Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations

Konstantinos D. Vogiatzis, Dongxia Ma, Jeppe Olsen, Laura Gagliardi, and Wibe A. de Jong

Citation: *The Journal of Chemical Physics* **147**, 184111 (2017); doi: 10.1063/1.4989858

View online: <https://doi.org/10.1063/1.4989858>

View Table of Contents: <http://aip.scitation.org/toc/jcp/147/18>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Single-reference coupled cluster theory for multi-reference problems](#)

The Journal of Chemical Physics **147**, 184101 (2017); 10.1063/1.5003128

[Tensor-structured coupled cluster theory](#)

The Journal of Chemical Physics **147**, 184113 (2017); 10.1063/1.4996988

[A note on the accuracy of KS-DFT densities](#)

The Journal of Chemical Physics **147**, 204103 (2017); 10.1063/1.5001939

[Incremental full configuration interaction](#)

The Journal of Chemical Physics **146**, 104102 (2017); 10.1063/1.4977727

[Non-orthogonal internally contracted multi-configurational perturbation theory \(NICPT\): Dynamic electron correlation for large, compact active spaces](#)

The Journal of Chemical Physics **147**, 174106 (2017); 10.1063/1.4999218

[Perspective: Multireference coupled cluster theories of dynamical electron correlation](#)

The Journal of Chemical Physics **149**, 030901 (2018); 10.1063/1.5039496

PHYSICS TODAY

WHITEPAPERS

ADVANCED LIGHT CURE ADHESIVES

Take a closer look at what these environmentally friendly adhesive systems can do

READ NOW

PRESENTED BY
MASTERBOND
ADHESIVES | SEALANTS | COATINGS

Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations

Konstantinos D. Vogiatzis,^{1,a),b)} Dongxia Ma,^{1,c)} Jeppe Olsen,² Laura Gagliardi,^{1,a)} and Wibe A. de Jong^{3,a)}

¹*Department of Chemistry, Minnesota Supercomputing Institute, and Chemical Theory Center, University of Minnesota, 207 Pleasant Street Southeast, Minneapolis, Minnesota 55455-0431, USA*

²*Department of Chemistry, Aarhus University, Langelandsgade 140, 8000 Aarhus C, Denmark*

³*Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

(Received 12 June 2017; accepted 20 October 2017; published online 14 November 2017)

A new large-scale parallel multiconfigurational self-consistent field (MCSCF) implementation in the open-source NWChem computational chemistry code is presented. The generalized active space approach is used to partition large configuration interaction (CI) vectors and generate a sufficient number of batches that can be distributed to the available cores. Massively parallel CI calculations with large active spaces can be performed. The new parallel MCSCF implementation is tested for the chromium trimer and for an active space of 20 electrons in 20 orbitals, which can now routinely be performed. Unprecedented CI calculations with an active space of 22 electrons in 22 orbitals for the pentacene systems were performed and a single CI iteration calculation with an active space of 24 electrons in 24 orbitals for the chromium tetramer was possible. The chromium tetramer corresponds to a CI expansion of one trillion Slater determinants (914 058 513 424) and is the largest conventional CI calculation attempted up to date. *Published by AIP Publishing.*
<https://doi.org/10.1063/1.4989858>

I. INTRODUCTION

The accurate calculation of near-degeneracy electron correlation effects for large orbital spaces is central in modern electronic structure theory. Many systems of interest cannot be quantitatively described by a single electronic configuration. Multireference effects, also referred to as static correlation, nondynamic correlation, left-right correlation, or strong correlation,^{1–3} can be captured by the full configuration interaction (full CI) expansion of the wave function. In full-CI theory, the wave function is a linear expansion of all the Slater Determinants (SDs) or spin-adapted configuration state functions (CSFs) that can be generated in a given one-electron basis. The exponential dependence of the number of SDs on the number of orbitals and electrons makes full-CI wave function applicable to only small- to modest-sized systems.

In multiconfigurational self-consistent field (MCSCF) theories, a full CI is employed on a selected orbital sub-space (*active* orbitals), while the remaining orbitals are kept either occupied (*inactive*) or empty (*virtual* or *secondary*), and orbitals are variationally optimized simultaneously with the configuration expansion coefficients.⁴ In recent MCSCF implementations, these two problems are usually decoupled and solved separately. In the inner loop (microiterations), the CI coefficients are optimized minimizing the energy. In outer

loops (macroiterations), the molecular orbitals are optimized by iteratively solving the Newton-Raphson equations using the first-order density matrix calculated from the CI expansion. A full-CI calculation in the active space is performed at every MCSCF iteration, and thus, considerable effort has been dedicated over the past 40 years to develop and implement efficient CI algorithms.^{4–8}

The CI-related methods have taken advantage of parallel architectures, and significant progress has been made in the last 20 years.^{4,9–16} However, new architectures with increased parallelism require algorithmic improvements of parallel CI and MCSCF implementations to take full advantage of these technological advances. Current parallel CI calculations are able to tackle expansions of a few billions of CSFs.⁴ To the best of our knowledge, the largest multireference CI (MRCI) calculation that has been reported is 2.8×10^9 CSFs (60×10^9 SDs),¹⁷ while the largest CI expansion in a full-CI calculation contains 10×10^9 determinants.⁹

Alternative expansions of the wave function have been proposed that allow access to larger active spaces while limiting the number of determinants in the CI expansion. The density-matrix renormalization group (DMRG)^{18–20} substitutes the exact diagonalization of large Hamiltonian matrices by encoding a sequential structure into the correlation. The DMRG wave function is built from local variational objects associated with the active orbitals of the system. The DMRG-SCF methodology allows the effective treatment of large molecular complexes and is gradually becoming a standard quantum chemical method.^{21,22} The variational two-electron reduced-density matrix (v2RDM) method and the corresponding v2RDM-CI and MCSCF variants have recently

^{a)}Authors to whom correspondence should be addressed: kvogiatz@utk.edu; gagliardi@umn.edu; and wadejong@lbl.gov

^{b)}Current address: Department of Chemistry, University of Tennessee, Knoxville, Tennessee 37996-1600, USA.

^{c)}Current address: Max Planck Institut für Festkörperforschung, Heisenbergstraße 1, 70569 Stuttgart, Germany.

been applied for solving strongly correlated systems.^{23–25} Stochastic approaches have been suggested as an efficient alternative to the standard Davidson CI eigensolver.^{26,27} The full-CI quantum Monte Carlo (FCIQMC)-MCSCF method has been applied to study transition metal complexes such as Fe porphyrins.^{28,29}

The restricted active space SCF (RASSCF),⁷ the generalized active space SCF (GASSCF),³⁰ and the occupation restricted multiple active space (ORMAS)³¹ methods provide a different approach to the reduction of the CI expansion by limiting the excitations within the active orbitals.^{7,30,32} In the GASSCF approach, multiple orbital spaces are chosen instead of one complete active space (CAS). The definition of the intra- and interspace electron excitations leads to an efficient elimination of negligible configurations from the configuration space, which effectively reduces the CI expansion. The spawning of multiple active spaces provides an approach to split the CI vector into smaller blocks (*vide infra*).^{7,33,34}

The most time-consuming step of a CI calculation is the construction of the σ vector needed in the Davidson algorithm. This step is also the most difficult part in terms of parallelization, as it is subject to load-balance, bandwidth, and memory constraints. In this study, we used the GAS framework for the development of a new implementation of a massively parallel MCSCF code. The new parallel MCSCF implementation, based on a serial version of the LUCIA³⁴ program, was efficiently parallelized and integrated into the NWChem program package.³⁵ This implementation allows us to perform large-scale CI calculations with a fast time-to-solution as well as allows us to explore active spaces beyond the limits of conventional MCSCF implementations.

The outline of this paper is as follows: In Sec. II, the foundations of MCSCF theory are discussed. In Sec. III, the technical aspects of the parallel MCSCF implementation are presented. The performance of this implementation for an active space with 20 electrons in 20 orbitals is presented in Sec. IV. In Sec. V, the applicability of the new parallel code to larger full-CI spaces and possible further improvements to the parallel performance are discussed. Finally, in Sec. VI we offer some conclusions.

