sgi®

Performance Report

Comparing the Performance of the SGI® Altix® 330 with Similar Dual Core Opteron-based Systems

Roberto Gomperts, principal scientist, SGI
Martin Hilgeman, chemical applications engineer, SGI

Table of Contents

1.0 Introduction

SGI Applications Engineering performed benchmark tests to compare the performance of the SGI® Altix® 330 using Intel® Itanium® 2 processors against comparable AMD® dual core Opteron™ based systems. The system performance evaluations were performed by running standard benchmark test cases from the most widely used computational chemistry applications available and comparing the results.

Standard benchmark test cases from the following chemistry applications were used:

- Gaussian 03 rev C.02 (Gaussian)
- DMol³ from Accelrys® Materials Studio® (Accelrys)
- CASTEP from Accelrys Materials Studio (Accelrys)
- VASP (University of Vienna),
- NAMD (University of Illinois at Urbana Champagne),
- Sander (Amber), and
- PMEMD (Amber)

Results of the tests revealed that the SGI Altix 330 using shared memory and unmatched I/O throughput achieved superior performance and ran faster than twice the number of cores in comparable AMD dual core Opteron-based systems.

In addition, a Parallel Throughput System test revealed that the SGI Altix 330 built on SGI NUMAflex™ architecture experienced only 1% degradation in performance when running Gaussian test397 on a fully loaded system (24.4 minutes on a fully loaded system compared to 24.1 minutes in a stand alone test) as compared to a 33% degradation in performance for the comparable AMD dual core Opteron-based system (71 minutes on a fully loaded system compared to 53 minutes in a stand alone test).

On a side note: there were some execution problems with running Gaussian, CASTEP and VASP on the Opteron-based systems. For example,

- Gaussian would not run with more than 4 cores on the Opteron-based system,
- CASTEP test case TiN_Cas would not run on a 1 core Opteron-based system, and
- VASP test case FeLiq288 would not run on a 16 core Opteron-based system.

Corresponding successful SGI Altix 330 runs were ignored for these data points when averaging test results.

2.0 Summary of Results

For the purposes of this paper, the following terms are used
Core = a single Central Processing Unit (CPU)
Socket = the physical connection on the mother board

To illustrate these definitions, the SGI Altix 330 uses Intel Itanium 2 processors (cores) and has multiple sockets (up to 16 per system) with a single core per socket. The AMD dual core Opteron-based system has multiple sockets (up to 8 per system) with 2 cores per socket.

The results of the benchmark testing are broken into 3 areas:

2.1 Core to Core comparison–Traditional way to evaluate performance, this is the comparison of running the chemistry test cases on the same number of Intel Itanium 2 cores as AMD Opteron cores.

2.2 Socket to Socket comparison–This is the comparison of running the chemistry test cases on twice the number of AMD Opteron cores as Intel Itanium 2 cores.

2.3 Parallel Throughput System Test results–This is The most realistic test performed compared to a production environment in that the test case was run in parallel with other jobs on a fully loaded system.

2.1 Core to Core comparison

This is the traditional way to evaluate performance. For this case, the results of running the benchmark test cases on the same number of Intel Itanium 2 cores were compared to the same number of AMD Opteron cores.

The results were spectacular for the Altix 330. On average, the chemistry codes tested ran twice as fast on the SGI Altix 330 compared to the Opteron-based systems.

The following graph illustrates these results when comparing the same number of Intel Itanium 2 cores in the SGI Altix 330 versus the same number of cores in the Opteron-based system. The average performance graphed is the combination of all test cases for each application run over all number of cores (1 core, 2 core, 4 core, 8 core and 16 core cases).

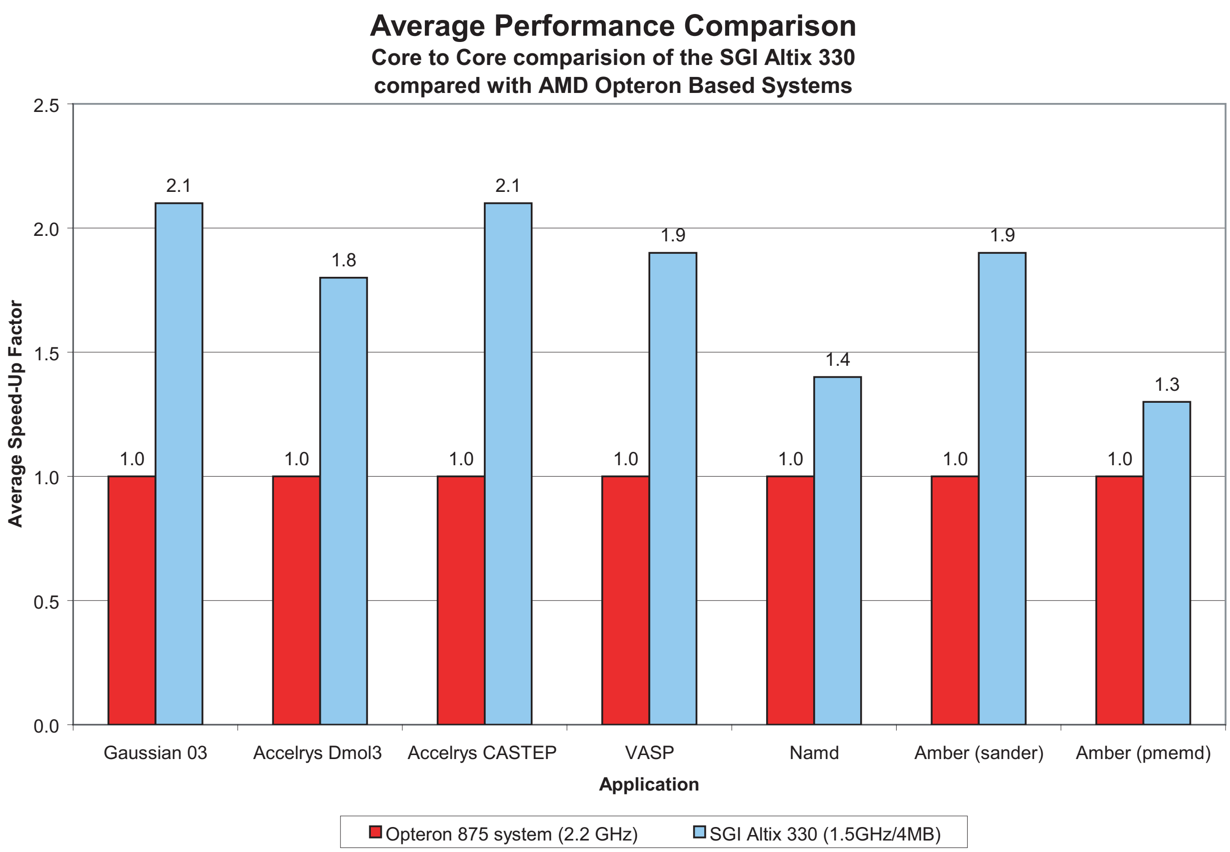

Source: SGI internal benchmark testing, December, 2005

2.2 Socket to Socket comparison

This is not the traditional way to evaluate performance but is a very interesting way to compare systems. To calculate this comparison, the results of running the benchmark test cases on twice the number of AMD Opteron cores are compared to the Intel Itanium 2 cores.

Again, the results were stunning for the Altix 330. The computational chemistry codes tested ran up to twice as fast on half the number of cores in the SGI Altix 330 system compared to the AMD Opteron-based system.

The following graph illustrates the results obtained when comparing twice the number of cores in the AMD Opteron-based system versus the Intel Itanium 2 cores in the SGI Altix 330.

2.3 Parallel Throughput System Test

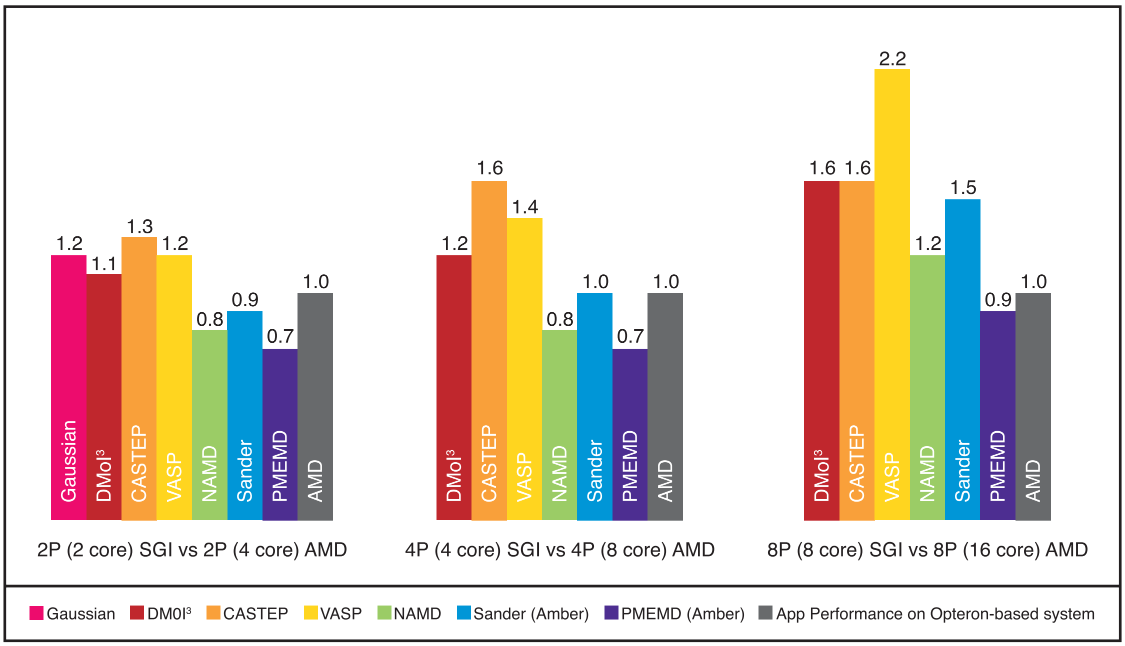

The Parallel Throughput System Test evaluated the performance of running a benchmark test case on a fully loaded system as opposed to running the same test case stand alone on the system.

For this test, all 7 Gaussian test cases were used. The throughput set (or the background jobs) consisted of 2 copies of apinefreq and 4 copies of the other 5 tests, for a total of 22 jobs. These jobs were constantly running simultaneously on the computer. In parallel, Gaussian test397 was started and measured.

The test was designed to effectively have Gaussian test397 run while a heterogeneous load is being applied to the computer system.

The results were very interesting.

From the figure above, it can be seen that the full test set finished on the Opteron-based system in 79 minutes. Note: test397 finished

Source: SGI internal benchmark testing, December, 2005

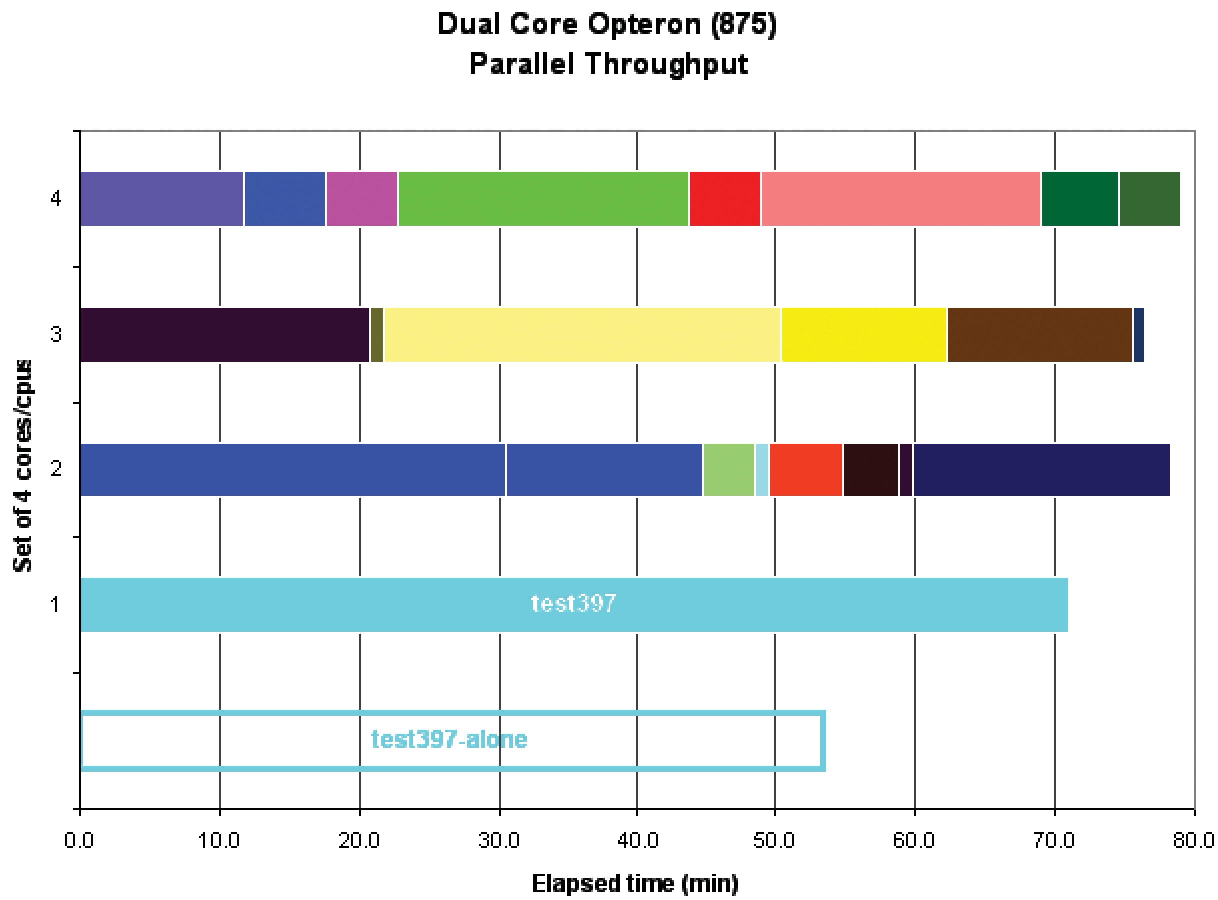

in 71 minutes.

This time is compared to the measured time of 53 minutes for test397 to run on another wise empty Opteron-based system.

The result was a system degradation of ~33% to run the test case on the fully loaded system.

Contrast this to the results obtained on the SGI Altix 330 1.6GHz/6MB

From the figure above, it can be seen that the full test set finished on the SGI Altix 330 in 25 minutes. Note: this is over 3 times as fast

Source: SGI internal benchmark testing, December, 2005

Altix 330 1.6GHz,6Mb
Parallel Throughput

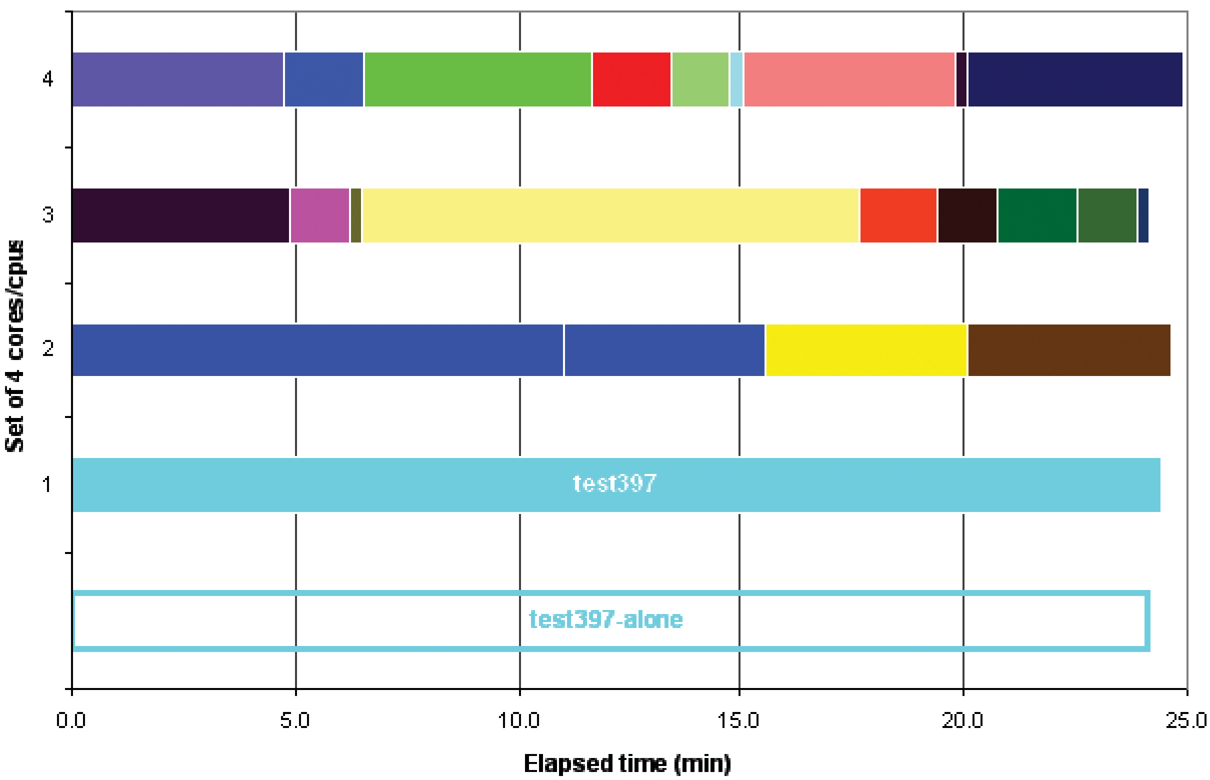

as the Opteron-based system which finished in 79 minutes. The time for test397 to finish was 24.4 minutes on the SGI Altix 330.

This time is compared to the measured time of 24.1 minutes for test397 to run on another wise empty SGI Altix 330 system.

The result was a system degradation of ~1% to run the test case on the fully loaded system which is much better than the ~33% degradation when running the same job on the fully loaded Opteron-based system.

