

PAWEŁ GEPNER*, DAVID L. FRASER**, MICHAŁ F. KOWALIK***,
KAZIMIERZ WAĆKOWSKI****

EVALUATING NEW ARCHITECTURAL FEATURES OF THE INTEL(R) XEON(R) 7500 PROCESSOR FOR HPC WORKLOADS

In this paper we take a look at what the Intel Xeon Processor 7500 family, code named Nehalem-EX, brings to high performance computing. We compare two families of Intel Xeon based systems (Intel Xeon 7500 and Intel Xeon 5600) and present a performance evolution of 16 node clusters based on these CPUs. We compare CPU generations utilizing dual socket platforms and a cluster across a number of HPC benchmarks and focused on different performance field and aspect. We will evaluate also technologies and features like Intels Hyper Threading Technology (HT) and Intel Turbo Boost Technology (Turbo Mode) and the performance implication of these technologies for HPC.

Keywords: HPC, benchmarking, cluster, performance evaluation

EWALUACJA NOWEJ ARCHITEKTURY PROCESORÓW INTEL XENON W OBLICZENIACH WYSOKIEJ WYDAJNOŚCI

W artykule przedstawiamy możliwości procesorów z rodziny Intel Xeon 7500 w obliczeniach wysokiej wydajności. Porównaniu poddano dwa 16-węzłowe klastry oparte na rodzinach procesorów Intel Xeon (7500 i 5600). Eksperyment przeprowadzono na klastrach zbudowanych w oparciu o platformę sprzętową wyposażoną w dwa gniazda procesorowe, wykorzystując popularne benchmarki z dziedziny HPC, koncentrując się na różnych aspektach wydajności. Przedstawiono również wpływ technologii Intel Hyper Threading oraz Intel Turbo Boost Technology na wydajność obliczeń.

Słowa kluczowe: obliczenia wysokiej wydajności, klastry, testy wydajnościowe

* EMEA Platform Architecture Specialist, Intel Corporation, pawel. gepner@intel.com

** EMEA Regional Applications Manager, Intel Corporation, david.l. fraser@intel.com

*** Market Analyst Intel Corporation, michal.f. kowalik@intel.com

**** University Professor Warsaw University of Technology, k.wackowski@chello.pl

1. Introduction

Based on the 37th edition of TOP500 processor family sub-list study we can say that Intel Xeon processors are the clear architecture of choice for the current supercomputers. Exactly 76% of all systems on the list are based on Intel Xeon processors but the most popular is Intel Xeon 5600 family with 169 systems on the list. The Eight-Core Intel Xeon 7500 is not so widely used but is the heart of the biggest supercomputer in Europe [1].

The first generation of Eight-Core Intel Xeon 7500 family processor with an integrated memory controller (IMC) not only represents the first member of the Intel Xeon Processor family for more than dual socket configurations but also brings some new mechanisms which improve overall performance and performance per watt characteristics. Intel Xeon 7500 family is based on the same micro-architectural principles as 5500 family Quad-Core Intel Xeon processors (Nehalem-EP) and has two independent memory controllers (SMI) and utilizes four Intel Quick Path Interconnect (QPI) point-to-point interconnect between processors and to the input/output (I/O) subsystem. Intel Xeon 7500 is dedicated to 2, 4, 8 and even up to 256 sockets platform and it has many Reliability, Accessibility and Serviceability (RAS) features not implemented in Intel Xeon 5500 or Intel Xeon 5600 family. Nehalem-EX delivers many of the attributes which position this product ideally for Mission Critical applications. A bigger cache and 8 cores, and larger memory support also makes this processor ideally positioned for many HPC applications. The new Intel Xeon 7500 family processors have been designed with all the advantages of 45 nm Hi-k metal gate silicon technology. This process technology uses a combination of Hi-k gate dielectrics, conductive materials, and a 45 nm lithography process to improve transistor size and properties such as reduce electrical leakage, chip size, power consumption, and manufacturing costs [2].

This study compares two generations of the Intel Xeon families based platform: Clusters build on the Intel Xeon 5600 with Intel Xeon 7500. We run typical HPC benchmarks and evaluate the performance of the new Intel Xeon 7500 family. We will also discuss some architecture aspects like a bigger cache; more memory channels as well QPI ports. Such technology features as Turbo Mode, HT will also be looked at with their performance implications for HPC. Intels HT technology allows two threads to execute on single core simultaneously utilizing each others unoccupied stages in the execution pipelines. Intel Turbo Mode technology allows processor cores to operate on faster frequency then the specified operating one if the CPU is running below signed specification limits of power, current and temperature.

