

NWChem: New Functionality

Theresa L. Windus, Eric J. Bylaska, Michel Dupuis, So Hirata, Lisa Pollack,
Dayle M. Smith, T.P. Straatsma, and Edoardo Aprà

Pacific Northwest National Laboratory,
906 Battelle Blvd.,
Richland, WA 99352

Abstract. NWChem is a computational chemistry code designed for efficient execution on massively parallel computers. Through effective use of the aggregate resources of such computers, NWChem enables much larger and more accurate production simulations than previously feasible. In this paper we briefly describe the NWChem architecture and highlight some of the new capabilities in NWChem.

1 Introduction

Parallel computing has been a part of computational chemistry for over 20 years [1,2] and most major chemistry packages support some degree of parallelism. However, very few computational chemistry packages scale well on large numbers of processors. The massively parallel computational package NWChem [3] is one such package and it was designed and implemented to scale to large numbers of processors.

Only a very brief overview of the capabilities of NWChem is provided here. For a more complete overview of NWChem, the reader is referred to previous articles [3] and to the NWChem web page [4]. NWChem provides many methods to compute the properties of molecular and periodic systems using standard quantum mechanical descriptions of the electronic wavefunction or density. These include Gaussian-orbital based methods such as self-consistent field (SCF) molecular orbital (Hartree-Fock, HF) theory, density functional theory (DFT), second order Møller-Plesset perturbation theory (MP2), coupled cluster (CC) theory, and multi-configuration SCF (MCSCF). They also include plane-wave based DFT methods used in Car-Parrinello simulations. In addition to these quantum chemical approaches, NWChem can also perform classical molecular dynamics and free energy simulations using classical force fields. These methodologies may be combined to perform quantum molecular dynamics and mixed quantum mechanics/molecular mechanics simulations.

The underlying support for the parallel capabilities in the functionality outlined above and in the following sections is provided by the Global Arrays toolkit (GA) [5, 6]. These tools include the Memory Allocator (MA) to access local memory, the Global Arrays (GA) to provide portable shared-memory programming tools, the Aggregate Remote Memory Copy Interface (ARMCI) [7] to provide general-purpose, portable and efficient remote memory copy operations (one-sided communication) optimized for noncontiguous (strided, scatter/gather, I/O vector) data transfers, and the Parallel I/O (ParIO) [8] tool to extend the non-uniform memory architecture model to disk. Libraries from the Global Arrays toolkit interoperate with MPI [9] for message passing.

In this paper, we will briefly discuss key new functionality in the NWChem. The following sections discuss aspects of orbital-dependent density functionals, of the projected augmented wave module, the electron transfer module, single excitation theories, the tensor contraction engine work, and the molecular dynamics module and its application to simulations of biological membranes. We conclude with a brief summary.

2 Orbital-Dependent Density Functionals

Density functional theory is an important tool in computational chemistry, due to its predictive capabilities. The various approximate functionals available for the exchange and correlation energy can be viewed in the context of the Jacob's ladder [10] idea proposed by Perdew and coworkers. Essentially, the ladder has various rungs, each representing a subset of density functionals, with improved chemical accuracy obtained by climbing to higher rungs. The two lowest rungs of the ladder are the LDA (Local Density Approximation) and the GGA (Generalized Gradient Approximation).

There is a general agreement that development on the LDA and GGA levels is basically complete, and that further improvement in chemical accuracy will be reached with higher rungs of the ladder, namely rungs that incorporate the kinetic energy density and/or exact (Hartree-Fock) exchange [10,11]. Some of the most popular functionals in quantum chemistry are the hybrid functionals, which mix in a fraction of exact exchange to GGA exchange. The success of the hybrid functionals indeed suggests that exact exchange contains essential information, and that inclusion of exact exchange could help in the solution of major challenges like self interaction error, time dependent and excited state phenomena, and in the improvement of the accuracy of the exchange-correlation potential [10].

Use of the kinetic energy density and/or exact exchange, means the functionals depend explicitly on the Kohn-Sham orbitals. However, calculating the exchange-correlation potential for an orbital dependent functional cannot be done in the usual manner, by taking the functional derivative of the exchange-correlation energy with respect to the density, although various methods exist to address this problem. Below we discuss which methods are in place, preparing NWChem for the upcoming generation of orbital functionals.

The rung above GGAs on the ladder of approximations is the meta-GGA, which uses the kinetic-energy density (or sometimes the Laplacian of the density) as an additional functional ingredient. The inclusion of the kinetic energy density means the meta-GGA is necessarily an orbital dependent functional and as such must be appropriately handled in the code. There are various ways to implement a meta-GGA.

