ISTITUTO NAZIONALE DI FISICA NUCLEARE

Sezione di Padova

INFN/CCR-07/9
November 30, 2007

CCR/16/07/P

PERFORMANCE OF DUAL PROCESSOR SERVER ON HEP CODE

Michele Michelotto[1)]

[1)] *INFN-Sezione di Padova, Via F. Marzolo, 8, I-35131 Padova, Italy*

Abstract

Evaluation of cpu architecture for dual processor Worker Node. The performances of dual processor WN are evaluated using some programs commonly used in HEP environment. Finally the performances are compared to the clock of the cpu to assess and to the Spec 2000 CpuInt (commonly called specint) taken from the spec.org site.

PACS.: 89.20-Ff (Computer Science and Technology)

*Published by **SIS–Pubblicazioni***
Laboratori Nazionali di Frascati

1 METHODOLOGY

If the dual processor machine has single core chip we run one and two instances of the code in parallel. If the dual processor machine has dual core chip we run up to four instances of the code in parallel.

The performances drop is always under 1% and sometimes performances showed a slightly enhancement meaning that those are statistical fluctuations and that one can generally assume that it is safe to run N copies of a programs if the computer has N cores.

2 WN SERVER UNDER TEST

- Dual Processor dual core Woodcrest 5160 3.0 GHz with 3GB of DDR memory (FMDIMM), SAS disk on loan from E4 thanks to Intel Italy. This is the most powerful machine available in this moment so I take is reference machine.
- Dual Processor AMD 265 dual core 1.80 GHz with 4GB of DDR1 memory, SATA disk available at INFN Padova computing centre.
- Dual Processor AMD single 248 2.2 GHz with 2 GB and with 8 GB of DDR1 memory and SATA disk available at INFN Padova computing centre.
- Dual processor dual core AMD 2218 (Socket F) 2.6 GHz with 8 GB DDR2 memory on loan from AMD Italy (actually this server had a pre-release version of the 2218 chip so performance could improve in a future release).

3 HEP PROGRAM USED

The programs used are taken from a suite prepared by Hans Wenzel from Fermilab and used to perform tests recently presented at a CHEP conference. Hans has been very helpful in solving the first problems I encountered to run the suite.

4 ROOTMARK

Root is an framework for the analysis of HEP events developped at CERN. In our case I measured the performances with a stress test that concentrates on mathematical calculation avoiding the graphical aspects. The result is a number called “Rootmark”. Biggest numbers mean better performances.

The performances running at 64 bit on a 64 bit (64/64) operating system are much better (from 127 to 146%) than running a 32bit on 64 bit o.s. (32/64) or 32/32. The Rootmarks are displayed on pale green background and is the number taken for a single core when the machine is fully loaded with all other core running the same program.

	clock GHz			SI2K base	rootmarks 64/64	rootmark 32/64	64 bit 32 bit ratio rootmark	rootmarks 64/64 clock ratio	rootmark 32/64 clock ratio	rootmarks 64/64 SI2K ratio	rootmark 32/64 SI2K ratio
Intel 5160	3.00	3G	4 cores	**3061**	2005.80	1584.60	**127%**	668.6000	528.2000	655.276	517.674
AMD 265	1.80	4G	4 cores	1168	1083.40	743.30	**146%**	601.8889	412.9444	927.568	636.387
AMD 248	2.20	2G	2 cores	1411	1273.70	884.60	**144%**	578.9545	402.0909	902.693	626.931
AMD 248	2.20	8G	2 cores	1411	1299.50	884.20	**147%**	590.6818	401.9091	920.978	626.648
AMD 2218	2.60	8G	4 cores	**1648**	1546.40	1079.60	**143%**	594.7692	415.2308	938.350	655.097
Intel 5160	3.00	3G	4 cores	100.00%	100.00%	100.00%	**100%**	100.00%	100.00%	100.00%	100.00%
AMD 265	1.80	4G	4 cores	38.16%	54.01%	46.91%	**115%**	90.02%	78.18%	141.55%	122.93%
AMD 248	2.20	2G	2 cores	46.10%	63.50%	55.82%	**114%**	86.59%	76.12%	137.76%	121.11%
AMD 248	2.20	8G	2 cores	46.10%	64.79%	55.80%	**116%**	88.35%	76.09%	140.55%	121.05%
AMD 2218	2.60	8G	4 cores	53.84%	77.10%	68.13%	**113%**	88.96%	78.61%	143.20%	126.55%

