Clusters and Grids in Chemistry and Reinsurance Calculations

Dr. Wibke Sudholt

Project Leader, Grid Computing Team, Baldridge Group,
Institute of Organic Chemistry, University of Zurich, Switzerland

wibke@oci.uzh.ch

http://ocikbws.uzh.ch/grid/

Overview

"Average user" utilization of clusters and grids for scientific and business applications

- Infrastructure building
- Software porting and application
- User interfaces
- Focus on computation

Two example domains

- Quantum chemical calculations
- Reinsurance natural catastrophe calculations

Presentation structure

- Cluster and grid infrastructures
- Example background
- User / application requirements
- Domain-specific issues and solutions
- Conclusions for HPC

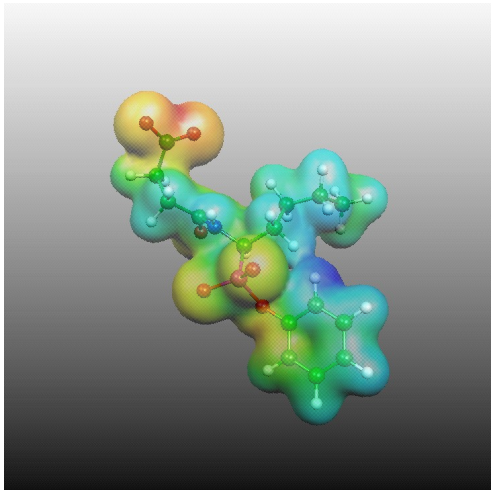

http://www.swissre.com/pws/research%20publications/risk%20and%20expertise/risk%20perception/natural%20catastrophes%20and%20reinsurance.html

Infrastructure Characteristics

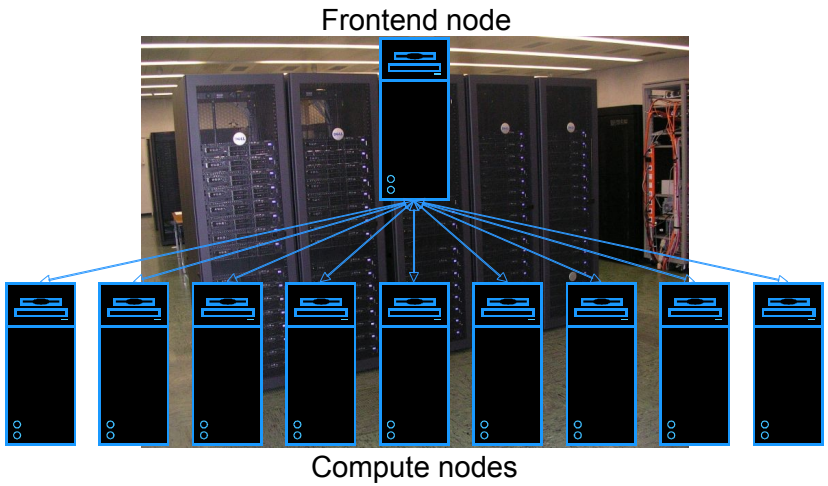

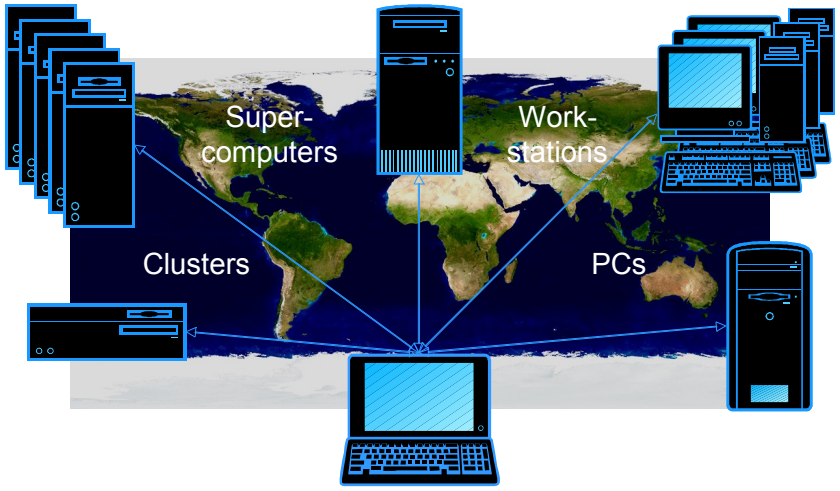

Clusters / supercomputers

- Homogenous hardware and software
- Tightly-coupled machines
- Usually reliable high-speed network
- One central location
- One administrative domain
- Usually secure
- Resource management system (SGE, PBS, etc.)
- Fast communication between parallel jobs
- Focus on high performance/capability computing

Grids

- Heterogeneous hardware and software
- Loosely-coupled machines
- Potentially unreliable low-speed network
- Several distributed locations
- Several administrative domains
- Potentially insecure
- Grid middleware (Globus, UNICORE, etc.)
- Best for embarrassingly parallel jobs
- Focus on high throughput/capacity computing

Group Infrastructure

Local Linux clusters

- Two small grid project Rocks / SuSE clusters
  - 10+ 2-CPU Xeon nodes
  - Gigabit interconnection
- One medium group production Rocks cluster
  - 100+ 2-CPU Xeon nodes
  - Gigabit interconnection
- One large university production cluster maintained by the UZH IT Services
  - 384 2-CPU Opteron nodes (30 group nodes)
  - Myrinet / Gigabit interconnection

Grid infrastructure projects

- PRAGMA (Pacific Rim Applications and Grid Middleware Assembly)
  - http://www.pragma-grid.net/
  - 29 institutions
  - Globus, Nimrod, Gfarm, etc. middleware
- SEPAC (Southern European Partnership for Advanced Computing)
  - http://www.sepac-grid.org/
  - 8 institutions
  - Globus, iGrid, GRB, etc. middleware
- SwiNG (Swiss National Grid Association)
  - http://www.swing-grid.ch/
  - 19 institutions
  - ARC, gLite, XtremWeb-CH, etc. middleware

