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

Joseph E. Subotnik *,Yihan Shao*, WanZhen Liang †,
Martin Head-Gordon*

June 4, 2004

Abstract

We present here three new algorithms (one purely iterative and two DIIS-like) to compute maxima of symmetric homogeneous functions of orthogonal matrices. These algorithms revolve around the mathematical lemma that, given an invertible matrix $\mathbf{A}$, the function $f(\mathbf{U})=\mathrm{Tr}(\mathbf{A}\mathbf{U})$ has exactly one local (and global) maximum for $\mathbf{U}$ special orthogonal (i.e. $\mathbf{U}\mathbf{U}^{\mathbf{T}} = 1$ 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

*Department of Chemistry, University of California, Berkeley, and Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720.

†Laboratory of Bond Selective Chemistry, University of Science and Technology of China, Hefei, Anhui 230026, P. R. China

the utility of the method, we provide below the ER orbitals of nitrated and nitrosated benzene, and we discuss the chemical implications.

Introduction

Although a set of occupied molecular orbitals can be rotated among themselves and still preserve their determinant, the associated n-electron wavefunction,[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[1], where one wants local orbitals so one may allow only local excitations in a correlated wavefunction.

Several sets of localized orbitals are popular today among chemists, including those of Boys[2][3], those of Edmiston-Ruedenberg[4], and those of Pipek-Mezey[5]. (See reference [5] 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" [6][7][8]. 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:

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

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

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

(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 and symmetric function of orthogonal matrices.[2] For example, in the ER case, see equation (4) in the next section. We will revisit this point below.

---

[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]Here, by homogenous and symmetric, we mean that a function that can be written

Though we expect our algorithm and results to be applicable to any set of localized orbitals which maximize some homogeneous and symmetric 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[9] 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 minimized the non-classical exchange energy. This statement is just the fact that, for any set of orthonormal vectors,

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

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[4], 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 [5].)

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 2-plane. One then performs this steps over all $n \cdot (n-1)/2$ 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

as: $f(\mathbf{U}) = \sum_{\substack{i_1 i_2 \ldots i_n \\ j_1 j_2 \ldots j_n}} H_{i_1 j_1 i_2 j_2 \ldots i_n j_n} U_{i_1 j_1} U_{i_2 j_2} \cdots U_{i_n j_n}$ (i.e. homogeneous) and

$H_{i_1 j_1 \ldots i_r j_r \ldots i_s j_s \ldots i_n j_n} = H_{i_1 j_1 \ldots i_s j_s \ldots i_r j_r \ldots i_n j_n}$ (i.e. symmetric)

3

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 after each step in a sweep, one must update the orbitals or the integrals. If we suppose that our algorithm saves the n^4 molecular orbital integrals $(ij|lk)$ (which is already problematic for big systems!), then after each step, one must update these integrals. 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 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 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* [22]). 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 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 symmetric homogeneous func-

4

tions of orthogonal matrices. Simlarly to direct minimization, our algorithm requires only a good-starting guess and an analogue 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.)

1 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: $\mathbf{U} = e^{\Delta}$ for Δ antisymmetric. For $\mathbf{U}$ close to the identity, there is a one-to-one map between $\mathbf{U}$ and Δ. In the equations below, Δ is **always** an antisymmetric matrix, and $\mathbf{U}, \mathbf{D} \in SO(m)$ **always**.

2 A basic iterative approach towards calculating the ER orbitals

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

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

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:

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

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

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

$$\frac{\partial \xi}{\partial \Delta_{ij}}\Big|_{\boldsymbol{\Delta}=0} = 4(R_{ij} - R_{ji}) \tag{5}$$

where $R_{ji} = \left(\chi_j^{(0)} \chi_i^{(0)} | \chi_i^{(0)} \chi_i^{(0)} \right)$. 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 $\left\{\chi_i^{(*)}\right\}$), we can work easily with the surrogate function

$$
\begin{aligned}
\eta(\mathbf{U}) &= \sum_i \left(\chi_i \chi_i^{(0)} | \chi_i^{(0)} \chi_i^{(0)} \right) = \sum_{ij} U_{ji} \left(\chi_j^{(0)} \chi_i^{(0)} | \chi_i^{(0)} \chi_i^{(0)} \right) \\
&= \sum_{ij} R_{ji} U_{ji} = \text{Tr}(\mathbf{R}^{\mathbf{T}} \mathbf{U})
\end{aligned}
$$