II. THEORY AND METHODOLOGY

In CASSCF theory, the size of the CI expansion is dictated by the size of the complete active space or CAS. The choice of the number of orbitals and electrons that compose the active space is usually system dependent and is based on the nature of the chemical problem under consideration. The number of SDs included in the CI expansion scales exponentially with the size of the active space, and active spaces larger than 18 electrons in 18 orbitals cannot be currently treated.³⁶

Restricting excitations between orbitals in the generation of the CI expansion leads to the reduction of the number of SDs or CSFs that need to be considered. Such restrictions can usually be rationalized by the chemistry of the molecular system and are system dependent but can lead to a simplification of the CI problem without significant loss of accuracy.

A. Determinant-based direct-CI

In full-CI theory, and for a given one-electron expansion, the exact solution of the Schrödinger equation may be written as a linear combination of all SDs that can be constructed in the N -electron Fock space. The MCSCF wave function in which the SD basis is constructed from a subspace (active space) of the full Fock space is expressed as

$$|\Psi_{\text{MCSCF}}\rangle = \exp(-\hat{k}) \sum_i C_i |\Psi_i\rangle, \quad (1)$$

where i is the total number of SDs, C_i are the variational CI coefficients, and $\exp(-\hat{k})$ is the orbital-rotation operator. The CI eigenvalue problem can be solved with the direct-CI approach,³⁷ in which the expansion coefficients are computed in operator form *directly* from the one- and two-electron integrals within an iterative scheme. The handling of the SDs is simplified if each SD is represented as a product of an alpha and a beta string,³⁸

$$|\Psi_i\rangle = |\alpha(I_\alpha)\beta(I_\beta)\rangle = \hat{\alpha}(I_\alpha)\hat{\beta}(I_\beta)|\text{vac}\rangle, \quad (2)$$

where $\hat{\alpha}(I_\alpha)$ and $\hat{\beta}(I_\beta)$ are ordered products of alpha and beta creation operators, respectively, and $|\text{vac}\rangle$ is the vacuum state. The MCSCF wave function (or CI expansion) can be written as

$$|\Psi_{\text{MCSCF}}\rangle = \sum_{I_\alpha, I_\beta} C(I_\alpha, I_\beta) |\alpha(I_\alpha)\beta(I_\beta)\rangle. \quad (3)$$

We are using the *modified inverted-Davidson* algorithm of LUCIA.^{7,34} The Davidson eigensolver³⁹ iteratively diagonalizes a subspace instead of the full Hamiltonian matrix.

In a direct-CI iteration,⁴⁰ the main computational cost is the construction of the sigma vector

$$\sigma(I_\alpha, I_\beta) = \sum_{J_\alpha, J_\beta} \langle \beta(J_\beta) \alpha(J_\alpha) | \hat{H} | \alpha(I_\alpha) \beta(I_\beta) \rangle C(J_\alpha, J_\beta) \quad (4)$$

or, in a matrix notation,

$$\sigma = HC. \quad (5)$$

The non-relativistic Hamiltonian of Eqs. (4) and (5) is expressed as

$$\hat{H} = \sum_{kl} h_{kl} \hat{E}_{kl} + \frac{1}{2} \sum_{ij,kl} (ij|kl) (\hat{E}_{ij} \hat{E}_{kl} - \delta_{jk} \hat{E}_{il}), \quad (6)$$

where $\hat{E}_{kl}$ is the one-electron excitation operator,

$$\hat{E}_{kl} = a_{k\alpha}^\dagger a_{l\alpha} + a_{k\beta}^\dagger a_{l\beta}. \quad (7)$$

By inserting the Hamiltonian of Eq. (6) in Eq. (5), we can rewrite the sigma vector as a sum of three terms,

$$\sigma(I_\alpha, I_\beta) = \sigma_1(I_\alpha, I_\beta) + \sigma_2(I_\alpha, I_\beta) + \sigma_3(I_\alpha, I_\beta), \quad (8)$$

where σ_1 is a column vector with only beta-beta contributions ($I_\alpha = J_\alpha$), σ_2 is a column vector with only alpha-alpha contributions ($I_\beta = J_\beta$), and σ_3 includes the alpha/beta couplings. For more details on the form of the three sigma vector terms and on the efficiency that this splitting introduces, see Ref. 7.

B. Orbital optimization

A second-order Newton-Raphson iterative procedure is applied for the orbital optimization step of the MCSCF macro-iteration, as implemented in LUCIA, and the variational orbital parameters of the vector κ are calculated according to Eq. (1). The energy can be expressed as

$$E(\kappa) = E(0) + \kappa \mathbf{g} + \frac{1}{2} \kappa^2 \mathbf{H}, \quad (9)$$

where $\mathbf{g}$ and $\mathbf{H}$ are the orbital gradient and orbital Hessian, respectively. The stationary points are obtained as solutions to the equation $\partial E / \partial \kappa_i = 0$. Orbital rotations between the inactive-active and active-virtual orbital spaces are allowed. The full orbital-orbital Hessian is used without any approximations.

III. IMPLEMENTATION

The Global Arrays Toolkit⁴¹ was used to facilitate the parallelization of LUCIA. The global arrays set of tools was co-developed with NWChem as a shared-memory programming interface for distributed data algorithms relevant to the field of computational chemistry. They allow ease of programming and lack of synchronization between cores, considering that the nonlocal data take more time to access and offer support for both task and data parallelism. The potentially very large CI and sigma vectors and smaller Fock matrices utilized in the LUCIA CI and MCSCF code are stored in global arrays that are distributed over the memory available on the allocated cores. A maximum of three vectors are stored in memory, whereas additional vectors needed in the Davidson iterative scheme are stored using parallel IO (ParIO) within the native framework in the Global Arrays Toolkit. LUCIA's local data memory registration routine was modified to utilize NWChem's memory allocation process.

An MCSCF calculation can be divided into four tasks: (1) the generation of the CI expansion, (2) the partial atomic orbital (AO)–molecular orbital (MO) integral transformation, (3) the CI eigenvalue problem, and (4) the solution of the second-order Newton-Raphson equations for the orbital optimization step. For small CI expansions, the AO-MO integral transformation is the most CPU-time demanding step. For large CI expansions, the CI eigensolver (usually the Davidson algorithm via the direct-CI method⁴⁰) dominates the computational time of the MCSCF calculation. Prior to this work, the limitation for the selected active space was 18 electrons in 18 orbitals for a singlet ($S = 0$) spin state. The implementation details about the parallelization of each of the four tasks are discussed in Secs. III A–III E.

A. Generation of CI expansion for parallel processing

The main strategy in the parallelization of the CI algorithm, to enable MCSCF calculations with large active spaces, is to distribute CI and sigma vectors into batches. The batches can be subsequently assigned as parallel tasks to different cores to obtain a good load-balance of work among the cores. With each core only having to store a subsection of the CI and sigma vectors, the memory footprint is significantly reduced.

For simplicity, NWChem's LUCIA version will store the CI and sigma vectors as SDs instead of the CSFs used by other codes.

The algorithm presented in this work is based on the GAS concept. In the GASSCF method, an arbitrary number of active spaces (GASs) are defined. Within each GAS, a full-CI expansion is considered, while intra-GAS excitations are limited and defined by the user. An example of intra-GAS excitations for a GASSCF with 4 GASs is the following: one-electron excitations between GAS 1 and GAS 2, no excitations from GAS 2 to GAS 3, and two-electron excitations between GAS 3 and GAS 4. The definition of the GASs and the intra-GAS excitations is based on physical criteria. The aim of the GASSCF is the reduction of the configurations included in the CI expansion. The GAS distribution of a parent, large CAS is used in our parallel implementation, but without applying any restrictions on the intra-GAS excitations. This means that no truncation of the CI problem is implicitly considered. The benefits of this approach are discussed in the following paragraphs.

LUCIA organizes the alpha and beta strings of Eq. (2) into blocks of SDs with the same occupation type (T) and spin symmetry (S). The occupation type is defined according to the distribution of electrons in each GAS. Spatial symmetry and spin symmetry are defined according to the occupation of the spin orbitals in each string. Therefore, each alpha string has a specific TS value,

$$\begin{aligned} \hat{\alpha}(I_\alpha) &= \prod_{i=1}^N \hat{\alpha}(I_{i\alpha}) = \prod_{i=1}^N \hat{\alpha}(I_{i\alpha}^{T_i S_i}) \\ &= \hat{\alpha}(I_\alpha^{T_1 T_2 \dots T_N S_1 S_2 \dots S_N}) = \hat{\alpha}(I_\alpha^{TS}), \end{aligned} \quad (10)$$

where N is the number of GASs. A similar expression holds for a beta string. Combinations of alpha and beta TS strings generate SDs with a specific TTSS definition. SDs with the same TTSS definition are grouped into TTSS blocks. The TTSS blocks are further grouped into TTSS batches.