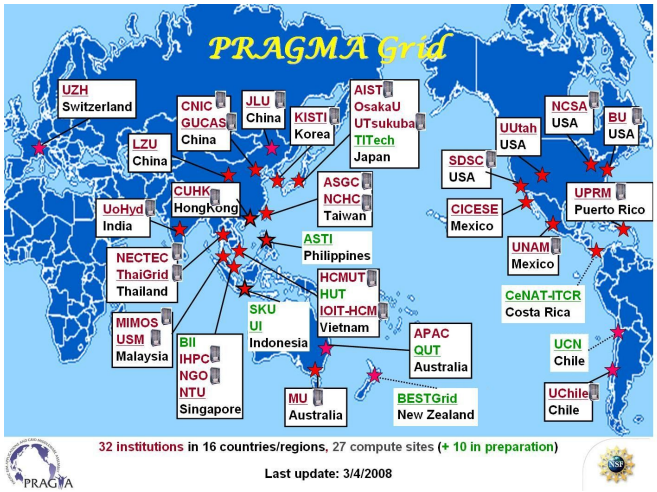

Quantum Chemical Calculations

Quantum Chemical Calculations

Approximate solutions to the Schrödinger equation of quantum mechanics:
$\mathbf{H}\,\Psi = E\,\Psi$

Many different methods
- Semiempirical: MNDO, AM1, PM3, etc.
- “Ab initio”: HF, DFT, MPn, CC, MCSCF, CI, etc.

Several legacy software packages
- Academic: GAMESS, MOLCAS, NWChem, etc.
- Commercial: Gaussian, Spartan, etc.

Used for the calculation of properties of molecular systems
- 3D structures, energies, electronic properties, spectra, reaction mechanisms, etc.
- Solvation, biological environments, materials, etc.

Advantages
- Cheaper, faster, and safer than laboratory experiments
- Deeper understanding of processes at the molecular level

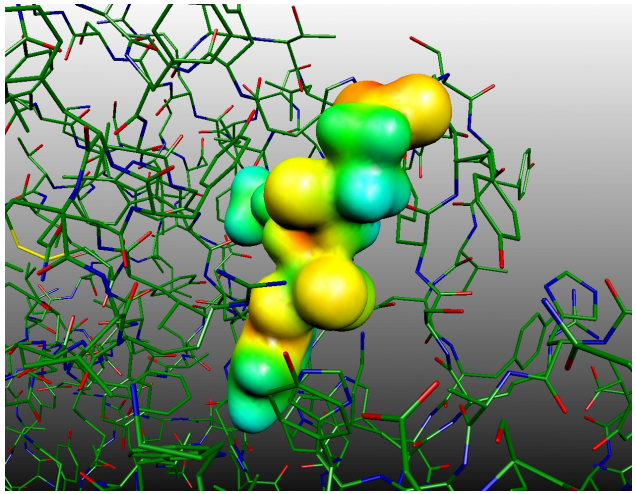

Challenges
- Size and dynamics of the systems
- Accuracy of the methods:
  Basis set, electron correlation, etc.
- Computationally expensive and data intensive:
  Two-electron integrals, conventional or direct SCF, etc.

User Requirements

Have all scientific methods reasonable for the molecular system of interest available at your fingertips, and invent new approaches

- Quantum chemical algorithms
- Computational science software development

Get calculations performed easily, quickly, and smoothly, with transparent access to the computational resources

- HPC systems, parallelization
- Performance, scalability, stability
- Grid infrastructures and middleware
- Support and education

Have advanced interfaces to and between different application software and data available to choose from

- CLIs, GUIs, APIs, etc.
- Data format standards
- Visualization
- Parametric modeling
- Workflow systems
- Interdisciplinary collaboration with biology, pharmaceutics, material science, physics, etc.