Source: SGI internal benchmark testing, December, 2005

3.0 Detailed Results

Description of Computer Systems

The following computer systems were used in this testing:

- **Opteron-based Systems**

  Dual Core AMD Opteron Processor 875

  16 Cores, 64GB of physical memory CPU MHz: 2.2GHz; off-chip cache size: 1MB I/O subsystem: /dev/sdb1 193Gb (SATA) /scratch SUSE LINUX Enterprise Server 9 (x86_64), Patchlevel 2 gcc version 3.3.3

- **Altix 330 systems**

  Altix 330 1.5GHz, 4MB Cache

  Intel Itanium 2 Processors (Madison) 16 CPUs, 64GB of physical memory CPU MHz: 1.5GHz, off-chip cache size: 4MB I/O subsystem: Linux 2.4.21-sgi306rp13 #1 SMP Wed Aug 3 14:14:32 PDT 2005 Red Hat Enterprise Linux AS release 3 (Taroon Update 4) SGI ProPack 3SP6 for Linux, Build 306rp27-0508031407

  Altix 330 1.6GHz, 6MB Cache

  Intel Itanium 2 Processors (Madison) 16 CPUs, 64 GB of physical memory CPU MHz: 1.6GHz, off-chip cache size: 6MB I/O subsystem: Linux 2.4.21-sgi306rp13 #1 SMP Wed Aug 3 14:14:32 PDT 2005 Red Hat Enterprise Linux AS release 3 (Taroon Update 4) SGI ProPack 3SP6 for Linux, Build 306rp27-0508031407

Summary

The following table summarizes for each application the relative performance of each machine over the Dual Core Opteron server. The speedups are calculated using the sum of all the elapsed times of all test cases on all CPUs or cores utilized.

Speedup of Altix 330 over Opteron DC 875, 2.2 GHz		
Application	**Altix 330 1.5GHz, 4M**	**Altix 330 1.6GHz, 6M**
Gaussian	2.1	2.2
DMol3	1.8	
Castep	2.1	2.3
VASP	1.9	2
Namd	1.4	
Sander	1.9	
Pmemd	1.3	

This methodology biases the results towards large calculations and small CPUs or cores counts.

Source: SGI internal benchmark testing, December, 2005

3.1 Test Descriptions

- Gaussian. Used official version of Gaussian 03 rev. C.02 for AMD64 and for Intel Itanium 2. All parallel tests used OpenMP.

The official AMD64 Gaussian binaries we received were built with the WorkStation version of the PGI compilers. They restrict the number of OpenMP threads to a maximum of 4. The Server version of these compilers bring the maximum to 16 OpenMP threads

	Gaussian 03 rev C.02
Test	**Description**
apinenefreq	$C_{10}H_{16}$, 182 basis functions, RB3LYP/6-31G, Frequency calculation.
apinenehf	$C_{10}H_{16}$, HF/6-311g(df,p), 346 basis functions
$C_3H_{11}O_4$ MP_2	Full Direct MP2 of $C_3H_{11}O_4^+$, C1 symmetry, 6-31G** basis set, 160 basis functions
$C_4H_{13}NO_4$ Freq	RHF Frequency NoRaman calculation of $C_4H_{13}NO_4$. C1 symmetry, 6-31G** (5D) basis set, 191 basis functions.
CH_6N_2 MP4	MP4 (SDTQ) test for Methyl Hydrazine: CH_6N_2, C1 symmetry. 6-311G(3d,3p) basis set, 156 basis functions
Taxol	HF single point calculation of taxol: $C_{47}H_{51}NO_{14}$, C1 symmetry, 3-21G basis set, 660 basis functions
Test397	Test from Gaussian's Inc. QA-suite. RB3LYP/3-21G Force calculation of $C_{54}H_{90}N_6O_{18}$, C1 symmetry, 882 basis functions

- CASTEP: Used official (licensed) version from Accelrys Materials Studio 3.2.
For the Opteron-based system, we used parallel execution with HP-MPI. On the Altix 330, we used the MPT linked version of the program

	CASTEP from Accelrys Materials Studio 3.2
Test	**Description**
TiN_old	Single Point Energy Calculation H atom on the metallic surface of TiN(100), $Ti_{16}N_{16}H$ (33 atoms) PW91 functional 8 k-points, 4 extra bands 20 SCF cycles
TiN_Cas	Single Point Energy Calculation H atom on the metallic surface of TiN(100), $Ti_{16}N_{16}H$ (33 atoms) PW91 functional 36 k-points, 25 extra bands 20 SCF cycles
indigo_1	Optical Properties Calculation $C_{32}O_4N_4H_{20}$ (60 atoms) PBE functional 10 k-points, 25 extra bands, 126 optics extra bands
Layer	Single Point Energy Calculation $Ti_6Ag_8O_{12}$ (26 atoms) PBE functional 13 k-points, 25 extra bands 20 SCF cycles

Source: SGI internal benchmark testing, December, 2005

- DMol³: Used official (licensed) version from Accelrys Materials Studio 3.2.
  For the Opteron-based system, we used parallel execution with HP-MPI.
  On the Altix 330, we used the MPT linked version of the program

	DMol³ from Accelrys Materials Studio 3.2
Test	**Description**
catalyst	One geometry optimization: Zirconium catalyst $ZrSiC_{43}H_{37}$ (82 atoms) Basis: dnp Functional: bp Integration grid: Medium Symmetry group of molecule: C1
TiN	One geometry optimization: Hydrogen adsorption and diffusion on TiN(100) surface $Ti_{10}N_{10}H$ (21 atoms) Basis: dnp Functional: bpe Integration grid: Medium Crystal Symmetry: tetragonal of order 4
zeolite AFY	One geometry optimization: Zeolite catalyst $Co_2Al_6P_8O_{32}$ (48 atoms) Basis: dnp Functional: bpe Integration grid: Medium Symmetry P1 for Periodic systems

Source: SGI internal benchmark testing, December, 2005

• Amber: sander and pmemd: Used official amber8 version from Scripps with July 5, 2005 patches. The program for Opteron was built with Intel compilers 9.0: Intel® Fortran Compiler for Intel® EM64T-based applications, Version 9.0 Build 20051020 Intel® C Compiler for Intel® EM64T-based applications, Version 9.0 Build 20051020 Parallel execution with MPICH: mpich-1.2.7-ch_shmem, built with ICC On the Altix 330 was built with 8.1 compilers for Itanium2, with MKL 7.2 and MPT

	Amber8
Test Case	**Description**
jac	"Joint Amber/Charmm" benchmark. DHFR protein solvated with TIP3 water in a periodic box. Number of atoms: 23,558. PME with a direct space cutoff of 9 Ang.
factor_ix	Protein solvated with TIP3 water in a periodic box Number of Atoms: 90,906 PME with direct space cutoff of 8 Ang.
hb	Hemoglobin benchmark. Protein solvated with TIP3 water in a periodic box Number of Atoms: 44,247 PME with direct space cutoff of 8 Ang.
rt	Reverse transcriptase complex Number of atoms: 14,1154 PME with direct space cutoff of 8 Ang.
trx	Thioredoxin , ff00 with extra points Number of atoms: 14,093 PME with direct space cutoff of 8 Ang. and polarizable potential
gb_alp	Alpha-lytic protease from 1qq4.pdb Number of atoms: 2,752 Generalized Born simulation with cutoff of 15 Ang., salt concentration 0.1 M, nrespa=4
gb_cox2	Thermalization Number of atoms: 18,056 Generalized Born simulation with cutoff of 12 Ang., salt concentration 0.2 M, nrespa=4
gb_mb	Myoglobin simulation. Number of atoms: 2,492 Generalized Born simulation with cutoff of 20 Ang., salt concentration 0.2 M, nrespa=4

Source: SGI internal benchmark testing, December, 2005

• NAMD 2.6b1: Used distributed versions of NAMD from the official NAMD site:
http://www.ks.uiuc.edu/Research/namd/

The Opteron version uses charmrun to run in parallel, the Altix version uses MPI (MPT)

	NAMD 2.6b1
Test Case	**Description**
jac	"Joint Amber/Charmm" benchmark. DHFR protein solvated with TIP3 water in a periodic box. Number of atoms: 23,558. PME with a direct space cutoff of 9 Ang.
apoa1	Dynamics 1 fs timestep, 500 steps Periodic boundary conditions. Number of Atoms: 92,224 12A cutoff + PME every 4 steps

• VASP 4.86. Built with the 9.0 compilers, GOTO Blas libraries, LAPACK from Intel and MPICH-1.2.7-ch_shmem

	VASP 4.6.26
Test Case	**Description**
Hg.bench	50 Hg atoms empty core
Pt/NAFIO	MD simulation: 3 ionic steps, 10 electronic steps 200 bands
FeLiq67	Liquid Fe, 67 atoms, 50 electronic steps, 1 k-point, 350 bands, gamma point only
FeLiq288	Liquid Fe, 288 atoms, 50 electronic steps, 1 k-point, 350 bands, gamma point only
Ta256	256 Ta atoms MD simulation, 770 bands, 1 k-point

Source: SGI internal benchmark testing, December, 2005

3.2. Detailed Test Results

3.2.1 Gaussian

The following figure shows that the Altix 330 is on average 2.1 faster than the Opteron-based system for the 7 cases measured. This refers to the single CPU (or core) runs only.

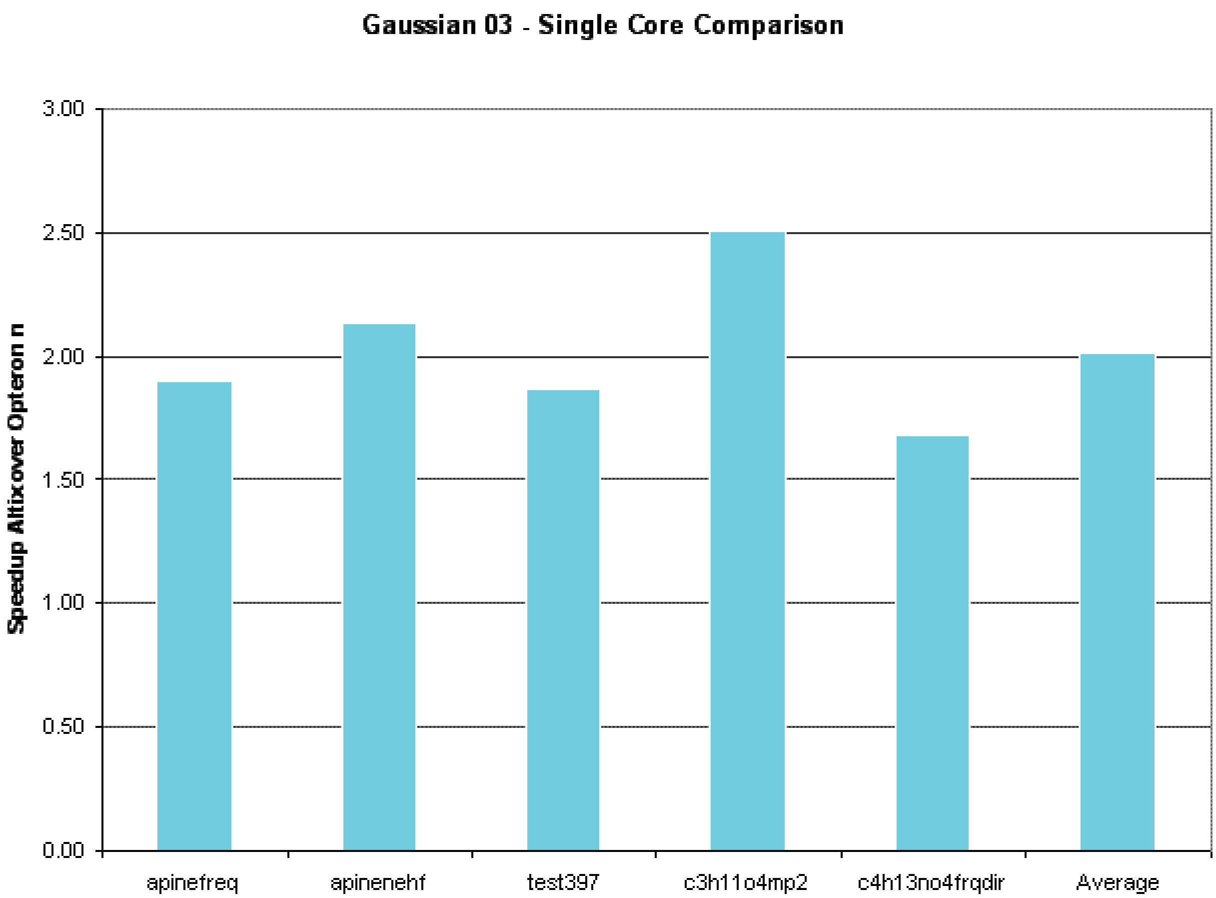

In the next graphs the parallel performance of the Altix 330 and the Opteron based system are showed side-by-side. The figures compare Gaussian Opteron performance against an Altix 330 1.5 GHz/4MB.

Note: currently it is not possible to run Gaussian's released binaries on more than 4 CPUs (or cores) using OpenMP on Opteron-based systems. The reason is that the rev C.02 binaries distributed by Gaussian, Inc. have been built with the Work Station version of the PGI compilers. The Server version of these compilers moves the limitation from 4 to 16 processes. Preliminary experiments with Gaussian 03 rev D.01 have shown that it is still not possible to use more than 11 CPUs (or cores).