To answer these research questions, we have organized the paper in the following way. Section 2 presents the platform architecture of the single node system and compare two generation of the platforms Intel Xeon X5670 based system with platform utilizing Intel Xeon X7560. We also outline the single node performance using LINPACK and STREAM benchmark. In section 3 we evaluate the two HPC clusters based on exactly the same configurations as those used for the platform evaluation

by running HPC Challenge benchmark. Section 4 we draw conclusions and outline future work.

2. Platform Architecture

The Intel Xeon 7500 family is the first generation of eight core Intel CPUs dedicated to more than dual socket servers, it is also the first eight core processor with integrated memory controller and QPI interface. Intel Xeon X7560 is a 45 nm eight core monolithic die with 24 MB of L3 cache, 4 channel integrated memory controller and integrated 4 ports of Quick Path Interconnect interface build out of 2.3 billion transistors. The Intel Xeon X7560 cache hierarchy has three levels, L1 and L2 staying fairly small and private to each core, while the L3 cache is much larger. Each core in Nehalem-EX has a private 32 KB first level L1 Instruction and 32 kB Data Cache. In addition the unified 256 KB L2 cache is 8 way associative and provides extremely fast access to data and instructions; the latency is typically smaller than 11 clock cycles. The Intel Xeon 7500 family has 24 MB L3 cache, which gives 3 MB per core. This ratio is better than Intel Xeon 5600 family offers. Figure 1 shows Intel Xeon 7500 family block diagram.




Fig. 1. Intel Xeon 7500 family block diagram

Intel Xeon X7560 has many RAS features and technology dedicated for Mission Critical computing. Some of these features might also be useful in HPC scenarios – especially to enhance the reliability of the system and implement more sophisticated checkpoint mechanisms. All these enhancements may have different performance im-

plications depending on the fields and areas of application however this is an area for further investigation and we did not evaluate this mechanism in this study.

The main focus of this section is to present a comparison of two generations of Intel Xeon processor based platforms and evaluate how a bigger L3 cache, more cores, Intels Hyper Threading Technology (HT) and Intel Turbo Boost Technology (Turbo Mode) can improve performance.

In this study for evaluation of CPU performance we will use LINPACK which is well-suited for parallel workload benchmarks. LINPACK is a floating-point benchmark that solves a dense system of linear equations in parallel. The metric produced is Giga-FLOPS (GFLOPS) or billions of floating point operations per second. LINPACK performs operations called LU Factorization. These are highly parallel and store most of their working data set on the processors cache. It makes relatively few references to memory for the amount of computation it performs. The processor operations are predominantly 64-bit floating-point vector operations and these use SSE instructions. This benchmark is used to determine the performance of world's fastest computers published at the website <http://www.top500.org/> [3].

In both processors the core has 3 functional units that are capable of generating 128 bit (2×64 bit) results per clock. In this case we may assume that a single processor core does two 64 bit floating-point ADD instructions and two 64 bit floating-point MUL instructions per clock. The third 128 bit instruction (or two 64 bit) could be e.g Shuffle but Shuffle is not taken in to account as the theoretical performance, calculated is the product of MUL and ADD executed in each clock, multiplied by frequency. This formula gives the following results [4]. For Intel Xeon X5670 theoretical performance we have = $2.93 \text{ GHz} \times 4 \text{ operations per clock} \times 6 \text{ cores} = 70.32 \text{ GFLOPS}$. For Intel Xeon X7560 theoretical performance we have = $2.266 \text{ GHz} \times 4 \text{ operations per clock} \times 8 \text{ cores} = 72.51 \text{ GFLOPS}$. This is theoretical performance only, and does not fully reflect the real life scenario. The Intel Xeon processor X5670 benefits from a higher CPU clock but Intel Xeon X7560 has more cores and a larger L3 cache per core. Processors characteristic is summarized in Table 1.

In this study we also evaluated what HT and Turbo Mode can bring to real performance and we verified the use of this technology. Turbo Mode in Westmere allows increase frequency about 4 steps ($4 \times 133 \text{ MHz}$) for a single core and for six cores it is two step (266 MHz). Nehalem-EX turbo scenario is not so aggressive and for 1 core we have increase of 3 steps but for 8 cores we get 1 step (133 MHz).

