

Performance of Parallel *Gaussian 94* on Cray Research Vector Supercomputers

Joe Ochterski, Carlos Sosa and John Carpenter, Applications Department, Cray Research, Inc., Eagan, MN 55121 USA

ABSTRACT: *Gaussian 94* is the latest in the *Gaussian* series of electronic structure programs. It is an integrated system to model a broad range of molecular systems under a variety of conditions, performing its calculations from the basic laws of quantum chemistry. This new version includes many new methods and algorithms that enhance the program's performance on Cray Research supercomputers. In this paper we discuss the multitasked version of the program and illustrate how it leverages the multiprocessors and large, real, shared memory of Cray Research, Inc. hardware systems.

Introduction

For more than twenty years, *Gaussian* [1] has been used for computational chemistry, the modeling and prediction of molecular systems, an increasingly common practice for exploring chemical phenomena to develop new pharmaceuticals and chemical-based products. *Gaussian* has been available on Cray systems for more than 15 years and is currently installed at over one hundred of Cray Research's customer sites, ranging from top leading chemical companies to supercomputing centers.

Gaussian is an integrated system of electronic structure computational chemistry programs for carrying out *ab initio*, Density Functional Theory (DFT), and semi-empirical calculations. Because of the nature of quantum mechanical calculations, most studies carried out on large systems are highly CPU intensive. Parallel processing long has been recognized as a potentially powerful tool for faster computing. On Cray Research vector supercomputers, *Autotasking* provides a way to partition programs among multiple processors[2]. The partition of the work among multiple processors can be done at the DO loop level or at higher levels of parallelism, such as parallel regions or subroutine boundaries. In other words, it enables users to run programs in parallel on multiprocessor CRAY J90, CRAY YMP, CRAY C90, and CRAY T90. *Autotasking* represents the third generation of parallel processing software on Cray Research vector platforms.

Copyright © Cray Research Inc. All rights reserved.

Features of *Gaussian 94*

Gaussian 94, as its predecessors, has evolved into a very friendly program. It has been developed to address a broad range of computational questions for different chemical systems under different conditions, carrying out calculations using quantum mechanical models[3].

Gaussian 94 calculates energies, molecular structures, vibrational frequencies and numerous molecular properties for systems, not only in the gas phase but also in solution, in which the solvent is modeled as a continuum and the solute is placed inside a cavity. Properties of a chemical system can be obtained in the ground state or excited states via the single excitations method (CI singles).

The latest version of *Gaussian* introduces parallelization and many new algorithms that enhance the program's performance on Cray Research supercomputers. Energies, optimized structures and frequencies can be computed in parallel using Hartree-Fock and a variety of Density Functional Theory (DFT) methods. Other features include MP3 and MP4(SDQ) analytical derivatives; automated computation of Gaussian 1 Theory, Gaussian 2 Theory, and G2(MP2) energies with a single keyword; and automated computation of Complete Basis Set (CBS) method energies. These calculations require the use of highly correlated methods and in the past have been limited to small molecules.

Parallel and Vector Capabilities

Gaussian 94 includes many performance enhancements over previous versions. These enhancements range from improvements to the direct algorithm[4], first introduced in *Gaussian*

88, to selecting the most efficient computational path based on resources available on a particular platform.

Currently, Hartree-Fock and DFT energies in conjunction with first and second derivatives have been parallelized in *Gaussian 94* for Cray Research supercomputers as illustrated in the previous section.

Energies

Electronic structure methods rely on a wide variety of mathematical approximations to compute molecular energies[5]. *Gaussian 94* is capable of computing molecular energies at different levels of accuracy and computational cost. Three major classes of methods are currently available in the *Gaussian* program:

- Semi-empirical methods[6], which avoid full evaluation of all the terms in the theoretical equations and rely on experimental parameters. These methods are fast enough and reliable enough to be applied to large systems.
- Density Functional Theory (DFT) methods[6], which are becoming popular amongst computational chemists. These methods include the effects of electron correlation at a computational cost similar to Hartree-Fock methods.
- Conventional *ab initio* methods[6]: in this category, *Gaussian* provides methods ranging from Hartree-Fock to highly accurate methods such as Quadratic Configuration Interaction (QCI) and Coupled Cluster methods (CC).

