

Benchmarking CMSSW on Intel and AMD single-core, dual-core and quad-core systems

Haiying Xu, Fengping Hu, David Braun, Preston Smith, Tom Hacker
Rosen Center for Advanced Computing,
Purdue University

We have benchmarked dual-processor quad-core AMD Opteron 2350 and 2356, dual-processor quad-core Intel Xeon E5345, single processor quad-core Intel Xeon X5472, dual-processor dual-core AMD Opteron 2214, dual-processor single-core Intel Xeon EM64T and single-processor single-core Intel Xeon EM64T systems using a CMSSW event simulation and reconstruction application. The results are presented in this note. We find the performance of Intel's and AMD's dual quad-core processors all impressive; however AMD quad-core Opteron 2350 processors were found to have a better performance to price ratio than the Intel Xeon processors.

1. Introduction

As a step in the acquisition of a new cluster computer system, we performed a series of benchmarking experiments using the CMSSW software on dual-processor, quad-core systems. The aim of our study was to assess the performance of Xeon and Opteron dual quad-core based systems to determine the effect of the different processor architectures on the performance of typical CMS event simulation and reconstruction jobs. We were especially interested in two things:

- i) How does the performance of the benchmark application running on these processors scale when executing up to 8 CMSSW jobs simultaneously?
- ii) How does the performance of an AMD Opteron dual quad-core system compare to the one of an Intel Xeon dual quad-core system for CMSSW jobs?

2. The CMS Benchmark Application

We wanted to run benchmarks that most closely reflected the computational workload of applications used by CMS physicists. We used the benchmark package [1] provided by Mario Kadastik. The CMSSW benchmark job developed by Kadastik is based on the 'FullChainExample' configuration and includes a Monte Carlo event generator step using PYTHIA (producing charged Higgs bosons), a full detector simulation step and an event reconstruction step. The benchmark job is based on CMSSW version 1_7_5. Kadastik performed benchmark experiments using this package on systems based on a dual-processor dual-core AMD Opteron 275, dual-processor quad-core AMD Barcelona containing 8 GB of memory, and a dual-processor quad-core AMD Barcelona containing 16 GB memory. Kadastik performed these benchmarking studies in November 2007, and found that the application scaled well up to seven cores. From running Kadastik's benchmarks, we found that the initiation time of the application before processing of the first event is approximately 328 seconds when running on a dual-processor quad-core Xeon based system. The primary memory and shared memory used for each benchmark experiment totaled 1.06 GB for each job. The benchmark used exactly the same configuration with the same seeds for the random number generator to ensure consistent conditions among all calculations. We ran the benchmark three times per system and calculated the mean values and standard deviations of the run time. The obtained results are shown in Figure 1.

We found that the installation of the CMSSW software for each new platform we assessed proved to be much simpler than past installations. After the package management tool APT is installed, the complete CMSSW package could be installed on a remote AMD Opteron 2350 based system within 90 minutes from a remote APT repository. All other systems used for our benchmarking tests were locally accessible and could use the local CMSSW installation at the Purdue Tier-2 center.

To ensure consistent results among the systems, we used Red Hat Enterprise Linux 4.6 in 64 Bit mode on all platforms.

3. Testing Procedures

On the dual single-core Xeon cluster, we used the Portable Batch System (PBS) job scheduler to submit and manage benchmarking jobs. The process we followed for benchmarking was as follows:

- 1) Request exclusive access to one node, which contains two processors:
`qsub -I -q cms_analysis -l nodes=1:ppn=2,walltime=24:00:00`
- 2) On the system reserved by the PBS scheduler, set up the CMSSW environment:
`source /apps/02/cmssoft/cms/cmsset_default.sh`
- 3) Create a CMSSW_1_7_5 project area in /tmp:
`cd /tmp`
`scramv1 project CMSSW CMSSW_1_7_5`
- 4) Set up the cmsRun environment:
`cd CMSSW_1_7_5/src`
`eval `scramv1 runtime -sh``
- 5) Copy the benchmarking package to the directory src, unpack it, and move a cfi file to a subdirectory test/:
`cp testpack_v2.tar .`
`tar -xvf testpack_v2.tar`
`mkdir test`
`mv dblh_normal.cfi test/`
- 6) Modify a parameter within the runTest.sh to spawn 2, 4 or 8 computational jobs for each system.
within runTest.sh:
`set NT=2 for dual single-core Xeon`
- 7) Start the benchmark:
`./runTest.sh`

On all other machines we benchmarked, we followed the same procedure, with the exception of omitting PBS scheduling when we had exclusive access to the system.

We measured the total execution time per job using Unix *time* utility. The obtained results of the benchmark jobs were stored on a local file system.

