

# Iterative stochastic subspace self-consistent field method

Pierre-François Loos<sup>1,2</sup>  · Jean-Louis Rivail<sup>3</sup> · Xavier Assfeld<sup>3</sup>

Received: 20 January 2017 / Accepted: 18 April 2017  
© Springer-Verlag Berlin Heidelberg 2017

**Abstract** We propose a new self-consistent field (SCF) algorithm based on an iterative, partially stochastic “*Divide & Conquer*”-type approach. This new SCF algorithm is a simple variant of the usual SCF procedure and can be easily implemented in parallel. A detailed description of the algorithm is reported. We illustrate this new method on one-dimensional hydrogen chains and three-dimensional hydrogen clusters.

**Keywords** Self-consistent field · Stochastic process · Divide & conquer · Parallel computing

---

**Electronic supplementary material** The online version of this article (doi: 10.1007/s00894-017-3347-3) contains supplementary material, which is available to authorized users.

---

✉ Pierre-François Loos  
loos@irsamc.ups-tlse.fr

Jean-Louis Rivail  
jean-louis.rivail@univ-lorraine.fr

Xavier Assfeld  
xavier.assfeld@univ-lorraine.fr

<sup>1</sup> Laboratoire de Chimie et Physique Quantiques, Université de Toulouse, Toulouse, CNRS, UPS, France

<sup>2</sup> Research School of Chemistry, Australian National University, Canberra, ACT 2601, Australia


<sup>3</sup> Théorie-Modélisation-Simulation, Université de Lorraine, CNRS SRSMC UMR 7565, Vandœuvre-lès-Nancy 54506, France

## Introduction

Moore’s law [1] predicts that the computer speed approximately doubles every two years. Therefore, we could expect that the system sizes accessible by conventional quantum chemistry *ab initio* methods will keep growing rapidly in the next few decades. However, since several years, Moore’s law is only fulfilled by an increasing number of cores per processor as CPU clock rate has stopped increasing. Indeed, today hardware is designed in a multi-core manner and one needs to write softwares which support multi-threading to take full advantage of the hardware. Consequently, it becomes more and more important to develop parallel computer codes and algorithms that are capable for exploiting high-performance parallel computing environments [2–8].

The computational bottleneck in Hartree-Fock (HF) [9] and density-functional theory (DFT) [10] methods, in their conventional formulations, is the computation of the two-electron integrals which scales quartically with the number of basis functions  $N$ . Fortunately, this formal  $O(N^4)$  scaling does reduce asymptotically to  $O(N^2)$  if one takes into account only the numerically significant integrals [11–15] by using standard Cauchy-Schwarz screening [16] or distance-dependent screening techniques [17].

The two main steps in a self-consistent field (SCF) calculation are the formation of the Fock or Kohn-Sham (KS) matrix [18] and the subsequent determination of the molecular orbital (MO) coefficients. The work required for the formation of the Fock or KS matrix can be done very efficiently by exploiting the local nature of chemistry. For example, the formation of the Coulomb part of the Fock or KS matrix scales linearly with the system size using



**Fig. 1** Schematic representation of the stochastic partition with  $N = 8$  and  $K = 2$

the continuous fast-multipole method developed by White and coworkers [19, 20]. Linear scaling can also be reached for the computation of the exact exchange matrix [21] and exchange-correlation potentials (see Refs. [14, 22], and references therein).

In the conventional formulation of HF and DFT methods, the MO coefficients are obtained by diagonalization of the Fock or KS matrix, respectively. This step scales cubically with the system size [9, 10], and may therefore become the time dominating step for large molecules.

The present method aims at reducing the  $O(N^3)$  cost of the matrix diagonalization by partitioning the MOs of the system into subsets and performing smaller diagonalizations in these subsets. The main advantage of the method is that each smaller diagonalization can be performed on a distinct core. The method is a MO-based variant of the SCF algorithm based on an iterative, partially stochastic “*Divide & Conquer*” strategy [23, 24]. This new algorithm that we have named *iterative stochastic subspace* SCF (I3SCF) is a

simple modification of the usual SCF procedure and can be easily implemented in parallel. Although major modifications and refinements have been made, the present method is inspired by the *ab initio* LSCF method [25–28] used in QM/MM methods [29–32].