Source: SGI internal benchmark testing, December, 2005

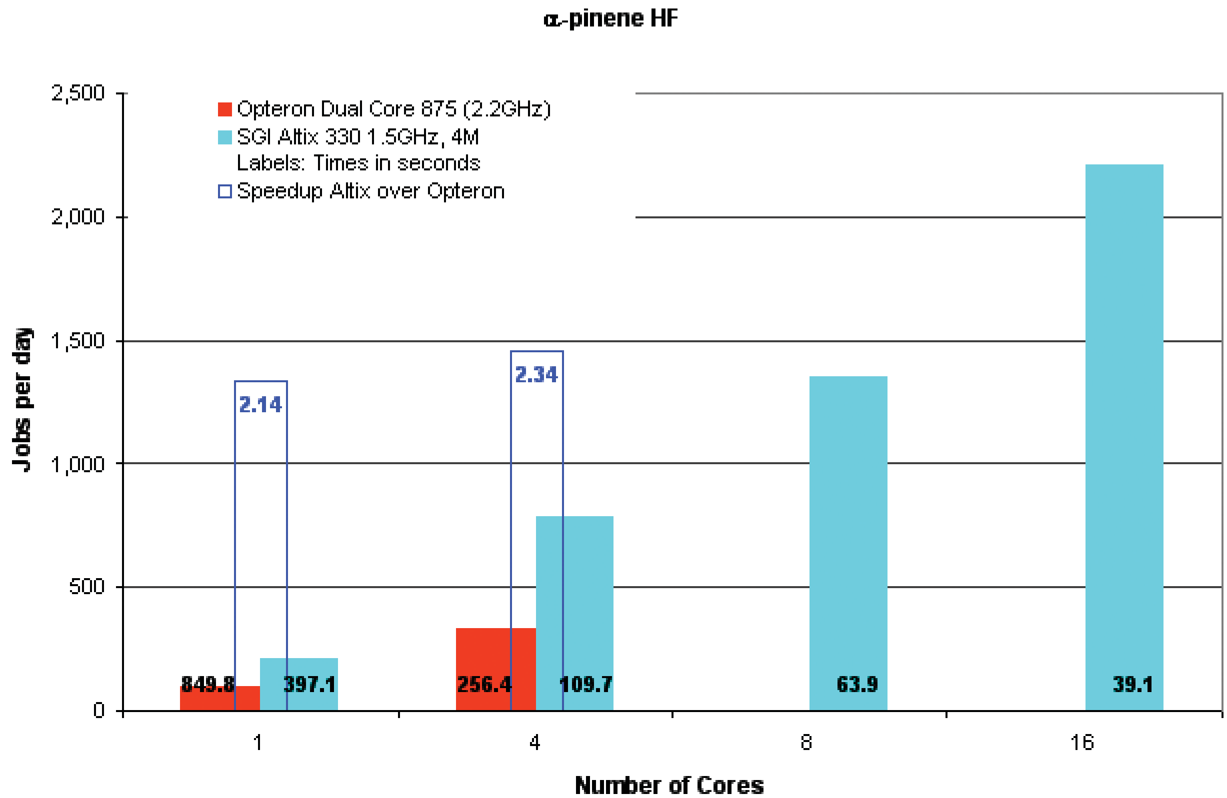

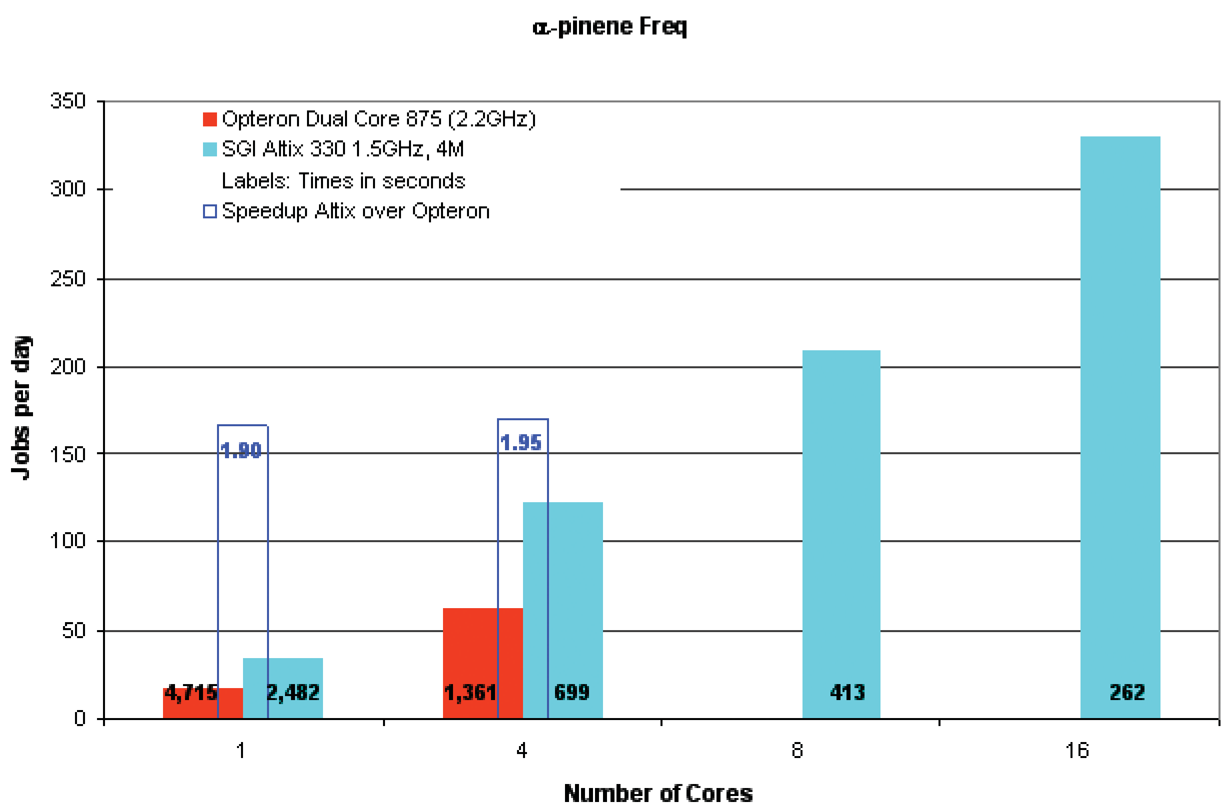

Source: SGI internal benchmark testing, December, 2005

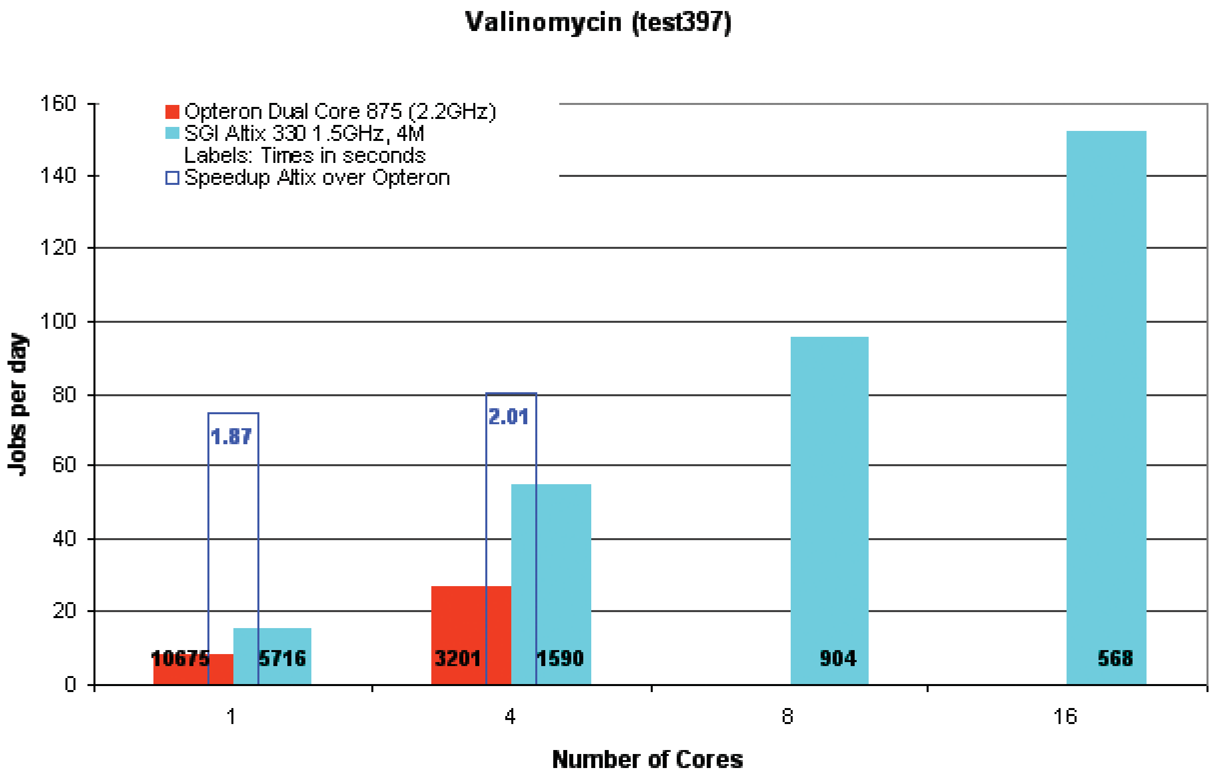

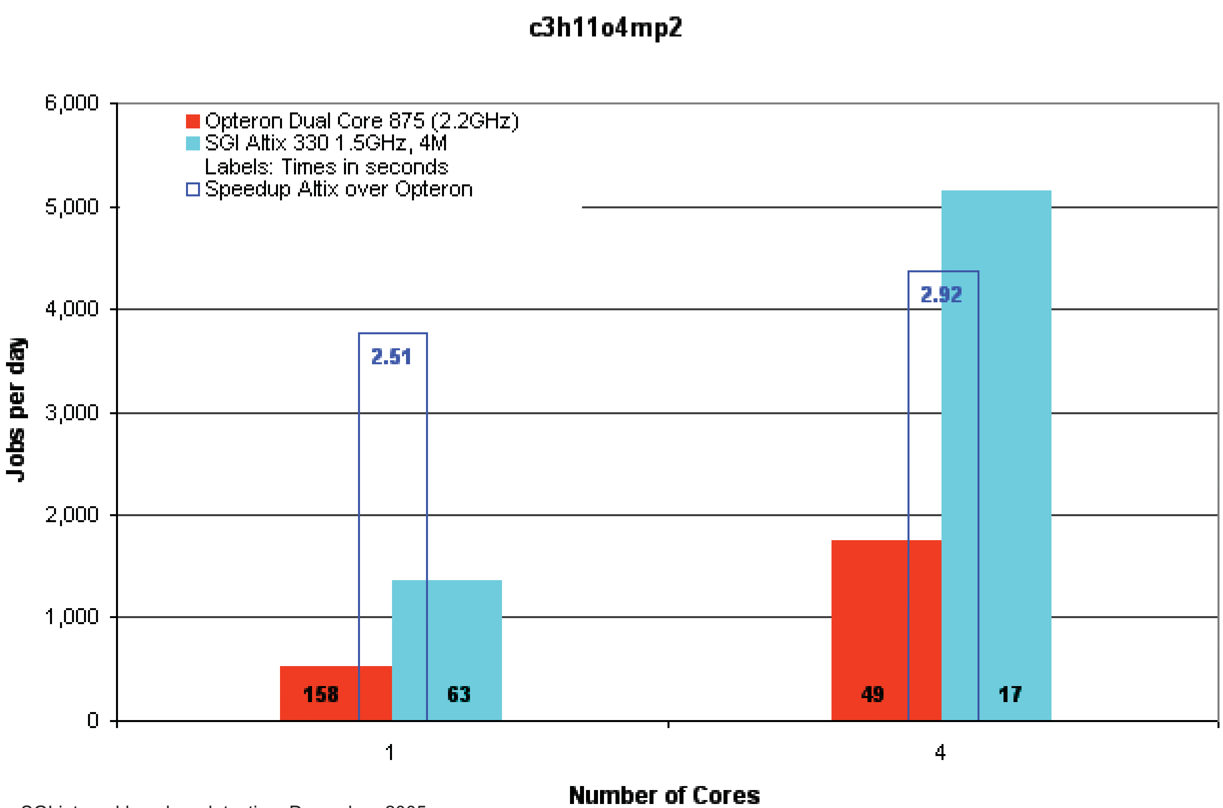

Source: SGI internal benchmark testing, December, 2005

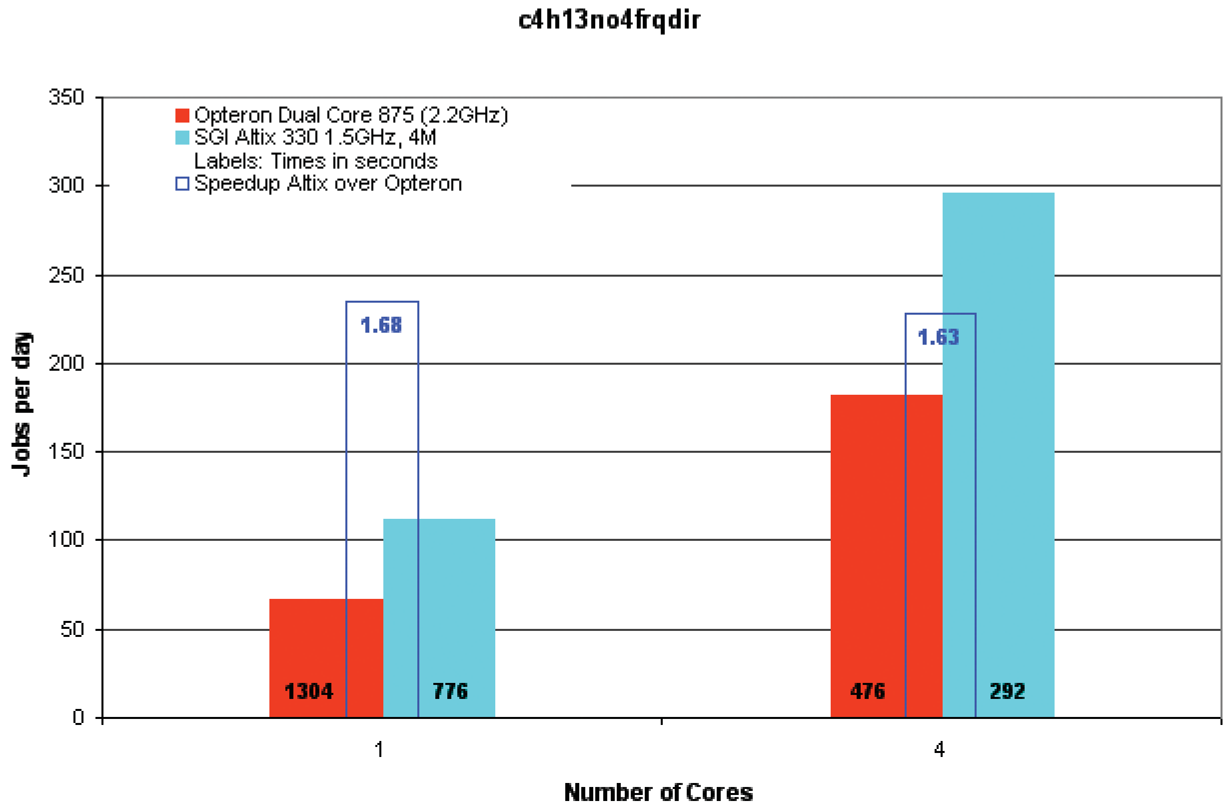

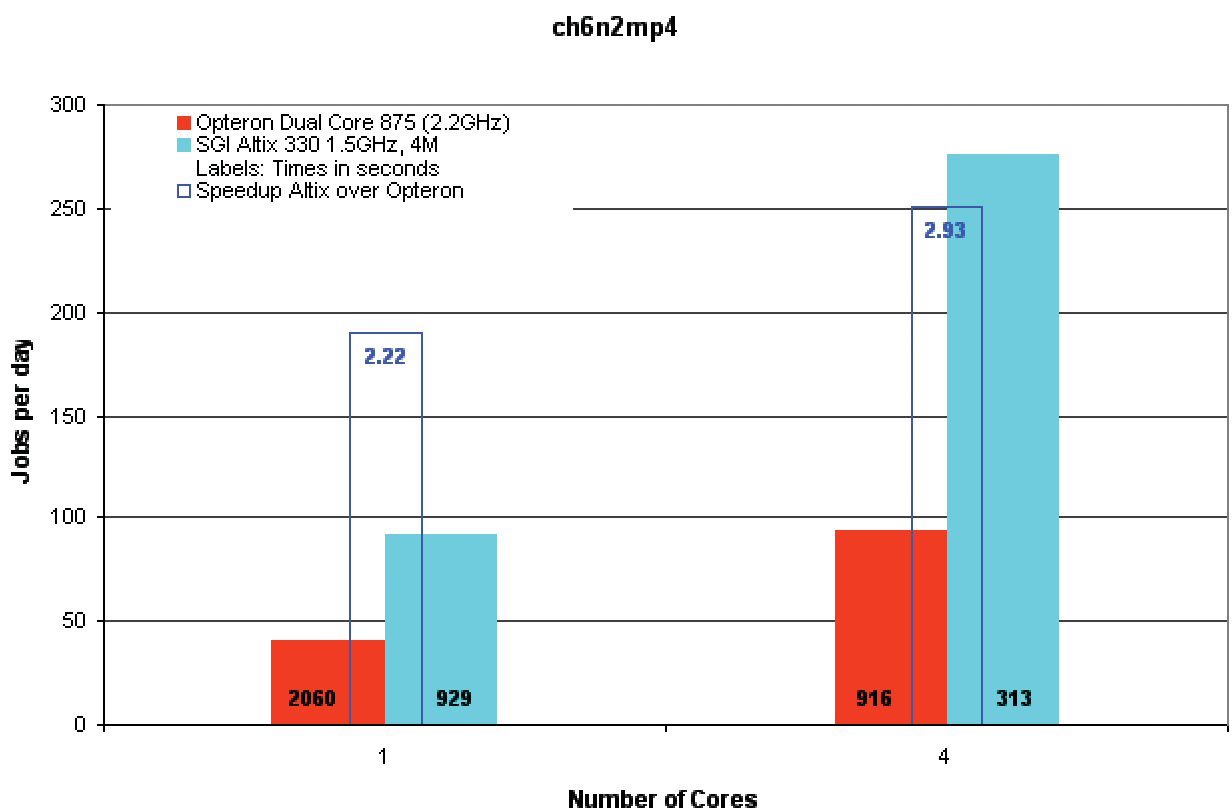

Source: SGI internal benchmark testing, December, 2005

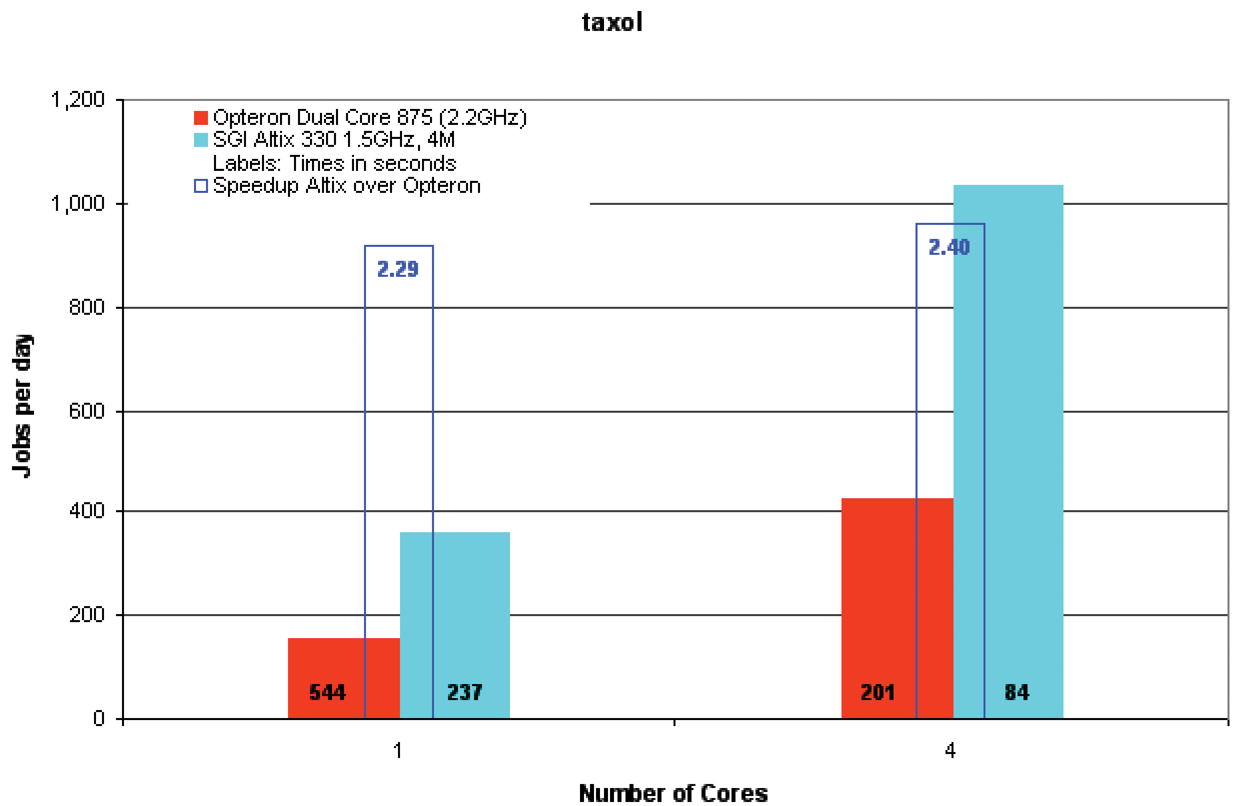

3.2.2 Parallel Throughput Test

The seven Gaussian tests were used to design a throughput test. The throughput set consisted of one copy of the test job (test397), 2 copies of apinefreq and 4 copies of the other 5 tests, for a total of 23 jobs. A parallel make ensured that 4 jobs were constantly running simultaneously on the computer, each job used 4 CPUs (or cores). The only constraints to the order of submission is to have test397 and one of the apinefreq jobs be part of the first batch of jobs submitted.

