

16 November 2009

The TOP500[®] Report

SUPERCOMPUTER SITES

80 GFlops + Tax

In its third run to knock “Roadrunner” off the top perch on the TOP500 list

ORNL’s Jaguar Claws its Way to Number One

MANNHEIM, Germany; KNOXVILLE, Tenn.; and BERKELEY, Calif.—In its third run to knock the IBM supercomputer nicknamed “Roadrunner” off the top perch on the TOP500 list of supercomputers, the Cray XT5 supercomputer known as Jaguar finally claimed the top spot on the 34th edition of the closely watched list.

The newest version of the TOP500 list, which is issued twice yearly, will be formally presented on Tuesday, Nov. 17, at the SC09 Conference to be held at the Oregon Convention Center in Portland.

Jaguar, which is located at the Department of Energy’s Oak Ridge Leadership Computing Facility and was upgraded earlier this year, posted a 1.75 petaflop/s performance speed running the Linpack benchmark. Jaguar roared ahead with new processors bringing the theoretical peak capability to 2.3 petaflop/s and nearly a quarter of a million cores. One petaflop/s



refers to one quadrillion calculations per second.

When the Roadrunner system at Los Alamos first appeared at the top of the June

2008 TOP500 list, it was the world’s first petaflop/s supercomputer. This time around, Roadrunner recorded a performance of 1.04 petaflop/s, dropping from 1.105 petaflop/s in

	MANUFACTURER/COMPUTER	LOCATION	COUNTRY	CORES	R _{max}
1	Jaguar, Cray XT5 6-core 2.6 GHz	DOE / OS / ORNL	USA	224162	1.76
2	Roadrunner, IBM BladeCenter QS22/LS21 Cluster, PowerXCell 3.2 Ghz / Opteron 1.8 GHz, Voltaire Iband	DOE / NNSA / LANL	USA	122400	1.04
3	Kraken, Cray XT5 6-core 2.6 GHz	NSF / U of Tennessee	USA	98928	.832
4	Jugene, IBM Blue Gene/P Solution	Forschungszentrum Juelich	Germany	294912	.826
5	Tianhe-1 Cluster, Xeon E5540/E5450 + AMD ATI 4870	National SuperComputer Center in Tianjin/NUDT	China	71680	.563

The Great Leap of the Jaguar

Hans W. Meuer, Prometheus & Universität Mannheim

At the third attempt, it finally made it. The supercomputer at the National Center for Computational Science in Oak Ridge, named Jaguar, now tops the 34th TOP500 list of supercomputers with 1.759 petaflops. In the previous two lists, the Jaguar had already gotten quite close to list leader Roadrunner from IBM (the first system worldwide to break the petaflops barrier in summer 2008) – the difference was less than

.. this is the very first time – after 34 TOP500 lists – that the Cray, well known for being specialized in supercomputers only, has achieved first place.

4 percent Linpack performance.

How did the Jaguar manage to leap to first place with 69 percent more Linpack performance than the Roadrunner? Since July, the Jaguar’s number of compute cores has been increased from 129600 to 224162 and – what’s substantial – it has been equipped with AMD’s new six-core processors ‘Istanbul’, which have only been available since August, and 2GB of memory

Tianhe-1 China's first Petaflop/s scale supercomputer

The Chinese sNatsional Univesrsity of the middle of 2009, and will be moved to Defense Technology (NUDT) recently unveiled China’s fastest supercomputer, also NSCC-TJ soon. The TH-1 system is a hybrid design with Intel Xeon processors and AMD GPUs. The TH-1 uses AMD GPUs as accelerators. Each

node consists of two AMD GPUs attached to two Intel Xeon processors. The TH-1 is made up of 80 compute cabinets including 2560 compute nodes and 512 operation nodes. There are two kinds of Intel processors used in the system, including 4096 Intel Xeon E5540 processors with a



node consists of two AMD GPUs attached to two Intel Xeon processors.