As we see in Table 2 single node system performance on LINPACK is 125.28 GFLOPS and 120.95 GFLOPS, respectively for the two platforms. The achieved performance is 86% of the theoretical performance for both the X7560 and the X5670 platforms. Turbo Mode does not bring a lot of benefits to LINPACK, in both cases it improves performance about 1%. For very well optimized threaded codes like LINPACK Turbo Mode will not bring a lot of benefits because the code itself stresses the CPU to the limit, there is no headroom left for the frequency increase as thermal design power (TDP) reaches the limit already.

Table 1
Processors characteristic

Processor type	Intel Xeon processor X7560	Intel Xeon processor X5670
Technology(nm)	45	32
Cores per socket	8	6
Intel Turbo Mode	Yes	
Intel HT	Yes	
Core frequency (MHz)	2266	2930
L1 cache size	32 KB Inst / 32 KB Data per core	
L2 cache size	No256 KB per core	
L3 cache size	24 MB	12 MB
Integrated memory controller	Yes (2 × SMI)	Yes
Front side bus FSB	No	
Memory transfer rate/socket (GB/s)	34	32
Intel Quick Path Interconnect	Yes – 4 Links (6.40, 5.86 or 4.80 GT/s)	Yes – 2 links (6.40, 5.86 GT/s)
Number of threads /core	2	
TLB Page Size	4 KB, 2 MB, 1 GB	4 KB, 2 MB, 1 GB
TDP (W)	130	95

As the second aspect of platform performance we have evaluated is a bandwidth. The benchmark we have used to measure the bandwidth is the STREAM benchmark. It estimates, both memory reads and memory writes. This gives us an indication of how effective the memory subsystem is. The STREAM benchmark gives only results of memory system efficiency ignoring the cache efficiencies of the tested platforms.

STREAM is also not an ideal benchmark for Intel Xeon 7500 family as Intel Xeon 7500 for coherency reasons must issue a read before the write instruction is executed so this extra read, isnt counted by the STREAM benchmark. The STREAM benchmark does a simple sequence of 2 reads and 1 write however the Intel Xeon 7500 has the capability to perform an extra read before the write instruction is executed. Unfortunately this extra read degrades the benchmark performance with this specific workload but real applications will see the benefit of this extra read cycle and as a result will deliver an additional 10 GB/s peak bandwidth compared to the result reported by the STREAM benchmark [5].

New platforms based on Intel Xeon 7500 also require the correct DIMM socket population. Maximum bandwidth will be delivered when both memory controllers are populated with 2 DIMMs per channel. Any other configuration will generate a drop in relative memory bandwidth available to the applications. This is completely different from how the Xeon processor 5600 series-based servers behave.

Table 2
Single node system performance

Processor type	Intel Xeon processor X7560	Intel Xeon processor X5670
Core frequency (MHz)	2266	2930
Core performance (GFLOPS)	9.06	11.72
CPU Peak Performance (GFLOPS)	72.51	70.32
Peak node performance (GFLOPS)	145.02	140.64
Turbo model	1/1/2/2/3/3/3/3/3	2/2/3/3/4/4
Core Turbo performance (GFLOPS)	10.66	13.85
Peak Turbo CPU performance (GFLOPS)	76.76	76.7
Single node LINPACK (GFLOPS)	125.28	120.95
Single Turbo node LINPACK (GFLOPS)	126.08	122.21
Single node LINPACK with HT-ON (GFLOPS)	123.54	119.53
Stream single node (GB/s)	43.7	42
Stream single node HT-ON (GB/s)	41.2	41.35

Intels Hyper Threading Technology has also been evaluated. The study clearly shows that when Intel Hyper Threading Technology is enabled it degrades LINPACK results by 1% and reduces STREAM results by 6% [6].