Hartree-Fock calculations are extensively used in computational chemistry[5]. They represent a first approximation to more accurate methods and in the case of large systems, it may be the only applicable method. Thus, vectorization and parallelization of Hartree-Fock calculations is of paramount importance in computational chemistry. Figure 2 gives a simple example of the effect of vectorization as a function of the basis sets and size of the system. Clearly, small systems with small basis sets tend to favor mainly scalar work. However, as the system and basis sets get larger and larger, so does the number of floating point operations which can be computed in a given amount of time. This performance improvement can be attributed to an increase in the average vector length as the system gets bigger.

Additional performance can be gained by making use of multiple processors. The first parallel version of the program was introduced with *Gaussian 92*. *Gaussian 94* for the first time offers a parallelized version using *autotasking* on Cray Research vector supercomputers.

It is well known that in molecular electronic structure calculations, the rate-limiting step is the computation of the two-electron integrals and their derivatives. In a self-consistent field (SCF) iterative scheme, these integrals are part of an NxN matrix termed *Fock* matrix[4,5],

$$F_{\mu\nu} = H_{\mu\nu} + \sum_{\lambda=1}^N \sum_{\sigma=1}^N P_{\lambda\sigma} \left[(\mu\nu|\lambda\sigma) - \frac{1}{2} (\mu\nu|\nu\sigma) \right]$$

where $H_{\mu\nu}$ represents the energy of a single electron in a field. The quantities $(\mu\nu|\lambda\sigma)$ are two-electron repulsion integrals. In Gaussian, these quantities are computed once and stored or recomputed many times as needed.

In a simple algorithm, the parallelization of the *Fock* matrix elements may be seen in Figure 1. At the end of the first loop over N_{CPUS} all the contributions to the *Fock* matrix have been computed, the last DO loop adds all the contributions together in serial block of code.

```

loop over NCPUS
  loop over total angular momentum
    loop over degrees of bra and ket contraction
      do integrals for 1/NCPUS of shell quartets
    end loop
  end loop
end loop
loop over NCPUS ( serial code )
  add 1/NCPUS contributions together
end loop

```

Figure 1: Parallelization of the *Fock* matrix.

Inside the loop over the number of CPUs, several loops generate the drivers needed to compute integrals and the shell quartets to look at. Each task looks at $1/N_{\text{CPUS}}$ of the shell quartets, discarding those which do not need to be done due to symmetry, cutoffs, or the fact that they have already been done. Then, the pertinent integrations are done. The last loop sums up the contributions to the *Fock* matrix, derivative matrices, or density matrices, depending on the type of integrals computed. Figure 3 illustrates the significance of spreading the work among several central processing units (CPUs) during a *Gaussian* calculation. Shown in Figure 3 are performance results for a model of taxol calculated on a CRAY J916 system with multiple CPUs. The benchmark was executed at the Hartree-Fock level with 660 basis functions.

Autotasking speedup for a large application can be difficult to predict. In general, a code needs to be parallelized 99% to show a perfect (or near perfect) speedup. Otherwise, the serial processing in the code will tend to dominate the overall performance. To look at the performance of the code, the estimated speedup has been computed as follows, $s(E) = t_0/t_N$, where t_0 corresponds to the elapsed time that took to run the job on 1 CPU and t_N is the elapsed time for N CPUs. In this example, the speedup of going from 1 to 2 CPUs is 1.7 or 85% parallelized according to Amdahl's law. This speedup for large systems can be substantial. The effect of going from 8 to 16 CPUs, for this particular case, is small.

In addition to Hartree-Fock, all the Density Functional Theory methods in *Gaussian 94* have also been parallelized. DFT offers a promising tool that may be applied to large systems. This type of methods include correlation effects in a

form in which the cost is formally N^3 , where N is the number of basis functions. In practice, the scaling can be lower.

The most accurate methods in *Gaussian 94* are generally highly CPU intensive. In the past, these methods have been limited to small systems. Methods such as Gaussian 1 Theory, Gaussian 2 Theory, and G2(MP2) as well as Complete Basis Set (CBS) method require the use of highly correlated methods. With the latest version of *Gaussian 94*, parallel kernels for MP4 and QCISD(T) leverage the multiple processors. Figure 4 illustrates MP4 calculations for methylhydrazine using multiple CPUs on a CRAY J916.

The speedups for this system are 1.7 and 2.7 for 2 and 4 CPUs, respectively. Most of the parallelization comes during the calculation of the triple excitations. Triples contributions to the MP4 energy can be implemented mainly in terms of BLAS routines which leverage multiple processors. Since triples in general account for more than 90% of the total CPU time, hence, they show good parallel speedups. This feature is particularly important when computing accurate thermochemical data.

