Computer Physik

Computational Physics

Ralf Klessen, Institut für Theoretische Astrophysik
Rainer Spurzem, Astronomisches Rechen-Institut
Zentrum für Astronomie, Universität Heidelberg

spurzem@ari.uni-heidelberg.de, klessen@uni-heidelberg.de
http://www.ari.uni-heidelberg.de/mitarbeiter/spurzem/
http://www.ita.uni-heidelberg.de/~ralf

Warum Computerphysik notwendig?

- Differentialgleichungen (DGL)–keine analytische Lösung!
  Zusammenhang: Nichtlineare Dynamik, Chaostheorie
- Gewöhnliche zeitabh. DGL – Teilchensimulationen
  Plasmaphysik, Moleküldynamik, Stellardynamik…
- Partielle zeitabh. DGL – Gittermethoden
  Hydrodynamik, Relativitätstheorie, Quantenmechanik
- Zeitunabh. DGL – Mittelwerte, stationäre Zustände
  Quantenmechanik, statistische Mechanik, Monte-Carlo Verfahren.
- Spezialfall Computermathematik-komplexe analytische Lösungen, konventionell praktisch unlösbar.
  Mathematica, Maple, MatLab, …

Geschichte

- Erik Holmberg (1908-2000)

Dissertation Univ. Lund (Schweden) (1937):
``A study of double and multiple galaxies´´
Galaxien oft in Gruppen und Paaren
Satellitengalaxien ungleichmäßig verteilt (Holmberg-Effekt)

- **Vater der numerischen Astrophysik....**
  - **...mit 200 Glühbirnen**

Geschichte

- The Astrophysical Journal, Nov. 1941

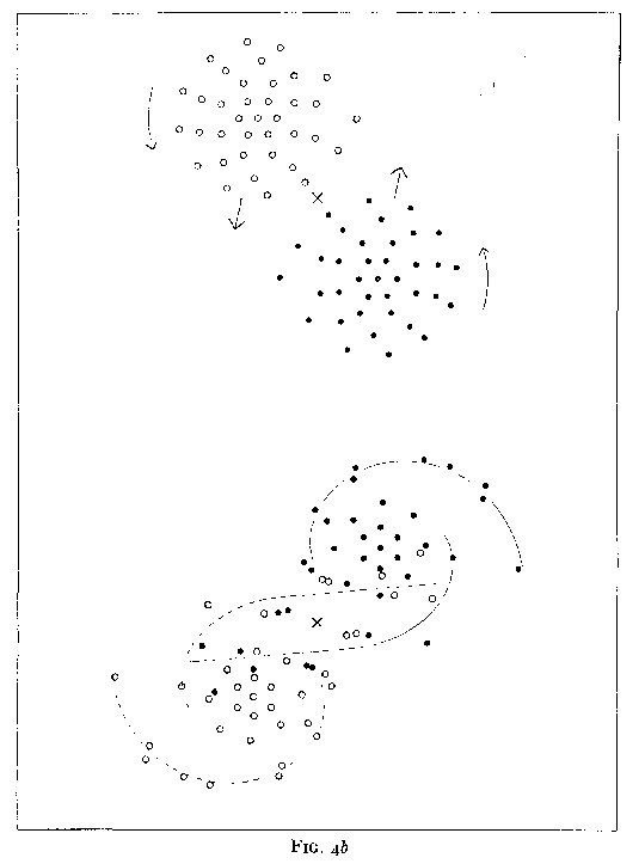

FIG. 4b

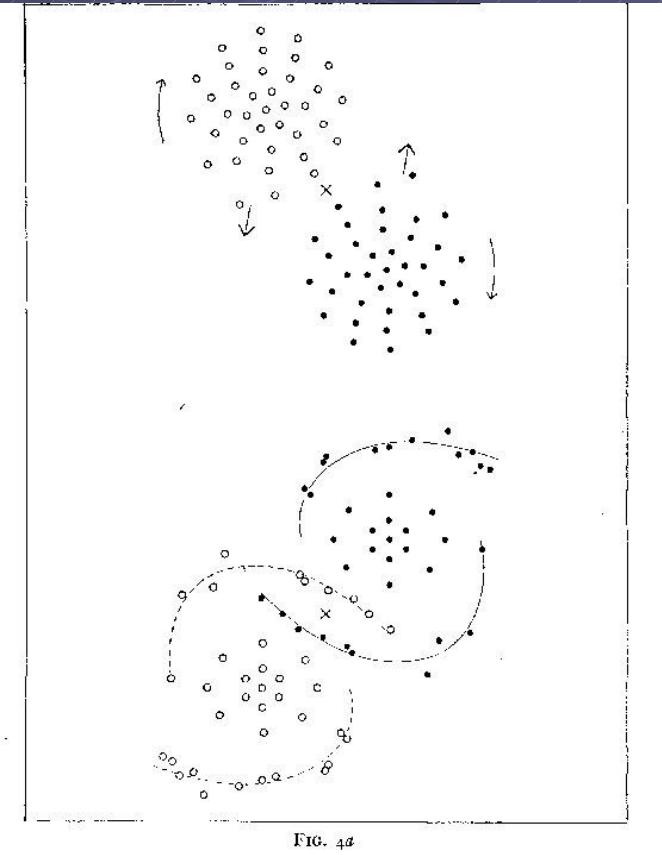

FIG. 4a

Geschichte

- Konrad Zuse (1910-1995) Berlin

Erfinder des frei programmierbaren Rechners

Z1 in der elterlichen Wohnung 1936

Schwingungen

Geschichte

0.03 Mflops

http://www.rtd-net.de/Zuse.html

Zuse Z4: 1944 Berlin, 1950 Zürich

1954 Frankreich

1959 Deutsches Museum München

Rechenanlage 0.03 MHz

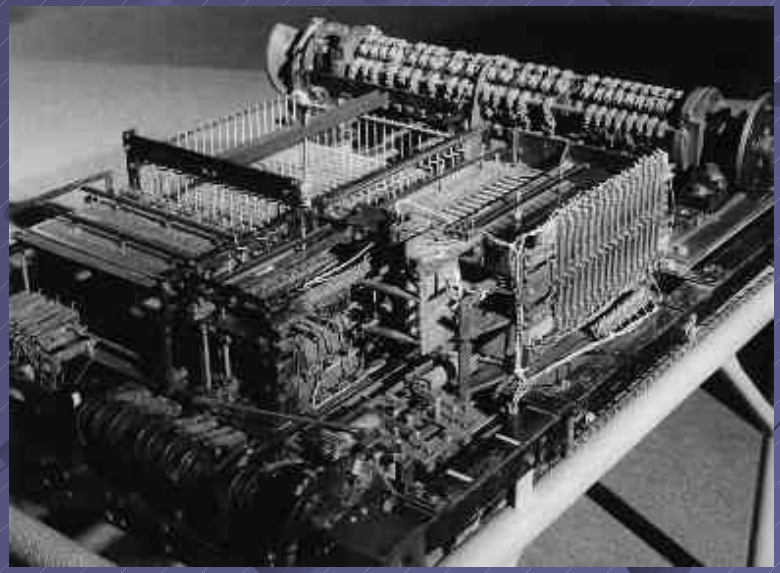

Speicher 256 byte

Geschichte

- Grundprinzipien des elektronischen. Rechners

Verwirklicht bei Zuse, Theorie von Neumann

Freie Programmierbarkeit
Binäres Zahlenformat
Speicher
Gleitkommaarithmetik

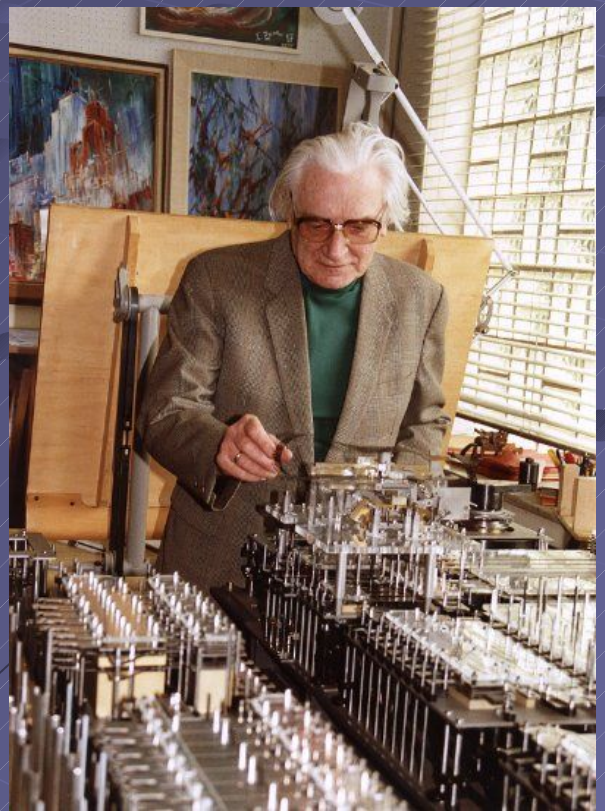

Geschichte

- John von Neumann (1903-1957)

Geb. Budapest, Dozent Berlin,

ab 1930 Princeton Univ. NJ USA Princeton

„Requirements for an electronic computing machine“ (1946)

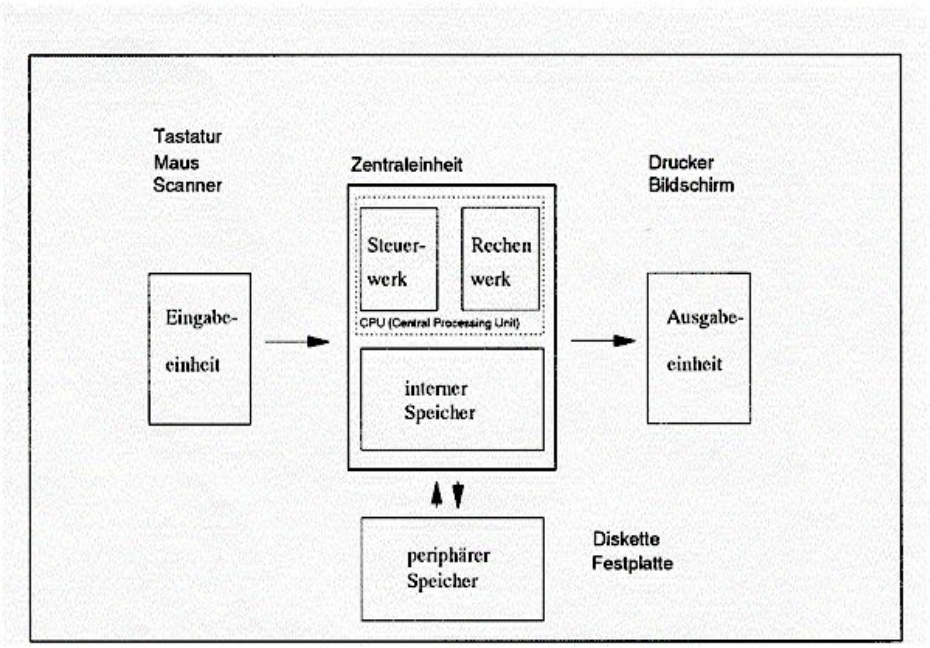

Geschichte

Astronomisches Rechen-Institut in Heidelberg
Mitteilungen Serie A Nr. 14

Die numerische Integration des *n*-Körper-Problemes für Sternhaufen I

Von

SEBASTIAN VON HOERNER

Mit 3 Textabbildungen

(Eingegangen am 10. Mai 1960)

Astronomisches Rechen-Institut in Heidelberg
Mitteilungen Serie A Nr. 19

Die numerische Integration des *n*-Körper-Problems für Sternhaufen, II.

Von

SEBASTIAN VON HOERNER

Mit 10 Textabbildungen

(Eingegangen am 19. November 1962)

S.v. Hoerner,
1919 – 2003

Z.f.Astroph. 1960, 63

Sebastian von Hoerner, 2001: How it all started, in: Dynamics of Star Clusters and the Milky Way, ASP Conference Series, Vol. 228. Edited by S. Deiters, B. Fuchs, R. Spurzem, A. Just, and R. Wielen. San Francisco: Astronomical Society of the Pacific. ISBN: 1-58381-060-9, 2001.

Geschichte

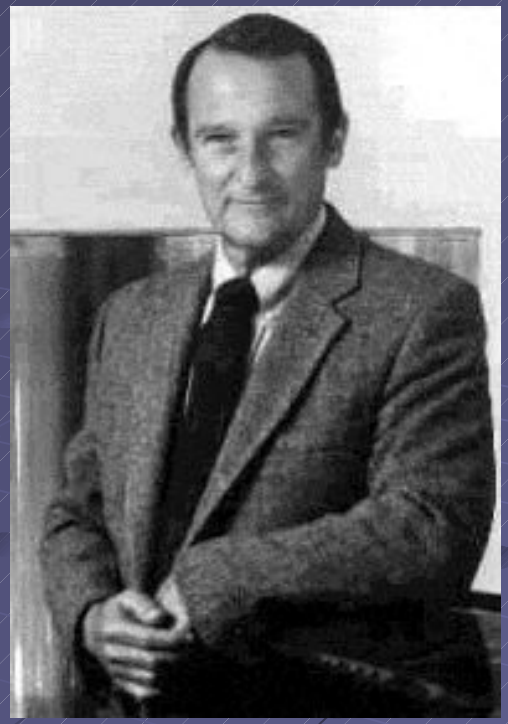

- Seymour Cray (1925-1996)
  - ``The father of supercomputing´´

CRAY1: Vektorregister (1976)
160 Mflop, 80 MHz, 8 MByte RAM
CRAY2: (1984)
1Gflop, 120MHz, 2GByte RAM

Einführung

Superrechner
JUGENE
IBM Blue Gene
Jülich
Supercomputing
Centre (JSC)
223 Tflop/s
...Petaflop/s...
2015: Exaflop/s
Prototype with IBM
2019: Exaflop/s System
(...Trillion...)

Eröffnet mit J. Rüttgers Juni 2008

Computational Science...

...after von Neumann...

Exaflop/s?