We would like to mention here that the present investigation is a “*proof of principle*” study. Obviously, a more detailed study would be necessary to fully understand the advantages and limitations of the present algorithm, as well as calculations on larger systems. We will report on this in the future.

The present article is organized as follows: in “**Theory**”, we give a detailed derivation of each step of the I3SCF algorithm, while illustrative examples are given in “**Applications**”. Atomic units are used throughout.


## Theory

### Stochastic partition

The purpose of the stochastic process is to partition the MOs  $\{|\psi_p\rangle\}_{1 \leq p \leq N}$  into  $K$  subsets of size  $L$  at each SCF cycle. For sake of simplicity, we assume that the  $K$  subsets have the same dimension  $L$ . However, the present process can be easily generalized for subsets of different sizes (see below). I3SCF calculations using  $K$  subsets are labelled I3SCF[ $K$ ]. The atomic orbital (AO) coefficients of the  $k$ th subset are gathered in the matrix  $\mathbf{C}_k$  of dimension  $N \times L$ . For  $K = 1$  (i.e.  $L = N$ ), the present algorithm is equivalent to the conventional SCF procedure. First, to illustrate the stochastic partition procedure, we set  $N = 2^m$ , where  $m \in \mathbb{N}^*$ . The general case is discussed below.

| Step   | I3SCF algorithm   | Step | SCF algorithm   |
|--------|---|------|---|
| 1.     | Obtain a set of $N$ orthogonalized MOs  | 1.   | Diagonalize $\mathbf{S}$ and compute $\mathbf{X} = \mathbf{S}^{-1/2}$                                 |
| 2.     | For each I3SCF iteration,   | 2.   | For each SCF iteration,   |
| 2.1    | Partition of the MOs into $K$ subsets $\mathbf{C}_k$ using stochastic process                               |      |   |
| 2.2    | For each subset $k$ of size $L$ ,   |      |   |
| 2.2.1. | Calculate $\mathbf{F}'_k = \mathbf{B}_k^\dagger \cdot \mathbf{F} \cdot \mathbf{B}_k$                        | 2.1  | Calculate $\mathbf{F}' = \mathbf{X}^\dagger \cdot \mathbf{F} \cdot \mathbf{X}$                        |
| 2.2.2. | Diagonalize $\mathbf{F}'_k$ to get $\mathbf{C}'_k$ and $\epsilon_k$   | 2.2. | Diagonalize $\mathbf{F}'$ to get $\mathbf{C}'$ and $\epsilon$   |
| 2.2.3. | Back-transform the eigenvectors to the original basis set $\mathbf{C}_k = \mathbf{B}_k \cdot \mathbf{C}'_k$ | 2.3. | Back-transform the eigenvectors to the original basis set $\mathbf{C} = \mathbf{X} \cdot \mathbf{C}'$ |
| 2.3    | Build the matrices $\mathbf{C}$ and $\epsilon$ and populate the orbitals                                    | 2.4. | Populate the orbitals   |
| 2.4.   | Exit test. If not satisfied, go back to 2. for a new I3SCF iteration  | 2.5. | Exit test. If not satisfied, go back to 2. for a new SCF iteration                                    |
| 3.     | End of I3SCF calculation  | 3.   | End of SCF calculation  |

**Fig. 2** Main steps in the I3SCF and SCF algorithms



**Fig. 3** Calculation time (in second) for performing step 2.2.2. of the I3SCF[K] algorithm as a function of N for various K values. The calculation time of step 2.2. in the conventional SCF algorithm is also reported for comparison