Source: SGI internal benchmark testing, December, 2005

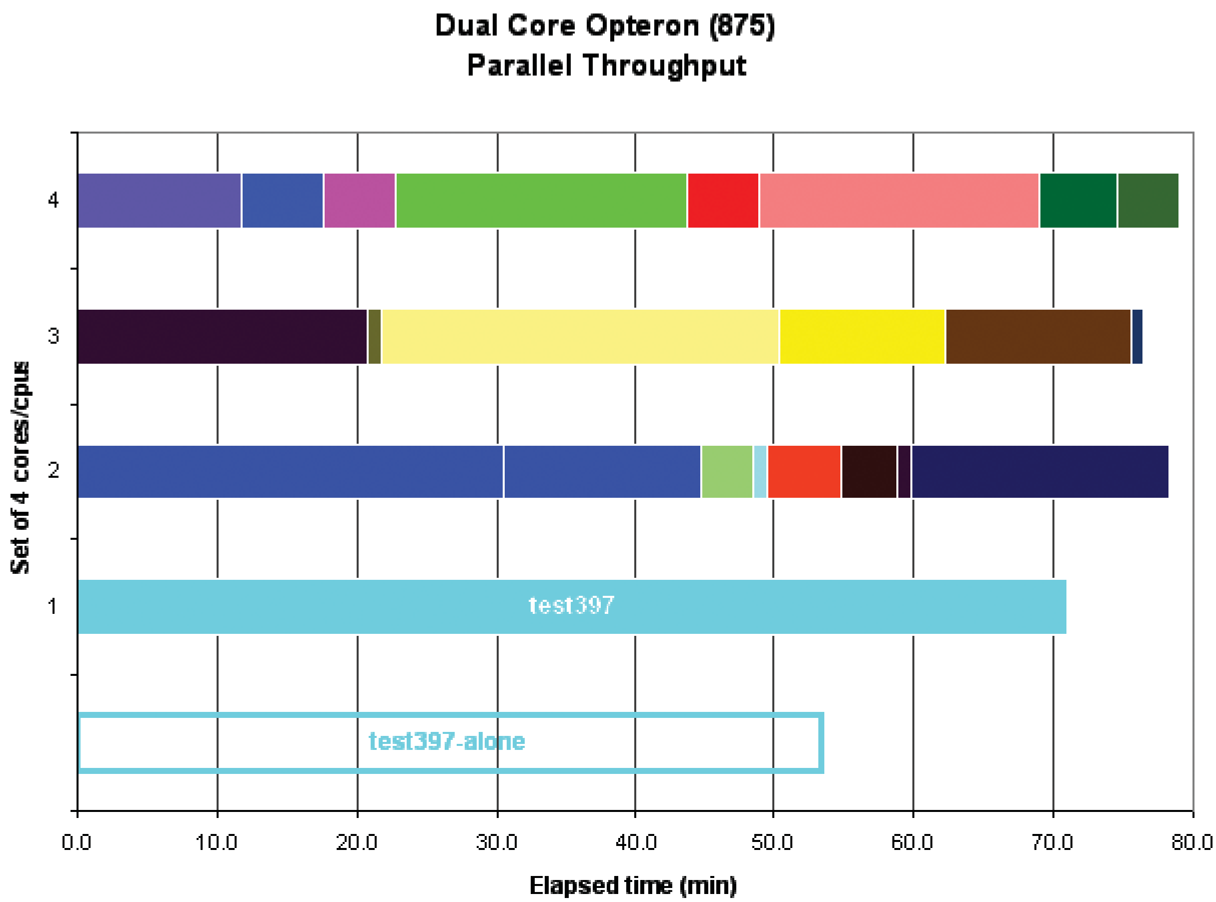

The test is designed to effectively have test397 run while a heterogeneous load is being applied to the computer system.
From the figure above can be seen that the full test set is done in 79 minutes, test397 finishes in 71 minutes. Included is also the time measured when test397 is run on another wise empty system, 53 minutes. This means that in a loaded machine test397 experiences a degradation of around 33%.

Contrast these results with those obtained on an Altix 330 1.6GHz/6MB

Source: SGI internal benchmark testing, December, 2005

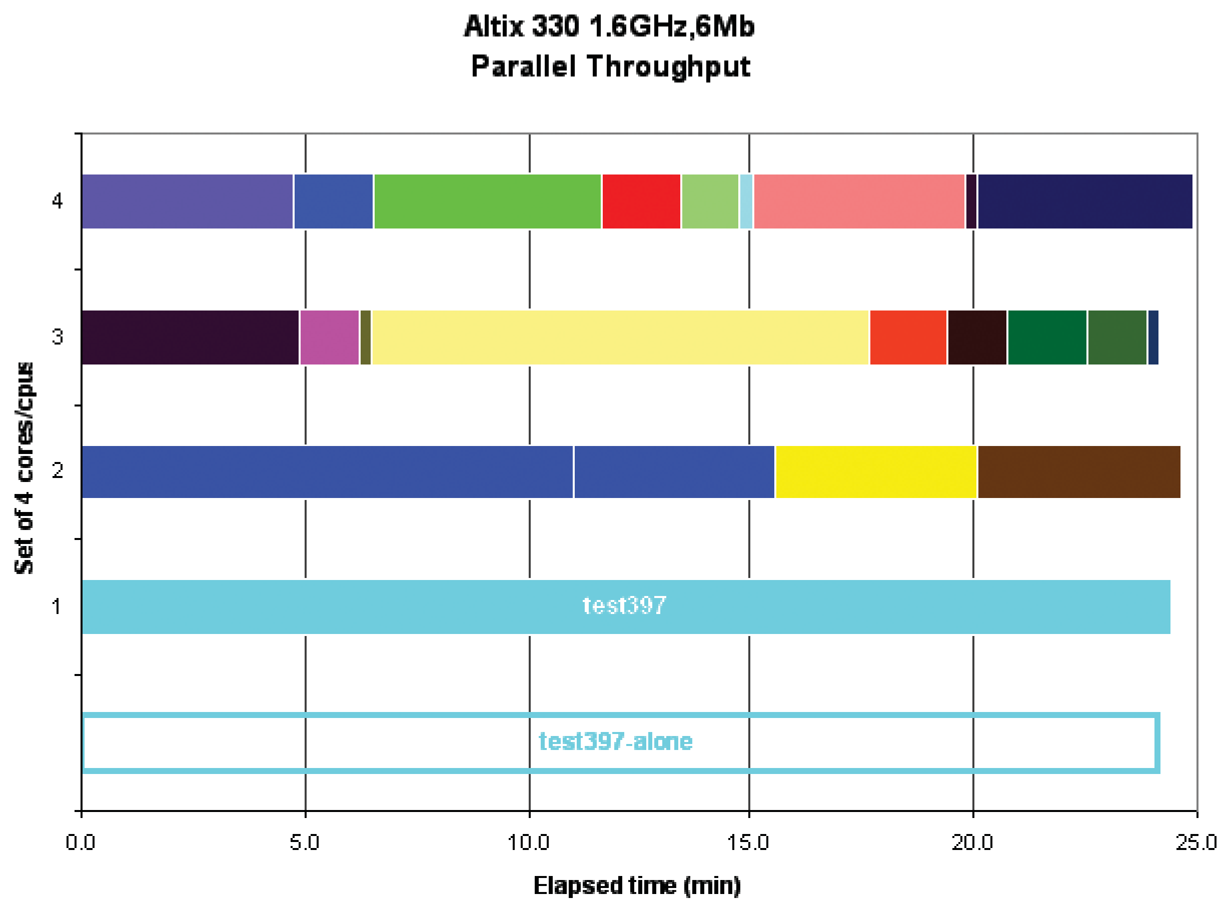

The total execution time for the 23 tests was 25 min, this is over 3 times faster than the Dual Core Opteron. Test397 finished in 24.4 minutes whereas the stand-alone time is 24.1 minutes, this is around 1% degradation.

The results of the same throughput test on an Altix 330 1.5GHz/4MB should be similar.

Source: SGI internal benchmark testing, December, 2005

3.2.3 CASTEP

The figures below compare CASTEP Opteron performance against an Altix 330 1.6 GHz/6MB.

CASTEP on the Altix shows very good performance relative to the dual core Opteron system. This is partially attributed to the memory bandwidth requirements of the application.

Note that the single core TiN_Cas failed to run on the Opteron-based machines systems.

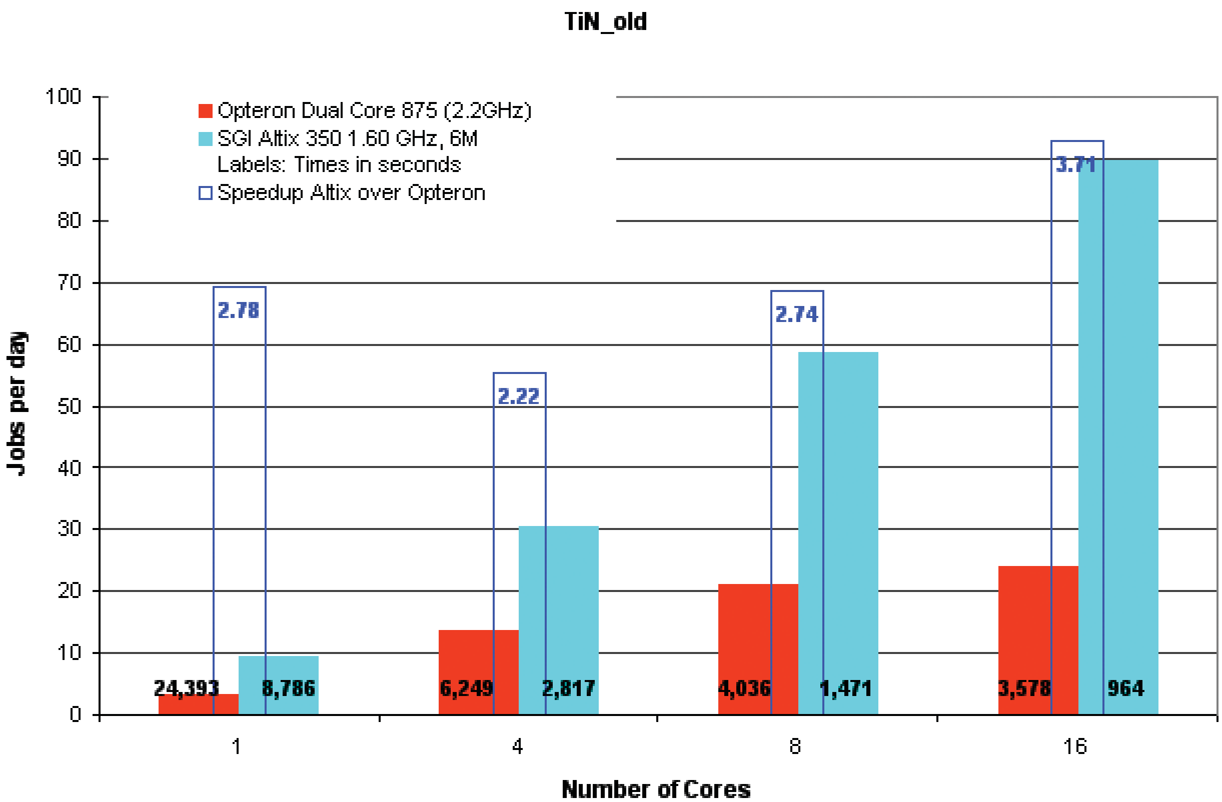

Source: SGI internal benchmark testing, December, 2005

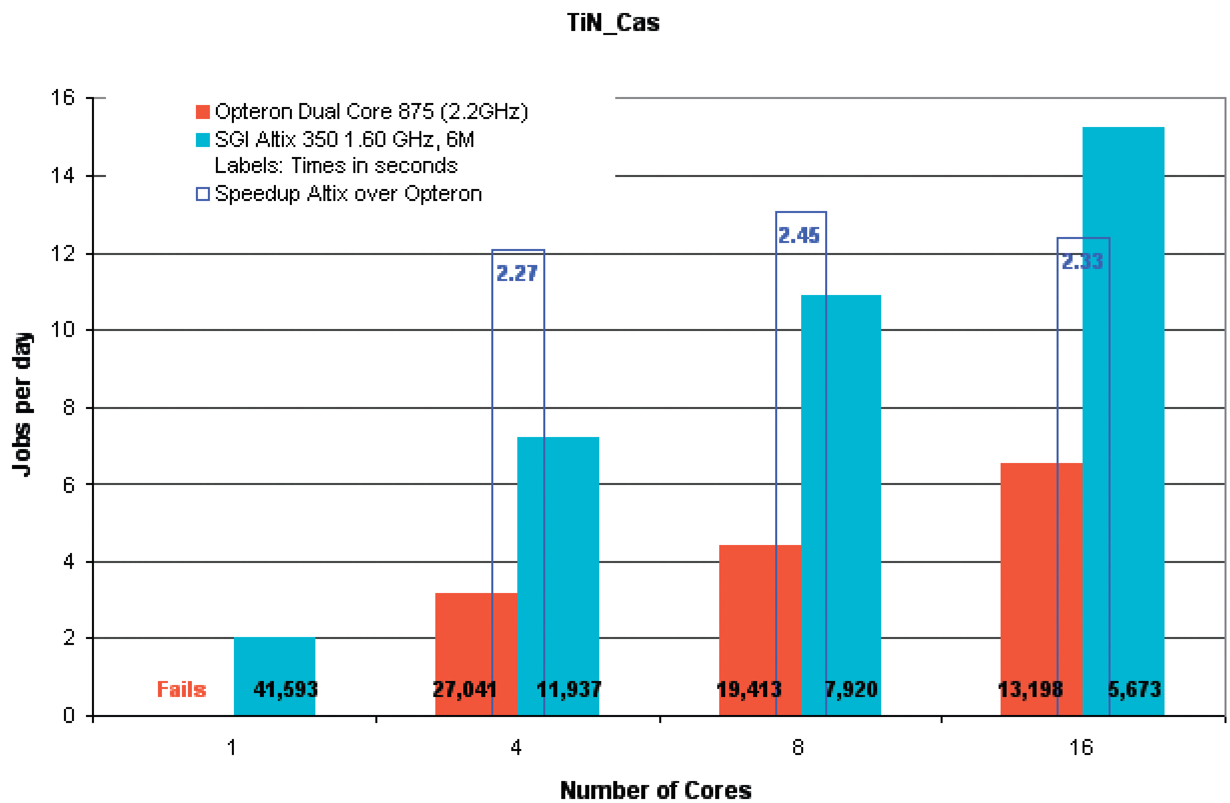

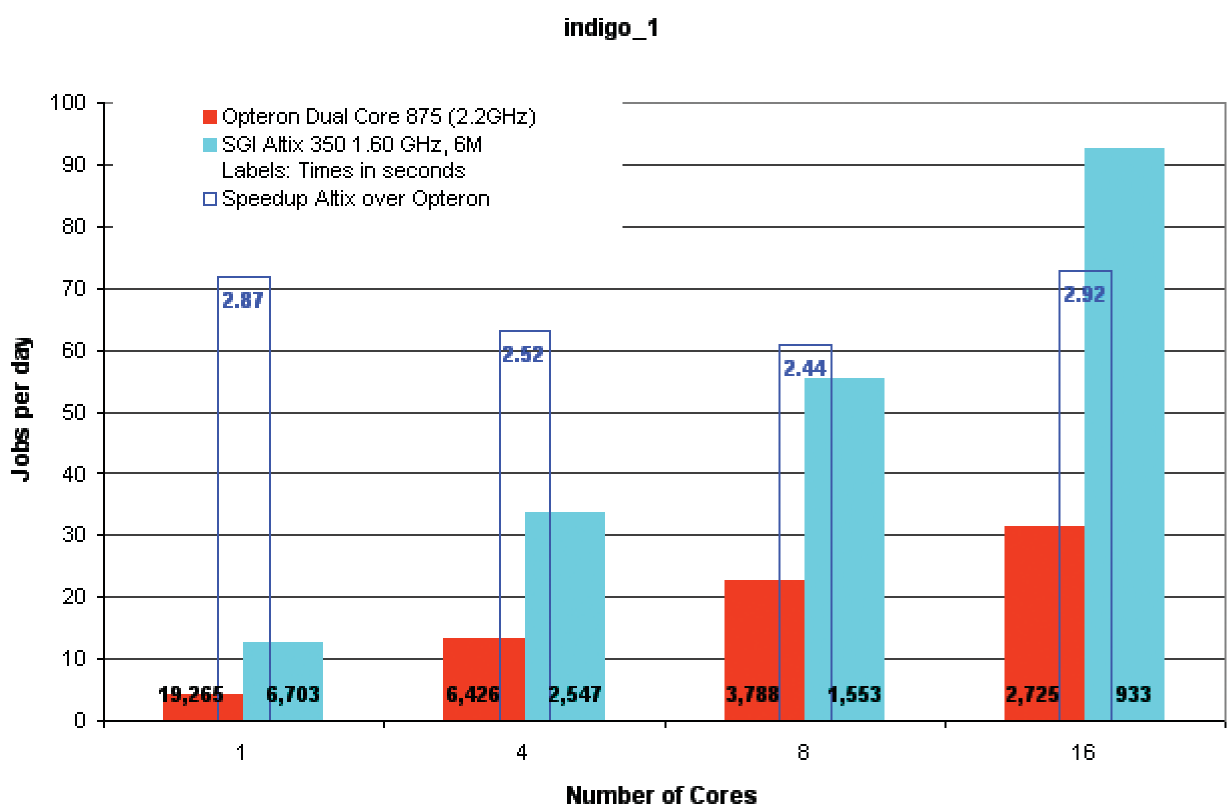

Source: SGI internal benchmark testing, December, 2005

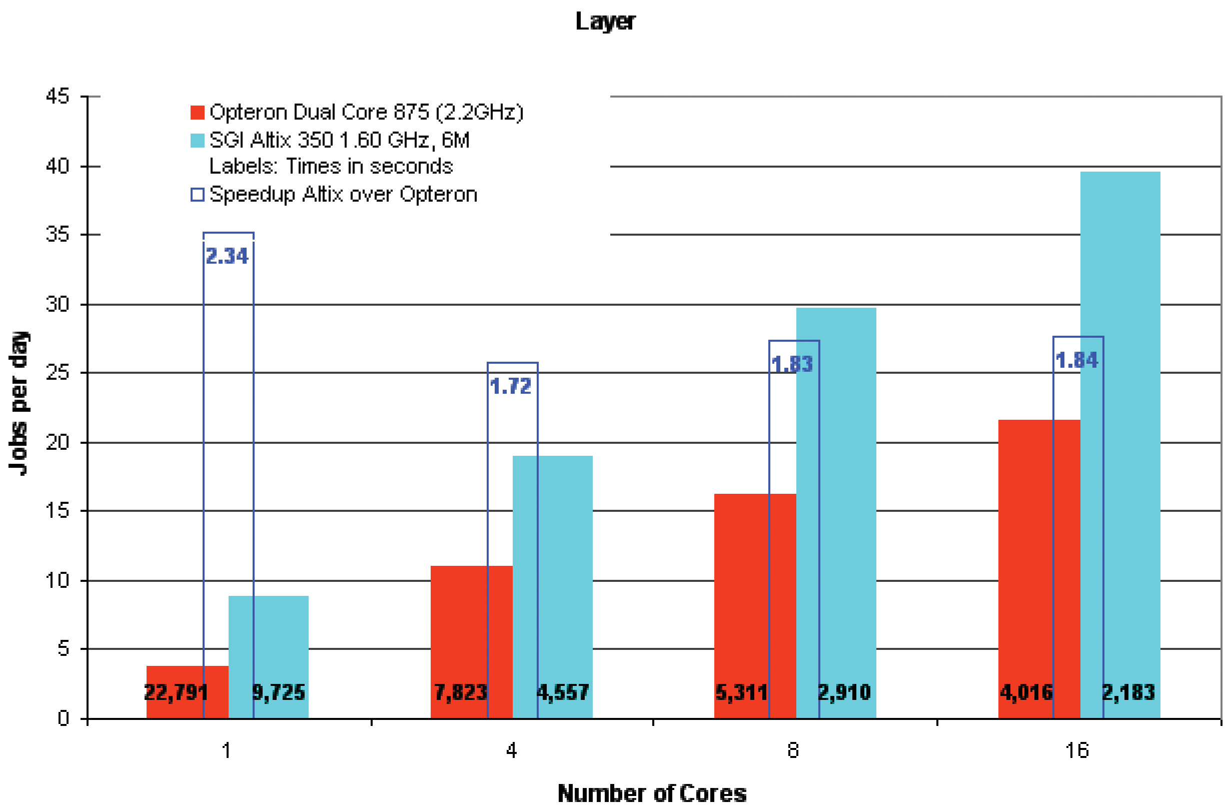

Source: SGI internal benchmark testing, December, 2005

3.2.4 DMol[3]

The figures below compare DMol[3] Opteron performance against an Altix 330 1.5GHz/4MB

