Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Linear-scaling evaluation of the local energy in quantum Monte Carlo

Permalink

https://escholarship.org/uc/item/2z06r8sf

Authors

Austin, Brian
Aspuru-Guzik, Alan
Salomon-Ferrer, Romelia
et al.

Publication Date

2006-02-11

Linear-scaling evaluation of the local energy in quantum Monte Carlo

Brian Austin[1,2], Alán Aspuru-Guzik[1], Romelia Salomón-Ferrer[1,2], and William A. Lester, Jr.[1,2]

[1]Kenneth S. Pitzer Center for Theoretical Chemistry, Department of Chemistry, University of California at Berkeley, Berkeley, CA 94720-1460
[2]Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Abstract

For atomic and molecular quantum Monte Carlo calculations, most of the computational effort is spent in the evaluation of the local energy. We describe a scheme for reducing the computational cost of the evaluation of the Slater determinants and correlation function for the correlated molecular orbital (CMO) ansatz. A sparse representation of the Slater determinants makes possible efficient evaluation of molecular orbitals. A modification to the scaled distance function facilitates a linear scaling implementation of the Schmidt-Moskowitz-Boys-Handy (SMBH) correlation function that preserves the efficient matrix multiplication structure of the SMBH function. For the evaluation of the local energy, these two methods lead to asymptotic linear scaling with respect to the size of the molecule.

In the quantum Monte Carlo (QMC) method [1–7], the expectation value of the Hamiltonian $\langle H \rangle$ is computed as a statistical average of the local energy of a trial wave-function, $\Psi_T(\mathbf{R})$, where $\mathbf{R}$ denotes the 3N coordinates of the N-particle system.

$$E_L(\mathbf{R}) \equiv \frac{\widehat{H}\Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R})}, \tag{1}$$

Averages of this quantity provide an estimator of the energy of the system,

$$E \approx \langle E_L \rangle = \lim_{N_s \to \infty} \frac{1}{N_s} \sum_i^{N_s} E_L(\mathbf{R}_i). \qquad (2)$$

where N_s is the number of sample points $\mathbf{R}_i$ and the local energy is evaluated during the random walk. The points $\mathbf{R}_i$ are sampled from a probability function $f(\mathbf{R}, \tau)$, such as Ψ_T^2 for variational Monte Carlo [1]. Presently, the more commonly used QMC version is the diffusion Monte Carlo (DMC) that yields time-independent solutions of the Schrödinger equation from an imaginary time stochastic formalism [8].

For most molecular applications, the correlated molecular orbital (CMO) wave function form is used.

$$\Psi_T = \mathcal{D} \times \mathcal{F} = \left[\sum_k c_k \cdot D_k^\alpha D_k^\beta \right] \times \left[e^{U(\mathbf{r}_{ij})} \right] \qquad (3)$$

Here $\mathcal{D}$ is a linear combination of Slater determinants and $\mathcal{F}$ is a correlation function that depends on inter-particle distances. The determinants D_i^σ, where σ denotes α or β spin, are usually constructed with molecular orbitals (MOs) obtained in a previous ab initio calculation such as Hartree-Fock (HF), multi-configuration self consistent field theory (MCSCF), or density functional theory (DFT).

Interactions between electrons in a molecule diminish exponentially as the distance between them increases [9]. In recent years, linear diffusion Monte Carlo (LDMC) methods [10–12] have been introduced to take advantage of this property to reduce the computational demands of QMC to linear dependence on system size. Williamson *et al.* [10] developed a linear-scaling method in a plane-wave basis expansion for periodic systems and applied the approach in several studies [13–15]. Alfè and Gillan [12] developed a method that employs non-orthogonal localized orbitals in a plane-wave basis expansion that further improves efficiency. Manten and Lüchow (ML) [11, 16] have presented a similar method that uses Gaussian-type orbitals (GTOs).

ML have also shown how parts of the DMC algorithm, including the evaluation of the Coulomb potential and inter-particle correlation function, can be made to scale linearly with system size[11, 16]. In the present work, we address the two most important contributions to the computational cost of wave function evaluation, namely, the evaluation of the Slater determinant [17] and that of the correlation function. We will describe a fast implementation of these two steps in the following two sections. Progress towards deeper understanding of results from QMC is described in the last section.

