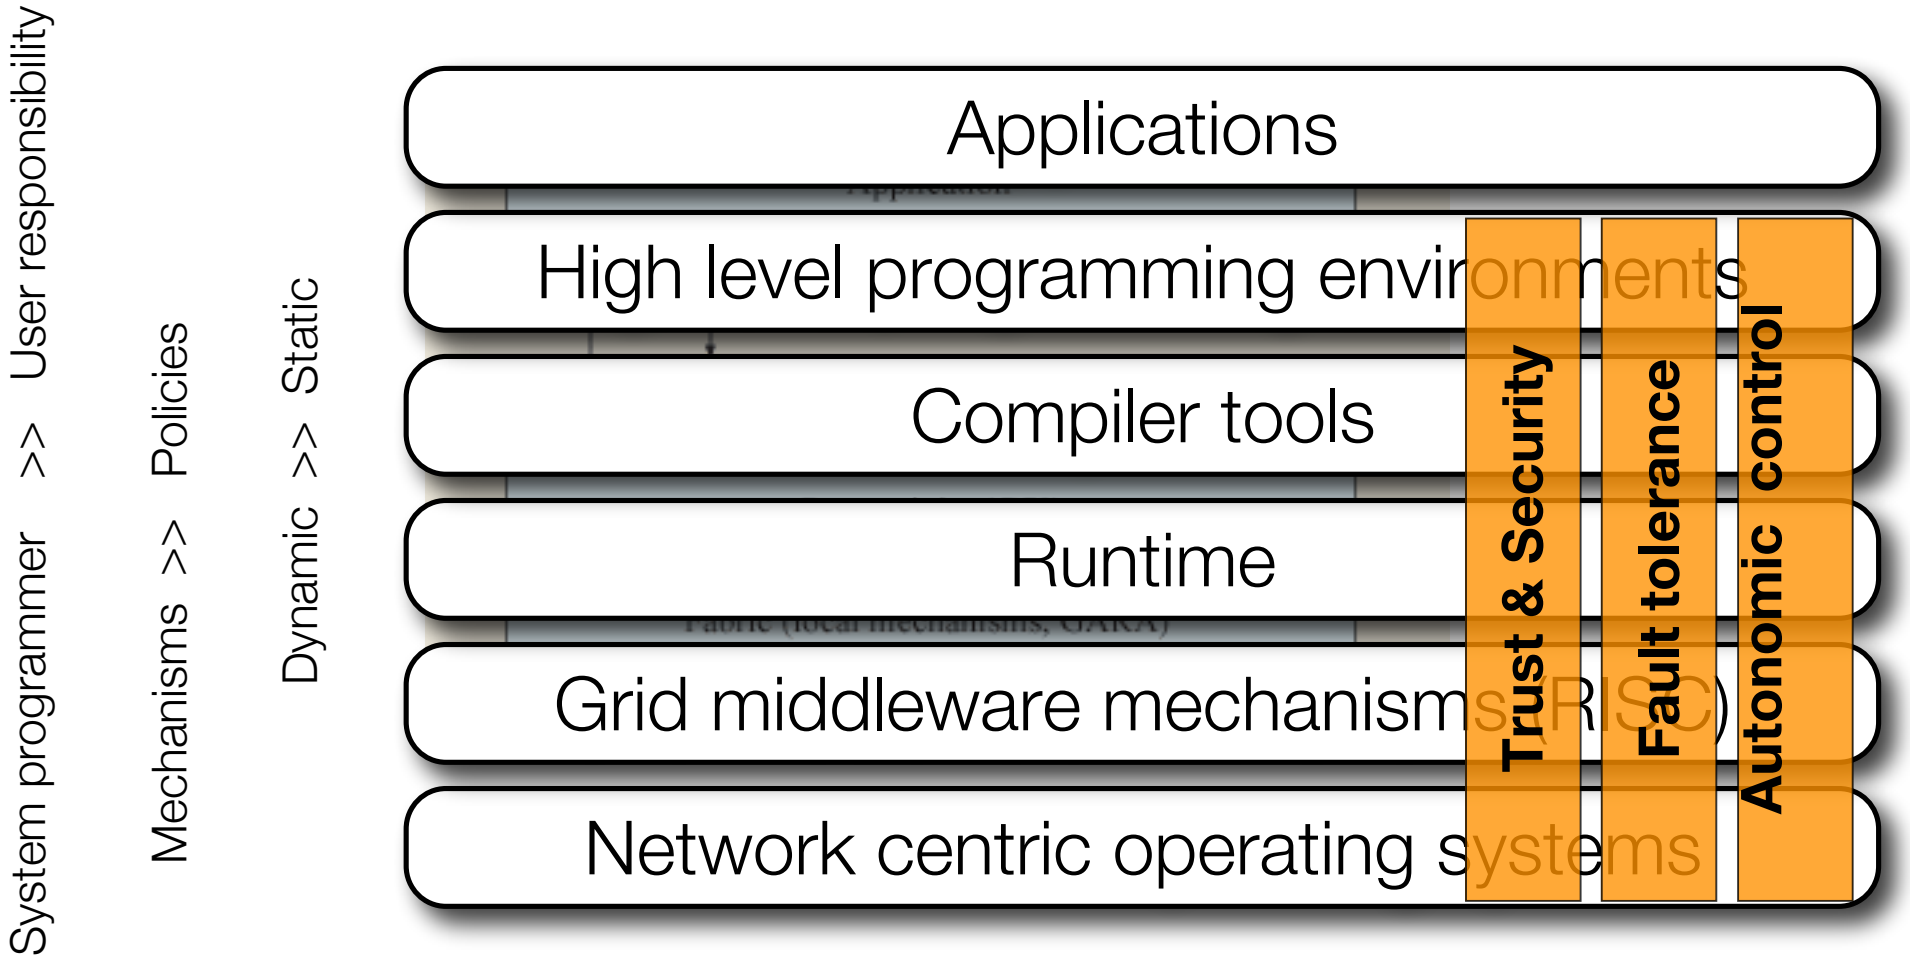

Plus ...

- Advanced programming models
  - providing predefined, customizable, efficient programming skeletons - design/coordinations patterns
  - move away from assembly language
  - implemented by compilers or RTS/libraries (exploiting layer hierarchy)
  - relieve programmers of unnecessary burden,
    - possibly free them !

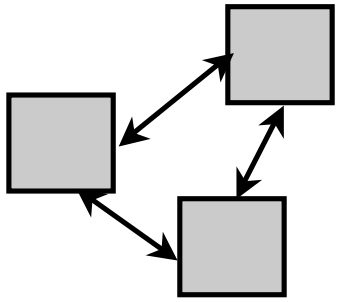

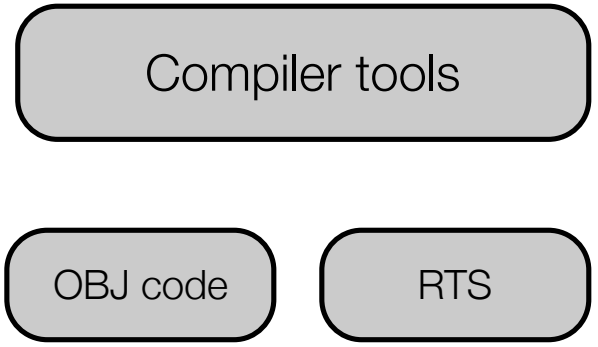

Is it possible ?

