

A black-box, general purpose quadratic self-consistent field code with and without Cholesky decomposition of the two-electron integrals

Tommaso Nottoli^a, Jürgen Gauss^b and Filippo Lipparini^a

^aDipartimento di Chimica e Chimica Industriale, Università di Pisa, Via G. Moruzzi 13, I-56124 Pisa, Italy; ^bDepartment Chemie, Johannes Gutenberg-Universität Mainz, Duesbergweg 10-14, D-55128 Mainz, Germany

ARTICLE HISTORY

Compiled October 18, 2021

ABSTRACT

We present the implementation of a quadratically convergent self-consistent field (QCSCF) algorithm based on an adaptive trust-radius optimization scheme for restricted open-shell Hartree-Fock (ROHF), restricted Hartree-Fock (RHF), and unrestricted Hartree-Fock (UHF) references. The algorithm can exploit Cholesky decomposition (CD) of the two-electron integrals to allow calculations on larger systems. The most important feature of the QCSCF code lies in its black-box nature – probably the most important quality desired by a generic user. As shown for pilot applications, it does not require one to tune the self-consistent field (SCF) parameters (damping, Pulay’s DIIS, and other similar techniques) in difficult-to-converge molecules. Also, it can be used to obtain a very tight convergence with extended basis sets – a situation often needed when computing high-order molecular properties – where the standard SCF algorithm starts to oscillate. Nevertheless, trouble may appear even with a QCSCF solver. In this respect, we discuss what can go wrong, focusing on the multiple UHF solutions of ortho-benzynes.

KEYWORDS

Hartree-Fock, self-consistent field, second-order, Levenberg-Marquardt, Cholesky decomposition

1. Introduction

Self-consistent field (SCF) methods are the starting point of virtually every quantum chemistry application. Kohn-Sham (KS) density-functional theory[1] (DFT), Hartree-Fock[2] (HF) and even semiempirical methods[3] require as a fundamental numerical step the variational optimization of the energy with respect to the orbitals, that are determined in the process. The standard algorithm to solve the optimization problem – the self-consistent field algorithm itself[4, 5] – is the iterative solution of a non-linear eigenvalue problem, which is solved using a fixed-point approach. As the problem is strongly non-linear, a simple iterative procedure is often not sufficient to achieve convergence. In the last several decades, various strategies have been developed to improve the reliability and stability of the SCF algorithm[6, 7], most notably Pulay’s direct inversion in the iterative subspace[8–10] (DIIS) and various refinement thereof[7, 11, 12], damping of the SCF iterations[13], level shifting[14], and combinations of the methods. Advanced strategies to compute accurate guesses of the initial density have also greatly improved the overall reliability of the method[15, 16], to the point that closed-shell systems seldom pose convergence problems in real life applications.

Unfortunately, things are not so easy when one is dealing with open-shell systems, or even with closed-shell systems with small energy gaps between the highest occupied and lowest unoccupied molecular orbital (HOMO, LUMO). Furthermore, even in well-behaved cases, achieving

a very tight convergence of the SCF density can be difficult, but mandatory for applications involving high-order response properties or geometrical derivatives, especially if a post-HF method is used. Being frustrated by problematic SCF convergence is therefore a common experience for computational chemists. For these reasons, alternative strategies and robust numerical procedures are still object of active investigation, despite SCF being possibly the most well-established technique in computational chemistry.

An alternative strategy for solving the SCF problem is to use a standard optimization technique, and, in particular, a second-order method[17, 18]. In second-order methods, the SCF energy is parametrized as a function of non-redundant orbital rotations and expanded up to second order in a Taylor series, obtaining thus a quadratic energy model. The latter is then optimized to find a step and the process is iterated until convergence. The straightforward Newton-Raphson (NR) method just described however suffers from a small convergence radius, which can cause an erratic or even divergent behavior if the optimization process is started far from a local minimum. It is possible to constrain the minimization so that the computed step is no larger than a user-defined trust radius. The trust-radius Newton method, known as Levenberg-Marquardt[19] (LM) optimization, can be further coupled with an adaptive choice of the trust radius, based on the agreement of the quadratic model with the actual energy. The global strategy, which has been originally proposed by Fletcher[19] (FLM), guarantees convergence to the closest local minimum for a well-behaved function and can therefore be used to implement a black-box SCF procedure. One of the most attractive features of such a procedure is that it exhibits a quadratic convergence rate, which makes it suited for applications where a very tight convergence of the SCF orbitals and density is required.

Computational cost is, of course, a fundamental aspect that one needs to keep in mind when designing a SCF code. A quadratically convergent SCF (QCSCF) scheme can be implemented in a direct fashion, where all the Hessian-vector products needed to compute the step are performed without assembling the full Hessian matrix and in the atomic orbital (AO) basis. The operations required to perform such a matrix-vector product (MVP) are computationally equivalent to the direct construction of a Fock matrix. Therefore, QCSCF and SCF exhibit the same computational scaling and can benefit both of Cauchy-Schwarz screening in an integral-direct implementation and even of linear-scaling techniques exactly in the same way. Nevertheless, QCSCF requires, in general, a larger number of Fock matrix builds and for well-behaved systems is always going to be more expensive than standard SCF. On the other hand, QCSCF does not require the $\mathcal{O}(N^3)$ diagonalization of the Fock matrix, which makes it in principle advantageous in the asymptotic regime.

Despite the increased computational cost, many second-order SCF implementations are available. An implementation based on the same algorithm as the one described in the present paper is present in the Dalton suite of programs[20] for restricted (RHF) and high-spin restricted open (ROHF) references, and a similar trust-region augmented Hessian implementation has been recently presented by Helmich-Paris[21] and implemented in the ORCA package[22] for restricted and unrestricted (UHF) references. In the Gaussian 16 suite of programs[23], a NR second-order method combined with a linear search when far from the quadratic regions is available for RHF and UHF. Other approaches, either based on a quasi-Newton update[24, 25] or on an orbital Hessian based preconditioned conjugate gradient[26], can be found in MOLCAS[27, 28], GAMESS,[29] and NWChem[30], respectively.

The QCSCF program described in this contribution has been implemented in the CFOUR suite of programs[31, 32]. CFOUR is a quantum chemistry package devoted to high-level post-HF calculations, and therefore the computational cost associated with solving the SCF equations is usually not a main concern for the typical application. Therefore, we do not pursue an integral-direct implementation, even though it would not present any additional difficulty, nor the use of linear-scaling techniques. To achieve some computational efficiency, we offer instead an implementation that can either proceed in a traditional fashion, reading pre-computed two-electron integrals from disk, or use their Cholesky decomposition[33–39] (CD). The latter possibility comes as a part of a long-term goal to deploy the CD machinery for subsequent post-HF calculations[32] that has been actively pursued by several developers of the CFOUR suite of

programs and that has recently been proposed for complete active space-SCF calculations[40] and for the calculation of NMR chemical shielding tensors at second-order Møller-Plesset perturbation theory (MP2) using gauge-including atomic orbitals[41]. On the other hand, having a robust, almost black-box SCF implementation is particularly attractive for the users of CFOUR that deal with open-shell systems, where the unrestricted (UHF) and high-spin restricted-open-shell HF (ROHF)[5] SCF equations can be particularly hard to converge. In other words, the main goal of this implementation is to save human time rather than machine time.

The paper is organized as follows. In section 2, we briefly recapitulate the norm-extended optimization algorithm and its application to the SCF problem. In section 3, we discuss the implementation of the various quantities required for the QCSCF procedure with and without CD of the two-electron integrals. In section 4, we present a few case studies, which represent prototypical applications of the QCSCF program and that illustrate possible problems and drawbacks that a user can encounter. We finally conclude the paper with a short summary.

