


Basics



Theme: Towards Reconfigurable HPC
Lecture 1


Basics

Iris Christadler
Leibniz Supercomputing Centre

Inverted CERN School of Computing, 3-5 March 2008

1
iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics




Introduction


- **Objectives:**
 - Explain why there is a major challenge in HPC
 - Identify its implications for hardware & software
 - Give a short overview of available accelerators
- **Content:**
 - HPC platforms
 - HPC users
 - The Petascale Area

2
iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics




LRZ's SGI Altix 4700 ("HLRB II")



	Phase 1 (until 03/2007)	Phase 2 (since 04/2007)
Overall characteristics:		
Total number of cores	4096	9728
Peak Performance	26.3 TF	62.3 TF
Linpack Performance	24.5 TF	56.5 TF
LRZ-Benchmark Perf.	8.2 TF	16.2 TF
Size of memory	17.5 TB	39 TB
Processor type	Intel Itanium2 Madison 9M	Intel Itanium2 Montecito Dual Core
Clock rate	1.6 GHz	1.6 GHz
Peak Performance	6.4 GF	6.4 GF
L3 Cache (per core)	6 MB	9 MB


3
iCSC2008, Iris Christadler, Leibniz Supercomputing Centre
Image: www.lrz.de

Basics



Juelich's Blue Gene/P (JUGENE)

- currently the fastest Computer in Europe
- ranked #2 in the top500 (11/2007)



Overall characteristics	JUGENE
Total number of cores	65536
Peak Performance	222.8 TF
Linpack Performance	167.3 TF
Size of memory	32 TByte
Processor type	PowerPC 450
Clock rate	850 MHz
Peak Performance	3.4 GF

4
iCSC2008, Iris Christadler, Leibniz Supercomputing Centre
Image: www.fz-juelich.de

Basics

iCSC
CERN
School of Computing

An introduction to... HIGH PERFORMANCE COMPUTING

5

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iCSC
CERN
School of Computing

Definition

- **High Performance Computing (HPC):**
[http://en.wikipedia.org/wiki/High-performance_computing]
 - use of parallel supercomputers or
 - compute clusters
 - usually mass-produced processors
 - linked together in a single system
 - with commercially available interconnects
 - systems above the teraflops-region
 - used for scientific research
- **High Productivity Computing**
[http://en.wikipedia.org/wiki/High-performance_computing]

“The more current and evolving definition of HPC refers to High Productivity Computing, and reflects the purpose and use model of the myriad of existing and evolving architectures, and the supporting ecosystem of software, middleware, storage, networking and tools behind the next generation of applications.”

6

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iCSC
CERN
School of Computing

top500

- consists of the top500 most powerful computing systems
- updated twice a year
 - in June at the European ISC
 - in November at the US SC
- www.top500.org

08/11/2007 <http://www.top500.org>

7

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre Image: www.top500.org

Basics

iCSC
CERN
School of Computing

top500 Statistics

Countries Share Over Time
1993-2007

Vendors Share Over Time
1993-2007

8

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre Images: www.top500.org

The actual top5 ...

No	Site	System	#Proc	Peak Perf	Linpack	Efficiency
1	LLNL, USA	IBM Blue Gene/L	212992	596378	478200	80%
2	FZJ, Germany	IBM Blue Gene/P	65536	222822	167300	75%
3	NMCCAC, USA	SGI Altix ICE 8200	14336	172032	126900	74%
4	India	Cluster Platform, HP	14240	170880	117900	69%
5	Sweden	Cluster Platform, HP	13728	146430	102800	70%
...
15	LRZ, Germany	SGI Altix 4700	9728	62259	56520	91%
16	Japan	Sun + ClearSpeed	11664	102021	56430	55%
17	EPCC, UK	HECToR, Cray XT4	11328	63436	54648	86%

Efficiency= Linpack Perf./Peak Perf.

9

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Having a closer look at the LRZ performance figures

MORE STATISTICS

10

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Who is using supercomputers?

Figures taken from the top500 and LRZ's SGI Altix (9728 cores)

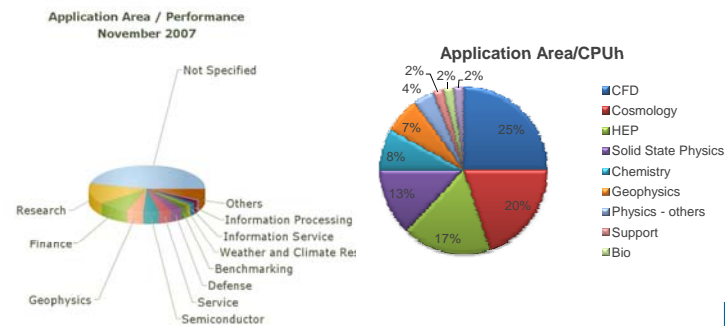


Image: www.top500.org

Image: LRZ

11

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

How are supercomputers used?