Fast evaluation of the Slater determinant

One purpose of LDMC is to decrease computational time by reducing computation associated with wave function evalution. The approach of Williamson *et al.* [10] has the potential of near-optimal efficiency because only a single three-dimensional spline (3DS) evaluation is required for a given molecular orbital (MO). Maximally-localized Wannier (MLW) functions of systems using effective core potentials (ECPs) are smooth functions that can be described accurately with 3DS. All-electron MOs expressed in a Slater or Gaussian basis set are rapidly-varying with distance, and their representation with 3DS is memory intensive due to the cubic scaling of 3DS memory requirements with MO complexity. On the other hand, Gaussian or Slater basis sets are more compact than plane waves and therefore amenable to being evaluated explicitly without substantially increasing computational cost as proposed by ML [11, 16]. We improve upon the work of ML by introducing a sparse representation of the Slater matrix and cutoff criteria that derive from the sparse representation. The method proposed here can be extended to 3DS-based methods to achieve further computational efficiency.

Evaluation of $\mathcal{D}$ of Eq. 3 requires the construction of Slater matrices, $\mathbf{D}_i^\sigma$, and the corresponding determinants $D_i^\sigma \equiv \det(\mathbf{D}_i^\sigma)$. Note that naive construction of the Slater matrix leads to cubic scaling with system size.

The main feature of the method is the use of a three-dimensional grid to store the non-zero contributions to the Slater matrix. The use of the grid does not affect the accuracy of the wave function because the orbitals are evaluated using accurate one-dimensional representations of Slater basis functions. The memory requirements for the grid are of the order of a few megabytes for the polyalanine systems studied [17].

Evaluation of determinants $\mathcal{D}$ involves an additional cubic step associated with the inversion of $\mathbf{D}_i^\sigma$ [18]. Although ML proposed using a sparse linear algebra routine for this purpose, an optimized algorithm for dense matrices [19] is used in the present approach because inversion of $\mathbf{D}_i^\sigma$ is not a computational bottleneck for the size of systems treated here.

Slater matrices require evaluation of MOs for all electrons of a given spin σ. An MO ϕ is expanded in a basis of atom-centered functions, typically GTOs or Slater-type orbitals (STOs),

$$\phi_i(\mathbf{r}) = \sum_{\mu=1}^{N_{bas}} C_{\mu i} \chi_\mu(\mathbf{r}) \tag{4}$$

The $C_{\mu i}$ elements form the coefficient matrix $\mathbf{C}$. The atom-centered functions χ are usually expressed as linear combinations of primitive functions φ,

$$\chi_i(\mathbf{r}) = \sum_{k=1}^{N_{cont}} c_{ki} \cdot \varphi_k(\mathbf{r}) \tag{5}$$

It may seem that construction of a Slater matrix is cubic scaling, because the evaluation of Eq. 5 requires iteration over basis functions, MOs and positions. However, MO evaluation can be made linear with system size by the use of localized MOs (LMOs) coupled with a sensible cutoff of LMO spatial extent. These two steps are essential for achieving linear scaling with system size.

Most ab initio computer programs employ orthogonal LMOs. The orthogonalization restriction has the disadvantage of producing "orthogonalization tails", i.e., small contributions to the LMOs at centers that lie far from the main contribution to the LMO [20]. To accelerate the onset of linear scaling, ML removed the orthogonalization tails [11]. In this study, we chose not to do so in order to preserve fully the nodal characteristics of the LMOs.

We next present a brief explanation of the steps involved in the calculation of the Slater matrix. The first step required for the sparse representation of the Slater matrix is to generate a pre-sorted coefficient matrix. For CMOs of the form of Eq. 3, Slater matrices are constructed separately for α and β electrons leading to sorted coefficient matrices, $\tilde{\mathbf{C}}_i^{\alpha}$ and $\tilde{\mathbf{C}}_i^{\beta}$, built from $\mathbf{C}_i^{\alpha}$ and $\mathbf{C}_i^{\beta}$, respectively.

The rows of $\tilde{\mathbf{C}}_i^{\alpha}$ run over the basis functions, while the columns run over the LMO coefficients sorted by their absolute value,

$$|\tilde{C}_{1i}| \geq |\tilde{C}_{2i}| \geq \ldots |\tilde{C}_{N_{mos}i}| \tag{6}$$

It is convenient to construct an auxiliary array that maps indexes of $\tilde{\mathbf{C}}_i^{\alpha}$ and $\tilde{\mathbf{C}}_i^{\beta}$ to the indexes of the Slater matrices $\mathbf{D}_i^{\sigma}$. This step facilitates filling the columns of the Slater matrices in the correct order.