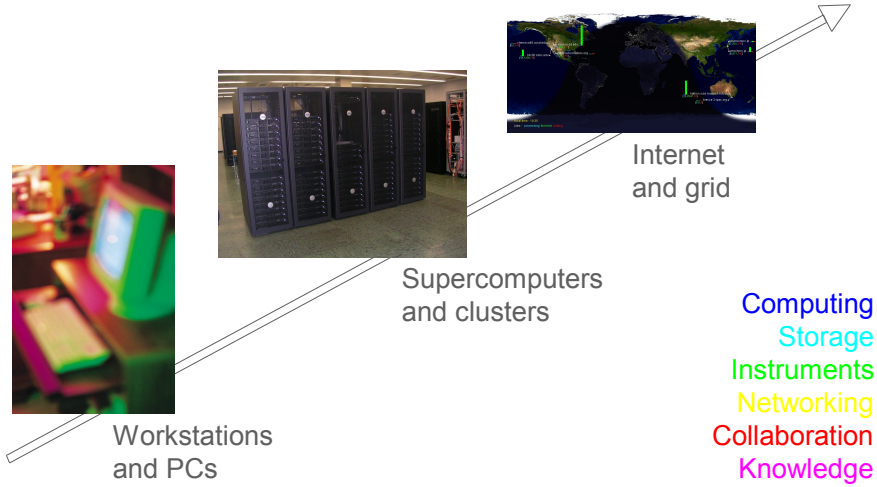

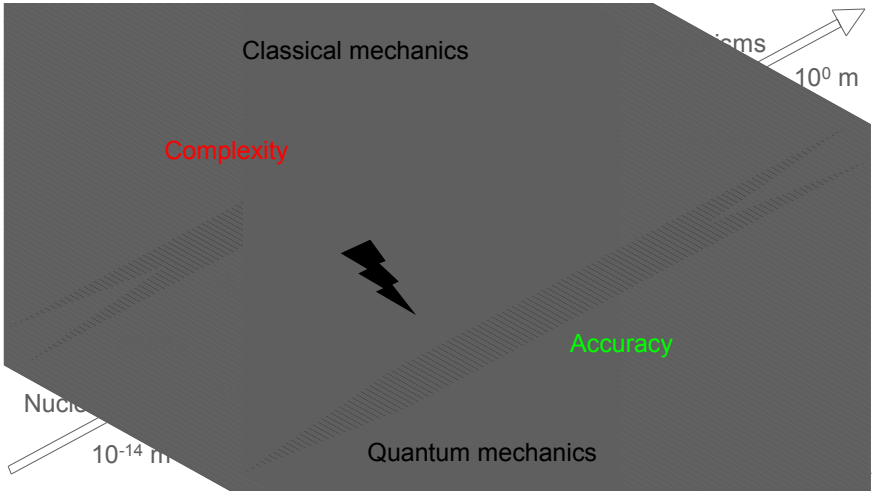

GAMESS Software

http://www.msg.chem.iastate.edu/gamess/

General Atomic and Molecular Electronic Structure System

Ab initio quantum chemistry program package

Large world-wide user and developer base

Adapted to many different hardware platforms / OS

Early versions date back to 1984

Mainly Fortran 77 and C code as well as shell scripts

Distributed Data Interface (DDI) message passing layer for parallel execution on top of sockets, MPI, or shared memory

SCFTYP=	RHF	ROHF	UHF	GVB	MCSCF
Energy	CDpF	CDp	CDp	CDp	CDpF
analytic gradient	CDpF	CDp	CDp	CDp	CDpF
numerical Hessian	CDpF	CDp	CDp	CDp	CDp
analytic Hessian	CDp	CDp	-	CDp	CDp
Mp2 energy	CDpF	CDp	CDp	-	CDp
Mp2 gradient	CDpF	Dp	CDp	-	-
CI energy	CDp	CDp	-	CDp	CDp
CI gradient	CD	-	-	-	-
CC energy	CDpF	-	-	-	-
EOM energy	CD	-	-	-	-
DFT energy	CDpF	CDp	CDp	-	-
DFT gradient	CDpF	CDp	CDp	-	-
MOpAC energy	yes	yes	yes	yes	-
MOpAC gradient	yes	yes	yes	-	-

