

Computational and Theoretical Advances in Accurately Modeling Molecular Wave Functions

by

Duy-Khoi T. Dang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Chemistry and Scientific Computing)
in the University of Michigan
2023

Doctoral Committee:

Professor Paul Zimmerman, Chair

Professor Eitan Geva

Professor Robert Krasny

Professor Roseanne Sension

Duy-Khoi T. Dang

khoidang@umich.edu

ORCID iD: 0000-0001-7530-0540

© Duy-Khoi T. Dang 2023

DEDICATION

This thesis is dedicated to all my friends and family who have been a constant source of love, encouragement, and support throughout my time in graduate school.

“We do this not because it is easy, but because we thought it would be easy.”

ACKNOWLEDGMENTS

The work in this thesis would not have been possible without the mentorship, support, and friendship of countless people. First, I would like to thank Paul Zimmerman who has played a key part in my intellectual and professional development over the last 5 years. I would also like to thank my committee members Eitan Geva, Roseanne Sension, and Robert Krasny for their feedback throughout my projects and milestones. The scientific expertise of my collaborators from the Zimmerman and Gavini Groups as well as at the Oak Ridge National Laboratory and University of Wrocław, Poland have been immensely valuable.

I would like to acknowledge the National Science Foundation (NSF) for their support through a NSF Graduate Research Fellowship (DGE 1841052) as well as for computational resources provided through the Extreme Science and Engineering Discovery Environment (XSEDE) program at the San Diego Supercomputing Center's Expanse (CHE220041). I have additionally been supported by the Department of Energy grant DE-SC0022241, the Rackham graduate school through the Rackham One-Term Dissertation Fellowship, the Chemistry department at University of Michigan, and the Michigan Institute for Computational Discovery and Engineering.

Besides my research advisor, Zach Minden is likely the only person who has read every document I have composed in graduate school. I deeply value our countless hours of video games and inane conversations. Leighton Wilson has been an important mentor and friend throughout my time in graduate school. I appreciate our many abandoned projects and conversations on science, technology, and markets.

I am also thankful to the many friends I've made in graduate school, especially from the Hill House—special thanks to Jill and Bryan Kinzer, Kristen and Owen Hong, and Mery Vet George De la Rosa for always hosting and introducing me to the group. I'm grateful for all of our adventures. Phoebe Hunter, Ellie Frydendall, Chloe Weiers, Kevin Nguyen, Kathy Law, and Joanne Kwon are deeply valued friends from before I started graduate school. I appreciate that we've stayed in touch all these years.

During my undergraduate education, many mentors helped guide my academic decisions and cultivate my interest in science and computation including Rein Kirss, Mary Jo Ondrechen, Steven Lopez, Kun Song, and Alan Mathiowetz.

Last, but not least, I thank Sushi Cow, my parents, and family for their love and support.

TABLE OF CONTENTS

Dedication	ii
Acknowledgments	iii
List of Figures	vi
List of Tables	xi
List of Algorithms	xiii
List of Appendices	xiv
List of Acronyms	xv
Abstract	xvii
Chapter	
1 Introduction	1
1.1 Overview of quantum chemistry	1
1.1.1 The Schrödinger Equation	1
1.1.2 The Variational Principle	3
1.1.3 Molecular Orbitals and Slater Determinants	3
1.1.4 Basis Sets	4
1.1.5 Electron Correlation	6
1.2 Correlated Methods in Quantum Chemistry	6
1.2.1 Coupled-Cluster	7
1.2.2 Configuration Interaction Methods	7
1.2.3 Multi-Configurational Self-Consistent Field	8
1.2.4 Incremental and Many-Body Expansion Methods	9
1.2.5 Summary	11
2 An Open-Shell Coronoid with Hybrid Chichibabin-Schlenk Conjugation	12
2.1 Introduction	12
2.2 Computational Details	14
2.3 Results and Discussion	15
2.4 Conclusions	20
3 Advances in Parallel Heat Bath Configuration Interaction	22

3.1	Introduction	22
3.2	Theory and Implementation Details	25
3.2.1	The Variational Step	25
3.2.2	The Perturbative Step	26
3.2.3	Data Structures	27
3.2.4	Parallelism and Algorithmic Optimization	28
3.3	Computational Details	33
3.4	Results and Discussion	34
3.5	Conclusions	41
4	Fully Variational Incremental CASSCF	43
4.1	Introduction	43
4.2	Theory	46
4.2.1	iCASSCF	46
4.2.2	Computational Effort	50
4.3	Computational Details	51
4.4	Small-molecule benchmarks	51
4.5	oxoMn(salen)Cl	53
4.6	Cyclobutadiene Automerization	57
4.7	Conclusions	59
5	The Numerical Evaluation of Slater Integrals on Graphics Processing Units 61	
5.1	Introduction	61
5.2	Theory and Computational Details	64
5.2.1	Resolution of the Identity	64
5.2.2	Grid Construction	66
5.2.3	Implementation on GPU	67
5.2.4	Computational Details	69
5.3	Results and Discussion	69
5.3.1	Performance Analysis	69
5.3.2	Mixed-Precision Evaluation	72
5.3.3	Hartree-Fock and HCI	73
5.4	Conclusions	76
6	Final Remarks	78
6.1	Conclusions	78
6.2	Future Directions	79
6.3	Outlooks in Quantum Chemistry	80
	Appendices	81
	Bibliography	119

LIST OF FIGURES

FIGURE

1.1	A schematic of the Slater determinant representation is shown.	4
1.2	Contracted Gaussian functions to approximate a Slater function are shown for $\zeta = 1$	5
1.3	All possible electron configurations for 2 electrons in 2 orbitals.	6
2.1	Chichibabin and Schlenk conjugation in organic oligoradicaloids. Fully open-shell configurations shown for all all oligoradicaloids. Substituents not shown for 4a,b and 5 . Unpaired electrons are presented as red dots or p orbitals.	13
2.2	A) SF energies and spins for the first eight states of $\mathbf{8}'$. The lowest-energy states for a given spin multiplicity is indicated with an arrow. The singlet geometry $\mathbf{18}'$ has a 6-fold degeneracy at 0.156 eV consisting of three singlets, two triplets, and one septet state. B) Relationship between ΔE_{SF} and the number of unpaired electrons n_U derived from natural orbital occupation numbers. C) Frontier natural orbitals for the $\mathbf{18}'$ - S_0 state (0.02 a.u. isosurface). D) SF odd-electron density for $\mathbf{18}'$ (0.002 a.u. isosurface). E) Definition of coupling constants used in an approximate spin Hamiltonian. J values are given for the $\mathbf{18}'$ state.	16
2.3	Polyradicaloid indices (γ_i) from RAS(6,6)-SF/cc-pVDZ are shown for the ground singlet electronic state at various optimized CASSCF(6,6)/cc-pVDZ geometries.	17
2.4	A) The spin Hamiltonian is shown. B) The odd-electron density (isovalue=0.005 a.u.) of the optimized $\mathbf{18}'$ geometry with the six radical sites labeled. C) The coupling diagram corresponding to the spin Hamiltonian is drawn. For the D_{3h} geometries, $\varepsilon_1 = \varepsilon_2 = 0$	19
3.1	Heatmaps of the CI Hamiltonian at each iteration in the variational step of a HCI calculation are shown. The submatrices of prior HCI iterations are boxed in green, red, and black for the first, second, and third iterations, respectively. Darker colored cells indicate larger magnitude Hamiltonian elements.	25
3.2	The heatmap shows the variational Hamiltonian (left) from Figure 3.1 as well as the terms $H_{ki}c_i$ for the perturbative determinants (right). On the left, each column represents a determinant in the variational space, while each column on the right represents a perturbative determinant. The inner sum in Equation 3.1 runs over columns in this figure. Darker colors indicate larger magnitudes.	26

3.3	This schematic shows an example of how an arbitrary determinant, d_k , is assigned to a MPI rank. In this example, there are 9 MPI ranks involved in the calculation. The determinant is hashed by its α - and β -strings separately and then assigned based on predetermined ranges or bins for both the α and β hashes.	31
3.4	Wall times (left) and speedups (right) for a HCI calculation of butadiene on up to 32 nodes are shown. Perfect speedup is shown as a solid black line. The HCI calculation included all valence electrons in the cc-pVDZ basis set with thresholds set to $\varepsilon_1 = 1 \times 10^{-4}$ and $\varepsilon_2 = 1 \times 10^{-7}$	35
3.5	Automerization of cyclobutadiene (top) is shown with relative full CI energies in kcal mol ⁻¹ plotted (below). Full CI energies were computed using a 4-point extrapolation from HCI calculations (See Table 3.4).	36
3.6	Linear extrapolation of ¹ D _{2h} , ³ D _{2h} , ¹ D _{4h} , and ³ D _{4h} cyclobutadiene energies to the full CI limit. Extrapolated energies and R^2 values are provided in the legend. HCI energies were computed in the cc-pVTZ basis with a (20e,172o) active space and ε_1 values of 0.10 mHa, 0.15 mHa, 0.20 mHa, and 0.25 mHa. The perturbative cutoff, ε_2 was set to 10 ⁻⁷ Ha for all cyclobutadiene calculations.	36
3.7	The 4-point FCI extrapolation of [FeO(NH ₃) ₅] ²⁺ with (22e,168o) is shown. All perturbative energies are computed with $\varepsilon_2 = 1 \times 10^{-7}$ Ha. The FCI extrapolated energy and R^2 are provided in the legend. Variational thresholds for ε_1 and the basis set are listed in the Computational Details.	39
3.8	NOs for the triplet state computed using the variational density from HCI with $\varepsilon_1 = 0.1$ mHa are shown with isosurface values of 0.8 a.u. NO occupations are listed below each orbital.	40
3.9	NOs for the quintet state computed using the variational density from HCI with $\varepsilon_1 = 0.1$ mHa are shown with an isosurface values of 0.8 a.u. NO occupations are listed below each orbital.	41
4.1	a. The full valence CASSCF active space is decomposed into b. pairs of localized bonding and anti-bonding orbitals which define the bodies in the many-body expansion. Solving the CAS-CI problem within these active spaces provides the 1-body interactions (ε_i) in Equation 4.1. c. The combination of any two active spaces in the 1-body terms generates the 2-body terms. Solving the CAS(4,4)-CI problem for these new active spaces returns the 2-body interaction energies (ε_{ij}) in Equation 4.1. Higher-order expansions are similarly defined (not shown). . .	47
4.2	The magnitude of the orbital gradient (see Equation 4.7) for formaldehyde is shown with single excitations (left) and without (right). The gradient is truncated so only values corresponding to the first twenty orbitals are shown. One can see that the magnitude of gradient terms increase as the redundant single excitations are removed.	49
4.3	Comparison of iCASSCF to CASSCF (a) bond angles, (b) bond distances, and (c) energies. Included with the energy comparisons are iCAS-CI energies which use converged iCASSCF orbitals.	53
4.4	Select iCASSCF($n = 3$) natural orbitals in the π space for the singlet state of the oxo-Mn(salen)Cl complex. The natural occupation numbers are displayed below each orbital and the HONO and LUNO are boxed.	56

4.5	Select iCASSCF($n = 3$) natural orbitals in the π space for the triplet state of the oxo-Mn(salen)Cl complex are shown. The occupation of each orbital is provided below the orbital. The SONOs are boxed.	56
4.6	Optimized reaction pathway via GSM and iCASSCF($n=3$)/6-31+G**. Single point energies from iCAS-CI($n=3$) are nearly identical, and are shown in red.	58
4.7	Optimized transition state geometry for cyclobutadiene automerization from GSM. The C-C bond lengths are provided in Å, and some key bond angles and dihedrals are shown.	59
5.1	The $(6H 6H)$ integral is scanned in the $(0.370, 0.370, 0.853)$ direction with the left center at the origin. Evaluations are in mixed precision using either the Laurent polynomial expansion of Equation 5.10 or the lower incomplete gamma function, where mixed precision is defined similarly to Equation 5.14. Both basis functions have $m = 0$ and $\zeta = 1$	65
5.2	The GPU speedups for integral evaluation over CPU for the 2080-Ti (top) and GV100 (bottom) are shown for various alkanes using the DZP basis from ADF. The speedups are partitioned into the various integrals. Speedups for mixed (left) and double (right) precision evaluations are also shown.	70
5.3	Multi-GPU speedups over single GPU and parallel efficiency for mixed (left) and double (right) precision evaluation of the 3-center ERIs for C_9H_{20} . There are 76,873,200 3-center integrals. Perfect scaling is plotted as a solid black line. All GPUs are co-located on a single compute node. Single GPU run times were 125s and 542s for mixed- and double-precision implementations, respectively.	71
5.4	The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.	72
5.5	The value of 2-center ERIs are evaluated in mixed precision. All basis functions have $\zeta = 1$ and the right basis is scanned radially away from the origin in various directions. The directions selected are provided in the legend and were selected using the 16 all-positive directions of an 86-point Lebedev grid. For all integrals shown here, $m = 0$. The legend entries are direction unit vectors.	73
5.6	Heatmaps of the relative error of HF energies when using mixed- vs double-precision integral evaluation are shown for various alkanes using different angular and radial grid sizes.	74
5.7	The S_N2 reaction for fluoride exchange of fluoromethane.	75
5.8	The Hartree-Fock geometric gradient projections (solid blue lines) of molecules with D_{3h} and T_d point groups are plotted as the A-X bond distance is scanned, where A=B,C and X=H,F. Gradients were computed in mixed precision. The mixed-precision (solid green) and double-precision (dotted red) Hartree-Fock energies at each point are also plotted. For CF_4 , the auxiliary basis for fluorine is extended with additional $2p$, $3d$, $4f$ and $5g$ functions.	76

A.1	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (S_0) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	82
A.2	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	83
A.3	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_2) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	84
A.4	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_3) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	85
A.5	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Q_{n_1}) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	86
A.6	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Q_{n_2}) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	87
A.7	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Sp_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	88
A.8	Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (S_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.	89
B.1	5-point HCI extrapolation to the FCI energy of butadiene in the ANO-L-VDZP basis. Energies were computed with ε_1 of 2.5×10^{-4} , 2.0×10^{-4} , 1.5×10^{-4} , 1.0×10^{-4} , and 5.0×10^{-5} with $\varepsilon_2 = 1.0 \times 10^{-8}$. The extrapolated energy and associated R^2 are listed in the legend.	91
B.2	The localized occupied and singly-occupied ROHF orbitals for the triplet state of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ included in the active space are depicted. Isosurface value of 0.1 a.u. is used for all figures.	92
C.1	Differences between canonical CASSCF, iCASSCF($n = 4$), iCAS-CI($n = 4$) and the original implementation of iCASSCF($n = 4$) are shown for ethane dissociation. The NPE for iCASSCF, iCAS-CI, and the original iCASSCF are 2.15 mHa, 0.23 mHa, and 0.14 mHa, respectively.	96
C.2	Small molecule benchmark results including results from previous studies.	97
C.3	Select iCAS-CI($n=3$) natural orbitals for the singlet state of oxo-Mn(salen)Cl are shown. HONO and LUNO are boxed.	99
C.4	Select iCAS-CI($n=3$) natural orbitals for the triplet state of oxo-Mn(salen)Cl are shown. SONOs are boxed.	99
D.1	Relative performance of Slater integration is plotted as memory at various memory and SM clock speeds. Integrals were computed for the C_5H_{12} molecule. The x -axis labels list the memory clock followed by the SM clock in MHz.	111

D.2	The LiF molecule is stretched from 1.4 Å to 3.0 Å. Energies (top) were computed using HBCI($\varepsilon_1 = 0.5$ mHa, $\varepsilon_2 = 0.1$ μHa) with the TZ2P basis. An additional point at the experimental bond distance (1.564 Å) was included. Energies are relative to complete dissociation of the LiH bond. The non-parallelity error (bottom) from using mixed-precision integrals is also plotted.	113
D.3	The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.	115
D.4	The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$. The left basis has $m = 0$ and the right basis has $m = n - 1$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.	116
D.5	The mixed-precision 2-center ERIs are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The ERIs are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.	117
D.6	The mixed-precision 2-center ERIs are plotted for various basis functions. All basis functions have $\zeta = 1$. The left basis has $m = 0$ and the right basis has $m = n - 1$. The ERIs are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.	118

LIST OF TABLES

TABLE

1.1	Number of Slater determinants for various small molecules for FCI in the cc-pVDZ basis.	6
2.1	Computational results for adiabatic spin states of $\mathbf{8}'$	15
2.2	Computed spin couplings for $\mathbf{8}'$ (eV). J_i are computed spin coupling constants, and ε_i are perturbations to spin couplings due to the C_{2v} geometry.	19
3.1	The notation for various data structures and their descriptions are listed below.	29
3.2	Key data structures and parallel implementation details.	33
3.3	Diatomic energy benchmarks of HCl computed in the cc-pVTZ basis. The perturbative cutoff is $\varepsilon_2 = 10^{-8}$ Ha for HCl and SHCl.	35
3.4	HCl energies for cyclobutadiene automerization. All calculations are done in the cc-pVTZ basis set with $\varepsilon_2 = 1 \times 10^{-7}$ Ha. Linear extrapolations using the tightest n ε_1 values are shown for each state.	37
3.5	Relative energies of cyclobutadiene computed using different methods. Energies are in kcal mol $^{-1}$	38
3.6	Quintet energy relative to the triplet energy of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ at various levels of theory. Energies are in kcal mol $^{-1}$	38
3.7	HCl energies for the $[\text{FeO}(\text{NH}_3)_5]^{2+}$ triplet and quintet states. Basis set is described in Computational Details. In the perturbative step, $\varepsilon_2 = 1 \times 10^{-7}$ Ha for all calculations. Linear extrapolation is used with the n smallest ε_1 values.	40
4.1	Total energies of the small molecule benchmark set from iCASSCF, iCAS-CI, and CASSCF using the 6-31G* basis set.	52
4.2	Norm of difference between analytical and numerical gradient (Ha Bohr $^{-1}$).	53
4.3	The singlet-triplet gap of oxoMn(salen)Cl from previous multi-configurational studies are listed.	54
4.4	Total energies of the oxoMn(salen)Cl complex from iCASSCF and iCAS-CI using the 6-31G* basis set.	56
5.1	Double-precision timing data (in seconds) for various alkanes. Each atom contributes 46,200 grid points.	70
5.2	Mixed-precision timing data (in seconds) for various alkanes. Each atom contributes 46,200 grid points.	70

5.3	HF energies computed for several small molecules are listed. The STO basis sets and grids are described in the Computational Details. The numbers in parenthesis in the header denotes the bits of precision used for integral evaluation.	74
5.4	HF and HCl activation energies (kcal mol ⁻¹) of CH ₃ F fluoride exchange at various grid sizes using single- and double-precision integral evaluations. The number of radial points and Lebedev order are provided for the radial and angular grids. The size of the angular grid is given in parenthesis next to the Lebedev order. .	75
5.5	Relative energies of cyclobutadiene at <i>D</i> _{2h} and <i>D</i> _{4h} geometries (kcal mol ⁻¹). . .	75
B.1	Energy benchmarks for butadiene in the ANO-L-VDZP basis set. The perturbative cutoff is $\varepsilon_2 = 10^{-8}$ Ha for HCl.	91
B.2	Butadiene energy benchmarks in the ANO-L-VDZP basis set.	91
B.3	Time to hash 5.6×10^8 α -strings of length 192 with 11 set bits.	93
C.1	Small molecule benchmark geometries and energies from iCASSCF and iCAS-CI using 6-31G* basis set. Conventional CASSCF from Molpro. iCAS-CI($n = k$) energies use iCASSCF($n = k$) geometries and orbitals. Bond distances, angles, and energies given in Å, degrees, and Hartrees, respectively.	98
C.2	Coordinates for the optimized ω B97XD/6-31G* geometry of oxo-Mn(salen)Cl. .	100
C.3	Coordinates for the optimized iCASSCF($n=3$)/6-31+G** geometry of the cyclobutadiene reactant.	100
C.4	Coordinates for the iCASSCF($n=3$)/6-31+G** geometry of the cyclobutadiene TS found from GSM	101
C.5	Coordinates for the <i>D</i> _{4h} -symmetrized cyclobutadiene TS.	101
C.6	Conventional CASSCF coordinates for optimized methanol geometry.	101
C.7	Conventional CASSCF coordinates for optimized hydrogen peroxide geometry. .	101
D.1	Timings (in seconds) for Slater integration at various memory and SM clock speeds. Integrals were computed for the C ₅ H ₁₂ molecule.	111
D.2	Hartree-Fock timings (in seconds) for various alkanes. Each atom contributes 46,200 grid points. Timings are recorded on a 2080-Ti GPU. Total time in HF does not include integral computation.	112

LIST OF ALGORITHMS

ALGORITHM

3.1	Pseudocode for adding new determinants in the variational step	30
3.2	Pseudocode for perturbative single excitations for a single OpenMP thread operating on hash range, R	31
3.3	Pseudocode of the perturbative step for mixed α/β double excitations for a single OpenMP thread operating on hash range, R	32
3.4	Pseudocode of a single OpenMP thread for removing perturbative contributions from determinants in the variational space.	32
5.1	GPU compute structure for generating 3-center ERIs	68
B.1	C++ code for FNV hash.	93
B.2	C++ code for hash function used in HCI implementation	94
D.1	Loop for computing the basis χ or potential V on a grid \mathbf{x} with <code>npts</code> . The function <code>f</code> refers to either χ , or V , and <code>fx</code> refers to <code>f</code> evaluated on the grid \mathbf{x} . .	113
D.2	Sample code for contracting V_P , μ , and ν into the 3-center integrals. Arrays corresponding to V_P , μ , and ν are assumed to have already been computed, and properly weighted.	114

LIST OF APPENDICES

APPENDIX

A Supporting Information for Chapter 2	81
B Supporting Information for Chapter 3	90
C Supporting Information for Chapter 4	95
D Supporting Information for Chapter 5	102

LIST OF ACRONYMS

AO atomic orbital

CAS complete active space

CC coupled-cluster

CCSD coupled-cluster singles and doubles

CCSDT coupled-cluster singles, doubles, and triples

CCSD(T) coupled-cluster singles, doubles, and perturbative triples

CI configuration interaction

CIS configuration interaction singles

CISD configuration interaction singles and doubles

EN Epstein-Nesbet

FCI full configuration interaction

GTO Gaussian-type orbital

HCI heat-bath configuration interaction

HF Hartree-Fock

HOMO highest occupied molecular orbital

HONO highest occupied natural orbital

iCAS incremental complete active space

iFCI incremental full configuration interaction

LUMO lowest unoccupied molecular orbital

LUNO lowest unoccupied natural orbital

MBE many-body expansion

MO molecular orbital

NO natural orbital

QMC quantum Monte Carlo

RAS restricted active space

SCF self-consistent field

SF spin-flip

STO Slater-type orbital

ABSTRACT

Throughout the past few decades, the use of theoretical methods has become standard practice in the chemical disciplines. This thesis addresses developments in quantum chemical methods for studying challenging electronic structure problems. Background in quantum chemistry and various correlated methods are presented throughout the first chapter. The remaining chapters are described in the following.

Chapter 2 is based on an experimental collaboration studying a novel, open-shell coronoid system. Using the spin-flip (SF) method combined with restricted active space (RAS) configuration interaction, the ground state is predicted to be a singlet with significant hexaradicaloid character ($\gamma_0 = 0.826, \gamma_1 = \gamma_2 = 0.773$). It has multiple high-spin, low-lying states (up to septet) that are found to be thermally accessible, with nearly uniform energy gaps between consecutive multiplicities. Using the results of RAS-SF, a spin-interaction Hamiltonian—generated to analyze the spin alignment of the molecule—finds predominantly antiferromagnetic coupling between radical site pairs.

In Chapter 3, computational advances in the heat-bath configuration interaction (HCI) method are presented. MPI+OpenMP are used to target improvements in speed, parallel efficiency, and memory requirements. The implementation introduces a hash function to distribute determinants in both the CI and perturbative spaces. These advances enable the study of the triplet-quintet gap in the $[\text{FeO}(\text{NH}_3)_5]^{2+}$ molecule using a (22e,168o) active space, which explicitly included 2.39×10^7 variational determinants and 8.95×10^{10} perturbative determinants. Benchmarks show up to 86% parallel efficiency of the perturbative step on 32 nodes (4096 cores) and total efficiency of 75%. The chapter also includes benchmarks for accuracy against prior studies.

The complete active space self-consistent field (CASSCF) method holds a central place in conceptualizing and practicing quantum chemistry. For application to realistic molecules, however, CASSCF must be approximated to circumvent its exponential scaling. Applying the many-body expansion—also known as the method of increments—to CASSCF (iCASSCF) has been shown to produce a polynomially scaling method retaining the accuracy of the parent theory while also being capable of treating substantially larger active spaces. However, the orbital parameters of the original iCASSCF implementation were not variationally optimized.

Chapter 4 details the theoretical advances to iCASSCF making the method fully variational. These advances enable the method to produce accurate nuclear gradients and optimize stable geometries as well as transition states. Demonstrations on challenging test cases, such as the oxoMn(salen)Cl complex with an active space of (84e,84o) and the automerization of cyclobutadiene show the power of fully variational iCASSCF for describing challenging molecular systems.

Finally, Chapter 5 introduces SlaterGPU, a GPU accelerated library to numerically evaluate the Slater-type orbital (STO) integrals. The electron repulsion integrals (ERIs) are computed under the RI approximation using the analytic Coulomb potential of the Slater basis function. To fully realize the performance capabilities of modern GPUs, the Slater integrals are evaluated in mixed-precision, resulting in speedups for the ERIs of over $80\times$. Parallelization on multiple GPUs allows for integral throughput of over 3 million integrals per second, placing STO integration within reach of single-threaded, conventional Gaussian integration schemes. Benchmarks highlighting the quality and speed of the integrals demonstrate the library's ability to generate the full set of integrals necessary for configuration interaction with up to $6h$ functions in the auxiliary basis.

CHAPTER 1

Introduction

Chemistry is a historically empirical discipline. However, theoretical methods have become widely used in chemical research throughout the past few decades. This thesis addresses developments in quantum chemical methods for studying challenging electronic structure problems. Quantum mechanics is notorious for problems which scale poorly with system size. However, advances in computer hardware and theory have enabled routine use of theoretical methods. These advances must balance the trade-off between accuracy and computational cost. Background in quantum chemistry is presented throughout the current chapter. Applications of electronic structure theory methods to an open-shell coronoid molecule are presented in Chapter 2, which acts to showcase the capabilities of previously developed quantum chemical methods. Computational advances in the heat-bath configuration interaction (HCI) method are discussed in Chapter 3. Theoretical advances in the incremental complete active space self-consistent field (iCASSCF) method are detailed in Chapter 4. Chapter 5 details the development of a GPU accelerated library to handle the molecular integrals needed to use Slater orbital basis sets in modern quantum chemical methods.

1.1 Overview of quantum chemistry

We will first discuss the key principles and equations used in quantum chemistry. This will only briefly summarize concepts that are explored in greater detail elsewhere.^{1,2} As the name suggests, quantum chemistry is built on top of quantum mechanics—utilizing quantum mechanical principles to provide theoretical descriptions of the electronic structure of molecules. As such, the term electronic structure theory is frequently used interchangeably with quantum chemistry.

1.1.1 The Schrödinger Equation

The first postulate of quantum mechanics states that a quantum mechanical system is completely specified by its wave function, frequently denoted ψ . From the wave function, we

can obtain any classically observable property, O , by applying its corresponding Hermitian operator, \hat{O} , to the wave function

$$\hat{O} |\psi\rangle = O |\psi\rangle. \quad (1.1)$$

For example, the operator for the non-relativistic energy, E , of a quantum system—referred to as the Hamiltonian operator and denoted by \hat{H} —can be written for a molecule as

$$\hat{H} = -\frac{1}{2} \sum_i \nabla_i^2 - \frac{1}{2} \sum_I \nabla_I^2 + \sum_{i>j} \frac{1}{|r_i - r_j|} + \sum_{I>J} \frac{Z_I Z_J}{|R_I - R_J|} - \sum_I \sum_i \frac{Z_I}{|r_i - R_I|}, \quad (1.2)$$

where Z_I is the charge of nucleus I , R_I is the coordinate of nucleus I , and r_i is the coordinate of electron i . The Hamiltonian contains kinetic energy contributions of the **electrons** and the **nuclei**, and potential energy contributions from the **electron-electron repulsion**, **nuclear-nuclear repulsion**, and **nuclear-electron attraction**. Substituting Equation 1.2 into Equation 1.1 yields the Time-Independent Schrödinger Equation

$$\hat{H} |\psi\rangle = E |\psi\rangle, \quad (1.3)$$

which we will refer to simply as the Schrödinger Equation.

In most cases, the Born-Oppenheimer approximation is invoked, which allows the electronic and nuclear degrees of freedom to be treated separately. In the Born-Oppenheimer approximation, the nuclei are treated classically due to the large difference in timescale of nuclear motion compared to electronic motion. Essentially, as the nuclei—which are several orders of magnitude more massive than electrons—move, the surrounding electrons are able to fully relax their distribution around the nuclei. The electronic wave function, ψ_{el} , is then solved in a field of nuclei at fixed positions. Thus, the new Hamiltonian of interest becomes the electronic Hamiltonian, \hat{H}_{el} , which omits terms dealing only with the nuclei:

$$\hat{H}_{el} = -\frac{1}{2} \sum_i \nabla_i^2 + \sum_{i>j} \frac{1}{|r_i - r_j|} - \sum_I \sum_i \frac{Z_I}{|r_i - R_I|}. \quad (1.4)$$

This sets up the electronic Schrödinger Equation

$$\hat{H}_{el} |\psi_{el}\rangle = E_{el} |\psi_{el}\rangle, \quad (1.5)$$

which is the primary equation of interest in quantum chemistry. Because this thesis deals in electronic structure theory, we will drop the subscript and use \hat{H} and ψ to refer to the electronic Hamiltonian and wave function, respectively, throughout the remainder of the text.

1.1.2 The Variational Principle

Closed-form analytic solutions to Equation 1.5 only exist for hydrogen-like atoms (i.e. one electron and one nucleus). For all other systems where such solutions are unavailable, we rely on the Variational Principle, which has the following mathematical statement

$$\frac{\langle \psi | \hat{H} | \psi \rangle}{\langle \psi | \psi \rangle} \geq E_0. \quad (1.6)$$

In other words, the computed energy of an approximate wave function will always be an upper bound to the ground state energy of the true ground state wave function. Thus, even without analytic solutions, wave functions for molecular systems are solvable via optimization of the energy functional

$$E[\phi] = \frac{\langle \phi | \hat{H} | \phi \rangle}{\langle \phi | \phi \rangle}, \quad (1.7)$$

where we start with some guess wave function ϕ . Many widely used electronic structure theory methods (including those developed in this thesis) take advantage of this principle.

1.1.3 Molecular Orbitals and Slater Determinants

To use the Variational Principle, a guess wave function must first be constructed, which describes an initial spatial distribution of electrons. In canonical quantum chemistry, the guess is typically a Slater determinant, which has the advantage of implicitly encoding exchange and Pauli exclusion into the wave function. The Slater determinant is often written as

$$\psi = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(\mathbf{r}_1) & \phi_2(\mathbf{r}_1) & \cdots & \phi_N(\mathbf{r}_1) \\ \phi_1(\mathbf{r}_2) & \phi_2(\mathbf{r}_2) & \cdots & \phi_N(\mathbf{r}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_N) & \phi_2(\mathbf{r}_N) & \cdots & \phi_N(\mathbf{r}_N) \end{vmatrix}, \quad (1.8)$$

where the ϕ_i are (spin) molecular orbitals (MOs) and \mathbf{r}_j are the coordinates for electron j . Another representation of the Slater determinant is shown in Figure 1.1. Here, the MOs are denoted by horizontal lines, and “up”- and “down”-spin electrons are represented by the up- and down-pointing arrows, respectively. There are generally more MOs than electron pairs leaving many MOs unoccupied, which are sometimes referred to as virtual orbitals. The MOs themselves are composed of a linear combination of one-electron, atom-centered orbital basis functions:

$$\phi_p = \sum_{\mu} c_{\mu p} \chi_{\mu}, \quad (1.9)$$




Figure 1.1: A schematic of the Slater determinant representation is shown.

where χ_μ is the atomic orbital (AO) basis function and $c_{\mu p}$ is the MO orbital coefficient. The use of a one-electron, atom-centered basis allows the representation of electrons in 3-dimensional space to be accurately and efficiently captured using only a finite number of functions.

The coefficients describing the MOs can be optimized using a self-consistent field (SCF) algorithm to provide the best single Slater determinant wave function. This procedure is called Hartree-Fock (HF) and the Slater determinant where only the most stable MOs are occupied is referred to as the HF state. The HF procedure typically yields a wave function that accounts for 99% of the electronic energy. However to achieve chemical accuracy, which is generally accepted to be $< 1 \text{ kcal mol}^{-1}$, the remaining electronic energy must be captured. Many post-HF methods have been developed to recover this missing energy—the so-called “correlation” energy—from the HF wave function.

1.1.4 Basis Sets

One-electron, atom-centered basis sets—of which there are multiple choices—are the most commonly used in quantum chemistry. The most natural choice are the Slater-type orbitals (STOs), which are hydrogen-like orbitals of the form

$$S(\zeta, n, l, m, r, \theta, \phi) = N^{\text{STO}} r^{n-1} e^{-\zeta r} Z_{lm}(\theta, \phi), \quad (1.10)$$

where ζ is the exponent, n, l, m are the atomic quantum numbers, r, θ, ϕ are spherical coordinates, N^{STO} is the normalization constant, and Z_{lm} are the spherical harmonics.²⁻⁴ In many post-HF methods, construction of the Hamiltonian matrix requires computing one-

and two-electron integrals of the form

$$O_{\mu\nu} = \langle \chi_\mu | \hat{O}_1 | \chi_\nu \rangle = \int \chi_\mu(\mathbf{r}) \hat{O}_1(\mathbf{r}) \chi_\nu(\mathbf{r}) d\mathbf{r}, \quad (1.11)$$

$$\begin{aligned} O_{\mu\nu\lambda\sigma} &= \langle \chi_\mu(1) \chi_\nu(1) | \hat{O}_2 | \chi_\lambda(2) \chi_\sigma(2) \rangle \\ &= \int \int \chi_\mu(\mathbf{r}_1) \chi_\nu(\mathbf{r}_1) \hat{O}_2 \chi_\lambda(\mathbf{r}_2) \chi_\sigma(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2. \end{aligned} \quad (1.12)$$

However, the necessary integrals, such as the electron-repulsion integrals ($\hat{O}_2 = \frac{1}{r_{12}}$), are not known analytically for STOs. As such, they require numerical integration, which is slow and creates a large computational bottleneck. This difficulty led to the expansion of STOs in terms of Gaussian-type orbitals (GTOs)⁵

$$G(\alpha, n, l, m, r, \theta, \phi) = N^{\text{GTO}} e^{-\alpha r^2} S_{lm}(r, \theta, \phi), \quad (1.13)$$

where S_{lm} are the real solid harmonics.² Figure 1.2 shows a STO approximated as a GTO. As one can see, there are some key differences between GTOs and STOs. For one, GTOs




Figure 1.2: Contracted Gaussian functions to approximate a Slater function are shown for $\zeta = 1$.

clearly do not have a cusp at the nucleus. Furthermore, they decay much more rapidly than the STO (Gaussian fast as opposed to exponentially fast). For accurate representations of the wave function, both the nuclear cusp and exponential decay are required. These two features strongly affect the calculation of nuclear magnetic resonance tensors^{6,7} and the HOMO energy⁸ as well as the inverse density functional theory problem.^{9,10} However, GTOs have fast analytic expressions for just about every integral of interest in quantum chemistry,

which is the primary reason for their wide adoption. With GTOs the world is your oyster, but it happens to not be a very good one. In this thesis, we will explore strategies for quickly evaluating STO integrals by leveraging graphics processing units (GPUs) in Chapter 5.

1.1.5 Electron Correlation

Electron correlation is frequently divided into two kinds: the first is strong (or static) correlation, which comes from several Slater determinants being close in energy, and the second is weak (or dynamic) correlation, which can be thought of as coming from instantaneous electron-electron repulsion. Strong correlation requires the wave function to be represented with the dominant Slater determinants, i.e. those with large population, while weak correlation requires the inclusion of as many determinants as possible. For proper treatment, one must explicitly include multiple Slater determinants to handle strong correlation, while the recovery of weak correlation can be dealt with using strategies such as perturbation theory.

1.2 Correlated Methods in Quantum Chemistry

The exact solution to the Schrödinger equation in a given basis can be obtained with full configuration interaction (FCI).² The FCI wave function represents the wave function as a linear combination Slater determinants formed from all possible electron occupancies in all available MOs. Figure 1.3 shows the possible Slater determinants—often referred to




Figure 1.3: All possible electron configurations for 2 electrons in 2 orbitals.

as electron configurations, or configurations for short—for a system with 2 electrons and 2 orbitals. While FCI is formally exact, the combinatorial scaling causes the problem size to

Table 1.1: Number of Slater determinants for various small molecules for FCI in the cc-pVDZ basis.

Molecule	C ₂	C ₂ H ₄	C ₄ H ₆	C ₆ H ₈	C ₈ H ₁₀
Valence electrons	8	12	22	32	42
Basis functions	28	48	86	124	162
Slater determinants	2.2×10^8	8.8×10^{13}	2.0×10^{26}	5.4×10^{38}	1.7×10^{51}

grow rapidly restricting its use to only the smallest systems. Table 1.1 shows the number of Slater determinants required in a FCI calculation for increasingly large hydrocarbons. As one can see, the problem size grows by ~ 13 orders of magnitude for every 2 carbon and 2 hydrogen atoms added. Furthermore, the world has only recently entered the exascale era, i.e. the fastest supercomputer (Frontier) at the time of writing is only capable of 1.6×10^{18} floating point operations per second.¹¹ This means it would take $\sim 10^8$ seconds (> 3 years) to finish a single row of the matrix-vector product between the Hamiltonian matrix and CI wave function vector for the butadiene (C_4H_6) molecule. Consequently, many alternatives have been formulated, which balance the trade-off between computational cost and accuracy. The remainder of the section briefly introduces the zoo of correlated wave function methods.

1.2.1 Coupled-Cluster

We start by briefly discussing coupled-cluster (CC) methods. Coupled-cluster methods can recover large amounts of electron correlation in reasonable wall time for small- to medium-sized molecules. CC also benefits from size-extensivity, which means the absolute electronic energy scales appropriately with system size. These methods have seen such success in quantum chemical calculations to the extent that CC with singles, doubles, and perturbative triples (CCSD(T)) is often referred to as the “gold standard” of quantum chemistry.^{2,12–14} But as we know, the U.S. dropped the gold standard when it ended the Bretton Woods system in the 1970s.¹⁵ Similarly, CC methods tend to break down in systems where multiple electron configurations are needed to accurately represent the wave function (i.e. strong correlation), for example in bond dissociation.^{2,14} Consequently, coupled-cluster methods see the most usage in computing the thermochemistry of predominantly closed-shell molecular systems, i.e. in the weakly correlated regime.

1.2.2 Configuration Interaction Methods

Configuration interaction (CI) methods construct wave functions as linear combinations of Slater determinants. The use of multiple Slater determinants allows CI to have more flexibility in the placement of electron density. This is especially useful in the strongly correlated regime—for example in a homolytic bond dissociation, a single Slater determinant is unable to place one electron on each atom.^{1,2} Due to the scaling of FCI (See Table 1.1), all CI methods used in practice only include a small subset of possible electron configurations.