As shown in the appendix, the function $\eta(\mathbf{U})$ has a unique local (and global) maximum at $\mathbf{U} = \mathbf{R} \left(\mathbf{R}^T \mathbf{R} \right)^{-1/2}$ 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 $k \geq 0$, construct the two-electron integrals above $R_{ji}^{(k)} = \left(\chi_j^{(k)} \chi_i^{(k)} | \chi_i^{(k)} \chi_i^{(k)} \right)$.

3. Construct the transformation $\mathbf{U}^{(k+1)} = \mathbf{R}^{(k)} \left(\left(\mathbf{R}^{(k)} \right)^T \mathbf{R}^{(k)} \right)^{-1/2}$.
   Define $\mathbf{D}^{(k+1)} = \mathbf{D}^{(k)} \cdot \mathbf{U}^{(k+1)} = \mathbf{U}^{(1)} \cdot \mathbf{U}^{(2)} \cdots \mathbf{U}^{(k)} \cdot \mathbf{U}^{(k+1)}$.

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 $\left\{\chi_i^{(*)}\right\}$ should be the desired ER orbitals, with $\mathbf{R}^{(*)}$ symmetric. If we have converged after N iterations, we define $\mathbf{D}^{(*)} = \mathbf{U}^{(1)} \cdot \mathbf{U}^{(2)} \cdots \mathbf{U}^{(N)}$ and note that $\chi_i^{(*)} = \sum_j \chi_j^{(0)} D_{ji}^{(*)}$.

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 $\left\{\chi_i^{(k)}\right\}$ 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, $\mathbf{D}^{(k)} = \mathbf{U}^{(1)} \cdot \mathbf{U}^{(2)} \cdots \mathbf{U}^{(k)}$ 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 $\left\{\chi_i^{(*)}\right\}$ 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

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

which implies that $\mathbf{R}^{(*)} = (\mathbf{R}^{(*)})^T$. Given equation (5) for the first derivative of ξ, $\left\{\chi_i^{(*)}\right\}$ must be a stationary point of ξ. But that is all we can say about the limit orbitals $\left\{\chi_i^{(*)}\right\}$, 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 $\xi(\mathbf{D}^{(1)}) \leq \xi(\mathbf{D}^{(2)}) \leq \cdots \leq \xi(\mathbf{D}^{(k)}) \leq$, it is true that, for small enough steps, the value of ξ must increase. This is a direct consequence of the proportionality

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

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

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

From a certain point of view, the reason we cannot guarantee that, for large steps, $\xi(\mathbf{D}^{(k+1)}) \geq \xi(\mathbf{D}^{(k)})$ 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.

3 An accelerated DIIS-like algorithm

The algorithm presented above, while stable, can be greatly sped up using a DIIS-like algorithm.[13, 14] Just as SCF theory seeks molecular orbital matrices which are orthogonal ($\mathbf{C}^T \mathbf{S} \mathbf{C} = \mathbf{Id}$) 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. $\chi_i = \sum_j \chi_j^{(0)} D_{ji}$. See eqn (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 $E_{ij}(\mathbf{D}) := R_{ij}(\mathbf{D}) - R_{ji}(\mathbf{D})$. 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):

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

2. For $k \geq 0$, compute $R_{ij}^{(k)}$ and calculate the error matrix
$E_{ij}^{(k)} := R_{ij}^{(k)} - R_{ji}^{(k)}$

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

8

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

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

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

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

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

      i. $\phi_i^{(k+1)} = \sum\limits_j \chi_j^{(0)} C_{ji}^{(k+1)}$

      ii. $\tilde{S}_{ij}^{k+1} = (\phi_i^{(k+1)} | \phi_j^{(k+1)})$

      iii. $\tilde{R}_{ji}^{(k+1)} = \left(\phi_j^{(k+1)} \phi_i^{(k+1)} | \phi_i^{(k+1)} \phi_i^{(k+1)} \right)$

