The Journal of Physical Chemistry A

 VOLUME 103, NUMBER 13, APRIL 1, 1999

FEATURE ARTICLE

Correlated ab Initio Electronic Structure Calculations for Large Molecules

Richard A. Friesner,*,† Robert B. Murphy,‡ Michael D. Beachy,† Murco N. Ringnalda,‡ W. Thomas Pollard,‡ Barry D. Dunietz,† and Yixiang Cao†

Department of Chemistry, Columbia University, New York, New York 10027, and Schrödinger, Inc., 121 SW Morrison, Suite 1212, Portland, Oregon 97204

Received: June 5, 1998; In Final Form: January 13, 1999

We discuss computational methods for carrying out correlated ab initio electronic structure calculations for large systems. The focus is on two types of methods: density functional theory (DFT) and localized orbital methods such as local MP2 (LMP2) and a multireference version based upon a generalized valence bond reference wave function, GVB-LMP2. The computational performance of both approaches using pseudospectral numerical methods is documented, and calculated thermochemical and conformational energetics are compared to experimental data.

1. Introduction

Solution of the electronic Schrodinger equation, or for relativistic systems, the Dirac equation, is fundamental to the application of theoretical methods to chemical problems. In the early days of quantum mechanics, accurate solutions were limited to atoms or the H_2 molecule; there was simply no way to calculate many electron wave functions to a high level of precision for anything more complicated. The development of fast digital computers opened the possibility of solving the electronic Schrodinger equation, albeit while utilizing approximations, for interesting chemical systems. Beginning with the Gaussian suite of programs in 1970, computer codes of increasing power and generality have been constructed and at present are employed ubiquitously in academia and industry to address a wide variety of questions concerning molecular structure, energetics, and properties.

However, the computational expense of solving the Schrodinger equation remains a significant barrier to the deployment of ab initio quantum chemical methods in the solution of complex chemical problems. A principal difficulty arises from the scaling of the calculations with the number of electrons N in the molecule. The formal scaling of the Hartree−Fock equations with N in the Roothaan−Hall[1] formulation is N^4, a rather steep dependence. This can be reduced to N^3 using numerical methods, to N^2 via integral cutoffs, and ultimately to $N \ln N$ (or possibly N) via fast multipole methods[2−6] (although this last reduction occurs only when the molecule is quite large). N^2 scaling allows Hartree−Fock calculations to be routinely carried out for systems with on the order of 100 atoms. For many (if not most) chemical problems of practical interest, though, the self-consistent field (SCF) approximation inherent in Hartree−Fock theory is not sufficiently accurate. A major goal of ab initio electronic structure theory has therefore been development of electron correlation methods that can be easily applied to large molecules.

Traditional methods for electron correlation scale much less favorably with system size than Hartree−Fock theory. The conventional implementation of second order Møller−Plesset perturbation theory, for example, has a formal scaling of N^5, and cutoffs appear to be substantially less effective than for Hartree−Fock theory in reducing this in practice. Furthermore, as problems grow in size, very large amounts of disk storage and memory are required if calculations are to be run with reasonable efficiency. Methods based upon coupled cluster theory, such as QCISD(T)[7] and CCSD(T),[8] scale as N^7 (and

† Department of Chemistry.
‡ Schrödinger, Inc.

10.1021/jp9825157 CCC: $18.00
Published on Web 03/16/1999

require very large amounts of disk space) and hence can presently be applied only to small molecules.

During the past decade, two new approaches to the electron correlation problem have been developed that have led to a great deal of progress towards the goal of large molecule applications. The first of these, density functional theory (DFT), has undergone an explosive surge of popularity in the last 5 years. The computational scaling of DFT with system size is roughly comparable to that of HF, although for DFT some numerical integration methodology is always required. Therefore, most of the focus in improving DFT methods has centered on achieving higher accuracy compared to experiment for chemical properties, and on reducing the absolute computational expense. Recently, however, there has been considerable interest in developing fast multipole methods for DFT,[2−6] and for very large systems the theoretical asymptotic linear scaling regime has been reached.[4]

The second approach is at present much less widely used in the chemical community. It is based upon the observation that the orbitals in a Hartree−Fock calculation can be effectively localized by a simple unitary transformation into chemically understandable bonds and lone pairs. Once the orbitals are localized, the scaling of electron correlation methods with system size can be dramatically reduced, as suggested by Pulay and co-workers in their work in the mid 1980s.[9,10] An improved, albeit somewhat more expensive, version of the theory can be developed based upon a multiconfigurational wave function employing localized orbitals. The reference wave function of choice is the generalized valence bond (GVB) approach, pioneered by Goddard and co-workers more than 20 years ago.[11,12]

The implementation of localized orbital methods in an efficient numerical algorithm, so that the large gains in computational efficiency available in theory could be realized in practice, has been an arduous process. Much of our own work[13,14] has been focused on this effort. We discovered that the replacement of conventional analytical methods for integral evaluation with numerical algorithms allows a remarkable reduction in the computational effort required to carry out localized correlation methods. The resulting methodology, which has only been fully realized in the past two years, reduces the scaling of powerful localized correlated methods to the $N^2 - N^3$, and is readily applicable to 50−100 atom systems using current computational hardware.

The objective of the present article is 2-fold. First, we discuss the different computational methods that can be used for DFT and correlated localized orbital calculations, and we present some timing results from our pseudospectral implementations.[15] Secondly, we survey the performance of both approaches for several important chemical properties, examining the effect of basis set size as well as electron correlation method. In the ab initio literature, the vast majority of systematic assessment of the performance of correlation methods has been concerned with small molecule bond energies, as in studies utilizing the Gaussian-2(G2) database of Pople and co-workers.[16] However, there are an enormous number of problems of practical interest, for example, ligand binding to proteins, the key problem of structure based drug design, in which no chemical bonds are made or broken. Rather, what is important are small energy differences, arising from intermolecular interactions and conformational changes. While these problems are in some sense "easier" than the computation of bond energies in that the qualitative nature of the wave function is not being altered by rupture of an electron pair, the level of precision demanded is considerably greater, and the performance of correlation methods on one type of problem may not carry over to the other. For this reason, it is important to bring the same level of statistical assessment to the performance of quantum chemical methods for this sort of energy differences as has been applied to bond energies. The results presented in this paper are only a beginning of this effort, but they already allow a much clearer evaluation of the accuracy of the various correlation methods than do anecdotal comparisons of theory and experiment for a few molecules.

The design of the paper, in which DFT and localized orbital methods are contrasted side by side, is intended to provide a balanced picture of the strengths and weaknesses of both approaches. There are clearly a large class of applications for which DFT is at present the method of choice, based upon computational efficiency and ease of use, and the number of such applications is likely to increase in the future as better functionals are developed. Nevertheless, there are problems for which DFT does not yet provide adequate accuracy, or, in some cases, fails quite badly. For many of these cases, localized orbital methods constitute a practical, cost-effective alternative. We have implemented versions of both DFT and localized correlation methodologies in the Jaguar suite of ab initio electronic structure codes. When combined with a graphical user interface allowing either type of calculation to be run conveniently, the availability of both methodologies constitutes an excellent starting point for the robust application of ab initio electronic structure calculations to large systems. The growth of available computing power, continued improvement in software implementations, and advances in theory guarantee that these tools will become considerably more powerful and reliable in the next decade, in our codes and in those developed in other laboratories.

The paper is organized as follows. In section 2 we briefly discuss the underlying formalisms of DFT and localized correlation methods. Our goal is to provide a qualitative overview of the structure of each approach, as opposed to a detailed exposition of the mathematics, which has been presented in many other publications. Sections 3 and 4 discuss computational methods for solving the relevant equations, emphasizing the numerical approach, based upon pseudospectral (PS) methods, that we have been developing for the past 10 years. Again, our goal is to evaluate performance as opposed to explicate the algorithms in detail, so we present an overview of the methods along with timing results for a variety of molecules, basis sets, and chemical problems. Section 5 examines the accuracy of the methods as compared to experimental data for thermochemistry and conformational energetics. Finally, in section 6, the Conclusions, we consider future directions.

2. Theoretical Formalisms

2.1. Density Functional Theory. Density functional theory originates from the work of Hohenberg and Kohn,[17] who proved that the energy of a many electron system can formally be expressed as a functional of the electron density. Given this functional, the ground state density $\rho(\mathbf{r})$ can be found via its minimization, subject to the constraint that the integral of the density over all space is unity. There are a variety of ways to recast the equations for $\rho(\mathbf{r})$. From a computational point of view, the Kohn−Sham equations,[18] in which $\rho(\mathbf{r})$ is expanded in an orbital basis set, are most convenient. This leads to a three-dimensional self-consistent field equation for the density orbitals $\phi_i(\mathbf{r})$ of the form

$$\left[\mathbf{h}(\mathbf{r}) + \int \frac{\rho(\mathbf{r}_2)}{|\mathbf{r} - \mathbf{r}_2|}\, d\mathbf{r}_2 + V_{xc}(\mathbf{r})\right]\phi_i(\mathbf{r}) = \epsilon_i\phi_i(\mathbf{r}) \qquad (1)$$

With $\mathbf{h}(\mathbf{r})$ the one-electron kinetic/nuclear attraction operator. $V_{xc}(\mathbf{r})$ is an operator derived from a universal exchange−correlation functional, f_{xc}, containing all of the information about the exchange and correlation energy, whose specification defines the particular density functional approximation that is being employed.

The advantages of density functional theory flow from the reduction of the many-electron Schrodinger problem to a three-dimensional field equation. The solution of eq 1 is apparently less complicated and expensive than conventional correlated ab initio formulations involving huge numbers of Slater determinants. The disadvantage of DFT is that one must guess the form of f_{xc}, which cannot be derived rigorously from first principles. The quality of any given guess, which controls the level of accuracy of the method, must be therefore determined by comparing the results with experiment.

For many years, density functional calculations were carried out predominantly using the local density approximation (LDA).[18] Within the LDA the function f_{xc}^{lda} is assumed to depend only upon the local electron density ($f_{xc}(\rho)$) with matrix elements between basis functions χ_μ, χ_ν of the corresponding exchange operator $V_{xc}(\mathbf{r})$ being given by

$$V_{\mu\nu}^{xc} = \int dr\left(\left(\frac{\partial f_{xc}[\rho]}{\partial\rho}\right)\chi_\mu(r)\chi_\nu(r)\right) \qquad (2)$$

In the most widely used versions, $f_{xc}^{lda}(\rho)$ is fit to reproduce the uniform electron gas energy, which can be accurately calculated from Monte Carlo calculations.[19]

The LDA is a qualitatively reasonable electronic structure theory and often provides results that are comparable in quality to Hartree−Fock theory, for example, for molecular geometries. From a statistical point of view, however, it is not a quantitatively accurate theory for a broad range of molecular properties. For example, bond energies are typically over bound by 50 kcal/mol, and hydrogen bonding energies and geometries are grossly in error. Thus, the LDA is not a competitor with high- level-correlated ab initio methods.

Throughout the past 2 decades, there have been numerous attempts to improve density functionals, primarily by incorporating corrections to the LDA based upon the gradient of the electron density. The first major breakthrough was made by Becke, who developed a new exchange functional[20] that, when combined with gradient corrected correlation functionals, yielded enormously improved bond energies for a data base of small molecule test cases. Subsequent work by Becke[21] and others,[22] including the Pople[23] group, led to an entire family of generalized gradient approximation (GGA) functionals, with more or less similar performance in comparison with experimental data. We consider a representative example of these functionals, the B-LYP (Becke exchange[20] with Lee−Yang−Parr correlation[24]) functional, below.

A second major breakthrough, in a rather different direction, was also made by Becke.[25,26] He devised a functional, based upon adiabatic connection arguments, which incorporated an admixture of Hartree−Fock exchange in a linear combination with the usual density functional ingredients. By fitting a small number of adjustable parameters, a low average error in predicting atomization energies of molecules in the G2 data base was obtained.[26] Other properties, such as equilibrium geometries and hydrogen bonding energies, were also improved. Again an entire class of such functionals have been developed by Becke and others. Below, we consider a representative example of this class of functionals as well, the B3-LYP functional (Becke three parameter adiabatic connection fit combined with the LYP correlation functional).[20,24,26] More recently, Becke introduced a new functional based on generalized correction terms.[27] It was shown that the expansion to third order, resulting in a 15 parameter functional, has further improved performance. This functional was fitted against the extended G2 database.[16]

2.2. Localized Electron Correlation Methods. *2.2.1. Local Second-Order Møller−Plesset Perturbation Theory.* Second-order Møller−Plesset perturbation theory (MP2) is the least expensive of the conventional ab initio electron correlation methods, scaling formally as N^5. It provides the great majority of the correlation energy and significantly improves the calculation of many molecular properties as compared to Hartree−Fock theory, e.g., dipole moments, equilibrium geometries, and conformational energy differences, if a large basis set is employed.