Textbooks commonly introduce CI methods by their excitation level, which indicates how many differences an electron configuration is allowed to have with a given reference determinant. For example, CI + single excitations (CIS) starts with a reference HF wave

function and includes all determinants which differ by one occupied orbital from the HF state.^{1,2} Truncation of the CI space in these ways results in non-size-extensive wave functions. Additionally, the results from truncated CI calculations are highly dependent on the choice of reference orbitals. However, many other variations of truncated CI methods exist.

One alternative is the spin-flip (SF) CI method, which utilizes a high-spin set of reference orbitals.^{16–21} The singly occupied MOs of the reference result in some orbital optimization of the first few virtual frontier orbitals and enables good descriptions of low-lying excited states.^{22,23} The restricted active space (RAS) SF method is used for this purpose in Chapter 2.

Other alternatives aim to provide systematic approximations that converge to FCI with simple parameter tuning. This includes quantum Monte Carlo (QMC) based CI methods,^{24–28} which randomly sample determinants from the FCI space, and select CI plus perturbation theory (SCI+PT) methods,^{29–46} which are discussed in detail in Chapter 3. In SCI+PT, electron correlation is recovered in two stages. The first stage “selects”—hence the name—the electron configurations that are needed to describe the most significant electronic interactions, i.e. the static correlation. This is done by iteratively selecting configurations that are considered “important” and including them in a variational wave function, ψ_{var} . The second stage applies perturbation theory (PT)—typically Epstein-Nesbet (EN) PT—to recover the remaining dynamic correlation. The EN correction is written

$$E_{\text{EN}} = \sum_k \frac{(\sum_i H_{ki} c_i)^2}{E_{\text{var}} - H_{kk}}, \quad (1.14)$$

where the index k is over determinants not in ψ_{var} , E_{var} is the variational energy, c_i is the coefficient of variational determinant i , H_{ki} is the Hamiltonian element connecting determinants k and i , and H_{kk} are the diagonal elements of the CI Hamiltonian. This equation is particularly difficult to parallelize due to the squaring of the inner sum over i , which forces the partial sums to be stored for each k . This is further compounded by the fact that there are, in practice, several orders of magnitude more perturbative determinants than variational ones. These issues introduce memory bottlenecks as well as difficulties with parallelization in the EN-PT problem. Computational strategies to address these difficulties are the topic of Chapter 3.

1.2.3 Multi-Configurational Self-Consistent Field

As previously mentioned, the results of CI calculations are highly dependent on the choice of reference orbitals. In multi-configurational self-consistent field (MCSCF) methods, the wave function is constructed using multiple Slater determinants, however the orbitals are also

variationally optimized together with the determinantal coefficients. The most commonly used MCSCF method is the complete active space self-consistent field (CASSCF) method.^{2,47–52} This method performs FCI in a subset of the MO space called the active space. Outside the active space, a HF density is used, which allows the problem of recovering electron correlation to be localized to the most important user-defined regions. The FCI expansion is also coupled with the orbital optimization procedure to ensure the best orbitals are used with the CI wave function. CASSCF is most useful for capturing strong correlation where multiple configurations are needed for a good description of the wave function, however it does not generally provide quantitative results and is restricted to maximum active spaces of 18 electrons in 18 orbitals, denoted (18e,18o).⁵²

Other MCSCF methods include the restricted active space SCF (RASSCF)^{53–57} and generalized active space SCF (GASSCF)^{58,59} methods, which further partition the orbitals into various active spaces with rules limiting the kinds of excitations allowed within and between active spaces. These are generalizations of the CASSCF method and allow electron correlation to be recovered in much larger overall active spaces, however the restriction on the primary active space size from CASSCF remains.

While RAS- and GAS-based methods aim to recover additional correlation outside the primary active space, some methods have been developed to approximate the underlying FCI expansion within a CAS. These include density matrix renormalization group (DMRG)^{60–66} and localized active space self-consistent field (LASSCF).^{67,68} With DMRG, the reduction in computational cost relies on some regular lattice-like structure to reduce the variational degrees of freedom, which does not generalize well. In LASSCF, the active space is split into localized, unentangled fragments allowing for each fragment to be treated independently. The exponential scaling then only applies to the individual fragment subspaces rather than the total active space of the system, however the possibility of fragments with large active spaces remains. The idea of splitting the orbital space into smaller subspaces is similarly used in incremental CASSCF (iCASSCF), where the subspaces are instead pairs of bonding/anti-bonding orbitals and the many-body expansion is used to compute incremental electron correlation corrections.^{69,70} The iCASSCF method has been shown to treat problems with active spaces containing over 80 active electrons and 80 active orbitals and is the topic of Chapter 4.

1.2.4 Incremental and Many-Body Expansion Methods

We will briefly discuss the many-body expansion (MBE), often called the method of increments when applied to quantum chemistry. The many-body expansion—especially in the context of FCI and CASSCF—has demonstrated very nice convergence properties toward

obtaining FCI and CASSCF energies as well as CASSCF gradients. In general, the MBE breaks down large problems into smaller, more manageable components and can be applied to calculate many physical quantities. For example, the MBE expression for the electronic energy can be written

$$E = E_{ref} + \sum_i \varepsilon_i + \sum_{j<i} \varepsilon_{ij} + \sum_{k<j<i} \varepsilon_{ijk} + \dots, \quad (1.15)$$

where indices i, j, k refer to individual bodies. Each ε_x in Equation 1.15 contains the correlation due to interactions among the bodies included in x . This approach can utilize any correlated method, such as CC,^{71,72} FCI,^{73–82} and CASSCF.^{69,70} The MBE applied to FCI has had multiple implementations, each with subtle differences. Gauss defined the bodies of the MBE to be virtual MOs in the system,^{76–78,83} Windus utilized groups of occupied valence orbitals to define individual bodies,⁸² and our group frequently uses valence bonding/anti-bonding pairs from perfect pairing to define the bodies.^{79–81} Regardless of choice of body, each approach computes the ε_x in Equation 1.15 by performing FCI within the orbitals contained in each body of x . While these approaches all converge to the FCI solution, the approach by Gauss requires expansion beyond 6-bodies to reach convergence.

The MBE is also fundamental to the iCASSCF method, however the previously mentioned MBE approaches in FCI all deal with the energy. Application of the MBE to CASSCF, on the other hand, requires some way of forming the generalized Fock matrix in order to compute the orbital gradient.^{69,70} The generalized Fock matrix can be defined by the 1- and 2-particle reduced density matrices (1-RDM, 2-RDM)

$$F_{mn} = \sum_q D_{mq} h_{nq} + \sum_{qrs} d_{mqrs} g_{nqrs}, \quad (1.16)$$

where D_{mn} and d_{mqrs} are the 1- and 2-RDMs, and h_{nq} and g_{nqrs} are the 1- and 2-electron integrals. Using this observation, the Fock matrix is generated by computing the 1- and 2-RDMs, which are computed incrementally via

$$D_{pq}^{tot} = D_{pq}^{ref} + \sum_i \Delta D_{pq}^i + \sum_{j<i} \Delta D_{pq}^{ij} + \dots, \quad (1.17)$$

where the ΔD_{pq}^x terms are changes in density due to interactions between the bodies in x . The analogous MBE for d_{pqrs} is similar. However, this construction of the Fock matrix in this way results in an electronic energy that is not invariant to active-active orbital rotations unlike the parent CASSCF theory. The original formulation of iCASSCF mended this problem by symmetrization of the active-active block of each incremental Fock matrix, which explicitly

enforced the expected symmetry.⁶⁹ However, this procedure effectively projects out some components of the orbital gradient, which is defined by

$$E_{pq}^o = 2(F_{pq} - F_{qp}), \quad (1.18)$$

thereby making the resultant iCASSCF orbital optimization not fully variational. Chapter 4 details the efforts to make iCASSCF a fully variational method.

1.2.5 Summary

The computational and theoretical advances in CI and MCSCF detailed in the following chapters will enable accurate, FCI-quality calculations at fractions of the cost. For example, a molecule like butadiene (C_4H_6)—which is beyond the reach of FCI—can be accurately computed in just a few hours with HCI as we will show in Chapter 3. Even larger systems, such as the oxoMn(salen)Cl molecule (containing 84 valence electrons!), can be handled in a matter of days on a single compute node by applying the method of increments. The advances in CI and MCSCF detailed in this thesis will be key to solving challenging electronic structure problems and reducing the necessary trade-off between accuracy and cost.

CHAPTER 2

An Open-Shell Coronoid with Hybrid Chichibabin-Schlenk Conjugation

This chapter is based on a previously published experimental collaboration²³ with emphasis on the computational results and methods. In this chapter, the spin-flip (SF) method is combined with a restricted active space (RAS) configuration interaction (CI) in order to compute low-lying excited states of an open-shell coronoid molecule. In addition, the spin alignment of the coronoid is analyzed using a spin Hamiltonian.

2.1 Introduction

Since their early discovery at the turn of the 20th century (**1-3**, Figure 2.1),⁸⁴⁻⁸⁶ open-shell organic molecules have attracted continued interest as a testing ground for theories of molecular bonding and electronic structure, and as an emerging class of organic materials.⁸⁷⁻⁹⁶ Polyradicals derived from π -extended aromatics play an important role in these developments, as evidenced by their recent use as semiconductors,⁹⁷ near-infrared dyes,^{98,99} cages,^{100,101} switches,^{102,103} and components for covalent self-assembly.¹⁰⁴⁻¹⁰⁹ Embedding radicaloid centers in a circular π -conjugated array provides a way of studying the interplay between the open-shell character and global (macrocyclic) aromaticity.^{96,110} The interactions between spins are particularly enhanced in fully fused cyclic systems, that is, open-shell circulenes¹¹¹ and coronoids.¹¹²⁻¹¹⁷ These systems are characterized by efficient p_z -orbital overlap and multiple conjugation pathways within their fused ring frameworks. They often combine unusual electronic structure characteristics with appreciable chemical stability. For example, Wu’s coronoid oligoradicals **4a,b** displayed unprecedented annulene-within-annulene (AWA) aromaticity that qualitatively depended on the size of the macrocycle.^{113,118} The fully conjugated [4]chrysaorene **5**, reported concurrently by Wu and co-workers¹¹⁴ and by our group,¹¹⁵ displayed rich redox chemistry coupled with anion binding in the macrocyclic cavity.

In the open-shell coronoids studied so far, the interactions of adjacent spins were either full equivalent, as in **4** and **5**, or topologically similar.^{116,117} It was reasoned that it might




Figure 2.1: Chichibabin and Schlenk conjugation in organic oligoradicaloids. Fully open-shell configurations shown for all all oligoradicaloids. Substituents not shown for **4a,b** and **5**. Unpaired electrons are presented as red dots or p orbitals.

be possible to create systems that alternate between two interaction types to produce spins arrays with a more complex internal structure. With this goal in mind, the coronoid system **8** was designed that combines two classic conjugation types, those found, respectively in Chichibabin^{85,119} and Schlenk-Brauns^{86,120} hydrocarbons (**2** and **3**, Figure 2.1). The Chichibabin conjugation (Kekulé-type⁹²) is characterized by an even-electron pathway between interacting spins, and is thus fundamentally different from the Schlenk conjugation (non-Kekulé), which features an odd-electron path. This difference, which is reflected in the ground states of **2** (singlet¹¹⁹) and **3** (triplet¹²⁰), creates an unusual structural dichotomy in the π -conjugated structure of **8**.

Hydrocarbons **2** and **3**, which contain only unfused benzene rings, have non-planar structures and are consequently relatively unstable. Robustness of such diradicaloids can be

enhanced by indene fusion,^{93,95} and steric protection as illustrated in indeno[2,1-*b*]fluorene **6**¹²¹ and in [1,2-*a*:2',1'-*i*]phenanthrene **7**,¹²² which feature, respectively the Schlenk- and Chichibabin-type conjugation embedded into a fused ring framework. In **8**, the radical centers are connected via alternating 2,7-phenanthrenylene and *meta*-phenylene subunits, which form a conjugation pattern that is simultaneously analogous to those found in **6** and in diindeno[2,1-*b*:1',2'-*h*]phenanthrene **9**. The latter ring system is an isomer of **7**, and likewise features a Chichibabin-like pathway.

The six radical sites of **8'** will result in near degeneracies in the six frontier orbitals (HOMO-2, ..., LUMO+2), which would be difficult for single Slater determinant methods to handle. However, multiconfigurational methods such as the complete active space self-consistent field (CASSCF) method are typically only used to treat static correlation. To accurately capture the energetics of the ground and low-lying excited states, recovery of dynamic correlation is also required. The restricted active space (RAS) spin-flip (SF) method,¹⁶⁻²¹ which utilizes a high spin (in this case septet) restricted open shell Hartree-Fock (ROHF) reference, is well-suited for handling the electronic structure of **8'**. The use of a high spin reference results in partial orbital optimization of the low-lying virtual orbitals. Furthermore, allowed excitations outside of the primary active space recovers some dynamical correlation. However, analytic gradients for RAS-SF do not exist, thus geometries must be obtained at some other level of theory.

The remainder of this chapter will be dedicated to the computational insights gathered regarding **8'**. Details of the synthesis and experimental characterization of **8'** can be found in reference 23.

2.2 Computational Details

Gas-phase geometries of the four feasible spin states of the substituent-free **8'**—singlet ¹**8'**, triplet ³**8'**, quintet ⁵**8'**, and septet ⁷**8'**—were optimized at two levels of theory: (a) dispersion¹²³ and range-corrected¹²⁴ GD3BJ-CAM-B3LYP/6-31G(d,p), and (b) CASSCF(6,6)/cc-pVDZ^{47,48} with the RICD approximation,^{125,126} denoted, respectively CAM and CAS in subsequent discussions. CAM calculations were run using Gaussian 16¹²⁷ while CAS calculations were run in OpenMolcas.¹²⁸ As a single reference method, CAM is not suitable for quantitative analysis of open-shell species, although it correctly predicts that different spin states of **8'** should have similar energies with preference for the singlet configuration. RAS(6,6)-SF/cc-pVDZ¹⁶⁻²¹ energies, denoted SF, were calculated for the eight lowest electronic states at each of the optimized CAS geometries.

Natural orbitals (NOs) and the corresponding occupation numbers (NOONs)¹²⁹ obtained from the SF densities are used to gather additional insight into the ground and excited states of $\mathbf{8}'$. From the SF densities, polyradicaloid indices (γ_i)¹³⁰ and number of unpaired electrons (n_U),¹³¹ derived from NOONs, were also computed. The spin alignments of $\mathbf{8}'$ were also analyzed by constructing a spin Hamiltonian.

2.3 Results and Discussion

Both CAM and CAS geometry optimizations converged to fully planar geometries for all spin states considered. At both levels of theory, the singlet and septet optimizations yielded D_{3h} -symmetric structures, while C_{2v} -symmetric structures were found for the triplet and quintet. Minor variations of C-C bond distances in ${}^3\mathbf{8}'$ and ${}^5\mathbf{8}'$ lead to reduction of point group symmetry from D_{3h} to C_{2v} , apparently caused by the Jahn-Teller effect.

Table 2.1: Computational results for adiabatic spin states of $\mathbf{8}'$.

State	${}^1\mathbf{8}'\text{-S}_0$	${}^3\mathbf{8}'\text{-T}_1$	${}^5\mathbf{8}'\text{-Qn}_1$	${}^7\mathbf{8}'\text{-Sp}_1$
Point group symmetry	D_{3h}	C_{2v}	C_{2v}	D_{3h}
ΔE_{CAM} (kcal mol ⁻¹)	0.0	1.6	0.8	0.1
ΔE_{SF} (eV)	0.000	0.043	0.084	0.129
ΔE_{SF} (kcal mol ⁻¹)	0.0	1.0	1.9	3.0
n_U	4.746	5.229	5.633	6.000

While CAM is not suitable for quantitative analysis of open-shell species, it correctly predicts the various spins of $\mathbf{8}'$ should have similar energies, with preference for a singlet configuration. At each CAS geometry, RAS-SF/cc-pVDZ energies (E_{SF}) were calculated at each of the CAS optimized geometries (See Table 2.1 and Figure 2.2A) predicting the ground state to always be a singlet (S_0 , $S^2 = 0$), followed consecutively by three triplets (T_1 through T_3 , $S^2 = 2$), two quintets (Qn_1 and Qn_2 , $S^2 = 6$), one septet (Sp_1 , $S^2 = 12$) and one singlet state (S_1).

For the D_{3h} -symmetric structures (${}^1\mathbf{8}'$ and ${}^7\mathbf{8}'$), the two lowest-energy triplet excited states T_1 and T_2 where degenerate with analogous degeneracies found for the Qn_1 and Qn_2 quintets. These degeneracies are indicative of spin frustration,¹³² and account for the Jahn-teller distortion found for the ${}^3\mathbf{8}'$ and ${}^5\mathbf{8}'$ geometries at the CAM and CAS levels.

For each multiplicity, the lowest SF energy was identified at the corresponding CAS geometry (S_0 at ${}^1\mathbf{8}'$, T_1 at ${}^3\mathbf{8}'$, etc., see Table 2.1 and Figure 2.2A) confirming the consistency between the SF and CAS approaches. Relative SF energies of the adiabatic states ${}^1\mathbf{8}'$, ${}^3\mathbf{8}'$, ${}^5\mathbf{8}'$, and ${}^7\mathbf{8}'$ are thus 0.0, 1.0, 1.9, and 3.0 kcal mol⁻¹, that is, the energy rises by nearly a constant value of ~ 1 kcal mol⁻¹ for each increase in multiplicity. This is smaller than the

SF-based estimates from ΔE_{ST} obtained for $\mathbf{5}'$ ($-6.6 \text{ kcal mol}^{-1}$)¹¹⁵ and $\mathbf{7}'$ (R=H, $-1.3 \text{ kcal mol}^{-1}$).¹²²




Figure 2.2: **A**) SF energies and spins for the first eight states of $\mathbf{8}'$. The lowest-energy states for a given spin multiplicity is indicated with an arrow. The singlet geometry $^1\mathbf{8}'$ has a 6-fold degeneracy at 0.156 eV consisting of three singlets, two triplets, and one septet state. **B**) Relationship between ΔE_{SF} and the number of unpaired electrons n_U derived from natural orbital occupation numbers. **C**) Frontier natural orbitals for the $^1\mathbf{8}'$ - S_0 state (0.02 a.u. isosurface). **D**) SF odd-electron density for $^1\mathbf{8}'$ (0.002 a.u. isosurface). **E**) Definition of coupling constants used in an approximate spin Hamiltonian. J values are given for the $^1\mathbf{8}'$ state.

Natural orbitals (NOs) and the corresponding occupation numbers (NOONs) are shown in (Figures A.1-A.8 in Appendix A). For all these states, amplitudes of the frontier NOs may be approximated as linear combinations of singly occupied molecular orbitals corresponding to the three embedded *m*-xylylene⁹² fragments. This feature, which is most evident in the HONO and LUNO of ¹**8'**-S₀ (Figure 2.2C), suggests a Schlenk-like behavior of the unpaired electrons in the system. A similar picture is provided by the SF odd-electron density determined for ¹**8'**-S₀ (Figure 2.2D). The NOONs of ¹**8'**-S₀ yielded high polyradicaloid indices of $\gamma_0 = 0.826$ and $\gamma_1 = \gamma_2 = 0.773$, confirming the open-shell character of the singlet state. In comparison, indices based on the CAM density were $\gamma_0 = 0.984$ and $\gamma_1 = \gamma_2 = 0.559$. The polyradicaloid indices for the singlet ground state of all geometries is shown in Figure 2.3. The polyradicaloid indices ($\gamma_0, \gamma_1, \gamma_2$) for the S₀ state of **8'** increase as the geometry of **8'** is optimized for increasing spin manifolds. Furthermore, for the *D*_{3h} geometries, $\gamma_1 = \gamma_2$ while the *C*_{2v} geometries have $\gamma_1 > \gamma_2$ indicating that symmetry has a strong effect on the calculation of polyradicaloid indices. Similarly, the number of unpaired electrons n_U is very

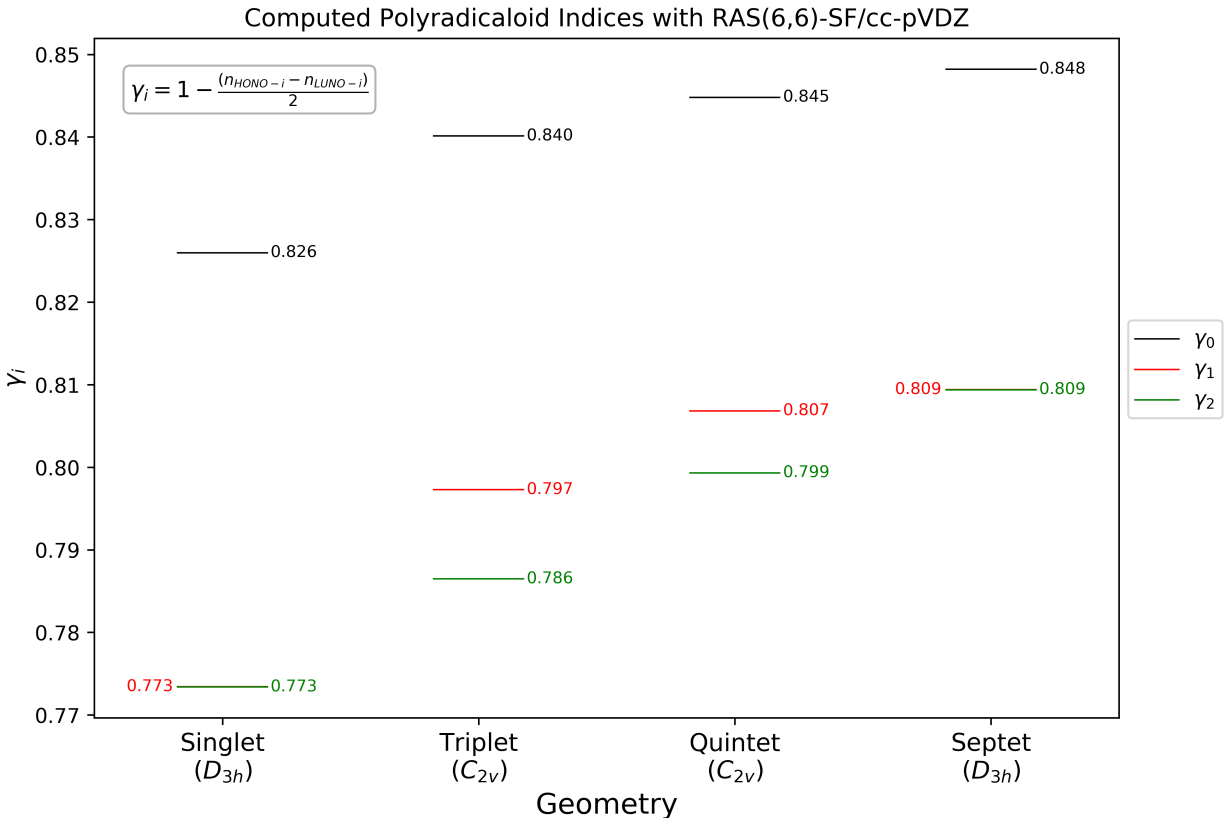


Figure 2.3: Polyradicaloid indices (γ_i) from RAS(6,6)-SF/cc-pVDZ are shown for the ground singlet electronic state at various optimized CASSCF(6,6)/cc-pVDZ geometries.

high for ¹**8'**-S₀ ($n_U = 4.746$, see Table 2.1). The n_U values of excited states increase with

their increasing multiplicity and are larger than the total spin ($n_U > 2S$), except for the septet (Sp_1) configurations. In particular, the n_U index differentiates excitations with the same S . For instance, n_U is 5.18 for ${}^1\mathbf{8}'\text{-T}_1$ and ${}^1\mathbf{8}'\text{-T}_2$, and 5.3 for ${}^1\mathbf{8}'\text{-T}_3$. Interestingly, for each CAS geometry, there is an excellent linear correlation between the SF energy, Δ_{SF} , and the n_U index (See Figure 2.2B). The slope of this dependence (0.124 eV/electron or 2.9 kcal mol⁻¹/electron for ${}^1\mathbf{8}'$) provides an alternative measure of electron pairing energy that considers the multiconfigurational nature of all states.

The spin alignment of each ground singlet electronic wave function were analyzed to explain the behavior of the energetics and polyradicaloid indices. This was done by analyzing the radical sites and spin Hamiltonian (See Figure 2.4). The radical sites were identified (labeled in Figure 2.4A,B) with the largest odd-electron density.

The analysis begins with the spin Hamiltonian for the D_{3h} geometries, where $\varepsilon_i = 0$ for all i . Diagonalization of the D_{3h} spin Hamiltonian produces the following analytic expressions for the eigenvalues:

$$E_0 = J_1 + J_2 + 2J_3 + J_4, \quad (2.1)$$

$$E_1 = E_2 = -J_3 - \sqrt{J_1^2 - J_1J_2 - J_1J_4 + J_2^2 - J_2J_4 + J_4^2}, \quad (2.2)$$

$$E_3 = E_4 = -J_3 + \sqrt{J_1^2 - J_1J_2 - J_1J_4 + J_2^2 - J_2J_4 + J_4^2}, \quad (2.3)$$

$$E_5 = -J_1 - J_2 + 2J_3 - J_4. \quad (2.4)$$

To compute J_1 and J_2 , the spin gap error between the D_{3h} eigenvalues and the RAS(6,6)-SF spin states are minimized, i.e. we minimize the following equation:

$$R = \sqrt{\frac{\sum_{i=1}^5 (E_i^{SF} - E_i)^2}{5}}, \quad (2.5)$$

where E_i^{SF} corresponds to the RAS(6,6)-SF energies and the ground state eigenvalue is assumed to be zero. Due to the difference in dimension of the spin Hamiltonian and the RAS(6,6)-SF spin Hamiltonian, it is important to select the appropriate RAS(6,6)-SF spin states which have the corresponding spin Hamiltonian eigenvalues. Since the non-degenerate triplet state (T_3) is higher in energy than the doubly degenerate set, it may be considered an excited triplet, which is consistent with the higher number of unpaired electrons, n_U . Furthermore, since there are two sets of doubly degenerate eigenvalues in the spin Hamiltonian (one corresponding to a set of triplets and the other a set of quintets), Equation 2.5 is minimized




Figure 2.4: **A**) The spin Hamiltonian is shown. **B**) The odd-electron density (isovalue=0.005 a.u.) of the optimized $1\mathbf{8}'$ geometry with the six radical sites labeled. **C**) The coupling diagram corresponding to the spin Hamiltonian is drawn. For the D_{3h} geometries, $\varepsilon_1 = \varepsilon_2 = 0$.

against the sets of doubly degenerate triplets and quintets from RAS(6,6)-SF in each D_{3h} geometry. The resulting couplings and RMSE are reported in Table 2.2.

Table 2.2: Computed spin couplings for $\mathbf{8}'$ (eV). J_i are computed spin coupling constants, and ε_i are perturbations to spin couplings due to the C_{2v} geometry.

Geometry	J_1	J_2	J_3	J_4	ε_1	ε_2	ε_{3a}	ε_{3b}	ε_4	RMSE
Singlet (D_{3h})	-0.0382	-0.0302	-0.0008	-0.0097	0	0	0	0	0	1.70×10^{-7}
Triplet (C_{2v})	-0.0322	-0.0262	-0.0011	-0.0081	0.0101	-0.0097	-0.0007	0.0003	0.0042	6.31×10^{-9}
Quintet (C_{2v})	-0.0295	-0.0240	-0.0011	-0.0076	-0.0053	-0.0095	-0.0004	0.0004	-0.0017	1.03×10^{-9}
Septet (D_{3h})	-0.0277	-0.0233	-0.0012	-0.0073	0	0	0	0	0	5.25×10^{-7}

One might consider the possibilities of J_i having different signs and could start the minimization of Equation 2.5 at various points to find other minima with different signs for J_i , however these result in RMSE several orders of magnitude higher than those reported in Table 2.2. This suggests that the J_i should all have the same sign. In the case of all the J_i being positive, the RMSE remains the same, however the eigenvalue $-J_1 - J_2 + 2J_3 - J_4$ corresponds to the septet state, which would consequently place the septet as the ground state. Therefore, the assignment of negative sign for J_i in the D_{3h} case gives antiferromagnetic

coupling for the ground state. The S_0 eigenvector of the D_{3h} spin Hamiltonian has uniform magnitudes indicating a single spin alignment.

A similar analysis can be applied to the C_{2v} geometries, but the eigenvalue decomposition of the spin Hamiltonian for the C_{2v} geometries must be done numerically. For the C_{2v} geometries, the two triplet states lowest in energy (See Figure 2.2A) are chosen for minimizing Equation 2.5. The C_{2v} Hamiltonian is minimized in two separate steps. First, it is treated with a D_{3h} Hamiltonian to obtain the initial J_i , then Equation 2.5 is minimized to compute the ε_i with fixed J_i . In other words, the C_{2v} Hamiltonian is treated as a perturbed D_{3h} Hamiltonian. The resulting J_i and ε_i are reported in Table 2.2 along with couplings from the D_{3h} geometries.

Unlike the D_{3h} case, the singlet eigenvector for each of the C_{2v} spin Hamiltonians has non-uniform magnitudes indicating multiple spin alignments are involved. By pairing the spins across the vertical mirror plane (i.e. pairs 1-2, 3-6, 4-5 with ordering from Figure 2.4B,C), the overall spin system can be treated as a trimer of paired spins and the singlet eigenvector for the two C_{2v} spin Hamiltonians can then be decomposed as a linear combination of multiple spin products.¹³³ The resultant singlet wave function at the triplet and quintet geometries are then, respectively,

$$\psi_3 = 0.600\psi_{\uparrow\downarrow\uparrow\downarrow\downarrow} + 0.292\psi_{\downarrow\uparrow\uparrow\downarrow\downarrow} + 0.212\psi_{\uparrow\downarrow\downarrow\uparrow\uparrow} + 0.095\psi_{\uparrow\downarrow\uparrow\uparrow\downarrow}, \quad (2.6)$$

$$\psi_5 = 0.607\psi_{\uparrow\downarrow\uparrow\downarrow\downarrow} + 0.136\psi_{\downarrow\uparrow\uparrow\downarrow\downarrow} + 0.198\psi_{\uparrow\downarrow\downarrow\uparrow\uparrow} + 0.274\psi_{\uparrow\downarrow\uparrow\uparrow\downarrow}, \quad (2.7)$$

where the radical site ordering follows that of Figure 2.4. In other words, the ground state singlet wave functions for the C_{2v} geometries are superpositions of multiple singlet spin alignments in this spin Hamiltonian model. This spin interaction model suggests that all spin pairs in $\mathbf{8}'$ are effectively antiferromagnetic in character.

2.4 Conclusions

This chapter describes the insights gathered from computational methods in the first example of an open-shell coronoid molecule in which Kekulé and non-Kekulé conjugation pathways are juxtaposed in a cyclic oligoradical array. Locally, this system shares features of the Schlenk and Chichibabin hydrocarbons, to which it is structurally related, but it also displays global characteristics resulting from macrocyclic conjugation. In particular, it has a singlet ground state with highly multiconfigurational character as well as three high-spin states—triplet, quintet, and septet—that have low energies and can be significantly populated at room temperature. In addition, the spin pairs of this new coronoid are predicted to

be predominantly antiferromagnetic at room temperature. These computational insights complement the experimental results gathered in reference 23.

CHAPTER 3

Advances in Parallel Heat Bath Configuration Interaction

The previous chapter showcased an application of the RAS-SF method, which can treat static correlation in a very limited active space (up to (8e,8o)) as well as some dynamic correlation from select excitations outside of the active space. This chapter explores the Heat-bath configuration interaction (HCI) method, which is a deterministic method that approaches the FCI limit at greatly reduced computational cost. In this chapter, computational improvements to the existing HCI algorithm are discussed targeting speed, parallel efficiency, and memory requirements. This new implementation introduces a hash function to distribute determinants and takes advantage of MPI and OpenMP for parallelism allowing for a (22e,168o) active space to be explicitly correlated. This chapter is based on work published previously in the *Journal of Physical Chemistry A*.¹³⁴

3.1 Introduction

Wave function simulations provide a wealth of essential insight into diverse electronic structures. While attaining qualitatively correct wave functions is usually possible, reaching chemically accurate energies can be intractably difficult.² The exact solution to the electronic Schrödinger equation can be constructed using full configuration interaction (FCI), which represents the wave function as a linear combination of all possible occupancies of electrons in the full molecular orbital space. The combinatorial number limits the application of conventional FCI to only the smallest chemical systems.

Alternatives to FCI have sought to maintain good accuracy and achieve tractable computational cost. For molecules where a single electron configuration dominates the wave function, computationally tractable, accurate alternatives exist.² These methods—such as coupled cluster theory—tend to break down in strongly correlated systems where multiple configurations are required for a qualitatively correct wave function, for example in bond dissociation.¹⁴ The complete active space (CAS)^{47–49} method, which performs a full CI expansion within the set of active orbitals, performs much better when strong correlation

is present, but is not chemically accurate in general. These limitations have led to the development of a number of alternatives, such as quantum Monte Carlo (QMC) CI,^{24–28} density-matrix renormalization group (DMRG),^{60–66} and select CI plus perturbation theory (SCI+PT) methods.^{29–39} Among the SCI+PT methods is heat-bath CI (HCI)⁴⁰ and its semi-stochastic variant (SHCI).^{41–46} These methods (QMC-CI, DMRG, and SCI+PT) all share the advantages of being systematically improvable approximations to FCI with much more bearable computational burden.

This chapter’s focus will be on HCI, though the insights herein apply equally well to all SCI+PT methods. SCI+PT employs two simple strategies to converge and extrapolate to the full CI energy.^{40,43} The first strategy is to include all electron configurations that are needed to describe the most significant electronic interactions. This is done by (iteratively) selecting configurations that are deemed “important” and including them in a variational wave function that exists within a subspace of the full Hamiltonian. The second strategy is to apply perturbation theory (PT), namely Epstein-Nesbet (EN) PT,^{135,136} to recover the missing weak correlation energy. The EN correction can be written

$$E_{\text{EN}} = \sum_k \frac{(\sum_i H_{ki} c_i)^2}{E_{\text{var}} - H_{kk}}, \quad (3.1)$$

where the index k is over determinants not present in the variational wave function, E_{var} is the variational energy, c_i is the coefficient of variational determinant i , H_{ki} is the Hamiltonian element connecting determinants i and k , and H_{kk} are the diagonal elements of the Hamiltonian. This equation is nontrivial to evaluate in parallel due to squaring the inner sum over i . The square forces the partial sums to be stored for each perturbative determinant, and there are orders of magnitude more of these than in the variational step.

One of the first SCI+PT methods, configuration interaction by perturbatively selecting iteratively (CIPSI),^{29,30} expanded the variational wave function by iteratively adding in configurations that are large contributors to the perturbative energy. In CIPSI, any configuration with a PT energy contribution above a threshold is added to the variational space. The related adaptive sampling CI (ASCI)^{35,36,38} similarly uses first-order perturbation coefficients to select electron configurations, but keeps only a fixed number to avoid growing the variational problem too large. HCI simplifies the variational search and selection criteria by only checking Hamiltonian elements of configurations connected to the wave function, instead of performing the full PT step.⁴⁰ A similar strategy is applicable to the (expensive) PT stage (see Theory and Implementation Details), providing an additional computational cost savings. Regardless, the high memory costs of the PT step led to introduction of (semi-)stochastic PT in HCI,^{41,42,44,46} which reduces memory burden at the cost of some sampling errors. Our

group has made practical use of HCI as a CI solver for many-body expansions of the FCI problem^{69,70,79–81} as well as a data source for inverse DFT.^{9,10} Here, the FCI total energy must be converged to tight tolerances ($\sim 10\mu\text{Ha}$), which motivated the continued development of HCI, rather than its stochastic variants, as a workhorse for these problems.

Significant efforts have been made in the SCI+PT realm to efficiently generate the variational wave function and compute the perturbative correction, and these inform the present developments in HCI. In the variational step of ASCI, a ranking algorithm—originally introduced in the context of FCI-QMC²⁵—was used to compute the first-order perturbation coefficients that are likely to generate important connections. Additionally, the first-order perturbation theory contributions are partially sorted for maximum cache efficiency during the variational stage.³⁸ In SHCI, hash tables paired with auxiliary index arrays are used to speed up the construction of the variational Hamiltonian.⁴⁴ For the perturbative step, ASCI uses a partial sorting algorithm that is similar to its variational stage to speed up the calculation of the PT energy.³⁸ In HCI, Equation 3.1 is approximated by only including contributions to the PT energy that are larger than some threshold in the inner sum.⁴⁰ In SHCI, the stochastic PT correction substantially reduces computational cost and memory overhead since the stochastic sampling only requires contributions from a small fraction of the perturbative determinants.⁴⁴ SHCI also uses hash tables to store the partial inner sums to avoid the computational cost of sorting as in ASCI. To further reduce memory footprint in SHCI, the determinants in the perturbative space are hashed allowing contributions to the ENPT energy to be split into independent chunks, which was also proposed, although not implemented, by Tubman *et al.*³⁸ It is difficult to compare the various computational strategies as these have been applied to different underlying algorithms. Nevertheless, these optimizations have enabled the accurate treatment of large active spaces in reasonable wall times.

However, even with these computational advances, all SCI+PT methods have high computational costs when enough electrons are in play. Therefore ASCI and SHCI have been optimized to run in parallel environments and take advantage of modern computational resources. In ASCI and SHCI, Hamiltonian construction is parallel across rows. ASCI also uses freely available parallel sorting libraries in its implementation.³⁸ In SHCI, the partial sums of Equation 3.1 are stored in lock-free hash tables to minimize communication overhead.⁴⁴ These parallel advances have allowed SHCI to treat wave functions with more than 10^9 variational determinants and its associated perturbative correction.

This work presents a reconceived implementation of deterministic HCI, which combines new algorithmic strategies inspired by those discussed by Li *et al.* and Tubman *et al.* The largest calculation showcases the $[\text{FeO}(\text{NH}_3)_5]^{2+}$ complex with an active space of (22e,168o).

This allowed HCI to explicitly treat 2.39×10^7 variational determinants and 8.95×10^{10} perturbative determinants. This calculation is the first attempt to treat a transition metal complex with an active space of this size using SCI+PT.

The organization of this chapter is as follows. In the Theory and Implementation Details, the HCI algorithm is briefly reviewed, followed by a discussion of key data structures and parallelization considerations. The Results and Discussion section presents parallel scaling and accuracy benchmarks, examines cyclobutadiene automerization, and computes the quintet-triplet spin gap of $[\text{FeO}(\text{NH}_3)_5]^{2+}$.

3.2 Theory and Implementation Details

The HCI algorithm in this work was implemented in a development version of QChem in C++ with MPI for multi-node parallelization and OpenMP for thread-level parallelization. HCI has been thoroughly described in previous works,^{40–46} thus we only briefly summarize the algorithm in this section before continuing to detail the computational implementation.

Like other SCI+PT methods, HCI consists of two steps: generation of a variational wave function and the computation of the perturbative energy correction. For HCI, two parameters (denoted ε_1 and ε_2) are used to tune accuracy.

3.2.1 The Variational Step




Figure 3.1: Heatmaps of the CI Hamiltonian at each iteration in the variational step of a HCI calculation are shown. The submatrices of prior HCI iterations are boxed in green, red, and black for the first, second, and third iterations, respectively. Darker colored cells indicate larger magnitude Hamiltonian elements.