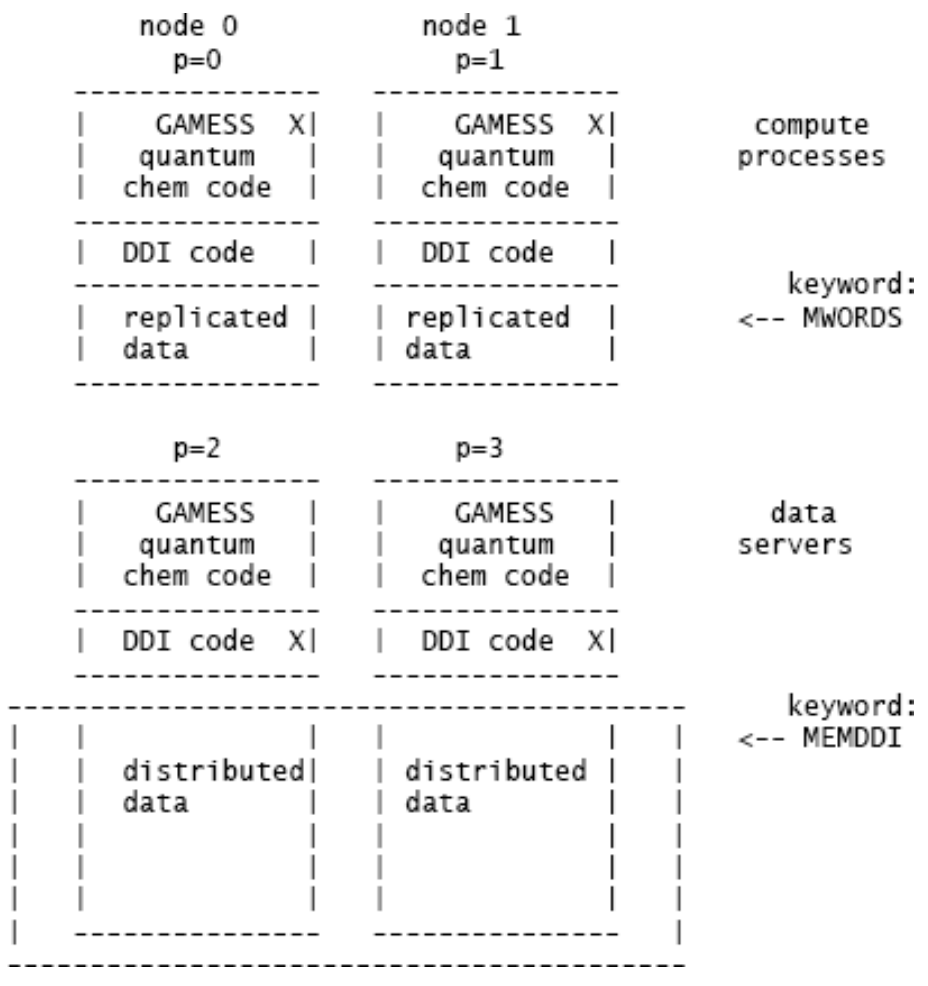

GAMESS Scalability on Clusters

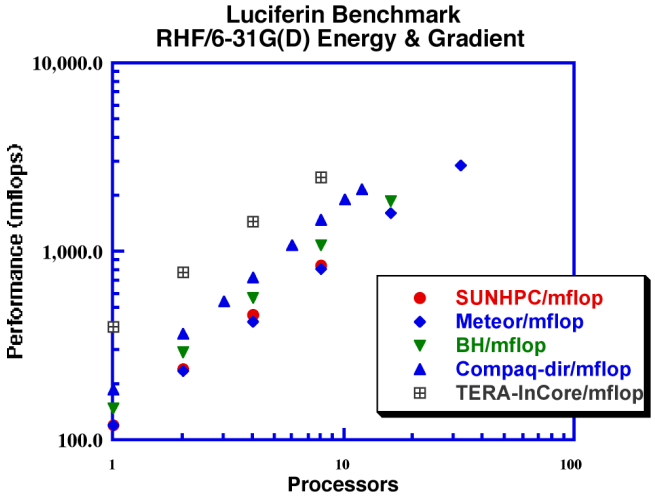

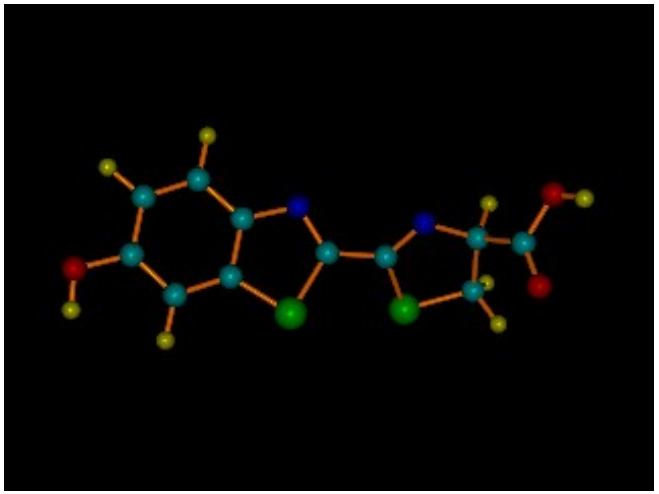

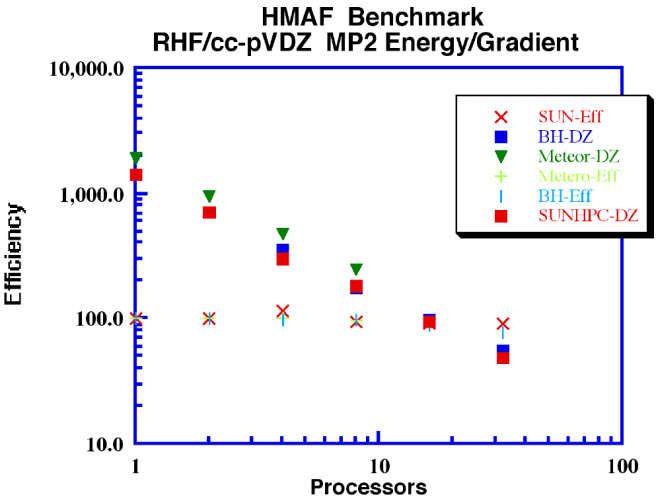

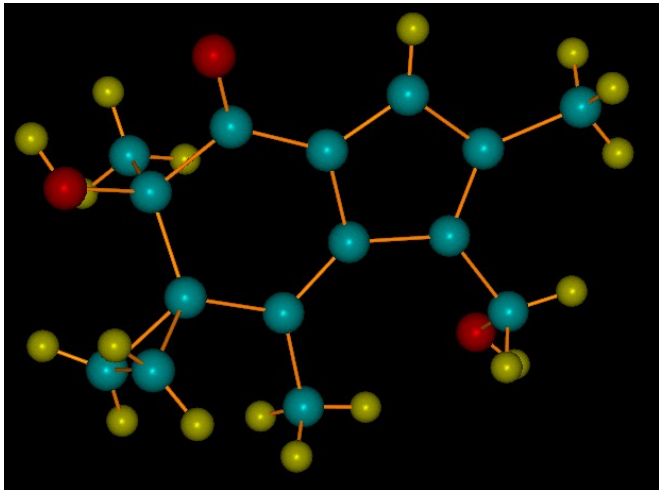

K. K. Baldridge et al. in "Parallel Computing for Bioinformatics and Computational Biology", A. Y. Zomaya (Ed.), Chapter 22, John Wiley & Sons, 2006

Size and Method Scaling

Hydrocarbon example

- Methane: CH_4
- Ethane: CH_3-CH_3
- Propane: CH_3-CH_2-CH_3
- n-Butane: CH_3-CH_2-CH_2-CH_3
- n-Pentane: CH_3-CH_2-CH_2-CH_2-CH_3

Scaling with increased electron correlation (N = number of basis functions)

- HF, DFT: N^3-N^4
- MP2: N^5
- MP4: N^6
- CC: N^7
- Etc.

⇒ Linear scaling approaches, plane wave methods, QM/MM, continuum electrostatics, etc.

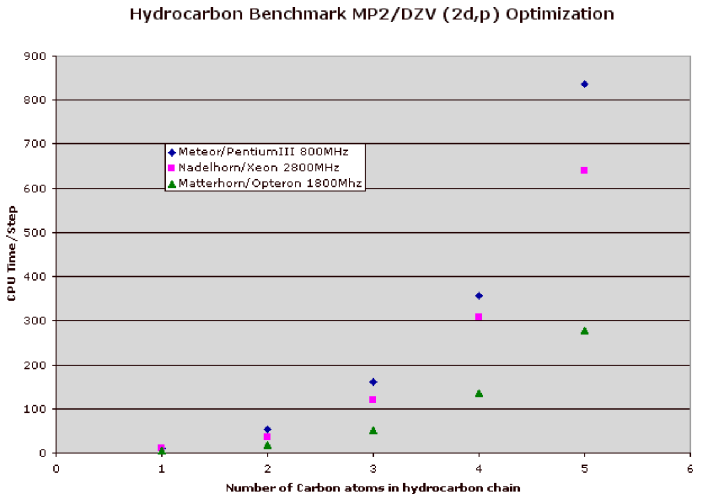

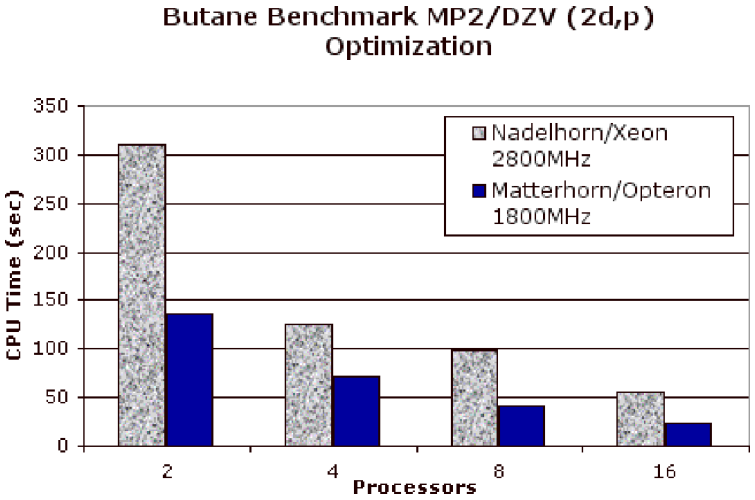

K. K. Baldridge et al. in "Parallel Computing for Bioinformatics and Computational Biology", A. Y. Zomaya (Ed.), Chapter 22, John Wiley & Sons, 2006