A localized version of Møller−Plesset perturbation theory was systematically constructed by Pulay and co-workers[9,10] in the mid 1980s. The formal development of the theory is straightforward. First, orbitals from a canonical Hartree−Fock calculation are localized, via either a Boys[28] or Pipek−Mezey[29] localization procedure, into bond or lone pair orbitals. Then, a local correlation space is defined for each electron pair excitation out of the occupied space. Each electron in the pair is taken from one lone pair or bond. The local space is defined to be the union of atomic basis functions from each of the atoms associated with the a lone pair or bond from which the excitation occurs. The atomic basis functions in the local correlation space are in all cases orthogonalized to the orbitals of the occupied space.

To carry out local second-order Møller−Plesset perturbation theory (local MP2, or LMP2), the first-order wave function $\Psi^{(1)}$ must be calculated. The first order wave function is represented as a linear combination of determinants Ψ_{ij}^{pq}, each of which has a double excitation from a pair of occupied orbitals ij to a pair of orbitals pq in the local correlation space:

$$\Psi^{(1)} = \sum_{i\geq j,pq} C_{ij}^{pq}\,\Psi_{ij}^{pq} \qquad (3)$$

Because neither the localized occupied orbitals or the correlating virtual orbitals are eigenfunctions of the Hartree−Fock Hamiltonian, it is necessary to obtain the expansion coefficients in the first order wave function by solving a linear system of equations of the form[10]

$$\mathbf{T}_{ij}^{(2)} = \mathbf{K}_{ij} + \mathbf{F}\mathbf{C}_{ij}\mathbf{S} + \mathbf{S}\mathbf{C}_{ij}\mathbf{F} - \mathbf{S}\sum_k [F_{ik}\mathbf{C}_{kj} + F_{kj}\mathbf{C}_{ik}]\mathbf{S} = 0 \qquad (4)$$

where $\mathbf{S}$ is the overlap matrix among the nonorthogonal local virtual orbitals and $\mathbf{F}$ is the Fock matrix. The quantity K_{ij}^{pq} is the exchange integral $(ip|jq)$.

The solution of this equation is readily accomplished via conjugate gradient methods, which typically converge in 4−8 iterations. Solution of the linear system scales as $\sim(N^2 - N^3)$ and does not require an inordinate amount of computational effort.

The use of a localized correlation space, as opposed to the complete set of virtual orbitals, means that the correlation energy obtained from an LMP2 calculation is slightly different (typically by 1−2%) than a conventional MP2 result. However, this

difference is actually an advantage, as LMP2 eliminates much of the basis set superposition error (BSSE) associated with canonical MP2.[30] Calculations carried out by Saebo and Pulay have explicitly demonstrated that BSSE is effectively removed via the LMP2 formalism.[30] LMP2 calculations converge much more quickly to the large basis set limit than do conventional MP2 calculations. Recently, work in our own laboratory suggests[31] that similar problems arise in intramolecular interactions (for example, internal hydrogen bonding in a dipeptide) and that the local MP2 results are again preferable.

The key question with regard to the LMP2 formalism is whether or not the reduction of the virtual space can be translated into a practical reduction in computational effort. In the work of Pulay and co-workers,[9,10,30] such reductions were not observed, because localization of the virtual space has minimal impact upon the four index transform, the principal computational expense in a conventional implementation of LMP2. However, using pseudospectral methods, the four index transform can be eliminated, and this yields large reductions in the scaling of the calculation with system size. Our implementation of this approach is described below in section 4.

2.2.2. Localized Multireference Perturbation Theory: The GVB-LMP2 Methodology. While LMP2 provides significant improvements over Hartree−Fock results for most molecular properties, and as we shall see below, is superior in some cases to DFT as well, it is not a quantitatively accurate electronic structure theory. Traditionally, there are two ab initio quantum chemical formalisms providing a route to high accuracy. The first of these is multireference configuration interaction and related methods;[32] the wave function optimized at the SCF level consists from the start of many determinants, and the resulting MC-SCF wavefunction is subsequently corrected for dynamical correlation effects via, for example, single and double excitations. The second approach, which retains a single determinant reference but models triple and quadruple excitations via a quite effective exponential ansatz, is based upon an approach originally designated coupled cluster theory.[33] Two rather similar variants of this method, CCSD(T),[8] and QCISD(T),[7] are presently in wide use.

For the past decade, there has been a steadily increasing preference of professional quantum chemists for coupled cluster methods as opposed to the multireference methods, for several reasons. First, the coupled cluster approaches can readily be structured as a well defined "model chemistry" in the sense of Pople and co-workers, whereas multireference methods require selection of a particular MC-SCF reference state that is often perceived as arbitrary (or at least difficult to determine without doing quite a few computational experiments). Secondly, significant effort has been expended in developing highly efficient CCSD(T) and QCISD(T) algorithms, the last of which is encoded in the widely used Gaussian[34] series of programs. The ACES II package[35] provides a full spectrum of coupled cluster methods, including CCSD, CCSD(T), and CCSDT on a variety of wave functions. The scaling of either coupled cluster method with system size is $\sim N^7$, which restricts applications to small molecules. However, MR-CI methods that are systematic, such as the CASSCF[36] approach (in which an "active space" is designated and all electrons in the active space can be distributed in all possible arrangements among the active space orbitals) scales exponentially with the size of the active space and, hence, are also restricted to small molecules. Finally, QCISD(T) methods have been shown to yield accurate and robust results in a wide variety of small molecule test cases. For example, bond energies can be computed to ∼1−2 kcal/mol average error via the use of a single adjustable parameter, in the G2 theory prescription of Pople and co-workers.[16,37] In contrast, few systematic studies over a large database of molecules have been carried out with MR-CI approaches.

The recent development of multireference methods in which the CI step is replaced with perturbation theory, such as the CASPT2 approach of Roos and co-workers,[38] has diminished the cost of the MR approach, although the exponential scaling with active space size remains, and the MP2 step scales as N^5. Nevertheless, CASPT2 has been applied very effectively to the calculation of excited state properties,[38] and it is presently one of the most reliable methods for obtaining these for small to medium sized molecules.

While conventional multireference perturbation algorithms represent a significant step forward, they are still not suitable for routine application to systems in the 30−100 atom range, where much of interesting chemistry takes place. For MR methods to address such problems, they must be coupled with localized correlation methods. The perturbation theory part of the methodology can be transplanted from the LMP2 formalism of Pulay and co-workers[10,39,40] in a straightforward fashion. However, the MC-SCF component of the algorithm has to be modified substantially from a CASSCF type of wave function. Fortunately, a robust and generally applicable localized MR method was developed more than 20 years ago by Goddard and coworkers.[12] This method, the generalized valence bond (GVB) approach, automatically produces orbitals localized on bonds and lone pairs in the context of a self-consistent wave function. Furthermore, the computational scaling with system size of the two main versions of the method−GVB-PP (perfect pairing)[12] and GVB-RCI (restricted configuration interaction)[41]—can be reduced to the N^2-N^3 range, as described in section 4. Finally, work by several groups, principally those of Messmer[42,43] and Pulay,[39,40] has shown that a GVB reference wave function delivers excellent performance in a MR perturbation theory; the most sophisticated GVB-RCI-MP2[42,43] approach yields results comparable to that of CASSCF at a much lower cost.

The combination of GVB self-consistent field methods with localized perturbation theory suggests that, in principle, MR-based methods cannot only be made competitive with coupled cluster approaches, but that they may actually be able to produce superior accuracy at a much lower computational cost. However, there are formidable practical problems to implementing GVB-LMP2 calculations in a sufficiently robust fashion that it can be truly viewed as an automatic model chemistry, which can be applied to an arbitrary molecular problem without expert user intervention. It is also technically quite challenging to realize the potential enhancements in computational performance without sacrificing accuracy; if there are significant random numerical errors, the exercise has lost its purpose, which is the development of chemically accurate methods capable of reliably treating large molecules. Despite these difficulties, the progress that will be described below is quite substantial. The results we present at the very least suggest that the goal stated above is within reach in the next few years. In contrast, it is not at all obvious how the scaling of the competitive QCISD(T) or CCSD(T) can be reduced much below N^7. Localized methods have been developed for these approaches,[44] but the performance reported to date indicates only a marginal bottom line improvement in total CPU time over conventional algorithms. Of course, this does not preclude a novel attack on the problem leading to a qualitative reduction in computational effort in the future.

3. Numerical Implementation of Density Functional Theory

3.1. Introduction. In the Kohn−Sham formulation of density functional theory,[18] the electron density $\rho(\mathbf{r})$ is represented by the occupied orbitals obtained from the solution of eq 1,

$$\rho_\alpha(\mathbf{r}) = \sum_i \phi_i^{\alpha *}(\mathbf{r})\phi_i^{\alpha}(\mathbf{r}) \tag{5}$$

In the B-LYP version of gradient corrected density functional theory,[20,24] the exchange-correlation potential $V_{xc}(\mathbf{r})$ has matrix elements of the form

$$(V_{xc}^{\alpha})_{\mu\nu} = \int f^{\rho\alpha}\phi_\mu\phi_\nu + [2f^{\gamma\alpha\alpha}\nabla\rho_\alpha + f^{\gamma\alpha\beta}\nabla\rho_\beta]\,\mathrm{d}\mathbf{r} \tag{6}$$

with a corresponding equation for β spin. This equation is a generalization of eq 2 with f_{xc} now a functional of the density and gradient of the density.

Most approaches to the solution of eq 1 involve expansion of the density orbitals in terms of atomic basis functions χ_j,

$$\phi_i(\mathbf{r}) = \sum_j C_{ji}\chi_j(\mathbf{r}) \tag{7}$$

In Gaussian,[34] DGAUSS,[45] and Jaguar,[15] the basis functions used are the standard contracted Gaussian functions of conventional quantum chemistry. Other programs, such as DMOL[46] and ADF,[47] employ Slater orbitals or numerical orbitals.

If an atomic orbital expansion of the density orbitals is substituted into eq 1, standard manipulations produce a set of equations very similar to those of the Roothaan−Hall formulation of Hartree−Fock theory. The effective Fock operator $\mathbf{F}^{DFT}$ is given by,

$$\mathbf{F}^{DFT}(\mathbf{r}) = \mathbf{h}(\mathbf{r}) + \mathbf{J}(\rho(\mathbf{r})) + V_{xc}(\rho(\mathbf{r}),\nabla\rho(\mathbf{r})) \tag{8}$$

with $\mathbf{h}$ and $\mathbf{J}$ the one-electron and Coulomb operators. The Kohn−Sham eq 1 is solved using standard iterative techniques, such as direct inversion in the iterative subspace, typically in 5−10 iterations (although DFT is more difficult to converge than HF, due to the smaller size of the HOMO−LUMO energy gap). The computationally demanding step in this process is assembly of the Coulomb and exchange−correlation operators from the new guess for the density orbitals at each iteration. We consider each of these steps below.

3.2. Overview of Computational Methods for Density Functional Calculations. *3.2.1. Assembly of the Coulomb Operator.* Efficient calculation of electrostatic interactions is one of the core problems of modern computational physics. It is therefore not surprising that several powerful algorithms have been developed for this purpose. The quantum chemical electronic structure problem requires computation of the interaction between continuous charge distributions, as opposed to the point charges that are typically relevant to classical molecular mechanics simulations. Nevertheless, the two problems have considerably similarities, and there has been significant borrowing from one problem to another.

If the density orbitals in the Kohn−Sham equation are expanded in a finite basis of AO functions (α,β,μ,ν), the Coulomb operator from eq 8 has the form

$$J_{\mu\nu} = \sum_{\alpha,\beta}\rho_{\alpha,\beta}\int\frac{\alpha(\mathbf{r}_1)\beta(\mathbf{r}_1)\mu(\mathbf{r}_2)\nu(\mathbf{r}_2)}{r_{12}} = \sum\rho_{\alpha\beta}(\alpha\beta|\mu\nu) \tag{9}$$

with $\rho_{\alpha,\beta}$ the AO density matrix.

Following is a brief listing of the major strategies that have used to evaluate eq 9: (1) The simplest approach is to explicitly evaluate all of the relevant two center electron repulsion integrals $(\alpha\beta|\mu\nu)$ and assemble $J_{\mu\nu}$ directly as written in eq 9. This is the method used in the Gaussian system[34] of programs. Formally, there are N^4 integrals to evaluate; however, the use of integral cutoffs readily reduces this asymptotically to $\sim N^2$.