The fastest processor is the Intel 5160 Woodcrest in absolute.

If we divide Rootmark by clock (light blue background) we can see that the Intel chip is still the most efficient. This means that an AMD chip extract 90% of Rootmark from a GHz with respect to the Intel.

If we divide Rootmark by the specint figure - published in the spec.org web site – we see that the AMD processor are more efficients (green background). This can be read that the AMD chip produces more Rootmark from a Specint compared to Intel (on 64/64).

5 PYTHIA

Pythia is an HEP events generator. In our case we create 300 SUSY events. More event/sec means better performances.

	clock GHz			SI2K base	pythia 64/64 evt/sec	pythia 32/64 evt/sec	64 bit 32 bit pythia ratio	pythia 64/64 clock ratio	pythia 32/64 clock ratio	pythia 64/64 SI2K ratio	pythia 32/64 SI2K ratio
Intel 5160	3.00	3G	4 cores	**3061**	189.15	152.86	**124%**	63.0500	50.9533	61.794	49.938
AMD 265	1.80	4G	4 cores	1168	110.78	92.40	**120%**	61.5444	51.3333	94.846	79.110
AMD 248	2.20	2G	2 cores	1411	135.29	116.70	**116%**	61.4955	53.0455	95.882	82.707
AMD 248	2.20	8G	2 cores	1411	135.48	116.69	**116%**	61.5818	53.0409	96.017	82.700
AMD 2218	2.60	8G	4 cores	**1648**	162.00	134.41	**121%**	62.3077	51.6962	98.301	81.559
Intel 5160	3.00	3G	4 cores	100.00%	100.00%	100.00%	**100%**	100.00%	100.00%	100.00%	100.00%
AMD 265	1.80	4G	4 cores	38.16%	58.57%	60.45%	**97%**	97.61%	100.75%	153.49%	158.42%
AMD 248	2.20	2G	2 cores	46.10%	71.53%	76.34%	**94%**	97.53%	104.11%	155.17%	165.62%
AMD 248	2.20	8G	2 cores	46.10%	71.63%	76.34%	**94%**	97.67%	104.10%	155.38%	165.61%
AMD 2218	2.60	8G	4 cores	53.84%	85.65%	87.93%	**97%**	98.82%	101.46%	159.08%	163.32%

Again the Woodcrest is the most powerful processor and again we can see that at 64/64 we runs 120% faster.

The number of events/sec per clock is more or less the same for the intel and the AMD chip (this is rather amazing). Of course this means that in terms of Event/sec per Specint the Intel chip that as very big number of Specint/Clock suffers from the comparison. It looks like an AMD chip give 150% more performance per Specint than Intel.

Obviously the experiments are interested in buying the worker nodes that give the best figure in term of event/sec per money taking in account also the power consumption that influences the Total Cost of Ownership (TCO) increasing the Electricity bill but also the cost of the Computing Centre (UPS, Power Distribution, Rack Cooling, Emergency Power Generator). In this case I didn’t take in account these costs since some machine have been bought in the past months while other were on loan from local computer vendor.

6 CMS_SW

This software for the moments runs only at 32 bit (natively 32/32 or 32/64).

