Industry Brief

SGI® Computational Chemistry Applications Performance Report

Spring 2002

Table of Contents

For detailed technical support information on computational chemistry applications running on SGI systems, please visit the SGI Web site at *www.sgi.com/solutions/sciences/chembio/tech_resources.html*.

You will find general information, technical information, porting notes, and known problems and fixes.

For more information on how SGI can help you solve your computational problems, visit *www.sgi.com/solutions/sciences/chembio/* or contact us at *chembio@sgi.com*. To contact your local SGI sales representative or distributor, please see *www.sgi.com/sales/contact_gen.html* for telephone numbers and office locations.

1.0 Introduction

We are pleased to present this collection of tests of computational chemistry applications to help you understand the computational performance of SGI computer systems.

These tests provide a small but representative sample of timing data to show the type of performance you might expect from the hundreds of chemistry and biology applications that are currently available on SGI computer systems. The combination of optimized and parallelized applications, high-performance graphics computers, and supercomputers from one company—SGI—offers an outstanding integrated solution for scientific research.

This report provides information on the performance of a number of leading computational chemistry applications running on SGI® Origin® family servers. The SGI team of computational chemistry professionals has worked with the software providers to optimize the performance of these applications. They take full advantage of the scalability of the shared-memory SGI Origin servers by incorporating parallelism into these applications.

Computational chemistry applications such as Amber, CHARMM, CNX, GAMESS, Gaussian®, Jaguar, NWChem, and MOPAC can be evaluated as to their computational performance from two perspectives:

- Throughput performance, which addresses the need to process a large batch of jobs in the shortest amount of time
- Turnaround performance, which measures how quickly an individual job is completed

Delivering high-throughput and fast-turnaround performance for the full range of computational chemistry applications is one of the main goals of the SGI computational chemistry applications team. This report highlights how SGI Origin family systems meet both types of computational demands.

High levels of application performance and scalability on SGI Origin family servers stem from highly tuned applications benefiting from the performance of the MIPS® RI2000™ and RI4000™ processors and the supporting memory and I/O subsystems of SGI Origin family servers. Production computing customers also benefit from a number of important features offered by the IRIX® operating system on SGI Origin family servers. These features help you make maximum use of the systems. IRIX features such as checkpoint restart and weightless threads are discussed in this report.

2.0 Relative Performance on Different Computer Systems

In the following pages, the relative single-processor performance of several applications will be shown on a variety of test cases. The SGI computer systems being compared all belong to the SGI Origin family of servers (table 1).

Table 1. SGI Servers Used in This Report

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin® 2000 Series	MIPS RI4000	500	8
Origin® 3000 Series	MIPS RI2000	400	8
Origin 3000 Series	MIPS RI4000	500	8
Origin 3000 Series	MIPS RI4000	600	8
Origin® 300	MIPS RI4000	500	2

The results presented in this report highlight both the qualities of the applications as well as the excellent performance of these compute servers. It is noteworthy that in most instances SGI systems achieve a higher percentage of theoretical peak performance on actual production calculations than alternative hardware systems can achieve, the difference being particularly true for commodity-based hardware. The combination of efficient RISC microprocessors (MIPS) optimized for floating-point performance and a scalable memory architecture (SGI® NUMA) that minimizes latencies and maximizes bandwidth ensures that the SGI Origin family of servers is the best-balanced system in the industry. All of these system properties are very relevant for computational chemistry applications.

The virtues of these systems come especially to the forefront when single-processor performance of different SGI servers is compared and, using Amdahl's Law, it is demonstrated that the parallel scaling of the problems is in accordance with the amount of parallelism present in the application. The principles associated with Amdahl's Law are explained in a separate section.

In this report we contrast the performance of applications on different memory architectures all equipped with the same MIPS microprocessor. The results help us understand which problems are memory bound and which are less so. The continuous technical evolution of SGI NUMA architecture, from the SGI Origin 2000 series to the SGI Origin 3000 series, becomes evident from the results.

We look at the seemingly more CPU-bound problems by comparing the performance of MIPS microprocessors

running at different speeds in the same memory subsystem (the SGI Origin 3000 series). In theory, the processing speed of computers that have identical memory architectures and are equipped with the same type of microprocessors but running at different clock rates should improve according to the frequency ratios of the processors. In practice, however, the improvement is lower than this ratio, largely due to bottlenecks in the memory subsystem. Most computer systems experience this slowdown, but, thanks to SGI NUMA architecture, the improvements in CPU-bound problems in SGI Origin servers are indeed very close to the theoretical expectations. This indicates that no performance bottlenecks exist and that there is enough room for further microprocessor speedups as SGI introduces future MIPS family products.

Similar conclusions can be drawn from the parallel scaling experiments in this report. Since the parallel speedups follow specific Amdahl's Law curves, it can be inferred that no extraneous parallel performance anomalies are introduced by the hardware. This is certainly not the case in other, less well-balanced systems. Extreme examples of poor scalability performance have been seen in parallel machines based on clusters, where the fabric of the connection network often quickly becomes "the" bottleneck.

3.0 SGI Solutions for Life and Chemical Sciences

SGI has made a long-term commitment to and has been a leader in the life and chemical sciences community for more than 15 years. This commitment is expressed in a variety of ways, including by providing our customers superior computational platforms and supporting bioinformatics and computational chemistry applications, as well as high-performance and 3D visualization solutions for the life and chemical sciences. The SGI commitment is to provide our customers with the tools they need to ensure the most efficient utilization of research computing resources. Our goal is to enable our customers to make the best research decisions and to ultimately reduce their time to achieve research insights.

The SGI worldwide team of scientists and engineers is focused on helping customers solve their hard computational problems. This team works with worldwide commercial and academic software developers to advance and implement algorithms that take full advantage of the high-performance computational and graphical environments available on SGI hardware.

4.0 Amdahl's Law

Amdahl's Law defines the increase in speed that can be gained by parallelizing an application. It states that performance improvements due to parallelization are limited by the fraction of the code that is not running in parallel. For example, if 80% of the run time of an application can be run in parallel (top portion of the bars), then as more processors are added, the nonparallel portion or sequential part of the application (bottom portion of the bars) dominates the performance. When running this application with eight processors, the nonparallel or sequential (or serial) run time dominates the overall run time and very little, if any, speed increase can be achieved by using additional processors. Even if an infinite number of processors would be used to solve this problem, the parallel speedup would not exceed five.

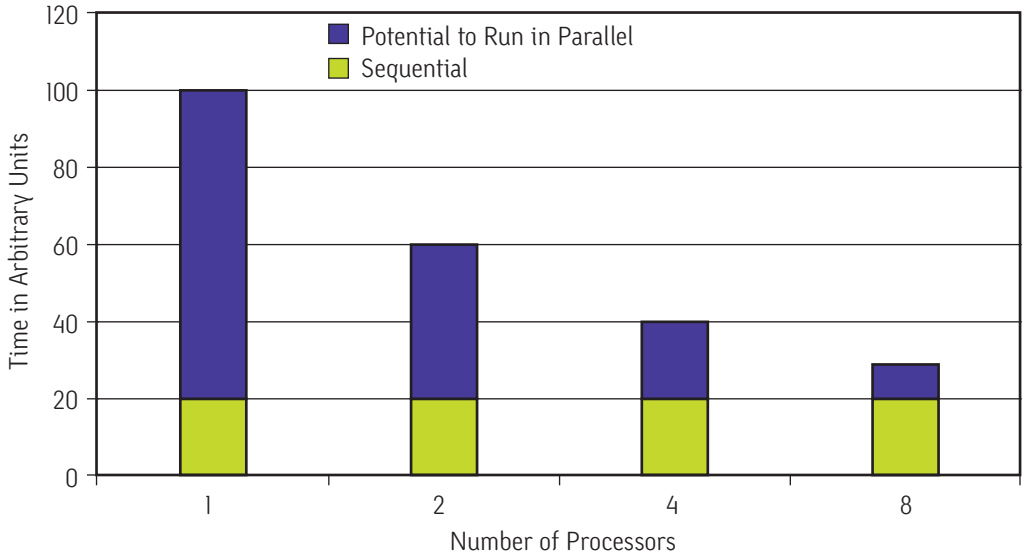

Fig. 1. Simple Amdahl's Law Scheme

4.1 Amdahl's Law Example

Amdahl's Law is of greatest importance in understanding and gauging parallel performance; thus, we include on each of the following graphs depicting parallel speedup the "best" theoretical performance curve predicted by Amdahl's Law. Large deviations from this curve could indicate performance anomalies. For example, in figure 2, almost 98% of the code could run in parallel. When 32 processors are used, the observed parallel speedup is close to 19. Furthermore, none of the measured parallel speedups show significant deviations from the curve. From this it can be inferred that no extraneous (hardware or other overheads) influences affect the parallel performance and that the parallel behavior is determined by the amount of parallelism in this particular path through the program.

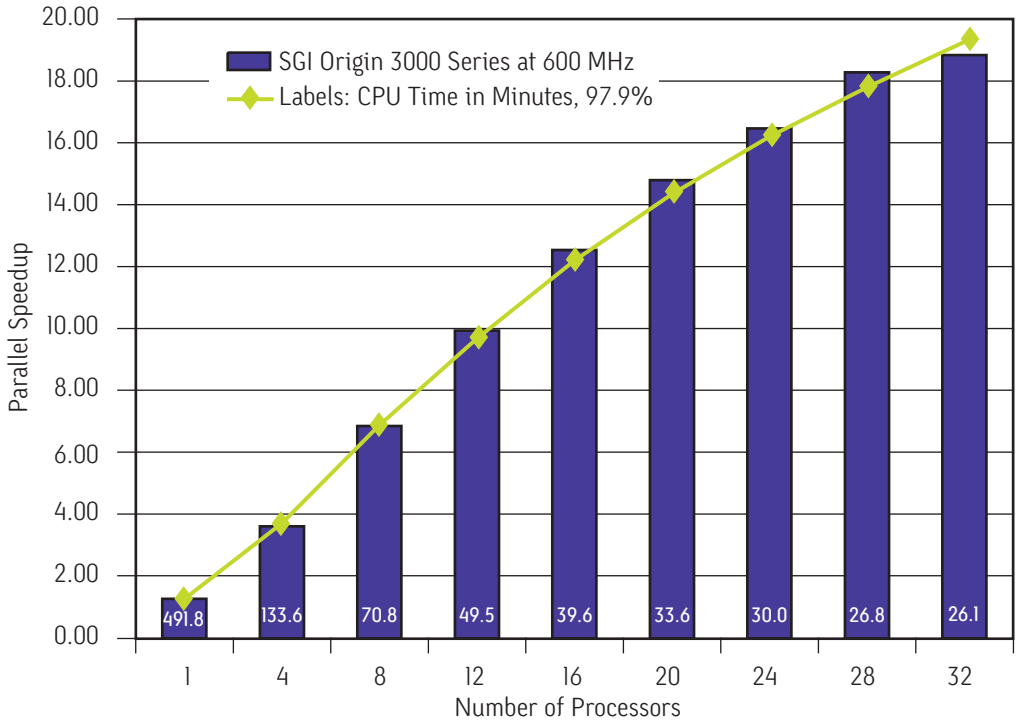

Fig. 2. Amdahl's Law Example

5.0 Quantum Chemistry and DFT Methods

For quantum chemistry calculations, the SGI family of servers offers outstanding scalability and performance. We draw examples from Gaussian® 98, GAMESS (U.S. version), Jaguar, NWChem, and MOPAC.

A number of different types of results are presented, including direct SCF, direct MP2, DFT, DFT with frequency, and Post Hartree-Fock methods.

An important feature of SGI servers is that, when using 64-bit versions of applications running on the 64-bit IRIX operating system, it is possible to address very large amounts of memory (tens of gigabytes) and to use very large files (hundreds of gigabytes).

6.0 Gaussian[1] 98 Revision A.11

Gaussian, Inc.'s Web site (*www.gaussian.com/g98broc.htm*) offers the following description of the program:

"Gaussian 98 is the latest in the Gaussian series of electronic structure programs. Designed to model a broad range of molecular systems under a variety of conditions, it performs its computations starting from the basic laws of quantum mechanics. Gaussian 98 is used by chemists, physicists and engineers for research in established and emerging areas of chemical interest, studying molecules and reactions of definite or potential interest, including both stable species and compounds which are difficult or impossible to observe experimentally: short-lived intermediates, transition structures and the like.