(2) The density $\rho(\mathbf{r})$ can be expanded in a set of auxiliary basis functions Φ_γ

$$\rho(r) = \sum_\gamma P_\gamma\Phi_\gamma(\mathbf{r}) \tag{10}$$

and the matrix element $J_{\mu\nu}$ can then be expressed as

$$J_{\mu\nu} = \sum_\gamma P_\gamma\int\frac{\mu(\mathbf{r}_1)\nu(\mathbf{r}_1)\Phi_\gamma(\mathbf{r}_2)}{r_{12}} \tag{11}$$

This is the approach taken in two codes that have their origin in the laboratory of Dennis Salahub, DEMON[48] (the current code from the Salahub laboratory), and DGAUSS[45] (now distributed by Oxford Molecular). This method has a formal scaling of N^3, which also reduces to N^2 by using cutoffs. However, the scaling prefactors, and rate of reaching the asymptotic N^2 limit, can be different for the two method. The three center Coulomb integrals in eq 11 are in principle considerably less expensive to evaluate per integral than the four center terms that dominate the evaluation of eq 9. On the other hand, there are issues concerning the accuracy and stability of the fitting procedure used to generate the density expansion. The performance of any given implementation of this approach must be judged by computational experiments.

(3) Solutions to Poisson's equation can be employed to generate the Coulomb field, and numerical integration methods then can be used to evaluate the integrals in eq 11. This is the approach taken in the DMOL[46] (distributed by MSI, Inc.) and ADF (developed in the research group of E. Baerends);[47] it is also used by Becke in his NUMOL[49] code, which has the additional feature of making no use of conventional basis set expansions. The preponderant use of grid-based methods can lead to an efficient code, particularly if cutoffs are properly implemented. The efficiency of the method depends upon how many grid points per atom are needed to achieve stable and accurate results. Again, this is difficult to specify in advance; the tradeoffs of speed versus accuracy have to be documented in numerical experiments, comparing against converged results.

(4) The pseudospectral construction of the Coulomb operator proceeds by calculating the Coulomb field $J(r)$ on a grid $\{g\}$, multiplying a basis function $R_\mu(g)$ by $J(g)$, fitting the product to an auxiliary basis expansion via least squares, and projecting the result onto the basis set. The fitting and projection back to basis function ν can be represented[50−54] by the operator $Q_\nu(g)$

$$J_{\mu\nu} = \sum_g Q_\nu(g)J(g)R_\mu(g) \tag{12}$$

The assembly of $J(g)$

$$J(g) = \sum_{\alpha\beta}\rho_{\alpha\beta}\int\frac{\chi_\alpha(\mathbf{r})\chi_\beta(\mathbf{r})}{\mathbf{r}_g - \mathbf{r}}\mathrm{d}\mathbf{r} = \sum_{\alpha\beta}\rho_{\alpha\beta}A_{\alpha\beta}(g) \tag{13}$$

is achieved using the three center one electron integrals $A_{\alpha\beta}$, which are very inexpensive to evaluate.

(5) All of the above methods exhibit N^2 scaling asymptotically. In order to reduce the scaling further, to N ln N or even N, multipole expansion methods must be used. Recent work[2−6,55] in the development of a fast multipole method specialized for continuous charge distributions has demonstrated the first practical implementation of this strategy. For very large systems, on the order of hundreds or thousands of atoms, a significant reduction in computational effort is provided. Multipole methods can be used in conjunction with any of the four strategies discussed above.

What is the best strategy for evaluation of the Coulomb operator? The individual approaches described above are not the only ones possible; the various techniques can be combined together to produce new algorithms as well. A formal analysis of the accuracy of any particular implementation is not feasible. The only way to determine performance is to carry out a set of test calculations for which converged benchmark results (produced using, for example, completely analytical methods or very large numerical grids) exist. It is critical to assess the quality of the results and efficiency with which they have been produced simultaneously; obtaining high efficiency is trivially accomplished simply by sacrificing accuracy or reliability (for example, by drastically reducing the number of grid points in a grid-based method).

3.2.2. Exchange-Correlation Functional: Local and Gradient Corrected Terms. For the B-LYP functional,[20,24] the only term required beyond the Coulomb and one electron terms is the matrix element of the exchange correlation potential $V^{xc}_{\mu\nu}$ in eq 6. This potential is defined by the exchange correlation functional $f_{xc}(\rho,\nabla\rho)$ a functional of the density ρ and gradient of the density in nonlocal versions as discussed above. There are two widely used approaches to the evaluation of V^{xc} matrix elements. The first is numerical integration, carried out by expressing the basis functions μ and ν and the functional on the grid, multiplying, and summing over grid weights. This requires design of an accurate numerical integration scheme (grid points and quadrature weights), of which several have been proposed in the literature. This approach is used in Gaussian,[34] DMOL,[46] ADF,[47] and Jaguar.[15] The second approach is expansion of the exchange-correlation functional in a set of auxiliary basis functions (typically Gaussians) followed by computation of three center overlap integrals. Methods along these lines are presently implemented in DGAUSS[45] and DEMON.[48]

3.2.3. Exchange-Correlation Functional: Hartree−Fock Exchange Terms. For the B3-LYP functional,[20,24,26] it is necessary to calculate matrix elements of the Hartree−Fock exchange operator $K_{\mu\nu}$ in addition to the V^{xc} terms discussed above. Several of the codes mentioned previously, DMOL, DGAUSS, DEMON, and ADF, do not possess this capability. There are numerous approaches to the evaluation of $K_{\mu\nu}$ presently in the literature. The first is straightforward summation over analytical integrals:

$$K_{\mu\nu} = \sum_{\alpha\beta} \rho_{\alpha\beta}(\mu\alpha|\nu\beta) \tag{14}$$

as is carried out in Gaussian and other conventional quantum chemistry codes. The second is the pseudospectral formulation of the exchange operator:

$$K_{\mu\nu} = \sum_g Q_\mu(g)K_\nu(g); \quad K_\nu(g) = \sum_\alpha A_{\nu\alpha}(g)\sigma_\alpha(g); \quad \sigma_\alpha(g) = \sum_\beta R_\beta(g)\rho_{\alpha\beta} \tag{15}$$

$Q_\mu(g)$, $A_{\nu\alpha}(g)$, and $R_\beta(g)$ are as defined above in the pseudospectral formulation of the Coulomb matrix. More recently, linearly scaling methods have been developed for the calculation of the exchange matrix.[56−59]

3.3. Pseudospectral Implementation of Density Functional Theory: Methods and Results. *3.3.1. Computational Methods.* Over the past several years, we have been developing a robust implementation of both gradient-corrected and adiabatic connection DFT methods. The energies and structures obtained have been extensively tested against results from Gaussian 92 and from our own codes using analytic integration, with agreement within a few tenths of a kcal/mol being obtained when dense exchange−correlation grids are used to converge the DFT operators in both approaches. The timing results we present below can therefore be taken to represent reliable results for the specific basis set in question (indeed, the errors due to basis set incompleteness and the inherent approximations of the DFT methods are in all cases much larger than any noise in the PS calculations). All calculations reported use Jaguar version 3.5.[15]

As stated above, we employ now standard methods for carrying out numerical quadratures to evaluate the exchange−correlation operators, based upon the original approach proposed by Becke.[20] However, we have made modifications in the details of the implementation that result in significant reductions in computational effort. The use of effective cutoffs is critical in obtaining optimal performance, and this must be combined with utilization of matrix multiply routines so as to drive assembly of the exchange-correlation functional at peak megaflop rates. The most important improvement we have made is the use of multigrid methods so that the exchange−correlation operator can be evaluated less frequently on the densest mesh. This algorithm will be described in detail in another publication.

Calculation of the Coulomb and Hartree−Fock exchange operators for the adiabatic connection methods is carried out using the standard PS algorithms as described above. However, for generalized gradient calculations we have accelerated the computations in two ways. First, the Hartree−Fock exchange assembly can be turned off, thus saving a significant amount of computation time (in contrast to conventional quantum chemical methods, where assembly is relatively trivial compared to integral generation). Secondly, we have implemented a fast multipole algorithm specifically tailored to the PS approach. Briefly, we use atom-centered multipoles and calculate the Coulomb field by expanding the $A_{\mu\nu}(g)$ three-center one-electron integrals in these multipoles. The use of atom-centered-only multipoles allows for significant reductions in CPU time to be achieved even for relatively small molecules (in contrast to the alternative cell multipole methods, which require large systems to display any effects). The reason is that there are a large number of exponent pairs contributing to the integral whose product center is very close to a particular atom (for example, all pairs formed from one large exponent and small exponent), and in this case the multipole expansion converges very rapidly and will yield accurate Coulomb fields as little as 2−3 Å from the product center. We have extensively tested the accuracy of the multipoles, and have set the tolerance to converge the energies better than 0.05 kcal/mol as compared to the results when multipoles are not employed. This methodology will be described in detail elsewhere.

3.3.2. Results. We present DFT timing results for a wide range of molecules (illustrated in Figure 1a−f) in Tables 1 and 2. The examples include a number of biologically interesting organic molecules (porphine, bacteriopheophytin (BPh), and paclitaxel), an inorganic Ge containing system, and an open-