The TH-1 is made up of 80 compute cabinets including 2560 compute nodes and 512 operation nodes. There are two kinds of Intel processors used in the system, including 4096 Intel Xeon E5540 processors with a

TOP500 HIGHLIGHTS

- ❖ The upgraded Jaguar system at Oak Ridge National Laboratory took the No. 1 spot from Roadrunner with 1.75 PF linpack performance.
- ❖ Roadrunner is now No. 2 with reduced linpack performance of 1.04 PF (1.105 PF in June) after repartitioning of the system.
- ❖ The upgraded Kraken system at the National Institute for Computational Sciences/University of Tennessee is then second Cray XT5 in the TOP3 with 832 TFlop/s.
- ❖ The most powerful system outside the U.S. is an IBM BlueGene/P system at the German Forschungszentrum Juelich (FZJ) at No. 4
- ❖ The new Tianhe-1 system installed at the National Super Computer Center in Tianjin, China is #5 and the highest ranked Chinese system ever.
- ❖ Intel dominates the high-end processor market with 80.4 percent of all systems and 76 percent of quad-core based systems.
- ❖ Intel Core i7 (Nehalem-EP) increases its presence in the list with 95 systems compared with 33 in the last list.
- ❖ Quad-core processors are used in 85.4 percent of the systems. Their use accelerates performance growth at all levels.
- ❖ The Tianhe-1 system at No. 5, which is a hybrid design with Intel Xeon processors and AMD GPUs. The TH-1 uses AMD GPUs as accelerators. Each node consists of two AMD GPUs attached to two Intel Xeon processors.
- ❖ A new Sun Blade system at No. 10 installed at Sandia National Lab, following another Sun Blade system installed at TACC.
- ❖ Hewlett-Packard kept a narrow lead in market share by total systems from IBM, but IBM still stays ahead by overall installed performance.

Highlights from the Top 10:

- ❖ After a close race last November when Roadrunner was almost surpassed by the second petaflop/s system ever - the Jaguar system installed at the DOE’s Oak Ridge National Laboratory, Jaguar now comes back with new processors bringing the total peak capability of that machine to 2.3 petaflop/s and nearly a quarter of a million cores. Jaguar is now at No. 1 with 1.75 petaflop/s.
- ❖ The Roadrunner system at DOE’s Los Alamos National Laboratory (LANL) was built by IBM and in June 2008 was the first system ever to break the petaflop/s Linpack barrier. Roadrunner