Parametric Modeling

Parameter sweeps and optimizations

⇒ Embarrassingly parallel, thus perfectly suited for grids

⇒ Typical variables: Method keywords, function values, molecular structures

Example experiment

- Development of Group Difference Pseudopotential (GDP)
- Nimrod tools for distributed parametric modeling
- PRAGMA grid test-bed

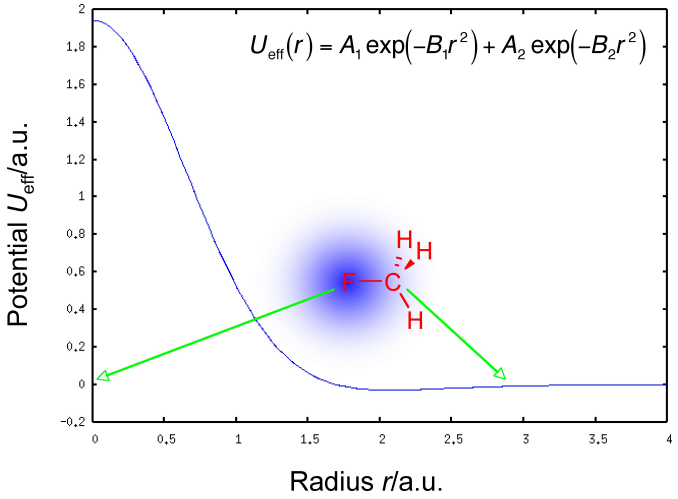

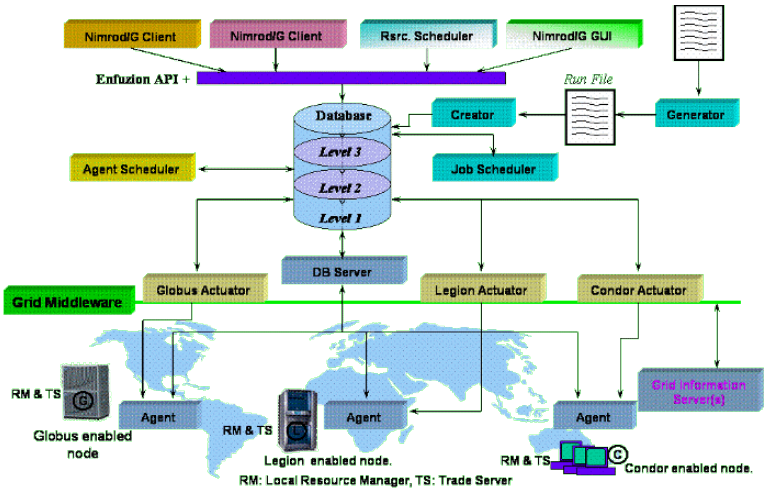

Resource	# of CPUs peak	# of CPUs reported by MDS	Total number of jobs	Total execution time	Average job execution time
brecca-2.vpac.org	29	186	4648	19 days 21:11:08	0:06:10
apbs.rocksclusters.org	2	12	143	15:05:57	0:06:20
slic00.sdsc.edu	7	148	443	2 days 21:08:46	0:09:22
erikson.ucsd.edu	32	76	3965	35 days 09:12:01	0:12:51
hathor.csse.monash.edu.au	36	36	2523	48 days 20:45:40	0:27:53
ume.hpcc.jp	57	64	3178	69 days 04:05:33	0:31:21
koume.hpcc.jp	4	8	255	5 days 21:37:46	0:33:19
chemcca40.ucsd.edu	13	36	721	18 days 00:17:15	0:35:58
		566	15876	200 days 17:24:06	0:18:12

http://www.csse.monash.edu.au/~davida/nimrod/nimrodg.htm
W. Sudholt et al., New Generation Computing 22 (2004) 137-146

Grid User Interfaces

Reasons

- Hide the complexity of the infrastructure and middleware, transparency
- Required for education, beginning users, setup, visualization, analysis, automation, interdisciplinarity, etc.

User interfaces

- Command line
- Libraries
- Web portals
- Rich clients
- Web services

Interface types

- Middleware specific (Nimrod, GRB, etc.)
- Grid specific
- VO / domain specific
- Application specific

Gemstone (Grid Enabled Molecular Science Through Online Networked Environments)

- http://gemstone.mozdev.org/
- Collaboration with SDSC
- Rich client based on Mozilla and XUL
- Molecular visualization (Garnet, etc.)
- Application specific web services (Opal)
- Strongly typed data schemas
- Globus middleware
- Rocks cluster infrastructure

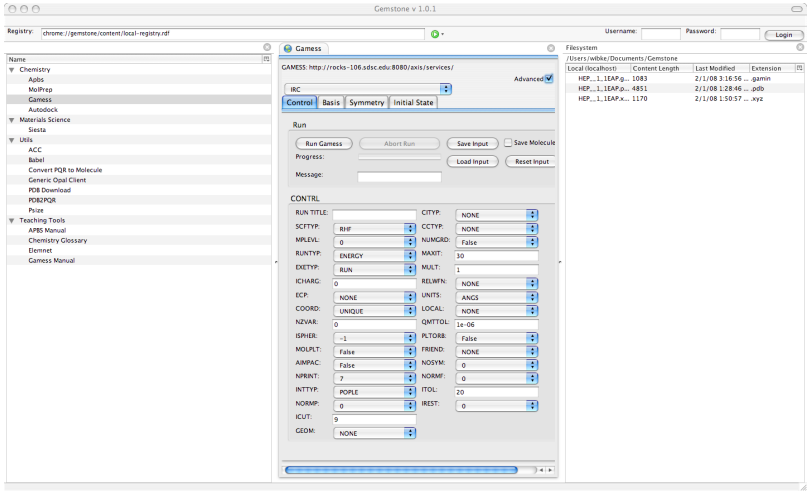

Scientific Workflow Systems

Universität Zürich

Advantages

- Interoperability across different scientific methods, scales, and domains, collaboration
- Automation of processes

Challenges

- Code and algorithm interoperability, data format standards

⇒ CML, QC-ml, Q5cost, etc.

