Local weak pairs spectral and pseudospectral singles and doubles configuration interaction

Gregg Reynolds, Todd J. Martinez, and Emily A. Carter
Department of Chemistry and Biochemistry, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, California 90095-1569

(Received 6 May 1996; accepted 15 July 1996)

A new approximate correlation method has been developed by application of the local weak pairs approximation of Sæbø and Pulay to pseudospectral singles and doubles configuration interaction (SDCI) as developed by Martinez and Carter. The combination of the localization and pseudospectral approximations attacks both the problems of two-electron integral storage on disk and CI vector storage in memory that, respectively, hinder nondirect local spectral and nonlocal pseudospectral SDCI calculations individually and provides a scaling advantage over even direct local spectral SDCI calculations. The reproduction of total energies to within a kcal/mol leads to speed increases with respect to nonlocal calculations that grow larger with increasing molecular size: little or no savings for ethane and a factor of 1.1–1.6 for larger molecules studied (glyoxal, glycine, C_6H_2, and C_8H_2). The prediction of conformational energy differences with the new method appears quite promising, since energy difference predictions accurate to within a kcal/mol of the exact energy differences are obtained even when the single-point total energies are individually many kcal/mol in error. The speed increases for energy difference predictions of both local spectral and pseudospectral SDCI also grow with molecular size: from a factor of 4 in ethane and glyoxal to a factor of 6 in glycine. Additionally, when compared to the exact spectral result, the fastest local pseudospectral prediction of the conformational energy difference in glyoxal is in error by 0.2 kcal/mol and saves a factor of 10 in CPU time, indicating the prospects of combining local correlation and pseudospectral methods. © *1996 American Institute of Physics.* [S0021-9606(96)02639-6]

I. INTRODUCTION

Since its inception quantum chemistry has sought to gain insights into ever larger systems. The finite basis set Hartree–Fock (HF) approximation is a standard method of obtaining a first guess of a single-reference (SR) wave function, but its deficiencies are well known. Accurate predictions of bond energies, descriptions of potential energy surfaces, and even single-point calculations of excited-state systems require a better description of the electron–electron interactions or generally of electron correlation, but correlation methods drastically increase the cost and ultimately the feasibility of a calculation.

There are two main limitations on computer calculations: time and storage space. Conventional correlation energy calculations place severe demands on both these resources. For the moment we will discuss nondirect implementations of correlation methods; direct implementations,[1] in which integrals are calculated repeatedly instead of stored, will be discussed later. In terms of time, the cheapest correlation energy method is second-order perturbation theory. Most commonly one uses the Møller–Plesset partitioning of the Hamiltonian[2] in perturbation theory and the calculation is labeled as MPn where n is an integer indicating the highest order of perturbation used; the cost of an MP2 calculation scales as N^5, where N is the number of basis functions. MP3 and single and double excitation configuration interaction (SDCI)[3] both scale as N^6. Other correlation methods which include triply or more highly excited electronic configurations (CSFs—configuration state functions) scale as N^7 or worse. In terms of storage, there are two issues: the amount of high-speed core memory and the amount of low-speed disk space used by a method. Møller–Plesset perturbation theory calculations are conducted with one of two standard choices for the correlating orbitals. Canonical MP calculations make use of the higher energy unoccupied molecular orbitals (MOs), which have the advantage of being orthonormal and which simplify the energy expression, to correlate the occupied orbitals. These unoccupied orbitals are also referred to as ''virtual'' or ''external'' orbitals. Atomic orbitals (AOs) can also be used as correlating orbitals for Møller–Plesset perturbation theory; this is actually advantageous for localizing correlation effects, but introduces the requirement that the equations be solved iteratively. Canonical MP2 stores the transformed two-electron integrals (ERIs—electron repulsion integrals) on disk, an N^4 storage requirement. Atomic orbital MP2 slightly reduces the necessary disk space, but efficient solution of the iterative equations mandates the storage of the coefficients for all doubly excited CSFs in core, which grows roughly as N^4 as well. MP3 and SDCI, which also store the transformed ERIs on disk and the doubles' coefficients in core, have similar constraints. Storage of the CI vector in memory becomes quite limiting for methods that include triply or more highly excited CSFs.

For large calculations, one can turn to direct correlation methods to avoid storing ERIs on disk. As seen more com-

monly in self-consistent field (SCF) methods, direct implementations[1] allow a calculation to be performed when lack of disk space would otherwise make it impossible, with the penalty of potentially having to calculate the same integrals repeatedly. The direct calculation of the ERIs needed in a correlation method is complicated somewhat due to the fact that partially transformed ERIs, i.e., transformation of one or more of the four indices, are used in all efficient formulations of correlation methods; thus, to reduce storage, the AO integrals must be calculated and transformed into molecular orbital (MO) integrals repeatedly in direct correlation methods. In general the formal performance of direct correlation methods will be somewhat less than that of nondirect implementations due to the need to recalculate integrals, but the penalty paid in doing large amounts of slow data retrieval from disk may offset this. Certainly the ability to do calculations which beforehand were impossible is reason enough to use direct correlation methods. Almlöf and Sæbø have written direct local MP2, MP3, and MP4 programs,[4,5] but applications to large molecules have been sparse so far.

We will examine the combination of two approximations to conventional correlation methods in order to combat these problems of treating large molecules. The first approximation we consider is localization. Sæbø and Pulay showed that localization can be used to reduce the size of the CI vector and thus the memory requirements and overall cost of MP2 and MP3;[6] MP4 including only singly, doubly, and quadruply excited CSFs;[7] coupled electron pair theory and single-reference, closed-shell SDCI.[8] Hampel and Werner has also recently applied localization to solve similar problems in coupled cluster theory including single and double excitations.[9] The second approximation is the pseudospectral (PS) approximation. The pseudospectral method in quantum chemistry comes from the application by Friesner[10] of ideas originating in fluid mechanics[11] to the calculation of ERIs. Basically, it involves approximating the integration over one electronic coordinate by quadrature on a grid. Friesner *et al.* developed this approximation in SCF methods—pseudospectral Hartree–Fock[12] and pseudospectral generalized valence bond (GVB).[13] Martinez, Mehta, and Carter applied the pseudospectral approximation to full CI,[14] and Martinez and Carter continued this work in developing pseudospectral doubles CI (PSDCI),[15] PS MP2 and MP3,[16] and pseudospectral multireference SDCI (PSMRSDCI).[17] Friesner and his co-workers have also recently developed two pseudospectral correlation methods: pseudospectral restricted CI[18] and local PSMP2.[19] In general, the pseudospectral approximation accelerates the calculation of ERIs, thus reducing the cost of the calculation. Also, it only requires the storage of intermediate quantities and hence reduces the amount of disk space needed for MP2, MP3, and SDCI to N^3.

Thus we see that localization and the pseudospectral approximation as applied to the calculation of correlation energies address separate limitations. Indeed, nondirect implementations of both pseudospectral and localized spectral methods find themselves constrained by the limitations they do not remedy; the storage of the all-virtual ERIs eventually kept Sæbø and Pulay[5,20] from treating larger systems and the storage of the CI vector becomes a problem for PS full CI, PSMP2/3, PSDCI, and PSMRSDCI. Therefore the combination of these two approximations should be complementary, as presaged by Martinez and Carter.[17] Additionally, it is superior to direct local spectral correlation methods in that it not only reduces disk and memory storage but also accelerates the calculation. Hence we follow the approach suggested by Martinez and Carter[17] and followed recently by Friesner *et al.* in developing localized PSMP2[19] by applying the local approximation to pseudospectral SDCI.

Finally, to avoid confusion some brief notes on the nomenclature to be used throughout are in order. For measures of the size of a molecule, N denotes the total number of basis functions (AOs), N_v the number of virtual or external orbitals, and n the number of occupied or internal orbitals. M represents the number of grid points used in a pseudospectral calculation. In labeling orbitals, a, b, c, and d all stand for virtual MOs and i, j, k, and l for occupied MOs.

II. THEORY

A. Localization

The theory of localization has a long history in quantum chemistry. The Fock operator is left unchanged by an arbitrary unitary transformation among the spin orbitals. For a single-determinant, closed-shell wave function, the fact that each spatial orbital is either doubly occupied or unoccupied (additional spin symmetry) means that an arbitrary unitary transformation among either the occupied or the unoccupied spatial orbitals does not change the Fock operator or the resulting energy of the wave function. Since the HF canonical orbitals (eigenfunctions of the Fock operator) are generally delocalized, exhibiting the full point-group symmetry of the molecule, chemists have for decades sought transformations of the orbitals which render them more similar to the atom-centered lone pairs and bicentric bonds that chemists are used to. Since a unitary transformation of the orbitals as discussed above does not change the energy of the wave function, the localized wave function is not in any way superior or more physical than the delocalized SCF wave function; localization is done only to aid the interpretation of the wave function and as a convenience. There are four localization procedures in common use for HF wave functions: Boys–Foster,[21] Edmiston–Ruedenberg,[22] Pipek–Mezey,[23] and von Niessen.[24] Due to the facts that Boys localization is inexpensive, that Sæbø and Pulay have generally used the Boys procedure to localize occupied orbitals, and that Sæbø and Pulay have shown that the local correlation approximation is relatively insensitive to the method used to localize the orbitals,[20] here we will describe only Boys localization.

The Boys procedure[21] is built around separation of charge in the occupied orbitals. By evaluating the expectation values of the position operators in Cartesian coordinates of the MOs, one may obtain the charge centroids of the occupied orbitals. These are then subjected to pairwise orthogonal transformations until their mutual distances are

maximized. Stated symbolically, Boys localization maximizes the expression:

$$\sum_{i>j} |\mathbf{R}_i - \mathbf{R}_j|^2,$$

where $\mathbf{R}_i$ is the vector to the charge centroid of occupied orbital i, all measured from a common origin. The $\{\mathbf{R}_i\}$ are determined by the integrals $\int \varphi_i^*(\mathbf{r})\mathbf{r}\varphi_i(\mathbf{r})d\mathbf{r}$. Thus the Boys procedure only requires an AO to MO transformation of the one-electron dipole integrals and scales as N^3.