More highlights and trends on page 2

TOP 25 Supercomputing Sites

	Site	Computer	Cores	Rmax (TFs)	Rpeak (TFs)	Power (kW)
1	Oak Ridge National Laboratory	Cray XT5-HE Opteron Six Core 2.6 GHz	224162	1759.00	2331.00	6950.60
2	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband	122400	1042.00	1375.78	2345.50
3	National Institute for Computational Sciences/University of Tennessee	Cray XT5-HE Opteron Six Core 2.6 GHz	98928	831.70	1028.85	
4	Forschungszentrum Juelich (FZJ)	Blue Gene/P Solution	294912	825.50	1002.70	2268.00
5	National SuperComputer Center in Tianjin/NUDT	NUDT TH-1 Cluster, Xeon E5540/E5450, ATI Radeon HD 4870 2, Infiniband	71680	563.10	1206.19	0.00
6	NASA/Ames Research Center/NAS	SGI Altix ICE 8200EX, Xeon QC 3.0 GHz/Nehalem EP 2.93 Ghz	56320	544.30	673.26	2348.00
7	DOE/NNSA/LLNL	eServer Blue Gene Solution	212992	478.20	596.38	2329.60
8	Argonne National Laboratory	Blue Gene/P Solution	163840	458.61	557.06	1260.00
9	Texas Advanced Computing Center/Univ. of Texas	SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband	62976	433.20	579.38	2000.00
10	Sandia National Laboratories / National Renewable Energy Laboratory	Sun Blade x6275, Xeon X55xx 2.93 Ghz, Infiniband	41616	423.90	487.74	0.00
11	DOE/NNSA/LLNL	Blue Gene/P Solution	147456	415.70	501.35	1134.00
12	Moscow State University - Research Computing Center	T-Platforms T-Blade2, Xeon 5570 2.93 GHz, Infiniband QDR	35360	350.10	414.42	
13	Forschungszentrum Juelich (FZJ)	Sun Constellation, NovaScale R422-E2, Intel Xeon X5570, 2.93 GHz, Sun M9/ Mellanox QDR Infiniband/Partec Parastation	26304	274.80	308.28	1549.00
14	KISTI Supercomputing Center	Sun Constellation, NovaScale R422-E2, Intel Xeon X5570, 2.93 GHz, Sun M9/ Mellanox QDR Infiniband/Partec Parastation	26232	274.80	307.44	1275.96
15	NERSC/LBNL	Cray XT4 QuadCore 2.3 GHz	38642	266.30	355.51	1150.00
16	Oak Ridge National Laboratory	Cray XT4 QuadCore 2.1 GHz	30976	205.00	260.20	1580.71
17	Sandia National Laboratories	Cray XT3/XT4	38208	204.20	284.00	2506.00
18	King Abdullah University of Science and Technology	Blue Gene/P Solution	65536	190.90	222.82	504.00
19	Shanghai Supercomputer Center	Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008	30720	180.60	233.47	
20	University of Edinburgh	Cray XT4, 2.3 GHz	22656	174.08	208.44	
21	Swiss Scientific Computing Center (CSCS)	Cray XT5 SixCore 2.4 GHz	22032	168.70	211.51	713.00
22	SciNet/University of Toronto	iDataPlex, Xeon E55xx QC 2.53 GHz, GigE	30240	168.60	306.03	869.40
23	Government	Cray XT5 QC 2.4 GHz	20960	165.60	201.22	
24	ERDC DSRC	SGI Altix ICR 8200 Enh. LX, Xeon X5560 2.8Ghz	15360	160.20	172.03	774.50
25	New Mexico Computing Applications Center (NMCAC)	SGI Altix ICE 8200, Xeon quad core 3.0 GHz	14336	133.20	172.03	861.63

Highlights...

- is now at No. 2 with 1.04 petaflop/s, downgraded from 1.105 petaflop/s after it was re-partitioned. Roadrunner is based on the IBM QS22 blades which are built with advanced versions of the processor in the Sony PlayStation 3. These nodes are connected with a commodity InfiniBand network.
- ❖ The TOP10 features 2 new systems and some upgrades (Jaguar) and downgrades (Roadrunner)
 - ❖ In the TOP10 only the No. 4 and 5 systems are installed outside the U.S. – in this case in Germany and China.
 - ❖ The No. 10 system called "Red Sky" is a new Sun Blade system, installed at the Sandia National Lab with a Linpack performance of 423 Tflop/s.
- General highlights since the last edition:**
- ❖ The last system on the newest list would have been listed at position 336 in the previous TOP500 just six months ago. This turnover rate is just below average after the TOP500 recorded the highest turnover in its history one year ago.
 - ❖ Total combined performance of all 500 systems has grown to 27.6 Pflop/s, compared to 22.6 Pflop/s six months ago and 16.9 Pflop/s one year ago.
 - ❖ The entry point for the top 100 increased in six months from 39.58 Tflop/s to 47.72 Tflop/s.
 - ❖ The average concurrency level in the TOP500 is 9,174 cores per system up from 8,210 six month ago and 6,240 one year ago.
 - ❖ The IBM Power processors are the second most common used processor family with 52 systems (10.4 percent), down from 55.
 - ❖ They are followed by the AMD Opteron family with 42 systems (8.4 percent), down from 43.
 - ❖ Multi-core processors are the dominant chip architecture. The most impressive growth showed the

