Theor Chem Acc (1997) 97:158–163

Theoretical
Chemistry Accounts

The impact of the resolution of the identity approximate integral method on modern ab initio algorithm development

Rick A. Kendall, Herbert A. Früchtl*

High Performance Computational Chemistry Group Theory Modeling and Simulation Program, Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory, P.O. Box 999, MSIN K1-90, Richland, WA 99352-0999, USA

Received: 14 January 1997 / Accepted: 11 March 1997

Abstract. The computation of the two-electron four-center integrals over gaussian basis functions is a significant component of the overall work of many ab initio methods used today. Improvements in the computational efficiency of the base algorithms have provided significant impact. Somewhat overlooked are methods that provide approximations to these integrals and their implementation in application software. A partial review of approximate integral techniques focused on the resolution of the identity (RI) four-center, two-electron integral approximation is given. The past and current uses of the RI algorithms are presented along with possibilities for further exploitation of the technology.

Key words: Resolution – Identity – Ab initio
Algorithm – Integral

1 Introduction

The ultimate goal (and thus rarely achieved) of computational chemistry is to produce chemically accurate molecular properties of interest to experimentalists and engineers at a low computational cost and in a relatively expeditious time frame. The reality of the situation is that a series of compromises, which optimize the cost accuracy turnaround functional for the given problem at hand, must be made. This series of compromises often will focus the computational chemist on problems that can be solved on the available computational resources, in terms of both hardware and software. There is a complex set of choices that includes basis sets, theoretical methodologies, the appropriate implementation of said methodologies, and computational resources. Therefore, even marginal increases in computing "power" can expand the scope of science being addressed in a particular research group [1, 2].

It is the synergistic role of the advances in computational hardware and infrastructure software and the advances in chemistry algorithms and implementations that was a driving force behind many of Professor Jan Almlöf's research interests. It is clear that using yesterday's (even 10-year-old) software on the currently available parallel or vector supercomputers would yield far less science than those chemistry applications that have either adapted to or been written for the changes in computer technology. The pioneering work on direct algorithms done in the Almlöf group was crucial to many of the computational chemistry community's first attempts to use massively parallel computing technology [3–8], especially since many of these first parallel supercomputers did not have balanced input/output (I/O) subsystems.

The computation of the two-electron four-center integrals over gaussian basis functions is a significant component of the overall cost of many ab initio algorithms. Improvements in the computational efficiency of the base integral evaluation algorithms have provided significant impact [9–16]. Other research efforts have focused on methods that provide approximations to these integrals and their use in application software. One method of approximation is to use the resolution of the identity (RI) in the calculation of the traditional four-center two-electron integrals. It is important to note here that this is one of many algorithms that approximate molecular integrals. All semi-empirical methods by definition use approximate integral technology (c.f., Refs. [17–21] and references therein). In addition, following the initial work of Friesner and coworkers [22, 23], recent developments of Martinez and Carter [24] review the use of pseudospectral methods that show great promise and possibly wider applicability than the RI technology presented here. There are and have been attempts to circumvent the gaussian functions completely using fully numerical

---

* *Present address*: Fujitsu European Centre for Information Technology, 2 Longwalk Road, Uxbridge,
Middlesex, UB11 1AB, UK
e-mail: fruechtl@fecit.co.uk

Correspondence to: R.A. Kendall
e-mail: ra-kendall@pnl.gov

schemes [25–31] and recently developed wavelet techniques [32–34]. The former requires a complete engineering of the theoretical methods/algorithms from the ground up and the computational costs are still high. The latter is still very new and exciting but will require the same software engineering of needed theoretical models. Again, the use of a particular method is a conscious choice of the tradeoffs between accuracy, the cost of the computation, and the cost of the implementation of the methodology.

Modern use of the RI method in ab initio electronic structure theory began with the work of Feyereisen et al. [35], who proposed this approximation for second order Möller-Plesset perturbation theory (MP2) calculations as a means of circumventing the sort and transformation step. Almlöf and coworkers [36] then provided an elegant extension to SCF calculations and a formal presentation of multiple RI algorithms. The goal of this paper is to give an overview of the RI technology as used in ab initio electronic structure theory. Derivative works based on these efforts are reviewed as well as the areas that need additional development.

