Numerische Simulationen im AstroGrid

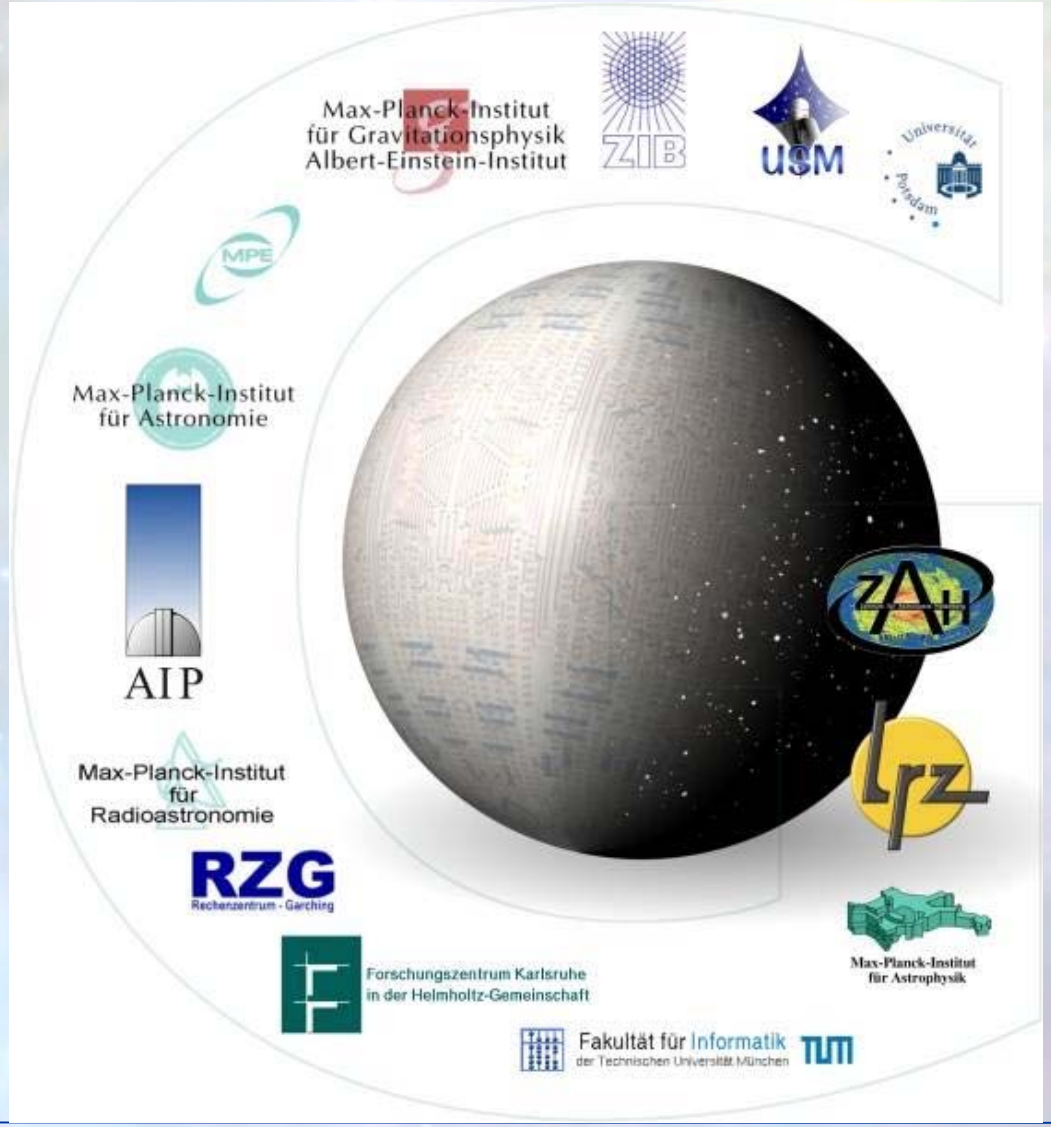

Übersicht

- Geschichte
  - Num. Rel. Simulation mit GTK 2.1
  - Tests mit kleinem Testbed
- AstroGrid-D:
  - ‚Embarrassingly' parallel : Dynamo + GEO600
  - MPI im D-Grid
  - Special Hardware Cluster
- Lessons learned
  - MPI im D-Grid
  - MPI über Internet

Astrophysikalische Simulationen

Jahr	Simulationen	Daten	CPUh
1990	Zusammenstoss zweier Galaxien 32^3 Partikel	~ 20 Megabyte	10 000
2001	Zusammenstoss zweier schwarzer Löcher (AMR): LSU,LRZ,AEI via Globus	~ 200 Megabyte	100 000
2009	Zusammenstoss von Galaxien 1024^3 Partikel	~100 Terabyte	1 Mio.
2009	Grossräumige Strukturen im Universum Lokale Auflösung bis 4096^3 Partikel	~1 Petabyte	20-50 Mio.
2009	Sterndynamos / Galaktische Magnetfelder (Magneto-Hydro-Dynamik)	~50 Gigabyte	100 000
	Black Hole Collisions (AMR)	~5 Terabyte	1 Mio.