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

7. Take a generalized η-step. (See subsection 8.5.) More precisely, we construct $\tilde{\mathbf{V}}^{(k+1)} = (\tilde{\mathbf{S}}^{(k+1)})^{-1} \tilde{\mathbf{R}}^{(k+1)} \left((\tilde{\mathbf{R}}^{(k+1)})^T (\tilde{\mathbf{S}}^{(k+1)})^{-1} \tilde{\mathbf{R}}^{(k+1)} \right)^{-1/2}$, we define $\mathbf{D}^{(k+1)} = \mathbf{C}^{(k+1)} \tilde{\mathbf{V}}^{(k+1)}$, and we set $\chi_i^{(k+1)} = \sum\limits_j \chi_j^{(0)} D_{ji}^{(k+1)}$.

8. Increment k by 1 and repeat steps 2-7. 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 $\mathbf{R}$ 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 subsection 8.5.) This process is iterated until convergence.

In the algorithm above there are two procedures for constructing the $\mathbf{R}$ matrix at the extrapolated coordinates $\mathbf{C}^{(k)}$: either by (a) explicit construction (which is slow but exact), or (b) by extrapolation of previous $\mathbf{R}$ 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.

4   Numerical Performance

4.1   DIIS vs. Jacobi Sweeps

The algorithms described above have been implemented into a development version of the Q-Chem program[18]. 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. [3] To do so, we made the following definitions:[4]

$$R_{ij}^{Boys} = <i|\mathbf{r}|i> \cdot <i|\mathbf{r}|j> \qquad error_{Boys} = \max_{ij} \left| R_{ij}^{Boys} - R_{ji}^{Boys} \right|$$

---

[3] As has been emphasized, our algorithm applies equally well to the Boys' orbitals or Pipek-Mezey orbitals as to the ER orbitals. In this case, we chose not to compare DIIS and Jacobi sweeps when generating ER orbitals because, as has been pointed out, ER orbitals are computationally expensive when computed by Jacobi sweeps. Furthermore, the amount of memory required to hold the necessary integrals was beyond our capacity for large basis sets. See Table 2.

[4] It can be shown that, just as in the ER case, the Boys function is maximized when $\mathbf{R}^{Boys}$ is symmetric.

We then computed (via Jacobi sweeps) "rough" Boys' orbitals, defined as Boys' orbitals with $error_{Boys} < 1.0 \times 10^{-2}$. With these rough orbitals as our initial guess, we counted separately both the number of Jacobi sweeps and the number of DIIS-like steps required before $error_{Boys} < 1.0 \times 10^{-5}$. Results for the simple molecule benzene (C_6H_6) are given in Table 1.

Recalling that each individual DIIS-step requires evaluation of just one or two $\mathbf{R}$ matrices, whereas Jacobi sweeps loop over all orbital pairs and compute integrals, we conclude that our algorithm competes with Jacobi sweeps for convergence.

4.2   DIIS vs. 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. For reference, we compared the speed of convergence of out DIIS-like algorithms with the ordinary η-step algorithm. As an initial guess of localized orbitals, $\{\chi^{(0)}\}$, we used the the Boys' orbitals from the occupied SCF space. We defined the ER error as $error_{ER} = \max_{ij} |R_{ij} - R_{ji}|$ and declared convergence when $error_{ER} < 1.0 \times 10^{-5}$.