In the variational step, a guess wave function is first generated, which can be a single Slater determinant or a small CI expansion. At each variational iteration of HCI, new determinants

are added using the heat-bath selection criteria,

$$\max(|H_{ki}c_i|) > \varepsilon_1, \tag{3.2}$$

where the index i runs over determinants in the variational wave function and c_i denotes the coefficients of the variational CI vector. In other words, a determinant is added if it is connected to a determinant in the current wave function by a significant coupling, weighted by the variational CI coefficient. The variational stage ends when the number of new determinants in a step is less than 1% of the number of determinants in the current wave function, although other termination criteria can be used.⁴³ The variational HCI iteration can be reduced to three primary steps:

1. Find all determinants not in the wave function where $|H_{ki}c_i| > \varepsilon_1$ for at least one determinant in the wave function,
2. Add these determinants to the wave function,
3. Augment the Hamiltonian and compute the new wave function,

until the termination criteria are satisfied (see Computational Details). Figure 3.1 shows a heatmap demonstrating how the Hamiltonian is augmented with every HCI iteration.

3.2.2 The Perturbative Step




Figure 3.2: The heatmap shows the variational Hamiltonian (left) from Figure 3.1 as well as the terms $H_{ki}c_i$ for the perturbative determinants (right). On the left, each column represents a determinant in the variational space, while each column on the right represents a perturbative determinant. The inner sum in Equation 3.1 runs over columns in this figure. Darker colors indicate larger magnitudes.

In HCI, the perturbative step uses a modification of EN PT to recover the remaining correlation energy outside of the variational space. This energy is closely related to Equation

3.1, but only includes terms larger than some threshold, ε_2 , in the inner sum,⁴⁰

$$E_{\text{EN}} \approx \sum_k \frac{\left(\sum_i^{|H_{ki}c_i| > \varepsilon_2} H_{ki}c_i\right)^2}{E_{\text{var}} - H_{kk}}. \quad (3.3)$$

Naturally, ε_2 is orders of magnitude smaller than ε_1 . Figure 3.2 shows a heatmap of the perturbative terms, $H_{ki}c_i$. Strategies to mitigate the computational challenge of enumerating over and storing contributions to Equation 3.3—where the memory requirement grows quickly with system size—are discussed in the remainder of this section.

3.2.3 Data Structures

The key data structures in this implementation are described for the variational and perturbative steps. The selection and design of these data structures take into account discussions by Tubman *et al.* regarding ASCII³⁸ and borrows ideas from prior work on SHCI by Li *et al.*; the differences will be noted.

The determinants are stored as bit strings using what Tubman *et al.* call the standard representation. In HCI, the α and β bit strings are stored separately within a C++ struct. This struct also contains a 32-bit integer denoting the index of the determinant in the Hamiltonian and coefficient vector. Li *et al.* likewise use bit-strings to represent determinants, though they stored the bit-strings as dynamic arrays of unsigned 64-bit integers to allow the orbital count to be assigned at run-time. The present implementation represents the bit-strings using `std::bitset` from the C++ standard template library (STL). While `std::bitset` has the downside of having a fixed length (and therefore fixed maximum number of orbitals) at compile time, a reasonable choice is unlikely to introduce memory or computational bottlenecks for small calculations, and the larger calculations are the main concern here. Furthermore, statically sized data structures can benefit from several compiler optimizations, especially when passed into a hash function, which is one of the most numerous operations in this article’s HCI implementation.

Similarly to Li *et al.*’s work, the Hamiltonian is stored in sparse format. Each node keeps its own local copy of the diagonal of the Hamiltonian. Since the Hamiltonian is symmetric, only the non-zero, upper triangular elements of the Hamiltonian are stored. For parallel memory balancing, each node stores Hamiltonian rows corresponding to determinants which have hashes falling within its assigned ranges. To do this, a modified version of the FNV hash function¹³⁷ is utilized (see Appendix B).

The use of hash tables to store the partial sums as well as partitioning the perturbative space into hash ranges are likewise borrowed from Li *et al.*⁴⁴ The hash table implementation

used in our work is the Parallel Hashmap.¹³⁸ In the prior SHCI algorithm, each node is required to store a hash table of partial sums for each of the other nodes, which necessitates an expensive merge of hash tables. With p hash tables per node, where p is the number of nodes, there is a possibility of multiple threads attempting to update the same hash table at the same time. To avoid these collisions, Li *et al.* utilize thread-local tables that are periodically merged with the process-level tables. While this improves parallel performance, it also introduces code complexity. In the present article’s implementation, this is avoided by making the problem embarrassingly parallel and avoiding merge altogether. Furthermore, the use of hash tables allows the variational Hamiltonian to be cleared from memory prior to beginning the perturbative step, resulting in a large memory savings. This will be discussed in the next section along with the parallelization details.

3.2.4 Parallelism and Algorithmic Optimization

MPI+OpenMP are used to parallelize this HCI implementation, and the key points for parallelization are detailed in this section. To facilitate discussion of the implementation, Table 3.1 lists the notation used to describe various data structures.

The variational step begins with an initial CAS wave function up to (8e,8o) in size. The Hamiltonian for the guess is stored in dense format and diagonalized exactly, which is done separately on each MPI process. With at most 4900 determinants for an active space of (8e,8o), the guess is inexpensive. Once the guess wave function is generated, the initial Hamiltonian is cleared from memory.

To find connected determinants, iterations over each determinant in the variational wave function generate and screen new determinants according to Equation 3.2. In this step, the determinants are statically distributed across MPI ranks in a round-robin fashion and dynamically load balanced within each MPI process across OpenMP threads. Each thread tracks its own set of new determinants, which are subsequently merged and sorted (removing duplicates) within each MPI processes. A call to `MPI_Allgatherv` then communicates all determinant lists between MPI ranks. Each MPI rank then once again sorts and removes duplicates from its own list. The new determinants are then merged into the wave function. At this point, each MPI rank has the same set of determinants that have the same indexing across processes. Algorithm 3.1 shows pseudocode demonstrating the load balancing and communication for adding new determinants as well as key algorithms used from the C++ STL. Notably, the loop over $\{d\}^{\text{var}}$ in Algorithm 3.1 only needs to occur over the new determinants from the previous HCI iteration.

Table 3.1: The notation for various data structures and their descriptions are listed below.

Notation	Description
$\{d\}^{\text{var}}$	The set of variational determinants
d	An arbitrary determinant.
d_i	An arbitrary determinant with index i .
h	Hash function. Takes in an α - or a β -string.
$H(d_i, d_j)$	Hamiltonian element between determinants d_i, d_j .
$c(d)$	CI coefficient for determinant d .
R	2D hash range. R_α and R_β will denote the range for α - and β -strings, respectively.
α, β	Arbitrary α - and β -string, respectively
${}^d\alpha, {}^d\beta$	α -string and β -string of determinant, d , respectively.
α_i^a, β_i^a	Singly excited α -string and β -string, respectively.
$\alpha_{ij}^{ab}, \beta_{ij}^{ab}$	Doubly excited α -string and β -string, respectively.
S^d	Set of singly excited determinants connected to d .
D^d	Set of doubly excited determinants connected to d .
S_α^d, S_β^d	Set of singly excited α -strings and singly excited β -strings for determinant d , respectively.
D_α^d, D_β^d	Set of doubly excited α -strings and doubly excited β -strings for determinant d , respectively.
d_i^a	Determinant d with α excitation from $i \rightarrow a$.
\bar{d}_j^b	Determinant d with β excitation from $j \rightarrow b$
M	Hash table.

Algorithm 3.1 Pseudocode for adding new determinants in the variational step

```
1: newdets = new vector<determinant>[nomp]
2: #pragma omp parallel for schedule(dynamic)
3: for  $d_i$  in  $\{d\}^{\text{var}}$  ▷ stride = MPI size
4:     for  $d_k$  in  $S^{d_i}, D^{d_i}$ 
5:         if  $|H(d_k, d_i)c(i)| > \varepsilon_1$ 
6:             newdets[n].push_back( $d_k$ )
7: merge(newdets)
8: sort(newdets)
9: unique(newdets)
10: MPI_Allgatherv(newdets)
11: sort(newdets)
12: unique(newdets)
13: merge(dets, newdets)
```

As previously mentioned, determinants are stored as pairs of α - and β -strings. This allows for hashing the α - and β -strings separately to produce a 2D hash, which in turn facilitates load balancing in the variational and perturbative steps. Figure 3.3 shows how a determinant gets assigned to a MPI rank based on the 2D hash for an example with nine MPI processes.

During the variational step, the Hamiltonian rows are assigned to MPI ranks using the 2D hash ranges described above. The Hamiltonian is then constructed by distributing rows across threads using dynamic scheduling within MPI processes. The subsequent matrix multiplication to form σ in the Davidson algorithm^{139,140} works across the same row division. Within a MPI process and during the sparse matrix-vector multiplication, each OpenMP thread stores its own σ to prevent collisions and to allow lock-free operations. These thread-local σ vectors are then merged and collected to the head MPI process which handles the (trivial) Davidson subspace diagonalization.

Like in the variational step, the determinants in the perturbative step are hashed over their α and β strings into a 2D hash, which is similarly distributed among MPI ranks based on predetermined α - and β -hash ranges. The 2D hash allows the PT step to be made embarrassingly parallel, as the PT determinants are predistributed by the hash across MPI ranks. Within each MPI rank, the 2D hash is further divided among the threads. Each thread then only needs to store the partial sums for determinants that lie in its assigned hash range. These smaller ranges are dynamically scheduled, which effectively distributes the outer sum of Equation 3.3 among the threads. Furthermore, partitioning of the perturbative step has the added advantage of allowing the perturbative space to be split into smaller bins as necessary to reduce the memory overhead for large calculations. The decomposition of the perturbative space into hash ranges is also used in SHCI,⁴⁴ however in SHCI, the perturbative

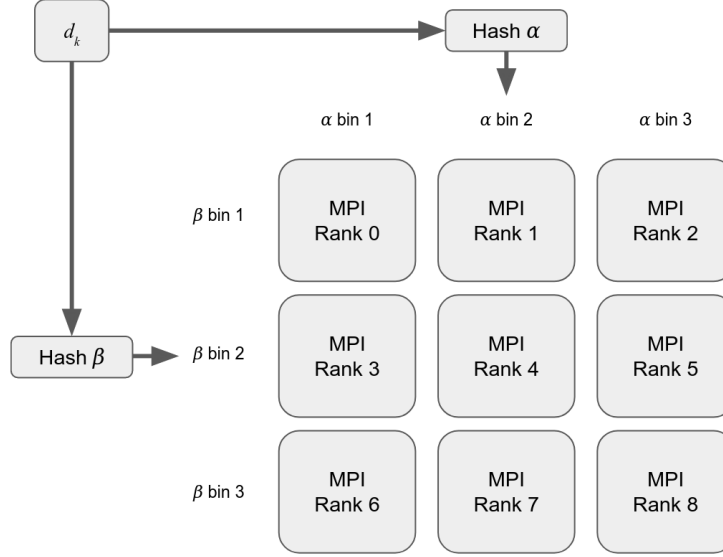


Figure 3.3: This schematic shows an example of how an arbitrary determinant, d_k , is assigned to a MPI rank. In this example, there are 9 MPI ranks involved in the calculation. The determinant is hashed by its α - and β -strings separately and then assigned based on predetermined ranges or bins for both the α and β hashes.

step is parallel within a hash range, i.e. all processes and threads work on the same hash range simultaneously before proceeding to the next.

Algorithm 3.2 Pseudocode for perturbative single excitations for a single OpenMP thread operating on hash range, R .

```

1: for  $d$  in  $\{d\}^{\text{var}}$ 
2:    $S_\alpha = \{\alpha_i^a : \alpha_i^a \in S_\alpha^d \text{ and } h(\alpha_i^a) \in R_\alpha\}$  ▷ cached for  $d$  with same  $\alpha$ 
3:   if  $h(d\beta)$  in  $R_\beta$ 
4:     for  $\alpha_i^a$  in  $S_\alpha$ 
5:       if  $|H(d_i^a, d)c(d)| > \varepsilon_2$ 
6:          $M(d_i^a) += H(d_i^a, d)$ 
7:   if  $h(d\alpha)$  in  $R_\alpha$ 
8:      $S_\beta = \{\beta_i^j : \beta_i^j \in S_\beta^d \text{ and } h(\beta_i^j) \in R_\beta\}$ 
9:     for  $\beta_i^j$  in  $S_\beta$ 
10:      if  $|H(d_i^j, d)c(d)| > \varepsilon_2$ 
11:         $M(d_i^j) += H(d_i^j, d)c(d)$ 

```

Algorithm 3.2 details pseudocode for a single OpenMP thread computing the perturbative partial sums due to singles excitations. The array S_α is only generated once per α string. This is possible since the variational determinant list is sorted in blocks allowing S_α to be kept until the α string changes in the outer loop of Algorithm 3.2. This lowers the number of times perturbative determinants must be generated. A similar algorithm is used for the α, α

and β, β doubles. Algorithm 3.3 shows pseudocode for generating the mixed α/β perturbative doubles, which can utilize a similar caching strategy for S_α as in Algorithm 3.2.

Algorithm 3.3 Pseudocode of the perturbative step for mixed α/β double excitations for a single OpenMP thread operating on hash range, R .

```

1: for  $d$  in  $\{d\}^{\text{var}}$ 
2:    $S_\alpha = \{\alpha_i^a : \alpha_i^a \in S_\alpha^d \text{ and } h(\alpha_i^a) \in R_\alpha\}$  ▷ cached for  $d$  with same  $\alpha$ 
3:    $S_\beta = \{\beta_i^j : \beta_i^j \in S_\beta^d \text{ and } h(\beta_i^j) \in R_\beta\}$ 
4:   for  $\alpha_i^a$  in  $S_\alpha$ 
5:     for  $\beta_i^j$  in  $S_\beta$ 
6:       if  $|H(d_{ij}^{a\bar{b}}, d)c(d)| > \varepsilon_2$ 
7:          $\mathbf{M}(d_{ij}^{a\bar{b}}) += H(d_{ij}^{a\bar{b}}, d)c(d)$ 

```

Vital to the speed and accuracy of these PT strategies, these algorithms do not need to check if the generated determinant is already in the variational space. Instead, after a thread has gone through Algorithm 3.2, the analogous algorithms for α, α and β, β doubles, and Algorithm 3.3, a single-index loop over the variational determinants zeroes all variational partial sums, as shown in Algorithm 3.4. This is another advantage of hash tables, as the operation to find and zero the partial sum of variational determinant, d , can be done in $O(1)$ time. This is key to allowing HCI to free memory from the variational Hamiltonian before starting the perturbative stage, where memory storage is limiting.

Algorithm 3.4 Pseudocode of a single OpenMP thread for removing perturbative contributions from determinants in the variational space.

```

1: for  $d$  in  $\{d\}^{\text{var}}$ 
2:    $\mathbf{M}(d) = 0$ 

```

This decomposition is key to removing the memory bottleneck in HCI without relying on semi-stochastic sampling. Furthermore, since each inner sum in Equation 3.3 is only computed by a single thread, no synchronization is needed until all hash ranges have been exhausted. At the end of the perturbative step, each MPI rank only needs to send its local correction for the final communication. This approach to the perturbative step of HCI limits the need for synchronization and network data transfers. Table 3.2 summarizes the key data structures and algorithmic advances.

Table 3.2: Key data structures and parallel implementation details.

Implementation	Computational Benefit
α - and β -strings stored as <code>std::bitset</code>	Facilitates 2D hash and allows compile-time optimizations of hash functions.
2D Hash	Allows the variational stage to be divided into evenly-sized chunks, facilitating load balancing. In the perturbative space, it also overcomes memory bottlenecks by allowing the space to be split into arbitrarily small chunks.
Hash Table	Allows for fast access of intermediate partial sums in the perturbative step. Also allows variational Hamiltonian to be cleared prior to the perturbative step.

3.3 Computational Details

The variational step of HCI is considered complete when the number of new determinants added via Equation 2 is less than 1% of the number in the current wave function. Throughout this study, a CAS-CI wave function (8 electrons in 8 orbitals or smaller) is used as the initial guess. Additionally, the threshold for singles excitations is set to $\varepsilon_1/5$, with ε_1 being the cutoff for doubles excitations. The RI approximation is used for 2-electron integrals, using the RIMP2-cc-pVDZ and RIMP2-cc-pVTZ auxiliary basis sets¹⁴¹, as appropriate.

Parallel scaling benchmarks were computed for *trans*-butadiene (See Appendix B for geometry) using up to 32 nodes on the Perlmutter Supercomputer at the National Energy Research Scientific Computing Center. Each run was launched with one MPI process per node and 128 OpenMP threads per process. The cc-pVDZ¹⁴² basis set was used with ε_1 of 1×10^{-4} Ha and ε_2 of 1×10^{-7} Ha. The full set of valence orbitals and the entire virtual space were correlated in these computations (22e,82o) using canonical RHF orbitals. The HCI space is composed of 3.2×10^6 variational determinants and 2.7×10^{10} perturbative determinants, selected from the FCI space of 2.0×10^{26} possible determinants.

To test for accuracy, the FCI energy of butadiene was computed using a linear extrapolation as done in prior heat-bath studies.⁴⁰⁻⁴³ Thresholds for ε_1 of 2.5×10^{-4} , 2.0×10^{-4} , 1.5×10^{-4} , 1.0×10^{-4} and 5.0×10^{-5} Ha were used with $\varepsilon_2 = 1.0 \times 10^{-8}$ Ha for all ε_1 's. To properly compare to prior benchmarks, the geometry for the accuracy test was taken from Daday *et al.*¹⁴³ The ANO-L-VDZP basis was used along with natural orbitals (NOs) taken from the

HCI variational density at $\varepsilon_1 = 1.0 \times 10^{-4}$. Since there is no RI basis optimized for ANO basis sets, the RIMP2-aug-cc-pVTZ¹⁴¹ auxiliary basis was used for this benchmark.

The accuracy of the new HCI implementation is also tested against diatomic benchmarks from prior SHCI studies.⁴¹ Specifically, the energies of C_2 , N_2 , O_2 , and F_2 are computed with the cc-pVTZ¹⁴² basis set. The HCI thresholds were set to $\varepsilon_1 = 1 \times 10^{-4}$ Ha and $\varepsilon_2 = 1 \times 10^{-8}$ Ha. RHF (ROHF for O_2) reference orbitals are used for a more direct comparison to the prior SHCI benchmark.

The singlet and triplet gaps in the cyclobutadiene automerization reaction were also studied with HCI. D_{2h} and D_{4h} geometries for cyclobutadiene were taken from a prior coupled-cluster study.¹⁴⁴ The HCI calculations included all valence electrons and were performed with the cc-pVTZ basis set resulting in a FCI space of (20e,172o).

The triplet and quintet states of the $[FeO(NH_3)_5]^{2+}$ complex were computed using geometries from previous studies.¹⁴⁵⁻¹⁴⁸ The cc-pVTZ basis set was used for Fe and O, the cc-pVDZ basis set for N, and the 6-31G¹⁴⁹⁻¹⁵¹ basis set for H. Polarization functions on H contribute 45 orbitals at the double-zeta level, and these were removed to keep the HCI calculation computationally tractable. ROHF orbitals were first obtained, followed by Pipek-Mezey localization of the valence space. The resultant local orbitals included in the active space are the Fe d -orbitals, O valence orbitals, and N lone-pair orbitals which donate to Fe (See Appendix B for localized orbitals). The remaining valence orbitals consist of N-H σ orbitals which are inactive. The resultant active space is (22e,168o).

For cyclobutadiene and $[FeO(NH_3)_5]^{2+}$, HCI extrapolation to the FCI limit was performed with ε_1 values of 2.5×10^{-4} , 2.0×10^{-4} , 1.5×10^{-4} , and 1.0×10^{-4} Ha. The perturbative step for the FCI extrapolation was computed with ε_2 of 1×10^{-7} Ha for all ε_1 's. Prior to performing HCI calculations for FCI extrapolation, natural orbitals for the states of interest were obtained using the density from the variational step of a HCI calculation with $\varepsilon_1 = 1 \times 10^{-4}$ Ha.

3.4 Results and Discussion

This section presents benchmarks of parallel scaling and accuracy, then demonstrates the new implementation's capabilities on strongly correlated systems.

Parallel scaling tests were run on butadiene as detailed in the Computational Details. The resultant total wall time on a single node was 5.5 hours. A total of 3.2×10^6 variational determinants and 2.7×10^{10} perturbative determinants were explicitly correlated in this calculation. Figure 3.4 shows the wall times and speedups for HCI variational and perturbative steps on up to 32 nodes. The variational step drops to $\sim 37\%$ parallel efficiency at 8 nodes,




Figure 3.4: Wall times (left) and speedups (right) for a HCI calculation of butadiene on up to 32 nodes are shown. Perfect speedup is shown as a solid black line. The HCI calculation included all valence electrons in the cc-pVDZ basis set with thresholds set to $\varepsilon_1 = 1 \times 10^{-4}$ and $\varepsilon_2 = 1 \times 10^{-7}$.

but the total scaling (including variational and perturbative steps) at 32 nodes remains above 74% at 23x speedup. This is due to the excellent scaling of the perturbative step, which is the more computationally intensive one and stays above 86% efficiency on 32 nodes. The current implementation is highly parallel allowing systematic approximations to FCI to be achieved with reasonable wall times.

Due to the orbitals, basis, and geometry chosen, the energy from the scaling benchmark is not directly comparable to prior energy benchmarks. Instead FCI extrapolation in the ANO-L-VDZP basis using the geometry from Daday *et al.*¹⁴³ is used to obtain an accuracy benchmark. A 5-point extrapolation results in a total energy of -155.55739 Ha which is only 3×10^{-4} Ha from a prior DMRG benchmark¹⁵² of -155.55718 Ha (See Appendix B for more detail).

Table 3.3: Diatomic energy benchmarks of HCI computed in the cc-pVTZ basis. The perturbative cutoff is $\varepsilon_2 = 10^{-8}$ Ha for HCI and SHCI.

Molecule	Method	S^2	ε_1	Energy (var)	Energy (total)
C ₂	SHCI ⁴¹	0.000	3×10^{-4}	-75.7738	-75.78463
N ₂		0.000	3×10^{-4}	-109.3608	-109.37486
O ₂		2.000	3×10^{-4}	-150.1130	-150.13078
F ₂ *		0.000	5×10^{-4}	-199.2787	-199.29727
C ₂	HCI	0.000	1×10^{-4}	-75.7803	-75.78493
N ₂		0.000	1×10^{-4}	-109.3689	-109.37524
O ₂		2.000	1×10^{-4}	-150.1221	-150.13062
F ₂		0.000	1×10^{-4}	-199.2854	-199.29715

*MP2 NOs were used. RHF/ROHF orbitals used for all others.

To additionally check HCI for accuracy, Table 3.3 shows the computed HCI and SHCI⁴¹ energies for C₂, N₂, O₂, and F₂ in the cc-pVTZ basis set. The total energies for all benchmarks differ by less than 0.3 mHa. Differences in the energies arise due to the use of RI for the two electron integrals as well as the use of deterministic versus stochastic perturbation theory.




Figure 3.5: Automerization of cyclobutadiene (top) is shown with relative full CI energies in kcal mol⁻¹ plotted (below). Full CI energies were computed using a 4-point extrapolation from HCI calculations (See Table 3.4).

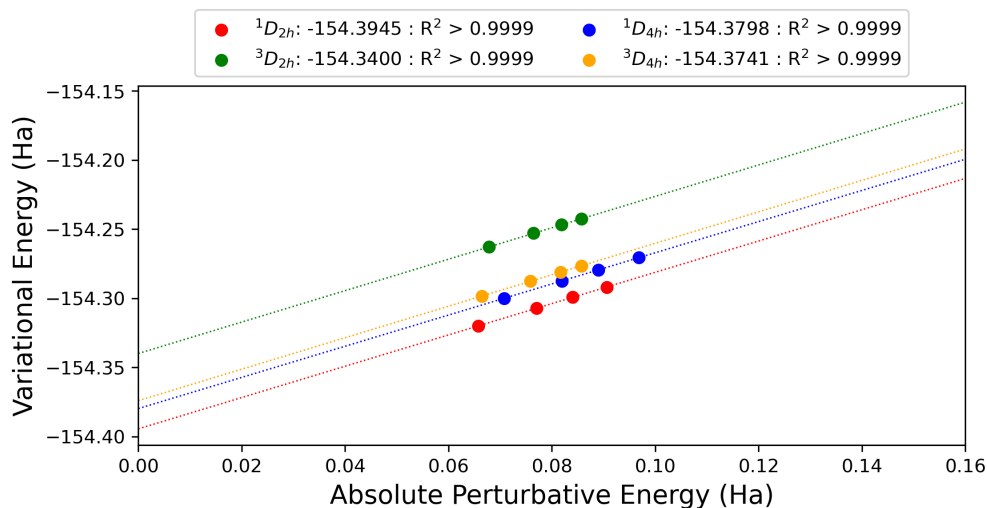


Figure 3.6: Linear extrapolation of ¹D_{2h}, ³D_{2h}, ¹D_{4h}, and ³D_{4h} cyclobutadiene energies to the full CI limit. Extrapolated energies and R² values are provided in the legend. HCI energies were computed in the cc-pVTZ basis with a (20e,172o) active space and ϵ_1 values of 0.10 mHa, 0.15 mHa, 0.20 mHa, and 0.25 mHa. The perturbative cutoff, ϵ_2 was set to 10⁻⁷ Ha for all cyclobutadiene calculations.

To further test the capabilities of the new HCI implementation, HCI is applied to the automerization of cyclobutadiene,^{70,144,153–161} a well-known strongly correlated system. The

reaction begins and ends at a D_{2h} geometry, and proceeds through a D_{4h} transition state that has a low-lying triplet electronic state. Figure 3.5 shows the automerization reaction (top) with computed barrier and spin gaps (bottom). The strong correlation stems from the near degeneracy of the HOMO and LUMO, which gives cyclobutadiene a biradical electronic structure.

Table 3.4: HCI energies for cyclobutadiene automerization. All calculations are done in the cc-pVTZ basis set with $\varepsilon_2 = 1 \times 10^{-7}$ Ha. Linear extrapolations using the tightest n ε_1 values are shown for each state.

State	S^2	ε_1 (mHa)	N_{var}	N_{PT}	E_{var}	E_{total}	
$^1D_{2h}$	0.000	0.25	1,681,015	3.6×10^{10}	-154.291913	-154.382519	
		0.20	2,402,347	4.2×10^{10}	-154.299285	-154.383326	
		0.15	3,915,304	5.2×10^{10}	-154.307161	-154.384214	
		0.10	8,261,698	6.7×10^{10}	-154.320075	-154.385807	
		extrapolation 4 points					-154.394473
		extrapolation 3 points					-154.394734
$^3D_{2h}$	2.000	0.25	1,925,018	4.0×10^{10}	-154.242656	-154.328342	
		0.20	2,608,106	4.6×10^{10}	-154.246864	-154.328736	
		0.15	4,048,046	5.4×10^{10}	-154.252969	-154.329419	
		0.10	8,374,320	7.0×10^{10}	-154.262967	-154.330764	
		extrapolation 4 points					-154.339966
		extrapolation 3 points					-154.340583
$^1D_{4h}$	0.000	0.25	1,735,690	3.6×10^{10}	-154.270659	-154.367513	
		0.20	2,379,890	4.1×10^{10}	-154.279442	-154.368444	
		0.15	3,703,206	4.9×10^{10}	-154.287426	-154.369361	
		0.10	7,692,782	6.5×10^{10}	-154.300110	-154.370827	
		extrapolation 4 points					-154.379803
		extrapolation 3 points					-154.380042
$^3D_{4h}$	2.000	0.25	1,796,992	3.5×10^{10}	-154.276501	-154.362263	
		0.20	2,464,584	4.0×10^{10}	-154.281035	-154.362738	
		0.15	3,913,936	5.0×10^{10}	-154.287727	-154.363521	
		0.10	8,312,068	6.6×10^{10}	-154.298450	-154.364918	
		extrapolation 4 points					-154.374077
		extrapolation 3 points					-154.374454

4-point energy extrapolations for cyclobutadiene are shown in Figure 3.6. The FCI extrapolation leads to an activation barrier for automerization on the singlet state of 9.21 kcal mol⁻¹, a singlet-triplet gap of 34.20 kcal mol⁻¹ at the D_{2h} geometry, and a singlet-triplet gap of 3.59 kcal mol⁻¹ at the D_{4h} geometry. Error estimates of the total energies can be made by comparing 4-point to 3-point extrapolation (see Table 3.4), which results in less than 1 mHa energy differences for each state and geometry. Using the two extrapolation schemes,

Table 3.5: Relative energies of cyclobutadiene computed using different methods. Energies are in kcal mol⁻¹.

Method	basis	$\Delta E(^1D_{4h} - ^1D_{2h})$	$\Delta E(^3D_{4h} - ^1D_{4h})$	$\Delta E(^3D_{2h} - ^1D_{2h})$
FCI 4-point extrapolation	cc-pVTZ	9.21	3.59	34.20
FCI 3-point extrapolation	cc-pVTZ	9.22	3.51	33.98
CCSDTQ ¹⁴⁴	aug-cc-pVTZ	8.93	3.32	33.05

the relative energies change only slightly, with the largest change being 0.22 kcal mol⁻¹ (See Table 3.5). Prior studies computed the "theoretical best estimate" for the singlet-triplet gaps of the D_{2h} and D_{4h} geometries and singlet barrier at the CCSDTQ/aug-cc-pVTZ level of theory.¹⁴⁴ These quantities are shown in Table 3.5 alongside the CI extrapolation energies. Some discrepancy arises from the difference in basis, though the largest relative energy difference is relatively small at 1.15 kcal mol⁻¹.

The $[\text{FeO}(\text{NH}_3)_5]^{2+}$ test case represents a challenging test for correlated electronic structure theories, as it contains a large number of electrons in close spatial proximity and two spin states that are close in energy. This complex has been studied with methods such as coupled-cluster and CASPT2.¹⁴⁵⁻¹⁴⁸ Feldt *et al.* applied CCSD(T) with reference orbitals from UHF, ROHF, UKS, ROKS, and CASSCF methods, which led to several different signs and magnitudes of the triplet-quintet spin gap.¹⁴⁵ In addition, Feldt *et al.* used DMRG-CASPT2 to compute valence

Table 3.6: Quintet energy relative to the triplet energy of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ at various levels of theory. Energies are in kcal mol⁻¹.

Method	Relative Quintet Energy
UKS-UCCSD(T) ¹⁴⁵	0.6
ROKS-UCCSD(T) ¹⁴⁵	1.2
ROKS-RCCSD(T) ¹⁴⁵	2.5
CASSCF(12,12)-UCCSD(T) ¹⁴⁵	-0.5
CASPT2/CC ¹⁴⁵	-0.5
DLPNO-CCSD(T ₀) ¹⁴⁸	-5.0
DLPNO-CCSD(T ₁) ¹⁴⁸	-5.4
CASPT2/CC ¹⁴⁸	0.4
HCI (this work)	-6.0

correlation and UKS-CCSD(T) to compute the semicore correlation (denoted CASPT2/CC), however, they noted the lack of a high quality benchmark left the energetics of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ an open question. Recent work by Drosou *et al.* used DLPNO-CCSD(T) and CASPT2/CC with similarly varying results.¹⁴⁸ Table 3.6 summarizes prior spin gaps, showing values from

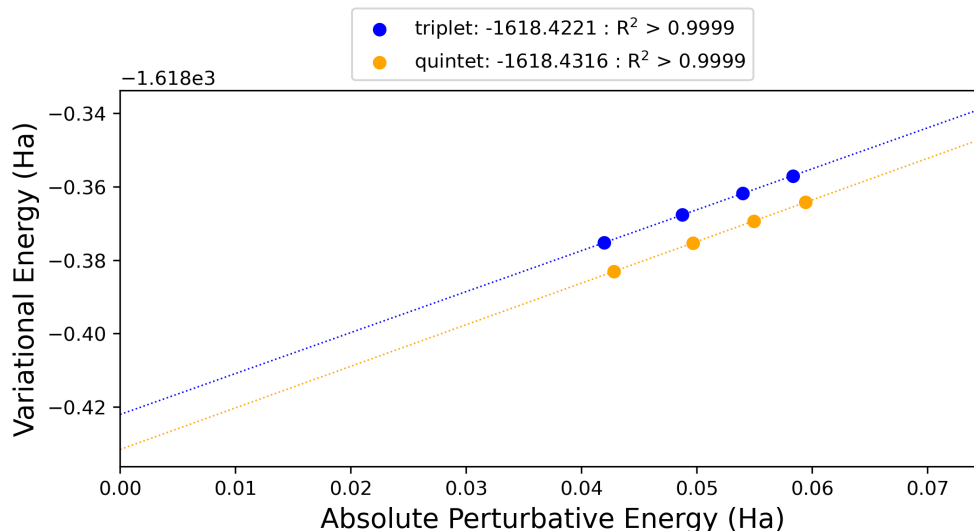


Figure 3.7: The 4-point FCI extrapolation of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ with (22e,168o) is shown. All perturbative energies are computed with $\varepsilon_2 = 1 \times 10^{-7}$ Ha. The FCI extrapolated energy and R^2 are provided in the legend. Variational thresholds for ε_1 and the basis set are listed in the Computational Details.

-5.4 kcal mol⁻¹ to 2.5 kcal mol⁻¹. The different predictions for the sign and magnitude of the spin gap demonstrate the challenging electronic structure of this complex.

While including all valence electrons and Fe semicore (3s, 3p) electrons is likely necessary to accurately refine the T-Q spin gap, such an active space (60e,187o) is still out of reach for HCI. The current active space considered reflects the most important valence contributions that are feasible in a reasonable wall time. An expanded HCI space including at least the Fe semicore orbitals may be visited in a future study.

A 4-point full CI extrapolation for the triplet and quintet states is shown in Figure 3.7. The most expensive calculation was the quintet with $\varepsilon_1 = 1 \times 10^{-4}$ Ha. The variational stage for this calculation took 1.75 hours on 16 nodes while the perturbative step took approximately 9.5 hours on 36 nodes. This extrapolation results in the quintet state being 5.98 kcal mol⁻¹ lower than the triplet state. Similarly to the cyclobutadiene calculation, an estimate of the error is computed by conducting a 3-point extrapolation, which slightly lowers the gap to 5.77 kcal mol⁻¹. In the 4- and 3-point extrapolations, R^2 for the fits remain above 0.9999. This suggests the spin gap for this active space is reliable, even though the active space itself is likely too small for quantitative accuracy.

Natural orbitals were computed using densities from the variational wave functions with $\varepsilon_1 = 0.1$ mHa. Figure 3.8 shows NOs for the triplet state and Figure 3.9 for the quintet. The singly-occupied orbitals for the triplet are the Fe=O π_{xz}^* and π_{yz}^* orbitals. These are also singly occupied in the quintet, along with the d_{xy} and $d_{x^2-y^2}$ orbitals. These occupations

Table 3.7: HCI energies for the $[\text{FeO}(\text{NH}_3)_5]^{2+}$ triplet and quintet states. Basis set is described in Computational Details. In the perturbative step, $\varepsilon_2 = 1 \times 10^{-7}$ Ha for all calculations. Linear extrapolation is used with the n smallest ε_1 values.

State	S^2	ε_1 (mHa)	N_{var}	N_{PT}	E_{var}	E_{total}	
triplet	2.000	0.25	4,574,604	5.7×10^{10}	-1618.357028	-1618.415337	
		0.20	6,383,114	6.3×10^{10}	-1618.361823	-1618.415781	
		0.15	9,917,081	7.0×10^{10}	-1618.367631	-1618.416375	
		0.10	18,900,066	8.0×10^{10}	-1618.375265	-1618.417225	
		extrapolation 4 points					-1618.422056
		extrapolation 3 points					-1618.422275
		quintet	6.000	0.25	6,215,652	6.76×10^{10}	-1618.364232
0.20	8,611,136			7.28×10^{10}	-1618.369339	-1618.424287	
0.15	13,120,044			7.97×10^{10}	-1618.375351	-1618.424992	
0.10	23,907,184			8.95×10^{10}	-1618.383049	-1618.425872	
extrapolation 4 points					-1618.431591		
extrapolation 3 points					-1618.431471		

are consistent with Feldt *et al.*'s investigation using DMRG-CASPT2.¹⁴⁵ This consistency indicates that the same qualitative states are being targeted between the studies, though HCI places the high spin quintet much lower than the triplet compared to coupled-cluster methods.^{145,148}




Figure 3.8: NOs for the triplet state computed using the variational density from HCI with $\varepsilon_1 = 0.1$ mHa are shown with isosurface values of 0.8 a.u. NO occupations are listed below each orbital.

From Figures 3.8 and 3.9, it can be seen that the quintet NOs have substantially less mixing between the Fe $d_{x^2-y^2}$ orbitals and N lone-pair orbitals relative to the triplet NOs, which indicates that the triplet state has more electron correlation stemming from interactions between the equatorial NH_3 lone pairs and the Fe d orbitals in the xy -plane. This suggests




Figure 3.9: NOs for the quintet state computed using the variational density from HCI with $\varepsilon_1 = 0.1$ mHa are shown with an isosurface values of 0.8 a.u. NO occupations are listed below each orbital.

that more correlation energy will be recovered from the triplet state than the quintet state if the N-H σ -bonds were included in the CI space. While these were excluded from this study due to high computational cost (resulting in an additional 30 electrons and 15 orbitals), their contribution is likely to lower the magnitude of the current spin gap, and possibly even reverse its sign.

3.5 Conclusions

Advances in the HCI method are demonstrated in a new highly parallel code. The most expensive step—the perturbative energy—exceeds 80% parallel efficiency on 32 nodes (4096 cores) on the Perlmutter Supercomputer. This strategy allowed fast evaluation of an active space size of (22e,168o), which explicitly includes 23 million variational determinants and 89 billion perturbative determinants.

Methods that approximate the FCI energy are especially useful for approaching strongly correlated problems. Therefore high quality activation barriers and gaps are computed for the cyclobutadiene automerization, which are consistent with prior CCSDTQ¹⁴⁴ calculations. Additionally, FCI extrapolations are obtained to compute the electron correlation from the Fe=O valence space along with N lone-pairs in $[\text{FeO}(\text{NH}_3)_5]^{2+}$. The triplet-quintet gap, however, is still unresolved due to the lack of correlation from N-H sigma bonds and iron semicore orbitals, which are each known to be important for this complex. These may be resolved using incremental FCI,^{39,69,70,79–81} which relies on the HCI algorithm to solve its individual CI energies.

The calculations presented in this work demonstrate the HCI implementation is highly scalable and can provide high-quality energy benchmarks. The 2D hash function combined with a new parallel implementation was particularly instrumental in overcoming the otherwise high memory requirements of HCI. Further development and optimization is expected to improve the performance and scalability of HCI, which will additionally facilitate deeper study of strongly correlated complexes such as $[\text{FeO}(\text{NH}_3)_5]^{2+}$.

CHAPTER 4

Fully Variational Incremental CASSCF

The previous chapter introduced computational improvements to HCI enabling the treatment of a (22e,168o) active space. While these improvements are impressive, the fact remains that HCI still scales exponentially albeit with a smaller exponential factor than FCI. Thus, theoretical advances to approximate FCI at reduced scaling—in particular, in polynomial time—is necessary to kick us off the gold standard. This chapter explores this possibility in the context of the complete active space self-consistent field (CASSCF) method.