By default, LUCIA generates only a limited number of batches, which restricts the number of cores among which the workload can be distributed. The approach used here to increase the number of TTSS blocks and batches, and subsequently the number of tasks available for parallel processing, is to fragment a parent (complete or general) active space and distribute the orbitals over additional GASs. The GAS partitioning approach within LUCIA offers an inherent block distribution of both the CI vector and the Hamiltonian to increase the number of tasks that can be distributed over the available cores. The technique is also used in serial calculations, to reduce the batch size and memory footprint.³⁴ The approach is schematically depicted in Fig. 1.

The algorithm that generates multiple TTSS blocks is based on the different distribution of electrons of a spin alpha or beta in different GASs (types, T). The number of possible types T that an alpha or a beta string can belong to is increased when additional GASs are introduced. This is shown schematically in Table I and analyzed with an example in the Appendix. The CI vector is generated by combining all possible alpha and beta strings. Only SDs with the correct spin (SS)

FIG. 1. Schematic representation of the distribution of the CI vector. Fragmentation of the CI vector (upper line) into a limited number of TTSS blocks for a given CAS (second line). Partitioning of the parent CAS into multiple GASs generates more TTSS blocks (third line), which can be grouped into TTSS batches. The TTSS batches are distributed into the available cores (bottom line).

are included in the CI expansion and they are grouped into different TTSS blocks based on their alpha and beta occupation type (TT).

Dividing the parent active space across multiple GASs can be accomplished in many ways. In the current implementation, two strategies have been automated, which are represented in Table II. The first strategy is to distribute the MOs of each irreducible representation into a separate GAS (Table II A). This will increase the number of TTSS batches but still has limitations, for example, when limited point group symmetry is available. The second strategy is an iterative scheme where one MO is moved into a new GAS (Table II B). We elected to move an MO from the irreducible representation with the largest number of orbitals. This process is repeated until the number of TTSS batches is larger than the number of cores and a distribution of the CI tasks is feasible. This approach has the additional advantage that it generates smaller TTSS batches that can easily fit in local memory, something that will be crucial for the execution of CI and MCSCF calculations with more than 60×10^9 SDs. It should be noted that these strategies still have little control over the size and computational intensity of the batches that are created, which will affect our ability to effectively load-balance the work over the cores.

Having divided the CI expansion into a large number of TTSS batches, the next step is to distribute the batches over the allocated cores and ensure that the computational work of each core is balanced and maximum parallelization is achieved. In the design of the parallel algorithm, the choice was made to

TABLE II. Representative examples of possible approaches to CAS(20,20) distribution over multiple GASs. (A) Each GAS holds all orbitals from one or more irreducible representations, generating 888 TTSS batches. (B) Individual orbitals from irreducible representations are stored in different GASs, generating 14 502 TTSS batches.

(A)	Irreducible representation							
	a_g	b_{3u}	b_{2u}	b_{1g}	b_{1u}	b_{2g}	b_{3g}	a_u
GAS 1	6	0	0	0	0	0	0	0
GAS 2	0	1	1	2	0	0	0	0
GAS 3	0	0	0	0	5	0	0	0
GAS 4	0	0	0	0	0	2	0	0
GAS 5	0	0	0	0	0	0	2	0
GAS 6	0	0	0	0	0	0	0	1
(B)								
GAS 1	3	1	1	2	2	2	2	1
GAS 2	1	0	0	0	0	0	0	0
GAS 3	0	0	0	0	1	0	0	0
GAS 4	1	0	0	0	0	0	0	0
GAS 5	0	0	0	0	1	0	0	0
GAS 6	1	0	0	0	0	0	0	0
GAS 7	0	0	0	0	1	0	0	0

only allow each core to access to the data of the sigma batches locally, while CI batches needed in the calculation were fetched using one-sided get operations.⁴¹ To ensure memory locality for the sigma vector, the allocation of the global arrays containing the CI expansion and sigma vectors is done such that no TTSS batch is distributed over the memory of multiple cores. The distribution of the TTSS batches over the cores can be non-uniform, depending on the computational complexity. This will be discussed later in this section. One-sided get operations enable access to, and retrieval of data from, memory of cores physically located on different nodes through remote direct memory access (RDMA) protocols. These get (fetch) operations do not interfere with the running process on the remote core and are also available in the new MPI-3 standard. The Global Arrays Toolkit can also perform non-blocking get calls, where the (pre-)fetch essentially gets posted and the core continues its compute and returns to the get operation when the data are needed. This would require the allocation of extra memory buffers to retrieve these data. The get operations were found to provide limited overhead, and therefore the pre-fetch approach was not implemented. All the computational work associated with a TTSS sigma batch is assigned to the core where the data reside. This approach significantly reduces the overall communication volume needed but does

TABLE I. Schematic representation of the distribution of alpha (or beta) strings in different types T by increasing the number of GASs. Each symbol corresponds to a different alpha (or beta) string, and different colors represent different supergroups.

Irreducible representation	One GAS	Two GASs	Three GASs
1	● ● ● ●	● ● ● ●	● ● ● ●
2	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●
3	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●
⋮	⋮	⋮	⋮

require the computed batches to be statically distributed at the beginning of the CI calculation. An alternative approach would be the use of a global task pool, requiring communication of both CI and sigma data blocks if non-locality is assumed. This alternative approach was implemented in an earlier prototype and demonstrated considerably poorer parallel performance than the current algorithm, primarily because of memory contention with many processors accumulating data into the same sigma data block, effectively serializing this operation.

The order of the TTSS batches in the CI vector is fixed, and it is conceivable that batches have drastically different computational time requirements [see, for example, Fig. 2(a)] and that computationally expensive batches are located at the beginning or end of the CI vector. One would prefer the number of parallel tasks to be many orders of magnitude larger than the number of cores to balance out the irregular batch sizes and associated computational work. The large number of GASs needed to create, for example, a batch to core ratio of 100:1, themselves generate a significant computational overhead, as will be shown in Sec. IV. Hence, the challenge in achieving parallel efficiency is finding the optimum balance between the number of GASs and parallel tasks/batches needed, given the number of cores.

The work associated with each TTSS batch needs to be estimated accurately and enough batches, i.e., parallel tasks, need to be available to ensure a balanced workload. Three intrinsic parameters of the TTSS (or CI) batches were initially considered to estimate the computational work, and they are plotted versus the timings of the corresponding sigma calculation tasks in Fig. 2. These three parameters are the number of SDs or size of the TTSS batch [Fig. 2(a)], the number of TTSS blocks per batch [Fig. 2(b)], and the number of connections between the CI vector and the Hamiltonian [Fig. 2(c)]. The connectivity between the CI vector and the Hamiltonian is defined by the electron difference between a TTSS block

(subgroup of a TTSS batch) and the Hamiltonian. The diagonal elements differ by zero electrons, while the off-diagonal elements can differ by one or more electrons. The cases that differ by more than two electrons do not couple, and they are not included in the connectivity calculation. All data of Fig. 2 were obtained from a CI expansion that includes 20 electrons in 20 orbitals or CAS(20,20) in a short-hand notation. The size of the CI vector is 2 133 595 282 SDs (linear Cr₃ molecule, *D*_{2h} symmetry with 8 irreducible representations, singlet A_g state). It is evident from Fig. 2 that there is no correlation between the size of a TTSS batch and the task time [Fig. 2(a)]. The time spent for batches of maximum size ranges from a few seconds to the maximum sigma TTSS batch task time (~1400 s). On the other hand, a reasonable correlation is observed between the number of blocks [Fig. 2(b)]/number of connections [Fig. 2(c)] and the timings of the sigma tasks.