Petaflop/s

Teraflop/s

Gigaflop/s

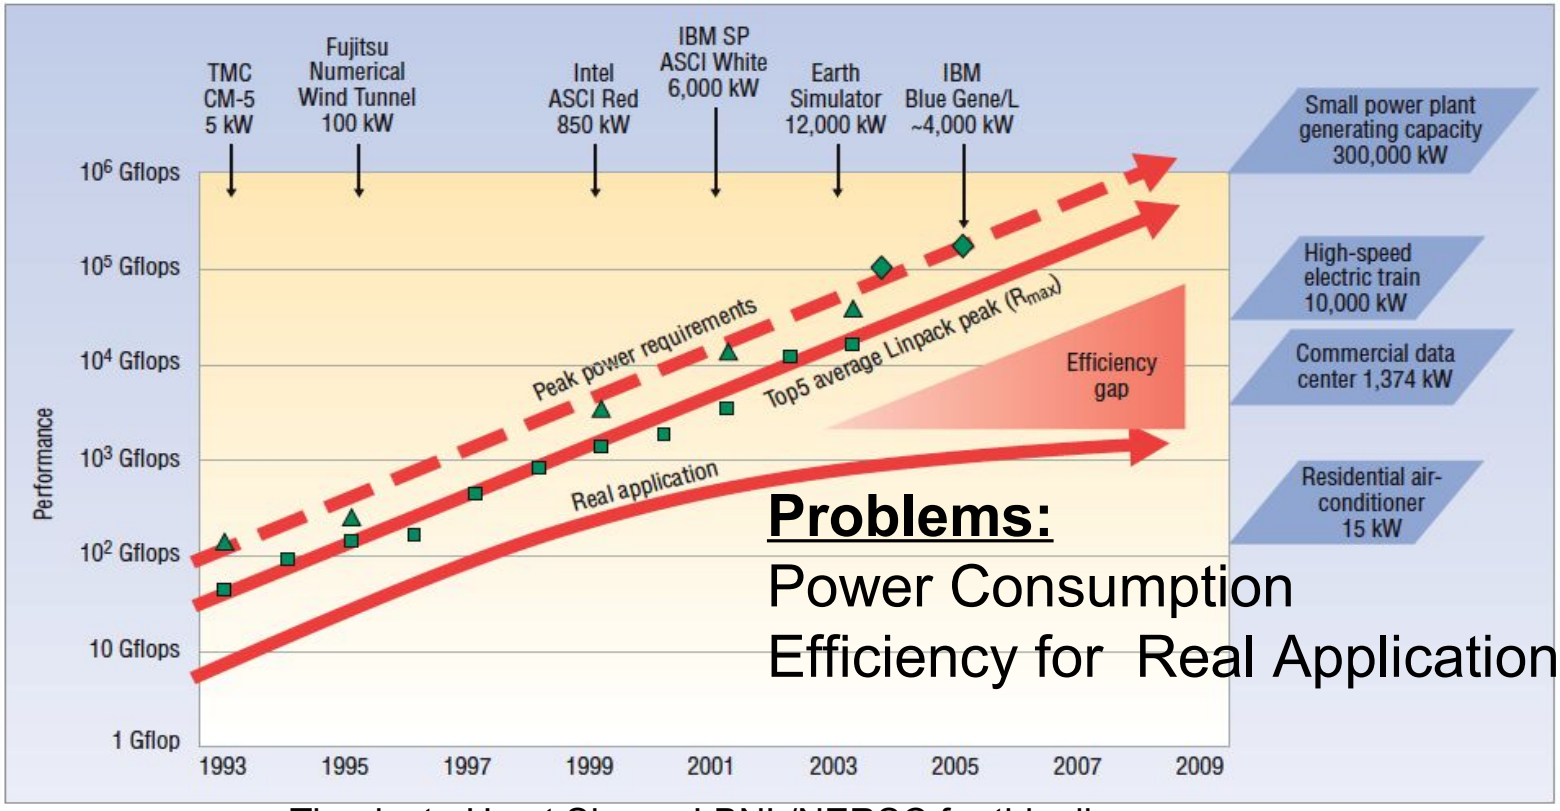

Problems:

Power Consumption

Efficiency for Real Applications

Thanks to Horst Simon, LBNL/NERSC for this diagram.

Figure 1. Rising power requirements. Peak power consumption of the top supercomputers has steadily increased over the past 15 years.

SPECIAL HARDWARE

CPUs

Central Processing Units

General Purpose oriented

1-12 Cores

Up to 4 pipes per core using Vector Units

Fully Programmable, many languages available

Very well studied

Max. 125W per processor

GPUs

Graphic Processing Units

Graphics oriented

16-512 Cores

Massively Parallel Architecture, specialized instructions for parallel processing

Fully programmable, but limited languages

Algorithms not fully explored

Max. 400W per card

FPGAs

Field Programmable Gate Arrays

Custom designs, best for processing streaming data

Programmable Logic, Architecture is custom-built for the required application

Requires extensive knowledge to program, development time is longer than CPUs and GPUs

Application interface is custom built on each case

Max. 60W per FPGA

ASICs

Application Specific Integrated Circuits

Fully custom designs, built for a specific application

Not flexible, cannot be changed once it is built

Development is even more specialized than FPGAs

Power consumption varies with the application, usually best performance per Watt

Slide: Guillermo Marcus

HITS

FPGA-Plattform MPRACE

GRACE

GRACE = GRAPE + RACE

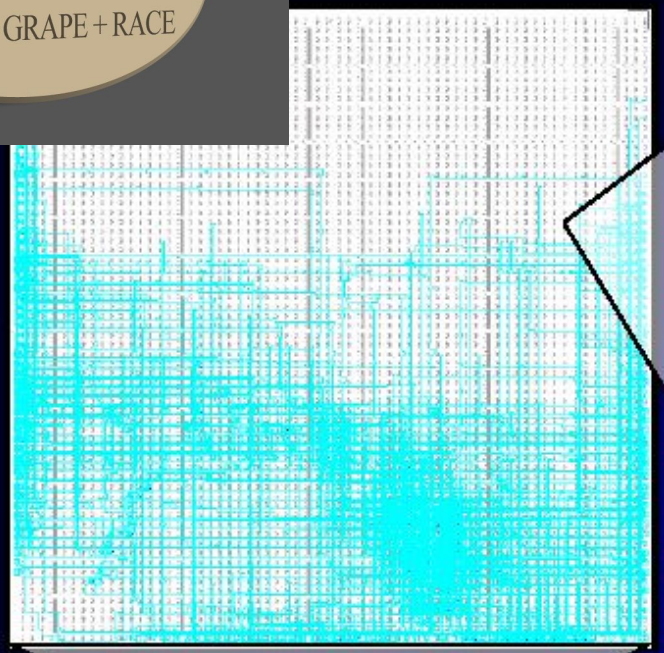

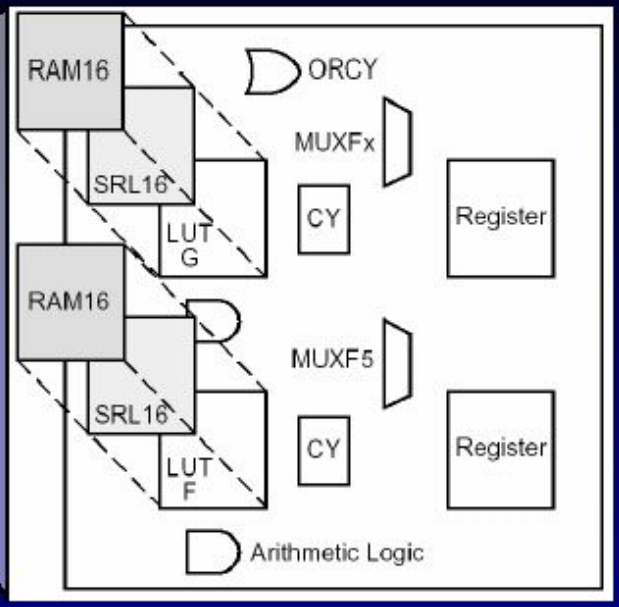

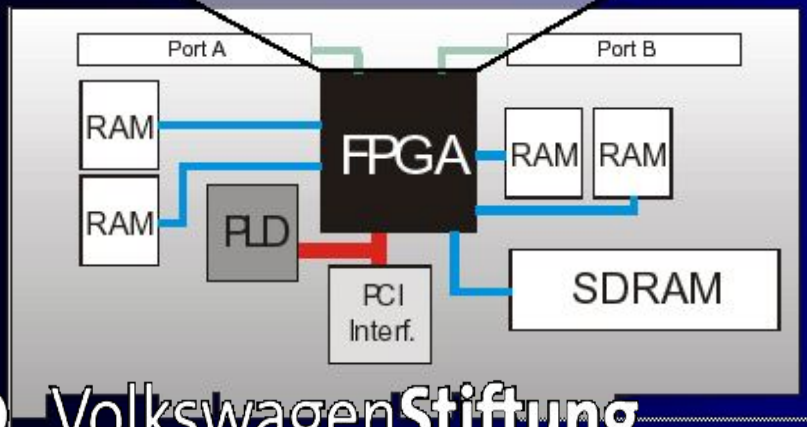

Volkswagen Stiftung

lienhart@ti.uni-mannheim.de

GPU Hardware

GeForce 680

2007: GeForce 8800 GTX, 128 SP, 768 MB
2008: GeForce 9800 GTX, 128 SP, 512 MB

2009: GeForce GTX 280, 240 SP, 1 GB
2010: GeForce GTX 480, 480 SP, 1.5 GB
2012: GeForce 680, 1400 SP, 4/6 GB

http://www.nvidia.com

http://gpgpu.org

HARDWARE

Custom PC Clusters

with manycore accelerator hardware:

Heidelberg (2x, GRAPE, FPGA, GPU)

Rochester, NY (only GRAPE)

Kiev, Ukraine (2x, GRAPE, GPU)

NAOC Beijing ***(CAS and Silk Road Project)***

Bottom Right: NAOC Cluster

86x8 Cores, 170 Tesla C1060

Top: Frontier Cluster

41x8=328 Cores, 40 Tesla GPU

Klessen, Banerjee,

Spurzem, Männer et al.

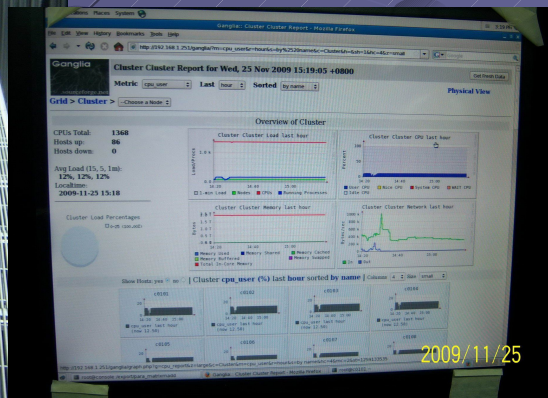

GPU Computing

PCI Express 2.0

NVIDIA TESLA with CUDA

General Purpose GPU Supercomputing (GPGPU)

http://www.nvidia.com

http://www.astrogpu.org

http://gpgpu.org

GPU

- Number of processor cores: 240
- Processor core clock: 1.296 GHz
- Voltage: 1.1875 V
- Package size: 45.0 mm × 45.0 mm 2236-pin flip-chip ball grid array (FCBGA)

Board

- Fourteen layer printed circuit board (PCB)
- PCI Express 2.0 ×16 system interface
- Physical dimensions: 4.376 inches × 10.50 inches, dual slot
- Board power dissipation: 187.8 W

External Connectors

- None

Internal Connectors and Headers

- One 6-pin PCI Express power connector
- One 8-pin PCI Express power connector
- 4-pin fan connector

NVIDIA Tesla C1060
in kolob cluster Heidelberg Univ.

Fermi-based GPU supercomputer IPE (2010.04.24)

Rpeak SP	2Pflops
Rpeak DP	1Pflops
Linpack:	207.3T (Top500 19th)
Mflops/Watt:	431 (Green500 8th)
Total RAM	17.2TB
Total VRAM	6.6TB
Total HD	360TB
Inst. Comm.	H3C GE
Data Comm.	Mellanox QDR IB
Occupied area	150 sq.m.
Weight	12.6 tons
Max Power	600kW(computing)
	200kW(cooling)
System	CentOS 5.4, PBS
Monitor	Ganglia, GPU monitor
Languages	C, C++, CUDA 3.1 , OpenCL