Geschichte

the globus alliance

FAQ
Site Map
Contact Us

Home | Globus Alliance | Globus Toolkit | Grid Software | Grid Solutions | dev.globus

Website Email Lists Search:

10.22.2009 **Important information on Globus events and plans** Learn more...
07.13.2009 **CoG JGlobus 1.7.0 Release** Learn more...
07.13.2009 **Spider, the world's biggest Lustre file system at ORNL's LCF is accessible via GridFTP** Learn more...
07.10.2009 **KnowARC Project Brings Grids to Debian** Learn more...
06.09.2009 **Nimbus in the news** Learn more...

XML What's this?

Archive of Globus Alliance news

Welcome to Globus®

The **Globus Alliance** is a community of organizations and individuals developing fundamental technologies behind the "Grid," which lets people share computing power, databases, instruments, and other on-line tools securely across corporate, institutional, and geographic boundaries without sacrificing local autonomy. Learn more...

The **Globus Toolkit** is an open source software toolkit used for building Grid systems and applications. It is being developed by the Globus Alliance and many others all over the world. A growing number of projects and companies are using the Globus Toolkit to unlock the potential of grids for their cause. Learn more...

The Globus Alliance is an active member in the community of **Grid Software** developers. Learn more... As partners in e-Science and e-Business projects, we've built **Grid Solutions** for a variety of challenges that come up when people share resources. Learn more...

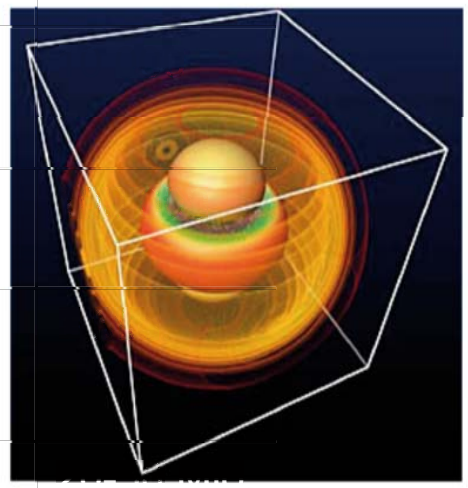

Physicists used the Globus Toolkit and MPICH-G2 to harness the power of multiple supercomputers to simulate the gravitational effects of black hole collisions. The team, which included researchers from Argonne National Laboratory, the University of Chicago, Northern Illinois University, and the Max Planck Institute for Gravitational Physics in Germany, was awarded a prestigious Gordon Bell prize for its work. Image courtesy of Max Planck Institute for Gravitational Physics.