Geometry Optimizations

Now, *Gaussian 94* can compute stationary points on a potential energy surface using multiple processors on a Cray Research supercomputer. In conjunction with the newly developed method for stationary point searches using redundant internal coordinates[6], optimizations for large systems can be reduced literally from weeks to days. Figure 5 compares *Gaussian 92* running an optimization using a single CPU against *Gaussian 94* taking advantage of multiple-CPU's and redundant internal coordinates.

Molecular Properties

In computational chemistry, predicting properties such as vibrational frequencies, magnetic properties, atomic charges, dipole and higher multipole moments, electrostatic potentials and electron densities is important to compare or provide a finger print to experiment. These properties can be computed using multiple-CPU's with the Hartree-Fock and Density Functional Theory methods using analytical second derivatives.

I/O Capabilities

The I/O requirements for *Gaussian 94* vary widely from very small to very intensive I/O. In the case of direct or incore SCF calculations the two-electron integrals are recomputed as needed or stored in memory, so I/O requirements are minimal. On the other hand, post-SCF calculations store partial lists of two-electron integrals on disk and perform intensive I/O.

Gaussian 94 offers I/O performance improvements. I/O optimization can provide significant savings in elapsed time (time to solution)[7]. The movement of data relative to the CPU can

be either synchronous or asynchronous. Synchronous is the most typical. This form of I/O returns control to the caller after all requested data is transferred to or from the program's variables. On the other hand, asynchronous I/O returns control to the caller as soon as the transfer of data has been initiated, but not necessarily completed. This of course, allows CPU and I/O to work in parallel. Asynchronous I/O has been implemented in the latest version of *Gaussian 94*. Cray Research supercomputers are the only ones which take advantage of this optimization.

Summary

Gaussian 94 introduces parallelization and many new algorithms that enhance the program's performance on Cray Research supercomputers, including the CRAY J90 series systems. Energies, optimized structures, and frequencies can be computed in parallel using the Hartree-Fock and a variety of Density Functional Theory methods. Also, parallel kernels for MP4 and QCISD(T) can take advantage of multiple processors and large, shared memories of Cray Research systems to compute new, more complex systems.

In general, SCF and DFT with and without gradients have been shown to run efficiently on all the Cray Research vector line of supercomputers. Scaling is clearly dependent on the size of the system and the number of basis functions. For more than 16 CPUs we have observed that the scaling tends to be small. *Gaussian 94* represents a major step forward to allow chemists to treat systems not possible in the past.

References

- [1] Frisch, M. J., G. W. Trucks, H. B. Schlegel, P. M. W. Gill, B. G. Johnson, M. A. Robb, J. R. Cheeseman, T. Keith, G. A. Petersson, J. A. Montgomery, K. Raghavachari, M. A. Al-Laham, V. G. Zakrzewski, J. V. Ortiz, J. B. Foresman, J. Cioslowski, B. B. Stefanov, A. Nanayakkara, M. Challacombe, C. Y. Peng, P. Y. Ayala, W. Chen, M. W. Wong, J. L. Andres, E. J. Replogle, R. Gomperts, R. L. Martin, D. J. Fox, J. S. Binkley, D. J. DeFrees, J. Baker, J. P. Stewart, M. Head-Gordon, C. Gonzalez, and J. A. Pople. 1995. *Gaussian 94* Rev. C.3. Gaussian Inc., Pittsburgh, PA.
- [2] 1989. *Autotasking: Automatic Parallel Processing on Cray Research Sytems*. Cray Channels Winter:2-3.
- [3] Foresman, J. B., and AE. Frisch. 1993. *Exploring Chemistry with Electronic Structure Methods: A Guide to Using Gaussian*. Gaussian Inc., Pittsburgh, PA.
- [4] Almlof, J., K. Faegri, and K. Korsell. 1982. Principles for the Direct SCF Approach to LCAO-MO Ab Initio Calculations. *J. Comp. Chem.* 3:385.
- [5] Hehre, W. J., L. Radom, P. v.R. Schleyer, and J. A. Pople. 1985. *Ab Initio Molecular Orbital Theory*. John Wiley & Sons, New York, New York.
- [6] Frisch, M. J., AE. Frisch, and J. B. Foresman. 1995. *Gaussian 94 User's Reference*. Gaussian Inc., Pittsburgh, PA.
- [7] 1995. *Cray Programming Environments: Optimizing Code on Cray PVP Systems, SG-2192*. Cray Research, Inc., Mendota Heights, MN.