Knecht *et al.*¹⁴ developed a computational work estimator, or weight factor, that is based on the connectivity and the size of the individual TTSS blocks. For the *i*th TTSS batch, the weight factor is calculated as

$$a(i) = \sum_j^{N_{\text{blocks}}(i)} c_j(i) l_j(i), \quad (11)$$

where $N_{\text{blocks}}(i)$ is the number of TTSS blocks of the *i*th TTSS batch, $c_j(i)$ is the number of connections of the *j*th TTSS block that belongs to the *i*th TTSS batch, and $l_j(i)$ is the number of SDs of the *j*th TTSS block. For the CAS(20,20) test case, no good correlation was found between the $a(i)$ weight factors and the sigma task timings [Fig. 2(d)]. In an attempt to better capture some of the outliers seen in the correlation graphs for the number of blocks and connections, an alternative weight factor $b(i)$ was considered, combining the connectivity and the number of TTSS blocks $N_{\text{blocks}}(i)$ in a batch. The factor b_i is defined as

FIG. 2. Upper line: Scatter plots between the timings of the sigma tasks for each batch and three parameters of the CI batches: (a) batch size in SDs, (b) number of blocks, and (c) number of connections. Lower line: scatter plots between two different TTSS weight factors: (d) Knecht *et al.*¹⁴ and (e) this work. All y axis values are normalized to 1. All data are obtained from a CI expansion of an active space that includes 20 electrons in 20 orbitals (2 133 595 282 SDs, 8 irreducible representations, singlet state).

$$b(i) = \left(\sum_j^{N_{\text{blocks}}(i)} c_j(i) \right) N_{\text{blocks}}(i). \quad (12)$$

The correlation of this weight with the sigma task timings [Fig. 2(e)] is similar to that of the number of blocks and number of connections separately though a few outliers now have weight factors that better reflect their computational workload.

An important step for the distribution of SDs to TTSS blocks is the organization of alpha/beta superstrings into occupation classes. This step involves a double loop over all alpha and beta occupation types (supergroups), where the full-CI vector is read for each alpha/beta superstring combination. No significant time is spent for a small number of GASs (i.e., small number of alpha/beta occupation types) or small CI expansions ($<10^8$). However, significantly more time is spent when multiple GASs (typically more than 6) or large active spaces are employed which will be discussed in Sec. IV. To reduce the amount of time spent in this generation, the double loop was decoupled and effectively parallelized by distributing the alpha occupation types over the available cores. The occupation alpha/beta connection map is built locally on each core and is stored in a global array using allocated, but not yet utilized, memory for storing the CI and sigma vectors. The connection map is updated, and in a second parallel loop, the alpha/beta supergroup combinations are sorted into occupation classes.

For wave functions with vanishing spin-projection and well-defined spin, the CI vectors are either symmetric or anti-symmetric under transposition, $C(I_\alpha, I_\beta) = -1^S C(I_\beta, I_\alpha)$, where S is the total spin. This may be used to reduce the storage and computational effort by a factor of two as originally suggested for singlet FCI states by Knowles and Handy.³⁸ In the present context, this transpositional symmetry relates TTSS blocks that have the type and symmetry of the alpha- and beta-strings permuted. When the type and symmetry of the alpha- and beta-strings differ, this symmetry relates two different TTSS blocks, allowing us to evaluate and store only one of these. When the type and symmetry of the alpha- and beta-strings are identical, the symmetry is used to evaluate and store only the lower half of the elements in this matrix-block.

B. The direct-CI algorithm

The evaluation of the three parts of the sigma vector in Eq. (8) constitutes the inner part of the CI iterations, and an efficient evaluation of these terms is central for the overall performance of the code. We will now briefly discuss the most important aspects of this evaluation. The point of departure is within the loops over TTSS blocks of the sigma and CI-vectors discussed in Sec. III A, so the computational task is the evaluation of the contribution from a given TTSS block of a CI-vector, $\mathbf{C}$, to a given TTSS block of the sigma vector. The code first examines whether the two TTSS blocks are connected by at most a double spin-orbital excitation. Provided that the two TTSS blocks differ by at most a double excitation, the code uses the occupation of the two TTSS blocks to select the algorithm that will be most efficient for this particular task. The code thus has several algorithms available for the construction of the various terms of the sigma vector of

Eq. (8). Furthermore, a given algorithm has several variants, as the code may use particle-hole rearrangements of the Hamiltonian operator and transpose the blocks of σ and $\mathbf{C}$ to reduce the computational complexity. We will here focus on the generally used algorithm for the evaluation of σ_3 , which is the most time-consuming part of the generation of the sigma vector. The algorithm is based on the introduction of the resolution of the identity to write this contribution as

$$\sigma_3(I_\alpha, I_\beta) = \sum_{J_\alpha J_\beta, i, j, k, l} \langle I_\alpha | a_{i\alpha}^\dagger \sum_{K_\alpha} | K_\alpha \rangle \langle K_\alpha | a_{j\alpha} | J_\alpha \rangle \times \langle I_\beta | a_{k\beta}^\dagger a_{l\beta} | J_\beta \rangle (ijkl) C(J_\alpha, J_\beta), \quad (13)$$

where the sums over i, j, k, l are over active orbitals and the resolution of the identity is over alpha-strings with one alpha-electron less than that in the reference state. This resolution of the identity is now used to organize the calculation as follows:

1. Evaluate

$$D(K_\alpha, j, J_\beta) = \sum_{I_\alpha} \langle K_\alpha | a_{j\alpha} | J_\beta \rangle C(I_\alpha, J_\beta). \quad (14)$$

2. Evaluate

$$T(K_\alpha, i, I_\beta) = \sum_{k, l, J_\beta} \langle I_\beta | a_{k\beta}^\dagger a_{l\beta} | J_\beta \rangle \left(\sum_j (ijkl) D(K_\alpha, j, J_\beta) \right). \quad (15)$$

3. Update

$$\sigma_3(I_\alpha, I_\beta) = \sigma_3(I_\alpha, I_\beta) + \sum_i \langle I_\alpha | a_{i\alpha}^\dagger | K_\alpha \rangle T(K_\alpha, i, I_\beta). \quad (16)$$

The program chooses between several variants of the above algorithm. As presented above, the algorithm uses a resolution of the identity in the space of alpha-strings. For some forms of the TTSS blocks, the calculation is faster if the resolution of the identity is over the beta-strings, so the code checks and compares the operation counts for these two choices to select the optimal route. As mentioned above, the code may also use particle-hole reorganizations of the second-quantized Hamiltonian of Eq. (6) so that it is not necessarily the annihilation operators that occur to the right in the operator. In this case, the resolution of the identity may be in a space with an additional electron, rather than with one electron less. The time-defining step of the algorithm is the matrix-multiplication in Eq. (15), which is performed over the orbital-index j to obtain the $\mathbf{T}$ from the integrals and $\mathbf{D}$. The advantages of the algorithm are as follows: (1) the overhead is minimal as the constructions of $\mathbf{D}$ from $\mathbf{C}$ in Eq. (14) and σ_3 from $\mathbf{T}$ in Eq. (16) scale as the number of active orbitals times the number of determinants, (2) the calculations may be blocked in batches of limited cost, and (3) the inner loop is a matrix multiplication. A less appealing aspect of the presented algorithm is that the matrix multiplication has a single orbital index as the inner loop, so for small active spaces with symmetry, this loop may be rather small.

The parallel distribution of tasks in the generation of the sigma vector proceeds as follows: each node receives a replicated copy of the one- and two-electron integrals, while the

large sigma and CI vectors are stored in global arrays and they are distributed as described in Sec. III A. To construct its part of the sigma vector, each core only works on the blocks of the sigma vector it has in local memory, with the interacting blocks of the CI vector fetched from the memory of remote cores using one-sided get operations and combined with the replicated integrals available locally on each node. Only non-zero sigma blocks are processed, i.e., the corresponding Hamiltonian and/or the CI coefficients are non-zero, which avoids spending computation time on redundant tasks.

C. Partial AO-MO integral transformation

The one- and two-electron integrals and the 4-index AO-MO transformation of the two-electron integrals are performed with the optimized parallel subroutines of NWChem.³⁵ The two-electron integrals in the MO basis are subsequently reordered to the minimal integral list that LUCIA needs, and the full set of integrals is replicated across the nodes used in the calculation. This reordering step has also been parallelized in the implementation. The choice to replicate the full set of integrals is driven by the inherent random single integral element access nature of the underlying algorithm. While each CI step may only need a subset of integrals, the full set is stored to allow for an easy 4-index transformation of the integrals needed in each MCSCF step after the molecular orbitals are rotated.