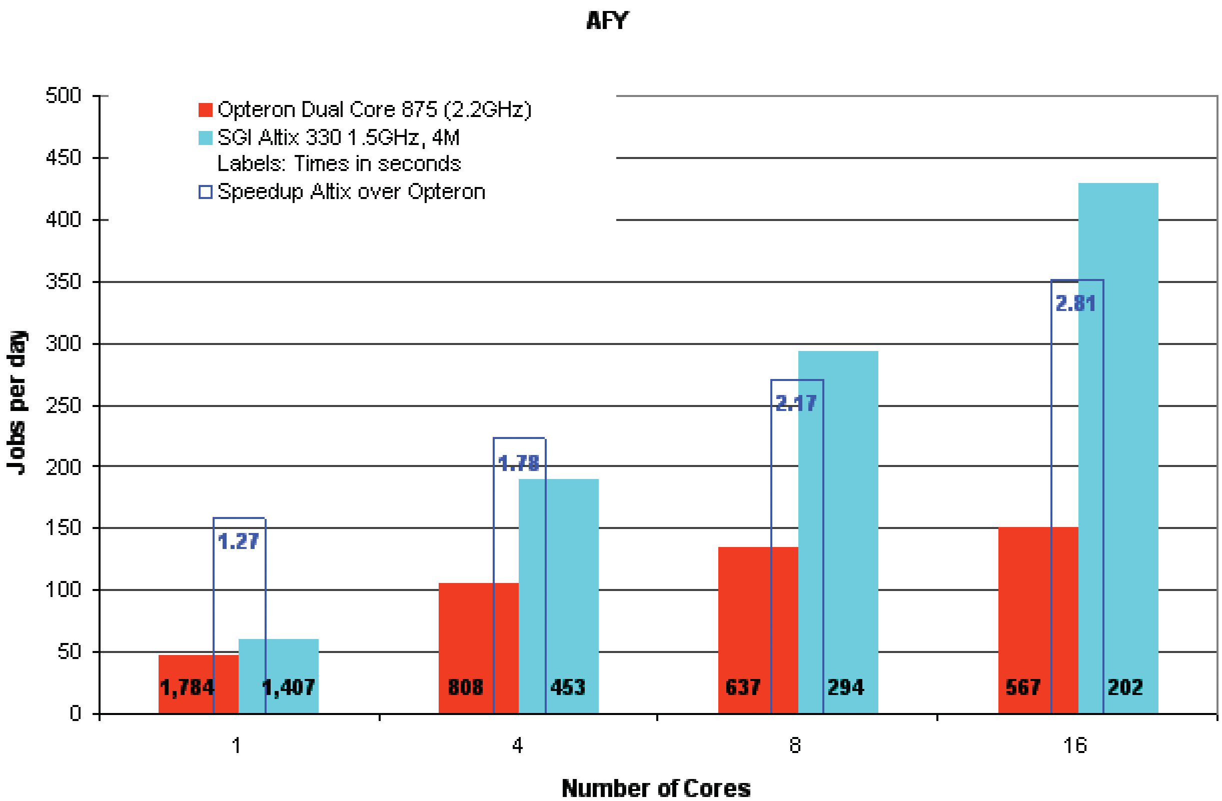

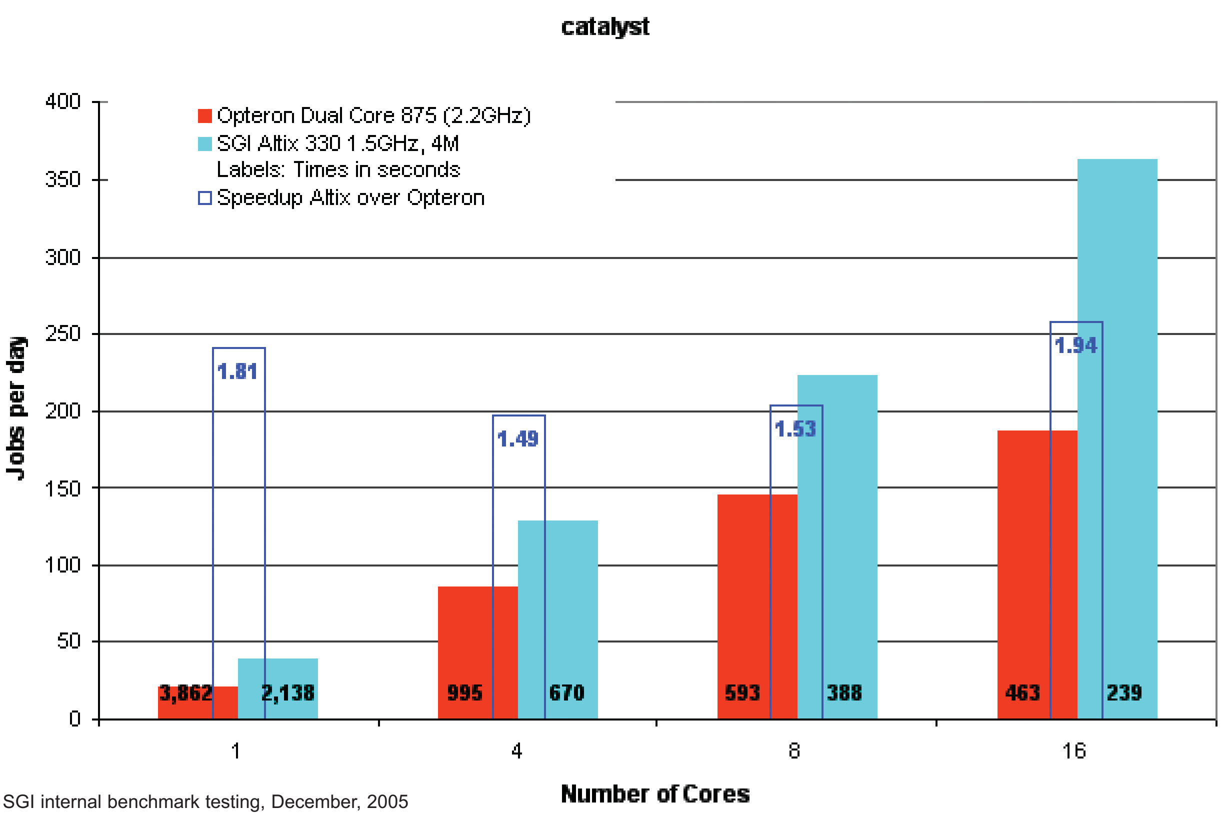

Source: SGI internal benchmark testing, December, 2005

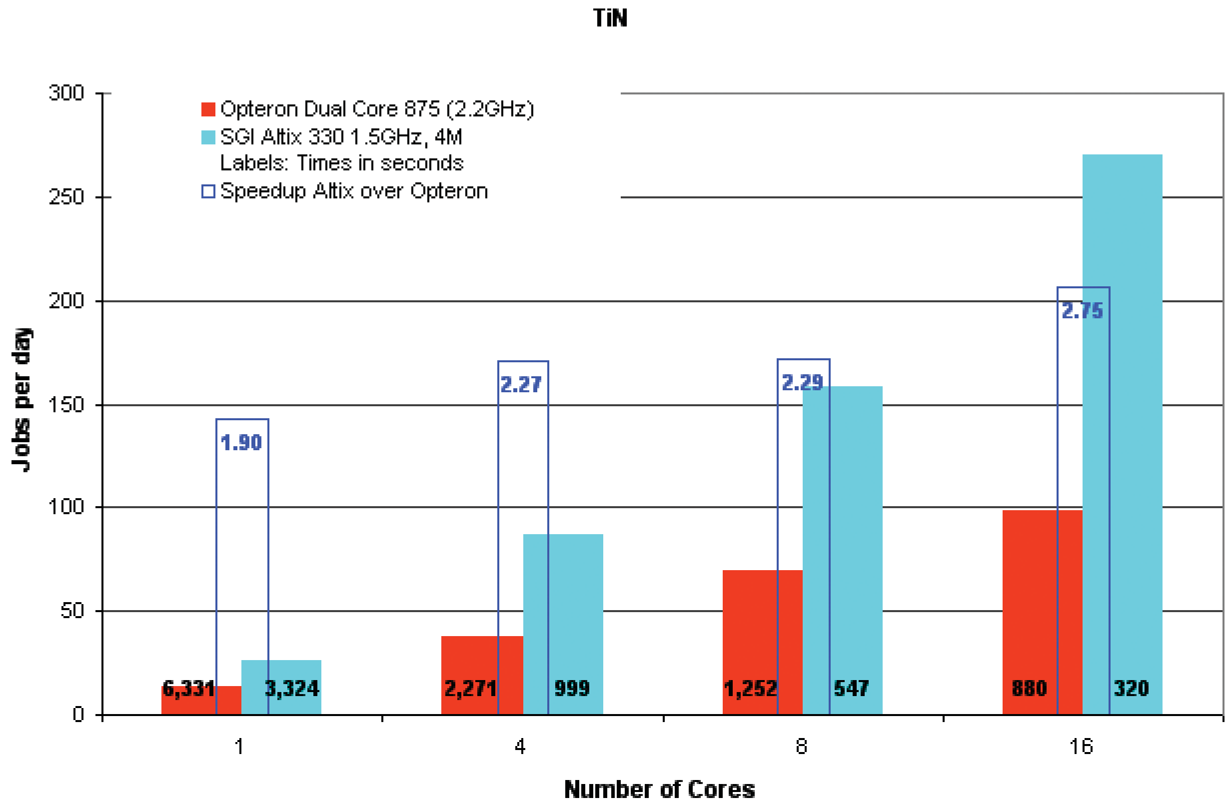

3.2.5 Amber 8 (sander)

The figures below compare Sander Opteron performance against an Altix 330 1.5GHz/4MB

The relative big degradation of the dual core Opteron in some cases when going from 8 to 16 cores is probably attributable to the latency characteristics of the application and of MPICH.

