2008 eScience 4th IEEE International Conference on e-Science

Classification of Different Approaches for e-Science Applications in Next Generation Infrastructures

Morris Riedel et al.
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

Distributed
European
Infrastructure for
Supercomputing
Applications

HELMHOLTZ ASSOCIATION

JÜLICH FORSCHUNGSZENTRUM

Outline

- e-Science and Grids
- e-Science in HPC-driven Grids
- Approach I. Simple Scripts & Control
- Approach II. Scientific Application Plug-ins
- Approach III. Complex Workflows
- Approach IV. Interactive Access
- Approach V. Interoperability
- Lessons Learned & Conclusions
- Announcement: OGF Production Grid Infrastructure (PGI) WG
- References

e-Science and Grids

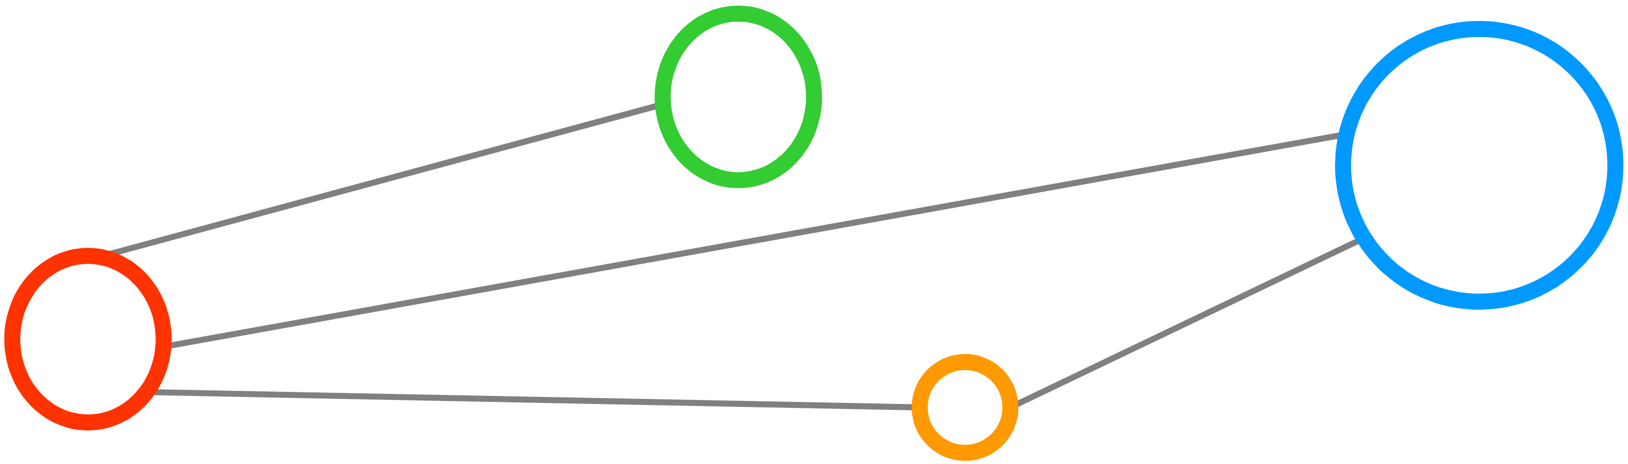

Scientific Computing

- *'Today, the natural sciences regard computational techniques as a third pillar alongside experiment and theory'*

[1] Lippert et al., *'IBM delivers Europe's biggest supercomputer'*

science - scientific innovation - understanding of earth fundamentals

I.
Theory
(and models)

III.
Computational
Techniques

II.
Experiment

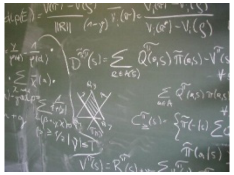
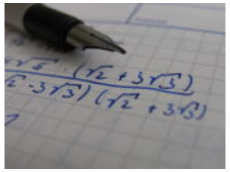



Computational Science & Simulations

- *'Today, the natural sciences regard computational techniques as a third pillar alongside experiment and theory'*

[1] Lippert et al., *'IBM delivers Europe's biggest supercomputer'*

III. Computational Techniques

- Computational techniques that enable computational science
- Some examples of computational science
  - e.g. large-scale data processing to screen for a new drug
  - e.g. simulation of systems with infinite variables in the fields of chemistry, physics, or computational biology
  - e.g. in many problems in solid state or high energy physics, only computer simulations can provide the missing links between empiricism and theory

Computational Science & Parallel Computing

III. Computational Techniques

- Scientists of different kinds of science…
  - Physics, chemistry, biology,…
  - They improve their work with computational science
- Parallel Computing has emerged as an effective technique to support scientific work
  - Execute scientific applications in parallel on more than one CPU/core
  - Implementations compliant with Message Passing Interface (MPI)
- Most often used programming languages…
  - Fortran, C, C++
  - Side-remark: Java basically is not a player here

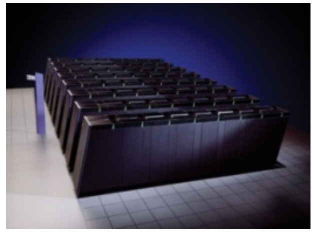

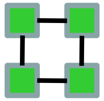

cpus

[2] '*The Message Passing Interface (MPI) Standard*'

What is e-Science (enhanced Science)?

[3] John Taylor
'The definition of e-Science'

- *"e-Science is about global collaboration in key areas of science and the next generation infrastructure that will enable it"*
- 'Next generation infrastructures' represented by Grids today (more recently, also partly represented by so-called 'clouds')
  - 'Services' for seamless and secure collaborations over networks and organizational boundaries

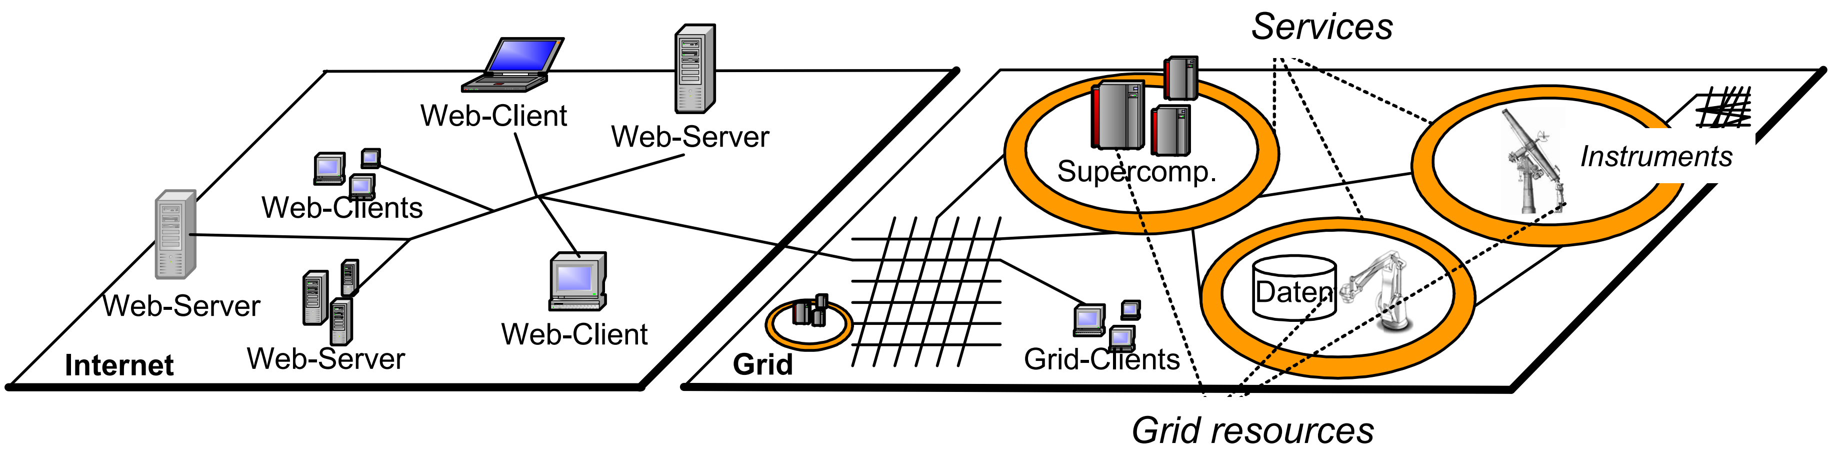

e-Science with Next Generation Infrastructures

- *“e-Science is about global collaboration in key areas of science and the next generation infrastructure that will enable it”*

[3] John Taylor
‘The definition of e-Science’

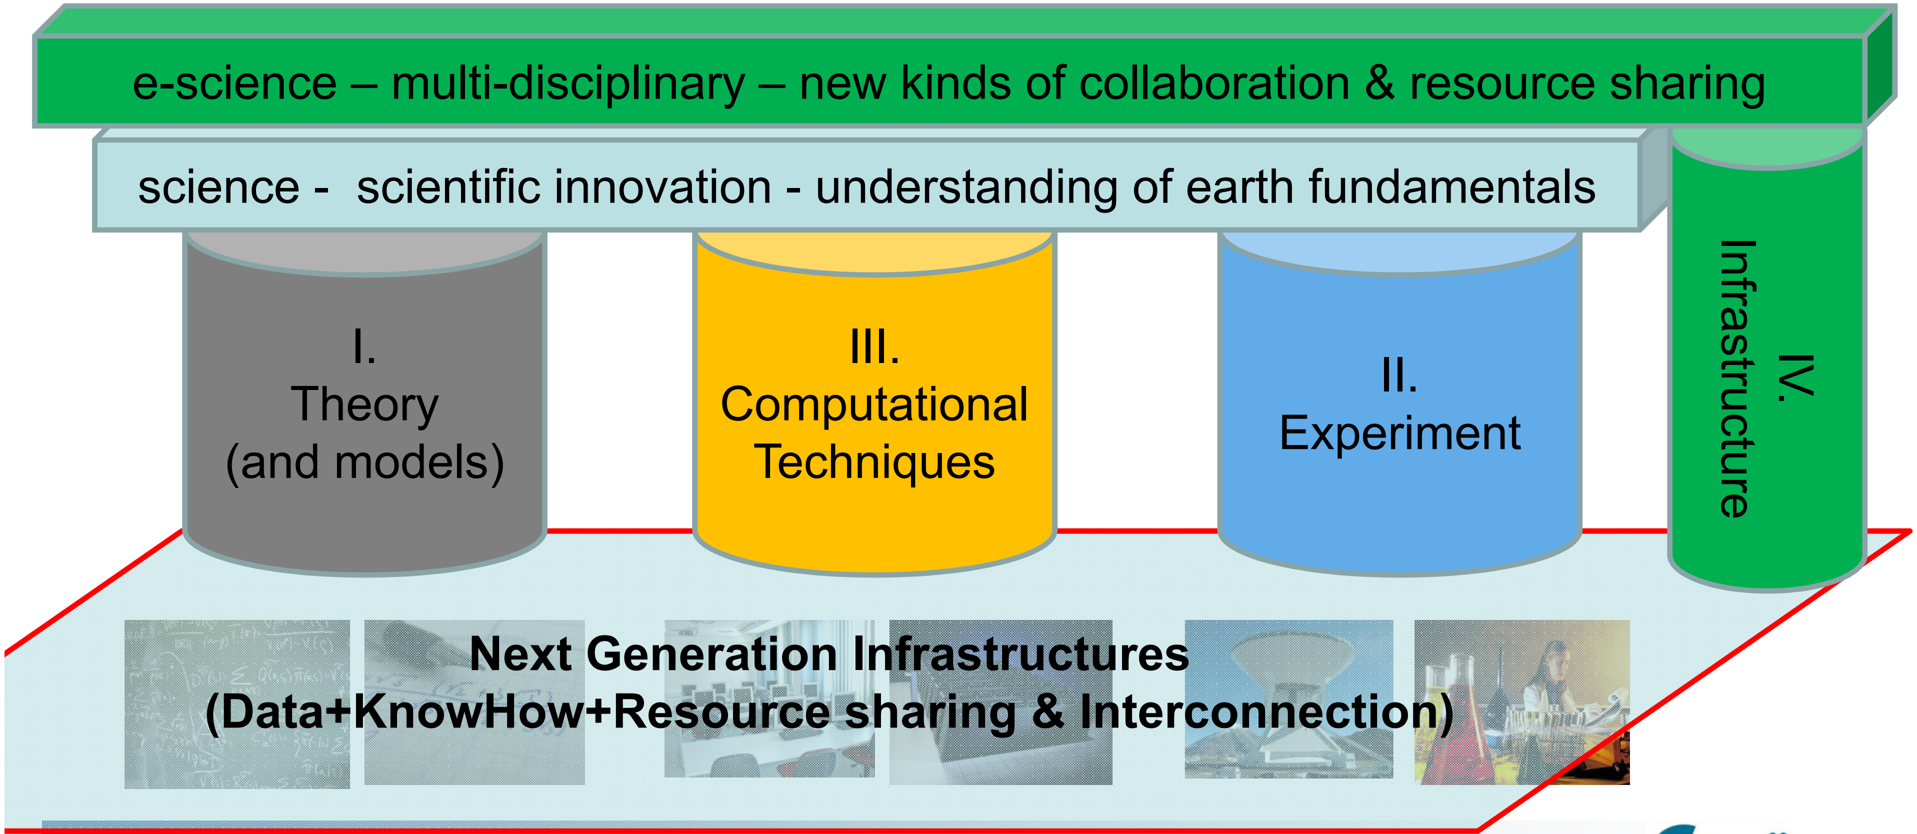

Next Generation Infrastructure Examples

- European Infrastructures
  - Distributed European Infrastructure for Supercomputing Applications (DEISA)
  - Enabling Grids for e-Science (EGEE)

- National Infrastructures
  - TeraGrid – USA
  - Open Science Grid (OSG) – USA
  - D-Grid – Germany
  - National Grid Service (NGS) United Kingdom
- And many others…

[4] DEISA

[5] EGEE

[6] TeraGrid

[28] OSG

[7] D-Grid

[8] NGS

e-Science & Key Areas of Science

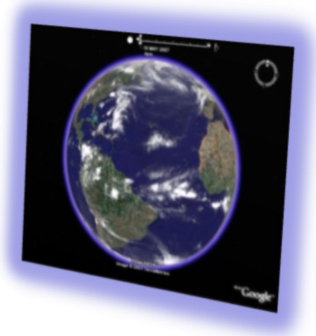

- Often solve “Grand Challenge Problems”
  - E.g. Protein folding, global weather prediction, …
- Major problems of science and society today (e.g. diseases)
  - Cannot be solved in a reasonable amount of time with computers that are broadly available today
  - e.g. Simple PC for protein folding?!

- Problems (e-Science applications) can in principle be tackled
  - Require an increase in computing power by orders of magnitude
  - Solutions can have a significant economic and/or social impact
- Next generation infrastructures provide such computing power