Black Hole Collision

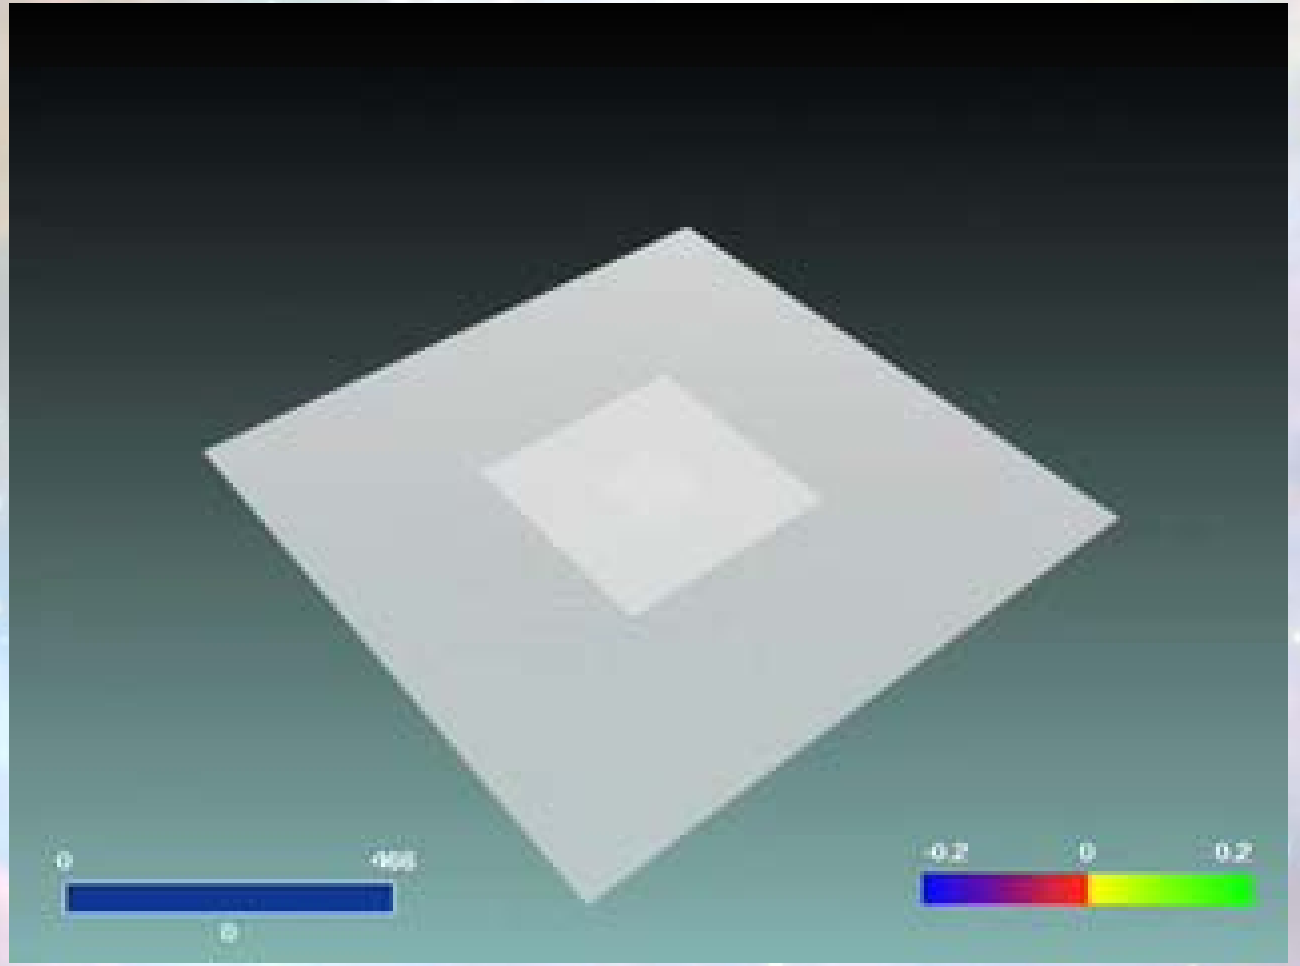

AEI
(AMR)

GAVO-GRID (testbed) mit Globus 2.3 (2004)

