

Benchmarking AMD64 and EMT64

Hans Wenzel*, Oliver Gutsche, FNAL, Batavia, IL 60510, USA
Mako Furukawa, University of Nebraska, Lincoln, USA

Abstract

We have benchmarked various single and dual core 64 Bit CPU's from AMD (AMD64) and Intel (EMT64) using various HEP and CMS specific applications. This results are compared with results obtained with 32 Bit (IA32) CPU's.

INTRODUCTION

There are now two 64-bit implementations in the Intel compatible processor marketplace: the Intel EM64T and AMD AMD64. For information about the architecture of the Opteron chip see [1],[2],[3]. In addition, dual-core CPU's are now available. They promise to nearly double the computing power while not increasing the infrastructure cost especially for cooling and electricity significantly compared to boxes equipped with single-core chips [4]. We should also benefit from: lower space requirements and need for fewer racks, fewer console connections and fewer network connections. Although the later might adversely affect the performance when one network interface is shared by 4 instead of 2 CPU-cores. We were especially interested in how the performance of the multicore processors scales when each CPU core is running a copy of the application simultaneously and how much power these chips consume. Table lists the various processor types that we evaluated.

OPERATION MODES

There are three distinct operation modes available for the AMD64 and EM64T architectures:

- **32-bit legacy mode (32Bit/32Bit)**

In this mode, both AMD64 and EM64T processors will act just like any other IA32 compatible processor. You can install your 32-bit OS on such a system and run 32-bit applications, however, you will not be able to make use of the new features such as the flat memory addressing above 4 GB or the additional General Purpose Registers (GPR).

- **Compatibility mode (64Bit/32Bit)**

This is an intermediate mode of the full 64-bit mode described below. Compatibility mode gives you the ability to run a 64-bit operating system while still being able to run unmodified 32-bit applications. Each

Processor	Speed	Architecture (cores)
Opteron 244	1.8 Ghz	AMD-64 (×1)
Opteron 265	1.8 Ghz	AMD-64 (×2)
Opteron 246	2.0 Ghz	AMD-64 (×1)
Opteron 246HE	2.0 Ghz	AMD-64 (×1)
Opteron 270	2.0 Ghz	AMD-64 (×2)
Athlon 64 3500+	2.2 Ghz	AMD-64 (×1)
Athlon 64 4200+	2.2 GHz	AMD-64 (×2)
Opteron 248	2.2 Ghz	AMD-64 (×1)
Opteron 275	2.2 Ghz	AMD-64 (×2)
Opteron 250	2.4 Ghz	AMD-64 (×1)
Xeon	3.4 Ghz	EM64T (×1)
Xeon	3.6 Ghz	EM64T (×1)
Xeon	2.4 Ghz	IA32 (×1)
Xeon	2.8 Ghz	IA32 (×1)

Table 1: *Processors tested. We used single CPU mainboards for the Athlon 64 CPU's, all the other processors were tested with dual CPU mainboards. The 246HE is a low power version of the Opteron performs just as well as the "normal" one with the same speed.*

32-bit application will still be limited to a maximum of 4 GB of physical memory. However, the 4 GB limit is now imposed on a per-process level, not at a system-wide level.

- **Full 64-bit mode (64Bit/64Bit)**

This is the full 64 Bit mode which means that a 64-bit operating system and 64-bit applications are used. In this mode, an application can have a virtual address space of up to 40-bits (1 TB of addressable memory). Applications that run in full 64-bit mode have access to the full physical memory range, and also have access to the new and expanded GPRs.

INSTALLATION AND ADMINISTRATION

64 Bit Linux distributions have been available for quite a while now. In principle the new platform behaves like any PC, installation and booting the system just works as we are used to. We also successfully used the Rocks Cluster Distribution [5] to install the OS. As operating system we used Scientific Linux 3.04 which is available in 32 Bit and 64 Bit and comes with the 2.4.21 kernel.