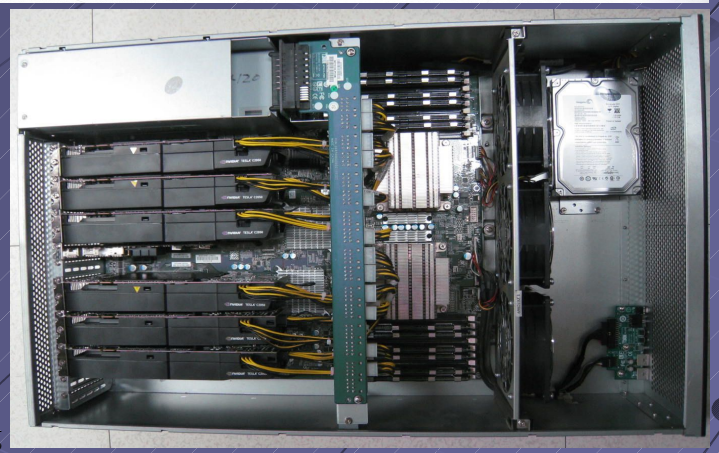

中国科学院过程工程研究所
Institute Of Process Engineering, Chinese Academy Of Sciences

Heidelberg

IPE CAS 372 node 6xC2050 cluster

2232 GPU = 2.2 Pflops SP / 1.1 Pflops DP

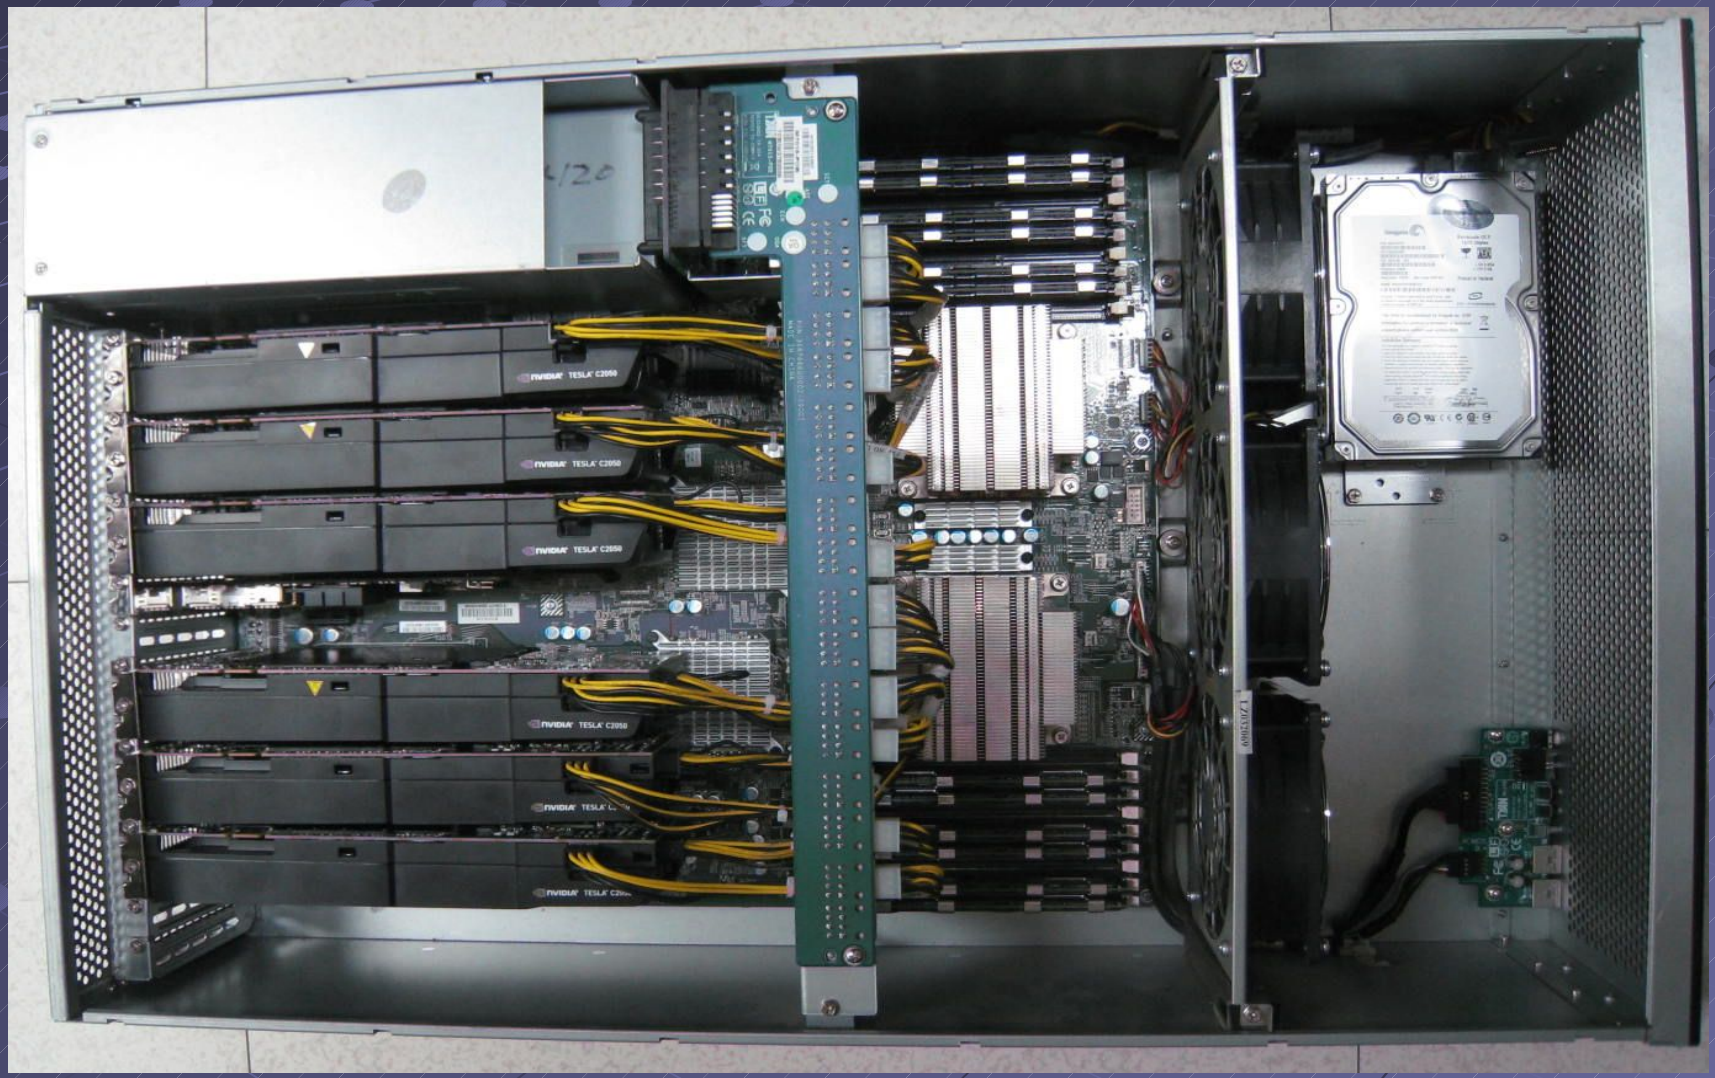

天河

“天河一号”超级计算机系统

TH-1 supercomputer

Landmark result of the important project "High Efficient Supercomputer and Grid Service Environment" supported by National 863 Program.

- Built by National University of Defense Technology, with the cooperation of National Supercomputer Center in Tianjin (NSCC –TJ) and Inspur (Beijing) Electronic Information Industry Co., Ltd.

Host system of NSCC-TJ, installed in Tianjin Binhai New Area.

A backbone node of the national grid of China.

Top 10 List June 2011

1	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect
2	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C
3	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz
4	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU
5	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows
6	Cielo - Cray XE6 8-core 2.4 GHz
7	Pleiades - SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband
8	Hopper - Cray XE6 12-core 2.1 GHz
9	Tera-100 - Bull bullx super-node S6010/S6030
10	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband

From www.top500.org - list of fastest supercomputers in the world...
... last year Nov. 2010:

China Grabs Supercomputing Leadership Spot in Latest Ranking of World's Top 500 Supercomputers

Thu, 2010-11-11 22:42

MANNHEIM, Germany; BERKELEY, Calif.; and KNOXVILLE, Tenn.—The 36th edition of the closely watched TOP500 list of the world's most powerful supercomputers confirms the rumored takeover of the top spot by the Chinese Tianhe-1A system at the National Supercomputer Center in Tianjin, achieving a performance level of 2.57 petaflop/s (quadrillions of calculations per second).

Now (June 2011) Japan gets to the top....
... with Fujitsu K Computer (8 Petaflop/s)

www.green500.org

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	1684.20	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype	38.80
2+	1448.03	National Astronomical Observatory of Japan	GRAPE-DR accelerator Cluster, Infiniband	24.59
2	958.35	GSIC Center, Tokyo Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.80
3	933.06	NCSA	Hybrid Cluster Core i3 2.93Ghz Dual Core, NVIDIA C2050, Infiniband	36.00
4	828.67	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	57.96
5	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
5	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
5	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
8	740.78	Universitaet Frankfurt	Supermicro Cluster, QC Opteron 2.1 GHz, ATI Radeon GPU, Infiniband	385.00
9	677.12	Georgia Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5660 2.8Ghz, nVidia Fermi, Infiniband QDR	94.40
10	636.36	National Institute for Environmental Studies	GOSAT Research Computation Facility, nvidia	117.15
11	635.15	National Supercomputing Center in Tianjin	NUDT YH Cluster, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C	4040.00
12	628.13	Lawrence Livermore National Laboratory	Appro GreenBlade Cluster Xeon X5660 2.8Ghz, nVIDIA M2050, Infiniband	160.00
13	555.50	CSIRO	Supermicro Xeon Cluster, E5462 2.8 Ghz, Nvidia Tesla s2050 GPU, Infiniband	94.60
14	492.64	National Supercomputing Centre in Shenzhen (NSCS)	Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	2580.00

NCSA director: GPU is future of supercomputing

by Brooke Crothers

The director of the National Center for Supercomputing Applications has seen the future of supercomputing and it can be summed up in three letters: GPU.

Thom Dunning, who directs the NCSA and the Institute for Advanced Computing Applications and Technologies at the famed supercomputing facilities on the campus of University of Illinois at Urbana-Champaign, says high-performance computing will begin to move toward graphics processing units or GPUs. Not coincidentally, **this is exactly what China has done to achieve the world's fastest speeds with its "Tianhe-1A"** supercomputer. That computer combines about 7,000 Nvidia GPUs with 14,000 Intel CPUs: the only hybrid CPU-GPU system in the world of that scale.

Thom Dunning directs the Institute for Advanced Computing Applications and Technologies and the NCSA.

"What we're really seeing in the efforts in China as well as the ones we have in the U.S. is that GPUs are what the future will look like," said Dunning in a phone interview Thursday. "What we're seeing is the beginning of something that's going to be happening all over the world."

NCSA already has a small CPU-GPU hybrid system. "It's something we have been working on for a number of years. We have a CPU-GPU cluster for the NCSA academic community. Made up of Intel CPUs and Nvidia GPUs. A 50 teraflop machine," he said. (Note that **Oak Ridge National Laboratories is also installing a hybrid system now**.)

Speedups using GPU vs. CPU

146X

Interactive visualization of volumetric white matter connectivity[1]

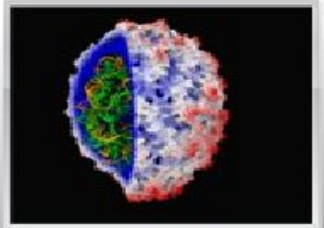

36X

Ionic placement for molecular dynamics simulation on GPU[2]

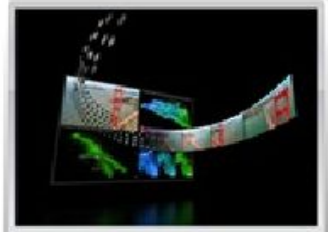

18X

Transcoding HD video stream to H.264 for portable video[3]

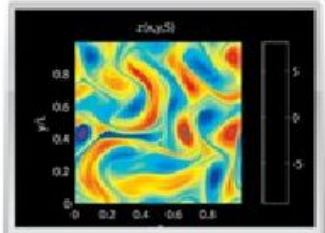

17X

Simulation in Matlab using .mex file CUDA function[4]

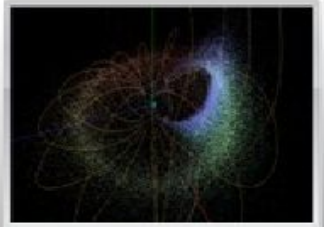

100X

Astrophysics N-body simulation[5]

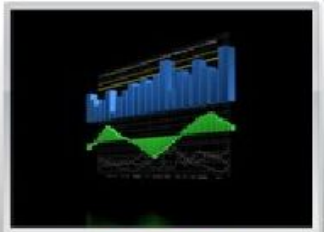

149X

Financial simulation of LIBOR model with swaptions[6]

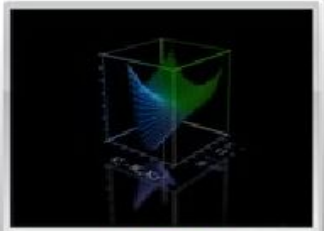

47X

GLAME@lab: M-script API for linear Algebra operations on GPU[7]

20X

Ultrasound medical imaging for cancer diagnostics[8]

24X

Highly optimized object oriented molecular dynamics[9]

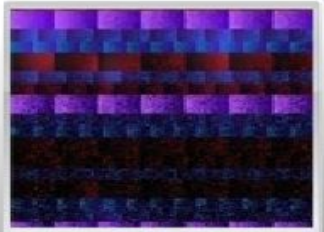

30X

Cmatch exact string matching - find similar proteins & gene sequences[10]

Hardware

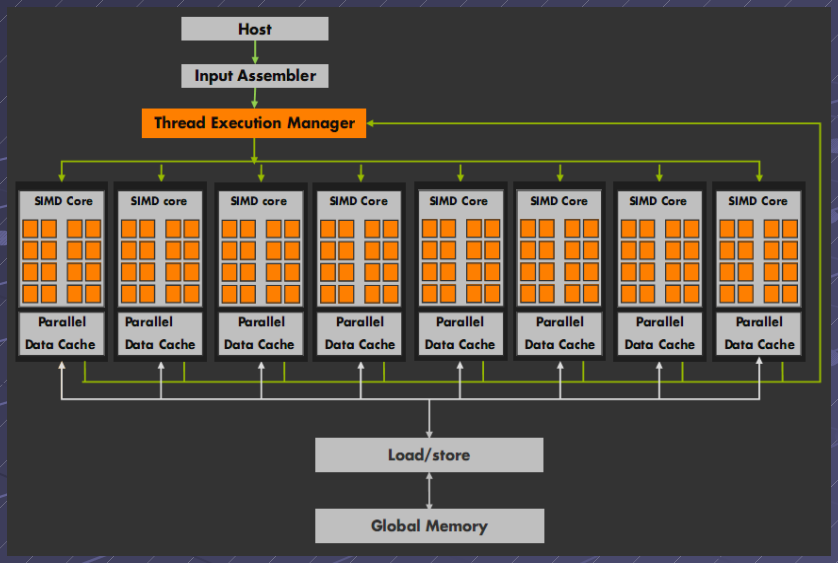

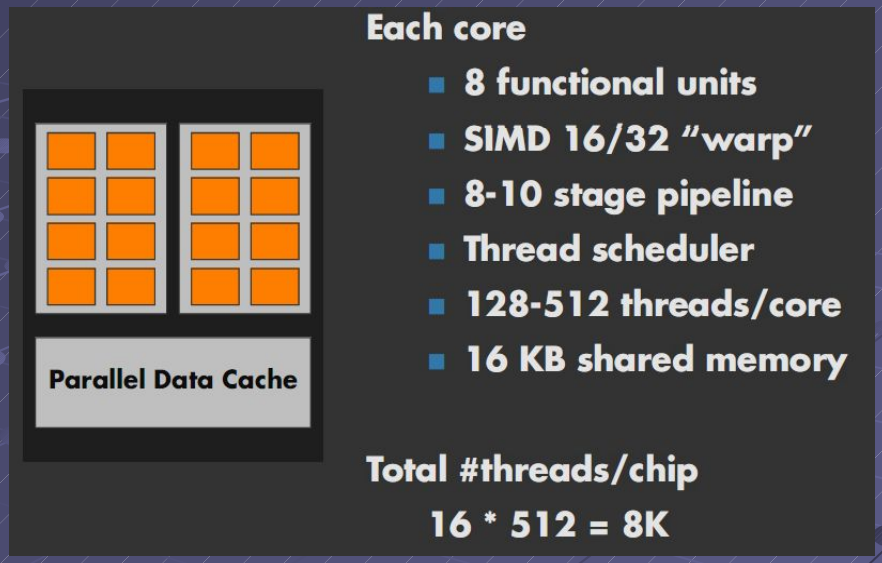

NVIDIA Tesla C1060 :