First, fully self-consistent orbitals and densities obtained from GGA calculations can be used as input to calculate meta-GGA energies. This capability was recently added to the code, as experience suggests that energies obtained this way are similar to fully self-consistent meta-GGA calculations.

Second, using a Gaussian basis, an orbital dependent exchange-correlation potential can be found by energy minimization with the Kohn-Sham matrix elements. This method [12] is already implemented in NWChem for exact exchange, and the incorporation of the kinetic energy density term should introduce little numerical instability.

Third, an exchange-correlation potential can be obtained from the optimized effective potential (OEP) integral equation, the preferred method of calculating self-consistent meta-GGA energies. Solution of the OEP integral equation is very difficult, and at present has very few implementations for molecules[13]. Recently, a new approach[14] makes the OEP much more practical to implement; a method has been developed to iteratively solve a system of partial differential equations coupled to the Kohn-Sham equations. This promises a more computationally efficient OEP opens up the possibility of fully self-consistent solutions using orbital dependent density functionals in NWChem.

3 Electron Transfer Module

The newly implemented electron transfer (ET) module of NWChem calculates the electronic coupling energy (also called the electron transfer matrix element) between ET reactant and product states. The electronic coupling (V_{AB}), activation energy (ΔG^*), and nuclear reorganization energy (λ) are all components of the electron transfer rate defined in Marcus' theory [15].

$$k_{et} = \frac{2\pi}{\hbar} V_{AB}^2 \frac{1}{\sqrt{4\pi\lambda k_B T}} \exp\left[-\frac{\Delta G^*}{k_B T}\right] \quad (1)$$

These energies are illustrated below in Figure 1, which shows the energy profile for a simple two-state ET system. Curve A shows how the reactant state's potential energy changes along the reaction coordinate, and curve B describes the product state.




Fig. 1. Energy profile for a simple two-state ET system.

If the system has adequate energy to overcome the activation energy (ΔG^*), and if the electron transfer coupling energy (V_{AB}) is sufficiently large, the system can proceed to ET products, thereby completing the reaction. For organometallic systems, ET is said to be "adiabatic" (100% crossover), when V_{AB} is greater than about 200 cm^{-1} .

The ET module of NWChem utilizes the method of Corresponding Orbital Transformation [16] to calculate V_{AB} from the MO vectors of the reactant and product states, and uses new and existing NWChem routines and GA tools. For instance, the two-electron

contribution to V_{AB} , which is by far the most time consuming, is calculated using the Fock matrix routine from the SCF module. The amount of CPU time required in the calculation of V_{AB} is the same as a single Hartree-Fock iteration, so the most expensive part of a V_{AB} calculation is converging the wavefunctions of the reactant and product states.

The ET module has been applied to a variety of systems. In a recent publication [17] we calculated V_{AB} for a hematite tetramer, which yielded a hematite electron mobility in accord with the experimentally measured value. Similar tests were performed with an aqueous iron dimer, quinone dimers, and several diatomic molecules. Currently the NWChem ET module is being applied to study the electron pathway through flavocytochrome c_3 fumarate reductase (I_{fc}₃), which participates in the iron respiration of bioremediation bacteria.

4 Single Excitation Theories

NWChem now includes several single excitation theories for vertical excitation energy calculations, such as configuration interaction singles (CIS) [18], time-dependent Hartree-Fock (TDHF, also known as the random-phase approximation), time-dependent density functional theory (TDDFT) [19,20], and Tamm-Dancoff TDDFT [21] within a single algorithmic framework based on Davidson's iterative subspace method [22] or its extension to a non-Hermitian matrix eigenvalue problem [23]. These implementations permit the selective calculations of singlet or triplet excited states (for closed-shell systems) and excited states of a specified spatial symmetry symbol. They also compute oscillator strengths, transition moments, and S^2 spin expectation values of each excited state to assist in band assignments. Although not very efficient, geometry optimizations and frequency calculations can be performed on an excited state potential energy surface by virtue of NWChem's universal finite-difference capability.