B. Local correlation

Pulay and Sæbø used localized orbitals in two ways to improve the performance of correlation energy calculations. First they localized the internal orbitals. The spatial compactness of the localized orbitals then allowed them to think meaningfully of distances between internal orbitals. The inclusion of only up to two-body terms in the Hamiltonian allows the correlation energy to be separated into discrete interactions between pairs of occupied orbitals. This led to treating pair interactions at different levels of approximation depending on the distance between the orbitals in a given pair. Sæbø and Pulay finally decided upon three classes of orbital pairs based on their estimated contribution to the correlation energy: strong, intermediate, and negligible.[20] The strong pairs were treated without approximation. The interaction between intermediate pairs was treated approximately, e.g., by omitting certain terms, the sum of which is approximately zero at large distance, from the energy expression or by estimating the interaction at a lower level of theory. Negligible pairs' interactions were omitted from the calculation altogether. Discarding these negligible pairs amounts to removing electronic configuration state functions (CSFs) resulting from excitations out of these pairs from the doubles portion of the CI vector. The omission of negligible pairs from the calculation effects a change in the number of doubly excited CSFs from $n^2N_v^2$ to $\tilde{n}^2N_v^2$, where $\tilde{n}^2$ is the number of strong pairs ($\tilde{n}^2 \leqslant n^2$)—the equality holding only when no pairs are discarded; thus the savings from this approximation amounts to a factor of $(n/\tilde{n})^2$.

The second benefit attained from localization by Pulay and Sæbø was the truncation of the external space. The virtual orbitals of the SCF calculation were not localized per se; instead Sæbø and Pulay used the atomic valence orbitals, appropriately projected against the internal space to maintain strong orthogonality between the two subspaces, as external correlating orbitals.[20] This helped in two ways. First, it has been long known that the correlating orbitals which lead to the fastest convergence in a CI calculation are located in the same region of space, with extra nodal planes, as the orbitals to be correlated; the delocalized HF virtuals are deficient in this regard. Thus virtuals resembling the AOs are natural choices for correlating the localized internal orbitals efficiently. Second, the use of AO-like externals allowed Pulay and Sæbø to restrict the effective external space for each internal orbital. Each localized occupied orbital typically involves only one or two atoms; therefore they associated with each internal orbital a ''domain''—the set of valence AOs centered on the atoms involved in the occupied orbital. Double excitations from a pair of internal orbitals were restricted to the union of the domains of the pair. This truncation of the virtual space, as Pulay and Sæbø have labeled it, drastically shortens the CI vector; the memory required to hold the doubles portion then scales as $n^2N_v^2$, where N_v ($\leqslant N_v$) is an average size of the local domains. It is important to realize that N_v no longer depends on the size of the system; N_v depends mostly on the basis set used. Because of this, the dominant memory requirement for the CI vector scales for large system as n^2 (with a prefactor determined by the basis set used or equivalently by N_v). The use of local domains eliminates the necessity in standard correlation methods of exciting electrons from each occupied orbital to every virtual orbital—a necessity which contradicts physical intuition concerning the range of electron–electron interactions in large molecules. It also appears to eliminate many of the interactions between basis functions on nearby atoms that lead to basis set superposition error (BSSE).[25]

Another attempt at treating correlation effects in a localized part of a larger system comes from the work of Kirtman and co-workers. Kirtman's local space approximation (LSA)[26] was originally developed to help expedite SCF calculations and is, thus, somewhat different in focus than Sæbø and Pulay's local methods. Kirtman advocated visualizing a large system as composed of arbitrary, smaller fragment subsystems. He then was interested in describing accurately the inherently local area of interaction when the fragments combined. As an example of his LSA-SCF procedure, he would obtain correct zeroth-order densities for separated substrate and adsorbate fragments and then describe the density change in some local interaction area when the two systems were brought together; he also allowed for an indirect density change in the surroundings to describe charge flow into or out of the local subsystem as a reaction to the local density change. As applied to the SDCI, coupled-cluster, and Møller–Plesset perturbation theory correlation methods,[27,28] Kirtman concentrated on describing the change in the singles and doubles coefficient matrix instead of the overall density in the local area. While Kirtman's approach and ideas are certainly both valid and important, their application to representative systems has been, unfortunately for our purposes, done mostly with HF[29] or with semiempirical methods.[30] The few examples he has provided look promising in terms of the errors incurred and of the theoretical scaling advantage of his methods; however, he has not yet provided any data of which we are aware on the actual CPU timings of his implementations of LSA.

Another important point about the advantages of localization is that the use of any local version of SDCI does not obviate the need to transform the original AO integrals. Efficient formulation of SDCI requires working in an underlying orthonormal basis; it is convenient to use the SCF MO basis, which is orthonormal by construction. Use of the weak pairs approximation alone requires a full transformation of all AO integrals to the MO basis. SDCI simplified by truncation of the virtual space requires transformation of occu-

pied orbital indices to the MO basis and virtual orbital indices to the projected atomic valence orbital basis of the domains. Even with Kirtman's LSA approach,[26] a transformation step from the AOs to at least the orbitals within each local interacting subsystem[27] cannot be escaped. The use of direct local methods[4,5] does not aid in this issue; they merely save disk space by performing the transformation repeatedly. The computational cost of the ERI transformation scales as N^5, a performance bottleneck in local SDCI for large systems. The optimal method for conducting the transformation in any of the local methods is still an active area of research.

In the present work, we implement localization in a way that combines the approaches of Sæbø and Pulay and Kirtman. First, we follow the implementation of Sæbø and Pulay in our grouping of pairs of occupied orbitals: we separate the pairs into two categories, corresponding to Sæbø and Pulay's strong and negligible classifications,[20] based upon the distance between the centroids of the two localized internal orbitals in a given pair; the CSFs resulting from excitations out of ''strong'' pairs are included in the calculation and those resulting from excitations out of ''negligible'' pairs are excluded. Our implementation differs from Sæbø and Pulay's in that our purpose is not necessarily to reproduce total energies to a given precision. In fact, for some of the calculations we performed, over 75% of the CSFs have been classified as negligible while experience with CI indicates, as our data bear out, that such a severe truncation of the CSF space is sure to have anything but a negligible effect on the total energy. Kirtman sought to describe the changes only in the CSF coefficients arising from excitations out of localized orbitals within an arbitrary restricted volume.[28] He did so to focus his computational efforts on describing the change in those CSFs most important to treating the interactions between his systems. We have selected a set of CSFs in a different manner, the changes in whose coefficients should adequately describe the changes in the system during the processes we examine. Just as Kirtman was not interested in the exact prediction of the CSF coefficients outside the boundary of his local space, we shall judge our work as successful, despite inaccuracy in its total energy predictions, so long as it predicts correct energy differences.

Also there is a second important difference between our local SDCI implementation and Sæbø and Pulay's.[20] We have made no attempt to truncate the virtual space for each CSF. As such, some of the potential benefits of localization will be lacking in our results. More of the gain from Sæbø and Pulay's localization comes from preventing the virtual space from growing as quickly as the square of the molecular size than from neglecting weak pairs; and, without truncation of virtuals, BSSE will be unaffected by localization. Our approach is hence the opposite of that of Friesner *et al.* in their localized PS MP2 code.[19] They have only implemented the truncation of the virtual space and have so far made no attempt to discard localized pairs. However, since unlike Friesner and Sæbø and Pulay, we need only keep those CSFs important in describing the particular process under consideration, we have been able to apply the weak pairs approximation much more aggressively than our predecessors. Thus the fact that they achieved only modest savings by its implementation does not automatically discount our current efforts. And the implementation of truncation of the virtuals promises that additional increases in performance will be realized.

C. Pseudospectral ERIs

The pseudospectral method involves approximations originally developed in fluid mechanics by Orszag[11] applied by Friesner and co-workers[10] to the problem of calculating ERIs in a SCF procedure. The basic idea behind the approximation is this: a two-electron integral

$$(ab|cd)=\int \frac{\varphi_a^*(\mathbf{r}_1)\varphi_b(\mathbf{r}_1)\varphi_c^*(\mathbf{r}_2)\varphi_d(\mathbf{r}_2)}{|\mathbf{r}_1-\mathbf{r}_2|}\,d\mathbf{r}_1 d\mathbf{r}_2$$

is approximated by performing the calculation in both physical (grid) space and in function or spectral space. The integration over the singular term in the denominator is best done analytically, since prohibitively dense grids would otherwise be necessary to obtain the desired accuracy in the final result. To this end the Coulomb potential is formulated in physical space as

$$A^{kl}(\mathbf{r}_g)=\int \frac{d\mathbf{r}\varphi_k^*(\mathbf{r})\varphi_l(\mathbf{r})}{|\mathbf{r}_g-\mathbf{r}|}.$$

Representing the Coulomb potential in this way allows one to take advantage of its locality, i.e., the rapidity of its decay, in physical space. Then one can evaluate the two-electron integral in the following way:

$$(ab|cd)\approx\sum_g Q_{ag}A^{cd}(\mathbf{r}_g)R_{gb},$$

where

$$R_{gb}=\varphi_b(\mathbf{r}_g).$$

Finally Q is a generalized inverse of R defined by a nonlinear least-squares solution[12] to the equation $QR=S$, where S is the analytic AO overlap matrix. This allows the PS approximation of ERIs to be interpreted in the following manner. First a transformation by R of the basis function b from function space to grid space is performed. This is then multiplied by the representation of the analytic integration over the Coulomb singularity in physical space and subsequently followed by a backtransformation by Q from grid space to function space. The use of a generalized inverse allows R to be rectangular, i.e., to have more grid points than basis functions. Q is also formed using extra higher angular momentum basis functions, called dealiasing functions, which help prevent spurious effects of the incompleteness of the representation of the grid space from occurring in the backtransformation.[12]

In correlation methods with up to double excitations, one of the dominant terms is the external exchange matrix that arises from interactions between CSFs with two electrons in the external space[31]

$$\sigma_{ij}^{ab} = \sum_{cd} C_{ij}^{cd}(ac|bd).$$

The term scales as $n^2N_v^4$. Applying the pseudospectral approximation to this expression one gets[17]

$$\sigma_{ij}^{ab} \approx \sum_{cd} C_{ij}^{cd} \sum_{g} Q_{ag} A^{bd}(g) R_{gc}.$$

Then rearranging terms

$$\sigma_{ij}^{ab} \approx \left[\sum_{g} Q_{ag} \left[\sum_{d} A^{bd}(g) \left[\sum_{c} C_{ij}^{cd} R_{gc} \right] \right] \right].$$

Each term in brackets scales as $n^2MN_v^2$. Typically M is on the order of $N_v(\approx N)$, so if each enclosed term is evaluated separately and stored, the overall scaling is $n^2N_v^3$, an approximate scaling advantage of N. The complete analysis is somewhat more complicated. To achieve the pseudospectral savings, a tradeoff must be made: a lower formal scaling with respect to N is realized at the cost of increasing the prefactor due to the added work of having to transform between the two spaces. Thus the scaling advantage is only seen when N is sufficiently large. For small values of N, the PS formulation of correlation methods can actually be slower; but beyond the break-even point, the performance of the PS method does improve relative to the spectral method.