The CASSCF method is a canonical electronic structure theory that holds a central place in conceptualizing and practicing first principles simulations. For applications to realistic molecules of interest, however, CASSCF must be approximated to circumvent its exponentially scaling computational costs. Applying the many-body expansion—also known as the method of increments—to CASSCF (iCASSCF) has been shown to produce a polynomially scaling method that retains much of the accuracy of the parent theory and is capable of treating full valence active spaces. Due to an approximation made in the orbital gradient, the orbital parameters of the original iCASSCF formulation were not variationally optimized, which limited the accuracy of its nuclear gradient. In this chapter, we explore a variational implementation of iCASSCF where all parameters are fully optimized during energy minimization. This chapter is based on work published previously in the *Journal of Chemical Physics*.⁷⁰

4.1 Introduction

Effective combinations of modern computer hardware and electronic structure theory have allowed quantum chemical calculations to become standard practice in chemical research. Many of these techniques are grounded in the Hartree-Fock (HF) method, a standard molecular orbital theory that captures approximately 99% of the total electronic energy.² To obtain chemical accuracy, however, the remaining 1% (the correlation energy) must be recovered. Electron correlation can be split into static correlation (correlation from nearly

degenerate electron configurations) and dynamic correlation (correlation to reduce electron repulsion).² One widely used, canonical strategy to capture (mostly) the static correlation effects is the complete active space self-consistent field (CASSCF) method.^{2,47–52} In CASSCF, single-particle orbitals are partitioned into inactive, active, and virtual spaces, with the active space accounting for correlation effects beyond the HF approximation. In particular, the complete set of electronic configurations is treated using a configuration interaction (CI) expansion within the active space. CASSCF therefore is a systematic approximation to full CI (FCI), where the latter is a CAS-CI with all orbitals included in the active space. While CASSCF is a powerful method, the exponentially scaling costs of solving the CAS space CI problem limits conventional implementations to around 18 electrons in 18 orbitals. Using highly parallel compute architectures, this limit has been pushed upward to 22 electrons in 22 orbitals.⁵²

Because the basic limitation in CASSCF is in the active space CI step, several strategies have been introduced to partition the active space and reduce computational costs. Some approaches include the restricted active space (RAS) SCF^{53–57} and the generalized active space (GAS) SCF^{58,59} methods. While these do reduce computational cost, RASSCF and GASSCF come at the price of greater ambiguity in selection of the active space. This ambiguity exists in CASSCF as well, though to a lesser extent—this will be discussed in the next section. Other methods have aimed specifically at reducing the computational effort of the CI problem without changing the active space selection. These include quantum Monte Carlo CI,²⁴ density-matrix renormalization group (DMRG),^{60–66} and select CI methods^{32–34} such as heat-bath CI (HCI).^{40,41,43–45,162} These techniques—when converged to chemical accuracy in the total electronic energy—remain exponentially scaling with growing system size. Related to these are variational 2-particle reduced density matrix (2-RDM) methods, which replace the CI wave function with a constrained optimization of the two-particle density,^{163–166} allowing lower-order scaling with slight decreases in accuracy.

CASSCF’s underlying CI problem may also be solved by collecting correlation energies from many small active spaces, and assembling these units into a close approximation of the exact result. This technique relies on the many-body expansion (MBE), and in this context is known as the method of increments. Stoll^{73–75} was the first to use MBE along with CI. Later work by the Dolg group applied the MBE to coupled-cluster methods^{71,72} and more recently, Gauss^{76–78} used MBE to approximate full CI. Our group has employed MBE to approach large full CI^{79–81} and CASSCF⁶⁹ problems. The MBE expression for the energy is

$$E = E_{ref} + \sum_i \varepsilon_i + \sum_{j < i} \varepsilon_{ij} + \sum_{k < j < i} \varepsilon_{ijk} + \dots, \quad (4.1)$$

where the indices i, j, k refer to individual bodies. Each term, ε_x , in Equation 4.1 adds correlation via interactions among the bodies included in x . Since this expansion derives the total correlation energy as the sum of many small terms (rather than via one large CI), the MBE provides a tractable approach to otherwise intractable electronic problems.

The incremental terms of Equation 4.1 are

$$\begin{aligned}\varepsilon_i &= E_i - E_{ref}, \\ \varepsilon_{ij} &= E_{ij} - E_{ref} - \varepsilon_i - \varepsilon_j, \\ \varepsilon_{ijk} &= E_{ijk} - E_{ref} - \varepsilon_{ij} - \varepsilon_{ik} - \varepsilon_{jk} - \varepsilon_i - \varepsilon_j - \varepsilon_k,\end{aligned}\tag{4.2}$$

and so on for higher order terms. The expansion therefore introduces correlation beyond the Hartree-Fock reference via CAS-CI calculations of $2n$ electrons in $2n$ orbitals. Since n is expected to stay small even as the number of correlated electrons increases, the exponential wall of the parent CAS problem is avoided. Prior work has shown the method of increments recovers electron correlation to a high degree of accuracy with tractable computational effort for full CI^{76–81} as well as for CASSCF.⁶⁹

The MBE need not be restricted to solving the CI problem.^{167–171} The definition of bodies and n -body interactions may also be extended to individual molecules or fragments of molecules, as in the fragment molecular orbital (FMO) framework.^{170–176} The FMO method has allowed for meaningful quantum chemical calculations of large molecular systems, well beyond what can be approached with conventional means. For these calculations to be most useful, however, reliable geometries—and hence analytic gradients—are needed. While the original report of analytic gradients for FMO were valid only when using minimal basis sets,¹⁷² the correct gradient has since been derived and implemented.^{173–176} This difference also lies at the core of the present study.

Recently, approximate gradients for the incremental CASSCF (iCASSCF) approach were introduced.⁶⁹ In order to accomplish this, iCASSCF assembles the generalized Fock matrix (required for the orbital gradient) from incremental 1- and 2-particle reduced density matrices (1- and 2-RDM). The usual CASSCF formula for geometric gradients⁵¹ can then be used once the orbitals are fully optimized. While elegant to a degree, one approximation was used that resulted in an inexact Fock matrix, and therefore an approximate gradient (see Theory Section for further details). While this resulted in reasonably accurate energies, the Hellman-Feynman theorem could not strictly be applied to this approximation.¹⁷⁷

Variational optimization of all wave function parameters is crucial for computing accurate nuclear gradients in CASSCF as well as iCASSCF. In turn, having access to accurate gradients allows exploration of potential energy surfaces (PESs) as well as obtaining many

other molecular properties, for example the rotational entropy.¹⁷⁷ In conventional CASSCF implementations, however, the necessity of using a truncated active space for virtually all molecules of interest greatly hinders its usefulness in practice. In particular for PES exploration, large errors in reaction energies are observed (in some cases $> 10 \text{ kcal mol}^{-1}$)^{178–180} and potential energy profiles can be discontinuous.^{181,182} While such discontinuity may sometimes be resolved by state-averaging (SA), SA-CASSCF also exhibits similar issues¹⁸³ due mostly to the incomplete treatment of static correlation via the truncated active space. On the other hand, iCASSCF is designed with full valence active space calculations in mind. Since there are many choices for a full valence active space, a well-defined procedure for choosing this space is important in making iCASSCF broadly useful and will be discussed in the Theory section. One additional challenge in iCASSCF is the original implementation suffers from a lack of variationality in the optimization procedure. An improvement here is necessary for iCASSCF to be fully useful—with correct gradients and ability to treat full valence active spaces—by modifying the iCASSCF procedure to permit fully variational parameter optimization.

In this chapter, we introduce a variational iCASSCF procedure for electronic structure simulations of chemically interesting systems. The variational iCASSCF will inherit the good properties of the previous method, including close reproduction of CASSCF-level results as well as ability to treat full valence active spaces. The method will also be demonstrated for the first time on transition metal complexes, showing iCASSCF to be sufficiently robust to tackle strongly-correlated systems with experimental relevance.

4.2 Theory

4.2.1 iCASSCF

Like CASSCF, iCASSCF is an iterative procedure consisting of an orbital optimization step and a CI step. The primary procedural difference lies in the fact that iCASSCF utilizes the many-body expansion to recover the correlation energy (Equation 4.1). Since the MBE can correlate all valence electrons, a full valence active space is a natural choice in iCASSCF. To define this active space, the generalized valence bond perfect pairing (PP) method^{184–188} is used to obtain the initial orbitals. In the valence space, PP creates localized orbital pairs for each electron pair—one bonding (or lone pair) orbital and one anti-bonding (or correlated) orbital. The full valence active space in iCASSCF is therefore composed of all valence orbitals from PP. The PP ansatz has the added benefit of recovering strong intrabond correlation effects.

In iCASSCF the bonding and anti-bonding orbital pairs from PP are then used to define the individual bodies of the MBE, and therefore the CAS-CI problem for the full CAS space is never explicitly constructed. Instead, the total correlation energy is computed by solving the CAS-CI problem using many smaller active spaces, as shown in Figure 4.1. Prior work has shown the 3-body expansion to be a good compromise between accuracy and cost for the method of increments applied to full CI as well as CASSCF.^{69,79–81} While studies of the MBE (in FMO) have shown that the errors of the MBE are extensive, they also report that error growth is small and negligible in practical calculations.^{167,168} This same behavior is expected from iCASSCF due to the extensivity of the energy that is intrinsic to MBE.




Figure 4.1: **a.** The full valence CASSCF active space is decomposed into **b.** pairs of localized bonding and anti-bonding orbitals which define the bodies in the many-body expansion. Solving the CAS-CI problem within these active spaces provides the 1-body interactions (ε_i) in Equation 4.1. **c.** The combination of any two active spaces in the 1-body terms generates the 2-body terms. Solving the CAS(4,4)-CI problem for these new active spaces returns the 2-body interaction energies (ε_{ij}) in Equation 4.1. Higher-order expansions are similarly defined (not shown).

In CASSCF, the energy can be expressed using the 1- and 2-RDMs by

$$E = \sum_{pq} D_{pq} h_{pq} + \sum_{pqrs} d_{pqrs} g_{pqrs}, \quad (4.3)$$

where h_{pq} and g_{pqrs} are the 1- and 2-electron integrals and the RDMs are defined using the usual expressions from second quantization, $D_{pq} = \langle \psi | \hat{a}_p^\dagger \hat{a}_q | \psi \rangle$ and $d_{pqrs} = \langle \psi | \hat{a}_p^\dagger \hat{a}_r^\dagger \hat{a}_s \hat{a}_q | \psi \rangle$. The generalized Fock matrix, which is essential for orbital optimization, can similarly be

defined by the 1- and 2-RDMs

$$F_{mn} = \sum_q D_{mq} h_{nq} + \sum_{qrs} d_{mqr} g_{nqrs}. \quad (4.4)$$

Using the method of increments, the 1-RDM can be computed by

$$D_{pq}^{tot} = D_{pq}^{ref} + \sum_i \Delta D_{pq}^i + \sum_{j<i} \Delta D_{pq}^{ij} + \dots, \quad (4.5)$$

where the ΔD_{pq}^x terms are changes in the density due to interactions between the bodies in x . Analogous to Equation 4.2,

$$\begin{aligned} \Delta D_{pq}^i &= D_{pq}^i - D_{pq}^{ref} \\ \Delta D_{pq}^{ij} &= D_{pq}^{ij} - D_{pq}^{ref} - \Delta D_{pq}^i - \Delta D_{pq}^j, \end{aligned} \quad (4.6)$$

and higher-order n -body densities are defined analogously. The incremental expansion of the 1-RDM applies similarly to the 2-RDM. Taken together with the expressions for E and F , this means that any quantity which depends on the RDMs may be constructed via increments. Since there is only one definition of the molecular orbitals (across all incremental terms), there is no complication in assembling the RDMs in this way. Since the orbital gradient is defined by the anti-symmetric part of the Fock matrix,

$$E_{pq}^o = 2(F_{pq} - F_{qp}), \quad (4.7)$$

orbital optimization can follow naturally from the many-body expansion.

There exists, however, a potentially complicating issue in incremental Fock matrix construction via Equation 4.4 in iCASSCF theory. Specifically, since the full CAS-CI is not performed within the total active space, the energy is not invariant to active-active orbital rotations.² At low orders of MBE (low n) this invariance is significant, though invariance can be restored at high n . In practice this lack of invariance can complicate iCASSCF orbital optimization compared to CASSCF, since CASSCF normally would avoid active-active rotations as they have zero gradients. To work around this issue, our prior strategy was to symmetrize the active-active block of each incremental Fock matrix⁶⁹ and explicitly restore much of the expected symmetry. At its core, this symmetrization projects out some components of the orbital gradient (Equation 4.7), and therefore the resulting iCASSCF orbital optimization is not fully variational. Energies and gradients could then be acquired, but the gradients were only accurate to around 2×10^{-4} a.u. In order to achieve fully variational parameter

optimization for the iCASSCF method, the full, unprojected orbital gradient must be used for orbital optimization.

Before detailing the variational iCASSCF procedure, a discussion of orbital rotations and the effect of single excitations is warranted. In CASSCF, orbital rotations of the form $\exp(\boldsymbol{\kappa})$ are used in the orbital optimization steps, where $\boldsymbol{\kappa}$ is an anti-symmetric matrix defined by

$$\boldsymbol{\kappa} = \sum_{p>q} \kappa_{pq} (\hat{a}_p^\dagger \hat{a}_q - \hat{a}_q^\dagger \hat{a}_p). \quad (4.8)$$

A Taylor series expansion of $\exp(\boldsymbol{\kappa})$ then results in first and second order terms containing single ($\hat{a}_p^\dagger \hat{a}_q$) and double ($\hat{a}_p^\dagger \hat{a}_q \hat{a}_r^\dagger \hat{a}_s$) excitations. Within the active space, these single excitations from orbital rotation coincide with single excitations within CI and are therefore redundant to a first-order approximation. Figure 4.2 shows this effect, where the orbital gradients within the active space increase when CI singles excitations are removed. Gradients outside of the active space are also affected, but to a much lesser extent. In practice, reduced redundancy in the iCASSCF parameter space is expected to permit more facile optimization of those parameters. This approach—removing the singles excitations from the CI—allows fully variational iCASSCF optimization.

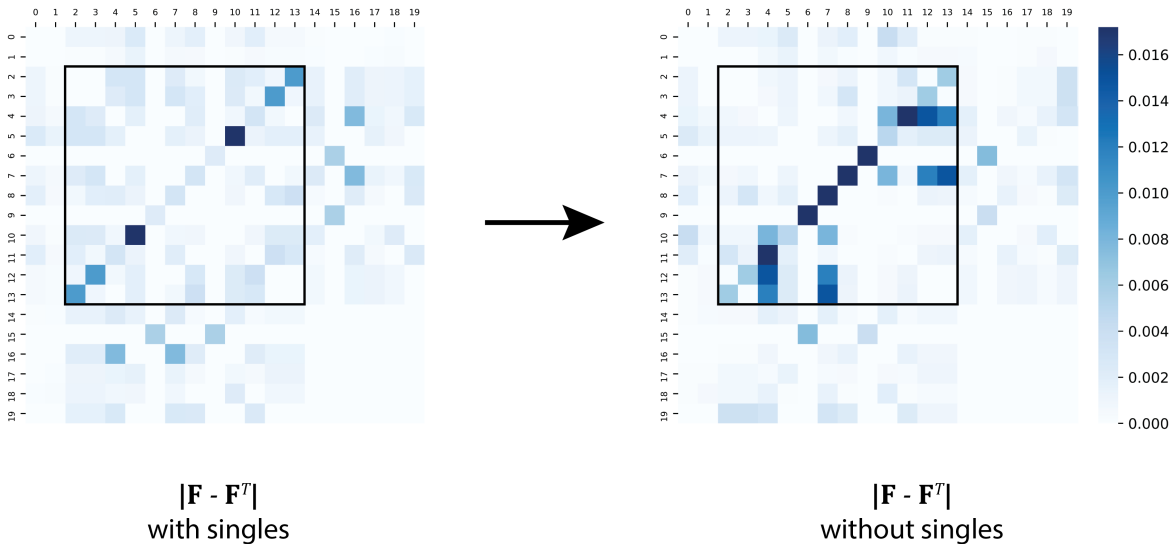


Figure 4.2: The magnitude of the orbital gradient (see Equation 4.7) for formaldehyde is shown with single excitations (left) and without (right). The gradient is truncated so only values corresponding to the first twenty orbitals are shown. One can see that the magnitude of gradient terms increase as the redundant single excitations are removed.

While the use of the Hellman-Feynman theorem¹⁷⁷ results in the correct geometric gradient, the underlying energy function is no longer strictly CAS-CI. This is due to the fact that

though the active-active singles excitations of the CI are largely redundant with the orbital rotations, they do not completely overlap. The truncated CI therefore does not precisely reproduce the CAS-CI result. Similar approximations in orbital-optimized coupled-cluster (CC) theory have shown the error due to missing singles to be small,^{189–191} and benchmarks in the present study will show the omission of single excitations from iCASSCF results in small differences from conventional CASSCF.

The missing correlation from single excitations may be recovered by reintroducing them once orbital optimization is completed. In other words, a regular iCAS-CI may be performed using singles-free iCASSCF optimized orbitals. For the remainder of the chapter, the singles-free procedure is denoted simply as iCASSCF. When discussing iCAS-CI, it is implied that iCASSCF orbitals are used.

4.2.2 Computational Effort

The calculation of all n -body interactions in iCASSCF requires solving

$$\binom{N}{n} = \frac{N!}{(N-n)!n!} \tag{4.9}$$

CAS-CI problems. The computation therefore scales with $O(N^n)$ for large N , where N refers to the total number of bodies (one half of the number of electrons in the active space). Since each CAS-CI problem requires more computational effort as n increases, it remains desirable to truncate the expansion to reduce the $O(N^n)$ scaling and prefactor. Herein, the iCASSCF expansion is limited to $n = 3$ or $n = 4$.

Computational effort can be further diminished by screening the higher order interactions.^{78,171} The screening prescription is as follows. Once the j th order of the expansion has been generated, the tuples involved in the $(j + 1)$ th order are generated. Prior to evaluation of the $(j + 1)$ th order interactions, the order of magnitude of the terms is predicted using the current (j th order) terms. Terms predicted to be small are then omitted, thus reducing the number of $(j + 1)$ -body calculations required. For example, when predicting the magnitude of a 3-body term, one would examine the three 2-body terms which generate the 3-body term. If the 2-body terms, ε_{ij} , are below a predetermined threshold, then the 3-body interaction is not computed. In this case, screening would avoid computing the full cubic number of 3-body terms. In general, the effect of screening propagates to higher order terms and allows scaling closer to $O(N^k)$ where $k < n$ for an n -body expansion. In principle, the significant $j + 1$ -body tuples could be estimated at every SCF step, however, in this study the selection of tuples is only performed once before any SCF iterations.

4.3 Computational Details

The variational iCASSCF method is implemented in a development version of Q-Chem 5.0.¹⁹² Results for variational iCASSCF were obtained using similar parameters as a previous study,⁶⁹ and differences are noted here. Simulations use the 6-31G* basis set, with 2-electron integrals computed using the resolution of the identity (RI) approximation^{193,194} and the RIMP2-VDZ auxiliary basis. Bodies are generally defined using bonding-antibonding pairs, initially from perfect pairing.^{184–188} Exceptions are noted where appropriate. The reference energy (E_{ref} in Equation 4.1) is computed using a Hartree-Fock reference configuration. iCASSCF is considered converged when the norm squared of the orbital gradient is below 4×10^{-8} . All active spaces are full valence.

Conventional CASSCF calculations used localized orbitals from CCSD as initial orbitals and were performed without density fitting. The MOLPRO 2012.1 electronic structure package¹⁹⁵ was employed for CASSCF. The same small-molecule benchmarks as in our previous study⁶⁹ were reexamined here using the new algorithm. In computing the energy for formaldehyde the CO π -system and O lone-pair orbitals are combined into a single body.

The singlet-triplet gap of oxoMn(salen)Cl, an intermediate in the Mn(III)(salen)-catalyzed epoxidation reaction,^{196,197} was also examined. The iCASSCF reference energy and configuration are from ROHF, and the calculations are considered converged when the norm squared of the orbital gradient is below 5×10^{-7} . For oxoMn(salen)Cl, ω B97X-D/6-31G* was used to obtain the geometry, starting from the truncated complex used in previous studies of this system.^{41,198–204} The model complex has a full valence active space of (84,84). Due to large degrees of correlation between them, O (oxo) lone pairs and Mn=O σ - and π -systems are clustered as a single body.

Cyclobutadiene automerization was explored using a Python implementation of the growing string method (PyGSM).^{205–208} A double-ended string was grown and optimized with 7 nodes along the reaction path. For this reaction, the 6-31+G** basis and RIMP2-cc-pVDZ auxiliary basis were used. The iCASSCF calculations for this reaction path use ROHF for the reference energy and configuration. The calculations are converged when the norm squared of the orbital gradient is below 2×10^{-8} . The many-body expansion goes out to $n = 3$. The four valence π orbitals are grouped as a single body in these calculations.

4.4 Small-molecule benchmarks

To validate the new iCASSCF procedure, comparisons are made to conventional CASSCF energies and geometries. The benchmark set⁶⁹ consists of small molecules with 10-14 valence electrons: ethylene, formaldehyde, hydrogen cyanide, hydrogen peroxide, and methanol.

Comparisons to results in our previous study can be found in Appendix C. For CASSCF as well as iCASSCF, full valence active spaces with an equal number of orbitals and electrons were used (starting from PP orbitals). To test the quality of the optimizations, iCAS-CI energy evaluation was used to recover correlation from missing singles, using the final iCASSCF orbitals.

Table 4.1 shows the variational iCASSCF and iCAS-CI energies compared to conventional CASSCF. The energy differences between iCASSCF($n = 3, n = 4$) and CASSCF differ by more than 2 mHa in all benchmarks and in some cases, iCASSCF($n = 4$) has larger deviation than the iCASSCF($n = 3$). However, when including single excitations via iCAS-CI, the energies are substantially improved. The largest error for iCAS-CI($n = 3$) is 0.9 mHa, and for iCAS-CI($n = 4$), the largest error is reduced to 0.5 mHa (see Figure 4.3c). Across all the benchmarks, iCAS-CI($n = 4$) resulted in a mean unsigned error of 0.3 mHa compared to conventional CASSCF. In some instances, the iCAS energy drops slightly below the CASSCF energy. While the MBE does tend towards the exact energy with increasing n , it is not guaranteed to do so from above. Thus iCASSCF, while variational with respect to optimization of its internal parameters, does not provide an upper bound to the parent CASSCF energy.

The characteristic bond distances and angles produced by the new iCASSCF procedure are in good agreement with conventional CASSCF. For $n = 3$ and $n = 4$, the maximum error in angles is 0.5° and this occurs in the H-O-O-H dihedral (Figure 4.3a). The bond distances are also satisfactory, though in some cases the $n = 4$ bond distances have slightly larger errors than $n = 3$. Regardless, the error remains below 0.006 \AA (0.6 pm) for both $n = 3$ and $n = 4$.

Table 4.1: Total energies of the small molecule benchmark set from iCASSCF, iCAS-CI, and CASSCF using the 6-31G* basis set.

Molecule	iCASSCF($n = 3$)	iCASSCF($n = 4$)	iCAS-CI($n = 3$)	iCAS-CI($n = 4$)	CASSCF
Ethylene	-78.1747	-78.1742	-78.178	-78.1774	-78.1777
Formaldehyde	-114.0453	-114.0451	-114.0506	-114.0504	-114.0509
Hydrogen Cyanide	-93.042	-93.0412	-93.0452	-93.0439	-93.0443
Hydrogen Peroxide	-150.989	-150.9895	-150.9944	-150.9951	-150.9953
Methanol	-115.2211	-115.2221	-115.2258	-115.2259	-115.2261

To directly test the variationality of the iCASSCF procedure, finite difference gradients were computed using a five-point central difference formula and compared to the analytical gradients. The norms of the differences between analytical and numerical gradients are provided in Table 4.2. At $n = 3$, the largest difference between numerical and analytical gradients is $1.05 \times 10^{-4} \text{ Ha Bohr}^{-1}$. At $n = 4$, this maximum has a similar value of $8.25 \times 10^{-5} \text{ Ha Bohr}^{-1}$. The agreement between numerical and analytical gradients is satisfactory, as a

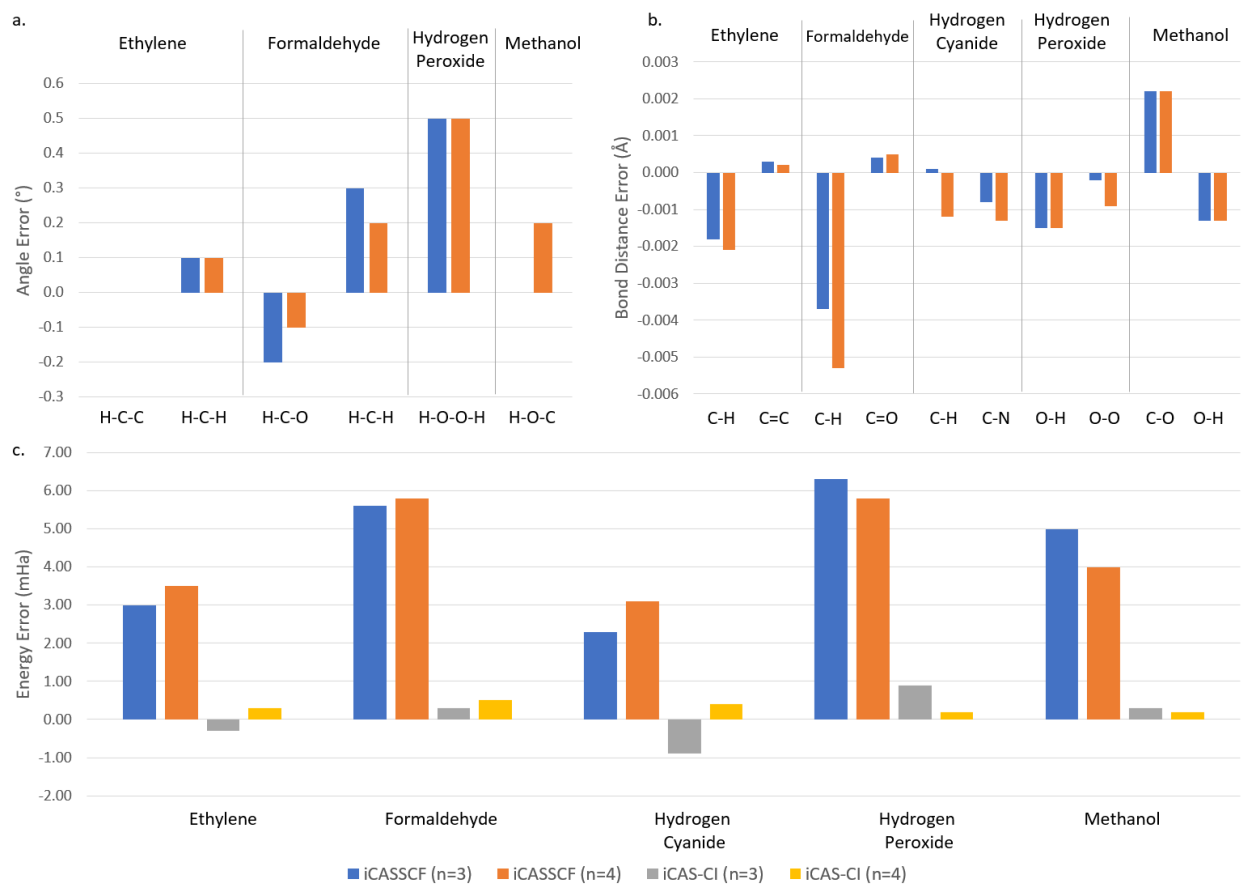


Figure 4.3: Comparison of iCASSCF to CASSCF (a) bond angles, (b) bond distances, and (c) energies. Included with the energy comparisons are iCAS-CI energies which use converged iCASSCF orbitals.

typical gradient threshold for a stationary geometry is on the order of 10^{-4} Ha Bohr $^{-1}$. The total error is similar to other analytical gradients relying on the RI approximation.^{209,210}

Table 4.2: Norm of difference between analytical and numerical gradient (Ha Bohr $^{-1}$).

calc	Ethylene	Formaldehyde	Hydrogen Cyanide	Hydrogen Peroxide	Methanol
iCASSCF(n=3)	5.59E-05	1.97E-05	6.43E-06	4.49E-05	1.05E-04
iCASSCF(n=4)	6.63E-06	2.36E-05	2.29E-06	8.25E-05	5.76E-05

4.5 oxoMn(salen)Cl

Transition metal complexes contain challenging electronic correlation effects due the interaction of d -electrons with the σ - and π -bonds of the ligands. In particular, the oxoMn(salen)Cl complex is a postulated intermediate in the Mn(III)(salen)-catalyzed epox-

idation reaction.^{196,197} This species is difficult to isolate and characterize due to its short lifetime, having a highly reactive oxo ligand. A truncated oxoMn(salen)Cl complex has been previously studied using density functional theory^{198,199} and multi-configurational methods including CASSCF,²⁰⁰ SA-CASSCF,²⁰¹ DMRG-SCF^{202–204} and stochastic HCI (SHCI).⁴¹ As will be discussed shortly, conclusions on the electronic structure of this complex are lacking, due to the large number of highly correlated electrons surrounding the central Mn=O bond. Fortunately, the full valence active space of 84 electrons in 84 orbitals is approachable using iCASSCF (timing information is available in Appendix C). This example will therefore be the first test of full valence iCASSCF on a transition metal complex to date.

Table 4.3: The singlet-triplet gap of oxoMn(salen)Cl from previous multi-configurational studies are listed.

Method	Basis	Delta,
CASSCF(12,11) ²⁰⁰	6-31G*	2.8
SA-CASSCF(12,11) ²⁰¹	6-31G*	3.0
DMRG-SCF(28,22) ²⁰²	6-31G*	-5.0
DMRG-SCF(28,22) ²⁰²	cc-pVDZ	-5.3
DMRG-SCF(26,21) ²⁰³	cc-pVDZ	0.4
DMRG-SCF(28,22) ²⁰⁴	cc-pVDZ	0.7
DMRG-SCF(28,22) ²⁰⁴	cc-pVTZ	1.9
SHCI(28,22)* ⁴¹	cc-pVDZ	0.9

*SHCI calculation used the DMRG-SCF(28,22) orbitals from ref 204

Several studies, starting with one by Linde et al., suggested that multiple spin states of oxoMn(salen)Cl are relatively close together in energy (Table 4.3). That study used density functional theory to predict that the singlet, triplet, and quintet are within 2.6 kcal mol⁻¹ of one another.¹⁹⁸ Alternatively, Abashkin et al. found the oxoMn(salen)Cl system to have a singlet state lower than the triplet state by 6.7 kcal mol⁻¹.¹⁹⁹ Using CASSCF(12,11), Ivanic et al. predicted the triplet to be more stable than the singlet by 2.8 kcal mol⁻¹, although they were unable to converge the CASSCF(12,11) triplet wave function due to a large gradient between the d_{xy} orbital in the active space and an inactive orbital on chlorine.²⁰⁰ Instead the d_{xy} orbital was omitted from the active space and a CASSCF(10,10) wavefunction was computed, followed by orbital localization and a CAS-CI(12,11) calculation so that the triplet energy could be compared. A later study by Sears and Sherill was able to converge the triplet wavefunction using a (12,11) active space, but required state averaging to do so.²⁰¹ This SA-CASSCF(12,11) found the triplet to be higher than the singlet by 3.0 kcal mol⁻¹. DMRG-SCF studies gave singlet-triplet gaps ($E_3 - E_1$) of around -5 kcal mol⁻¹²⁰² and 0.4 kcal mol⁻¹.²⁰³ Another DMRG-SCF study by Sharma et al. predicts a singlet-triplet gap of 0.7

kcal mol⁻¹ and 1.9 kcal mol⁻¹ when using the cc-pVDZ and cc-pVTZ basis sets, respectively. Perturbative (PT) corrections to these energies flip the sign of the singlet-triplet gap in each basis set.²⁰⁴ The DMRG study by Stein et al. used a (26,21) active space,²⁰³ while the other DMRG studies used a (28,22) active space. SHCI (which utilized DMRG-SCF orbitals) resulted in a singlet-triplet gap of 0.9 kcal mol⁻¹.⁴¹ The conflicting predictions (Table 4.3) for the sign of the singlet-triplet gap—and its value—demonstrate the challenges of applying electronic structure theory in this strongly correlated system. All of these spin gaps could easily be perturbed by additional correlation from out-of-active-space orbitals, as found by the DMRG-PT results, leaving the correct description of correlation in this complex yet to be resolved.

Any active space smaller than the full valence limit is possibly insufficient to provide an accurate depiction of the electronic structure of the oxoMn(salen)Cl complex. The full valence active space of the truncated oxoMn(salen)Cl model has 84 electrons in 84 orbitals, which would result in a CAS involving 10⁴⁷ electronic configurations. By comparison, a (12,11) active space—used in the conventional CASSCF studies^{200,201}—only has 6 × 10⁴ electron configurations, while a (28,22) active space—used in the DMRG studies^{41,202,204}—has approximately 10¹⁰ configurations. On the other hand, the full valence active space of the oxoMn(salen)Cl complex can be explored with iCASSCF without bias from active space truncation.

In this study, the O lone pair orbitals and Mn-O σ - and π -systems of oxo-Mn bond are highly correlated and therefore combined as a single body. This selection was made by grouping the initial (simple) bodies that had the largest 2-body correlation energies. As with prior examples in this article, the electronic total energy is computed for the oxoMn(salen)Cl model by obtaining orbitals at the iCASSCF level and using the converged orbitals in a subsequent iCAS-CI calculation. The total energies and singlet-triplet gaps are listed in Table 4.4. At the highest level of iCASSCF theory, $n = 4$, iCAS-CI additionally recovers 65.0 mHa of correlation for the singlet and 55.2 mHa for the triplet state. Although iCASSCF($n = 4$) predicts the triplet to be lower in energy than the singlet by 1.0 mHa (0.6 kcal mol⁻¹), inclusion of the single excitations in iCAS-CI($n = 4$) results in a singlet ground state with the triplet 8.8 mHa (5.5 kcal mol⁻¹) higher in energy. The $n = 4$ gaps agree with the $n = 3$ values to a close margin, suggesting that the n -body expansion of iCAS is adequately converged. Since Table 4.1 shows that the iCAS-CI energies are closer to CASSCF than those of iCASSCF, a singlet ground state appears to be preferred by the oxo-Mn(salen)Cl.

Figures 4.4 and 4.5 show the iCASSCF($n = 3$) natural orbitals (NOs) for the π spaces of the singlet and triplet, respectively. These orbitals are local in nature, without the expected mixing that usually characterizes NOs. This lack of mixing is due to lack of singles excitations

Table 4.4: Total energies of the oxoMn(salen)Cl complex from iCASSCF and iCAS-CI using the 6-31G* basis set.

State	iCASSCF (n=3)	iCAS-CI (n=3)	iCASSCF (n=4)	iCAS-CI (n=4)
Singlet	-2,252.2588	-2,252.3224	-2,252.2590	-2,252.3240
Triplet	-2,252.2599	-2,252.3144	-2,252.2600	-2,252.3152
Gap	-0.0011	0.0081	-0.0010	0.0088




Figure 4.4: Select iCASSCF($n = 3$) natural orbitals in the π space for the singlet state of the oxo-Mn(salen)Cl complex. The natural occupation numbers are displayed below each orbital and the HONO and LUNO are boxed.

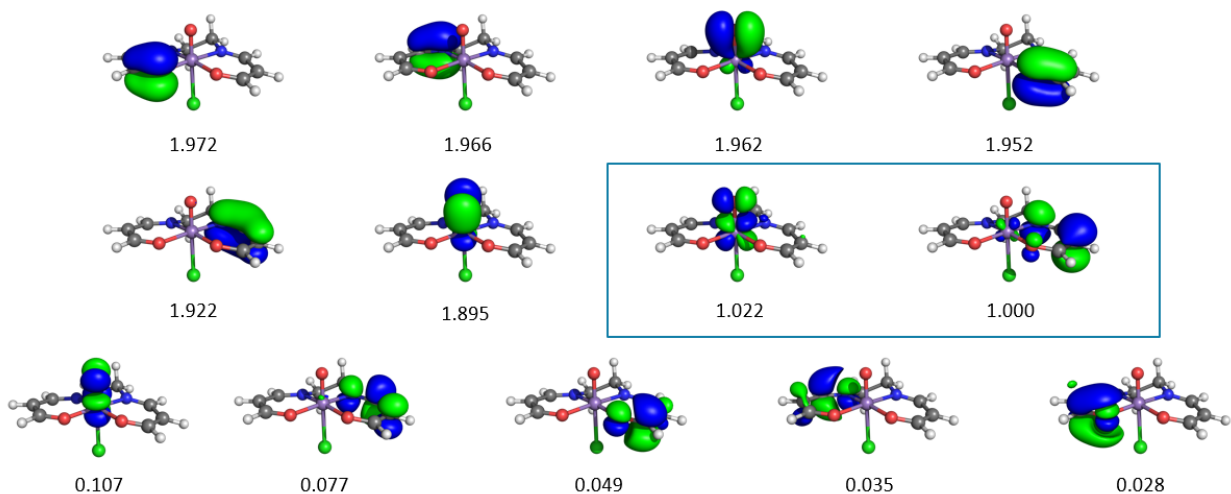


Figure 4.5: Select iCASSCF($n = 3$) natural orbitals in the π space for the triplet state of the oxo-Mn(salen)Cl complex are shown. The occupation of each orbital is provided below the orbital. The SONOs are boxed.

at the iCASSCF level of theory, and more conventional NOs emerge when using iCAS-CI instead (see Appendix C, Figures C.3 and C.4). Close inspection of the triplet NOs (Figure 4.5) reveals that the triplet state is a ligand-to-metal charge transfer (LMCT) state.

The identification of a LMCT triplet state instead of a $d_{xy}-\pi^*$ (Mn-O) triplet is surprising, given that the latter is predicted by all prior studies.^{200–204} Two reasons are suggested to explain this finding: 1) The full valence active space captures more correlation, allowing the LMCT state to drop lower in energy than the $d_{xy}-\pi^*$ state, or 2) either state is the optimized SCF state, but not necessarily the true lowest-energy triplet. Either way, the LMCT state is predicted to be close in energy to the singlet ground state by iCAS, suggesting it is accessible under relevant experimental conditions. Prior studies have indeed found LMCT states to be common for metal-salen complexes and metal-salen analogues,^{211–214} which supports this possibility. Additional studies involving multi-state iCAS (or other highly correlated levels of theory) will be needed to better resolve this challenging situation.

4.6 Cyclobutadiene Automerization

The automerization of cyclobutadiene (see Figure 4.6) from one D_{2h} symmetry rectangle to another is a challenging, strongly correlated electronic structure problem. In the reactant state the rectangular D_{2h} starting geometry can be described by Jahn-Teller effect, producing a lower symmetry than the square planar D_{4h} geometry due to anti-aromaticity of the latter structure. The four π orbitals are strongly correlated due to near degeneracy of the HOMO/LUMO, giving cyclobutadiene a biradical electronic state.^{215,216} While a majority of prior studies have assumed the TS takes a square planar shape with D_{4h} symmetry,^{153–160} one RASSCF study predicted the TS to be an isosceles trapezoid with C-C-C bond angles of 90.16° and 89.84° .¹⁶¹ These prior theoretical studies predict automerization barriers ranging between 4.0 and 21.0 kcal mol⁻¹,^{153–161} while experimental results place the barrier between 1.6 and 10 kcal mol⁻¹.²¹⁷ To provide insight from iCASSCF into this reaction, the Growing String Method (GSM) was applied to compute a reaction pathway and exact transition state (i.e. saddle point).