The first and crudest level of partition (level 1 in Fig. 1) consists in creating occupied-virtual MO pairs by randomly choosing an occupied MO  $\psi_i$  and associating it with the virtual MO  $\psi_a$  which has the largest coefficient

$$f_{ia} = \|\mathbf{c}_i^\dagger \cdot \mathbf{F} \cdot \mathbf{c}_a\|_{\mathcal{F}}, \tag{1}$$


where  $\mathbf{F}$  is the Fock matrix, the vectors  $\mathbf{c}_i$  and  $\mathbf{c}_a$  contain the AO coefficients of  $\psi_i$  and  $\psi_a$ , and

$$\|\mathbf{Z}\|_{\mathcal{F}} = \sqrt{\text{Tr}(\mathbf{Z}^\dagger \cdot \mathbf{Z})} \tag{2}$$

is the Frobenius norm.

This maximizes the probability of decreasing the energy by mixing the occupied-virtual MO pairs with large

**Fig. 4** Structure of the three-dimensional hydrogen clusters: H<sub>16</sub> (left), H<sub>32</sub> (center) and H<sub>64</sub> (right)



off-diagonal Fock elements. We anticipate that these “level 1” calculations will exhibit the slowest convergence rate due to the restricted number of MOs per subset ( $L = 2$ ). At level 1, the stochastic nature of the algorithm is due to the random choice of the occupied MOs. It is interesting to note that the level 1 partition has some similarities with the well-known Jacobi sweep technique [33].

The next step of the stochastic process (level 2 in Fig. 1) associates two occupied-virtual MO pairs to create larger subsets (if required). This is done by randomly picking a subset  $k$  and associating it with the subset  $k'$  having the largest Frobenius norm [34]

$$\mathbf{f}_{kk'} = \|\mathbf{C}_k^\dagger \cdot \mathbf{F} \cdot \mathbf{C}_{k'}\|_{\mathcal{F}}. \tag{3}$$

At level 2, the stochastic nature of the algorithm is due to the random choice of the subset  $k$ .

This process is repeated until the required number of subsets is created. This stochastic partition is schematically illustrated in Fig. 1. In the case where  $N \neq 2^m$ , the level 1 partition can be modified in order to build the target number of subsets. For example, two virtual MOs can be associated with one occupied MO to create triplets. Then, the level 2 partition associates these triplets to create larger subsets.

### Diagonalization

The next step consists in computing and diagonalizing the transformed Fock matrix


$$\mathbf{F}'_k = \mathbf{B}_k^\dagger \cdot \mathbf{F} \cdot \mathbf{B}_k \tag{4}$$

of size  $L \times L$  to obtain the eigenvectors  $\mathbf{C}'_k$  and eigenvalues  $\epsilon_k$ . We use the orthogonalization matrix [9]

$$\mathbf{B}_k = \mathbf{C}_k, \tag{5}$$

where  $\mathbf{C}_k$  in Eq. 5 refers to the MO coefficients from the previous iteration. The matrix  $\mathbf{B}_k$  fulfills the desired orthogonality condition, i.e.

$$\mathbf{B}_k^\dagger \cdot \mathbf{S} \cdot \mathbf{B}_k = \mathbf{I} \tag{6}$$



**Fig. 5** Energy convergence (in hartree) of the  $H_{16}$  chain for 25 different random seeds using the I3SCF algorithm with  $L = 8$

(where  $\mathbf{I}$  is the identity matrix), which must be satisfied if the transformed orbitals are to form an orthogonal set. Note that the procedure requires an orthogonal set of  $N$  MOs as initial guess.

