HTML ABSTRACT + LINKS

An efficient method for calculating maxima of homogeneous functions of orthogonal matrices: Applications to localized occupied orbitals

Joseph E. Subotnik and Yihan Shao
Department of Chemistry, University of California, Berkeley, California and Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720

WanZhen Liang
Laboratory of Bond Selective Chemistry, University of Science and Technology of China, Hefei, Anhui 230026, People's Republic of China

Martin Head-Gordon
Department of Chemistry, University of California, Berkeley, California and Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720

(Received 15 March 2004; accepted 19 July 2004)

We present here three new algorithms (one purely iterative and two DIIS-like [Direct Inversion in the Iteractive Subspace]) to compute maxima of homogeneous functions of orthogonal matrices. These algorithms revolve around the mathematical lemma that, given an invertible matrix **A**, the function $f(\mathbf{U})=\mathrm{Tr}(\mathbf{AU})$ has exactly one local (and global) maximum for **U** special orthogonal [i.e., $\mathbf{UU}^{\mathbf{T}}=1$ and $\det(\mathbf{U})=1$]. This is proved in the Appendix. One application of these algorithms is the computation of localized orbitals, including, for example, Boys and Edmiston-Ruedenberg (ER) orbitals. The Boys orbitals are defined as the set of orthonormal orbitals which, for a given vector space of orbitals, maximize the sum of the distances between orbital centers. The ER orbitals maximize total self-interaction energy. The algorithm presented here computes Boys orbitals roughly as fast as the traditional method (Jacobi sweeps), while, for large systems, it finds ER orbitals potentially much more quickly than traditional Jacobi sweeps. In fact, the required time for convergence of our algorithm scales quadratically in the region of a few hundred basis functions (though cubicly asymptotically), while Jacobi sweeps for the ER orbitals traditionally scale as the number of occupied orbitals to the fifth power. As an example of the utility of the method, we provide below the ER orbitals of nitrated and nitrosated benzene, and we discuss the chemical implications. © *2004 American Institute of Physics.* [DOI: 10.1063/1.1790971]

I. INTRODUCTION

Although a set of occupied molecular orbitals can be rotated among themselves and still preserve their determinant, the associated n-electron wave function,[1] one cannot but help looking for some physical meaning in one electron, individual orbitals. At the very least, the choice of ground state orbitals affects one's chemical intuition of how individual electrons behave. In a more computational context, the choice of occupied orbitals is important in local correlation calculations,[2] where one wants local orbitals so one may allow only local excitations in a correlated wave function.

Several sets of localized orbitals are popular today among chemists, including those commonly attributed to Boys,[3–6,16] those of Edmiston-Ruedenberg,[6,16] and those of Pipek-Mezey.[7] (See Ref. 7 for a brief but thorough introduction to these different sets of orbitals and the Jacobi sweep method used to calculate them.) Furthermore, in solid-state physics, localized orbitals have been computed as "maximally localized Wannier functions."[8–10] Thus far, most sets of localized orbitals in common practice are found by maximizing some single-valued function of the orbitals $\{\chi_1,\ldots,\chi_n\}$. For the three methods listed above and often used by chemists, these functions have the following form:

$$\zeta_{Boys}(\chi_1,\ldots,\chi_n)=\sum_i |\langle\chi_i|\mathbf{r}|\chi_i\rangle|^2, \tag{1}$$

$$\zeta_{ER}(\chi_1,\ldots,\chi_n)=\sum_i \langle\chi_i\chi_i|\frac{1}{r}|\chi_i\chi_i\rangle, \tag{2}$$

$$\zeta_{PM}(\chi_1,\ldots,\chi_n)=\sum_{i=1}^{n}\sum_{A=1}^{Numb.Atoms} |\langle\chi_i|P_A|\chi_i\rangle|^2. \tag{3}$$

(Here, P_A projects a given orbital onto the space of atomic orbitals centered at atom A.) Note that each of these functions can be written as a homogenous function of orthogonal matrices.[11] For example, in the ER case, see Eq. (4) in the following section. We will revisit this point below.

Though we expect our algorithm and results to be applicable to any set of localized orbitals which maximize some homogeneous function of orthogonal matrices, for the sake of concreteness, in this paper we shall focus on the ER orbitals. The defining property of the ER orbitals was first suggested by Lennard-Jones and Pople[12] when the two were looking to combine orbitals of different irreducible representations of symmetric molecules and form "equivalent orbitals." The authors noted that they came upon sets of orbitals that presumably maximized orbital self−repulsion, and thus

 © 2004 American Institute of Physics

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

minimized the nonclassical exchange energy. This statement is just the fact that, for any set of orthonormal vectors,

$$\mathrm{Tr}\left(\frac{1}{r}\right)=\sum_{ij}(ij|\frac{1}{r}|ij)=\sum_{i}(ii|\frac{1}{r}|ii)+\sum_{i\neq j}(ij|\frac{1}{r}|ij)$$

is constant in any basis. Hence, when the first term (self-repulsion) is maximized, the second term (exchange) is minimized. A decade later, in their paper of 1963 (Ref. 6), Edmiston and Ruedenberg made this construction formal (irregardless of any symmetry) and motivated the use of these orbitals as localized orbitals spanning the same space as the occupied space. Of all of the localization schemes, we consider these (the ER) orbitals the most natural objects because they have such a simple physical interpretation. (A side benefit is that, unlike Boys orbitals, ER and Pipek-Mezey orbitals preserve $\sigma-\pi$ separation. See Ref. 7.)

Unfortunately, the ER orbitals are also the most computationally expensive set to compute. In their original paper, Ruedenberg and Edmiston proposed two methods for calculating the ER orbitals.

(1) Jacobi sweeps. Here one moves a single step by rotating in the plane spanned by any two coordinate vectors, maximizing the ER function in that two plane. One then performs this steps over all $n(n-1)/2$ pairs of coordinate vectors, a process called a sweep. Sweeps are repeated until self-consistency is reached.

(2) Direct minimization following the gradient of the ER function (i.e., steepest descent).

Declaring that Jacobi sweeps were ''simpler'' than direct minimization, Edmiston and Ruedenberg focused on the former as the primary tool for computing ER orbitals. After all, every Jacobi sweep is guaranteed to increase the ER function; furthermore, no line minimizations are required, as each step in a single sweep has an exact distance to follow. Unfortunately, though, Jacobi sweeps do not lend themselves easily to algorithms designed to improve computational efficiency. Essentially, the problem is that an efficient Jacobi sweep algorithm must save the n^4 molecular orbital integrals $(ij|lk)$ (which is already problematic for big systems!), and then update these integrals after each step. If we have just done a step in the ij plane, then the most difficult integrals to update are of the form $(ip|qr)$, a process which scales as n^3. Because a sweep contains on the order of n^2 steps, each sweep scales as n^5. Any attempt to improve the scaling of this algorithm will require efficient computation of the integrals $(ij|kl)$. Though this should be possible once the orbitals are well-enough localized, efficient computation will depend on the extent to which these orbitals localize, and the speed of such localization. Such an algorithm would not directly take advantage of the sparsity of the atomic orbital (AO) basis; furthermore, one still has n^2 steps in each sweep, making linear scaling impossible. As such, the problem seems unnecessarily difficult—if one seeks improved scaling, one really should take steps closer to the gradient.

The second approach towards optimization suggested by Edmiston and Ruedenberg does exactly that–it steps directly along the gradient, doing a line minimization in the gradient direction. Because each step requires only the gradient of the ER function, the algorithm requires only the integrals $(ij|jj)$. Starting from the AO basis, this procedure scales naively as nN^4. However, using the sparsity of the AO basis, and given initial guess orbitals that are well localized, such integrals can be generated in linear time starting from the AO integrals (see below). After all, each localized electron only interacts with a bounded number of other localized electrons around it. However, there are two concerns in doing direct minimization: (i) line minimizations are costly; and (ii) direct minimization does not converge quickly near a stationary point. If one seeks a quickly converging direct minimization scheme, one must invoke conjugate-gradient methods. We have indeed implemented conjugate-gradient minimization over the orthogonal group (following Edelman *et al.*[13]). However, even though this method is reliable and does converge to a stationary point of the ER function, we have found the need for accurate line minimizations far more costly than the alternative algorithm we present below.

In this paper, we propose a third, iterative approach towards calculating the ER orbitals, one which is faster than past algorithms and which generalizes to other localized orbitals coming from homogeneous functions of orthogonal matrices. Simlarly to direct minimization, our algorithm requires only a good-starting guess and an analog of the first derivative of the function to be maximized. Most importantly, we demonstrate in this paper that our algorithm scales roughly quadratically in the number of occupied orbitals when we have $n\approx 100-200$, provided that our initial guess is good. (The algorithm should scale cubicly asymptotically as n goes to infinity.)