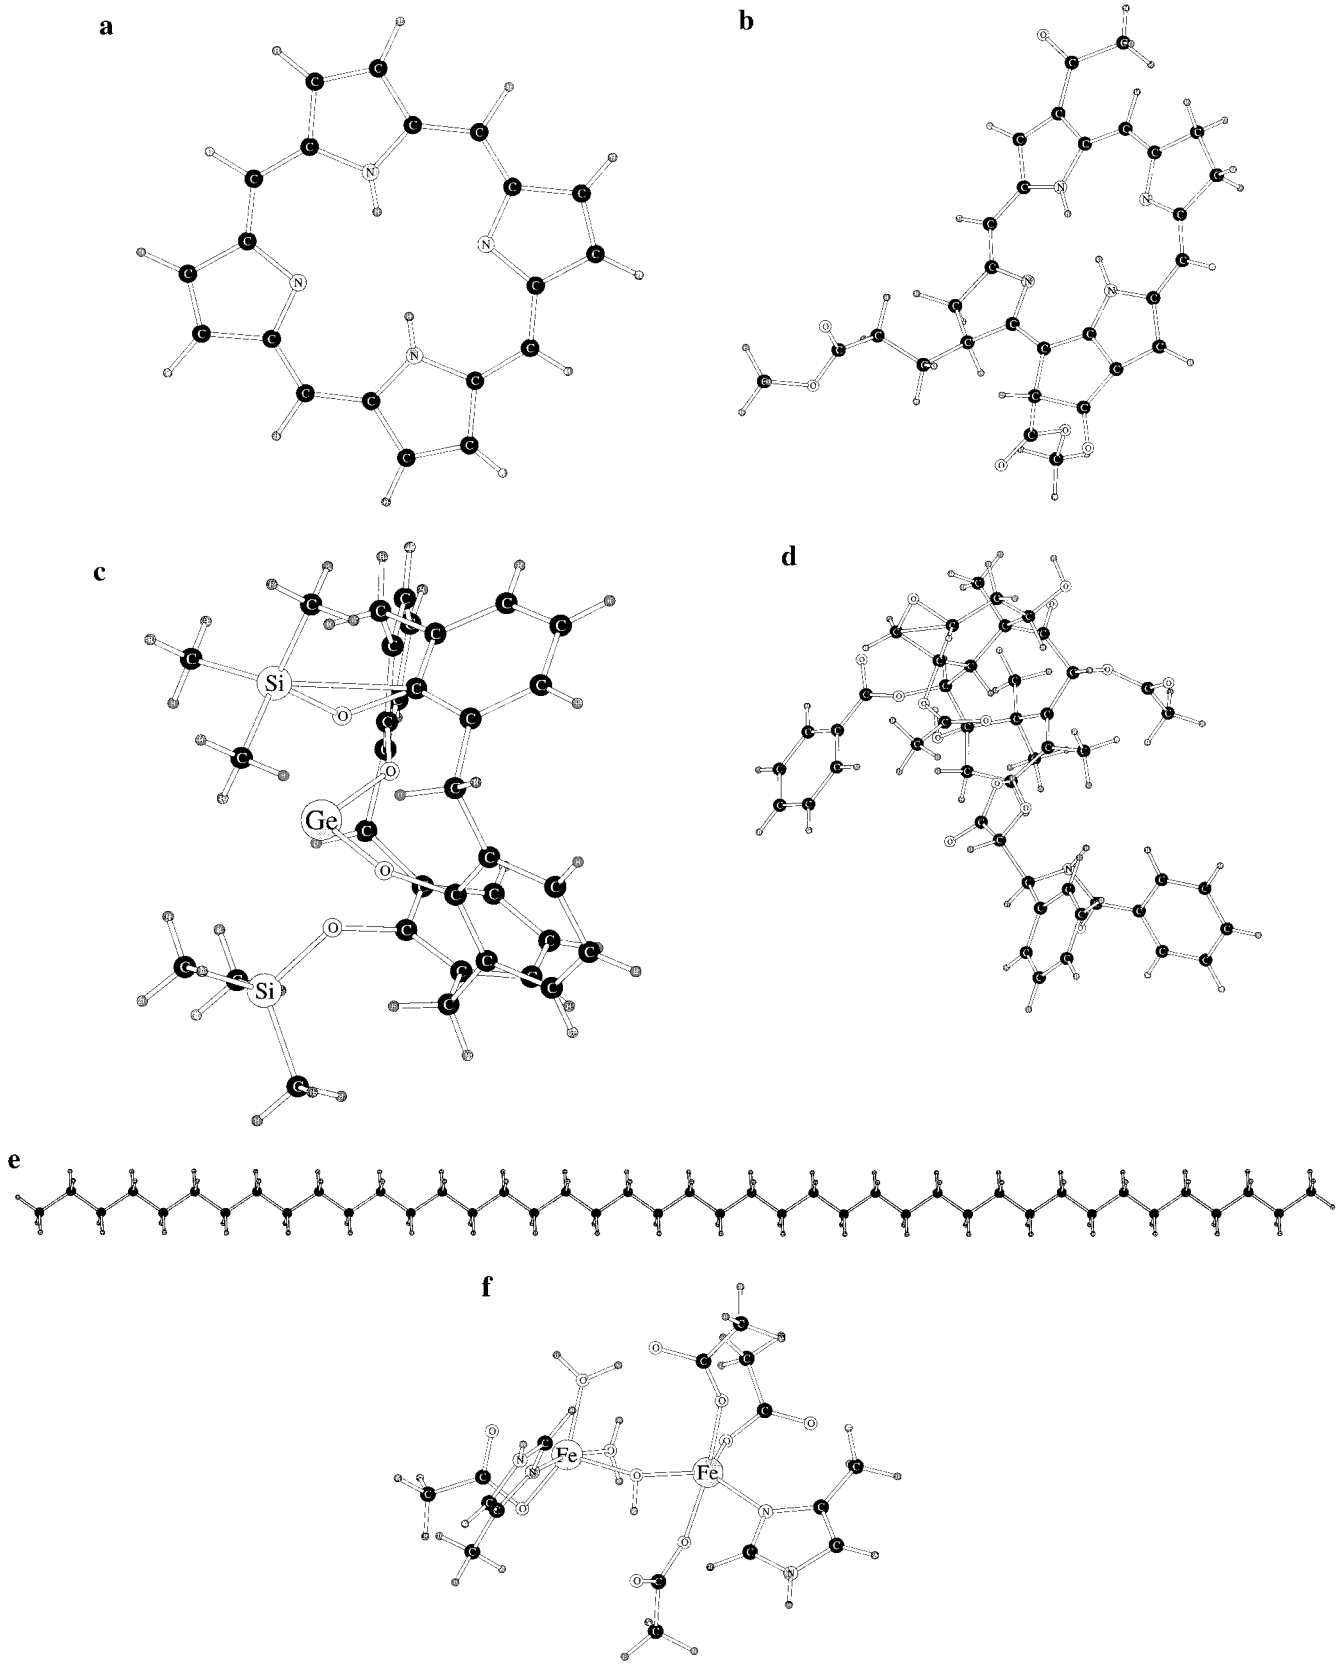

Figure 1. (a) Porphine. (b) Bacteriopheophytin. (c) $C_{34}H_{38}O_4Si_2Ge$. (d) Paclitaxel. (e) $C_{42}H_{48}$. (f) Methane monooxygenase active site model.

shell (spin state 11) diiron system that is a model for the active site of soluble methane monooxygenase (sMMO). As discussed above, we have chosen the B-LYP method as representative of the gradient corrected DFT approaches and B3-LYP as similarly representative of the adiabatic connection methods. Single-point energy and analytical gradient timings are presented in Table 1, and for a subset of molecules, analytical second derivative timings are presented in Table 2.

For the single-point energy and gradient calculations, timings for two basis sets are given: the 6-31G* basis set of Pople and co-workers and the cc-pVTZ(-f) (without f functions) basis of Dunning.[60] The former, a double-ζ plus polarization (DZP) basis

TABLE 1: CPU Times (Minutes) of Single-Point Energy and Average Energy + Gradient Times for Density Functional Calculations with Jaguar v 3.5 on an SGI Irix62-r10k Workstation

molecule	basis set	*N* basis	SCF iterations	SCF-time	average SCF + gradient time
			B-LYP		
porphine	6-31G*	388	10	23	27
sMMO active site	6-31G**	623	19	217	167
$C_{34}H_{38}O_4Si_2Ge$	6-31G*	653	10	95	100
$C_{42}H_{86}$	6-31G*	802	8	46	61
paclitaxel	6-31G*	1032	11	207	216
			B3-LYP		
porphine	6-31G*	388	10	39	43
sMMO active site	6-31G**	623	17	337	259
$C_{34}H_{38}O_4Si_2Ge$	6-31G*	653	9	179	210
$C_{42}H_{86}$	6-31G*	802	8	78	105
paclitaxel	6-31G*	1032	10	421	472
			B-LYP single point		
porphine	cc-pVTZ(-f)	678	6	106	
BPh	cc-pVTZ(-f)	1172	7	616	
$C_{42}H_{86}$	cc-pVTZ(-f)	1740	8	394	
paclitaxel	cc-pVTZ(-f)	1885	10	1428	
			B3-LYP single point		
porphine	cc-pVTZ(-f)	678	5	191	
BPh	cc-pVTZ(-f)	1172	7	950	
$C_{42}H_{86}$	cc-pVTZ(-f)	1740	6	546	
paclitaxel	cc-pVTZ(-f)	1885	8	2674	

TABLE 2: CPU Times (Hours) for Density Functional Frequency Calculations with Jaguar v 3.5 on an SGI Irix62-r10k Workstation. All Molecules Were Run with C1 Symmetry

molecule	*N* basis	basis	method	time
porphine	388	6-31G*	B-LYP	27.2
porphine	388	6-31G*	B3-LYP	33.5
BPh	656	6-31G*	B-LYP	114.9

TABLE 3: Comparisons of CPU Times for Jaguar and Gaussian 92. All Timings Are in Hours Using a 125 MHz HP 735 Workstation

molecule	basis set	basis functions	method	Jaguar	Gaussian 92
alanine tetrapeptide	6-31G*	344	B-LYP	0.7	9.2
			B3-LYP	1.2	9.0
porphine	6-31G*	388	B-LYP	1.0	13.0
			B3-LYP	1.6	11.8
porphine	cc-pVTZ(-f)	678	B-LYP	4.7	84.4
			B3-LYP	9.6	81.9
paclitaxel	6-31G*	1032	B-LYP	1.0	25.0
			B3-LYP	1.7	25.0

set, is typically used for geometry optimizations and for approximate calculation of relative energetics. The latter, containing highly contracted functions at the triple-ζ, double polarization (TZDP) level, provides high quality thermochemical energy differences. All calculations were converged to standard Jaguar convergence (less than 5.0×10^{-5} hartree energy change between iterations, and less than 5.0×10^{-6} RMS density matrix element change between iterations). Single-point timings begin the SCF iteration cycle from the standard initial guess routines in Jaguar. While the cc-pVTZ(-f) calculations used converged 6-31G* wave functions as initial guesses, the times presented include the time for 6-31G* convergence. All molecules were run with C1 symmetry except for $C_{42}H_{48}$, which was run with C_s symmetry. The LACVP** basis was used for Fe and the LACVP* basis was used for Ge.

The total CPU time for geometry optimization depends upon the initial guess and the quality of the geometry optimizer as well as the average computation effort to evaluate a single gradient cycle. Because our focus in the present paper is on the efficiency of the electronic structure algorithms, we therefore present only average CPU timings for a gradient and energy cycle for each molecule considered. We should point out that obtaining convergence with any sort of approximate numerical method (which includes any of the numerical integration techniques used in DFT, as well as pseudospectral technology) requires considerable care in the design of the optimizer, as some level of noise is invariably present in the gradient. We have invested a considerable effort in ensuring that this problem is handled in our code, the details of which will be presented elsewhere. Here, we note only that preliminary results indicate that the number of geometry optimization steps required in our implementation is comparable to that observed in analytically based programs such as Gaussian.

A number of interesting observations can be made concerning the timing results presented in Table 1. First, the B-LYP calculations are roughly twice as fast as the B3-LYP calculations for the same molecule. This is due principally to two factors that differentiate the B-LYP computations: elimination of the Hartree−Fock exchange assembly and use of the fast multipole routines to assemble the Coulomb operator. For ultralarge systems (which we have not examined here), the latter factor would lead to significantly larger timing differentials. Second, there is a substantial dependence of the timing results not only on molecular size and number of basis functions, but also on molecular shape: the CPU time for $C_{42}H_{86}$, a linear carbon chain, is substantially less than that for paclitaxel (a quasi, three-dimensional molecule) despite the fact that the number of basis functions in each calculation are comparable. This is due to the increased number of neighboring basis functions at higher dimensionality that are not eliminated by cutoffs. Such effects must be incorporated into any scaling analysis of performance (which we do not attempt here). Third, the average SCF + gradient time is comparable to the single point CPU time. This apparently counterintuitive result is explained by the fact that for a typical geometry optimization step, the wave function from the previous geometry is used as an initial guess, so that the number of SCF iterations is roughly half of what is required when converging from the initial guess.

Timings for second derivative calculations are presented in Table 2. The vibrational frequencies obtained from our analytical second-derivative methods uniformly agree with standard

methods[61] to within 1−2 cm^{-1}. Details of this methodology will be described elsewhere.

The most important point of the data in Tables 1 and 2 is that DFT single-point, geometry optimization, and frequency calculations can now be carried out routinely for large molecules using high quality basis sets with a relatively inexpensive workstation. The widespread extension of high quality quantum chemical methods to systems in the 30−200 atom range is therefore possible. While calculations on ultralarge systems, facilitated by fast multipole methods, are certainly feasible, and represent an important research direction, there is a lot of interesting chemistry that is now accessible in these more modest regimes that was previously out of range of most laboratories. The next 5 years will see the exploitation of this capability, in both our codes and those of others, and we believe this will have a qualitative impact upon the usefulness of ab initio methods, for example in problems of industrial interest.

In Table 3, we present comparisons of computational efficiency of the DFT methods in Jaguar with Gaussian 92/DFT. Note that we have been unable to make comparisons with Gaussian 94 due to lack of access to this program. The Jaguar timings are approximately 5−15 times faster than those obtained from Gaussian for adiabatic connection based (B3-LYP) calculations and 10−25 times faster for gradient corrected (B-LYP) calculations The greater ratio for the latter arises primarily from the elimination of the exact exchange calculation, which adds relatively more CPU time in PS methods than in conventional ones, and the multipole methodology described above. All of the timings presented, other than for paclitaxel, are for single point SCF calculations, using the default single-point cutoffs of 5×10^{-5} hartrees in both programs. The "fine" DFT grid option in Gaussian was employed which is necessary for accurate energies of large molecules (a comparable setting is used in Jaguar for the DFT grid). For paclitaxel, we were unable to converge Gaussian using the default cutoffs. The timings shown are for tight cutoffs and represent average CPU time per SCF iteration, as opposed to total SCF time. We have reported the data in this fashion because when tight cutoffs are imposed, Gaussian requires a considerable number of additional SCF iterations, which would bias the comparison. Finally, we have run one comparison of average SCF plus gradient timings, for the alanine tetrapeptide, resulting in CPU ratios of 8.0 for B3-LYP and 15.8 for B-LYP, which are comparable to the single-point ratios presented.