- number of systems using the Intel Harpertown, Clovertown, and Nehalem-EP quad core chips, which grew from 336 last November to 336 systems in June and now 380 systems with the addition of the Nehalem-EP processor to Intel’s quad core lineup.
- ❖ 417 systems are labeled as clusters, making this the most common architecture in the TOP500 with a stable share of 83.4 percent.
 - ❖ Gigabit Ethernet is still the most-used internal system interconnect technology (259 systems), due to its widespread use at industrial customers, followed by InfiniBand technology with 181 systems.
 - ❖ IBM and Hewlett-Packard continue to sell the bulk of systems at all performance levels of the TOP500.
 - ❖ HP kept a narrow lead in systems with 210 systems (42 percent) over IBM with 185 systems (37 percent). HP had 212 systems (42.4 percent) six months ago, compared to IBM with 188 systems (37.6 percent).
 - ❖ In the system category, Cray, SGI, and Dell follow with 3.8 percent, 3.8 percent and 3.2 percent respectively.
 - ❖ In the performance category, the manufacturers with more than 5 percent are: Cray (15.9 percent of performance) and SGI (6.6 percent), each of which benefits from large systems in the TOP10.
 - ❖ HP (192) and IBM (114) sold together 306 out of 313 systems at commercial and industrial customers and have had this important market segment clearly cornered for some time now.
 - ❖ Dominant countries in Asia are China with 21 systems (unchanged), Japan with 16 systems (up from 15), and India with 3 systems (down from 6).
 - ❖ In Europe, UK remains the No. 1 with 45 systems (44 six months ago). Germany and France share the No. 2 spot with 2 systems 27 systems each.
 - ❖ Highlights from the Top 50:

- ❖ The entry level into the TOP50 is at 94.7 Tflop/s
- ❖ The U.S. has a lower percentage of systems (44 percent) in the TOP50 than in the TOP500 (55.4 percent).
- ❖ The dominant architectures are custom-built massively parallel systems MPPs with 60 percent ahead of commodity clusters with 40 percent.
- ❖ IBM leads the TOP50 with 30 percent of systems and 33.6 percent of performance.
- ❖ No 2 is Cray with a share of 18 percent of systems and 29.4 percent of performance.
- ❖ SGI is third with 16 percent of systems and 10.5 percent of performance.
- ❖ HP has 10 percent of systems and 4.0 percent of performance..
- ❖ There is only a single system using Gigabit Ethernet in the TOP50.
- ❖ Cray’s XT is the most-used system family with 9 systems (18 percent), followed by IBM’s BlueGene with 8 systems (16 percent).
- ❖ Intel processors are used in 38 percent of systems, AMD processors in 30 percent and IBM’s Power processors in 28 percent.
- ❖ The average concurrency level is 44,338 cores per system – up from 40,871 cores per system six month ago and 30,490 one year ago.

Jaguar...

per core. Each compute node features two Opterons with 12 cores and 16GB of shared memory. The whole system has 300TB of memory and 10PB of hard disc space.

Another factor was that two of the Roadrunner's CUs (connecting units) have been dedicated to other tasks reducing the system's Linpack performance from 1105 teraflops in June to 1042 teraflops now.

In June, Oak Ridge and Cray had hinted that they hadn't wasted valuable computer resources on the Linpack, but had instead used the Jaguar for work on substantially scientific problems – in retrospect, a clever (bluff) move. They passed on a close first place in June in order to surprise everyone with a truly conclusive

when the first TOP500 list was released in June 1993, 41 percent of the listed systems were from Cray.

advantage over the Roadrunner now in November – nobody doubts that the Jaguar could have topped the list in June. Accordingly, Adolff Hoisie of Los Alamos described the circumstance of the Roadrunner placing first instead of the Jaguar as the “miracle of Hamburg” at the ISC'09 Conference. The Jaguar now achieved 1.759 petaflops of computing power solving a linear system of equations with $n = 5474272$. It finished calculating after 17 hours and 17 minutes.

It's worth mentioning that this is the very first time – after 34 TOP500 lists – that the company Cray (“The Supercomputer Company”), well known for being specialized in supercomputers only, has achieved first place. During the early stages of TOP500 and even before that, since 1976, Cray had been the market leader for supercomputers. Consequently, when the first TOP500 list was released in June 1993, 41 percent of the listed systems were from Cray. That's an unbelievably high percentage because the list's second, Fujitsu, only had a market share of 13.8 percent back then.