2 Background of RI methods

The modern use of the RI methodology as applied to ab initio methods is not new. The basic approach to all these methods is the factorization of the four-center integral into two parts:

$$(ij|kl) = \sum_{\gamma}^{N} L_{ij\gamma} R_{\gamma kl}. \qquad (1)$$

The details vary in each application but the ultimate goal is qualitatively the same, e.g., the reduction in computational cost. As discussed below (c.f., Sect. 3) and elsewhere [24], this factorization is beneficial for any Coulomb interaction summation. The cost is now $O(n^3)$, as opposed to the $O(n^4)$ for traditional integral evaluation. Unfortunately, this factorization does not formally reduce the scaling of an exchange interaction summation. The time to compute the quantities used in the Coulomb or exchange interaction summations will vary but the cost of the summation remains $O(n^4)$ for the exchange term.

The basic idea of approximating four-center integrals by two- and three-center quantities calculated from auxiliary functions was first introduced by Boys and Shavitt [37]. Van Alsenoy [38] gives a review of these methods up to about a decade ago. This or a similar technique is also used in the scattering theory to simplify the computation of the requisite integrals to compute the bound-state-free-electron Coulomb matrix elements [39]. Other methods that reduce the overall work for two-electron integrals include the work of Beebe and Linderberg [40], who used a full Cholesky decomposition of the two-electron integral super-matrix. The details of the Cholesky decomposition show that the inner projection used by Beebe and Linderberg is similar to that of the work by Vahtras et al. [36]. O'Neal and Simons [41] extended this to include integral derivative methods. They showed the accuracy, cutoff-based savings for the computational costs and requisite disk space minimization for both integral and gradient based algorithms. This allows the application user to tune the application based on the desired accuracy and cost.

Outside those efforts reviewed by Van Alsenoy [38], the first use of an RI method was most likely that of Baerends et al. [42], who used an RI algorithm to fit the density matrix. The second implementation and the most widely known instance was the utilization of the Dunlap fit in the development of density functional theory (DFT). Dunlap and coworkers [43] fit the Coulomb potential using essentially an identical algorithm as implemented, later, by Vahtras et al. (the detailed Almlöf work outlined in Sect. 3). The wide applicability and use of the DFT methodologies alone easily demonstrate the usefulness of the RI technology.

3 The RI two-electron integral approximation

In the RI approximation the reduction (c.f. Eq. 1) is formally introduced by inserting a resolution of the identity into the two-electron integrals:

$$(ij|kl) \approx \sum_{t} (ijt)(t|kl). \qquad (2)$$

Unless the auxiliary basis $|t)$ spans the whole space of products $|ij)$, this expansion introduces an error that has to be minimized. Vahtras et al. [36] showed that by inserting the resolution of the identity more than once and minimizing different properties of the residual function

$$R_{ij}(r) = |ij) - \sum_{t} c_t|t), \qquad (3)$$

three different three-center approximations for the four-center two-electron integrals result:

$$(ij|kl) = \sum_{tuvw} (ijt)\ S^{-1}_{tu} V_{uv} S^{-1}_{vw} (klw) \qquad (4)$$

$$(ij|kl) = \sum_{tu} (ijt)\ S^{-1}_{tu} (kl|u) \qquad (5)$$

$$(ij|kl) = \sum_{tu} (ij|t) V^{-1}_{tu} (kl|u). \qquad (6)$$

Depending on the term between the three-center integrals, they are named SVS-, S-, and V-approximation, which correspond to Eqs. (4), (5), and (6), respectively. The two-center quantities are the overlap ($S_{tu} = (tu)$) and electron repulsion ($V_{tu} = (t|u)$) integrals using the auxiliary functions or the RI fitting basis. Preliminary comparisons in Ref. [36] show that self-consistent field (SCF) energies and energy differences obtained with the V-approximation are on average more accurate by two orders of magnitude than the others. Früchtl et al. [44] showed results where this approximation and a moderate-sized expansion basis (4–6 times the size of the atomic orbital basis) yields good results.

This is not surprising, as the V-approximation is used extensively in DFT theory to approximate the Coulomb contribution. Dunlap et al. [43] showed that for this purpose it is variational and an optimized expansion basis would give the best possible three-center approximation. This in principle also holds for the exchange contribution, but as the sign of Coulomb and exchange part are different, the sum is not variational. In addition, the optimal basis for the two contributions is different, and using the same basis takes away the "best approximation" argument. Nevertheless it gives further indication that the V-approximation is the more natural one to use.

4 Impact on algorithms

Besides the "traditional" use in DFT, the RI approximation has been implemented for SCF [36, 44], multi-configuration self-consistent field (MCSCF) [45, 46], MP2 [35], [47], and coupled cluster singles and double (CCSD) with and without perturbative triples [48].