- ASSIST high level, structured, parallel programming environment
- compiler + run time + grid middleware
- heterogeneous node handling + adaptivity (performance contracts)
- up & running since 2001 (advanced features in 2005)

Exp 2: Data-Parallel(STP)

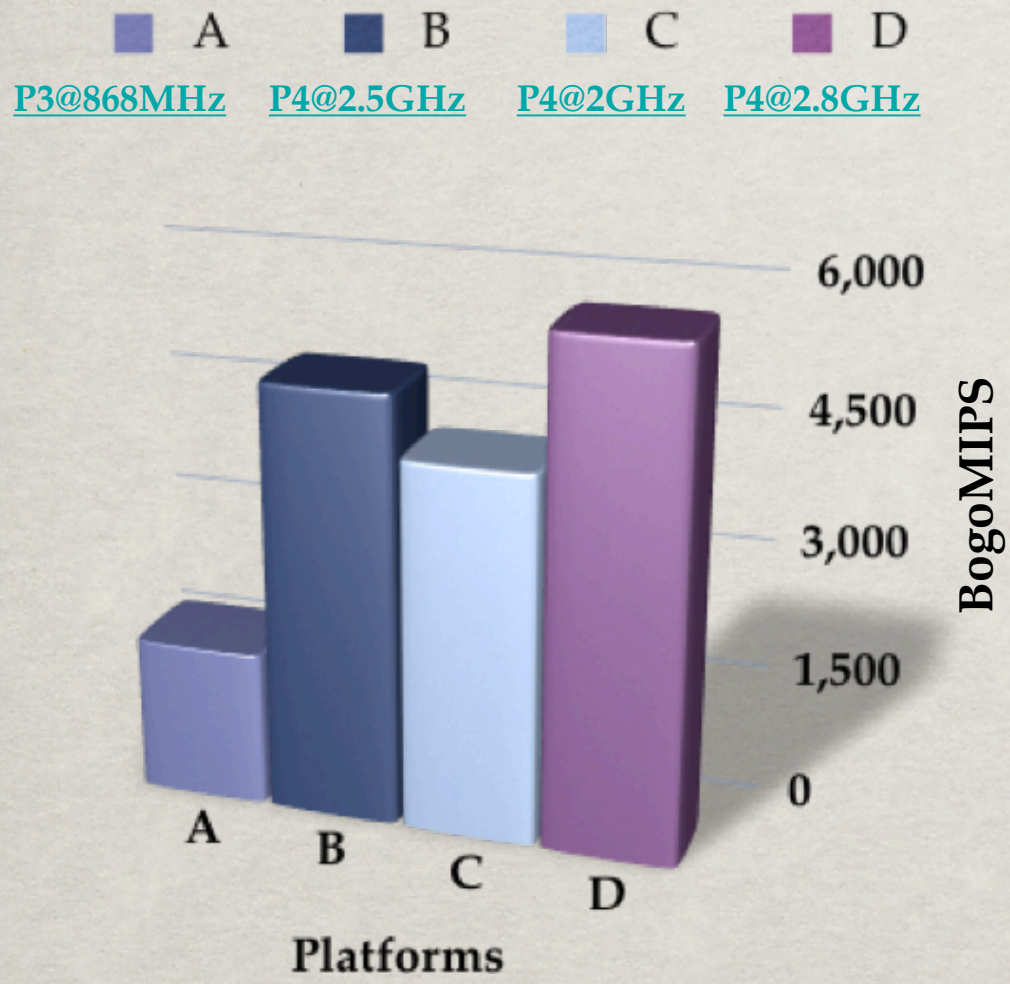

Expected work balance among platforms

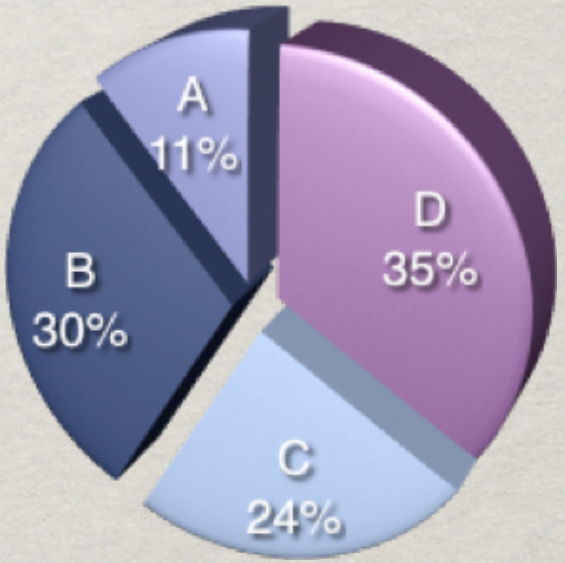

Exp 2: Data-Parallel(STP)

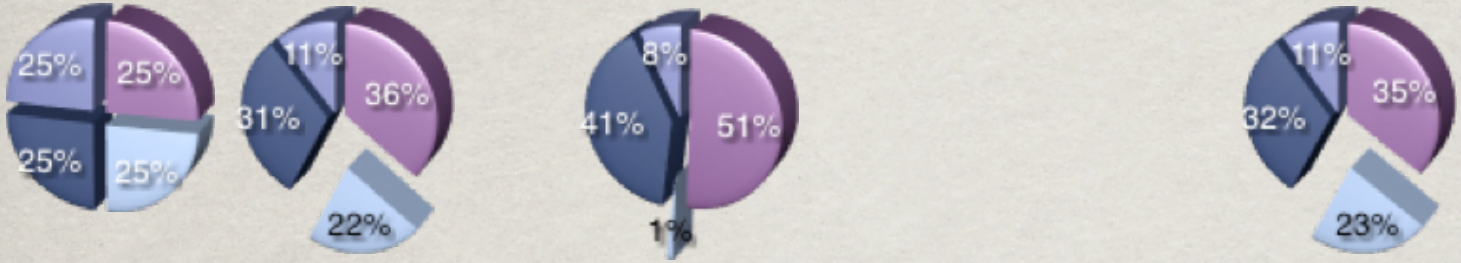

Distribution of load among platforms (n. of VPs)