1.3 GHz *240 processors * 3 flopclock = 930 Gflops

CUDA

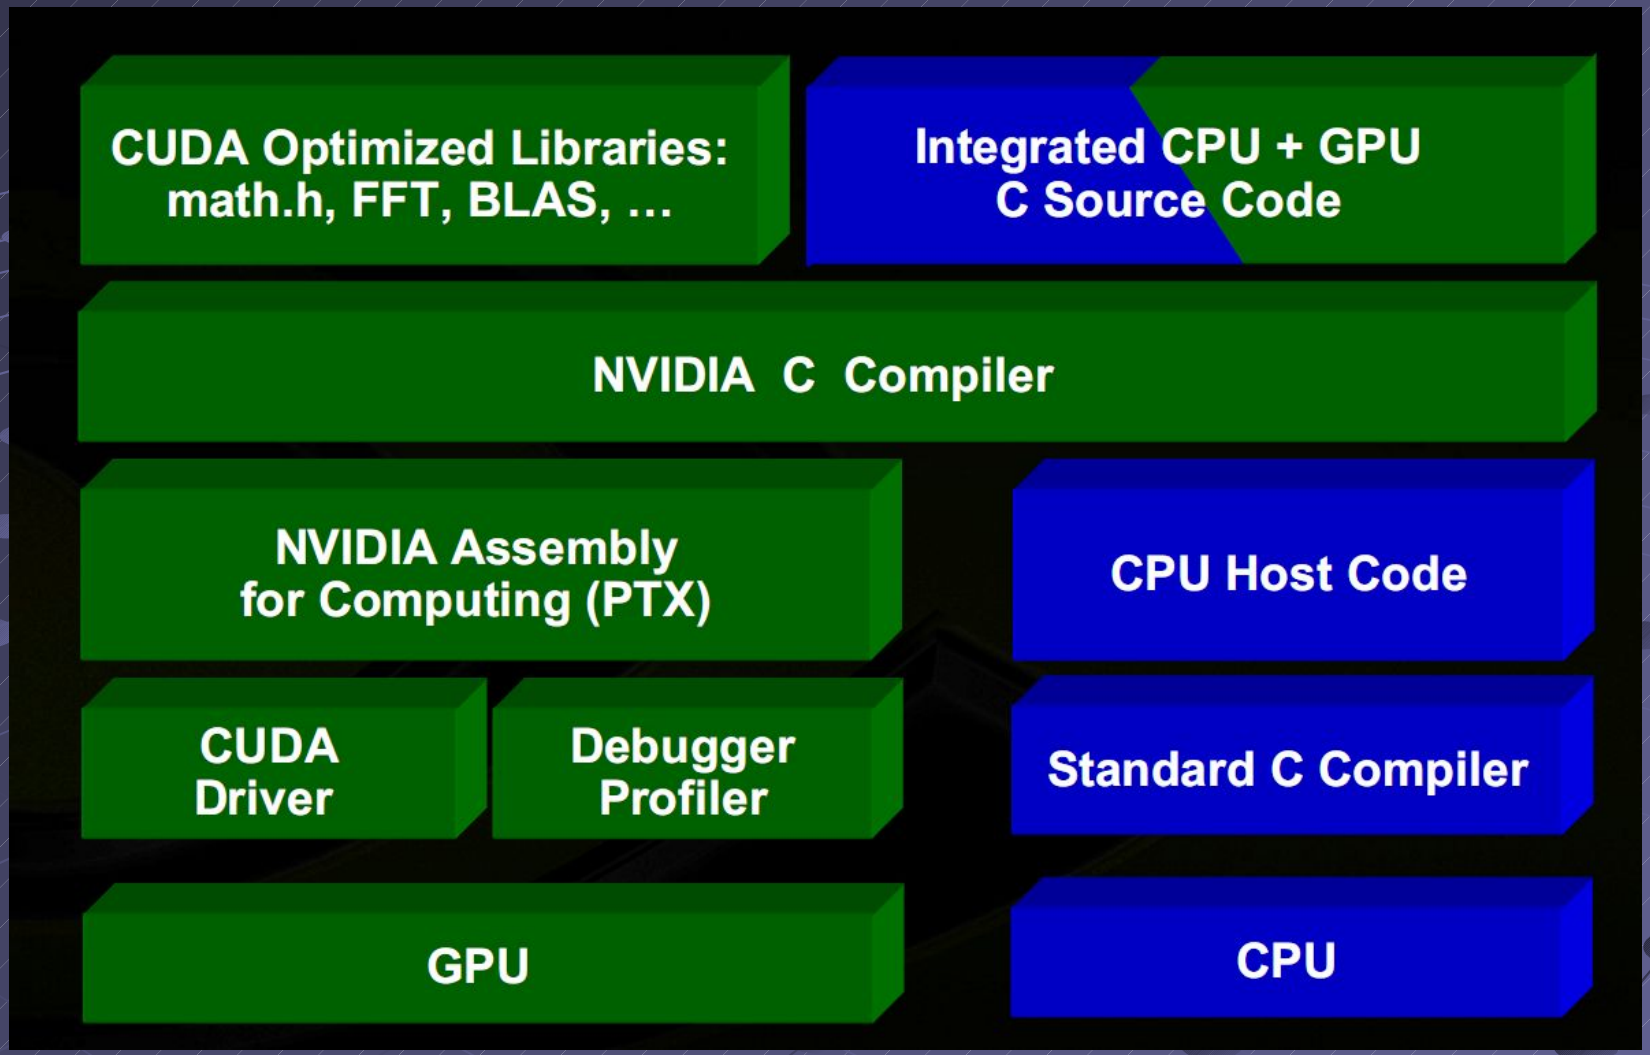
CUDA Optimized Libraries:
math.h, FFT, BLAS, …
Integrated CPU + GPU
C Source Code
NVIDIA C Compiler
NVIDIA Assembly
for Computing (PTX)
CPU Host Code
CUDA
Driver
Debugger
Profiler
Standard C Compiler
GPU
CPU

Einführung

Computer-Physik

- Traditionell:

  Theorie **Experiment/Beobachtung**

- Modern:

  Theorie 

  Beobachtung

Überprüfung von Theorien!

Vorhersage von Beobachtungen!

Computerphysik

http://www.fz-juelich.de/nic/Publikationen/Broschuere/Broschuere-d.html

- Projekte am Jülich Supercomputing Centre (JSC)

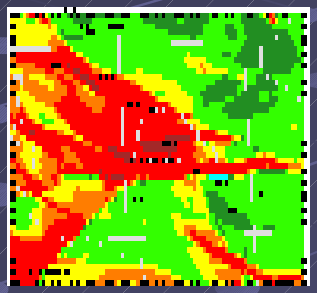

Umwelt

Elementar-Teilchen

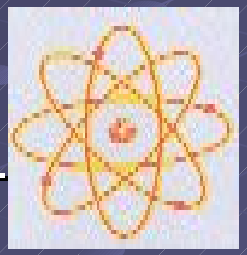

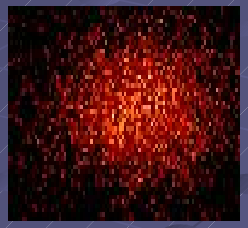

Astrophysik

Chemie

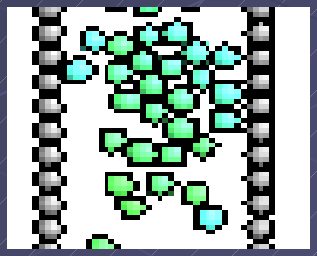

Vielteilchen

Polymere

http://www.fz-juelich.de/ias/jsc/DE/Forschung/forschung_node.html

Forschung

Die Forschungs- und Entwicklungsarbeiten im JSC konzentrieren sich auf mathematische Modellierung und numerische, parallele Algorithmen für verschiedene Fachgebiete. In der Informatik liegt der Schwerpunkt auf Cluster-Computing, Leistungsanalyse paralleler Programme, Visualisierung und Grid-Computing.

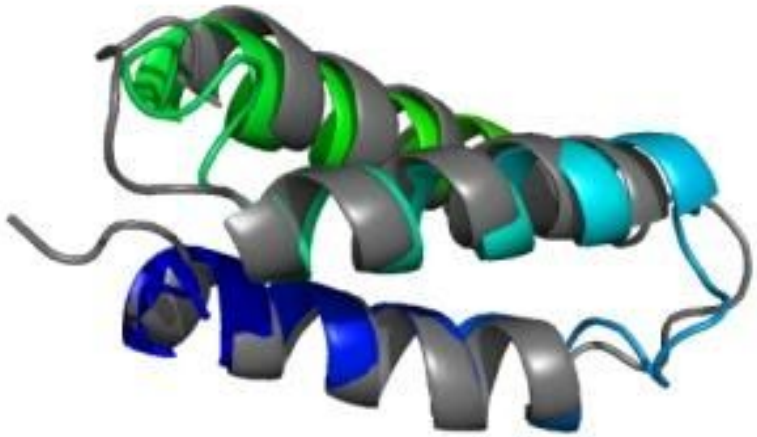

Modellierung und Simulation

Das JSC beschäftigt sich mit wissenschaftlichen Simulationen auf Supercomputern und der Entwicklung bzw. Weiterentwicklung der zugrundeliegenden mathematischen Modelle. Beispiele für Anwendungen sind:

- Computergestützte Plasmaphysik
- Proteinfaltung
- Quanteninformationsverarbeitung
- Fußgängerdynamik

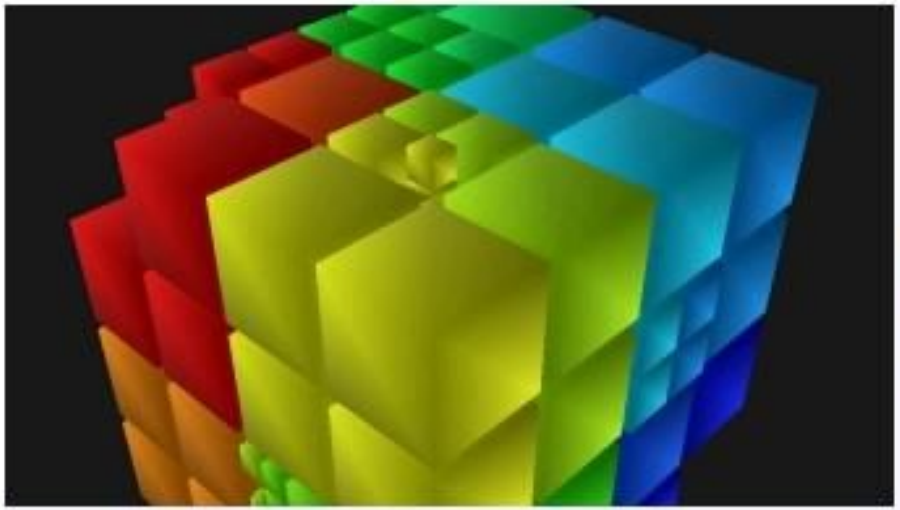

Algorithmen und Methoden

Für effiziente Simulationen braucht man leistungsfähige Algorithmen und Methoden. Am JSC steht die Entwicklung der folgenden Methoden im Mittelpunkt: → Mehr

- Eigenwert-Löser
- Coulomb-Löser
- Paralleles I/O

Computerphysik - Astrophysik

Kosmologie

Computerphysik - Astrophysik

- Strukturbildung im Universum

Im Jahre 100.000....