The usefulness of the RI-SCF method, as first implemented by Vahtras et al. [36] and recently developed for parallel computers by Früchtl et al. [44], is restricted to a limited range of chemical systems. As noted by Martinez and Carter [24], an exchange-type sum does not reduce the $O(n^4)$ scaling of the calculation, but this does not necessarily mean that a reduction of work cannot be accomplished. In a disk-based implementation [44], the storage as well as the I/O requirement scale only with the third power of the number of basis functions, $O(n^3)$. Also, in a fully direct implementation the integral calculation, normally the most time-consuming step, would not only take advantage of the cheaper three-center integrals, but also only scale with $O(n^3)$. The $O(n^4)$ part consists only of matrix multiplications, which are in most cases computationally much less expensive than integral calculations. In addition, linear algebra operations are a very general problem, so one can expect that efficient implementations are available for most computer architectures.

In the implementation by Früchtl et al., the matrix V^{-1} in Eq. (6) is not stored, but contracted into the three-center integrals. Following Rendell and Lee [48], the integrals are transformed with $V^{-1/2}$

$$(ij|kl) = \sum_{stu} (ij|s)\ V_{st}^{-1/2} V_{tu}^{-1/2}\ (u|kl) \tag{7}$$

$$= \sum_{v} (ij|v)\ (v|kl), \tag{8}$$

and the transformed integrals $(ij|v)$ are used in the requisite summations.

This simplifies the calculation of the Coulomb and exchange contributions to the Fock matrix, but cannot be used for a straightforward implementation of a direct scheme for the calculation. We therefore store the transformed three-center integrals for repeated use. The implementation within NWChem [49, 50] chooses between a distributed in-core method, where the integrals are stored in a global array [51] (possibly distributed over the distributed memory of a parallel computer), and a disk-based version using the Disk Resident Array library [52] if the memory is not sufficient.

The exchange contribution

$$K_{ij} = \sum_{vkl} (ik|v)(v|jl)D_{kl}, \tag{9}$$

which is the least scalable part of the computation of the Fock matrix, can be calculated via two matrix multiplications and a sum over the results:

$$T_{jk}^{v} = \sum_{l} (v|jl)D_{kl} \tag{10}$$

$$W_{ij}^{v} = \sum_{l} (ik|v)T_{jk}^{v} \tag{11}$$

$$K_{ij} = \sum_{v} W_{ij}^{v}. \tag{12}$$

In the case of restricted closed shell SCF calculations and if the molecular orbital (MO) vectors are available, the number of operations can be further reduced by inserting the definition of the density matrix and using the MO vectors instead. In the second order SCF procedure as implemented in NWChem [49, 53, 54] the MO vectors are available during the energy and gradient calculations, but not during the line search algorithm, which in typical calculations is about 50% of the instances when a Fock matrix build is needed. In a DIIS-based Restricted Hartree Fock (RHF) or SCF procedure these savings in computation time could be used for every Fock build.

With C_{ij} being the MO vectors, the expression for the exchange contribution is then

$$K_{ij} = \sum_{vkl} \sum_{\sigma} (ik|v)(v|jl)C_{k\sigma}C_{l\sigma}, \tag{13}$$

and instead of Eqs. (10) and (11) we can use

$$X_{i\sigma}^{v} = \sum_{l} (il|v)C_{l\sigma} \tag{14}$$

$$W_{ij}^{v} = \sum_{\sigma} X_{i\sigma}^{v} X_{j\sigma}^{v}. \tag{15}$$

The sum over σ involves only the occupied orbitals, resulting in a reduction in floating point operations by a factor of n_{occ}/n.

The benchmark calculation by Früchtl et al., [44] on a Cray T3D show a speedup of 13 or 20 over the NWChem direct SCF (depending on required accuracy and the size of the RI expansion basis) [49] for benzonitrile with Dunning's cc-pVTZ basis set [55] and an even-tempered, uncontracted RI expansion basis. Sparsity can currently only be used for calculating, storing, and retrieving the three-center integrals. The matrix multiplications are carried out over full matrices, so that integral sparsity does not reduce the overall scaling of $N \cdot n^3$ for the density-based method[1] and $N \cdot n^2 \cdot n_{\mathrm{occ}}$ using the MO vectors. The strength of this method is thus restricted to

---

[1] N corresponds to the number of functions in the RI fitting basis and n is the size of the atomic basis. Since the size of the fitting basis is a multiple of the atomic basis, $O(N \cdot n^3)$ is also $O(n^4)$. See the discussion below about expansion basis sets, Sect. 5

spatially compact molecules, where other methods also cannot make much use of integral sparsity.