The iCASSCF energy profile for the $D_{2h} \rightleftharpoons D_{2h}$ pathway is shown in Figure 4.6. The automerization barrier is predicted to be approximately 11 kcal mol⁻¹ at the iCASSCF (geometry optimization) and iCAS-CI (single point energy) levels of theory, both of which are just outside of the estimated experimental bounds.²¹⁷ The singles excitations account for between 8.7 mHa and 9.1 mHa of correlation along the pathway, suggesting that the errors due to omitting singles largely cancel out.




Figure 4.6: Optimized reaction pathway via GSM and iCASSCF($n=3$)/6-31+G^{**}. Single point energies from iCAS-CI($n=3$) are nearly identical, and are shown in red.

Figure 4.7 shows that the TS takes on an almost square structure, with the C-C-C bond angles being within 0.01° of right angles and the C-C distances between equivalent to within 0.001 \AA . The TS geometry is not fully symmetric, however, as the carbon atoms are slightly non-planar with a C-C-C-C dihedral angle of 0.12° and the hydrogen atoms come out of plane at angles between 0.5° and 3.1° . To test the energetic impact of this distortion against the expected fully symmetric TS geometry, the energy was computed for the D_{4h} structure closest to the GSM-optimized TS. Using C-C and C-H bond lengths of 1.4672 \AA and 1.0930 \AA respectively, the TS energy goes down slightly, by less than 0.1 mHa at iCASSCF and iCAS-CI levels of theory. While GSM and the previous RASSCF study do find slightly non-symmetric structures for the TS, these are likely just the result of the PES being nearly flat near the D_{4h} geometry. Either way, iCASSCF geometry optimization was sufficiently accurate to place cyclobutadiene in a nearly D_{4h} configuration at the TS, which was within a negligible margin (0.1 mHa) of energy from the fully symmetric TS.




Figure 4.7: Optimized transition state geometry for cyclobutadiene automerization from GSM. The C-C bond lengths are provided in Å, and some key bond angles and dihedrals are shown.

4.7 Conclusions

The iCASSCF method is designed to bypass the exponential wall of the full CI problem and permit CAS-like electronic structure simulations with all valence electrons correlated. In the proposed fully variational iCASSCF procedure, single excitations are omitted from the wave function, which allows orbital optimization to make the energy fully stationary with respect to active-active rotations. This variational procedure not only has a well-defined gradient, but it also permits geometry optimization to the same level of accuracy as its parent theory, conventional CASSCF. Accurate energies are easily evaluated by performing an iCAS-CI calculation *with* single excitations, using the optimized orbitals from iCASSCF.

iCASSCF has the ability to treat full valence active spaces for relatively large molecules, which results in unambiguous active space choices when valence bond orbitals are used as the MBE bodies. This facilitated the incorporation of iCASSCF with the growing string method, allowing for reaction path optimization of strongly correlated systems such as cyclobutadiene automerization. A TS search by GSM results in a TS which is not fully symmetric, however the D_{4h} geometry is stabilized by less than 0.1 mHa relative to the GSM TS. In other words, the PES near the D_{4h} geometry is nearly flat with respect to small out-of-plane bending of the hydrogen.

While our previous study had shown the ability of iCASSCF to treat large active spaces (116,116),⁶⁹ the variability of the new procedure has a profound effect. Using our previous implementation—which relied on a numerical prescription to deal with the lack of invariance to active-active rotations—orbital optimization could be troublesome for some chemical systems, for instance the oxoMn(salen)Cl complex. The new combined iCASSCF/iCAS-CI procedure is able to tame the strong correlation present in the oxoMn(salen)Cl complex with

an (84,84) active space, predicting a singlet ground state with a triplet state about 5 kcal mol⁻¹ higher in energy. Furthermore, the low-lying triplet state was surprisingly predicted to be a LMCT state, in contrast to the $d_{xy}-\pi^*$ (Mn-O) state predicted by previous studies. In summary, this article's modifications to iCASSCF retain its ability to treat large full valence active spaces, and open the door to examining the challenging electronic structure of transition metal complexes.

Introducing variationality into iCASSCF therefore has made the method more robust in computing strongly correlated systems. Further development and optimization is expected to continue expanding the scope of systems and chemical processes that may be studied with the iCASSCF family of methods.

CHAPTER 5

The Numerical Evaluation of Slater Integrals on Graphics Processing Units

The previous chapters have showcased computational and theoretical advances in quantum chemistry methods aimed at recovering electron correlation for large electronic systems. In this chapter, we instead target the basis set, which is essential for accurately representing the electron distribution. In most quantum chemical calculations, atom-centered Gaussian-type orbitals (GTOs) are used due to the availability of fast, analytic electron integrals. However, GTOs have incorrect short-range and long-range behaviors, thus they fail to capture the electron density properly. A better choice of basis is the Slater-type orbital (STO), which can represent the expected behaviors of molecular wave functions. However, the lack of analytic electron integrals have hindered its wide-spread adoption. This chapter presents SlaterGPU, a GPU accelerated library that uses OpenACC to numerically compute Slater-type orbital (STO) integrals and is based on previously published work from the *Journal of Computational Chemistry*.²¹⁸

5.1 Introduction

Advances in quantum chemistry and computer hardware have facilitated the routine use of electronic structure simulations for chemical applications. Some of the most widely used theories make use of one-electron, atom-centered basis functions² to represent the electron density. The simplest wave function that approximately solves the Schrödinger equation is Hartree-Fock (HF), which represents the wave function using a single Slater determinant. While HF is not a quantitatively accurate theory, it forms the basis for more sophisticated theories. In many canonical, post-HF methods, evaluation of Hamiltonian elements in the Schrödinger picture requires computing integrals of the form

$$O_{\mu\nu} = \langle \chi_\mu | \hat{O}_1 | \chi_\nu \rangle = \int \chi_\mu(\mathbf{r}) \hat{O}_1(\mathbf{r}) \chi_\nu(\mathbf{r}) d\mathbf{r}, \quad (5.1)$$

$$\begin{aligned} O_{\mu\nu\lambda\sigma} &= \langle \chi_\mu(1) \chi_\nu(1) | \hat{O}_2 | \chi_\lambda(2) \chi_\sigma(2) \rangle \\ &= \int \int \chi_\mu(\mathbf{r}_1) \chi_\nu(\mathbf{r}_1) \hat{O}_2 \chi_\lambda(\mathbf{r}_2) \chi_\sigma(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2. \end{aligned} \quad (5.2)$$

The first equation denotes 1-electron quantities such as the overlap $\hat{O}_1 = 1$, the kinetic energy $\hat{O}_1 = -\frac{1}{2}\nabla^2$, and the nuclear attraction $\hat{O}_1 = \frac{Z_A}{R_{1A}}$ operators. 2-electron operators include the Coulomb repulsion $\hat{O}_2 = \frac{1}{r_{12}}$, where r_{12} is the distance between electrons 1 and 2. Derivatives of these terms, for example with respect to nuclear position, are also quantities of interest.

Amongst these integrals, the electron repulsion integrals (ERIs) are the most difficult (and numerous) to evaluate, being 2-electron quantities that require six-dimensional integration. In addition, the $\frac{1}{r_{12}}$ operator contains a singularity at every point in three-dimensional space, further challenging their integration. Consequently, the choice of basis is important for not only accurate representation of the molecular wave function, but also for computational evaluation of integrals.

One physically-motivated choice are Slater-type orbitals (STOs), which are hydrogen-like orbitals of the form

$$S(\zeta, n, l, m, r, \theta, \phi) = N^{\text{STO}} r^{n-1} e^{-\zeta r} Z_{lm}(\theta, \phi), \quad (5.3)$$

where ζ is the exponent, n, l, m are the usual atomic quantum numbers, r, θ, ϕ are spherical coordinates, N^{STO} is the normalization constant, and Z_{lm} are the spherical harmonics.²⁻⁴ The STOs satisfy the Kato cusp and exponential decay of atomic wave functions,^{2,219,220} making them a natural basis choice for quantum chemical calculations. However, the ERIs over STOs do not have a known general analytic form.

The difficulties of STO integration led to the expansion of STOs in terms of Gaussian-type orbitals⁵ (GTOs)

$$G(\alpha, n, l, m, r, \theta, \phi) = N^{\text{GTO}} e^{-\alpha r^2} S_{lm}(r, \theta, \phi), \quad (5.4)$$

where the S_{lm} are the real solid harmonics.² The GTOs benefit from the Gaussian product rule, i.e. the product of two GTOs is again a GTO, which simplifies Equation 5.2 from a 4-center, 2-electron integral to 2-center, 2-electron. The 2-center, 2-electron integral can be evaluated over the Coulomb potential of one GTO reducing a 6-dimensional integral to 3-dimensions. These nice analytical properties of GTOs facilitated the development of fast analytical integral evaluation.²²¹⁻²²⁵

While GTOs can be quickly evaluated using modern integral libraries, they do not contain the correct short- and long-range behaviors expected in molecular wave functions.⁹ For example, the cusp near the nucleus is important for computing properties such as nuclear magnetic resonance shifts and polarizabilities,^{6,7} but the cusp is not present in the GTO basis, and only crudely treated by using contracted sets of GTOs. Exponential decay of the wave function for an accurate description is required for precise quantification of the HOMO energy, but this behavior is also absent in GTOs.⁸

The imperfections of GTO basis sets have left room for the continued development and use of STOs for quantum chemical applications. Several schemes have been developed to compute general STO integrals. One approach is to expand each STO in a very large number of GTOs and compute the GTO integrals analytically.^{226–228} Additionally, Monte Carlo has been used to correct integrals over Gaussian expansions to evaluate the Slater quantity.²²⁹ These schemes are prohibitively expensive for routine use. While the focus of this article is on the use of STOs in integrals such as Equations 5.1 and 5.2, STOs have seen frequent use in quantum Monte Carlo (QMC) wave functions, where the 1- and 2-electron integrals are not important.^{230–233}

An attractive alternative to explicit evaluation of STO ERIs involves density fitting—in particular the resolution-of-the-identity (RI) approximation (see Theory and Computational Details)—which allows Equation 5.2 to be approximated as a tensor product of 2- and 3-index ERIs. Within the RI approximation, one of the two electrons is described by a single basis function. This facilitates the use of a Coulomb potential to represent one electron without relying on an explicit basis set product rule—which does not exist for STOs—to condense multiple centers. This simplification, which is only necessary for systems with at least four distinct atomic centers, allows STO integration of ERIs to be amenable to numerical quadrature schemes. The Amsterdam Density Functional (ADF) package and other density functional theory (DFT) codes implement a density fitting approach to use STOs in DFT.^{4,6,234,235} Other density fitting frameworks have allowed STOs to be used in approximate MP2²³⁶, double-hybrid DFT²³⁷, and Green’s function methods^{238,239}. These previous STO studies, however, did not generate the full complement of ERIs required for multiconfigurational methods such as those based on configuration interaction,^{40,79–81} multiconfigurational self-consistent field^{47,48,69,70} and coupled cluster.^{72,240–242}

This chapter introduces and benchmarks a graphics processing unit (GPU) library for evaluating STO integrals for wave function theories. The chapter will show that these can be accurately and efficiently evaluated using numerical integration by combining the RI approximation with the STO Coulomb potential. The large number of processing cores and high memory bandwidth make modern GPUs the architecture of choice for evaluating

and summing numerical grids. For additional performance, the integrals are also computed using mixed-precision evaluation. Timings suggest that this library allows STOs to be useful alongside strongly correlated wave function theories. Accuracy benchmarks indicate minimal loss in accuracy from using mixed-precision relative to double-precision. The resulting code, called SlaterGPU,²⁴³ is the first reported library to use GPUs to accelerate STO integrals and evaluate the full set of 1- and 2-electron STO integrals up to the $6h$ subshell as well as $5g$ for first derivatives for the auxiliary basis.

5.2 Theory and Computational Details

The present STO integral scheme relies on numerical integration over atom-centered grids. Grid-based integration can make use of single instruction, multiple data (SIMD) parallelism and therefore can leverage GPU hardware for acceleration. Even with this acceleration, the 6-dimensional ERIs remain too costly for routine computations. The dimensionality of integration can be reduced, however, by employing the resolution-of-the-identity (RI) approximation,^{244–246} where the Coulomb potentials for the auxiliary basis functions are known analytically. The various components of the STO integral algorithm are explained in the following sections: the Resolution of the Identity, Grid Construction, Implementation on GPU, and Computational Details.

5.2.1 Resolution of the Identity

This section focuses on simplifying the challenging ERIs for numerical evaluation. The expressions for numerically evaluating the 1-electron integrals are listed in Section D.1 of Appendix D. In the RI approximation, the 4-index ERIs $(\mu\nu|\lambda\kappa)$ are decomposed into tensor products of 2- and 3-index integrals by representing the density in terms of an auxiliary basis. Using the Coulomb metric, the integral can be approximated with the expression^{244–246}

$$(\mu\nu|\lambda\kappa) \approx \sum_{PQ} (\mu\nu|P)(PQ)^{-1}(Q|\lambda\kappa) = \sum_Q B_{\mu\nu}^Q B_{\lambda\kappa}^Q, \tag{5.5}$$

where

$$B_{\mu\nu}^Q = \sum_P (\mu\nu|P)(PQ)^{-1/2}. \tag{5.6}$$

In a numerical integration scheme, Equation 5.5 not only reduces the count of numerical integrals for a given basis set size (N) from $O(N^4)$ to $O(N^3)$, it also has a secondary

consequence that is useful in the context of STO basis functions. Specifically, the integral

$$(P|\mu\nu) = \int \int \chi_P(\mathbf{r}_1) \frac{1}{r_{12}} \chi_\mu(\mathbf{r}_2) \chi_\nu(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2 \quad (5.7)$$

can be simplified to

$$(P|\mu\nu) = \int V_C^P(\mathbf{r}) \chi_\mu(\mathbf{r}) \chi_\nu(\mathbf{r}) d\mathbf{r} \quad (5.8)$$

by using the known analytical form of the single-Slater Coulomb potential. In spherical coordinates, this potential has the form⁴

$$V_C(\zeta, n, l, m, r, \theta, \phi) = \frac{4\pi(2\zeta)^{n+(1/2)}}{\sqrt{(2n)!(2l+1)}} Z_{lm}(\theta, \phi) I_{nl}(r), \quad (5.9)$$

where

$$I_{nl}(r) = r^{-l-1} \int_0^r (r')^{n+l+1} e^{-\zeta r'} dr' + r^l \int_r^\infty (r')^{n-l} e^{-\zeta r'} dr'. \quad (5.10)$$

I_{nl} has analytic expressions using finite Laurent polynomials for each n, l of interest (See Section D in Appendix D).

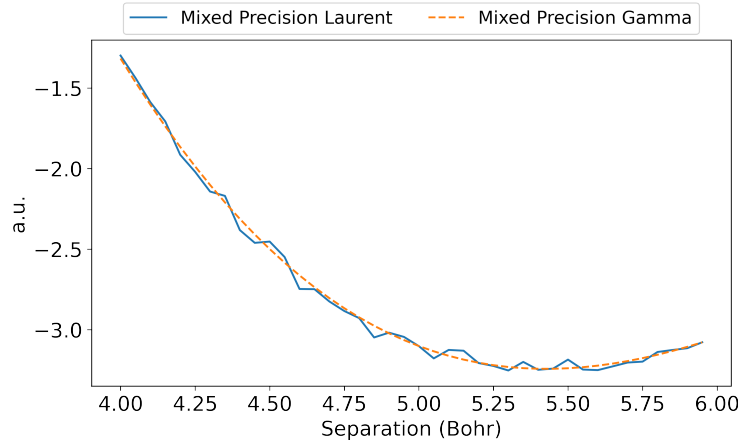


Figure 5.1: The $(6H|6H)$ integral is scanned in the $(0.370, 0.370, 0.853)$ direction with the left center at the origin. Evaluations are in mixed precision using either the Laurent polynomial expansion of Equation 5.10 or the lower incomplete gamma function, where mixed precision is defined similarly to Equation 5.14. Both basis functions have $m = 0$ and $\zeta = 1$.

For large angular momentum l , the Laurent expressions (See Section D.2 in Appendix D)—and especially their derivatives—exhibit numerical instability, especially when using mixed-precision arithmetic, which is essential for high performance integral evaluation. This can result in non-smooth integrals as shown in Figure 5.1, which can in turn result in non-smooth or discontinuous energies. Instead of applying the Laurent expressions, Equation 5.10

can be evaluated using lower incomplete gamma functions, which have fast, numerically precise implementations²⁴⁷. The final form of Equation 5.10 used in the current implementation of SlaterGPU is

$$I_{nl}(r) = r^{-l-1} \zeta^{-l-n-2} \{ (r\zeta)^{2l+1} [(-l+n)! - \gamma(-l+n+1, r\zeta)] + \gamma(l+n+2, r\zeta) \}, \quad (5.11)$$

where $\gamma(s, x)$ is the lower incomplete gamma function,

$$\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt. \quad (5.12)$$

After evaluation of all 2- and 3-center Coulomb integrals, the full set of 4-index ERIs can be reconstructed using Equation 5.5. SlaterGPU therefore uses the RI approximation for Slater integrals, similar to prior implementations for DFT applications,^{234,235} but further provides all 4-index integrals, $(ij|kl)$, which are not generated or required for DFT. This allows the SlaterGPU library to be useful for wave function theories, which require a larger set of ERIs. In particular, while prior codes demonstrated applicability to $l \leq 3$,^{235,239} SlaterGPU is shown here to be useful for $l \leq 5$.

5.2.2 Grid Construction

When numerically evaluating integrals over atomic orbitals (AO), the choice of grid is important. The atom-centered grids used here borrow their core concept from prior studies, especially those involving integration of DFT functionals.^{234,248–252} The accepted route for integrating the exchange-correlation energy is to build atomic grids as products of radial and angular grids, then reweight these using Voronoi polyhedra centered about the nuclei. The atom-centered grids are necessary to capture the spherical harmonics and radial decay of atomic orbitals, while partitioning 3-dimensional space into polyhedra divides the grid into volumes centered around each nucleus. The Voronoi boundaries are smoothed and reweighted to avoid double counting of volume elements.^{249,250,252} This same framework is used in SlaterGPU, though only a maximum of three atom-centered grids are required for any given integral, since the ERIs only involve up to three centers at a time in the RI approximation. In a polyatomic system, this greatly simplifies the form of the integration grid, keeping each integral grid small enough to be efficiently evaluated. The grids chosen for this implementation are the "Log3" grid from Mura and Knowles²⁵² for the radial component, and the Lebedev grid²⁵³ for the angular component. Both grids are widely used in electronic structure codes. Once each atom-centered grid is generated, and the Becke partitioning scheme²⁴⁹ is applied, the 3-center integral $(\mu\nu|P)$ can be evaluated over the grid points x

and grid weights $w(x)$ as

$$(\mu\nu|P) = N_{V_C}^{\text{STO}} N_{\chi_\mu}^{\text{STO}} N_{\chi_\nu}^{\text{STO}} \sum_x \bar{V}_C(x) \bar{\chi}_\mu(x) \bar{\chi}_\nu(x) w(x), \quad (5.13)$$

where \bar{V}_C , $\bar{\chi}_\mu$, $\bar{\chi}_\nu$ are the Coulomb potential and basis functions with their respective normalization constants, $N_{V_C}^{\text{STO}}$, $N_{\chi_\mu}^{\text{STO}}$, and $N_{\chi_\nu}^{\text{STO}}$, factored out. While Lebedev and Mura-Knowles grids are used with Becke weights, any quadrature grid and weighting scheme can be used in Equation 5.13. The 2-index integrals $(P|Q)$ are evaluated in a similar manner.

5.2.3 Implementation on GPU

All integral code in this chapter is written in C++ using OpenACC for GPU acceleration, which has the advantage of being based on pragma directives allowing the same code base to be compiled to run on CPUs or GPUs. When evaluating Equation 5.10, a modified version from the Cephys library²⁵⁴ was used for the lower incomplete gamma function, noting that OpenACC allows these implementations to be used directly. Most GPUs contain more single precision compute units than double precision, so mixed precision operations are an attractive choice in a practical implementation.^{255–258} For example, the 2080-Ti contains $\frac{1}{32}$ the double precision units compared to single precision units, while the GV100 contains $\frac{1}{2}$ the double precision units compared to single precision units. In SlaterGPU, mixed precision is available, where evaluations over the grid are performed using single-precision arithmetic, and the final summation occurs in double precision. In mixed precision, Equation 5.13 becomes

$$(\mu\nu|P)_{64} = N_{V_C}^{\text{STO}} N_{\chi_\mu}^{\text{STO}} N_{\chi_\nu}^{\text{STO}} \sum_x \bar{V}_P(x)_{32} \bar{\chi}_\mu(x)_{32} \bar{\chi}_\nu(x)_{32} w(x)_{32}, \quad (5.14)$$

where the subscript refers to the bits of precision of the quantity. Factorization of the normalization constants reduces the number of floating point operations required, which is essential for high performance. In addition to making use of the greater quantity of single-precision compute units in GPUs, single precision also reduces storage and memory bandwidth demands by a factor of 2. A generalization of the mixed precision procedure would be to adaptively determine which integrals to evaluate at each level of precision, as has been done in (analytic) GTO integration.²⁵⁸ This is not done here; instead, the accuracy of the mixed precision approach is evaluated in comparison to double precision integration.

Algorithm 5.1 GPU compute structure for generating 3-center ERIs

```
1: #pragma omp parallel for schedule(dynamic)
2: for  $A, B, C$  in Atom List //  $A, B$  over all Atoms,  $C \geq B$ 
3:   Generate  $x_A, x_B, x_C, w(x_A), w(x_B), w(x_C)$ 
4:    $x \leftarrow x_A \cup x_B \cup x_C$ 
5:    $w(x) \leftarrow w(x_A) \cup w(x_B) \cup w(x_C)$ 
6:   for  $P_i$  in Aux( $A$ )
7:     Compute  $V_{P_i}(x)$ 
8:   for  $\chi_{\mu_j}$  in Basis( $B$ )
9:     Compute  $\chi_{\mu_j}(x)$ 
10:  for  $\chi_{\nu_k}$  in Basis( $C$ )
11:    Compute  $\chi_{\nu_k}(x)$ 
12:  for  $P_i \in \text{Aux}(A), \mu_j \in \text{Basis}(B), \nu_k \in \text{Basis}(C)$ 
13:     $(P_i|\mu_j\nu_k) \leftarrow \sum_x V_{P_i}(x)\chi_{\mu_j}(x)\chi_{\nu_k}(x)w(x)$ 
```

Aux(\cdot) and Basis(\cdot) denote the set of auxiliary and main basis functions centered at \cdot , respectively.

In the GPU computing framework, data transfers between CPU and GPU incur large overhead penalties, and thus it is necessary to minimize these transactions for maximum performance. As such, all quantities in Equation 5.14 are generated and evaluated directly on GPU. The grid x and its weights $w(x)$ only depend on the set of atoms and not the basis functions, so these are generated once for each unique triad of atoms as described in Algorithm 5.1. Additional computations may be avoided by evaluating each V_P , χ_μ and χ_ν on the grid only once per triad of atoms. In other words, when evaluating $(P_i|\mu_j\nu_k)$, the quantities $V_P(x)$, $\chi_\mu(x)$ and $\chi_\nu(x)$ are all computed and stored as GPU arrays for all P_i on atom A , μ_j on atom B and ν_k on atom C to avoid duplicating computations. These arrays can then be contracted all at once in a single tensor operation as indicated in lines 11-12 of Algorithm 5.1. With OpenACC, the contraction on line 12 can be handled using a single *pragma* directive containing the parallel and reduction clauses. Sample OpenACC code is provided in Section D.8 of Appendix D. As the grid and weights are generated directly on GPU, reuse of the grid benefits from the high memory bandwidth of the GPU (~ 600 GB/s on the 2080-Ti). Once the integrals are computed on GPU, a single data transfer step returns the integrals to CPU memory. The code for numerically computing the STO integrals is freely available on GitHub under an Apache 2.0 license with Commons Clause as noted in the Data Availability Statement.

Multi-GPU parallelization is also implemented for a single node, using OpenMP to manage the GPU processes. Each OpenMP thread is assigned a GPU, and a manager-worker scheme is used for load balancing, where the work is partitioned using sets of atoms to take advantage

of grid/weight reuse. The parallelization occurs over the loop in Line 2 of Algorithm 5.1 and can be accomplished with a single *pragma* directive, shown in Line 1.

5.2.4 Computational Details

An all-electron double-zeta STO basis set with polarization functions²⁵⁹ was used (denoted DZP) as the primary atomic orbital basis. The auxiliary basis sets were taken from the same source. Full specification for the primary and auxiliary basis sets are provided in Section D.3 of Appendix D. Unless otherwise specified, the integration grid was a direct product of 60 radial points and 770 angular points (Lebedev order 18). Hartree-Fock and heat-bath configuration interaction^{40,43–45,134,162} (HCI) were used as representative electronic structure methods. The HCI parameters used are detailed in the following section. For GTOs, the 6-31G* basis with the RI-cc-pVTZ auxiliary basis was used. All GTO integral evaluation was performed using the Libcint library.²⁶⁰ Molecules were placed in standard nuclear orientations.²⁶¹ The Nvidia HPC SDK 20.7 compiler suite with CUDA 11.0 was used to compile all code. CPU code was run on Intel Xeon Gold 6242 processors clocked at 2.8 GHz and GPU code was executed using the Nvidia RTX 2080-Ti and GV-100 GPUs.

5.3 Results and Discussion

5.3.1 Performance Analysis

High throughput integral evaluation is necessary for any electronic structure theory code, regardless of basis set type. Grid-based numerical integration, however, requires orders of magnitude more floating point operations than analytical integration. To achieve the integral performance required, GPUs are used in this study for numerical integration of STO integrals. These integrals include all of the common 2-center integrals (overlap, electron-nuclear attraction, and kinetic energy) as well as the 2- and 3-center Coulomb integrals needed for the RI approximation. The relative speedup for numerical GPU integration compared to CPU integration is visualized in Figure 5.2. In double precision, the 2080-Ti can achieve over 30× speedup and the GV100 achieves ~ 70× speedup, allowing for tractable wall times for the integrals as listed in Table 5.1. Even further performance can be gained by utilizing mixed precision, showing speedups of over 60× and overall integral throughput increasing by a factor of ~ 4 for the 2080-Ti (see Table 5.2). The speedup relative to CPU drops slightly to ~ 55× on the GV100. The performance behavior is a consequence of the hardware configuration and the integral kernels being compute bound, with a detailed analysis provided in Section D.5 of Appendix D.

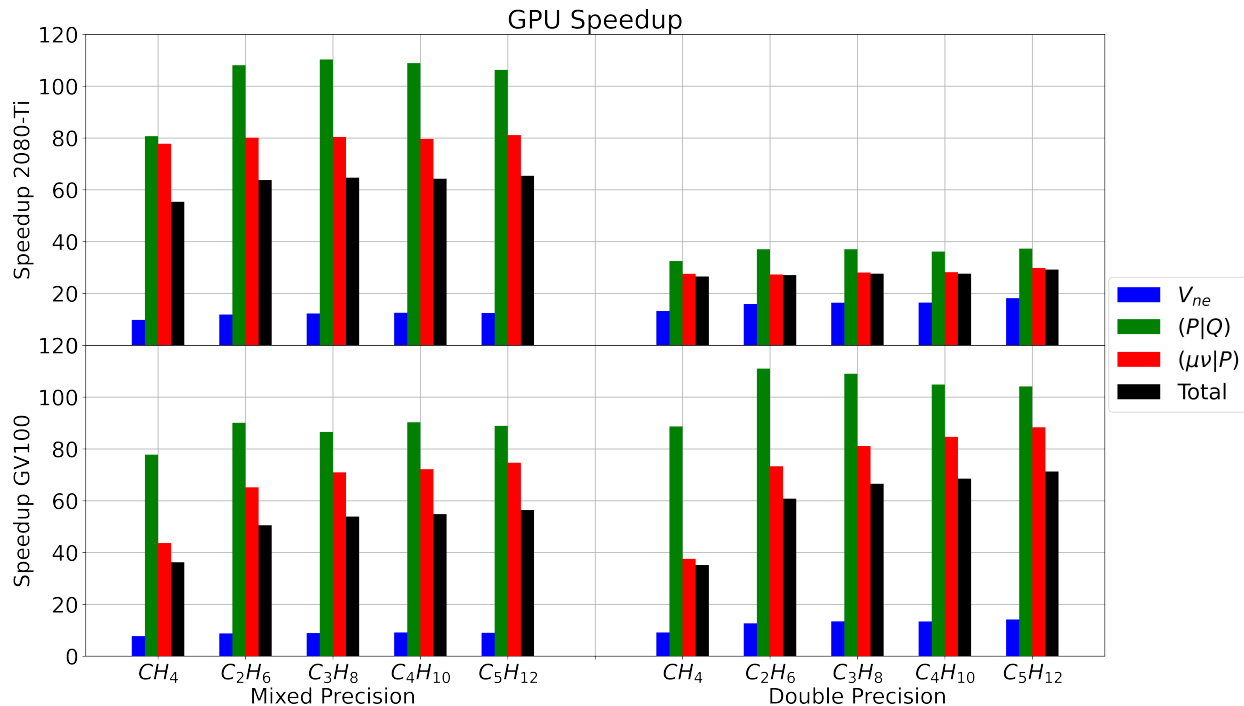


Figure 5.2: The GPU speedups for integral evaluation over CPU for the 2080-Ti (top) and GV100 (bottom) are shown for various alkanes using the DZP basis from ADF. The speedups are partitioned into the various integrals. Speedups for mixed (left) and double (right) precision evaluations are also shown.

Table 5.1: Double-precision timing data (in seconds) for various alkanes. Each atom contributes 46,200 grid points.

Molecule	Basis size		CPU Time			2080-Ti Time			V100 Time		
	Main	Aux	V_{ne}	$(P Q)$	$(\mu\nu Q)$	V_{ne}	$(P Q)$	$(\mu\nu Q)$	V_{ne}	$(P Q)$	$(\mu\nu Q)$
CH ₄	35	224	3.282	14.02	56.88	0.2483	0.4316	2.062	0.3583	0.1581	1.513
C ₃ H ₈	85	516	40.28	67.98	786.2	2.452	1.835	28.00	2.991	0.6237	9.693
C ₅ H ₁₂	135	808	166.1	163.5	3286	9.125	4.385	110.1	11.71	1.570	37.20

Table 5.2: Mixed-precision timing data (in seconds) for various alkanes. Each atom contributes 46,200 grid points.

Molecule	Basis size		CPU Time			2080-Ti Time			V100 Time		
	Main	Aux	V_{ne}	$(P Q)$	$(\mu\nu Q)$	V_{ne}	$(P Q)$	$(\mu\nu Q)$	V_{ne}	$(P Q)$	$(\mu\nu Q)$
CH ₄	35	224	2.095	8.686	34.34	0.2139	0.1076	0.4415	0.2708	0.1117	0.7867
C ₃ H ₈	85	516	24.95	42.34	484.8	2.037	0.3838	6.029	2.795	0.4892	6.832
C ₅ H ₁₂	135	808	95.62	98.58	1957	7.680	0.9275	24.12	10.58	1.109	26.21

Faster integral evaluation is also possible by distributing the workload across multiple GPUs. To test multi-GPU scaling, 3-center ERIs were evaluated for C_9H_{20} , which has 76.9 million ERIs, taking 125s to compute in mixed precision and 542s in double-precision. Figure 5.3 shows the strong scaling performance when evaluating the 3-center ERIs for C_9H_{20} , which maintains parallel efficiency greater than 75% on up to 5 GPUs for mixed-precision

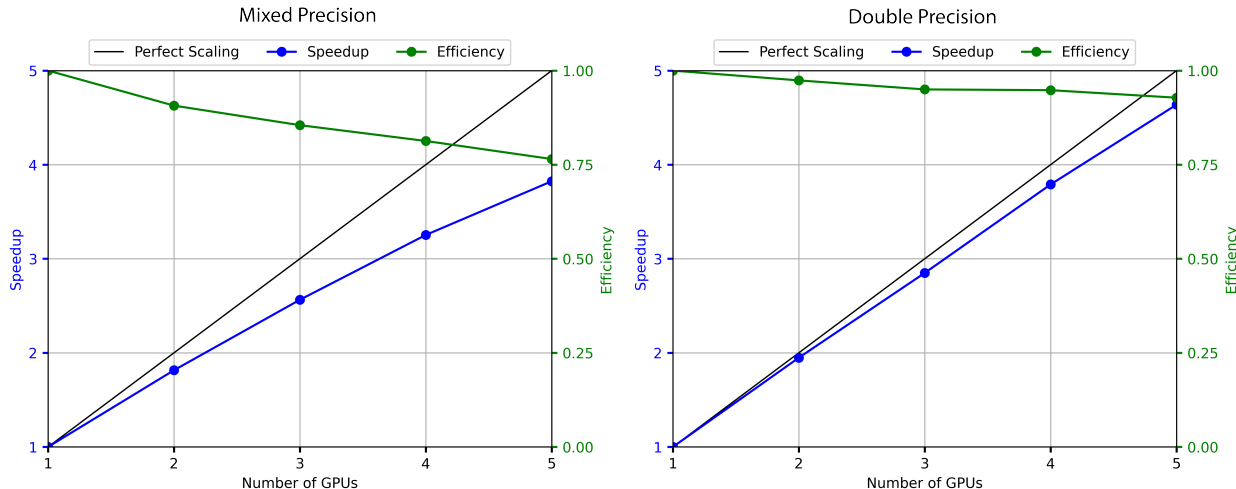


Figure 5.3: Multi-GPU speedups over single GPU and parallel efficiency for mixed (left) and double (right) precision evaluation of the 3-center ERIs for C_9H_{20} . There are 76,873,200 3-center integrals. Perfect scaling is plotted as a solid black line. All GPUs are co-located on a single compute node. Single GPU run times were 125s and 542s for mixed- and double-precision implementations, respectively.

evaluation and greater than 90% for double-precision evaluations in this benchmark. Due to the reduced computational demand of mixed-precision integration, the serial components and communication overhead take up a proportionally larger amount of computational time. Consequently, the parallel efficiency for mixed-precision integral evaluation drops off more rapidly than for double precision in strong scaling tests. However, this parallelization scheme still allows STO integration to achieve greater than 75% parallel efficiency and overall integral throughput greater than 3 million integrals per second in mixed precision on 5 GPUs. For comparison, Sun reported a throughput of approximately 6-8 million explicitly calculated integrals per second per thread with Libcint,²⁶⁰ thus placing STO integral throughput within reach of analytical integral evaluation for GTOs. While the Libcint performance was reported for 4-center ERIs, the comparison demonstrates the feasibility of Slater integration under the RI approximation. Additional developments in code optimization²⁶² and screening protocols²⁵⁶ may narrow this gap further. Other grid-based STO integral evaluation implementations^{6,235-239} do not report timings nor do they use GPU acceleration. The closest available performance comparison is an example where a 9-Gaussian

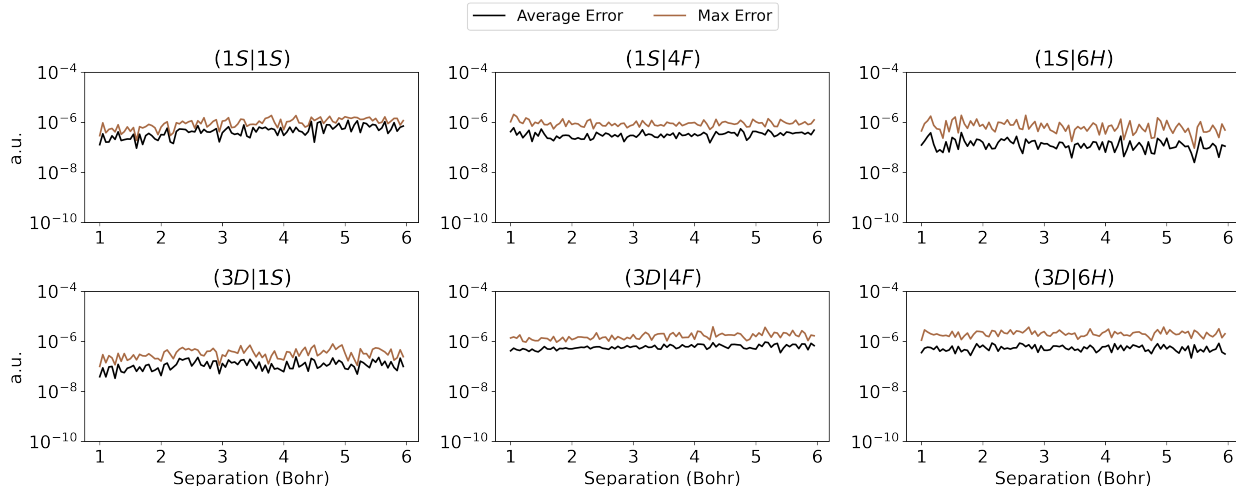


Figure 5.4: The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.

expansion was used to approximate STO integrals,²²⁸ which would reduce throughput, relative to GTOs, by approximately a factor of 700 under the RI approximation.

The large performance gain ($\sim 4\times$ speedup) when using mixed precision on the 2080-Ti units necessarily comes with some loss in accuracy compared to double precision arithmetic. Therefore tests of the mixed-precision integral evaluation are needed, in order to gauge the quantitative trade-off between accuracy and speed.

5.3.2 Mixed-Precision Evaluation

Numerical evaluation of integrals, whether done in single, double, or mixed precision, will necessarily contain some residual error with any finite grid. While this is expected with grid-based integration, estimates of the error and smoothness of the resulting integrals are necessary to test the accuracy of the procedure. First, a selection of 2-center ERIs were evaluated to determine the relative loss in precision when using the mixed-precision procedure. For each $(P|Q)$, the center Q was scanned radially away from center P in 16 directions corresponding to all-positive vectors of an 86-point Lebedev grid. Figure 5.4 plots the max and average absolute errors between mixed- and double-precision integrals at each distance. This indicates that the error of individual integrals are similarly sized across various distances, with the errors all being less than 8×10^{-5} .

The next measure of performance for mixed-precision integration is to evaluate the smoothness of the integrals with respect to changes in nuclear position. Therefore 2-center

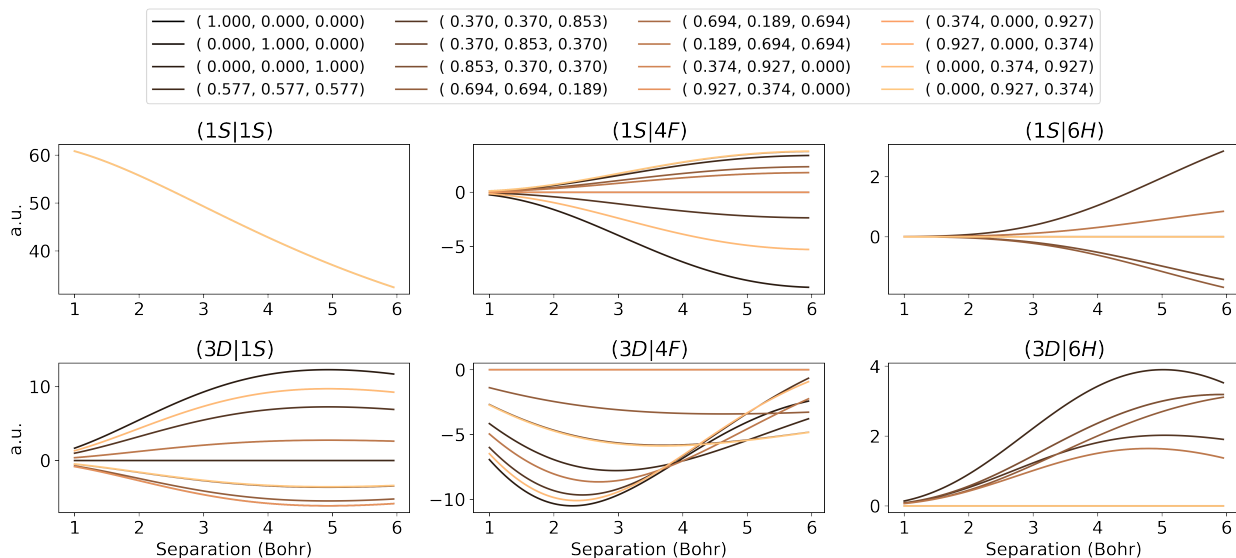


Figure 5.5: The value of 2-center ERIs are evaluated in mixed precision. All basis functions have $\zeta = 1$ and the right basis is scanned radially away from the origin in various directions. The directions selected are provided in the legend and were selected using the 16 all-positive directions of an 86-point Lebedev grid. For all integrals shown here, $m = 0$. The legend entries are direction unit vectors.