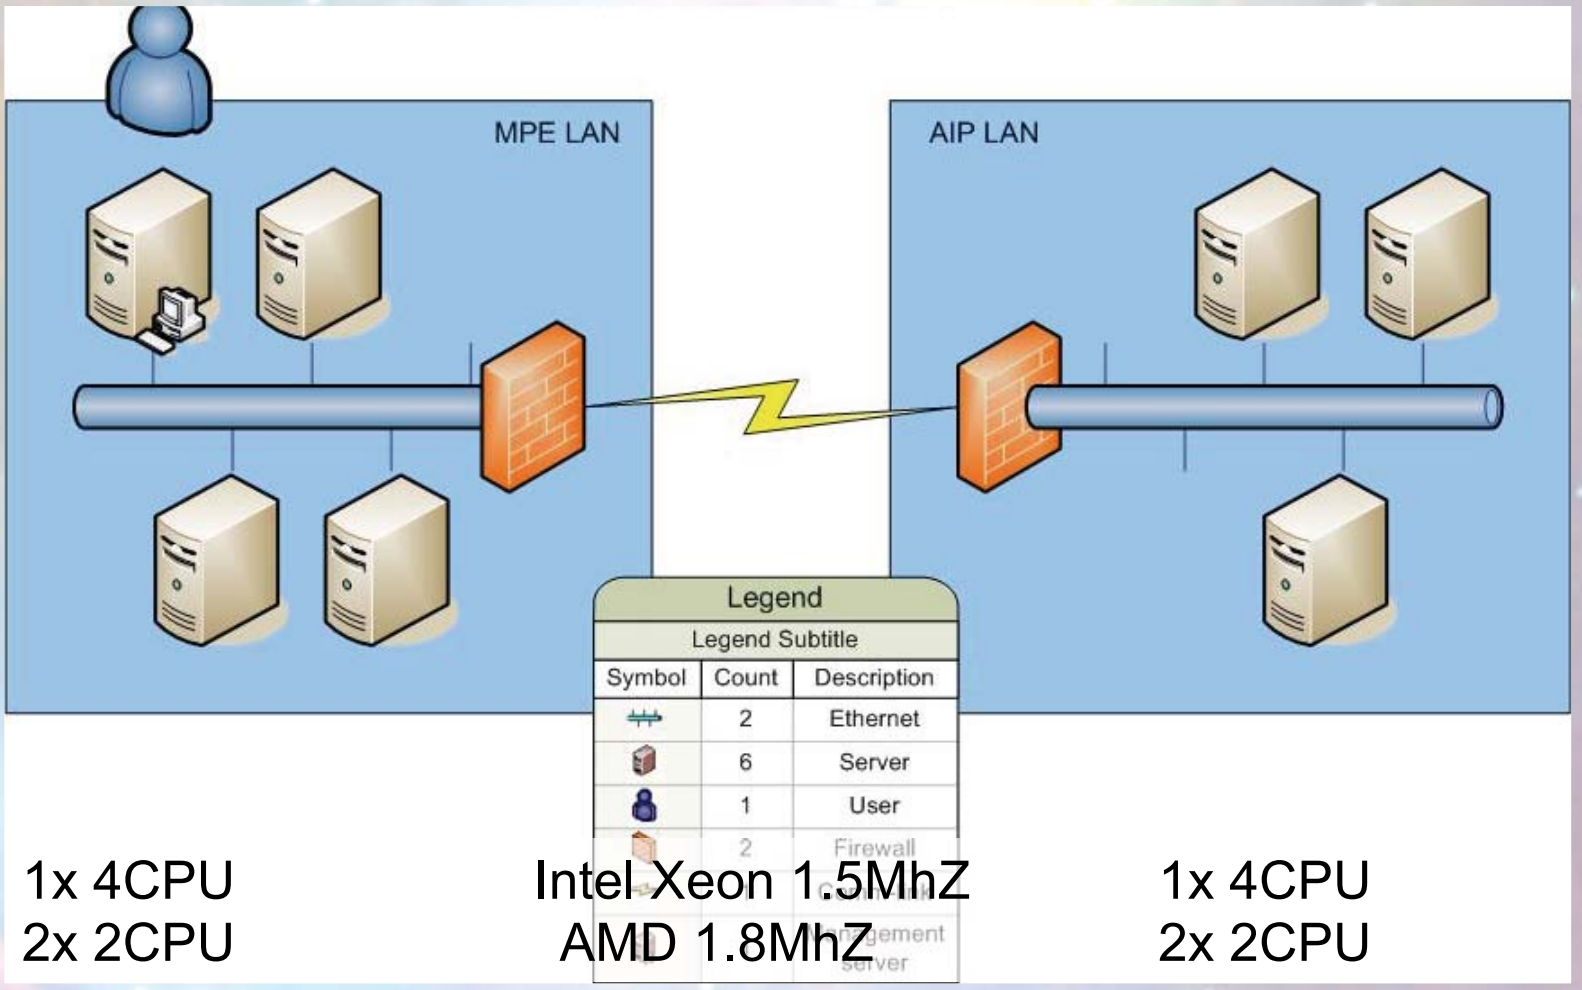

GAVO-GRID: MPI across Sites, Code GADGET (v 1.5), MPICH Globus

uniform nodes	Wall time	Comm time
On board	167,05	0,11
Across site	179,69	1,05
(np = 8)	105,3	1,30

GAVO-GRID: MPI across Sites, Code GADGET (v 1.5), MPICH Globus

heterogenous nodes			
	Np	Wall time	Comm time
On site	2/2 AMD/ XEON	162,03	0,11
Across site	4 AMD	97,7	1,16
On site	4/4 AMD/ XEON	85,9	0,15
Across site	4/4 AMD/ XEON	105,33	1,18
Across site	8/8 AMD/ XEON	77,63	1,96