2. Norm-extended optimization of the SCF energy

In this section, we discuss the main principles of a QCSCF implementation based on the norm-extended optimization (NEO) algorithm. The NEO scheme, originally formulated and implemented by Jensen and coauthors for multiconfigurational SCF wavefunctions[42, 43], is an elegant and efficient practical realization of the FLM second-order procedure that allows for a direct implementation. In this section, we focus our discussion on the high-spin ROHF optimization problem. In the following, we omit the high-spin specification and refer to the method simply as ROHF. RHF and UHF can be easily derived from the more general ROHF case. The ROHF determinant is parametrized in terms of N_r orbital rotations

$$|\Phi\rangle = e^{-\hat{\kappa}}|0\rangle, \quad (1)$$

where $|0\rangle$ is a reference determinant,

$$\hat{\kappa} = \sum_{ix} \kappa_{ix} \hat{E}_{ix}^- + \sum_{ia} \kappa_{ia} \hat{E}_{ia}^- + \sum_{xa} \kappa_{xa} \hat{E}_{xa}^- \quad (2)$$

is the elementary orbital rotation operator, where the rotations mix internal (i, j, \dots , doubly occupied) and active (x, y, \dots , singly occupied), internal and external (a, b, \dots , empty), and active and external orbitals, and $\hat{E}_{pq}^- = \hat{E}_{pq} - \hat{E}_{qp}$, where \hat{E}_{pq} is a singlet excitation operator (p, q, \dots , generic orbitals). The matrix elements of $\hat{\kappa}$ introduce a complete, non-redundant parametrization of orbital rotations that can connect the reference determinant $|0\rangle$ to any non-orthogonal determinant. This fundamental result, known as Thouless' theorem,[44] is the basis of direct optimization SCF techniques. We define a quadratic model of the SCF energy by expanding the expectation value of the Hamiltonian up to second order in κ :

$$\mathcal{Q}(\kappa) = E_0 + \sum_{pq} \kappa_{pq} g_{pq} + \frac{1}{2} \sum_{pqrs} G_{pq,rs} \kappa_{pq} \kappa_{rs}, \quad (3)$$

where E_0 is the reference energy. The orbital-rotation gradient $g \in \mathbb{R}^{N_r}$ is

$$g_{pq} = \langle 0 | [\hat{E}_{pq}^-, \mathcal{H}] | 0 \rangle \Big|_{\kappa=0}, \quad (4)$$

and the orbital-rotation Hessian $G \in \mathbb{R}^{N_r \times N_r}$ is given by

$$G_{pq,rs} = \frac{1}{2} (1 + P_{pq,rs}) \langle 0 | [\hat{E}_{pq}^-, [\hat{E}_{rs}^-, \mathcal{H}]] | 0 \rangle \Big|_{\kappa=0}, \quad (5)$$

where $P_{pq,rs}$ permutes the indices pairs pq and rs . The FLM procedure is an iterative algorithm that computes an optimization step by minimizing the quadratic model in eq. 3 under the constraint that the norm of the step is not larger than a user-defined trust radius R_t . This is achieved by introducing the constraint using a Lagrange multiplier ν :

$$\mathcal{L}(\kappa, \nu) = \mathcal{Q}(\kappa) - \frac{1}{2}\nu(\|\kappa\|^2 - R_t^2). \quad (6)$$

The resulting Euler-Lagrange equations, also known as the LM equations, are

$$\begin{cases} (G - \nu I)\kappa = -g, \\ \|\kappa\|(\nu) = R_t. \end{cases} \quad (7)$$

The trust radius is updated dynamically during the optimization based on the agreement between the energy and its quadratic model. Let ΔE be the actual energy variation after a LM step, i.e., $\Delta E = E(\kappa) - E_0$ and let ΔQ be the predicted variation using the quadratic model, i.e., $\Delta Q = Q(\delta) - E_0$, where κ is the solution to the LM equations. Let $r = \Delta E/\Delta Q$ be the ratio of the variations. If the ratio is negative, the energy is rising and the step is rejected. The trust radius is reduced by a factor (0.66 in our implementation) and a new step is computed. If the ratio is positive the step is accepted. If $0 < r \leq 0.25$, the agreement of the quadratic model with the energy is poor and the trust radius is reduced (again, in our implementation, by a factor 0.66). If $0.25 < r \leq 0.75$ the trust radius is left unchanged, while if $r > 0.75$ the trust radius is increased (by a factor 1.2 in our implementation). The algorithm is robust with respect to the choice of the parameters used to adapt the trust radius, and it can be proven that convergence to a local minimum is always achieved[19].

In principle, solving the LM equations requires the knowledge of at least the lowest eigenvalue λ_1 of the Hessian, as it can be proved that, to get to a minimum, the constraint equation needs to be solved in the range $\nu \in (-\infty, \lambda_1)$. In other words, one would need to compute the lowest eigenvalue of the Hessian and then to solve the LM equations. The NEO algorithm is an efficient, combined realization of the two steps that is achieved by introducing a gradient-scaled, augmented Hessian $L(\alpha) \in \mathbb{R}^{N_r+1, N_r+1}$, defined as

$$L(\alpha) = \begin{pmatrix} G & \alpha g \\ \alpha g^\dagger & 0 \end{pmatrix}. \quad (8)$$

Let P be an orthogonal projector such that $PL(\alpha)P = G$. The NEO step is given by

$$\kappa = \frac{1}{\alpha s}Py, \quad (9)$$

where

$$L(\alpha)y = \nu_1 y, \quad (10)$$

i.e., y is the eigenvector of $L(\alpha)$ associated to its lowest eigenvalue, $s = (1 - P)y$, and α is obtained by solving numerically the one-dimensional equation $\|\kappa\| = R_t$. It can be shown that the NEO step solves the LM equations with a level-shift parameter $\nu = \nu_1$. The Hylleraas-Undheim-MacDonald theorem guarantees that, as $G = PL(\alpha)P$, $\nu_1 \leq \lambda_1$. Therefore, the NEO algorithm converges to a local minimum[42].

The NEO algorithm requires one to compute the lowest eigenvalue of the augmented Hessian L , which can be done in a direct fashion using an iterative algorithm such as Davidson diagonalization. It can also be shown that if the vector $(0, \dots, 0, 1)$ is kept into the subspace, it is possible to compute a new step, in case the one computed is rejected, without having to solve the eigenvalue problem in eq. 10 again. As a final note, we remark that, as soon as the optimization has reached a local region and the quadratic approximation becomes valid, the

NEO step becomes fully equivalent to a standard NR step. Therefore, in the final stage of the optimization, we switch from solving the NEO equation to the plain NR ones.

3. Implementation

Equations for the energy, the gradient, and the MVP can be conveniently obtained by introducing the generalized Fock matrix, F , whose elements are obtained as follows

$$F_{ip} = 2(F_{ip}^I + F_{ip}^A), \quad (11)$$

$$F_{xp} = Q_{xp} + F_{xp}^I, \quad (12)$$

$$F_{ap} = 0, \quad (13)$$

with F_{pq}^I , F_{pq}^A , and Q_{xp} being the inactive Fock matrix, the active Fock matrix, and the Q matrix respectively. The latter can be effectively computed in the AO basis

$$F_{\mu\nu}^I = h_{\mu\nu} + \sum_{\rho\sigma} P_{\rho\sigma}^I \left[(\mu\nu|\rho\sigma) - \frac{1}{2}(\mu\sigma|\rho\nu) \right], \quad (14)$$

$$F_{\mu\nu}^A = \sum_{\rho\sigma} P_{\rho\sigma}^A \left[(\mu\nu|\rho\sigma) - \frac{1}{2}(\mu\sigma|\rho\nu) \right], \quad (15)$$

$$Q_{\mu\nu} = \sum_{\rho\sigma} P_{\rho\sigma}^A [(\mu\nu|\rho\sigma) - (\mu\sigma|\rho\nu)], \quad (16)$$

where we have used Mulliken notation for the two-electron integrals and Greek indices to refer to the AOs. Here, $P_{\mu\nu}^I = 2 \sum_i C_{\mu i} C_{\nu i}$ and $P_{\mu\nu}^A = \sum_u C_{\mu u} C_{\nu u}$ are the inactive and active one-body density matrices written in the AO basis, respectively. As it is evident from the equations above, the actual implementation requires minor modifications to the customary routine that assembles the RHF Fock matrix.