But now, 16 years after the start of TOP500, the dynamic HPC market has changed completely: in the course of the years, Cray has lost its position as market leader and has become a niche computer company serving the needs of high-end government laboratories and academic users. Other companies, first and foremost IBM and HP, have clearly taken the lead in the list. Cray had undergone a severe crisis because of the total collapse of the vector computer market, a crisis it still hasn't recovered from. In the current list – Nov 2009 – Cray only has 19 systems, which equals a share of 3.8 percent. However, 14 of these systems are in the TOP100 and two of them even hold excellent positions in the TOP3. So, it seems possible the former leader Cray is on the rise again, especially now, motivated by the success of topping the TOP500 list. In any case, Cray will keep its first place until the next TOP500 list is released at the ISC'10 (www.isc10.org) in Hamburg next year, between May 31 and June 3.

Tianhe-1...

frequency of 2530MHz and 1024 Intel Xeon E5450 processors with a frequency of 3000MHz. The L1 cache size of the E5540 and E5450 processors is 128KB, the L2 cache sizes are 1MB and 12MB respectively. The E5540 processor has a L3 cache with the size of 8MB. Each compute node has two Intel Xeon processors with 32GB of memory. ATI Radeon HD 4870x2 GPUs are connected via PCI-E connections on each compute node. The maximum power recorded during the execution of Linpack was 0.58KW per node. The Tianhe system used 20480 CPU cores (2 cpu*4 cores*2560 nodes) and 4096000 SPU (1600 Stream Processing Units* 2560 nodes) while executing the Linpack benchmark. The cores of nodes not in operation whilst running the Linpack benchmark were not included in the measurements. Each operation node also has two Intel Xeon E5450 processors and 32GB memory. The theoretical peak performance and the total memory of the whole system, which includes compute nodes and operation nodes, are 1.206PFlops and 98,304GB.

The compute nodes with E5540 processors are connected to 9 first-stage Infiniband switches. Each first-stage switch is connected to each second-stage switch through 18 uplinks, which makes a total of 72 uplink connections for 4 second-stage switches. The compute nodes with E5450 processors are connected by 64 Infiniband switch modules in the cabinets. Each switch module is connected to the second-stage switches through 8 uplinks.

The DGEMM and DTRSM, which are the core procedures of Linpack benchmark, are accelerated by cooperations of CPU and GPU on TH-1 system.

There are two processes located on each node. Each process is executed by a CPU and GPU pair. The three cores of CPU are used to execute part of the computing tasks, and the left core is used to control the GPU to participate the execution of computing tasks. The part of program executing on CPU uses the Intel MKL-10.2.1.017 library and the part of GPU uses the AMD ACMLG-1.0 library especially optimized by NUDT. The optimization algorithms used to achieve a better result of Linpack benchmark are as follows.

First, dynamic load balance technique is used when allocating the tasks between CPU and GPU.

Second, the instruction of streaming load/store is adopted to reduce the conflict between CPU and GPU's data accessing. The third, software-pipelining technique is used to overlap the execution of GPU and the transmission of data between GPU and CPU. The forth, affinity-scheduling technique is used to reduce the performance fluctuation by utilizing processor cores' computing and controlling ability. The fifth, we optimize the functions of DGEMM in AMD ACMLG library and DTRTRI and DTRMM in Intel MKL library to speed up the performance of DTRSM.


For the sake of the stabilization, the frequency of GPU core is decreased from 750MHz to 575MHz. Besides this, the frequency of GPU's memory is also decreased from 900MHz to 650MHz.

Number One...

June 2009 due to a repartitioning of the system. In both November 2008 and June 2009, Jaguar came close but couldn't dislodge Roadrunner from the top slot.

Kraken, another upgraded Cray XT5 system at the National Institute for Computational Sciences/University of Tennessee, claimed the No. 3 position with a performance of 832 teraflop/s (trillions of calculations per second).

At No. 4 is the most powerful system outside the U.S. -- an IBM BlueGene/P supercomputer located at the



Forschungszentrum Juelich (FZJ) in Germany. It achieved 825.5 teraflop/s on the Linpack benchmark and was No. 3 in June 2009.