With the new platform system management becomes a little bit more complex. When running with a 64 Bit OS in addition to the 64 Bit libraries also the 32 Bit Environment

* wenzel@fnal.gov

must be provided and maintained to allow applications to run in compatibility mode. Some of the tools (e.g. APT) are not multi platform aware and for other tools (YUM, RPM) the support for the new environment could be improved.

BENCHMARK APPLICATIONS

We wanted to run benchmarks related to HEP and CMS physics. We used the analysis tool ROOT [6], the MonteCarlo generator Pythia [7], the CMS detector simulation program OSCAR and the CMS event reconstruction programs ORCA and ORCA. The CMS applications OSCAR and ORCA could only be benchmarked in (32Bit/32Bit) and (64Bit/32Bit) mode since no 64 Bit version exist. The CMS software was installed using DAR [9]. In [10] one can find results Benchmarking results for a different set of HEP Astroparticle Physics applications. Root and PYTHIA were compiled as 32 Bit and 64 Bit application. For the compilation we used the g++ (ROOT) and g77 (Pythia) compilers based on gcc version 3.2.3 that is provided with Scientific Linux 3.04. The CMS applications and ROOT are dynamically linked C++ programs while PYTHIA is a statically linked FORTRAN program.

- **ROOT:** We used the 4.02/00 production version. We used the stress benchmark which is a suite of programs that tests the essential parts of Root.
- **Pythia:** For the benchmark we generate 100.000 supersymmetry events at $\sqrt{s} = 14$ TeV (the Pythia main65.f example). The code is compiled with g77 with the O2 option: `g77 -O2 -o main65 main65.f pythia6227.o`.
- **OSCAR.3.7.0:** The GEANT 4 based CMS detector simulation program. Here we simulate 300 single pion events of 50 GeV/c P_t .
- **ORCA.8.7.1:** The OSCAR output events are then digitized and reconstructed.

Table shows the matrix of the different applications and operation modes that we used to benchmark the performance.

Application (up to 4 copies)	Operating mode:		
	32-Bit leg. (32/32)	Compatib. (64/32)	64-Bit (64/64)
ROOT: stress	Yes	Yes	Yes
Pythia	Yes	Yes	Yes
OSCAR	Yes	Yes	N/A
ORCA Digi.	Yes	Yes	N/A
ORCA Reco.	Yes	Yes	N/A

Table 2: Matrix of various benchmarks run in the different operating modes.

BENCHMARKING RESULTS

In this section we present the various benchmark results. The Opteron 250 that we tested in (64/64) Bit mode runs ROOT stress/Pythia about 2.4/2.2 times faster than the 2.4 GHz Xeons in (32/32). The Opteron 275 which allows to run 2 applications simultaneous per CPU runs ROOT stress/Pythia about 4.4/4.1 times faster than the 2.4 GHz Xeons in legacy mode. The Pythia results are summarized in Table . For the dual core CPUs we see basically no drop in performance when running up to 4 copies of Pythia simultaneously (1 per core). We observe a 20 % boost in performance when running in (64/64) bit mode compared to (64Bit/32Bit) compatibility mode for both AMD64 and EMT64. There is a small drop in performance ($\approx 1\%$) for the Opterons when running in compatibility mode compared to legacy mode.

Table summarizes the Root Stress Benchmarking results. We observe that this application also scales very well when running up to 4 copies (1 per core) on the dual core machines. There is only a 1.6% performance drop when running 4 copies of stress simultaneously on the dual-cores. We observe that this application runs about 6% slower in compatibility mode compared to legacy mode. Also in this case the Opteron runs more than 40 % faster in 64Bit mode compared to (32/32) compared to the 64 Bit Xeons which only gain approximately 20% .