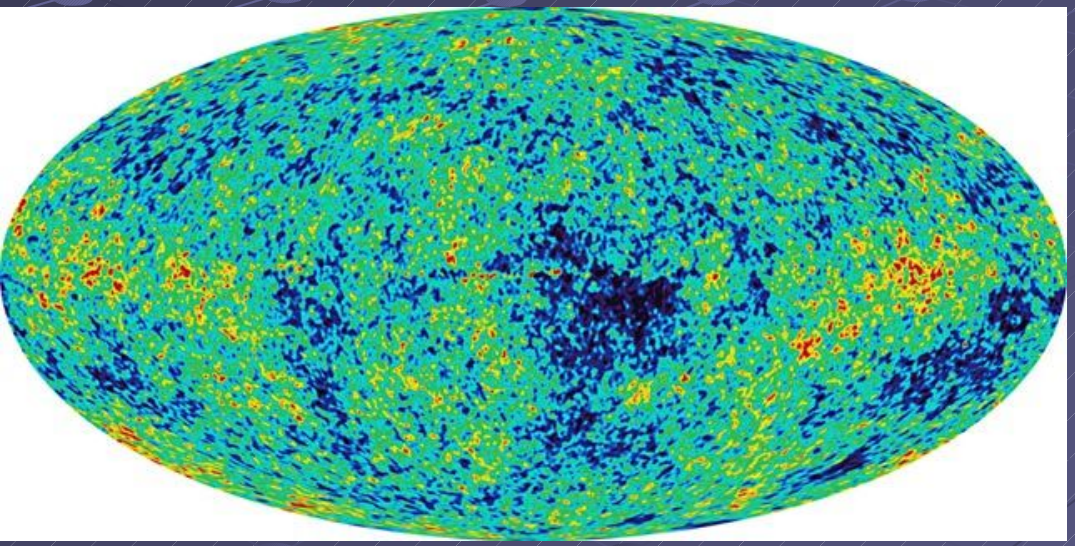

Wilkinson Microwave Anisotropy Probe (WMAP)

(Kosmischer Mikrowellenhintergrund)

...und ``heute´´

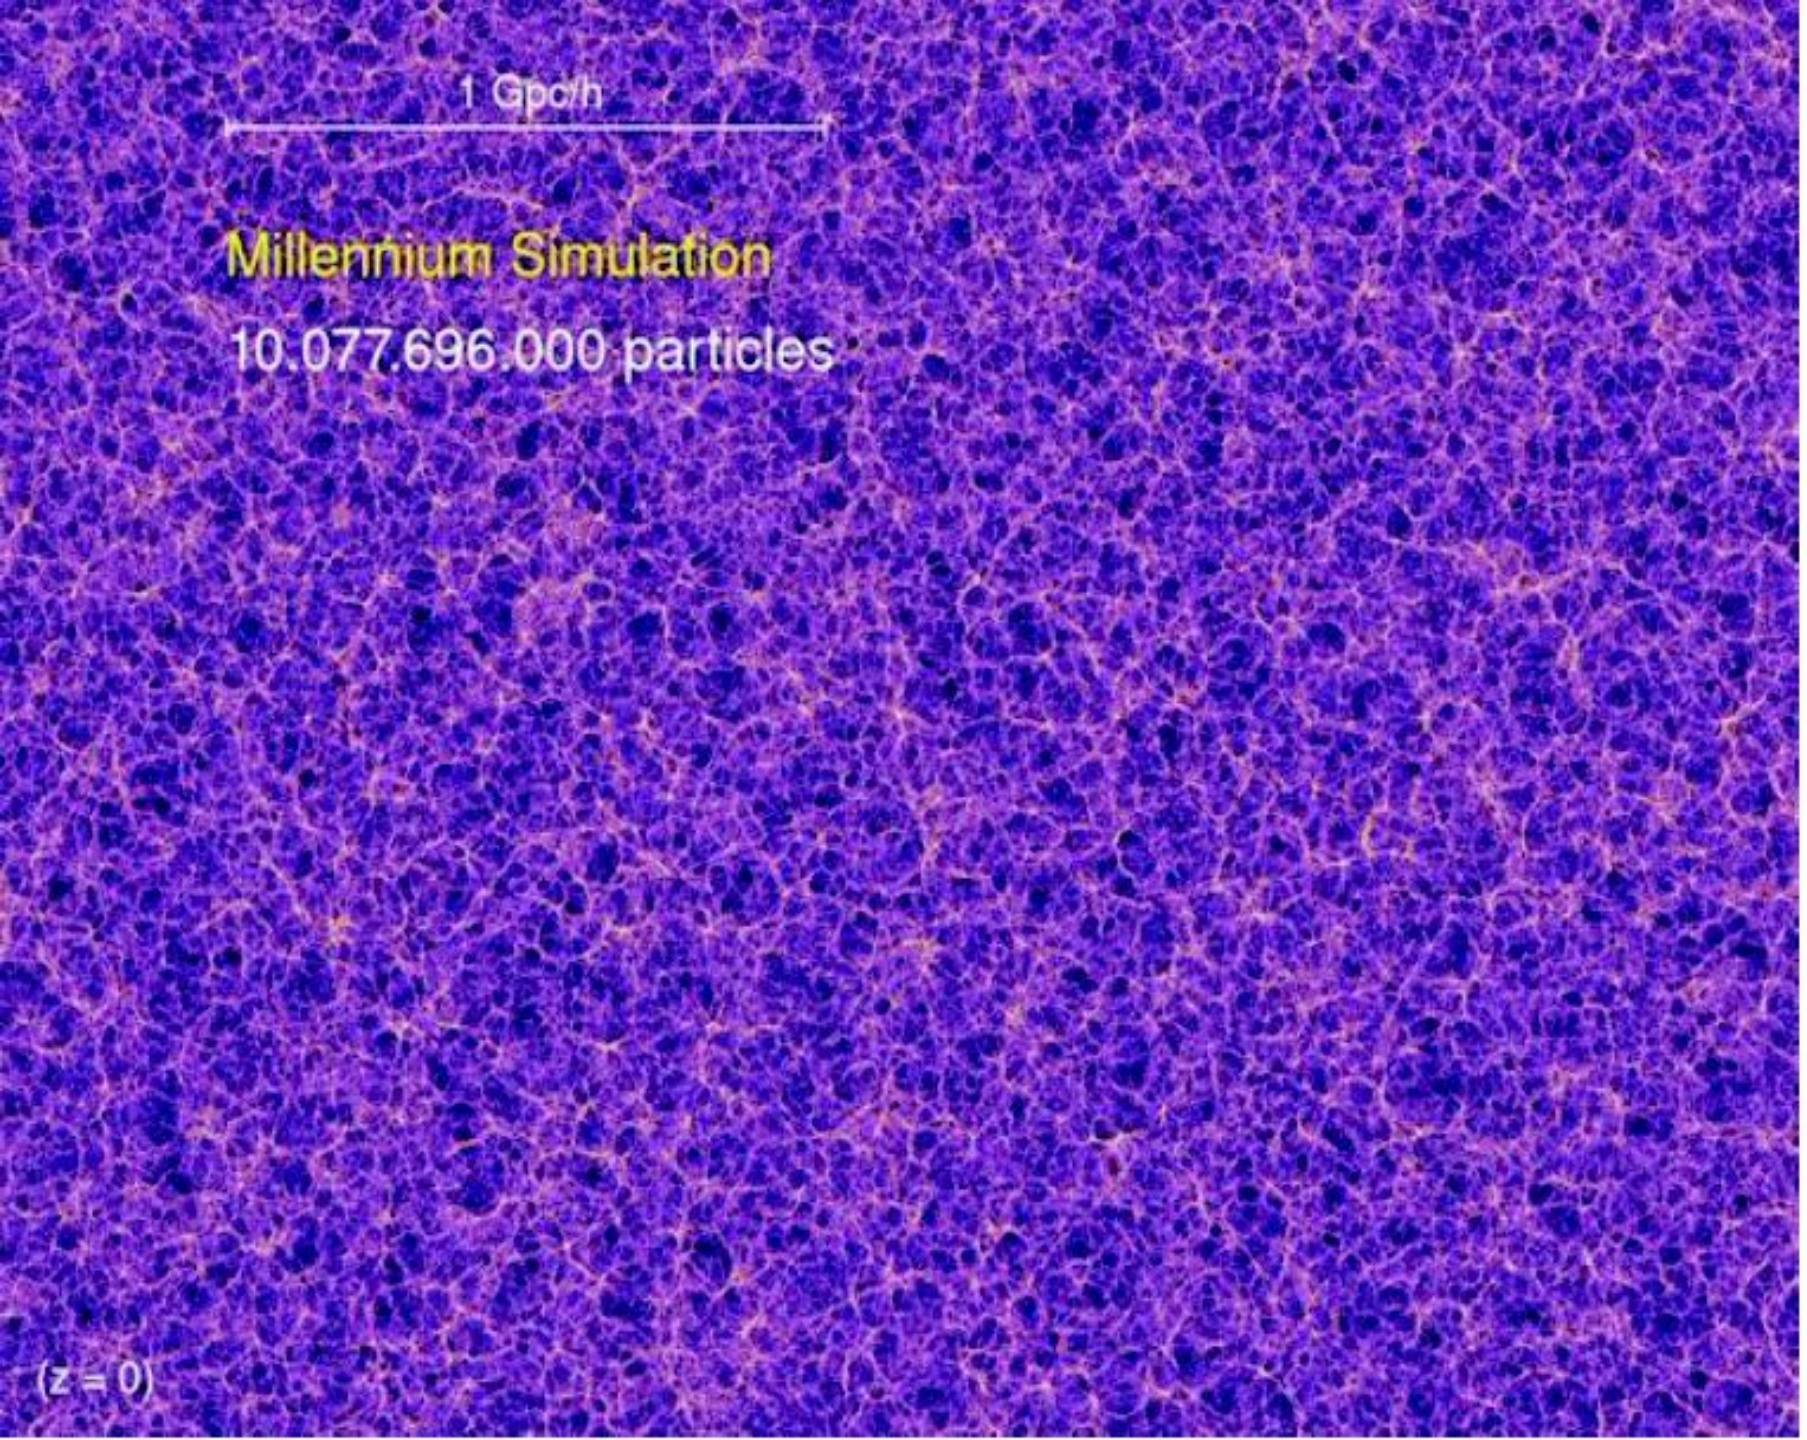

Millenium Simulaiton (Springel et al.)

Millenium Simulaiton (Springel et al.)