The generalized Fock matrix is transformed into the MO basis and then used to calculate the ROHF energy

$$E_{\text{ROHF}} = \sum_i \left(h_{ii} + \frac{1}{2} F_{ii} \right) + \frac{1}{2} \sum_u (h_{uu} + F_{uu}). \quad (17)$$

Furthermore, its anti-symmetric part is used to compute the gradient as follows

$$g_{pq} = 2(F_{pq} - F_{qp}), \quad (18)$$

where the only relevant rotations are the ones mixing orbitals belonging to different classes (i.e., internal, active, external).

The eigenvalue problem in eq. 10 is solved via Davidson diagonalization, while the NR linear system is solved using an iterative preconditioned conjugate-gradient (PCG) solver. Both algorithms are implemented in a matrix-free spirit, that is, they only require one to perform MVPs, and not to build and store in memory the full Hessian matrix. It is important to stress that the overall algorithm works in the MO basis, as this allows us to exploit the diagonal dominant character of the MO rotation Hessian. In the MO basis, the MVP can be written as

$$b_{pq} = 2(\tilde{F}_{pq} - \tilde{F}_{qp}) + \frac{1}{2} \sum_r (\kappa_{pr} g_{rq} - g_{pr} \kappa_{qr}), \quad (19)$$

where κ is a trial vector in Davidson's algorithm and where we introduce the one-index transformed generalized Fock matrix \tilde{F} , which is defined as in eq. 11 but using intermediate matrices

computed with density matrices dressed with the trial vector. To avoid transforming the two-electron integrals into the MO basis, we compute the MVP in the AO basis. In particular, we define transformed and symmetrized one-body density matrices:

$$\tilde{P}_{\mu\nu}^I = 2 \sum_{iq} C_{\mu i} \kappa_{iq} C_{\nu q} + 2 \sum_{iq} C_{\nu q} \kappa_{qi} C_{\sigma i}, \quad (20)$$

$$\tilde{P}_{\mu\nu}^A = \sum_{uq} C_{\mu u} \kappa_{uq} C_{\nu q} + \sum_{uq} C_{\nu q} \kappa_{qu} C_{\sigma u}, \quad (21)$$

and use them to assemble the one-index transformed internal and active Fock matrices and Q matrix. These are in turn used to build the one-index transformed generalized Fock matrix, which is then finally transformed back to the MO basis. Therefore, computing the required MVP exhibits a computational cost similar with the one of a standard SCF iteration and, more importantly, the same scaling with respect to the system size. Specifically, the leading term operation in both the SCF and QCSCF algorithms scale as $\mathcal{O}(N^4)$, where N is the number of basis functions. However, each QCSCF iteration requires the solution of either the NEO or the NR equations with an iterative solver thus increasing the prefactor of a QCSCF calculation with respect to the standard SCF one.

When using the CD, the two-electron integrals matrix – written in the AO basis – is approximated as follows

$$(\mu\nu|\rho\sigma) \simeq \sum_K^{N_{\text{ch}}} L_{\mu\nu}^K L_{\rho\sigma}^K, \quad (22)$$

where $L_{\mu\nu}^K$ is a Cholesky vector and N_{ch} is the rank of the decomposition. In this framework, it is convenient to compute the Coulomb and exchange contributions to the various matrices separately[38, 45]. The Coulomb contribution can be computed by performing a straightforward contraction of a Cholesky vector with the one-body density matrix and then multiplying the resulting factor by a Cholesky vector. On the other hand, in order to compute the exchange contributions, it is convenient to first half-transform the Cholesky vectors into the MO basis, and then assemble the exchange contribution. Considering the inactive Fock matrix as an example we have

$$F_{\mu\nu}^I = h_{\mu\nu} + \sum_K^{N_{\text{ch}}} \left(\sum_{\rho\sigma} P_{\rho\sigma}^I L_{\mu\nu}^K L_{\rho\sigma}^K - \sum_i L_{\mu i}^K L_{\nu i}^K \right) \quad (23)$$

where, $L_{\mu i}^K = \sum_{\nu} C_{\nu i} L_{\mu\nu}^K$. A similar procedure is applied also for the calculation of the one-index transformed matrices.

The restricted and unrestricted SCF implementations can be trivially obtained as sub-cases of the ROHF one. Second-order RHF is simply derived by neglecting the contributions of active orbitals, that is, by setting the active density matrix to zero. Under these circumstances, only the inactive Fock matrix has to be considered. Regarding the second-order implementation of UHF, we have to take into account two different bases – one for the alpha and one for the beta electrons. As a result, we have two different set of orbital rotation parameters

$$\hat{\kappa} = \sum_{ia} \kappa_{ai}^{\alpha} \hat{E}_{ai}^{\alpha} + \sum_{ia} \kappa_{ai}^{\beta} \hat{E}_{ai}^{\beta}, \quad (24)$$

where, $\hat{E}_{pq}^{\alpha} = \hat{a}_{p\alpha}^{\dagger} \hat{a}_{q\alpha}$ and $\hat{E}_{pq}^{\beta} = \hat{a}_{p\beta}^{\dagger} \hat{a}_{q\beta}$. Furthermore, two different density matrices can be obtained namely, ${}^{\alpha}P_{\mu\nu}^I = \sum_i C_{\mu i}^{\alpha} C_{\nu i}^{\alpha}$ and ${}^{\beta}P_{\mu\nu}^I = \sum_i C_{\mu i}^{\beta} C_{\nu i}^{\beta}$. These are used to assemble the

alpha and beta Fock matrices

$${}^{\alpha}F_{\mu\nu}^I = h_{\mu\nu} + \sum_{\rho\sigma} {}^{\alpha}P_{\rho\sigma}^I [(\mu\nu|\rho\sigma) - (\mu\sigma|\rho\nu)] + \sum_{\rho\sigma} {}^{\beta}P_{\rho\sigma}^I (\mu\nu|\rho\sigma), \quad (25)$$

$${}^{\beta}F_{\mu\nu}^I = h_{\mu\nu} + \sum_{\rho\sigma} {}^{\beta}P_{\rho\sigma}^I [(\mu\nu|\rho\sigma) - (\mu\sigma|\rho\nu)] + \sum_{\rho\sigma} {}^{\alpha}P_{\rho\sigma}^I (\mu\nu|\rho\sigma). \quad (26)$$

Accordingly, ${}^{\alpha}F^I$ and ${}^{\beta}F^I$ are used to compute the alpha and beta parts of the gradient respectively. A summary of the main equations for the RHF and UHF references can be found in Table 1.

Reference	Energy	Gradient	MVP
RHF	$\sum_i (h_{ii} + F_{ii}^I)$	$g_{ai} = -4F_{ia}^I$	$b_{ai} = \tilde{g}_{ai}$
UHF	$\sum_i (h_{ii} + {}^{\alpha}F_{ii}^I + {}^{\beta}F_{ii}^I)$	$g_{ai}^{\alpha} = -4{}^{\alpha}F_{ia}^I$ $g_{ai}^{\beta} = -4{}^{\beta}F_{ia}^I$	$b_{ai}^{\alpha} = \tilde{g}_{ai}^{\alpha}$ $b_{ai}^{\beta} = \tilde{g}_{ai}^{\beta}$

Table 1. Leading equations used in the second-order implementation of RHF and UHF. We have used a shorthand notation for \tilde{g}_{ai} that means a gradient built with one-index transformed matrices; moreover, this is the only contribution to the MVP since the commutator-like term of eq. 19 vanishes.