Table summarizes the OSCAR Simulation Benchmarking results. We observe that OSCAR scales very well resulting in a performance drop of only 0.9%/2.2%/1.6% on the 265/270/275 running up to 4 copies (1 per core) of OSCAR simultaneously. OSCAR is a very CPU and memory intensive application but it doesn't require a lot of io

Finally Table summarizes the CMS ORCA Benchmarking results. As the Applications become more IO intense we observe that the efficiency drop increases as the processes compete for disk IO bandwidth. For the ORCA digitization the drop is 0.6 %/3.3 %/2.8 % on the 265/270/275 running 4 copies simultaneously. For the ORCA reconstruction application we observe a significant drop in performance of 1.4 %/8.8 %/21.8 % for the 265/270/275 respectively.

CMS software runs in 64 Bit compatibility mode and 32 Bit native mode without any modifications to the code. We observe OSCAR and ROOT applications to run around 5% slower in compatibility mode than in legacy mode. The effect is even more pronounced for the digitization step where the loss in performance is around 15 %. When running in 64 Bit native mode we observe speed increases of 21% (Pythia) and 40% (ROOT) compared to 32 Bit native mode. The CMS applications OSCAR and ORCA will be replaced by a new framework which will be ported to the new 64 Bit platform.

POWER CONSUMPTION

Nowadays the amount of computing power that can be provided is very often limited by the amount of cooling and

Pythia Benchmarking results			
CPU	Operating mode:		
	legacy (32/32)Bit	Compatibility (64/32)Bit	64-bit (64/64)Bit
	(Evts/sec) per core for ×1/×2 processes	(Evts/sec) per core for ×1/×2/×3/×4 processes	(Evts/sec) per core for ×1/×2/×3/×4 processes
Opteron 244	91.0/91.0	90.3/90.4/-/-	110.1/110/-/-
Opteron 265	-/-	91.98/92.11/92.1/92.03	110.9/110.91/110.8/110.38
Opteron 246	101.3/101.4	100.9/101.1/-/-	121.2/121.5/-/-
Opteron 270	-/-	102.36/102.44/102.35/102.31	122.95/122.95/123.02/123.07
Athlon 64 3500+	-/-	111.06/-/-/-	129.8/-/-/-
Athlon 64 ×2 4200+	-/-	106.7/102.9/-/-	128.6/130/-/-
Opteron 248	113.0/112.9	112.3/112.3/-/-	136.4/136.1/-/-
Opteron 275	-/-	112.67/112.75/112.6/112.5	135.3/135.3/135.2/134.9
Opteron 250	121.3/121.3	120.8/120.7/-/-	146.9/146.8/-/-
XEON 3.4	88.5/88.7	88.1/88.3/-/-	108.6/108.6/-/-
XEON 3.6	93.5/94.21	93.4/93.4/-/-	115.6/115.4/-/-
XEON 2.4	65.5/65.5	nA	nA
XEON 2.8	75.9/75.9	nA	nA

Table 3: Summary of Pythia Benchmarking results. We see basically no drop in performance when running up to 4 copies of Pythia simultaneously. There is a 20 % boost in performance when running in (64/64) bit mode compared to (64Bit/32Bit) compatibility mode. There is a small drop in performance ($\approx 1\%$) for the Opterons when running in compatibility mode compared to legacy mode.

electrical power that the computing center can provide. In addition the expenses for cooling and electricity contribute significantly to the operating costs of a computing center. Therefore the power consumption is one of the main criteria when selecting worker nodes for computer farms.

Power Consumption		
CPU	Idle	Loaded
242	1.2A	1.4A
248	1.4A	1.6A
270	1.2A	1.7A

Table 7: Summary of power consumption based on idle load running SL 3.04 (64 Bit) and fully loaded CPUs using a calculation of π and running "dd".

As shown in Table , the dual core CPUs use about the same amount of power as the single core CPUs. This shows that we in fact get twice the performance using dual-core CPUs for the same cost in power and cooling. The low power Opteron 246HE was measured to use 7% less power when idle and 11% less under load compared to the regular 246. We found the single core Intel processors that we tested significantly less energy efficient than the AMD Opterons. This makes the dual core Opteron chips very attractive for running a cluster of high-performance machines.