- Job and data distribution, metadata

⇒ Application in QSAR, molecular docking, image processing, etc.

Resurgence (RESearch sURGe ENabled by CyberinfrastructurE)

- http://ocikbws.uzh.ch/resurgence/
- Collaboration with SDSC
- Based on Kepler

Chemomentum

- http://www.chemomentum.org/
- EU FP6 project with 9 partners from 7 countries
- Based on UNICORE

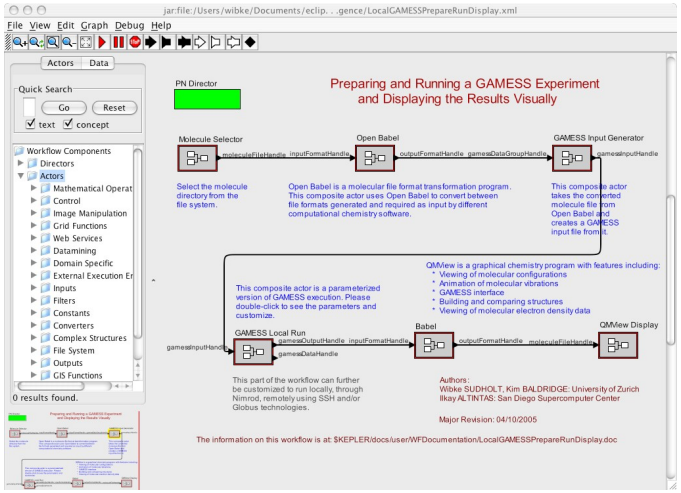

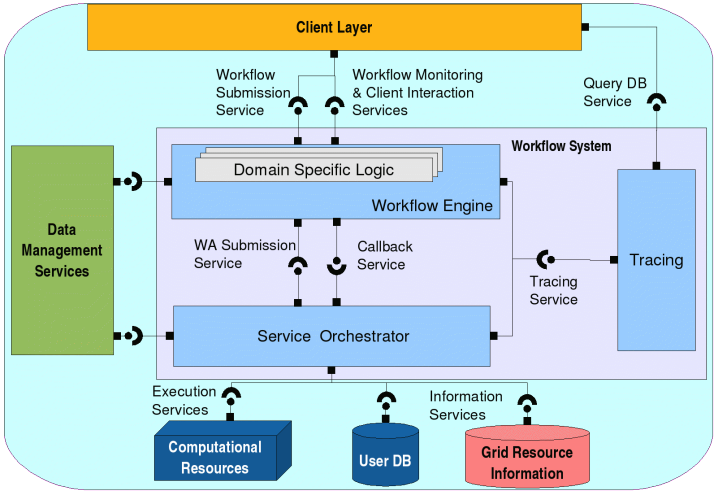

W. Sudholt et al., ICCS 2006, LNCS 3993 (2006) 69-76
B. Schuller et al., UNICORE Summit 2007, LNCS, to appear

Reinsurance Natural Catastrophe Calculations

Reinsurance Natural Catastrophe Calculations

Universität Zürich

Research collaboration with Swiss Re

Background

- Probabilistic modeling of financial losses from natural catastrophes (earthquakes, storms, etc.)
- Based on actual and simulated events
- Applied to portfolios of insured assets (houses, factories, etc.)
- Losses calculated from parameterized vulnerability curves
- Application of (re)insurance conditions and summation of losses over a portfolio and event

NatCat application and test-bed

- Multi-threaded Java server and client code
- Oracle database
- Set of typical test cases
- IBM WebSphere Application Server ND cluster
- Intel Xeon dual processor machines

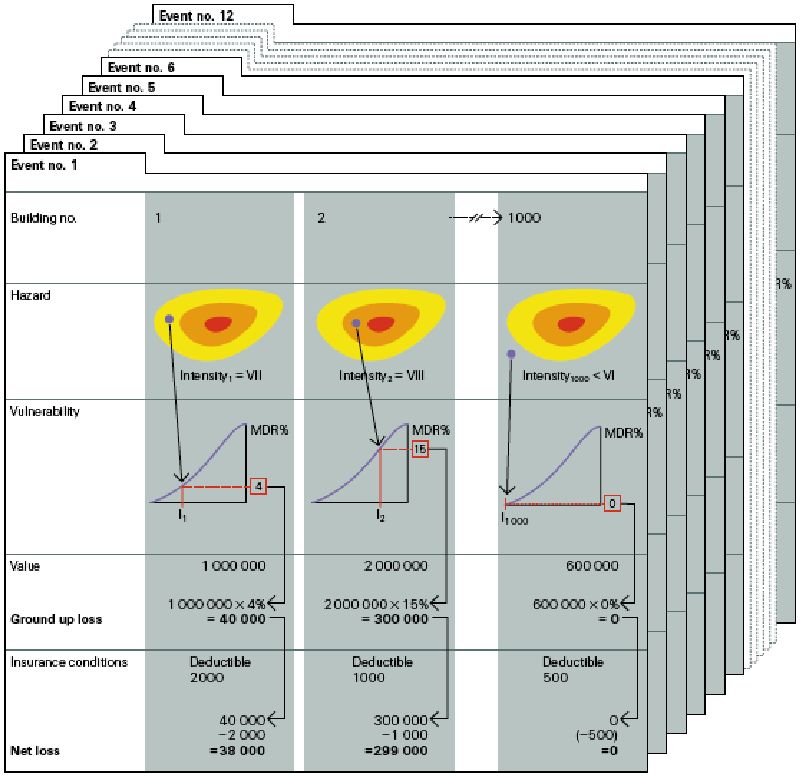

Event loss

No. 1 =	(38 000	+	299 000	+	.....	+	.....	+	0)	=	23.5 m
No. 2 =	(0	+	25 000	+	.....	+	.....	+	54 000)	=	42.5 m
No. 3 =	(75 000	+	36 000	+	.....	+	15 000	+	0)	=	74.8 m
No. 4 =	(0	+	0	+	.....	+	.....	+	63 000)	=	8.9 m
No. 5 =	(.....	+	.....	+	.....	+	.....	+	.....)	=	13.1 m
No. 6 =	(.....	+	.....	+	.....	+	.....	+	.....)	=	69.6 m
↓												
No. 12 =	(0	+	0	+	.....	+	.....	+	63 000)	=	58.6 m