In Table 2, we record the total number of iterations required to start from the Boys' orbitals and converge to the ER orbitals. We provide data from three separate calculations, one using only η iterations and the other (almost only) two DIIS-like iterations. In all of these calculations, when the algorithm had found a stationary point, a full second-derivative check was employed to make sure that the stationary point was a maximum. When a saddle point was encountered, we stopped our count of the iterations, and did a one line-minimization in the direction v_i, where v_i is the eigenvector of the Hessian with the most positive eigenvalue. Subsequently, we ran twenty (uncounted) η iterations in order to improve our starting point. At that point, we turned the count back on, monitoring the number of iterations (either η or DIIS-like) until convergence.

The data in Table 2 demonstrates how much faster calculations become when using DIIS-like interpolations instead of pure η 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 behavior 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 $\eta-$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 (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 don't 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.

4.3 Scaling with System Size: Alkanes

The bottleneck in doing the calculations above is the evaluation of the matrix elements $(ij|jj)$. However, using the well Boys' orbitals as our initial guess, we have generated these integrals in sub-quadratic (and potentially linear) time. We computed these integrals by first making n matrices of the form $K_{\mu\nu}^{(i)} = (\mu i|\nu i)$, 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. [16]. However, when we exploit the locality of our guess orbitals (and hence the sparsity of the density matrix $P_{\mu\nu}^i$), we can form $\left\{K_{\mu\nu}^{(i)}\right\}$ in a ("linear") time proportional to N.[17]

In Figures 1a and 1b, 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 $K_{\mu\nu}^{(i)}$, 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 $K_{\mu\nu}^{(i)}$, and this time grows cubicly and prohibitively. For the "linear-scaling" algorithm, by contrast, the integrals are generated in linear time (as ex-

pected). Furthermore, the digestion of the integrals into $K_{\mu\nu}^{(i)}$ 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 asymptotic 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, given 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.

5   A chemical application: Nitration and Nitrosation of Benzene

The highly different reactivies of nitronium (NO_2^+) and nitrosonium (NO^+) toward benzene has been investigated from a theoretical perspective by Gwaltney et al.[19]. There has also been a recent comprehensive DFT study on benzene-NO_2^+[20]. 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 of NO^+ as vertically aligned over the benzene plane–see Figure 2a). 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 pre-

vious 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 Figure 2a. 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 Figure 3a.[5] (This geometry corresponds to Structure 1 in the paper by Olah *et al* [20].) 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 Figure 2a and Figure 3a (i.e. NO^+ 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 Figure 2b. 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 Figure 3a and a two-dimensional contour map is given in Figure 3b. (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 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

---

[5]We note that the geometries of both Figure 2a and Figure 3a are saddle points of the B3LYP functional, rather than minima.

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 MO's 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 Figure 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 Figure 3a, we note that this geometry, like Figure 3a, is also a saddle point of the energy function (rather than a minimum)[19][20].[6] 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.

6 Discussion

Though the methods presented here can be powerful tools for solving symmetric 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 $\mathbf{R}$ matrices. For the ER function, $R_{ij}^{ER} = (ij|jj)$, while for the Boys function, $R_{ij}^{Boys} = <i|\mathbf{r}|i> \cdot <i|\mathbf{r}|j>$. Stationary points occur when $\mathbf{R}$ is symmetric, and each step requires inversion of $\mathbf{RR}^T$. As such, there are obvious difficulties when $\mathbf{R}$ becomes singular. Any possible physical significance of singular $\mathbf{R}$ matrices is unclear and certainly is specific to the function being maximized. As a rule of thumb, we have had no problems with $\mathbf{R}$-singularity when working with the ER function. However, the

---

[6]The geometry of Figure 4 corresponds to Structure 7 in the paper by Olah *et al.*

15

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 [23][24]) 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 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, like 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.

7 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 best 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 wavefunctions. If they can can be computed quickly, the ER orbitals may well be used as building blocks in future local-correlation work because, by minimizing $\sum_{i \neq j} (ij|ij)$, 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 symmetric 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.

Acknowledgements

8 Appendix

8.1 A Mathematical Lemma

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

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

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, $g(O) = \text{Tr}(\mathbf{AUO}) = \text{Tr}(\mathbf{BO}) = \sum_{i,j} B_{ij} O_{ji}$, where $\mathbf{B} = \mathbf{AU}$.

