



**More than 10,000 Itanium Cores**  
**High-Performance Computing at Leibniz Computing Centre**  
**Munich, Germany**



LRZ High Performance Systems Department



www.lrz.de

## National High-End System: SGI Altix 4700



© 2007 LRZ

Overall characteristics for both installation phases:	Phase 1 (until 03/2007)	Phase 2 (since 04/2007)
Total number of cores	4096	9728
Peak Performance (entire system)	26.2 TFlop/s	62.3 TFlop/s
Linpack Performance	24.5 TFlop/s	<b>&gt; 56 TFlop/s</b>
LRZ-Benchmark Performance	8.2 TFlop/s	> 14.5 TFlop/s
Size of memory (entire system)	17.5 TByte	39 TByte
Direct Attached Disks	300 TByte	600 TByte
Network Attached Disks	40 TByte	60 TByte
Processor type	Intel Itanium2 Madison 9M	Intel Itanium2 Montecito Dual Core
Clock rate	1.6 GHz	1.6 GHz
L3 Cache (per core)	6 MByte	9 MByte
Memory per core	4 GByte	4 GByte per Core
Clock rate of frontside bus (FSB)	533 MHz	533 MHz
Peak bandwidth to local memory	8.5 GByte/s per core	8.5 GByte/s (now shared between 2 or 4 cores)



www.lrz.de

## Linux Cluster for regional users

### **SGI Altix 3700 Bx2**

128 Itanium2 Madison cores  
(6 MB L3, 1.6 GHz,  
2 cores share 6.4 GB/s memory bus)  
0.8 TFlop/s peak performance  
0.5 TByte main memory

#### **Interconnect**

Numalink 4  
Bisection Bandwidth > 50 GB/s  
MPI Latency 3.1  $\mu$ s

#### **Filesystem**

10 TByte XFS scratch disk storage

### **IA64**

83 nodes  
200 processor cores  
(3 MB L3, 1.6/1.3 GHz)  
1.2 TFlop/s peak performance  
0.9 TByte main memory

#### **Interconnect**

Gigabit Ethernet (Myrinet)  
Parastation MPI  
MPI Latency 14 (10)  $\mu$ s

#### **Filesystem**

50 TByte Lustre parallel file system

### **EM64T/x86\_64**

26 nodes  
72 processor cores  
0.4 TFlop/s peak performance  
0.2 TByte main memory

### **IA32**

130 nodes  
138 processor cores  
0.8 TFlop/s peak performance  
0.3 TByte main memory






www.lrz.de

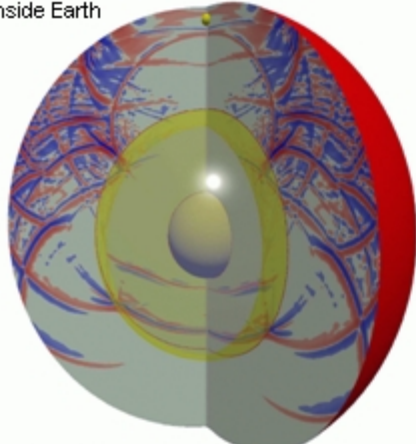
## SGI Altix 4700: Scientific Areas of Usage

The Leibniz Computing Centre (LRZ) is a national supercomputing center. We offer our services to academic researchers throughout Germany. The SGI Altix and the Linux Cluster are accessible via UNICORE and Globus to grid users from Europe. Our customers are from many different fields of research. As a consequence, there is not only one „standard user“; instead, we have to satisfy a large spectrum of requirements.