http://www.swissre.com/pws/research%20publications/risk%20and%20expertise/risk%20perception/natural%20catastrophes%20and%20reinsurance.html

Application Requirements

Enhance the performance, scalability, and stability of the calculations compared to the current system

Find solutions to increase the throughput and maintain fairness of job distribution in case of high load

⇒ Parallelization to distribute both computational as well as memory load

⇒ Minimization of data transfer and usage

⇒ Optimization of code and calculations

⇒ Access to additional computation resources

Data Decoupling and Distribution

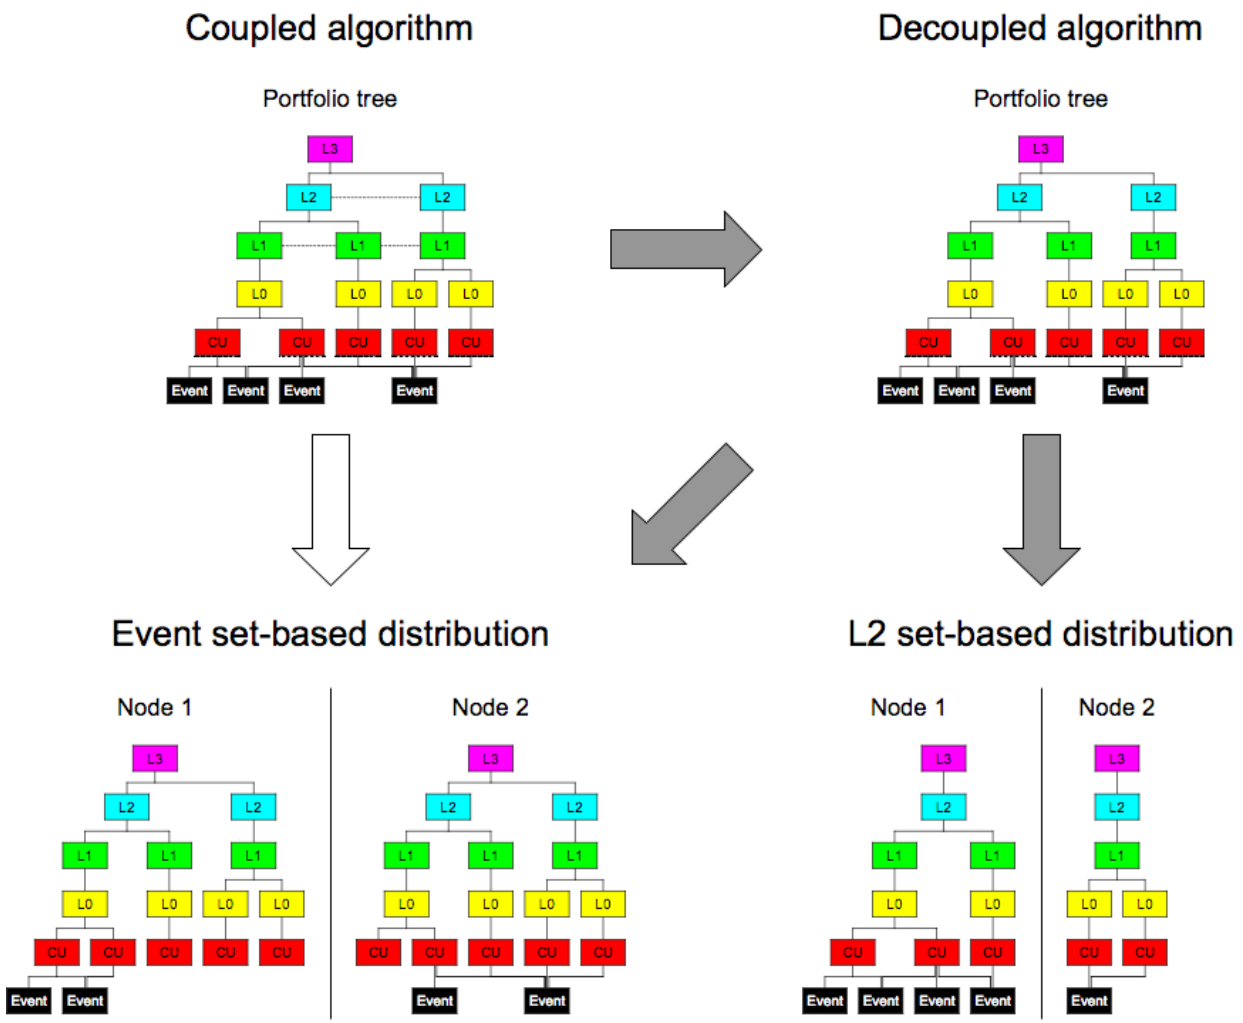

Performance and Scalability Enhancements

Event set-based distribution

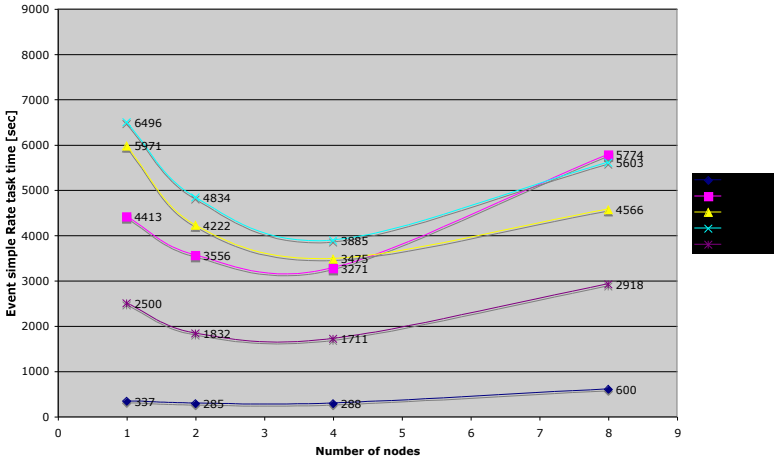

L2 set-based distribution

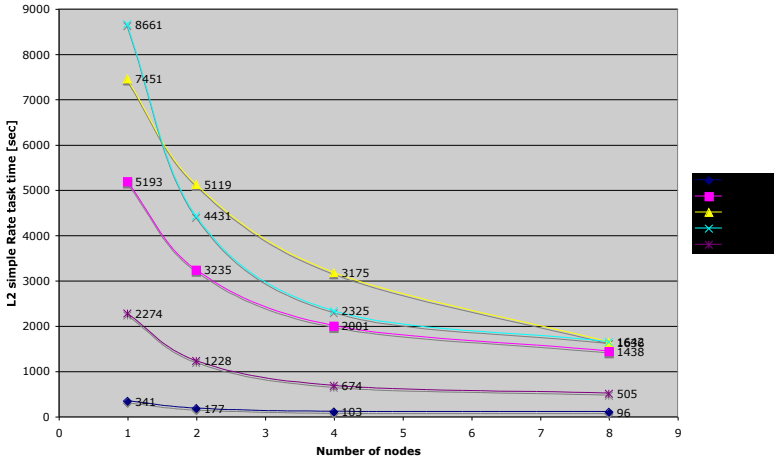

Data and processing optimization

- Combination of sub-processes to limit back-and-forth data transfer
- Substitution of remote database requests by local file access
- Removal of duplicate data
- Cluster-wide caching
- Improvement of I/O
- Testing of different grid middleware
- Etc.

Porting to Grid Infrastructures

Computing becomes more and more important in the financial sector

- Risk estimation
- Monte Carlo simulations for stock pricing
- Etc.

Issues with current grid infrastructures include

- Porting of applications
- Usage and maintenance of middleware
- Availability and complexity of the systems
- Security and policies

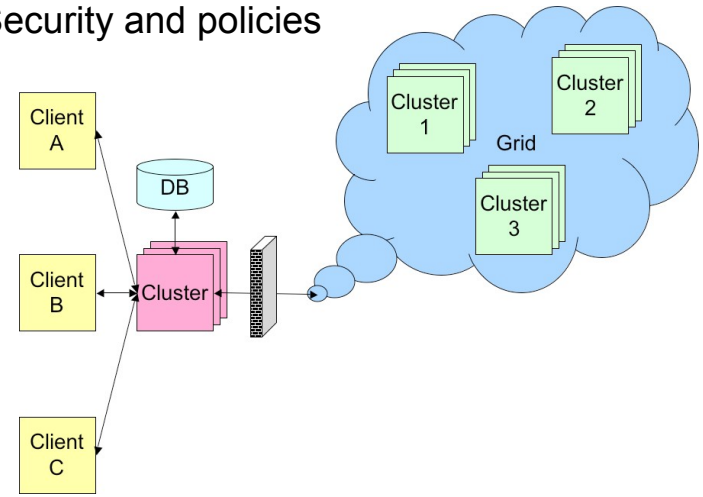

Concepts of application porting

- Keep current cluster infrastructure