Another important advantage of the pseudospectral method is the savings in disk storage space required. The formation of the external exchange part of the Hamiltonian, for example, requires the calculation and storage of ERIs over all external MOs. This storage requirement scales approximately as N_v^4. This requirement is in fact the main limiting factor for localized correlation methods. For all but the smallest cases, the number of integrals precludes storing them in main memory. The time required to write this file out to disk and to read it in to form the external exchange matrix also eventually becomes prohibitive. In PS SDCI, only the $A^{ab}(\mathbf{r}_g)$ need be stored, which scales as MN_v^2 or approximately as N_v^3. This storage scaling advantage can prove critical as N_v increases, but eventually even this greatly reduced requirement will prove constraining. The final solution will be direct on-the-fly computation of parts of the $A^{ab}(\mathbf{r}_g)$ as needed. Again, as with spectral SDCI, there is also a necessary transformation stage, but here the dominant cost is for the conversion of $A^{ab}(\mathbf{r}_g)$ from the AO to the MO basis. The cost of this transformation is N^4, a savings of a factor of N over spectral SDCI for large molecules.

It should also be pointed out that since $M < N_v^2$, the pseudospectral representation of the backtransformation to spectral space and hence the two-electron integrals cannot be exact. Thus PS versions of variational methods will no longer give energies that are strictly upper bounds to the true ground state energy. However, as long as the grids are dense enough and standard correction techniques are used,[12] the relative energy difference between the variational method and its PS counterpart should be within chemical accuracy (approximately 1 mihartree=0.6 kcal/mol). Since localized PS SDCI is essentially an approximation to PS SDCI, it too will not reliably lead to upper bounds on ground state energies. Localized spectral SDCI, though, is basically just SDCI with a smaller CSF basis; so it will give energies that are upper bounds to both the true energy in the given basis and to the SDCI energy.

Here we will present an approximation to SDCI that will address both the disk and memory storage problems of the method. We will localize the occupied orbitals to limit their spatial extent. Then we will use the weak or negligible pairs approximation[20] to eliminate portions of the CI vector, or equivalently the wave function, that should not change significantly during the particular process investigated. As discussed above, this reduces the doubles portion of the CI vector by a factor of $(n/\tilde{n})^2$, where $\tilde{n}^2$ is the number of undiscarded pairs, and the concomitant memory requirements. Finally, in addition, the pseudospectral approximation will be used to accelerate the processing and transformation of the ERIs, speeding both up by a factor of N, and reduce the external disk storage needed for the ERIs similarly by a factor of N. This combination of local and pseudospectral methods will allow the new approximation method to be much faster than conventional spectral SDCI and should ultimately allow the assessment of correlation effects on large molecules.

III. CALCULATIONAL DETAILS

To begin, we note various details concerning the stages leading up to the actual CI calculations or affecting the calculations generally. All calculations presented here were conducted entirely with the 6-31G** basis sets of Pople *et al.*[32] Equilibrium geometries were found at the restricted Hartree–Fock (RHF) level using the GAUSSIAN 92 suite of codes.[33] The SCF wave functions are all closed-shell HF and were calculated using the GVB suite[34] of programs developed by Goddard *et al.* The grids for PS quantities were generated by PSGVB V1.00.[35] The PS calculations were all conducted on the ''coarse'' grids except for C_8H_2, where convergence problems forced the use of the ''medium'' grids. The coarse and medium grids contain approximately 70 and 100 points per atom, respectively. Finally, in the integral transformation, no orbitals were kept as doubly occupied core orbitals excepting the carbon $1s$ orbitals in C_8H_2.

SDCI, due to its single-configuration reference and size-consistency problems, is best suited to the investigation of processes where neither bonds are broken nor large geometric changes occur. For these reasons, we have specifically chosen to examine only conformational energy differences of ground state molecules in this work. Investigating such simple systems is also consistent with the work done by Sæbø and Pulay and Friesner. Sæbø and Pulay focused their investigations on conformational energy differences,[36] relative stabilities of isomers and tautomers,[37] and ring puckering[38] for closed-shell molecules. Friesner, in developing LPSMP2,[19] primarily investigated a series of conformational energy differences along with some examples of small geometrical distortions in ground-state molecules and the dimerization of water. Following a similar philosophy, none

TABLE I. Local pseudospectral SDCI calculations on C_6H_2.

R_{cutoff}[a] (Å)	ΔE[b] (kcal/mol)	Time per iteration[c] (s)	Speedup[d]	No. of pairs cut[e]	No. of CSFs[f]	% CSFs cut[g]
1.5	13.0	243.5	3.58	115	432 055	63.6
2.0	12.7	298.6	2.92	103	510 787	57
2.5	12.6	326.9	2.67	97	550 153	53.6
3.0	2.4	458.8	1.90	72	714 178	39.8
4.0[h]	**0.6**	**632.8**	**1.38**	**41**	**917 569**	**22.7**
5.0	0.6	686.4	1.27	31	983 179	17.1
6.0	0.2	810.6	1.08	10	1 120 960	5.5
7.0	0.0	855.4	1.02	3	1 166 887	1.7
N/A	0.0	871.8	1.00	0	1 186 570	0

[a]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.

[b]$\Delta E = E_{local} - E_{nonlocal}$.

[c]Average sum of CPU and system time per CI iteration, excluding first iteration. The initial guess of the CI vector has all singles and doubles coefficients as zero, so all terms involving coupling with singles and doubles CSFs are not calculated until the second iteration.

[d]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00.

[e]Number of occupied orbital pairs, all double excitations (simultaneous single excitations from each orbital in the pair) from which are excluded from the calculation.

[f]Number of CSFs included in the calculation.

[g]Percent of CSFs excluded in a local calculation relative to the number included in the nonlocal calculation.

[h]Rows in boldface correspond to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closest to 1 kcal/mol.

of the spectral or pseudospectral calculations involved any size-consistency corrections. This may affect the accuracy of our results on the larger systems when compared to those of size-consistent methods or to experimental measurements, but does not invalidate our overall conclusions; the purpose of this current work is to reproduce SDCI predictions more efficiently, not to correct all the inherent deficiencies of SDCI. The investigation of more interesting chemical processes such as bond breaking or significant geometrical distortions such as those that occur in transition states are best explored by multireference methods and so will be deferred to a future publication.

Finally, concerning the SDCI calculations themselves, all SDCI calculations, both spectral and pseudospectral, were conducted using the program MARTCI written by Martinez and Carter. This program is capable of SDCI calculations from any multireference wave function. All the CI coupling coefficients were generated by the GUGA2 program.[39] The Davidson method[40] was used to iteratively solve the direct CI equations for spectral calculations. For PS SDCI, the Hamiltonian matrix is not perfectly symmetric, and the modifications to the Davidson method by Rettrup[41] were used. Several corrections were used for PS SDCI calculations: the $(ij|kl)$ integrals, the $(ai|jk)$ integrals, and the $(ai|bj)$ integrals for coupling between valence and doubly excited CSFs were all generated spectrally to ensure accuracy. As noted previously,[17] this does not affect the method's performance since the number of such integrals is small compared to the $(ab|cd)$ integrals.

Finally, localization was implemented in the following manner. First the occupied orbitals were localized by the Boys method and the centers of the localized orbitals computed. Then, the distances between each pair of localized orbitals were calculated and compared with a single cutoff radius. Those pairs whose separation is greater than the cutoff radius were deemed negligible, and doubles CSFs resulting from excitations from this pair of orbitals were excluded from the calculation. The sensitivity of the results to the cutoff radius will be examined in detail later. All other CSFs were included in the calculation and no further localizing approximations were made.

IV. RESULTS

This work was designed to be an approximation primarily to pseudospectral SDCI, since Sæbø and Pulay have already developed local spectral SDCI;[8] spectral results are given merely to show the transferability of the effects of localization, to give the reader more confidence in the accuracy of the pseudospectral results, and lastly to show how much more efficient the local pseudospectral calculations are relative to their exact spectral counterparts. There are two alternative ways of judging the success of our methods. The first and more stringent criterion is absolute accuracy, i.e., to how many decimal places a single-point localized calculation reproduces the corresponding nonlocal result (i.e., in the latter, no attempt was made to shorten the CI vector—equivalent to an infinite cutoff radius). This is the criterion used by Sæbø and Pulay[20] to assess their localized methods and by Friesner *et al.* to judge their local PSMP2 technique.[19] It has also generally been the way of proving the efficacy of the applications of the pseudospectral approxima-

TABLE II. Local pseudospectral SDCI calculations on C_8H_2.