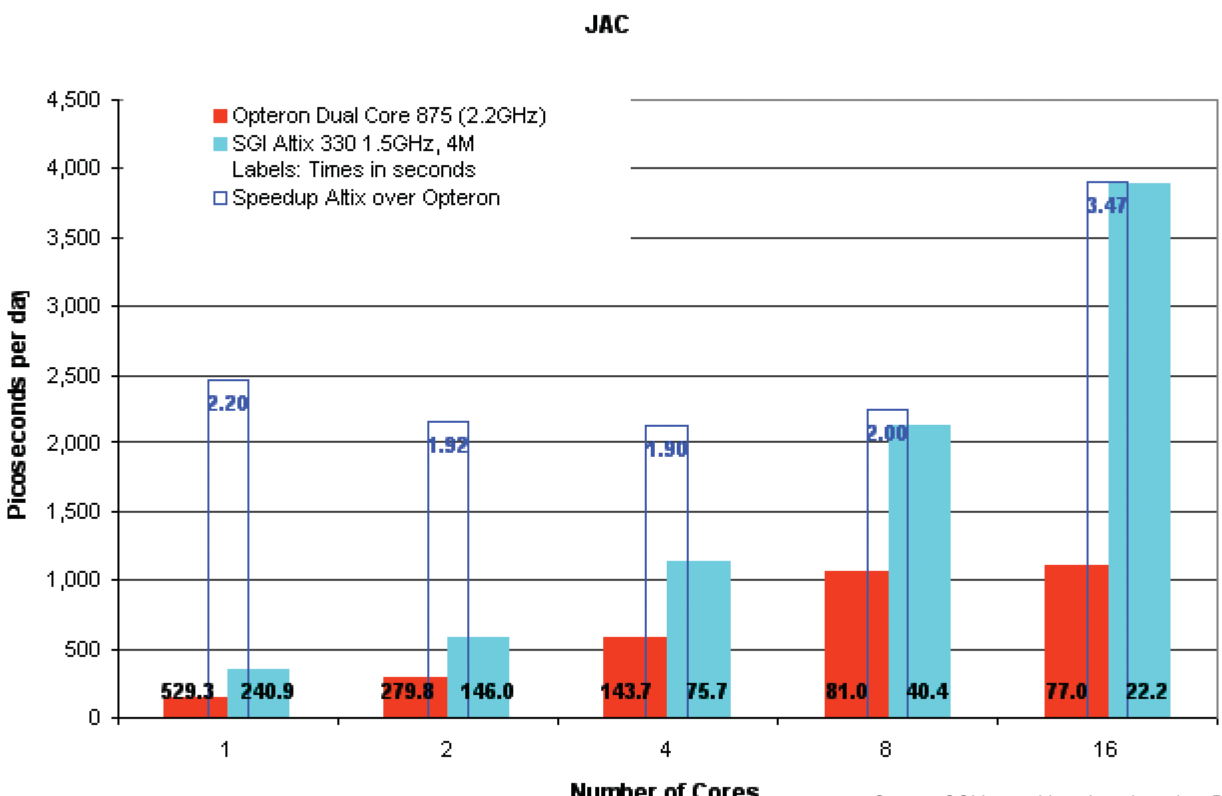

Source: SGI internal benchmark testing, December, 2005

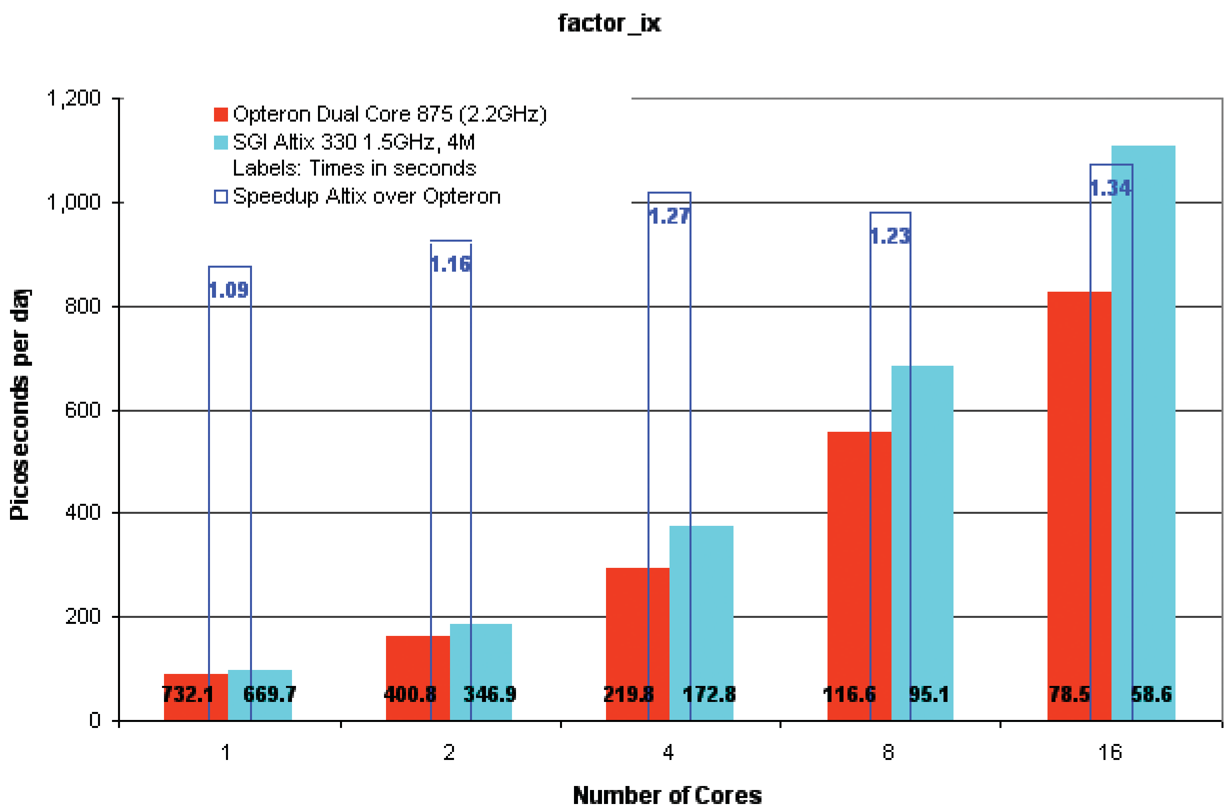

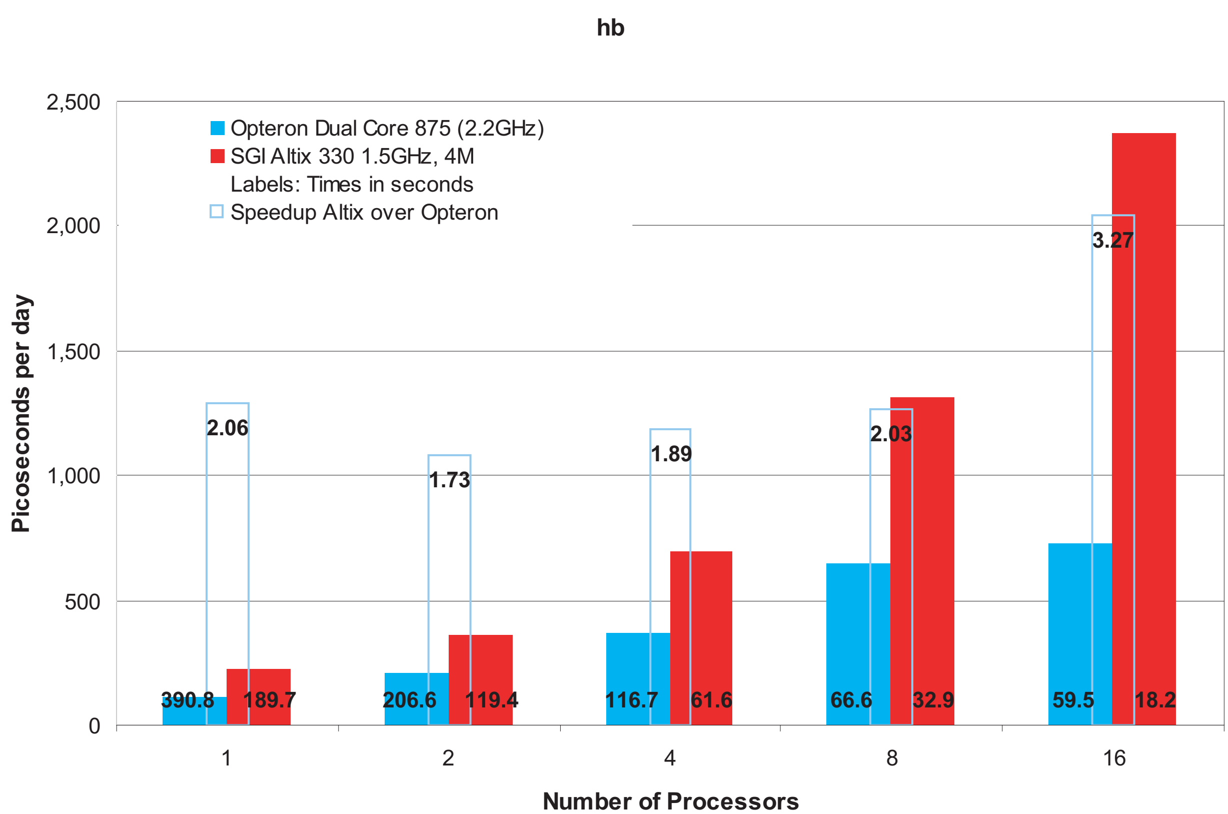

Source: SGI internal benchmark testing, December, 2005

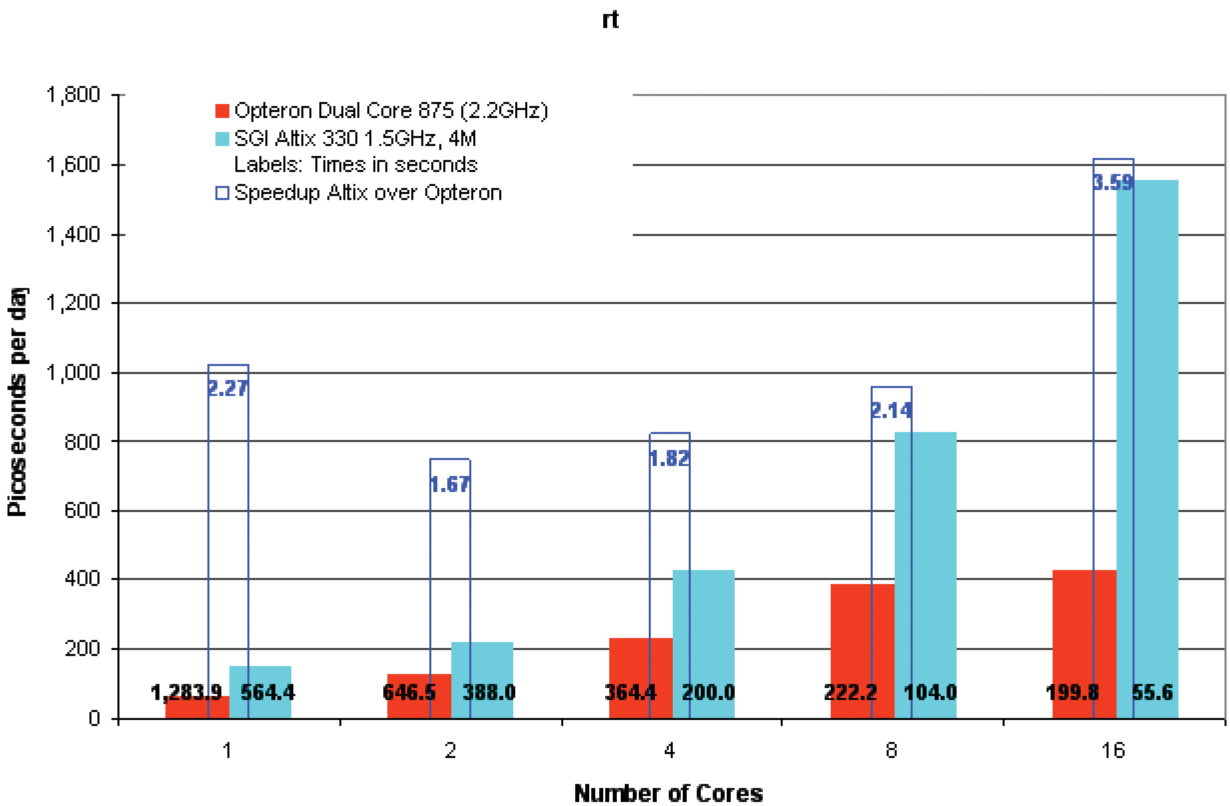

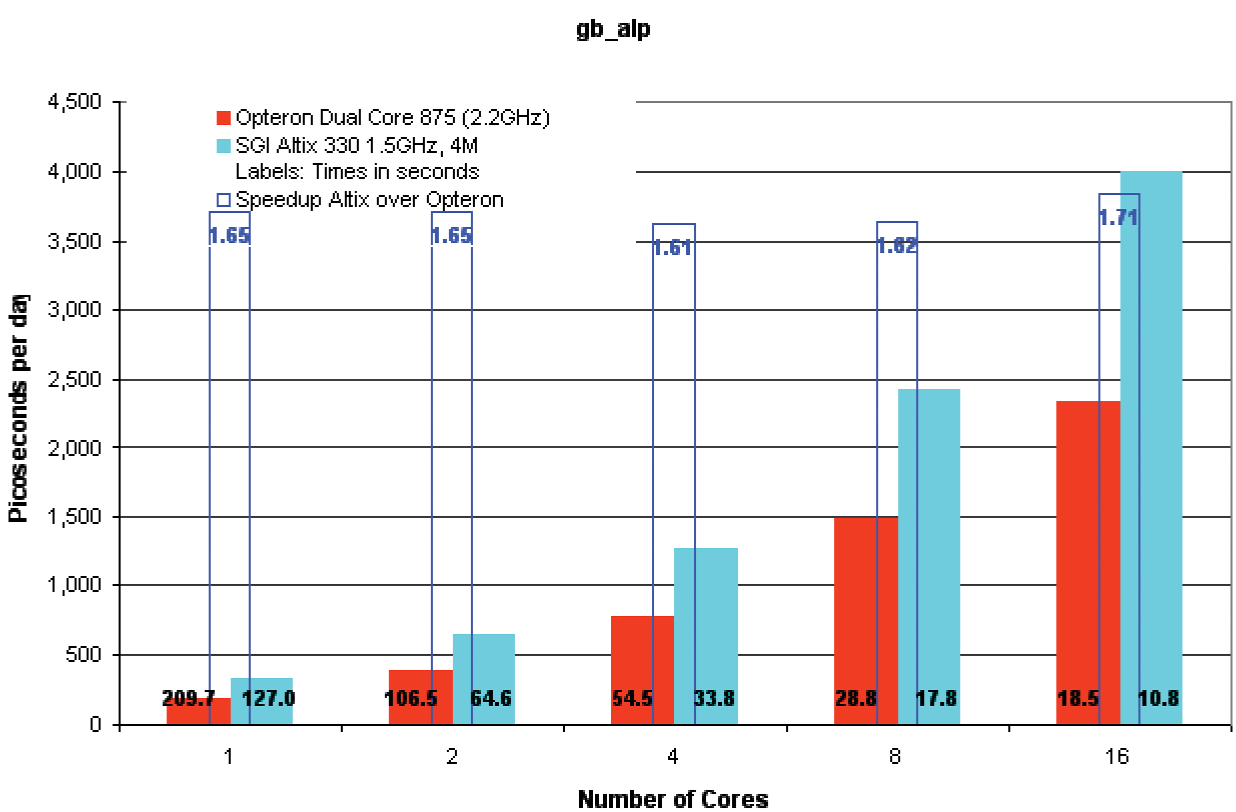

Source: SGI internal benchmark testing, December, 2005

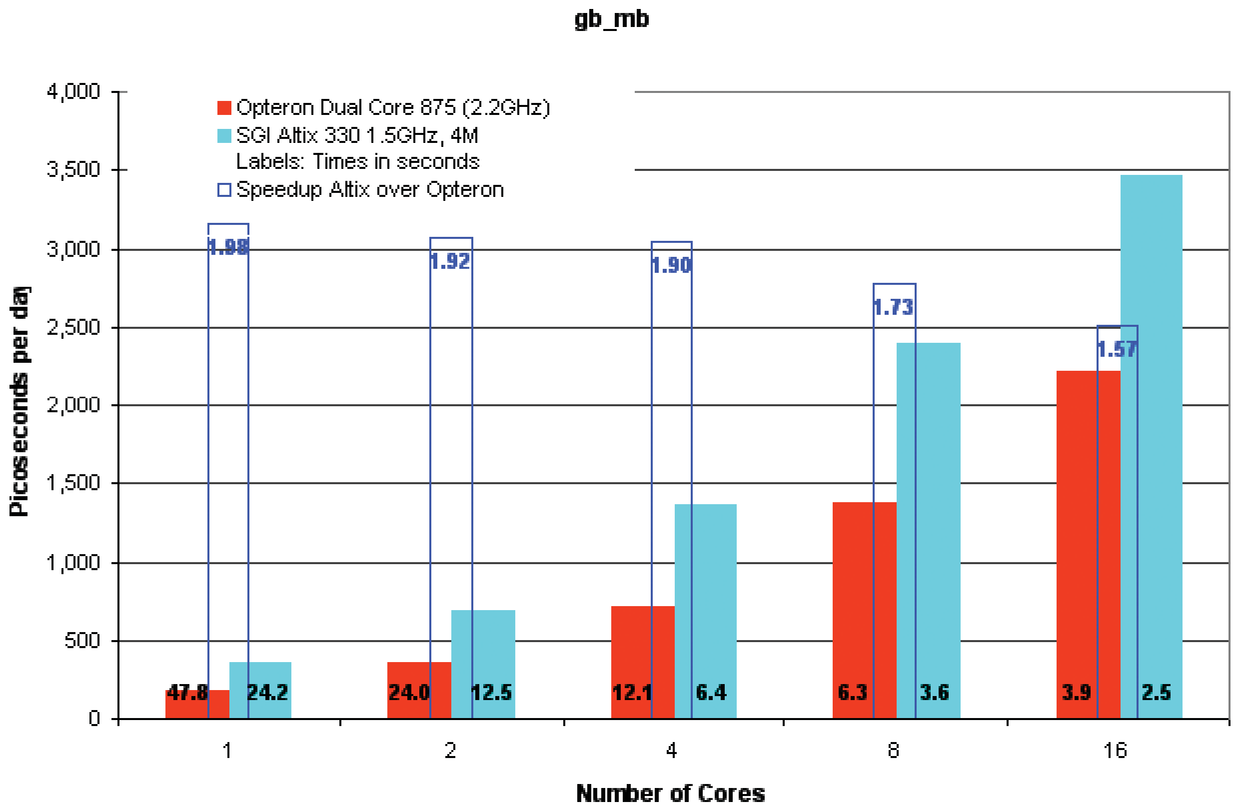

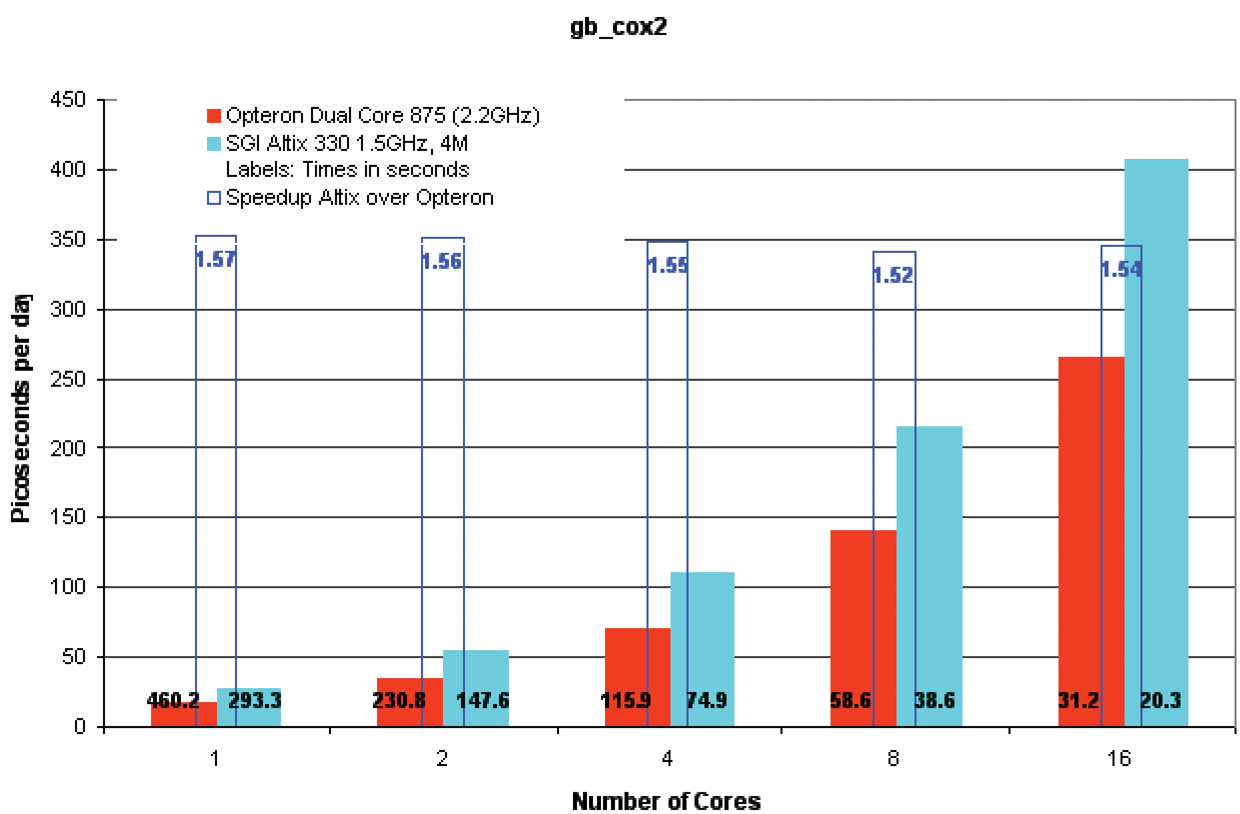

Source: SGI internal benchmark testing, December, 2005

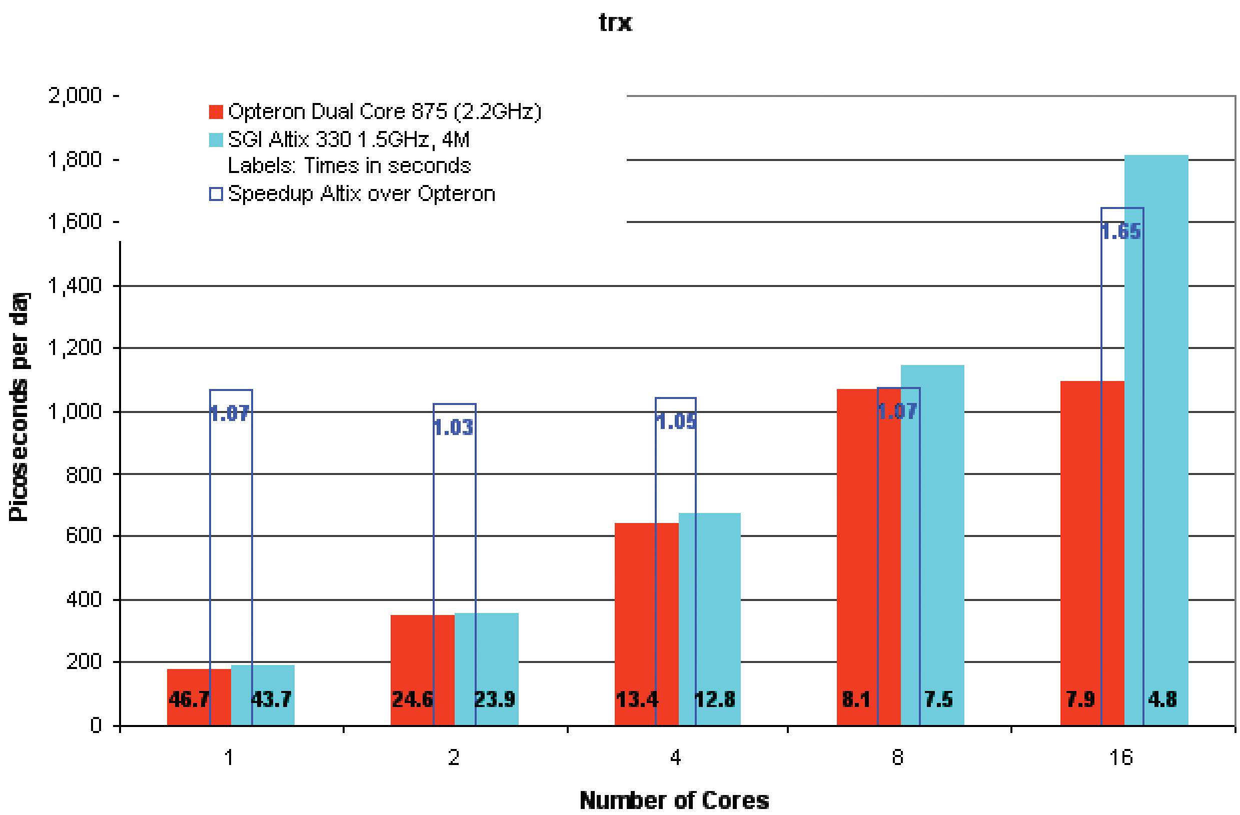

3.2.6 Amber8 (PMEMD)

The figures below compare PMEMD Opteron performance against an Altix 330 1.5GHz/4MB