4. Localized Orbital Methods. Computational Implementation

4.1. Local MP2. The key to reducing the computational scaling with system size of LMP2 calculations lies in the construction of the two electron integrals over molecular orbitals, $(ip|jp)$ in eq 16 below. In conventional quantum chemical methods, this is accomplished via a four, index transform from integrals over AO basis functions $(\mu\nu|\alpha\beta)$

$$K_{ij}^{pq} = (ip|jq) = \sum_{\mu\nu\alpha\beta} C_{\mu i} C_{\nu p} C_{\alpha j} C_{\beta q} (\mu\nu|\alpha\beta) \qquad (16)$$

which is in practice carried out in four steps:

$$K_{i\nu\alpha\beta} = \sum_{\mu} C_{\mu i} (\mu\nu|\alpha\beta) \qquad (17)$$

$$K_{ij\nu\beta} = \sum_{\alpha} C_{\alpha j} K_{i\nu\alpha\beta} \qquad (18)$$

$$K_{ijp\beta} = \sum_{\alpha} C_{p\alpha} K_{ij\alpha\beta} \qquad (19)$$

$$K_{ij}^{pq} = \sum_{\beta} C_{\beta q} K_{ijp\beta} \qquad (20)$$

The formal scaling of this algorithm is nN^4, where n is the number of occupied orbitals to be correlated and N is the basis set size. Asymptotically, cutoffs can reduce the first step of the transform to nN^2 and the second step to n^2N^2. However, the asymptotic limit is reached very slowly in practice, particularly for large basis sets, which are necessary if highly accurate evaluation of the correlation energy is to be carried out. The use of a localized virtual space has no impact on the first two steps of eqs 17 and 18, although it can be used to reduce the scaling of the last two steps. These are typically smaller in magnitude in practical problems, however, due to the fact that $n \ll N$.

In the PS formulation of LMP2,[13] the two electron matrix elements over MOs are computed directly via the formula:

$$(ip|jq) = \sum_{g} Q_i(g) R_p(g) A_{jq}(g) \qquad (21)$$

In a canonical MP2 context, this has a formal n^2N^3 scaling, no better than that of conventional methods. However, LMP2 allows the virtual indices pq to be restricted to the size of the localized correlation space N_v, leading to a scaling of n^2NN_v2. Since the last of these does not grow with system size M, the overall formal scaling is $\sim M^3$. Formation of the intermediate quantities in eq 21, such as $A_{ij}(g)$, also scales in this regime. In practice, the algorithm does somewhat better than the formal analysis, scaling as $\sim M^{2.5}$.

Once the two electron integrals over MOs are formed and stored on disk (disk storage is drastically minimized due to the local correlation space restriction; the scaling with basis set size is only N^2), eq 4 must be solved iteratively to yield the first-order wave function. Conveniently, this process also has a formal scaling of M^3, and typically converges quite rapidly (6−8 iterations) if appropriate numerical techniques, again based upon methods suggested by Pulay and co-workers,[9,10] are implemented.

The numerical accuracy of the PS implementation of LMP2 is a serious issue, which we have examined in considerable detail. As in the case of HF or DFT calculations, PS methods need to be supplemented with analytical calculation of the largest terms. Here, those terms are selected on the basis of which functions have large amplitudes in the appropriate localized wave functions. Furthermore, it is necessary to optimize the grids and auxiliary fitting basis to yield accurate results for the terms that appear in MP2, as opposed to HF, calculations. Testing of the methodology has been carried out by comparisons with the MOLPRO code of Werner and co-workers[62] for a series of test cases that include intermolecular binding energies and conformational energy differences. For a total of 100 such test cases, errors are typically less than 0.03 kcal/mol. This, along with the successes in comparing with experimental data described below, leads to a high degree of confidence in the robustness of the numerical methods. In contrast, there have been other proposals for accelerating MP2 calculations via numerical methods (all of which involve auxiliary fitting basis sets)[63] where a few reasonable test results for small molecules have been presented, but systematic demonstration of the reliability of the numerical approximations have not yet been given. In our experience, a large suite of comparisons is essential

TABLE 4: CPU Times (minutes) for Pseudospectral HF and LMP2 Compared to Gaussian 92 HF and MP2. N_{bas} is the Number of Basis Functions and N_{occ} is the Number of Occupied Orbitals Correlated

molecule	N_{bas}	N_{occ}	T-G92HF	T-PSHF	T-G92MP2	T-LMP2
			cc-pVTZ basis IBM-370[a]			
1,3-butadiene	146	11	25	10	50	15
methyl acetate	169	15	51	18	134	23
butanone	197	15	72	35	210	30
piperidine	237	18	170	43	1037	55
			6-31G** basis IBM-580[b]			
alinine	160	23	13	5	34	17
leucine	200	27	26	10	90	30
arginine	250	35	37	17	234	50
Si_4Me_{10}	392	43	120	30	1740	154
porphine	430	56	167	60	1800	270

[a] G92-MP2 and PS-LMP2 scaling exponents for this set are 4.9 and 2.6, respectively. [b] G92-MP2 and PS-LMP2 scaling exponents for this set are 4.0 and 2.8, respectively.

TABLE 5: Pseudospectral LMP2 Parallel Performance. Times in Seconds, with Speedups in Parentheses Following[a]

molecule	basis set	basis functions	1	2	4	8	16
caffeine	6-31G**	260	2388(1.0)	1250(1.9)	681(3.5)	372(6.4)	225(10.6)
alanine pentapeptide	6-31G**	510	26409(1.0)	12992(2.0)	6535(4.0)	3353(7.9)	1827(14.5)
caffeine	cc-pVTZ(-f)	412	11225(1.0)	5584(2.0)	2876(3.9)	1560(7.2)	912(12.3)
alanine tetrapeptide	cc-pVTZ(-f)	658	45110(1.0)	23797(1.9)	11527(3.9)	6041(7.5)	3398(13.3)
alanine decapeptide	cc-pVTZ(-f)	1010			82108(1.0)	41259(2.0)	21662(3.8)

[a] Speedups are in comparison to the one-node case for all except alanine decapeptide, which is in comparison to the four-node case.

TABLE 6: CPU Times (Minutes) for Pseudospectral GVB (T-GVBg for 6-31G** Basis, T-GVBcc for cc-pVTZ(-f) Basis Using the 6-31G** Initial Guess), Exchange Integral Generation (T-Kij), and Iterative Solve (T-Solv) on a Single IBM-SP2 390 Thin Node**

molecule	N_{bas}	N_{pair}	T-GVBg**	T-GVBcc	T-Kij	T-Solv
alanine dipeptide	338	29	222	793	213	393
methylcyclohexane	287	21	76	550	107	103
cyclohexane	246	18	43	340	67	81
methylvinyl ether	146	12	12	83	16	30

TABLE 7: Basis Size, CPU Hours (SGI-R10000 Single Processor), Disk Usage (Gigabytes) for J2 Components of the Cage C_{20} Isomer

method	basis functions	CPU time (h)	disk use (GB)
LMP2 cc-pVTZ(-f)	460	5.7	0.7
LMP2 ccpVTZ++	1040	59.8	7.2
GVB-LMP2 cc-pVTZ(-f)	460	95.5	4.5

before the method can be used on real problems where the answer is not known in advance.

Once the numerical accuracy of the method is established, the computational efficiency can be examined. In addition to the single-node workstation version of LMP2, a parallel implementation has been developed[64] for the IBM SP2, which allows calculations to be carried out on very large systems. Tables 4−7 present timing results for the serial and parallel versions, examining a range of molecular sizes and basis sets and comparing with an efficient canonical MP2 implementation in Gaussian 92.[34] Even for a relatively small molecule, such as piperidine (237 basis functions), a factor of 10 reduction in computational effort is obtained as compared to a Gaussian 92 calculation. For larger molecules, using estimations based upon the scaling with basis set size established in ref 13 and presented in Table 4, a CPU ratio of 2−3 orders of magnitude is approached rapidly. These results demonstrate that new areas of chemistry−molecules in the 30−100 atom range−are opened up to the MP2 level of theory with large basis sets. The utility of this level of theory for various chemical problems will be examined below.

4.2. LMP2 Gradients. We have implemented[15,65] the LMP2 gradients associated with PS-LMP2 described above. The gradient formulation follows the same principle of producing local derivative integrals and the right-hand side of the CPHF equation as in the energy formulation. In addition, the pseudospectral CPHF equations have been implemented using PS-HF and PS-HF gradient techniques. LMP2 gradients have the unique capability of optimizing systems for which BSSE effects can obscure the correct trends, such as in intermolecular geometries.

4.3. GVB-LMP2. We have developed[14] the formalism of GVB-LMP2 based upon a synthesis of the theoretical developments of Pulay and co-workers.[10,39] First, a GVB-PP wave function is obtained using the usual SCF algorithms; orbitals not correlated at the GVB level are then localized via Boys or Pipek−Mezey localization. Then, doubly excited configurations starting from the GVB-PP wavefunction are constructed. Each of these configurations contains 2^N determinants; this large scale internal contraction of the wave function retains the efficiency of the LMP2 method while yielding higher accuracy. Assembly of the matrix elements among these excited GVB configurations is more complicated than for the usual HF configurations, but does not greatly add to the computational cost per integral. Similarly, the equations for the first-order wave function, again solved iteratively, are analogous in structure to those of LMP2, although different in details. The major increases in computational expense as compared to LMP2 come from (a) the necessity to compute the GVB-PP wave function rather than HF, (b) the doubling of the number of occupied orbitals, due to the fact that in GVB-PP methods there are two orbitals per electron pair rather than one, (c) the greater number of iterations required to converge the solver, due to the higher complexity of the wavefunction. These factors lead to a factor of 5−10 slowdown, as is shown in Table 6, and an increase in the scaling

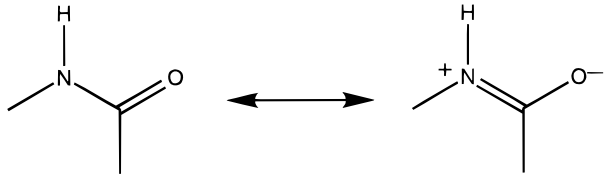

Figure 2. Resonance in the amide group occurs via the transformation of the nitrogen lone pair to a bond pair N−C and subsequent charge transfer to the oxygen of the carbonyl group.

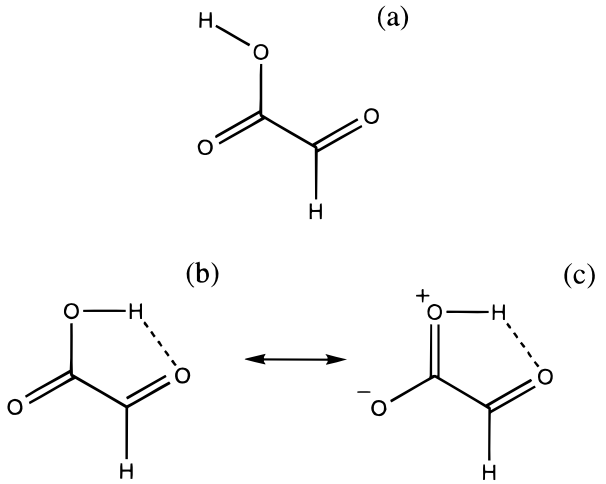

Figure 3. Two conformers of glyoxylic acid. The hydrogen bonded ground state. (c) A resonance structure of (b) which stabilizes the internal hydrogen bond.

from $N^{2.5}$ to N^3. Timing and and disk usage for GVB-LMP2 calculations using 460 basis functions are provided in Table 7.

A significant improvement in the theory can be made by replacing the GVB-PP reference wave function with the superior GVB-RCI wave function,[41] in which excitations within GVB pairs are allowed. GVB-RCI is particularly important in dealing with open-shell systems as it allows correct treatment of spin couplings. Work on this version of the theory is currently in progress.

4.4. Resonance and Delocalization of the Virtual Space. Localized orbital methods work best when chemical bonds and lone pairs are clearly defined for every electron in the molecule. However, some chemical structures are best thought of as containing an admixture of several bonding structures. Benzene is an obvious example, but resonance is ubiquitous in chemistry. For example, the amide (peptide) group has a strong resonance with a quaternary nitrogen configuration as is shown in Figure 2; the carboxylic acid group has a less important resonance depicted in Figure 3. Finally, transition state structures possess inherent ambiguity concerning which atoms are bonded to which other ones. In such cases, one has to be extremely careful to make sure that localized orbital methods provide an accurate description of electron correlation, and also that this description remains constant for all calculations. Furthermore, whatever schemes are implemented have to be automated, otherwise the methods will be accessible only to expert users. What is really needed is an expert system, able to examine any chemical structure and make the appropriate modifications of the theory.