II. MATHEMATICAL PRELIMINARIES

In the derivations and algorithms given below, we will make reference to $SO(m)$, which is the manifold of orthogonal matrices with determinant 1. [Formally, $\mathbf{U}\in SO(m)$ if and only if $\mathbf{U}\mathbf{U}^T=\mathbf{Id}$ and $\det(\mathbf{U})=1$.] We parametrize the orthogonal group (around the identity) using the exponential map: $\mathbf{U}=e^{\mathbf{\Delta}}$ for $\mathbf{\Delta}$ antisymmetric. For $\mathbf{U}$ close to the identity, there is a one-to-one map between $\mathbf{U}$ and $\mathbf{\Delta}$. In the equations below, $\mathbf{\Delta}$ is always an antisymmetric matrix, and $\mathbf{U},\mathbf{D}\in SO(m)$ always.

III. A BASIC ITERATIVE APPROACH TOWARDS CALCULATING THE ER ORBITALS

We are given a set of m orthonormal occupied molecular orbitals $\{\chi_i^{(0)}\}_{i=1}^m$ and we seek the (orthonormal) molecular orbitals $\{\chi_i^{(*)}\}$ which are maximally self-interacting. In other words, we want to maximize the ER sum,

$$\xi(\chi_1,\chi_2,\ldots,\chi_m)=\sum_i(\chi_i\chi_i|\chi_i\chi_i),$$

where $\{\chi_i\}$ is any set of orthonormal orbitals spanning the occupied space. Expressed as a function of an orthogonal matrix, the ER function is homogeneous of degree 4:

$$\xi(\mathbf{U})=\sum_{ijklr}U_{ji}U_{ki}U_{li}U_{ri}(\chi_j^{(0)}\chi_k^{(0)}|\chi_l^{(0)}\chi_r^{(0)}). \tag{4}$$

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

We want to maximize this function for $\mathbf{U}\in SO(m)$, and hence compute the ER orbitals: $\chi_i^{(*)}=\Sigma_j\chi_j^{(0)}U_{ji}$.

If we parametrize $SO(m)$ by its antisymmetric generators, i.e., we write $\mathbf{U}=e^{\mathbf{\Delta}}$, and differentiate with respect to Δ_{ij}, we find

$$\left.\frac{\partial\xi}{\partial\Delta_{ij}}\right|_{\mathbf{\Delta}=0}=4(R_{ij}-R_{ji}), \tag{5}$$

where $R_{ji}=(\chi_j^{(0)}\chi_i^{(0)}|\chi_i^{(0)}\chi_i^{(0)})$. In other words, the stationary points of ξ are those orthonormal orbitals $\{\chi_i\}$ for which $(\chi_i\chi_j|\chi_j\chi_j)=(\chi_j\chi_i|\chi_i\chi_i)$.

Although we cannot analytically maximize ξ (nor find the optimal $\{\chi_i^{(*)}\}$), we can work easily with the surrogate function

$$\begin{aligned}\eta(\mathbf{U})&=\sum_i\ (\chi_i\chi_i^{(0)}|\chi_i^{(0)}\chi_i^{(0)})\\&=\sum_{ij}\ U_{ji}(\chi_j^{(0)}\chi_i^{(0)}|\chi_i^{(0)}\chi_i^{(0)})\\&=\sum_{ij}\ R_{ji}U_{ji}=\mathrm{Tr}(\mathbf{R}^TU).\end{aligned}$$

As shown in the Appendix, the function $\eta(\mathbf{U})$ has a unique local (and global) maximum at $\mathbf{U}=\mathbf{R}(\mathbf{R}^T\mathbf{R})^{-1/2}$. If we can assume that our initial guess orbitals, $\{\chi^{(0)}\}$, are ''close'' to the ER orbitals (a point of discussion later on), then this suggests a basic iterative scheme to calculate the ER orbitals.

(1) Start with some orthonormal set of orbitals $\{\chi^{(0)}\}$ which is localized and hopefully close to the ER orbitals.

(2) For $k\geqslant 0$, construct the two-electron integrals above $R_{ji}^{(k)}=(\chi_j^{(k)}\chi_i^{(k)}|\chi_i^{(k)}\chi_i^{(k)})$.

(3) Construct the transformation

$$\mathbf{U}^{(k+1)}=\mathbf{R}^{(k)}[(\mathbf{R}^{(k)})^T\mathbf{R}^{(k)}]^{-1/2}.$$

Define

$$\mathbf{D}^{(k+1)}=\mathbf{D}^{(k)}\cdot\mathbf{U}^{(k+1)}=\mathbf{U}^{(1)}\cdot\mathbf{U}^{(2)}\ldots\mathbf{U}^{(k)}\cdot\mathbf{U}^{(k+1)}.$$

(4) Build $\{\chi^{(k+1)}\}$ as

$$\chi_i^{(k+1)}=\sum_j\ \chi_j^{(k)}U_{ji}^{(k+1)}=\sum_j\ \chi_j^{(0)}D_{ji}^{(k+1)}. \tag{6}$$

(5) Set $k=k+1$. Repeat steps 2–4 until $\mathbf{R}^{(k)}$ is sufficiently close to symmetric and the process appears to have converged. The limit orbitals $\{\chi_i^{(*)}\}$ should be the desired ER orbitals, with $\mathbf{R}^{(*)}$ symmetric. If we have converged after N iterations, we define $\mathbf{D}^{(*)}=\mathbf{U}^{(1)}\cdot\mathbf{U}^{(2)}\ldots\mathbf{U}^{(N)}$ and note that $\chi_i^{(*)}=\Sigma_j\chi_j^{(0)}D_{ji}^{(*)}$.

We will call the iterative steps defined above η steps.

We emphasize that at no time in the above algorithm do we ever produce a set of orbitals $\{\chi_i^{(k)}\}$ which is not orthonormal. Alternatively, we may say that $\forall k,\mathbf{U}^{(k)}\in SO(m)$ where m is the number of occupied orbitals. Hence, for every k, $\mathbf{D}^{(k)}=\mathbf{U}^{(1)}\cdot\mathbf{U}^{(2)}\ldots\mathbf{U}^{(k)}$ is orthogonal, where $\mathbf{D}^{(k)}$ is the matrix expressing $\{\chi_j^{(k)}\}$ in terms of the orbitals $\{\chi_i^{(0)}\}$. (This constraint will be relaxed in the second algorithm we propose.) So the geometric picture in the algorithm above is that, within the space of all $m\times m$ matrices, $\mathbb{R}^{m^2}$, we walk along the surface $SO(m)$ until we come to a matrix $\mathbf{D}^{(*)}$ for which $\mathbf{R}^{(*)}$ is symmetric. This $\mathbf{D}^{(*)}$ often represents the ER orbitals.

We say above that the orbitals $\{\chi_i^{(*)}\}$ are ''often'' the ER orbitals, but that need not always be the case. The convergent orbitals can be maxima, minima, or saddle points of the function $\xi(\mathbf{D})$—of course, $\mathbf{D}$ for the ER orbitals must be the (global) maximum of ξ. One can see that this is a problem as follows: when the iteration convergences, we must have

$$\lim_{k\to\infty}\mathbf{U}^{(k)}=\lim_{k\to\infty}\mathbf{R}^{(k)}[(\mathbf{R}^{(k)})^T\mathbf{R}^{(k)}]^{(-1/2)}=\mathbf{Id},$$

which implies that $\mathbf{R}^{(*)}=(\mathbf{R}^{(*)})^T$. Given Eq. (5) for the first derivative of ξ, $\{\chi_i^{(*)}\}$ must be a stationary point of ξ. But that is all we can say about the limit orbitals $\{\chi_i^{(*)}\}$, for there are no conditions on the second derivative.