Earthquake Scenario  
in Beijing Metropolitan Area



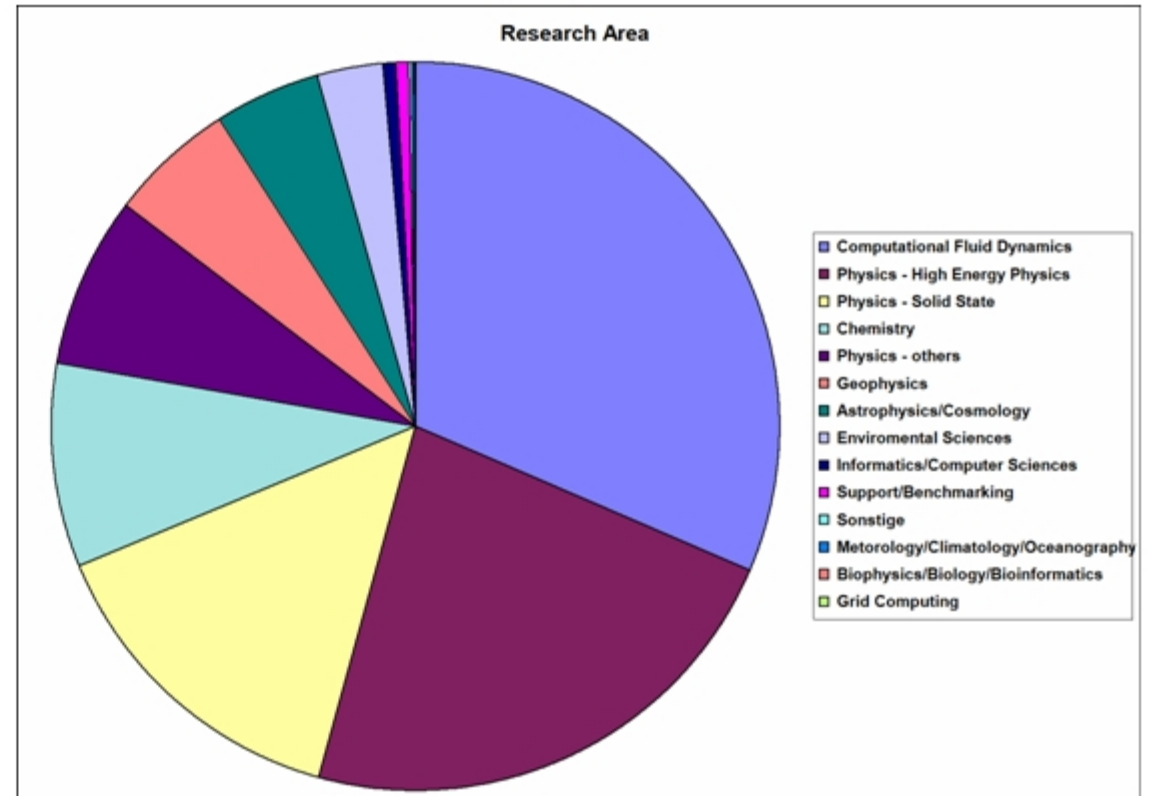
3D wavefield of seismic velocities  
inside Earth