We conclude this section with an important remark. Both Davidson diagonalization and the PCG solver require a preconditioner, the most common choice being the inverse diagonal of the matrix. This is a good choice, as the MO rotation Hessian is diagonally dominant in the MO basis. However, assembling the exact diagonal of the Hessian can be expensive. This is not a problem in neither RHF nor UHF, as the leading term of the electronic Hessian’s diagonal is given by the difference $F_{aa}^I - F_{ii}^I$: the two-electron integral contributions to the diagonal can therefore be safely neglected. However, the same is, unfortunately, not the case for ROHF. In particular, the diagonal elements associated with the inactive-active and active-external rotations are poorly approximated by the diagonal elements of the inactive Fock matrices. A much better approximation can be obtained by noting that some of the two-electron contributions to the aforementioned diagonal elements are given by Q_{ii} and Q_{aa} . From the definition of the Q matrix such blocks would not exist; nevertheless, we have direct access to them since we are computing Q in the AO basis. In this way, a good approximation to the Hessian’s diagonal can be computed with no additional cost. The inclusion of such terms is fundamental in order to achieve a good convergence for both the Davidson and PCG algorithms.

4. Numerical results

In this section, we present numerical results obtained with the quadratic SCF implementation described in the paper. First, we apply the strategy to medium-sized, hard to converge cases to show the numerical stability of the algorithm. Then, we apply the CD version of the code to larger molecules, to show that thanks to the compression operated by the CD technique, calculations can be performed efficiently even for very large systems. Finally, we present a discussion on the possible numerical issues that can arise even when using a quadratically convergent implementation and give some suggestions to rationalize and troubleshoot them. All the geometries are listed in the Supplemental Material.

The QCSCF algorithm has been implemented in the CFOUR suite of programs[31, 32], which has been used for all the following QCSCF calculations. In our setup, a few standard SCF iterations are performed to generate a reasonable starting guess for the second-order algorithm. While the second-order optimization can be used for the overall calculation, this is seldom a

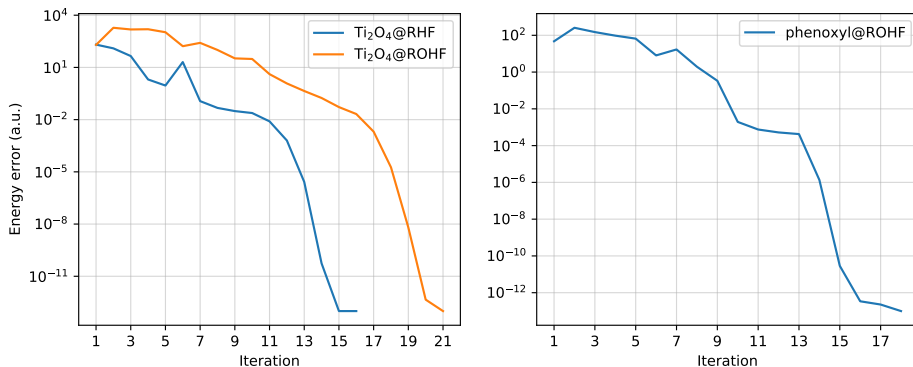


Figure 1. Convergence profile for Ti_2O_4 , both for the RHF (blue) and ROHF (orange) references, and of the phenoxyl radical respectively on the left and right panel. The QC algorithm started at iteration 16 for triplet Ti_2O_4 , 10 for singlet Ti_2O_4 , and 12 for the phenoxyl radical. The final electronic energies for the singlet and triplet Ti_2O_4 are $-1996.144\ 738\ 811$ and $-1996.100\ 447\ 937\ E_h$ respectively while the one for the phenoxyl radical is $-304.953\ 646\ 044\ E_h$.

good idea, as it would require a large number of expensive quadratic steps to reach the local region. On the other hand, a small number of standard SCF iterations are usually enough to get to a good starting point, from which the strong and fast convergence of a second-order scheme can be efficiently exploited.

Specifically, starting from a guess generated by diagonalizing the core-Hamiltonian, we perform up to 30 SCF iterations applying a damping of 0.5 to the first 5 iterations and then using Pulay’s DIIS[8] to accelerate convergence. We switch to the quadratic algorithm as soon as the root mean square (RMS) deviation of the density-matrix increment is lower than 10^{-1} and its maximum deviation is lower than 1.

4.1. Quadratically convergent calculations on small and large molecules

In order to show the black-box nature of the proposed algorithm, we tested it on two systems that exhibit problematic convergence patterns. The first system is Ti_2O_4 in its D_{2h} geometry, where the two titanium atoms lie on one of the three C_2 axes, the second is the phenoxyl radical. The first system is used as a template to troubleshoot SCF converge issues¹ in the Amsterdam Density Functional (ADF) quantum-chemistry package[46]. We compute, using Dunning’s aug-cc-pVTZ basis set[47], both a singlet and a triplet wavefunction, the latter using a ROHF reference. The second system, a doublet in its ground state, is usually well behaved, but we choose a particular geometry at which two electronic states are nearly intersecting. Again, we use a ROHF reference and we employ Dunning’s cc-pVDZ basis set[48]. We report in Figure 1 the convergence profile of the calculation on Ti_2O_4 (left panel) and of the phenoxyl radical (right panel). Both calculations converge reasonably smoothly in a limited number of iterations. It is interesting to note that we could not get the phenoxyl radical ROHF calculation to converge using the standard SCF code in CFOUR using the algorithm described in Ref. [49], despite various attempts using several combinations of DIIS and damping parameters. In other words, using a QCSCF program can turn a labor intensive, possibly fruitless activity into a simple, routine application, at the cost of increased computer time.

As a second example of standard use of a QCSCF implementation, we optimize the RHF wavefunction of a small organic molecule, parnitroaniline (PNA), using Dunning’s aug-cc-pVDZ basis set[48]. This is not a problematic system, as convergence can be easily achieved with a standard SCF code, but becomes an issue for applications where a very tight SCF convergence is required. This is the case when one is interested in computing high-order molecular properties using a post-HF method. A typical example is the calculation of accurate anharmonic force

¹See https://www.scm.com/doc/ADF/Examples/SCF_Ti2O4.html

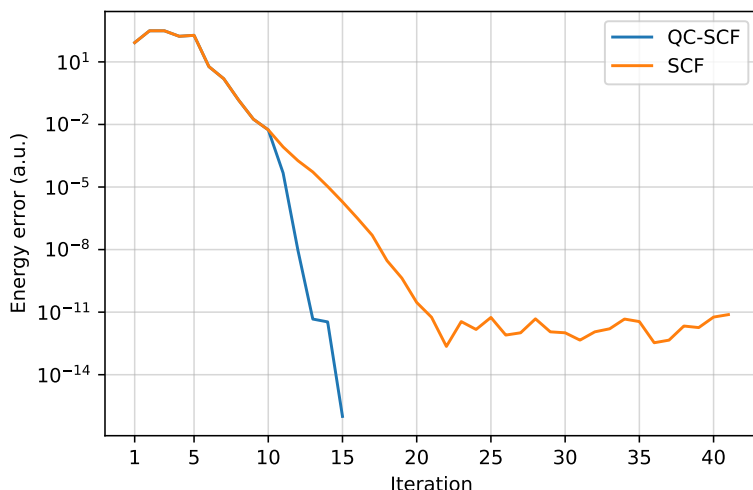


Figure 2. Convergence profile for the PNA molecule for the SCF (orange) and QC-SCF (blue) code using Dunning’s aug-cc-pVDZ basis set. The QCSCF code performed 10 regular SCF iterations before switching to the QC algorithm. The final electronic energy is $-489.277\ 111\ 404\ E_h$.