"Gaussian 98 can predict energies, molecular structures, vibrational frequencies—along with the numerous molecular properties that are derived from these three basic computation types—for systems in the gas phase and in solution, and it can model them in both their ground state and excited states. Chemists apply these fundamental results to their own investigations, using Gaussian 98 to explore chemical phenomena like substituent effects, reaction mechanisms, and electronic transitions."

Gaussian is used to perform fundamental research in chemical, pharmaceutical, and material sciences in academic and commercial settings.

6.1 Gaussian 98 Tests Description

The test cases described in table 2 serve to characterize the single-processor performance of the computer systems described below.

Table 2. Test Set for Gaussian 98

Test Case	Empirical Formula	Basis Set/Basis Functions	Type of Calculation
Test 178QA-suite	$C_6H_6N_6O_6$	6-31G**/300 basis functions	RHF, single point direct SCF calculation of TATB with extended population analysis
$C_3H_{11}O_4MP_2$	$C_3H_{11}O_4^+$	6-31G**/160 basis functions	Full direct MP2, C1 symmetry
$C_6H_6MP_4$	C_6H_6	6-31G(d)/102 basis functions	MP4(SDTQ) test for benzene, D6H symmetry
CISInc	$C_8H_8O_2$	6-31++G/154 basis functions	Direct CI singles in core force calculation, C1 symmetry, 255 MW
CIS	$C_8H_8O_2$	6-31++G/154 basis functions	Direct CI singles force calculation, C1 symmetry, 8 MW
Taxol	$C_{47}H_{51}NO_{14}$	3-21G/660 basis functions	HF single point calculation of taxol, C1 symmetry
$C_4H_{13}NO_4$	$C_4H_{13}NO_4$	6-31G** (5d)/191 basis functions	RHF frequency NoRaman calculation, C1 symmetry
$CH_6N_2MP_4$	CH_6N_2	6-311G(3d,3p)/156 basis functions	MP4 (SDTQ) test for methyl hydrazine, C1 symmetry
α-pinene	$C_{10}H_{16}$	6-31G/182 basis functions	RB3LYP frequency calculation
Test 397QA-suite	$C_{54}H_{90}N_6O_{18}$	3-21G/882 basis functions	RB3LYP force calculation, C1 symmetry

6.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of Gaussian 98 is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 3.

Table 3. Servers Used to Measure Gaussian 98 Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

6.2.1 SGI Origin 500 MHz systems

The cases shown in table 4 were run on MIPS R14000 processors running at 500 MHz installed on different memory fabrics, the SGI Origin 2000 series, SGI Origin 300, and the SGI Origin 3000 series. The differences in performance can be explained by the differences in memory architecture and secondary cache sizes.

The speedup comparison of the SGI Origin 2000 series with the SGI Origin 3000 series clearly shows the effects of improved memory architecture in the SGI Origin 3000 series system. These are most evident in the memory-bound problems like the Post-SCF calculations. In addition, the InCore CIS case stands out because of its use of a relatively large amount of memory and lack of cache blocking in the used algorithm.

The influence of the size of the secondary cache on performance is best seen in the column that compares the SGI Origin 3000 series with SGI Origin 300. This effect is less marked in the more memory-bound problems (MP4) and the InCore test. In these cases the improved memory architecture of SGI Origin 300, compared to the SGI Origin 2000 series, helps overcome the slowdown caused by a smaller secondary cache.

Table 4. Gaussian 98: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
test178	118.3	116.8	112.1	1.01	1.06	1.04
$C_3H_{11}O_4MP_2$	332.4	321.1	290.6	1.04	1.14	1.10
$C_6H_6MP_4$	356.3	374.8	340.3	0.95	1.05	1.10
CISInc	573.6	684.1	544.6	0.84	1.05	1.26
CIS	1,936.7	1,822.7	1,814.1	1.06	1.07	1.00
taxol	2,563.2	2,406.7	2,362.9	1.07	1.08	1.02
$C_4H_{13}NO_4$	3,727.8	3,365.3	3,276.0	1.11	1.14	1.03
$CH_6N_2MP_4$	3,731.0	4,323.7	3,660.0	0.86	1.02	1.18
α-pinene	12,626.4	11,767.7	11,624.6	1.07	1.09	1.01
test397	39,002.2	36,055.6	35,271.9	1.08	1.11	1.02
Average				**1.01**	**1.08**	**1.08**

6.2.2 SGI Origin 3000 Series Servers

Table 4 illustrates that Gaussian execution time not only depends on the processor speed but is also influenced by secondary cache size and memory architecture. Table 5 focuses more on how the processor speed affects the performance of the different types of Gaussian problems we have chosen.

Table 5 presents the execution CPU times of the test set on three SGI Origin 3000 series systems, each equipped with processors running at, respectively, 400, 500, and 600 MHz. The 400 MHz processors are MIPS R12000 chips, whereas the other two frequencies are found in MIPS R14000 microprocessors. In general, the measured speedups are close to what could be expected based on the frequency ratio of the different processors. The most noticeable exception is the CIS InCore case. Memory access rate clearly dominates the performance of this problem, and therefore it does not fully benefit from increased processor speed.

Table 5. Gaussian 98: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A™ 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
test178	131	112	94	1.17	1.39	1.19
$C_3H_{11}O_4MP_2$	342	291	247	1.18	1.38	1.17
$C_6H_6MP_4$	415	340	285	1.22	1.46	1.20
CISInc	606	545	492	1.11	1.23	1.11
CIS	2,125	1,814	1,513	1.17	1.40	1.20
taxol	2,740	2,363	1,975	1.16	1.39	1.20
$C_4H_{13}NO_4$	3,872	3,276	2,743	1.18	1.41	1.19
$CH_6N_2MP_4$	4,232	3,660	3,158	1.16	1.34	1.16
α-pinene	13,830	11,625	9,724	1.19	1.42	1.20
test397	40,656	35,272	29,298	1.15	1.39	1.20
Average				**1.17**	**1.38**	**1.18**

6.3 Gaussian 98 Single-Processor Performance—New Workstations

Recently SGI has introduced two powerful workstations that can be of interest to customers running Gaussian. These are described in table 6.

Table 6. Workstations Used to Measure Gaussian Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Silicon Graphics Fuel™	MIPS R14000	600	4
OEM IPF-Based System	Intel® Itanium™	733	4

The IPF-based system represents the first SGI OEM workstation based on the Intel 64-bit Itanium Processor Family (IPF). These workstations run a 64-bit Linux® operating system.

The performance of the same Gaussian test set (except the CIS InCore) is presented in table 7.

Table 7. Gaussian 98: Performance of Workstations. Silicon Graphics Fuel and IPF-Based System

Silicon Graphics Processor Frequency	Silicon Graphics Fuel R14000A 600 MHz Time in Seconds	OEM Itanium 733 MHz Time in Seconds
test178	95.1	79.6
c3h11o4mp2	263.9	189.3
c6h6mp4	288.3	308.3
CIS	1,553.5	1,204.2
taxol	2,044.4	1,709.0
c4h13no4	2,902.1	2,565.9
ch6n2mp4	3,124.5	2,069.6
a-pinene	10,706.5	8,316.8
test397	31,224.5	27,159.6

6.4 Gaussian 98 Parallel Performance

Two examples have been chosen to illustrate the parallel scaling that can be obtained on SGI Origin servers.

6.4.1 α-pinene ($C_{10}H_{16}$)

B3-LYP Frequency Calculation
basis: 6-31G(d)
182 basis functions

This moderate-sized test case exhibits good parallel speedups on an SGI Origin 300 server equipped with MIPS R14000 microprocessors running at 500 MHz. The acceleration with respect to one processor is over 14 when the test is ran on 28 processors, which is in accordance with over 96% of the path through the code being parallel. This means that the execution time is reduced from over two hours and a half on one processor to just 11 minutes on 28 processors. Members of the SGI Origin 3000 series exhibit slightly better parallel performance for this test case.

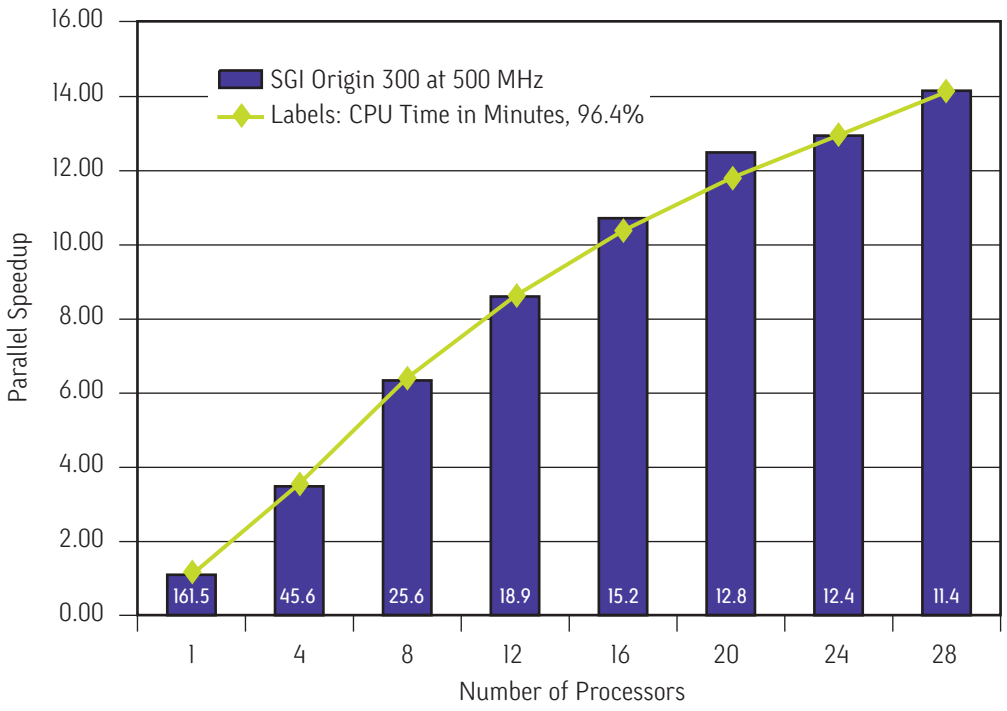

Fig. 3. Gaussian 98 Parallel Performance: α-pinene

6.4.2 Valinomycin ($C_{54}H_{90}N_6O_{18}$)

Test 397 from Gaussian's QA-suite:
B3-LYP Force Calculation
basis: 3-21G
882 basis functions

This test case, when run on an SGI Origin 3000 series system with MIPS R14000 600 MHz processors, showcases the benefits of a well-balanced multiprocessing system. A significant reduction of the execution time is achieved when 32 processors are used to solve this problem: from over eight hours on one processor to under half an hour on 32. This job now runs almost 19 times faster.

Similar parallel speedups can be expected from the other members of the SGI Origin 3000 series and the Origin 300 server.