Figures taken from LRZ's SGI Altix (9728 cores)



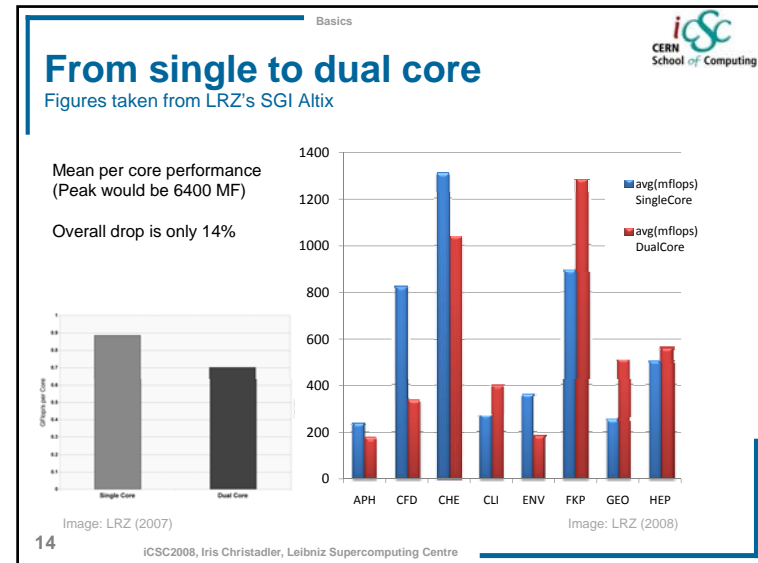



Image: LRZ

12

ICSC2008, Iris Christadler, Leibniz Supercomputing Centre



Basics

iSC
CERN
School of Computing

"The free lunch is over"

- **Frequency scaling is now prevented by physical constraints**
 - Heat (too much of it and too hard to dissipate)
 - Power Consumption (too high)
 - Current leakage problems
- **Future performance gains will come from**
 - Hyperthreading
 - Multicore
 - Cache
- **This requires better software**

Software is the main issue!

"But if you want your application to benefit from the continued exponential throughput advances in new processors, it will need to be a well-written concurrent application. And that's easier said than done, because not all problems are inherently parallelizable and because concurrent programming is hard." [taken from the article "The free lunch is over"]

17 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iSC
CERN
School of Computing

Problems exponentiate in HPC

- **Power Consumption of the system**
- **Power Consumption of the cooling system**
- **Footprint**
- **Mean Time Between Failure (MTBF)**

"I can hear the howls of protest: "Concurrency? That's not news! People are already writing concurrent applications." That's true. Of a small fraction of developers."

[<http://www.getw.ca/publications/concurrency-ddj.htm>]

18 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iSC
CERN
School of Computing

PETASCALING

Sounds interesting? So let's start with...

19 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iSC
CERN
School of Computing

Petascale...

- **is on it's way both in Europe and in the US**
 - First US petascale system estimated for 2008 (called "Roadrunner"):
 - Hybrid Opteron-Cell system
 - use Cell BE as "accelerators"
 - built at Los Alamos National Laboratory
 - A European petascale supercomputer is expected in 2010
- **Are we ready for petascale?**
- **Are our codes ready for petascale?**
- **Strong interest towards "accelerators" in HPC**

20 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

The Alternatives

- **Multicore**
- **GPGPU**
- **Cell**
- **FPGA**
- **ClearSpeed**
- ...

➤ **We will have a closer look on all of them today!**

21 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

Multicore

- **Standard CPUs are now multicore CPUs**
- **available since 2005**
- **de-facto standard**
- **used in many supercomputers**
- **up to 8 cores today**
- **hundreds of cores tomorrow?**

22 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

GPGPUs

- **General-purpose computing on graphics processing units**
- **Using the GPU as “vector CPU” by (ab)using the programmable vertex shaders**
- **Available since 2000**
- **Becoming more and more popular**
- **Use stream processing to exploit the extremely parallelism**
- **CUDA tutorial at the Supercomputing Conference 2007**
- A nice introduction can be found at:
[<http://en.wikipedia.org/wiki/GPGPU>]

23 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

Cells

- **A heterogeneous architecture developed for the PS3**
- **Combines a PowerPC with co-processing elements to accelerate multimedia and vector processing applications**
- **Software controlled memories**
- **Available since 2005**
- **Many research CELL clusters/projects**
- **A hybrid Opteron-Cell cluster will become the first petaflop system**
- A nice overview can be found at:
[[http://en.wikipedia.org/wiki/Cell_\(microprocessor\)](http://en.wikipedia.org/wiki/Cell_(microprocessor))]

24 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

FPGAs

- Field-programmable gate array
- “Adjust the architecture to the needs of your algorithm”
- Invented 1984
- Used heavily in embedded and real-time systems
- Used in supercomputers like Cray XD1, SGI RASC Blades
- Programmability!
- An overview can be found at:
[http://en.wikipedia.org/wiki/Field-programmable_gate_array]