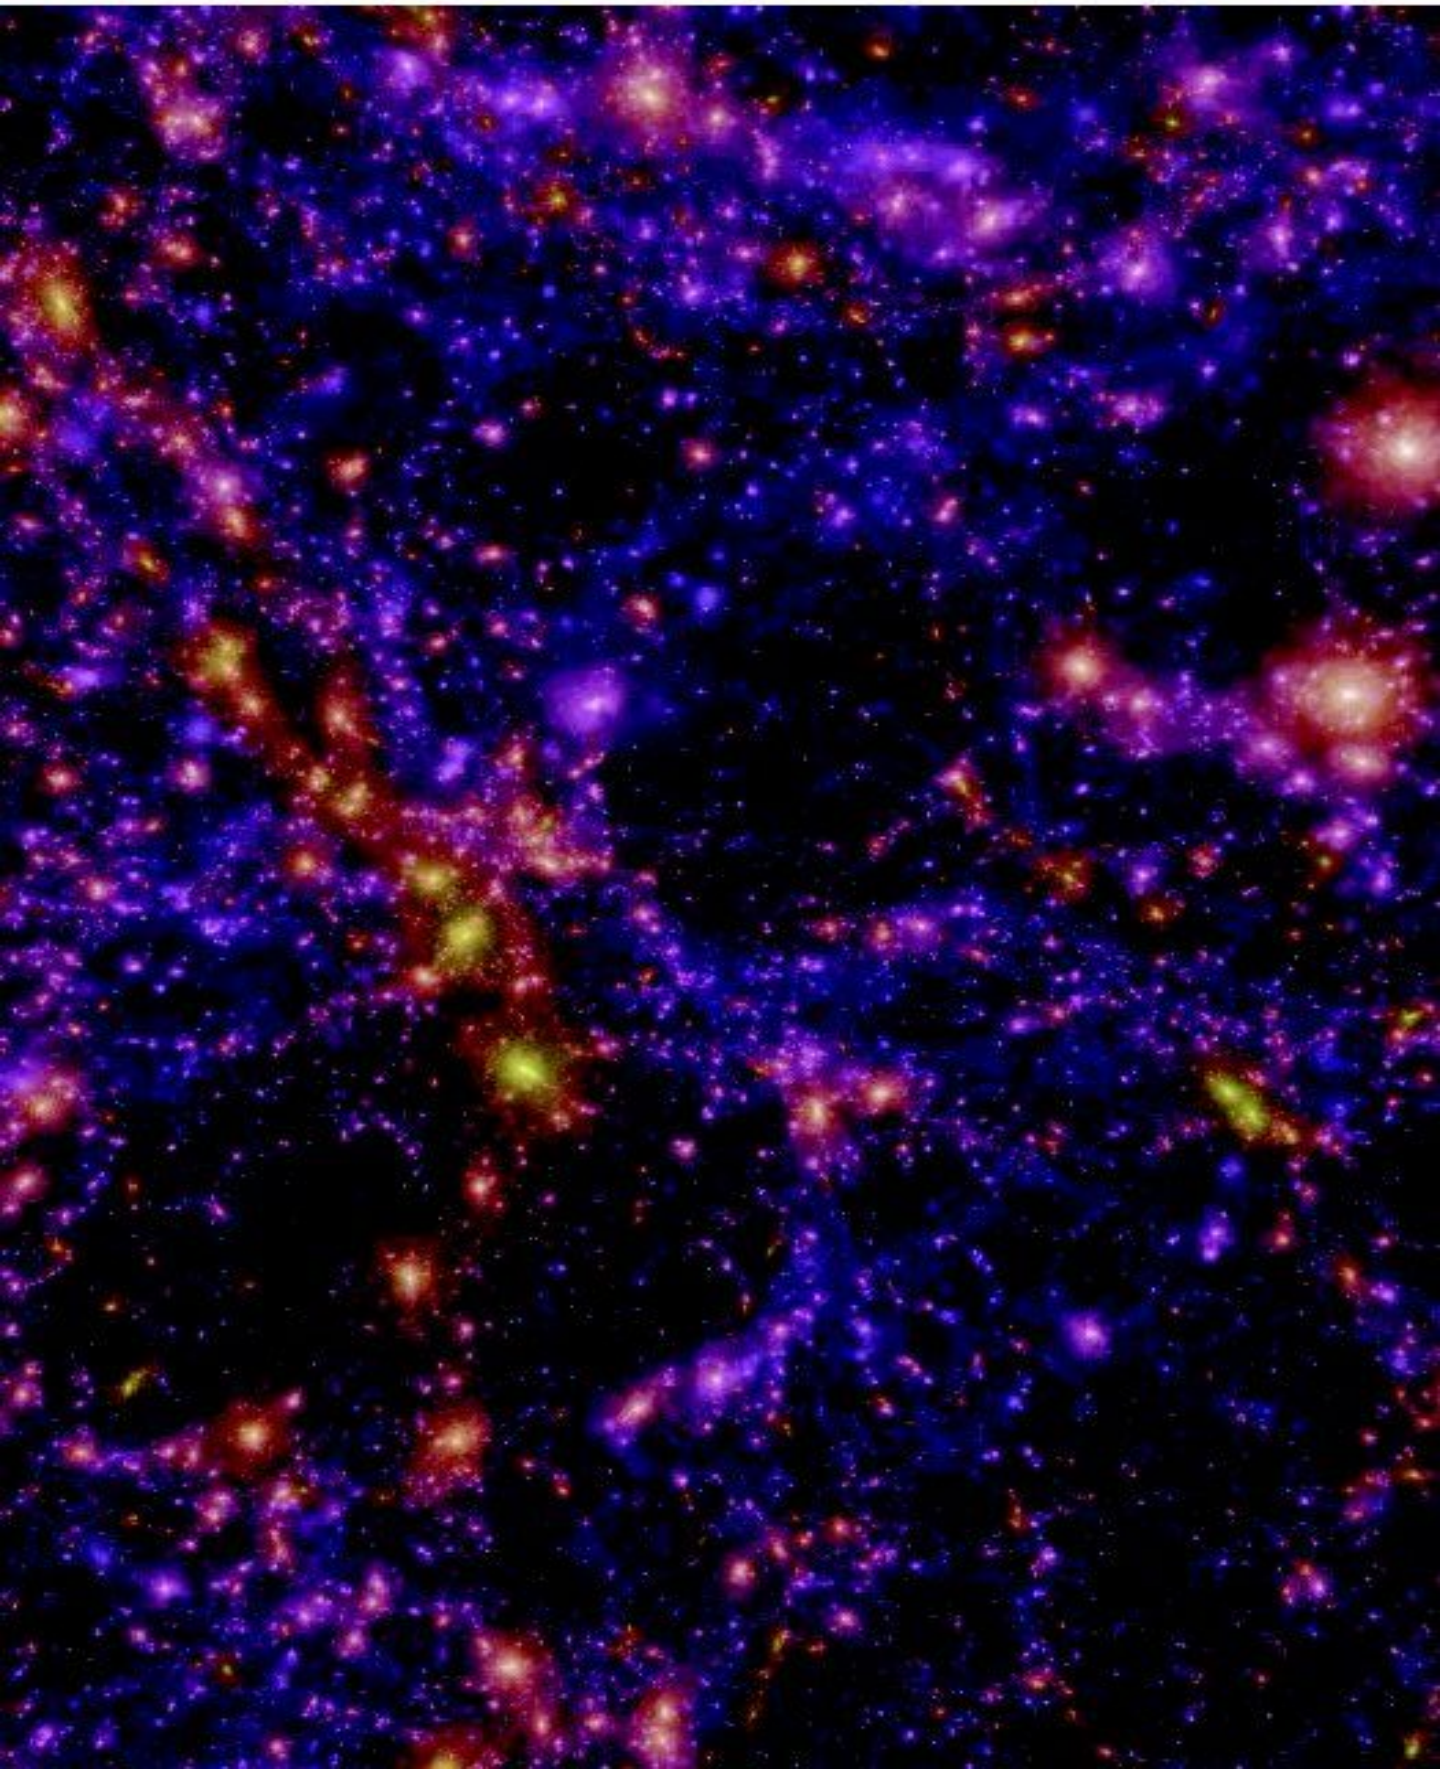

Challenges: mehr Physik

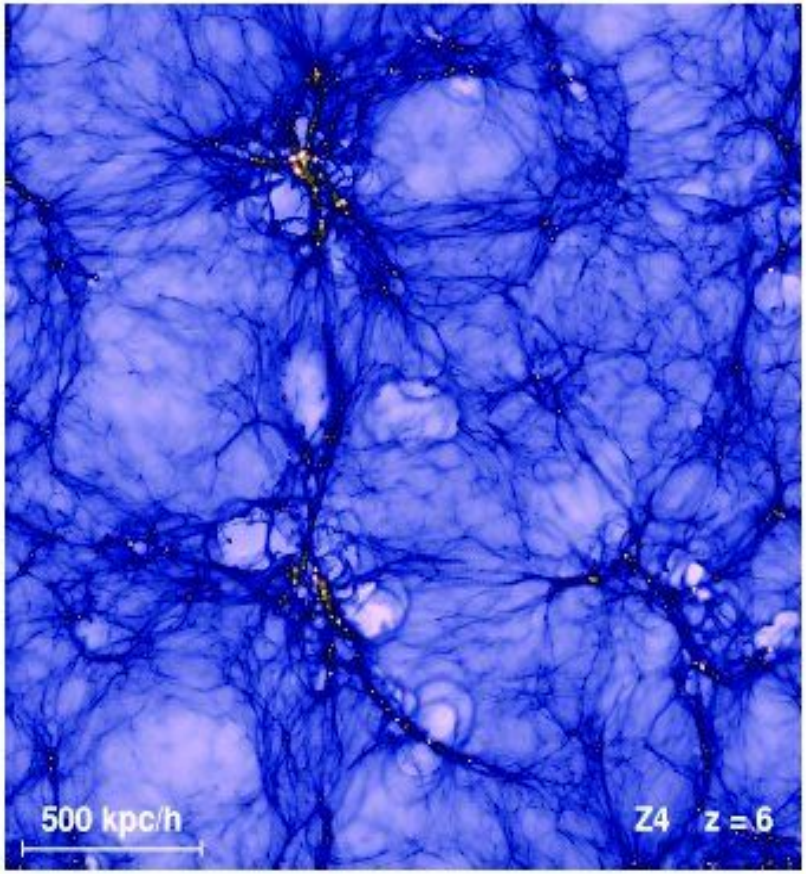

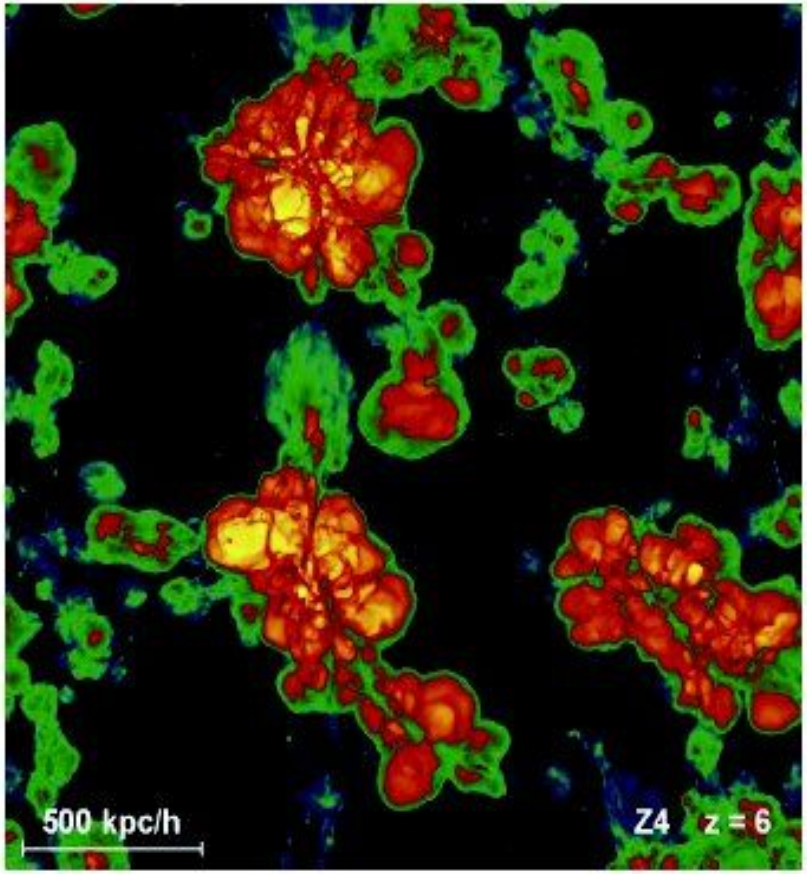

Hot bubbles of ionized gas around small primordial galaxies (Springel & Hernquist)

Challenges: mehr Physik

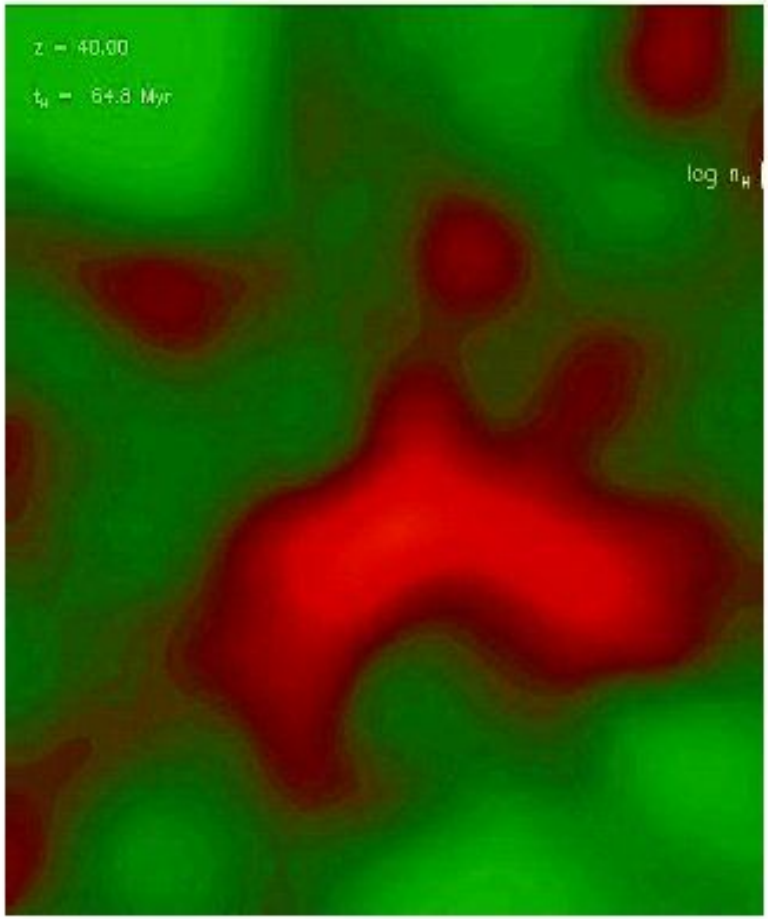

From our own group: Formation of the first galaxy with gas cooling and chemistry (Greif et al. 2008)

Computerphysik - Astrophysik

Galaxien

Galaxien-Entwicklung

NGC1232:

Nucleus
Bulge
Disk
Wave Pattern
Star Formation
Satellites

Computerphysik - Astrophysik

Modelle von John Dubinski, Canadian Institute for Theoretical Astrophysics

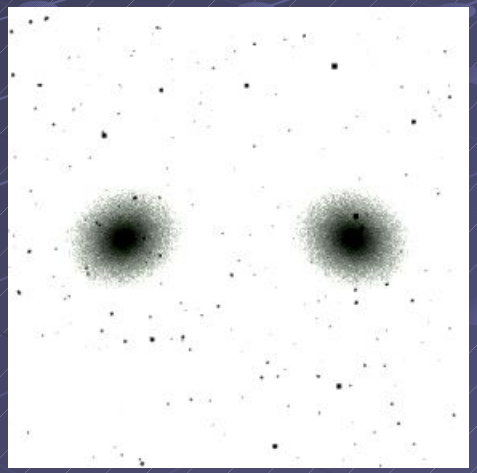

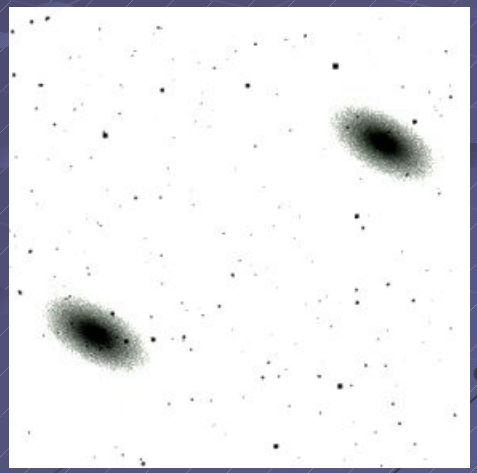

Wie wird's gemacht? -> Vorlesung!