The eigenvectors are then back-transformed to the original basis to obtain a new set of MO coefficients:

$$\mathbf{C}_k = \mathbf{B}_k \cdot \mathbf{C}'_k. \quad (7)$$

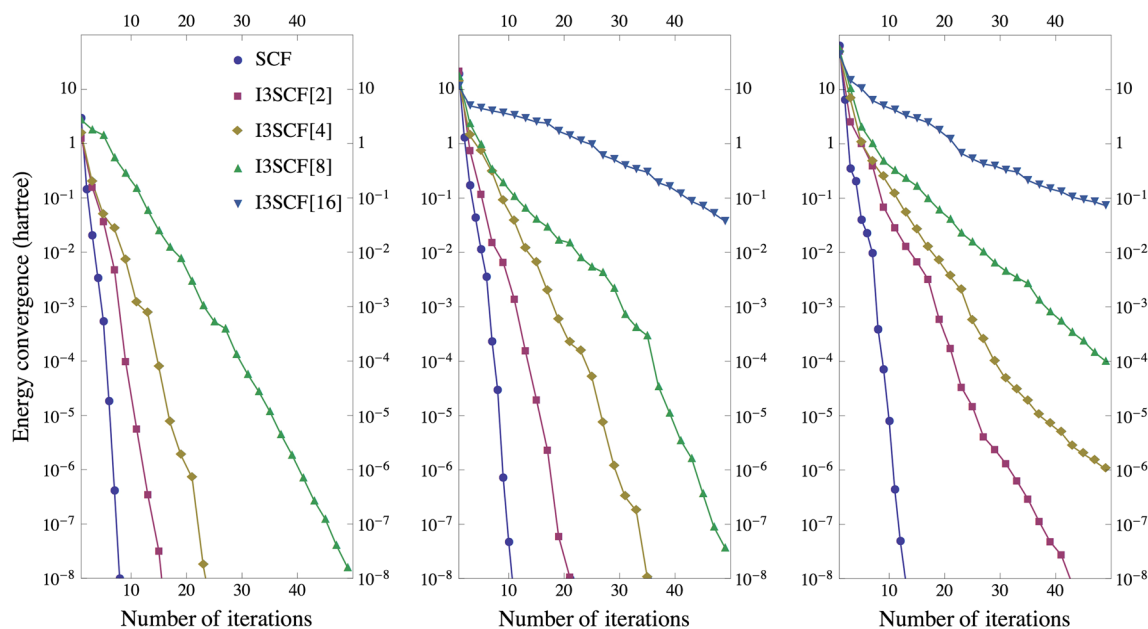
After this diagonalization step has been performed on each subset, we form the matrices

$$\mathbf{C} = \bigcup_{k=1}^K \mathbf{C}_k, \quad \boldsymbol{\epsilon} = \bigcup_{k=1}^K \boldsymbol{\epsilon}_k, \quad (8)$$

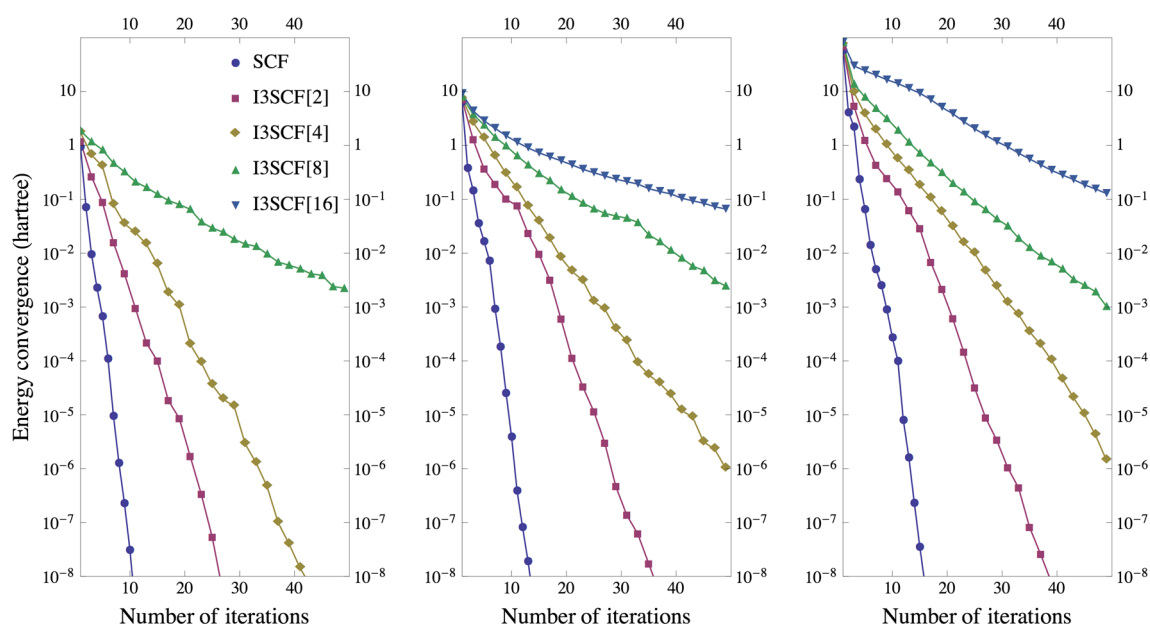
where the union of two matrices  $\mathbf{Z}_1$  and  $\mathbf{Z}_2$  returns a matrix with the rows of  $\mathbf{Z}_1$  followed by the rows of  $\mathbf{Z}_2$ . The orbitals are then populated using the Aufbau principle [9].