Rounding out the top 5 positions is the new Tianhe-1 (meaning River in Sky) system installed at the National Super Computer Center in Tianjin, China and to be used to address research problems in petroleum exploration and the simulation of large aircraft designs. The highest ranked Chinese system ever, Tianhe-1 is a hybrid design with Intel Xeon processors and AMD GPUs used as accelerators. Each node consists of two AMD GPUs attached to two Intel Xeon processors.

The only other new system in the top 10 is a Sun Microsystems Blade system at No. 10 and installed at Sandia National Laboratories. Called “Red Sky,” the new Sun Blade system posted a Linpack performance of 423 teraflops.

Here are some other highlights from the latest list showing changes from the June 2009 edition:

- ▶ The entry level to the list moved up to the 20 teraflop/s mark on the Linpack benchmark from 17.1 teraflop/s six months ago. The last system on the newest list would have been listed at position 336 in the previous TOP500 just six months ago.
- ▶ Quad-core processor based systems have taken over the TOP500 quite rapidly. Already 427 systems are using them, with 59 systems using dual-core processors and only four systems still use single core processors. Six systems use IBM's advanced Sony PlayStation 3 processor with 9 cores and two systems at Cray are using the new six-core Shanghai AMD Opteron processors.
- ▶ A total of 402 systems (80.4 percent) are now using Intel processors. This is slightly up from six months ago (399 systems, 79.8 percent). Intel continues to provide the processors for the largest share of TOP500 systems. The IBM Power processors are the second most common used processor family with 52 systems (10.4 percent), down from 55. They are followed by the AMD Opteron family with 42 systems (8.4 percent), down from 43.
- ▶ IBM and Hewlett-Packard continue to sell the bulk of systems at all performance levels of the TOP500. HP kept a narrow lead in

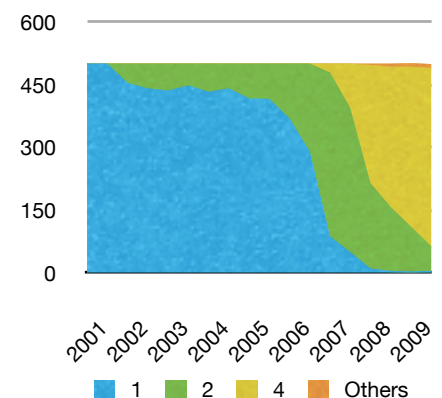
systems with 210 systems (42 percent) over IBM with 186 systems (37.2 percent). HP had 212 systems (42.4 percent) six months ago, compared to IBM with 188 systems (37.6 percent). In the system category, Cray, SGI, and Sun follow with 3.8 percent, 3.8 percent and 2 percent respectively.

- ▶ IBM remains the clear leader in the TOP500 list in performance with 35.1 percent of installed total performance (down from 39.4 percent), compared to HP with 23.0 percent (down from 25.1 percent). In the performance category, the manufacturers with more than 5 percent are: Cray (16.1 percent of performance) and SGI (6.6 percent), each of which

benefits from large systems in the TOP10.

- ▶ The U.S. is clearly the leading consumer of HPC systems with 277 of the 500 systems (down from 291). The European share (153 systems – up from 145) is still substantially larger than the Asian share (50 systems – up from 49). In Europe, UK remains the No. 1 with 45 systems (44 six months ago). Germany and France share the No. 2 spot

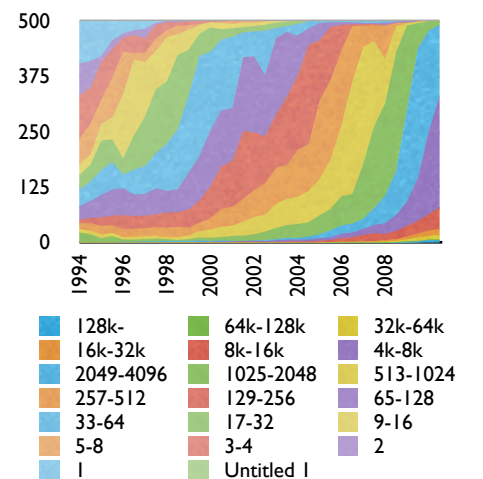
Multi Core / Systems



with 2 systems 27 systems each. Dominant countries in Asia are China with 21 systems (unchanged), Japan with 16 systems (up from 15), and India with 3 systems (down from 6).