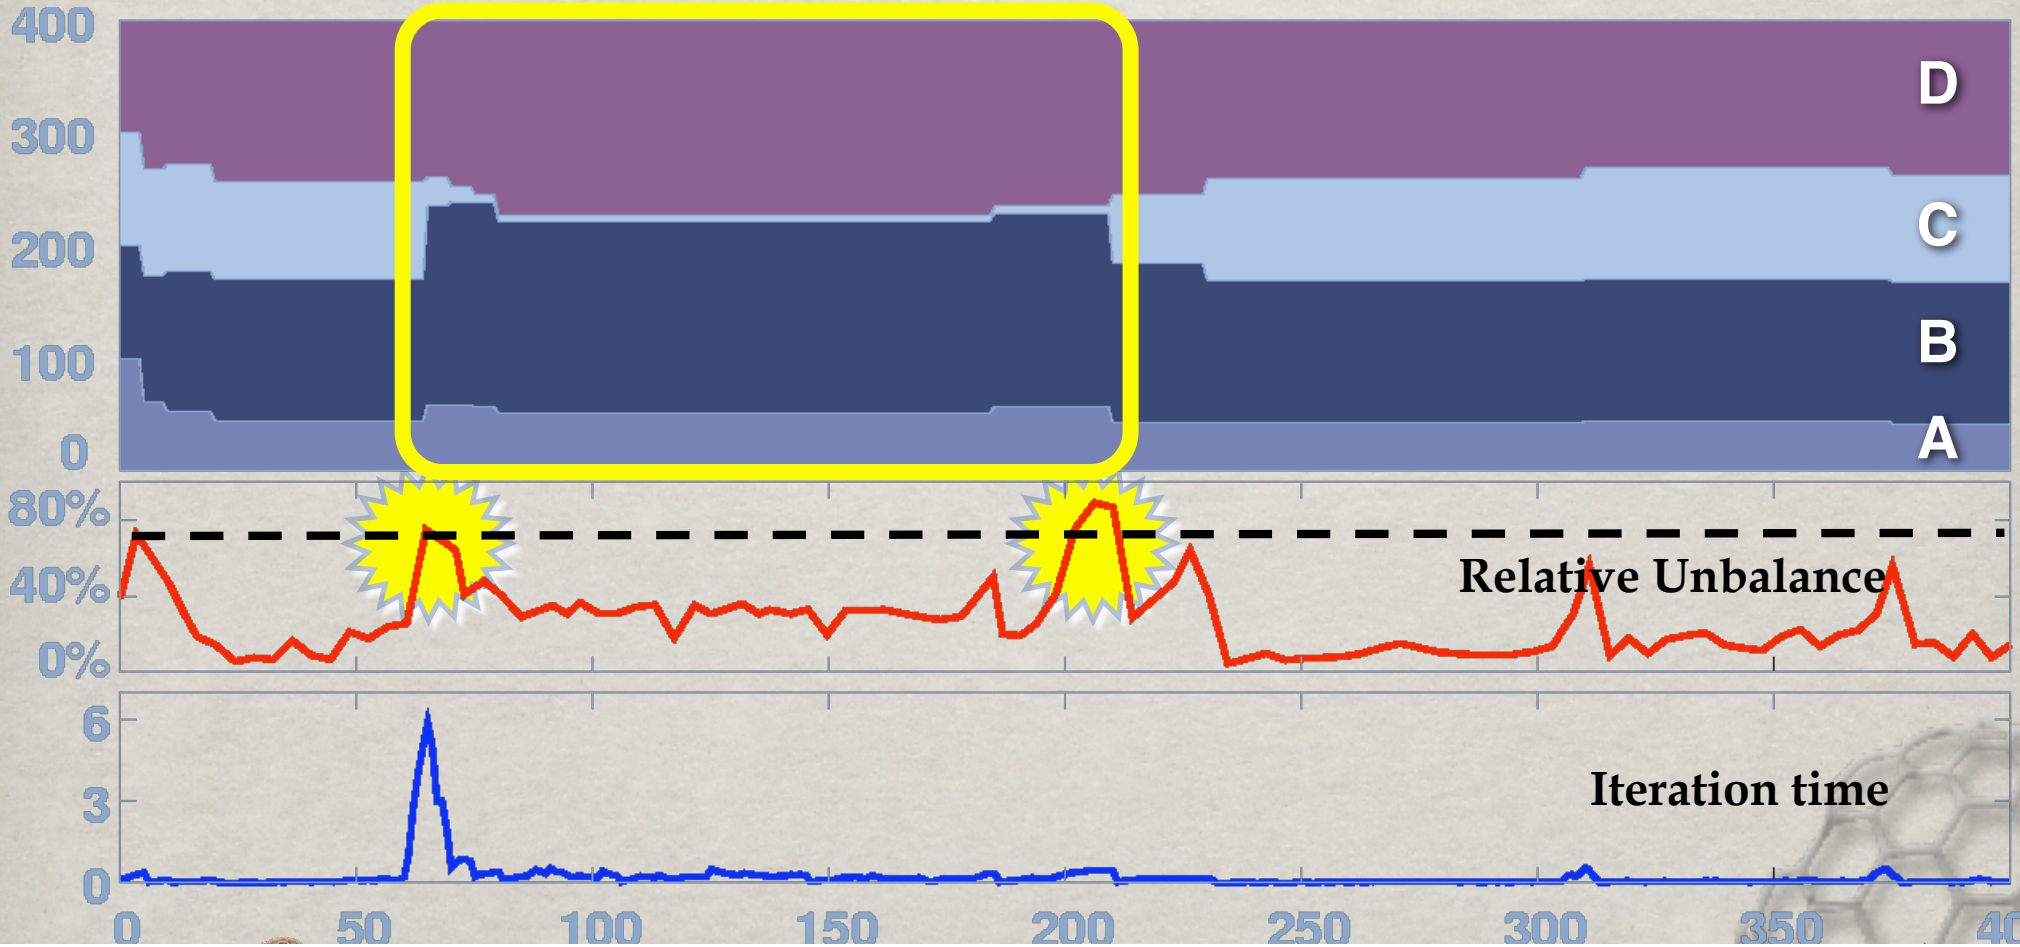

The Programming model Institute perspective

- GCM: the Grid Component Model

  - hierarchical composition
  - RPC + data/stream + event ports + structured interaction patterns
  - component autonomic control (non functional features, including performance, and grid management)
  - reference implementation in GridCOMP Effective Components for the Grids
  - interoperable (WS),
    advanced programming models on top of GCM (skeletons)