Grid Generation

For an electron located at $\mathbf{r}$, the value of a LMO ϕ_i is a sum of products $p_{\mu i}$ of basis functions and LMO coefficients for each LMO,

$$\phi_i(\mathbf{r}) = \sum_{\mu=1}^{N_{bas}} p_{\mu i} = \sum_{\mu=1}^{N_{bas}} C_{\mu i}\chi_{\mu}(\mathbf{r}). \tag{7}$$

In the present LDMC approach, we only sum over products that are greater than a numerical threshold ε_c. This threshold is directly connected to the numerical precision of the truncated trial wave function. This is the single most important parameter for the onset of linear scaling.

When an electron is evaluated at position $\mathbf{r}$, the nearest point from a 3D grid is referenced. Each 3D grid element is associated with a compressed representation of the Slater matrix that describes the non-zero elements of a column of the Slater matrix evaluated at the particular volume element. The compression scheme is depicted in Figure 1.
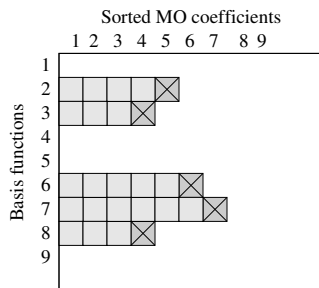
Figure 1: The compressed matrix representation used for this work. The non-zero basis function index is stored along with the index of the smallest LMO coefficient that multiplies the basis function. These pairs of indexes are designated by an x and for the example correspond to (2,5)(3,4)(6,6)(7,7)(8,4).

Besides the numerical threshold ε_c, a second parameter is introduced. It is the fraction (f) of products $p_{\mu i}$ retained. The parameters f and ε_c are related. For small thresholds, a higher fraction of products can be discarded without change in the value of the wave function. At higher thresholds, a smaller fraction of products can be removed from the calculation.

Determinant Evaluation Algorithm

In DMC electrons can be moved one at a time [8], or all at the same time [21]. For this discussion, we assume that the electrons are moved one at a time. On this basis we present a procedure for the construction of a column of a Slater matrix corresponding to an electron located at a position **r**.

With CMO wave functions, electrons have a given spin σ throughout the simulation. As soon as one specifies the electron index, and therefore spin, one can use the grid and sorted coefficient matrix associated with the particular spin. With the electron spin specified, the grid element corresponding to the spatial location is identified, and then a sequence of doublets $(\mu_1, i_1), (\mu_2, i_2) \ldots (\mu_n, i_n)$ is read. Here, $\mu_1, \mu_2, \ldots, \mu_n$ are the indexes of the n non-zero basis functions, and $i_1, i_2, \ldots, i_m$ are the maximum LMO coefficient indexes needed to evaluate the products $p_{\mu i}$. The Slater matrix columns are updated in iterations over electrons. A pseudo-code representation of the algorithm is illustrated in Figure 2.

Rapid Evaluation of Primitive Basis Functions

STOs have been shown to be advantageous for QMC owing to their correct asymptotic behavior at short and long electron-nucleus distances. To employ wave func-

```
Require: x, y, z ⊂ Grid {ensure that the point is inside the grid}
   DoubletList = Grid^σ(x, y, z)
   SortedCoefficients = PresortedCoefficients^σ
   SortKey = SortKey^σ
   SlaterMatrix = SlaterMatrix^σ
   while Doublet from DoubletList: do
      (μ, i_max) = Doublet
      Evaluate χ_μ(x, y, z)
      for i = 0 to i = i_max do
         c = SortedCoefficients(μ, i)
         j = SortKey(μ, i)
         SlaterMatrix(j, k) = φ_j(x, y, z) = φ_j(x, y, z) + c * χ_μ(x, y, z)
      end for
   end while
```

Figure 2: Algorithm for the evaluation of a column k of the Slater matrix for an electron of spin σ at a point $\mathbf{r} = [x, y, z]$.

tions expanded in an STO basis, we use the Amsterdam Density Functional (ADF) package [22].