One more point should be made. Although it does not seem that we are analytically guaranteed that the value of ξ will increase over the course of several iterations, i.e., that $\xi(\mathbf{D}^{(1)})\leqslant\xi(\mathbf{D}^{(2)})\leqslant\cdots$, it is true that, for small enough steps, the value of ξ must increase. This is a direct consequence of the proportionality

$$\left.\frac{\partial\eta}{\partial\Delta_{ij}}\right|_{\mathbf{\Delta}=0}=\frac{1}{4}\left.\frac{\partial\xi}{\partial\Delta_{ij}}\right|_{\mathbf{\Delta}=0}=R_{ij}-R_{ji}$$

(which is really what ensures that stationary points of η are stationary points of ξ). Hence, for a small step, $\mathbf{U}\approx\mathbf{Id}+\mathbf{\Delta}$,

$$\begin{aligned}\xi(\mathbf{U})&=\xi(\mathbf{Id}+\mathbf{\Delta})\\&=\xi(\mathbf{Id})+\sum_{i<j}\left.\frac{\partial\xi}{\partial\Delta_{ij}}\right|_{\mathbf{\Delta}=0}\cdot\Delta_{ij}\\&=\xi(\mathbf{Id})+4\cdot\sum_{i<j}\left.\frac{\partial\eta}{\partial\Delta_{ij}}\right|_{\mathbf{\Delta}=0}\cdot\Delta_{ij}\approx\xi(\mathbf{Id})+4\cdot\underbrace{\delta\eta}_{\geqslant 0}.\end{aligned}$$

From a certain point of view, the reason we cannot guarantee that $\xi(\mathbf{D}^{(k+1)})\geqslant\xi(\mathbf{D}^{(k)})$ for large steps is because ξ and η have different second and higher derivatives. So, even though η is maximized in each step, ξ may decrease. However, we have found empirically that, if we operate not too far from a maximum/minimum, step sizes are not too big and ξ always increases.

IV. AN ACCELERATED DIIS-LIKE ALGORITHM

The algorithm presented above, while stable, can be greatly sped up using a DIIS-like algorithm.[14,15] Just as self-consistent field (SCF) theory seeks molecular orbital matrices which are orthogonal ($\mathbf{C}^T\mathbf{SC}=\mathbf{Id}$) and for which the corresponding Fock matrices are block diagonal, we seek $m\times m$ matrices ($\mathbf{D}$) which are orthogonal [i.e., $\mathbf{D}\in SO(m)$] and for which the corresponding $\mathbf{R}$ matrices are symmetric. [Here, we again think of the columns of $\mathbf{D}$ as the coefficients of the ER orbitals in the original basis given to us, i.e., $\chi_i=\Sigma_j\chi_j^{(0)}D_{ji}$. See Eq. (6).] For any matrix $\mathbf{D}$, we define the error of $\mathbf{D}$ as the lack of symmetry of $\mathbf{R}(\mathbf{D})$, that is,

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

$E_{ij}(\mathbf{D}) \equiv R_{ij}(\mathbf{D}) - R_{ji}(\mathbf{D})$. The assumption of any DIIS-like algorithm is that, for small changes in D, the changes in the error matrix are linear.

With this in mind, we implemented the following two DIIS-like procedures (which we call DIIS-1 and DIIS-2 and which very much parallel SCF DIIS):

(A) Start with some orthonormal set of orbitals $\{\chi^{(0)}\}$ which is localized and hopefully close to the ER orbitals.

(B) For $k \geqslant 0$, compute $R_{ij}^{(k)}$ and calculate the error matrix $E_{ij}^{(k)} \equiv R_{ij}^{(k)} - R_{ji}^{(k)}$.

(C) Construct the DIIS **B**-matrix, i.e., for $1 \leqslant i,j \leqslant k$, we define $B_{ij} = \langle E^{(i)} | E^{(j)} \rangle = \Sigma_{r,s=1}^{m} E_{rs}^{(i)} E_{rs}^{(j)}$.

(D) Set up the standard DIIS equations and solve for $\{(c_i)\}$:

$$\begin{pmatrix} B_{11} & B_{12} & \cdots & B_{1k} & -1 \\ B_{21} & B_{22} & \cdots & B_{2k} & -1 \\ \vdots & \vdots & \ddots & \vdots & -1 \\ B_{k1} & B_{k2} & \cdots & B_{kk} & -1 \\ -1 & -1 & \cdots & -1 & 0 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \cdots \\ c_k \\ \lambda \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \cdots \\ 0 \\ -1 \end{pmatrix}.$$

(E) Construct an extrapolated matrix $\mathbf{C}^{(k+1)} = \Sigma_{a=1}^{k} c_a \mathbf{D}^{(a)}$.

(F) At the extrapolated $\mathbf{C}^{(k+1)}$, construct an extrapolated **R**. This can be done in two ways:

(a) For DIIS-1: Most formally and precisely, we may define:

(i) $\phi_i^{(k+1)} = \Sigma_j \chi_j^{(0)} C_{ji}^{(k+1)}$.
(ii) $\tilde{S}_{ij}^{k+1} = (\phi_i^{(k+1)} | \phi_j^{(k+1)})$.
(iii) $\tilde{R}_{ji}^{(k+1)} = (\phi_j^{(k+1)} \phi_i^{(k+1)} | \phi_i^{(k+1)} \phi_i^{(k+1)})$.

(b) For DIIS-2: Alternatively, we may note that, to first order around the identity, $\tilde{R}_{ji}^{(k+1)} = \Sigma_{a=1}^{k} c_a R_{ji}^{(a)}$.

(G) Take a generalized η step. (See part 5 of the Appendix). More precisely, we construct

$$\tilde{\mathbf{V}}^{(k+1)} = (\tilde{\mathbf{S}}^{(k+1)})^{-1} \tilde{\mathbf{R}}^{(k+1)} [(\tilde{\mathbf{R}}^{(k+1)})^{T} (\tilde{\mathbf{S}}^{(k+1)})^{-1} \times \tilde{\mathbf{R}}^{(k+1)}]^{-1/2},$$

we define $\mathbf{D}^{(k+1)} = \mathbf{C}^{(k+1)} \tilde{\mathbf{V}}^{(k+1)}$, and we set

$$\chi_i^{(k+1)} = \Sigma_j \chi_j^{(0)} D_{ji}^{(k+1)}.$$

(H) Increment k by 1 and repeat steps (B)–(G). Iterate until convergence, which is when $\mathbf{R}^{(k)}$ is sufficiently symmetric and the error matrix is close to zero.

Here the geometric picture in $\mathbb{R}^{m^2}$ is as follows. Just as before, we want to find an $m \times m$ matrix which is both on the manifold $SO(m)$ and for which the corresponding **R** matrix is symmetric. We start out on $SO(m)$ at the identity, $\mathbf{D}^{(0)} = \mathbf{Id}_{m \times m}$. We next make one η step and move to $\mathbf{C}^{(1)} = \mathbf{D}^{(1)}$, which is also on $SO(m)$. From this point on, extrapolation occurs. Given $k+1$ points on $SO(m)$, $\{\mathbf{D}^{(0)}, \mathbf{D}^{(1)}, \ldots, \mathbf{D}^{(k)}\}$, we assume that ξ is locally a quadratic function near $\mathbf{D}^{(k)}$, and (just as in Pulay's DIIS algorithm) we estimate the point $\mathbf{C}^{(k+1)}$, lying in the k-dimensional subspace spanned by $\{\mathbf{D}^{(k)} - \mathbf{D}^{(0)}, \mathbf{D}^{(k)} - \mathbf{D}^{(1)}, \ldots, \mathbf{D}^{(k)} - \mathbf{D}^{(k-1)}\}$, which should have the minimum error by extrapolation. This takes us off $SO(m)$, i.e., $\mathbf{C}^{(k+1)} \notin SO(m)$. In other words, the extrapolated orbitals are not orthogonal. We then do a generalized η step, computing $\mathbf{D}^{(k+1)}$, which brings us back to $SO(m)$, so we again have orthogonal orbitals. (A generalized η step maximizes the surrogate function η just like the usual η step, only it recognizes that the starting orbitals are not orthonormal, but rather have overlap matrix $\tilde{\mathbf{S}}$. See part 5 of the Appendix.) This process is iterated until convergence.

In the algorithm above there are two procedures for constructing the **R** matrix at the extrapolated coordinates $\mathbf{C}^{(k)}$: either by (a) explicit construction (which is slow but exact), or (b) by extrapolation of previous **R** matrices (which is fast, but efficient only when we are working close to the identity). With either method, the program jumps on and off of $SO(m)$, and converges significantly faster than pure η iterations (see below). With regards to the difference between DIIS-1 and DIIS-2, we expect DIIS-1 to be more broadly applicable than DII-2; i.e., DIIS-2 is faster than DIIS-1, but should presumably require a better initial guess (as it is formally correct only near the identity). The question of how good an initial guess is required will be investigated in a future paper.

V. NUMERICAL PERFORMANCE

A. DIIS versus Jacobi sweeps

The algorithms described above have been implemented into a development version of the Q-CHEM program.[17] A simple assessment of the speed of our DIIS-like algorithm was made by comparison to the method of Jacobi sweeps in the computation of Boys' orbitals. Results for the simple molecule benzene (C_6H_6) are given in Table I of the supplementary material.[18]

Recalling that each individual DIIS step requires evaluation of just one or two **R** matrices, whereas Jacobi sweeps loop over all orbital pairs and updates dipole matrices, we conclude that our algorithm competes with Jacobi sweeps for convergence.