4. Assessed Platforms

For our analysis, we assessed the following platforms:

<i>System Name</i>	<i>Processor</i>	<i>Processor Speed</i>	<i>Architecture</i>	<i>System Memory</i>
Lear	Xeon EM64T	3.2 GHz	Intel-64 Single-Processor, single core	4 GB
Lear	Xeon EM64T	3.2 GHz	Intel-64 Dual-processor, single-core	4 GB
Intel	Xeon X5472	3.0 GHz	Intel-64 Single-processor, quad-core	16 GB
CMS	Opteron 2214	2.2 GHz	AMD-64 Dual-processor, dual-core	4 GB
AMD	Opteron 2350	2.0 GHz	AMD-64 Dual-processor, quad-core	16 GB
AMD	Opteron 2356	2.3 GHz	AMD-64 Dual-processor, quad-core	8 GB
Intel	Xeon E5345	2.33 GHz	Intel-64 Dual-processor, quad-core	8 GB

Table 1: Systems tested

5. Benchmarking Results

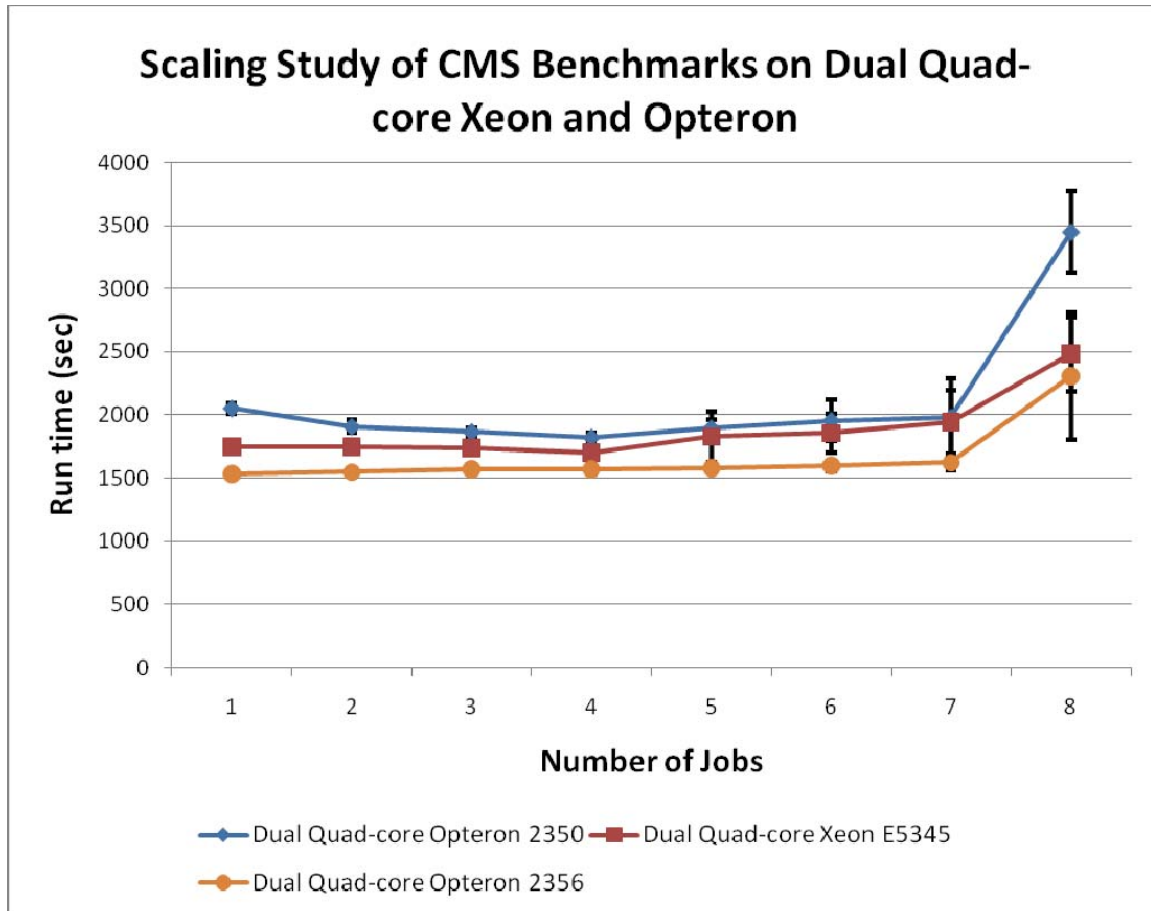


Figure 1. Benchmark results of Dual Quad-core Xeon and Opteron only

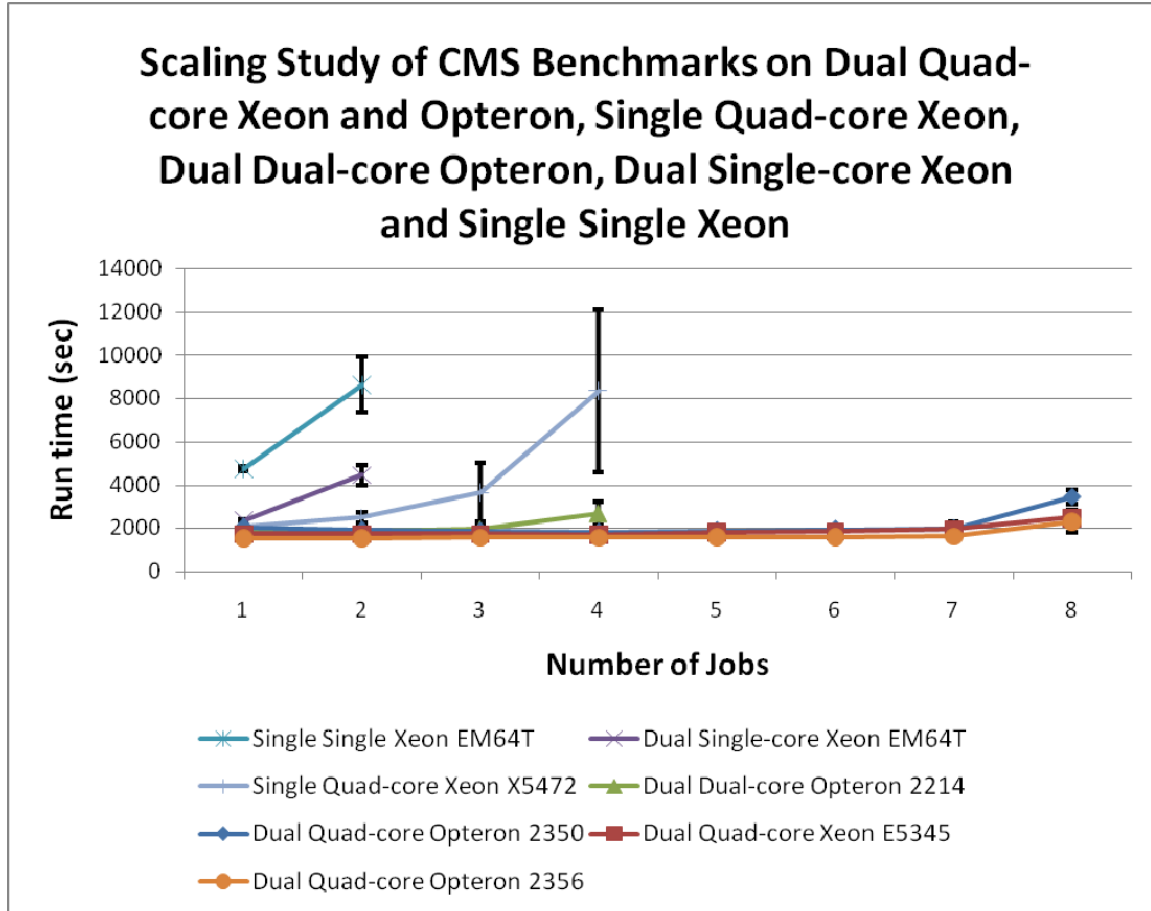


Figure 2: Benchmark results of all seven systems

As shown in Figure 1, the benchmark performance scaled very well when running up to 7 CMSSW jobs simultaneously. The execution time per job increased when eight jobs were running simultaneously, which we believe is caused by the interference of the operating system with the application while competing for CPU resources. We also checked the effect of having dCache running on a worker node set up for a typical resilient dCache installation. When dCache is running, we found that the benchmark cannot scale very well when 2 or 3 jobs are running simultaneously. As shown in Figure 2, on the dual-processor dual-core AMD Opteron, contention for the final core when running four jobs with the operating system cause a sharp increase in execution time, which is also the cause of the increased running time for two jobs on the dual-processor single-core Xeon. Overall, we found that dedicating one core