The CIS and TDHF excitation energies of closed-shell systems are comparable to each other in accuracy, and are normally considered a zeroth-order description of the excitation process. There are indications that these methods are well balanced in describing Rydberg excited states, in contrast to TDDFT. However, for open-shell systems, the errors in the CIS and TDHF excitation energies are often excessive, primarily due to the multi-determinantal character of the ground and/or excited state wave functions of open-shell systems in a HF reference [24]. The scaling of the computational cost of a CIS or TDHF calculation per state with respect to the system size is the same as that for a HF calculation for the ground state, since the critical step of both methods is the Fock build, namely, the contraction of two-electron integrals with density matrices. In each cycle of Davidson's iterative solution of an eigenvalue problem, we usually need to process more than one density matrix; each density matrix corresponds to a trial vector of Davidson's algorithm. In the implementation of CIS and TDHF within NWChem, we process as many density matrices as fit in the available memory simultaneously; so, the number of two-electron integral evaluations, which are the most expensive in the Fock build, can be minimized. This strategy (multiple Fock build) is standardly adopted in HF and DFT Hessian calculations and the subroutine for this task can be reused in CIS and TDHF with no modification [25].

The accuracy of TDDFT varies with the exchange-correlation functional. In general, the exchange-correlation functionals that are widely used today and are implemented in NWChem work well for low-lying valence excited states. However, for high-lying diffuse excited states and Rydberg excited states in particular, TDDFT employing these conventional functionals breaks down and the excitation energies are substantially underestimated. This is because of the fact that the exchange-correlation potentials generated from these functionals decay too rapidly (exponentially) as opposed to the slow $-1/r$ asymptotic decay of the true potential [26] Unlike CIS or TDHF, TDDFT provides a balanced description of radical excited states [27], and recently NWChem's TDDFT module was used to study 51 species of the radical ions of polycyclic aromatic hydrocarbons with a mean absolute error of ca. 0.3 eV in excitation energies [28]. The computational cost per state of TDDFT calculations scales as the ground state DFT calculations, although the prefactor of the scaling may be much greater in the former. Thus, TDDFT is applicable to fairly large molecules. In addition to yielding excitation energies in general good accord with experiment, it offers the advantage of the orbital picture familiar to chemists when discussing spectral properties.

5 Projected Augmented Wave

In conjunction with the group of Professor John H. Weare (UCSD), we have recently implemented into NWChem a massively parallel version of Blöchl's Projected Augmented Wave (PAW) method [29]. The PAW method is an all-electron density functional theory method that combines the adaptability of linear-augmented-plane-wave methods with the computational advantages of Car-Parrinello methods. As illustrated in Figure 2, with this method the all-electron Kohn-Sham orbital are defined in terms of a smooth wave function (expanded in plane-waves) and a deficit normalization part (expanded using a radial grid about each atom). The computational advantage is that all variational calculations are performed on smooth wave functions, while retaining all-electron information with the deficit normalization part (or augmentation charges).




Fig. 2. Using the PAW transformation, the rapid variations around the nuclei are projected onto a local basis set.

The PAW method has become the approach of choice for calculating systems that include transition metals, compared to methods that use pseudopotentials with plane-wave basis sets. Indeed transition metals require extremely "hard" pseudopotentials which necessitate the use of very large and/or complicated basis sets. We note again here the fundamental advantage of plane-wave approaches is the ease of calculating first derivatives compared to approaches using Gaussian atom-centered basis set. With plane-waves, the Pulay first derivative terms (except for a small one-electron part from augmentation charges) do not exist, and their calculations are the most computational demanding step in force calculations. As a result, plane-wave approaches are easily combined with molecular dynamics (time-dependent) calculations.

We have performed simulations on a variety of molecular and solid-state systems using the parallel PAW code [30,31]. Results obtained for solids are found to be of good quality and simulations for a series of first-row transition metal dimers, monoxides and dioxides show good agreement with Gaussian DFT simulations and with experiment. These small transition metal-oxide molecules, which are well characterized experimentally, are a challenging test of the reliability of an ab initio method because they have a large number of low lying states, many with high spin multiplicity. These and other results have shown that

- The accuracy of the PAW is similar to the density functional calculations based on local basis sets.
- The convergence with respect to the plane wave basis set leads to practical calculations, even for very difficult systems (F, O, Fe, Cr, etc).
- The method is stable with respect to the choice of the local atomic basis set.
- For a given plane wave basis set size the execution times of PAW calculations are very similar to those of pseudopotential plane-wave methods.
- The size of the plane wave basis is smaller compared to norm-conserving pseudopotential plane-wave methods.