The use of RI integrals for second order MP2 was first implemented by Feyereisen et al. [35] and later developed for parallel computers by Bernholdt and Harrison [47]. While the first implementation used the SVS approximation, Bernholdt implemented all three methods described above (Eqs. 4–6) and concluded that the V-approximation gives the best results. His results show that, especially if energy differences and not absolute energies are of interest, an expansion basis of moderate size gives nearly exact results in about 5% of the computation time. The savings are, as in the case of SCF, due to transforming the part of the computation with the highest scaling to simple matrix algebra. Here the $O(n^5)$ four-index transformation is replaced by a three-index transformation scaling with $O(n^4)$, while the exact energy evaluation now becomes $O(n^5)$, but can be carried out via matrix multiplications.

Ten-no and Iwata [45, 46] have used the RI approximation for their implementation of a multi-configuration SCF procedure. Their linear combination of atomic-electron distributions MCSCF (LCAD-MCSCF) yields energies of extremely high precision with significantly shorter computation time (e.g., the energy of the $^1A'$ state of HNO with cc-pVTZ basis is correct to 5×10^{-5} kcal/mol and takes only 25% of the computation time). The computational efficiency as well as accuracy improve with larger basis sets.

Rendell and Lee [48] used the RI approximation for the calculation of all integrals involving one or zero occupied orbitals [$(ab|cd)$ and $(ab|ci)$ in their notation] for their implementation of CCSD and CCSD(T). While adding some computational cost at the $O(n^5)$ level, it greatly reduces the amount of disk storage required, thus pushing the limit on the size of chemical systems for which such calculations can be carried out. In their benchmark calculations of energy differences, the error introduced by the RI approximation was always smaller than the error resulting from incompleteness of the atomic orbital (AO) basis, which they take as a sufficient criterion.

5 RI expansion basis sets

An area that will definitely need some more exploration is the choice of expansion basis sets. The most natural choice (and the easiest to implement, as most existing integral packages can only handle one basis set at a time) would be to use the atomic orbitals as an expansion basis as well. While this gives reasonable results for coupled cluster methods [48], it is much too inaccurate for most other methods reviewed here.

The other extreme would be to take the full gaussian product basis. The LCAD approach [46] starts from there, but takes into account only products of functions that are located on the same atom. It then removes product functions that lead to linear dependencies.

Most other investigations used basis sets that are somewhere between these two extremes in terms of size and sophistication. Feyereisen et al. [35] used standard AO basis sets of higher quality as expansion basis. Bernholdt and Harrison [47] uncontracted their AO basis and added some functions from the product basis set.

Vahtras et al. [36] took the boundaries for exponent and angular momentum from the product basis and constructed an uncontracted and even-tempered basis set that fulfilled these conditions. Früchtl et al. [44] used a more heuristic scheme based on similar considerations.

The expansion or fitting basis sets [56] developed for the DGauss and DeMon density functional programs that are widely used for DFT calculations are based on the same theme. A mostly (all except the highest exponent) even-tempered set, starting with double the lowest exponent of the atomic orbital basis set, is constructed with as many exponents as possible shared between the different angular momenta to simplify integral calculation. Although the small number of functions, and especially the lack of functions with high angular momentum (even for second-row atoms and transition metals they only go up to d functions), seems rather limiting, the large number of successful applications seems to indicate that chemically reasonable results can be obtained with them.

One significant systematic optimization of a basis set for use as an RI expansion basis is the work of Eichkorn et al. [57] on an expansion basis for use in DFT calculations. They require the exponents of their expansion basis to fulfill a dependency similar to the definition of an even-tempered basis

$$\eta_{i+1} = \eta_i \beta \left(1 + \gamma \frac{i^2}{(n+1)^2} \right) \tag{16}$$

and optimize $\eta_0, \beta,$ and γ to maximize the Coulomb energy, using its variational property. The number of functions and highest angular momentum are chosen to fulfill an accuracy criterion for the charge density with a minimal number of expansion functions.

A more thorough investigation of this problem is clearly necessary. It may turn out that basis sets have to have different properties for different methods, depending on if Coulomb or exchange interaction contributions are more important, or if the core or the valence electrons have to be described more accurately.