From a number of numerical experiments, we have developed a set of heuristic rules concerning how to effectively deal with resonance and other forms of delocalization. In some cases, for example, the benzene ring, the results are surprisingly unaffected by the use of a localized virtual space, despite the strength of the underlying resonance. For example, a geometry optimization of benzene using local MP2 methods breaks symmetry, localizing the double bonds and yielding alternating unequal bond lengths. However, the difference is quite small. Furthermore, if the asymmetric structure is substituted into a canonical MP2 program (which preserves symmetry), the energy difference from the symmetric structure is less than 0.05 kcal/mol. In a real system, the benzene ring would be surrounded by other molecules, creating an inhomogeneous field, and hence breaking symmetry.

In other cases, proper treatment of resonances is crucial. A good example is the lone pair on the nitrogen in the peptide group (Figure 2). In any localized orbital implementation, it is necessary to examine each occupied localized orbital and define the atoms to which it is localized. Some sort of numerical criterion has to be used to distinguish bonds from lone pairs. We examine the ratio of the wavefunction amplitude on the two atoms with the largest coefficients. However, for the nitrogen "lone pair" orbital, this ratio changes by a factor of ∼10 as one changes the structure of the peptide group from planar (where the resonance form in Figure 2 is important) to nonplanar (where it is less important). If the virtual space associated with an orbital changes as a function of geometry, qualitative inconsistencies are introduced into the relative energetics. Consequently, it is necessary to define this orbital as always delocalized between the C and N atoms, independent of geometry. This is easily done with software that recognizes chemical groups. The delocalization to the peptidic carbon is computationally inexpensive and does not introduce superposition error, as the atoms are at fixed distance in different conformations. The major difficulty is identifying the set of chemical groups for which problems exist and taking the appropriate action. While other methods for automatically choosing the localization scheme exist,[66] they have not been tested extensively on a wide variety of molecules, and we do not believe that they would be immune to the sort of problem discussed here.

The most difficult cases for which to specify localization are transition state structures, where we have found it essential to delocalize bonds that are being broken and formed over a minimum of three centers. Again, the key point is that the localization scheme must remain consistent between any calculations comparing relative energies. In some cases, the relevant delocalization patterns can be identified by numerical amplitudes. However, this is still an ongoing area of research. The results we have obtained to date for transition states follow well defined rules and hence indicate that development of a successful scheme is possible. However, there is no guarantee that what has been developed will work for every new case that comes up. This is the price that one pays for using localized orbital approaches, with their attendant benefits.

For GVB-LMP2, additional difficulties are manifested; these are discussed in detail in ref 14. Even greater delocalization of the virtual space than in LMP2 is necessary to obtain an accurate treatment of hydrogen bonding and resonance. Again, automatic protocols based upon chemical groups have been developed, and have been successful to date as is shown below. However, exploration of these effects is just beginning and there will no doubt be more surprises to come in the future.

5. Benchmark Studies of Chemical Energetics: A Comparison of DFT, LMP2, and GVB-LMP2

The standards for assessment of quantum chemical methods against experimental data were established by the Pople group more than 20 years ago. First, Pople insisted upon using

quantum chemical methods that could be defined as a model chemistry:[67] an automatic protocol that could be applied to an arbitrary molecule without further human intervention. Methods such as Hartree−Fock, Møller−Plesset perturbation theory, coupled cluster based methods, and the various versions of density functional theory, in conjunction with a specified basis set, clearly qualify as model chemistries. More complex procedures built out of these methods, such as G2 theory[16,37] or the CBS procedures of Petterson and co-workers,[68] have also been proposed as model chemistry, although the quantum chemical calculations are supplemented by a small number of adjustable parameters.

Given a model chemistry, the next step is to evaluate performance for a well defined chemical property, by comparing with experimental data. Such tests are invariably carried out for small molecules, as reliable experimental results are easier to come by in quantity for small molecules than for large. A critical mass of data is necessary to draw substantive conclusions, and the experimental data must itself be of sufficient accuracy to make the comparison meaningful.

In this section, we consider two important energetic properties: bond energies and conformational energies. For the first of these we compare[69] to the extended G2 data base that was recently studied with G2 and B3-LYP theories.[16] In addition we have recently studied[70] a larger system, isomers of C_{20}, that we compare to quantum Monte Carlo results. For conformational energies, there is no standard test suite that is as well accepted as the G2 database. However, Halgren and coworkers have recently assembled[71] a 36 molecule database of small and medium sized organic molecules for which experimental data is available. This database can serve the same role as the G2 database for conformational energetics, which is to provide an initial assessment of the statistical performance of quantum chemical methods. We use this database as a standard in the results presented below.

5.1. Thermochemistry: The J2 Model. The calculation of heats of formation poses a challenge to quantum-chemistry (QC) methods. The G2 theory uses quadratic configuration interaction (QCI)[7] and several basis set extension calculations at various levels of Møller−Plesset (MP) perturbation theory. The G2[16,37] theory uses experimental thermochemical values to generate enthalpies of formation to be tested against reliable experimental data. The current G2 data base[16] consists of 148 molecules, which were used to test G2 theory as well as various DFT methods.

The Jaguar-2(J2) model uses the localized perturbation approach, which scales computationally with system size much lower than the calculations involved with the G2 theory. In the J2 theory, atomization energies are calculated using the generalized valence bond−localized Møller−Plesset (GVB-LMP2) wave function.[14] GVB-LMP2 uses the GVB-PP wave function as the reference for the localized perturbation(LMP2)[9,10] treatment. We have implemented GVB-LMP2 in conjunction with the pseudospectral(PS) quantum mechanical methodology. The PS version of a GVB-LMP2 calculation scales with system size as N^3 as opposed to the MP4 and QCISD(T) calculations of the G2 theory which scale as N^7. In addition, the J2 theory has a single basis set extension calculation using LMP2.

In order to obtain reasonable agreement of the theory with experiment, all the theories examined here need to use a parameterized empirical correction. G2 theory uses a single correction parameter for each electron pair. In J2 theory the use of GVB-LMP2 introduces a somewhat bigger imbalance for treatment of the σ bond relative to π bond. The π bonding is underestimated relatively to σ bonding. This imbalance suggests using an additional correction for π bonds. Thus J2P3 theory uses a three-parameter correction composed from σ bond and π bond parameters and an additional parameter to account for the difference of lone pairs between the molecule and the separated atoms. This correction method, which has a total of three adjustable parameters, was first suggested by Martin.[72] An extended correction scheme, based on the results, is introduced in our final J2 theory.[69] In this extended correction scheme 2 additional parameters are being used(J2P5). The J2P5 offers a significant improvement of the statistical performance. The additional parameters have been introduced by careful examination of the J2P3 results. The two parameters account for further imbalance of treatment of bonds in strained systems compared to "regular" bonds. Thus, the J2P5 model introdues a correction by using a parameter for carbon-carbon σ bond in three-member rings and a parameter for π bonds in three-, four-, and five-member rings. These are well-defined parameters.

The J2 theory involves a GVB-LMP2/cc-pVTZ(-f) calculation. A single-basis-set correction term is added by a LMP2 calculation at the cc-pVTZ++ level. The J2 energy is given by

$$E_{J2} = E[\text{GVB} - \text{LMP2/cc} - \text{pVTZ}\ (-\text{f})] + E[\text{LMP2/cc} - \text{pVTZ} + +] - E[\text{LMP2/cc} - \text{pVTZ}\ (-\text{f})] + \text{EC} \quad (22)$$

EC is a parameterized correction term, given by

$$\text{EC}\ (\text{n1, n2, n3; } P_{\text{pair}}, P_\sigma, P_\pi) = \text{n1}*P_{\text{pair}} + \text{n2}*P_\sigma + \text{n3}*P_\pi \quad (23)$$

where P_{pair}, P_σ, and P_π are 3 correction parameters, and n1, n2, and n3 describe the frequency of appearance of the pair types in the system considered. The atomization energy is obtained by

$$\text{Energy (molecule)} = E_{J2}\ (\text{molecule}) - \sum E_{J2}\ (\text{atoms}) \quad (24)$$

The J2 performance offers a significant improvement in accuracy over the current version of G2 theory, despite qualitatively lower computational cost. The deviation from experiment in enthalpies of formation for the J2, G2, and B3-LYP theories are compared in Table 9. The B3-LYP is the most accurate DFT method reported in the G2 reference. Becke[27] has recently optimized a generalized gradient-corrected exchange−correlation functional to the G2 test set. The new functional was shown to perform better than the B3-LYP, with a MAD of 1.78 kcal/mol and maximum deviation of 8.9 kcal/mol compared to 1.58 and 8.2 of the G2 theory; however, we do not present these results here.

Tables 8 and 9 summarize the performance of three different computational methods: G2, J2[69], and B3-LYP for calculating heats of formations (at 298K) from atomization energies as applied to the extended G2 data base.[16] The standard enthalpies of formation and temperature factors were taken from ref 16. We present results for 67 closed-shell molecules containing first row atoms. Preliminary results indicate that the cc-pVTZ basis sets used in our method need to be enhanced for the second row. Also only closed shell systems have been considered, because of the use of GVB-PP. (The COF_2 molecule has been excluded from the considerations, attributing the anomalously large errors displayed by all methods for this molecule to poor experimental data.) The GVB-PP uses only the PP spin eigenfunction(SEF), the inclusion of other SEF is important for

TABLE 8: Errors in Heats of Formation (kcal/mol) as Calculated by G2, J2, and B3-LYP Theories

	energy deviation (kcal/mol)			
molecule's name	J2P3	J2P5	G2	B3LYP
CH_4	0.54	0.39	0.7	1.6
NH_3	0.58	0.91	−0.2	3.5
H_2O	0.75	0.97	0.3	−1.3
HF	1.52	1.63	1.0	−1.6
C_2H_2	0.73	0.36	−1.6	−2.5
C_2H_4	1.93	1.55	−0.2	0.6
C_2H_6	0.65	0.26	0.5	0.6
HCN	0.43	0.55	0.3	0.0
H_2CO	0.46	0.45	2.0	0.4
CH_3OH	−1.30	−1.34	1.4	0.1
N_2	−0.69	−0.10	−1.3	−1.4
N_2H_4	−0.73	−0.17	−0.9	6.3
H_2O_2	0.71	1.17	−0.2	−1.8
F_2	1.45	1.69	−0.3	−2.6
CO_2	−0.14	−0.03	2.7	−0.2
CF_4	0.78	0.64	5.5	−4.5
NF_3	−2.49	−1.91	3.7	4.0
OF_2	−1.15	−0.67	0.5	−0.4
C_2F_4	2.38	2.01	8.2	3.2
CF_3CN	−0.25	−0.39	4.8	−3.7
propyne	0.99	0.37	−1.5	−1.9
allene	2.10	1.48	−0.9	1.9
cyclopropene	−2.13	0.65	−2.9	−3.2
propene	1.37	0.74	−0.5	−0.6
cyclopropane	−0.91	0.16	−0.9	−2.2
propane	0.36	−0.29	0.4	−1.5
trans-1,3-butadiene	0.81	−0.06	−1.7	−1.5
dimethylacetylene	−0.05	−0.92	−2.1	−2.4
methylenecyclopropane	1.29	2.13	0.3	0.0
bicyclo[1.1.0]butane	−3.64	−1.70	−3.0	−7.1
cyclobutene	−1.53	−0.84	−2.9	−6.1
cyclobutane	−0.18	−1.19	−0.2	−5.2
isobutene	0.88	−0.01	−0.6	−3.1
trans-butane	0.03	−0.87	0.4	−3.7
isobutane	−0.18	−1.08	0.3	−4.8
spiropentane	−2.57	−0.28	−1.4	−5.4
benzene	1.41	−0.06	−3.9	−4.5
difluoromethane	1.16	1.02	3.1	0.0
trifluoromethane	1.24	1.10	4.3	−2.2
methylamine	1.07	1.16	0.0	3.2
acetonitrile	0.48	0.33	−0.1	0.6
methylnitrite	0.43	0.75	2.7	1.3
formic acid	−0.46	−0.36	2.0	−0.9
methyl formate	0.93	0.77	3.8	−0.2
acetamide	−1.41	−1.46	0.2	1.6
aziridine	−0.66	−0.34	−0.3	1.0
cyanogen	−1.61	−1.37	−1.5	−0.4
dimethylamine	−0.49	−0.66	0.3	2.0
trans-ethylamine	1.05	0.88	0.8	2.2
ketene	−0.81	−1.07	0.8	2.4
oxirane	1.49	1.70	1.3	−1.4
acetaldehyde	0.66	0.39	1.3	−0.3
glyoxal	1.50	1.36	2.9	−1.6
ethanol	0.89	0.60	1.0	−1.9
dimethylether	−0.02	−0.31	2.0	0.0
vinyl fluoride	1.33	0.95	1.7	1.5
cyanoethylene	−1.50	−1.88	−2.7	−2.0
acetone	1.01	0.49	1.1	−2.0
acetic acid	0.13	−0.02	1.5	−2.6
acetyl fluoride	−1.20	−1.47	2.0	−1.5
isopropyl alcohol	0.93	0.40	1.2	−4.5
methylethyl ether	−0.64	−1.18	2.3	−1.5
trimethyl amine	−0.53	−0.95	1.4	0.2
furan cyclic	−1.15	1.37	−1.0	−4.2
pyrrole planar	−3.90	−1.28	−2.2	−0.8
pyridine cyclic	2.42	1.42	−2.2	−0.2
H2	0.86	0.97	1.1	1.0

open-shell systems. The restricted configuration interaction(RCI) includes all relevant spin couplings.[41] It is the intention to use GVB-RCI-LMP2 instead of the GVB-LMP2 in order to extend the applicability to the open shell systems. It is predicted to improve the results of the closed shell cases as well. Work in this direction is currently in progress.