6 Tensor Contraction Engine

With the advent of massively parallel supercomputers, highly accurate many-electron theory calculations are becoming feasible or routine for a wider range of chemical systems. However, such calculations are often hindered not just by the limitation in computer resources but by the complexity of deriving the necessary equations and implementing them into efficient computer programs that support both sequential and parallel executions. Although complex, these symbolic manipulation processes bear much in common across different models of many-electron theory. Specifically, the vast majority of them are defined in the language of second-quantized creation and annihilation operators and the working equations invariably take the form of tensor algebraic expressions that are implemented into a sequence of matrix multiplications [32]. In this light, we have developed a symbolic manipulation program (Tensor Contraction Engine or TCE) automating the entire processes of formula derivation and parallel implementation of various models of many-electron theory [33]. TCE not only expedites the derivation and implementation processes, but it also facilitates parallelization and other laborious optimizations,

enhances portability, maintainability, and extensibility of synthesized programs, and helps design and test new models of many-electron theory.

Provided a definition of many-electron theory as quasi-vacuum expectation values of normal-ordered second-quantized operators, TCE (an interpreted, interactive, object-oriented program written in Python programming language) [33] performs valid contractions of creation and annihilation operators according to Wick's theorem, consolidates identical terms, and reduces the expressions into the form of multiple tensor contractions acted upon by index permutation operators. In various stages of these processes, TCE canonicalizes [34] the expressions of tensors, index permutation operators, etc., to perform rapid pattern matching operations. TCE subsequently performs strength reduction (determines the binary contraction orders with minimal operation and memory costs), factorization (eliminates common binary contractions and defines the intermediate tensors), and identifies reusable intermediates.

The resulting ordered list of binary tensor contractions, additions, and index permutations is translated into an optimized program that is in turn interfaced to NWChem. The programs automatically synthesized by TCE take advantage of spin symmetry, spatial (Abelian point-group) symmetry, and index permutation symmetry simultaneously to minimize the operation cost and storage requirement, adjust the peak local memory usage by index range tiling, and support multiple parallel I/O interfaces and dynamic load balancing for parallel executions. The index range tiling algorithm achieves the best compromise between the conflicting demand for the minimization of the number of arithmetic operations and the minimization of memory requirement. Consequently, the programs synthesized by TCE can compete with carefully hand-coded programs in operation and memory cost.

As an initial application, we invoked TCE to automatically derive the formulae of various models of spin-unrestricted configuration interaction theory (CISD, CISDT, CISDTQ), spin-unrestricted coupled-cluster theory (CCD, LCCD, CCSD, LCCSD, QCISD, CCSDT, CCSDTQ), and spin-unrestricted non-canonical many-body perturbation theory [MBPT(2), MBPT(3), MBPT(4)] and implement them into optimized parallel programs within NWChem. For many of these high-end correlation treatments, this constitutes the first reported automatic or parallel implementation, which implies that TCE can potentially change the way second-quantized many-electron theories are studied. Let us consider the example of CCSD, which has undergone perhaps the most thorough algorithmic refinement over the years. According to Stanton and coworkers [35], the number of arithmetic operations of their unrestricted CCSD program (the implementation in ACES II) is approximately $5/4 O^2V^4 + 20 O^3V^3 + 5/2 O^4V^2$ in the leading order, where O and V are the number of α -spin occupied and virtual orbitals, respectively. The operation count of the TCE-synthesized CCSD program is approximately $5/4 O^2V^4 + 45/2 O^3V^3 + 25/2 O^4V^2$. Remembering $V \gg O$ in most applications and hence $O^2V^4 \gg O^3V^3 \gg O^4V^2$, the TCE-generated CCSD program can therefore compete with, albeit not outperform, the equivalent, carefully hand-coded program. A list of enhancements that are being considered for TCE's program generator includes storage space minimization via the loop fusion and space-time tradeoff techniques [36], and the development of a scalable parallel I/O interface for a global file system. TCE also offers an immediate avenue to automatically derive and implement existing and new models

of many-electron theory, which may include various relativistic many-electron theories, equation-of-motion coupled-cluster theory for excited states, and combined perturbation and coupled-cluster theories. With the adequate interface programs already in place, we believe that these and other models of many-electron theory can be most conveniently studied with NWChem and TCE.

7 Molecular Dynamics Simulations of Biological Membranes

The molecular dynamics (MD) simulation of biological membranes is generally more difficult to setup compared to simulations of other biological systems such as proteins, protein complexes and DNA that typically start from available three-dimensional structures from protein X-ray crystallography or NMR solution structure determination, and that typically consider only a single molecule or molecule complex in the model system. Biological membranes, however, derive their structural and functional properties from the fact that they are made up from a specific arrangement of many molecules for which little or no experimental data is available. Molecular modeling of such systems, therefore, relies on initial molecular assemblies that need to be constructed from primary molecular structures and application of computational relaxation methods.