AstroGrid-D: Vorhaben

sanssouci.aip.de

Galaxien-Kollision
Postprocessing

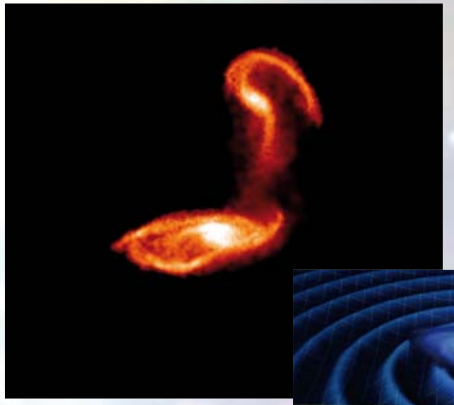

LISA

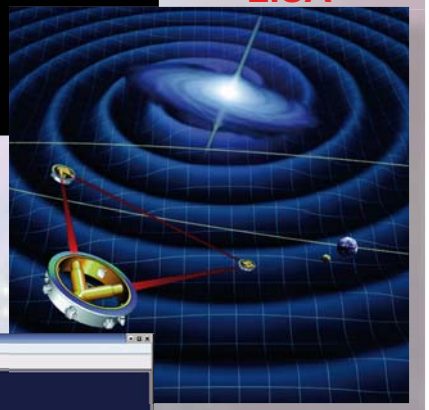

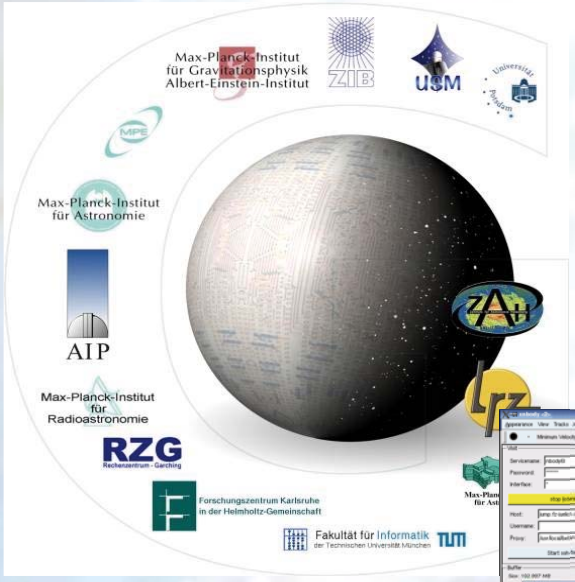

LOFAR

X-NBODY6++
Visualisierung

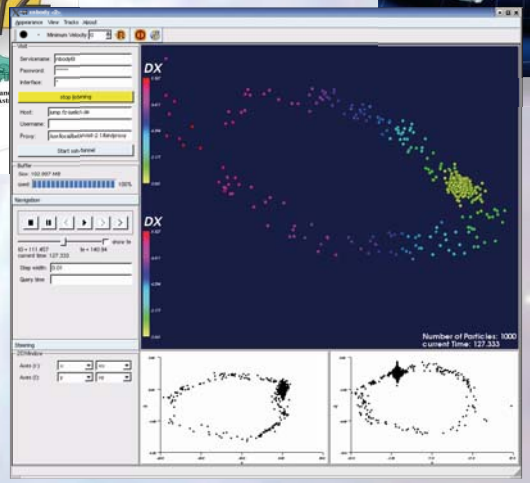

Embarrassingly Parallel:
GEO600 (mit BOINC)

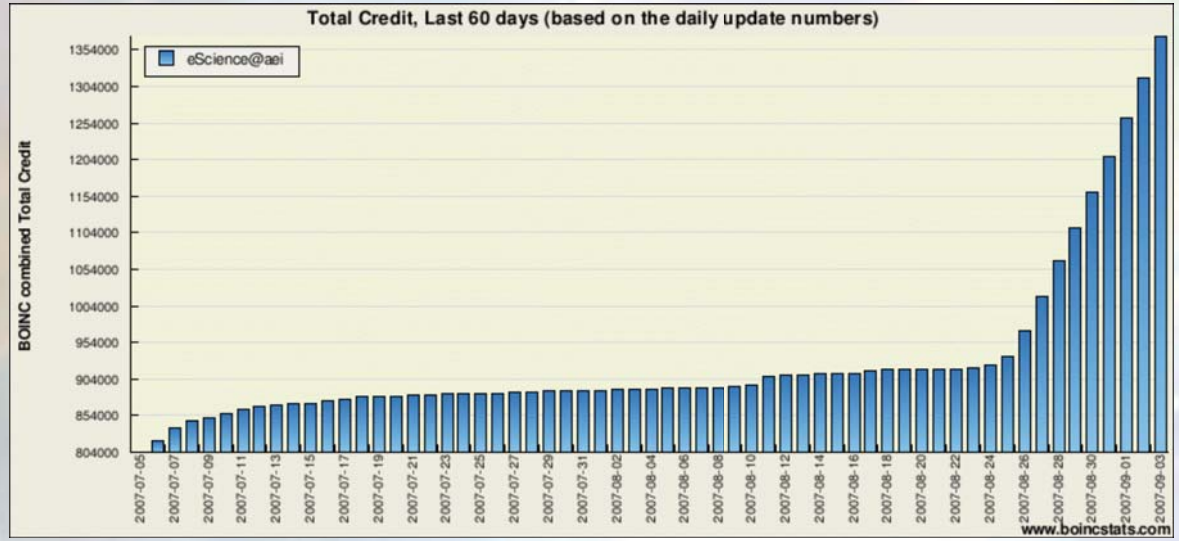

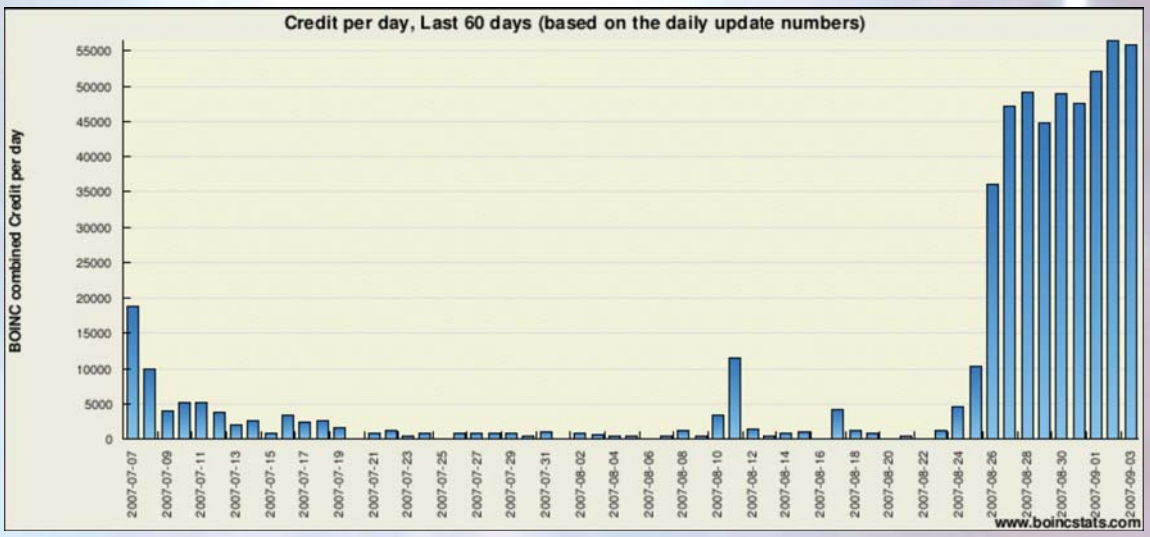

Embarrassingly Parallel:
GEO600 (mit BOINC)