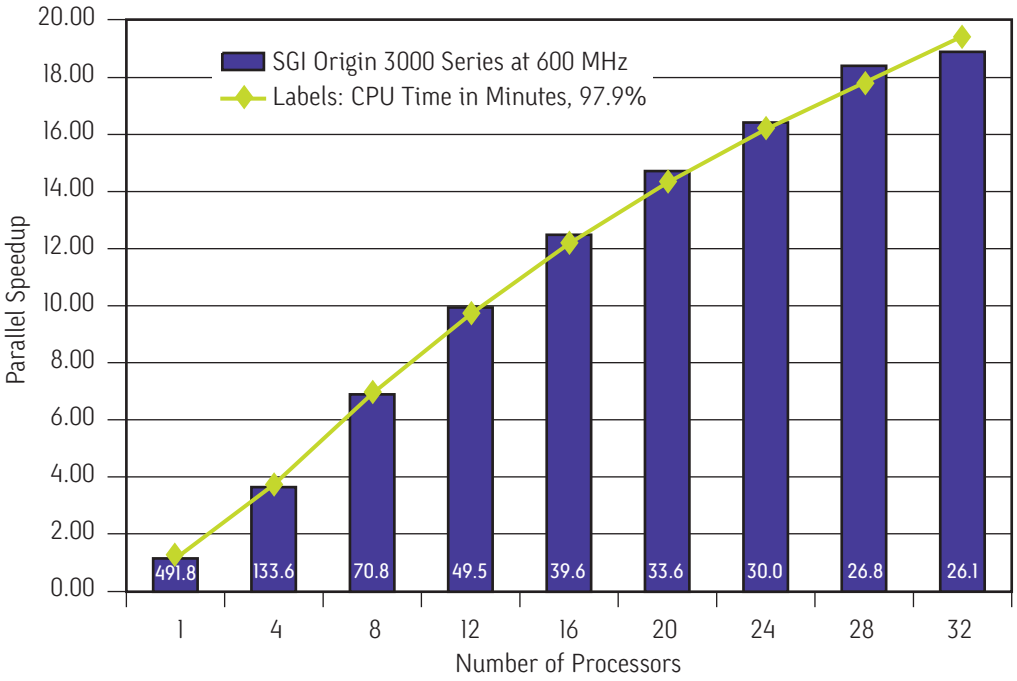

Fig. 4. Gaussian 98 Parallel Performance: Valinomycin

7.0 GAMESS[2] Version June 25, 2001

GAMESS's Web site [*www.msg.ameslab.gov/GAMESS/summary.html*] describes the program as follows:

"GAMESS [General Atomic and Molecular Electronic Structure System] is a program for ab-initio quantum chemistry. Briefly, GAMESS can compute wavefunctions ranging from RHF, ROHF, UHF, GVB, and MCSCF, with CI and MP2 energy corrections available for some of these. Analytic gradients are available for these SCF functions, for automatic geometry optimization, transition state searches, or reaction path following. Computation of the energy hessian permits prediction of vibrational frequencies. A variety of molecular properties, ranging from simple dipole moments to frequency dependent hyperpolarizabilities may be computed. Many basis sets are stored internally, and together with effective core potentials, all elements up to Radon may be included in molecules. Several graphics programs are available for viewing of the final results. Many of the computational functions can be performed using direct techniques, or in parallel on appropriate hardware."

GAMESS is used to perform fundamental research in chemical, pharmaceutical, and material sciences in academic and commercial settings.

For additional information about GAMESS running on SGI computer systems, visit *www.sgi.com/solutions/sciences/chembio/resources/gamess_us/*.

7.1 GAMESS Tests Description

The tests are used to characterize GAMESS performance are described in table 8.

Table 8. Test Set for GAMESS

Test Case	Empirical Formula	Basis Set/Basis Functions	Type of Calculation
Crown_HF	C6H12O3	3-21G/210 basis functions	Direct SCF of crown ether, CI symmetry
Crown_MP2	C6H12O3	3-21G/210 basis functions	Direct MP2 of crown ether, CI symmetry

The input files were supplied by Pacific Northwest National Lab [*www.emsl.pnl.gov:2080/docs/tms/abinitio/cover.html*].

7.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of GAMESS is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI

Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 9.

Table 9. Servers Used to Measure GAMESS Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

7.2.1 SGI Origin 500 MHz Systems

Table 10 shows that at the Hartree-Fock level, the memory architecture and secondary cache sizes have only a small influence on the execution times. Different machines equipped with the same microprocessors run at virtually the same speed. The MP2 example shows that some small gains are possible when the superior architecture of the SGI Origin 3000 series is used. On the other hand, the effects of cache size and memory architecture seem to cancel each other out as evidenced in the column that compares the SGI Origin 2000 series with SGI Origin 300.

Table 10. GAMESS: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
Crown_HF	234	229	232	1.02	1.01	0.99
Crown_MP2	816	809	765	1.01	1.07	1.06
Average				**1.02**	**1.04**	**1.02**

7.2.2 SGI Origin 3000 Series Servers

The table below validates the previous conclusions. Only the MP2 case shows some small effects of secondary cache size and memory architecture. Otherwise, the performance improvements follow very closely the frequency ratio of the MIPS processors.

Table 11. GAMESS: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
Crown_HF	285	232	192	1.23	1.48	1.21
Crown_MP2	903	765	646	1.18	1.40	1.18
Average				**1.20**	**1.44**	**1.20**

7.3 GAMESS: Parallel Performance

7.3.1 Crown_MP2 ($C_6H_{12}O_3$)

Even for this small case, the SGI Origin 3000 series server with MIPS R14000 processors running at 600 MHz provides a reasonable parallel scaling. According to Amdahl's Law, around 98.7% of the execution path is running in parallel. This translates into a speedup of over 13 on 16 processors.

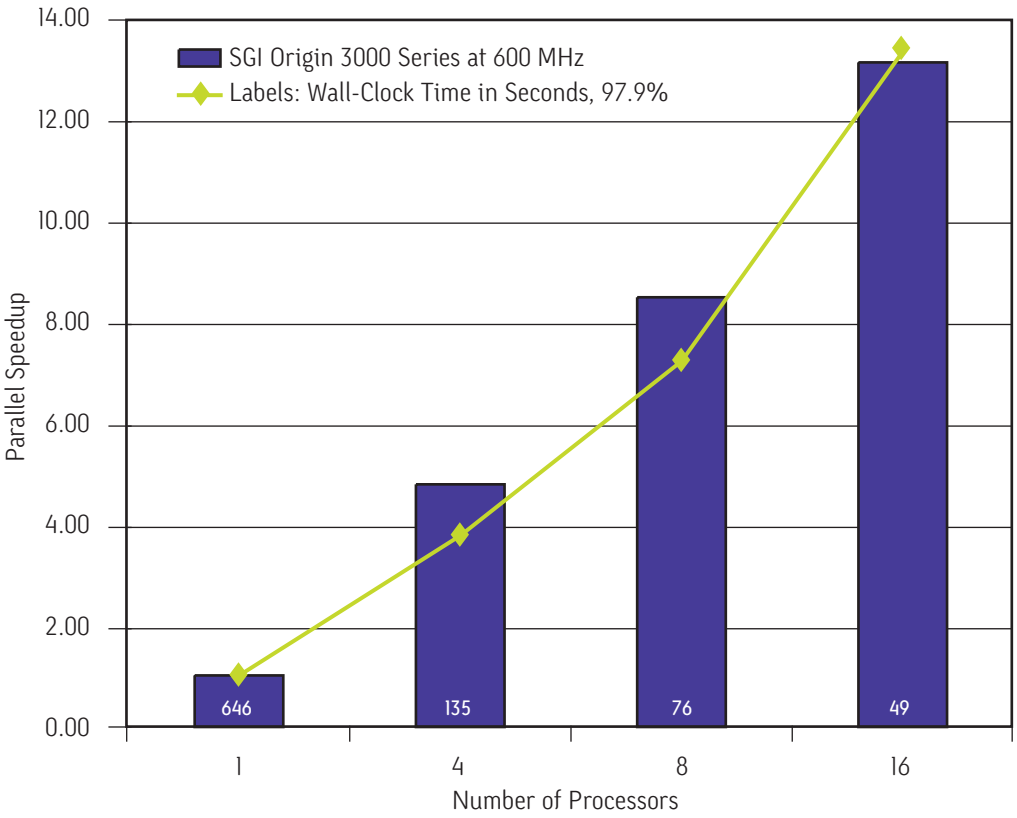

Fig. 5. GAMESS Parallel Performance: Crown_MP2

8.0 Jaguar[3] 4.2

Schrödinger's Web site (*www.schrodinger.com*) describes Jaguar as follows:

"Jaguar was designed to increase the speed of ab-initio calculations in order to accelerate basic and applied research projects and to enable calculations at a higher level of theory. Jaguar's speed and power make it possible to study larger systems than ever before, or to study many more systems than previously possible, within a reasonable time frame."

Jaguar is used to perform fundamental research in chemical, pharmaceutical, and material sciences in academic and commercial settings.

8.1 Jaguar Tests Description

The test set MJAG99, provided by Schrödinger, Inc., consists of six calculations using different methods and no symmetry. In addition, a larger LMP2 (Class4a.lmp2) calculation is included. This test serves to highlight Jaguar's parallel performance and is described in table 12.

Table 12. Test Set for Jaguar

Test Case	Empirical Formula	Basis Set/Basis Functions	Type of Calculation
Blypeg	$C_{27}OH_{46}$	6-31G** basis set/650 basis functions	BLYP 1 step of geometry optimization
Blypfreq	$C_5H_8O_2$	6-31G** basis set/145 basis functions	BLYP vibrational frequencies
Ccpvtz	$C_{12}H_{14}O_3F_2$	cc-pvtz(-f) basis set/517 basis functions	B3LYP 1 step of geometry optimization
Hfeg	$C_{27}OH_{46}$	6-31G** basis set/650 basis functions	HF 1 step of geometry optimization
Hffreq	$C_5H_8O_2$	6-31G** basis set/145 basis functions	HF vibrational frequencies
Lmp2	$C_{12}H_{14}O_3F_2$	6-31G** basis set/325 basis functions	LMP2 energy
Class4a.lmp2	$C_{15}H_{20}O_3F_2$	6-31G** basis set/400 basis functions	LMP2 energy, no symmetry

8.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of Jaguar is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 13.

Table 13. Servers Used to Measure Jaguar Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

8.2.1 SGI Origin 500 MHz Systems

The results in table 14 suggest that secondary cache size does affect the frequency calculations as both the SGI Origin 3000 series and SGI Origin 2000 series show significant speedups over SGI Origin 300. The improvement in memory architecture of the SGI Origin 3000 series over the SGI Origin 2000 series positively influences the performance of the post-SCF LMP2 calculation.

Table 14. Jaguar: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
Blypeg	2,143	1,949	1,903	1.10	1.13	1.02
Blypfreq	1,838	1,590	1,555	1.16	1.18	1.02
Ccpvtz	3,440	3,116	3,061	1.10	1.12	1.02
Hfeg	2,752	2,517	2,451	1.09	1.12	1.03
Hffreq	1,670	1,427	1,354	1.17	1.23	1.05
Lmp2	1,718	1,751	1,549	0.98	1.11	1.13
Class4a.lmp2	3,622	3,568	3,149	1.02	1.15	1.13
Average				**1.09**	**1.15**	**1.06**

8.2.2 SGI Origin 3000 Series Servers

The relative speedups of various MIPS microprocessors on the same SGI Origin 3000 series servers are uniform. Only the post-SCF LMP2 tests are further away from the frequency ratio, probably due to a larger dependency on memory bandwidth than the rest of the cases shown in table 15.

Table 15. Jaguar: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
Blypeg	2,185	1,903	1,591	1.15	1.37	1.20
Blypfreq	1,787	1,555	1,287	1.15	1.39	1.21
Ccpvtz	3,579	3,061	2,566	1.17	1.39	1.19
Hfeg	2,842	2,451	2,056	1.16	1.38	1.19
Hffreq	1,527	1,354	1,134	1.13	1.35	1.19
Lmp2	1,585	1,549	1,320	1.02	1.20	1.17
Class4a.lmp2	3,573	3,149	2,707	1.13	1.32	1.16
Average				**1.13**	**1.34**	**1.19**

8.3 Jaguar 4.2 Parallel Performance

8.3.1 Class4a.lmp2 ($C_{15}H_{20}O_3F_2$)

SCF: 6-31G** No Symmetry 400 Basis functions 40 Atoms.

This medium-sized SCF calculation, performed with Jaguar on an SGI Origin 3000 series system with MIPS R14000 600 MHz processors, shows an Amdahl's Law parallelization percentage of over 95% corresponding to a parallel speedup of more than 9 on 16 processors.

The parallel scaling of this test case is similar to that of the other members of the SGI Origin server family, including SGI Origin 300.