The MD simulation preparation module of NWChem has a number of features to facilitate the construction of bio-molecular membranes. The typical sequence of steps consists of aligning of the lipid tails of the molecule that will form the membrane, placing copies of these aligned molecules on a two-dimensional grid perpendicular to the alignment axis, randomization of the orientation along the molecular alignment axis, 'collapsing' of the membrane by moving the molecules in the membrane plane to within a small, predefined distance from each other, and placement of counter-ions for lipid head groups that contain charged groups, or for lipopolysaccharides with charged saccharide moieties. Although several iterations may be required to tweak many of the parameters, this process can be executed in an automated way in NWChem.

Not only the setup, but also the equilibration of a membrane system is more complicated than many other bio-molecular systems. In order to obtain a molecular structure with the proper membrane integrity it is necessary to apply a number of restraints, especially during the equilibration phase of the simulation. NWChem allows for simulations with many restraining potentials that can be gradually turned on or off, for example to keep lipid tails of individual molecules perpendicular to the plane of the membrane, and to keep all head groups in the membrane plane. Since these membranes typically involve many identical molecules, specifying the details of these restraining potentials would be very tedious and prone to errors without the tools implemented in NWChem that automatically recognize identical molecules. These tools can be used to duplicate a single specified restraint potential or apply a single restraint potential to all other identical molecules in the system.

The equilibration is first performed on the constructed membrane in vacuo with restraining potentials applied to obtain and maintain membrane integrity. Once the molecules that form the model membrane are sufficiently dense, the system can be solvated. Special tools are available in NWChem to remove solvent molecules that are placed or diffuse into selected regions such as between the lipid tails. An important fea-

ture required for these simulations is that anisotropic pressure can be applied, since the compressibility of the membrane will be different in the plane compared to perpendicular to the plane of the membrane.

The membrane setup procedures can include the incorporation of other molecules, such as trans-membrane proteins. This procedure has been applied to an simulation of a lipopolysaccharide microbial membrane consisting of the lipopolysaccharide layer, a lipid layer, and an embedded trans-membrane protein [37].

8 Conclusions

In this paper, we have highlighted several of the key new functionalities in the computational chemistry code NWChem. These new functionalities range from tools to improve simulations of biological membranes to new orbital-dependent density functionals to the new projected augmented wave, electron transfer, and excitation modules to the tensor contraction engine for generating highly correlated methods. Future work in NWChem includes improvements to these capabilities as well as new functionality such as proton hopping in molecular dynamics simulations, an extended properties analysis module, excited state gradients and additional linear scaling methods. The authors encourage other developers to add functionality to NWChem. Further information concerning this may be obtained by sending an e-mail to nwchem-support@emsl.pnl.gov.

Acknowledgements. This research was performed in part using the Molecular Science Computing Facility (MSCF) in the William R. Wiley Environmental Laboratory at the Pacific Northwest National Laboratory (PNNL). The MSCF is funded by the Office of Biological and Environmental Research in the U.S. Department of Energy. The TCE work is funded by the U.S. Department of Energy, the Division of Basic Energy Science, Office of Science, as part of the SciDAC program. PNNL is operated by Battelle for the U.S. Department of Energy under contract DE-AC06-76RLO 1830.

References

1. Ostlund, N.S.: *Int. J. Quant. Chem. Symp.* **13** (1979) 15
2. Harrison, R.J., Shepard, R.: *Annu. Rev. Phys. Chem.* **45** (1994) 623
3. Kendall, R.A., Aprà, E., Bernholdt, D., Bylaska, E.J., Dupuis, M., Fann, G.I., Harrison, R.J., Ju, J., Nichols, J.A., Nieplocha, J., Straatsma, T.P., Windus, T.L., Wong, A.T.: *Computer Phys. Comm.* **128** (2000) 260.
Bernholdt, D.E., Aprà, E., Fruchtl, H.A., Guest, M.F., Harrison, R.J., Kendall, R.A., Kutteh, R.A., Long, X., Nicholas, J.B., Nichols, J.A., Taylor, H.L., Wong, A.T., Littlefield, R.J., Nieplocha, J.: *Int. J. Quantum Chem. Symposium* **29** (1995) 475
4. <http://www.emsl.pnl.gov/pub/docs/nwchem>
5. <http://www.emsl.pnl.gov:2080/docs/parsoft/>
6. Nieplocha, J., Harrison, R.J., Littlefield, R.J.: *The Journal of Supercomputing*, **10** (1996) 197