**Department for Geosciences**

University of Munich

G. Brietzke, H. Igel, G. Jahnke et al.

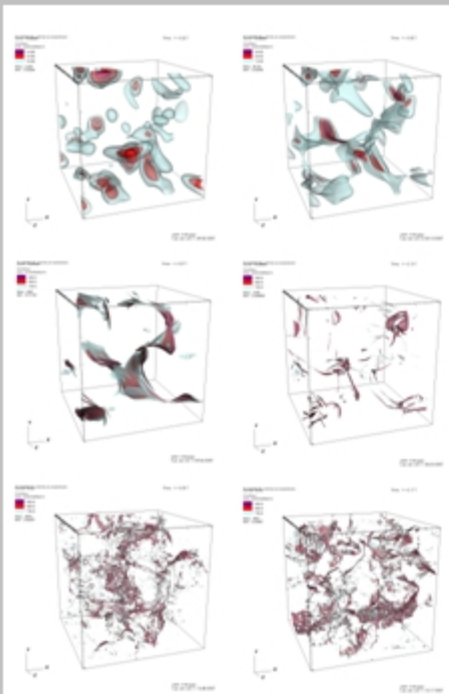




www.lrz.de

## SGI Altix 4700: Sustained System Performance (Phase 1)

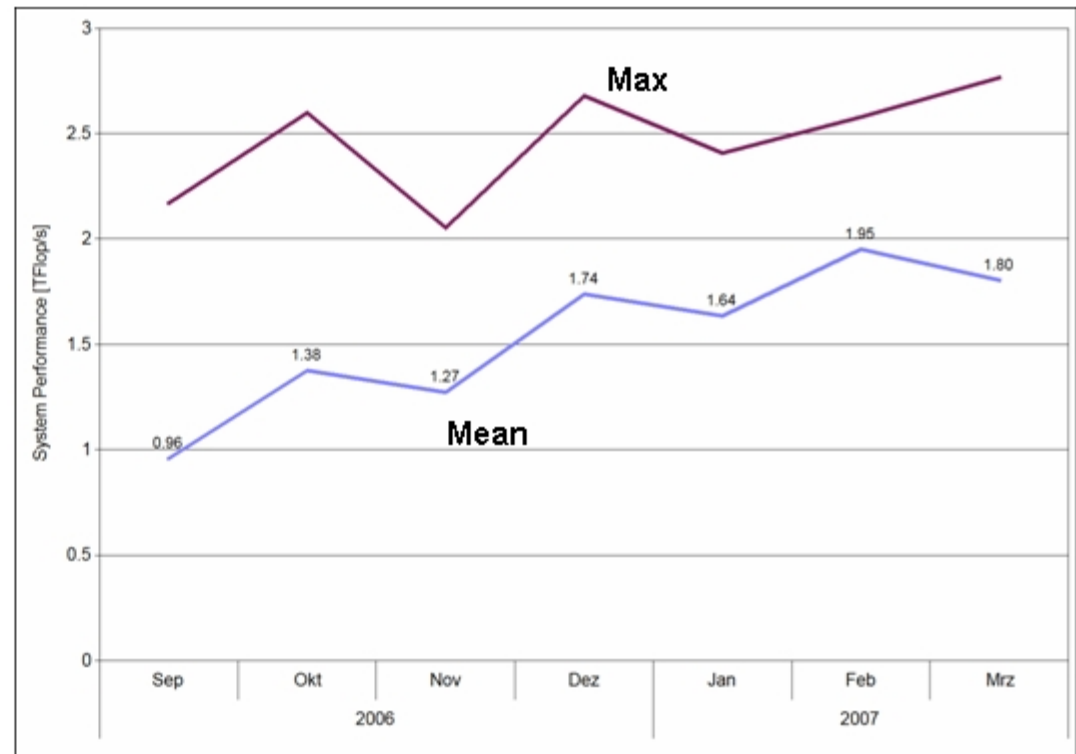
The many different Itanium PMEs gave us the opportunity to closely monitor the system as well as the user code's performance. At regular intervals we sample ~40 performance counter on a system-wide basis. The performance data is stored in a database and allows us to draw conclusions on various performance aspects.



### AMR for Turbulent Astrophysical Flows

Institute of Astronomy  
University of Würzburg  
W. Schmidt, J. Niemeyer et. al.

The image to the right shows the overall system performance for the 4096 Madison processors during the seven months the system was in user operation. In February the system reached a mean performance of 1.95 TFlop/s (7.5% of peak).






www.lrz.de

Chair for Computer Engineering  
Technical University Munich  
P. Wenisch, S. Kühner, E. Rank  
O. Wenisch (LRZ)

Computational Steering  
Interactive Airflow Simulation



© 2007 LRZ

## Software Stack

We offer our customers a wide range of software and tools for developing, debugging and optimizing their codes.