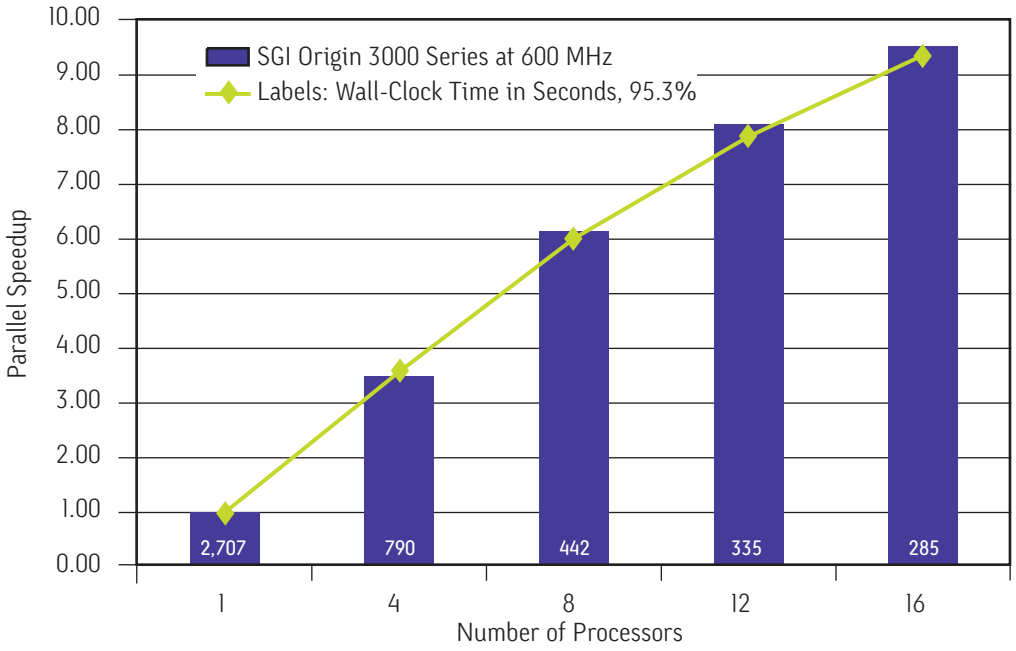

Fig. 6. Jaguar Parallel Performance: Class4a.lmp2

9.0 NWChem[4] 4.1

NWChem's homepage (*www.emsl.pnl.gov:2080/docs/nwchem/nwchem.html*) describes the program as follows:

"NWChem is a computational chemistry package that is designed to run on high-performance parallel supercomputers as well as conventional workstation clusters. It aims to be scalable both in its ability to treat large problems efficiently and in its usage of available parallel computing resources. NWChem has been developed by the Molecular Sciences Software group of the Theory, Modeling & Simulation program of the Environmental Molecular Sciences Laboratory (EMSL)

at the Pacific Northwest National Laboratory (PNNL). Most of the implementation has been funded by the EMSL Construction Project.

"The suite utilizes parallel-programming tools developed by PNNL staff. Most of the tool and scalable algorithm development has been funded by the High Performance Computing and Communications Initiative (HPCCI) grand-challenge software program and the DOE-2000 ACTS Tools project."

NWChem is used to perform fundamental research in chemical, pharmaceutical, and material sciences in academic and commercial settings.

9.1 NWChem Tests Description

The following tests demonstrate the performance of NWChem LDA DFT calculations of two zeolite fragments (347 and 1,687 basis functions) on SGI Origin 300 and Origin 3000 series systems, all of them with 4GB of main memory per 4-CPU module.

The input files were retrieved from *www.emsl.pnl.gov/pub/docs/nwchem/benchmarks/dft/siosi3.nw* and *www.emsl.pnl.gov/pub/docs/nwchem/benchmarks/dft/siosi6.nw* and were modified only in the "memory" line to ensure that all three-center two-electron integrals were held in-core.

See table 16 for a more detailed test description.

Table 16. Test Set for NWChem

Test Case		SiOSi3	SiOSi6
Empirical Formula		$Si_8H_{18}O_7$	$Si_{28}H_{30}O_{67}$
AO Basis	Number of Functions	347	1,687
	Number of Shells	160	716
CD Basis	Number of Functions	832	3,928
	Number of Shells	335	1,527
Convergence on Energy		0.10D-05	0.10D-05
Schwarz Screening/ accCoul		0.10D-09	0.10D-09
Number of Grid Points		39,8567	134,7516

9.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of NWChem is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 17.

Table 17. Servers Used to Measure NWChem Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

9.3 NWChem Single-Processor Performance

All timings are wall-clock seconds as reported by NWChem. The same binary was used for all runs on the various computer systems tested. Results are consistently better than those previously published for NWChem 4.0.1 on the NWChem homepage (*www.emsl.pnl.gov:2080/docs/nwchem/*).

9.3.1 SGI Origin 500 MHz Systems

For small test cases like SiOSi_3, the size of the secondary cache seems to have little effect (5%) on the performance of NWChem. The situation is somewhat different for the larger case where the SGI Origin 3000 series system, with its larger cache, is almost 10% faster than the SGI Origin 300 server.

Table 18. NWChem: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 3000 Series over Origin 300
SiOSi_3	1,091.3	1,037.8	1.05
SiOSi_6	38,435.3	35,127.0	1.09
Average			**1.07**

9.3.2 SGI Origin 3000 Series Servers

Like table 18, table 19 leads to the conclusion that for small test cases, NWChem performance is not bound by the memory subsystem. The SGI NUMA architecture of the SGI Origin 3000 series servers permits a perfect scaling with the frequency of the microprocessor. The larger test case, which requires in excess of 16GB of memory, shows the effect of having to address non-local memory. The excellent combination of low latency and high bandwidth in the SGI Origin 3000 series

results in only a small deviation from the theoretical speedup.

Table 19. NWChem: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 600 over 500 MHz
SiOSi_3	1,037.8	866.1	1.20
SiOSi_6	35,127.0	30,311.7	1.16
Average			**1.18**

9.4 NWChem: Parallel Performance

9.4.1 SiOSi_3 ($Si_8H_{18}O_7$)

The SGI Origin 300 500 MHz server shows a good parallel speedup for this test case. Based on Amdahl's Law, the path through the program is parallel for 97.8%. This results in an acceleration of close to 17 when running on 24 processors, compared with the single-processor execution time.

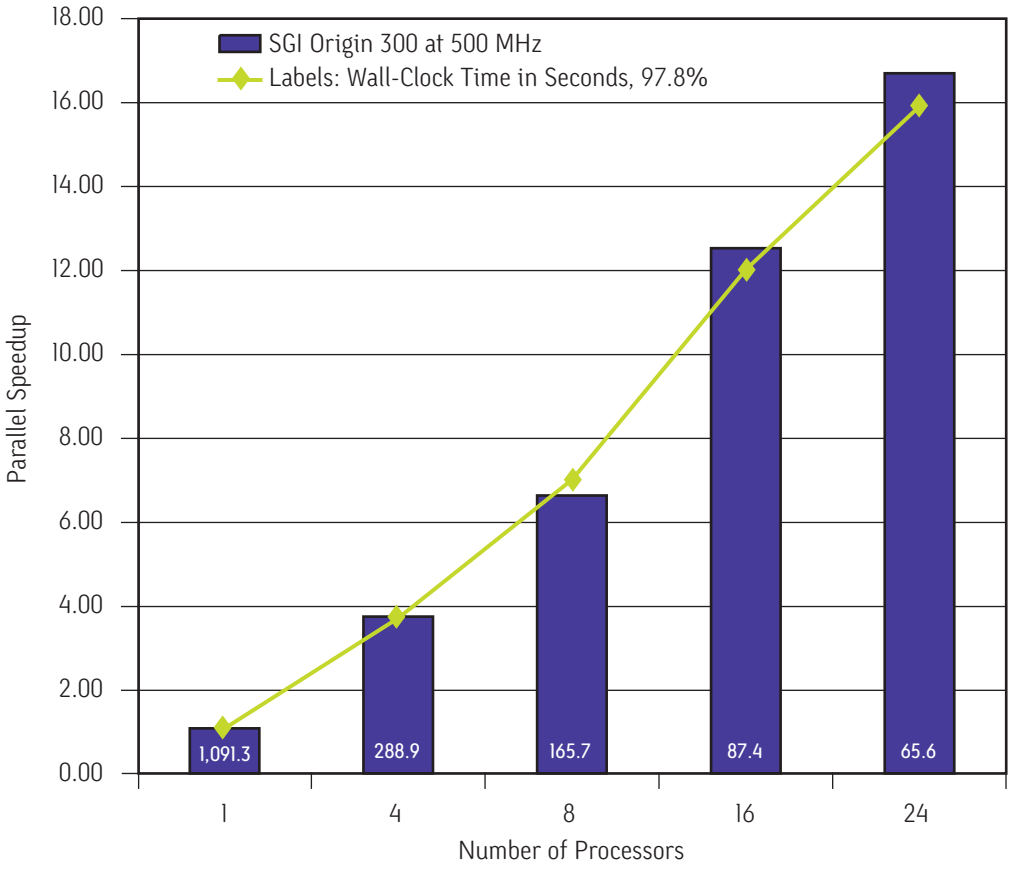

Fig. 7. NWChem Parallel Performance: SiOSi_3

9.4.2 SiOSi_6 ($Si_{28}H_{30}O_{67}$)

This test case shows an exceptional parallel performance as can be seen in figure 8 for the SGI Origin 3000 series 600 MHz server. The code is over 99% parallel, which translates into a speedup of more than 50 on 88 processors, compared with the single-processor execution time. This job takes close to eight hours and 25 minutes of wall-clock time to complete on one processor and around 10 minutes on 88 processors.

Other members of the SGI Origin 3000 series and the SGI Origin 300 server perform similarly.

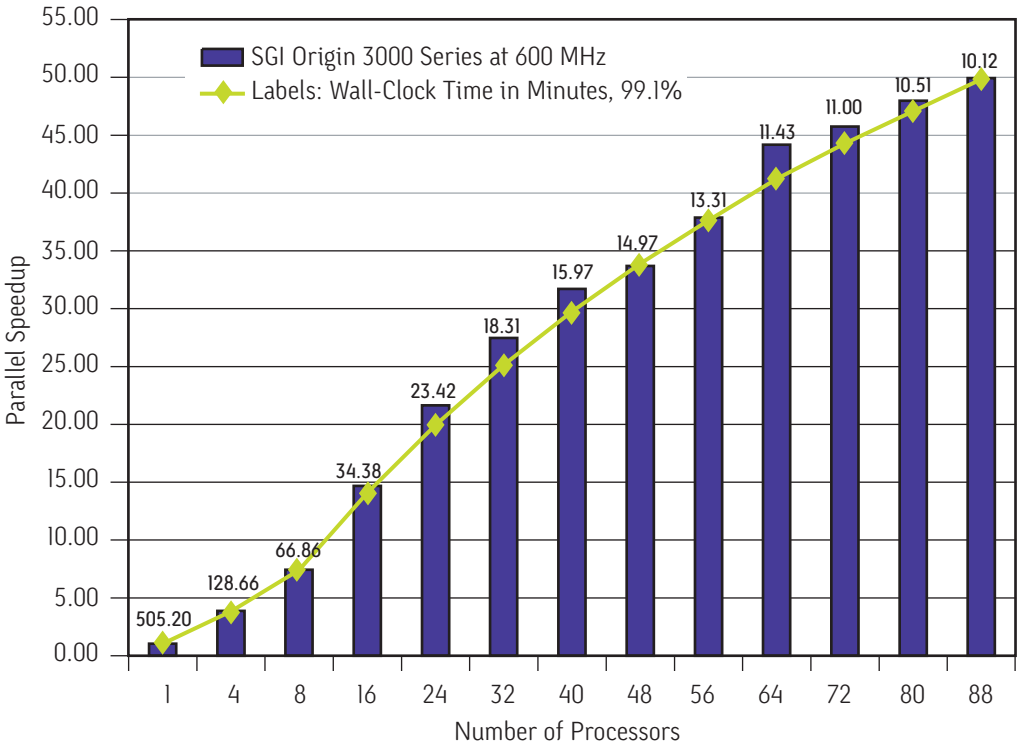

Fig. 8. NWChem Parallel Performance: SiOSi_6

10.0 MOPAC[5] 2002 1.0

Both Fujitsu's CAChe Group [*www.cachesoftware.com/mopac/index.shtml*] and Schrödinger, Inc. [*www.schrodinger.com/Products/mopac.html*] describe MOPAC as follows:

"MOPAC is a general-purpose semiempirical quantum mechanics package for the study of chemical properties and reactions in gas, solution, or solid-state. MOPAC is used to directly predict numerous chemical and physical properties, such as Gibbs free energies, activation energies, reaction paths, dipole moments, non-linear optical properties, and infrared spectra. It is also used as the basis of quantitative structure-property [or activity] relationships, to predict a wide variety of biological and other properties including carcinogenicity, vapor pressure, water solubility, and reaction rates."