e-Science in HPC-driven Grids

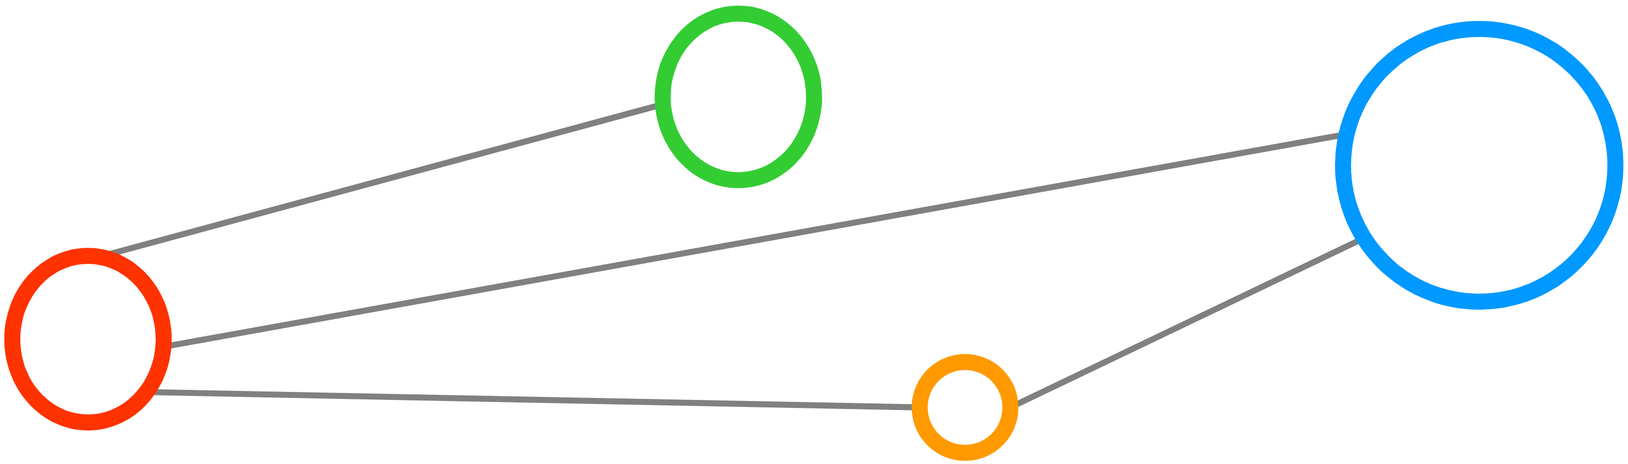

Scientific Applications as Grid Jobs

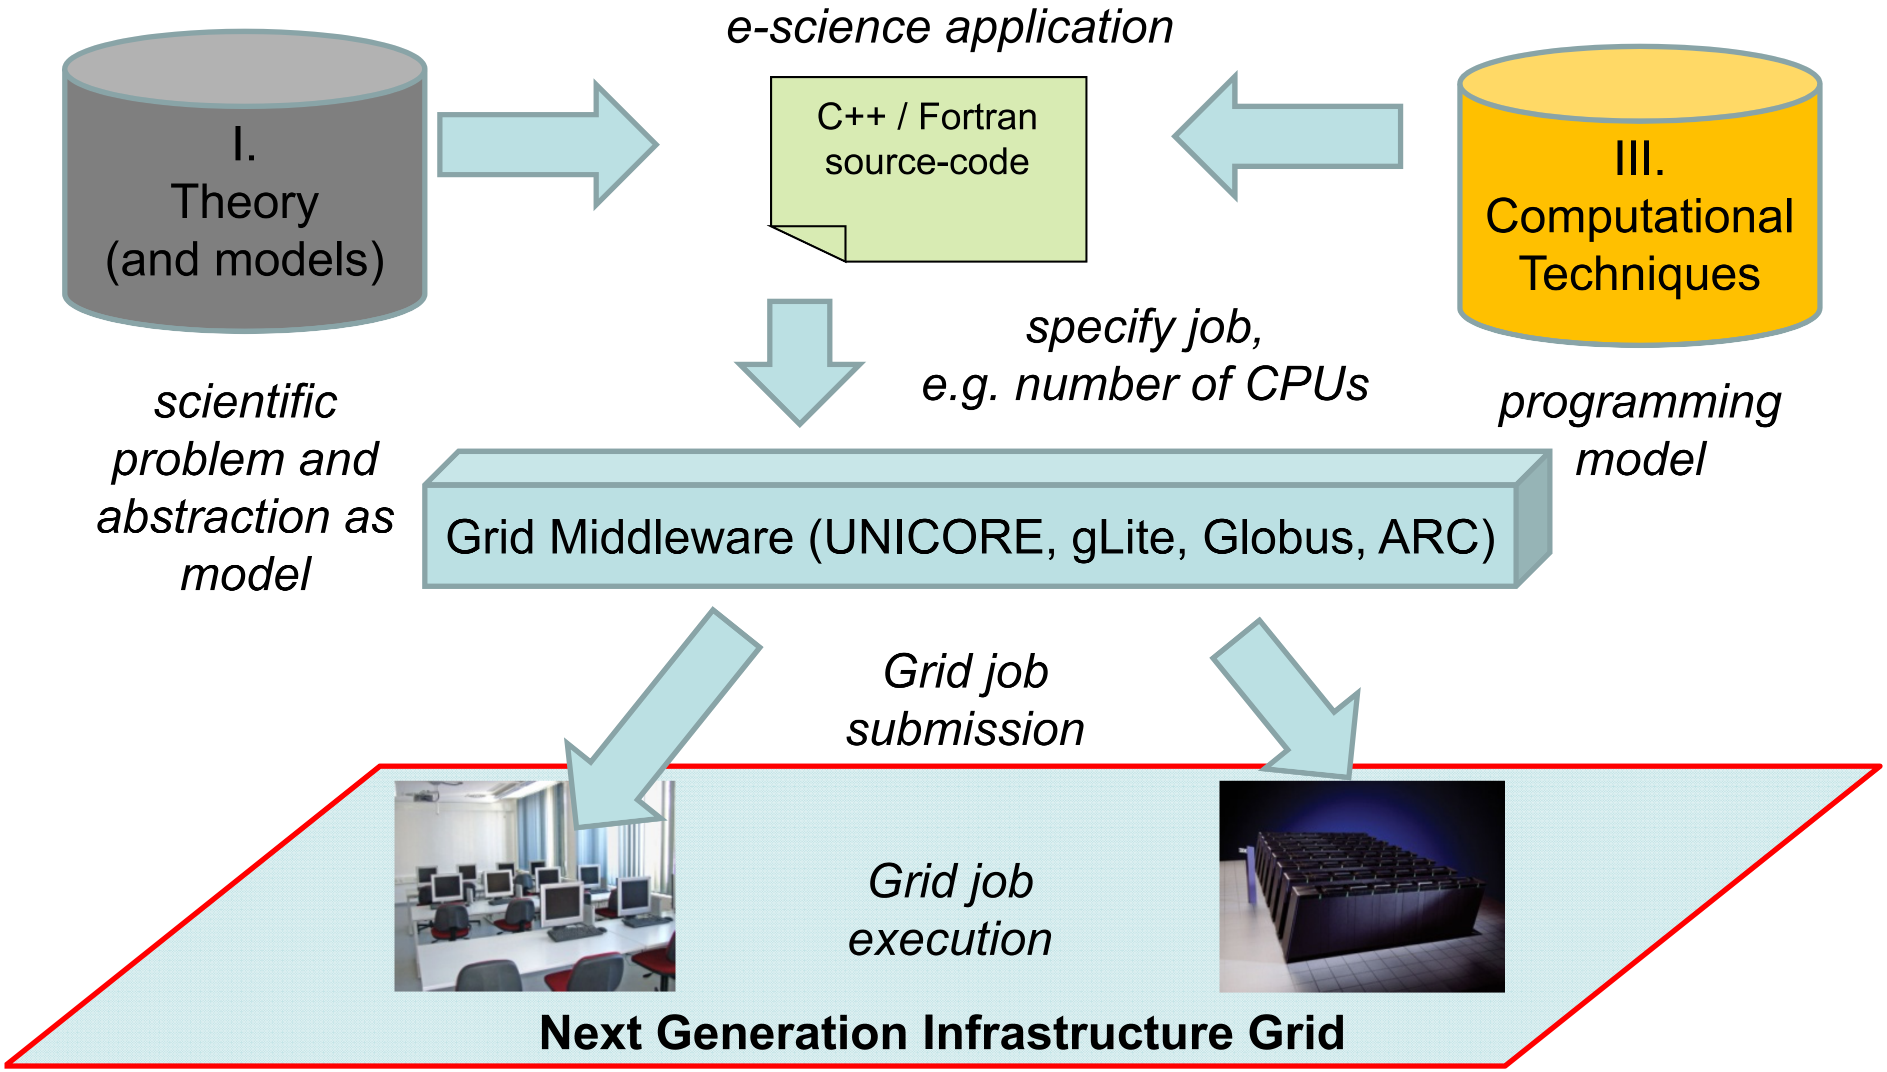

HPC-based e-Science with UNICORE

- e-Science applications are used with Grid middleware
  - Here: HPC-driven Grid Infrastructures running UNICORE
  - High Performance Computing (HPC) → massively parallel applications
  - Massively parallel: n CPUs/cores interact with each other

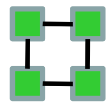

cpus

- HPC-driven Grid Infrastructure with multiple (different) supercomputers as ‘Grid resources’
  - Supercomputing Grid Infrastructure DEISA use Grid Middleware UNICORE

Distributed European Infrastructure for Supercomputing Applications

  - Future: Partnership for Advanced Computing in Europe (PRACE)

[9] PRACE

    - New HPC systems with unprecedented scale (i.e. petaflop/s range)

How do e-Scientists get computational time?

scientific problem and expert in the field has an idea (e-Scientist)

project proposal – apply for computational time (get CPU time)

evaluation from a scientific committee – (project granted)

project granted – (computational time available)

- Necessary because HPC resources are very costly and highly demanded – there are typically no ‘idle supercomputers’
  - Such systems are typically even ‘overbooked’

[29] DECI

- Examples: DEISA Extreme Computing Initiative (DECI)
  - e-Science applications requiring high amounts of CPUs and memory using the DEISA Supercomputing Grid infrastructure

Examples of DECI e-Science Projects

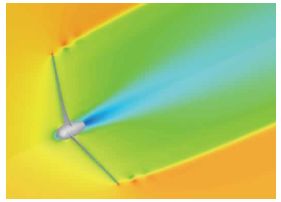

- 3C4WTS → Large scale Wind Turbines
  - Computational Fluid Dynamics (CFD), Wind Turbines, Aerodynamic Design

[19] DECI Projects 2008

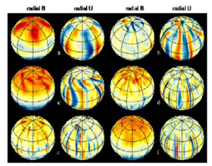

- 3DEarth
  - Numerical modeling of geodynamical processes
- AntiEflx → Bacterial resistance against antibiotics research
  - Computational Biophysics

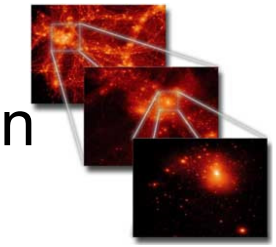

- AQUILA → Hydrodynamical cosmological simulation
  - Earth System Science
- Dratchet → Particle transport studies
  - fluid-structure interactions, multi-physics simulations, parallel algorithms, theoretical physics, particle sorting on a microscale
- …and many others…

Classification of Different Approaches

- Once computational time is available the question arises how exactly e-scientists use an HPC-based Grid infrastructure
- There are different approaches how e-scientists use these infrastructures in general and UNICORE in particular

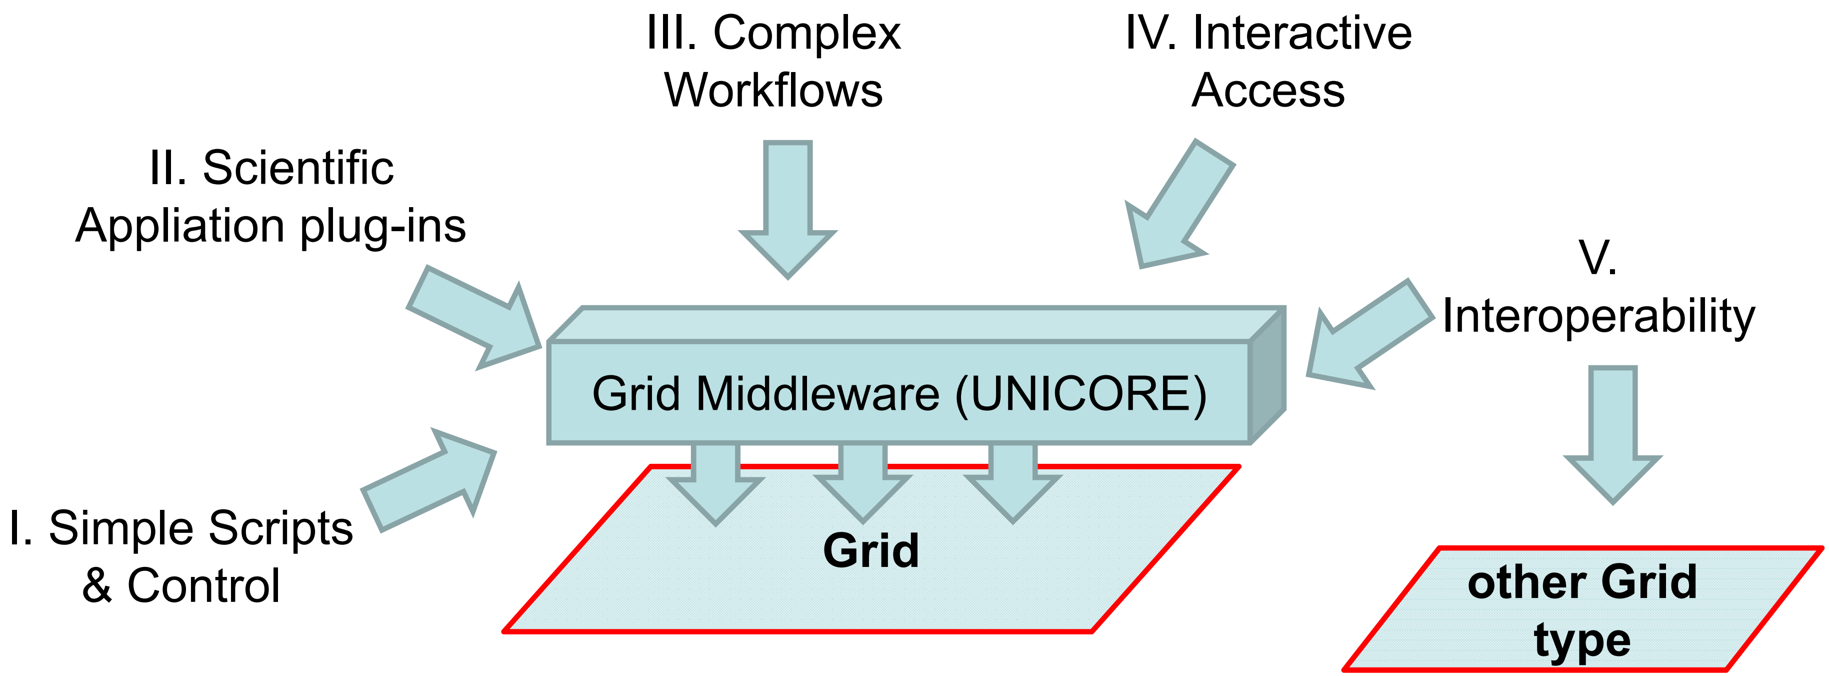

Approach I. Simple Scripts & Control

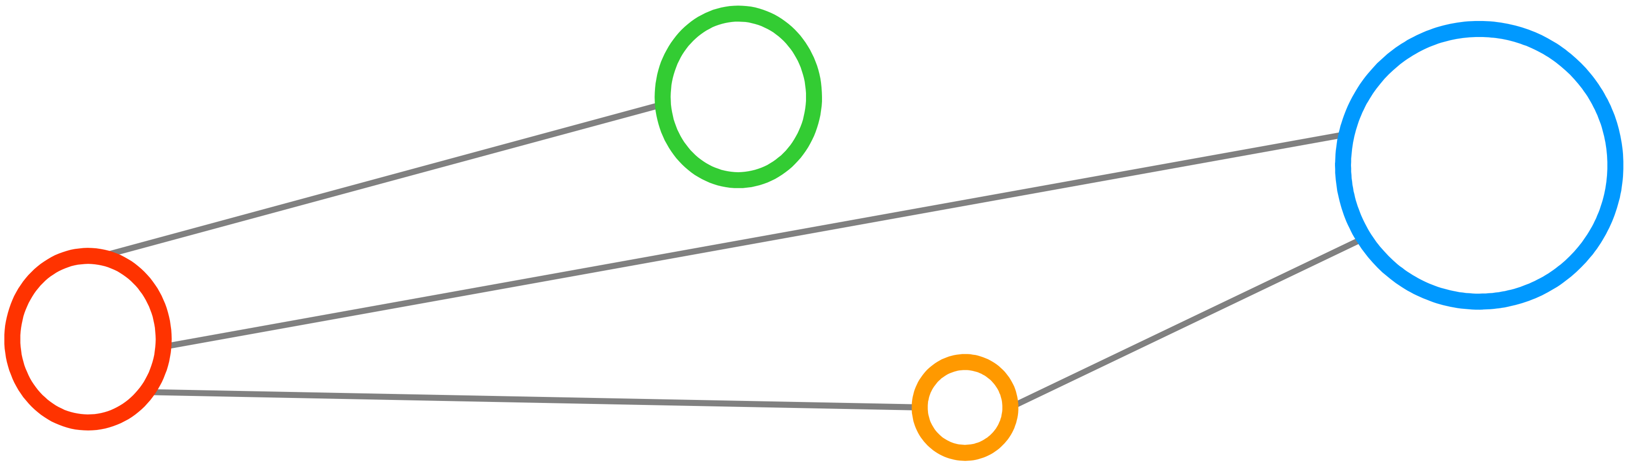

Approach Simple Scripts & Control

- Scientific application is submitted as a Grid job
- Grid job is a simple UNIX-based script calling an executable
- Computation is influenced by control functionalities