8.2 Stationary Points

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

$$
\begin{aligned}
\left(\frac{\partial g}{\Delta_{pq}}\right)|_{\Delta=0} &= \sum_{i,j} B_{ij} \frac{\partial O_{ji}}{\Delta_{pq}}|_{\Delta=0} \\
&= \sum_{i,j} B_{ij} \left(\delta_{jp}\delta_{iq} - \delta_{jq}\delta_{ip}\right) \\
&= B_{qp} - B_{pq} \\
&= \sum_i A_{qi} U_{ip} - A_{pi} U_{iq}
\end{aligned}
$$

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: $\mathbf{A} = \mathbf{N}\tilde{\mathbf{U}}$ 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
$$\mathbf{B} = \mathbf{NV} = (\mathbf{NV})^T = \mathbf{V}^T\mathbf{N} = \mathbf{B}^T \tag{7}$$

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 eqn (7) 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

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

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 $t_{ii} = \sum_j w_{ij} w_{ji}$ and $\sum_j w_{ij}^2 = 1$, it follows by Cauchy-Schwarz that $|t_{ii}| < 1$. Moreover, $\mathrm{Tr}(\mathbf{\Lambda}) = \mathrm{Tr}(\mathbf{W}\mathbf{\Lambda}\mathbf{W}) = \mathrm{Tr}(\mathbf{W}^2\mathbf{\Lambda})$ 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, eqn (7) 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 into in terms of $\mathbf{A}$, using the polarization formulae $\mathbf{A} = \mathbf{N}\tilde{\mathbf{U}}$ and $\mathbf{N} = \sqrt{\mathbf{A}\mathbf{A}^T}$ (here we mean the postive square root). Then, $\tilde{\mathbf{U}} = \mathbf{N}^{-1}\mathbf{A} = \left(\mathbf{A}\mathbf{A}^T\right)^{-1/2}\mathbf{A}$. It follows that

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

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.

8.3   Maxima

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

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

Hence,

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

To check for negative-definiteness, note that:

$$
\begin{aligned}
\sum_{pqrs} M_{pq}\frac{\partial^2 g}{\partial \Delta_{pq}\partial \Delta_{rs}}\big|_{\mathbf{\Delta}=0}M_{rs} &= M_{pr}B_{sp}M_{rs}+M_{sq}B_{qr}M_{rs} \\
&\quad -M_{ps}B_{rp}M_{rs}-M_{rq}B_{sq}M_{rs} \\
&= \mathrm{Tr}\left(\mathbf{M}^T\mathbf{B}\mathbf{M}^T+\mathbf{M}\mathbf{B}\mathbf{M}-\mathbf{M}^T\mathbf{B}\mathbf{M}-\mathbf{M}\mathbf{B}\mathbf{M}^T\right) \\
&= \mathrm{Tr}\left(\mathbf{B}\left(\mathbf{M}^T\mathbf{M}^T+\mathbf{M}\mathbf{M}-\mathbf{M}^T\mathbf{M}-\mathbf{M}\mathbf{M}^T\right)\right)
\end{aligned}
$$

It is easy to see that

$$
\begin{aligned}
Y &= \mathbf{M}^T\mathbf{M}^T+\mathbf{M}\mathbf{M}-\mathbf{M}^T\mathbf{M}-\mathbf{M}\mathbf{M}^T \\
&= -\left(\mathbf{M}-\mathbf{M}^T\right)\cdot\left(\mathbf{M}-\mathbf{M}^T\right)^T
\end{aligned}
$$

is a non-positive matrix. Because we are interested only in antisymmetric $\mathbf{M}$ (the tangent space of the orthogonal group), it follows $\mathbf{Y}\neq 0$. 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 last section, $\mathbf{B}=\mathbf{N}\mathbf{V}=\mathbf{D}^T\mathbf{\Lambda}\mathbf{D}$ where $\mathbf{\Lambda}=\mathbf{\Lambda}_N\mathbf{\Lambda}_V$.