With OSCAR we reports the performances in terms of events/sec simulating 300 single pions events. With DIGIS and DST we run Digitization and Reconstructions of 100 HCAL events.

	clock GHz			SI2K base	CMS DST evt/sec	CMS DIGIS evt/sec	CMS OSCAR evt/sec	CMS DST clock ratio	CMS DIGIS clock ratio	CMS OSCAR clock ratio	CMS DST SI2K ratio	CMS DIGIS SI2K ratio	CMS OSCAR SI2K ratio
Intel 5160	3.00	3G	4 cores	**3061**	0.5997	0.0703	0.1713	0.1999	0.0234	0.0571	0.1959	0.0230	0.0560
AMD 265	1.80	4G	4 cores	1168	0.2850	0.0340	0.0909	0.1583	0.0189	0.0505	0.2440	0.0291	0.0778
AMD 248	2.20	2G	2 cores	1411	0.3829	0.0503	0.1083	0.1740	0.0229	0.0492	0.2714	0.0356	0.0768
AMD 248	2.20	8G	2 cores	1411	0.3865	0.0503	0.1084	0.1757	0.0229	0.0493	0.2739	0.0356	0.0768
AMD 2218	2.60	8G	4 cores	**1648**	0.4258	0.0486	0.1311	0.1638	0.0187	0.0504	0.2584	0.0295	0.0796
Intel 5160	3.00	3G	4 cores	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
AMD 265	1.80	4G	4 cores	38.16%	47.52%	48.36%	53.06%	79.21%	80.61%	88.44%	124.55%	126.75%	139.07%
AMD 248	2.20	2G	2 cores	46.10%	63.85%	71.55%	63.22%	87.07%	97.57%	86.21%	138.51%	155.22%	137.15%
AMD 248	2.20	8G	2 cores	46.10%	64.45%	71.55%	63.28%	87.88%	97.57%	86.29%	139.81%	155.22%	137.28%
AMD 2218	2.60	8G	4 cores	53.84%	71.00%	69.13%	76.53%	81.93%	79.77%	88.31%	131.88%	128.41%	142.15%

Again the results are similar to those obtained with Rootmarks with a clear leadership of the Intel chip in absolute terms or dividing by the clock. While the AMD chip is 30% more efficient with respect to the declared Specint.

This advantage is more evident on the old single cores chip but remains clearly visible also in the dual core chip that are very similar (AMD265 vs AMD2218).

7 CONCLUSIONS

At the beginning of 2006 (Hepix Spring in Rome). AMD had a clear technological leadership with the first processors that ware fully dual core. Those chip permitted to double the throughput performances at roughly the same price and power consumption of single cores Intel and AMD processors. The new Woodcrest family gave the leadership back to Intel thanks to a completely new chip design and a reduced feature size. The 65nm technology permits higher clock (up to 3.0 GHz) keeping power consumption very low, lower than AMD comparing the processor and about the same when looking the fully configured machine.

On the other hand the performances of the Woodcrest are lower than one could expect when extrapolating from the published specint number.

A sensible increase of performance can be obtained when code is ported to 64 bit. I could not see differences when running on machines with different amounts of physical memory (this means that up to 1GB per core there is no drop of performances).

This means that when a laboratory or an agency is buying new server for computing farms they should take in account not only the specint (measured or declared) but also the real codes that will run on the computing nodes.

8 FUTURE ACTIVITIES

I can see several place where this paper could be improved:

1. Understand if the AMD 2218 (socket F) in its final releases will improve the performance.
2. Measuring the Energy efficiency in terms of performances per Volt-Ampere (obviously the whole server not only the chip).
3. Taking in account the cost of the machine.
4. Finding other benchmark from other experiments (Alice, Atlas, Babar, LHCB).
5. Testing the new quad-core chip (53xx Clovertown) to check the presence of bottleneck or performances drop due to memory.
6. Testing the quad-core chip from AMD.
7. Testing the quad-core for desktop Kentsfield from Intel.
8. Testing the performance with compiler more efficient than gcc.