The TOP500 list is compiled by Hans Meuer of the University of Mannheim,

Processor Count



Germany; Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory; and Jack Dongarra of the University of Tennessee, Knoxville. For more information, visit www.TOP500.org.

The International Supercomputing Conference (ISC) turns 25 in 2010

SC09, Portland, Oregon - November 16, 2009 – Prometheus GmbH, the organizer of the International Supercomputing Conference (ISC), Europe's premier and world's oldest HPC trade show, proudly announces that it will be celebrating its 25th anniversary at ISC'10 from May 31 – June 3, 2010 in Hamburg, Germany.

"In the spirit of 25-years of steady growth and renowned tradition, ISC'10 will be marking its success with a power-packed conference.

Over 220 invited speakers from research and industry will be sharing their expert knowledge with about 2000 participants," remarked Martin Meuer, ISC's Executive Director.

For the first time in ISC's history, four Keynote talks will be delivered by vertical industry experts and prominent researchers, exploring the latest trends affecting the global supercomputing industry as well as the IT market. Distinguished speakers include Kirk B. Skaugen, Vice President, Intel Architecture Group & General Manager, Data Center Group, Intel USA; Prof. Dr. Helmut Merkel, Manager & Shareholder, EurAsia Global



Concept, Peking, China; Prof. Dr. Thomas Sterling, Arnaud & Edwards Professor of Computer Science, Louisiana State University, USA and Dr.-Ing. habil. Horst Zuse, TU Berlin/FH Lausitz, Germany. Incidentally, Zuse, the son of Konrad Zuse – the creator of world's first fully automatic, program controlled and freely programmable working computer, Z3 – will talk on the "TOP500 List of the early Computers".

The four day conference sessions will also be coupled with a three day comprehensive exhibition participated by 140 businesses and research institutions, showcasing all leading technologies and applications in the supercomputing industry today. ISC'10 is most popular among global IT decision-makers, researchers and HPC enthusiasts from around the world.

Adding more spirit to the celebration, ISC'10 will be offering a 20 percent early bird registration discount for the entire duration of SC09. Registration forms are obtainable at Booth #454.

Conference Highlights of the 25th ISC

In 2010, ISC will not only celebrate its 25th anniversary but also the centennial of Konrad Zuse's birth, the creator of the first fully automatic, program controlled and freely programmable working computer, Z3.


Conference Topics

The four-day conference will highlight many cutting-edge presentations divided into various sessions focusing on innovations and applications of supercomputing.


The areas of emphasis amongst others will be:

- ▶ High Performance Computational Life Sciences – The Challenge for HPC Systems
- ▶ Evolving New HPC Markets: China, Middle East and Russia
- ▶ Fault Tolerance for Manycore Systems
- ▶ HPC: Future Technology Building Blocks
- ▶ High Performance Computing & Networking
- ▶ Storage/Flash Technology in HPC
- ▶ Focusing Linpack: The TOP500 Yardstick
- ▶ Innovative Applications (Cryptography, Gaming, Aerospace, Finance, Defense)
- ▶ Presentation of the 35th TOP500 List of the World's Fastest Supercomputers
- ▶ Panel on Exascale Computing
- ▶ Scientific, BoF, Hot Seat and Poster Sessions, Award Ceremonies, Exhibitor Forums and numerous networking opportunities and social events will complete the 2010 conference program.


ARCHITECTURES



INSTALLATION TYPE



CHIP TECHNOLOGY



Daily Raffle: Win a Kindle

Where? ISC & TOP500 Booth #454

**When? 2:45 p.m.
on Tue, Wed and Thu**



HPLINPACK

A Portable Implementation of the High Performance Linpack Benchmark for Distributed Memory Computers

Algorithm: recursive panel factorizations, multiple lookahead depths, bandwidth reducing swapping

Easy to install, only needs MPI + BLAS or VSIBL

Highly scalable and efficient from the smallest cluster to the largest supercomputers in the world

FIND OUT MORE AT <http://icl.eecs.utk.edu/hpl/>