MOPAC is used to perform fundamental research in chemical, pharmaceutical, and material sciences in academic and commercial settings.

10.1 MOPAC Tests Description

Fujitsu, Inc., supplied the test set described in table 20.

Table 20. Test Set for MOPAC

Test Case	Empirical Formula	Total No. Atoms	Type of Calculation	Molecule
bRhodopsin	$C_{1216}H_{1877}N_{273}O_{311}S_9$	3,686	MOZYME RESIDUES ISCF PL CUTOF2=6 XENO=(20,0,0,0,RETINAL) MEMORY=260	Bacteriorhodopsin
barnase	$C_{550}H_{841}N_{151}O_{168}$	1710	MOZYME RESIDUES ISCF PL RELSCF=10 CHARGE=2 MEMORY=170	Barnase
Crambin	$C_{202}H_{314}N_{55}O_{64}S_6$	641	MOZYME RESIDUES ISCF PL CHARGE=-1	Crambin from PDB
isomerase	$C_{1196}H_{1893}N_{320}O_{365}S_2$	3,776	MOZYME RESIDUES ISCF PL CUTOF1=9 CUTOF2=6 SCFCRT=0.2 MEMORY=310	Triose Phosphate Isomerase
lipase	$C_{1306}H_{1991}N_{337}O_{406}S_8$	2,698	MOZYME RESIDUES ISCF PL CUTOF1=8 CUTOF2=6 GEO-OK CHARGE=0 MEMORY=240	Rhizomucor Mehei Lipase
papain	$C_{1065}H_{1613}N_{295}O_{310}S_7$	3,290	MOZYME ISCF PL SCFCRT=0.1 SYMMETRY NLMO=400 MEMORY=500	Papain
C960	C_{960}	960	MOZYME ISCF PL SCFCRT=0.1 SYMMETRY NLMO=400 MEMORY=500	C960

10.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of MOPAC is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 21.

Table 21. Servers Used to Measure MOPAC Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

All timings are wall-clock seconds as reported by MOPAC. The same binary was used for all runs on the various computer systems tested.

10.2.1 SGI Origin 500 MHz Systems

The results from table 22 indicate that MOPAC's execution time is somewhat dependent on the memory architecture. The 30% speedup of the SGI Origin 3000 series over SGI Origin 300 for the c960 case is quite remarkable.

Table 22. MOPAC: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 3000 Series over Origin 300
bRhodopsin	1,357.1	1,205.0	1.13
barnase	569.4	509.1	1.12
crambin	135.9	117.6	1.16
isomerase	3,577.2	3,109.1	1.15
lipase	1,553.5	1,359.2	1.14
papain	1,174.8	1,042.4	1.13
c960	3,224.0	2,480.7	1.30
Average			**1.16**

10.2.2 SGI Origin 3000 Series Servers

Within the same memory architecture MOPAC's speed scales very closely to the frequency ratio, once again demonstrating the excellent characteristics of SGI NUMA architecture.

Table 23. MOPAC: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
bRhodopsin	1,447.0	1,205.0	1,009.3	1.20	1.43	1.19
barnase	571.0	509.1	408.0	1.12	1.40	1.25
crambin	139.3	117.6	98.4	1.18	1.42	1.20
isomerase	3,710.6	3,109.1	2,628.6	1.19	1.41	1.18
lipase	1,615.9	1,359.2	1,143.4	1.19	1.41	1.19
papain	1,240.9	1,042.4	875.2	1.19	1.42	1.19
c960	2,856.9	2,480.7	2,230.5	1.15	1.28	1.11
Average				**1.18**	**1.40**	**1.19**

10.3 MOPAC Parallel Performance

10.3.1 C960 (C_{960})

The 500 MHz SGI Origin 300 server shows a modest parallel speedup for this test case. Based on Amdahl's Law, the path through the program is parallel for 91.7%. This results in an acceleration of over 7 when running on 16 processors compared to the single-processor execution time.

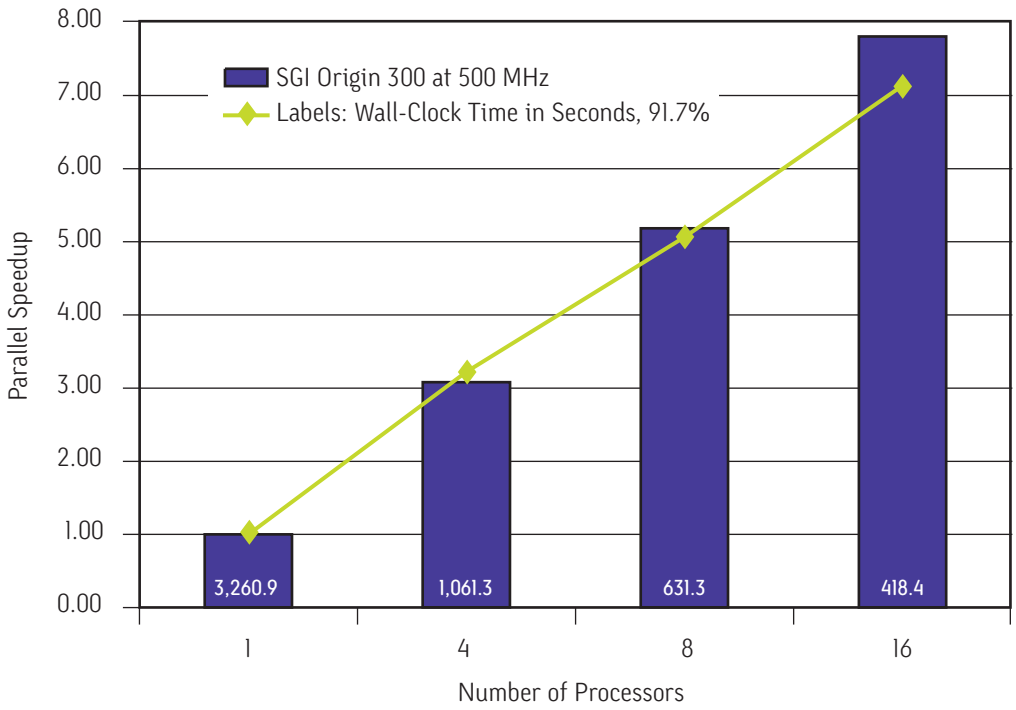

Fig. 9. MOPAC Parallel Performance: C960

11.0 Force-Field-Based Calculations

In the areas of molecular mechanics and molecular dynamics, the range of performance for single-processor and multiprocessor SGI servers is demonstrated for Amber, CHARMM, and CNX. The outstanding scaling efficiency of the SGI Origin family is showcased by the new versions of Amber and CHARMM. The same versions of the applications were run sequentially and in parallel on the SGI Origin 2000 series, SGI Origin 3000 series, and SGI Origin 300 servers using shared-memory parallel directives or message-passing library calls.

12.0 Amber[6] 7

Amber's homepage (*www.amber.ucsf.edu/amber/amber.html*) describes the program as follows:

"Amber refers to two things: a molecular mechanical force field for the simulation of biomolecules (which is in general use in a variety of simulation programs); and a package of molecular simulation programs which includes source code and demos. The current version of this package is AMBER version 7, which is sold by UCSF subject to a licensing agreement as described below. The code is written in Fortran and C and requires approximately 65MB of disk space.

Amber is developed in an active collaboration of the Kollman group at UCSF, Dave Case at The Scripps Research Institute, Ken Merz at Penn State, Thomas Cheatham at the University of Utah, Carlos Simmerling at SUNY-Stony Brook, Tom Darden at NIEHS, and Dave Pearlman at Vertex Pharmaceuticals."

Amber is used to perform fundamental research in the area of life sciences in academic and commercial environments.

12.1 Amber 7 Tests Description

In the next pages, results are shown from a set of timing tests published by the Amber Group at UCSF. The tests can be characterized as shown in table 24.

The Amber development group has recently posted a summary of results for some of these benchmarks. SGI Origin servers and other computer systems are included. Visit *www.amber.ucsf.edu/amber/amber7.bench4.html*.

Table 24. Test Set for Amber

Test Case	Number of Atoms	Type of Calculation
Dhfr	22,930	Dihydrofolate reductase in explicit solvent, pme simulation: 100 steps of MD
Gb_mb	2,492	100 steps of MDGeneralized Born myoglobin simulation. This protein has 2492 atoms and is run with a 20 Ang. cutoff and a salt concentration of 0.2 M, with nrespa=4 (long-range forces computed every 4 steps).
Jac	23,558	"Joint Amber/CHARMM" dhfr benchmark. Note: even though this is also on the protein dhfr, it is *not* the same benchmark as the dhfr one listed above. Parameters here are set to allow as close as possible a comparison of CHARMM and Amber (and other programs such as NAMD). 100 steps of MD. This is the protein DHFR, solvated with TIP3 water, in a periodic box. There are 23,558 total atoms, and PME used with a direct space cutoff of 9 Ang.
Rt_polymerase	141,154	Large-ish protein in water: 100 steps of MD

12.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of Amber is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin 3000 series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers support microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 25.

Table 25. Servers Used to Measure Amber Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

12.2.1 SGI Origin 500 MHz Systems

In general, neither differences in the memory fabric of the system nor the secondary cache sizes of the microprocessors seem to have a significant effect on the single-processor performance of Amber 7.

Table 26. Amber: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
dhfr	118.86	113.31	110.70	1.05	1.07	1.02
gb_mb	143.25	141.87	141.61	1.01	1.01	1.00
jac	149.67	143.72	140.80	1.04	1.06	1.02
rt_polymerase	849.21	853.45	816.52	1.00	1.04	1.05
Average				**1.02**	**1.05**	**1.02**

12.2.2 SGI Origin 3000 Series Servers

A comparison of the relative performance of different microprocessors in the same computer system validates the conclusion drawn from the previous table. Owing to the outstanding latency properties of SGI NUMA memory architecture, the improvements that are measured on faster chips are very close to the frequency ratio of the microprocessors.

Table 26. Amber: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
dhfr	131.97	110.70	92.62	1.19	1.42	1.20
gb_mb	173.55	141.61	118.32	1.23	1.47	1.20
jac	165.16	140.80	116.76	1.17	1.41	1.21
rt_polymerase	965.79	816.52	686.16	1.18	1.41	1.19
Average				**1.19**	**1.43**	**1.20**

12.3 Amber 7 Parallel Performance

12.3.1 RT_polymerase

The 500 MHz SGI Origin 300 server shows a good parallel speedup for this test case. Based on Amdahl's Law, the path through the program is parallel for 98.5%. This results in an acceleration of the "non-setup" portion of Amber of around seven when running on eight processors compared with the single-processor execution time.