CONCLUSION

It has been shown that for the considered applications the available 64-bit commodity computers from AMD and Intel are a viable alternative to comparable 32-Bit systems. Today's 32 Bit applications like the CMS software tested here can be executed on 64 Bit machines in either 32 bit mode or in 64 Bit compatibility without any modification of the code. This makes the transition to the new platform painless. A significant gain in performance might be achieved porting to the new processor architecture platform in 64 Bit mode. We saw a significant increase in performance (Pythia: 21 % ROOT: 40%) when applications were compiled for 64 Bit. Based on the investigations reported here we chose AMD Opterons for our latest farms purchase and run a 64 Bit OS. We found that Dual-core processors scale well as we run a copy of the application on each available core. The power efficiency makes the AMD dual-core processors a very attractive choice when running farms of high-performance machines.

REFERENCES

- [1] CMS IN 2005/012, Hans Wenzel, "Benchmarking AMD Opteron (AMD64) and Intel (EM64T) systems".
- [2] "What's So Great for Developers About the AMD64?" David Kaplowitz, June 23, 2003 <http://www.devx.com/amd/Article/16018>.
- [3] "64-bit Computing with Intel EM64T and AMD AMD64" <http://www.redbooks.ibm.com/abstracts/tips0475.html>.

ROOT Stress			
Operating mode:			
CPU	legacy (32/32)Bit	Compatibility (64/32)Bit	64-bit (64/64)Bit
	ROOTMks per core for ×1/×2 processes	ROOTMks per core for ×1/×2/×3/×4 processes	ROOTMks per core for ×1/×2/×3/×4 processes
Opteron 244	723.6/715.35	672.5/673.3/-/-	1059.9/1050.75/-/-
Opteron 265	-	719.3/718.7/719.6/715.65	1093.3/1085.35/1077.1/1075.9
Opteron 246	797.1/793.45	751.6/748.25/-/-	1176.1/1165.65/-/-
Opteron 270	-	827.3/819.3/819.77/819.33	1201.6/1201.75/1190.9/1182.8
Athlon 64 3500+	-	891.4/-/-/-	1295.9/-/-/-
Athlon 64 ×2 4200+	-	738.6/788/-/-	1120/1142/-/-
Opteron 248	894.2/888.4	834.7/835.75/-/-	1320.8/1305.95/-/-
Opteron 275	-	910.2/907.1/904.2/903	1327.96/1335.1/1311.1/1306
Opteron 250	957.6/956.7	898.3/891.35/-/-	1402.6/1395.55/-/-
XEON 3.4	964.6/950.05	909.3/903.8/-/-	1145.9/1135.8/-/-
XEON 3.6	1034.8/1020.6	969.7/958.7/-/-	1208.6/1194.4/-/-
XEON 2.4	590.7/583.75	nA	nA
XEON 2.8	731/720.05	nA	nA

Table 4: Summary of Root Stress Benchmarking results. We observe that this application scales very well. There is only a 1.6% performance drop when running one copy of the benchmark per core simultaneously. We observe that the application runs about 6% slower in compatibility mode compared to (32/32) mode. The Opterons run more than 40 % faster in 64Bit mode compared to (32/32) mode while the 64 Bit XEONS only gain about 20% .

- [4] CMS IN 2005/030, Hans Wenzel, Spencer Langley, Mako Furukawa, "Benchmarking dual-core AMD Opteron (AMD64) systems".
- [5] <http://www.rocksclusters.org>.
- [6] <http://root.cern.ch>.
- [7] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, Computer Phys. Commun. 135 (2001) 238 (LU TP 00-30, hep-ph/0010017).
- [8] <http://cmsdoc.cern.ch/cms00/cms00.html>.
- [9] DAR (**D**istribution **A**fter **R**elease) was developed by Natalia Ratnikova. DAR provides a complete runtime environment for various CMS applications but doesn't allow for recompilation. <http://home.fnal.gov/~natasha/DAR/dar.html>.
- [10] Proceedings of CHEP 2004 Interlaken, Switzerland, P. Wegner, S. Wiesand, "64-Bit Opteron systems in High Energy and Astroparticle Physics".