25

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

ClearSpeed Boards

- Accelerator boards specially developed for scientific computing and the needs of the HPC community
- Only accelerated platform in the current top500
- Advertisement claims:
 - “World’s highest performance processor” (80.64 GF per board)
 - “World’s highest performance per watt” (2 GF/Watt)
- Linpack results:

System Specification	Linpack Result
4 nodes (16GB) w/o Advance boards	136.0 GF
4 nodes (16GB) w/ 2 x Advance boards each	364.2 GF
1 node (16GB) w/o Advance boards	34.0 GF
1 node (16GB) w/ 2 x Advance boards	90.1 GF

[<http://www.clearspeed.com/acceleration/performance/benchmarks/>]

26

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Overview of available Accelerators

	GPU	FPGA	ClearSpeed	Cell
Price	\$	\$\$-\$\$\$	\$\$\$	\$\$-\$\$\$
Power	high	low	medium	medium
Good at	graphics, 32bit	integer	64bit	graphics, 32bit
64bit?	no	yes	yes	yes
64bit Perf.	-	low	high	high in 2008
IEEE-754	no	expensive	yes	nearly

Source: White Paper from HP
[<http://www.hp.com/techservers/hpccn/hpccollaboration/ADCatalyst/downloads/accelerators.pdf>]

27

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Our answer to the challenge in HPC?

RECONFIGURABLE COMPUTING

28

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iCSC
CERN
School of Computing

Definition

- **Reconfigurable Computing (RC)**
 - Idea of reconfiguring a computer to your **current** needs
 - Use FPGAs for the reconfiguration
- **Concept exists since 1960s (Paper by Gerald Estrin)**
 “Unfortunately this idea was far ahead of its time in needed electronic technology.”
- **Renaissance in the 80s/90s**
 “The world’s first commercial reconfigurable computer, the Algotronix CHS2X4, was completed in 1991. It was not a commercial success.”
 Quotes are taken from [http://en.wikipedia.org/wiki/Reconfigurable_computing]
- **Reconfigurable HPC (RHPC)**
- **Pros and Cons discussed during the next two days!**

29

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iCSC
CERN
School of Computing

Alternative Platforms need ... ALTERNATIVE PROGRAMMING MODELS

30

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

iCSC
CERN
School of Computing

My favorite “Dongarra Slide”

Real Crisis With HPC Is With The Software

- ♦ **Programming is stuck**
 - Arguably hasn't changed since the 70's
- ♦ **It's time for a change**
 - Complexity is rising dramatically
 - highly parallel and distributed systems
 - From 10 to 100 to 1000 to 10000 to 100000 of processors!
 - multidisciplinary applications
 - ♦ **A supercomputer application and software are usually much more long-lived than a hardware**
 - Hardware life typically five years at most.
 - Fortran and C are the main programming models
 - ♦ **Software is a major cost component of modern technologies.**
 - The tradition in HPC system procurement is to assume that the software is free.

[<http://www.netlib.org/utk/people/JackDongarra/SLIDES/dongarra-isc2004.pdf>]

29

Basics

iCSC
CERN
School of Computing

Parallel Programming Languages

- **Partitioned Global Address Space (PGAS) languages**
 Unified Parallel C (UPC), Co-Array Fortran (CAF), Titanium
- **HPCS languages**
 Fortress, Chapel, X10
- **Data-stream languages**
 Brook, CUDA, RapidMind
- **Others**
 Ct, STM, ...

➢ **All of them will be introduced during our lectures!**

32

iCSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

Operating system Family / Systems
November 2007

Linux
Others
Windows
Unix
Mixed

Last but not least, the (scary) Microsoft ideas:
“HOW TO MAKE MONEY IN HPC”
[\[http://www.eetimes.com/showArticle.jhtml?articleID=201200019\]](http://www.eetimes.com/showArticle.jhtml?articleID=201200019)

33 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre Image: www.top500.org

Basics

Conclusion

- Expected major change in the basic hardware architecture
- Lead to the necessity of new programming models
- Ideas exist but are mostly research projects
- With new programming models reconfigurable computing is becoming more and more interesting
- This workshop will give an introduction of alternatives that are currently investigated
- You, as programmers, have to choose and let the dream of easy-to-use, easy-to-change, errorless code “come true”

34 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre

Basics

References and further reading

- **The Landscape of Parallel Computing Research: A View from Berkeley**
<http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>
- **The Free Lunch Is Over**
<http://www.gotw.ca/publications/concurrency-ddj.htm>
- **Accelerators For HPC, White Paper, HP**
<http://www.hp.com/techservers/hpccn/hpccollaboration/ADCatalyst/downloads/accelerators.pdf>
- **M'soft: Parallel programming model 10 years off**
<http://www.eetimes.com/showArticle.jhtml?articleID=201200019>

35 ICSC2008, Iris Christadler, Leibniz Supercomputing Centre