- Spill over jobs to grid in case of high load
- Package data with computations in a secure way

Advantages of architecture

- Complete transparency of grid to users
- Almost no change to current infrastructure
- Grid usage just on demand
- Fallback to cluster possible
- Minimization of data transport
- Not much data exploitation possibilities

Grid infrastructures to be used potentially

- Other local resources
- Existing grids (PRAGMA, SEPAC, etc.)
- Emerging grids (SwiNG, EGI, etc.)
- Cloud computing (Amazon EC2, etc.)
- Etc.

Conclusions

The challenge of using cluster and grid systems is not so much the availability of the hardware, but the **complexity of the middleware** and the efforts necessary to **port and interface the application software**

Data become an increasingly important issue, not only their **amount and distribution**, but even more their format **complexity and incompatibility**

We not only have to satisfy the small number of **"high end users"**, but also a large number of **"middle class users"**

⇒ Investment into **human resources and training**, in particular to promote software development and to build a bridge between computational scientists and infrastructure providers, is necessary

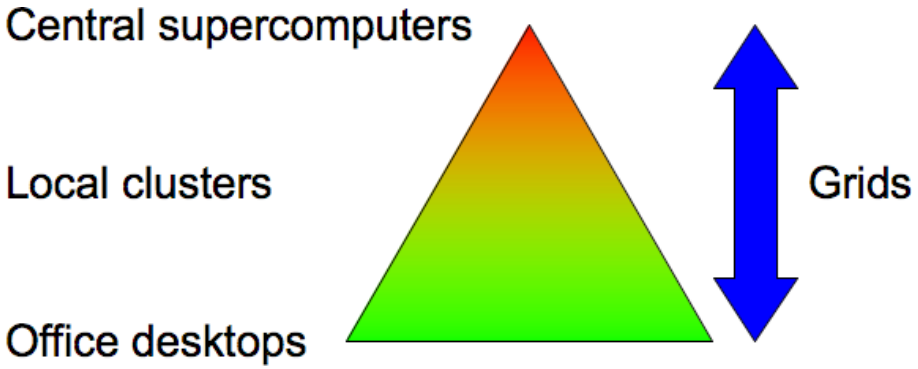

⇒ A **continuum of computational resources**, including remote supercomputers and local clusters, interconnected by grids, needs to be provided and accessible to computational scientists

⇒ The access to these resources should be steered both by **scientific merit** as well as by **appropriateness of usage** for a particular type of resource