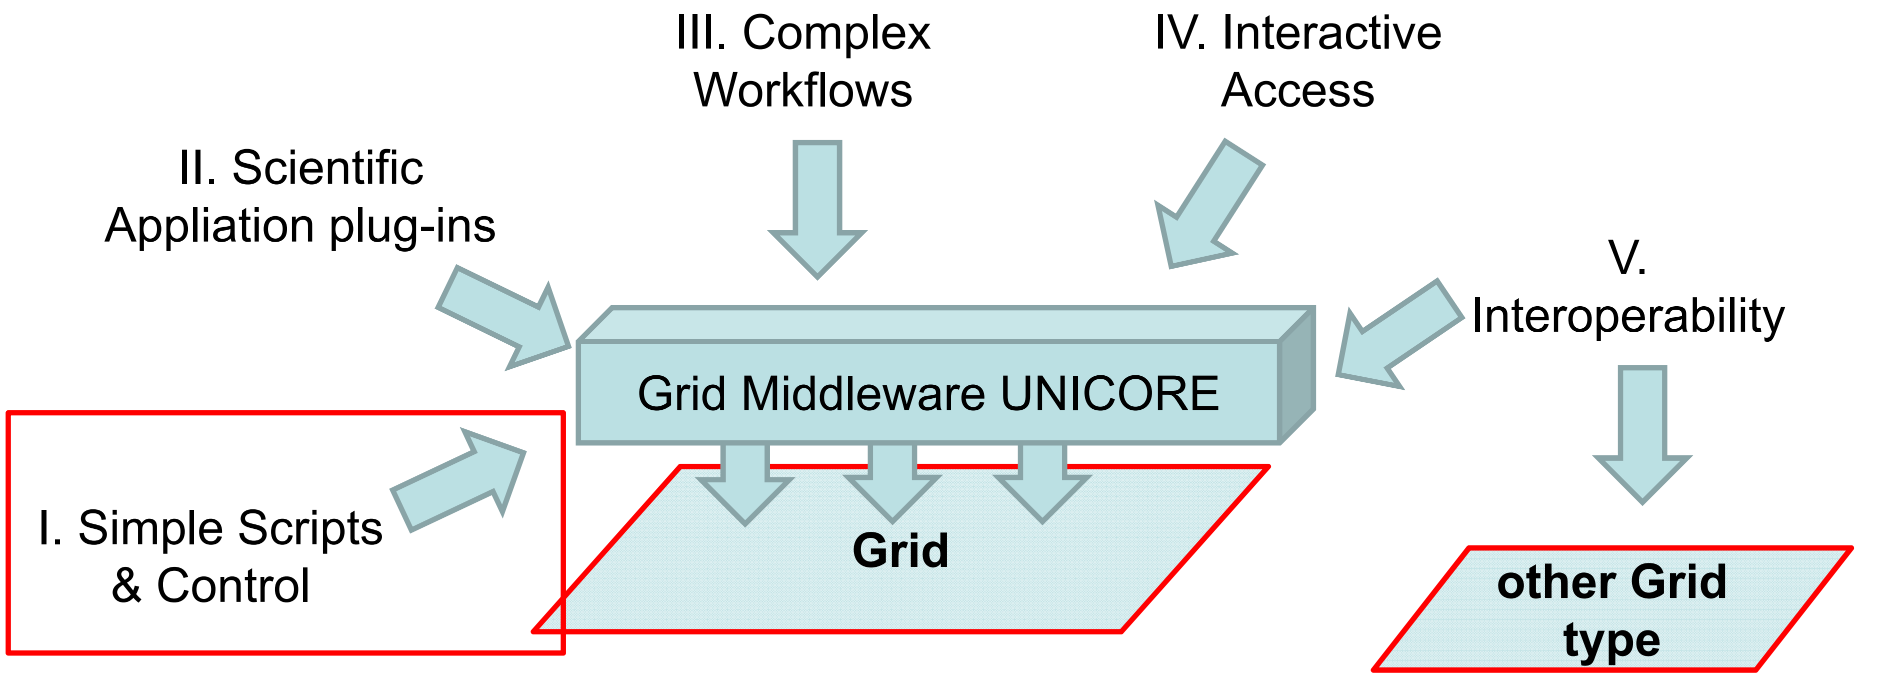

Control Functionalities

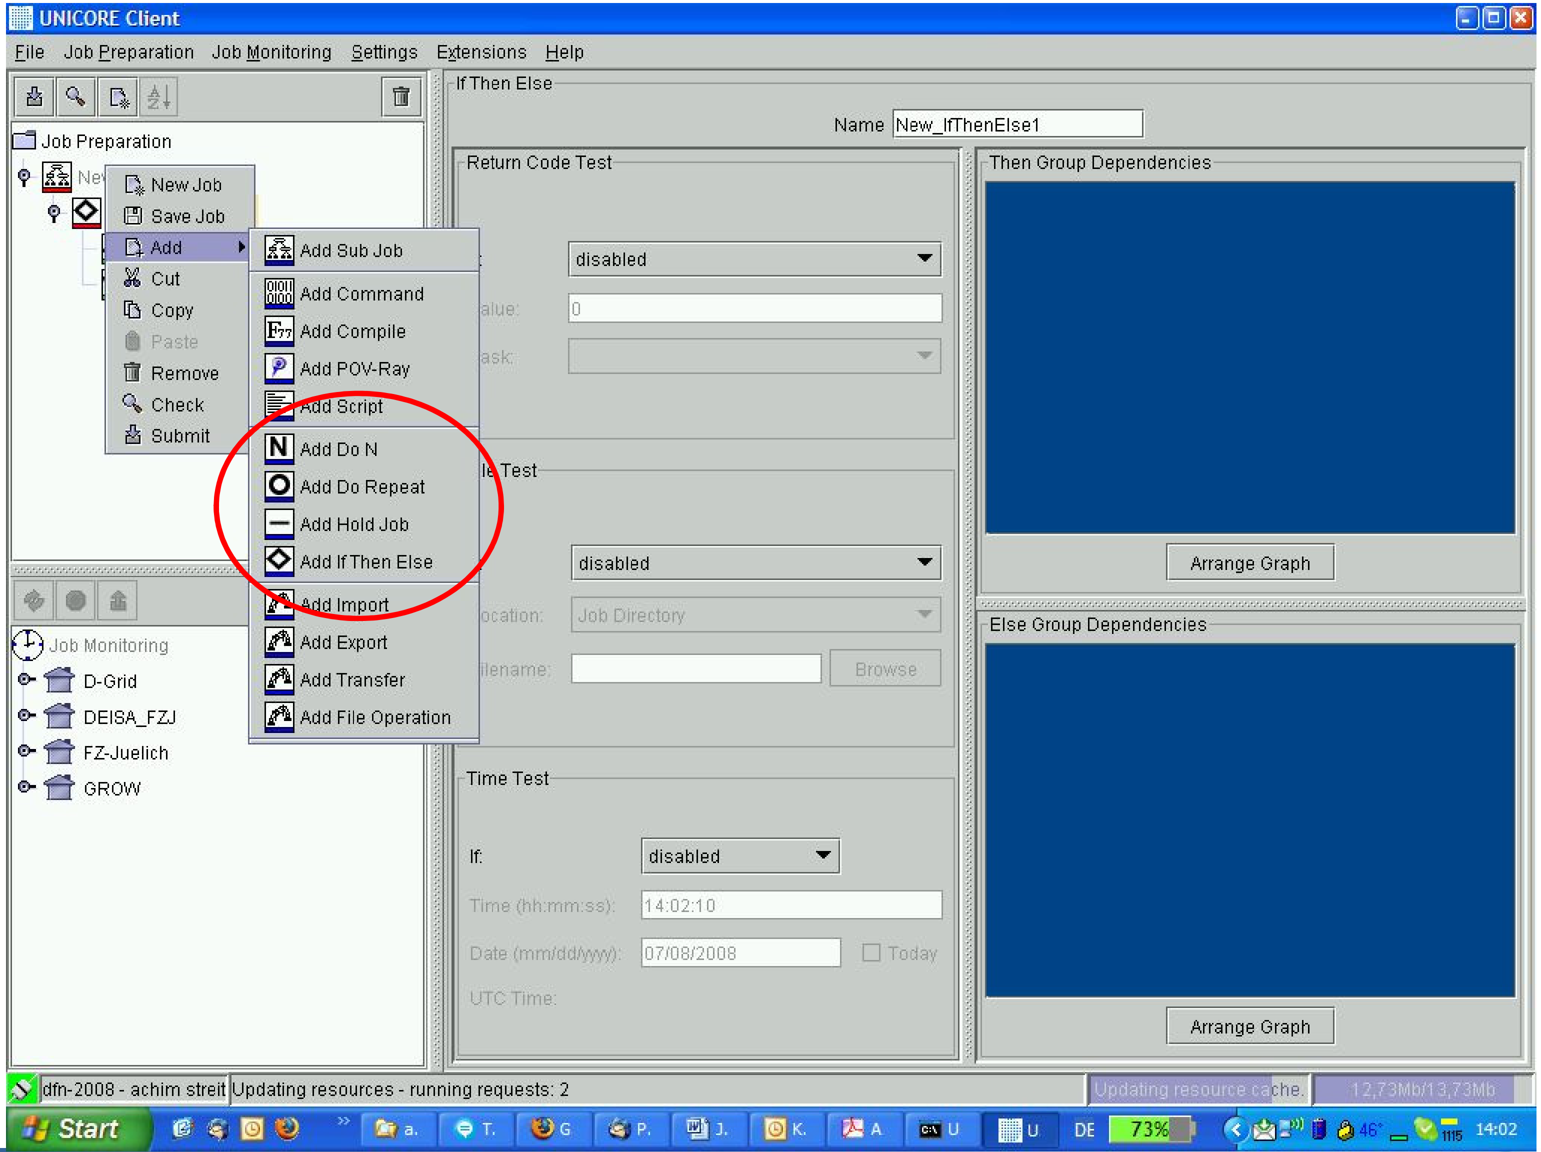

e-Science Example - Hydrodynamics

- Scientific background: Hydrodynamics
  - Study of liquids in motion, sub-discipline of fluid dynamics
- Hydrodynamics of Active Biological Systems → Sperm
  - Study Sperm cluster size dependence for 2D and 3D systems
  - Experiments have revealed an interesting swarm behavior of sperm, when the sperm concentration of the system is high
  - The mechanism behind the experimental phenomenon is still not clear
  - Even after 2D Simulations not clear

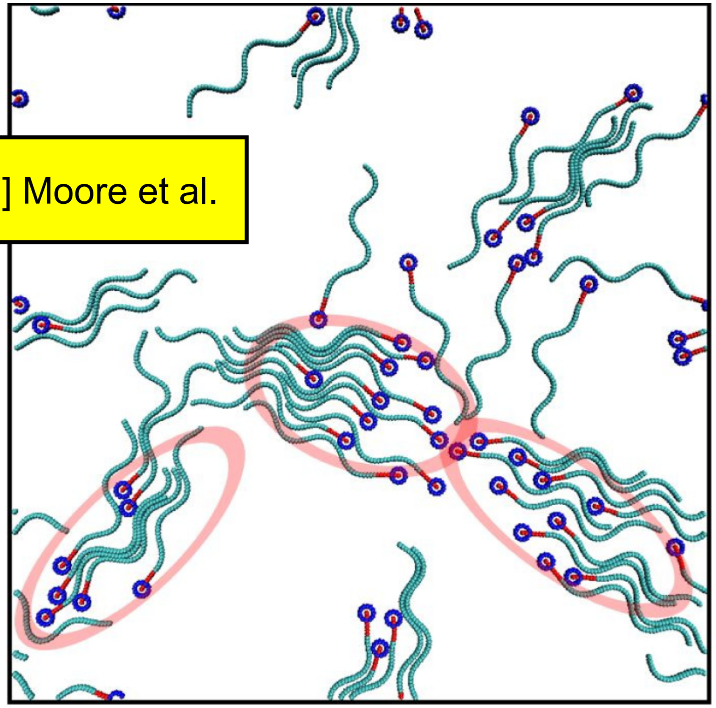

[10] Moore et al.

Hydrodynamics using UNICORE (1)

- 2D studied – now study 3D systems
  - Goal: Study the hydrodynamics interaction between sperm and explain its importance to the cooperation behavior
  - Simulations in 3D are very time consuming so that a systematic study is not yet done
- UNICORE is used for job submission (c++ executable) using a simple UNIX-style script and the LOOP DO-N control functionality

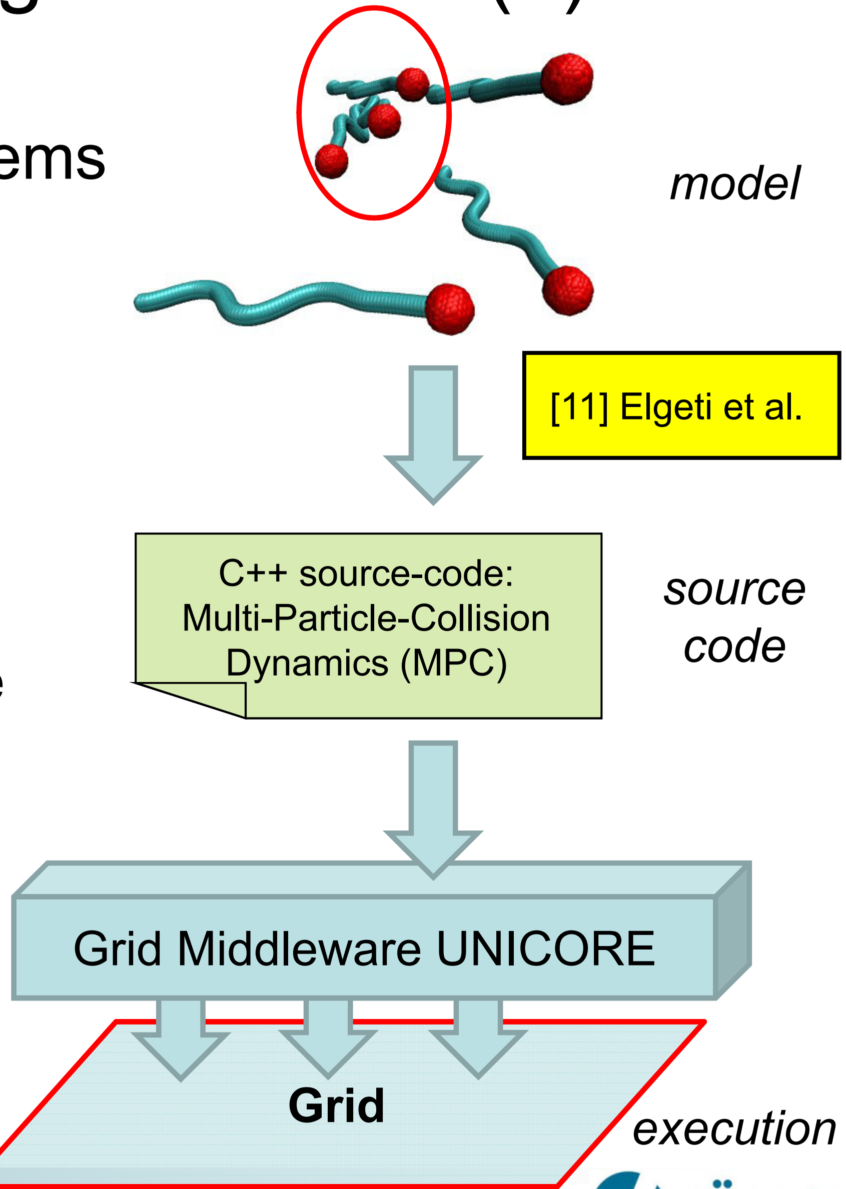

Hydrodynamics using UNICORE (2)

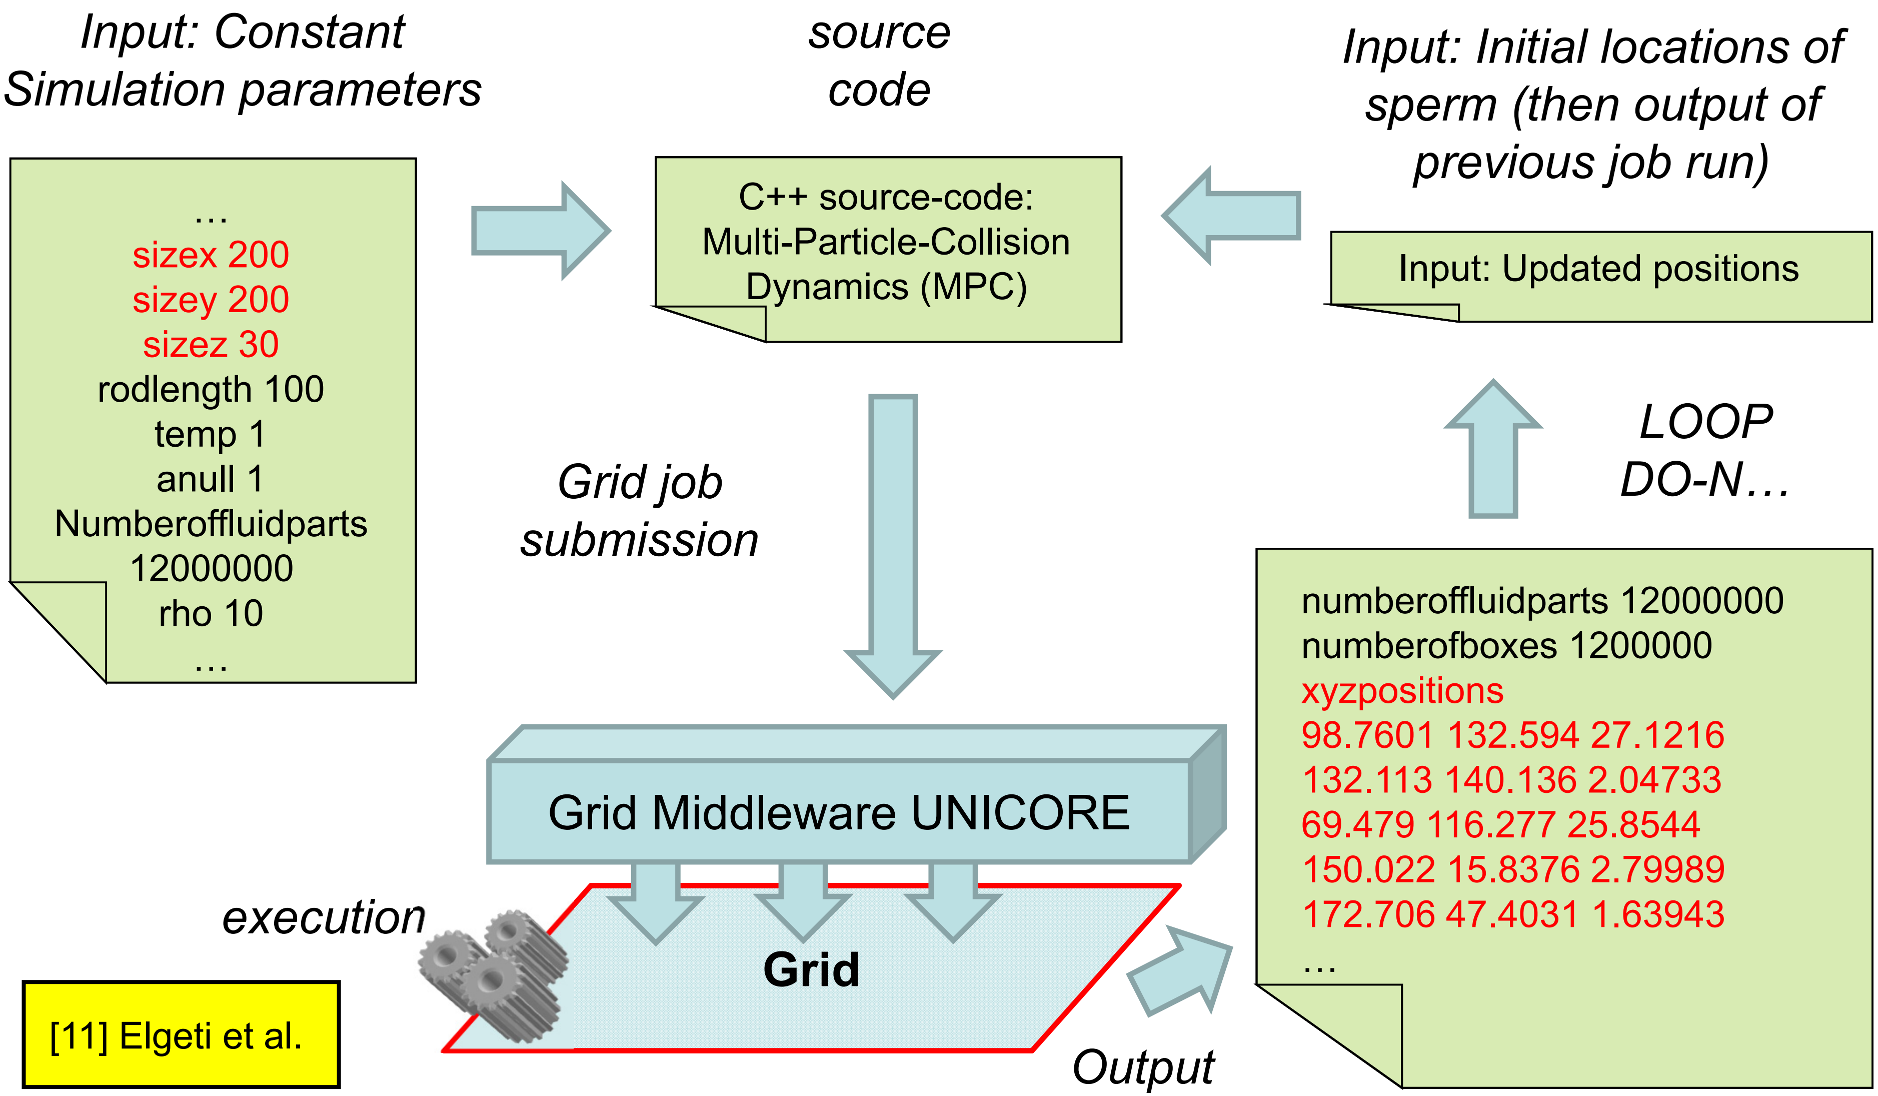

Approach II. Scientific Application Plug-ins

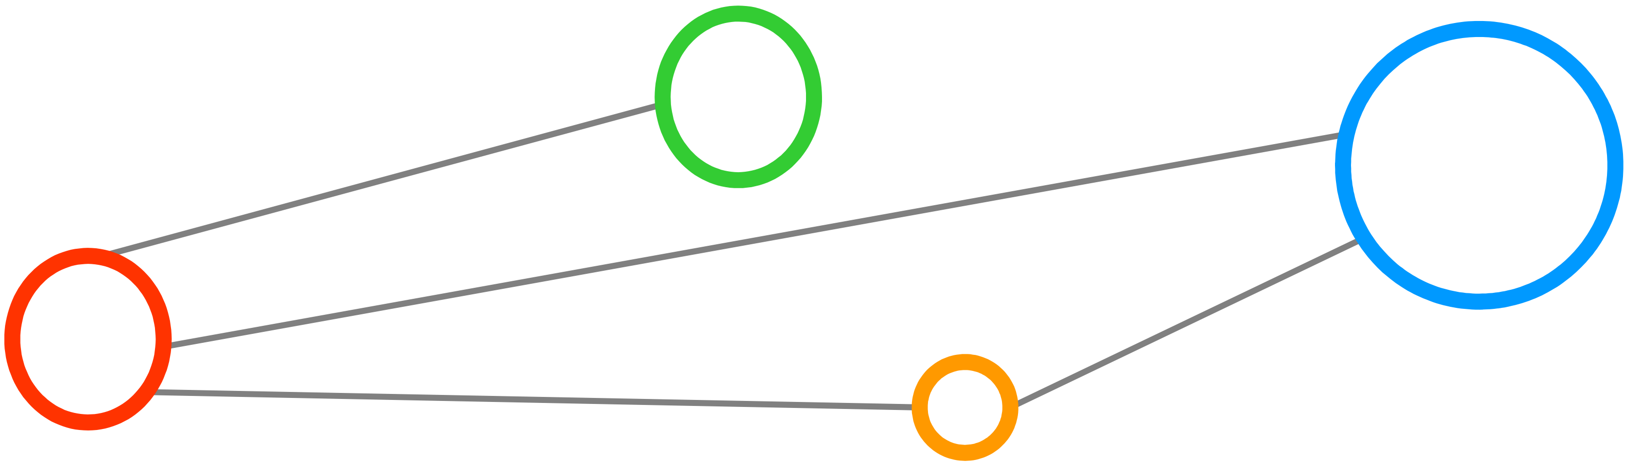

Approach Scientific Application Plug-ins

- Scientific application is specifically supported in the UNICORE client/portal as scientific domain-specific plug-in
- Plug-ins (may) support a scientific application that is used in a wide variety of different e-science projects