D. CI eigensolver

The Davidson algorithm is an efficient procedure for the calculation of the lowest few eigenvalues by iteratively diagonalizing a subspace of a large sparse matrix. The steps that are followed are (1) the construction and diagonalization of an initial subspace matrix, which involves the calculation of an initial sigma vector as outlined in Sec. III B, (2) calculation of the preconditioner, (3) multiplication with the inverse Hessian and diagonalization to all previous vectors, and (4) construction of the new sigma vector and diagonalization of the updated projected matrix. If convergence is not reached, return to step 2. The update of the Davidson subspace involves vector operations that are naturally parallelized along the distributed CI and sigma vectors within the global array framework. As stated earlier, only a subset of the CI and sigma vectors is kept in memory. The additional vectors are stored on a disk using the parallel I/O (ParIO) tools of the Global Arrays Toolkit.

E. Orbital optimization step and outer MCSCF iteration

The underlying vector-vector and vector-matrix multiplications of the orbital optimization and the kappa update [Eq. (9)] are parallelized via the NWChem tools and take little time in the overall MCSCF calculation. As already mentioned in Sec. III C, each MCSCF iteration requires a transformation of the one- and two-electron integrals to the new MO basis. Each node does this transformation locally based on a transformation matrix describing the rotation from the previous iteration MOs to the current MOs. All necessary Fock matrices needed are stored globally in a global array and matrix elements are updated by the core that has them in local memory, thereby minimizing any communication.

F. Memory requirements

Memory is the main driving force of the new MCSCF implementation for performing calculations with large active spaces, where the storage of large CI and sigma vectors is needed. A larger number of cores increase the total allocated memory available for a specific calculation. Therefore, memory requirements are decreased with increasing number of cores and are roughly estimated from the ratio (number of SDs)/cores. Memory allocated for integrals is constant per core since they are stored in local memory on each core. Eventually, this could become a bottleneck for larger molecules or basis sets. This issue can be resolved by storing integrals once per node using the global arrays framework with processor groups.

IV. PARALLEL PERFORMANCE

The linear chromium trimer, Cr₃, was used to assess the performance of the new parallel MCSCF implementation. The Cr–Cr distance was arbitrarily set to 1.5 Å, a singlet ground state was computed, and the 6-31G* basis set was used to describe the molecular orbitals. Larger basis sets increase both the local memory requirements and the computational time of the CI microiterations. A systematic examination of the basis set size effect is beyond the scope of this work. The CI expansion was constructed using an active space of 20 electrons in 20 orbitals, CAS(20,20). An MCSCF calculation of this size is computationally not feasible for a serial code and has very large memory requirements. All 3d4s orbitals of the three chromium atoms were included in the active space, augmented with one occupied and one unoccupied molecular orbital to obtain the CAS(20,20) target active space. The choice of the correlated orbitals inside the active space is somehow arbitrary since these calculations were performed only for demonstrating the parallel scaling performance of the MCSCF implementation. All calculations were performed with the highest abelian point group (D_{2h}), which limits the number of SDs to about 4×10^9 (4 267 005 808). Within this symmetry, the 20 CAS orbitals are distributed over the irreducible representations as 6 a_g , 1 b_{3u} , 1 b_{2u} , 2 b_{1g} , 5 b_{1u} , 2 b_{2g} , 2 b_{3g} , and 1 a_u . In the benchmark results, the timings of a single MCSCF macroiteration with 20 Davidson CI microiterations and one extra CI iteration, used in creating the first-order density matrix, are reported. All calculations were performed using the Intel Haswell nodes (processors) on the Cori supercomputer located at the National Energy Research Scientific Computing Center (NERSC). Each node has 128 GB of memory and two 16-core Haswell processors running at 2.3 GHz, for a total of 32 cores per node. The developed version of NWChem was compiled with the Intel 16.0 compiler version, and the Global Arrays Toolkit delivering the parallel infrastructure was compiled with the MPI-PR setting. For each calculation, 17 cores were used per processor, 16 for the computation, and one to support the MPI communication.

Table III and Fig. 3 show the results for the CAS(20,20) case with 6 GASs from 32 to 512 cores. These extra GASs were created by moving all orbitals from one irreducible representation into a different GAS (as demonstrated in Table II A), which resulted in 888 TTSS batches available for parallel

TABLE III. Individual times (in sec) of the MCSCF steps for one macroiteration with CAS(20,20) for the chromium trimer. The CAS(20,20) is divided into 6 GASs moving all orbitals belonging to one irreducible representation into a different GAS (as demonstrated in Table II A).

Number of cores	32	64	128	256	512
Time (s)					
CI generation	56	38	24	18	12
MO-AO transformation	24	18	12	8	4
Integral evaluation	132	80	43	23	18
CI eigensolver (20 iterations)	30 026	17 040	10 331	5 815	4 285
1e/2e density matrices	3 897	2 219	1 442	1 062	856
Total time per MCSCF iteration	34 135	19 395	11 852	6 926	5 175
Parallel speedup	×1	×1.8	×2.9	×4.9	×6.6

processing. The introduction of GASs does not truncate the CI expansion; it only generates TTSS batches that can be efficiently distributed in different cores, as described in Sec. III. Each step discussed in Table III is performed once in a MCSCF macroiteration, which involves 20 CI microiterations.

As stated before, the most time-consuming step (over 85% of the computational time) is the CI eigensolver that constructs the sigma vector in the iterative Davidson algorithm. Reasonable parallel performance was obtained scaling the CAS(20,20) calculation to 512 cores, with a speedup of 6 going from 32 to 512 cores. No perfect load-balance was achieved, which can be attributed to the small set of 888 TTSS batches or parallel tasks, each of different size and computational complexity. To increase the number of TTSS batches that can be processed in parallel, and allow for better static load balancing, additional GASs or different orbital distributions need to be introduced. In Table IV, the timing data of the CI eigensolver component for the 6 GASs are compared to those of calculations using 8 GASs, where each GAS contains all orbitals of one irreducible representation. The two additional GASs, each containing one orbital, generate few additional batches, and these few additional batches come at a significant computational cost. There are several factors contributing to the large increase in the computational timings when the number of GASs is increased. First of all, the inner loops of the calculation of the sigma vector are organized as loops over

TABLE IV. Comparison of CI eigensolver timings (in sec) in one macroiteration with a CAS(20,20) for the chromium trimer using different distributions of orbitals and GASs.

Number of cores	32	64	128	256	512
6 GASs (888 batches)	30 026	17 040	10 331	5 815	4 285
8 GASs (896 batches)	42 644	22 764	13 559	8 736	4 752

types of double excitations where a type specifies the four GASs of the double excitation. By increasing the number of GASs from 6 to 8, the number of times these loops are executed increases from 1296 to 4096. Although the number of operations executed in a given pass through the loop is reduced, the total number of operations is increased significantly. Furthermore, the use of GASs with one or very few orbitals per symmetry leads to rather inefficient matrix multiplication due to matrices with one very small dimension, being equal to the number of orbitals or orbital pairs with given symmetry and GAS.

The approach of placing orbitals of each irreducible representation in a different GAS will be able to generate at most 896 TTSS batches for the CAS(20,20) calculation considered here. As such there are not enough batches and tasks to utilize more than 896 processes. Note that if lower symmetry would be applied, the ability to generate batches for parallel processing would be limited even further. The alternative approach that was utilized is to iteratively generate new GASs with only a single orbital until the number of TTSS batches significantly exceeds the number of cores available for the calculation, i.e., distributing orbitals of a single irreducible representation over multiple GASs as depicted in Table II B. The one orbital per GAS approach enabled calculations beyond 512 cores. In Table V, the timings of calculations up to 2048 cores are presented.

Table V shows that while for smaller core counts, the one orbital per new GAS approach increases the computational cost, it does enable the overall algorithm to achieve reasonable speedup all the way to 2048 cores. Therefore, the encoded GAS distribution approach first attempts to distribute the orbitals in each irreducible representation into different GASs, and only when it still does not have enough work for all available cores, it switches to distributing one orbital into a different GAS.