ERIs were evaluated in mixed precision as center Q is scanned radially for the same 16 directions as before. These yield qualitatively smooth plots, as seen in Figure 5.5. Additional plots for other basis set pairings are provided in Section D.6 of Appendix D and show the same qualitative behavior as this figure.

5.3.3 Hartree-Fock and HCI

Two levels of wave function-based electronic structure theory were selected to provide practical tests for the Slater GPU integrals. First, the Hartree-Fock energies for a set of benchmark molecules were computed and these are listed in Table 5.3 (see Table D.2 of Appendix D for timing information). The DZP basis set, corresponding auxiliary basis sets, and grid described in the Computational Details were used for these tests. Energies using the 6-31G* and RI-cc-pVTZ auxiliary GTO basis sets are also reported, to provide a baseline for comparison. The HF results for alkanes (C_nH_{2n+2}) show a slight increase in the mixed-precision error as the chain length increases. This is shown in Figure 5.6, which depicts the relative error of the HF energy when using mixed- and double-precision at various grid sizes. The roughly constant relative error as system size grows suggests that the mixed-precision error is size extensive. Combined with Figures 5.4 and 5.5, Figure

5.6 indicates that errors due to using mixed-precision integrals may largely result in error cancellation.

Table 5.3: HF energies computed for several small molecules are listed. The STO basis sets and grids are described in the Computational Details. The numbers in parenthesis in the header denotes the bits of precision used for integral evaluation.

Molecule	DZP (32)	DZP (64)	6-31G*
CH ₄	-40.199728	-40.199730	-40.194806
C ₂ H ₆	-79.232585	-79.232593	-79.227194
C ₃ H ₈	-118.269381	-118.269391	-118.261168
C ₄ H ₁₀	-157.305952	-157.305972	-157.294705
C ₅ H ₁₂	-196.342295	-196.342322	-196.328158
BH ₃	-26.395615	-26.395617	-26.390665
BF ₃	-323.166750	-323.166775	-323.142633
CF ₄	-435.667561	-435.667608	-435.642948
Cr(CO) ₆	-1714.832816	-1714.832901	-1714.469310




Figure 5.6: Heatmaps of the relative error of HF energies when using mixed- vs double-precision integral evaluation are shown for various alkanes using different angular and radial grid sizes.

The small error margins for STO integrals—as measured at the HF level of theory—suggest that thermochemical properties can be precisely evaluated. To test this hypothesis, an S_N2 reaction involving fluoride exchange in fluoromethane was evaluated (Figure 5.7). Since the HF level of theory is not expected to be quantitative, activation energies were computed not only with HF, but also with the heat-bath configuration interaction (HCI) method, with ε_1 set to 1.0 mHa and ε_2 set to 1.0 μ Ha. HCI provides a close approximation to full CI, and importantly, is tractable for the 20e⁻ in 60 orbital system of interest here. The activation energies of the exchange reaction using various grids are reported in Table 5.4. At the grids considered, the change in activation energy at the HF level is negligible between mixed and




Figure 5.7: The S_N2 reaction for fluoride exchange of fluoromethane.

double precision as well as between grid sizes. At the HCI level, more integrals contribute to the total energy. Consequently, the variation in the activation energy is larger for HCI relative to HF. However, the range of activation barriers for HCI is still less than half a kcal mol^{-1} .

Table 5.4: HF and HCI activation energies (kcal mol^{-1}) of CH_3F fluoride exchange at various grid sizes using single- and double-precision integral evaluations. The number of radial points and Lebedev order are provided for the radial and angular grids. The size of the angular grid is given in parenthesis next to the Lebedev order.

		Double Precision			Mixed Precision		
		HF					
Radial	Angular	17(590)	18(770)	19(974)	17(590)	18(770)	19(974)
	50	18.4	18.4	18.4	18.4	18.4	18.4
	60	18.4	18.4	18.4	18.4	18.4	18.4
		HCI					
	50	13.7	13.6	13.8	13.6	13.8	13.7
	60	13.9	13.8	13.7	13.8	13.6	13.6

Another test using the HCI method was the calculation of the singlet-triplet gaps of cyclobutadiene at its D_{2h} and D_{4h} geometries using a triple- ζ polarized (denoted TZP) basis set. Cyclobutadiene has a multireference singlet ground state, due to its degenerate π orbitals in the D_{4h} geometry. The results, for HCI parameters of $\varepsilon_1 = 1.0$ mHa and $\varepsilon_2 = 0.1$ μHa , are given in Table 5.5. These demonstrate mixed-precision errors of less than 0.1 kcal mol^{-1} . Furthermore, the singlet-triplet gap at the D_{4h} is consistent with prior full-CI using GTOs.⁷⁹

Table 5.5: Relative energies of cyclobutadiene at D_{2h} and D_{4h} geometries (kcal mol^{-1}).

	Double Precision		Mixed Precision	
	D_{2h}	D_{4h}	D_{2h}	D_{4h}
Singlet	0.0	9.4	0.0	9.3
Triplet	36.2	14.2	36.2	14.2
Gap	36.2	4.8	36.2	4.9

One final test will further show the utility of the Slater GPU integrals in quantum chemistry. Specifically, the geometric gradients—which are essential in studying chemical reactions—were evaluated using analytical nuclear derivatives of the quantities V_P , μ , and ν in Equation 5.13. As a benchmark, the fully symmetric BF_3 , BH_3 , CF_4 , and CH_4 molecules were

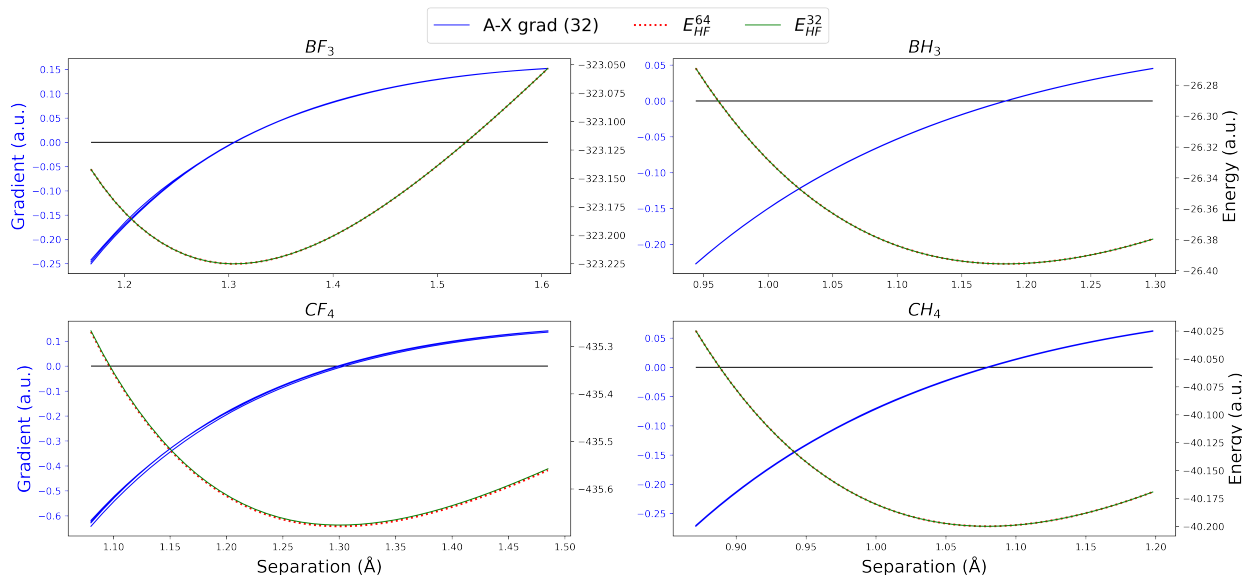


Figure 5.8: The Hartree-Fock geometric gradient projections (solid blue lines) of molecules with D_{3h} and T_d point groups are plotted as the A-X bond distance is scanned, where A=B,C and X=H,F. Gradients were computed in mixed precision. The mixed-precision (solid green) and double-precision (dotted red) Hartree-Fock energies at each point are also plotted. For CF_4 , the auxiliary basis for fluorine is extended with additional $2p$, $3d$, $4f$ and $5g$ functions.

symmetrically stretched. The HF energies as well as the projection of the mixed-precision HF geometric gradient onto each A-X bond (A=B,C;X=H,F) are plotted in Figure 5.8. As before, the mixed- and double-precision energies overlap with one another. As for the gradient, the magnitude along each A-X bond should be identical for all distances. This is largely achieved in these test cases, however, there is some variation when fluorine is present. For CF_4 , using the ADF fitting basis led to large gradient errors, thus the auxiliary basis of fluorine was extended with additional functions (see Section D.3 of Appendix D for additional details). Since this addition resulted in substantially improved gradients, the remaining variations for BF_3 and CF_4 are attributed to an incomplete RI auxiliary basis. While this work has not examined the choice of RI basis in detail, this subject will need to be revisited in a future study.

5.4 Conclusions

The SlaterGPU integral code is herein shown capable of evaluating the full complement of ERIs needed for HF and post-HF theories. Modern computer architectures combined with the RI approximation have allowed STO integrals to be feasible even though analytic expressions are currently unavailable. The use of mixed-precision integration allows further performance gains—achieving speedups greater than $80\times$ for the ERIs—with minimal loss to accuracy. In

the future, computing select integrals in double precision may mitigate errors due to using mixed-precision integrals. The combination of GPU acceleration, multi-GPU parallelization, and mixed-precision integration make SlaterGPU competitive with single-threaded GTO integration with the possibility of tuning SlaterGPU for additional performance.

The current implementation and basis sets are adequate for performing correlated electronic structure computations at the full CI level, however room for improvement remains in the STO RI gradients, where the available auxiliary basis sets appear to be inadequate. Further development of auxiliary basis sets will be required before STO integrals are generally useful for gradient computations.

CHAPTER 6

Final Remarks

6.1 Conclusions

In quantum chemistry, the goal is to compute accurate wave functions with as little computational resource as possible. These criteria are at odds and must be balanced. We played with this balance throughout this thesis.

In Chapter 2, we were able to computationally characterize an open-shell (and strongly correlated) coronoid system using the spin-flip restricted active space configuration interaction (RAS-SF) method. The coronoid possessed highly multiconfigurational character with multiple low-lying spin states. Additional analysis of the RAS-SF results using a spin Hamiltonian also determined that the system featured predominantly antiferromagnetic coupling between the radical site pairs. Due to the localization of electron correlation to primarily six radical sites, the RAS-SF method is highly suitable for the **8'** molecule. However, RAS-SF is still limited in active space size and the amount of electron correlation it can recover.

Chapter 3 introduced computational advances to the heat-bath configuration interaction (HCI) method, thus extending the size of active space that can be treated accurately. Conventional CI methods can handle around 18 electrons. With the advances in HCI, extrapolations to the full configuration interaction (FCI) solutions can be achieved for active spaces containing around of 20 to 30 electrons, which would be limit calculations to small complexes such as dimers of transition metals or small organic molecules containing up to 5 heavy atoms when the full set of valence electrons is required. In our case, we were able to explicitly correlate and extrapolate to FCI an active space of (22e,168o) by reducing the full Fock space containing $\sim 10^{33}$ Slater Determinants to only $\sim 10^7$ determinants. However, the calculation required over 4000 CPU cores on a super computer, and HCI is still an exponentially scaling method, thus remains beyond the reach of the typical quantum chemistry user.

We introduced improvements to the method of increments as applied to the complete active space self-consistent field (iCASSCF) method in Chapter 4 making iCASSCF variational.

The iCASSCF method approximates CASSCF, which performs FCI within the active space. The use of the incremental expansion produced a highly accurate, polynomially scaling approximation to CASSCF enabling treatment of large active spaces containing over 80 electrons, corresponding to medium-sized organic molecules with up to ~ 20 heavy atoms, or small transition metal complexes. Using iCASSCF, we were able to treat a (84e,84o) active space, which has a Fock space containing $\sim 10^{48}$ Slater determinants. The iCASSCF method took 8 days for a single-point calculation in the (84e,84o) active space on 12 CPU cores. This wall/CPU time pushes the limits of what might be considered reasonable for a standard quantum chemistry user.

The previously mentioned chapters focused on the ability of configuration interaction (CI) methods to recover electron correlation. However, Chapter 5 addresses the accuracy of the wave function at the basis set level. While Slater-type orbitals (STOs) are able to better represent electron density than Gaussian-type orbitals (GTOs), the difficulty of solving their electron integrals has historically hindered their adoption, especially in correlated wave function methods. With the development of the SlaterGPU package, the gap between GTO- and STO-integral evaluation is narrowed by utilizing the hardware capabilities of graphics processing units (GPUs).

6.2 Future Directions

There are several natural approaches one can take to enhance the capabilities of the methods developed throughout this thesis. We discuss these approaches below.

In HCI, a simple variation of the FNV hash is used which means perturbative determinants are hashed after they are generated. The development of a hash which is informed by the underlying structure of the CI problem could dramatically improve both the overall performance and the parallel efficiency of the algorithm.

For the method of increments, high accuracy has been established for complexes containing on the order of 80 to 120 electrons at the 4-body—significantly out of reach to other correlated methods including HCI. However, at a fixed expansion order, the errors of many-body expansion methods are known to be size extensive, i.e. they scale with system size. Thus, for active spaces with significantly more electrons—on the order of hundreds of electrons or more—one needs to first establish the appropriate truncation order. In addition, CASSCF is often used to study photochemical processes. For iCASSCF to be generally useful in photochemistry, the excited state theory still needs to be developed.

To achieve higher accuracy in STO integration (or any numerical integration generally), one simply increases the fineness of the grid. In SlaterGPU, a direct-product grid is used

meaning one can only increase the size of the angular and radial grids, which increases the number of points across the entire 3D grid. Another approach would be to utilize adaptive grids where the grid resolution is only increased for volume components that significantly affect the accuracy of integration.

Finally, quantum chemistry is most practically applied to provide insights to experimental data. Most molecular experiments are performed in solution phase, however the methods presented in this thesis are gas phase methods. Additionally, relativistic effects can play an important role in the energetics of many reactions, especially where transition metal catalysts are used. Thus, the incorporation of solvation and relativistic effects into the HCI and iCASSCF implementations are necessary to more accurately model experimental conditions.

6.3 Outlooks in Quantum Chemistry

Continued developments in quantum chemistry—both from computational and theoretical perspectives—will be needed to extend its reach to increasingly larger systems and to improve the accuracy of its methods. On the computational side, we must tailor our software to leverage the advances in computer hardware. For example, the majority of the compute capability enabled by the latest and upcoming exascale supercomputers come from accelerators as opposed to raw CPU horsepower. To some extent, this effort has been ongoing in the quantum chemistry community, including in this thesis with the introduction of SlaterGPU. These computational improvements will be necessary to address increasingly larger and more difficult problems.

On the theoretical side, the challenge remains to develop methods that can provide systematic approximations to FCI with favorable scaling characteristics that can run without highly specialized infrastructure. Methods like HCI scale exponentially, which will likely preclude their widespread adoption as standalone methods. However, with computational advances, they can act as the underlying engine in fragmented methods like iFCI or iCASSCF and can also provide valuable benchmark data for the development of other methods, such as density functional theory (DFT) or those based in machine learning. Finally, the further development of STO integration libraries and the STO basis sets themselves will be needed to truly solve many quantum chemistry problems. Artefacts from using an improper—i.e. Gaussian—basis, for example, has long misguided developers of density functionals.

In summary, future electronic structure theorists will need to be well-versed in computation and theory, and be familiar with contemporary computer hardware. By combining expertise in all three, one can design future methods where the computational cost required for more accuracy need not be so expensive.

APPENDIX A

Supporting Information for Chapter 2

The natural orbitals for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) optimized geometries are shown in the figures that follow.




Figure A.1: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (S_0) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.2: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.3: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_2) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.4: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (T_3) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.5: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Qn_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.6: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Q_{n_2}) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.7: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (Sp_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.




Figure A.8: Natural orbitals for the RAS(6,6)-SF/cc-pVDZ ground state (S_1) at CASSCF(6,6)/cc-pVDZ geometries optimized for the singlet ($^1\mathbf{8}'$), triplet ($^3\mathbf{8}'$), quintet ($^5\mathbf{8}'$), and septet ($^7\mathbf{8}'$) configurations.

APPENDIX B

Supporting Information for Chapter 3

B.1 Butadiene Geometry

Geometry for Scaling Studies

C	-3.11156756	0.23231420	0.0
H	-2.57840381	-0.69539072	0.0
H	-4.18156756	0.23231420	0.0
C	-2.43629325	1.40729149	0.0
H	-2.96945700	2.33499641	0.0
C	-0.89629325	1.40729149	0.0
H	-0.36312950	0.47958657	0.0
C	-0.21404985	2.57823598	0.0
H	-0.74169232	3.50909226	0.0
H	0.85593126	2.57187868	0.0

B.2 Butadiene Energies

The HCI energies for the butadiene accuracy benchmark are listed in Table B.1. The geometry for the accuracy benchmark is taken from Daday *et al.*¹⁴³ Figure B.1 shows a 5-point extrapolation which results in a total energy of -155.55720 Ha. Table B.2 compares 5-point and 4-point extrapolated energies from this work to prior studies. The extrapolated energies are in agreement with DMRG results¹⁵² within 3×10^{-4} Ha.

Table B.1: Energy benchmarks for butadiene in the ANO-L-VDZP basis set. The perturbative cutoff is $\varepsilon_2 = 10^{-8}$ Ha for HCI.

ε_1	Energy (var)	Energy (total)
2.5×10^{-4}	-155.46483	-155.54458
2.0×10^{-4}	-155.46921	-155.54515
1.5×10^{-4}	-155.47498	-155.54597
1.0×10^{-4}	-155.48335	-155.54710
5.0×10^{-5}	-155.50018	-155.54947
5-point extrapolation		-155.55739
4-point extrapolation		-155.55744

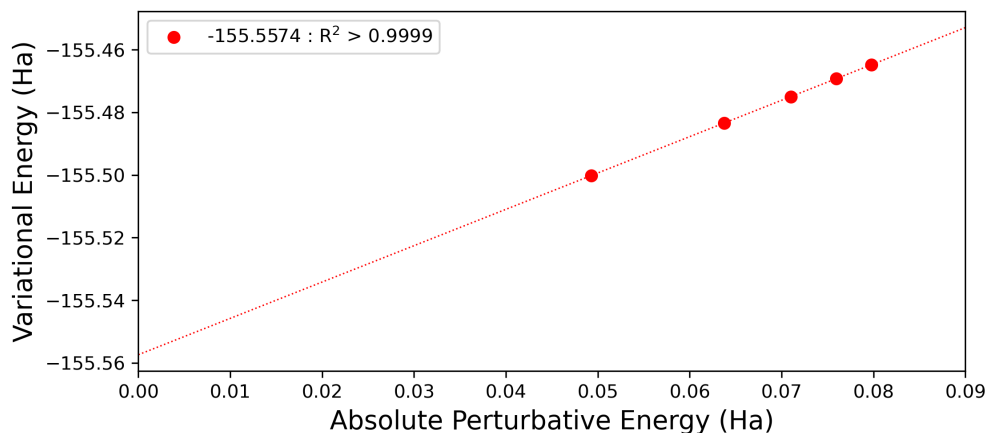


Figure B.1: 5-point HCI extrapolation to the FCI energy of butadiene in the ANO-L-VDZP basis. Energies were computed with ε_1 of 2.5×10^{-4} , 2.0×10^{-4} , 1.5×10^{-4} , 1.0×10^{-4} , and 5.0×10^{-5} with $\varepsilon_2 = 1.0 \times 10^{-8}$. The extrapolated energy and associated R^2 are listed in the legend.

Method	Energy (Ha)
i-FCIQMC ¹⁴³	-155.54914
DMRG ¹⁵²	-155.55718
SHCI (extrapolated) ⁴³	-155.55821
HCI (5-point extrapolation)	-155.55739
HCI (4-point extrapolation)	-155.55744

Table B.2: Butadiene energy benchmarks in the ANO-L-VDZP basis set.




Figure B.2: The localized occupied and singly-occupied ROHF orbitals for the triplet state of $[\text{FeO}(\text{NH}_3)_5]^{2+}$ included in the active space are depicted. Isosurface value of 0.1 a.u. is used for all figures.

B.3 $[\text{FeO}(\text{NH}_3)_5]^{2+}$ Orbitals

Figure B.2 shows the Pipek-Mezey localized occupied and singly-occupied which were included in the HCI active space. Similar orbitals for the quintet state were used. The localized orbitals were taken from ROHF references.

B.4 Hash Functions

A hash function is a function that takes an arbitrary input and maps it to an unsigned integer. The FNV hash¹³⁷ can be written in C++ as shown in Algorithm B.1. The FNV hash loops over `len` bytes performing XOR (\wedge) and integer multiplication. In the HCI implementation, the FNV hash is modified to Algorithm B.2. Rather than looping over every byte, the loop is over every 8 bytes. Since `size_t` is 8 bytes wide, a stride of 8 bytes still captures information from every bit in the XOR operation and also reduces operation count in the hash. Furthermore, Algorithm B.2 hashes a fixed size bit-string with `MAXORBS` bits, which is set to 192 in the current study. Using a fixed size bit-string allows Algorithm B.2 to utilize the "unroll-loops" optimization printed in line 1 of Algorithm B.2. Algorithm B.2 also has right bit-shift of 37 bits in order to capture more significant bits from the last multiplication operation in the loop. The bit-shift does not practically affect the number of

ability to distribute determinants especially since determinants are binned according to both α - and β -string hashes. The timings of these hash functions are discussed next.

Table B.3: Time to hash 5.6×10^8 α -strings of length 192 with 11 set bits.

Hash function	Total time (s)	Avg Time per hash (ns)
<code>std::hash</code>	4.46	7.87
Algorithm B.1	13.3	23.5
Algorithm B.2	1.88	3.32
Algorithm B.2 without loop unroll	2.08	3.67

Algorithm B.1 C++ code for FNV hash.

```

1 size_t Fnv_hash(const void* ptr, int len) {
2     size_t hash = 0xcbf29ce484222325;
3     const char* cptr = static_cast<const size_t*>(ptr);
4
5     for (;len;--len) {
6         hash ^= static_cast<size_t>(*cptr++);
7         hash *= static_cast<size_t>(0x000001000000001B3);
8     }
9     return hash;
10 }
```

Table B.3 lists the timings for `std::hash` from the C++ STL, Algorithm B.1, Algorithm B.2, and Algorithm B.2 without the "unroll-loops" optimization. The compiler used in these timing tests was GCC 12 with the `-O2` flag enabled, which does not automatically include loop unrolling in its optimizations. By using a stride of 8 bytes, Algorithm B.2 is about 7x faster than Algorithm B.1 and more than twice as fast as `std::hash`. It can also be seen that without loop unrolling, the time per hash increases by approximately 10%, which has significant impact on the performance of the perturbative step of the HCI implementation.

Furthermore, Algorithm B.2 distributes determinants in a consistent manner across the hash. Taking the HCI calculation of the $[\text{FeO}(\text{NH}_3)_5]^{2+}$ quintet state with $\varepsilon_1 = 1 \times 10^{-4}$ Ha as an example, there would be the 8.95×10^{10} perturbative determinants. These determinants were distributed across 144 MPI processes (4 MPI processes per node on 36 nodes). Each MPI process had an average of 6.2×10^8 determinants with standard deviation 8.6×10^6 determinants and a range of 3.7×10^7 determinants. This corresponds to a standard deviation and range of only 1.3% and 5.9% of the average, which leads to consistent memory requirements for each MPI process.

Algorithm B.2 C++ code for hash function used in HCI implementation

```
1 __attribute__((optimize("unroll-loops")))
2 short Hash_bytes(const void* ptr)
3 {
4     size_t hash = 0x01000193;
5     const size_t* cptr = static_cast<const size_t*>(ptr);
6     int len = MAXORBS / 64;
7
8     for (;len;--len)
9     {
10        hash ^= static_cast<size_t>(*cptr++);
11        hash *= static_cast<size_t>(2870177450012600261ULL);
12    }
13    hash = hash >> 37;
14    return hash;
15 }
```

APPENDIX C

Supporting Information for Chapter 4

C.1 oxoMn(salen)Cl

Dual socket Intel Xeon E5-2683 v3 CPUs were used for the iCASSCF computations of this work. Using 12 cores on these CPUs, the integral evaluation as well as generation and computation of the increments took 0.46 hours (1653 s) of total wall time for iCASSCF($n = 3$). Subsequent orbital optimization at $n = 3$ required a total wall time of approximately 8 days 3 hours (701571 s). For $n = 4$, the resulting orbitals at $n = 3$ were used and the subsequent orbital optimization required 20.1 hours (72310 s).

C.2 Parallelity

The error between conventional CASSCF and the method of increments is computed for the PES of ethane dissociation (see Figure C.1). The conventional CASSCF computed bond dissociation energy is 0.1653 Ha (103.7 kcal mol⁻¹). The non-parallelity error (NPE) for iCASSCF($n = 4$), iCAS-CI($n = 4$), and the previous implementation of iCASSCF($n = 4$) are 2.15 mHa, 0.23 mHa, and 0.14 mHa, respectively. Single excitations are required for precisely describing bond dissociation, hence the (relatively) large error in in the singles-free iCASSCF implementation. However, as with other examples, the missing correlation is recovered when using iCAS-CI.

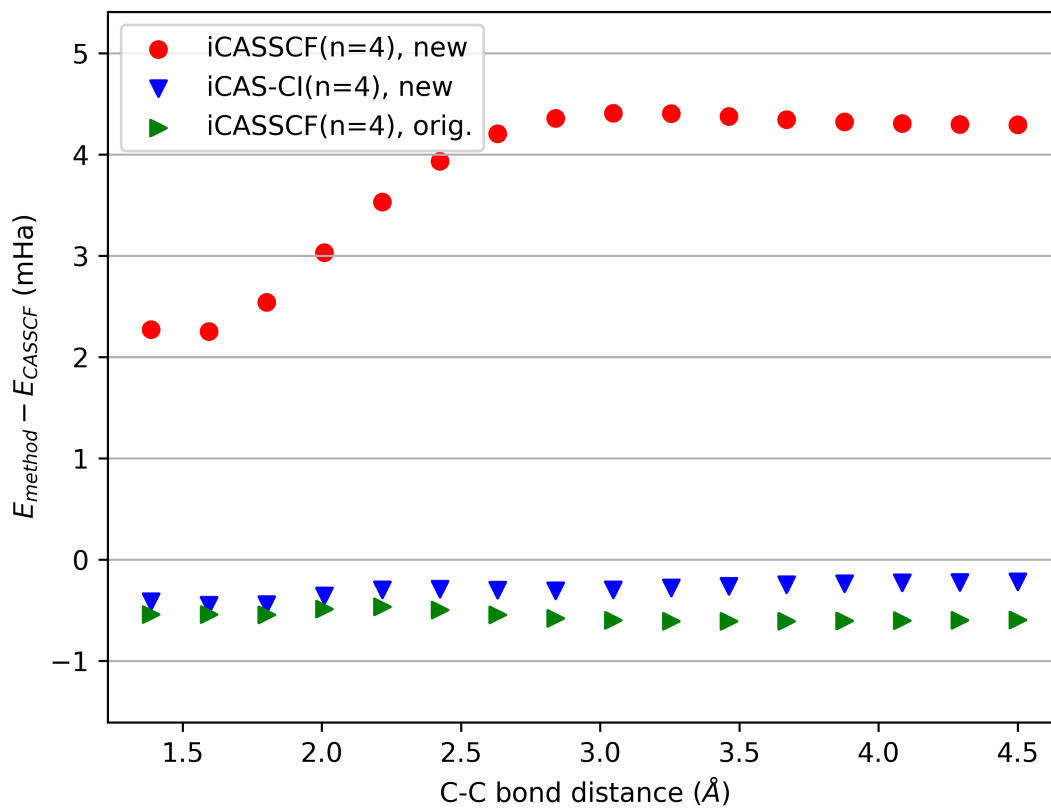


Figure C.1: Differences between canonical CASSCF, iCASSCF($n = 4$), iCAS-CI($n = 4$) and the original implementation of iCASSCF($n = 4$) are shown for ethane dissociation. The NPE for iCASSCF, iCAS-CI, and the original iCASSCF are 2.15 mHa, 0.23 mHa, and 0.14 mHa, respectively.

C.3 Additional Data



Figure C.2: Small molecule benchmark results including results from previous studies.

Table C.1: Small molecule benchmark geometries and energies from iCASSCF and iCAS-CI using 6-31G* basis set. Conventional CASSCF from Molpro. iCAS-CI($n = k$) energies use iCASSCF($n = k$) geometries and orbitals. Bond distances, angles, and energies given in Å, degrees, and Hartrees, respectively.

	Bond Distance		Bond Angle		Energy
	C=C	C-H	H-C-H	H-C-C	
<i>Ethylene</i>					
iCASSCF (n=3)	1.3525	1.0977	116.6	121.7	-78.1747
iCASSCF (n=4)	1.3524	1.0974	116.6	121.7	-78.1742
iCAS-CI (n=3)					-78.1780
iCAS-CI (n=4)					-78.1774
CASSCF (12,12)	1.3522	1.0995	116.5	121.7	-78.1777
	Bond Distance		Bond Angle		Energy
	C=O	C-H	H-C-H	H-C-O	
<i>Formaldehyde</i>					
iCASSCF (n=3)	1.2156	1.1144	115.6	122.2	-114.0453
iCASSCF (n=4)	1.2157	1.1128	115.5	122.3	-114.0451
iCAS-CI (n=3)					-114.0506
iCAS-CI (n=4)					-114.0504
CASSCF (12,12)	1.2152	1.1181	115.3	122.4	-114.0509
	Bond Distance		Energy		
	C-N	C-H			
<i>Hydrogen Cyanide</i>					
iCASSCF (n=3)	1.1689	1.0796	-93.0420		
iCASSCF (n=4)	1.1684	1.0783	-93.0412		
iCAS-CI (n=3)			-93.0452		
iCAS-CI (n=4)			-93.0439		
CASSCF (10,10)	1.1697	1.0795	-93.0443		
	Bond Distance		Dihedral Angle	Energy	
	O-O	O-H	H-O-O-H		
<i>Hydrogen Peroxide</i>					
iCASSCF (n=3)	1.4630	0.9733	118.3	-150.9890	
iCASSCF (n=4)	1.4623	0.9733	118.3	-150.9895	
iCAS-CI (n=3)				-150.9944	
iCAS-CI (n=4)				-150.9951	
CASSCF (14,14)	1.4632	0.9748	117.8	-150.9953	
	Bond Distance		Bond Angle	Energy	
	O-H	C-O	C-O-H		
<i>Methanol</i>					
iCASSCF (n=3)	0.9683	1.4315	107.6	-115.2211	
iCASSCF (n=4)	0.9683	1.4315	107.8	-115.2221	
iCAS-CI (n=3)				-115.2258	
iCAS-CI (n=4)				-115.2259	
CASSCF (14,14)	0.9696	1.4293	107.6	-115.2261	




Figure C.3: Select iCAS-CI($n=3$) natural orbitals for the singlet state of oxo-Mn(salen)Cl are shown. HONO and LUNO are boxed.




Figure C.4: Select iCAS-CI($n=3$) natural orbitals for the triplet state of oxo-Mn(salen)Cl are shown. SONOs are boxed.

Table C.2: Coordinates for the optimized ω B97XD/6-31G* geometry of oxo-Mn(salen)Cl.

1	Mn	0.00426853	0.01193109	0.00516565
2	O	-1.07904774	1.02407160	-1.29257396
3	O	-0.38568147	-1.57420571	-1.08201933
4	N	0.73795289	1.58032839	0.78132122
5	N	1.31212351	-0.88100686	1.04879369
6	C	-1.26610124	2.26797128	-1.24074706
7	C	-0.65407825	3.18037381	-0.37324817
8	C	0.34200739	2.81019719	0.52148226
9	C	0.22132065	-2.67532723	-0.98288661
10	C	1.23549686	-3.01016466	-0.08096144
11	C	1.71367906	-2.12181813	0.87688968
12	C	1.96257143	1.35466613	1.54148960
13	C	1.85745984	-0.04961373	2.11342121
14	H	0.86149270	3.60566714	1.05998860
15	H	2.49343703	-2.47401977	1.55504499
16	H	2.81286294	1.41062034	0.85157883
17	H	1.15975852	-0.05809563	2.96170481
18	H	2.08238684	2.10907682	2.32756115
19	H	2.82951715	-0.42140951	2.45358312
20	H	-1.97877721	2.65753526	-1.97620543
21	H	-0.09971460	-3.44024001	-1.69785436
22	Cl	1.67965883	0.29802390	-1.66949117
23	O	-1.23699337	-0.22205126	0.84999273
24	H	-0.91766654	4.22747250	-0.45342641
25	H	1.65965399	-4.00483737	-0.12510851

Table C.3: Coordinates for the optimized iCASSCF(n=3)/6-31+G** geometry of the cyclobutadiene reactant.

1	C	1.829129	0.415737	-0.396948
2	C	3.415272	0.396100	-0.355189
3	C	3.427557	1.747942	-0.164081
4	C	1.839964	1.767412	-0.205660
5	H	1.056274	-0.347902	-0.525581
6	H	4.186968	-0.374687	-0.442217
7	H	4.201979	2.509792	-0.034599
8	H	1.070108	2.539996	-0.118645

Table C.4: Coordinates for the iCASSCF(n=3)/6-31+G** geometry of the cyclobutadiene TS found from GSM

1	C	1.941501	0.347782	-0.374239
2	C	3.408507	0.336234	-0.349552
3	C	3.416179	1.788872	-0.143371
4	C	1.949197	1.799616	-0.164995
5	H	1.164623	-0.412608	-0.489600
6	H	4.174132	-0.440912	-0.416305
7	H	4.191111	2.542146	0.019179
8	H	1.182472	2.574430	-0.084291

Table C.5: Coordinates for the D_{4h} -symmetrized cyclobutadiene TS.

1	C	1.941501	0.347782	-0.374239
2	C	3.416199	1.789067	-0.143343
3	C	3.408502	0.336892	-0.352636
4	C	1.949199	1.799957	-0.164946
5	H	1.164683	-0.411435	-0.495867
6	H	4.177209	-0.433798	-0.451504
7	H	4.193017	2.548284	-0.021715
8	H	1.180491	2.570647	-0.066078

Table C.6: Conventional CASSCF coordinates for optimized methanol geometry.

1	C	-0.0327780667	0.6630756802	0.0000000020
2	O	-0.0711064172	-0.7657519727	0.0000000079
3	H	-1.0810849290	1.0086056226	-0.0000000042
4	H	0.4716079663	1.0726542296	0.9016920551
5	H	0.4716079686	1.0726542246	-0.9016920517
6	H	0.8448534780	-1.0838377843	-0.0000000091

Table C.7: Conventional CASSCF coordinates for optimized hydrogen peroxide geometry.

1	O	0.0227950330	0.7312580176	-0.0636550734
2	O	-0.0227950330	-0.7312580176	-0.0636550734
3	H	0.8503508636	0.8685890402	0.4328550734
4	H	-0.8503508636	-0.8685890402	0.4328550734

APPENDIX D

Supporting Information for Chapter 5

D.1 One-electron integrals

The one electron integrals are of the form

$$O_{\mu\nu} = \langle \chi_\mu | \hat{O}_1 | \chi_\nu \rangle = \int \chi_\mu(\mathbf{r}) \hat{O}(\mathbf{r}) \chi_\nu(\mathbf{r}) d\mathbf{r}, \quad (\text{D.1})$$

where operators of interest are the overlap ($\hat{O} = 1$), kinetic energy ($\hat{O} = \frac{1}{2}\nabla^2$), and nuclear attraction ($\hat{O} = \frac{Z_A}{R_{1A}}$) operators. For a set of grid points x and grid weights $w(x)$, the integral in Equation D.1 can be evaluated by

$$O_{\mu\nu} = N_{\chi_\mu} N_{\chi_\nu} \sum_x \bar{\chi}_\mu(x) \hat{O}(\bar{\chi}_\nu(x)) w(x), \quad (\text{D.2})$$

where $\bar{\chi}_\lambda$ is the basis function with normalization constant N_{χ_λ} factored out.

D.2 Radial component of the Slater Coulomb Potential

The radial component of the Slater Coulomb potential has the form⁴

$$I_{nl}(r) = r^{-l-1} \int_0^r (r')^{n+l+1} e^{-\zeta r'} dr' + r^l \int_r^\infty (r')^{n-l} e^{-\zeta r'} dr', \quad (\text{D.3})$$

which has analytic forms for each n, l of interest in quantum chemistry. The expressions for I_{nl} up to $n = 3$ are listed below. As n, l increase, the expression demonstrates numerical instability resulting in non-smooth STO integration, as discussed in Chapter 5.

$$\begin{aligned}
I_{1,0}(r) &= \frac{(-\zeta r + 2e^{\zeta r} - 2) e^{-\zeta r}}{\zeta^3 r} \\
I_{2,0}(r) &= \frac{(-\zeta^2 r^2 - 4\zeta r + 6e^{\zeta r} - 6) e^{-\zeta r}}{\zeta^4 r} \\
I_{2,1}(r) &= \frac{3(-\zeta^3 r^3 - 4\zeta^2 r^2 - 8\zeta r + 8e^{\zeta r} - 8) e^{-\zeta r}}{\zeta^5 r^2} \\
I_{3,0}(r) &= \frac{(-\zeta^3 r^3 - 6\zeta^2 r^2 - 18\zeta r + 24e^{\zeta r} - 24) e^{-\zeta r}}{\zeta^5 r} \\
I_{3,1}(r) &= \frac{3(-\zeta^4 r^4 - 6\zeta^3 r^3 - 20\zeta^2 r^2 - 40\zeta r + 40e^{\zeta r} - 40) e^{-\zeta r}}{\zeta^6 r^2} \\
I_{3,2}(r) &= \frac{5(-\zeta^5 r^5 - 6\zeta^4 r^4 - 24\zeta^3 r^3 - 72\zeta^2 r^2 - 144\zeta r + 144e^{\zeta r} - 144) e^{-\zeta r}}{\zeta^7 r^3}
\end{aligned}$$

D.3 Basis set specifications

The STO subshell and exponent are listed for the main and auxiliary basis for each atom considered in the main paper. Basis sets are modified versions of those provided in the ADF package²⁵⁹.