R_{cutoff}(Å)[a]	ΔE[b] (kcal/mol)	Time per iteration[c] (s)	Speedup[d]	No. of pairs cut[e]	No. of CSFs[f]	% CSFs cut[g]
1.5	17.9	631.4	4.18	100	493 291	69.1
2.0	17.9	629.4	4.19	100	493 291	69.1
2.5	17.9	629.9	4.19	100	493 291	69.1
3.0	3.5	1184.8	2.23	69	835 066	47.7
4.0	1.3	1538.3	1.72	51	1 033 516	35.2
5.0[h]	**1.3**	**1541.2**	**1.71**	**51**	**1 033 516**	**35.2**
6.0	0.4	2001.0	1.32	30	1 265 041	20.7
7.0	0.2	2266.0	1.16	18	1 397 341	12.4
8.0	0.2	2488.9	1.06	7	1 518 616	4.8
9.0	0.0	2623.9	1.01	1	1 584 766	0.7
10.0	0.0	2654.6	0.99	0	1 595 791	0
N/A	0.0	2639.6	1.00	0	1 595 791	0

[a]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.

[b]$\Delta E = E_{local} - E_{nonlocal}$.

[c]Average sum of CPU and system time per CI iteration, excluding first iteration. The initial guess of the CI vector has all singles and doubles coefficients as zero, so all terms involving coupling with singles and doubles CSFs are not calculated until the second iteration.

[d]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00.

[e]Number of occupied orbital pairs, all double excitations (simultaneous single excitations from each orbital in the pair) from which are excluded from the calculation.

[f]Number of CSFs included in the calculation.

[g]Percent of CSFs excluded in a local calculation relative to the number included in the nonlocal calculation.

[h]Rows in boldface correspond to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closest to 1 kcal/mol.

tion by Martinez and Carter[14–17] and Friesner *et al.*[12,13,18,19] However, it is not apparent that this is always the most fruitful criterion to judge an approximation to a quantum chemical method. After all, quantum chemistry is not generally interested in the absolute energies produced by a calculation, but rather with the difference in energy between two calculations; indeed the basis sets used in calculations are often deficient in an absolute sense, producing errors an order of magnitude or more greater than chemical accuracy relative to numerical or HF-limit energy calculations. But these seemingly large errors are roughly insensitive to geometry and thus cancel out to give much more accurate energy differences. This suggests a second criterion not necessarily linked to the first by which one can measure these calculations: the accuracy with which the method reproduces energy differences. As explained earlier, we expect many of the calculations presented here to have large errors in their total energy predictions due solely to the large percentage of pairs discarded relative to the nonlocal calculations. We shall, therefore, examine accuracy in total energy and energy difference predictions separately; we will discuss the former briefly in a series of single-point calculations with varying cutoff radii for orbital pairs, and concentrate on the latter with calculations of conformational energy differences.

A. Total energies

First we examine the absolute accuracy of local PSSDCI. In Tables I–V are listed data from a series of local SDCI calculations on the linear poly-ynes C_6H_2 and C_8H_2, glyoxal (OCH-HCO), glycine (NH_2CH_2COOH), and ethane, respectively. In general, each table presents a set of calculations on the same molecule at varying cutoff radii (N/A indicates that the calculation was done without the exclusion of any pairs). Information on the CPU times, number of pairs excluded from the calculation, number of CSFs excluded from the calculation, and the energy difference in kcal/mol relative to the exact spectral SDCI result are also provided. In all the tables, we see, as expected, that all the local calculations with cutoff radii shorter than 3 Å predict total energies that are in error by more than 1 kcal/mol (roughly chemical accuracy). The errors in total energy increase monotonically as the cutoff radius is decreased becoming as large as 74 kcal/mol in ethane at a cutoff radius of 1 Å. For all but one of the molecules, there is a cutoff radius at or larger than which the total energy prediction will be within chemical accuracy of the nonlocal result. For ethane this cutoff radius is greater than 2.0 Å (the largest cutoff radius shown). Ethane is such a small molecule that one cannot achieve chemical accuracy in the local total energy prediction by means of a single cutoff radius. We will return to ethane in our examination of conformational energy differences below. For glyoxal the cutoff radius at which the local total energy prediction is within chemical accuracy occurs between 2.5 and 3.0 Å; for glycine, roughly 3 Å; for C_6H_2, between 3 and 4 Å; and for C_8H_2, slightly larger than 5 Å. Thus we see that the cutoff radius that predicts a total energy within chemical accuracy increases with the system size,

TABLE III. Local pseudospectral SDCI calculations on glycine[a] at equilibrium geometry. (b) Local pseudospectral SDCI calculations on glycine[a] with C–C–N–H torsion angle at 0°.

R_{cutoff}(Å)[b]	ΔE[c] (kcal/mol)	Time per iteration[d] (s)	Speedup[e]	No. of pairs cut[f]	No. of CSFs[g]	% CSFs cut[h]
(a)						
1.0	58.5	193.5	6.03	152	309 601	75.9
1.012	58.6	195.8	6.02	152	309 601	75.9
1.5	18.5	280.1	4.17	134	424 801	66.9
2.0	10.4	456.3	2.56	103	623 201	51.4
2.5	3.7	702.6	1.66	64	872 801	31.9
3.0[i]	**1.2**	**938.9**	**1.24**	**31**	**1 084 001**	**15.5**
4.0	0.1	1150.4	1.02	3	1 263 201	1.5
5.0	0.0	1176.1	0.99	0	1 282 401	0
N/A	0.0	1167.5	1.00	0	1 282 401	0
(b)						
1.0	64.3	194.5	6.06	153	303 201	76.4
1.012	58.5	193.4	6.04	152	309 601	75.9
1.5	18.5	288.5	4.09	134	424 801	66.9
2.0	10.7	454.2	2.60	104	616 801	51.9
2.5	3.4	742.7	1.59	59	904 801	29.4
3.0[i]	**1.1**	**945.3**	**1.25**	**31**	**1 084 001**	**15.5**
4.0	0.1	1168.1	1.01	2	1 269 601	1
5.0	0.0	1183.4	1.00	0	1 282 401	0
N/A	0.0	1179.3	1.00	0	1 282 401	0

[a]The chemical formula of glycine is NH_2CH_2COOH.
[b]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.
[c]$\Delta E = E_{local} - E_{nonlocal}$.
[d]Average sum of CPU and system time per CI iteration, excluding first iteration. The initial guess of the CI vector has all singles and doubles coefficients as zero, so all terms involving coupling with singles and doubles CSFs are not calculated until the second iteration.
[e]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00.
[f]Number of occupied orbital pairs, all double excitations (simultaneous single excitations from each orbital in the pair) from which are excluded from the calculation.
[g]Number of CSFs included in the calculation.
[h]Percent of CSFs excluded in a local calculation relative to the number included in the nonlocal calculation.
[i]Rows in boldface correspond to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closest to 1 kcal/mol.

which will moderate the speed increases one can obtain in accurately predicting correlation energies with this method. The other interesting phenomenon is how much faster the calculation is completed when the total energy is reproduced to within chemical accuracy. Again for glyoxal, the speed increase is either a factor of 1.10 or 1.17 for the *cis* and *trans* conformers, respectively; for glycine, roughly a factor of 1.2; for C_6H_2, greater than a factor of 1.4; and C_8H_2, roughly a factor of 1.6 (though the interpolation is admittedly inexact). Thus we see that although the cutoff radius to achieve chemical accuracy in the total energy increases with system size, the time savings relative to the nonlocal calculation also increase with system size. This trend is true even though the PSSDCI calculations on C_8H_2 were performed on a denser grid, which slows the calculation down relative to the coarser grid used for all the other molecules. Also, the speed increases and errors incurred are similar for both local spectral and local pseudospectral SDCI calculations relative to the exact spectral and pseudospectral calculations, respectively; localization accelerates both methods equally and causes similar energy deviations as implemented. Thus, to conclude this discussion, our data quantify what was largely implicit in earlier work. In the predictions of total energies to 1 kcal/mol, the negligible pairs approximation provides no speed increase for molecules on the size of ethane, moderate speed gains (factors of 1.10–1.40) for the two medium sized molecules we studied, and a speed increase of roughly a factor of 1.6 for the largest molecule examined. Since the increase in performance grows with the system size, factors of 2 or more in CPU time may be saved for molecules with ten or more atoms. This means that the weak pairs approximation will likely prove useful even for the prediction of correlation energies of large systems.

B. Energy differences

Now we turn to the more promising measure of the worth of the local PS SDCI method—reproduction of energy

TABLE IV. (a) Local pseudospectral SDCI calculations on *cis* conformer of glyoxal. (b) Local spectral SDCI calculations on *cis* conformer of glyoxal. (c) Local pseudospectral SDCI calculations on *trans* conformer of glyoxal. (d) Local spectral SDCI calculations on *trans* conformer of glyoxal.