The GCM features: collective communications

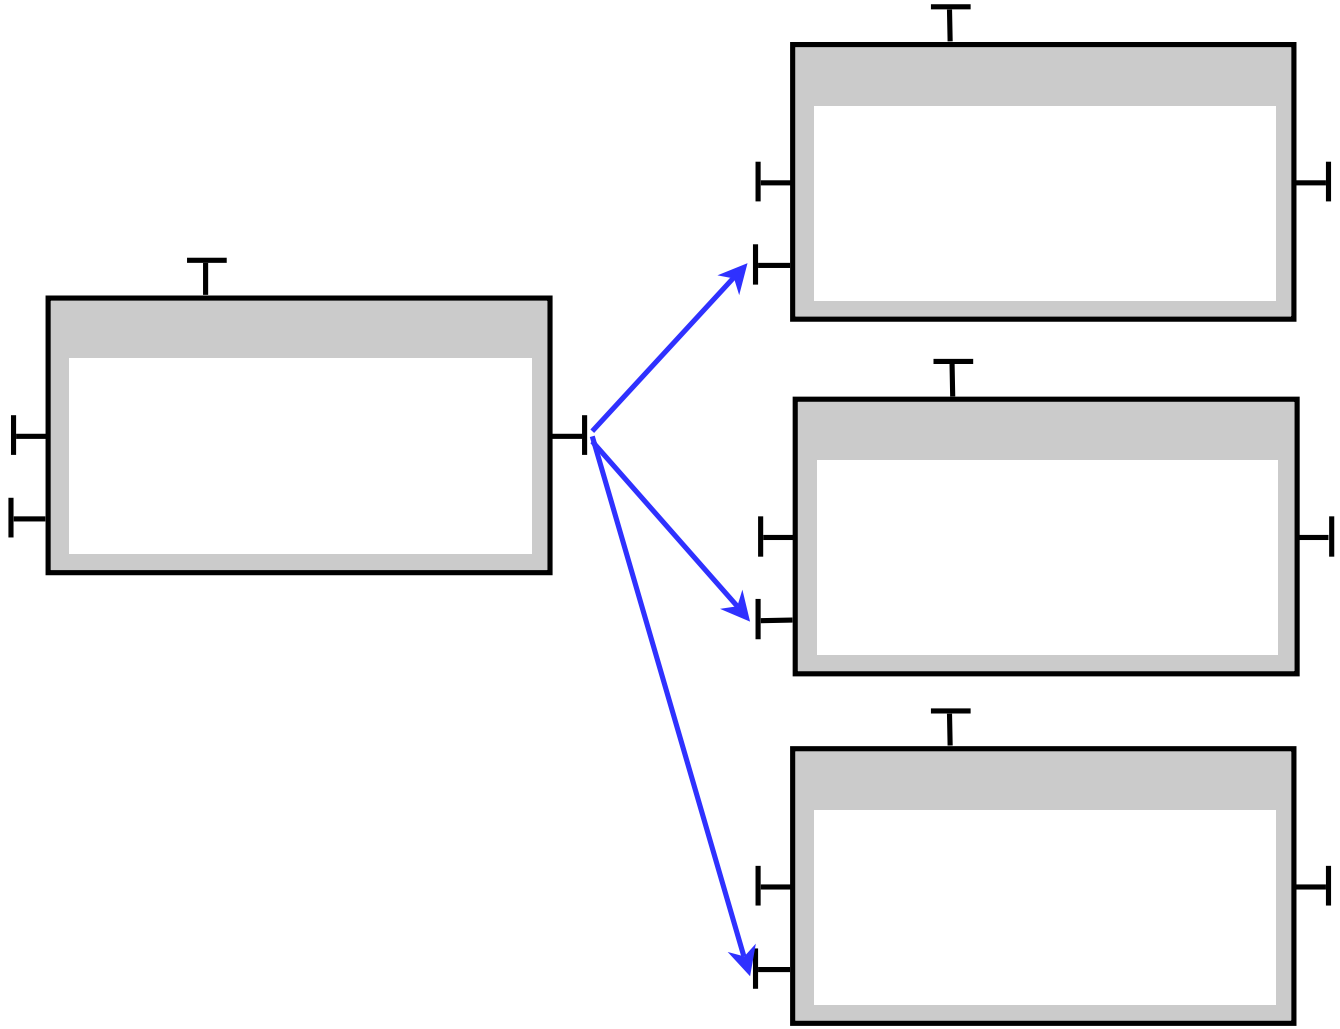

The GCM features: hierarchical composition

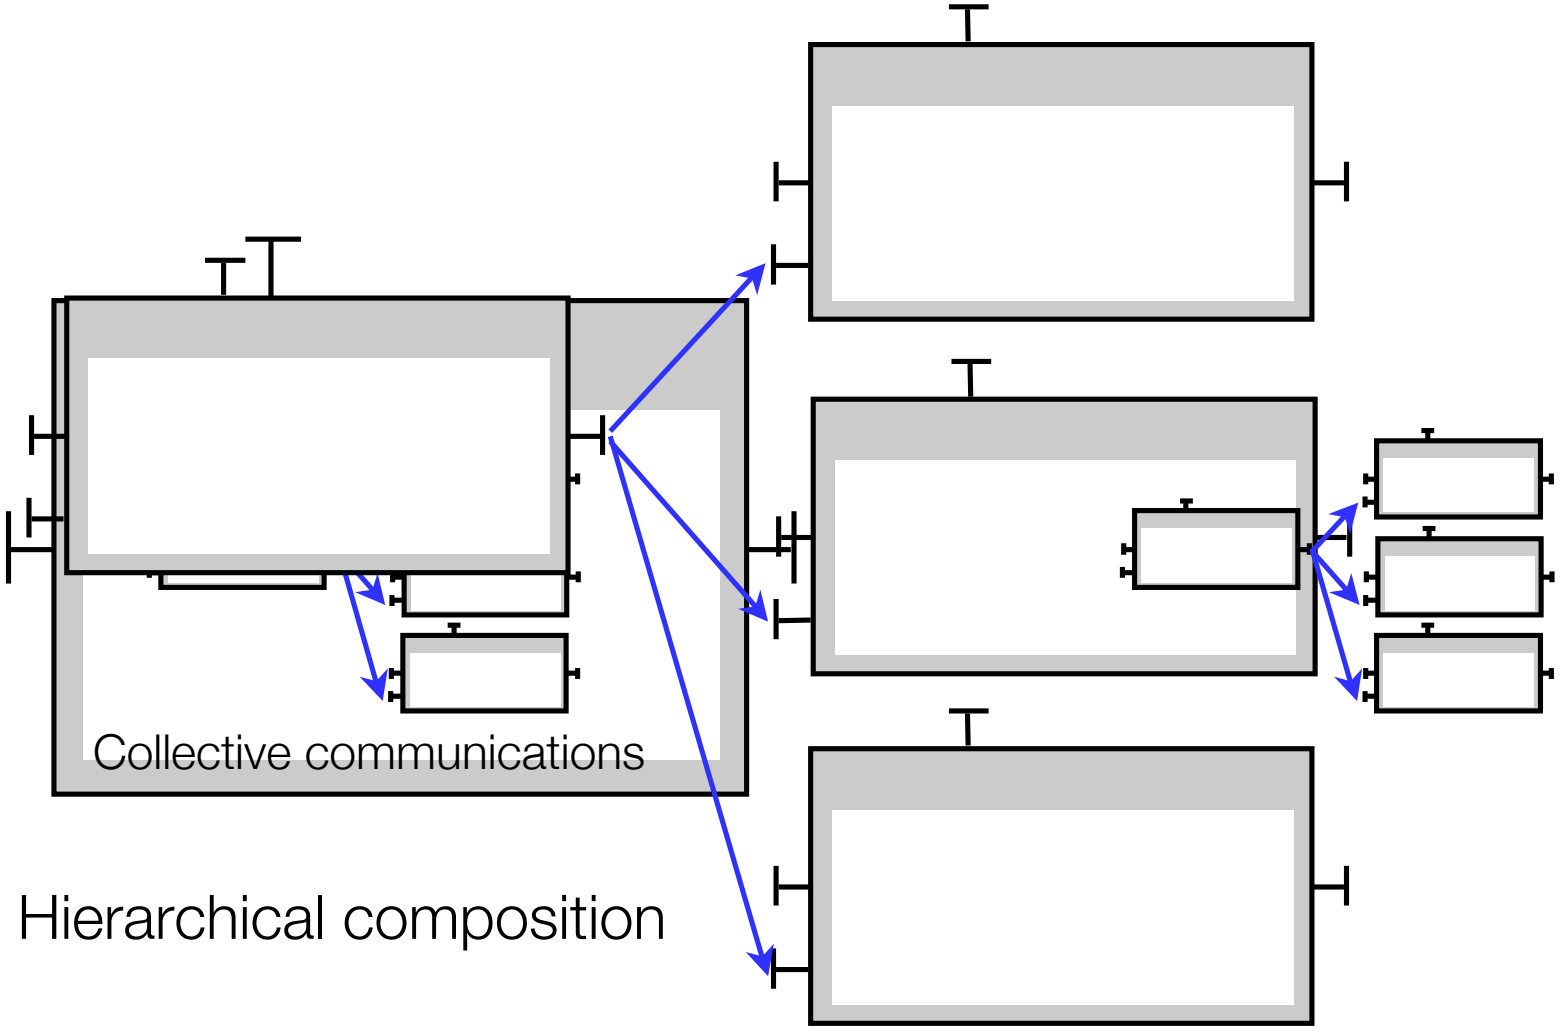

The GCM features: autonomic management