6 Future work and conclusions

It is clear that the RI approximation as used in computational chemistry algorithms, especially in DFT, has been and will be widely used in the future. The expansion of the high performance computing market into parallel supercomputers of O(100) nodes has led to continued development of algorithms that take advantage of these unique hardware and software capabilities. The RI algorithm is a natural method for expanding the scope of applications via specific algorithms that involve Coulomb interaction summations; fortunately, there are many such summations in computational chemistry (e.g., also in analytic gradient algorithms). It is also clear from Sect. 5 above that further exploration and optimization of expansion basis sets is warranted by the community.

This is not only for expansion basis sets for specific atoms with a given atomic basis set but methods need to be developed that will provide an appropriate recipe to generate a basis set for a given methodology as well. A DFT expansion or fitting basis set is likely to be different in nature that one for CCSD.

RI technology, although the focus of this short review is not necessarily the definitive algorithm of choice; it is one of many useful tools. One lesson learned from the migration to parallel supercomputers is that multiple algorithms must be implemented to afford scalability with both machine size and chemical system size. Like many software applications and algorithms, the user must be aware of the inherent nature and applicability of the RI algorithms being used.

Acknowledgements. We would like to thank Jan Almlöf for championing the efforts of cooperation among the computational chemistry community and computational science communities. This work was performed under the auspices of the Environmental Molecular Sciences Laboratory Construction Project of the Office of Health and Environmental Research, U.S. Department of Energy and the High Performance Computing and Communications Program of the Office of Scientific Computing, U.S. Department of Energy, under contract DE-AC06-76RLO 1830, with the Battelle Memorial Institute, which operates the Pacific Northwest National Laboratory. The authors would like to acknowledge the assistance and support of the High Performance Computational Chemistry Group in the development of some of the work described here. H.A.F. acknowledges the support of Associated Western Universities, Inc., Northwest Division under contract DE-FG06-89ER-75522. We would also like to thank Dr. Mary Brewster for providing preprints of the wavelet research prior to publication and Dr. Peter Taylor for a review of an early version of the manuscript.