7. Nieplocha, J., Carpenter, C.: ARMCI: A Portable Remote Memory Copy Library for Distributed Array Libraries and Compiler Run-time Systems. Proc. 3rd Workshop on Runtime Systems for Parallel Programming (RTSPP) of International Parallel Processing Symposium IPPS/SPDP '99, Lecture Notes in Computer Science, Vol. 1586. Springer-Verlag, Berlin Heidelberg New York (1999)
8. Nieplocha, J., Foster, I., Kendall, R.: ChemIO: High-performance parallel I/O for computational chemistry applications, The International Journal of Supercomputer Applications and High Performance Computing, **12** (1998) 260–266
9. Message Passing Interface Forum. MPI: A Message-Passing Interface standard (version 1.1). Technical report, 1995. <http://www.mpi-forum.org>.
10. Perdew, J., Schmidt, K., Ed. Van Doren, V.E., Van Alsenog, K., Geerlings, P.: Density Functional Theory and its applications to Materials (AIP, Melville NY) (2001)
11. Boese, A. D., Handy, N.: J. Chem. Phys. **116** (2002) 9559
12. Neumann, R., Nobes, R., Handy, N.: Mol. Phys. **87** (1996) 1
13. Yang, W., Wu, Q.: Phys. Rev. Lett. **89** (2002) 143002
14. Kümmel, K., Perdew, J.: Phys. Rev. Lett. **90** (2003) 43004
15. Marcus, R. A., Sutin, N.: Biochimica Biophysica Acta **35** (1985) 437
16. Farazdel, A., Dupuis, M., Clementi, E., Aviram, A.: J. Am. Chem. Soc. **112** (1990) 4206
17. Rosso, K. M., Smith, D. M. A., Dupuis, M.: J. Chem. Phys. (accepted for publication)
18. Foresman, J.B., Head-Gordon, M., Pople, J.A., Frisch, M.J.: J. Phys. Chem. **96** (1992) 135
19. Jamorski, C., Casida, M.E., Salahub, D.R.: J. Chem. Phys. **104** (1996) 5134
20. Bauernschmitt, R., Ahlrichs, R.: Chem. Phys. Lett. **256** (1996) 454
21. Hirata, S., Head-Gordon, M.: Chem. Phys. Lett. **314** (1999) 291
22. Davidson, E.R.: J. Comput. Phys. **17** (1975) 87
23. Olsen, J., Jensen, H.J.Aa., Jørgensen, P.: J. Comput. Phys. **74** (1988) 265
24. Maurice, D., Head-Gordon, M.: J. Phys. Chem. **100** (1996) 6131
25. Windus, T., et al. (unpublished)
26. Casida, M.E., Jamorski, C., Casida, K.C., Salahub, D.R.: J. Chem. Phys. **108** (1998) 4439
27. Hirata, S., Head-Gordon, M.: Chem. Phys. Lett. **302** (1999) 375
28. Hirata, S., Head-Gordon, M., Szczepanski, J., Vala, M. (unpublished)
29. Blöchl, P.E.: Physical Review B, **50** (1994) 17953
30. Valiev, M., Bylaska, E.J., Gramada, A., Weare, J.H.: Ab Initio Molecular Dynamics Simulations using Density-Functional Theory, Reviews In Modern Quantum Chemistry: A Celebration Of The Contributions Of R. G. Parr, Ed. K.D. Sen (World Scientific, Singapore) (2002) 1684
31. Bylaska, E.J., Valiev, M., Kawai, R., Weare, J.H.: Computer Physics Communications, **143**, (2002) 11
32. Head-Gordon, M., Maslen, P.E., White, C.A.: J. Chem. Phys. **108** (1998) 616
33. Hirata, S. (unpublished)
34. Janssen, C.L., Schaefer III, H.F.: Theor. Chim. Acta **79** (1991) 1
35. Stanton, J.F., Gauss, J., Watts, J.D., Bartlett, R.J.: J. Chem. Phys. **94** (1991) 4334
36. Baumgartner, G., Bernholdt, D.E., Cociorva, D., Harrison, R., Hirata, S., Lam, C., Nooijen, M., Pitzer, R., Ramanujam, J., Sadayappan, P.: Proceedings of Supercomputing 2002 (2002)
37. Shroll, R.M., Straatsma, T.P.: Biopolymers **65** (2002) 395