3. Cluster Performance

The platform level evaluation and LINPACK and STREAM benchmark results do not provide the full picture of the system capability in HPC scenarios. In this section we evaluate the two HPC clusters based on exactly the same configurations as those used for the platform evaluation. Keeping in mind the best evaluation approach we recognize that application performance is the ultimate measure of system capability. However analyzing HPC Challenge benchmark results allow us to predict cluster level behavior in typical HPC operational condition. The HPC Challenge benchmark is a suite of tests that examines the performance of HPC architectures in a more challenging way than LINPACK and STREAM by the way LINPACK and STREAM are also part of HPC Challenge. The suite is a collection of several well known computational kernels. The HPC Challenge benchmark test suite stresses not only the processors, but also the memory system and interconnects. It is a better indicator of how HPC system will perform across a spectrum of real-world applications.

The HPC Challenge benchmark consists of basically 7 tests: HPL, DGEMM, STREAM, PTRANS, FFT, Random Order Ring Bandwidth and Random Ordered Ring Latency [7].

In this section we will evaluate 16 node clusters based on the Intel Xeon X7560 vs. Intel Xeon X5670 across all the benchmark listed above. Cluster configuration details are listed in Table 3.

Table 3
Cluster configuration details

System name	IBM X3690 X5	Urbanna/Westmere
Processor type	Intel Xeon processor X7560	Intel Xeon processor X5670
Number of nodes	16	16
Number of cores/node	16	12
Local memory/node (GB)	64	48
Total memory on 16 nodes	1024	768
Memory speed (MHz)	1066	1333
Memory type	DDR3	
Memory transfer rate/socket (GB/s)	34	32
Interconnect type	ConnectX QDR Infiniband	
Operating system	Platform Pmm1.2 (rhel 5.3 2.6.18-128.el5)	
System manufacture	IBM	Intel
Intel C compiler	l_cproc_p.11.1.064	
MKL Library	intel-cpromkl064-11.1-1	
MPI	intel-mpi-em64t-3.2.2p-006	

Since the HPL benchmark is very well optimize and mainly depends on clock frequency, in consequence the faster system is this one with the higher clock CPUs. In our experiment Urbanna/Westmere system has the faster clock CPUs 2.93 GHz and HPL benchmark runs on this system faster than on IBM X3690 X5 where clock of installed CPUs is 2.26 GHz. Difference in clock is 23% almost identical as performance difference measured by HPL.




Fig. 2. Performance of HPL

Figure 2 shows performance of HPL for both systems. With 112 cores it was 880 GFLOPS, 1147 GFLOPS, IBM X3690 X5 and Urbanna/Westmere respectively. Urbanna/Westmere based cluster utilized 192 cores only and archived 2030 GFLOPS but IBM X3690 X5 has 256 cores and achieved 2013 GFLOPS when utilizing all the installed CPUs. The 16 nodes IBM X3690 X5 and cluster configuration with the same

physical size thermal and power envelope like Urbanna/Westmere benefits from 64 more cores.

As the IBM X3690 X5 each node has a bigger installed memory 64 GB (4GB/core) vs. 48 GB (4GB/core), but each core in the system has the same ratio memory per core. To make a fair comparison the size of data sets has been selected to minimize size of memory effects.




Fig. 3. Performance of DGEMM

Figure 3 shows performance of parallel DGEMM (matrix-matrix multiplication) for the two systems plus evaluation of Turbo capability implication for this benchmark. Urbanna/Westmere has the highest theoretical one-core peak of 11.72 GFLOPS. The IBM X3690 X5 cluster has theoretical one-core peak performance of 9.04 GFLOPS, but with enabled Turbo capability this gives 13.85 GFLOPS and 10.64 GFLOPS respectively. Overall the Urbanna/Westmere with Turbo ON enabled has obtained the best performance 12.31GFLOPS for 1 core and 11.12 GFLOPS for 192 cores, followed by Urbanna/Westmere with Turbo OFF 11.10 GFLOPS for 1 core 10.7 GFLOPS for 192 cores followed by IBM X3690 X5 with Turbo ON 9.04 GFLOPS for 1 core 8 GFLOPS for 256 core and IBM X3690 X5 with Turbo OFF 8.04 GFLOPS for 1 core 8 GFLOPS for 256 cores. The achieved performance was 88%, 94%, 84% and 89% of the peak performance on, Urbanna/Westmere with Turbo ON and Urbanna/Westmere with Turbo OFF and IBM X3690 X5 with Turbo ON and IBM X3690 X5 with Turbo OFF respectively.