B. DIIS versus η steps

Because our algorithm was designed to compute not just Boys' orbitals, but also more computationally expensive orbitals, we calculated the ER orbitals for benzene (C_6H_6) with increasing basis size. The results are given in Table II of the supplementary material.[18]

The data in Table II demonstrates how much faster calculations become when using DIIS-like interpolations instead of pure η steps. The DIIS algorithm converges in 15 steps, where the η steps take on the order of 400 steps. This difference is enormous when saddle points are encountered because DIIS-like interpolations allow you to move to nearby stationary points quickly, whereas η steps move slowly and cautiously (though more reliably) towards new stationary points. The quick convergence of the DIIS-like routines here reflects the parabolic and well-isolated nature of the ER maximum in the case of benzene. When the be-

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

havior of the ER function around the maximum is not as simple, the DIIS-like algorithms can fail to converge, and one must rely on other methods (Jacobi sweeps, conjugate gradient, or the steady, simple η procedure) to find a maximum. (See the application to benzene nitrosation and nitration below for such an example.)

The saddle point encountered in these benzene calculations can be described as follows: In the smallest basis [Slater-type orbital (STO)-3G], the ER orbitals are (qualitatively) the same as the Boys' orbitals, mixing σ and π orbitals, forming so-called "banana" bonds. As the basis size grows, however, this set of orbitals becomes a saddle point of the ER function, while the true maxima of the function becomes the usual ER orbitals (which do not mix σ and π orbitals). This unusual state of affairs, where saddle points and maxima switch (and the qualitative properties of ER orbitals change) with basis size, emphasizes that one must calculate second derivatives in order to be sure one has found a valid set of orbitals.

C. Scaling with system size: Alkanes

The bottleneck in doing the calculations above is the evaluation of the matrix elements $(ij|jj)$. However, using the Boys' orbitals as our initial guess, we have generated these integrals in subquadratic (and potentially linear) time. We computed these integrals by first making n matrices of the form $K^{(i)}_{\mu\nu}=(\mu i|\nu i)$, which are then transformed into $(ij|jj)$. Most naively, this computation can be done in a ("cubic") time proportional to nN^2, if we exploit only the sparsity of the AO basis and then transform.[19] However, when we exploit the locality of our guess orbitals (and hence the sparsity of the density matrix $P^i_{\mu\nu}$), we can form $\{K^{(i)}_{\mu\nu}\}$ in a ("linear") time proportional to N.[20]

In Figs. 1(a) and 1(b), we show the CPU time required for the computation of the ER orbitals of alkanes of increasing size (in a STO-3G basis) using the two ("cubic" and "linear-scaling") algorithms mentioned above. Note the difference in time scale. For each algorithm, we break down the CPU time into the time needed to (a) gather the AO integrals, (b) digest the integrals into $K^{(i)}_{\mu\nu}$, and (c) perform the linear algebra and manipulations required by our algorithm. This last component (for manipulations and linear algebra) has not been optimized. These graphs show that the time for the "cubic" algorithm is dominated by the digestion of the relevant AO integrals into $K^{(i)}_{\mu\nu}$, and this time grows cubicly and prohibitively. For the "linear-scaling" algorithm, by contrast, the integrals are generated in linear time (as expected). Furthermore, the digestion of the integrals into $K^{(i)}_{\mu\nu}$ is subquadratic and should become linear asymptotically. The only problematic component of the linear-scaling algorithms is that the time required for the linear algebra and memory manipulation appears to grow cubicly, though with a small prefactor. As stated above, this code has not been optimized and improvements can be made by cleaning up the interfaces between different blocks of our own code and the way we manipulate memory storage. Notwithstanding this optimization, however, we do expect the diagonalization of $\mathbf{RR}^T$ and the inversion of $\tilde{\mathbf{S}}$ to scale cubicly with a relatively large prefactor. For larger and larger systems, these effects must dominate and will need to be addressed to find a truly asymptotically quadratic (or better) algorithm. Perhaps in the future, the sparsity $\mathbf{RR}^T$ will allow better than cubic diagonalization, and given that $\tilde{\mathbf{S}}$ is usually close to the identity, a power series for inversion will suffice. For the moment, however, for the sizes of molecules treated today by quantum mechanics, the linear-scaling DIIS-like algorithms presented here are a big advance in the computation of ER orbitals.

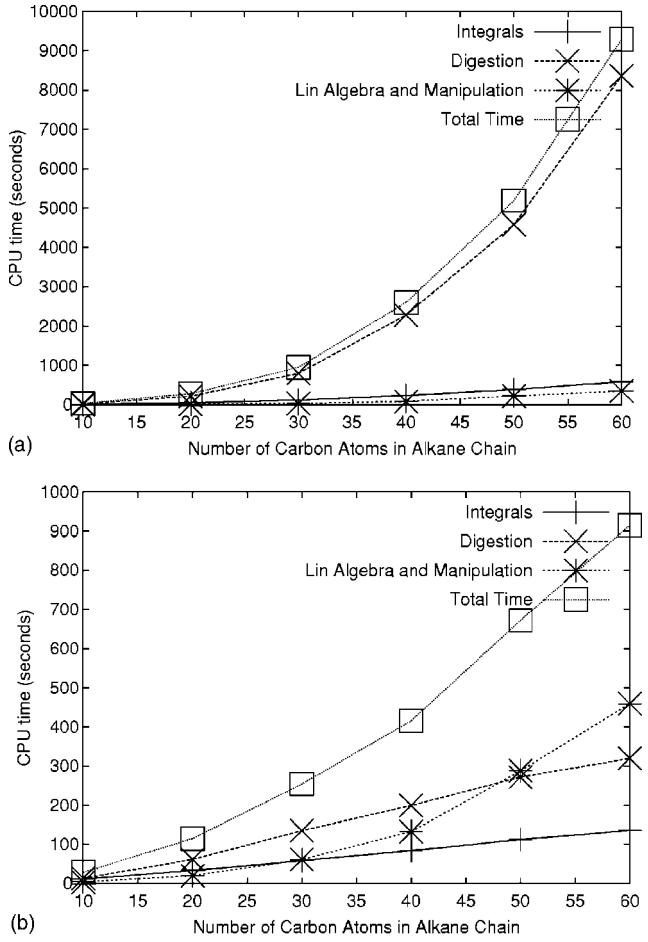


FIG. 1. (a) The CPU times (seconds) required in the calculation of ER orbitals of alkanes of increasing size. The AO basis is STO-3G. Here, the "cubic" algorithm was employed to generate the integrals $(ij|jj)$ from the AO integrals, see text. Note the change in time between this (a) and (b). The dominant effect by far in this graph is digestion, which scales cubicly. For each point on this graph (i.e., for every alkane), exactly seven DIIS-2 iterations were required for convergence (which implies that DIIS-2 converges exactly as well for all alkanes of different sizes). (b) The CPU times (seconds) required in the calculation of ER orbitals of alkanes of increasing size. The AO basis is STO-3G. Here, the "linear-scaling" algorithm was employed to generate the integrals $(ij|jj)$ from the AO integrals, see text. Note the change in time between (b) and (a). Also, note that for systems larger than $C_{50}H_{102}$, the linear algebra and memory manipulation of the algorithm require more CPU time than the integral formation and digestion. [Again, as in (a), for each point on this graph, exactly seven DIIS-2 iterations were required for convergence.]