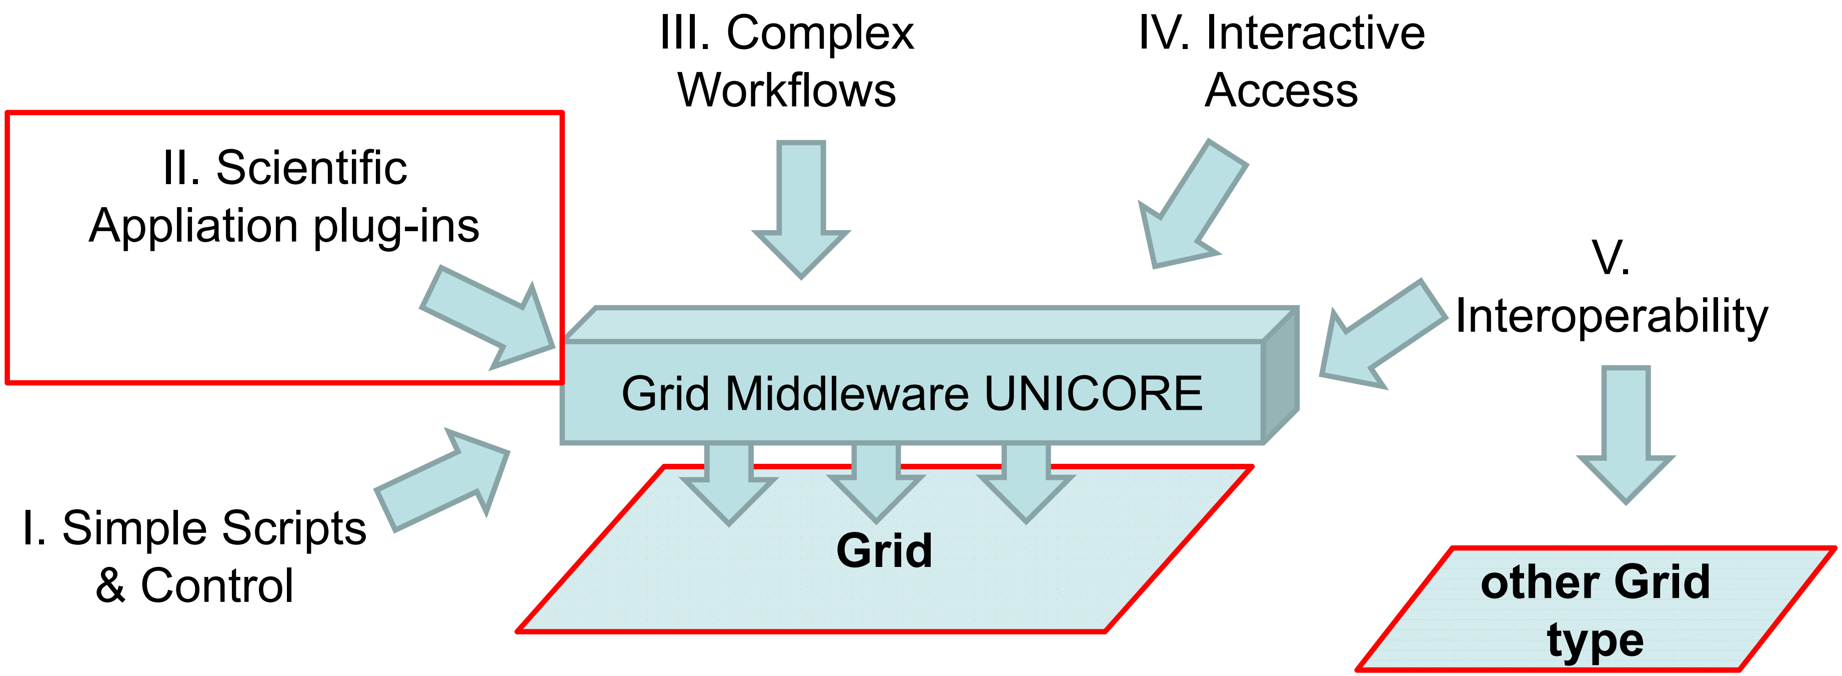

e-Science Example – Gaussian (1)

[25] Gaussian

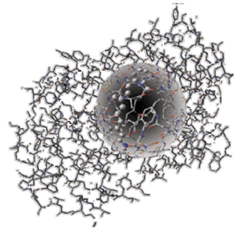

- Scientific background: Gaussian
  - Gaussian03 is an electronic structure program (e.g. chemistry)
- Gaussian Plug-in for one of the UNICORE clients
  - Usable in many different e-science applications
  - Specific job definition for Gaussian jobs
- Other plug-ins available
  - Assisted Model Building for Energy Refinement (AMBER) package

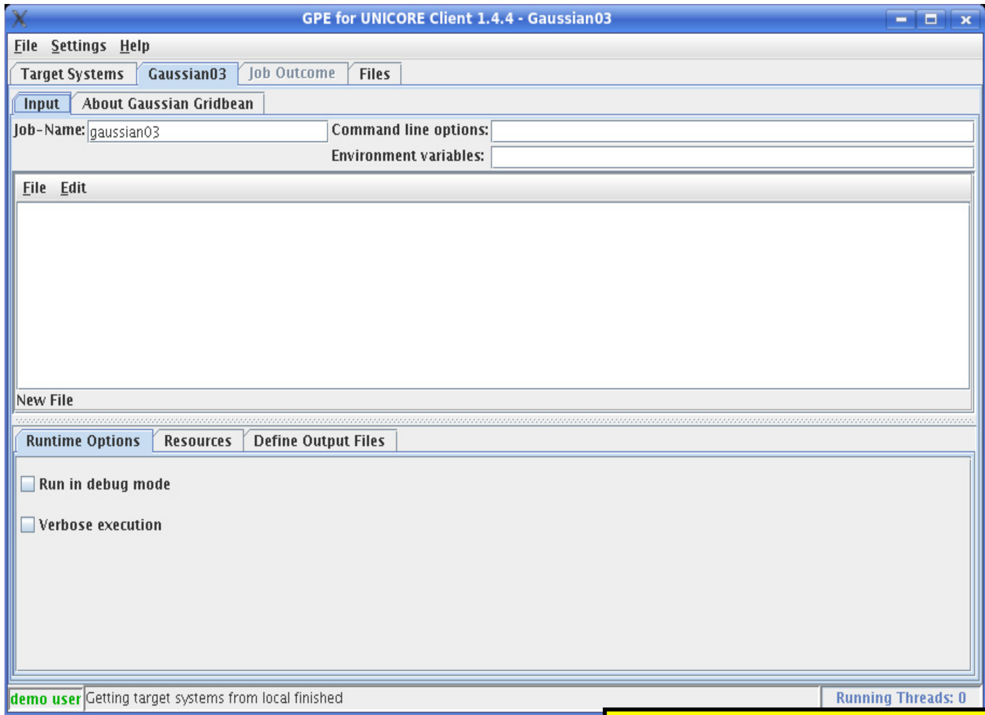

work in progress

[13] AMBER

[12] Gaussian Plug-in

e-Science Example – Gaussian (2)

- Specify precise requirements for computational resources

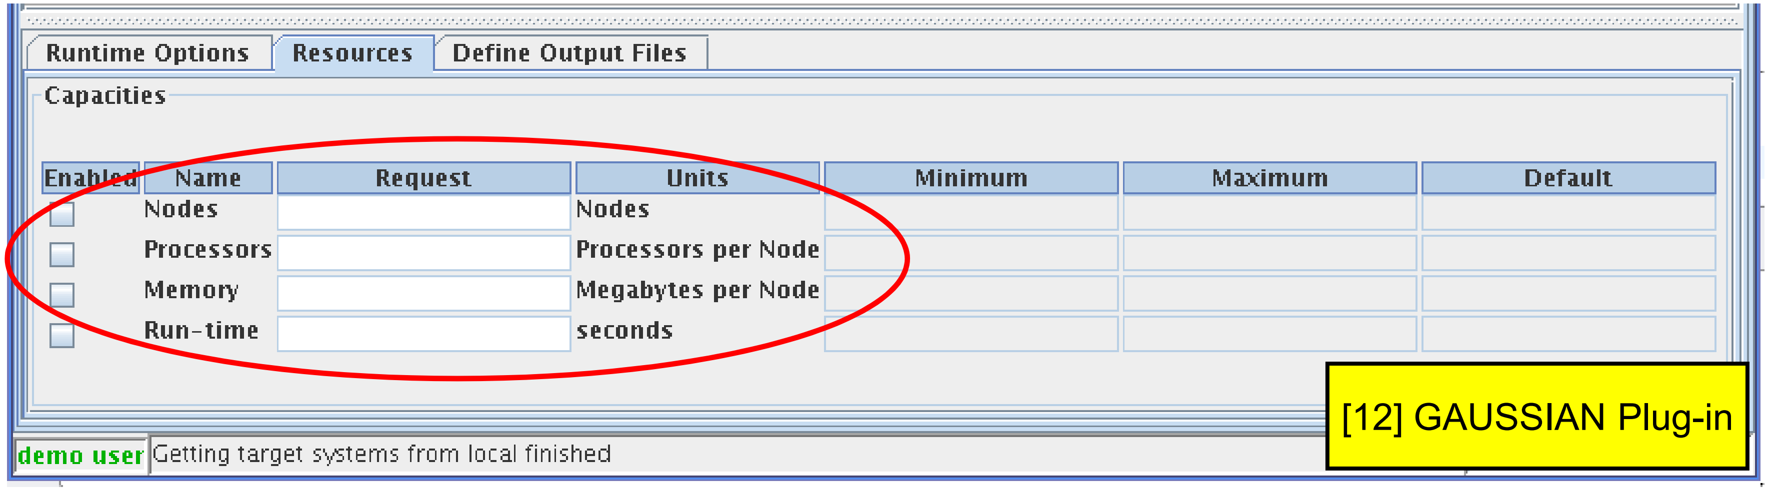

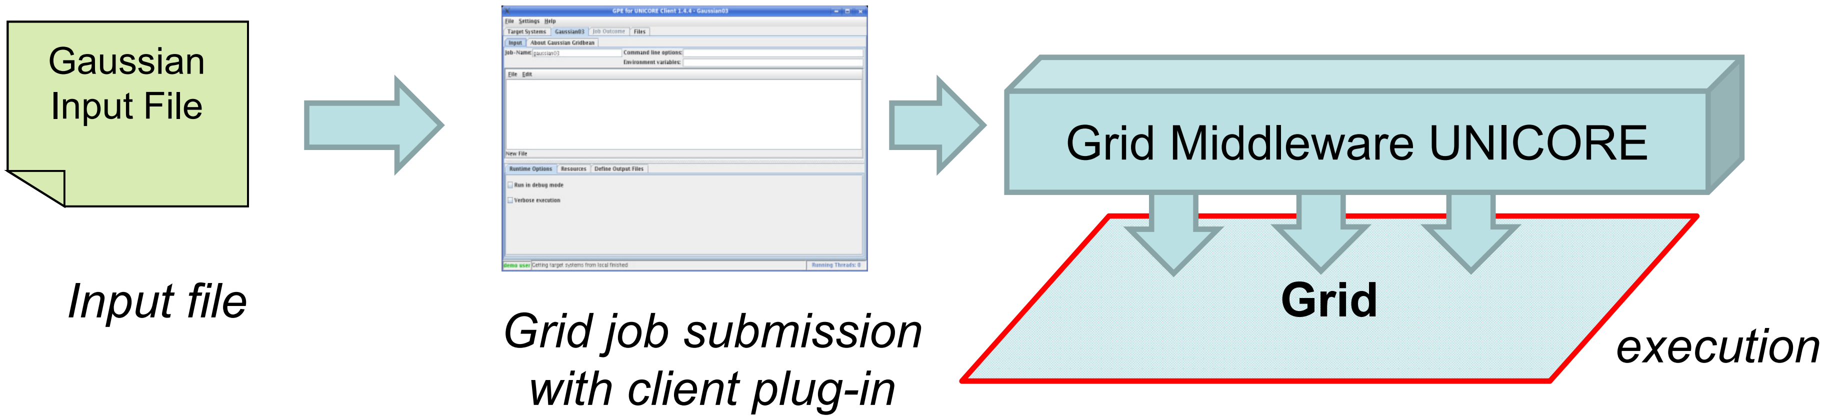

Approach III. Complex Workflows

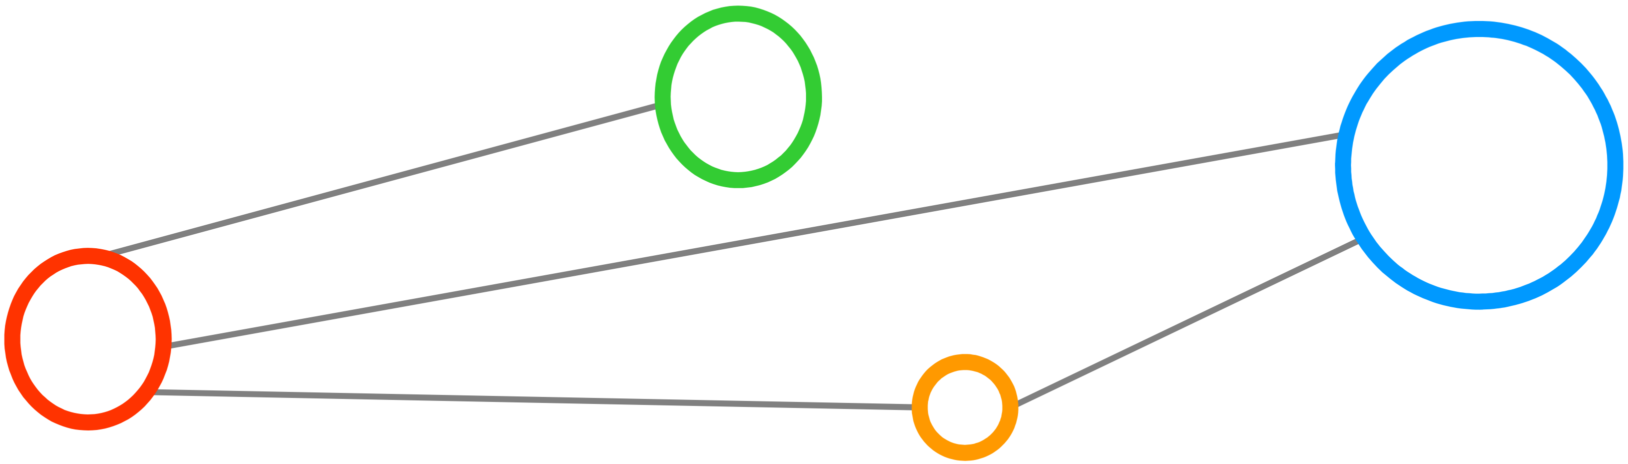

Approach Complex Workflows

- Different Grid Tasks together as one Grid workflow

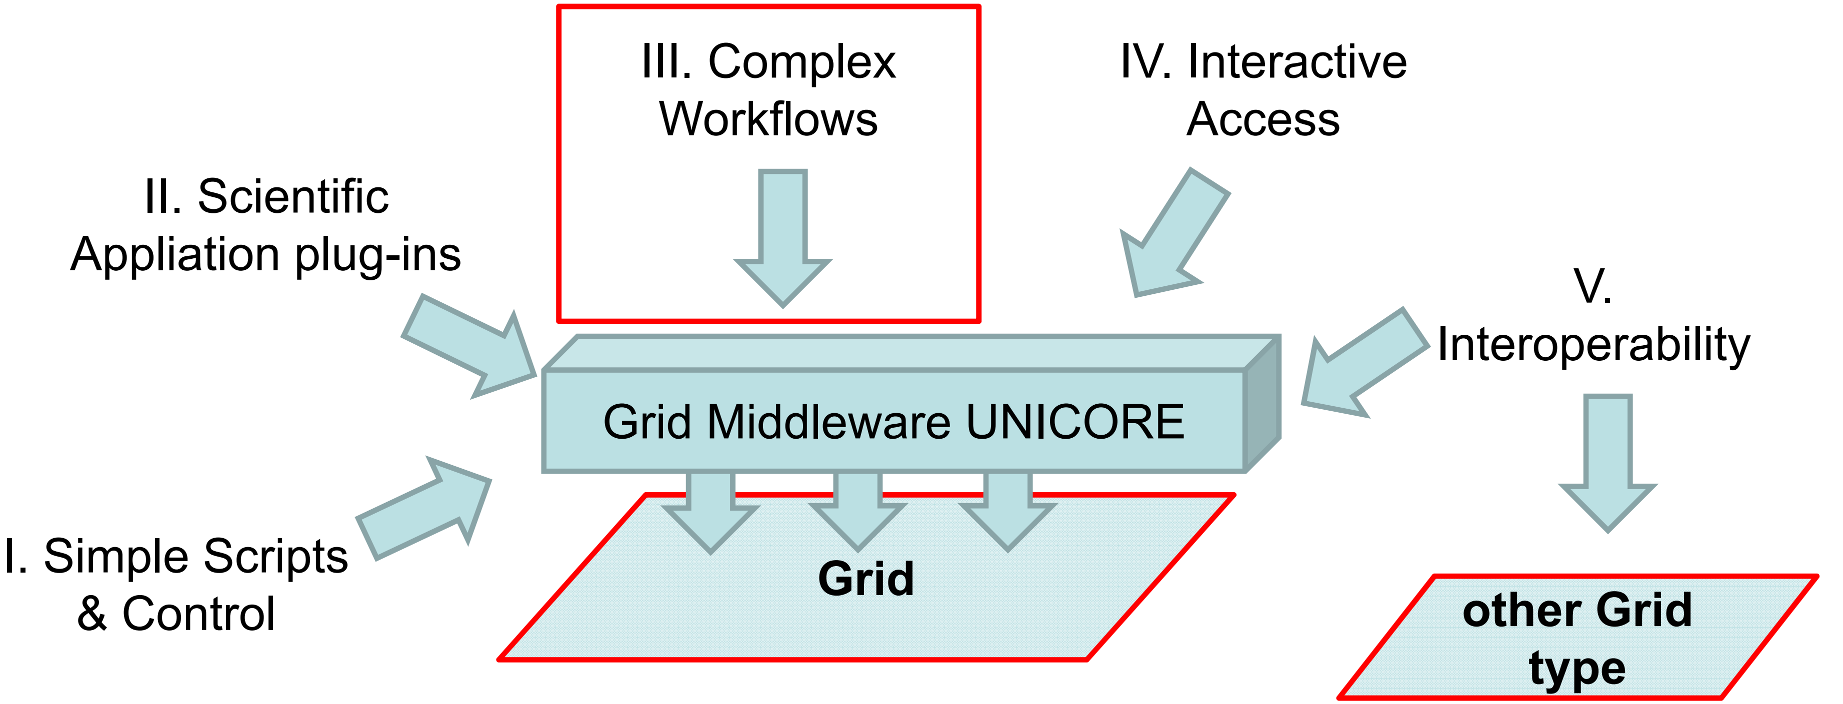

e-Science Example – QSAR Workflow (1)

- Scientific background: QSAR for Regulatory Purposes
  - Current regulatory framework: Registration, Evaluation, Authorisation, and Restriction of Chemical substances (REACH)
  - Goal: improve the protection of the human health/environment through the characterization of intrinsic properties of chemicals
  - Technique/computational method: Quantitative Structure-Activity Relationships (QSAR)
- UNICORE with QSAR workflows
  - Combine different QSAR applications in workflows
  - Access to existing databases
  - Get results with their documentation

[14] Casalegno et al.