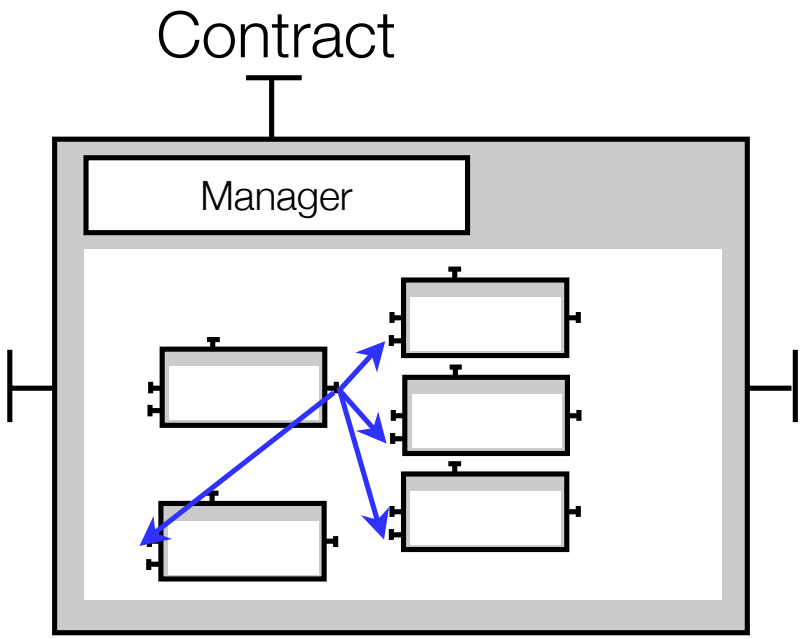

The GCM features: portability

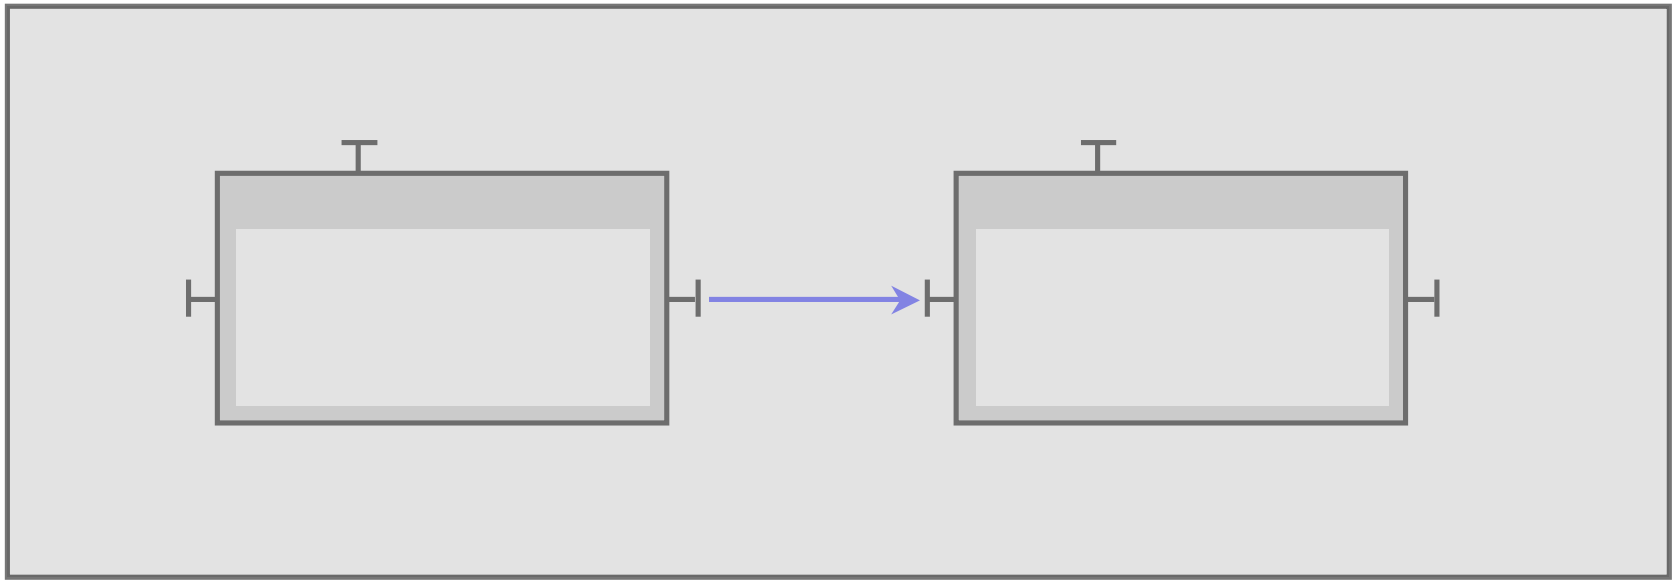

The GCM features: interoperability

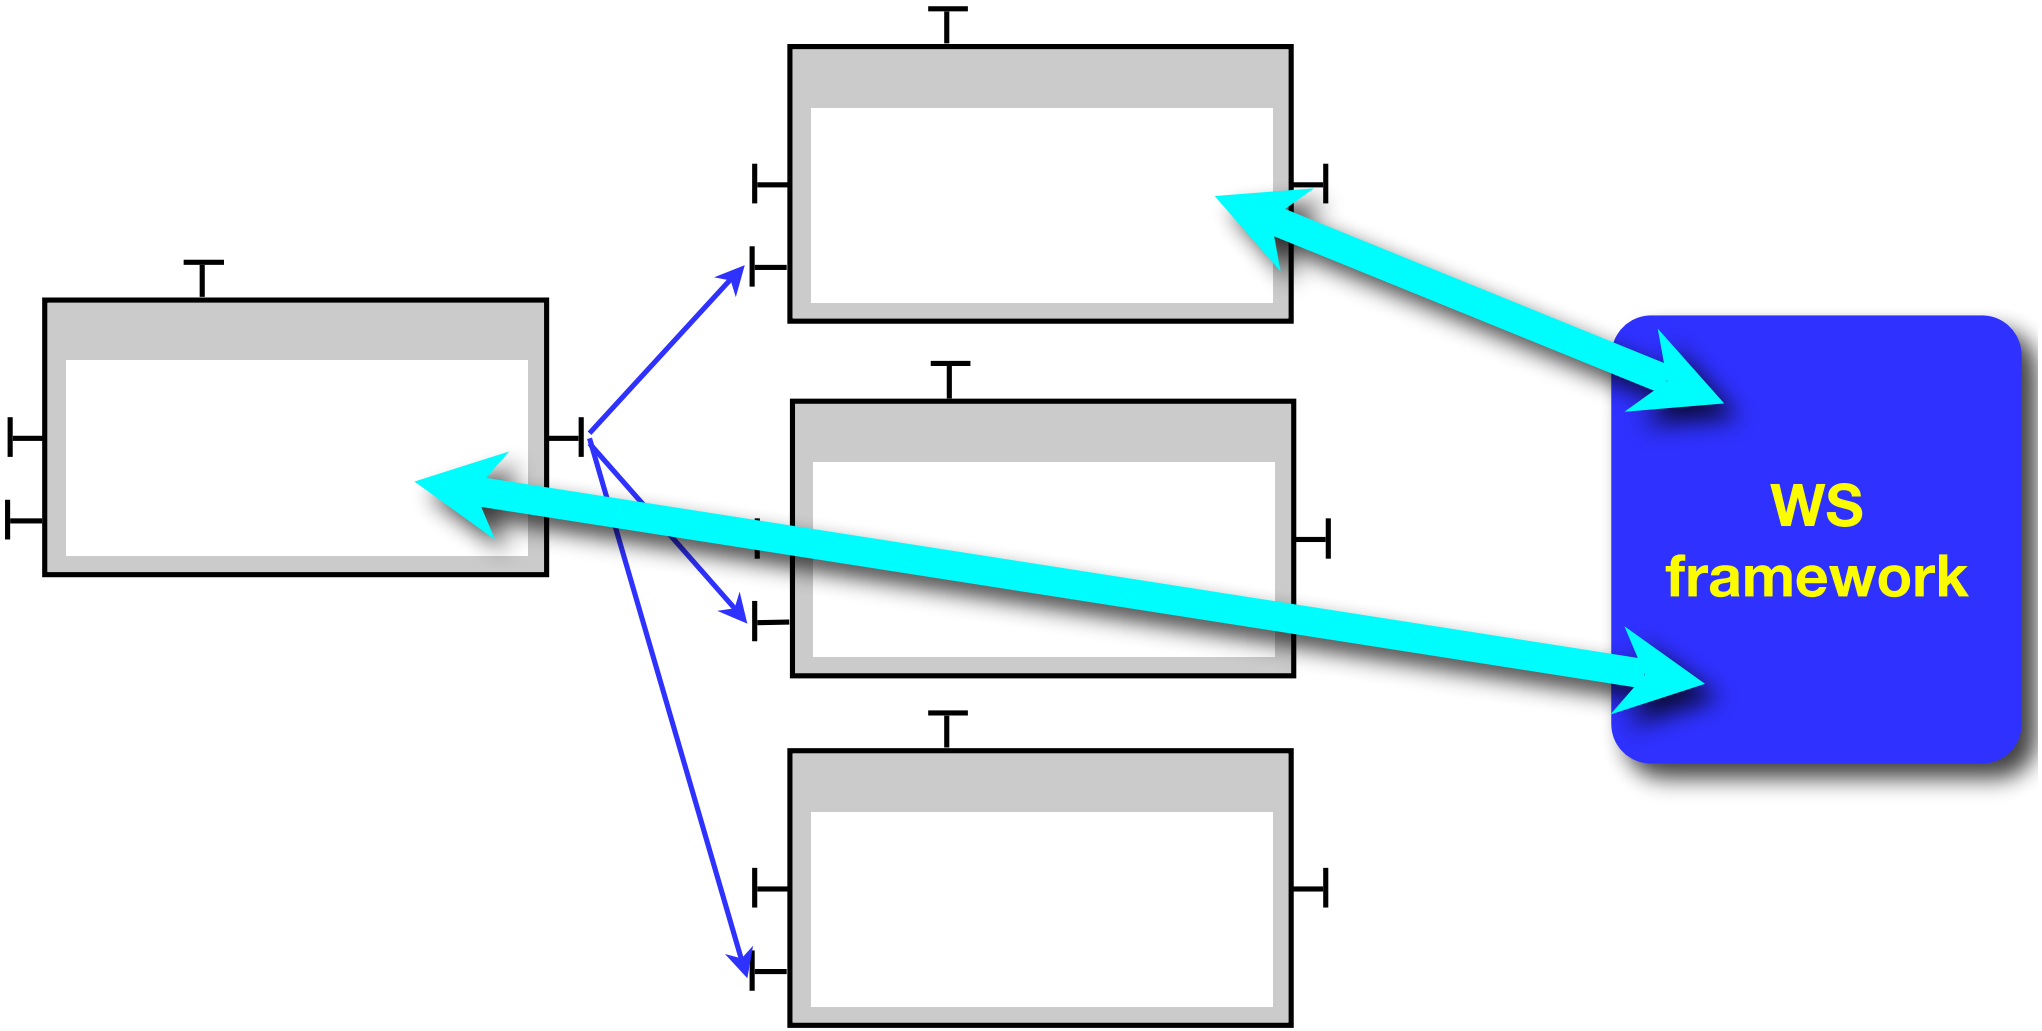

The challenges: “invisible” vs “aware” models