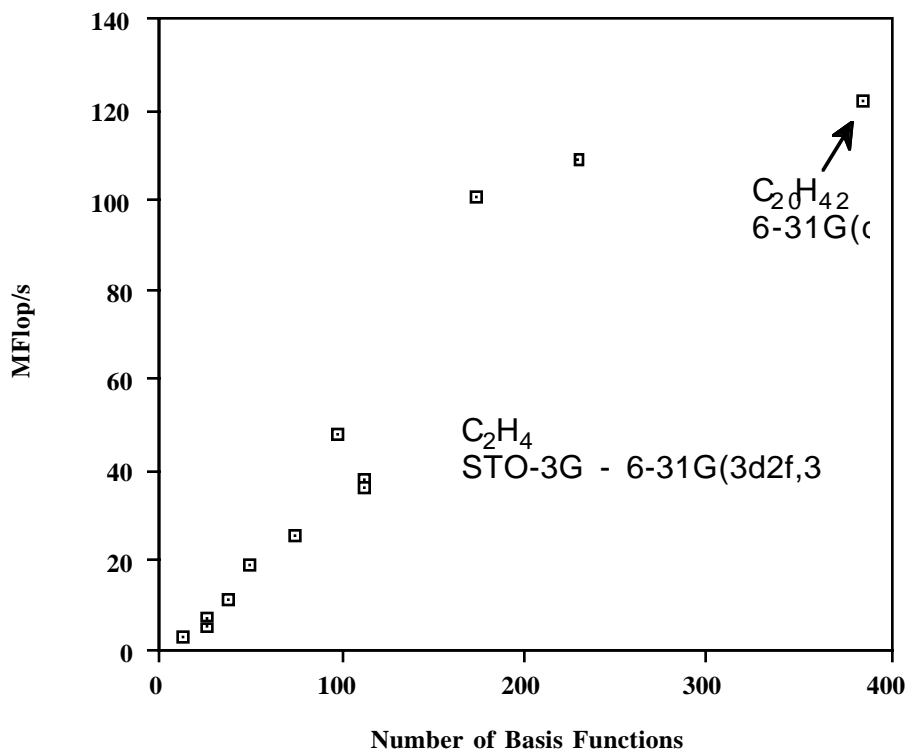


Figure 2. Single-CPU timings on a CRAY C90 for Hartree-Fock calculations as a function of the system size and basis sets.




Figure 3. Multiple-CPU benchmark at the Hartree-Fock level (single point) with 660 basis functions. Taxol is an anticancer drug. This calculation was carried out on a CRAY J916.




Figure 4. Performance results of methylhydrazine calculations executed at the MP4(SDTQ) level with 156 basis functions. This calculation was carried out on a CRAY J916 system.

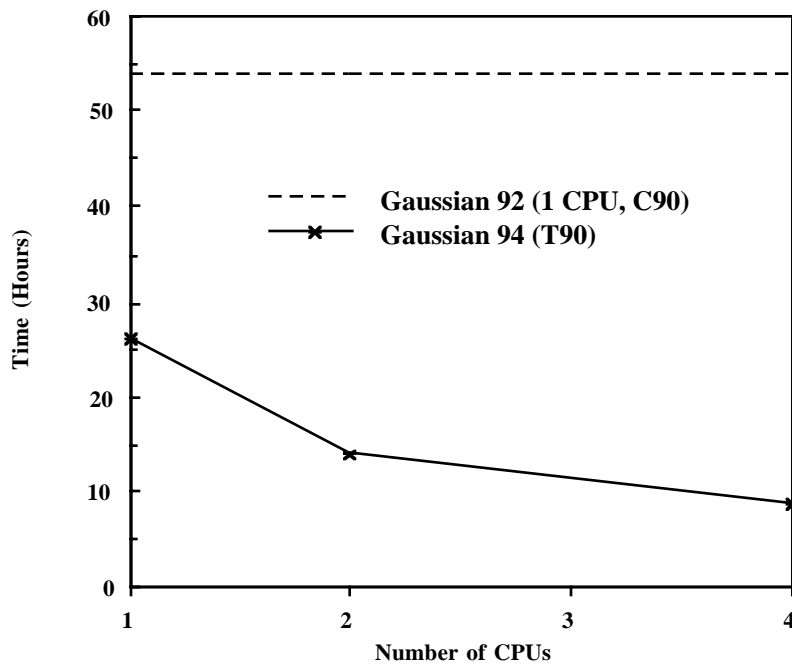


Figure 5. Comparison between *Gaussian 92* running on a single-CPU C90 versus *Gaussian 94* running on