R_{cutoff}(Å)[a]	ΔE[b] (kcal/mol)	Time per iteration (s)[c]	Speedup[d]	No. of pairs cut[e]	No. of CSFs[f]	% CSFs cut[g]
(a)						
1.0	49.6	51.7	4.36	77	108 626	68.2
1.5	11.7	80.3	2.81	61	157 026	54
2.0	4.4	111.3	2.03	46	202 401	40.7
2.5	2.3	152.8	1.48	28	256 851	24.8
3.0[h]	**0.8**	**205**	**1.10**	**7**	**320 376**	**6.2**
3.5	0.0	222	1.02	0	341 551	0
4.0	0.0	222.3	1.02	0	341 551	0
N/A	0.0	225.6	1.00	0	341 551	0
(b)						
1.0	49.9	144.9	3.81	77	108 626	68.2
1.5	11.7	220.6	2.50	61	157 026	54
2.0	4.4	294.8	1.87	46	202 401	40.7
2.5	2.3	391.2	1.41	28	256 851	24.8
3.0[h]	**0.8**	**500.4**	**1.10**	**7**	**320 376**	**6.2**
3.5	0.0	542.9	1.02	0	341 551	0
4.0	0.0	543	1.02	0	341 551	0
N/A	0.0	551.6	1.00	0	341 551	0
(c)						
1.0	49.7	51.5	4.28	77	108 626	68.2
1.5	11.6	80.3	2.75	61	157 026	54
2.0	4.4	111.1	1.99	46	202 401	40.7
2.5	2.7	147.3	1.50	30	250 801	26.6
3.0[h]	**0.8**	**189.2**	**1.17**	**13**	**302 226**	**11.5**
3.5	0.2	210.2	1.05	5	326 426	4.4
4.0	0.0	222.2	0.99	0	341 551	0
N/A	0.0	220.5	1.00	0	341 551	0
(d)						
1.0	49.9	146.5	3.70	77	108 626	68.2
1.5	11.6	214.6	2.53	61	157 026	54
2.0	4.4	286.6	1.89	46	202 401	40.7
2.5	2.7	369.3	1.47	30	250 801	26.6
3.0[h]	**0.8**	**465.1**	**1.17**	**13**	**302 226**	**11.5**
3.5	0.2	514.7	1.05	5	326 426	4.4
4.0	0.0	544.8	1.00	0	341 551	0
N/A	0.0	542.2	1.00	0	341 551	0

[a]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.

[b]$\Delta E = E_{\text{local}} - E_{\text{nonlocal}}$.

[c]Average sum of CPU and system time per CI iteration, excluding first iteration. The initial guess of the CI vector has all singles and doubles coefficients as zero, so all terms involving coupling with singles and doubles CSFs are not calculated until the second iteration.

[d]Speedup=(Time$_{\text{nonlocal}}$/Time$_{\text{local}}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00.

[e]Number of occupied orbital pairs, all double excitations (simultaneous single excitations from each orbital in the pair) from which are excluded from the calculation.

[f]Number of CSFs included in the calculation.

[g]Percent of CSFs excluded in a local calculation relative to the number included in the nonlocal calculation.

[h]Rows in boldface correspond to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closest to 1 kcal/mol.

TABLE V. Local spectral and pseudospectral SDCI calculations on ethane.

Type[a]	Conformer	R_{cutoff} (Å)[b]	ΔE[c] (kcal/mol)	Time per iterations (s)[d]	Speedup[e]	No. of CSFs[f]	% CSFs cut[g]
SP	Staggered	1.0	74.2	46.7	2.82	33 202	68.7
PS	Staggered	1.0	73.9	14.2	9.27	33 202	68.7
SP	Staggered	2.0	3.8	101.4	1.30	82 621	22.1
PS	Staggered	2.0	3.6	40.3	3.27	82 621	22.1
SP	Staggered	N/A	0.0	131.6	1.00	106 030	0.0
PS	Staggered	N/A	−0.3	55.2	2.38	106 030	0.0
SP	Eclipsed	1.0	74.2	46.9	2.81	33 202	68.7
PS	Eclipsed	1.0	73.9	13.9	9.48	33 202	68.7
SP	Eclipsed	2.0	3.6	101.7	1.30	82 621	22.1
PS	Eclipsed	2.0	3.5	39.8	3.31	82 621	22.1
SP	Eclipsed	N/A	0.0	131.7	1.00	106 030	0.0
PS	Eclipsed	N/A	−0.2	55.0	2.40	106 030	0.0

[a]SP indicates a spectral SDCI calculation; PS indicates a pseudospectral calculation.
[b]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.
[c]For a given conformer, energy of the calculation relative to the nonlocal spectral energy for that conformer.
[d]Average sum of CPU and system time per CI iteration, excluding first iteration. The initial guess of the CI vector has all singles and doubles coefficients as zero, so all terms involving coupling with singles and doubles CSFs are not calculated until the second iteration.
[e]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00.
[f]Number of CSFs included in the calculation.
[g]Percent of CSFs excluded in a local calculation relative to the number included in the nonlocal calculation.

differences. The first problem we consider is the rotational barrier in ethane. In Table VI, we show the results of spectral, pseudospectral, and localized spectral and pseudospectral calculations of the energy difference between eclipsed and staggered ethane. These local calculations are somewhat different from the others because one cutoff radius, 2 Å, used in them was chosen specifically to exclude pairs that give rise to CSFs that should not affect the conformational energy difference; the pairs consisting of a C–H bond on both carbon atoms were discarded since the CSFs resulting from excitations from these pairs represent modifications to the electronic density similar to dispersion interactions between those bonds that should have minor effects on the rotational barrier. Immediately, one can see that the energy difference is insensitive to the various approximations employed; the exact spectral and pseudospectral barriers are within 0.1 kcal/mol, and the local spectral and pseudospectral barrier predictions are within 0.2 kcal/mol of the exact spectral result. It should be remembered (see Table V) that some of the single-point calculations using the localization approximation were in error by more that 70 kcal/mol, yet the energy difference is over 100 times as accurate. Thus, while the calculations are deficient with respect to absolute reproduction of total energies, these large errors are balanced over the two points in the potential energy surface that we have explored, leading to a quite accurate energy difference.

Another point of interest is the speed increases obtained for the local prediction of the energy difference at various cutoff radii. In general the speed increase grows larger as the cutoff radius decreases. There are two comparisons of speed that are particularly helpful. Comparing the times of a local spectral calculation and its corresponding nonlocal spectral calculation or those of a local pseudospectral calculation and its corresponding nonlocal pseudospectral calculation shows the speed increase due to localization alone. In Table VI, the speed increase due solely to localization is a factor of 1.3 at 2 Å and a factor of 2.8 at 1 Å. The other interesting com-

TABLE VI. Local spectral and pseudospectral SDCI predictions of eclipsed-staggered barrier of ethane.

Type[a]	R_{cutoff} (Å)[b]	ΔE[c] (kcal/mol)	Speedup[d]
HF[e]	N/A	2.4	N/A
SP	1.0	2.5	2.82
PS	1.0	2.6	9.38
SP	2.0	2.4	1.30
PS	2.0	2.5	3.29
SP	N/A	2.6	1.00
PS	N/A	2.6	2.39

[a]SP indicates a spectral SDCI calculation; PS indicates a pseudospectral calculation.
[b]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.
[c]$\Delta E = E_{eclipsed} - E_{staggered}$, with both energies calculated either spectrally or pseudospectrally and both with the same cutoff radius as indicated.
[d]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00. All speedups here are averages of those of the two single-point calculations measured relative to the nonlocal spectral calculation.
[e]HF refers to the energy difference prediction of calculations conducted using the Hartree–Fock approximation. In the present context, this would be equivalent to using a negative cutoff radius (thereby eliminating all excited CSFs).

TABLE VII. Local pseudospectral SDCI predictions of glycine[a] C–N rotational energy differences.

R_{cutoff} (Å)[b]	ΔE_{conf}[c] (kcal/mol)	% Error[d]	Δ pairs[e]	Δ CSFs[f]	Speedup[g]
HF[h]	2.3	N/A	N/A	N/A	N/A
1.0	8.1	251.2	1	6 400	6.05
1.012	2.3	2.0	0	0	6.03
1.5	2.3	−1.6	0	0	4.13
2.0	2.5	10.2	1	6 400	2.58
2.5	2.1	−10.2	−5	−32 000	1.62
3.0[i]	**2.2**	**−2.3**	**0**	**0**	**1.25**
4.0	2.3	−1.4	−1	−6 400	1.01
5.0	2.3	0.0	0	0	1.00
N/A	2.3	0.0	0	0	1.00

[a]The chemical formula of glycine is NH_2CH_2COOH.
[b]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.
[c]$\Delta E_{conf}=E_{\tau=0°}-E_{equilibrium}$, where τ is the C–C–N–H torsion angle. $\tau_{equilibrium}=300°$.
[d]Percent error in the local conformational energy difference prediction relative to the nonlocal result.
[e]Difference in number of orbital pairs excluded between the two single-point conformer calculations. Positive values indicate fewer pairs were excluded at the equilibrium geometry and thus should lead to an energy difference prediction that is larger than the actual energy difference.
[f]Difference in the number of CSFs included between the two single-point conformer calculations. Positive values indicate more CSFs were included at the equilibrium geometry.
[g]Speedup=(Time$_{nonlocal}$/Time$_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00. All speedups here are averages of the speedups of the single-point calculations.
[h]HF refers to the energy difference prediction of calculations conducted using the Hartree–Fock approximation. In the present context, this would be equivalent to using a negative cutoff radius (thereby eliminating all excited CSFs).
[i]Rows in boldface correspond to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closest to 1 kcal/mol.

parison of speed occurs when one compares a local pseudospectral calculation to a corresponding nonlocal spectral calculation. Here the speed increase is due to the combined advantage of both localization and the pseudospectral approximation. From Table VI, the combined speed increase from both localization and the pseudospectral approximation grows from over a factor of 3 at 2 Å to over a factor of 9 at 1 Å. Also, as the speed increase of a nonlocal pseudospectral calculation over its corresponding nonlocal spectral counterpart is a factor of 2.4 for this molecule, we see that the combined advantage of local and pseudospectral is somewhat more than multiplicative ($9.4>2.4\times2.8$). We shall see similar speed increases due to the combination of local and pseudospectral methods later when we examine glyoxal.