Figure 4 illustrates memory bandwidth for each cluster using STREAM Triad benchmark. For a single core, the measured bandwidths were 11.1 GB/s, and 10.72 GB/s for Urbanna/Westmere and IBM X3690 X5. For 8 cores, we have seen degradation to level of 6.8 GB/s and 5.2 GB/s respectively. For 48 cores, it was 4.8 GB/s and 5.2 GB/s for Urbanna/Westmere and IBM X3690 X5. For 192 and 256 cores, it was 4.7 GB/s and 4.9 GB/s for Urbanna/Westmere and IBM X3690 X5. The

different implementation of Integrated Memory Controller in both platforms as well as different memory type and specially the nature of the STREAM benchmark do not allow for doing a real comparison between two platforms and demonstrates an advantage of Intel Xeon X7560 based platform.




Fig. 4. Performance of STREAM

In Figure 5, we show performance of the PTRANS benchmark. The PTRANS benchmark performance depends on the network and on memory bandwidth. The two systems use QDR IB and the same topology of interconnect. Performance was almost identical on both systems due the same interconnect and equal sustained-memory bandwidth of the tested platforms. The memory subsystem based on DDR3-1333 MHz based platform gave less than 1% versus DDR3-1066 MHz based platform. At 96 cores, bandwidth was 16.3 GB/s and 16.1 GB/s on Urbanna/Westmere and IBM X3690 X5 respectively. At 192 cores, bandwidth was 29.65 GB/s and 29.68 GB/s on Urbanna/Westmere and IBM X3690 X5 respectively. Total utilization of 256 cores on IBM X3690 X5 offers bandwidth of 54 GB/s.




Fig. 5. Performance of PTRANS

Figure 6, illustrates achieved performance of the Random Access benchmark. The benchmark uses SANDIA_OPT2 algorithm as Giga Updates per second (GUP/s). Both systems used identical QDR IB and almost equal memory bandwidth so results on both clusters are also the same [8]. Until 96 cores, the results were 0.48 GUP/s 0.47 GUP/s for Urbanna/Westmere and IBM X3690 X5 respectively. The Urbanna/Westmere system at 192 cores achieved 0.96 GUP/s and at 256 cores IBM X3690 X5 obtained 1.25 GUP/s.




Fig. 6. Performance of Random Access

Figure 7 shows performance of the FFT benchmark. Both systems used Intel MKL and intel-cpromkl064-11.1-1. The benchmark performs as mixture of flops, memory, and network bandwidth. The higher clock and faster memory subsystem DDR3-1333MHz deliver 10%–15% better results for Urbanna/Westmere versus IBM X3690 X5 cluster. In all two systems we observed almost perfect scaling. At 64 cores; performances were 53.2 GFLOPS and 48 GFLOPS on Urbanna/Westmere and IBM X3690 X5 respectively.




Fig. 7. Performance of FFT

At 96 cores; performances were 86.7 GFLOPS and 77 GFLOPS on Urbanna/Westmere and IBM X3690 X5 respectively. For 192 cores, performances were 143.8 GFLOPS and 122.23 GFLOPS on Urbanna/Westmere and IBM X3690 X5 respectively. The additional cores installed in the IBM X3690 X5 cluster (256 cores) compensate for the lower clock and slower memory subsystem and performance on fully utilized cores cluster gives the best results 157 GFLOPS.




Fig. 8. Performance of Random Order Ring Bandwidth

In Figure 8, we show the Random Order Ring bandwidth for the two systems. This benchmark measures conflict in the network and reports bandwidth achieved per core in a ring communication model. The Random Order Ring bandwidth measures the accumulated bandwidth of the communication network of parallel computing systems. The algorithm uses an average, short and long messages transferred with different bandwidth values. Because two systems have identical interconnect and topology of communication network the results are very comparable. At 48 cores, Random Order Ring bandwidth for Urbanna/Westmere and IBM X3690 X5 was 391 MB/s and 398 MB/s, respectively, at 96 cores 325MB/s and 337MB/s and at 192 cores 290 MB/s 321 MB/s.

Figure 9 shows Random Ordered Ring (ROR) latency for the two systems.