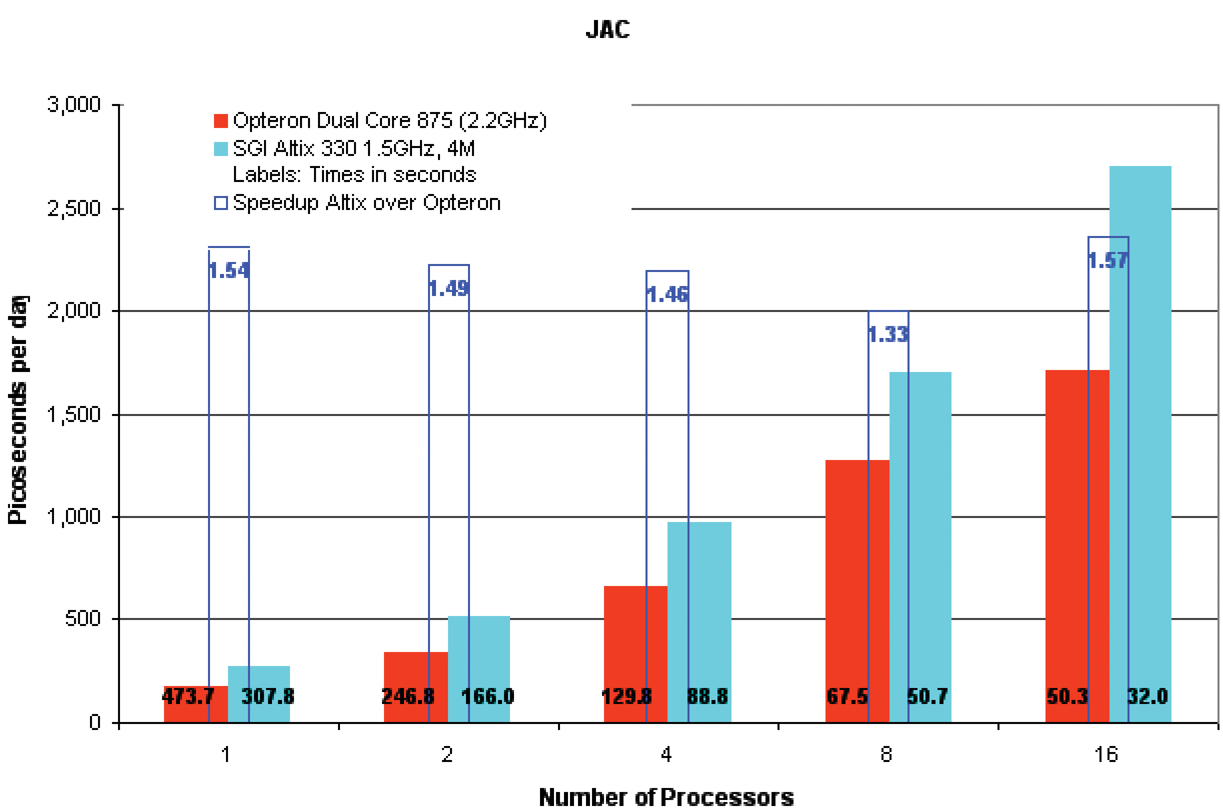

Source: SGI internal benchmark testing, December, 2005

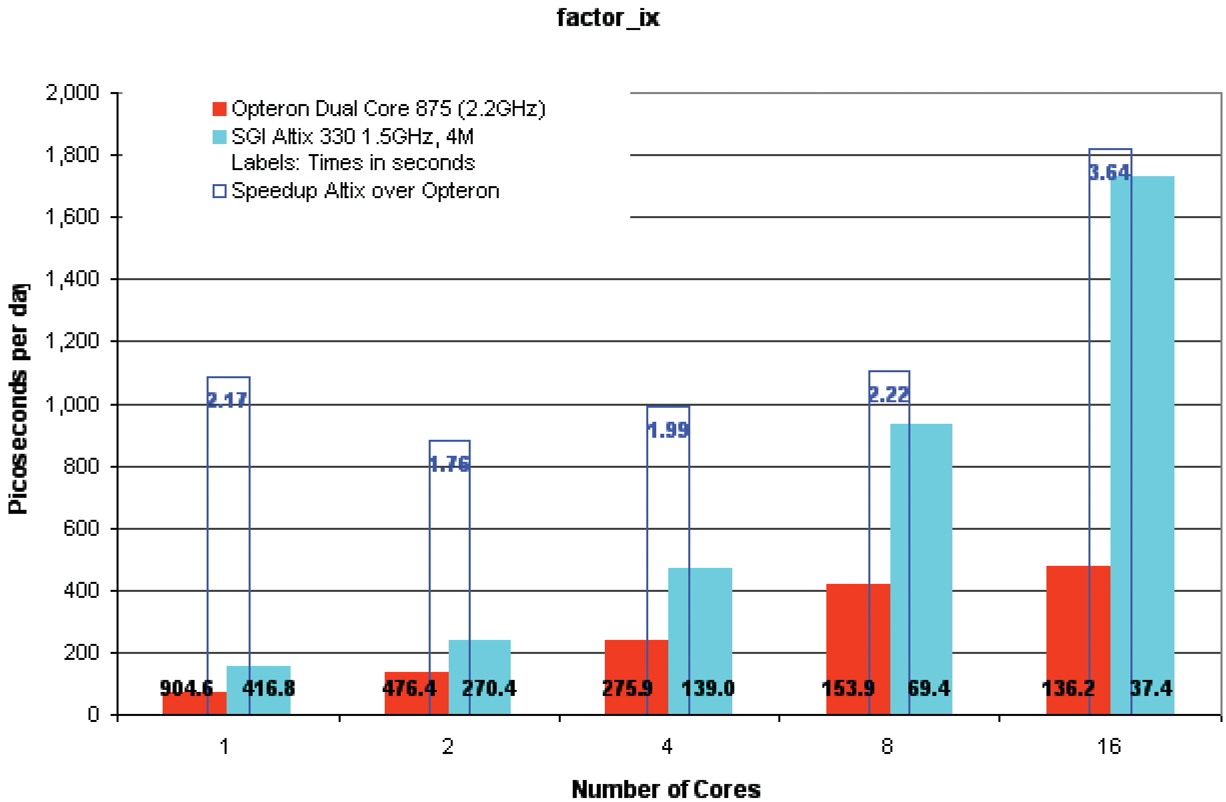

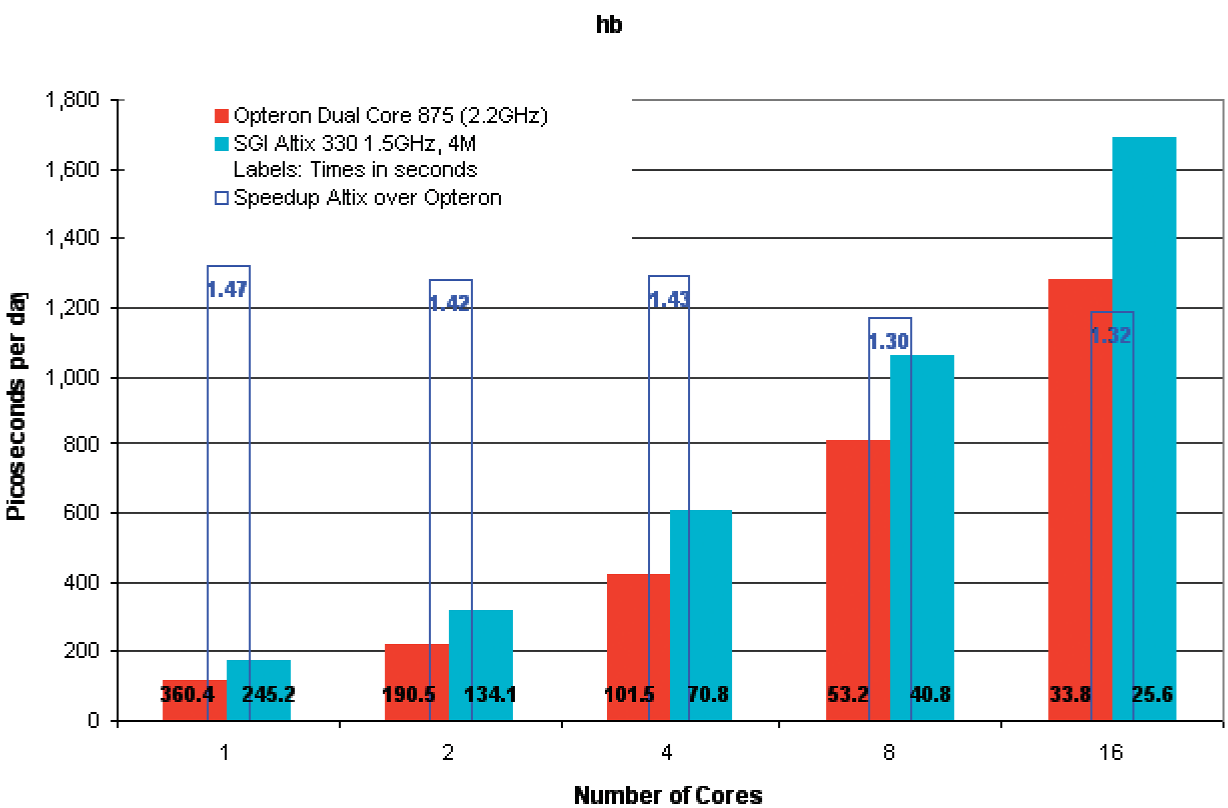

Source: SGI internal benchmark testing, December, 2005

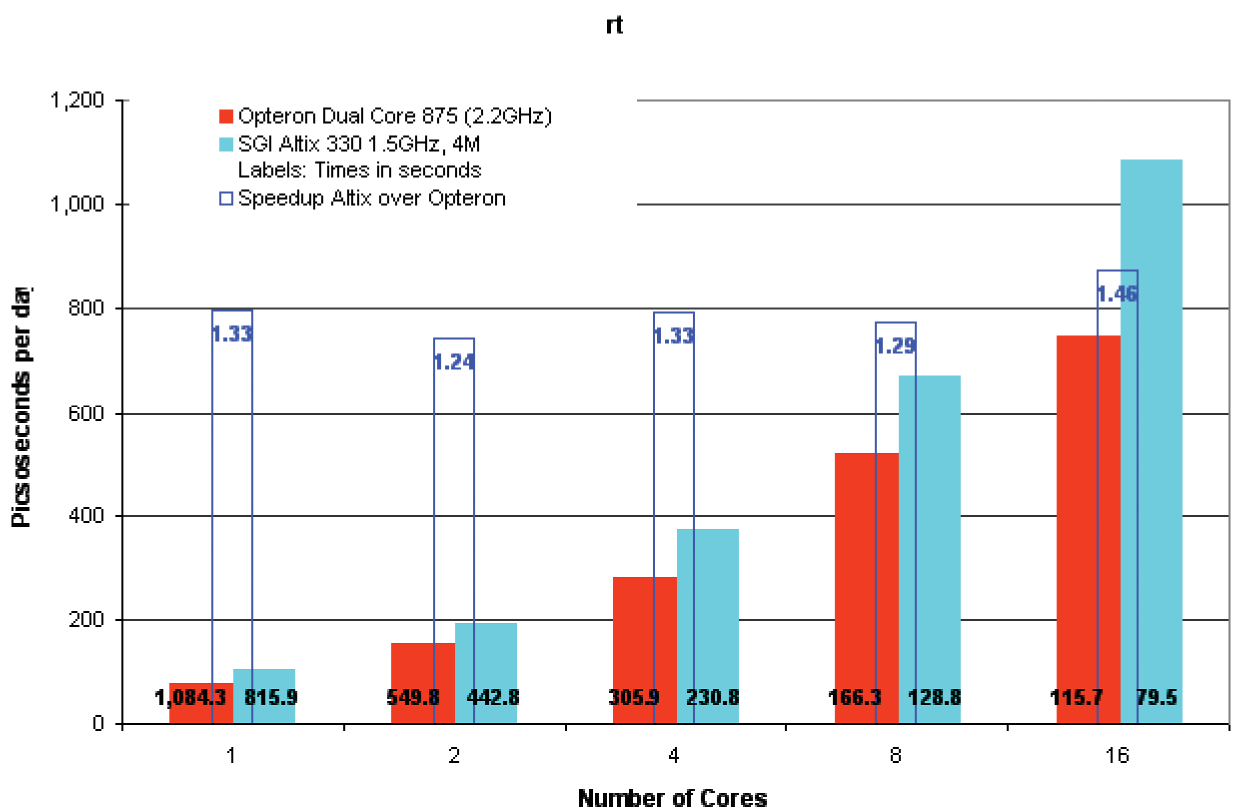

3.2.7 NAMD

The figures below compare NAMD Opteron performance is against an Altix 330 1.5GHz/4MB

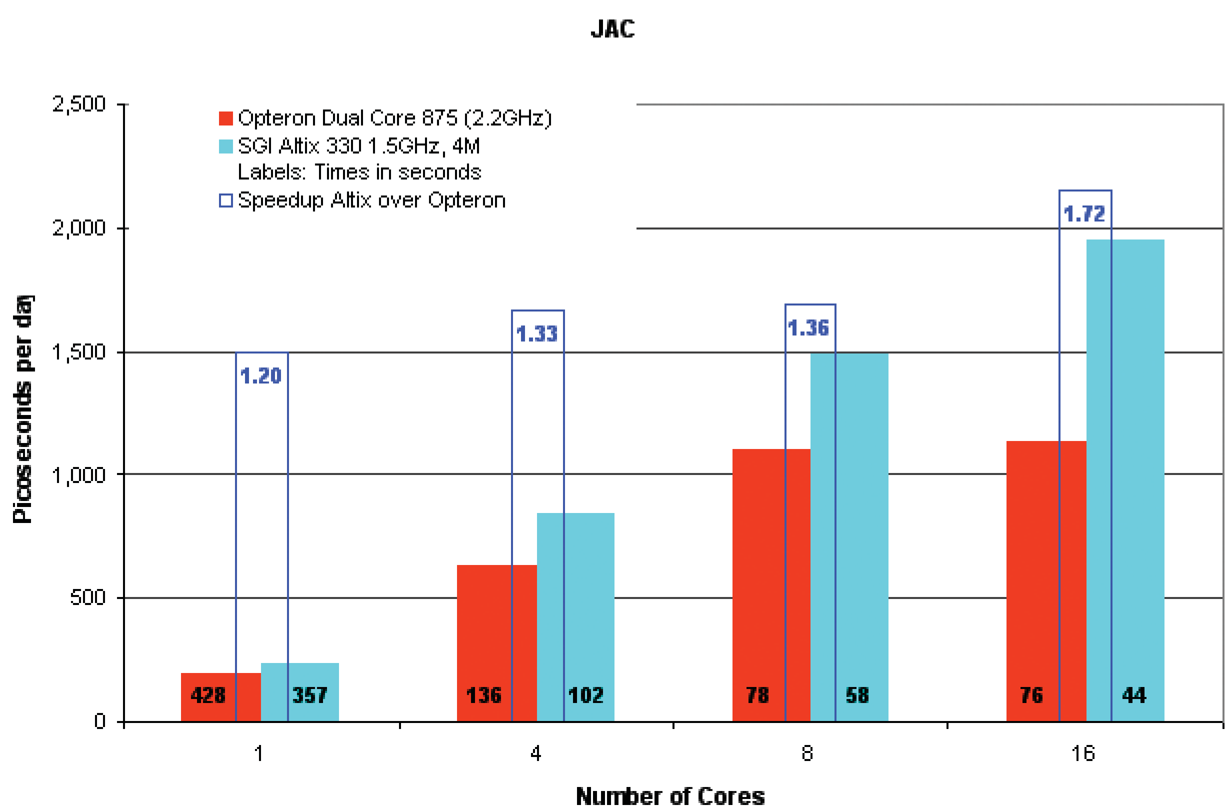

Source: SGI internal benchmark testing, December, 2005

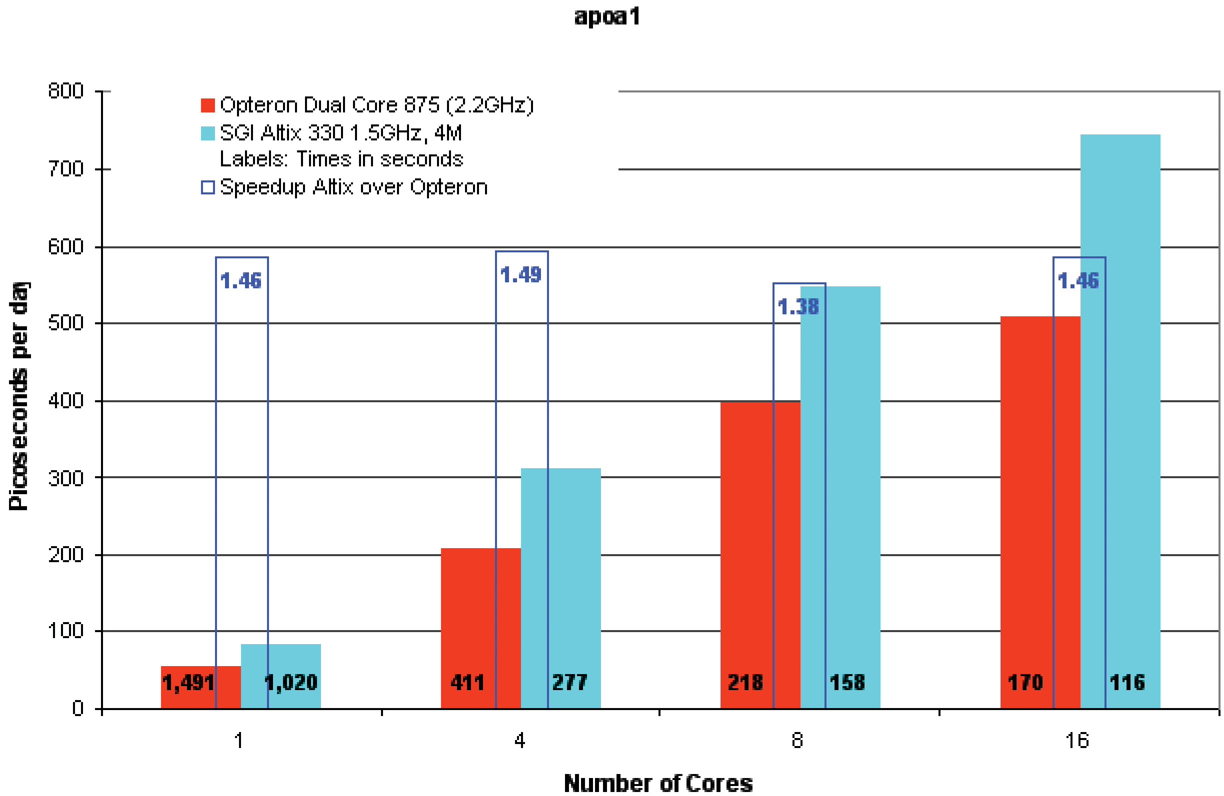

3.2.8 VASP

The figures below compare VASP Opteron performance against an Altix 330 1.5 HGz/4MB