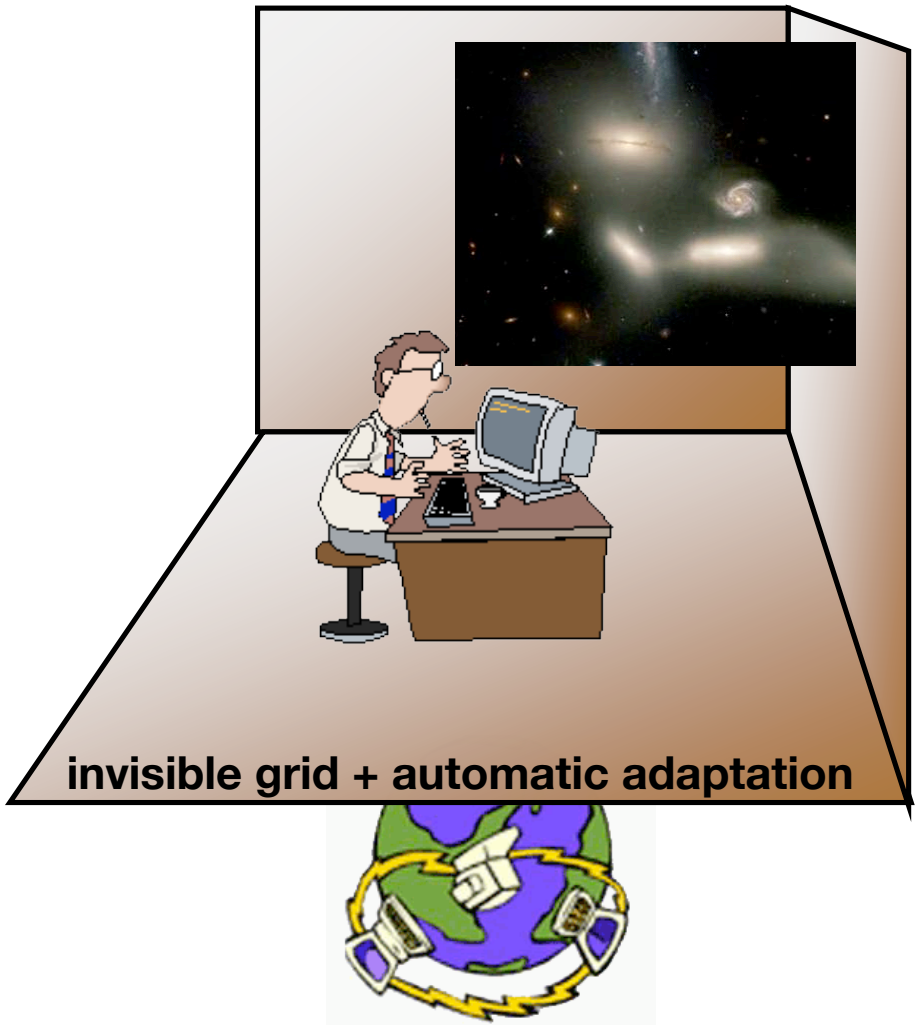

- Actually
  - Vision and awareness of grid peculiarities and problems (system programmer level)
  - Invisible grid á la power grid (application programmer level)