TABLE 9: J2, G2, and B3-LYP Average and Maximum Deviation (kcal/mol) from Experimental Heats of Formation. The COF_2 Case, where Experiment is Questioned, Has Been Excluded

method	average deviation	maximum deviation
J2P3-GVB-LMP2	1.07	3.9
J2P5-GVB-LMP2	0.87	2.1
G2-QCISD(T)	1.63	8.2
B3-LYP	2.10	7.1

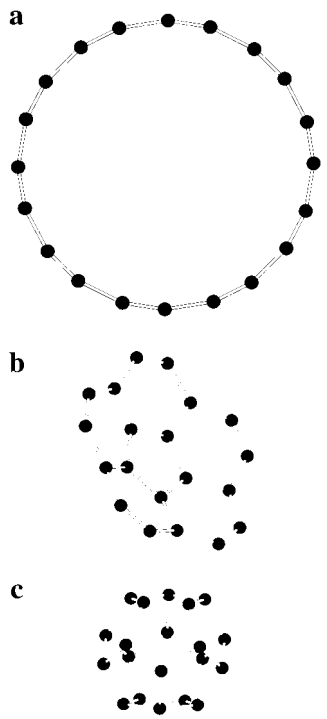

Figure 4. (a) Ring isomer of C_{20}. (b) Bowl isomer of C_{20}. (c) Cage isomer of C_{20}.

TABLE 10: Relative Energies of C_{20} Isomers (kcal/mol), Bowl as Reference, among ab Initio Components of the J2 Theory, the J2 Composite Results, QMC, and DFT

method	bowl	ring	cage
HF cc-pVTZ(++)	0.0	−32.6	69.2
LMP2 cc-pVTZ(-f)	0.0	21.9	61.6
LMP2 cc-pVTZ(++/3d1f)	0.0	39.3	42.6
LMP2 cc-pVTZ(++/3d2f)	0.0	41.7	32.8
GVB cc-pVTZ(-f)	0.0	−5.0	99.1
GVB-LMP2 cc-pVTZ(-f)	0.0	23.8	63.0
J2 cc-pVTZ(++/3d2f)	0.0	26.8	51.0
QMC	0.0	23 ± 4.6	52 ± 4.6
B3-LYP cc-pVTZ(-f)	0.0	−5.8	37.4
B-LYP cc-pVTZ(-f)	0.0	−20.9	45.7

The B3-LYP results are reasonably accurate with an average error of 2.1 kcal/mol and maximum error of 7.1 kcal/mol. Relative to correlated ab initio methods, the low cost and ease of applicability across the periodic table are powerful motivations for selecting this approach. However, as we discuss below, the success of B3-LYP may not extend uniformly to larger systems.

The average 1.07 kcal/mol error of the J2P3 method displayed in Table 9 is superior to that of G2 theory, 1.63 kcal/mol, a result that we attribute primarily to elimination of large errors in fluorine compounds arising from BSSE[69] by the localized treatment incorporated in J2 theory. The J2P5 theory introduces

TABLE 11: Relative Conformational Energies (kcal mol) from DFT (NLSDA), HF (HF), local MP2 (LMP2), GVB, and GVB-LMP2 Calculations with a cc-pVTZ(-f) Basis Set

molecule	ZPE correction	NLSDA[a]	HF	LMP2	GVB	GVB-LMP2	exptl
acrolein (c−t)	0.02	2.25	2.37	1.93	1.84	1.89	1.70
1,3-butadiene (g−t)	−0.02	3.78	3.42	2.92	2.05	2.65	2.89
butane (g−t)	0.09	0.76	1.14	0.73	1.09	0.80	0.67
butanone (s−c)	0.20	1.98	1.52	1.30	1.30	1.25	1.07
1-butene (c−s)	0.10	0.22	0.78	0.47	0.94	0.41	0.22
2-butene (c−t)	0.11	1.04	1.72	1.31	1.72	1.56	1.20
chloropropane (a−g)	−0.04	−0.33	−0.38	−0.09	−0.49	−0.23	0.09
cyclohexane (tb−ch)	−0.03	6.12	6.85	6.10	6.17	5.63	5.50
cyclohexanol (ax,C1−eq,C1)	0.14	1.10	0.92	0.57	0.88	0.73	0.58
cyclohexylamine (ax−eq)	0.14	1.60	1.36	0.75	1.27	0.97	1.15
1,2-dichloroethane (g−a)	−0.03	1.51	1.93	1.57	1.78	1.40	1.08
diethyl ether (a,g−a,a)	0.03	1.16	1.86	1.39	1.81	1.54	1.14
1,2-difluoroethane (t−g)	0.01	1.15	0.14	0.56	−0.17	0.50	0.56
2,3-dimethylbutane (g−a)	−0.01	0.35	−0.11	0.02	−0.11	−0.09	0.05
dimethyl dioxane (eq,ax−eq,eq)	0.15	1.31	1.40	0.79	1.24	0.98	0.92
ethanol (g−a)	0.01	−0.33	0.19	0.07	0.22	0.17	0.12
ethyl formate (c,g−c,t)	0.11	0.49	0.80	0.26	1.10	0.56	0.19
fluoropropane (a−g)	−0.07	−0.02	−0.06	0.09	−0.08	0.09	0.35
formic acid (t−c)	−0.23	4.38	4.76	4.36	4.07	4.41	3.90
glyoxylic acid (t,c−t,t)	−0.17	1.25	0.28	0.93	−0.41	0.79	1.20
isoprene (g−t)	−0.23	3.11	2.51	2.56	1.06	2.13	2.65
isopropanol (a−g)	−0.02	−0.09	0.25	0.22	0.26	0.24	0.28
isopropylamine (g−a)	−0.05	0.29	0.28	0.31	0.27	0.30	0.45
methoxycyclohexane (ax,C1−eq,C1)	0.15	0.51	0.92	0.19	0.90	0.43	0.55
2−methoxytetrahydropyran (eq−ax)	−0.24	0.87	0.31	1.20	0.45	1.12	1.05
N−methylacetamide (c−t)	0.19	2.52	2.50	2.28	0.65	2.19	2.30
methyl acetate (c−t)	−0.08	7.12	8.81	7.60	8.11	7.82	8.50
methylcyclohexane (ax −eq)	0.18	2.59	2.66	1.87	2.41	1.94	1.75
methyl ethyl ether (s−t)	0.02	1.26	1.80	1.38	1.75	1.48	1.50
N−methylformamide (c−t)	0.12	1.70	1.01	1.12	1.16	1.43	1.40
methyl formate (t−c)	−0.48	4.25	5.03	4.76	4.93	4.97	4.75
N−methyl piperidine (ax−eq)	0.03	3.18	4.11	3.73	3.99	3.62	3.15
methyl vinyl ether (s−c)	−0.53	1.94	1.02	2.13	−0.07	1.60	1.70
piperidine (ax−eq)	−0.08	0.56	0.74	0.71	0.69	0.72	0.53
propionaldehyde (s−c)	−0.03	0.79	0.85	0.91	0.54	0.77	0.67
absolute average error		0.34	0.49	0.20	0.56	0.21	
RMS error		0.42	0.57	0.26	0.73	0.25	

[a] All values are corrected with HF/6-31G* zero point energies calculated at HF/6-31G* minima. HF, LMP2, GVB, and GVB-LMP2 calculations used the cc-pVTZ(-f) basis; the NLSDA calculations used a mixed double and triple ζ plus polarization basis. Mean absolute and RMS errors exclude methyl acetate, for which the experimental error is ± 1 kcal/mol. Conformation abbreviations are as follows (g) gauche, (a) anti, (t) trans, (c) cis, (s) skew, (eq) equatorial, (ax) axial.

a further improvement with an average error of only 0.87 kcal/mol. The higher reliability of J2 theories over the G2 is demonstrated with the much lower max. deviation. J2P5 theory has a 2.1 kcal/mol for the max. deviation while the G2 theory has 8.2 kcal/mol (J2P3 has 3.9 kcal/mol for maximum deviation). The computational effort of GVB-LMP2 based J2 theory scales as N^3 rather than N^7 in G2 and, hence, as illustrated below, is applicable to large molecules using present day computers. Furthermore, if one is interested in the bond energy of only one bond in the molecule (as opposed to the atomization energy), correlation at the GVB level can be restricted to a small part of the molecule, thus saving considerable computational effort and reducing the scaling further.

We have applied[70] J2 theory to the problem of determining the relative energies of ring, bowl, and cage shaped isomers of C_{20} displayed in Figure 4. Timings for the most expensive case, the cage isomer for which cutoffs are the least effective, are presented in Table 7. The relative energies of the clusters as predicted by J2 theory in Table 10 are in quantitative agreement with the highly accurate quantum Monte Carlo results of ref 73. In contrast, as first pointed out in ref 73 and displayed in Table 10, the DFT methods differ dramatically from these results. The close agreement of the highly correlated J2 and quantum Monte Carlo methods suggests that the DFT methods have substantial errors for this problem.

Further studies are necessary to determine if this large disparity with DFT is due to the size of the system or the treatment of relative correlation between isomers. Preliminary results for the calculation of heats of formation using isodesmic reactions[74] suggest that the problem is that DFT makes large errors for pi bonds. Thus a tentative explanation is that since the C_{20} isomers differ in the number of π bonds the DFT errors become apparent for this problem.

5.2. Conformational Energies. Table 11 presents results for 35 conformational energy differences for molecules in the database assembled by Halgren and co-workers,[71] computed via HF, DFT, LMP2,[13] and GVB-LMP2[14] methods. Since the electron pair number is conserved in all these cases, and the molecules are closed shell, no corrections are required for the GVB-LMP2 calculations. MP2/6-31G* geometries were used for the HF, LMP2, GVB, and GVB-LMP2 calculations, while the DFT results reproduced from ref 71 used DFT geometries. All values in Table 11 have been corrected for zero-point energies using 6-31G* Hartree−Fock frequencies calculated at MP2/6-31G* geometries. Note that the GVB-LMP2 values are slightly different than those published previously[14] since in these results we have delocalized the lone pairs to atoms bound to the lone pair atom. In addition we have refined the grids and dealiasing sets. The experimental numbers used are the same

as those in ref 75 except for chloropropane, for which we chose the newer experimental number of ref 76.

Hartree−Fock calculations for individual molecules are often reasonable, and for molecules of this size are invariably within 1−2 kcal/mol of the experimental results. However, this is not a particularly good performance when the experimental numbers are also of this magnitude. For many applications, such as structure-based drug design, a higher level of accuracy is desirable. Furthermore, for larger molecules, such as peptides, errors in conformational energies of individual functional groups can add, leading to highly erroneous results for a conformational analysis.

The DFT results provide a statistical improvement over HF, but there are still a significant number of large errors (e.g., greater than 0.5 kcal/mol) and some cases in which the sign of the energy difference is incorrect. LMP2 results provide a further improvement when a large basis set is used. 6-31G** LMP2 results are actually no better than Hartree−Fock, but when the cc-pVTZ(-f) basis is used, performance improves noticeably. However, there are still a number of large discrepancies with the experimental data, such as for cyclohexane and methyl vinyl ether. The zero-point-energy-corrected GVB-LMP2 results, however, bring nearly all of the results into good agreement with experiment.