AstroGrid-D use case GEO600

GEO600 was identified as an ideal science use case to benefit from Grid computing. In the AstroGrid-D project we developed a grid-enabled framework around Einstein@Home. It uses the Globus 4 toolkit as standard Grid middleware. After the software has automatically deployed itself on a new Grid machine, an automated mechanism then takes care of continuously submitting jobs to available resources and maintaining their status.

Since October 2007 our GEO600 grid application has been running in production mode on a large number of supercomputers, clusters, and PCs within D-Grid in Germany and on the Open Science Grid (OSG) in the U.S. Achieving a throughput of several thousands of jobs per day, we now rank within the top 5 teams world wide to contribute to the Einstein@Home project. A GEO600 use case statistics is presented in the AstroGrid-D portal where you can view the number of job submissions and CPU hours spent on GEO600 data analysis over the last months.

Embarrassingly Parallel: Dynamo

DEMO

Special Hardware in computational science I

FPGA

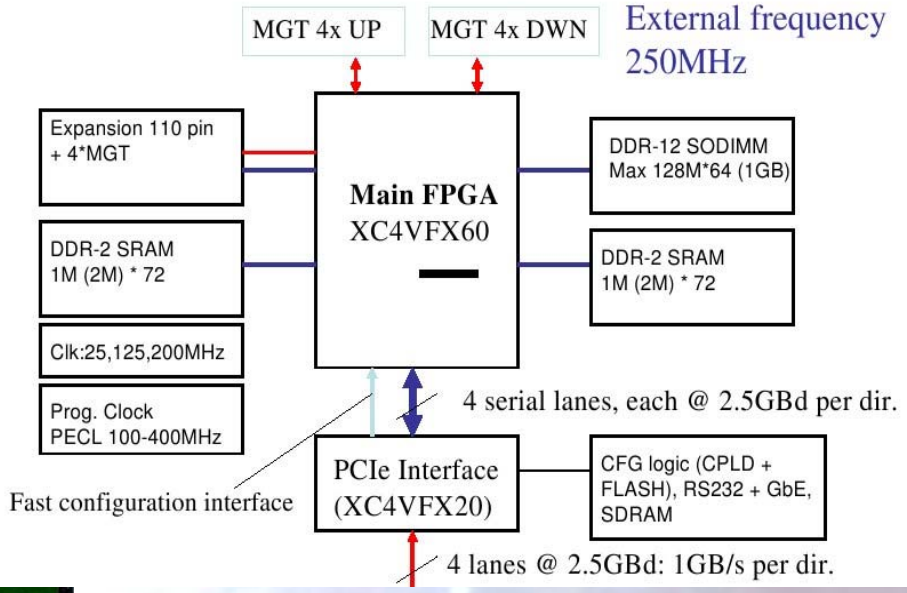

Courtesy of R. Spurzem

Special Hardware in computational science II

2007…
GeForce 8800 GTX, 128 Stream Proc., 768 MB
GeForce 8800 GTS, 128 Stream Proc., 512 MB
GeForce 8800 GT, 112 Stream Proc., 512 MB

2008…
GeForce 9800 GTX, 128 Stream Proc., 512 MB
GeForce 9800 GX2, 256 Stream Proc., 1 GB
GeForce 9800 GT, 64 Stream Proc., 512 MB

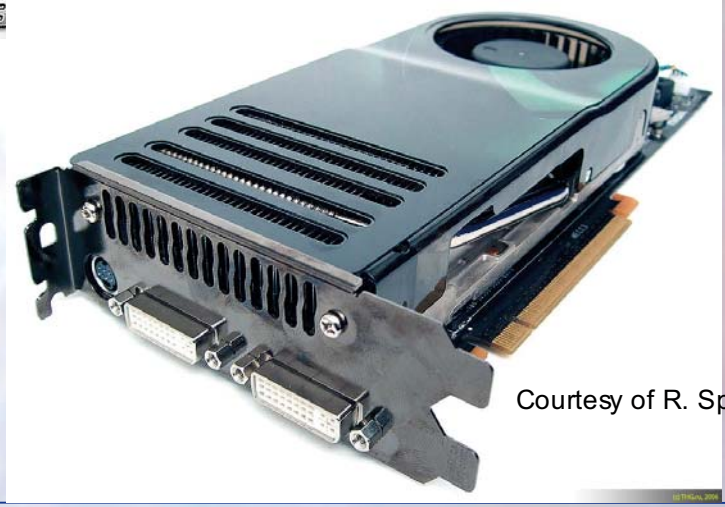

Courtesy of R. Spurzem

Special Hardware in computational science III

Beyond GRAPE solutions....