Fig. 9. Performance of Random Order Ring Latency

Inside a node (16 or 12 cores), latency of all two systems was similar $0.7 \mu\text{s}$ for Urbanna/Westmere and $0.8 \mu\text{s}$ for IBM X3690 X5. Those values increased when the communication went out of the single node. At 48 cores, latency was $2.5 \mu\text{s}$, $2.4 \mu\text{s}$ for Urbanna/Westmere and IBM X3690 X5 respectively. At 96 cores, latency was $3.4 \mu\text{s}$ and $3.5 \mu\text{s}$ for Urbanna/Westmere and IBM X3690 X5 respectively. At 192 cores on Urbanna/Westmere achieved a level $4.5 \mu\text{s}$ and for 256 cores IBM X3690 X5 reached $5.6 \mu\text{s}$

4. Conclusion

In this paper we have compared two families of clusters on typical HPC benchmarks; the new Eight-Core Intel Xeon X7560 processor based system and installation utilizing Six-Core Intel Xeon X5670 processor. We have found that compared platforms and clusters based on them are almost equal and they behave as we have been expecting taking in to account theoretical performance of both CPUs.

We see that more cores implemented in Intel Xeon X7560 processor compensate lack of clock frequency. On single node level there is only a 5 GFLOPS difference which represents 4% of the tested performance. On 16 nodes clusters both installations perform equal based on HPL benchmark and the difference is less than 1%. In our study we have clearly indicated that for benchmarks which are scaling with number of cores cluster based on Intel Xeon X7560 processor has advantage but if the benchmark does not scale with cores the cluster with Intel Xeon X5670 processor is equal to Eight-Core Xeon or even performing better.

This performance advantage generated by bigger number of cores benefits not only HPL benchmark but also PTRANS benchmark as well Random Access benchmark and FFT benchmark. In all those benchmarks we clearly observe that more cores deliver performance improvements. The reverse situation can be seen on the DGEMM benchmark, STREAM benchmark, ROR benchmark and RORL benchmark. In such selection of benchmark more cores do not bring a performance increase or even generate degradation of benchmark value. These benchmarks definitely favor higher clock and clearly Intel Xeon X5670 based cluster achieved better scores.

We also examined technologies; Intel Turbo Boost Technology and Intel Hyper Threading Technology and their impact for HPC. Intel Turbo Boost Technology delivers some performance improvements to compute-intensive codes like HPL or DGEMM the impact is smaller when the application is memory-intensive. Intel Hyper-Threading Technology for HPC applications has not been seen to offer any performance improvement in this segment and even may generate same performance degradation.

In summary, we can state that new Intel Xeon X7560 based platform brings more scalability and performance for core scaling oriented application and can be consider as good choice for many of the HPC installations.

Acknowledgements

We gratefully acknowledge the help and support provided by Jamie Wilcox and Victor Gamayunov from Intel EMEA Technical Marketing HPC Lab.

References

- [1] *TOP 500 Sublist generator*. <http://top500.org/sublist/results>
- [2] Gepner P., Fraser D.L., Kowalik M.F.: *Multi-Core Processors: Second generation Quad-Core Intel Xeon processors bring 45 nm technology and a new level of performance to HPC applications*. ICCS. 2008, pp. 417–426.
- [3] Dongarra J., Luszczek P., Petitet A.: *Linpack Benchmark: Past, Present, and Future*. <http://www.cs.utk.edu/~luszczek/articles/hplpaper.pdf>.
- [4] Gepner P., Kowalik M.F.: *Multi-Core Processors: New Way to Achieve High System Performance*. [in:] PARELEC 2006, pp. 9–13.
- [5] Barker K. J., Davis K., Hoisie A., Kerbyson D. J., Lang M., Pakin S., Sancho J. C.: *A Performance Evaluation of the Nehalem Quad-core Processor for Scientific Computing*. *Parallel Processing Letters*, vol. 18, No. 4, 2008, pp. 453–469.
- [6] Eggers S. J., Emer J. S., Levy H. M., Lo J. L., Stamm R. L., Tullsen D. M.: *Simultaneous multithreading: A platform for next generation processors*. *Proc. IEEE* 17, 1977, pp. 1219.
- [7] *HPC Challenge Benchmarks*. <http://icl.cs.utk.edu/hpcc/>.
- [8] Saini S., Talcott D., Jespersen D., Djomehri J., Jin H., Biswas H.: *Scientific Application-based Performance Comparison of SGI Altix 4700, IBM Power5+, and SGI ICE 8200 Supercomputers*. *Proc. of the 2008 ACM/IEEE Conference on Supercomputing*, Austin, Texas, November 15–21, 2008.