The I3SCF algorithm is summarized in Fig. 2. For comparison, we also report the main steps of the usual SCF procedure. Loop 2.2. over the  $K$  subsets can be efficiently parallelized as each subset can be allocated to a single core. The main advantage of the I3SCF algorithm is that the expensive  $O(N^3)$  diagonalization step of the “total” Fock matrix is reduced to smaller diagonalizations. Because each of these diagonalizations only cost  $O(L^3)$  (step 2.2.2.), the total computational cost (i.e. over the  $K$  subsets) is reduced to  $K \times O(L^3) = O(NL^2)$  for the I3SCF algorithm. This is a significant reduction in cost compared to the  $O(N^3)$  diagonalization cost of the usual SCF algorithm.

We illustrate this point in Fig. 3, where we compare the calculation time as a function of  $N$  for step 2.2 in the conventional SCF algorithm and for step 2.2.2 in the I3SCF algorithm. As one can see, for any value of  $K$ , the “Divide & Conquer” diagonalization strategy of the I3SCF algorithm eventually becomes faster than the “full” diagonalization required by the conventional SCF, with a cost growing near-linearly with  $N$ . In comparison, the usual SCF method exhibits a steep increase of the calculation time.



**Fig. 6** Energy convergence (in hartree) of one-dimensional hydrogen chains:  $H_{16}$  (left),  $H_{32}$  (center) and  $H_{64}$  (right)



**Fig. 7** Energy convergence (in hartree) of three-dimensional hydrogen clusters:  $H_{16}$  (left),  $H_{32}$  (center) and  $H_{64}$  (right)

## Applications

### Computational details

To illustrate the advantages and disadvantages of the I3SCF algorithm, we have computed the restricted HF ground-state energies of linear chains of  $n$  equally-spaced hydrogen atoms for  $n = 16, 32$  and  $64$  with the same number of spin-up and spin-down electrons (closed-shell systems). The distance  $R_{\text{HH}}$  between neighboring atoms has been set to 1.8 bohr [35]. These systems have recently attracted considerable interest due to their strong correlation character and metal-insulator transition [35–38].

To test the performance of the I3SCF algorithm for non-symmetric structures, we have also computed the restricted HF ground-state energies of three-dimensional hydrogen clusters for  $n = 16, 32$  and  $64$ . These clusters are obtained by considering the regular  $2 \times 2 \times 4$ ,  $2 \times 4 \times 4$  and  $4 \times 4 \times 4$  lattices with  $R_{\text{HH}} = 1.8$  bohr and applying a random Gaussian displacement of standard deviation  $\sigma = R_{\text{HH}}/4$  in each direction. These are represented in Fig. 4, and the cartesian coordinates of the corresponding structures are reported as [Supplementary Material](#).