The next system we consider concerns energy differences of conformers generated by rotation around the C–N bond of glycine. The values are listed in Table VII for the PS SDCI calculation and several local PS SDCI calculations with varying cutoff radii. We see that, with the notable exception of the 1 Å cutoff, the conformational energy difference is relatively insensitive to the approximations made. The maximum error, disregarding the erroneous energy difference prediction of 8.1 kcal/mol at 1 Å, with respect to the exact pseudospectral result is 0.2 kcal/mol. The errors in total energies of the calculations with cutoffs between 1 and 3 Å are several kcal/mol, but the energy difference predictions are much more accurate, easily within chemical accuracy. At 1.5 Å, the energy difference prediction is virtually the same as the exact calculation despite total energy errors of nearly 20 kcal/mol but is generated four times more quickly. It is worth mentioning that the local energy difference predictions are not monotonic in any sense with regard to the size of the cutoff radius. The energy difference prediction is 2.3 kcal/mol at a cutoff radius of 1.5 Å, increases to 2.5 kcal/mol at 2.0 Å, decreases to 2.1 kcal/mol at 2.5 Å, then increases again. The explanation of these fluctuations is inconsistencies in discarding CSFs between the two calculations at the two different geometries. These inconsistencies have increasingly large effects on the energy difference prediction as the cutoff radius is shortened and explain the largest error in our energy difference predictions.

The catastrophic error, roughly 6 kcal/mol, in the energy difference prediction conducted with a cutoff radius of 1 Å deserves special mention. One might be tempted to think that at such a small cutoff, many CSFs important for describing any energetic property of the molecule are being discarded. While this hypothesis should be tested further on larger and more complex systems, we present another interpretation of the data which is consistent with the results presented later for glyoxal. The single-point calculations are done independently of one another; for a given calculation, only the geometry of the isomer or conformer in question determines which pairs are to be excluded by a given cutoff radius. Thus, at different points on the potential energy surface of a molecule, it is possible that different numbers of pairs, or

equivalently CSFs, will be discarded using the same cutoff criterion.

To examine this effect, we have included data in Table VII indicating the difference in the number of pairs and CSFs discarded between the two single-point calculations. Again we first look at the data for cutoff radii greater than 1 Å. At 1.5 Å the same number of pairs are discarded from both single-point calculations, and the energy difference prediction is within 2% of the nonlocal PS value; at 3.0 Å the number of pairs included is again the same and the error is again approximately 2%. The maximum errors, one in either direction, occur at cutoffs of 2 and 2.5 Å, and here there are differences in the number of pairs included. Note that the magnitude of the differences of pairs discarded does not correlate with the size of the error. The magnitude of the errors at 2.0 and 2.5 Å are the same even though the latter calculation has a disparity five times as large in the number of pairs discarded. In general, the importance of having the same number of pairs across points on the potential energy surface decreases as the distance between the pairs increases. Thus, at 4 Å, the difference of one pair between the calculations leads to an error one-half the size as at 3 Å where the two calculations discarded the same number of pairs. At a cutoff of 2.0 Å, a mismatch of one pair corresponds to a 0.2 kcal/mol error; at 2.5 Å, five pairs correspond again to 0.2 kcal/mol; but at 4.0 Å, a one pair mismatch leads to less than 0.1 kcal/mol error. Returning again to the 1 Å case, we see that there is a difference of one pair discarded between the two single-point calculations which may account for the discrepancy in the energy difference prediction. Therefore there appears to be a link between these pair-inclusion inconsistencies and resulting energy difference errors in the local calculations.

We contend that this one pair mismatch between the two calculations at so small a cutoff distance is responsible for the tremendous error in the energy difference prediction. To further prove the point, we examined the pairs included in both local PSSDCI single-point calculations at 1.0 Å. After finding the pair that was not consistently included, we repeated both calculations with a slightly larger cutoff radius, 1.012 Å, which caused the inconsistently treated pair to be kept in both calculations but affected the inclusion of no other pairs. The energy difference prediction of this corrected calculation is 2.3 kcal/mol, within 0.1 kcal/mol of the nonlocal PSSDCI prediction. This is essentially forcing both calculations to include the union of the two sets of included pairs at the original cutoff radius. For simple processes such as we are examining it does not seem to matter whether the union or the intersection of the set of included pairs for all single-point calculations on the potential energy surface are included at short cutoff radii; but the union of included pairs is most certainly the safest choice. The issue of consistency, though, as mentioned above, becomes much less critical for cutoff radii above 2 Å for this size of molecule. We present this procedure as a way of correcting pair mismatches in successive local SDCI calculations. *Thus, even at short cutoff radii, when care is taken to consistently include the same pairs in all calculations, accurate energy differences can be obtained many times faster compared to the nonlocal calculations.*

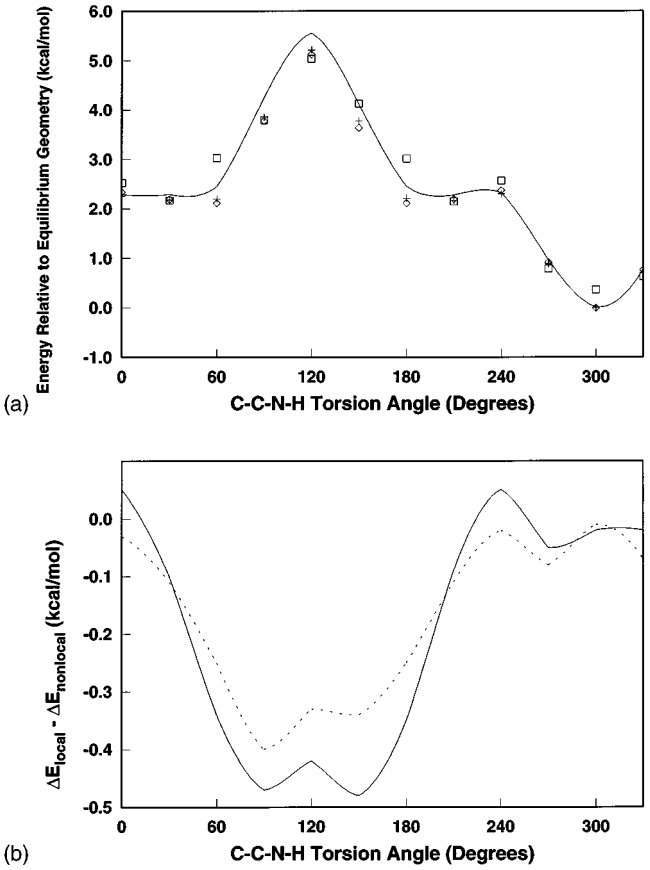

FIG. 1. (a) Nonlocal and local PSSDCI predictions of the C–N rotational barrier in glycine. The solid line indicates the nonlocal PSSDCI prediction of the energies of the conformers relative to that of the equilibrium geometry ($\tau_{\text{equilibrium}}=300°$), points indicate local PSSDCI predictions with different cutoff radii: diamonds indicate a cutoff radius of 1.0 Å, crosses indicate a cutoff radius of 1.5 Å, and squares indicate a cutoff of 2.0 Å. The 1.0 Å cutoff local calculations were actually done with a cutoff radius of 1.015 Å to ensure that the same set of orbital pairs was used at all geometries. (b) Errors in local PSSDCI predictions of the C–N rotational barrier in glycine. Both curves are errors in local PSSDCI predictions of the energy differences between the conformer energy and glycine's equilibrium geometry energy ($\tau_{\text{equilibrium}}=300°$). The errors are measured from the nonlocal PSSDCI prediction of the corresponding energy difference. The solid line indicates a local calculation was performed with a cutoff radius of 1.0 Å, the dashed line indicates a cutoff radius of 1.5 Å was used. The 1.0 Å cutoff local calculations were actually done with a cutoff radius of 1.015 Å to ensure that the same set of orbital pairs was used at all geometries.

The speed increases in glycine increase monotonically as the cutoff radius decreases. Thus, at 4 Å, the local method requires essentially as much CPU time as the nonlocal calculation. By 2.0 Å the speed increase has more than doubled to a factor of 2.57. Finally, with the corrected cutoff radius of 1.012 Å, where the excluded CSF spaces are consistent in both single-point calculations, the speed increase is slightly over a factor of 6. Thus being able to maintain accurate energy difference predictions at small cutoff radii provides the reward of quite large speed increases.

In Fig. 1(a), we present a plot of energy as a function of the C–C–N–H torsional angle in glycine relative to the equi-

TABLE VIII. (a) Local spectral SDCI predictions of glyoxal *cis–trans* energy difference. (b) Local pseudospectral SDCI predictions of glyoxal *cis–trans* energy difference.

R_{cutoff} (Å)[a]	$\Delta E_{cis-trans}$[b] (kcal/mol)	% Error[c]	% Error from spectral[d]	Δ pairs[e]	Δ CSFs[f]	Speedup[g]	Speedup from spectral[h]
			(a)				
HF[i]	5.6	N/A		N/A	N/A	N/A	
1.0	4.9	0.2		0	0	3.75	
1.5	5.0	1.8		0	0	2.51	
2.0	4.9	−0.2		0	0	1.88	
2.5	4.4	−9.1		−2	−6 050	1.44	
3.0[j]	**4.8**	**−1.5**		**−6**	**−18 150**	**1.13**	
3.5	4.7	−4.3		−5	−15 125	1.03	
4.0	4.9	0.0		0	0	0.99	
N/A	4.9	0.0		0	0	1.00	
			(b)				
HF[i]	5.6	N/A	N/A	N/A	N/A	N/A	N/A
1.0	4.7	−0.8	−3.8	0	0	4.32	10.60
1.5	4.8	1.6	−1.5	0	0	2.78	6.81
2.0	4.7	−0.8	−3.7	0	0	2.01	4.92
2.5	4.3	−9.7	−12.4	−2	−6050	1.49	3.64
3.0[j]	**4.7**	**−1.6**	**−4.5**	**−6**	**−18150**	**1.13**	**2.78**
3.5	4.6	−4.4	−7.2	−5	−15125	1.03	2.53
4.0	4.8	0.0	−3.0	0	0	1.00	2.46
N/A	4.8	0.0	−3.0	0	0	1.00	2.45

[a]If the distance between the centroids of charge of two localized occupied orbitals is greater than R_{cutoff} then all CSFs resulting from simultaneous single excitations of electrons from both of these orbitals will be discarded. N/A indicates a nonlocal calculation.

[b]Prediction of the energy difference between the *cis* and *trans* conformers with the indicated cutoff radius.