For rapid evaluation, the radial components of the contractions χ_j are interpolated using cubic splines, and cut-off at an appropriate distance from their center. Cut-off values are determined so that the local energy of a fixed sample of walkers changes by less than a prescribed percentage. Splines make possible the representation of a linear combination of all STOs belonging to a given contraction by a single spline function, with the result of reducing the computational cost. For the remainder of this discussion, we will refer to such representations of contracted STOs as basis functions.

Linear-Scaling Evalation of the Correlation Function

Rapid evaluation of the correlation function is equally important to obtain a linear scaling QMC algorithm. The correlation function of Schmidt and Moskowitz (SM) [23], which stems from the form suggested by Boys and Handy (BH) [24] is displayed in Equation 8.

$$U = \sum_{\mu} \sum_{A}^{atoms} c_{\mu A} \sum_{i<j}^{electrons} U^{\mu}_{Aij} \tag{8}$$

where,

$$U^{\mu}_{Aij} = (\bar{r}^{l_\mu}_{Ai} \bar{r}^{m_\mu}_{Aj} + \bar{r}^{l_\mu}_{Aj} \bar{r}^{m_\mu}_{Ai}) \bar{r}^{n_\mu}_{ij}; \qquad \bar{r}(r) = \frac{r}{1+br}$$

An important step in accelerating evaluation of the SMBH correlation function was taken by ML. These authors changed the scaled distance function to $\bar{r} = 1 - \exp[-\alpha r]$ leading to a shorter cutoff distance in their linear scaling algorithm.

A further change to $\bar{r}$ is required for the present approach. This change is motivated by examining the asymptotic behavior of the three-body terms in the SMBH expansion. Consider the (lmn)=(011) term:
$U_{ijA} = (1 - \exp[-\alpha r_{ij}])(1 - \exp[-\alpha r_{jA}])$.

Because ML's $\bar{r}$ approaches unity at large distances, U_{ijA} reduces to the two-body (lmn)=(010) term when electron i is separated from the jA pair, leaving many nonzero terms in the SMBH functions. To reduce the number of nonzero terms, we simply shift $\bar{r}$ so that its asymptotic value is zero: $\bar{r} = -\exp[-\alpha r]$.

The consequences of this change are simple to interpret. Although physical arguments would seem to favor an ansatz where three-body terms reduce to two-body terms, these extra terms are merely redundant descriptions of the two-body correlations included in the (lmn)=(010) term. For two-body terms, this merely changes the normalization of the wavefunction. Linear dependence among all the terms is also reduced. Our modification to $\bar{r}$ may also be alternatively viewed as a generalization of the Sun-Lester correlation function[25], which is included as the (lmn)=(001) term of the SMBH expansion.

The next step in the development of our linear scaling algorithm requires recognition that the SMBH function can be rewritten as a trace over a matrix product,

$$U = \sum_{\mu} \sum_{A} \sum_{i \neq j} c_{\mu A} \bar{r}^{l_\mu}_{Ai} \bar{r}^{n_\mu}_{ij} \bar{r}^{m_\mu}_{jA} \tag{9}$$

Simpler terms where l, m, or n are zero can be handled by deleting the appropriate matrix from Equation 9. The term where m_μ and l_μ are swapped in Equation 8 is unneccessary in the matrix form because the sum over i and j is no longer lower triangular. It is easy to exclude terms where $i = j$ by setting diagonal elements of the $\bar{r}_{ij}$ matrix to zero.

Because BLAS [26] libraries make matrix operations particulary efficient, we observed speedups of at least a factor of ten relative to our initial implementation after switching to a matrix-based implementation. Matrix multiplication, however, scales cubically with system size. To achieve linear scaling, we take advantage of the sparsity in the $\bar{r}$ matrices that was created by our modification to the scaled distance function.

```
for each atom, A do
    Make a list of electrons near A.
    Build the vectors of e-n distances $\bar{r}_i^{m[A]}$ and $\bar{r}_j^{l[A]}$.
    Build the matrix of distances $\bar{r}_{ij}^{[A]}$ for electrons near A.
    Perform the matrix operations: $U_A^\mu = \sum_{ij} \bar{r}_i^{m[A]} \bar{r}_{ij}^{n[A]} \bar{r}_j^{l[A]}$.
    Increment $U$: $U = U + c_{\mu A} U_A^\mu$.
end for
```

Figure 3: Algorithm for computing one term in the SMBH expansion using sparse matrices.