VI. A CHEMICAL APPLICATION: NITRATION AND NITROSATION OF BENZENE

The highly different reactivities of nitronium (NO_2^+) and nitrosonium (NO^+)[23] toward benzene have been investi-

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

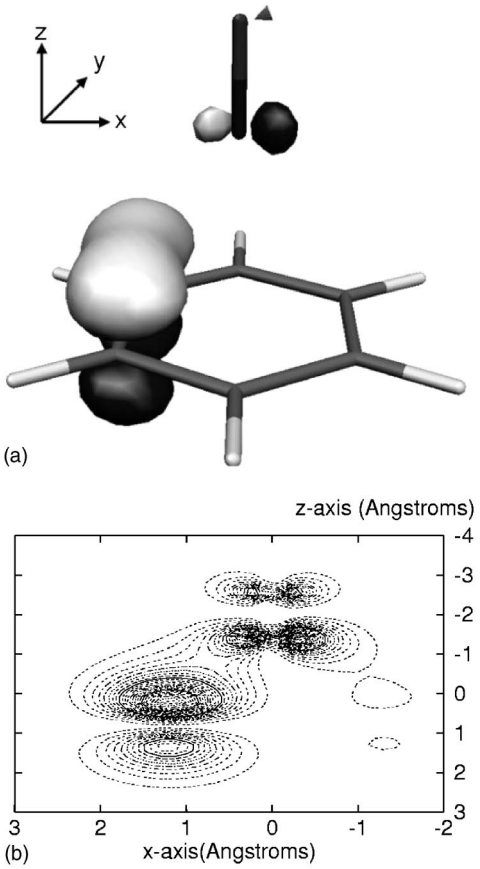


FIG. 2. (a) Nuclear geometry of benzene-NO^+, over which is plotted a three-dimensional spatial representation of a delocalized donor-acceptor ER orbital. The orbital shown is one of three equivalent (symmetrically related) orbitals. This structure was optimized with N-O vertically aligned over the center of the benzene ring. Perpendicular distance from N to benzene is 2.16 Å. (b) Two-dimensional contour plot of the delocalized donor-acceptor ER orbital in benzene-NO^+. The plane of the contour is perpendicular to the benzene plane, cutting across the midpoints of carbon-carbon bonds (rather than through any carbon nucleus).

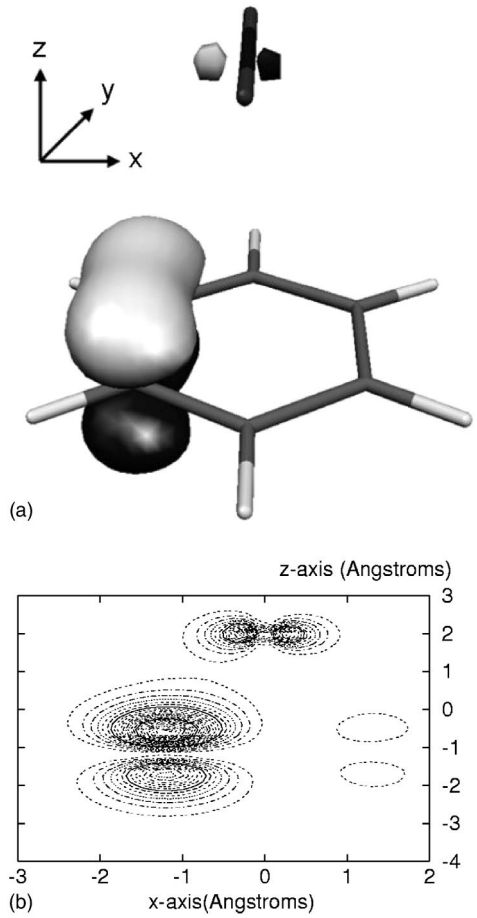


FIG. 3. (a) Nuclear geometry of benzene-NO_2^+, over which is plotted a three-dimensional spatial representation of the one (unique) delocalized donor-acceptor ER orbital. This structure was established by placing NO_2^+ horizontally over the benzene ring and optimizing. The geometry here corresponds to structure 1 in the paper by Olah (Ref. 22). The distance from the nitrogen atom to the benzene plane is 3.1 Å. (b) Two-dimensional contour plot of the delocalized donor-acceptor ER orbital in benzene-NO_2^+. The plane of the contour is perpendicular to both the benzene plane and O-N-O^+, passing directly through two carbon nuclei of benzene and the nitrogen nucleus of NO_2^+.

gated from a theoretical perspective by Gwaltney *et al.*[21] There has also been a recent comprehensive density-functional theory (DFT) study on benzene-NO_2^+.[22] The experimental observation that must be explained is that NO^+ forms a stable π complex with benzene, but NO_2^+ does not. (Instead, NO_2^+ adds directly and rapidly to benzene to form a σ adduct). The canonical explanation of the benzene-NO^+ bonding is that the π orbitals of benzene mix with the π^* orbitals of NO^+ [here we conceive NO^+ as vertically aligned over the benzene plane–see Fig. 2(a)]. One wonders why NO_2^+ is incapable of forming such a bond, given the similar electronic properties of the two (e.g., the two species have very similar ionization potentials in the gas phase)? Gwaltney *et al.* mapped out the potential energy surfaces of both benzene-NO^+ and benzene-NO_2^+ and argued that the different reactivities could be explained by the presence of different stationary points (which would be intermediates in a reaction pathway). These findings matched well with conclusions based upon the application of Marcus-Hush theory.

Since localized orbitals might qualitatively explain the differences in stability between benzene-NO^+ and benzene-NO_2^+, we computed the localized (ER) orbitals of these π complexes using the algorithms described in the previous sections. For benzene-NO^+, the nuclear geometry was fixed by placing the nitrogen and oxygen atoms directly above the center of the benzene plane, and then minimizing the energy of the complex using restricted DFT with the B3LYP functional, leading to the geometry shown in Fig. 2(a). For the sake of convenience, we will identify the benzene plane with the xy plane. For benzene-NO_2^+, the linear cation (NO_2^+) was initially placed horizontally above the benzene plane (with nitrogen above the center of the benzene ring, and the oxygen atoms above carbon nuclei in benzene); subsequently, the restricted B3LYP energy was minimized leading to the geometry shown in Fig. 3(a).[24] (This geometry corresponds to structure 1 in the paper by Olah and co-workers)[22] We shall call the plane, which incorporates the NO_2^+ molecule and is perpendicular to the benzene plane, the yz plane. Although the geometries of benzene-NO^+ and benzene-NO_2^+ are not directly comparable in Figs. 2(a) and 3(a) (i.e., NO^+

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

is vertically aligned while NO_2^+ is horizontally aligned over the benzene plane) in both cases one expects donation of electrons from benzene to the nearby cation, resulting in a stable π complex.

In the case of benzene-NO^+, five ER orbitals are distributed tightly over the NO^+ cation, exactly as expected by a simple count of electrons. In addition, three equivalent delocalized donor-acceptor ER orbitals connect the nitrogen atom of NO^+ to the benzene ring. These orbitals are mixtures of localized π orbitals belonging to benzene with unoccupied orbitals belonging to the nitrogen atom of NO^+. A contour map of one delocalized ER orbital (in the xz plane) is shown in Fig. 2(b). The shape of this delocalized ER orbital shows charge transfer, whereby benzene shares electronic density with the cation NO^+.

In the case of benzene-NO_2^+, the ER orbitals look quite different. For this geometry, only one delocalized ER orbital is spread out over both benzene and NO_2^+. A three-dimensional representation of this orbital is given in Fig. 3(a) and a two-dimensional contour map is given in Fig. 3(b). (The contour map is in the xz plane.) This unique donor-acceptor ER orbital is formed by the mixing of a localized π orbital of benzene with an unoccupied orbital (with π^* character) of NO_2^+. Aside from this donor-acceptor orbital, the other two localized π orbitals of benzene are unable to mix with an unoccupied orbital of NO_2^+ because of geometric concerns: the problem is that the π_x^*-LUMO (LUMO—lowest unoccupied molecular orbital) of NO_2^+ changes the sign of its phase across the yz plane. As such, only one localized π orbital of benzene (with density on the $x<0$ region) can mix constructively with NO_2^+ unoccupied orbitals. The second and third localized π orbitals of benzene have density on both sides of the $x=0$ plane and, therefore, cannot mix constructively with the unoccupied orbitals of NO_2^+. This example demonstrates the utility of localized orbitals, whereby conditions of symmetry can be simplified (i.e., the symmetric properties of the canonical MOs are more cumbersome).

Accordingly, the essential lesson drawn from these ER orbitals is that, unlike NO^+, NO_2^+ simply does not have the correct geometry (and symmetry) to form a stable π complex with benzene, provided that the nitrogen atom be over the center of the benzene ring. Hence, it is not surprising that unconstrained minimization breaks symmetry, bends the NO_2^+ cation (as charge is transferred to it), and moves it horizontally from a position above the center of the benzene ring to a position above the edge of the benzene ring. For completeness, in Fig. 4 we show the donor-acceptor ER orbital of benzene-NO_2^+ in this more stable geometry. Here, exactly one localized π orbital of benzene mixes strongly with one unoccupied orbital of NO_2^+, creating a stationary π complex. Even though this structure is energetically favorable compared to the geometry of Fig. 3(a), we note that this geometry, like Fig. 3(a), is also a saddle point of the energy function (rather than a minimum).[22,25] Thus, this structure is also unlikely to be experimentally observed; most likely, it evolves quickly into the σ adduct of NO_2^+ on benzene.