[c]Percent error in the local conformational energy difference prediction relative to the nonlocal result.

[d]Percentage error in the conformational energy difference prediction of the pseudospectral calculations relative to the nonlocal spectral calculation.

[e]Difference in number of orbital pairs excluded between the two single-point conformer calculations. Positive values indicate fewer pairs were excluded at the equilibrium geometry and thus should lead to an energy difference prediction that is larger than the actual energy difference.

[f]Difference in the number of CSFs included between the two single-point conformer calculations. Positive values indicate more CSFs were included at the equilibrium geometry.

[g]Speedup=($\text{Time}_{nonlocal}/\text{Time}_{local}$). If a local calculation takes half the time of the corresponding nonlocal calculation, the speedup is 2.00. All speedups here are averages of the speedups of the single-point calculations.

[h]Speedup of the pseudospectral calculations relative to the nonlocal spectral calculation.

[i]HF refers to the energy difference prediction of calculations conducted using the Hartree–Fock approximation. In the present context, this would be equivalent to using a negative cutoff radius (thereby eliminating all excited CSFs).

[j]Rows in boldface correspondence to those local calculations with the cutoff radius which reproduces the total energy of the nonlocal calculation to within an amount closes to 1 kcal/mol.

librium total energy of glycine. For reference, the equilibrium geometry occurs at a torsion angle of 300°; the other single-point calculation done above is at 0° (see footnote c of Table VII). The solid line represents the exact PSSDCI results and the data points are the results of the local PSSDCI calculations with cutoff radii of 1.0, 1.5, and 2.0 Å. As in the two single-point calculations discussed earlier, we have corrected for pair mismatches in the local calculations done with a cutoff radius of 1 Å, in order to get reliable predictions, but have left all other local calculations uncorrected. Immediately we see that all the local calculations provide a generally correct prediction of the conformational energy differences; the maximum and minimum occur at the same angles and the barrier height predictions are all similar. To quantify this further, in Fig. 1(b) we show the errors in energy differences of the local PSSDCI calculations with cutoff radii of 1 and 1.5 Å relative to the exact PSSDCI calculation for the same data points. The speed increases are similar to those obtained earlier with the same cutoff radii in glycine. The energy difference predictions are always within ±0.5 kcal/mol throughout this portion of the potential energy surface. Therefore we see that the success of local SDCI in energy difference predictions does not rely on being at any particular angle; it treats correlation effects consistently across all points investigated in this rotational barrier. Furthermore, neither cutoff radius predicts energy differences significantly better than the other; and both local calculations underestimate the amount of destabilization around the barrier (120°). On average, the energy predictions with a 1.0 Å cutoff appear to have slightly larger errors. Thus, in the C–C–N–H torsion angle of glycine, the local PSSDCI provides energy differences accurate to within 0.5 kcal/mol throughout and

up to six times faster than conventional PSSDCI.

Finally, we examine local spectral and pseudospectral predictions of the glyoxal *cis–trans* conformational energy difference in Tables VIII(a) and VIII(b), respectively. Here for all cases, we see that the errors incurred in both the spectral and pseudospectral calculations are less than 1 kcal/mol. That the energy difference predictions are virtually the same as the nonlocal result, despite single-point total energy errors of up to 50 kcal/mol, reinforces the earlier hypothesis that the large errors in the total energy occurring at small cutoff radii remain constant at different geometries so long as the same pairs are included. Indeed the same number of pairs are discarded at 1 Å for both the spectral and pseudospectral cases. In the both the spectral and pseudospectral calculations, the largest error is 0.5 kcal/mol; again it occurs in the calculations with the smallest cutoff radius where a mismatch in the number of discarded pairs appears. There is also once more a rough decline in the importance of equality in the number of pairs discarded as the cutoff radius increases; at a 2.5 Å cutoff a mismatch of two pairs corresponds to an energy error of 0.5 kcal/mol, at 3.0 Å a six pair mismatch corresponds to energy difference errors of less than 0.1 kcal/mol, and at 3.5 Å a five pair differential leads to an error of 0.2 kcal/mol. One can also see that the trend while general is not monotonic. That two calculations at different geometries discard the same number of CSFs at a given cutoff radius does not ensure that they both include the same CSFs. Thus it may be that the calculation at a 3.5 Å cutoff has a greater disparity in the CSF spaces than the difference in just the number of CSFs excluded would indicate.

Whatever the reason for these minute issues, the main trend is that the local SDCI methods are quite successful on a molecule even as small as glyoxal. The maximum error in the *cis–trans* energy difference predictions within the local pseudospectral and spectral calculations is 0.5 kcal/mol ($<$10% of the exact spectral or pseudospectral energy difference). The local pseudospectral calculations are within 0.6 kcal/mol (12.4%) of the exact spectral answer for all cutoff radii. More important, for the smallest cutoff examined, 1 Å, localization alone (local spectral results compared to nonlocal spectral results or local pseudospectral results compared to nonlocal pseudospectral results) predicts energy differences within 0.1 kcal/mol of the exact spectral and pseudospectral answers with speed increases of roughly a factor of 4. Finally for a cutoff of 1 Å, the local pseudospectral result is within 0.2 kcal/mol (4%) of the exact spectral answer with a speed increase of over a factor of 10, showing the promise of combining the advantages of local and pseudospectral methods.

Thus we see in general that the weak pairs approximation performs remarkably well in the prediction of conformational energy differences. In the three cases examined, cutoff radii of 2 Å were adequate to predict conformational energy differences to within chemical accuracy, even though the single-point total energies are dozens of kcal/mol in error. In glyoxal, the smallest cutoff radius leads to the most accurate energy difference prediction, within 0.01 kcal/mol of the exact difference, while only costing a quarter of the CPU time of the exact results. Similarly, in glycine, the smallest cutoff radius, 1.0 Å, leads to nearly the best predicted energy difference ($<$0.1 kcal/mol—roughly 2%—error from the exact result) at also less than a sixth the cost of the exact result. Also it should be noted in glyoxal and ethane, the local speed increases were the same for both spectral and pseudospectral calculations. If a negligible amount of preprocessing CPU time is spent in assuring the same pairs are included in all calculations for a given system, use of the weak pairs approximation by itself in predicting energy differences in large systems is quite promising, since the speed increases grow with system size. Finally, since PSSDCI is faster in general than spectral SDCI, the time increases for local PSSDCI relative to nonlocal spectral SDCI can be reasonably impressive, a factor of 10 for the best glyoxal calculations, with errors well within chemical accuracy.

V. DISCUSSION

The neglect of weak pairs, as developed by Sæbø and Pulay,[20] was never thought to provide great speed increases while predicting acceptable total energies. Indeed, in his first paper on the subject, Pulay[42] ignored this approximation entirely and Friesner did not need the approximation to prove the worth of his local pseudospectral MP2 program.[19] However, Pulay and Sæbø were limited in the system size they could consider since their spectral methods were hampered by the need to store the two-electron integral file on disk. Also they were philosophically inclined toward replicating nonlocal results to great precision which further restricted the usefulness of the approximation. Our results show that the weak pairs approximation alone can reproduce even total energies to within chemical accuracy in molecules of medium size with speedups of roughly a factor of 1.5. Combined with the pseudospectral approximation, not only are calculations on larger molecules possible, but the inherent advantage of pseudospectral SDCI over spectral SDCI increases the speedup attained to nearly a factor of 3 in a molecule as small as glyoxal. Similar speed increases should be attainable by applying both the weak pairs and pseudospectral approximations to other correlation methods as well. Since the efficacy of the weak pairs approximation increases with molecular size, as further developments allow for electron correlation calculations on larger systems, the ability of the weak pairs approximation to provide answers within chemical accuracy while saving considerable CPU time may allow it to be useful even in predicting single-point correlation energies.

What is encouraging and perhaps surprising about the current work is the insight gained into choice of an appropriate CSF space with which to explore a potential energy surface of a molecule. First we see that for the conformational changes investigated here energy differences can be successfully predicted in a severely truncated CSF space; so long as the same orbital pairs are included within all calculations, the exclusion of all localized pairs which are separated even by a distance less than most chemical bonds does

not seem to significantly affect energy difference predictions. In fact, using very short cutoff radii may be inherently advantageous. Since at short cutoff radii one is only including double excitations out of pairs located on the same atom and out of bonds that share a common atom, small geometrical distortions or conformational changes as we have investigated are unlikely to cause large deviations that would greatly affect the basic connectivity of a molecule. Therefore the pairs included at short cutoff radii select out the energetically most crucial pairs in a calculations; and slight geometry changes are less likely to affect pairs within such a tight radius, thus building in some amount of consistency in the pair/CSF space generated. Second, earlier work has stressed the importance of keeping the configurations included in a CI calculation constant across all geometries. In fact, Sæbø and Pulay themselves rejected an adaptive scheme of constructing domains of virtual orbitals in truncating the local virtual space precisely due to the fact that such a scheme often led to differences in the configurations included at varying geometries.[43] We have shown in our calculations the prospect of partial liberation from this restriction. As the separation between the two occupied orbitals from which electrons are excited to form doubles CSFs increases, the contributions to the electron density that the CSFs resulting from such pairs describe become less important; *thus as the cutoff radius increases, the impact of inconsistencies in the single-point CSF spaces decreases sharply*. We have obtained energy differences accurate to within 0.5 kcal/mol compared to the exact results with differences in the sizes of the CSF spaces as large as tens of thousands of CSFs. This is hardly an argument for keeping CSF bases constant to within a single configuration.