Computerphysik - Astrophysik

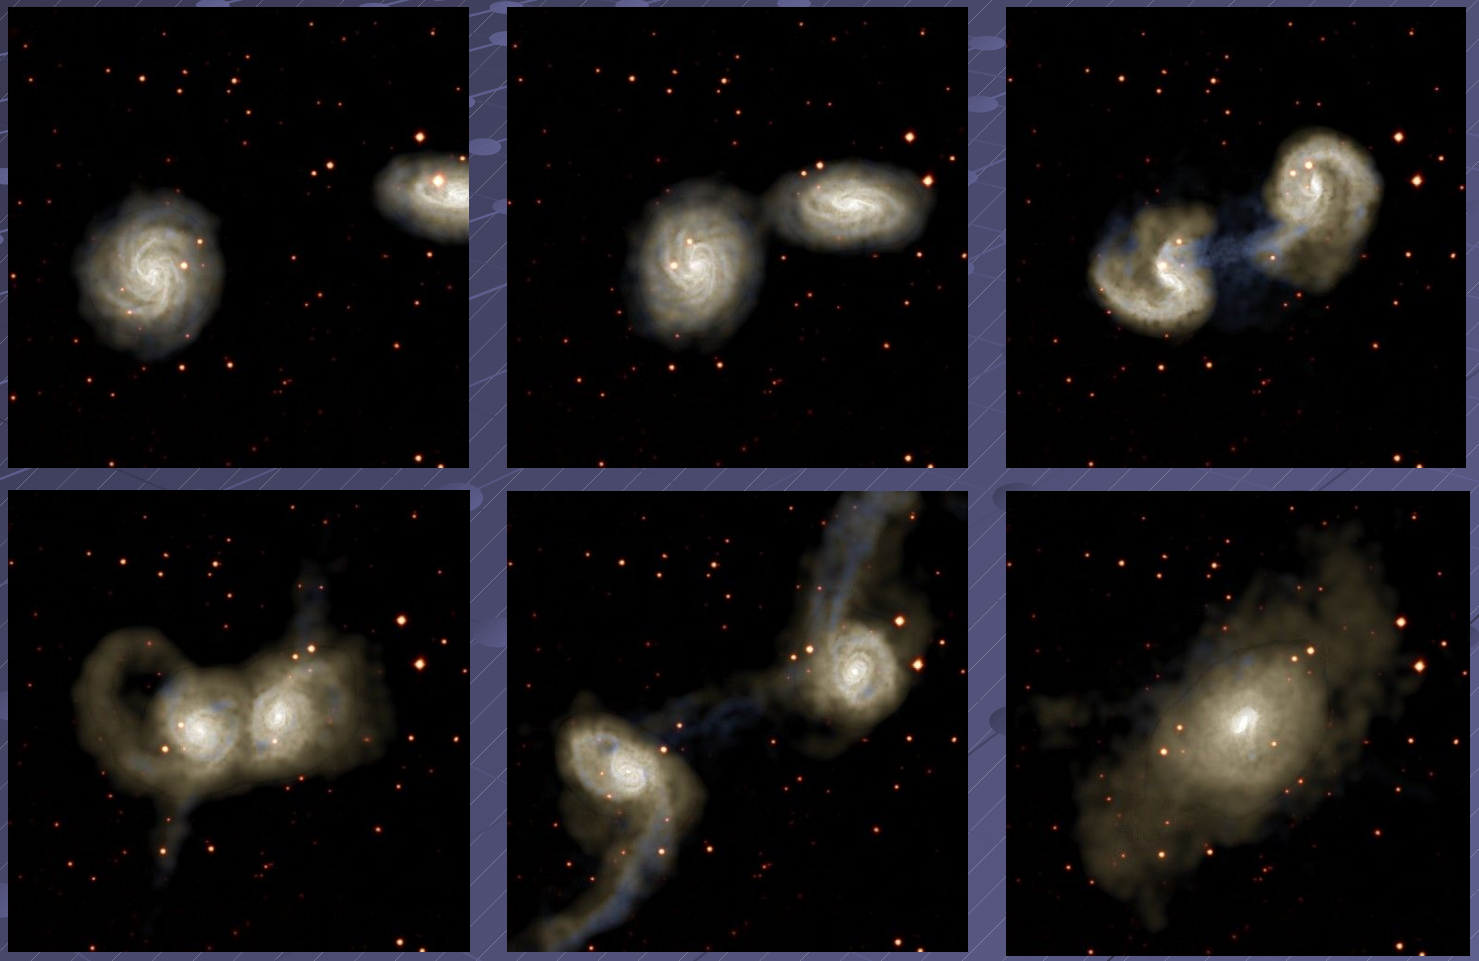

Gadget

(see webpage Of Virgo)

(By V. Springe

Gas Dynamics Also treated By Particles (SPH)

Computerphysik - Astrophysik

- **Theorie:**
  Differentialgleichung
- **Modell:**
  Diskretisierung
- **Simulation:**
  Zeitablauf
- **Ergebnisse:**
  Vergleich mit Beob-achtung oder Theorie

$$\vec{a}_0 = \sum_j G m_j \frac{\vec{R}_j}{R_j^3} \quad ; \quad \vec{a}_0 = \sum_j G m_j \left[\frac{\vec{V}_j}{R_j^3} - \frac{3(\vec{V}_j \cdot \vec{R}_j)\vec{R}_j}{R_j^5} \right]$$

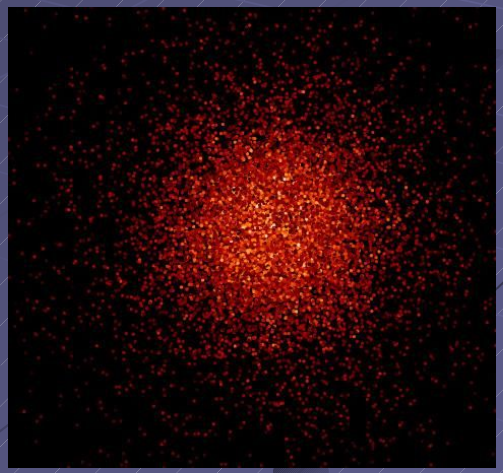

Computerphysik - Relativitätstheorie

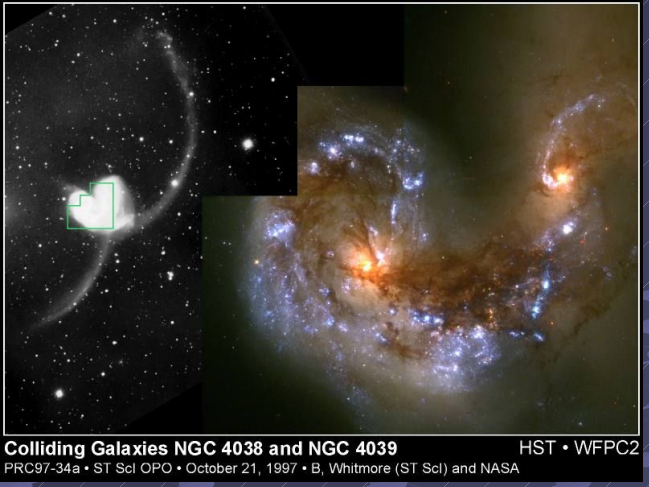

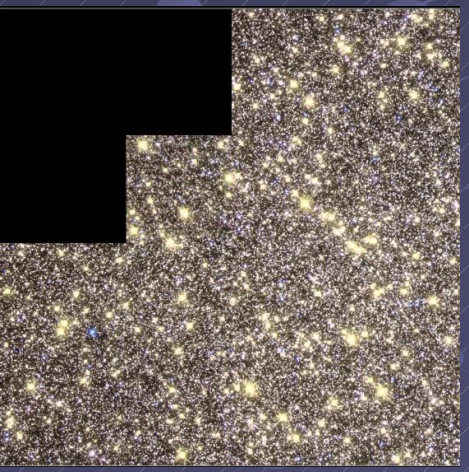

Verschmelzung
Massereicher
Schwarzer
Löcher
Gravitationswellen

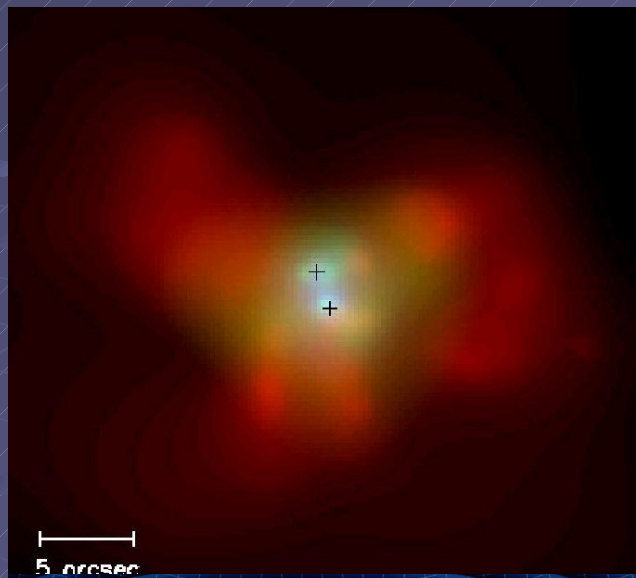

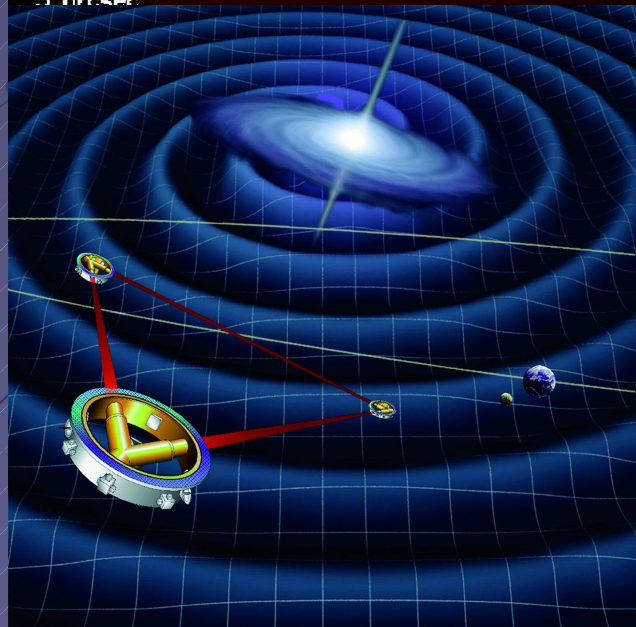

Direct Proof

still waiting …..

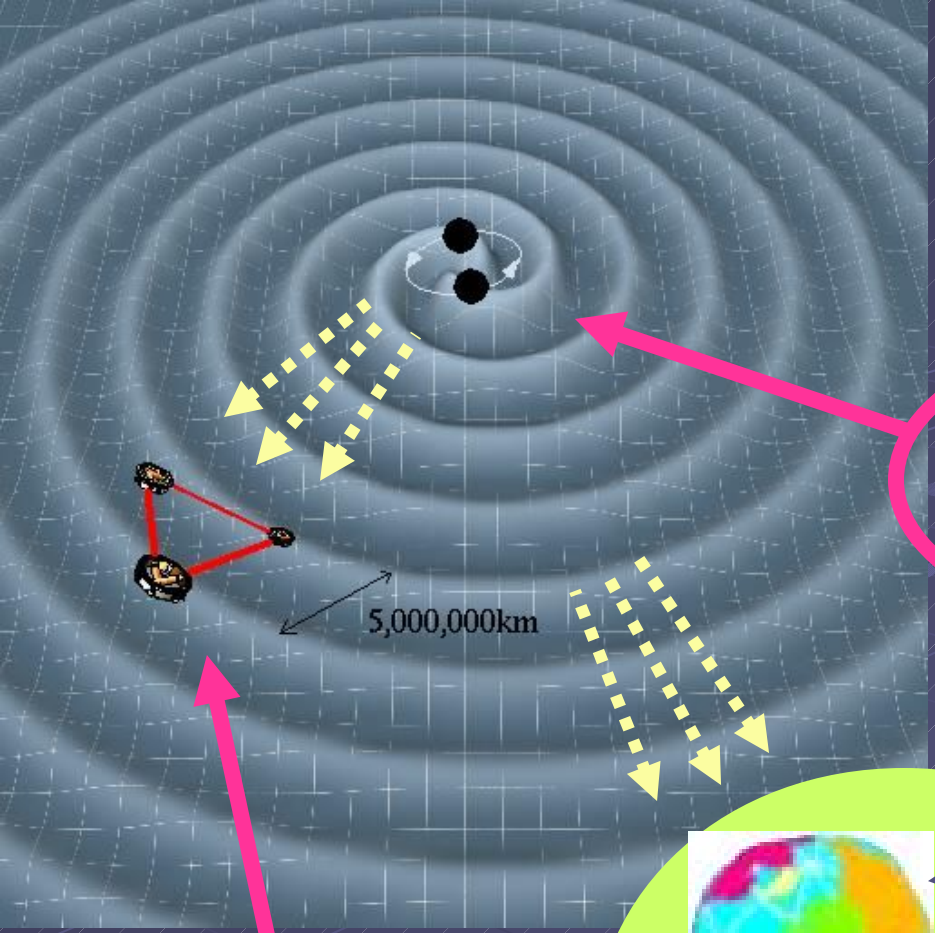

Astrophysical Sources

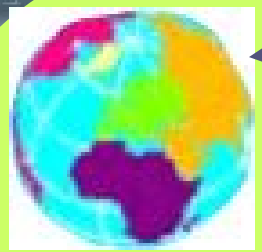

Terrestrial Detectors

VIRGO, LIGO, TAMA, AIGO

Space detectors

LISA

Example: VIRGO Detector in Cascina near Pisa, Italy

Computerphysik - Relativitätstheorie

VIRGO – Pisa 3km
LIGO – Livingston, LA
Hanford, WA
1km
GEO600 – Hannover
600m
AIGO – Australien
(planned, 5 km)

http://www.ligo-la.caltech.edu/
http://www.ego-gw.it
http://www.geo600.uni-hannover.de

Outreach to 50 Millionen light years (Neutron Stars)

PENNSTATE

Gravitationswellen-Spur der Verschmelzung Zweier Schwarzer Löcher – direkte Computer -Lösung der Einstein‘schen Feldgleichungen!