In summary, we conclude that, unlike NO^+, NO_2^+ does not form a stable π complex because first, basic geometric factors push the cation to the side of the benzene ring while, second, more subtle features of the potential energy surface push benzene-NO_2^+ towards the σ adduct structure, preventing a π complex above the benzene edge from being energetically stable.

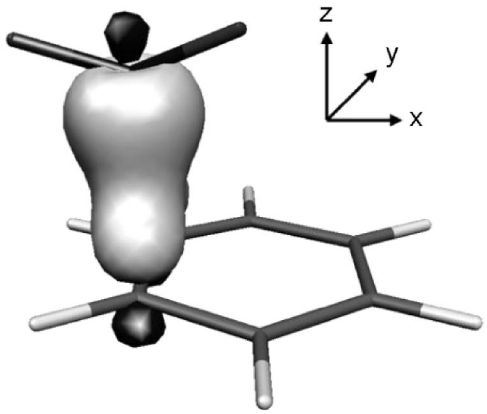


FIG. 4. ER-localized donor-acceptor orbital for benzene-NO_2^+, in which the NO_2^+ cation is bent and is situated above the edge of the benzene ring, over a carbon-carbon bond. This structure is the MP2-optimized structure discussed in Ref. 21. The interaction here is stronger than the interaction in Fig. 3(a) because symmetry here allows a strong mixing of one localized π orbital of benzene with an unoccupied orbital of NO_2^+. The distance from the nitrogen atom to the benzene plane is 2.17 Å.

VII. DISCUSSION

Though the methods presented here can be powerful tools for solving homogeneous equations of orthogonal matrices, specifically those designed to compute localized orbitals, several caveats need to be explicitly stated. A first and obvious problem faced by our algorithms is the problem of invertibility of the **R** matrices. For the ER function, $R_{ij}^{ER}=(ij|jj)$, while for the Boys function, $R_{ij}^{Boys}=\langle i|\mathbf{r}|i\rangle\cdot\langle i|\mathbf{r}|j\rangle$. Stationary points occur when **R** is symmetric, and each step requires inversion of $\mathbf{RR}^T$. As such, there are obvious difficulties when **R** becomes singular. Any possible physical significance of singular **R** matrices is unclear and certainly is specific to the function being maximized. As a rule of thumb, we have had no problems with **R** singularity when working with the ER function. However, the Boys' function is apparently less well behaved, as, for example, **R** becomes singular when localizing the SCF orbitals of (geometrically optimized) decane, $C_{10}H_{22}$. Again, we are unaware of any physical significance behind this singularity.

Second, we repeat that our fast algorithms do converge to saddle points just as they converge to maxima. This demands that one check second derivatives to confirm that one is indeed at a maxima. Of course, one need not compute all of the eigenvalues of the second-derivative matrix. Several techniques exist (such as those of Davidson[26,27]) to compute the smallest or largest eigenvalue of a symmetric matrix; such techniques will be incorporated in future implementations of these algorithms. Even with such a test to conclude the algorithm, future work should seek to quickly find paths to the basin of the maxima directly, without passing through saddle points. Once there, we may successfully turn on the

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

algorithms presented in this paper. In the examples given, we have relied on the (cheaply computed) Boys localized orbitals as our initial guess. But, as demonstrated above, these orbitals can correspond to a saddle point of the ER function, creating the ineffecient effect of forcing us to a saddle point, from which we must then jump off. Perhaps a few conjugate gradient steps may do the trick. However, we expect that any future algorithm designed to quickly find the basin of a maxima must rely on second-derivative information. In the case of the ER function, these derivatives involve terms with matrix elements $(jk|ii)$ and $(ij|ik)$. We speculate that these should not be that much more computationally expensive to compute than the terms from the first derivative, $(ij|jj)$. Though the coulomb-like integral $(jk|ii)$ will be harder to bound, each localized orbital still does interact with only a bounded number of other localized orbitals.

Third, and potentially most difficult, the DIIS-like algorithms presented above fail when the error near the maximum does not change linearly with small changes in the choice of orbitals. This occurs when the function is quartic near the maximum or, in practice, when the Hessian has a small eigenvalue in the region surrounding the purported maximum. For example, in the case of nitrated benzene, the ER function has one rather flat direction in the vicinity of the maximum. This fact complicated our procedure for calculating ER orbitals, as we were forced to use a combination of quicker, but more dangerous, DIIS-like steps and steadier, but slower, η steps. Finally, throughout this paper, we have not considered the case of multiple local maxima, which is a reality at the very least when the nuclear geometry, as in the case of benzene, has certain symmetries. However, such a global problem has no easy solution and was not investigated here.

Notwithstanding these liabilities, the algorithms presented here do potentially allow the calculation of ER orbitals for larger systems than those treated before, provided that the defining function is well behaved near its maximum. Given that the matrix elements $(ij|jj)$ are computed in a timely, linear fashion, we expect that this algorithm will be very useful in helping to search for the best local picture of electronic orbitals. And as the benzene nitration/nitrosation example above demonstrates, such localized orbitals can help us understand the quantum chemistry of bonding.

VIII. SUMMARY

Localized orbitals are and will continue to be an essential tool in the future of quantum chemistry, as on the one hand, they lend themselves to chemical interpretation and, on the other hand, they may be helpful in facilitating more efficient calculations. For instance, the choice of localized occupied orbitals, upon which virutal excitations are made, is necessarily a crucial ingredient in any recipe for post Hartree-Fock local-correlation wave functions. If they can be computed quickly, the ER orbitals may well be used as building blocks in future local-correlation work because, by minimizing $\Sigma_{i\neq j}(ij|ij)$, we expect the two-electron integrals generated from these orbitals to be very sparse, helping to speed up computation. With that in mind, this paper has presented algorithms which very quickly compute the ER orbitals, as well as other orbitals coming from homogeneous functions of orthogonal matrices, provided that we start from a good initial guess. It remains to integrate this algorithm with a good global optimizer, which can take one to the basin of the ER function's maximum in an optimal amount of time. If a decent initial guess can be made, though, this algorithm (or a derivative thereof) will likely find use as computational chemists seek to find and exploit localized orbitals.

ACKNOWLEDGMENTS

We thank Greg Beran, Anthony Dutoi, and Luis Seijo for useful discussions. J.E.S. was supported by the Fannie and John Hertz Foundation. Additionally, this work was supported by the Department of Energy through the Computational Nanosciences program. M.H.G. is a part owner of Q-CHEM.

APPENDIX

1. A mathematical lemma

In this paper, we use the following lemma: Given an invertible matrix A and a function $f:SO(n)\rightarrow\mathbb{R}$ defined by

$$f(\mathbf{U})=\sum_{i,j} A_{ij}U_{ji}=\mathrm{Tr}(\mathbf{AU})$$

then f has one unique local (and global) maximum.

Proof: Let $\mathbf{U}$ be an orthogonal matrix for which f is maximal, and consider, for O orthogonal, $g(O)=\mathrm{Tr}(\mathbf{AUO})=\mathrm{Tr}(\mathbf{BO})=\Sigma_{i,j}B_{ij}O_{ji}$, where $\mathbf{B}=\mathbf{AU}$.

2. Stationary points

We think of the orthogonal group as generated by the vector space of antisymmetric matrices ($\mathbf{O}=e^{\mathbf{\Delta}}$), and differentiate in those directions:

$$\left.\left(\frac{\partial g}{\Delta_{pq}}\right)\right|_{\Delta=0}=\sum_{i,j} B_{ij}\left.\frac{\partial O_{ji}}{\Delta_{pq}}\right|_{\Delta=0}$$

$$=\sum_{i,j} B_{ij}(\delta_{jp}\delta_{iq}-\delta_{jq}\delta_{ip})$$

$$=B_{qp}-B_{pq}=\sum_{i} A_{qi}U_{ip}-A_{pi}U_{iq}.$$

So g is stationary whenever $\mathbf{AU}=(\mathbf{AU})^T$, i.e., $\mathbf{AU}$ is symmetric.

To find $\mathbf{U}$, we write $\mathbf{A}$ in polar form:[30] $\mathbf{A}=\mathbf{N}\tilde{\mathbf{U}}$ where $\mathbf{N}$ is a positive definite Hermitian form and $\tilde{\mathbf{U}}$ is unitary. If we write $\mathbf{V}=\tilde{\mathbf{U}}\mathbf{U}$, also unitary, then we require

$$\mathbf{B}=\mathbf{NV}=(\mathbf{NV})^T=\mathbf{V}^T\mathbf{N}=\mathbf{B}^T. \tag{A1}$$

To find $\mathbf{V}$, we diagonalize $\mathbf{N}=\mathbf{C}^T\mathbf{\Lambda}\mathbf{C}$ where $\mathbf{\Lambda}$ is positive along the diagonal and $\mathbf{C}$ is unitary. We rewrite Eq. (A1) as $\mathbf{C}^T\mathbf{\Lambda}\mathbf{C}\mathbf{V}=\mathbf{V}^T\mathbf{C}^T\mathbf{\Lambda}\mathbf{C}$, or $\mathbf{\Lambda}=\mathbf{W}\mathbf{\Lambda}\mathbf{W}$, where $\mathbf{W}=\mathbf{C}\mathbf{V}^T\mathbf{C}^T$ is unitary.