D.3.1 DZP basis set specifications

H basis specification

Main basis		Auxiliary basis					
1s	0.76	1s	3.16	2s	1.50	2p	1.75
1s	1.28	1s	2.09	2p	4.00	3d	4.00
2p	1.25	1s	1.38	2p	2.65	3d	2.50
4f	3.00	5g	4.00				

B basis specification

Main basis		Auxiliary basis					
1s	6.50	1s	13.00	3s	2.56	3p	1.68
1s	4.08	2s	14.79	3s	1.87	3d	6.08
2s	1.00	2s	10.16	3s	1.36	3d	3.69
2s	1.56	2s	6.98	2p	9.10	3d	2.24
2p	1.70	2s	4.80	2p	5.17	3d	1.36
2p	0.76	3s	4.83	3p	4.36	4f	5.00

3d	1.50	3s	3.52	3p	2.71	4f	3.50
		5g	4.50	5g	5.50		

C basis specification

Main basis		Auxiliary basis					
1s	7.68	1s	15.36	3s	3.08	3p	2.06
1s	5.00	2s	17.53	3s	2.25	3d	7.20
2s	1.98	2s	12.07	3s	1.64	3d	4.40
2s	1.24	2s	8.31	2p	9.88	3d	2.69
2p	2.20	2s	5.73	2p	5.80	3d	1.64
2p	0.96	3s	5.78	3p	5.05	4f	5.40
3d	2.20	3s	4.22	3p	3.23	4f	3.55
		5g	4.50				

O basis specification

Main basis		Auxiliary basis					
1s	9.80	1s	19.60	3s	4.15	3p	2.84
1s	6.80	2s	22.60	3s	3.05	3d	8.80
2s	1.70	2s	15.68	3s	2.24	3d	5.58
2s	2.82	2s	10.89	2p	12.86	3d	3.53
2p	3.06	2s	7.56	2p	7.68	3d	2.24
2p	1.30	3s	7.69	3p	6.78	4f	6.20
3d	2.00	3s	5.65	3p	4.39	4f	3.70
		5g	4.50				

F basis specification

Main basis		Auxiliary basis					
1s	10.88	1s	21.76	3s	3.95	3p	3.42
1s	7.70	2s	24.39	3s	2.85	3d	9.70
2s	3.22	2s	16.56	3s	2.05	3d	6.16
2s	1.92	2s	11.24	2p	1.48	3d	3.91
2p	3.52	2s	7.63	2p	14.40	3d	2.48
2p	1.48	3s	7.60	3p	7.53	4f	6.50
3d	2.00	3s	5.48	3p	5.90	4f	3.75
		5g	4.50				

Additional Functions in extended Auxiliary Basis

2P	4.76	3D	0.62
2P	2.64	3D	18.71
2P	1.26	4F	2.97
2P	0.65	4F	1.44
3D	18.71	5G	3.12
3D	1.52		

Cr basis specification

Main basis		Auxiliary basis					
1S	27.25	1S	54.50	7S	2.68	6D	3.62
1S	21.70	2S	60.57	7S	2.13	7D	2.60
2S	9.20	2S	40.83	7S	1.70	4F	15.60
2S	6.05	3S	40.40	2P	40.45	5F	8.92
2P	13.20	3S	28.95	3P	27.34	5F	4.30
2P	8.25	4S	27.44	4P	18.68	6F	2.54
3S	5.25	4S	20.46	5P	12.96	5G	9.75
3S	3.30	4S	15.25	5P	7.53	5G	5.17
3P	4.65	5S	14.16	6P	5.30	5G	2.74
3P	2.80	5S	10.86	6P	3.22		
3D	5.70	5S	8.33	7P	2.30		
3D	2.70	6S	7.65	3D	29.95		
3D	1.24	6S	6.00	4D	20.59		
4S	1.75	6S	4.70	5D	14.37		
4S	1.00	6S	3.68	5D	8.39		
4P	1.30	7S	3.36	6D	5.94		

D.3.2 TZ2P basis set specifications

Li basis specification

Main basis		Auxiliary basis					
1s	4.24	2p	0.60	1s	8.48	3p	0.92
1s	2.26	2p	1.20	1s	5.43	3d	4.72
2s	2.36	3d	1.20	2p	6.60	3d	2.08
2s	0.68	4f	1.80	2p	2.98	3d	0.92

2s	0.46			2p	1.34		4f	5.00
				4f	3.50		5g	3.50

F basis specification

Main basis				Auxiliary basis			
1s	10.88	2p	4.54	1s	21.76	3s	2.85
1s	7.70	2p	2.30	2s	24.39	3s	2.05
2s	3.24	2p	1.24	2s	16.56	3s	1.48
2s	1.94	3d	2.00	2s	11.24	2p	14.40
2s	0.74	4f	3.00	2s	7.63	2p	7.53
				3s	7.60	3p	5.90
				3s	5.48	3p	3.42
				3s	3.95	3p	1.98
				3d	9.70	3d	6.15
				3d	3.91	3d	2.48
				4f	6.50	4f	3.75
				5g	4.50		

D.3.3 TZP basis set specifications

For the following TZP basis sets, the same auxiliary basis was used as specified in the DZP basis.

C basis specification

1S	7.68
1S	5.00
2S	1.28
2S	2.10
2S	4.60
2P	0.82
2P	1.48
2P	2.94
3D	2.20

H basis specification

1S	0.69
1S	0.92
1S	1.58
2P	1.25

D.4 Geometries

CH₄

C	-0.00000	0.00000	0.00000
H	0.35665	-1.00881	0.00000
H	0.35667	0.50440	0.87365
H	0.35667	0.50440	-0.87365
H	-1.07000	0.00001	0.00000

C₂H₆

C	-0.25667	-0.36298	-0.62870
H	0.09998	-1.37179	-0.62870
H	0.10000	0.14142	-1.50235
H	-1.32667	-0.36296	-0.62870
C	0.25667	0.36298	0.62870
H	-0.09839	1.37235	0.62772
H	-0.10160	-0.14029	1.50235
H	1.32667	0.36127	0.62968

C₃H₈

C	-0.16408	-0.97376	-0.82820
H	0.19257	-1.98257	-0.82820
H	0.19259	-0.46936	-1.70185
H	-1.23408	-0.97375	-0.82820
C	0.34926	-0.24780	0.42920
H	-0.00901	-0.75107	1.30285
H	1.41926	-0.24951	0.43018
C	-0.16175	1.20494	0.42779
H	-1.23175	1.20665	0.42644

H	0.19461	1.70923	1.30163
H	0.19682	1.70831	-0.44567

BH₃: B-H bond length 1.18000 Å

BF₃: B-F bond length 1.46000 Å

CF₄: C-F bond length 1.35000 Å


C₄H₁₀

C	-0.25772	-1.45229	-1.25677
H	0.09893	-2.46110	-1.25677
H	0.09895	-0.94789	-2.13042
H	-1.32772	-1.45227	-1.25677
C	0.25562	-0.72633	0.00063
H	-0.10265	-1.22960	0.87428
H	1.32562	-0.72804	0.00161
C	-0.25539	0.72641	-0.00077
H	-1.32539	0.72812	-0.00213
H	0.10318	1.22979	-0.87423
C	0.25751	1.45221	1.25690
H	1.32750	1.45043	1.25830
H	-0.09748	2.46161	1.25588
H	-0.10114	0.94888	2.13036

C₅H₁₂

C	-0.20131	-2.03735	-1.50082
H	0.15534	-3.04616	-1.50082
H	0.15536	-1.53295	-2.37447
H	-1.27131	-2.03734	-1.50082
C	0.31203	-1.31139	-0.24341
H	-0.04624	-1.81466	0.63024
H	1.38203	-1.31310	-0.24243
C	-0.19899	0.14135	-0.24482
H	-1.26898	0.14305	-0.24617
H	0.15959	0.64472	-1.11828
C	0.31391	0.86715	1.01285
H	1.38391	0.86537	1.01425

H	-0.04473	0.36382	1.88632
C	-0.19700	2.31993	1.01139
H	-1.26700	2.32170	1.00998
H	0.15935	2.82421	1.88523
H	0.16165	2.82326	0.13793



Cr	-0.44964	-0.05396	0.00000
C	-0.44964	1.84604	0.00000
C	1.45036	-0.05396	0.00000
C	-0.44964	-0.05396	1.90000
C	-0.44964	-0.05396	-1.90000
C	-2.34964	-0.05396	0.00000
O	2.56576	-0.05396	0.00000
O	-0.44964	2.96144	0.00000
O	-0.44964	-0.05396	-3.01540
O	-0.44964	-0.05396	3.01540
O	-3.46504	-0.05396	0.00000
C	-0.44964	-1.95396	0.00000
O	-0.44964	-3.06936	0.00000



C	0.33338	-0.23574	0.00000
F	1.48429	-1.04955	0.00000
H	0.34198	0.37985	0.87916
H	0.34198	0.37985	-0.87916
H	-0.53718	-0.86347	0.00000
F	-1.70265	1.20396	0.00000



C	-0.00151	-0.00151	-0.00151
F	1.12836	1.12836	1.12836
H	-0.44307	-0.44307	0.87143
H	-0.44307	0.87143	-0.44307

H	0.87143	-0.44307	-0.44307
F	-1.12662	-1.12662	-1.12662

Cyclobutadiene D_{2h}

C	2.02818	-0.10367	-0.56968
C	3.38504	-0.11014	-0.53143
C	3.38648	1.46992	-0.30910
C	2.02955	1.47649	-0.34724
H	1.25399	-0.85949	-0.69767
H	4.15799	-0.87319	-0.61720
H	4.16066	2.22574	-0.18090
H	1.25658	2.23961	-0.26231

Cyclobutadiene D_{4h}

C	1.94150	0.34778	-0.37424
C	3.41620	1.78907	-0.14334
C	3.40850	0.33689	-0.35264
C	1.94920	1.79996	-0.16495
H	1.16468	-0.41143	-0.49587
H	4.17721	-0.43380	-0.45150
H	4.19302	2.54828	-0.02171
H	1.18049	2.57065	-0.06608

D.5 Performance and Kernel Analysis

The performance behaviors of the 2080-Ti and GV100 can be explained by comparing the underlying hardware. Specifically, the GV100 has $\sim 40\%$ higher theoretical memory bandwidth than the 2080-Ti and $\sim 17\%$ more streaming multiprocessors (SM) than the 2080-Ti, however, the 2080-Ti benefited from 27% higher clock speed throughout the timing tests. The greater performance of the 2080-Ti in mixed precision integral evaluation therefore suggests that the GPU kernels are compute bound. The performance inversion observed in full double precision integration is primarily a result of the substantially fewer double precision execution units in the 2080-Ti compared to the GV100.

To further verify whether that the SlaterGPU kernels are compute bound, the integrals for C_5H_{12} were computed using a Nvidia A4000 GPU at various locked memory and SM clock

speeds. Figure D.1 shows the relative performance loss by lowering the memory and SM clock speeds separately. Full timing details are provided in Table D.1. In mixed and double precision, lowering the SM clock speed by 14% (1560 MHz to 1335 MHz) and 30% (1560 MHz to 1095 MHz) leads to ERI performance loss of $\sim 12\%$ and $\sim 25\%$, respectively. On the other hand, lowering the memory clock by 24% only leads to $\sim 5\%$ and $< 1\%$ performance loss in mixed and double precision, respectively. Thus, the current implementation of the ERI GPU kernels are indeed compute bound.




Figure D.1: Relative performance of Slater integration is plotted as memory at various memory and SM clock speeds. Integrals were computed for the C_5H_{12} molecule. The x -axis labels list the memory clock followed by the SM clock in MHz.

Table D.1: Timings (in seconds) for Slater integration at various memory and SM clock speeds. Integrals were computed for the C_5H_{12} molecule.

Precision	Memory clock (MHz)	SM Clock (MHz)	$S + T$	V_{ne}	(P Q)	($\mu\nu$ Q)	Total
mixed	6500	1560	0.467	16.08	2.358	63.97	82.88
mixed	6500	1335	0.4986	17.38	2.646	72.76	93.28
mixed	6500	1095	0.5507	19.46	3.095	86.06	109.2
mixed	5000	1560	0.5113	17.38	2.446	67.15	87.49
double	6500	1560	0.2852	8.92	6.474	167.5	183.2
double	6500	1335	0.3016	9.541	7.433	190.7	208.0
double	6500	1095	0.326	10.5	8.889	225.3	245.0
double	5000	1560	0.3119	9.542	6.622	167.9	184.3

Hartree-Fock (HF) timings are also provided in Table D.2. Compared to the GPU timings for the integrals, the HF timings take less than 10% of the overall compute time across all system sizes considered in mixed and double precision. These timings do not represent an optimized HF implementation.

Table D.2: Hartree-Fock timings (in seconds) for various alkanes. Each atom contributes 46,200 grid points. Timings are recorded on a 2080-Ti GPU. Total time in HF does not include integral computation.

Mixed Precision					
Operation/Quantity	CH ₄	C ₂ H ₆	C ₃ H ₈	C ₄ H ₁₀	C ₅ H ₁₂
Main basis size	35	60	85	110	135
Aux basis size	224	370	516	662	808
Time for S and T integrals (s)	4.45E-02	6.40E-02	9.75E-02	1.40E-01	1.95E-01
Time for V_{N_e} integrals (s)	2.03E-01	7.37E-01	1.94E+00	4.06E+00	7.32E+00
Time for 3-center integrals (s)	4.22E-01	2.05E+00	5.79E+00	1.25E+01	2.34E+01
Avg time for G Tensor formation (s)	1.51E-03	5.94E-03	1.62E-02	4.30E-02	8.01E-02
Avg time for AO Fock formation (s)	1.30E-05	1.48E-05	1.95E-05	3.95E-05	3.80E-05
Avg time for DIIS (s)	4.69E-04	6.22E-04	9.34E-04	1.10E-03	1.38E-03
Avg time for Fock diagonalization (s)	1.77E-04	2.57E-04	3.78E-04	4.74E-04	6.34E-04
Number of HF iterations	11	12	15	16	17
Total integral time (s)	7.67E-01	3.05E+00	8.20E+00	1.73E+01	3.18E+01
Total time in HF (s)	1.01E-01	2.72E-01	6.61E-01	1.43E+00	2.58E+00
Double Precision					
Operation/Quantity	CH ₄	C ₂ H ₆	C ₃ H ₈	C ₄ H ₁₀	C ₅ H ₁₂
Main basis size	35	60	85	110	135
Aux basis size	224	370	516	662	808
Time for S and T integrals (s)	4.98E-02	8.05E-02	1.12E-01	1.63E-01	2.27E-01
Time for V_{N_e} integrals (s)	2.35E-01	8.82E-01	2.33E+00	4.86E+00	8.70E+00
Time for 3-center integrals (s)	1.98E+00	9.56E+00	2.69E+01	5.79E+01	1.07E+02
Avg time for G Tensor formation (s)	1.50E-03	5.93E-03	1.57E-02	3.39E-02	7.91E-02
Avg time for AO Fock formation (s)	1.32E-05	1.49E-05	1.95E-05	2.25E-05	3.60E-05
Avg time for DIIS (s)	4.71E-04	6.03E-04	9.03E-04	1.07E-03	1.35E-03
Avg time for Fock diagonalization (s)	1.76E-04	2.54E-04	3.78E-04	4.79E-04	6.35E-04
Number of HF iterations	11	12	15	16	17
Total integral time (s)	2.66E+00	1.15E+01	3.11E+01	6.58E+01	1.20E+02
Total time in HF (s)	9.88E-02	2.70E-01	6.32E-01	1.24E+00	2.67E+00

D.6 Mixed-Precision Error

Additional plots due to mixed-precision errors are provided in Figures D.3 and D.4. Figures D.5 and D.6 plot the values of the 2-electron integrals as the internuclear distance between centers is scanned along various coordinate directions. Figures D.3 and D.5 use basis functions where $m = 0$, while in Figures D.4 and D.6, $m = 0$ for the left basis and $m = n - 1$ for the right basis. In each of these figures, ζ is set to 1 for the left and right basis functions.

D.7 Lithium Fluoride Parallely Test

Another test to check the mixed-precision error was performed with the HBCI method ($\varepsilon_1 = 0.5$ mHa, $\varepsilon_2 = 0.5$ μ Ha) on the lithium fluoride molecule. A triple- ζ doubly-polarized basis, denoted TZ2P, which contains a total of 48 basis functions and 128 auxiliary basis functions was used with no frozen core orbitals. The LiF bond was stretched from 1.4 Å to 3.0 Å. A plot of the LiF energies and non-parallelity error shown in Figure D.2 demonstrates a smooth potential energy curve with negligible errors from using mixed-precision integrals on the order of μ Ha.

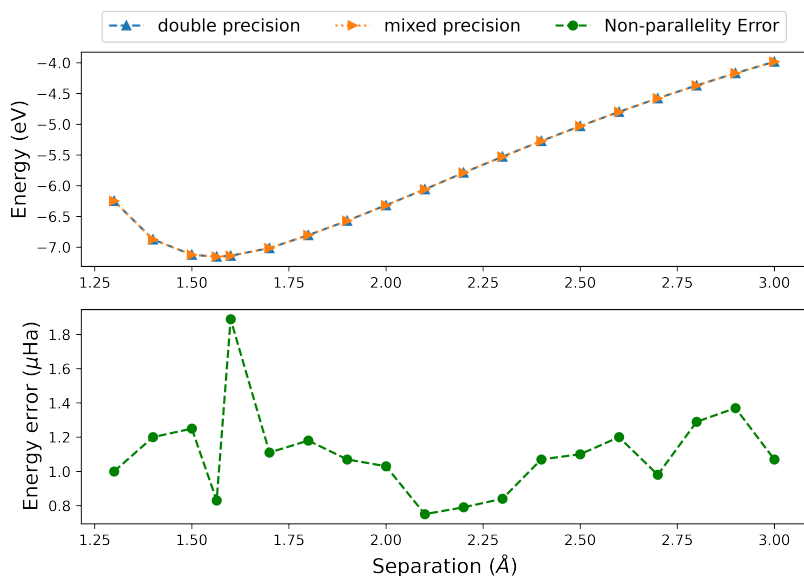


Figure D.2: The LiF molecule is stretched from 1.4 Å to 3.0 Å. Energies (top) were computed using HBCI($\varepsilon_1 = 0.5$ mHa, $\varepsilon_2 = 0.1$ μ Ha) with the TZ2P basis. An additional point at the experimental bond distance²⁶³ (1.564 Å) was included. Energies are relative to complete dissociation of the LiH bond. The non-parallelity error (bottom) from using mixed-precision integrals is also plotted.

D.8 OpenACC Code Examples

Algorithm D.1 Loop for computing the basis χ or potential V on a grid x with $npts$. The function f refers to either χ , or V , and fx refers to f evaluated on the grid x .

```
1 #pragma acc parallel loop present (fx[0:npts],x[0:npts])
2 for (int i = 0; i < npts; i++)
3     fx[i] = f(x[i]);
```

Algorithm D.2 Sample code for contracting V_P , μ , and ν into the 3-center integrals. Arrays corresponding to V_P , μ , and ν are assumed to have already been computed, and properly weighted.

```
1 int N2a = N*N*Naux;
2 #pragma acc parallel loop collapse(3) present(\
3   V_p[0:nbas_atm1][0:npts], mu[0:nbas_atm2][0:npts], nu[0:nbas_atm3][0:
   npts], Vmunu[0:N2a])
4 for (int i = atom1_start; i < atom1_end; i++) { // loop over basis
   functions on atom 1
5   for (int j = atom2_start; j < atom2_end; j++) { // loop over basis
     functions on atom 2
6     for (int k = atom3_start; k < atom3_end; k++) { // loop over basis
       functions on atom 3
7       // shift i,j,k to properly index V_p, mu, nu from 0 to atom 1,2,3
8       int i_shift = i - atom1_start;
9       int j_shift = j - atom2_start;
10      int k_shift = k - atom3_start;
11
12      double sum = 0.;
13      #pragma acc loop reduction(+:sum)
14      for (int l = 0; l < npts; l++) {
15        sum += v_p[i_shift][l] * mu[j_shift][l] * nu[k_shift][l];
16      }
17      Vmunu[i*N*N + j*N + k] = sum;
18    }
19  }
20 }
```

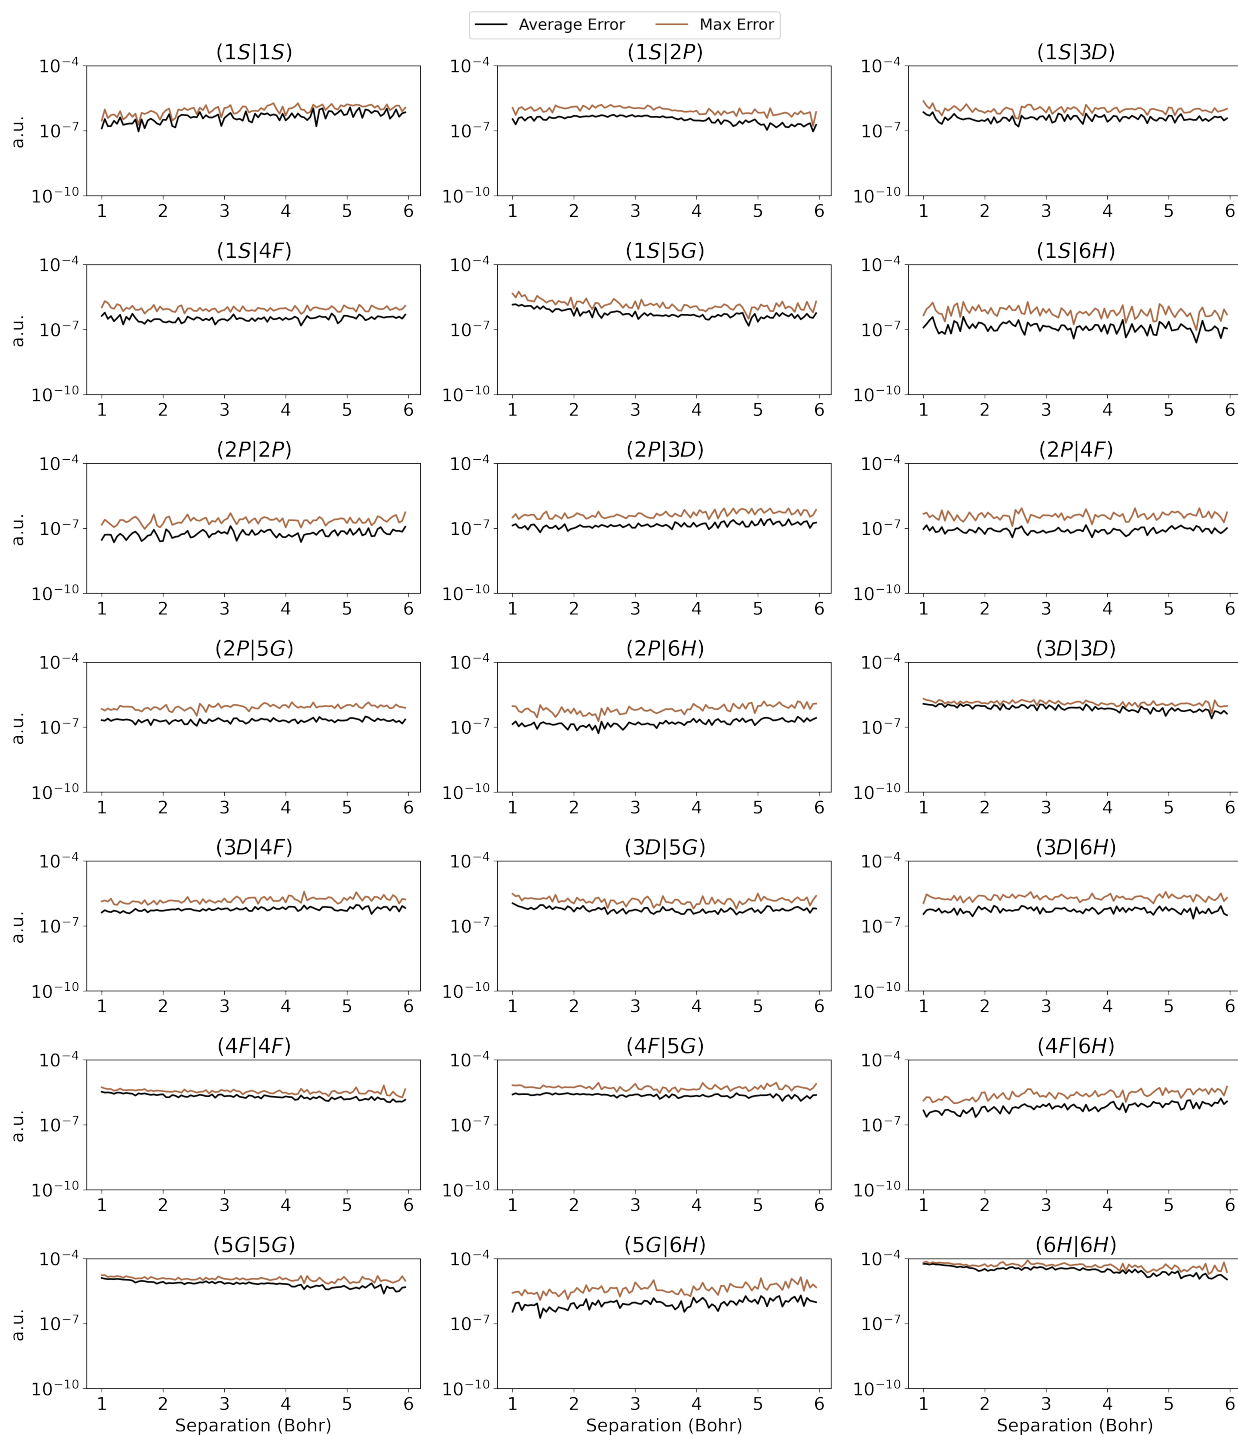


Figure D.3: The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.

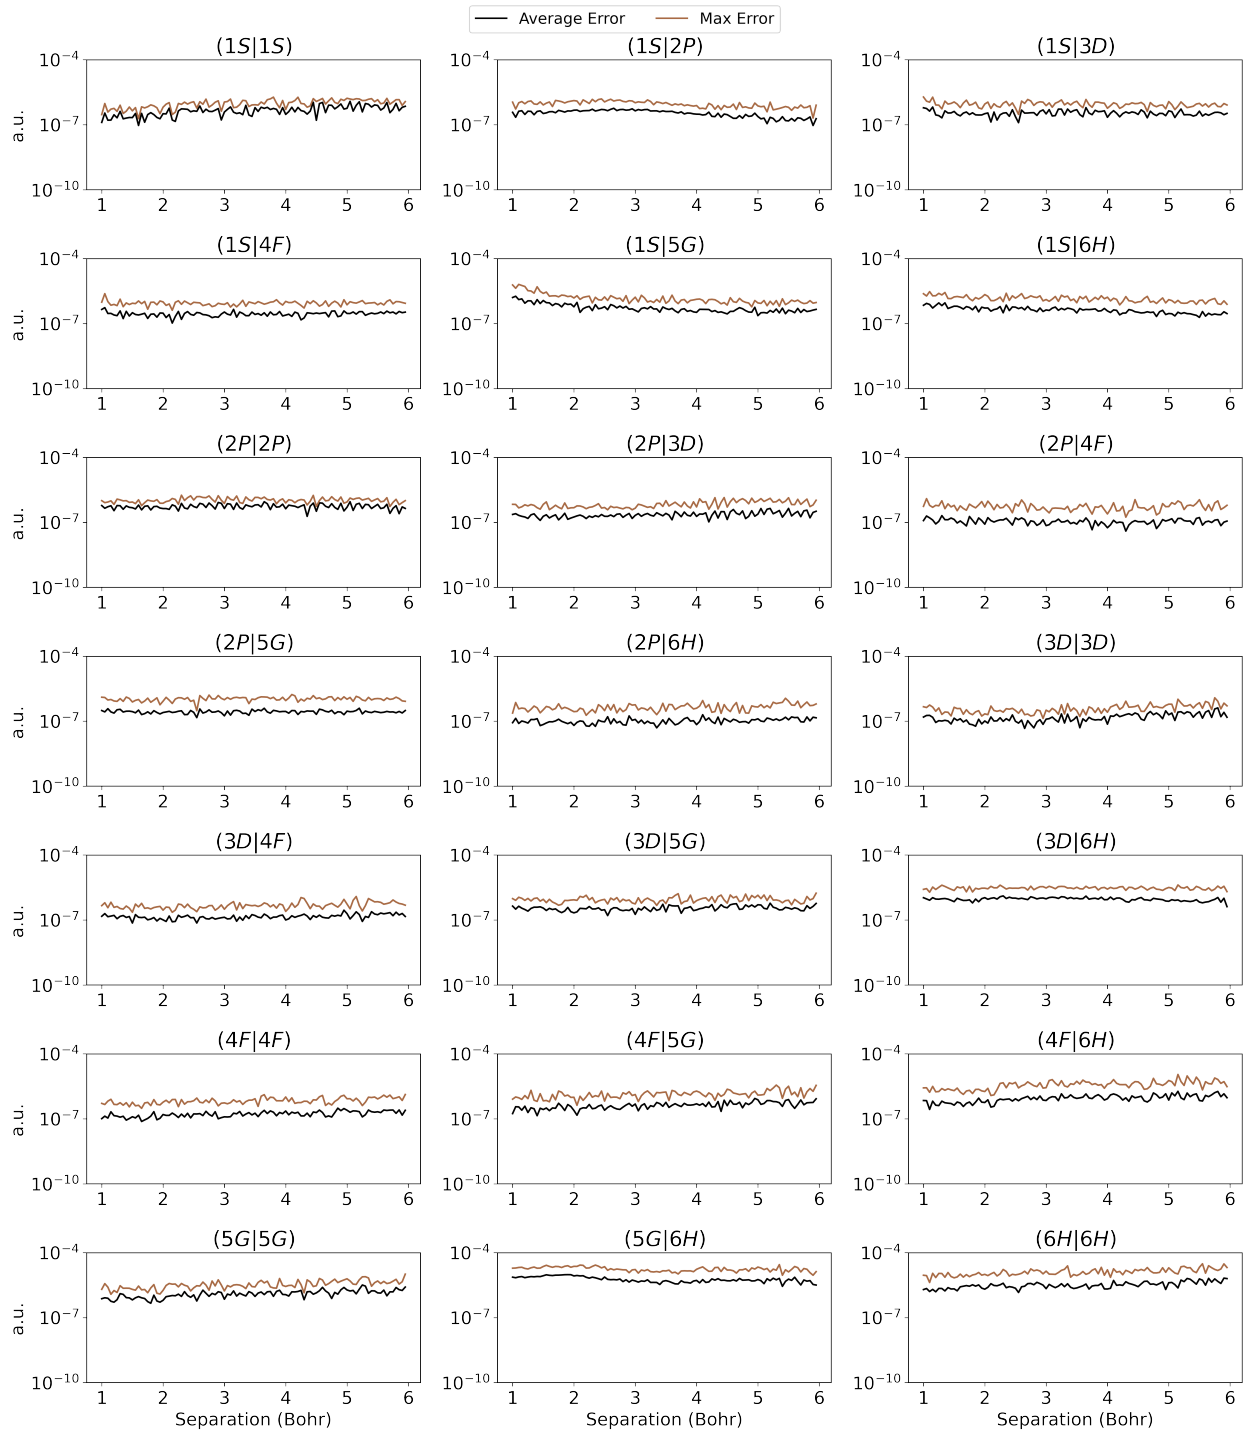


Figure D.4: The max and average errors between mixed- and double-precision integral evaluation are plotted for various basis functions. All basis functions have $\zeta = 1$. The left basis has $m = 0$ and the right basis has $m = n - 1$. The max and average errors are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.

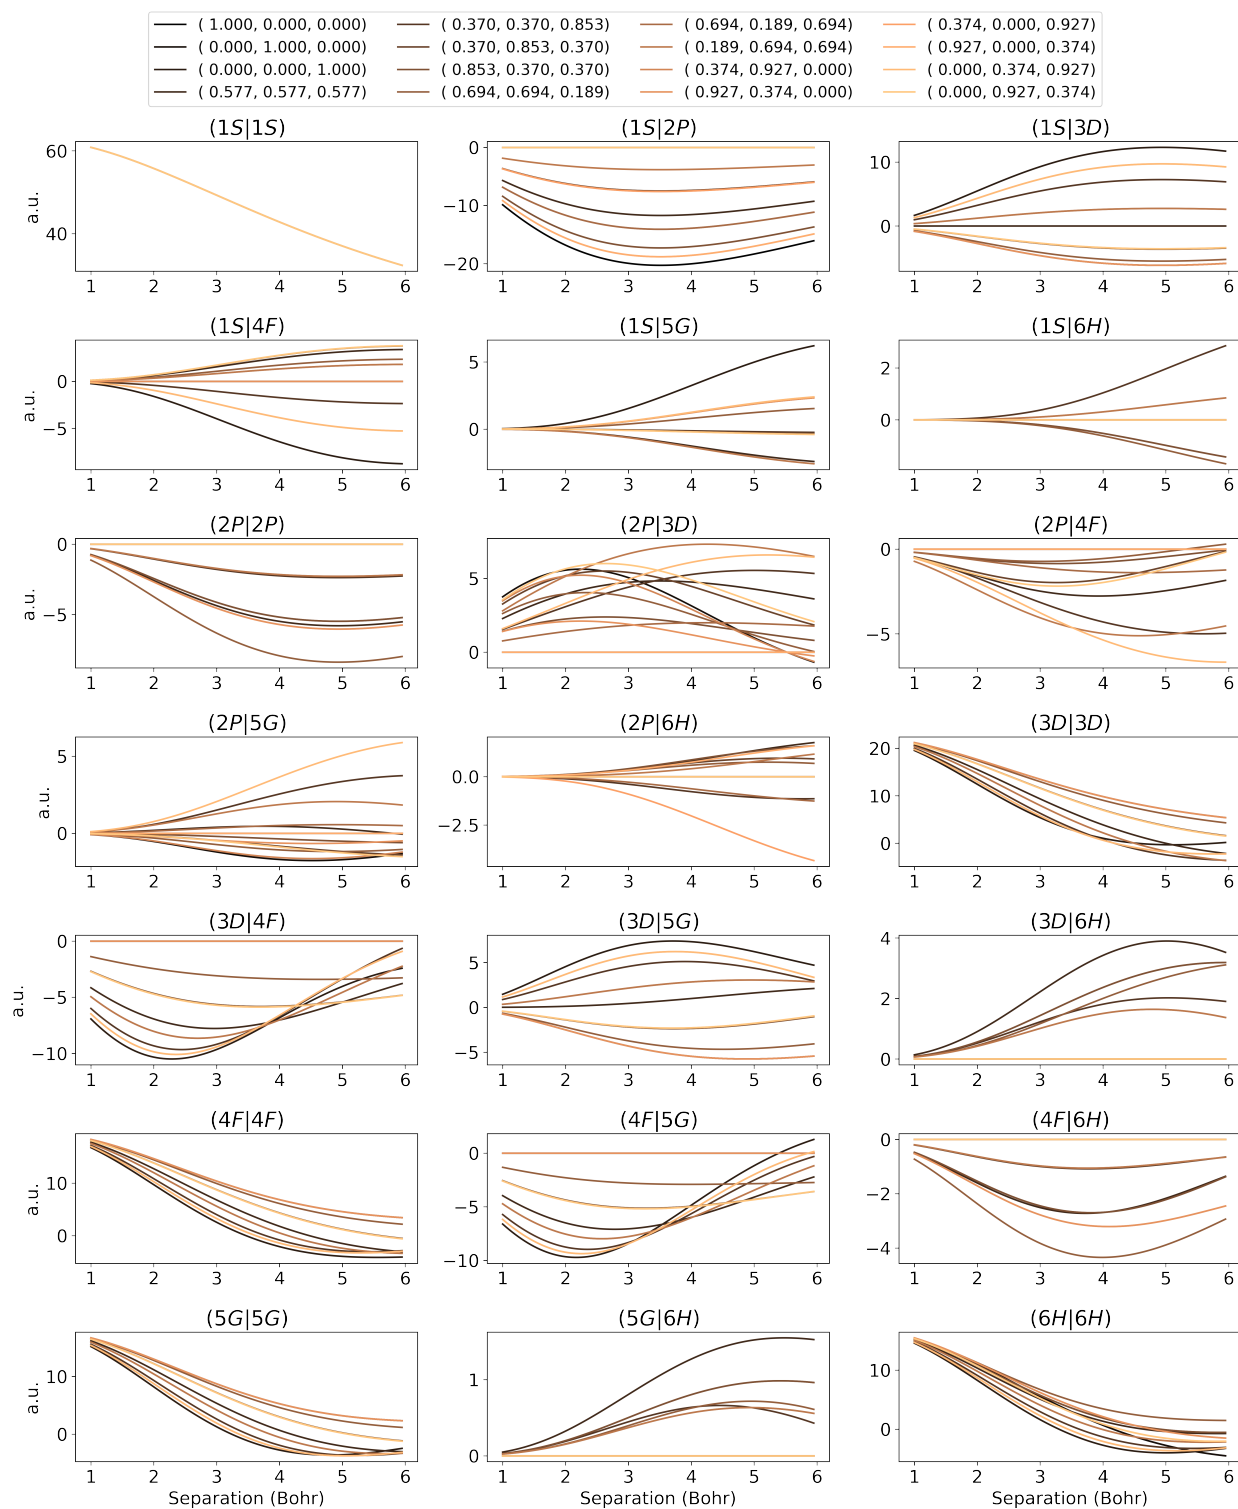


Figure D.5: The mixed-precision 2-center ERIs are plotted for various basis functions. All basis functions have $\zeta = 1$ and $m = 0$. The ERIs are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.