fields, as the ones required for the treatment of anharmonicity. In such applications, cubic and quartic force fields are in general computed by numerically differentiating analytical Hessians, the latter computed for instance using coupled-cluster (CC) theory. To achieve a good numerical accuracy, a tight convergence of the SCF and CC amplitude equations, as well as the various coupled-perturbed equations, is paramount: this can be difficult for a regular SCF code when using extended basis sets. In Figure 2, we compare the convergence profile of the standard SCF code in CFOUR with QCSCF. The standard SCF code has no issue achieving a reasonable (10^{-7} to 10^{-8} in the RMS norm of the density increment) convergence, but then stagnates and oscillates. On the other hand, the quadratically convergent code achieves effortlessly the required tight convergence (10^{-11} in the RMS norm of the gradient).

The systems proposed so far are medium-sized, and provide examples of applications that are typical for the users of the CFOUR suite of programs. The standard implementation that relies on precomputed two-electron integrals written on disk has been used in all cases.

molecule	$^1A_{1g}$	It.	3A_u	It.
Ti_2O_4	-1996.144 738 811	6	-1996.100 447 937	5

Table 2. Results for RHF and ROHF calculation on Ti_2O_4 with the aug-cc-pVTZ basis set. The energies are in Hartree units, next to them the number of iterations required by the second-order optimization. The two calculations were done without CD and exploiting point-group symmetry.

As the next examples, we tested the algorithm on three larger systems. Here, CD of the two-electron integrals has been exploited using a decomposition threshold equal to 10^{-4} . The first molecule is an aqua thiolate iron(III) porphyrin complex ($HSFe^{III}OH_2$) in its doublet state used as a model system for the active site of the cytochrome P450 in Ref. [50]. The second calculation was done on the triplet state of a binuclear copper magnet (CUAQUACO2) whose geometry has been taken from Ref. [51]. Finally, we optimized the singlet state of a chlorophyll molecule where the phytol tail has been substituted with a hydrogen atom to reduce the computational cost. The geometry was optimized with B3LYP/6-31G(d)[52, 53], using the Gaussian 16[23] suite of programs. In Figure 3, a representation of the three systems under study is shown. The optimization of the open-shell systems, i.e., $HSFe^{III}OH_2$ and CUAQUACO2, was carried out with ROHF. For all the calculations we adopted the cc-pVTZ basis set[48]. In Table 3,

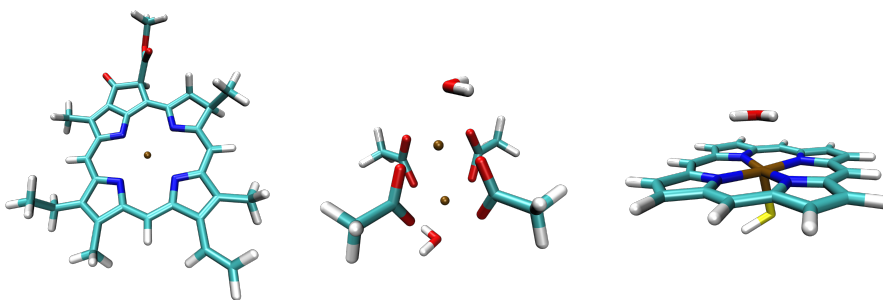


Figure 3. Molecular representations of the three molecules used to test the CD based QCSCF algorithm. From the left: chlorophyll, CUAQUACO2, and $\text{HSFe}^{\text{III}}\text{OH}_2$.

we report the number of basis functions, the number of iterations required to converge the second-order algorithm, the time spent to assemble the Fock matrix (F^I , F^A , and Q for the ROHF calculations), and the total CPU wall time for the whole calculation (i.e., preliminary SCF iterations and quadratically convergent ones). All the calculations have been performed on a computer node equipped with four Intel Xeon Gold 6140M CPUs running at 2.30GHz sharing the work among 28 cores. For all the systems the Cholesky vectors are kept in-core. The time spent performing the MVP is not reported since it is similar with the time needed to compute the Fock matrix.

molecule	N_b	It.	Fock (s)	Total (min)
$\text{HSFe}^{\text{III}}\text{OH}_2$	1062	9	19	26.1
CUAQUACO2	900	9	11	15.8
chlorophyll	1624	6	27	28.4

Table 3. Large systems results. For each of them, we report the number of basis functions, the number of iterations required by the second-order optimization, the time needed to build a Fock matrix (in seconds), and the total CPU wall time for the whole SCF calculation (in minutes).

While a comprehensive analysis of the computational advantages introduced by CD is out of the scope of this paper, it is worth pointing out that, using a traditional, non integral-direct SCF implementation such as the one available in CFOUR, the calculations reported in table 3 would probably be out of reach, or at least very demanding. A qualitative idea of the savings introduced by CD can be deduced by comparing the timings for the traditional and CD implementations on the Ti_2O_4 system reported in Table 2. For such a system, using Abelian point-group symmetry, the restricted QCSCF calculation took about 2 minutes on a single core using the traditional SCF implementation, with a Fock-matrix formation requiring about 2 seconds. The same calculation using CD, always exploiting point-group symmetry, only required 13 seconds, with about 0.2 seconds spent for a Fock-matrix formation. While a CD-based implementation is not going to be competitive with an integral-direct one for very large systems, especially when advanced linear-scaling techniques are used, it is already advantageous even for small and medium systems with respect to a traditional one.

4.2. What can go wrong and how to deal with that?

Despite the rigorous and sound convergence properties of a quadratic optimization algorithm, there are still a few issues that can arise in a calculation. First, one must mention that the

convergence properties of the FLM algorithm hold in infinite precision. While this is usually not an issue on double-precision machines, the effects of finite arithmetic precision become apparent when trying to achieve very tight convergence of the SCF, especially when using very large basis sets comprising diffuse functions. This can be rationalized in terms of overall conditioning of the problem. If the basis set presents near linear dependencies, i.e., if the overlap matrix has small eigenvalues, the numerical precision of the computed quantities degrades. As a rule of thumb, one cannot expect to achieve a convergence tighter than the product of the machine precision times the ratio between the largest and smallest eigenvalues of the overlap matrix. In practice, this means that it is possible to converge the SCF up to 10^{-11} – 10^{-12} depending on the basis set, which is usually more than sufficient for high-precision applications.

That being said, the fact that the QCSCF will converge does not imply that it will always converge to the desired solution. There are two possible, common scenarios that can present. First, the QCSCF code will converge to a state of the same symmetry as its starting point, but not the desired state. For well-behaved systems, the few SCF iterations performed in the beginning are usually enough to establish, using the Aufbau principle, the right occupation. The user should however be wary that there is no guarantee that this will work automatically: it is therefore a good practice to specify explicitly the occupation, or the symmetry of the state, to ensure the convergence of the optimization procedure to the desired wavefunction.

A second case that requires some attention by the user is, in general, any UHF calculation. It is well known that the solution to the UHF problem is not unique and that many solutions with different levels of spin contamination can exist. A minimization algorithm will converge to the closest local minimum, so there is no guarantee that the QCSCF solution will be the global minimum. It is the experience of the authors that, with respect to a standard SCF code, a QCSCF optimization tends to converge to the solution lowest in energy and with highest spin contamination. After a UHF calculation, the user needs to check whether the obtained solution is acceptable.