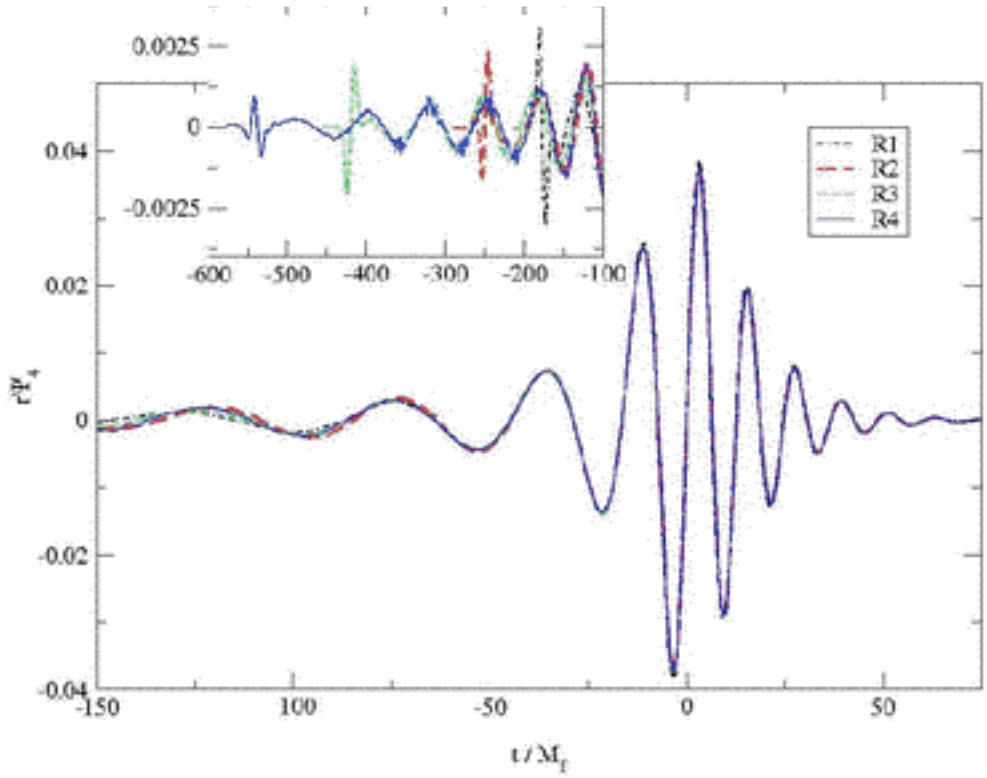

Initial separations:
R1 = 6.5 M
R2 = 7.6 M
R3 = 8.5 M
R4 = 9.6 M

NASA-GSFC
Baker, Centrella, Choi, Koppitz, van Meter
Phys.Rev. D73 (2006) 104002

3-Body Problems

Burrau‘s Problem

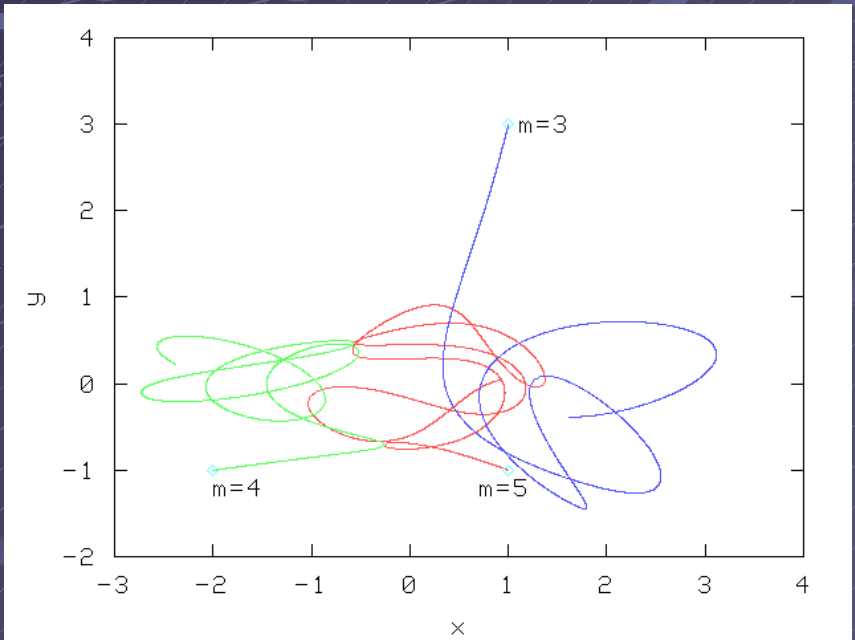

„The Eight“

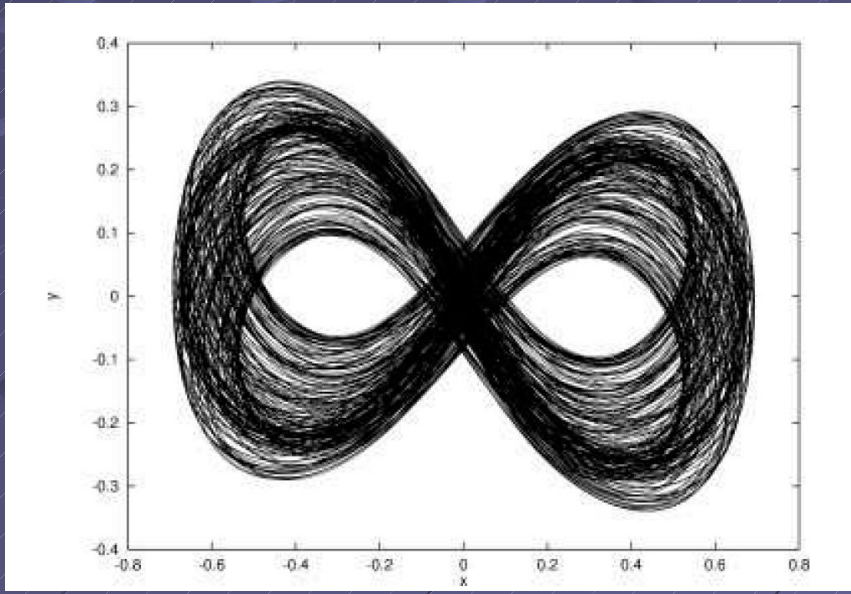

Lorenz-Attraktor

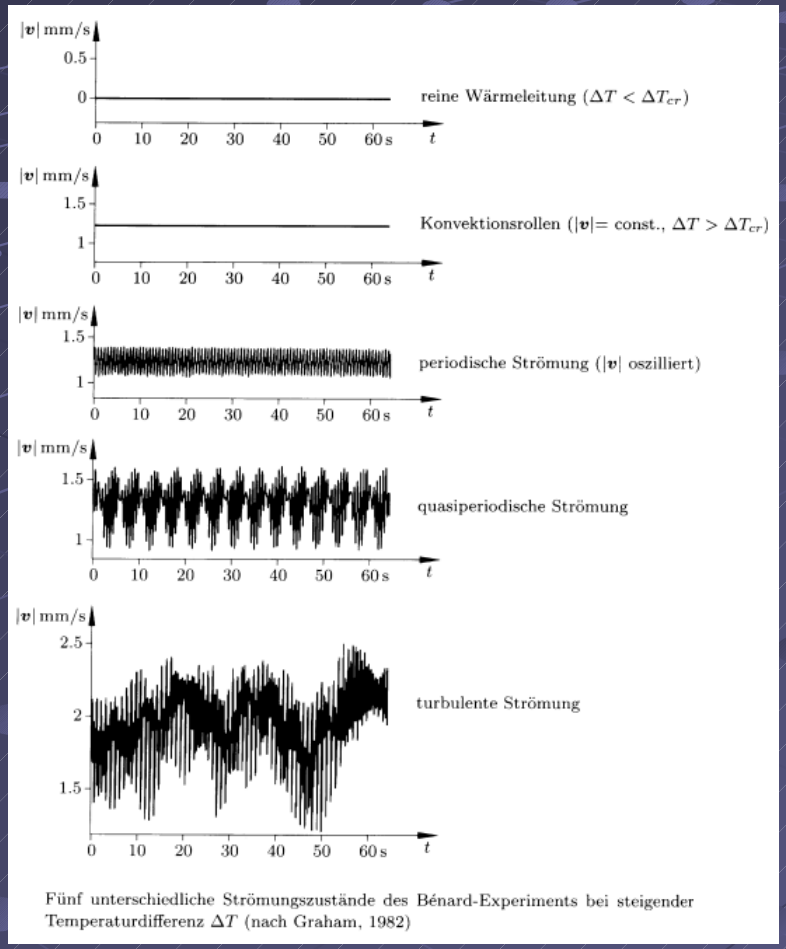

Fünf unterschiedliche Strömungszustände des Bénard-Experiments bei steigender Temperaturdifferenz ΔT (nach Graham, 1982)

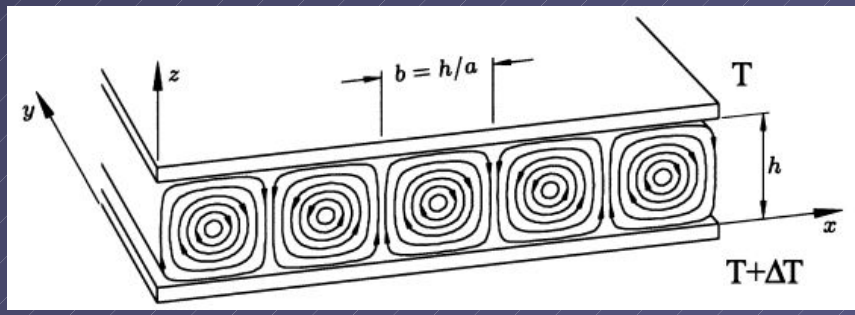

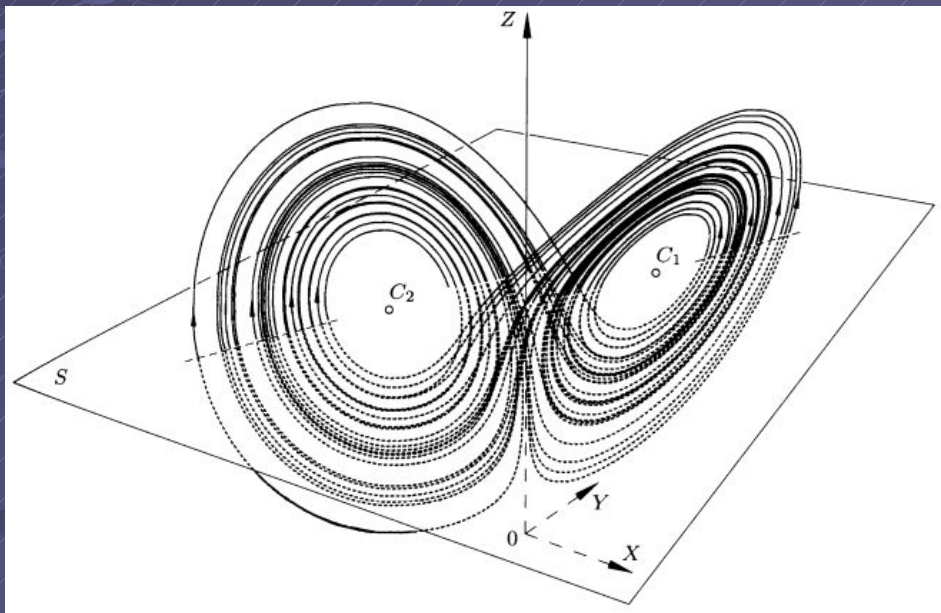

Lorenz-Attraktor für $r = 28, \sigma = 10$ und $b = 8/3$.
Der Trajektorienbereich, den die Ebene $Z = r - 1 = 27$ verdeckt, ist punktiert (Lanford, 1977)

Logistische Abbildung

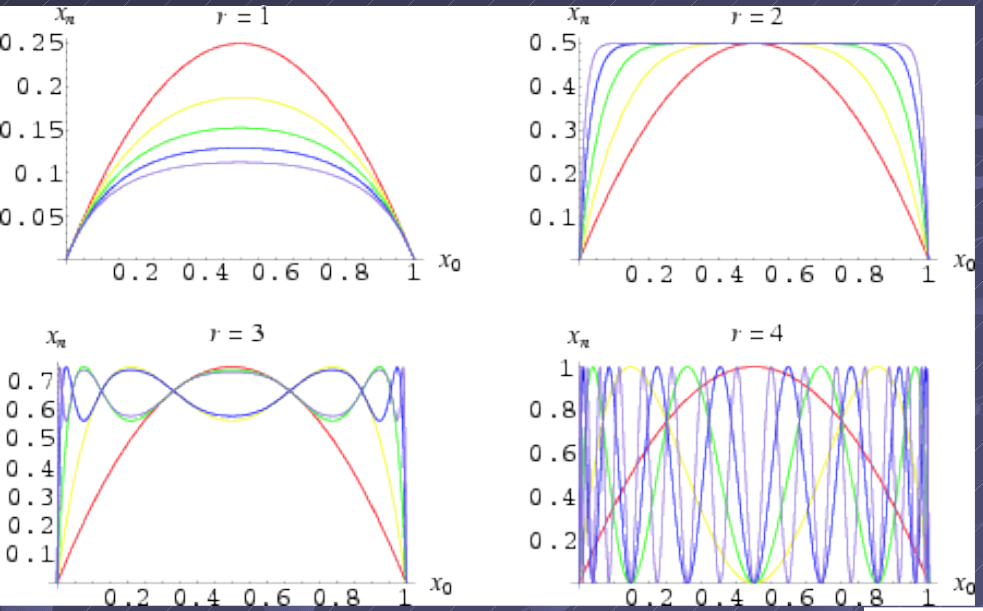

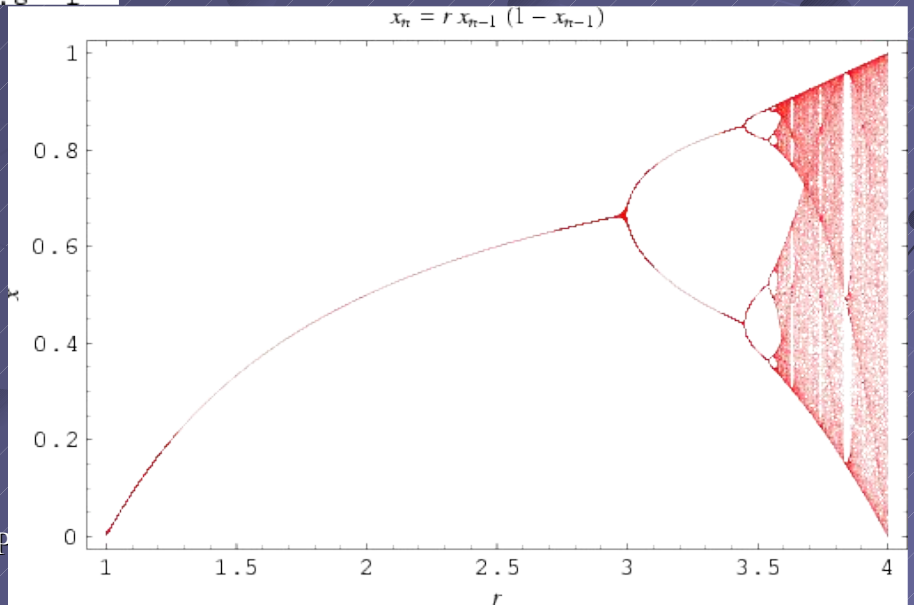

Computational Physics

* Physical Problems – Numerical Algorithms
* Numerical Solution of Ordinary Diff. Equations
* Linear Algebra
* Non-Linear Dynamics and Chaos
* Monte Carlo Methods