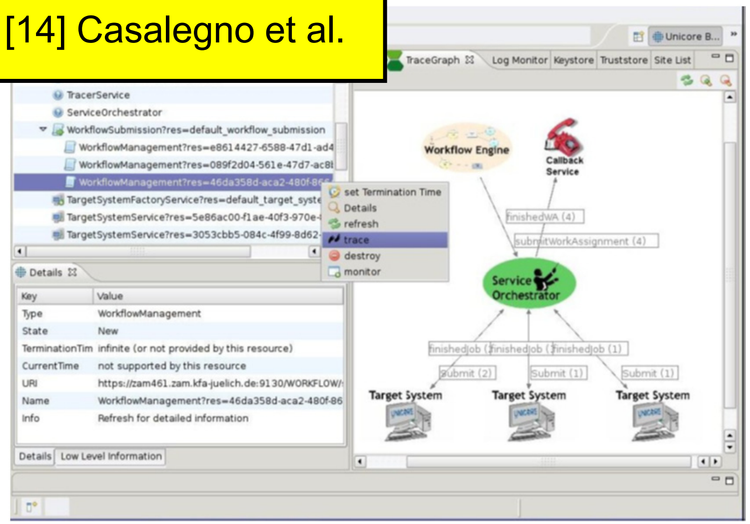

e-Science Example – QSAR Workflow (2)

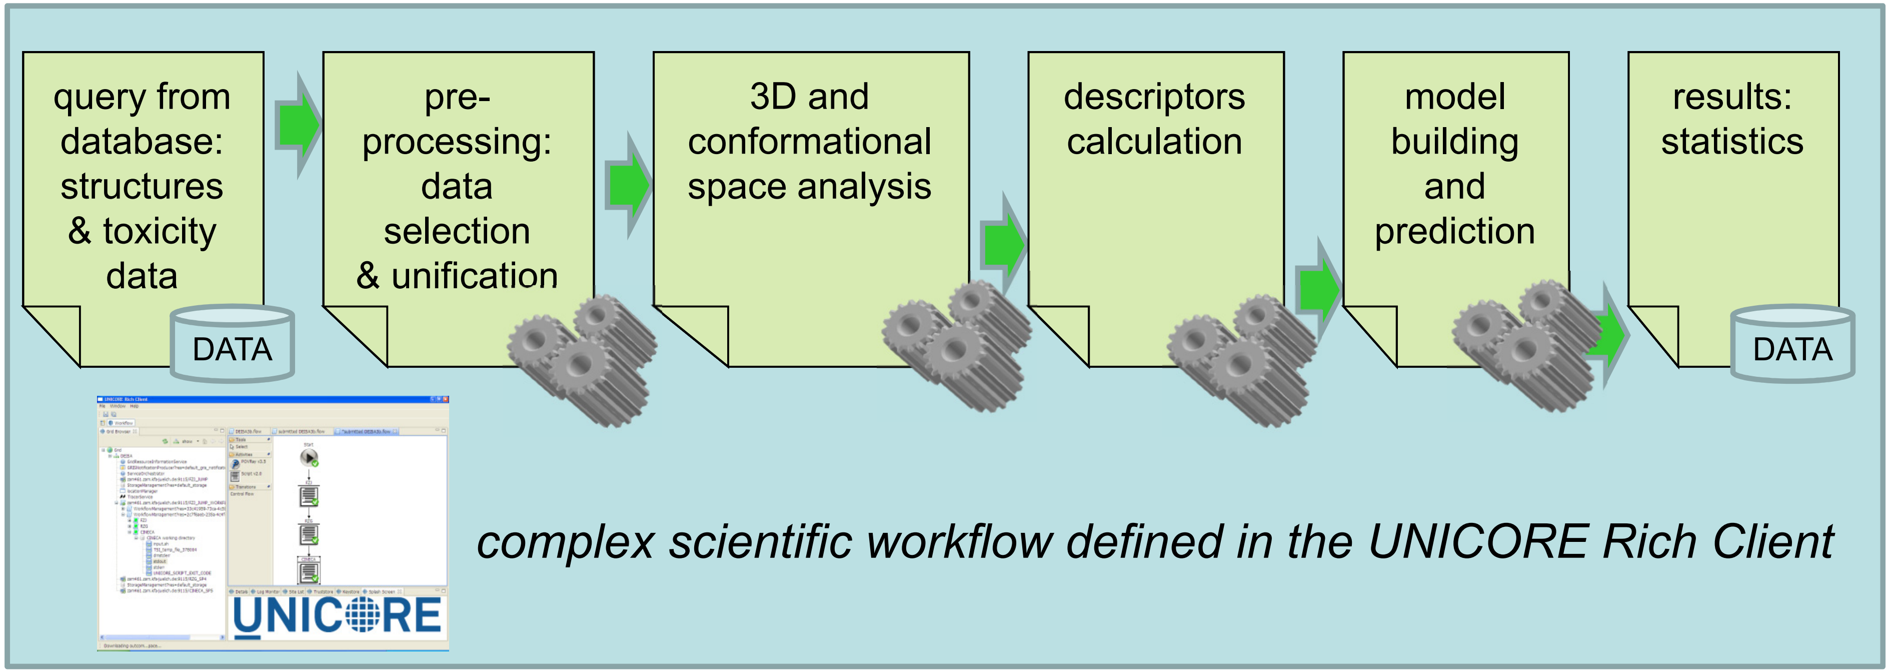

Grid job submission using workflow capabilities in UNICORE

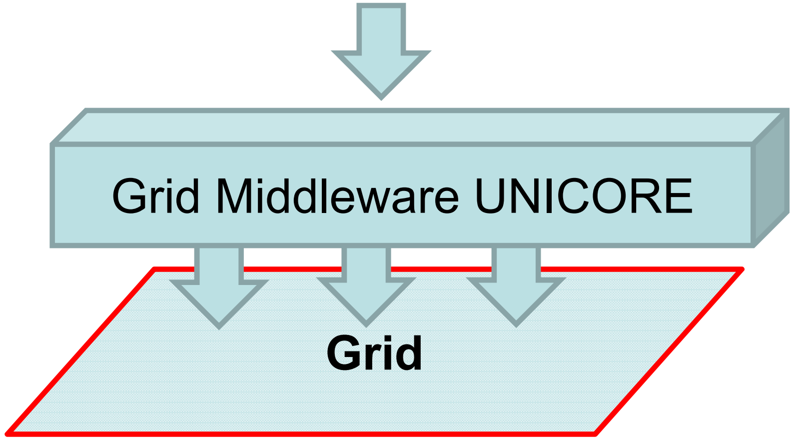

[15] Chemomentum

Chemomentum

Approach IV. Interactive Access

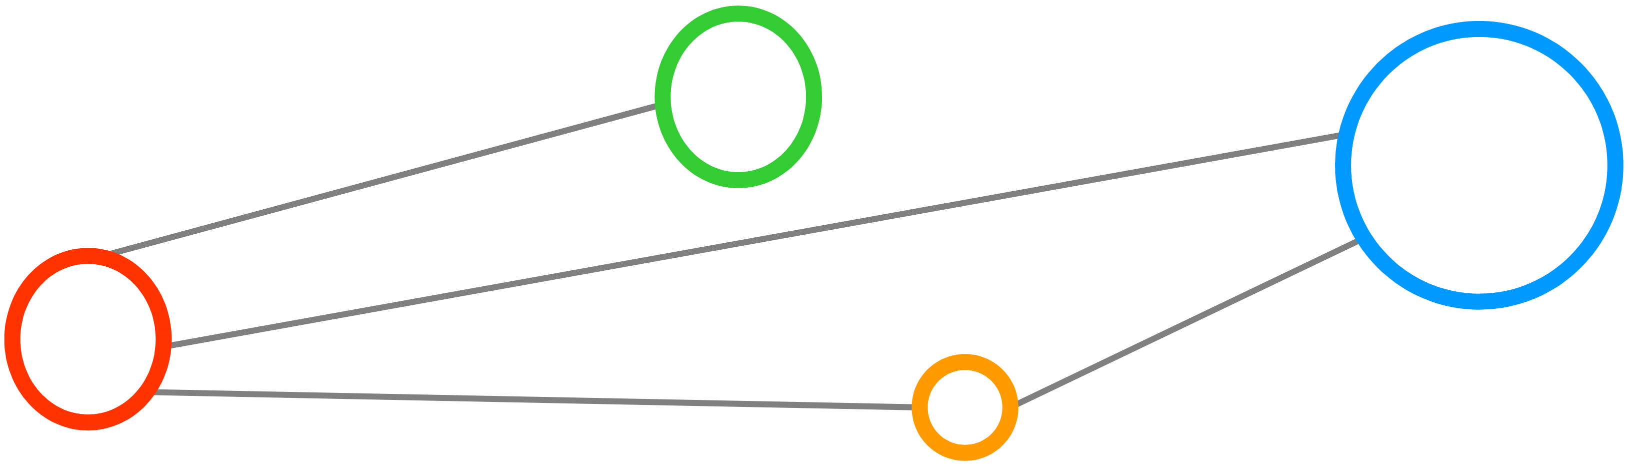

Approach Interactive Access

- Several e-scientists require direct access to the job execution directory to check intermediate results → SSH
- Some e-scientists use computational steering techniques and feedback visualization for their HPC-driven Grid applications

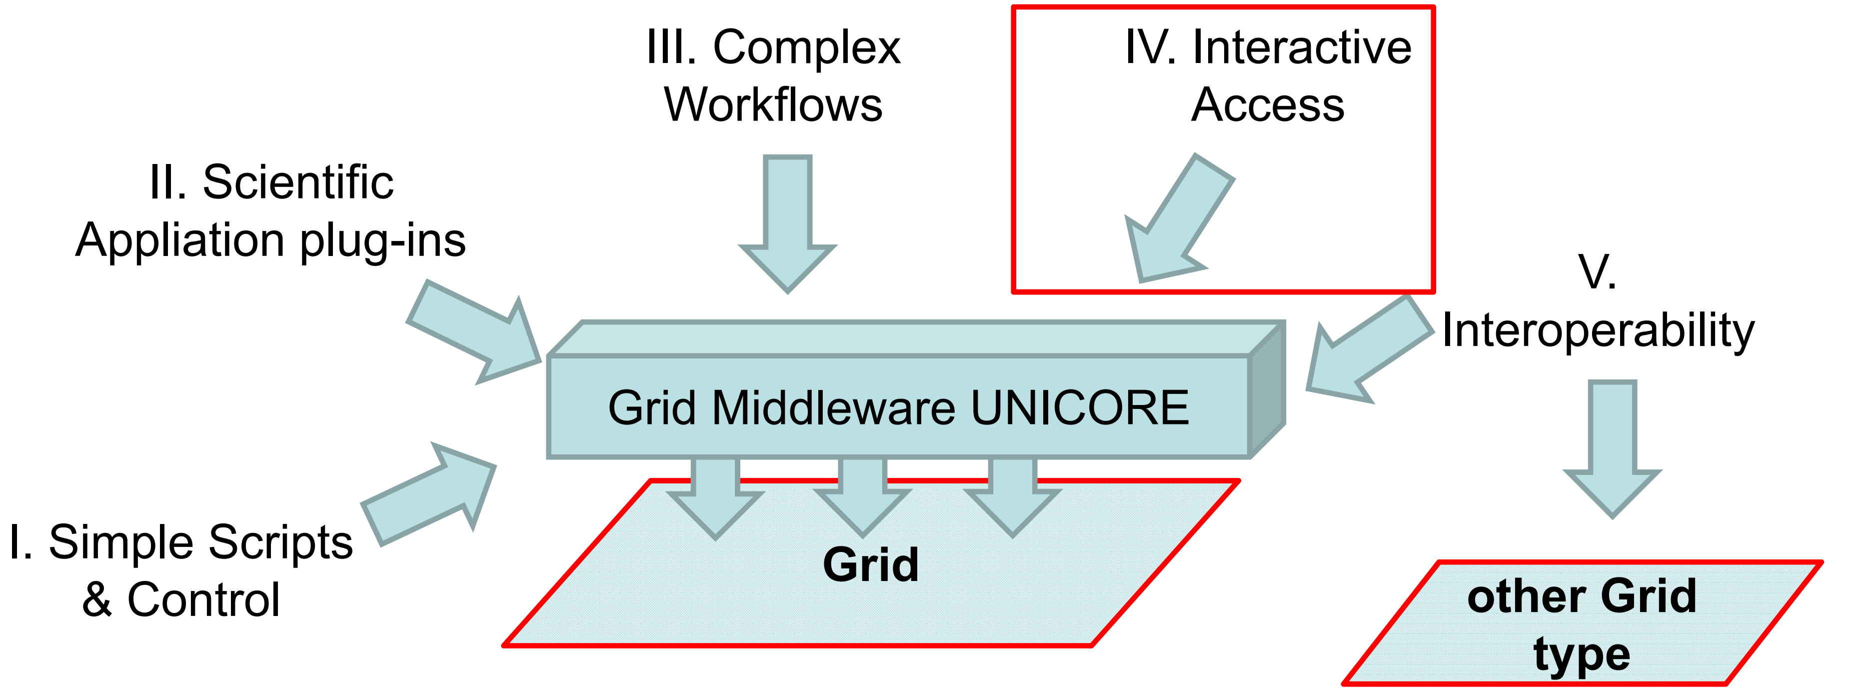

Interactive access with SSH

- SSH with single-sign on – benefit no password needed to get access to numerous sites, access to the job directory