A third, not common, case that can be encountered is convergence to an unstable solution. This is something that typically happens when performing a calculation on a symmetric system without enforcing point-group symmetry and is somewhat similar to what has already been discussed above about convergence to a state with the wrong occupation. The second-order code should in principle always converge to a minimum; however, there are numerical reasons that may leave the optimization stuck into a saddle point. This is a consequence of the parametrization choice, as Thouless’ theorem[44] specifically states that the determinant one can obtain with a rotation parametrized as in eq. 1 cannot be orthogonal to the reference. Therefore, if the QCSCF optimization starts from a determinant orthogonal to the minimum, whether due to imposed or to numerical symmetry constraints, it is prevented from reaching the minimum itself. If the minimum has the same symmetry of the starting point – including if no point-group symmetry is enforced – this can easily be avoided by perturbing, at the beginning of the second-order optimization, the MO gradient with random noise. On the other hand, if symmetry is enforced and there exists a broken-symmetry solution, this cannot be reached with a second-order method. The latter case can be diagnosed doing a stability analysis and resolved, if the broken-symmetry solution is of interest, enforcing symmetry in a lower subgroup or by completely removing the enforcement of symmetry.

We illustrate all the discussed problematic behaviors using the triplet state of *ortho*-benzyne (o-benzyne) as an example. Such a molecule exhibits C_{2v} point-group symmetry and has a singlet ground electronic state. This molecule has also been used by Tsuchimochi and Scuseria as a test case for their constrained UHF method[54, 55]. We follow their procedure and optimize the geometry using density-functional theory, namely, the B3LYP functional[52] in conjunction with Pople’s 6-31G(d) basis set[53]. The geometry optimization has been performed using the Gaussian 16[23] suite of programs. On the optimized geometry, we compute the UHF wavefunction with Dunning’s cc-pVTZ basis set[48] using four different setups, namely i) enforcing point-group symmetry, but without specifying the occupation numbers; ii) enforcing point-group symmetry and specifying the occupation numbers; iii) without symmetry, using the normal setup; iv) without symmetry, adding a small perturbation (0.01 times the gradient’s

norm times a uniformly distributed random number between -0.5 and 0.5) to the gradient at the beginning of the optimization. The results, together with a short comment, are reported in Table 4. In the first setup, the initial SCF iterations are tasked with computing the occupation

Setup	Energy	Spin contamination	Comment
i	-229.261 071	0.0204	wrong occupation
ii	-229.468 767	0.0277	symmetric solution
iii	-229.261 071	0.0204	unstable
iv	-229.471 139	0.4156	symmetry broken

Table 4. Results of UHF calculations for *o*-benzynes with four different setups. Setup i): point-group symmetry enforced, no initial occupation given. Setup ii): point-group symmetry enforced, initial occupation given in input. Setup iii): no symmetry enforced. Setup iv): no symmetry enforced, the gradient is perturbed at the beginning of the second-order optimization. We report the electronic energy (in Hartree), the spin contamination of the wavefunction, and a short comment that describes the solution found.

of the wavefunction using the Aufbau principle. When convergence of the preliminary SCF is reached, the resulting electronic configuration is the following: 10 doubly occupied a_1 orbitals, 1 doubly and one singly occupied b_1 orbitals, 8 doubly occupied b_2 orbitals and a singly occupied a_2 orbital. QCSCF then converges without problems the given state, which is however not the ground triplet state. In the second setup, we specify in input the correct occupations (9 doubly and one singly occupied a_1 orbitals, 2 doubly occupied b_1 orbitals, 7 doubly and one singly occupied b_2 orbitals and one doubly occupied a_2 orbitals), which results in the correct behaviour. We note that, while the states that result from the first two setups both have B_2 symmetry, the occupations make them strictly orthogonal. In other words, no symmetry-allowed orbital rotation can link the two states, making therefore it impossible for the QCSCF code to converge from one to the other.

In the third setup, no symmetry is enforced; however, the initial SCF iterations make the QC procedure start very close to an unstable stationary point, to which the optimization gets stuck. It is interesting to note that the unstable solution found with this setup is exactly the same found with setup i), i.e., enforcing symmetry but without specifying the occupation. This suggests that the QC optimizer started from a wavefunction that is orthogonal to the actual minimum. In other words, even though symmetry was not enforced, it was still present numerically, which explains the observed behavior.

In the fourth setup, we repeated the calculation without symmetry, but added a small, random perturbation to the gradient at the beginning of the QCSCF process. The small amount of random noise allowed the QCSCF optimizer to connect to the orthogonal subspace where the minimum lies and converge to it without particular effort. It is interesting to note that the optimization found a solution that has a lower energy than the symmetric one, and a much larger spin contamination. This is the same solution found by Tsuchimochi and Scuseria. A comparison between the UHF solutions found with the various setups is particularly interesting to understand the behavior of the QCSCF optimization. Performing a stability analysis, the symmetric solution from the second setup is found to exhibit an instability towards a symmetry-broken UHF solution, which is exactly the one found with the fourth setup. The latter is in turn stable. All these results confirm numerically what was discussed above. In setups i) and iii), QCSCF converges to a solution with the wrong occupation – which is explicitly enforced in the first case, and is a numerical consequence of the preliminary SCF iterations in the third. In the second setup, the minimum within the symmetry is found. In the fourth setup, as no symmetry is enforced and the symmetry is artificially broken with a small, random perturbation, a true minimum is found, which corresponds to the stable, symmetry broken solution. It is worth commenting that the lowest energy, stable solution exhibits a remarkably large spin contamination, besides being symmetry broken: whether this is the solution of interest, or whether the symmetry one found with the second setup is preferable, is ultimately a choice left

to the user.

We conclude this section with some considerations concerning computational efficiency. The choice of parameters that we adopt by default, namely, the parameters that control the preliminary SCF iterations and the initial trust radius, represent a good compromise between robustness and efficiency. This is, of course, not something that applies in general to every molecular system. It is therefore likely that there exists, for any system, a specific combination of parameters that minimizes the number of overall iterations, and therefore the computational cost. Tuning these parameters is not required to achieve convergence, but may speed up the computation, which is particularly relevant for applications involving large and very large systems. Nevertheless, the spirit of a quadratically convergent SCF is to minimize human effort with respect to computational effort: it is up to the user to decide whether to invest time into the numerical optimization of the procedure for a specific application. Let us remark once again that, even if an optimal setup is found, a QCSCF calculation is inevitably going to be more expensive, heuristically, about twice to three times as expensive, than a well-behaving regular SCF one.

5. Conclusions

In this contribution, we presented a quadratically convergent self-consistent field program that can achieve convergence of the SCF iteration in a black-box manner for restricted, unrestricted, and (high-spin) restricted-open-shell references. The implementation is based on the Fletcher-Levenberg-Marquardt trust-radius Newton method, in its direct formulation known as norm-extended optimization. All the operation are performed in a direct fashion, without the need of assembling explicitly the MO rotation Hessian nor of transforming the two-electron integrals into the MO basis. The implementation is also general, as it can deal with open-shell systems using either a UHF or a ROHF reference. The resulting algorithm is suited to be used in conjunction with integral-direct, and even linear-scaling techniques to increase its computational performances. As this implementation is meant for high-accuracy applications on small- to medium-sized molecules, we explore a different strategy to reduce computational cost, namely, the use of Cholesky decomposition of the two-electron integrals. We showed with numerical examples how the QCSCF code can be helpful not only with hard-to-converge cases, but also with regular cases where a very tight convergence is needed. Using the CD of the two-electron integrals, larger systems become easily treatable.