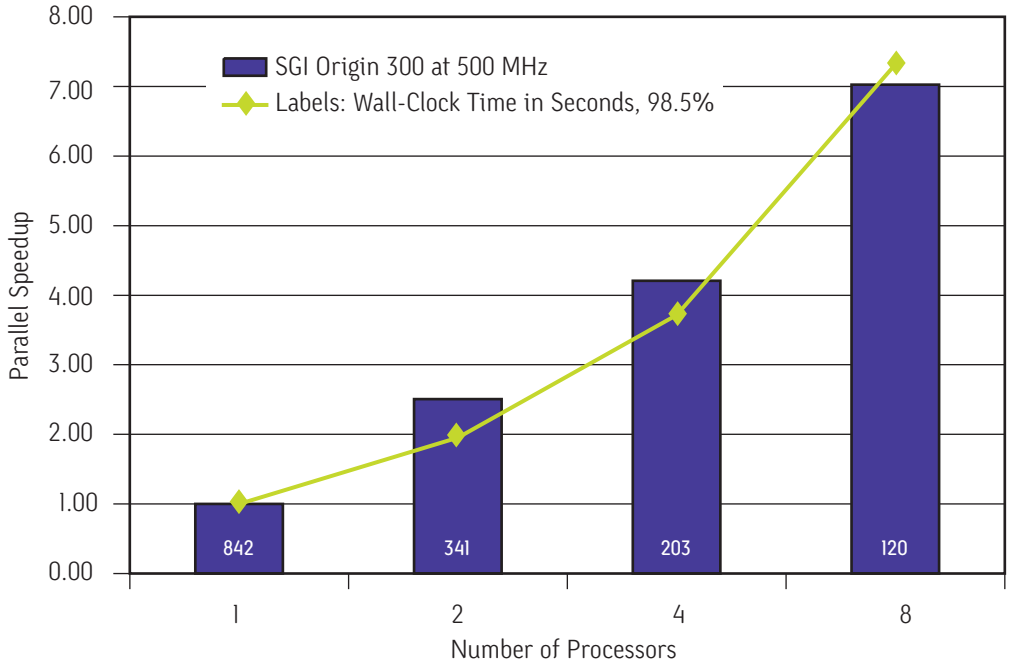

Fig. 10. Amber Parallel Performance: RT_polymerase

12.3.2 GB_MB

This test case shows an even better parallel performance as can be seen in the next graph for the SGI Origin 3000 series 600 MHz server. The code is for 99% parallel, which translates into a speedup of close to 24 on 32 processors, compared with the single-processor execution time. Other members of the SGI Origin 3000 series and SGI Origin 300 perform similarly.

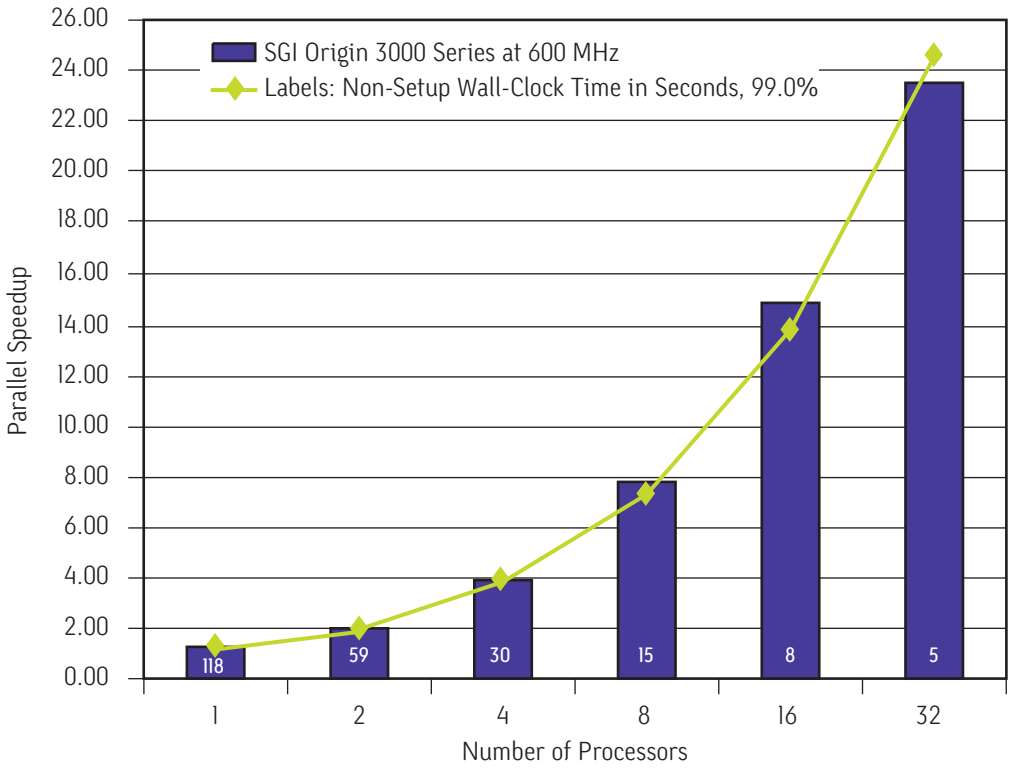

Fig. 11. Amber Parallel Performance: GB_MB

13.0 CHARMM[7] 28.b1

CHARMM's Web site at the Scripps Research Institute (*http://mmtsb.scripps.edu/software/charmm.html*) describes the program as follows:

"CHARMM (Chemistry at HARvard Macromolecular Mechanics) is a program for macromolecular dynamics and mechanics. It performs standard molecular dynamics in many different ensembles (e.g., NVE, NVT, NPT) using state-of-the-art algorithms for time-stepping, long-range force calculation and periodic images. CHARMM can be used for energy minimization, normal modes and crystal optimizations as well. The potential energy functions available for use with CHARMM have been extensively parameterized for simulations of proteins, nucleic acids and lipids. Free energy methods for chemical and conformational free energy calculations are also fully developed and available in CHARMM. Many other novel tools have been developed and are available in CHARMM, these include: replicas (multiple copies); many types of restraints and constraints, including fixed atoms, atomic, NOE, dihedral and internal coordinate restraints and generalized SHAKE for bonds and arbitrary internal coordinates; minimum energy path following and transition state optimization; etc."

CHARMm (the commercial version of CHARMM) is available to scientists in industry and academic researchers who wish to have full user support, through Accelrys.

CHARMm and CHARMM are used to perform fundamental research in the area of life sciences in academic and commercial environments.

13.1 CHARMM Tests Description

The test cases in table 28 were used to characterize the single-processor performance of the computer systems described above.

Table 28. Test Set for CHARMM

Test Case	Number of Atoms	Type of Calculation
Myoglobin	14,032	MbCO, with 3,830 water molecules. 1,000 steps of MD.
Cubic_dhfr	23,558	Dihydrofolate reductase in a cubic box of solvent using PBC with PME. There are 7,023 water molecules. 1,000 steps of MD.

13.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of CHARMM is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performance of the SGI Origin 3000 series of servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 29.

Table 29. Servers Used To Measure CHARMM Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

13.2.1 SGI Origin 500 MHz Systems

In general, CHARMM performance is not affected significantly by the underlying memory architecture or by the size of the secondary caches of the microprocessors.

13.2.2 SGI Origin 3000 Series Servers

Within the SGI Origin 3000 series of servers the speedups we measured are very close to the ratio of the speeds of the different MIPS processors. This is achieved thanks to the outstanding latency properties of the SGI NUMA memory system.

Table 30. CHARMM: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
Myoglobin	1,640.60	1,648.90	1,627.70	0.99	1.01	1.01
Cubic_dhfr	2,317.40	2,296.50	2,210.30	1.01	1.05	1.04
Average				**1.00**	**1.03**	**1.03**

Table 31. CHARMM: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
Myoglobin	1,999.40	1,627.70	1,363.00	1.23	1.47	1.19
Cubic_dhfr	2,669.20	2,210.30	1,854.00	1.21	1.44	1.19
Average				**1.22**	**1.45**	**1.19**

13.3 CHARMM 28.b1 Parallel Performance

13.3.1 DHFR Cubic PME

The SGI Origin 300 500 MHz server performance scales following the 95.8% Amdahl's Law curve. This translates into a parallel speedup of more than 10 when running on 16 processors. SGI Origin 3000 series servers show slightly better scaling.

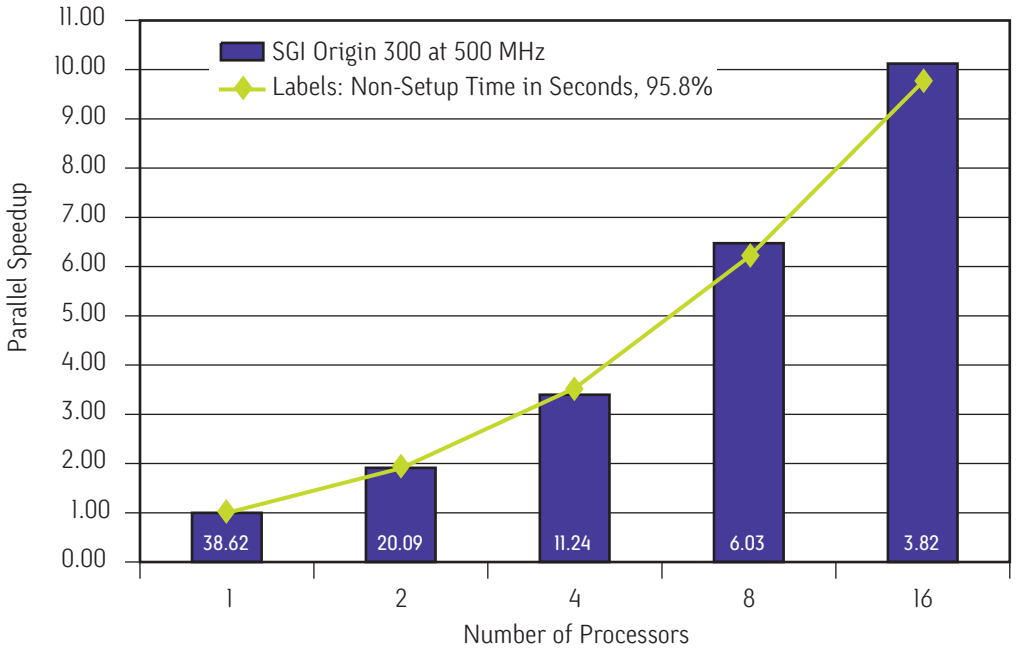

Fig. 12. CHARMM Parallel Performance: DHFR Cubic PME

13.3.2 Myoglobin

SGI Origin 3000 series 600 MHz server
MbCO] with 3,830 Water Molecules (14,026 atoms).
1,000 Steps of Dynamics

The effects of improved memory architecture of the SGI Origin 3000 series come to light when running in parallel. In this test case, the parallel scaling follows the 98.8% Amdahl's Law curve. Consequently, a 22-fold speedup over single-processor execution time is seen when the job is run on 32 processors.

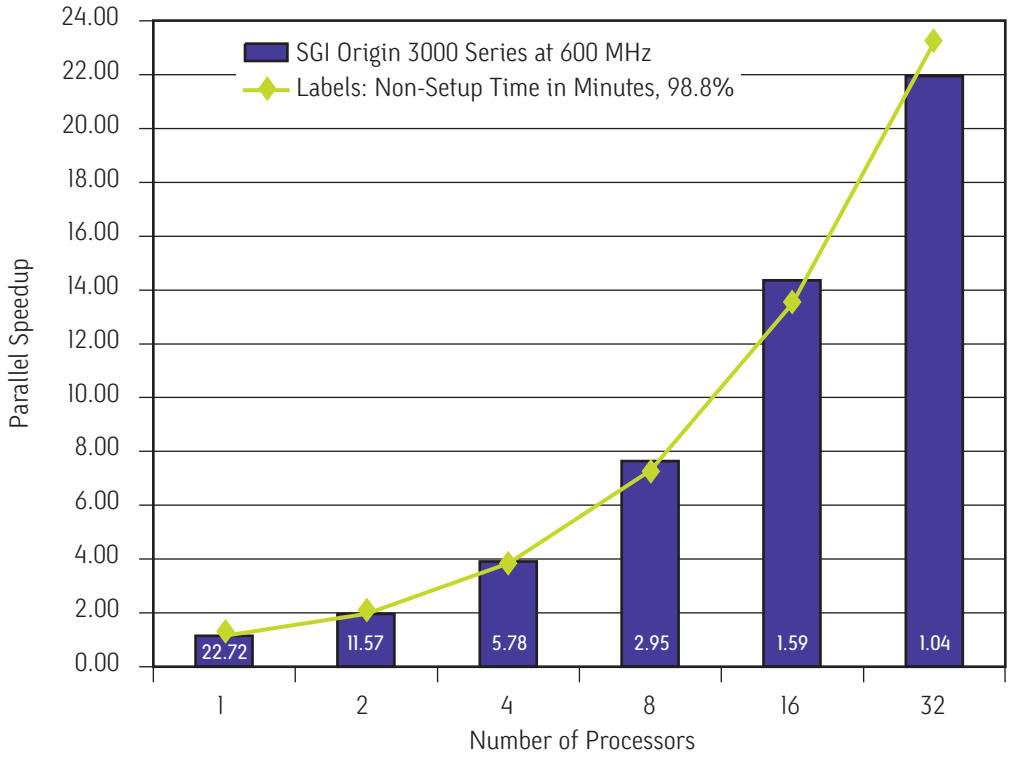

Fig. 13. CHARMM Parallel Performance: Myoglobin

14.0 CNX® 2000.1

Accelrys, Inc., [*www.accelrys.com*] describes CNX as follows:

"CNX Provides Tools for X-ray and NMR Structure Determination

"CNX is based on the widely used programs X-PLOR and CNS. This program integrates X-ray diffraction and NMR spectroscopic data with molecular mechanics, dynamics, and energy minimization to aid in the solution of three-dimensional molecular structures. In addition, CNX offers state of the art algorithms such as maximum likelihood refinement for X-ray refinement and residual dipolar coupling restraints for NMR structure determination. In addition, CNX uses several supported force fields including the Engh-Huber and CHARMm force fields."