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

$$
\mathrm{Tr}(\mathbf{B}\mathbf{Y})=\mathrm{Tr}(\mathbf{D}^T\mathbf{\Lambda}_+\mathbf{D}\mathbf{Y})=\mathrm{Tr}(\mathbf{\Lambda}_+\mathbf{D}\mathbf{Y}\mathbf{D}^T)<0
$$

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

$$
\begin{aligned}
S_{\alpha\beta} &= \delta_{\alpha i}\delta_{\beta j}-\delta_{\alpha j}\delta_{\beta i} \\
\left(S^2\right)_{\alpha\beta} &= -\delta_{\alpha i}\delta_{\beta i}-\delta_{\alpha j}\delta_{\beta j}
\end{aligned}
$$

Then the requirement

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

demands that $\forall i,j \quad \Lambda_{ii} + \Lambda_{jj} > 0$. So we can have at most one negative eigenvalue of $\mathbf{\Lambda}$. But for $SO(n)$, we further demand that $\Lambda_{11}\Lambda_{22}\cdots\Lambda_{nn} > 0$ 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 $\mathbf{B} = \mathbf{D}^T \mathbf{\Lambda} \mathbf{D} = \mathbf{D}^T \mathbf{\Lambda}_N \mathbf{\Lambda}_V \mathbf{D}$ 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

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

where here we take the positive square root. $\qquad\qquad \square$.

8.4 Carlson-Keller Corollary

One should note that the theorem above provides immediate (re)proof of the Carlson-Keller theorem[10], that $\mathbf{S}^{-1/2}$ is the transformation constructing orthogonal orbitals that most resemble a set of initial non-orthogonal orbitals. In that case, if $\chi_i = \sum_j \phi_j C_{ji}$ and $\mathbf{C} = \mathbf{S}^{-1/2}\mathbf{U}$ for $\mathbf{U}$ orthogonal, one wants to minimize the function:

$$
\begin{aligned}
h(\mathbf{C}) &= \sum_i \left(\chi_i - \phi_i | \chi_i - \phi_i\right) = \text{constant} - 2\sum_i \left(\chi_i | \phi_i\right) \\
&= \text{constant} - 2\sum_{ij} C_{ji}\left(\phi_j | \phi_i\right) = \text{constant} - \text{Tr}(\mathbf{S}^{1/2}\mathbf{U})
\end{aligned}
$$

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 [11], that h has exactly one minimum, and that is $\mathbf{C} = \mathbf{S}^{-1/2}$.

8.5 Generalized $\eta-$step

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

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

21

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,

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

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

References

[1] Saebo, S. and Pulay, P. *Annu. Rev. Phys. Chem.* **44**, 213-236 (1993).

[2] S.F. Boys. *Rev. Mod. Phys.*, **32** (2),296-302 (1960).

[3] S.F. Boys. *Quantum Theory of Atoms, Molecules and the Solid State*, ed. P.O. Lowdin. (Academic Press, NY: 1966). pp. 253-262.

[4] C. Edmiston and K. Ruedenberg. *Rev. Mod. Phys.*, **35**(3), 457-465 (1963).

[5] J. Pipek and P.G. Mezey. *J. Chem. Phys.*, **90**, (19),4916-4926 (1989).

[6] N. Marzari, I. Souza, and D. Vanderbilt. *Highlight of the Month, Psi-K Newsletter*, **57**, 129-168 (2003).
http://psi-k.dl.ac.uk/psi-k/newsletters.html.

[7] I. Souza, N. Marzari, and D. Vanderbilt. *Phys. Rev. B*, **65** (3), 035109/1-13 (2002).

[8] G. Berghold, C.J. Mundy, A.H. Romero, J. Hutter, and M. Parrinello. *Phys. Rev. B* **61** (15), 10040-8 (2000).

[9] J.E. Lennard-Jones and J.A. Pople, *Proc. Roy. Soc. London*, **202**, 166-180 (1950).