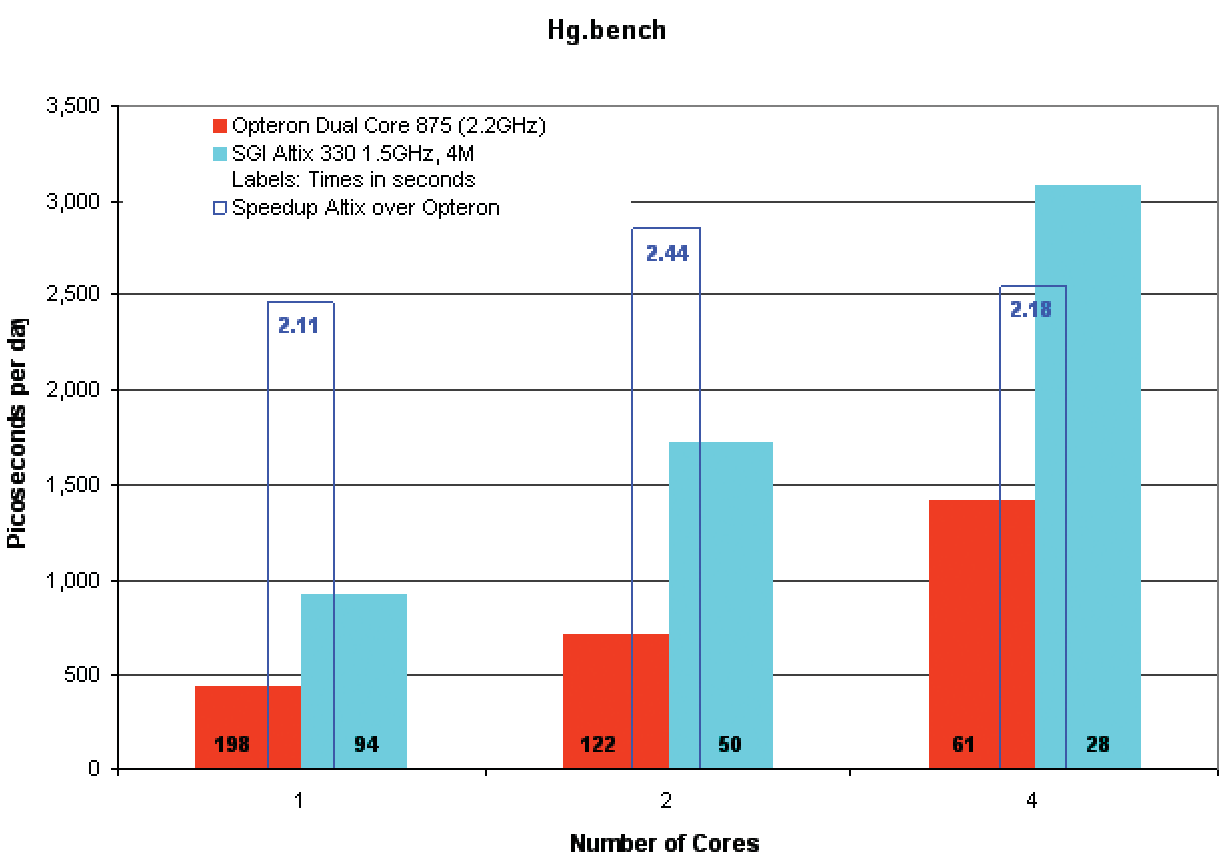

Source: SGI internal benchmark testing, December, 2005

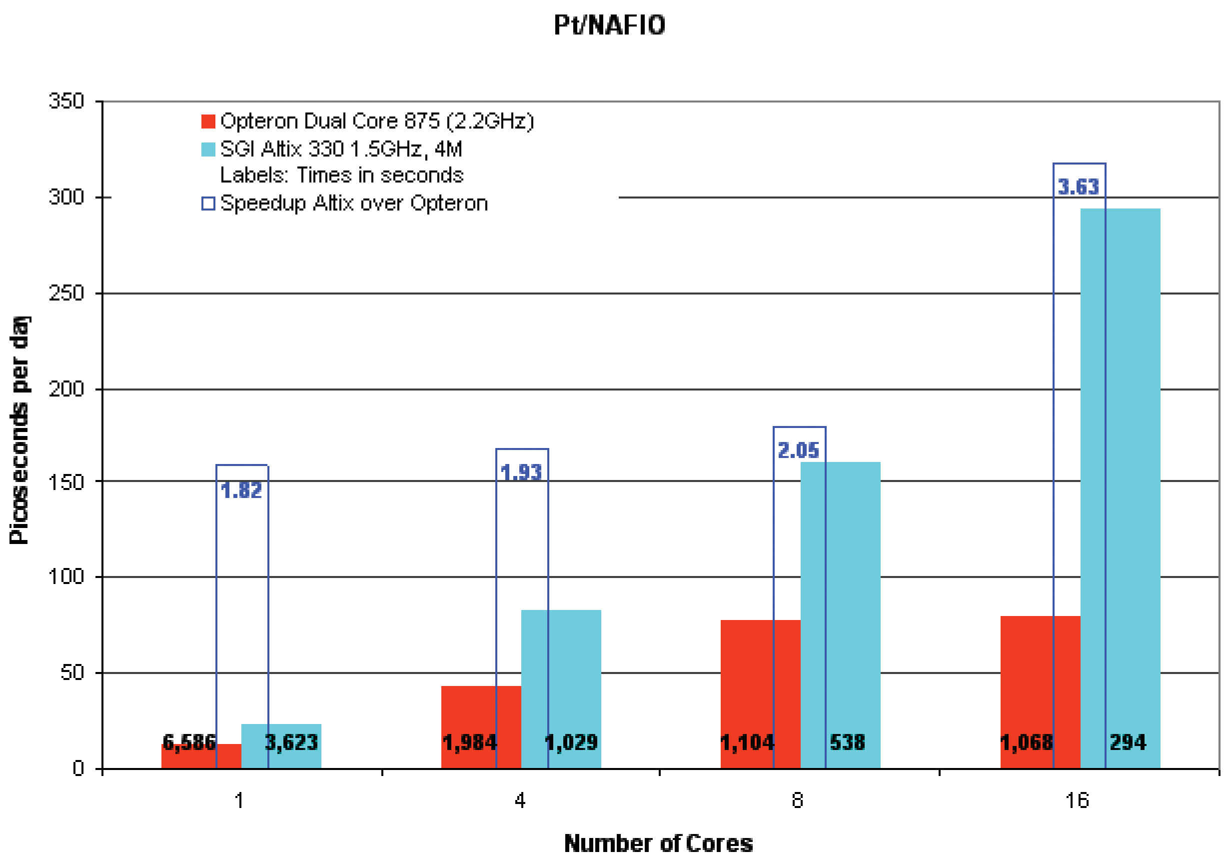

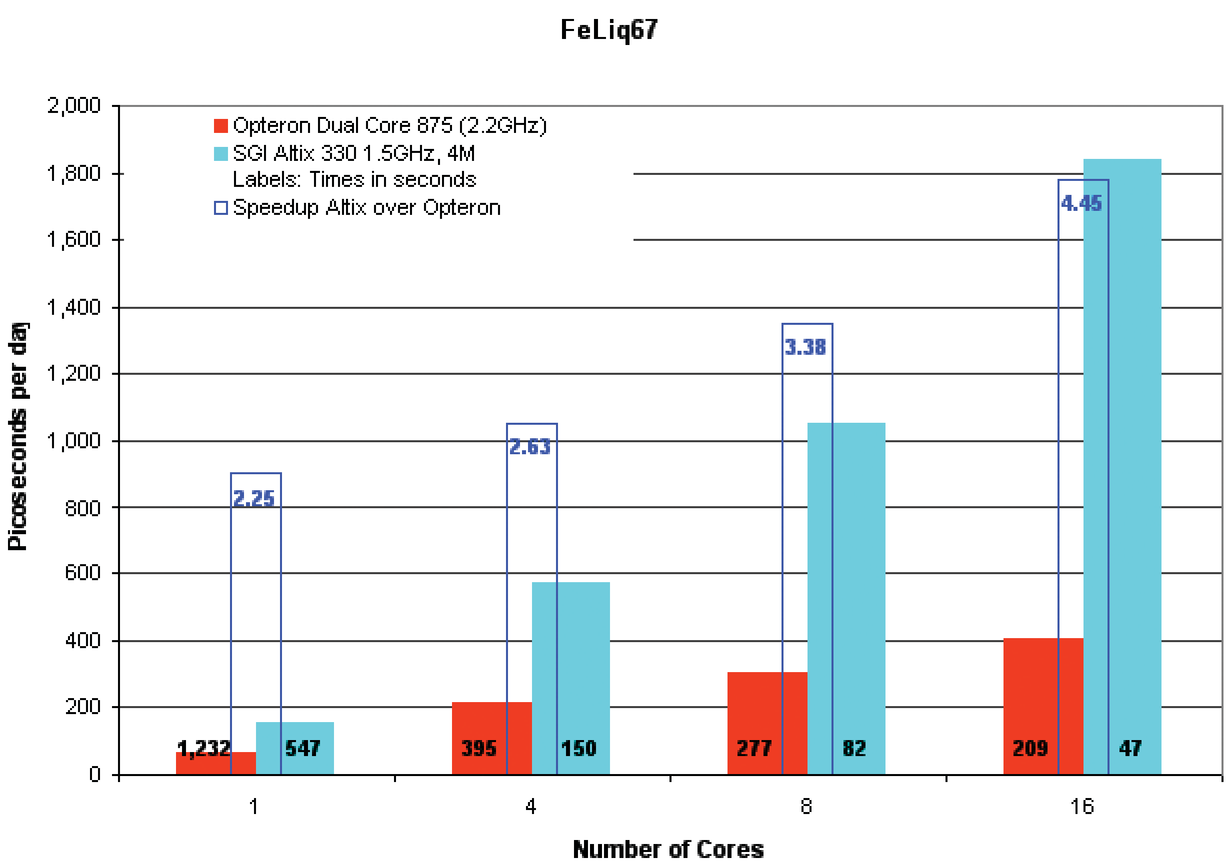

Source: SGI internal benchmark testing, December, 2005

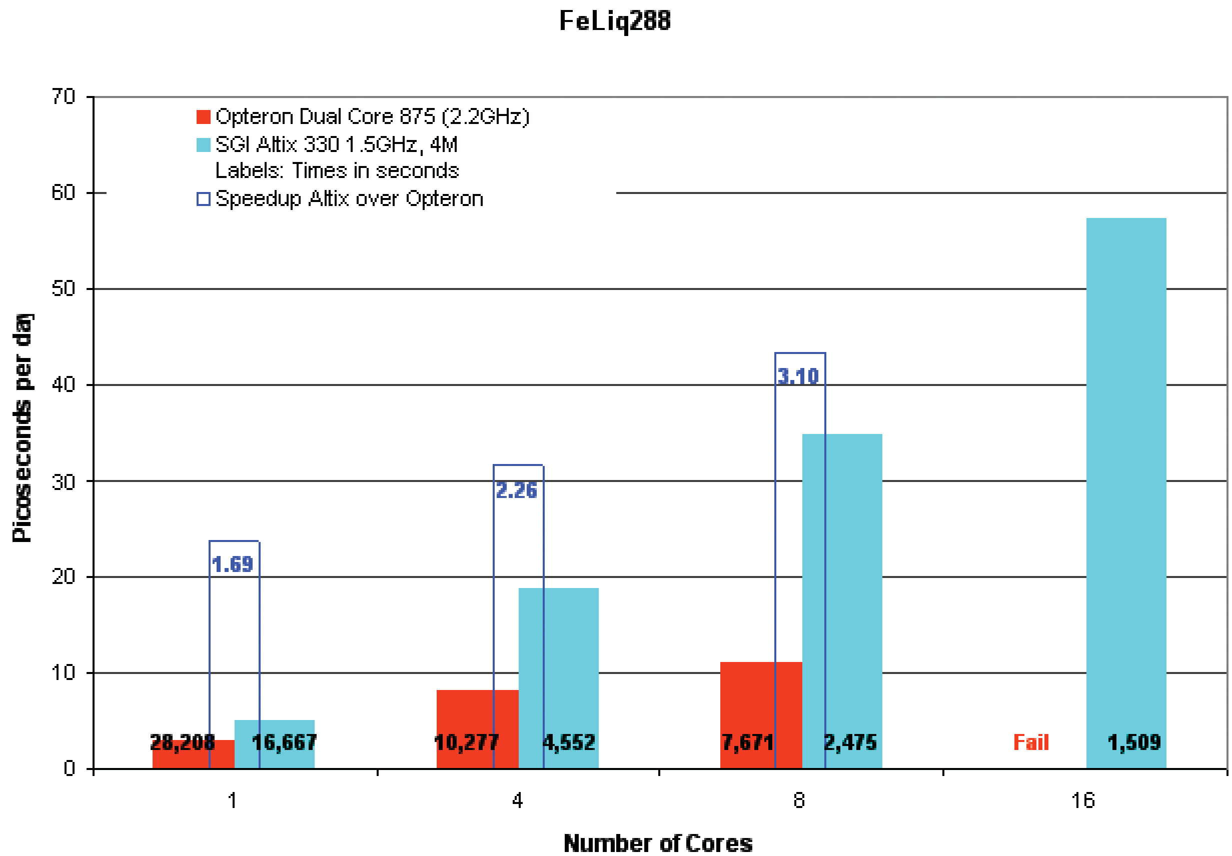

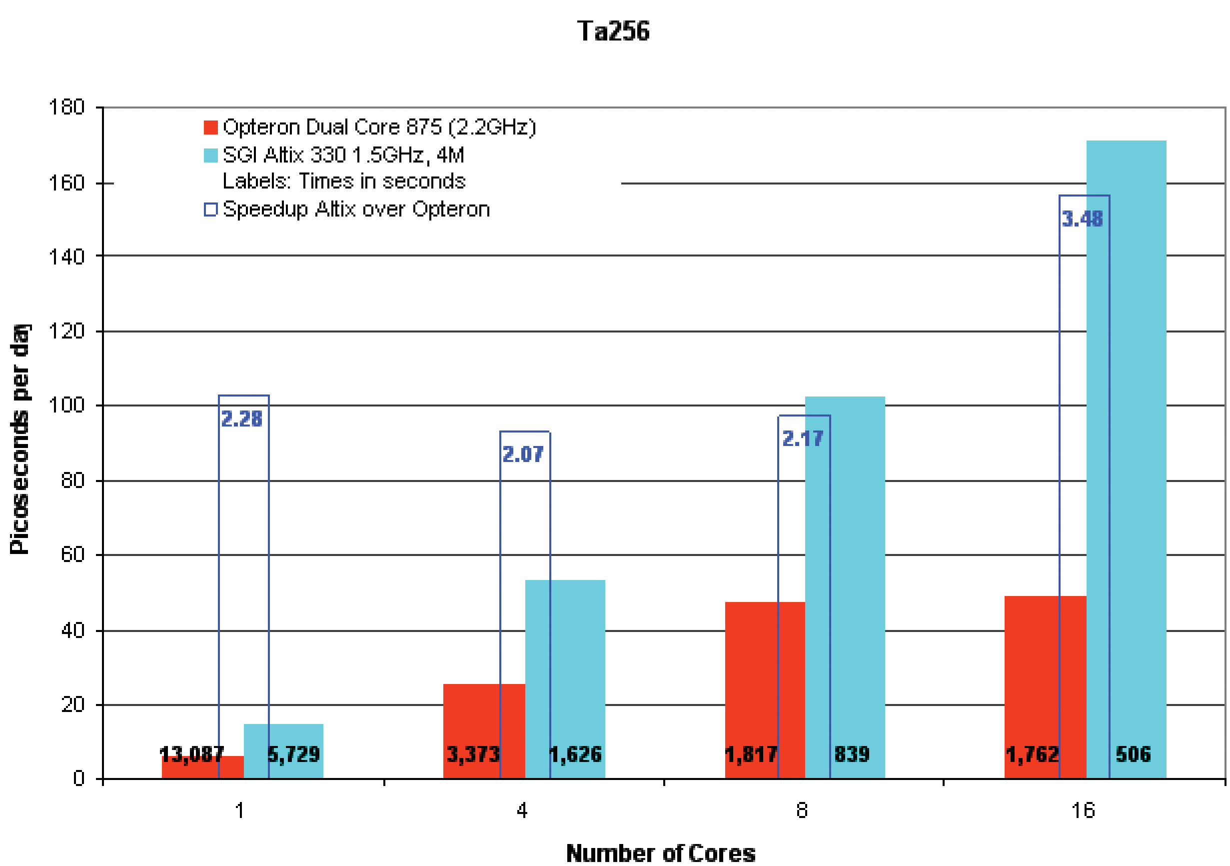

Source: SGI internal benchmark testing, December, 2005

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of SGI products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.

Source: SGI internal benchmark testing, December, 2005

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3905 [02.02.2006] J15113