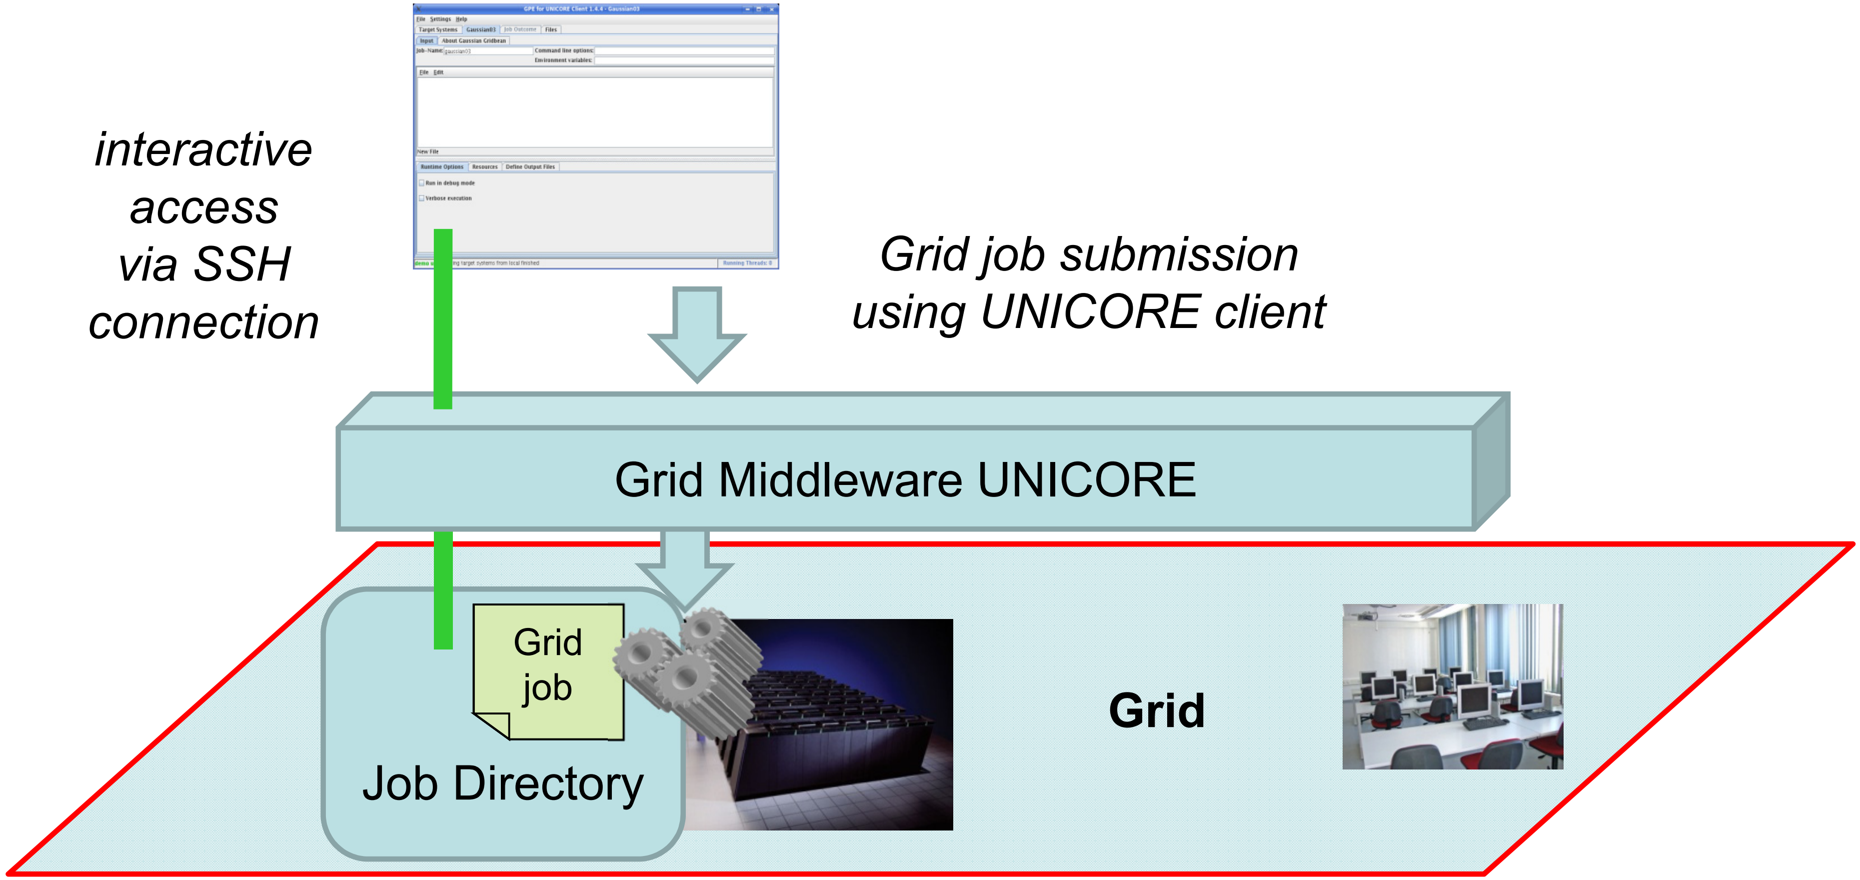

Interactive access with computational steering

- Computational steering to change parameters on the fly

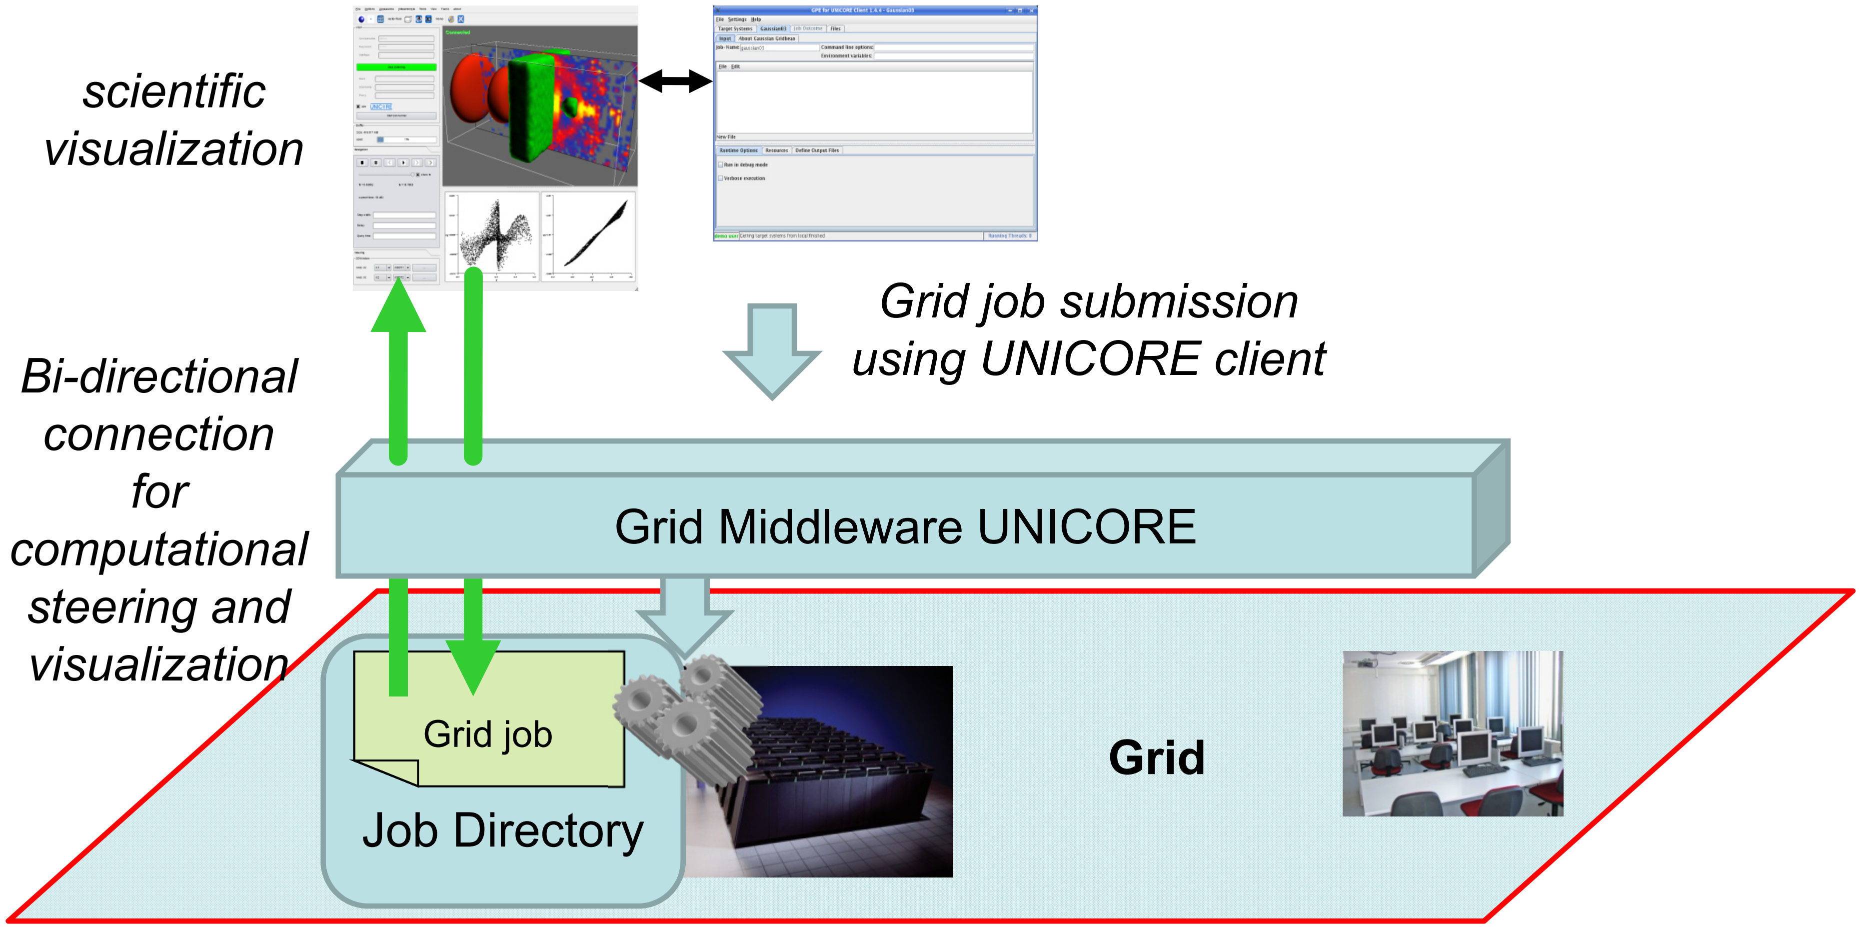

Example: n-body problems in e-science

- Xnbody: Light-weight 3D visualization application based on VTK
- Interactive display of particle and volumetric data
- Collaborative mode: simultaneous display at several sites via COVS

[18] Xnbody

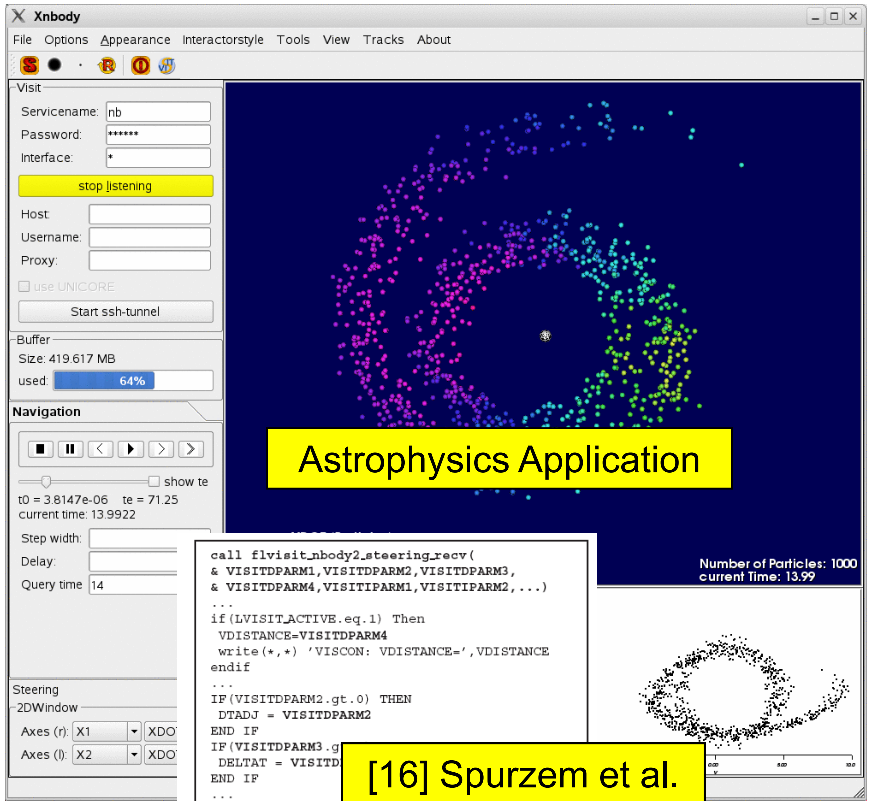

[16] Spurzem et al.

```
call flvisit_nbody2_steering_recv(
& VISITDPARM1,VISITDPARM2,VISITDPARM3,
& VISITDPARM4,VISITIPARM1,VISITIPARM2,...)
...
if(LVISIT_ACTIVE.eq.1) Then
 VDISTANCE=VISITDPARM4
 write(*,*) 'VISCON: VDISTANCE=',VDISTANCE
endif
...
IF(VISITDPARM2.gt.0) THEN
 DTADJ = VISITDPARM2
END IF
IF(VISITDPARM3.g
 DELTAT = VISITD
END IF
...
CALL MPI_BCAST(DTADJ,1,MPI_DOUBLE_PRECISION,
    0,& MPI_COMM_WORLD,ierr)
CALL MPI_BCAST(VISITDPARM3,1,
    MPI_DOUBLE_PRECISION,0,& MPI_COMM_WORLD,...)
```

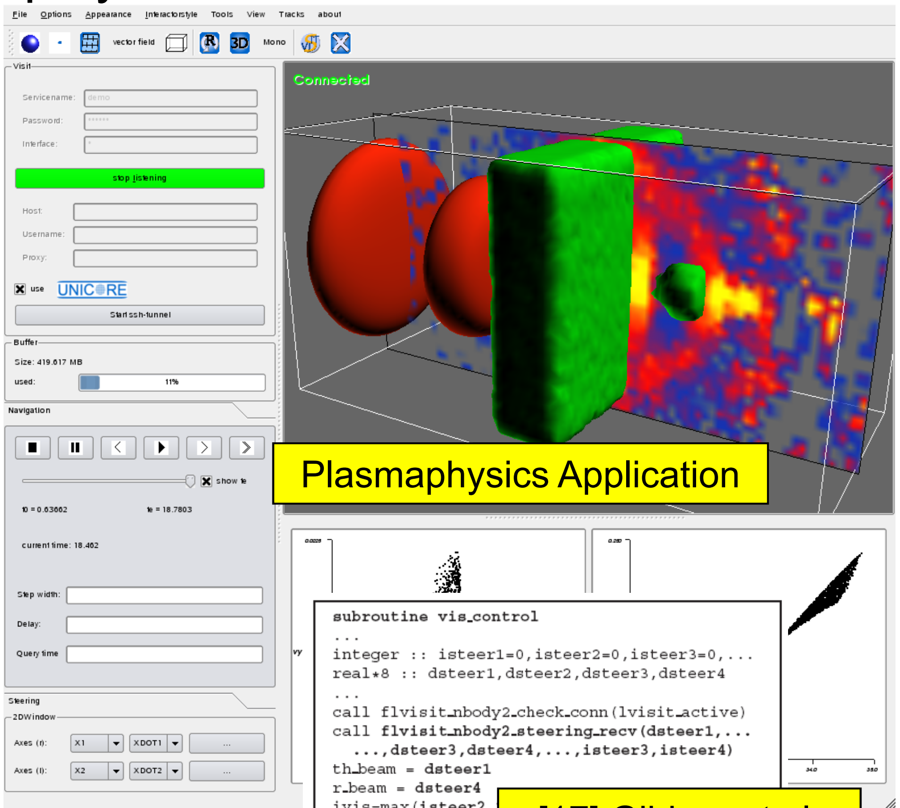

[17] Gibbon et al.

```
subroutine vis_control
...
integer :: isteer1=0,isteer2=0,isteer3=0,...
real*8 :: dsteer1,dsteer2,dsteer3,dsteer4
...
call flvisit_nbody2_check_conn(lvisit_active)
call flvisit_nbody2_steering_recv(dsteer1,...
   ...,dsteer3,dsteer4,...,isteer3,isteer4)
th_beam = dsteer1
r_beam = dsteer4
ivis=max(isteer2,
ivis_fields=max(is
...
call MPI_BCAST( th_beam, 1, MPI_REAL, 0,
    MPI_COMM_WORLD,ierr)
...
```


FORSCHUNGSZENTRUM

Grid Infrastructure enables Collaboration with COVS

- Collaborative Online Visualization and Computational Steering (COVS)

[20] Riedel et al.

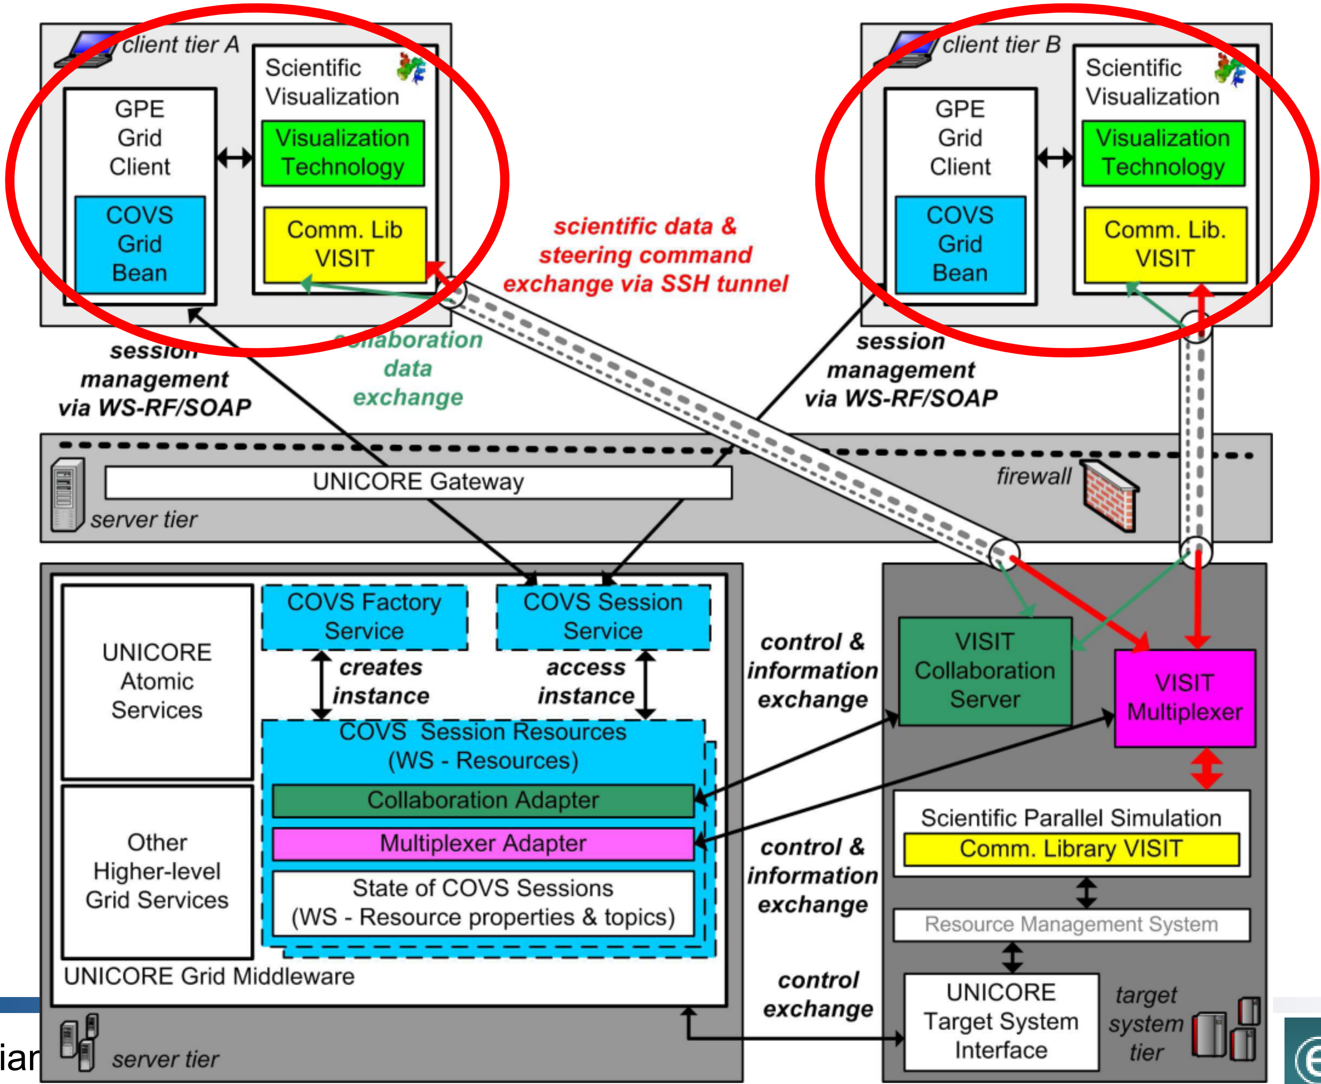

Approach V. Interoperability

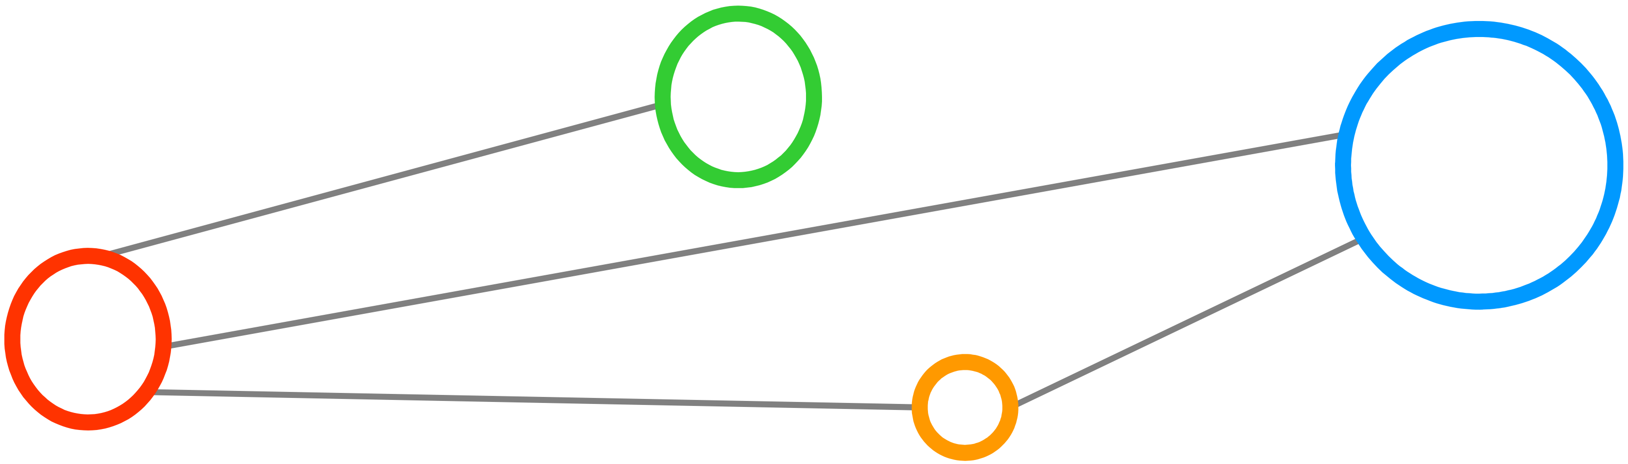

Approach Interoperability

- Increasing complexity of e-science applications that embrace multiple physical models (i.e. multi-physics) & larger scale
  - Creating a steadily growing demand of compute power
  - Demand for a ‘United Federation of world-wide Grids’