To be fair, the few points on the potential energy surfaces of the molecules considered cannot by any means constitute an exhaustive search. Also the energy differences in the processes we have investigated, bond rotations and *cis–trans* conformation, are all predicted to within 1 kcal/mol by HF calculations [see Tables VI–VIII(b)]. Thus, while our results show that the weak pairs approximation treats individual contributions to the correlation energy in a consistent manner even when 70% of the CSFs are eliminated, electron correlation is not crucial to predict these energy differences. Whether true isomerizations, processes requiring the inclusion of resonance for accurate description, or processes involving bond breaking would fare as well under the current treatment remains unanswered. It will also be the subject of further scrutiny in our current development of localized multireference methods. Finally, no evidence has been provided on the accuracy of other properties. Even in the regime where energy differences seem unaffected by inconstancy of the CSF spaces in the single-point calculations, there are no guarantees that accurate predictions of energy differences imply accurate predictions of other properties; in fact as with the total energy, some predictions of properties are inherently single-point entities and would likely deteriorate rapidly as CSFs were excluded from a calculation. Further research is warranted in the development of local correlation method properties, and further exploration of Kirtman's LSA ideas[26] would be helpful here.

In summation, the first step in constructing localized multireference spectral and pseudospectral methods has been completed; Sæbø and Pulay's localization[20] has been partially implemented in spectral and pseudospectral SDCI by the exclusion of doubly excited CSFs derived via substitutions from pairs of occupied orbitals whose separation is greater than some specified cutoff distance. The results are promising for conformational energy differences in the systems considered here; chemical accuracy in the energy differences of conformers is generally achieved at all cutoff distances despite the large errors in single-point total energies with cutoff radii below 2 Å. The main caveat in applying the method is that at small cutoff distances, the same CSFs must be excluded across all the points explored on the potential energy surface. This requirement becomes less important as the cutoff radius is increased. For conformational energy differences, we recommend using a cutoff radius equal to the longest bond distance in the molecule under consideration, so as to include all double excitations from orbitals on nearest-neighbor atoms. (This would correspond to a cutoff radius of roughly 1.5 Å in all three of the cases presented here; thus this recipe may be a little conservative.) Then, if inconsistencies in CSF exclusion occur in the single-point calculations, the cutoff radius should be increased by the minimum amount necessary to ensure that a constant number of CSFs are included at all points on the potential energy surface being investigated.

The speed increases for our implementation of the weak pairs approximation range from a factor of 2 in ethane to a factor of 6 in glycine, when comparing a localized calculation to its corresponding exact calculation (e.g., comparing local pseudospectral timings to pseudospectral timings). In glyoxal, the CPU time is reduced by a factor of 10 due to the combined advantage of the local and pseudospectral approximations (comparing localized pseudospectral to the exact spectral result) with a conformational energy difference error of 0.2 kcal/mol. The comparison of local pseudospectral timings to spectral timings has not been fully explored since the spectral portion of the code has not been as fully optimized as the pseudospectral part. Finally, this attempt differs from previous localized correlation methods in that localization was built into a code capable of handling multireference wave functions. Thus the current inability to conduct multireference local SDCI calculations is only a technical matter and rectifying this is the subject of our current research.

ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research. E.A.C. also is grateful to the Camille and Henry Dreyfus and Alfred P. Sloan Foundations for support via Teacher-Scholar and Research Fellow Awards, respectively. G.G.R. expresses his gratitude to the National Science Foundation for a predoctoral fellowship. T.J.M. thanks the University of California Office of the President and the Fulbright

Foundation for postdoctoral fellowships. Finally, we are indebted to Dr. Antonio J. R. da Silva for many helpful discussions.

[1] J. Almlöf, K. Faegri, and K. Korsell, J. Comput. Chem. **3**, 385 (1982).
[2] C. Møller and M. S. Plesset, Phys. Rev. **46**, 618 (1934).
[3] I. Shavitt, in *Methods of Electronic Structure Theory*, edited by H. F. Schaefer (Plenum, New York, 1977), p. 180.
[4] S. Sæbø and J. Almöf, Chem. Phys. Lett. **154**, 83 (1989).
[5] (a) S. Sæbø, Int. J. Quantum Chem. **38**, 644 (1990); (b) **42**, 217 (1992).
[6] P. Pulay and S. Sæbø, Theor. Chim. Acta **69**, 357 (1986).
[7] (a) S. Sæbø and P. Pulay, J. Chem. Phys. **86**, 914 (1987); (b) **88**, 1884 (1988).
[8] S. Sæbø and P. Pulay, Chem. Phys. Lett. **113**, 13 (1985).
[9] C. Hampel and H.-J. Werner, J. Chem. Phys. **104**, 6286 (1996).
[10] R. A. Friesner, Annu. Rev. Phys. Chem. **42**, 341 (1991).
[11] S. A. Orszag, Stud. Appl. Math. **51**, 253 (1972).
[12] M. N. Ringnalda, M. Belhadj, and R. A. Friesner, J. Chem. Phys. **93**, 3397 (1990).
[13] J. M. Langlois, R. P. Muller, T. R. Coley, W. A. Goddard III, M. N. Ringnalda, Y. W. Won, and R. A. Friesner, J. Chem. Phys. **92**, 7488 (1990).
[14] T. J. Martinez, A. Mehta, and E. A. Carter, J. Chem. Phys. **97**, 1876 (1992).
[15] T. J. Martinez and E. A. Carter, J. Chem. Phys. **98**, 7081 (1993).
[16] T. J. Martinez and E. A. Carter, J. Chem. Phys. **100**, 3631 (1994).
[17] T. J. Martinez and E. A. Carter, J. Chem. Phys. **102**, 7564 (1995).
[18] R. B. Murphy, R. A. Friesner, M. N. Ringnalda, and W. A. Goddard III, J. Chem. Phys. **101**, 2986 (1994).
[19] R. B. Murphy, M. D. Beachy, R. A. Friesner, and M. N. Ringnalda, J. Chem. Phys. **103**, 1481 (1995).
[20] S. Sæbø and P. Pulay, Annu. Rev. Phys. Chem. **44**, 213 (1993).
[21] S. F. Boys, in *Quantum Theory of Atoms, Molecules, and the Solid State*, edited by P.-O. Löwdin (Academic, New York, 1966), p. 263.
[22] C. Edmiston and K. Ruedenberg, Rev. Mod. Phys. **35**, 457 (1963).
[23] J. Pipek and P. G. Mezey, J. Chem. Phys. **90**, 4916 (1989).
[24] W. von Niessen, J. Chem. Phys. **56**, 4290 (1972).
[25] S. Sæbø, W. Tong, and P. Pulay, J. Chem. Phys. **98**, 2170 (1993).
[26] (a) B. Kirtman and C. de Melo, J. Chem. Phys. **75**, 4592 (1981); (b) B. Kirtman, J. Phys. Chem. **86**, 1059 (1982).
[27] B. Kirtman and C. E. Dykstra, J. Chem. Phys. **85**, 2791 (1986).
[28] B. Kirtman, Int. J. Quantum Chem. **55**, 103 (1995).
[29] K. A. Robbins and B. Kirtman, J. Chem. Phys. **99**, 6777 (1993).
[30] (a) B. Kirtman, J. Chem. Phys. **79**, 1835 (1983); (b) B. Kirtman and C. P. DeMelo, Int. J. Quantum Chem. **29**, 1209 (1986); (c) J. Chem. Phys. **86**, 1624 (1987); (d) M. Matos, B. Kirtman, and C. P. DeMelo, *ibid*. **88**, 1019 (1988).
[31] P. E. M. Siegbahn, J. Chem. Phys. **72**, 1647 (1980).
[32] W. J. Hehre, R. Ditchfield, and J. A. Pople, J. Chem. Phys. **56**, 2257 (1972).
[33] GAUSSIAN 92, Revision B, M. J. Frisch, G. W. Trucks, M. Head-Gordon, P. M. W. Gill, M. W. Wong, J. B. Foresman, B. G. Johnson, H. B. Schlegel, M. A. Robb, E. S. Replogle, R. Gomperts, J. L. Andres, K. Raghavachari, J. S. Binkley, C. Gonzalez, R. L. Martin, D. J. Fox, D. J. DeFrees, J. Baker, J. J. P. Stewart, and J. A. Pople, Gaussian, Inc., Pittsburgh, PA, 1992.
[34] R. A. Bair, W. A. Goddard III, A. F. Voter, L. G. Rappé, L. G. Yaffe, F. W. Bobrowicz, W. R. Wadt, P. J. Hay, and W. J. Hunt, GVB2P5 program (unpublished). See: (a) R. A. Bair, Ph.D. thesis, California Institute of Technology, 1980; (b) F. W. Bobrowicz, and W. A. Goddard III, in Ref. 3, p. 77; (c) W. J. Hunt, P. J. Hay, and W. A. Goddard III, J. Chem. Phys. **57**, 738 (1972).
[35] PS-GVB V1.00, M. N. Ringnalda, J. M. Langlois, B. H. Greely, T. V. Russo, R. P. Muller, B. Marten, Y. Won, R. E. Donnelly, Jr., W. T. Pollard, G. H. Miller, W. A. Goddard III, and R. A. Friesner, Schrödinger, Inc., 1993.
[36] (a) S. Sæbø, Chem. Phys. **113**, 383 (1987); (b) S. Sæbø and H. Sellers, J. Phys. Chem. **92**, 4266 (1988); (c) S. Sæbø and K. Kavana, J. Mol. Struct. (Theochem.) **235**, 447 (1991).
[37] (a) S. Sæbø, J. E. Boggs, and K. Fan, J. Phys. Chem. **96**, 9268 (1992); (b) J. Boughton and P. Pulay, Int. J. Quantum Chem. **47**, 49 (1993).
[38] (a) H. Sellers, S. Sæbø, and P. Pulay, Chem. Phys. Lett. **132**, 29 (1986); (b) H. Sellers, J. Almlöf, S. Sæbø, and P. Pulay, J. Phys. Chem. **91**, 4216 (1987).
[39] P. E. M. Siegbahn (unpublished).
[40] E. R. Davidson, J. Comput. Phys. **17**, 87 (1975).
[41] S. Rettrup, J. Comput. Phys. **45**, 100 (1982).
[42] P. Pulay, Chem. Phys. Lett. **100**, 151 (1983).
[43] S. Sæbø and P. Pulay, J. Chem. Phys. **86**, 914 (1987).