Figure 3 sketches our sparse matrix multiplication routine. A distance-based cutoff is used to determine which $\bar{r}$ elements need to be computed. By limiting operations to electrons near each atom, we implicitly compress the nonzero elements of $\bar{r}_{iA}$ and $\bar{r}_{ij}$ into a set of three $O(1)$ blocks ($\bar{r}_i^{[A]m}$, $\bar{r}_{ij}^{[A]n}$ and $\bar{r}_j^{[A]l}$) for each atom. The cost of the matrix operations for each atom is therefore $O(1)$, yielding an overall linear scaling algorithm.

Application to Biological Systems

As an example of the capability of treating large systems, we show some results for calculations on two biomolecules present in Photosystem II, bacteriochorophyll (Bchl) and spheroidene (Spo), containing 304 and 314 electrons respectively. Under high light conditions, the generation of chlorophyll triplet states and singlet oxygen, $^1O_2^*$, can increase dramatically due to the fact that the light flux exceeds the amount of light that the reaction center is able to process (turn-over capacity). Carotenoids help protect the organism by non-photochemical quenching (NPQ), which dissipates the excess excitation energy of singlet oxygen. If carotenes are absent, chlorophyll is readily oxidized, leading to the organism's death. Several scientific questions about the photoprotection mechanism remain unanswered. The ground-state to lowest-triplet-state excitation energies of both molecules as well as the energy transfer rate between the systems have not been determined. At present, DFT is not sufficiently accurate to answer these questions. A rigorous electronic structure method that accurately treats electron correlation is needed, but, as noted, basis-set correlated methods scale too steeply with system size to be applicable. DMC presently offers the only feasible option for the accurate calculation of these quantities.

To understand chemical processes, it is useful to have information besides total energies. Electron localization methods provide a insight on the behavior of electrons in molecules. Properties such as electron density, spin density and the electron pair localization function (EPLF) [27] can routinely be computed by post-processing. The EPLF provides a quantitative description of electron pairing in molecular systems and has similarities to the electron localization function (ELF) of Becke and Edgcombe [28]. The QMC method is a particularly well suited approach for obtaining such information because the simple and general definition of EPLF is easily evaluated in QMC.

In the next section, some details on the speedup in calculating the local energy obtained by the application of the techiques mentioned above are presented. Results for spheroidene as well as other molecules of biological importance are given.

Results and Discussion

Local energies for the spheroidene molecule were computed with a series of different grid sizes. Our tests show a factor of 10 speedup with respect to a standard calculation for a grid of 2660 elements. The basis-function cut offs were in the 11.0-15.0 a.u range, following the criteria for STOs given above. The grid memory requirements are modest; for spheroidene, a 20520 element grid requires 3.4 MB of RAM.

CPU timings for a sample of eight molecules of biological interest have been obtained. We chose the series of 2-,4-,6- and 8-polyalanine in an alpha helix conformation. We also selected four other molecules with different conformations to show dependence on geometrical factors. For the calculation of the eight molecules, we used a grid element length of 2 a.u. and a cutoff threshold of $\varepsilon_c = 1x10^{-12}$. The whole fraction ($f = 1.0$) of products was preserved for all calculations.

Geometrical effects can be compared as well. With fewer basis functions per volume element, linear molecules require less computer time than 2D or 3D arrangements. For example, the spheroidene molecule required 75% of the CPU time of 8-alanine, which has the same number of electrons. A model bacteriochlorophyll with ten fewer electrons required 124% of the time 8-alanine.

The CPU timings for the evalution of the Slater determinant for 8 molecules are given in Figure 4. After around 80 electrons, the CPU time for filling the Slater matrices scales linearly with the number of electrons.