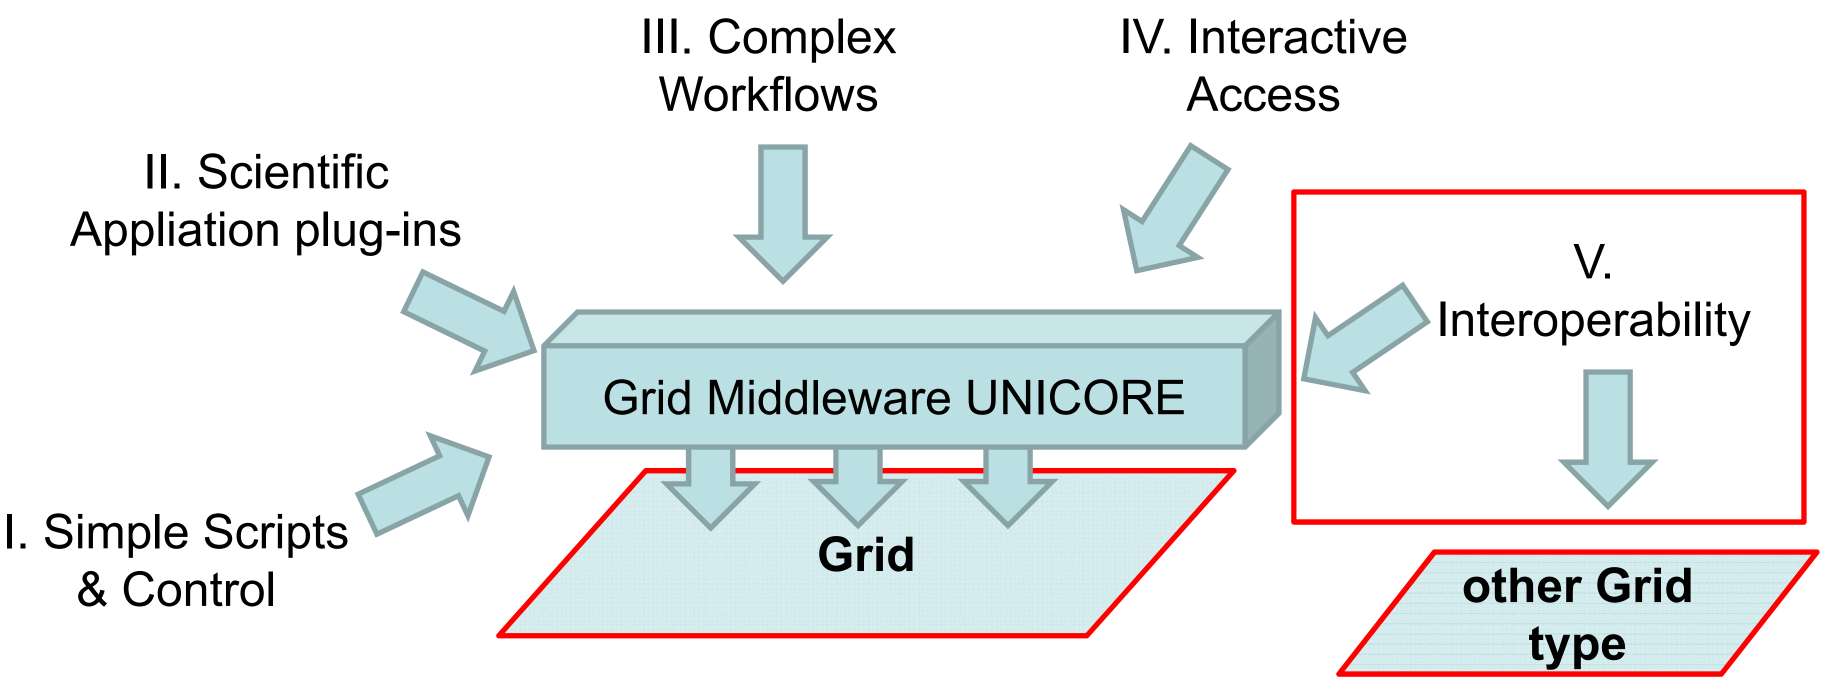

Interoperability between HTC and HPC

- Utilizing different kinds of Grids for one scientific use case
- High Throughput Computing (HTC) oriented Infrastructures

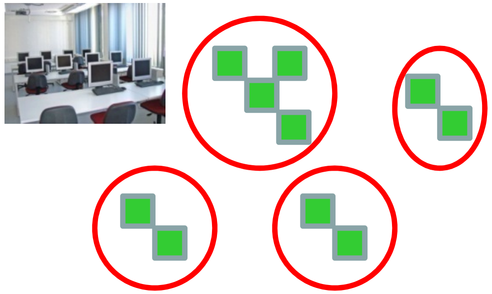

cpus/cores not well interconnected and physically distributed, higher availability, relatively cheap compared to HPC resources

- High Performance Computing (HPC) driven infrastructures

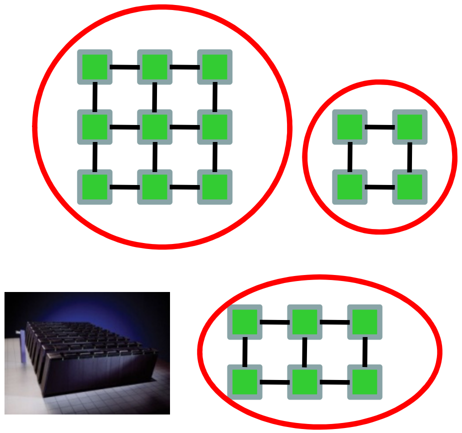

cpus/cores well interconnected, physically often concentrated in centres,'overbooked' good interconnection is costly

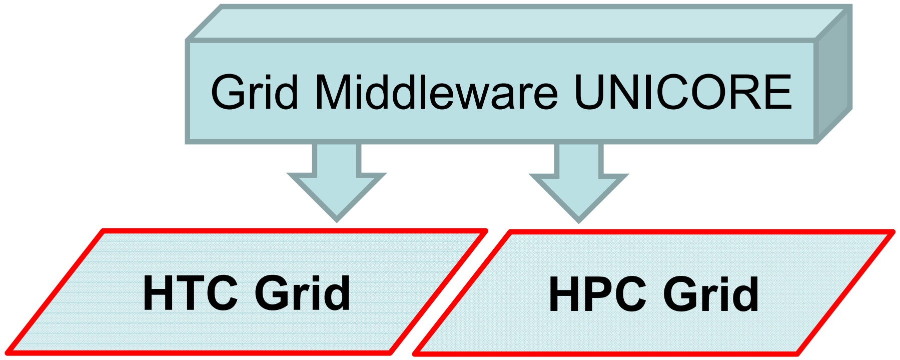

Grid Infrastructure Islands in Europe

- DEISA Grid (Supercomputing/HPC community)
  - Non WS-based UNICORE 5: Proprietary jobs (AJO/UPL)
  - No Virtual Organization Membership Service (VOMS), Full X.509
  - Suitable for massively parallel scientific jobs (MPI, much interactions)
- EGEE Grid (mainly HEP community + others)
  - Non WS-based gLite: Proprietary jobs (JDL)
  - Proxy-based X.509 security, but proprietary VOMS support
  - Suitable for embarrassingly parallel scientific jobs (less interactions)
- Both Grids are currently not technical interoperable
  - Scientists cannot use one middleware to access both
  - **Both Middleware's had less adoption of open standards**
  - **More recently a move towards standards integrations**

Distributed European Infrastructure for Supercomputing Applications

OMII-Europe [24]

Example: The WISDOM Project(s)

- Wide In Silico Docking On Malaria (WISDOM)
  - Developing new drugs for neglected and emerging diseases with a particular focus on malaria
  - Accelerated Research & Development for emerging and neglected diseases
  - Reduced Research & Development costs

WISDOM [21]

WISDOM
Initiative for grid-enabled drug discovery against neglected and emergent diseases

- Three large calculations:
  - WISDOM-I (Summer 2005)
  - Avian Flu (Spring 2006)
  - WISDOM-II (Autumn 2006)

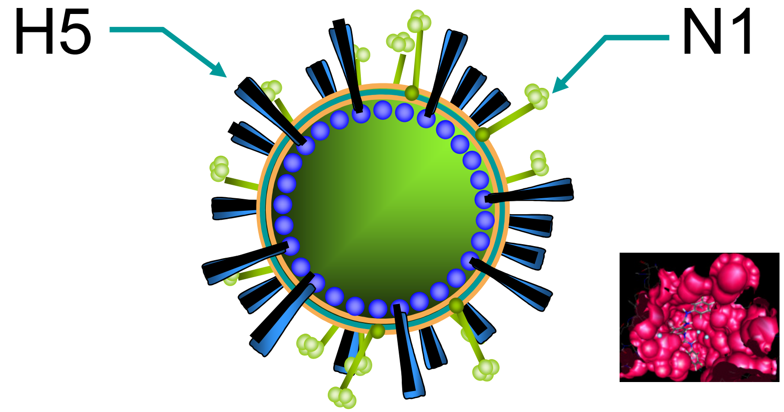

[23] N. Jacq

In silico drug discovery

A human grid initiative.

- **Problem: development of a drug takes 12 to 15 years and costs approximately 800 million dollars**

[22] N. Jacq

Target discovery | Lead discovery

Target Identification → Target Validation → Lead Identification → Lead Optimization → **Clinical Phases (I-III)**

in vitro experimentation

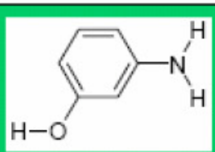

in silico experimentation

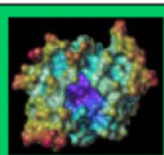

In vitro: ‚experiment in a controlled environment outside of a living organism’

In silico: ‚performed in the computer‘

Example: e-Science in WISDOM (1)

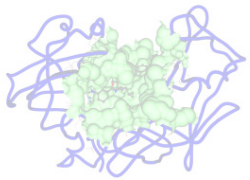

- WISDOM in the context of DEISA & EGEE
  - WISDOM also aims at developing new drugs for Malaria
  - WISDOM uses EGEE for large scale in silico docking
    - Comp. method for prediction of whether one molecule will bind to another
- Using AutoDock and FlexX provided via gLite in EGEE
  - Output is a list of best chemical compounds (potential drugs)
    - That is not the final solution, only a potential list of drugs
- Refine best compound list via molecular dynamics (MD)
  - Fast MD computations use highly scalable AMBER in DEISA
    - AMBER (Assisted Model Building with Energy Refinement) , version 9
- Goal: Accelerate drug discovery using EGEE + DEISA

[26] AutoDock [27] FlexX [13] AMBER

Example: e-Science in WISDOM (2)

- FleXX and Autodock software are being used by WISDOM e-Scientists

[22] N. Jacq

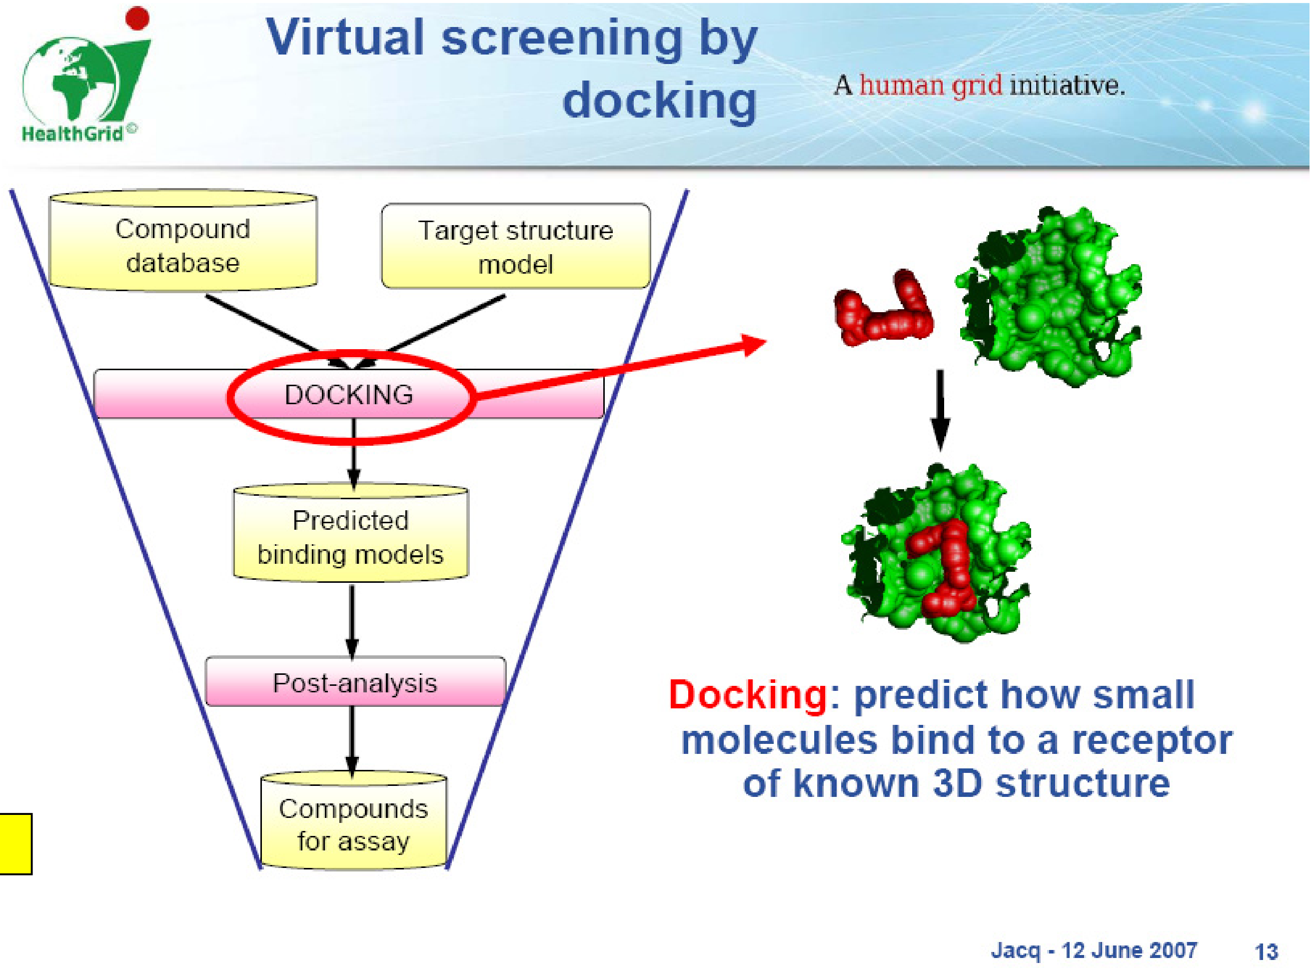

Example: e-Science in WISDOM (3)

- Dataflow and workflow in a virtual screening
  - Docking with FleXX and AutoDock in EGEE
  - Molecular Dynamics (MD) simulation with AMBER in DEISA

[22] N. Jacq

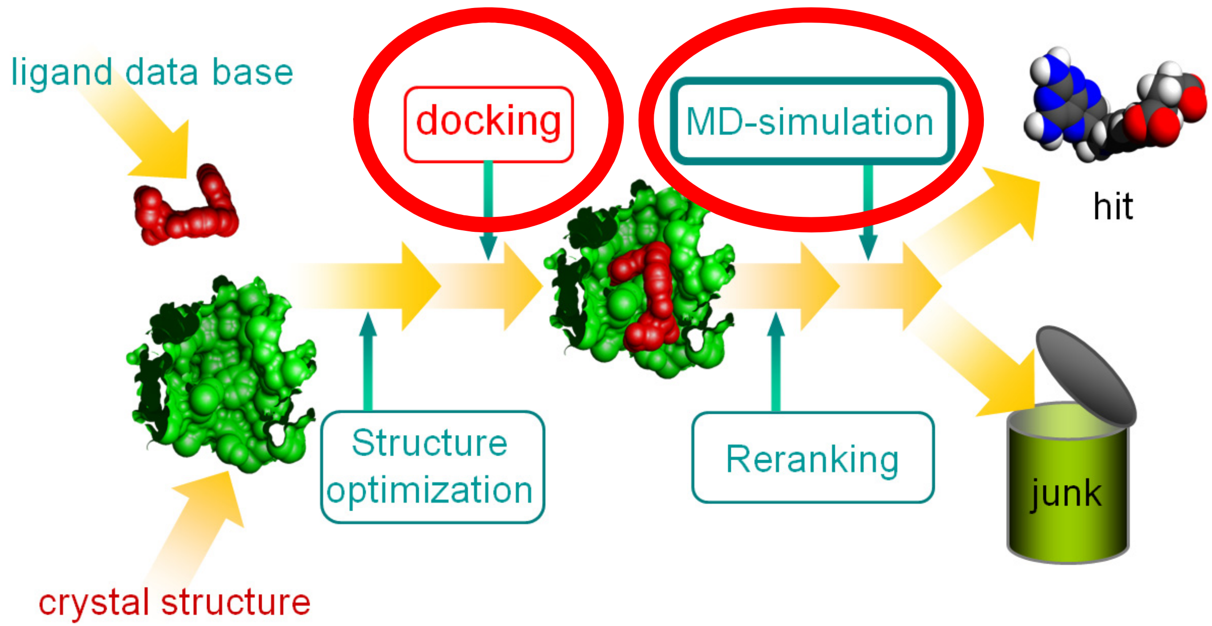

Molecular dynamics (MD): a computer simulation in which atoms and molecules interact for a period of time under known laws of physics providing a view of the motion of the atoms

Example: e-science in WISDOM (4)