The conventional SCF algorithm as well as the I3SCF algorithm have been implemented in MATHEMATICA [39]. The AOs basis consists of a single Gaussian function of exponent 0.4, which has been obtained in order to reproduce the STO-6G basis [40, 41]. All the calculations uses the core Hamiltonian as a guess Fock matrix, and Pulay's DIIS method [42, 43] is applied to accelerate convergence.

### Hydrogen chains and clusters

First, we have tested the dependence of the I3SCF algorithm with respect to the stochastic partitioning by performing I3SCF[2] calculations on the  $H_{16}$  chain using 25 different random seeds. In Fig. 5, we have plotted the energy convergence (in hartree) defined as the difference between the energy at a given iteration and the converged energy. As one can see, the final result is largely independent of the choice of the random seed.

The energy convergence of the  $H_{16}$ ,  $H_{32}$  and  $H_{64}$  chains is represented in Fig. 6. To reach a given energy accuracy, the I3SCF algorithm requires more iterations than the conventional SCF method. For a given value of  $K$ , the number of iterations seems to increase linearly with the system size.<sup>1</sup> The same behavior is observed for the SCF algorithm. This means that the present method does not shift the cost of each diagonalization step to the number of diagonalization steps. Because each diagonalization can be done on a distinct core, a parallel implementation of the I3SCF algorithm could be competitive with the conventional SCF algorithm, especially for large systems. The computation cost of the I3SCF method could be further reduced by employing localized orbitals [44] and sparse algebra routines [45, 46].

<sup>1</sup>In all the cases studied in the work, we have observed a smooth convergence of the SCF energy. Therefore, the curves reported in Figs. 6 and 7 can be easily extrapolated to the desired convergence threshold.

For small  $L$ , the I3SCF algorithm needs a larger number of iterations to achieve the desired convergence threshold, due to the small number of MOs per subset. For example, the I3SCF[8] calculation on  $H_{16}$  has only two MOs per subset. However, for any value of  $K$ , the I3SCF algorithm converges reliably to the SCF limit.

The same conclusions can be drawn for the three-dimensional hydrogen clusters, as shown in Fig. 7. Even for non-symmetric systems in higher dimensions, the I3SCF algorithm converges reliably to the SCF limit. Moreover, we observe that, for a given value of  $K$ , the number of iterations required to reach a given energy accuracy is roughly independent of the system size. Again, the same observation can be made for the SCF algorithm. Also, we note that, the convergence is getting slower when the number of subsets  $K$  increases. However, this could be fixed by using localized orbitals that will impose a spatial constraint in the stochastic partition in order to avoid having weakly interacting MOs in the same subset. We will investigate this possibility in a forthcoming paper. In some cases the overall computation gain can be rather limited (as typically in I3SCF[2] calculations). It is therefore important to find the best compromise between the size of the Fock matrices to diagonalize and the level of stochastic partition (see Fig. 1).

## Conclusions

In this article, we have described and studied a new SCF algorithm that we have dubbed *iterative stochastic subspace SCF* (I3SCF). This new method, which is a simple variant of the usual SCF algorithm, is based on a “*Divide & Conquer*” strategy which partitions the MOs of the system into subsets. It can be parallelized efficiently on modern parallel computers. The I3SCF algorithm has been tested on one-dimensional and three-dimensional hydrogen systems, for which it has shown promising performances. We hope to report results for larger molecules of biological interest in the near future.

**Acknowledgments** PFL thanks the NCI National Facility for a generous grant of supercomputer time, and the Australian Research Council for funding (Grant DP140104071) and a Discovery Early Career Researcher Award (Grant DE130101441). XA would like to thank the Research School of Chemistry of the Australian National University for a visiting fellowship during the construction of this manuscript.