A series of linear alkanes have been used to determine CPU scaling for evaluating the SMBH correlation function. Figure 5 presents timing data using a matrix
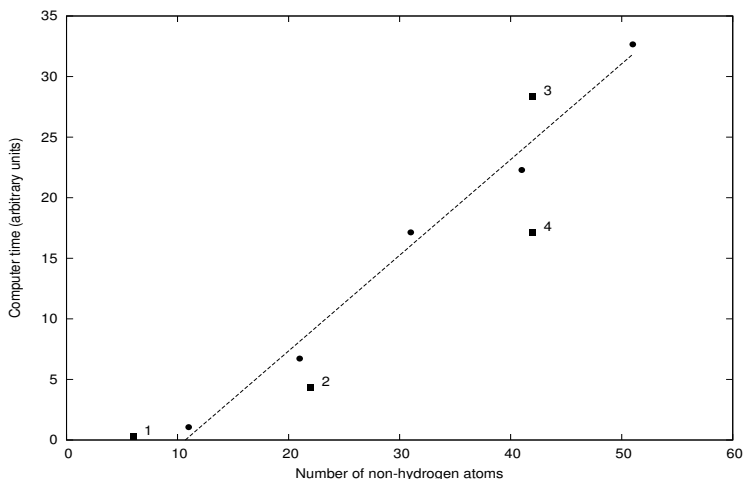
Figure 4: Computer time for filling the Slater matrices, $\mathbf{D}_i^\alpha$ and $\mathbf{D}_i^\beta$, as a function of the number of non-hydrogen atoms for different biological molecules. Data for a series of poly-alanine molecules (2-,4-,6- and 8-polyalanine) is plotted as circles. (1) hexatriene, (2) capsaicin, (3) bacteriochlorophyll and (4) spheroidene are designated by squares. The CPU time for polyalanine chains scales linearly with system size. Data are from reference [17].

based algorithm with dense $\bar{r}$ matrices and the linear scaling algorithm with the present $\bar{r}$. For these tests, we set $\alpha = 4.0$ and determined empirically that a cutoff of 4.0 a.u. affected local energies by less than 10^{-7} Hartrees. The sparse algorithm is linear scaling for all molecules studied and is faster than the dense algorithm for molecules with more than three carbon atoms.

Figure 6 presents slices of the EPLF function for spheroidene in the ground state singlet and the first excited triplet state in the $z = 0$ plane. The different values represent the different levels of electron pair localization for that particular region, larger values indicate higher degree of localization. Most of the EPLF values obtained for this system are less than 0.1, which indicates a high level of delocalization in this system due to its conjugated polyene structure. Important differences can be seen between the singlet and the triplet state. While the singlet appears to have no spin-polarized (gray) regions, the triplet does possess such regions near the outer boundary of the molecule. Grey regions show where excess alpha electron density is paired.
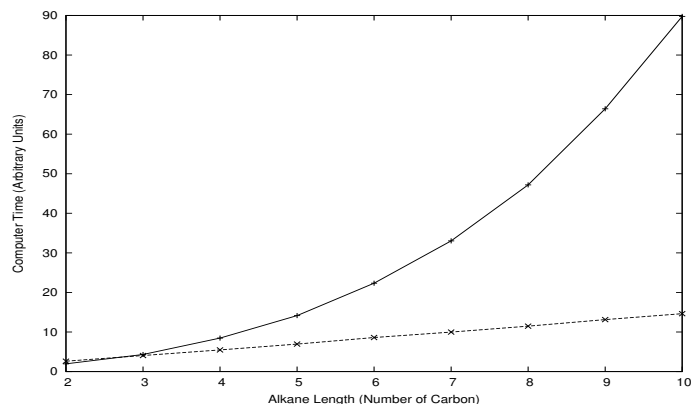
Figure 5: Computer time for evaluating the SMBH correlation function versus alkane length. The dense matrix-based algorithm (solid line) scales roughly cubically with system size. The sparse matrix based algorithm (dashed line) scales linearly.

Conclusions

In this study, we have presented and demonstrated a novel approach in quantum Monte Carlo for the evaluation of the local energy of molecular systems described by atom-centered basis functions. Our approach exploits the structure of the wave function representation and evaluates the wavefunction using sparse linear-algebra methodologies.

The product of basis functions and molecular orbital coefficients $p_{\mu i}$, truncated at $\varepsilon_c = 10^{-12}$ seems to be a good choice for retaining the high accuracy of the wave function while obtaining a significant speedup in the computation for the systems studied.

For increasingly large molecules, the number of operations needed to evaluate the wave function and the local energy at each point of the simulation increases. In the past few years, major advances have been achieved in this laboratory that have made possible the treatment of systems considerably larger than those that could be addressed few years earlier. Development of more efficient codes and, in particular, codes that scale linearly with system size, will have significant impact on the kind of systems that will become feasible to study with QMC.