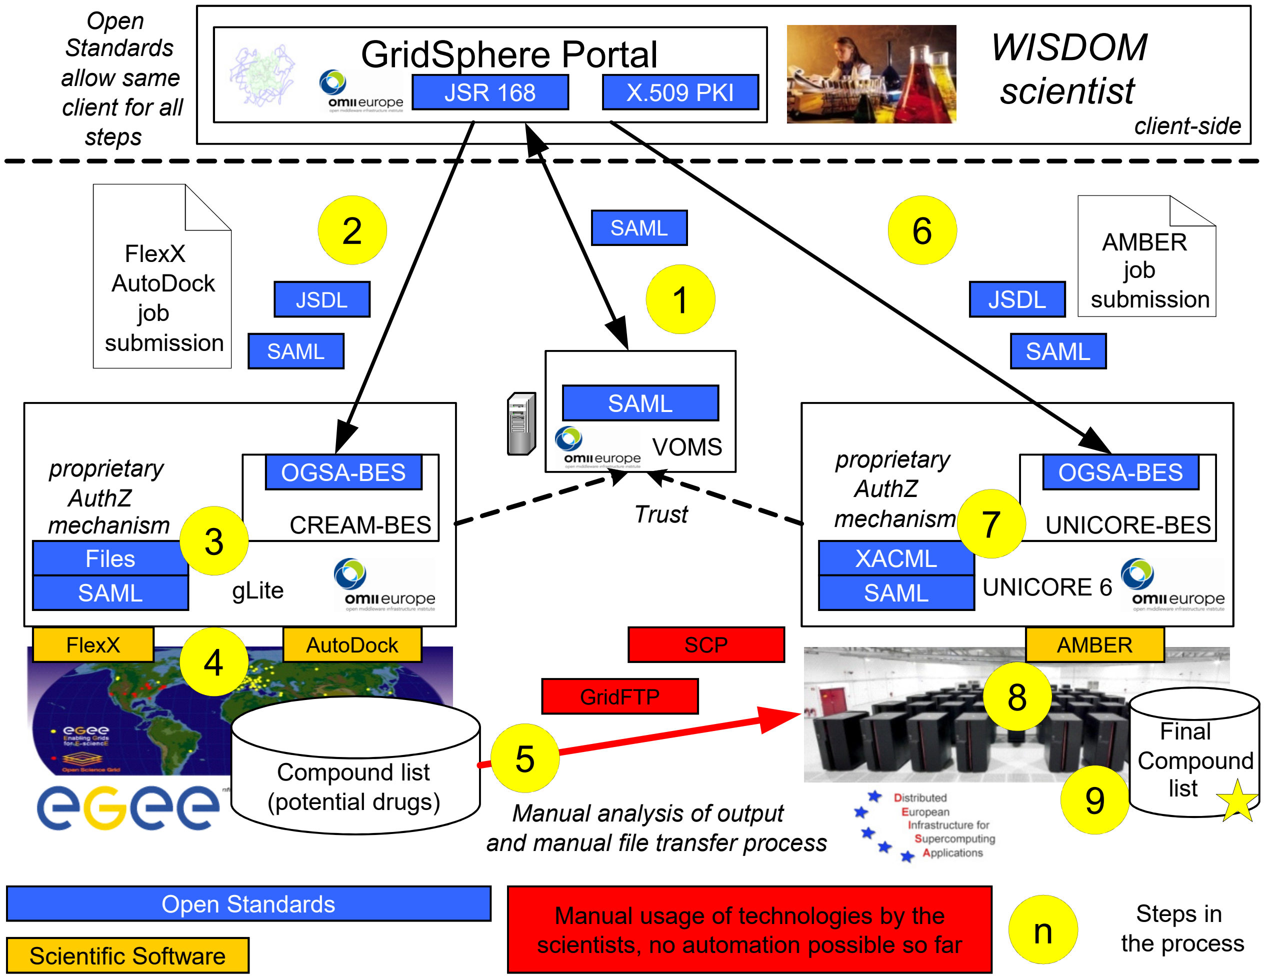

work in progress

Lessons Learned & Conclusions

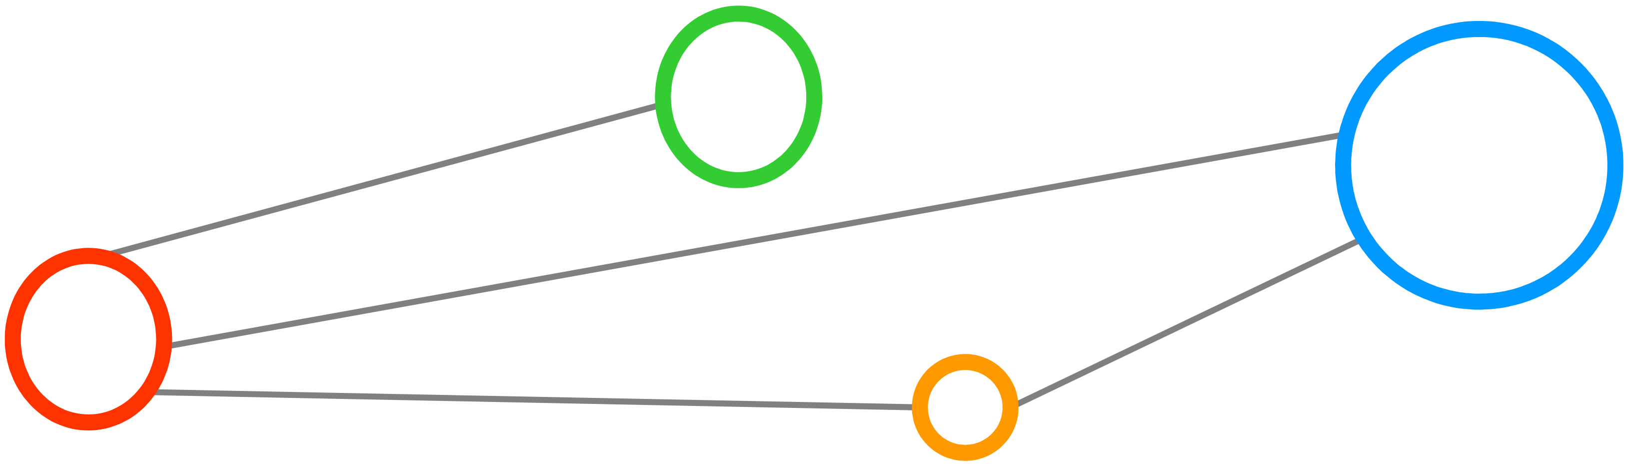

Some Benefits using Grids in e-Science (1)

- Central availability
  - Scientific jobs are kept on server-side
  - Using Desktop PC (office), Laptop (mobile), or a PC at a conference makes no difference
- Monitoring area of jobs very helpful
  - Learn about the actual Grid job status
  - Foundation for copy & paste to fastly create similiar jobs
- Control functionality often used (e.g. LOOP DO-N)
  - Submit scientific applications automatically as defined
  - Allow more possibilities how applications are executed
- ‚Re-submit' if job queue down (not cope with this problem)
  - UNICORE keep track of status of batch queuing system

Some Benefits using Grids in e-Science (2)

- High extensibility of UNICORE graphical clients
  - Scientific domain-specific application plug-ins improve the efficiency of e-Scientists, portals ease the access and use of Grids
  - Works well in daily work with commonly used applications
  - Idea: Commonly used applications can be specifically supported via dedicated plug-ins
  - Such plug-ins seamlessly offering all kinds of parameters/inputs available for a scientific application
- Workflow Support in UNICORE
  - Numerous simple operations are not carried out manually
  - Workflows allow to combine applications in different ways
  - Every step can be fully documented (e.g. important in REACH)
  - Reproducability of results (storing job definition)
  - Even scientific domain-specific language extensions in UNICORE 6

Some Benefits using Grids in e-Science (3)

- Interactive Access often necessary
  - Check intermediate steps in computational process
  - Influence simulation behaviour with computational steering by changing parameters on the fly allows much flexibility (‚events‘)
  - Full feedback visualization with COVS framework easy to use

- Interoperability between next generation infrastructures
  - Multi-physics/Multi-scale creating growing demand for compute power
  - Option to satisfy increasing e-science application demands
  - Using different kinds of infrastructures together (HTC and HPC)
  - Harness a ‚united federation of world-wide Grids‘
  - World-wide efforts in the Grid Interoperation Now (GIN) Community Group of the Open Grid Forum (OGF)

[30] OGF GIN

ANNOUNCEMENT:
OGF Production Grid Infrastructure (PGI) WG

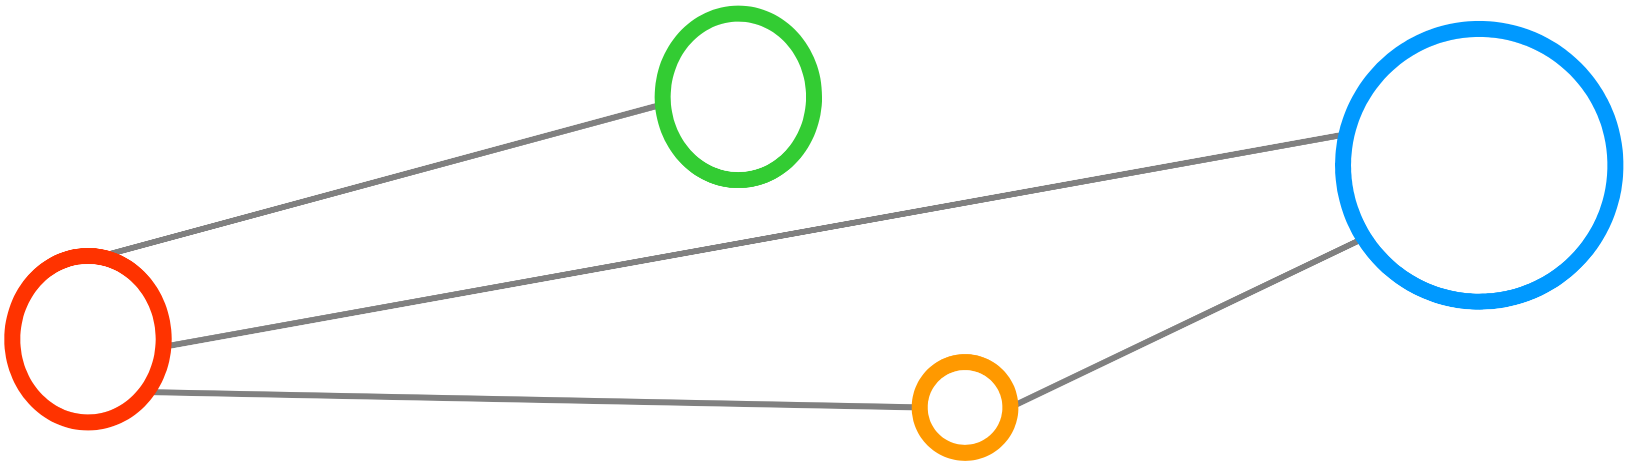

New OGF Production Grid Infrastructure WG

- OGF Grid Interoperation Now (GIN) Community Group
  - e-Science applications using standards deployed in Grids
  - Real practical applications that require interoperability
- OGF Production Grid Infrastructure (PGI) Working Group
  - Profile/Standardize a well-defined set of open standards (i.e. OGSA-BES, SRM, GridFTP, JSDL,
  - Focus on missing links between emerging open standards
  - Take production experience from GIN into account
- Get Involved!!!
  - Email: pgi-wg@ogf.org
  - Page: http://forge.ogf.org/projects/pgi-wg

References

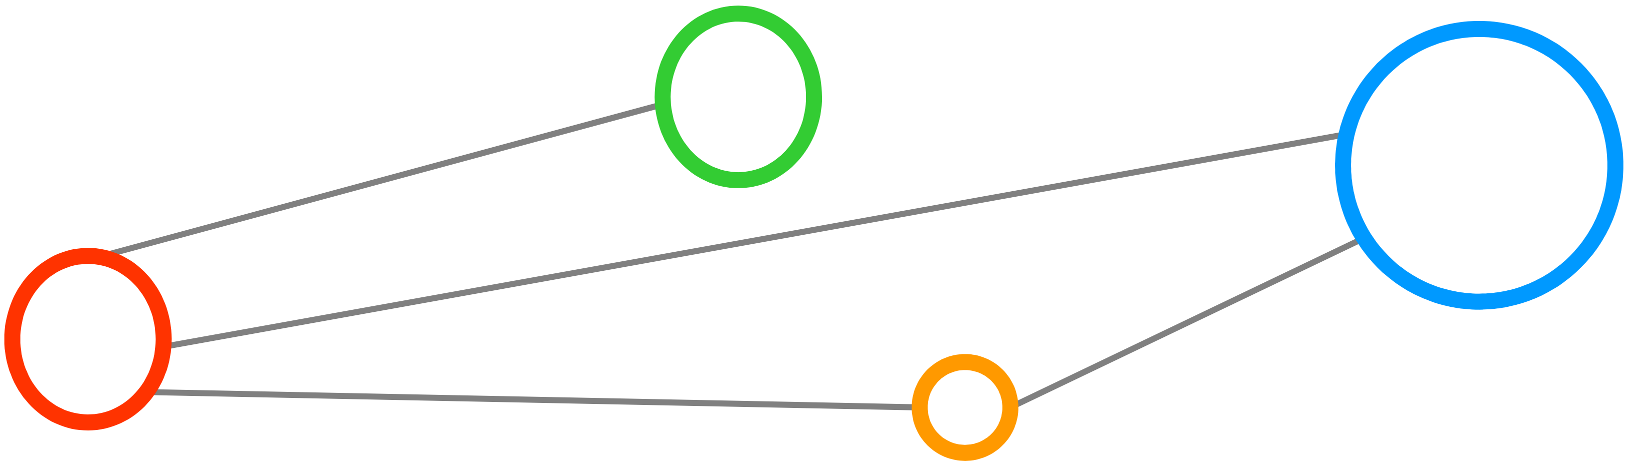

References (1)

[1] Lippert et al., ‘IBM delivers Europe's biggest supercomputer’,
Online: http://news.zdnet.co.uk/hardware/0,1000000091,39146680,00.htm

[2] Message Passing Interface (MPI), http://www-unix.mcs.anl.gov/mpi/

[3] John Taylor, ‘The Definition of e-Science’,
Online: http://www.lesc.ic.ac.uk/admin/escience.html

[4] Distributed European Infrastructure for Supercomputing Applications (DEISA),
Online: http://www.deisa.org

[5] Enabling Grids for e-Science (EGEE), Online: http://public.eu-egee.org/

[6] TeraGrid, Online: http://www.teragrid.org

[7] D-Grid, Online: http://www.d-grid.de/

[8] NGS, Online: http://www.grid-support.ac.uk/

[9] PRACE, Online: http://www.prace-project.eu/

[10] H. D. M. Moore, K. Dvorakova, N. Jenkins and W. G Breed. Exceptional sperm coorperation in the wood mouse. Nature 418, 174 (2002).

[11] Elgeti et al., J. Elgeti and G. Gompper. Hydrodynamics of Active Mesoscopic Systems. NIC Series 39, 53 (2008)

[12] Gaussian Plug-in, Online:
http://sourceforge.net/project/showfiles.php?group_id=102081&package_id=276123

References (2)

[13] AMBER, Online: http://amber.scripps.edu/

[14] Casalegno et al., Abstracts of QSAR-related Publications: Multivariate Analysis,QSAR & Combinatorial Science, Volume 26, No. 5, pages 724-725, 2007,WILEY-VCH Verlag

[15] Chemomentum, Online: http://www.chemomentum.org/c9m

[16] R. Spurzem et al. Nbody6 features. 2003. ftp://ftp.ari.uniheidelberg.de/pub/staff/spurzem/nb6mpi/nbdoc.tar.gz.

[17] Paul Gibbon et al. *Performance Analysis and Visualization of the N-body tree code PEPC on massively parallel computers*. In Proc. of the ParCo, Malaga, Spain, 2005.

[18] Xnbody Visualization, Online: www.fz-juelich.de/zam/xnbody

[19] DECI Projects 2008, Online: http://www.deisa.eu/science/deci/projects2007-2008

[20] M. Riedel, Th. Eickermann, W. Frings , S. Dominiczak, D. Mallmann, Th. Düssel, A. Streit, P. Gibbon, F. Wolf, W. Schiffmann, and Th. Lippert *Design and Evaluation of a Collaborative Online Visualization and Steering Framework Implementation for Computational Grids,* In Proc. of the 8th IEEE/ACM International Conference on Grid Computing (Grid 2007), Austin

References (3)

[21] WISDOM Project, http://wisdom.eu-egee.fr/

[22] Dr Nicolas Jacq, International Symp.on Grids for Science and Business 12 June 2007, *http://events.ibbt.be/grid2007/pdf/Jacq.pdf?PHPSESSID=51537714797e1362c6c1e33f554e2dd5*

[23] Dr Nicolas Jacq, World-wide in silico drug discovery against neglected and emerging diseases on grid infrastructures, GridPP, 2007

[24] Open Middleware Infrastructure Institute for Europe (OMII-Europe), Online: http://www.omii-europe.org

[25] Gaussian 03 Web-site, Online: http://www.gaussian.com/

[26] AutoDock, Online: http://autodock.scripps.edu/

[27] FlexX, Online: http://www.biosolveit.de/flexx/

[28] Open Science Grid, Online: http://www.opensciencegrid.org/

[29] DEISA DECI, Online: http://www.deisa.eu/science/deci/

[30] Open Grid Forum (OGF) – Grid Interoperation Now (GIN) Community Group, Online: http://forge.ogf.org/sf/projects/gin

Acknowledgements

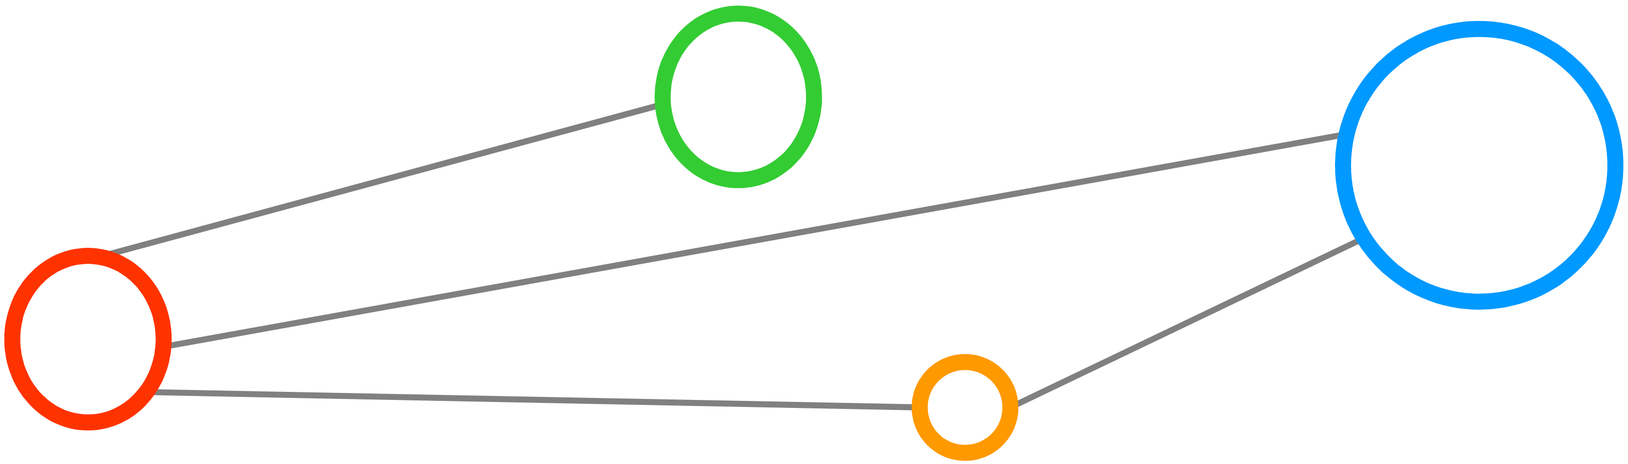

Morris: Acknowledgements

- Thanks for the Participation in the IGIIW Workshop!
- Morris Travel and Participation in OGF is funded by…
- Distributed European Infrastructure for Supercomputing Applications (DEISA)
  - DEISA2 is funded by the European Commission in FP7 under grant agreement RI-222919
- Jülich Supercomputing Centre (JSC) of Forschungszentrum Jülich (FZJ) in the HELMHOLTZ association