OSCAR Simulation		
CPU	Operating mode	
	legacy (32/32)Bit (Evts/sec) per core for ×1/×2 processes	Compatibility (64/32)Bit (Evts/sec) per core for ×1/×2/×3/×4 processes
Opteron 244	0.0915/0.0923	0.0866/0.0867/-/-
Opteron 265	-	0.0920/0.0918/0.0916/0.0912
Opteron 246	0.1021/0.1028	0.0956/0.0959/-/-
Opteron 270	-	0.1016/0.1012/0.0999/0.0994
Athlon 64 3500+	-	0.1076/-/-/-
Athlon 64 ×2 4200+	-	0.096/0.098/-/-
Opteron 248	0.1147/0.1139	0.1067/0.1074/-/-
Opteron 275	-	0.1089/0.1088/0.1081/0.1073
Opteron 250	0.1186/0.1221	0.1154/0.1156/-/-
XEON 3.4	0.1067/0.1060	0.0938/0.0929/-/-
XEON 3.6	0.1066/0.105	0.1012/0.1012/-/-
XEON 2.4	0.0603/0.0604	nA
XEON 2.8	0.0715/0.0702	nA

Table 5: Summary of OSCAR Benchmarking results. Like Root stress results we observe that the application runs about 5 % slower in compatibility mode compared to legacy mode. We observe that OSCAR scales very well resulting in a performance drop of only 0.9%/2.2%/1.6% on the 265/270/275 running up to 4 copies (1 per core) of OSCAR simultaneously. OSCAR is a very CPU and memory intensive application but it doesn't require a lot of io.

CPU	legacy (32/32)Bit	Compatibility (64/32)Bit	
	Digitization (Evts/sec) for ×1 processes	Digitization (Evts/sec) per core for ×1/×2/×3/×4 processes	Reconstruction (Evts/sec) per core for ×1/×2/×3/×4 processes
Opteron 244	0.1171	0.0989/0.098/-/-	0.809/0.806/-/-
Opteron 265	-	0.1096/0.108/0.1092/0.1089	0.9263/0.9268/0.9058/0.9131
Opteron 246	0.1293	0.1092/0.108/-/-	0.901/0.895/-/-
Opteron 270	-	0.1311/0.1290/0.1267/0.1268	1.055/1.057/0.974/0.962
Athlon 64 3500+	-	0.1421/-/-/-	1.118/-/-/-
Athlon 64 ×2 4200+	-	-	-
Opteron 248	0.1430	0.1210/0.1196/-/-	0.974/0.988/-/-
Opteron 275	-	0.142/0.140/0.140/0.138	1.106/1.089/0.946/0.864
Opteron 250	0.1562	0.1300/0.1296/-/-	1.039/1.035/-/-
XEON 3.4	0.1403	0.1251/0.1238/-/-	0.929/0.890/-/-
XEON 3.6	0.1427	0.1332/0.1310/-/-	1.026/1.019/-/-
XEON 2.4	0.0868	nA	nA
XEON 2.8	0.0996	nA	nA

Table 6: Summary of CMS ORCA Benchmarking results. As the Applications become more IO intense we observe that the efficiency drop increases as the processes compete for disk IO bandwidth. For the ORCA digitization the drop is 0.6 %/3.3 %/2.8 % on the 265/270/275 running 4 copies simultaneously. For the ORCA reconstruction application we observe a significant drop in performance of 1.4 %/8.8 %/21.8 % for the 265/270/275 respectively.