Let $\mathbf{W}=(w_{ij})$ and

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

$$\mathbf{W}^2=\begin{pmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ t_{21} & t_{22} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ t_{n1} & t_{n2} & \cdots & t_{nn} \end{pmatrix}$$

and

$$\mathbf{\Lambda}=\begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{pmatrix}.$$

Because $t_{ii}=\Sigma_j w_{ij}w_{ji}$ and $\Sigma_j w_{ij}^2=1$, it follows by Cauchy-Schwarz that $|t_{ii}|<1$. Moreover, $\mathrm{Tr}(\mathbf{\Lambda})=\mathrm{Tr}(\mathbf{W\Lambda W})=\mathrm{Tr}(\mathbf{W}^2\mathbf{\Lambda})$ says that

$$\lambda_1+\lambda_2+\cdots+\lambda_n=t_{11}\lambda_1+t_{22}\lambda_2+\cdots+t_{nn}\lambda_n.$$

Given that $\forall i,\lambda_i>0$, this can only be satisfied if, $\forall i,t_{ii}=1$, i.e., $\mathbf{W}^2=\mathbf{Id}$. Accordingly $\mathbf{W}$ must be symmetric, which implies $\mathbf{V}$ is symmetric (and still orthogonal).

Finally, because $\mathbf{V}$ is symmetric, Eq. (A1) implies that $\mathbf{NV}=\mathbf{VN}$, which in turn implies that $\mathbf{N}$ and $\mathbf{V}$ can be simultaneously diagonalized: $\mathbf{N}=\mathbf{D}^T\mathbf{\Lambda}_N\mathbf{D}$ and $\mathbf{V}=\mathbf{D}^T\mathbf{\Lambda}_V\mathbf{D}$ where $\mathbf{\Lambda}_V$ is all ±1 on the diagonal. To solve for $\mathbf{U}$, we recall $\mathbf{V}=\tilde{\mathbf{U}}\mathbf{U}$ hence $\mathbf{U}=\tilde{\mathbf{U}}^T\mathbf{D}^T\mathbf{\Lambda}_V\mathbf{D}$. This can be put in terms of $\mathbf{A}$, using the polarization formulas $\mathbf{A}=\mathbf{N}\tilde{\mathbf{U}}$ and $\mathbf{N}=\sqrt{\mathbf{AA}^T}$ (here we mean the postive square root). Then, $\tilde{\mathbf{U}}=\mathbf{N}^{-1}\mathbf{A}=(\mathbf{AA}^T)^{-1/2}\mathbf{A}$. It follows that

$$\mathbf{U}=\mathbf{A}^T(\mathbf{AA}^T)^{-1/2}\mathbf{D}^T\mathbf{\Lambda}_V\mathbf{D}.$$

This formula gives all of the stationary values of g, which include the maximal values of f. $\mathbf{D}$ here is the matrix of eigenvalues of $\mathbf{N}$, not necessarily unique. As it turns out, only one of these points is a maximum.

3. Maxima

Again expanding an arbitary unitary matrix $\mathbf{O}$ around the identity, it follows that, to second order,

$$\left.\frac{\partial^2 O_{wi}}{\partial\Delta_{pq}\partial\Delta_{rs}}\right|_{\mathbf{\Delta}=0}=\frac{1}{2}[\delta_{wp}(\delta_{qr}\delta_{is}-\delta_{qs}\delta_{ir})+\delta_{qi}(\delta_{wr}\delta_{ps}-\delta_{ws}\delta_{pr})-\delta_{wq}(\delta_{pr}\delta_{is}-\delta_{ps}\delta_{ir})-\delta_{pi}(\delta_{wr}\delta_{qs}-\delta_{ws}\delta_{qr})].$$

Hence,

$$\left.\frac{\partial^2 g}{\partial\Delta_{pq}\partial\Delta_{rs}}\right|_{\mathbf{\Delta}=0}=\sum_{i,j}B_{iw}\frac{\partial^2 O_{wi}}{\partial\Delta_{pq}\partial\Delta_{rs}}=\frac{1}{2}[(B_{sp}+B_{ps})\delta_{qr}+(B_{qr}+B_{rq})\delta_{ps}-(B_{rp}+B_{pr})\delta_{qs}-(B_{sq}+B_{qs})\delta_{pr}]=(B_{sp}\delta_{qr}+B_{qr}\delta_{ps}-B_{rp}\delta_{qs}-B_{sq}\delta_{pr}).$$

To check for negative definiteness, note that

$$\sum_{pqrs}M_{pq}\left.\frac{\partial^2 g}{\partial\Delta_{pq}\partial\Delta_{rs}}\right|_{\mathbf{\Delta}=0}M_{rs}=M_{pr}B_{sp}M_{rs}+M_{sq}B_{qr}M_{rs}-M_{ps}B_{rp}M_{rs}-M_{rq}B_{sq}M_{rs}=\mathrm{Tr}(\mathbf{M}^T\mathbf{BM}^T+\mathbf{MBM}-\mathbf{M}^T\mathbf{BM}-\mathbf{MBM}^T)=\mathrm{Tr}[\mathbf{B}(\mathbf{M}^T\mathbf{M}^T+\mathbf{MM}-\mathbf{M}^T\mathbf{M}-\mathbf{MM}^T)].$$

It is easy to see that

$$Y=\mathbf{M}^T\mathbf{M}^T+\mathbf{MM}-\mathbf{M}^T\mathbf{M}-\mathbf{MM}^T=-(\mathbf{M}-\mathbf{M}^T)\cdot(\mathbf{M}-\mathbf{M}^T)^T$$

is a nonpositive matrix. Because we are interested only in antisymmetric $\mathbf{M}$ (the tangent space of the orthogonal group), it follows $\mathbf{Y}\neq0$. Hence, $\mathbf{Y}$ is negative definite. We claim the sum above is negative $\forall\mathbf{M}$ iff $\mathbf{B}$ is positive definite. Recall from the preceding section, $\mathbf{B}=\mathbf{NV}=\mathbf{D}^T\mathbf{\Lambda D}$ where now we define $\mathbf{\Lambda}=\mathbf{\Lambda}_N\mathbf{\Lambda}_V$.

To check, we first suppose that $\mathbf{B}=\mathbf{D}^T\mathbf{\Lambda D}$ is positive definite. Then

$$\mathrm{Tr}(\mathbf{BY})=\mathrm{Tr}(\mathbf{D}^T\mathbf{\Lambda}_+\mathbf{DY})=\mathrm{Tr}(\mathbf{\Lambda}_+\mathbf{DYD}^T)<0,$$

since $\mathbf{DYD}^T$ is still negative definite. Conversely, choose $\mathbf{M}=\mathbf{D}^T\mathbf{SD}$ and $\mathbf{Y}=4\mathbf{D}^T\mathbf{S}^2\mathbf{D}$ where,

$$S_{\alpha\beta}=\delta_{\alpha i}\delta_{\beta j}-\delta_{\alpha j}\delta_{\beta i},$$

$$(S^2)_{\alpha\beta}=-\delta_{\alpha i}\delta_{\beta i}-\delta_{\alpha j}\delta_{\beta j}.$$

Then the requirement

$$\tfrac{1}{4}\mathrm{Tr}(\mathbf{BY})=\mathrm{Tr}(\mathbf{D}^T\mathbf{\Lambda S}^2\mathbf{D})=\mathrm{Tr}(\mathbf{\Lambda S}^2)=-\Lambda_{ii}-\Lambda_{jj}<0$$

demands that $\forall i,j\ \Lambda_{ii}+\Lambda_{jj}>0$. So we can have at most one negative eigenvalue of $\mathbf{\Lambda}$. But for $SO(n)$, we further demand that $\Lambda_{11}\Lambda_{22}\ldots\Lambda_{nn}>0$ so, in fact, $\forall i,\Lambda_{ii}>0$. $\mathbf{B}$ must be positive definite.