CNX is used to perform fundamental research in the area of life sciences in academic and commercial environments.

14.1 CNX Single-Processor Tests Description

The test cases in table 32 were used to characterize the single-processor performance of the computer systems described above.

Table 32. Test Set for CNX

Test Case	Number of Atoms and Reflections	Type of Calculation
bm-xray_min100	2,026 atoms 17,646 reflections	100-step X-ray minimization refinement
Direct FFT	558 atoms 3,788 reflections	1,000 steps of MD with direct FFT

14.2 Relative Single-Processor Performance on Different Computer Systems

In the following pages, the relative single-processor performance of CNX is shown on several test cases. The comparisons are grouped as follows. First, all systems with MIPS R14000 processors running at 500 MHz are compared with each other. These comparisons give information on the effect of the memory subsystems on performance. Subsequently, the performances of SGI Origin series servers with microprocessors running at different frequencies are compared with each other. These results help us understand how well the SGI NUMA architecture in the SGI Origin 3000 series servers supports microprocessors with higher clock rates.

The SGI computer systems being compared belong to the SGI Origin family. These are described in table 33.

Table 33. Servers Used to Measure CNX Performance

System	Processor	Frequency (MHz)	L2 Cache (MB)
Origin 2000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R12000	400	8
Origin 3000 Series	MIPS R14000	500	8
Origin 3000 Series	MIPS R14000	600	8
Origin 300	MIPS R14000	500	2

14.2.1 SGI Origin 500 MHz Systems

Table 34 makes clear that the single-processor performance of CNX is not influenced by the memory architecture of the computer system, or by the size of the secondary cache of the microprocessor.

Table 34. CNX: Relative Performance of SGI Origin Systems Equipped with MIPS R14000 500 MHz Processors

	Origin 300 Time in Seconds	Origin 2000 Series Time in Seconds	Origin 3000 Series Time in Seconds	Speedup Origin 2000 Series over Origin 300	Speedup Origin 3000 Series over Origin 300	Speedup Origin 3000 Series over Origin 2000 Series
bm-xray_min100	143.90	139.60	134.50	1.03	1.07	1.04
Direct FFT	69.60	67.30	67.40	1.03	1.03	1.00
Average				**1.03**	**1.05**	**1.02**

14.2.2 SGI Origin 3000 Series Servers

The fact that the speedups shown in the next table are essentially the same as the frequency ratio of the chips used corroborates the earlier conclusion that, for these relatively small test cases, CNX performance is largely CPU bound. In addition, SGI NUMA's excellent latency and bandwidth characteristics are able to sustain these CPU-bound performance improvements.

Table 35. CNX: Relative Performance of MIPS Processors in an SGI Origin 3000 Series System

Processor Frequency	R12000 400 MHz Time in Seconds	R14000 500 MHz Time in Seconds	R14000A 600 MHz Time in Seconds	Speedup 500 over 400 MHz	Speedup 600 over 400 MHz	Speedup 600 over 500 MHz
bm-xray_min100	163.90	134.50	113.30	1.22	1.45	1.19
Direct FFT	83.40	67.40	56.00	1.24	1.49	1.20
Average				**1.23**	**1.47**	**1.20**

14.3 CNX 2000.1 Parallel Performance

14.3.1 bm-xray_min100

We use figure 14, which shows the parallel speedup measured for an SGI Origin 300 500 MHz server, to illustrate the parallel performance of CNX. For this test case, a parallel acceleration of almost 3.5 over the single-processor execution is measured when the job is run on four processors. The SGI Origin 3000 series of servers exhibits a similar behavior.

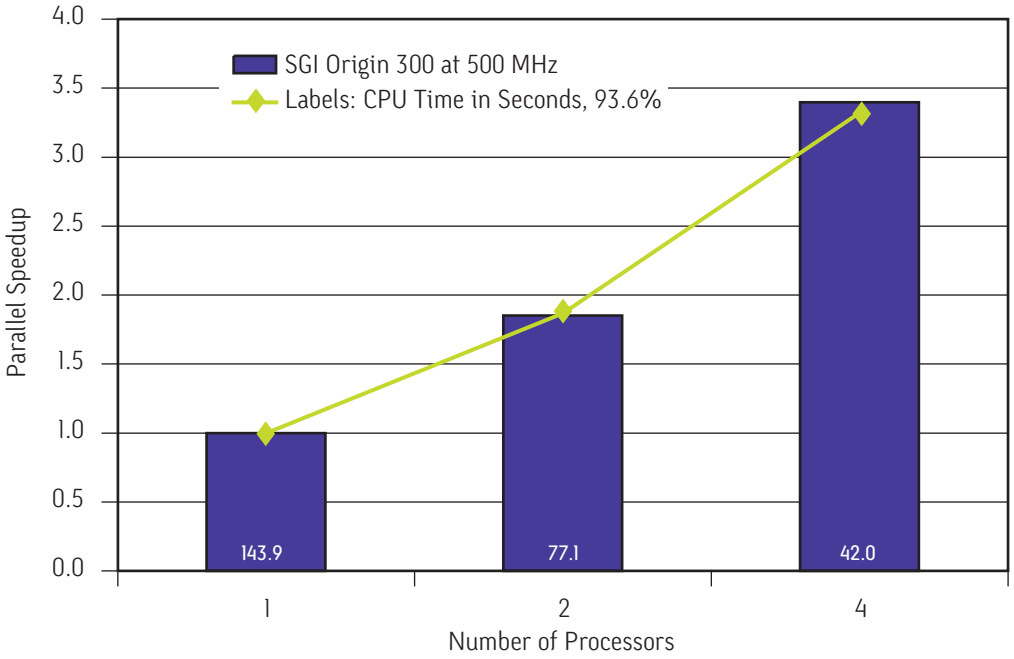

Fig. 14. CNX Parallel Performance: bm-xray_min100

15.0 High-Performance Computing in Today's Environment

High-performance computing demands for computational chemistry are co-evolving with broader industry HPC trends. Large quantum mechanical calculations once were the exclusive domain of traditional supercomputers, but today's chemists are demanding faster turnaround as well as faster throughput to address both the large volume of problems and the large problem sizes they are facing.

With today's large multiprocessor systems, such as the SGI Origin servers, there are two ways to approach the faster turnaround/faster throughput puzzle. One approach is to achieve a fast turnaround time for a single job. The other is to achieve high-throughput rates for multiple jobs.

The future trend will be to minimize turnaround times and maximize throughput rates of critical chemistry applications at the same time.

15.1 Computational Chemistry and HPC

Computational chemistry includes the simulation and modeling of chemical structures and reactions, allowing scientists to investigate molecular structures and properties that cannot always be obtained through experimental methods.

One widely used computational chemistry program is Gaussian 98, a connected system of programs for performing a variety of semiempirical and ab initio molecular orbital quantum mechanical calculations. In many instances, Gaussian 98 performance demands the features of a scalable, large-memory, single system image architecture to achieve reasonable calculation times for running very large jobs. Yet, chemists are also finding it useful to deploy large volumes of calculations simultaneously in a high-throughput model, especially when the scale of each individual calculation can be reduced.

15.2 High Throughput vs. Quick Turnaround

SGI Origin servers make use of a scalable SGI NUMA architecture that can be modularly grown while still maintaining a single system image to the user. Multiprocessor single-image systems can manage hundreds or even thousands of processors effectively when they are managed by a mature operating system. Good parallel efficiency on a multiprocessor system allows the user to achieve fast turnaround times on individual jobs, enabling runs to be made in a much shorter amount of time than on a single processor.

For large problems, peak performance can be solved by directly addressing the turnaround rate of a large job by utilizing efficient parallelization strategies for the application. But it is often not enough to be able to run one large job quickly—it's also important to have efficient throughput, that is, a steady workflow for numerous jobs that makes optimal use of computational resources. In high-throughput mode, a user can reduce the total solution time of running multiple jobs by simultaneously launching the jobs on multiple processors.

For years SGI has been leading the industry with tools that deliver high-throughput capability for popular bioinformatics and chemistry applications. High-throughput computing (HTC) drivers for many bioinformatics applications are made available on the SGI Web site (*www.sgi.com/solutions/sciences/chembio*).

Similar HTC services for deploying popular chemistry applications, like Gaussian 98, will also be available to allow customers to achieve high-throughput capability for both the single system image and the distributed system design of the SGI Origin family. Application performance has been demonstrated to achieve very high multiprocessing efficiency, resulting in a many-fold improvement in throughput rates.

For these applications and other technical applications requiring fast turnaround and high-throughput capabilities, SGI Origin servers are ideal. Not only can the Origin servers achieve near-linear speedups and near-perfect throughput measures, but they also can greatly simplify system and data management because they are shared-resource, single-image systems.

The SGI Origin server family is well suited to changing, heterogeneous user demands, and it is an excellent system for use as a central computing resource in a large, evolving organization.

16.0 SGI Application Support

One of the most important areas of focus for SGI in the life and chemical sciences has been the implementation of scalable, high-throughput drivers for critical applications commonly used in high-volume production settings, some of which have been described in this report. An example was the release in 1999 of a high-throughput version of the leading sequence-comparison application BLAST (*www.sgi.com/chembio/resources/blast*). This technology allows high volumes of genetic sequences to be compared with other genetic sequences. HT-BLAST is used at more than 200 labs worldwide. The improved algorithm, developed by SGI, is able to take full advantage of the hardware scalability of SGI Origin family servers. HT-BLAST achieves speedups of more than 240 times on 256 processors. SGI has provided algorithmic enhancements yielding comparable performance improvements on other key bioinformatics software, such as parallel and high-throughput implementations of CLUSTAL W (*www.sgi.com/chembio/resources/CLUSTAL W*).

The SGI Applications Engineering Team provides significant value to customers in both bioinformatics and chemistry. For example, Dr. Roberto Gomperts, principal scientist at SGI, has helped implement the initial parallelization in the Gaussian application (*www.gaussian.com*), for which he was named a co-author. In addition, the application engineering team focused on optimizing the performance of computational chemistry applications such as GAMESS, CHARMM, AMBER, Dmol, and CNX. The team also provides guidance on the incorporation of leading visualization algorithms to improve the utility of molecular visualization applications, including those available from Tripos, Accelrys, and the Chemical Computing Group, to name a few. The SGI commitment to improving the efficiency of applications has resulted in significantly shortening the cycle time for life sciences computing.

17.0 SGI Origin Family Servers for Bioinformatics and Chemistry

Building on the robust SGI NUMA architecture that has made the SGI Origin family of servers the most modular and scalable in the industry, the new SGI Origin 3000 series delivers flexibility, resiliency, investment protection, and superior performance. Now taking modularity a step further, you can scale CPU, storage, and I/O components independently within each system. Complete multidimensional flexibility allows organizations to deploy, upgrade, service, expand, and redeploy system components in every possible dimension to meet any business demand.

NUMAflex™ is a revolutionary snap-together server system concept that allows you to configure and reconfigure systems brick by brick to meet the exact demands of your applications. It allows you to upgrade CPUs to keep apace of innovation and to isolate and service I/O interfaces on the fly. You pay only for the computation, data processing, or communications power you need, and you expand and redeploy systems with ease as new technologies emerge.