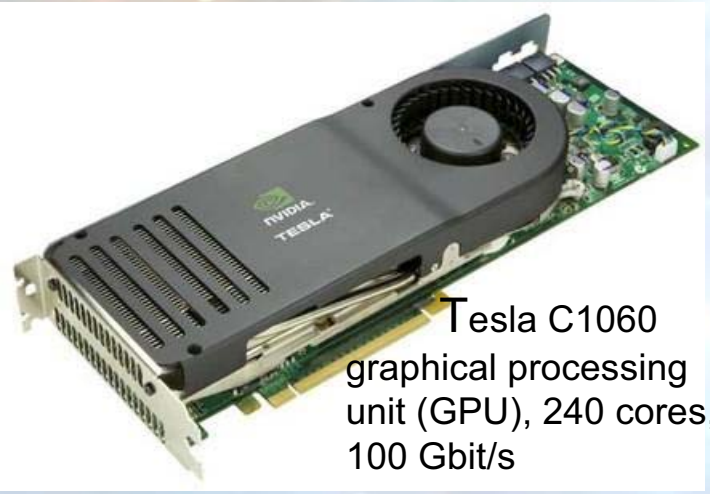

Tesla C1060 graphical processing unit (GPU), 240 cores, 100 Gbit/s

GPU 1 Tflop/s peak ~ 0.4 Watt / Gflop/s

IBM BlueGene ~ 2.5 Watt / Gflop/s

Standard PC ~ 25 Watt / Gflop/s

Real Codes Need More Power:

Efficiency N-Body ~ 80 % ~ 0.32 W/Gflop/s

Efficiency SPH ~ 4 % ~ 10 W/Gflop/s

Reconfigurable Hardware (FPGA)

Univ. Heidelberg MPRACE-1 Board

Efficiency SPH ~ 4 W/Gflop/s

Courtesy of R. Spurzem

Special Hardware in computational science IV

Versus GPU for SPH

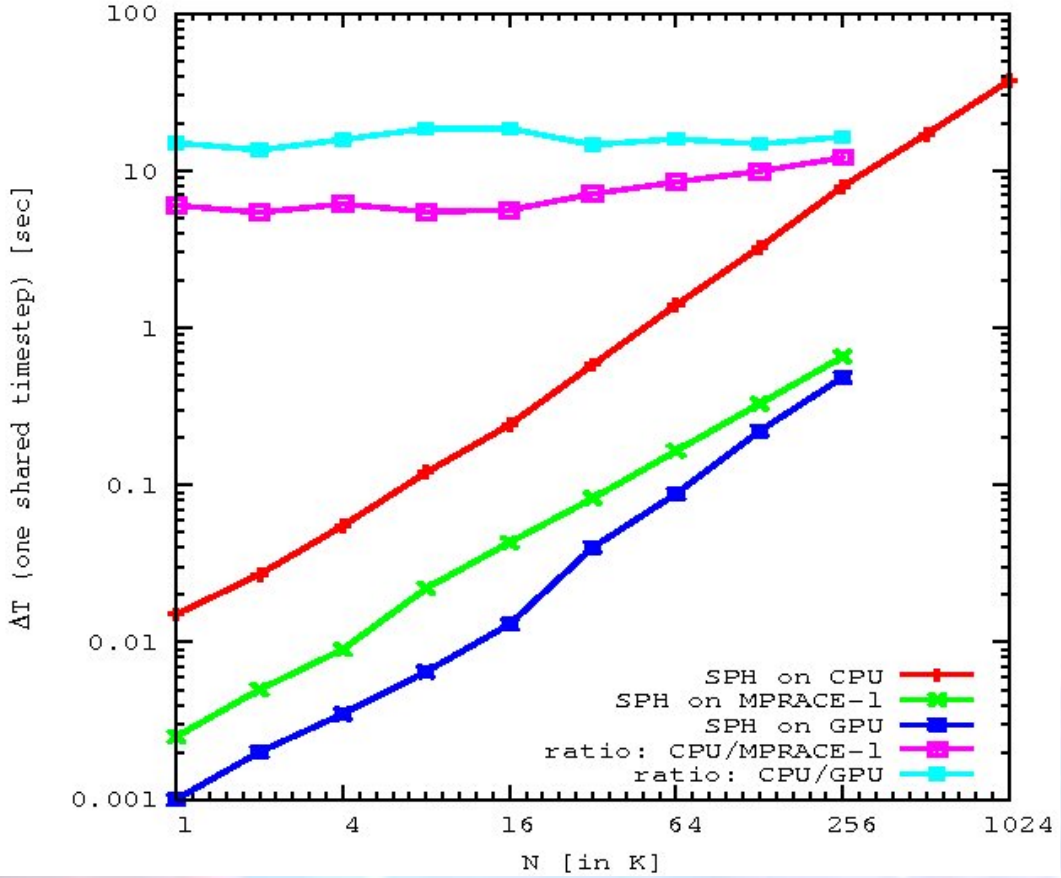

GeForce 8800 GTX (NVIDIA)
Using CUDA Library
Special Interfaces and API from GRACE project ported.