FIG. 3. Total CPU time (in sec) and the individual contributions of one MCSCF iteration for the chromium trimer with an active space of 20 electrons in 20 orbitals and with different number of cores. The CAS(20,20) is distributed over 6 GASs by moving all orbitals from one or more irreducible representations (see Table II A as an example).

TABLE V. Individual times (in sec) of the MCSCF steps for one macroiteration with a CAS(20,20) for the chromium trimer. The number of GASs is also given for all cases.

Number of cores	64	128	256	512	1024	1024	2048
Number of GASs	4	4	5	5	6	7	6
Number of TTSS batches	288	288	1076	1076	3963	14502	3963
CI generation	2	2	4	4	29	114	27
MO-AO transformation	19	13	8	5	2	4	1
Integral evaluation	4	4	8	8	15	34	14
CI eigensolver	40 967	26 478	14 131	8 065	4 873	6 059	2 868
1e/2e density matrices	3 947	2 051	1 844	941	803	872	555
Total time per MCSCF iteration	44 939	28 548	15 995	9 023	5 722	7 083	3 465

V. EXPLORING CI EXPANSIONS BEYOND CAS(20,20)

Additional test calculations were performed on larger CI expansions created by a CAS(22,22) and a CAS(24,24) to get insight into the performance of the parallel MCSCF code and to assess the feasibility of these classes of calculations. The test calculations were run on 4096 and 8192 cores of the NERSC Cori machine. Only a single CI iteration was performed, which provides the essential information about the building of the most computationally expensive sigma vector in the new parallel CI eigensolver. To ensure maximum connectivity between sigma and CI vectors and the computational work, representative of one of the last iterations in a CI calculation without zero-valued CI blocks that would be skipped, the coefficients of all SDs in the CI expansion were given a normalized value of $1.0/(\text{number of SDs})$. Given potential sparsity in the CI wave function, the calculated results are the upper bound for the computational time needed.

CASSCF calculations were performed on linear tetracene and pentacene, which are relevant species to design organic semiconductors.^{23,42–44} For tetracene, a full-CI expansion that includes the complete $\pi\pi^*$ system corresponds to a CAS(18,18). A $\pi\pi^*$ complete active space in pentacene corresponds to 22 electrons in 22 orbitals and a CI of 497 634 306 624 SDs. By applying symmetry restrictions (D_{2h} point group), $\pi\pi^*$ are transformed only by four different irreducible representations, and the 1A_g singlet ground state would require a CI expansion of 124 408 640 160 SDs. The number of SDs is about $30\times$ larger than the CAS(20,20) used in the performance benchmarks in Sec. IV. Our initial approach of moving all orbitals of a given irreducible representation into their own GAS did not create enough TTSS batches relative to the number of cores we intended to use for the calculation. Hence, we followed the approach discussed at the end of Sec. IV even though this will result in a significant computational overhead. For a calculation with 4096 cores, 7 GASs were needed to generate 7788 TTSS batches, providing each core with one or two batches to compute. Olsen and co-workers showed that the CI algorithm used here scales approximately linearly with the number of SDs,³⁴ which provides an excellent measure to understand possible performance degradation. The pentacene CAS(22,22) single CI iteration required 4212 seconds on 4096 cores. An additional eighth GAS was needed to generate enough TTSS batches

(28 607) that would utilize all 8192 cores. Only a slight reduction of time was achieved by doubling the number of cores, with the CI iteration requiring 4085 s. This means that a full MCSCF calculation with 10 macroiterations, each with 20 CI iterations, would still require 9 days of computational time.

The chromium dimer and trimer structures are typical benchmark systems for new electronic structure theory methods. It is generally accepted that all molecular orbitals composed by the 3d and 4s atomic orbitals of the chromium atoms should be included in the active space of MCSCF calculations for Cr_2 and Cr_3 . This gives rise to active spaces of CAS(12,12) and CAS(18,18) sizes, respectively. The chromium tetramer MCSCF with all 3d 4s atomic orbitals included would require a CAS(24,24) active space consisting of 7.3 trillion SDs (7 312 459 672 336 SDs), reduced by D_{2h} symmetry and a 1A_g ground state to almost one trillion SDs (914 058 513 424 SDs). Distributing the 24 orbitals in 8 GASs generates 57 538 TTSS batches, with each core receiving seven or eight batches. A single CAS(24,24) CI iteration on 8192 cores required 50 136 s. While this setup would allow the use of 20 000–30 000 cores, clearly, a full iterative MCSCF calculation at this scale is not yet feasible, unless the overall performance can be significantly improved.

VI. SUMMARY AND CONCLUSION

With the parallel implementation presented in this work, MCSCF calculations with active spaces composed by 20 electrons in 20 orbitals can be performed routinely in small computer clusters, while calculations with almost one trillion SDs can be executed in supercomputers with thousands of cores. Such large active spaces and CI expansions are beyond the limits of current MCSCF implementations. The methodology presented in this work is based on (1) the application of the native NWChem tools for the parallelization of the MCSCF steps and (2) the generation and distribution of computational work in assembling the sigma vector of the Davidson algorithm on available cores. Fragmentation of a parent active space (CAS or GAS) into multiple GASs was used to generate enough CI/sigma tasks to be distributed over cores. For active spaces as large as the CAS(20,20), good scalability up to 2048 cores has been demonstrated. For larger active spaces, such as CAS(22,22) and CAS(24,24), this approach

significantly increases the number of CI tasks that have to be performed so that a smaller number of GASs are preferred. It should be noted that this fragmentation approach does generate additional overhead in the calculations. Our approach shows that CI calculations with an active space composed by 22 electrons in 22 orbitals have become accessible on large parallel computing platforms, while a chromium tetramer with an active space composed by 24 electrons in 24 orbitals would require a significant increase in parallel performance.

As the seminal publication of Olsen *et al.*³⁴ marked the new era in multiconfigurational theory for the calculation of “exact,” non-relativistic electronic energies for systems with CI expansions of one billion elements, this work sets the foundations for the “exact” treatment of larger active spaces. The new parallel implementation of MCSCF can provide benchmark results for the calibration of novel non-conventional CI methods such as DMRG or FCIQMC.

SUPPLEMENTARY MATERIAL

See [supplementary material](#) for Cartesian coordinates and input examples of Cr₃, Cr₄, and pentacene molecules.

ACKNOWLEDGMENTS

This work was supported by the (U.S.) Department of Energy (DOE), Office of Basic Energy Sciences, under SciDAC Grant No. DE-SC0008666. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. One of the authors (J.O.) acknowledges support from the Danish Council of Independent Research/Natural Sciences through Grant No. DFF-4181-00 537.

APPENDIX: EXAMPLE OF MULTIPLE TTSS BLOCK GENERATION WITH THE GAS SCHEME

Let us consider a system with C_{2v} symmetry (4 irreducible representations—irreps). We would like to perform a calculation for the 1A_1 state with a complete active space (CAS) composed by 4 electrons in 8 orbitals (CAS(4,8)) which are transformed according to the following irreps: $1a_1$, $2a_1$, $1a_2$, $2a_2$, $1b_1$, $2b_1$, $1b_2$, and $2b_2$ (Table VI). The CAS(4,8) can be considered as a GAS calculation with only one GAS (GAS 1).

The CI expansion is evaluated in combinations (CMs), rather than Slater determinants (SDs), which can also be constructed by an alpha string $\alpha(I_\alpha)$ and a beta string $\beta(I_\beta)$.

TABLE VI. Composition of the parent CAS(4,8).

Irrep	a_1	a_2	b_1	b_2
GAS 1	2	2	2	2
Total active	2	2	2	2

TABLE VII. Distribution of alpha and beta strings in supergroups for CAS(4,8).

Irrep	Type of alpha supergroups	Type of beta supergroups
Number of supergroups	1	1
Distribution of electrons per GAS	{2}	{2}
a_1	4	4
a_2	8	8
b_1	8	8
b_2	8	8

TABLE VIII. Number of CMs per TTSS block for a CAS(4,8).