## References

- Moore GE (1965) Electron Mag 38
- Nielsen IM, Janssen CL (2008) Sci Prog 16:277
- de Jong WA, Bylaska E, Govind N, Janssen CL, Kowalski K, Muller T, Nielsen IMB, van Dam HJJ, Veryazov V, Lindh R (2010) Phys Chem Chem Phys 12:6896
- Yasuda K (2008) J Comput Chem 29:334
- Yasuda K (2008) J Chem Theory Comput 4:1230
- Ufimtsev IS, Martinez TJ (2008) Comput Sci Eng 10:26
- Ufimtsev IS, Martinez TJ (2009) J Chem Theory Comput 5:1004
- Ufimtsev IS, Martinez TJ (2009) J Chem Theory Comput 5:2619
- Szabo A, Ostlund NS (1989) Modern quantum chemistry. McGraw-Hill, New York
- Parr RG, Yang W (1989) Density-functional theory of atoms and molecules. Clarendon Press, Oxford
- Gill PMW, Johnson BG, Pople JA (1994) Chem Phys Lett 217:65
- Gill PMW (1994) Adv Quantum Chem 25:141
- Dyczmans V (1973) Theor Chim Acta 28:307
- Kussmann J, Beer M, Ochsenfeld C (2013) WIREs Comput Mol Sci 3:614
- Barca GMJ, Loos PF, Gill PMW (2016) J Chem Theory Comput 12:1735
- Haser M, Ahlrichs R (1989) J Comput Chem 10:104
- Maurer SA, Lambrecht DS, Flaig D, Ochsenfeld C (2012) J Chem Phys 136:144107
- Challacombe M, Schwegler E (1997) J Chem Phys 106:5526
- White CA, Johnson BG, Gill PMW, Head-Gordon M (1994) Chem Phys Lett 230:8
- White CA, Head-Gordon M (1994) J Chem Phys 101:6593
- Schwegler E, Challacombe M (1996) J Chem Phys 105:2726
- Goedecker S (1999) Rev Mod Phys 71:1085
- Yang W (1991) Phys Rev Lett 66:1438
- Yang W, Lee T-S (1995) J Chem Phys 103:5674
- Assfeld X, Rivail JL (1996) Chem Phys Lett 263:100
- Fornili A, Loos P.-F., Assfeld X, Sironi M (2006) Chem Phys Lett 427:236
- Loos PF, Assfeld X (2007) J Chem Theory Comput 3:1047
- Monari A, Rivail J.-L., Assfeld X (2013) Acc Chem Res 46:596
- Moreau Y, Loos PF, Assfeld X (2004) Theor Chem Acc 112:228
- Ferre N, Assfeld X, Rivail J-L (2002) J Comput Chem 23:610
- Loos P-F, Preat J, Laurent A, Michaux C, Jacquemin D, Perpete E, Assfeld X (2008) J Chem Theory Comput 4:637
- Loos PF, Dumont E, Laurent AD, Assfeld X (2009) Chem Phys Lett 475:120
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (2007) Numerical recipes: the art of scientific computing. Cambridge University Press, New York
- Golub GH, van Loan CF (2013) Matrix computations, 4th edn. The Johns Hopkins University Press, Baltimore
- Tsuchimochi T, Scuseria GE (2009) J Chem Phys 131:121102
- Sinitskiy AV, Greenman L, Mazziotti DA (2010) J Chem Phys 133:014104
- Sancho-Garcia JC (2011) Chem Phys Lett 511:172
- Stella L, Attaccalite C, Sorella S, Rubio A (2011) Phys Rev B 84:245117
- Wolfram Research, Inc. (2013) Mathematica 9
- Hehre WJ, Stewart RF, Pople JA (1969) J Chem Phys 51:2657
- Hehre WJ, Ditchfield R, Stewart RF, Pople JA (1970) J Chem Phys 52:2769
- Pulay P (1980) Chem Phys Lett 73:393
- Pulay P (1982) J Comput Chem 3:556
- Guo Y, Li W, Li S (2011) J Chem Phys 135:134107
- Challacombe M (1999) J Chem Phys 110:2332
- Tewarson RP (1973) Sparse matrices. Academic Press Inc., New York