With its high bandwidth, superior scalability, and efficient resource distribution, the new generation of SGI Origin servers—SGI® Origin® 3200, SGI® Origin® 3200C, SGI® Origin® 3400, and SGI® Origin® 3800—are performance leaders. Leveraging the next generation of SGI NUMA architecture and the IRIX 6.5 operating system, SGI Origin 3000 series servers work with your existing application software and are compatible with other SGI IRIX OS-based servers and workstations.

Among the key features of the SGI Origin 3000 series are increased memory bandwidth and lower memory latencies. The increase in memory bandwidth leads to improved processor performance, as more data is able to move from memory and I/O to the processors. The lower memory latency leads to increased processor, I/O, and memory scalability due to faster communications.

The robustness of the SGI IRIX operating system enables customers to make maximum use of these key hardware improvements. Among the software features that benefit computational chemists are:

- Binary compatibility with SGI Origin 2000 series servers, which ensures continued application availability
- A 64-bit operating system, which ensures accessibility to huge amounts of memory needed to solve large quantum chemical calculations or to load large databases into memory and large files and filesystems
- Weightless processes and threads—a feature that ensures the highest level of interactivity, while maintaining maximum system load
- IRIX checkpoint and restart capabilities, which ensure the success of long-running calculations

These IRIX features are described in the next section.

18.0 IRIX Features Supporting Large Job Capabilities

18.1 Weightless Processes and Threads

Achieving a balance between interactive and noninteractive usage is an important challenge in managing production computing environments to ensure optimal efficiency in system utilization. In an ideal situation, jobs run interactively will always receive the highest (UNIX®) prioritization and they will be run as if they were the only job on the system. Performance benchmarks, including all the results in this report, are usually run in this highly interactive, one-job-per-machine mode.

In production environments, most systems are severely overloaded with jobs. In an oversubscribed system, interactive performance can be quite poor, even when some of the jobs have been designated to run in the background with lower OS priorities. Research systems supporting staff in multiple sites and multiple countries have a strong need to run jobs interactively on the most heavily utilized systems. The outcome is that machine response can degrade quickly if the situation is not addressed.

A unique feature available under IRIX 6.5.9[9] helps to alleviate this problem. Weightless prioritization of processes and threads allows the important, interactive jobs to run at the highest level of priority, while background "weightless" jobs are essentially "sleeping" and in no way affect the turnaround time of the interactive jobs. The system manager can designate the relative importance of different jobs and can ensure that sleeping or weightless jobs are constantly ready to take advantage of available cycles not being utilized by the higher-priority, interactive jobs.

Here is an example of this system feature. An eight-processor SGI Origin family server was used for the following experiments:

- As background activity, eight single-processor FASTA jobs are running all the time, keeping all eight processors busy
- The interactive job is modeled by a short parallel CLUSTAL W job running in parallel on eight processors (24 seconds on the stand-alone machine)
- When both foreground and background jobs are running, the server is essentially oversubscribed by a factor of 2; if no special measures are taken with respect to the FASTA jobs, the execution time for the CLUSTAL W job doubles (48 seconds)
- However, when the FASTA jobs are submitted with weightless priority, the turnaround time for the CLUSTAL W job is the same as when the job is run on an otherwise idle machine, 24 seconds

Weightless processes and threads are very powerful tools to ensure full utilization of computer systems in production environments. As seen in the example above, interactive jobs get full control of the computational resources over weightless processes. As the processes of the interactive jobs free system resources, lower-priority jobs immediately make use of these relinquished resources.

18.2 Checkpoint and Restart

Another critical need for production computing is ensuring that long-running analytical jobs that take days, weeks, or even months are able to finish. Abnormal termination of these jobs due to electrical, system, or network outages or due to the needs of other users can result in weeks or months of wasted cycles. Applications software often takes into account these possibilities and on a regular basis writes out restart files. For those situations where either the application does not have this recovery mechanism and/or the checkpoint intervals are very large, the checkpoint and restart feature of IRIX 6.5 can help alleviate this problem.

An excerpt of the IRIX man page follows:

Name cpr - checkpoint and restart processes;

Description
IRIX Checkpoint and Restart (CPR) offers a set of user-transparent software management tools, allowing system administrators, operators, and users with suitable privileges to suspend a job or a set of jobs in midexecution, and restart them later on. The jobs may be running on a single machine or on an array of

networking connected machines. CPR may be used to enhance system availability, provide load and resource control or balancing, and to facilitate simulation or modeling.

Checkpoint and restart capabilities allow users to save the state of long-running jobs. This allows for on-demand access to the entire machine for high-priority jobs without losing the state of jobs in progress. Long-running jobs may be protected against failure by periodic checkpointing. Using checkpoint and restart as part of a regularly occurring IRIX cron utility enables long-running jobs to suspend, write an image of the job to disk, and then resume execution. If, during the ensuing period between cron jobs, the system fails, the job can be resumed from the checkpoint file. Thus, during a month-long job, the only part of a calculation that fails would be from the time between cron jobs. Scheduling the checkpoint period at 24 hours, for example, would ensure that less than one day's calculation is lost.

19.0 Acknowledgments

Professor Charles L. Brooks, III (*brooks@scripps.edu*)
CHARMM
mmtsb.scripps.edu/software/charmm.html

David A. Case (*case@scripps.edu*)
Amber
www.scripps.edu/case

Accelrys, Inc. (*support-us@accelrys.com*
support-eu@accelrys.com)
CHARMm
CNX
www.accelrys.com

Fujitsu Systems Business of America Inc. CAChe Group
(*sales@cachesoftware.com*)
MOPAC
www.cachesoftware.com/mopac/index.shtml

Gaussian, Inc. (*info@gaussian.com*)
Gaussian98
www.gaussian.com

Harvard University (*chmgr@tammy.harvard.edu*)
CHARMM
http://yuri.harvard.edu/

Iowa State University Chemistry Department (*gamess-users@Glue.umd.edu*)
GAMESS Version 25 June, 2001
www.msg.ameslab.gov/GAMESS/GAMESS.html

Pacific Northwest National Laboratory (PNNL)
NWChem
www.emsl.pnl.gov:2080/docs/nwchem/nwchem.html

Schrödinger, Inc. (*help@schrodinger.com*)
Jaguar
MOPAC
www.schrodinger.com

University of California San Francisco
(*willa@cgl.ucsf.edu*)
Amber
www.amber.ucsf.edu/amber/amber.html

20.0 References

[1] Gaussian 98, Revision A.11.3, M.J. Frisch, G.W. Trucks, H.B. Schlegel, G.E. Scuseria, M.A. Robb, J.R. Cheeseman, V.G. Zakrzewski, J.A. Montgomery, Jr., R.E. Stratmann, J.C. Burant, S. Dapprich, J.M. Millam, A.D. Daniels, K.N. Kudin, M.C. Strain, O. Farkas, J. Tomasi, V. Barone, M. Cossi, R. Cammi, B. Mennucci, C. Pomelli, C. Adamo, S. Clifford, J. Ochterski, G.A. Petersson, P.Y. Ayala, Q. Cui, K. Morokuma, P. Salvador, J.J. Dannenberg, D.K. Malick, A.D. Rabuck, K. Raghavachari, J.B. Foresman, J. Cioslowski, J.V. Ortiz, A.G. Baboul, B.B. Stefanov, G. Liu, A. Liashenko, P. Piskorz, I. Komaromi, R. Gomperts, R.L. Martin, D.J. Fox, T. Keith, M.A. Al-Laham, C.Y. Peng, A. Nanayakkara, M. Challacombe, P.M.W. Gill, B. Johnson, W. Chen, M.W. Wong, J.L. Andres, C. Gonzalez, M. Head-Gordon, E.S. Replogle, and J.A. Pople, Gaussian, Inc., Pittsburgh PA, 1998.

[2] "General Atomic and Molecular Electronic Structure System." M.W. Schmidt, K.K. Baldridge, J.A. Boatz, S.T. Elbert, M.S. Gordon, J.H. Jensen, S. Koseki, N. Matsunaga, K.A. Nguyen, S. Su, T.L. Windus, M. Dupuis, J.A. Montgomery J. Comput. Chem., 14, 1347-63(1993).

[3] Jaguar 4.2, Schrödinger, Inc., Portland, Oregon, 2000.

[4] NWChem Version 4.1, as developed and distributed by Pacific Northwest National Laboratory, P. O. Box 999, Richland, Washing-ton 99352 USA, and funded by the U. S. Department of Energy, was used to obtain these results.

[5] "MOPAC 2000", J.J.P. Stewart, Fujitsu Limited, Tokyo, Japan (1999).

[6] Amber Authors: David A. Case, David A. Pearlman, James W. Caldwell, Thomas E. Cheatham III, Junmei Wang, Wilson S. Ross, Carlos Simmerling, Tom Darden, Kenneth M. Merz, Robert V. Stanton, Ailan Cheng, James J. Vincent, Mike Crowley, Vickie Tsui, Holger Gohlke, Randall Radmer, Yong Duan, Jed Pitera, Irina Massova, George L. Seibel, U. Chandra Singh, Paul Weiner, and Peter A. Kollman

References: D.A. Case, D.A. Pearlman, J.W. Caldwell, T.E. Cheatham III, W.S. Ross, C.L. Simmerling, T.A. Darden, K.M. Merz, R.V. Stanton, A.L. Cheng, J.J. Vincent, M. Crowley, V. Tsui, R.J. Radmer, Y. Duan, J. Pitera, I. Massova, G.L. Seibel, U.C. Singh, P.K. Weiner and P.A. Kollman (1999), AMBER 6, University of California, San Francisco.

The basic description of the methods incorporated in Amber is in: D.A. Pearlman, D.A. Case, J.W. Caldwell, W.S. Ross, T.E. Cheatham, III, S. DeBolt, D. Ferguson, G. Seibel, and P. Kollman. AMBER, a package of computer programs for applying molecular mechanics, normal mode analysis, molecular dynamics and free energy calculations to simulate the structural and energetic properties of molecules. Comp. Phys. Commun. 91, 1-41 (1995).

[7] CHARMM: A Program for Macromolecular Energy, Minimization, and Dynamics Calculations, J. Comp. Chem. 4, 187-217 (1983), by B.R. Brooks, R.E. Bruccoleri, B.D. Olafson, D.J. States, S. Swaminathan, and M. Karplus.

CHARMM: The Energy Function and Its Parameterization with an Overview of the Program, in The Encyclopedia of Computational Chemistry, 1, 271-277, P.v.R. Schleyer et al., editors (John Wiley & Sons: Chichester, 1998), by A.D. MacKerell, Jr., B. Brooks, C.L. Brooks, III, L. Nilsson, B. Roux, Y. Won, and M. Karplus.

[8] CNX: Crystallography & NMR eXplorer (X-PLOR) Written by: A.T. Brunger, P.D. Adams, G.M. Clore, W.L. DeLano, P. Gros, R.W. Grosse-Kunstleve, J.-S. Jiang, J. Kuszewski, M. Nilges, N.S. Pannu, R.J. Read, L.M. Rice, T. Simonson, G.L. Warren, J. Junker, C. Griesinger and Molecular Simulations Inc. (J. Badger, D. Berard, R.A. Kumar, S. Szalma, P. Yip).

References:

1. Burling, F.T., Weis, W.I., Flaherty, K.M., Brünger, A.T. Science, 1996, 271, 72-77.
2. Read, R.J. Acta Cryst., 1992, A42, 140-149; Hodel, A., Kim, S.-H., Brünger,A.T. Acta Cryst., 1992, A48, 851-859.
3. Pannu, N.S., Read, R.J. Acta Cryst., 1996, A52, 659-668; Adams, P.D., Pannu, N.J., Read, R.J., Brünger, A.T. Proc. Natl. Acad. Sci. USA, 1997, 94, 5018-5023.
4. Bonvin, A.M., Brünger, A.T.J. Biomol. NMR, 1996, 7, 72-76.
5. Bonvin, A.M., Brünger, A.T.J. Mol. Biol. 1995, 250, 80-93.

[9] Weightless threads are available under IRIX 6.5.9 and 6.5.10 with a patch; incorporated as part of 6.5.11.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

3315 (06/17/2002)

J14030