The challenges: bunch of services vs. structured

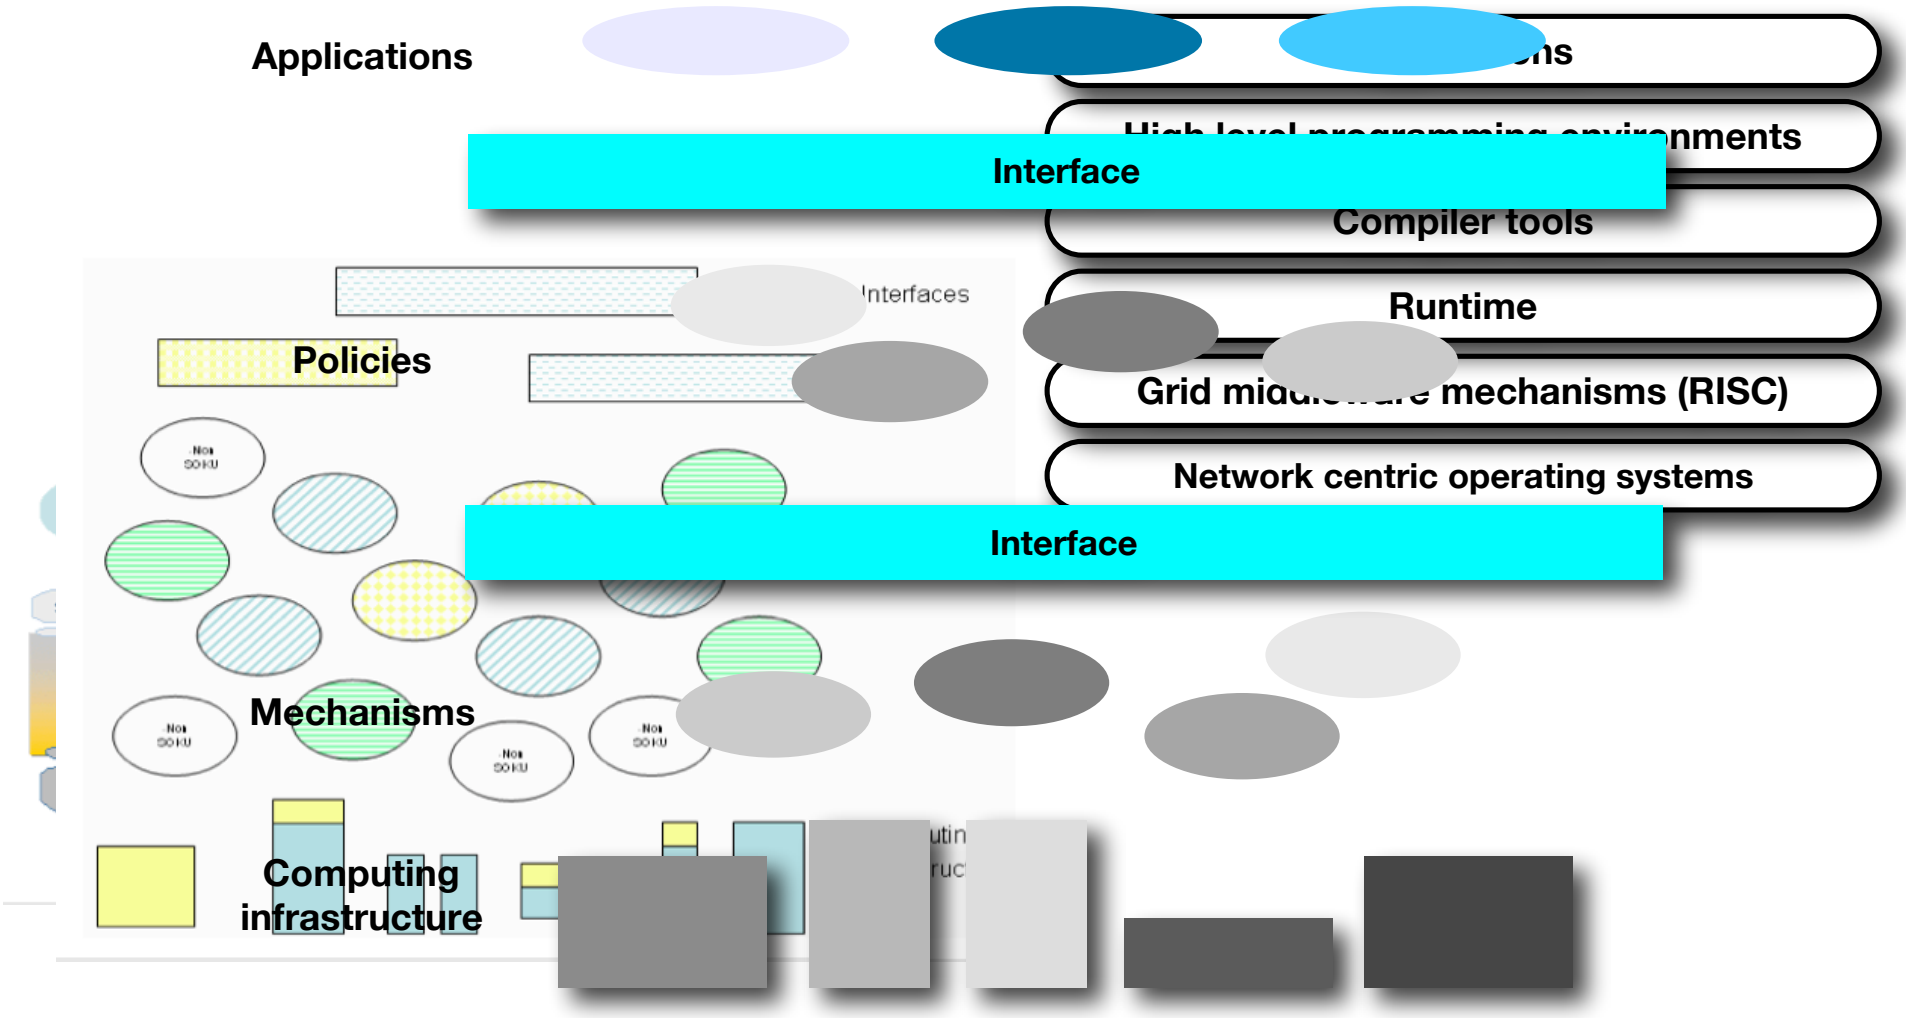

The challenges: AOP vs. everywhere concerns

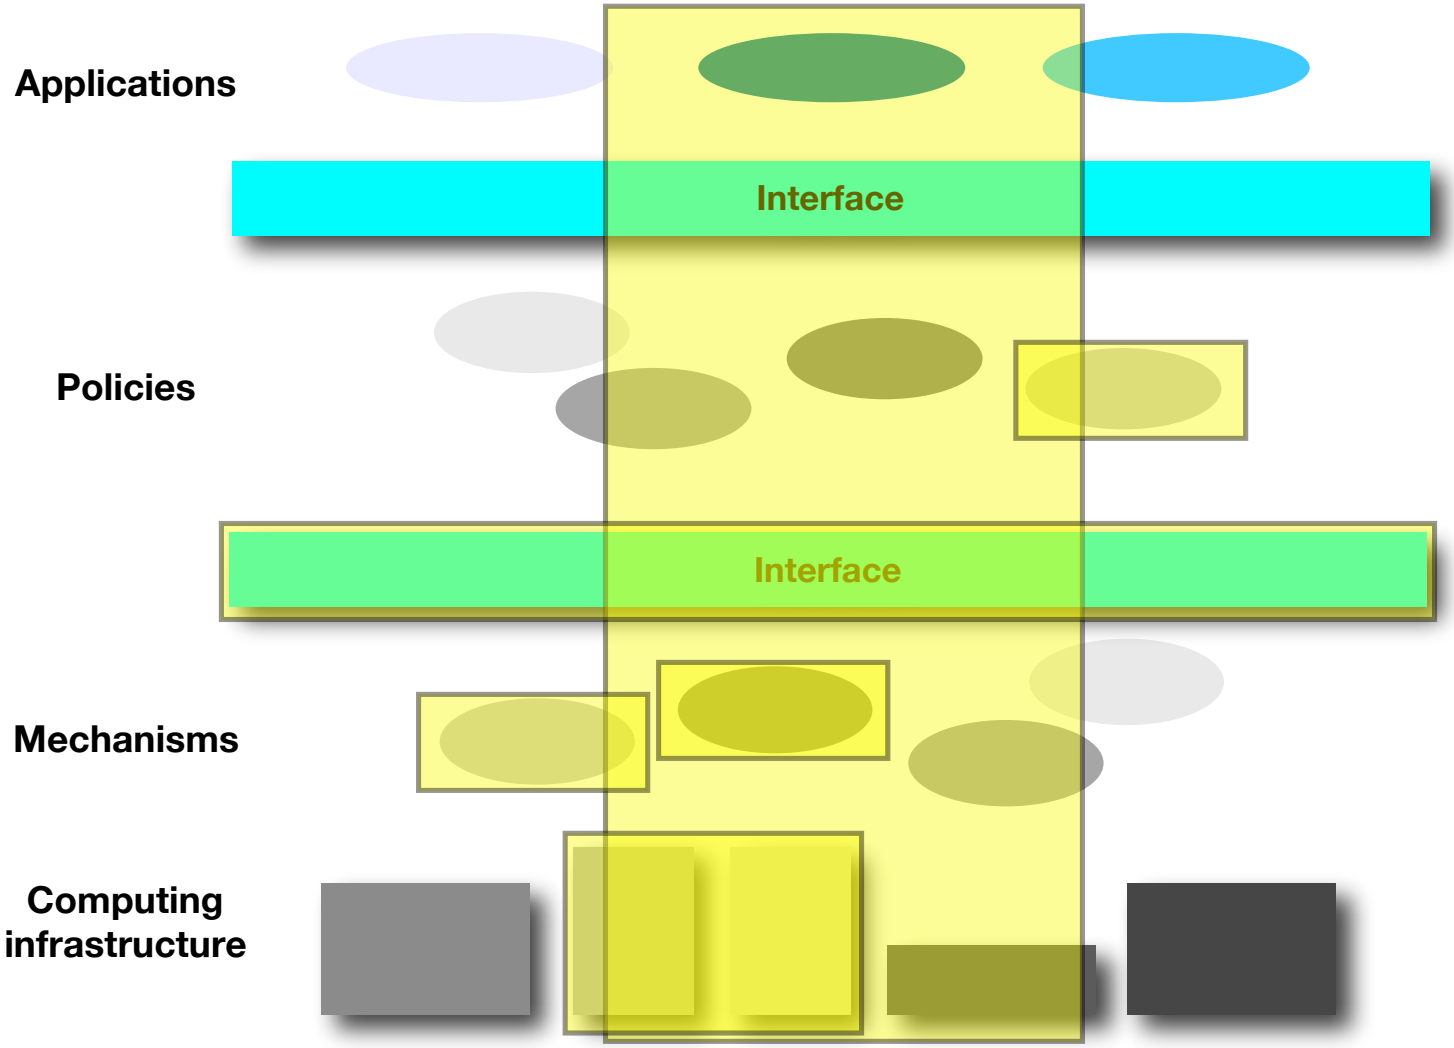

The challenges: global vs. grid computing ???

- is there any difference among global and grid computing ?
  - if we get rid of terminology issues
  - if we choose each time a common abstraction level
  - if take the more general/abstract viewpoint possibile