6. Conclusions

When the methodological advances discussed above are combined with the explosive increase in performance of computational hardware, a revolution in the applicability of first principles electronic structure methods to complex chemical problems is well within sight in the next 5 years. Even with the limitations of current functionals, DFT methods are proving to be very useful in a wide variety of applications, for example materials science problems in which estimates of thermochemical pathways are required. For problems where greater precision is required in energetics, DFT may still be the method of choice for carrying out the necessary geometry optimization. Furthermore, we can expect new functionals to be developed that will overcome many of the problems associated with those that are presently available. The size and complexity of molecular structures that can now be addressed via DFT calculations on a routine basis is truly remarkable.

On the other hand, there are a significant number of problems for which current DFT methods are not sufficient accurate to answer the chemical questions of interest. For example, conformational energetics are crucial to modeling biological systems; DFT calculations do not have the required level of precision, and furthermore are qualitatively deficient in that they do not properly compute dispersion interactions. For these problems, localized correlation methods provide a promising and cost effective alternative. Additionally, one can combine the two approaches, for example, by calculating geometries via DFT methods and then performing single-point LMP2 or GVB-LMP2 calculations. This protocol is in fact likely to be the method of choice in many cases, although tests need to be carried out to establish levels of accuracy for specific protocols.

A great deal of additional experimentation needs to be carried out to understand the performance of both types of methods for a wide range of problems and molecules. For example, we have not discussed here the calculation of transition state energetics, a more demanding application than either bond energies or conformational energies, and one which poses qualitative difficulties for all electronic structure methods. There is some evidence that performance for reaction energies, as opposed to atomization energies, does not follow the results presented in section 5.1; it is also far from clear that one can transfer estimations of accuracy from the small molecules in the G2 database to larger systems where there are strong resonance effects, interactions of multiple functional groups, and dispersion contributions to the total energy. Transition metal chemistry has its own challenges, not the least of which is the greater difficulty in obtaining high quality experimental data for benchmark comparisons. All of these challenges will be addressed during the next 5 years, at which point we should have a much better idea of the limits of performance of the methods discussed above, as well as other competing methods.

Acknowledgment. This work was supported in part by grants to RAF from the NIH (Grants RR-06892 and GM-40526) and NSF (Grant ASC-9217368).

References and Notes

(1) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover: Mineola, New York, 1996.
(2) Kutteh, R.; Apra, E.; Nichols, J. *Chem. Phys. Lett.* **1995**, *238*, 173−179.
(3) White, C.; Johnson, B.; Gill, P.; Head-Gordon, M. *Chem. Phys. Lett.* **1996**, *253*, 268−278.
(4) Challacombe, M.; Schwegler, E. *J. Chem. Phys.* **1997**, *106*, 5526−5536.
(5) Strain, M.; Scuseria, G.; Frisch, M. *Science* **1996**, *271*, 51−53.
(6) White, C.; Head-Gordon, M. *J. Chem. Phys.* **1994**, *101*, 6593−6605.
(7) Pople, J.; Head-Gordon, M.; Raghavachari, K. *J. Chem. Phys.* **1987**, *87*, 5968.
(8) Raghavachari, K.; Trucks, G. W.; Pople, J. A.; Head-Gordon, M. *CPL* **1989**, 157, 479−483.
(9) Pulay, P.; Saebø, S.; Meyer, W. *J. Chem. Phys.* **1984**, *81*, 1901−1905.
(10) Saebø, S.; Pulay, P. *J. Chem. Phys.* **1987**, *86*, 914−922.
(11) Ladner, R.; Goddard, W., III *J. Chem. Phys*. **1969**, *51*, 1073.
(12) Bobrowicz, F.; Goddard, W., III In *Methods of Electronic Structure Theory*; Schaefer, H. F., I, Ed.; Plenum: New York, 1977.
(13) Murphy, R. B.; Beachy, M. D.; Friesner, R. A.; Ringnalda, M. N. *J. Chem. Phys.* **1995**, *103*, 1481−1490.
(14) Murphy, R. B.; Pollard, W. T.; Friesner, R. A. *J. Chem. Phys.* **1997**, *106*, 5073−5084.
(15) Jaguar, v 3.5; Schrödinger, Inc.: Portland, OR, 1998.
(16) Curtiss, L. A.; Raghavachari, K.; Redfern, P. C.; Pople, J. A. *J. Chem. Phys.* **1997**, *106*, 1063−1079.
(17) Hohenberg, P.; Kohn, W. *Phys. Rev. A* **1964**, *136*, 864.
(18) Kohn, W.; Sham, L. *J. Phys. Rev. A* **1965**, *140*, 1133.
(19) Ceperley, D. M.; Alder, B. *Phys. Rev. Lett.* **1980**, *45*, 566.
(20) Becke, A. D. *Phys. Rev. A* **1988**, *38*, 3098.
(21) Becke, A. *Can. J. Chem.* **1966**, *74*, 995−997.
(22) Perdew, J. P.; Chevary, J. A.; Vosko, S. H.; Jackson, K. A.; Pederson, M. R.; J., S. D.; Fiolhais, C. *Phys. Rev. B* **1992**, *46*, 6671−6687.
(23) Johnson, B. G.; Gill, P. M. W.; Pople, J. A. *J. Chem. Phys.* **1993**, *98*, 5612−5626.
(24) Lee, C.; Yang, W.; Parr, R. G. *Phys. Rev. B* **1988**, *37*, 785.
(25) Becke, A. D. *J. Chem. Phys.* **1993**, *98*, 1372−1377.
(26) Becke, A. D. *J. Chem. Phys.* **1993**, *98*, 5648−5652.
(27) Schimder, H. L.; Becke, A. D. *J. Chem. Phys.* **1998**, *108*, 9624.
(28) Boys, S. F. Localized orbitals and localized adjustment functions. In *Quantum Theory of Atoms, Molecules, and the Solid State*; Löwdin, P.-O., Ed.; Academic Press: New York, 1968.
(29) Pipek, J.; Mezey, P. G. *J. Chem. Phys.* **1989**, *90*, 4916−4926.
(30) Saebø, S.; Tong, W.; Pulay, P. *J. Chem. Phys.* **1993**, *98*, 2170−2175.
(31) Beachy, M. D.; Chasman, D.; Murphy, R. B.; Halgren, T. A.; Friesner, R. A. *J. Am. Chem. Soc.* **1997**, *119*, 5908−5920.
(32) Shephard, R. *Adv. Chem. Phys.* **1987**, *69*, 63.
(33) Bartlett, R. *J. Phys. Chem.* **1989**, *93*, 1697−1708.
(34) Frisch, M. J.; et al. *Gaussian 92*, Revision A; Gaussian, Inc.: Pittsburgh, PA, 1992.
(35) Stanton, J. F.; Gauss, J.; Watss, J. D.; Lauderdale, W. J.; Bartlett, R. J. *Int. J. Quantum. Chem.: Quantum Chem. Symp.* **1992**, *26*, 879 (ACES II is a quantum chemical program package especially designed for CC and MBPT energy and gradient calculations).

(36) Roos, B. *Adv. Chem. Phys.* **1987**, *69*, 399.
(37) Pople, J.; Head-Gordon, M.; Fox, D.; Raghavachari, K.; Curtiss, L. *J. Chem. Phys.* **1989**, *90*, 5622−5629.
(38) Roos, B. O.; Andersson, K.; Fülscher, M. P.; Malmqvist, P.-Å.; Serrano-Andrés, L.; Pierloot, K.; Merchan, M. *Adv. Chem. Phys.* **1996**, *93*, 219.
(39) Wolinski, K.; Pulay, P. *J. Chem. Phys.* **1989**, *90*, 3647.
(40) Wolinski, K.; Sellers, H. L.; Pulay, P. *Chem. Phys. Lett.* **1987**, *140*, 225.
(41) Murphy, R. B.; Friesner, R. A.; Ringnalda, M. N.; Goddard, W. A., III *J. Chem. Phys.* **1994**, *101*, 2986.
(42) Murphy, R. B. Messmer, R. P. *Chem. Phys. Lett.* **1991**, *183*, 443−448.
(43) Murphy, R. B.; Messmer, R. P. *J. Chem. Phys.* **1992**, *97*, 4974−4985.
(44) Hampel, C.; Werner, H.-J. *J. Chem. Phys.* **1996**, *104*, 6286.
(45) Andzelm, J.; Wimmer, E. *Phys. B* **1991**, *172*, 307−317.
(46) Delley, B. *J. Chem. Phys.* **1990**, *92*, 508−517.
(47) Osinga, V. P.; van Gisbergen, S.; Snijders, J. G.; Baerends, E. J. *J. Chem. Phys.* **1997**, *106*, 5091−5101.
(48) St. Amant, A.; Salahub, D. R. *Chem. Phys. Lett.* **1990**, *169*, 387−392.
(49) Becke, A. D.; Dickson, R. M. *J. Chem. Phys.* **1988**, *89*, 2993−2997.
(50) Friesner, R. A. *Chem. Phys. Lett.* **1985**, *116*, 39−43.
(51) Friesner, R. A. *J. Chem. Phys.* **1986**, *85*, 1462−1468.
(52) Friesner, R. A. *J. Chem. Phys.* **1986**, *86*, 3522−3531.
(53) Friesner, R. A. *J. Chem. Phys.* **1986**, *85*, 3091−3096.
(54) Langlois, J.-M.; Muller, R. P.; Coley, T. R.; Goddard, W. A., III; Ringnalda, M. N.; Won, Y.; Friesner, R. A. *J. Chem. Phys.* **1990**, *92*, 7488.
(55) White, C. A.; Johnson, B. G.; Gill, P. M. W.; Head-Gordon, M. *Chem. Phys. Lett.* **1994**, *230*, 8−16.
(56) Schwegler, E.; Challacombe, M. *J. Chem. Phys.* **1996**, *105*, 2726−2734.
(57) Schwegler, E.; Challacombe, M.; Head-Gordon, M. *J. Chem. Phys.* **1997**, *106*, 9708−9717.
(58) Burant, J. C.; Scuseria, G. E.; Frisch, M. J. *J. Chem. Phys.* **1996**, *105*, 8969−8972.
(59) Ochsenfeld, C.; White, C. A.; Head-Gordon, M. *J. Chem. Phys.* **1998**, *109*, 1663−1669.
(60) Dunning, Jr., T. H. *J. Chem. Phys.* **1989**, *90*, 1007−1023.
(61) Johnson, B. G.; Frisch, M. J. *Chem. Phys. Lett.* **1993**, *216*, 133−140.
(62) MOLPRO is a package of ab initio programs written by H.-J. Werner and P. J. Knowles, with contributions from J. Almlof, R. D. Amos, M. J. O. Deegan, S. T. Elbert, C. Hampel, W. Meyer, K. Peterson, R. Pitzer, A. J. Stone, P. R. Taylor, and R. Lindh.
(63) Bernholdt, D. E.; Harrison, R. J. *Chem. Phys. Lett.* **1996**, *250*, 477−484.
(64) Beachy, M. D.; Chasman, D.; Friesner, R. A.; Murphy, R. B. *J. Comput. Chem.* **1998**, *19*, 1030−1038.
(65) Murphy, R. B.; Cao, Y.; Friesner, R. A. To be published; is currently implemented in the Jaguar code.
(66) Boughton, J. W.; Pulay, P. *J. Comput. Chem.* **1993**, *14*, 736−740.
(67) Pople, J. A.; Binkley, S.; Seeger, R. *Int. J. Quantum Chem.: Quantum Chem. Symp.* **1976**, *10*, 1−9.
(68) Ochterski, J. W.; Petersson, G. A.; Montgomery, J., Jr. *J. Chem. Phys.* **1996**, *104*, 2598−2619.
(69) Dunietz, B. D.; Murphy, R. B., Friesner, R. A. *J. Chem. Phys.* **1999**, *110*, 1921−1930.
(70) Murphy, R. B.; Friesner, R. A. *Chem. Phys. Lett.* **1998**, *288*, 403−407.
(71) St.-Amant, A.; Cornell, W. D.; Kollman, P. A.; Halgren, T. A. *J. Comput. Chem.* **1995**, *16*, 1483−1506.
(72) Martin, J. *J. Chem. Phys.* **1992**, *97*, 5012−5018.
(73) Grossman, J. C.; Mitas, L.; Raghavachari, K. *Phys. Rev. Lett.* **1995**, *75*, 3870−3873.
(74) Raghavachari, K.; Stefanov, B. B.; Curtiss, L. A. *J. Chem. Phys.* **1997**, *106*, 6764−6767.
(75) Halgren, T. A. *J. Comput. Chem.* **1998**. Submitted for publication.
(76) Herrebout, W. A.; van der Veken, B. J. *J. Phys. Chem.* **1996**, *100*, 9671−9677.