While a QCSCF calculation is in general more expensive than its linear counterpart, we showed that it is possible to achieve the same scaling with respect to the size of the system as in regular SCF. Furthermore, as regular SCF and QCSCF share as the leading computational operation the construction of Fock matrices, they can exploit the same techniques to accelerate the computation, including integral-direct implementations, linear-scaling methods and low-rank approximations of the two-electron integrals matrix. As an example, we showed an implementation based on the Cholesky decomposition of the two-electron integrals, as the use of such a technique is part of a widespread effort amongst the developers of the CFOUR suite of programs.

A second-order optimization algorithm has the remarkable property of being completely predictable. Its black-box nature is shown in its ability to always converge to a solution, which, however may not be the desired one. We discussed and rationalized the main reasons why QCSCF can converge to either a solution with the wrong symmetry or to an unstable solution, and showed how one can overcome such difficulties by either guiding the optimizer to the right symmetry state by specifying an occupation or, when working enforcing either no or a reduced symmetry point-group, by adding to the gradient with a small random perturbation.

In conclusion, a second-order SCF code can be a useful tool to converge problematic cases to a very tight threshold in an almost automated way, without the need of tweaking and tuning various parameters to achieve the desired accuracy. We hope that the QCSCF program that will be made available in the next release of CFOUR will provide the community with a useful

tool.

An interesting development that we intend to pursue in the future is the coupling of the FLM algorithm with the use of an approximated Hessian, in a quasi-Newton spirit. While, to the best of our knowledge, there are no convergence results as strong as the ones the FLM method relies on for quasi-Newton trust-radius methods, Fletcher[19] suggests that the same conclusions are likely to apply. This would lead to an approach that is, at the same time, computationally very efficient and as robust as possible.

Acknowledgments

This work is dedicated to Prof. John F. Stanton – quite possibly the biggest fan of the QCSCF code in CFOUR – on the occasion of his 60th birthday.

Supplemental Material

Molecular geometries of the systems used to test the QCSCF solver.

References

- [1] Kohn, W.; Sham, L.J. Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **1965**, *140*, A1133–A1138.
- [2] Hartree, D.R. *The calculation of atomic structures*; Wiley, London, 1957.
- [3] Thiel, W. *Semiempirical Methods*, Grotendorst, J., Ed.; John von Neumann Institute, Jülich, 2000; pp 261–283.
- [4] Roothaan, C.C.J. New developments in molecular orbital theory, *Rev. Mod. Phys.* **1951**, *23*, 69–89.
- [5] Roothaan, C.C. Self-consistent field theory for open shells of electronic systems, *Rev. of Mod. Phys.* **1960**, *32*, 179–185.
- [6] Rabuck, A.D.; Scuseria, G.E. Improving self-consistent field convergence by varying occupation numbers, *J. Chem. Phys.* **1999**, *110*, 695–700.
- [7] Cancès, E. *SCF algorithms for HF electronic calculations*; Springer, Berlin, Heidelberg, 2000; pp 17–43.
- [8] Pulay, P. Convergence acceleration of iterative sequences - the case of SCF iteration, *Chem. Phys. Lett.* **1980**, *73*, 393–398.
- [9] Pulay, P. Improved SCF convergence acceleration, *J. Comput. Chem.* **1982**, *3*, 556–560.
- [10] Hamilton, T.P.; Pulay, P. Direct inversion in the iterative subspace (DIIS) optimization of open-shell, excited-state, and small multiconfiguration SCF wave functions, *J. Chem. Phys.* **1985**, *84*, 5728–5734.
- [11] Sellers, H. The C²-DIIS Convergence Acceleration Algorithm, *Int. J. Quant. Chem.* **1993**, *45*, 31–41.
- [12] Kudin, K.N.; Scuseria, G.E.; Cancès, E. A black-box self-consistent field convergence algorithm: One step closer, *J. Chem. Phys.* **2002**, *116*, 8255–8261.
- [13] Karlström, G. Dynamical damping based on energy minimization for use ab initio molecular orbital SCF calculations, *Chem. Phys. Lett.* **1979**, *67*, 348–350.
- [14] Saunders, V.R.; Hillier, I.H. A “Level-Shifting” method for converging closed shell Hartree–Fock wave functions, *Int. J. Quant. Chem.* **1973**, *7*, 699–705.
- [15] Vacek, G.; Perry, J.K.; Langlois, J.M. Advanced initial-guess algorithm for self-consistent-field calculations on organometallic systems, *Chem. Phys. Lett.* **1999**, *310*, 189–194.
- [16] Lehtola, S. Assessment of initial guesses for self-consistent field calculations. Superposition of atomic potentials: Simple yet efficient, *J. Chem. Theory Comput.* **2019**, *15*, 1593–1604.
- [17] Bacskay, G.B. A quadratically convergent Hartree-Fock (QC-SCF) method. Application to closed shell systems, *Chem. Phys.* **1981**, *61*, 385–404.
- [18] Bacskay, G.B. A quadratically convergent Hartree-Fock (QC-SCF) method. Application to open shell orbital optimization and coupled perturbed Hartree-Fock calculations, *Chem. Phys.* **1982**, *65*, 383–396.
- [19] Fletcher, R. *Practical Methods of Optimization*, 2nd ed.; Wiley: New York, 1999.