Hence, we conclude that the only orthogonal matrix $\mathbf{U}$ for which f is maximal is when $\mathbf{B}=\mathbf{D}^T\mathbf{\Lambda D}=\mathbf{D}^T\mathbf{\Lambda}_N\mathbf{\Lambda}_V\mathbf{D}$ is positive definite. Since $\mathbf{\Lambda}_N$ is positive definite, this requires $\mathbf{\Lambda}_V=\mathbf{Id}$, and it follows that the only maximum is at

$$\mathbf{U}=\mathbf{A}^T(\mathbf{AA}^T)^{-1/2},$$

where here we take the positive square root. □.

4. Carlson-Keller corollary

One should note that the theorem above provides immediate (re)proof of the Carlson-Keller theorem,[28] that $\mathbf{S}^{-1/2}$ is the transformation constructing orthogonal orbitals that most resemble a set of initial nonorthogonal orbitals. In that case, if $\chi_i=\Sigma_j\phi_jC_{ji}$ and $\mathbf{C}=\mathbf{S}^{-1/2}\mathbf{U}$ for $\mathbf{U}$ orthogonal, one wants to minimize the function:

$$h(\mathbf{C})=\sum_i(\chi_i-\phi_i|\chi_i-\phi_i)=\mathrm{const.}-2\sum_i(\chi_i|\phi_i)=\mathrm{const.}-2\sum_{ij}C_{ji}(\phi_j|\phi_i)=\mathrm{const.}-\mathrm{Tr}(\mathbf{S}^{1/2}\mathbf{U}).$$

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp

Here, of course, $S_{ij}=(\phi_j|\phi_i)$ is the overlap matrix. From the theorem above, it is clear that h is minimized for $\mathbf{U}=(\mathbf{S}^{1/2})^T[\mathbf{S}^{1/2}(\mathbf{S}^{1/2})^T]^{-1/2}=\mathbf{Id}$ since $\mathbf{S}$ is symmetric. Hence, $\mathbf{C}=\mathbf{S}^{-1/2}$, which is the Carlson-Keller result. Furthermore, the theorem above also (re)proves the result of Aiken *et al.*,[29] that h has exactly one minimum, and that is $\mathbf{C}=\mathbf{S}^{-1/2}$.

5. Generalized η step

Suppose we are given nonorthogonal orbitals, $\{\phi_i\}$, and we seek the orthonormal orbitals $\chi_i=\Sigma_j\phi_j C_{ji}$ which maximize the η function,

$$\eta(\mathbf{C})=\sum_i\ (\chi_i\phi_i|\phi_i\phi_i)=\sum_{ij}\ C_{ji}(\phi_j\phi_i|\phi_i\phi_i).$$

Then, we define $S_{ij}=(\phi_i|\phi_j)$, $R_{ij}=(\phi_i\phi_j|\phi_j\phi_j)$, and we enforce orthogonality by writing $\mathbf{C}=\mathbf{S}^{-1/2}\mathbf{U}$ for $\mathbf{U}$ orthogonal. Then,

$$\eta(\mathbf{U})=\sum_{ij}\ R_{ji}S^{-1/2}_{jk}U_{ki}=\mathrm{Tr}(\mathbf{R}^T\mathbf{S}^{-1/2}\mathbf{U}).$$

Application of the lemma above shows that η is maximized for $\mathbf{U}=\mathbf{S}^{-1/2}\mathbf{R}(\mathbf{R}^T\mathbf{S}^{-1}\mathbf{R})^{-1/2}$ and $\mathbf{C}=\mathbf{S}^{-1}\mathbf{R}(\mathbf{R}^T\mathbf{S}^{-1}\mathbf{R})^{-1/2}$.

[1] Note that here and for the rest of this paper, n is the number of electrons in the system and N is the size of the atomic orbital basis chosen by the computational chemist.

[2] S. Saebo and P. Pulay, Annu. Rev. Phys. Chem. **44**, 213 (1993).

[3] As was pointed out to us, credit for the ''Boys'' orbitals should actually be attributed both to Edmiston and Ruedenberg, as well as Foster and Boys. Foster and Boys suggested in 1960 (Ref. 3) the criteria of maximizing the (nonlinear) product of the distances between orbital centroids, while Edmiston and Ruedenberg improved on this criteria in 1963 (Ref. 4), proposing to maximize the (linear) sum of the distances between orbital centroids. The latter definition is the criteria used today for ''Boys'' orbitals, see Eq. (1). Furthermore, Edmiston and Ruedenberg were the first to suggest a formal technique for computing these orbitals exactly (rather than approximately, as Foster and Boys had proposed).

[4] S. F. Boys, Rev. Mod. Phys. **32**, 296 (1960).

[5] S. F. Boys, *Quantum Theory of Atoms, Molecules and the Solid State*, edited by P. O. Lowdin (Academic, New York, 1966), pp. 253.

[6] C. Edmiston and K. Ruedenberg, Rev. Mod. Phys. **35**, 457 (1963).

[7] J. Pipek and P. G. Mezey, J. Chem. Phys. **90**, 4916 (1989).

[8] N. Marzari, I. Souza, and D. Vanderbilt, *Highlight of the Month, Psi-K Newsletter* **57**, 129 (2003); available at http://psi-k.dl.ac.uk/psi-k/newsletters.html

[9] I. Souza, N. Marzari, and D. Vanderbilt, Phys. Rev. B **65**, 035109 (2002).

[10] G. Berghold, C. J. Mundy, A. H. Romero, J. Hutter, and M. Parrinello, Phys. Rev. B **61**, 10040 (2000).

[11] Here, by homogenous, we mean that a function that can be written as:

$$f(\mathbf{U})=\sum_{\substack{i_1i_2\ldots i_n\\ j_1j_2\ldots j_n}} H_{i_1j_1i_2j_2\ldots i_nj_n}U_{i_1j_1}U_{i_2j_2}\ldots U_{i_nj_n}.$$

[12] J. E. Lennard-Jones and J. A. Pople, Proc. R. Soc. London **202**, 166 (1950).

[13] A. Edelman, T. A. Arias, and S. T. Smith, SIAM J. Matrix Anal. Appl. **20**, 303 (1998).

[14] P. Pulay, Chem. Phys. Lett. **73**, 393 (1980).

[15] P. Pulay, J. Comput. Chem. **3**, 556 (1982).

[16] D. A. Kleier, T. A. Halgren, J. H. Hall, Jr., and W. N. Lipscomb, J. Chem. Phys. **61**, 3905 (1974).

[17] J. Kong *et al.*, J. Comput. Chem. **21**, 1532 (2000).

[18] See EPAPS Document No. E-JCPSA6-121-306438 for results for the simple molecule benzene (C_6H_6). A direct link to this document may be found in the online article's HTML reference section. The document may also be reached via the EPAPS homepage (http://www.aip.org/pubservs/epaps.html) or from ftp.aip.org in the directory //epaps. See the EPAPS homepage for more information.

[19] M. J. Frisch, M. Head-Gordon, and J. A. Pople, Chem. Phys. **141**, 189 (1990).

[20] W. Liang, Y. Shao, C. Ochsenfeld, A. T. Bell, and M. Head-Gordon, Chem. Phys. Lett. **358**, 43 (2002).

[21] S. Skokov and R. A. Wheeler, J. Phys. Chem. **103**, 4261 (1999).

[22] P. M. Esteves, J. W. W. Carneiro, S. P. Cardoso, A. G. H. Barbosa, K. K. Laali, G. Rasul, G. K. S. Prakash, and G. A. Olah, J. Am. Chem. Soc. **125**, 4836 (2003).

[23] S. R. Gwaltney, S. V. Rosokha, M. Head-Gordon, and J. K. Kochi, J. Am. Chem. Soc. **125**, 3273 (2003).

[24] We note that the geometries of both Figs. 2(a) and 3(a) are saddle points of the B3LYP functional, rather than minima.

[25] The geometry of Fig. 4 corresponds to structure 7 in the paper by Olah *et al.*

[26] E. R. Davidson, J. Comp. Physiol. **17**, 87 (1975).

[27] M. L. Leininger, C. D. Sherrill, W. D. Allen, and H. F. Schaefer, III, J. Comp. Chem. **22**, 1574 (2001).

[28] B. C. Carlson and J. K. Keller, Phys. Rev. **105**, 102 (1957).

[29] J. Aiken, J. Erdos, and J. Goldstein, Int. J. Quantum Chem. **18**, 1101 (1980).

[30] K. Hoffman and R. Kunze, *Linear Algebra* (Prentice-Hall, New Jersey, 1971).

Downloaded 04 Nov 2004 to 128.32.198.31. Redistribution subject to AIP license or copyright, see http://jcp.aip.org/jcp/copyright.jsp