The ability to complement the information given by the energy with electron distribution and electron localization functions obtained from a high accuracy method such as QMC should prove useful in providing further insight on chemical properties and processes.
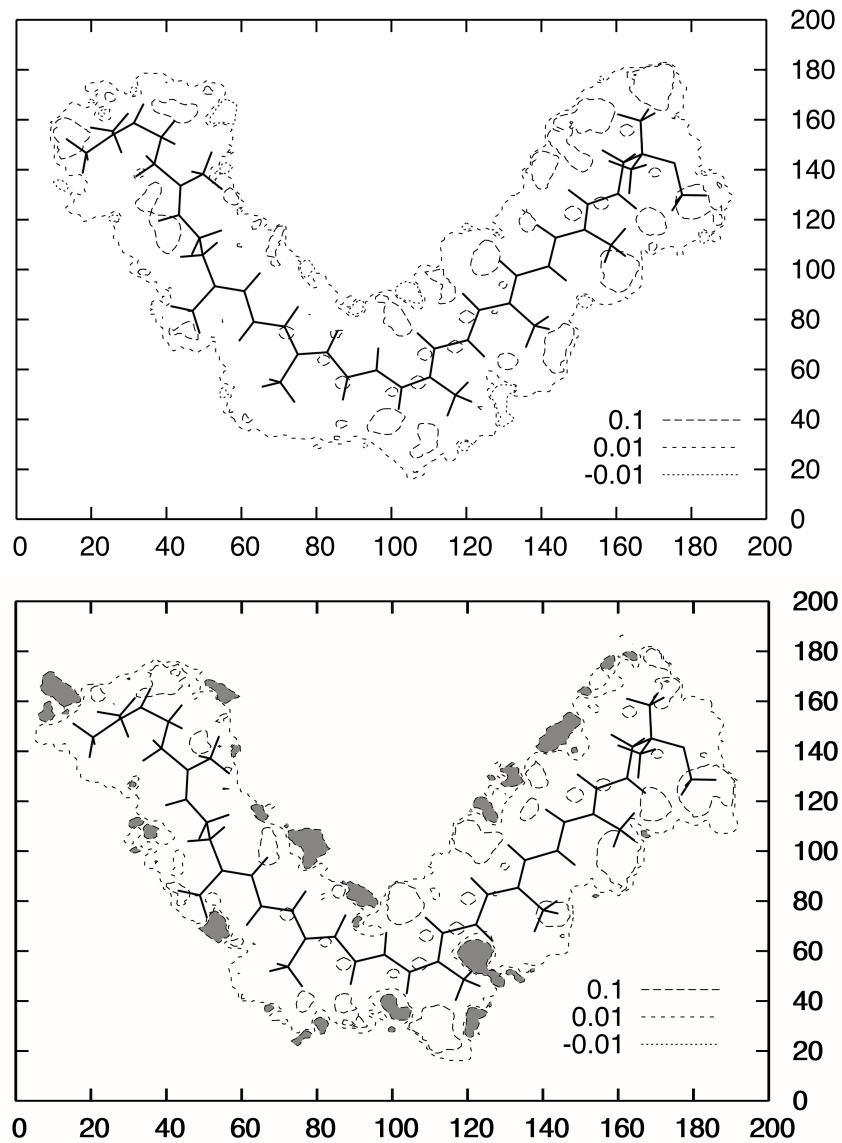
Figure 6: Plot of the $z = 0$ plane of the EPLF function values for Spo in the ground state singlet (top), and triplet (bottom). The triplet state shows spin-polarized regions (dark grey) that are absent in the singlet state. The EPLF domain for both plots is [-0.1:0.1]

References

[1] B. L. HAMMOND, W. A. LESTER, JR., and P. J. REYNOLDS, *Monte Carlo Methods in Ab Initio Quantum Chemistry*, World Scientific, Singapore, 1994.

[2] D. M. CEPERLEY and L. MITAS, Quantum Monte Carlo Methods in Chemistry, in *New Methods in Computational Quantum Mechanics*, edited by I. PRIGOGINE and S. A. RICE, volume XCIII of *Adv. Chem. Phys.*, John Wiley and Sons, 1996.

[3] J. B. ANDERSON, Quantum Monte Carlo: Atoms, Molecules, Clusters, Liquids and Solids, in *Reviews in Computational Chemistry*, edited by K. B. LIPKOWITZ and D. B. BOYD, volume 13, p. 133, John Wiley and Sons, New York, 1999.