- [20] Aidas, K.; Angeli, C.; Bak, K.L.; Bakken, V.; Bast, R.; Boman, L.; Christiansen, O.; Cimiraglia, R.; Coriani, S.; Dahle, P.; et al. The Dalton quantum chemistry program system, *WIREs Comput. Mol. Sci.* **2014**, *4*, 269–284.
- [21] Helmich-Paris, B. A trust-region augmented Hessian implementation for restricted and unrestricted Hartree–Fock and Kohn–Sham methods, *J. Chem. Phys.* **2021**, *154*, 164104.
- [22] Neese, F.; Wennmohs, F.; Becker, U.; Riplinger, C. The ORCA quantum chemistry program package, *J. Chem. Phys.* **2020**, *152*, 224108.
- [23] Frisch, M.J.; Trucks, G.W.; Schlegel, H.B.; Scuseria, G.E.; Robb, M.A.; Cheeseman, J.R.; Scalmani, G.; Barone, V.; Petersson, G.A.; Nakatsuji, H.; et al. Gaussian 16 Revision A.03, 2016. Gaussian Inc. Wallingford CT.
- [24] Fischer, T.H.; Almlöf, J. General methods for geometry and wave function optimization, *J. Phys. Chem.* **1992**, *96*, 9768–9774.
- [25] Chaban, G.; Schmidt, M.; Gordon, M. Approximate second order method for orbital optimization of SCF and MCSCF wavefunctions, *Theor. Chem. Acta* **1997**, *97*, 88–95.
- [26] Wong, A.T.; Harrison, R.J. Approaches to large-scale parallel self-consistent field calculations, *J. Comput. Chem.* **1995**, *16*, 1291–1300.
- [27] Karlström, G.; Lindh, R.; Malmqvist, P.Å.; Roos, B.O.; Ryde, U.; Veryazov, V.; Widmark, P.O.; Cossi, M.; Schimmelpfennig, B.; Neogrady, P.; et al. MOLCAS: A program package for computational chemistry, *Comput. Mat. Sci.* **2003**, *28*, 222–239.
- [28] Aquilante, F.; Autschbach, J.; Carlson, R.K.; Chibotaru, L.F.; Delcey, M.G.; De Vico, L.; Fdez. Galván, I.; Ferré, N.; Frutos, L.M.; Gagliardi, L.; et al. MOLCAS 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table, *J. Comput. Chem.* **2016**, *37*, 506–541.
- [29] Barca, G.M.J.; Bertoni, C.; Carrington, L.; Datta, D.; De Silva, N.; Deustua, J.E.; Fedorov, D.G.; Gour, J.R.; Gunina, A.O.; Guidez, E.; et al. Recent developments in the general atomic and molecular electronic structure system, *J. Chem. Phys.* **2020**, *152* (15), 154102.
- [30] Aprà, E.; Bylaska, E.J.; De Jong, W.A.; Govind, N.; Kowalski, K.; Straatsma, T.P.; Valiev, M.; Van Dam, H.J.; Alexeev, Y.; Anchell, J.; et al. NWChem: Past, present, and future, *J. Chem. Phys.* **2020**, *152*, 184102.
- [31] Stanton, J.F.; Gauss, J.; Cheng, L.; Harding, M.E.; Matthews, D.A.; Szalay, P.G. CFOUR, Coupled-Cluster techniques for Computational Chemistry, a quantum-chemical program package. With contributions from A.A. Auer, R.J. Bartlett, U. Benedikt, C. Berger, D.E. Bernholdt, S. Blaschke, Y. J. Bomble, S. Burger, O. Christiansen, D. Datta, F. Engel, R. Faber, J. Greiner, M. Heckert, O. Heun, M. Hilgenberg, C. Huber, T.-C. Jagau, D. Jonsson, J. Jusélius, T. Kirsch, K. Klein, G.M. Kopper, W.J. Lauderdale, F. Lipparini, T. Metzroth, L.A. Mück, D.P. O’Neill, T. Nottoli, D.R. Price, E. Prochnow, C. Puzzarini, K. Ruud, F. Schiffmann, W. Schwalbach, C. Simmons, S. Stopkowitz, A. Tajti, J. Vázquez, F. Wang, J.D. Watts and the integral packages MOLECULE (J. Almlöf and P.R. Taylor), PROPS (P.R. Taylor), ABACUS (T. Helgaker, H.J. Aa. Jensen, P. Jørgensen, and J. Olsen), and ECP routines by A. V. Mitin and C. van Wüllen. For the current version, see <http://www.cfour.de>, Last accessed 15 November 2020.
- [32] Matthews, D.A.; Cheng, L.; Harding, M.E.; Lipparini, F.; Stopkowitz, S.; Jagau, T.C.; Szalay, P.G.; Gauss, J.; Stanton, J.F. Coupled-cluster techniques for computational chemistry: The CFOUR program package, *J. Chem. Phys.* **2020**, *152*, 214108.
- [33] Beebe, N.H.F.; Linderberg, J. Simplifications in the generation and transformation of two-electron integrals in molecular calculations, *Int. J. Quantum Chem.* **1977**, *12*, 683–705.
- [34] Røeggen, I.; Wisløff-Nilssen, E. On the Beebe-Linderberg two-electron integral approximation, *Chem. Phys. Lett.* **1986**, *132*, 154–160.
- [35] Røeggen, I.; Johansen, T. Cholesky decomposition of the two-electron integral matrix in electronic structure calculations, *J. Chem. Phys.* **2008**, *128*, 194107.
- [36] Koch, H.; Sánchez De Merás, A.; Pedersen, T.B. Reduced scaling in electronic structure calculations using Cholesky decompositions, *J. Chem. Phys.* **2003**, *118*, 9481–9484.
- [37] Weigend, F.; Kattannek, M.; Ahlrichs, R. Approximated electron repulsion integrals: Cholesky decomposition versus resolution of the identity methods, *J. Chem. Phys.* **2009**, *130*, 164106.
- [38] Aquilante, F.; Boman, L.; Boström, J.; Koch, H.; Lindh, R.; de Merás, A.S.; Pedersen, T.B. *Cholesky decomposition techniques in electronic structure theory*, Zalesny, R., Papadopoulos, M.G., Mezey, P.G., Leszczynski, J., Eds.; Springer, Dordrecht, 2011; Vol. 13, pp 301–343.
- [39] Folkestad, S.D.; Kjøenstad, E.F.; Koch, H. An efficient algorithm for Cholesky decomposition of electron repulsion integrals, *J. Chem. Phys.* **2019**, *150*, 194112.

- [40] Nottoli, T.; Gauss, J.; Lipparini, F. A Second-order CASSCF algorithm with the Cholesky decomposition of the two-electron repulsion integrals, *J. Chem. Theory Comput.* **2021**, *xx*, submitted.
- [41] Burger, S.; Lipparini, F.; Gauss, J.; Stopkowitz, S. NMR chemical shift computations at second-order Møller-Plesset perturbation theory using gauge-including atomic orbitals and Cholesky-decomposed two-electron integrals, *J. Chem. Phys.* **2021**, *xx*, submitted.
- [42] Jensen, H.J.Aa.; Jørgensen, P. A direct approach to second-order MCSCF calculations using a norm extended optimization scheme, *J. Chem. Phys.* **1984**, *80*, 1204–1214.
- [43] Jensen, H.J.Aa.; Ågren, H. A Direct, restricted-step, second-order MC SCF program for large scale ab initio calculations, *Chem. Phys.* **1986**, *104*, 229–250.
- [44] Thouless, D. Stability conditions and nuclear rotations in the Hartree-Fock theory, *Nuc.Phys.* **1960**, *21*, 225–232.
- [45] Aquilante, F.; Pedersen, T.B.; Lindh, R. Low-cost evaluation of the exchange Fock matrix from Cholesky and density fitting representations of the electron repulsion integrals, *J. Chem. Phys.* **2007**, *126*, 194106.
- [46] te Velde, G.; Bickelhaupt, F.M.; Baerends, E.J.; Fonseca Guerra, C.; van Gisbergen, S.J.A.; Snijders, J.G.; Ziegler, T. Chemistry with ADF, *J. Comput. Chem.* **2001**, *22*, 931–967. <http://dx.doi.org/10.1002/jcc.1056>.
- [47] Kendall, R.A.; Dunning, T.H.; Harrison, R.J. Electron affinities of the first-row atoms revisited. Systematic basis sets and wave functions, *J. Chem. Phys.* **1992**, *96*, 6796–6806.
- [48] Dunning, Jr., T.H. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen, *J. Chem. Phys.* **1989**, *1007*, 4572–4585.
- [49] Binkley, J.; Pople, J.A.; Dobosh, P.A. The calculation of spin-restricted single-determinant wavefunctions, *Mol. Phys.* **1974**, *28*, 1423–1429.
- [50] Groenhof, A.R.; Swart, M.; L; Ehlers, A.W.; Lammertsma, K. Electronic ground states of iron porphyrin and of the first species in the catalytic reaction cycle of cytochrome P450s, *J. Phys. Chem. A* **2005**, *109*, 3411–3417.
- [51] Feng, X.; Epifanovsky, E.; Gauss, J.; Krylov, A.I. Implementation of analytic gradients for CCSD and EOM-CCSD using Cholesky decomposition of the electron-repulsion integrals and their derivatives: Theory and benchmarks, *J. Chem. Phys.* **2019**, *151*, 014110.
- [52] Becke, A.D. A new mixing of Hartree-Fock and local density-functional theories, *J. Chem. Phys.* **1993**, *98*, 1372–1377.
- [53] Hehre, W.J.; Ditchfield, R.; Pople, J.A. Self — Consistent molecular orbital methods . XII . Further extensions of Gaussian — type basis sets for use in molecular orbital studies of organic molecules, *J. Chem. Phys.* **1972**, *56*, 2257–2261.
- [54] Tsuchimochi, T.; Scuseria, G.E. Communication: ROHF theory made simple, *J. Chem. Phys.* **2010**, *133*, 141102.
- [55] Tsuchimochi, T.; Scuseria, G.E. Constrained active space unrestricted mean-field methods for controlling spin-contamination, *J. Chem. Phys.* **2011**, *134*, 064101.