LA-UR- 09-1380

Approved for public release;
distribution is unlimited.

Title: A Look at the Performance of AMD Shanghai

Author(s): Kevin Barker, CCS-1
Marion Kei Davis, CCS-1
Adolfy Hoisie, CCS-1
Darren Kerbyson, CCS-1
Michael K. Lang, CCS-1
Scott Pakin, CCS-1
Jose Sancho, CCS-1

Intended for: Electronic Distribution

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

A Look at the Performance of AMD Shanghai

Kevin Barker, Kei Davis, Adolfo Hoisie, Darren Kerbyson, Mike Lang, Scott Pakin, Jose Sancho
Performance and Architecture Lab (PAL)
Centre for Advanced Architectures and Useable Supercomputing (CAAUS)
Los Alamos National Laboratory
February 2009

This document summarizes a performance comparison between a four processor Shanghai node, a four-processor Barcelona node, and a two processor Nehalem node. All of these processors contain four cores. A comparison of achievable memory bandwidths, and application performance on a single-core, a single-processor (4-cores), and a single node are detailed. The characteristics of each processor are listed in Table 1. below.

Table 1. Summary of the processor characteristics

	Barcelona (AMD)	Shanghai (AMD)	Nehalem (Intel)
Architecture	Opteron	Opteron	Xeon
Clockspeed (GHz)	2.0	2.5	2.8
L1 Cache / core	128KB	128KB	64KB
L2 Cache / core	512KB	512MB	256KB
L3 Cache (shared)	2MB	6MB	8MB
Main memory	2x DDR2x667	2x DDR2-667	3x DDR3-1333

The analysis shows:

- 1) The memory bandwidth of Shanghai on a node has increased by ~30% compared to Barcelona (slightly above the increase in processor clock-speeds as measured), but remains significantly lower than Nehalem.
- 2) Application performance on Shanghai is higher than on Barcelona on average by 40%, 35% and 42% on a single-core, single-processor, and single-node, respectively. Again this is above the increase in processor clock-speeds.
- 3) Application performance on a 4-processor Shanghai node is between 0.6x and 1.35x the speed of a 2-processor Nehalem node. Some applications are faster on the Shanghai node and some are slower.
- 4) When idle the Shanghai node consumes 285W which increases up to be between 430W and 478W when using all 16-cores depending on the application.

1. Memory Bandwidth

To examine the memory bandwidth per core we ran the MPI version of Streams using the triad test. In Figure 1(a) the aggregate memory bandwidth is shown for two cases: when using a single processor and when using all processors in the node. The number of cores per processor used in both cases is varied from one to four. The bandwidth observed on a single Shanghai processor was less than on Barcelona. However the bandwidth observed is higher, by 30%, when using all 16 cores in a node. It can also be seen that the two-processor Nehalem significantly outperforms both AMD nodes in all cases.

Figure 1(b) is based on the same data as Figure 1(a) but presents the observed memory bandwidth per core. As shown, the per-core bandwidth decreases from 4.4GB/s to 1.1GB/s for the Barcelona (a factor of four decrease), from 2.55GB/s to 1.55GB/s for the Shanghai (a factor of 1.65 decrease), and from 12GB/s to 4.6GB/s for Nehalem (a factor of 2.6 decrease).

(a) Aggregate memory bandwidth

(b) Memory bandwidth per core

Figure 1. Streams Memory Bandwidth

2. Application performance

The suite of applications utilized here includes many large-scale production applications that are currently used within the U.S. Department of Energy. The testing for each application consisted of

- 1) comparing the performance on a single core of the Shanghai, Barcelona, and Nehalem;
- 2) examining the scaling behavior for two cases: (a) using only one processor, and (b) using all processors in a node;
- 3) determining the configuration (processors and cores per processor) that yields the best performance.

All of the applications are typically run in a weak-scaling mode. That is, the global problem size grows with the number of processors in the system. All available memory is typically used for increased fidelity in the physical simulations or for simulating larger physical systems. Our approach mimics typical usage by fixing the sub-problem per processor no matter how many cores per processor are used. This can be stated succinctly as doing strong scaling within a processor and weak scaling across processors.

The applications are summarized below, and the input-decks used are summarized in Table 2.

MILC – is a Lattice QCD application from Indiana University. It uses a 4-D space-time lattice to calculate the coupling between subatomic particles such as quarks and gluons.

S3D – high-fidelity 3D simulation of turbulent combustion that includes detailed chemistry. It originates from Sandia National Laboratory.

SAGE – an adaptive mesh (AMR) hydrodynamics code used for the simulation of shock-waves. Developed jointly by Los Alamos and SAIC.