Spurzem et al. 2007,
Jl.Phys.Conf.Ser.
Berczik et al. 2008,
Marcus et al. 2008
(SPHERIC)
Spurzem et al. 2009
(ISC 09 Procs., Springer)

Courtesy of R. Spurzem

Collaborative Grid-Infrastructure for specialized Hardware

AstroGrid-D currently offers: software to access the grid:

- Web Portal to access the grid for common applications
- Gridway / Grid-Gateway for classical job submission
- *First International Link to MAO Kiev, Ukraine*

```
spurzem@hydra:~> gwhost
HID PRIO OS              ARCH   MHZ %CPU MEM(F/T)     DISK(F/T)       N(U/F/T) LRMS HOSTNAME
0   1    Linux2.6.16.33- x86_6 2411  200 3076/7939 197805/220515       0/2/2 Fork astrodata04.gac-grid.org
1   1    Linux2.6.11.4-2 x86_6 3200  300  866/3974   67177/79742       0/3/4 Fork titan.ari.uni-heidelberg
2   1    Linux2.4.27-2-3 x86    996   99    6/502    7150/19366       0/1/2 Fork buran.aei.mpg.de
3   1    Linux2.6.9-42.0 x86_6 2411   98 1399/7970 289588/311777       0/1/2 Fork astrodata05.gac-grid.org
4   1                            0    0     0/0         0/0         0/0/0       astrodata07.gac-grid.org
```

Stellaris

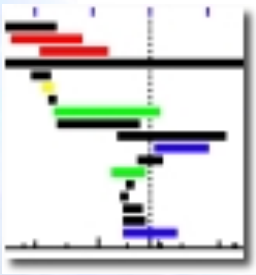

Client

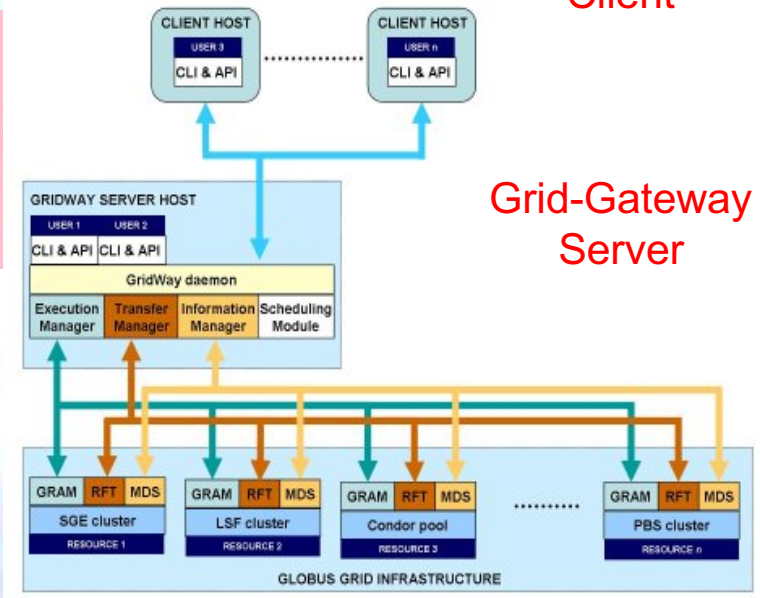

Grid-Gateway Server

- Data Processing Pipelines e.g. for Dome A ... (new project)

Information Systems Gridmap Grid Timeline Astrogrid-D Resources

Courtesy of R. Spurzem

MPI Jobs auf Grid Clustern

- MPI nicht vorhanden/konfiguriert
- Nicht standardisiert
  - (vendor/open implementation issues)
- Batch Queues oft nicht für exclusive Nutzung durch MPI-Jobs konfiguriert
- (zu wenig Möglichkeiten, eigene Installationen vorzunehmen)

www.astrogrid-d.org/project-documents/deliverables/wp1.html

Interconnect-Clusters für MPI Jobs

Konfiguration / Hardware

- Nodes mit unterschiedlicher Architektur
- Zeitliche Synchronisation
- Batch Queues
- Advanced Reservation / Scheduling

MPI bezogene Probleme:

- Bandwidth (heute IB ~10-20GBit/s)
- Latency (min. 10x grösser)
- Heterogenität der Cores

www.astrogrid-d.org/project-documents/deliverables/wp1.html

Grid und Numerische Simulationen

- Ideal für
  „embarrassingly parallel simulations“
  - Erforschung von Parameter- Räumen
  - Monte Carlo Rechnungen
  - Data Mining mit hohem numeerischen Aufwand und kleinen Datensätzen
- Gut für MPI-Jobs
  - mit geringen Ansprüchen an Version, Hardware
- Gut für Special Hardware MPI-Jobs
- Mässig für Integration von Clustern in Instituten für MPI-Jobs
- Ungeeignet für Integration von Clustern über Institutsgrenzen hinweg

Grid und Numerische Simulationen