TTSS block	Size
1	10
2	36
3	36
4	36
CI vector	118
Max size	36

Since we want to calculate a singlet state ($S = 0$), we need to construct CMs with two alpha and two beta electrons. The alpha string is composed by two creation operators,

$$\alpha(I_\alpha) = \alpha_{i\alpha}^\dagger(n)\alpha_{j\alpha}^\dagger(m),$$

$$i, j = 1a_1, 2a_1, 1a_2, 2a_2, 1b_1, 2b_1, 1b_2, \text{ or } 2b_2,$$

with each of them adding an alpha electron to the vacuum state. The indices n and m correspond to the GAS that the operator creates an electron ($n = m = 1$ for this example). For obtaining the type (T) of the alpha string $\alpha(I_\alpha)$, i.e., the spatial symmetry of $\alpha(I_\alpha)$, we have to consider all possible unique combinations

TABLE IX. Composition of the parent CAS(4,8) split into two GASs.

Irrep	a_1	a_2	b_1	b_2
GAS 1	1	2	2	2
GAS 2	1	0	0	0
Total active	2	2	2	2

TABLE X. Distribution of alpha and beta strings in supergroups for CAS(4,8) split into two GASs.

Irrep	Type of alpha supergroups		Type of beta supergroups	
	2	2	2	2
Number of supergroups	2	2	2	2
Distribution of electrons per GAS	{2 0}	{1 1}	{2 0}	{1 1}
a_1	3	1	3	1
a_2	6	2	6	2
b_1	6	2	6	2
b_2	6	2	6	2

TABLE XI. Number of CMs per TTSS block for a CAS(4,8) split into two GASs.

TTSS block	Size
1	6
2	21
3	21
4	21
5	3
6	12
7	12
8	12
9	1
10	3
11	3
12	3
CI vector	118
Max size	21

TABLE XII. Composition of the parent CAS(4,8) split into three GASs.

Irrep	a_1	a_2	b_1	b_2
GAS 1	1	1	2	2
GAS 2	1	0	0	0
GAS 3	0	1	0	0
Total active	2	2	2	2

between different irreps, i.e.,

$$\begin{aligned} &\alpha_{1a_1\alpha}^\dagger(1)\alpha_{2a_1\alpha}^\dagger(1), \\ &\alpha_{1a_1\alpha}^\dagger(1)\alpha_{1a_2\alpha}^\dagger(1), \\ &\alpha_{1a_1\alpha}^\dagger(1)\alpha_{2a_2\alpha}^\dagger(1), \\ &\text{etc.} \end{aligned}$$

In total, we have 28 different combinations. Let us define as a *supergroup* the distribution of electrons of a spin alpha or beta in different GASs. For this example, since we have only one GAS that both alpha electrons can occupy, there is only one alpha supergroup that we will call $\{2\}$ (i.e., two alpha electrons in GAS 1). Next, we group the strings of each supergroup according to their symmetry by calculating the corresponding direct products (Table VII), i.e., for the supergroup $\{2\}$, we have 4 strings with a_1 symmetry, 8 strings with a_2 , 8 strings with b_1 , and 8 strings with b_2 . Similar work is followed for the beta string ($\beta(I_\beta) = \alpha_{i\beta}^\dagger(n)\alpha_{j\beta}^\dagger(m)$).

TABLE XIII. Distribution of alpha and beta strings in supergroups for CAS(4,8) split into three GASs.

Irrep	Type of alpha supergroups				Type of beta supergroups			
Number of Supergroups	4				4			
Distribution of electrons per GAS	$\{2\ 0\ 0\}$	$\{1\ 1\ 0\}$	$\{1\ 0\ 1\}$	$\{0\ 1\ 1\}$	$\{2\ 0\ 0\}$	$\{1\ 1\ 0\}$	$\{1\ 0\ 1\}$	$\{0\ 1\ 1\}$
a_1	2	1	1	0	2	1	1	0
a_2	5	1	1	1	5	1	1	1
b_1	4	2	2	0	4	2	2	0
b_2	4	2	2	0	4	2	2	0

TABLE XIV. Number of CMs per TTSS block for a CAS(4,8) split into three GASs.

TTSS block	Size	TTSS block	Size
1	3	16	8
2	15	17	1
3	10	18	1
4	10	29	4
5	2	20	4
6	5	21	1
7	8	22	1
8	8	23	3
9	1	24	3
10	1	25	5
11	3	26	1
12	3	27	1
13	2	28	1
14	5	CI vector	118
15	8	Max size	15

By evaluating the spin of each supergroup product between alpha and beta types, we calculate the number of different TTSS blocks and the number of CMs that each block contains (Table VIII). For this example, where we want to calculate a root with A_1 symmetry, only four direct products between alpha and beta supergroups give CMs with the correct symmetry,

$$\begin{aligned} &a_1(\text{alpha}) \otimes a_1(\text{beta}), \\ &a_2(\text{alpha}) \otimes a_2(\text{beta}), \\ &b_1(\text{alpha}) \otimes b_1(\text{beta}), \\ &b_2(\text{alpha}) \otimes b_2(\text{beta}). \end{aligned}$$

Each direct product forms a TTSS block. In total, there are 4 TTSS blocks (Table VIII).

Let us now consider a case where we move one orbital (e.g., of a_1 symmetry) from GAS 1 to a new GAS (GAS 2, Table IX):

The number of supergroups per spin has increased to 2: two electrons in GAS 1 and zero in GAS 2 (termed as $\{2\ 0\}$), and the first electron in GAS 1 and the second in GAS 2 (termed as $\{1\ 1\}$).

By following this approach, we have increased the number of alpha types to eight (Table X). By considering all possible combinations between alpha and beta supergroups, avoiding

double counting of CMs and keeping those with the correct spin, we construct 12 TTSS blocks (Table XI).

Further splitting of the parent CAS(4,8) into three GASs (Table XII) increases the number of supergroups since we have now four different possibilities to distribute 2 α (or β) electrons in three GASs ($\{2\ 0\ 0\}$, $\{1\ 1\ 0\}$, $\{1\ 0\ 1\}$, and $\{0\ 1\ 1\}$). This increases the number of different strings (Table XIII) that eventually increases the number of TTSS blocks (Table XIV).