[4] W. A. LESTER, JR. and B. L. HAMMOND, *Annu. Rev. Phys. Chem.* **41**, 283 (1990).

[5] L. MITAS, *Comp. Phys. Comm.* **97**, 107 (1996).

[6] M. FOULKES, L. MITAS, R. NEEDS, and G. RAJAGOPAL, *Rev. Mod. Phys.* **73**, 33 (2001).

[7] A. ASPURU-GUZIK and W. A. LESTER, JR., Quantum Monte Carlo methods for the Solution of the Schrödinger equation for molecular systems, in *Computational Chemistry*, edited by C. LE BRIS, volume X of *Handbook of Numerical Analysis*, p. 485, Elsevier, Amsterdam, The Netherlands, first edition, 2003.

[8] P. J. REYNOLDS, D. M. CEPERLEY, B. ALDER, and W. A. LESTER, JR., *J. Chem. Phys.* **77**, 5593 (1982).

[9] P. MASLEN, C. OCHSENFELD, C. WHITE, M. LEE, and M. HEAD-GORDON, *J. Phys. Chem. A* **102**, 2215 (1998).

[10] A. WILLIAMSON, R. HOOD, and J. GROSSMAN, *Phys. Rev. Lett.* **8724** (2001).

[11] S. MANTEN and A. LÜCHOW, *J. Chem. Phys.* **119**, 1307 (2003).

[12] D. ALFE and M. GILLAN, *J. Phys. Cond. Mat.* **16**, L305 (2004).

[13] A. J. WILLIAMSON, J. C. GROSSMAN, R. Q. HOOD, A. PUZDER, and G. GALLI, *Phys. Rev. Lett.* **89** (2002).

[14] A. PUZDER, A. J. WILLIAMSON, J. C. GROSSMAN, and G. GALLI, *Phys. Rev. Lett.* **88** (2002).

[15] A. PUZDER, A. J. WILLIAMSON, J. C. GROSSMAN, and G. GALLI, *J. Chem. Phys.* **117**, 6721 (2002).

[16] S. MANTEN and A. LÜCHOW, Improved Scaling in Diffusion Quantum Monte Carlo with Localized Molecular Orbitals, in *Recent Advances in Quantum Monte Carlo Methods, Part II*, edited by S. M. ROTHSTEIN, W. A. LESTER, JR., and S. TANAKA, Singapore, 2002, World Scientific.

[17] A. ASPURU-GUZIK, R. SALOMÓN-FERRER, B. AUSTIN, and W. A. LESTER, JR., *J. Comp. Chem.* **26**, 708 (2005).

[18] D. CEPERLEY, G. V. CHESTER, and M. H. KALOS, *Phys. Rev. B.* **16**, 3081 (1977).

[19] R. C. WHALEY, A. PETITET, and J. J. DONGARRA, *Parallel Computing* **27**, 3 (2001), Also available as University of Tennessee LAPACK Working Note #147, UT-CS-00-448, 2000 (www.netlib.org/lapack/lawns/lawn147.ps).

[20] S. LIU, J. PEREZ-JORDA, and W. YANG, *J. Chem. Phys.* **112**, 1634 (2000).

[21] C. J. UMRIGAR, M. P. NIGHTINGALE, and K. J. RUNGE, *J. Chem. Phys.* **99**, 2865 (1993).

[22] G. VELDE, F. BICKELHAUPT, E. BAERENDS, C. GUERRA, S. V. GISBERGEN, J. SNIJDERS, and T. ZIEGLER, *J. Comp. Chem.* **22**, 931 (2001).

[23] K. E. SCHMIDT and J. W. MOSKOWITZ, *J. Chem. Phys.* **93** (1990).

[24] S. F. BOYS and N. C. HANDY, *Proc. R. Soc. London Ser. A* **310**, 63 (1969).

[25] Z. SUN, R. K. OWEN, P. J. REYNOLDS, and W. A. LESTER, JR., *Theo. Chim. Acta.* **75**, 353 (1989).

[26] C. LAWSON, R. HANSON, D. KINCAID, and F. KROGH., *ACM Trans. Math. Software* **5**, 308 (1979).

[27] A. SCEMAMA, P. CHAQUIN, and M. CAFFAREL, *J. Chem. Phys* **121**, 1725 (2004).

[28] A. D. BECKE and K. E. EDGECOMBE, *J. Chem. Phys.* **92**, 5397 (1990).