to the operating system can result in an overall higher performance (shorter execution time per running job). We measured the overall memory usage during the execution of a single job and found the memory profile of the benchmark process was: virtual memory 1.2 GB, resident 0.9 GB and shared memory 0.1 GB for a single job. We also tested binding one job with one processor and the second job with the other processor, and obtained an identical memory profile and execution time for the dual-processor single-core Xeon. The benchmark performance cannot scale very well on the single-processor quad-core Xeon, however, we measured the overall memory usage and obtained an identical memory profile as others. We will find out the explanation in the further research and study.

6. Conclusions

We tested various systems based on AMD quad-core Opteron 2350, 2356 and Intel quad-core Xeon E5345 processors using CMSSW_1_7_5 specific applications as benchmarks. We found that the performance of Xeon dual quad-core processors measured as the total execution time of the CMS benchmark job was statistically equivalent to the Opteron 2350 and 2356 based system when running up to seven jobs per node. When the CMSSW application contends with the operating system for the last core, the Intel processor based system and Opteron 2356 based system delivered better performance than the Opteron 2350 based system. However, when comparing costs, Opteron 2350 processor is ~\$389, Xeon E5345 processor is ~\$485, and Opteron 2356 processor is not available yet. Thus the AMD Opteron dual quad-core processors 2350 proved to have a better performance to price ratio than the Xeon based system.

[1] *AMD Barcelona testing*, Mario Kadastik, Nov 23, 2007
<http://hypernews.cern.ch/HyperNews/CMS/get/t2/57.html>