References

1. Feller DF (1993) The MSRC ab initio methods benchmark suite: a measurement of hardware and software performance in the area of electronic structure methods. Technical Report PNL-8670. This technical report is available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161
2. Feller DF, Kendall RA, Brightman MJ (1995) The EMSL ab initio methods benchmark report: a measure of hardware and software performance in the area of electronic structure methods, version 2. PNL Technical Report PNL-10481. This technical report is available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161
3. Guest MF, Harrison RJ, van Lenthe JH, van Corler LCH (1987) Theor Chim Acta 71: 117
4. Guest MF, van Lenthe JH, Sherwood P (1990) Supercomputer 36: 8
5. Harrison RJ, Kendall RA (1991) Theor Chim Acta 79: 337
6. Luethi HP, Mertz JE, Feyereisen MW, Almlöf JE (1992) J Comput Chem 13: 160
7. Colvin ME, Janssen CL, Whiteside RA, Tong CH (1993) Theor Chim Acta 84: 301
8. Feyereisen MW, Kendall RA (1993) Theor Chim Acta 84: 289
9. Helgaker T, Taylor PR (1995) In: Yarkony R (ed) Modern electronic structure theory. World Scientific Publishing Co. River Edge, New Jersey, p725. This includes a very good review of integral technology of interest to all who develop integral codes
10. Obara S, Saika A (1986) J Chem Phys 84: 3963
11. Head-Gordon M, Pople JA (1988) J Chem Phys 89: 5777
12. Gill PMW, Head-Gordon M, Pople JA (1990) J Phys Chem 94: 5564
13. Gill P, Pople JA (1991) Int J Quantum Chem 40: 753
14. Johnson BG, Gill PMW, Pople JA (1993) Chem Phys Lett 206: 229
15. Hamilton TP, Schaefer HF (1991) Chem Phys 150: 163
16. Lindh R, Ryu U, Liu B (1991) J Chem Phys 95: 5889
17. Pople JA, Beveridge DL (1970) Approximate molecular orbital theory. McGraw-Hill, New York
18. Dewar MJS, Thiel W (1977) Theo Chim Acta 46: 89
19. Dewar MJS, Thiel W (1977) J Am Chem Soc 99: 4899
20. Stewart JJP (1990) In: Lipkowitz KB, Boyd DB (eds) Reviews in computational chemistry. VCH, New York, pp 45–81
21. Zerner MC (1991) In: Lipkowitz KB, Boyd DB (eds) Reviews in computational chemistry, vol II. VCH, New York, pp 313–365
22. Friesner RA, (1987) J Chem Phys 86: 3522
23. Ringnalda M, Belhadj M, Friesner RA (1990) J Chem Phys 93: 3397
24. Martinez TJ, Carter EA (1995) In: Yarkony DR (ed) Modern electronic structure theory. World Scientific Publishing Co., River Edge, New Jersey, p1132
25. Fischer CF (1977) The Hartree-Fock method for atoms. Wiley, New York
26. Adamowicz L, McCullough EA Jr. (1984) J Phys Chem 86: 2178
27. Adamowicz L, McCullough EA Jr. (1984) Int J Quantum Chem 26: 373
28. Adamowicz L, Barlett RJ, McCullough EA Jr. (1985) Phys Rev Lett 54: 426
29. McCullough EA Jr., Morrison J, Richman K (1984) Faraday Symp Chem Soc 19: 165
30. McCullough EA Jr. (1986) Comput Phys Rep 4: 265
31. Fischer CF, Chen D (1989) THEOCHEM 58: 61
32. Fischer P, Defranceschi M (1994) Appl Comput Harmonic Anal 1: 232
33. Brewster ME, Fann GI (1996) In: Unser MA, Aldroubi A, Laine AF (eds) Proceedings of SPIE Annual Meeting, Denver, Colorado, August 1996. The International Society for Optical Engineering, Bellingham, Wash 2825: 919
34. Brewster ME, Fann GI, Zhiyun Yang, Private Communication to R.A.K., preprint "Wavelets for electronic structure calculations" to be submitted to J Math Chem
35. Feyereisen M, Fitzgerald G, Komornicki A (1993) Chem Phys Lett 208: Nr. 5, 6, 359
36. Vahtras O, Almlöf JE, Feyereisen MW (1993) Chem Phys Lett 213: Nr. 5, 6, 514
37. Boys S F, Shavitt I (1959) University of Wisconsin, Rep WIS-AF-13
38. Van Alsenoy C (1988) J Comp Chem 9: 620–626
39. McCurdy CW, Rescigno TN (1992) Phys Rev A 46: 255
40. Beebe NHF, Linderberg J, (1977) Int J Quant Chem, 12: 683
41. O'Neal DW, Simons J, (1989) Int J Quant Chem 36: 673
42. Baerends EJ, Ellis DE, Ros R, (1973) Chem Phys 2: 41
43. Dunlap BI, Connolly JWD, Sabin JR, (1979) J Chem Phys 71: 3369
44. Früchtl HA, Kendall RA, Harrison RJ, Dyall KG (1997) Int J Quant Chem (in press)
45. Ten-no S, Iwata S, (1996) J Chem Phys 105: 3604
46. Ten-no S, Iwata S, (1996) Int J Quant Chem Symp 30: 107
47. Bernholdt DE, Harrison RJ (1996) Chem Phys Lett 250: 477
48. Rendell AP, Lee TJ (1994) J Chem Phys 101: 400
49. High Performance Computational Chemistry Group (1995) NWChem, a computational chemistry package for parallel computers, version 2.0. Pacific Northwest National Laboratory, Richland, WA 99352, USA. See URL: "http://www.emsl.pnl.-gov: 2080/docs/nwchem/nwchem.html" on the EMSL WWW server
50. Bernholdt DE, Apra E, Früchtl HA, Guest MF, Harrison RJ, Kendall RA, Kutteh RA, Long X, Nicholas JB, Nichols JA, Taylor HL, Fann GI, Littlefield RJ, Nieplocha J (1995) Int J Quant Chem Symp 29: 475

51. Nieplocha J, Harrison RJ, Littlefield RJ (1995) Supercomputing '95 Proceedings, December 1995, San Diego, Calif., IEEE Computer Society Press, Los Alamitos, Calif., p 340
52. Nieplocha J, Foster I (1996) Proceedings, Frontiers '96. The Sixth Symposium on The Frontiers of Massively Parallel Computation, October 1996, Annapolis, MD. IEEE Computer Society Press, Los Alamitos, Calif., p 196
53. Wong AT and Harrison RJ, (1995) J Comp Chem 16: 1291
54. Shepard R, (1993) Theor Chim Acta 84: 343
55. Dunning TH Jr. (1989) J Chem Phys 90: 1007
56. Godbout N, Salahub DR, Andzelm J, Wimmer E (1992) Can J Chem 70: 560
57. Eichkorn K, Treutler O, Öhm H, Häser M, Ahlrichs R (1995) Chem Phys Lett 240: 283