Compilers	Intel Compilers Gnu Compilers
Parallel Environments	SGI MPT shmem API Intel MPI OpenMP
Tools	VTune Intel Tracing Tools Intel Threading Tools DDT Totalview MemoryScape Papi histx gprof
Libraries	MKL FFTW NAG (SMP and MPI parallel) Petc ScaLAPACK
Grid-Middleware	UNICORE Globus



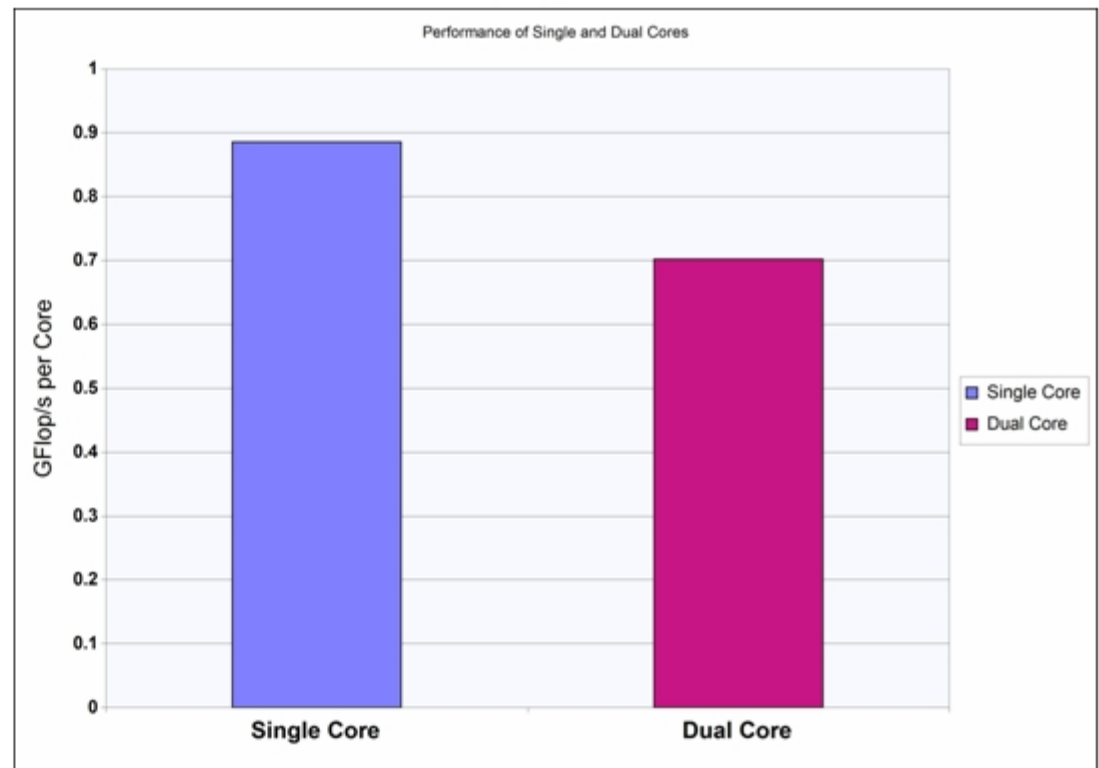
www.lrz.de

## SGI Altix 4700: Switching from Single to Dual Core

In March 2007 all Madison processors were replaced by Montecito Dual Core Sockets, thereby upgrading the system to a peak performance of 62.3 TFlop/s. The size of the L3 cache per core increased from 6 to 9 MByte. This reduces the pressure on the memory channel, which is now shared by two cores. First measurements indicate that overall per-core performance dropped by only ~20% notwithstanding the worse memory balance.



The data for the single core system was gathered in approx. 7 months, the data for the dual cores is based on constant performance measures over 10 days only.







www.lrz.de

## Switching from single to dual-core: Performance Data for different research areas

Most research areas keep up their per core performance. Essentially only the very memory intensive CFD codes suffer from the update to dual core.



© 2007 LRZ





www.lrz.de

## Switching from single to dual-core: Performance resolved by User

Here the first per-user impact of the dual-core update is displayed. It can be seen that highly optimized codes still perform well.



© 2007 LRZ