- ¹D. K. W. Mok, R. Neumann, and N. C. Handy, "Dynamical and nondynamical correlation," *J. Phys. Chem.* **100**(15), 6225–6230 (1996).
- ²N. C. Handy and A. J. Cohen, "Left-right correlation energy," *Mol. Phys.* **99**(5), 403–412 (2001).
- ³J. W. Hollett and P. M. W. Gill, "The two faces of static correlation," *J. Chem. Phys.* **134**(11), 114111 (2011).
- ⁴P. G. Szalay, T. Müller, G. Gidofalvi, H. Lischka, and R. Shepard, "Multiconfiguration self-consistent field and multireference configuration interaction methods and applications," *Chem. Rev.* **112**(1), 108–181 (2012).
- ⁵P. E. M. Siegbahn, "A new direct CI method for large CI expansions in a small orbital space," *Chem. Phys. Lett.* **109**(5), 417–421 (1984).
- ⁶N. C. Handy, "Multi-root configuration interaction calculations," *Chem. Phys. Lett.* **74**(2), 280–283 (1980).
- ⁷J. Olsen, B. O. Roos, P. Jørgensen, and H. J. A. Jensen, "Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces," *J. Chem. Phys.* **89**(4), 2185 (1988).
- ⁸S. Zarrabian, C. R. Sarma, and J. Paldus, "Vectorizable approach to molecular CI problems using determinantal basis," *Chem. Phys. Lett.* **155**(2), 183–188 (1989).
- ⁹W. A. de Jong, E. Bylaska, N. Govind, C. L. Janssen, K. Kowalski, T. Müller, I. M. B. Nielsen, H. J. J. van Dam, V. Veryazov, and R. Lindh, "Utilizing high performance computing for chemistry: Parallel computational chemistry," *Phys. Chem. Chem. Phys.* **12**(26), 6896–6920 (2010).
- ¹⁰H. Dachsel, H. Lischka, R. Shepard, J. Nieplocha, and R. J. Harrison, "A massively parallel multireference configuration interaction program: The parallel COLUMBUS program," *J. Comput. Chem.* **18**(3), 430–448 (1997).
- ¹¹A. J. Dobbyn, P. J. Knowles, and R. J. Harrison, "Parallel internally contracted multireference configuration interaction," *J. Comput. Chem.* **19**(11), 1215–1228 (1998).
- ¹²M. Klene, M. A. Robb, M. J. Frisch, and P. Celani, "Parallel implementation of the CI-vector evaluation in full CI/CAS-SCF," *J. Chem. Phys.* **113**(14), 5653 (2000).
- ¹³Z. Gan, Y. Alexeev, M. S. Gordon, and R. A. Kendall, "The parallel implementation of a full configuration interaction program," *J. Chem. Phys.* **119**(1), 47 (2003).
- ¹⁴S. Knecht, H. J. A. Jensen, and T. Fleig, "Large-scale parallel configuration interaction. I. Nonrelativistic and scalar-relativistic general active space implementation with application to (Rb–Ba)+," *J. Chem. Phys.* **128**(1), 014108 (2008).
- ¹⁵R. Ansaldi, G. L. Bendazzoli, S. Evangelisti, and E. Rossi, "A parallel Full-CI algorithm," *Comput. Phys. Commun.* **128**(1–2), 496–515 (2000).
- ¹⁶T. L. Windus, M. W. Schmidt, and M. S. Gordon, "Parallel algorithm for integral transformations and GUGA MCSCF," *Theor. Chim. Acta* **89**(1), 77–88 (1994).
- ¹⁷T. Müller, "Large-scale parallel uncontracted multireference-averaged quadratic coupled cluster: The ground state of the chromium dimer revisited," *J. Phys. Chem. A* **113**(45), 12729–12740 (2009).
- ¹⁸S. R. White, "Density matrix formulation for quantum renormalization groups," *Phys. Rev. Lett.* **69**(19), 2863–2866 (1992).
- ¹⁹U. Schollwöck, "The density-matrix renormalization group," *Rev. Mod. Phys.* **77**, 259–315 (2005).
- ²⁰G. K.-L. Chan and S. Sharma, "The density matrix renormalization group in quantum chemistry," *Annu. Rev. Phys. Chem.* **62**, 465–481 (2011).
- ²¹D. Zgid and M. Nooijen, "The density matrix renormalization group self-consistent field method: Orbital optimization with the density matrix renormalization group method in the active space," *J. Chem. Phys.* **128**(14), 144116 (2008).
- ²²D. Ghosh, J. Hachmann, T. Yanai, and G. K.-L. Chan, "Orbital optimization in the density matrix renormalization group, with applications to polyenes and β -carotene," *J. Chem. Phys.* **128**(14), 144117 (2008).
- ²³G. Gidofalvi and D. A. Mazziotti, "Active-space two-electron reduced-density-matrix method: Complete active-space calculations without diagonalization of the N-electron Hamiltonian," *J. Chem. Phys.* **129**(13), 134108 (2008).
- ²⁴W. Poelmans, M. Van Raemdonck, B. Verstichel, S. De Baerdemacker, A. Torre, L. Lain, G. E. Massaccesi, D. R. Alcoba, P. Bultinck, and D. Van Neck, "Variational optimization of the second-order density matrix corresponding to a seniority-zero configuration interaction wave function," *J. Chem. Theory Comput.* **11**(9), 4064–4076 (2015).
- ²⁵J. Fosso-Tande, T.-S. Nguyen, G. Gidofalvi, and A. E. DePrince III, "Large-scale variational two-electron reduced-density-matrix-driven complete active space self-consistent field methods," *J. Chem. Theory Comput.* **12**(5), 2260–2271 (2016).
- ²⁶G. H. Booth, A. J. W. Thom, and A. Alavi, "Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space," *J. Chem. Phys.* **131**(5), 054106 (2009).
- ²⁷N. S. Blunt, S. D. Smart, J. A.-F. Kersten, J. S. Spencer, G. H. Booth, and A. Alavi, "Semi-stochastic full configuration interaction quantum Monte Carlo: Developments and application," *J. Chem. Phys.* **142**(18), 184107 (2016).
- ²⁸R. E. Thomas, Q. Sun, A. Alavi, and G. H. Booth, "Stochastic multiconfigurational self-consistent field theory," *J. Chem. Theory Comput.* **11**(11), 5316–5325 (2015).
- ²⁹G. Li Manni, S. D. Smart, and A. Alavi, "Combining the complete active space self-consistent field method and the full configuration interaction quantum Monte Carlo within a super-CI framework, with application to challenging metal-porphyrins," *J. Chem. Theory Comput.* **12**(3), 1245–1258 (2016).
- ³⁰D. Ma, G. Li Manni, and L. Gagliardi, "The generalized active space concept in multiconfigurational self-consistent field methods," *J. Chem. Phys.* **135**(4), 044128 (2011).
- ³¹J. Ivanic, "Direct configuration interaction and multiconfigurational self-consistent-field method for multiple active spaces with variable occupations. I. Method," *J. Chem. Phys.* **119**(18), 9364 (2003).
- ³²K. D. Vogiatzis, G. Li Manni, S. J. Stoneburner, D. Ma, and L. Gagliardi, "Systematic expansion of active spaces beyond the CASSCF limit: A GASSCF/SplitGAS benchmark study," *J. Chem. Theory Comput.* **11**(7), 3010–3021 (2015).
- ³³S. Knecht, H. J. A. Jensen, and T. Fleig, "Large-scale parallel configuration interaction. II. Two- and four-component double-group general active space implementation with application to BiH," *J. Phys. Chem.* **132**(1), 014108 (2010).
- ³⁴J. Olsen, P. Jørgensen, and J. Simons, "Passing the one-billion limit in full configuration-interaction (FCI) calculations," *Chem. Phys. Lett.* **169**(6), 463–472 (1990).
- ³⁵M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. J. Van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, and W. A. de Jong, "NWChem: A comprehensive and scalable open-source solution for large scale molecular simulations," *Comput. Phys. Commun.* **181**(9), 1477–1489 (2010).
- ³⁶F. Aquilante, J. Autschbach, R. K. Carlson, L. F. Chibotaru, M. G. Delcey, L. De Vico, I. F. Galvan, N. Ferre, L. M. Frutos, L. Gagliardi, M. Garavelli, A. Giussani, C. E. Hoyer, G. Li Manni, H. Lischka, D. Ma, P. Malmqvist, T. Muller, A. Nenov, M. Olivucci, T. B. Pedersen, D. Peng, F. Plasser, B. Pritchard, M. Reiher, I. Rivalta, I. Schapiro, J. Segarra-Martí, M. Stenrup, D. G. Truhlar, L. Ungur, A. Valentini, S. Vancoullie, V. Veryazov, V. P. Vysotskiy, O. Weingart, F. Zapata, and R. Lindh, "MOLCAS 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table," *J. Comput. Chem.* **37**(5), 506–541 (2016).
- ³⁷B. Roos, "A new method for large-scale CI calculations," *Chem. Phys. Lett.* **15**(2), 153–159 (1972).
- ³⁸P. J. Knowles and N. C. Handy, "A new determinant-based full configuration interaction method," *Chem. Phys. Lett.* **111**(4–5), 315–321 (1984).
- ³⁹E. R. Davidson, "The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices," *J. Comput. Phys.* **17**(1), 87–94 (1975).
- ⁴⁰B. O. Roos and P. E. M. Siegbahn, in *Methods of Electronic Structure Theory*, edited by H. F. Schaefer III (Plenum Press, New York, 1977), Vol. 3, p. 277.

- ⁴¹J. Nieplocha, B. Palmer, V. Tipparaju, M. Krishnan, H. Trease, and E. Apra, "Advances, applications and performance of the global arrays shared memory programming toolkit," *Int. J. High Perform. Comput. Appl.* **20**(2), 203–231 (2006).
- ⁴²J. Hachmann, W. Cardoen, and G. K.-L. Chan, "Multireference correlation in long molecules with the quadratic scaling density matrix renormalization group," *J. Chem. Phys.* **125**(14), 144101 (2006).
- ⁴³S. Horn, F. Plasser, T. Müller, F. Libisch, J. Burgdörfer, and H. Lischka, "A comparison of singlet and triplet states for one- and two-dimensional graphene nanoribbons using multireference theory," *Theor. Chem. Acc.* **133**(8), 1511 (2014).
- ⁴⁴J. Lee, D. W. Small, E. Epifanovsky, and M. Head-Gordon, "Coupled-cluster valence-bond singles and doubles for strongly correlated systems: Block-tensor based implementation and application to oligoacenes," *J. Chem. Theory Comput.* **13**(2), 602–615 (2017).