SPaSM – the Scalable Parallel Short-range Molecular dynamics code from Los Alamos. Used to study material fracture and deformation properties.

Sweep3D – a code-kernel from Los Alamos that implements deterministic S_N transport. The computation is in the form of wavefronts that originate at the corners of a 3-D physical space.

VH1 – the Virginia Hydrodynamics code simulates an ideal inviscid compressible flow of gas hydrodynamics and is capable of simulating three-dimensional turbulent stellar flows.

VPIC – a Particle-In-Cell code from Los Alamos used to model particle flow within a plasma.

Table 2. Summary of the input decks used for each application

	Input deck	Problem per processor	Problem per node	Memory use per processor	Processing characteristic
GTC	1Dwedge	6.2M particles	24.8M particles	320MB	Particle based
MILC	SU3_RMD	8x8x8x32	16x16x16x16	32MB	Memory intensive
S3D	typical	50x50x50	100x100x50	140MB	Memory/Compute intensive
SAGE	timing_h	140K cells	560K cells	280MB	Memory intensive
SPaSM	BCC	64x64x64	128x128x64	150MB	Compute intensive
Sweep3D	Pencil	20x10x400	40x20x400	8MB	Kernel, small memory footprint
VH1	Shock_tube	200x200x200	400x400x200	900MB	Memory/Compute intensive
VPIC	3D-HOT	4M particles	16M particles	256MB	SSE optimized, memory intensive

2.1. Single-core performance comparison

The performance of each application on a single core of each of the three processors is shown in Figure 2(a). The metric *time* denotes the iteration time of the main computational loop (without I/O) for all applications except Sweep and Partisn, for which it is 10 iterations, and VPIC, for which it is 100 iterations, to aid visual comparison. The relative runtime, between Shanghai and Barcelona, and Shanghai and Nehalem, is shown in Figure 2(b). A value of one indicates the same performance between two processors, and a value of two indicates a 2x performance advantage of Shanghai. As shown, a Shanghai core is between 1.1 and 1.9 times faster than a

Barcelona core, and between 0.45 and 0.97 the performance of a Nehalem core. Note that for MILC and SPaSM a 2.2GHz Barcelona node was used for all performance data.

2.2. Single-processor performance comparison

The performance of each application on a single processor is shown in Figure 3(a). The relative runtimes between Shanghai and Barcelona, and Shanghai and Nehalem are shown in Figure 3(b). As shown, the Shanghai processor achieves performance between 1.1 and 2.0 times higher than a Barcelona processor, and between 0.44 and 0.9 the performance of a Nehalem processor.

2.3. Node performance comparison

The relative performance between the Shanghai and Barcelona nodes, and between the Shanghai and Nehalem nodes is shown in Figure 4. Note that the relative performance was calculated using the effective processing rate for each application since the core count and hence the problem size for each application differed between the nodes. As shown, the Shanghai node achieves performance between 1.13 and 2.0 times higher than the Barcelona node, and between 0.6 and 1.35 the performance of the Nehalem node. Note that this data is based on a four-processor Shanghai node vs. a two-processor Nehalem node and naturally biases this analysis towards the larger node.

Figure 4 – Shanghai performance advantage (Node)

4. Application scalability analysis

We followed a strict process in order to analyze the performance of using multiple cores: the problem size per processor was held constant for all tests. This can be considered as weak-scaling across processors and strong-scaling within processors as more cores per processor are used. For each application we show the performance when using between one and four cores per processor for the case of using one processor, shown in Figure 5, and all the processors in a node, shown in Figure 6.

The scalability within a single processor is reasonably consistent for the three processors on all applications. The scalability within a node is also reasonably consistent between Shanghai and Barcelona. Notable exceptions are GTC and S3D that show improved scaling, and SAGE which shows worse scaling. Also note that the maximum speedup on both the Shanghai and Barcelona nodes is 16 whereas it is 8 on the Nehalem.

Figure 5 – Application speedup on a single processor

Figure 6 – Application speedup when using all processors in a node.

5. Power

The power consumption of the Shanghai node was also measured at the same time as measuring the performance of each of the applications. A summary of the power average power consumed for each application when using a single core, a single processor (4 cores), and a single node (16 cores) is shown in Figure 5.

As shown, there is some variation in the average power consumption depending on application, between 430W and 478W. However, the first order effect is in the number of cores used. When using only a single-core, the average power consumption is between 294W and 308W. Indeed when no application was running, the background power consumption was observed to be 285W.

Figure 5 – Average Power Consumption of the Shanghai node for each application.

Acknowledgements

We are thankful to Greg Branch (AMD) for loaning a Shanghai node for this testing.