[10] B.C. Carlson and J.K. Keller. *Phys. Rev.*, **105**(1), 102-103 (1957).

[11] J. Aiken,J. Erdos, and J. Goldstein. *Intl. J. Quant. Chem.*, **18**, 1101-1108 (1980).

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

[13] P. Pulay. *Chem. Phys. Lett.*, **73**(2), 393-398 (1980).

[14] P. Pulay. *J. Comp. Chem.*, **3**(4), 556-560 (1982).

[15] D.A. Kleier, T.A. Halgren, J.H. Hall,Jr., and W.N. Lipscomb. *J. Chem. Phys.*, **61**, 3905-3919 (1974).

[16] M.J.Frisch, M.Head-Gordon and J.A.Pople, *Chem. Phys.* **141** 189-196 (1990).

[17] W. Liang, Y. Shao, C. Ochsenfeld, A.T. Bell, and M. Head-Gordon. *Chem. Phys. Lett.*, **358**, 43-50 (2002).

[18] J. Kong *et al. J. Comp. Chem*, **21**, 1532-1548 (2000).

[19] S. R. Gwaltney, S. V. Rosokha, M. Head-Gordon, and J.K. Kochi. *JACS*, **125**, 3273-3283 (2003).

[20] 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. *JACS*, **125** (16), 4836-4849 (2003).

[21] S. Skokov, and R. A. Wheeler. *J. Phys. Chem.*, **103**, 4261-4269 (1999).

[22] A. Edelman, T.A. Arias, and S.T.Smith. *SIAM J. Matrix Anal. Appl.*, **20** (2), 303-353 (1998).

[23] E.R. Davidson. *J. Comp. Phys.*, **17**, 87-94 (1975).

[24] M.L. Leininger, C.D. Sherrill, W.D. Allen, and H.F. Schaefer III. *J. Comp. Chem.*, **22**, 1574-1589 (2001).

Tables

Table 1: Number of Jacobi sweeps and DIIS-like iterations necessary to converge Boys orbitals for benzene (see text for convergence criteria and initial guess.)

Basis	Basis Size	Num. Jacobi sweeps before convergence	Num. DIIS-like steps before convergence (DIIS-1 and DIIS-2)
STO-3G	36	11	14
6-31G	66	10	17
6-31G*	102	10	15
6-31G**	120	10	13
6-311G**	144	10	15

Table 2: Number of η-steps and DIIS-like steps necessary to converge to ER orbitals for benzene (see text for convergence criteria and initial guess.)

| Basis	Basis Size	Num. Saddle Points	Num. DIIS-1 steps before convergence	Num. DIIS-2 steps before convergence	Num. of η steps before convergence
STO-3G	36	0	8	8	19
6-31G	66	1	16	16	486
6-31G*	102	1	16	15	406
6-31G**	120	1	16	15	395
6-311G**	144	1	15	15	366

Figures

Figure 1a: 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 Figure and Figure 1b. 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).

Figure 1b: 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 this Figure and Figure 1a. 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 Figure 1a, for each point on this graph, exactly seven DIIS-2 iterations were required for convergence.)

Figure 2a: 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. (Green = Carbon, White = Hydrogen, Red = Oxygen, Blue = Nitrogen). This structure was optimized with N-O vertically aligned over the center of the benzene ring. Perpendicular distance from N to benzene is 2.16Å.

Figure 2b: 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).

Figure 3a: Nuclear geometry of Benzene-NO_2^+, over which is plotted a three dimensional spatial representation of the one (unique) delocalized donor-acceptor ER orbital. (Green = Carbon, White = Hydrogen, Red = Oxygen, Blue = Nitrogen) 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[20]. The distance from the Nitrogen atom to the benzene plane is 3.1Å.

Figure 3b: 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^+.

Figure 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. (Green = Carbon, White = Hydrogen, Red = Oxygen, Blue = Nitrogen) This structure is the MP2-optimized structure discussed in [19]. The interaction here is stronger than the interaction in Figure 3a 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 2.17Å.