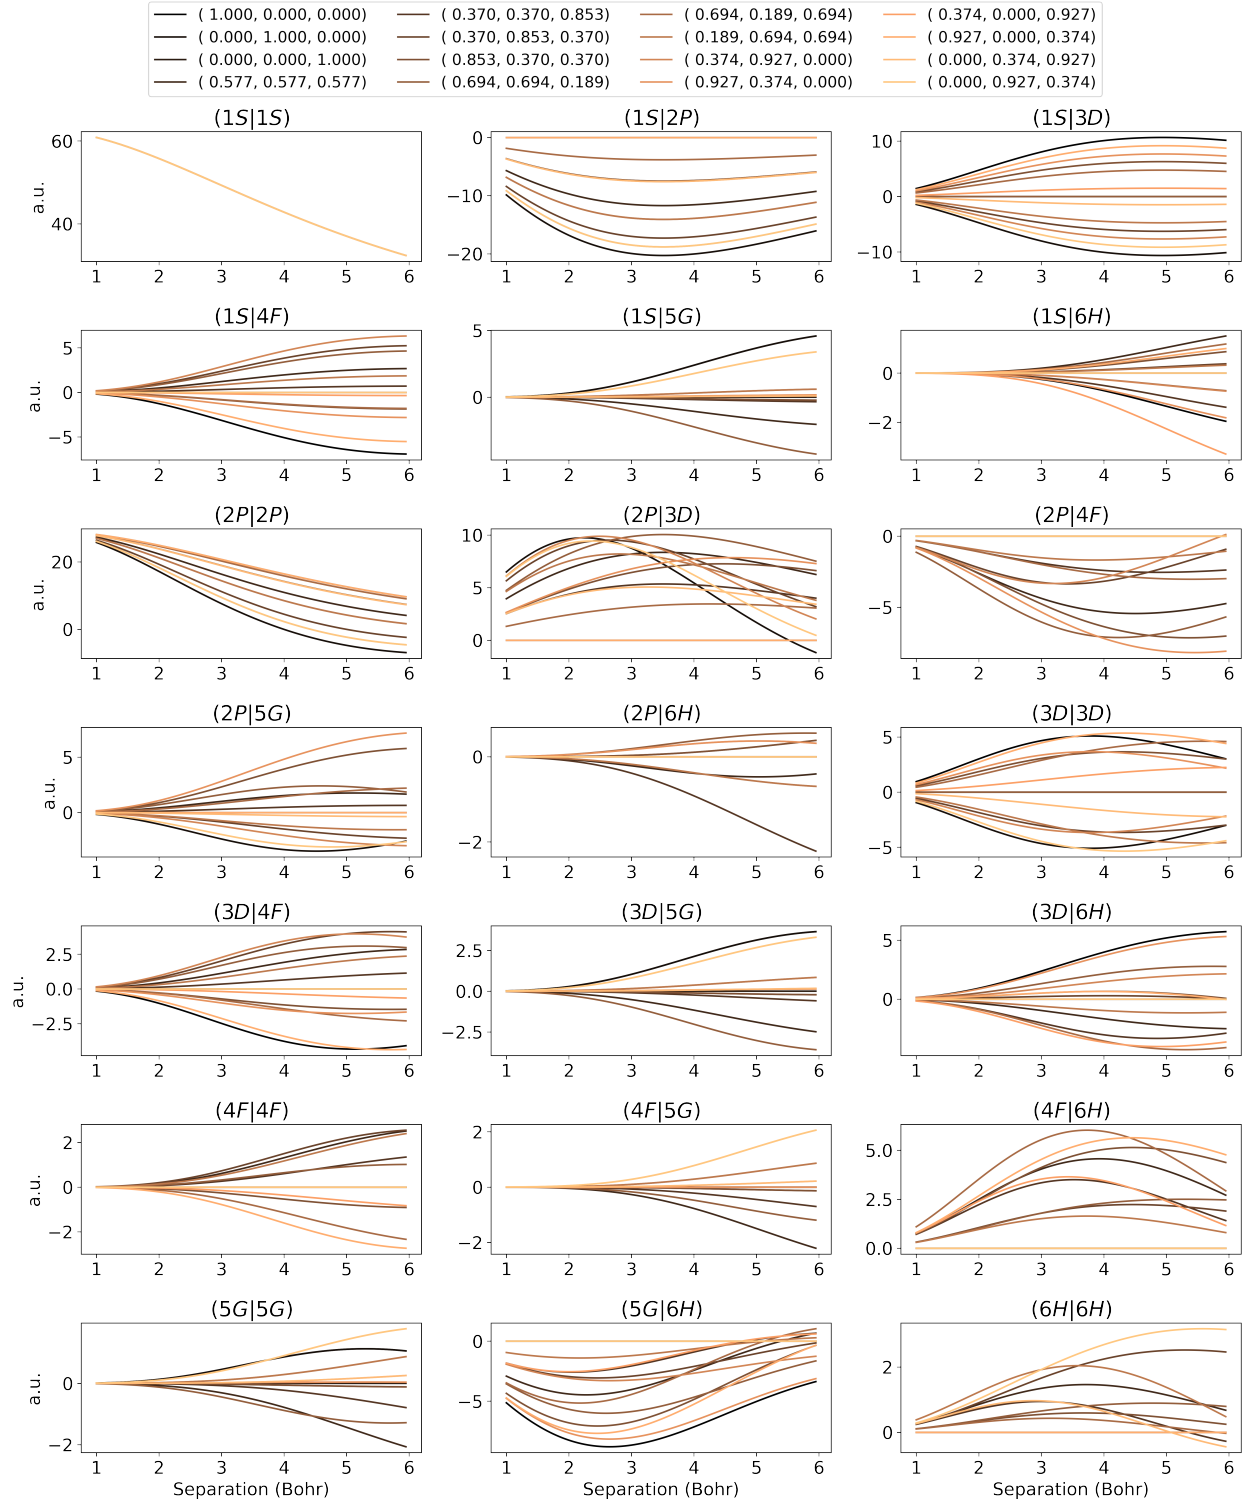


Figure D.6: The mixed-precision 2-center ERIs are plotted for various basis functions. All basis functions have $\zeta = 1$. The left basis has $m = 0$ and the right basis has $m = n - 1$. The ERIs are computed over internuclear distance scans based on the 16 all-positive directions of a Lebedev grid.

BIBLIOGRAPHY

- [1] A. Szabo and N. S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Publications, Inc., Mineola, first edition, 1996.
- [2] T. Helgaker, P. Jørgensen, and J. Olsen. *Molecular Electronic Structure Theory*. John Wiley & Sons, LTD, 2000.
- [3] J. C. Slater. Central Fields and Rydberg Formulas in Wave Mechanics. *Phys. Rev.*, 31: 333–343, Mar 1928.
- [4] A. J. Cohen and N. C. Handy. Density functional generalized gradient calculations using Slater basis sets. *The Journal of Chemical Physics*, 117(4):1470–1478, 2002.
- [5] S. F. Boys and A. C. Egerton. Electronic wave functions - I. A general method of calculation for the stationary states of any molecular system. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 200(1063):542–554, 1950.
- [6] M. A. Watson, N. C. Handy, A. J. Cohen, and T. Helgaker. Density-functional generalized-gradient and hybrid calculations of electromagnetic properties using Slater basis sets. *The Journal of Chemical Physics*, 120(16):7252–7261, 2004.
- [7] P. E. Hoggan. Choice of atomic orbitals to evaluate sensitive properties of molecules: An example of NMR chemical shifts. *International Journal of Quantum Chemistry*, 100(2):214–220, 2004.
- [8] P. W. Ayers, R. C. Morrison, and R. G. Parr. Fermi-Amaldi model for exchange-correlation: atomic excitation energies from orbital energy differences. *Molecular Physics*, 103(15-16):2061–2072, 2005.
- [9] B. Kanungo, P. M. Zimmerman, and V. Gavini. Exact exchange-correlation potentials from ground-state electron densities. *Nature Communications*, 10(1):4497, 2019. ISSN 2041-1723.
- [10] B. Kanungo, P. M. Zimmerman, and V. Gavini. A Comparison of Exact and Model Exchange–Correlation Potentials for Molecules. *The Journal of Physical Chemistry Letters*, 12(50):12012–12019, 2021.

- [11] D. Schneider. The Exascale Era is Upon Us: The Frontier supercomputer may be the first to reach 1,000,000,000,000,000 operations per second. *IEEE Spectrum*, 59(1): 34–35, 2022.
- [12] M. Urban, J. Noga, S. J. Cole, and R. J. Bartlett. Towards a full CCSDT model for electron correlation. *The Journal of Chemical Physics*, 83(8):4041–4046, 1985.
- [13] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon. A fifth-order perturbation comparison of electron correlation theories. *Chemical Physics Letters*, 157(6):479–483, 1989. ISSN 0009-2614.
- [14] V. Rishi, A. Perera, and R. J. Bartlett. Assessing the distinguishable cluster approximation based on the triple bond-breaking in the nitrogen molecule. *The Journal of Chemical Physics*, 144(12):124117, 2016.
- [15] B. Eichengreen. *Globalizing Capital: A History of the International Monetary System*. Princeton University Press, 3rd edition, 2019. ISBN 9780691193908.
- [16] A. I. Krylov. Size-consistent wave functions for bond-breaking: the equation-of-motion spin-flip model. *Chemical Physics Letters*, 338(4):375–384, 2001. ISSN 0009-2614.
- [17] J. S. Sears, C. D. Sherrill, and A. I. Krylov. A spin-complete version of the spin-flip approach to bond breaking: What is the impact of obtaining spin eigenfunctions? *The Journal of Chemical Physics*, 118(20):9084–9094, 2003.
- [18] D. Casanova and M. Head-Gordon. The spin-flip extended single excitation configuration interaction method. *The Journal of Chemical Physics*, 129(6):064104, 2008.
- [19] P. M. Zimmerman, F. Bell, M. Goldey, A. T. Bell, and M. Head-Gordon. Restricted active space spin-flip configuration interaction: Theory and examples for multiple spin flips with odd numbers of electrons. *The Journal of Chemical Physics*, 137(16):164110, 2012.
- [20] F. Bell, P. M. Zimmerman, D. Casanova, M. Goldey, and M. Head-Gordon. Restricted active space spin-flip (RAS-SF) with arbitrary number of spin-flips. *Phys. Chem. Chem. Phys.*, 15:358–366, 2013.
- [21] A. D. Chien and P. M. Zimmerman. Recovering dynamic correlation in spin flip configuration interaction through a difference dedicated approach. *The Journal of Chemical Physics*, 146(1):014103, 2017.
- [22] H. Jiang and P. M. Zimmerman. Charge transfer via spin flip configuration interaction: Benchmarks and application to singlet fission. *The Journal of Chemical Physics*, 153(6):064109, 2020.
- [23] B. Prajapati, D.-K. Dang, P. J. Chmielewski, M. A. Majewski, T. Lis, C. J. Gómez-García, P. M. Zimmerman, and M. Stepień. An Open-Shell Coronoid with Hybrid Chichibabin–Schlenk Conjugation. *Angewandte Chemie International Edition*, 60(41): 22496–22504, 2021.

- [24] G. 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. *The Journal of Chemical Physics*, 131(5):054106, 2009.
- [25] E. Giner, A. Scemama, and M. Caffarel. Using perturbatively selected configuration interaction in quantum Monte Carlo calculations. *Canadian Journal of Chemistry*, 91(9):879–885, 2013.
- [26] T. P. Kelly, A. Perera, R. J. Bartlett, and J. C. Greer. Monte Carlo configuration interaction with perturbation corrections for dissociation energies of first row diatomic molecules: C₂, N₂, O₂, CO, and NO. *The Journal of Chemical Physics*, 140(8):084114, 2014.
- [27] A. Scemama, T. Applencourt, E. Giner, and M. Caffarel. Quantum Monte Carlo with very large multideterminant wavefunctions. *Journal of Computational Chemistry*, 37(20):1866–1875, 2016.
- [28] Y. Garniron, A. Scemama, P.-F. Loos, and M. Caffarel. Hybrid stochastic-deterministic calculation of the second-order perturbative contribution of multireference perturbation theory. *The Journal of Chemical Physics*, 147(3):034101, 2017.
- [29] B. Huron, J. P. Malrieu, and P. Rancurel. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *The Journal of Chemical Physics*, 58(12):5745–5759, 1973.
- [30] S. Evangelisti, J.-P. Daudey, and J.-P. Malrieu. Convergence of an improved CIPSI algorithm. *Chemical Physics*, 75(1):91–102, 1983. ISSN 0301-0104.
- [31] F. A. Evangelista. Adaptive multiconfigurational wave functions. *The Journal of Chemical Physics*, 140(12):124114, 2014.
- [32] T. Zhang and F. A. Evangelista. A Deterministic Projector Configuration Interaction Approach for the Ground State of Quantum Many-Body Systems. *Journal of Chemical Theory and Computation*, 12(9):4326–4337, 2016.
- [33] J. B. Schriber and F. A. Evangelista. Communication: An adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *The Journal of Chemical Physics*, 144(16):161106, 2016.
- [34] J. B. Schriber and F. A. Evangelista. Adaptive Configuration Interaction for Computing Challenging Electronic Excited States with Tunable Accuracy. *Journal of Chemical Theory and Computation*, 13(11):5354–5366, 2017.
- [35] N. M. Tubman, J. Lee, T. Y. Takeshita, M. Head-Gordon, and K. B. Whaley. A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *The Journal of Chemical Physics*, 145(4):044112, 2016.

- [36] N. M. Tubman, D. S. Levine, D. Hait, M. Head-Gordon, and K. B. Whaley. An efficient deterministic perturbation theory for selected configuration interaction methods. *Condensed Matter, Strongly Correlated Electrons*, 2018.
- [37] Z. Wang, Y. Li, and J. Lu. Coordinate Descent Full Configuration Interaction. *Journal of Chemical Theory and Computation*, 15(6):3558–3569, 2019.
- [38] N. M. Tubman, C. D. Freeman, D. S. Levine, D. Hait, M. Head-Gordon, and K. B. Whaley. Modern Approaches to Exact Diagonalization and Selected Configuration Interaction with the Adaptive Sampling CI Method. *Journal of Chemical Theory and Computation*, 16(4):2139–2159, 2020.
- [39] V. Abraham and N. J. Mayhall. Selected Configuration Interaction in a Basis of Cluster State Tensor Products. *Journal of Chemical Theory and Computation*, 16(10):6098–6113, 2020.
- [40] A. A. Holmes, N. M. Tubman, and C. J. Umrigar. Heat-Bath Configuration Interaction: An Efficient Selected Configuration Interaction Algorithm Inspired by Heat-Bath Sampling. *Journal of Chemical Theory and Computation*, 12(8):3674–3680, 2016.
- [41] S. Sharma, A. A. Holmes, G. Jeanmairet, A. Alavi, and C. J. Umrigar. Semistochastic Heat-Bath Configuration Interaction Method: Selected Configuration Interaction with Semistochastic Perturbation Theory. *Journal of Chemical Theory and Computation*, 13(4):1595–1604, 2017.
- [42] A. A. Holmes, C. J. Umrigar, and S. Sharma. Excited states using semistochastic heat-bath configuration interaction. *The Journal of Chemical Physics*, 147(16):164111, 2017.
- [43] A. D. Chien, A. A. Holmes, M. Otten, C. J. Umrigar, S. Sharma, and P. M. Zimmerman. Excited States of Methylene, Polyenes, and Ozone from Heat-Bath Configuration Interaction. *The Journal of Physical Chemistry A*, 122(10):2714–2722, 2018.
- [44] J. Li, M. Otten, A. A. Holmes, S. Sharma, and C. J. Umrigar. Fast semistochastic heat-bath configuration interaction. *The Journal of Chemical Physics*, 149(21):214110, 2018.
- [45] K. R. Brorsen. Quantifying Multireference Character in Multicomponent Systems with Heat-Bath Configuration Interaction. *Journal of Chemical Theory and Computation*, 16(4):2379–2388, 2020.
- [46] Y. Yao and C. J. Umrigar. Orbital Optimization in Selected Configuration Interaction Methods. *Journal of Chemical Theory and Computation*, 17(7):4183–4194, 2021.
- [47] B. O. Roos. The complete active space SCF method in a fock-matrix-based super-CI formulation. *International Journal of Quantum Chemistry*, 18(S14):175–189, 1980.

- [48] B. O. Roos, P. R. Taylor, and P. E. Sigbahn. A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach. *Chemical Physics*, 48(2):157–173, 1980. ISSN 0301-0104.
- [49] H. Werner and P. J. Knowles. A second order multiconfiguration SCF procedure with optimum convergence. *The Journal of Chemical Physics*, 82(11):5053–5063, 1985.
- [50] M. W. Schmidt and M. S. Gordon. THE CONSTRUCTION AND INTERPRETATION OF MCSCF WAVEFUNCTIONS. *Annual Review of Physical Chemistry*, 49(1):233–266, 1998.
- [51] P. G. Szalay, T. Müller, G. Gidofalvi, H. Lischka, and R. Shepard. Multiconfiguration Self-Consistent Field and Multireference Configuration Interaction Methods and Applications. *Chemical Reviews*, 112(1):108–181, 2012.
- [52] K. D. Vogiatzis, D. Ma, J. Olsen, L. Gagliardi, and W. A. de Jong. Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations. *The Journal of Chemical Physics*, 147(18):184111, 2017.
- [53] J. Olsen, B. O. Roos, P. Jørgensen, and H. J. A. Jensen. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *The Journal of Chemical Physics*, 89(4):2185–2192, 1988.
- [54] P. A. Malmqvist, A. Rendell, and B. O. Roos. The restricted active space self-consistent-field method, implemented with a split graph unitary group approach. *The Journal of Physical Chemistry*, 94(14):5477–5482, 1990.
- [55] M. J. Bearpark, F. Ogliaro, T. Vreven, M. Boggio-Pasqua, M. J. Frisch, S. M. Larkin, M. Morrison, and M. A. Robb. CASSCF calculations for photoinduced processes in large molecules: Choosing when to use the RASSCF, ONIOM and MMVB approximations. *Journal of Photochemistry and Photobiology A: Chemistry*, 190(2):207–227, 2007. ISSN 1010-6030.
- [56] P. A. Malmqvist, K. Pierloot, A. R. M. Shahi, C. J. Cramer, and L. Gagliardi. The restricted active space followed by second-order perturbation theory method: Theory and application to the study of CuO₂ and Cu₂O₂ systems. *The Journal of Chemical Physics*, 128(20):204109, 2008.
- [57] V. Sauri, L. Serrano-Andrés, A. R. M. Shahi, L. Gagliardi, S. Vancoillie, and K. Pierloot. Multiconfigurational Second-Order Perturbation Theory Restricted Active Space (RASPT2) Method for Electronic Excited States: A Benchmark Study. *Journal of Chemical Theory and Computation*, 7(1):153–168, 2011.
- [58] D. Ma, G. Li Manni, and L. Gagliardi. The generalized active space concept in multiconfigurational self-consistent field methods. *The Journal of Chemical Physics*, 135(4):044128, 2011.

- [59] G. Li Manni, D. Ma, F. Aquilante, J. Olsen, and L. Gagliardi. SplitGAS Method for Strong Correlation and the Challenging Case of Cr2. *Journal of Chemical Theory and Computation*, 9(8):3375–3384, 2013.
- [60] S. R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, 1992.
- [61] S. R. White. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B*, 48:10345–10356, 1993.
- [62] S. R. White and R. L. Martin. Ab initio quantum chemistry using the density matrix renormalization group. *The Journal of Chemical Physics*, 110(9):4127–4130, 1999.
- [63] U. Schollwöck. The density-matrix renormalization group. *Rev. Mod. Phys.*, 77:259–315, Apr 2005.
- [64] D. Zgid and M. Nooijen. On the spin and symmetry adaptation of the density matrix renormalization group method. *The Journal of Chemical Physics*, 128(1):014107, 2008.
- [65] U. Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011. ISSN 0003-4916.
- [66] G. K.-L. Chan and S. Sharma. The Density Matrix Renormalization Group in Quantum Chemistry. *Annual Review of Physical Chemistry*, 62(1):465–481, 2011.
- [67] M. R. Hermes and L. Gagliardi. Multiconfigurational Self-Consistent Field Theory with Density Matrix Embedding: The Localized Active Space Self-Consistent Field Method. *Journal of Chemical Theory and Computation*, 15(2):972–986, 2019.
- [68] M. R. Hermes, R. Pandharkar, and L. Gagliardi. Variational Localized Active Space Self-Consistent Field Method. *Journal of Chemical Theory and Computation*, 16(8):4923–4937, 2020.
- [69] P. M. Zimmerman and A. E. Rask. Evaluation of full valence correlation energies and gradients. *The Journal of Chemical Physics*, 150(24):244117, 2019.
- [70] D.-K. Dang and P. M. Zimmerman. Fully variational incremental CASSCF. *The Journal of Chemical Physics*, 154(1):014105, 2021.
- [71] A. B. Mukhopadhyay, M. Dolg, and C. Oligschleger. Ab initio many-body investigation of structure and stability of two-fold rings in silicates. *The Journal of Chemical Physics*, 120(18):8734–8739, 2004.
- [72] J. Friedrich, M. Hanrath, and M. Dolg. Fully automated implementation of the incremental scheme: Application to CCSD energies for hydrocarbons and transition metal compounds. *The Journal of Chemical Physics*, 126(15):154110, 2007.
- [73] H. Stoll. The correlation energy of crystalline silicon. *Chemical Physics Letters*, 191(6):548–552, 1992. ISSN 0009-2614.

- [74] H. Stoll. Correlation energy of diamond. *Phys. Rev. B*, 46:6700–6704, Sep 1992.
- [75] H. Stoll. On the correlation energy of graphite. *The Journal of Chemical Physics*, 97(11):8449–8454, 1992.
- [76] J. J. Eriksen, F. Lipparini, and J. Gauss. Virtual Orbital Many-Body Expansions: A Possible Route towards the Full Configuration Interaction Limit. *The Journal of Physical Chemistry Letters*, 8(18):4633–4639, 2017.
- [77] J. J. Eriksen and J. Gauss. Many-Body Expanded Full Configuration Interaction. I. Weakly Correlated Regime. *Journal of Chemical Theory and Computation*, 14(10):5180–5191, 2018.
- [78] J. J. Eriksen and J. Gauss. Many-Body Expanded Full Configuration Interaction. II. Strongly Correlated Regime. *Journal of Chemical Theory and Computation*, 0(0):null, 2019.
- [79] P. M. Zimmerman. Incremental full configuration interaction. *The Journal of Chemical Physics*, 146(10):104102, 2017.
- [80] P. M. Zimmerman. Strong correlation in incremental full configuration interaction. *The Journal of Chemical Physics*, 146(22):224104, 2017.
- [81] P. M. Zimmerman. Singlet-Triplet Gaps through Incremental Full Configuration Interaction. *The Journal of Physical Chemistry A*, 121(24):4712–4720, 2017.
- [82] J. S. Boschen, D. Theis, K. Ruedenberg, and T. L. Windus. Correlation Energy Extrapolation by Many-Body Expansion. *The Journal of Physical Chemistry A*, 121(4):836–844, 2017.
- [83] J. J. Eriksen, T. A. Anderson, J. E. Deustua, K. Ghanem, D. Hait, M. R. Hoffmann, S. Lee, D. S. Levine, I. Magoulas, J. Shen, N. M. Tubman, K. B. Whaley, E. Xu, Y. Yao, N. Zhang, A. Alavi, G. K.-L. Chan, M. Head-Gordon, W. Liu, P. Piecuch, S. Sharma, S. L. Ten-no, C. J. Umrigar, and J. Gauss. The Ground State Electronic Energy of Benzene. *The Journal of Physical Chemistry Letters*, 11(20):8922–8929, 2020.
- [84] M. Gomberg. Triphenylmethyl, ein Fall von dreiwertigem Kohlenstoff. *Berichte der deutschen chemischen Gesellschaft*, 33(3):3150–3163, 1900.
- [85] A. Tschitschibabin. Über einige phenylierte Derivate des p, p-Ditolyls. *Berichte der deutschen chemischen Gesellschaft*, 40(2):1810–1819, 1907.
- [86] W. Schlenk and M. Brauns. Zur frage der metachinoide. *Berichte der deutschen chemischen Gesellschaft*, 48(1):661–669, 1915.
- [87] D. A. Dougherty. Spin control in organic molecules. *Accounts of chemical research*, 24(3):88–94, 1991.
- [88] A. Rajca. Organic diradicals and polyradicals: from spin coupling to magnetism? *Chemical reviews*, 94(4):871–893, 1994.

- [89] Z. Sun, Z. Zeng, and J. Wu. Zethrenes, extended p-quinodimethanes, and periacenes with a singlet biradical ground state. *Accounts of Chemical Research*, 47(8):2582–2591, 2014.
- [90] Z. Zeng, X. Shi, C. Chi, J. T. L. Navarrete, J. Casado, and J. Wu. Pro-aromatic and anti-aromatic π -conjugated molecules: an irresistible wish to be diradicals. *Chemical Society Reviews*, 44(18):6578–6596, 2015.
- [91] T. Kubo. Phenalenyl-Based Open-Shell Polycyclic Aromatic Hydrocarbons. *The Chemical Record*, 15(1):218–232, 2015.
- [92] N. M. Gallagher, A. Olankitwanit, and A. Rajca. High-spin organic molecules. *The Journal of Organic Chemistry*, 80(3):1291–1298, 2015.
- [93] C. K. Frederickson, B. D. Rose, and M. M. Haley. Explorations of the indenofluorenes and expanded quinoidal analogues. *Accounts of Chemical Research*, 50(4):977–987, 2017.
- [94] J. Casado. Para-quinodimethanes: A unified review of the quinoidal-versus-aromatic competition and its implications. *Physical Organic Chemistry of Quinodimethanes*, pages 209–248, 2018.
- [95] Y. Tobe. Quinodimethanes incorporated in non-benzenoid aromatic or antiaromatic frameworks. *Physical Organic Chemistry of Quinodimethanes*, pages 107–168, 2018.
- [96] C. Liu, Y. Ni, X. Lu, G. Li, and J. Wu. Global aromaticity in macrocyclic polyradicaloids: Huckel’s rule or baird’s rule? *Accounts of chemical research*, 52(8):2309–2321, 2019.
- [97] G. E. Rudebusch, J. L. Zafra, K. Jorner, K. Fukuda, J. L. Marshall, I. Arrechea-Marcos, G. L. Espejo, R. Ponce Ortiz, C. J. Gomez-Garcia, L. N. Zakharov, et al. Diindeno-fusion of an anthracene as a design strategy for stable organic biradicals. *Nature Chemistry*, 8(8):753–759, 2016.
- [98] Y. Ni, S. Lee, M. Son, N. Aratani, M. Ishida, A. Samanta, H. Yamada, Y.-T. Chang, H. Furuta, D. Kim, et al. A Diradical Approach towards BODIPY-Based Dyes with Intense Near-Infrared Absorption around $\lambda = 1100$ nm. *Angewandte Chemie International Edition*, 55(8):2815–2819, 2016.
- [99] K. Okino, D. Sakamaki, and S. Seki. Dicyanomethyl radical-based near-infrared thermochromic dyes with high transparency in the visible region. *ACS Materials Letters*, 1(1):25–29, 2019.
- [100] X. Gu, T. Y. Gopalakrishna, H. Phan, Y. Ni, T. S. Heng, J. Ding, and J. Wu. A Three-Dimensionally π -Conjugated Diradical Molecular Cage. *Angewandte Chemie*, 129(48):15585–15589, 2017.
- [101] Y. Ni, T. Y. Gopalakrishna, H. Phan, T. Kim, T. S. Heng, Y. Han, T. Tao, J. Ding, D. Kim, and J. Wu. 3D global aromaticity in a fully conjugated diradicaloid cage at different oxidation states. *Nature Chemistry*, 12(3):242–248, 2020.

- [102] P. Ravat, T. Solomek, D. Haussinger, O. Blacque, and M. Jurícek. Dimethylcethrene: a chiroptical diradicaloid photoswitch. *Journal of the American Chemical Society*, 140(34):10839–10847, 2018.
- [103] D. Shimizu, Y. Ide, T. Ikeue, and A. Osuka. Aminyl Radical Bridged NiIII-Porphyrin Dimer: Coordination-Induced Spin-State Switching between Doublet and Sextet States. *Angew. Chem., Int. Ed.*, 58:5023–5027, 2019.
- [104] D. Beaudoin, O. Levasseur-Grenon, T. Maris, and J. D. Wuest. Building Giant Carbocycles by Reversible C- C Bond Formation. *Angewandte Chemie International Edition*, 55(3):894–898, 2016.
- [105] K. Oda, S. Hiroto, and H. Shinokubo. NIR mechanochromic behaviours of a tetracyanoethylene-bridged hexa-peri-hexabenzocoronene dimer and trimer through dissociation of C–C bonds. *Journal of Materials Chemistry C*, 5(22):5310–5315, 2017.
- [106] H. Yokoi, S. Hiroto, and H. Shinokubo. Reversible σ -bond formation in bowl-shaped π -radical cations: the effects of curved and planar structures. *Journal of the American Chemical Society*, 140(13):4649–4655, 2018.
- [107] L. Yuan, Y. Han, T. Tao, H. Phan, and C. Chi. Formation of a Macrocycles-in-a-Macrocycle Superstructure with All-gauche Conformation by Reversible Radical Association. *Angewandte Chemie*, 130(29):9161–9165, 2018.
- [108] L. Moshniaha, M. Zyla-Karwowska, P. J. Chmielewski, T. Lis, J. Cybinska, E. Gonka, J. Oswald, T. Drewello, S. M. Rivero, J. Casado, et al. Aromatic nanosandwich obtained by σ -dimerization of a nanographene π -radical. *Journal of the American Chemical Society*, 142(7):3626–3635, 2020.
- [109] K. Sahara, M. Abe, H. Zipse, and T. Kubo. Duality of Reactivity of a Biradicaloid Compound with an o-Quinodimethane Scaffold. *Journal of the American Chemical Society*, 142(11):5408–5418, 2020.
- [110] H. Miyoshi, S. Nobusue, A. Shimizu, and Y. Tobe. Non-alternant non-benzenoid kekulenes: the birth of a new kekulene family. *Chemical Society Reviews*, 44(18):6560–6577, 2015.
- [111] S. Nobusue, H. Miyoshi, A. Shimizu, I. Hisaki, K. Fukuda, M. Nakano, and Y. Tobe. Tetracyclopenta [def, jkl, pqr, vwx] tetraphenylene: a potential tetraradicaloid hydrocarbon. *Angewandte Chemie International Edition*, 54(7):2090–2094, 2015.
- [112] S. Das, T. S. Herng, J. L. Zafra, P. M. Burrezo, M. Kitano, M. Ishida, T. Y. Gopalakrishna, P. Hu, A. Osuka, J. Casado, et al. Fully fused quinoidal/aromatic carbazole macrocycles with poly-radical characters. *Journal of the American Chemical Society*, 138(24):7782–7790, 2016.
- [113] C. Liu, M. E. Sandoval-Salinas, Y. Hong, T. Y. Gopalakrishna, H. Phan, N. Aratani, T. S. Herng, J. Ding, H. Yamada, D. Kim, et al. Macrocyclic polyradicaloids with unusual super-ring structure and global aromaticity. *Chem*, 4(7):1586–1595, 2018.

- [114] X. Lu, T. Y. Gopalakrishna, H. Phan, T. S. Heng, Q. Jiang, C. Liu, G. Li, J. Ding, and J. Wu. Global Aromaticity in Macrocyclic Cyclopenta-Fused Tetraphenanthrylene Tetraradicaloid and Its Charged Species. *Angewandte Chemie*, 130(40):13236–13240, 2018.
- [115] H. Gregolinska, M. Majewski, P. J. Chmielewski, J. Gregolinski, A. Chien, J. Zhou, Y.-L. Wu, Y. J. Bae, M. R. Wasielewski, P. M. Zimmerman, et al. Fully conjugated [4] chrysaorene. Redox-coupled anion binding in a tetraradicaloid macrocycle. *Journal of the American Chemical Society*, 140(43):14474–14480, 2018.
- [116] X. Lu, T. Y. Gopalakrishna, Y. Han, Y. Ni, Y. Zou, and J. Wu. Bowl-shaped carbon nanobelts showing size-dependent properties and selective encapsulation of C70. *Journal of the American Chemical Society*, 141(14):5934–5941, 2019.
- [117] X. Lu, D. An, Y. Han, Y. Zou, Y. Qiao, N. Zhang, D. Chang, J. Wu, and Y. Liu. A cyclopenta-fused dibenzo [b, d] thiophene-co-phenanthrene macrocyclic tetraradicaloid. *Chemical science*, 12(11):3952–3957, 2021.
- [118] M. Stepień. An aromatic riddle: decoupling annulene conjugation in coronoid macrocycles. *Chem*, 4(7):1481–1483, 2018.
- [119] L. K. Montgomery, J. C. Huffman, E. A. Jurczak, and M. P. Grendze. The molecular structures of Thiele’s and Chichibabin’s hydrocarbons. *Journal of the American Chemical Society*, 108(19):6004–6011, 1986.
- [120] G. Kothe, K.-H. Denkel, and W. Sümmermann. Schlenk’s Biradical.—A Molecule in the Triplet Ground State. *Angewandte Chemie International Edition in English*, 9(11):906–907, 1970.
- [121] A. Shimizu, R. Kishi, M. Nakano, D. Shiomi, K. Sato, T. Takui, I. Hisaki, M. Miyata, and Y. Tobe. Indeno [2, 1-b] fluorene: A 20- π -Electron Hydrocarbon with Very Low-Energy Light Absorption. *Angewandte Chemie*, 125(23):6192–6195, 2013.
- [122] M. A. Majewski, P. J. Chmielewski, A. Chien, Y. Hong, T. Lis, M. Witwicki, D. Kim, P. M. Zimmerman, and M. Stepień. 5, 10-Dimesityldiindeno [1, 2-a: 2, 1-i] phenanthrene: a stable biradicaloid derived from Chichibabin’s hydrocarbon. *Chemical science*, 10(11):3413–3420, 2019.
- [123] S. Grimme, S. Ehrlich, and L. Goerigk. Effect of the damping function in dispersion corrected density functional theory. *Journal of computational chemistry*, 32(7):1456–1465, 2011.
- [124] T. Yanai, D. P. Tew, and N. C. Handy. A new hybrid exchange–correlation functional using the Coulomb-attenuating method (CAM-B3LYP). *Chemical physics letters*, 393(1-3):51–57, 2004.
- [125] F. Aquilante, T. B. Pedersen, and R. Lindh. Low-cost evaluation of the exchange Fock matrix from Cholesky and density fitting representations of the electron repulsion integrals. *The Journal of chemical physics*, 126(19):194106, 2007.

- [126] J. Bostrom, F. Aquilante, T. B. Pedersen, and R. Lindh. Ab initio density fitting: Accuracy assessment of auxiliary basis sets from cholesky decompositions. *Journal of Chemical Theory and Computation*, 5(6):1545–1553, 2009.
- [127] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, G. A. Petersson, H. Nakatsuji, X. Li, M. Caricato, A. V. Marenich, J. Bloino, B. G. Janesko, R. Gomperts, B. Mennucci, H. P. Hratchian, J. V. Ortiz, A. F. Izmaylov, J. L. Sonnenberg, D. Williams-Young, F. Ding, F. Lipparini, F. Egidi, J. Goings, B. Peng, A. Petrone, T. Henderson, D. Ranasinghe, V. G. Zakrzewski, J. Gao, N. Rega, G. Zheng, W. Liang, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, K. Throssell, J. A. Montgomery, Jr., J. E. Peralta, F. Ogliaro, M. J. Bearpark, J. J. Heyd, E. N. Brothers, K. N. Kudin, V. N. Staroverov, T. A. Keith, R. Kobayashi, J. Normand, K. Raghavachari, A. P. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, J. M. Millam, M. Klene, C. Adamo, R. Cammi, J. W. Ochterski, R. L. Martin, K. Morokuma, O. Farkas, J. B. Foresman, and D. J. Fox. Gaussian~16 Revision C.01, 2016.
- [128] I. Fdez. Galván, M. Vacher, A. Alavi, C. Angeli, F. Aquilante, J. Autschbach, J. J. Bao, S. I. Bokarev, N. A. Bogdanov, R. K. Carlson, L. F. Chibotaru, J. Creutzberg, N. Dattani, M. G. Delcey, S. S. Dong, A. Dreuw, L. Freitag, L. M. Frutos, L. Gagliardi, F. Gendron, A. Giussani, L. González, G. Grell, M. Guo, C. E. Hoyer, M. Johansson, S. Keller, S. Knecht, G. Kovačević, E. Källman, G. Li Manni, M. Lundberg, Y. Ma, S. Mai, J. P. Malhado, P. Malmqvist, P. Marquetand, S. A. Mewes, J. Norell, M. Olivucci, M. Oppel, Q. M. Phung, K. Pierloot, F. Plasser, M. Reiher, A. M. Sand, I. Schapiro, P. Sharma, C. J. Stein, L. K. Sørensen, D. G. Truhlar, M. Ugandi, L. Ungur, A. Valentini, S. Vancoillie, V. Veryazov, O. Weser, T. A. Wesolowski, P.-O. Widmark, S. Wouters, A. Zech, J. P. Zobel, and R. Lindh. OpenMolcas: From Source Code to Insight. *Journal of Chemical Theory and Computation*, 15(11):5925–5964, 2019.
- [129] D. Doehnert and J. Koutecky. Occupation numbers of natural orbitals as a criterion for biradical character. Different kinds of biradicals. *Journal of the American Chemical Society*, 102(6):1789–1796, 1980.
- [130] M. Nakano. Electronic structure of open-shell singlet molecules: Diradical character viewpoint. *Topics in Current Chemistry*, 375(2):1–67, 2017.
- [131] M. Head-Gordon. Characterizing unpaired electrons from the one-particle density matrix. *Chemical physics letters*, 372(3-4):508–511, 2003.
- [132] O. Kahn. Competing spin interactions and degenerate frustration for discrete molecular species. *Chemical physics letters*, 265(1-2):109–114, 1997.
- [133] D. Dai and M.-H. Whangbo. Spin exchange interactions of a spin dimer: Analysis of broken-symmetry spin states in terms of the eigenstates of Heisenberg and Ising spin Hamiltonians. *The Journal of Chemical Physics*, 118(1):29–39, 2003.

- [134] D.-K. Dang, J. A. Kammeraad, and P. M. Zimmerman. Advances in Parallel Heat Bath Configuration Interaction. *The Journal of Physical Chemistry A*, 127(1):400–411, 2023.
- [135] P. S. Epstein. The Stark Effect from the Point of View of Schroedinger’s Quantum Theory. *Phys. Rev.*, 28:695–710, Oct 1926.
- [136] R. K. Nesbet. Configuration interaction in orbital theories. *Proc. R. Soc. Lond. A*, 230: 312–321, 1955.
- [137] Fowler-Noll-Vo hash function. https://en.wikipedia.org/wiki/Fowler%E2%80%9393Noll%E2%80%9393Vo_hash_function.
- [138] The Parallel Hashmap. <https://github.com/greg7mdp/parallel-hashmap>.
- [139] E. R. Davidson. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *Journal of Computational Physics*, 17(1):87–94, 1975. ISSN 0021-9991.
- [140] A. D. Chien and P. M. Zimmerman. Iterative submatrix diagonalisation for large configuration interaction problems. *Molecular Physics*, 116(1):107–117, 2018.
- [141] F. Weigend, A. Köhn, and C. Hättig. Efficient use of the correlation consistent basis sets in resolution of the identity MP2 calculations. *The Journal of Chemical Physics*, 116(8):3175–3183, 2002.
- [142] T. H. Dunning. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.*, 90:1007–1023, 1989.
- [143] C. Daday, S. Smart, G. H. Booth, A. Alavi, and C. Filippi. Full Configuration Interaction Excitations of Ethene and Butadiene: Resolution of an Ancient Question. *Journal of Chemical Theory and Computation*, 8(11):4441–4451, 2012.
- [144] E. Monino, M. Boggio-Pasqua, A. Scemama, D. Jacquemin, and P.-F. Loos. Reference Energies for Cyclobutadiene: Automerization and Excited States. *The Journal of Physical Chemistry A*, 126(28):4664–4679, 2022.
- [145] M. Feldt, Q. M. Phung, K. Pierloot, R. A. Mata, and J. N. Harvey. Limits of Coupled-Cluster Calculations for Non-Heme Iron Complexes. *Journal of Chemical Theory and Computation*, 15(2):922–937, 2019.
- [146] Q. M. Phung, C. Martín-Fernández, J. N. Harvey, and M. Feldt. Ab Initio Calculations for Spin-Gaps of Non-Heme Iron Complexes. *Journal of Chemical Theory and Computation*, 15(8):4297–4304, 2019.
- [147] M. Feldt, C. Martín-Fernández, and J. N. Harvey. Energetics of non-heme iron reactivity: can ab initio calculations provide the right answer? *Phys. Chem. Chem. Phys.*, 22: 23908–23919, 2020.

- [148] M. Drosou, C. A. Mitsopoulou, and D. A. Pantazis. Reconciling Local Coupled Cluster with Multireference Approaches for Transition Metal Spin-State Energetics. *Journal of Chemical Theory and Computation*, 18(6):3538–3548, 2022.
- [149] R. Ditchfield, W. J. Hehre, and J. A. Pople. Self-Consistent Molecular-Orbital Methods. IX. An Extended Gaussian-Type Basis for Molecular-Orbital Studies of Organic Molecules. *J. Chem. Phys.*, 54:724–728, 1971.
- [150] W. J. Hehre, R. Ditchfield, and J. A. Pople. 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.*, 56:2257–2261, 1972.
- [151] P. C. Hariharan and J. A. Pople. The influence of polarization functions on molecular orbital hydrogenation energies. *Theor. Chim. Acta*, 28:213–222, 1973.
- [152] R. Olivares-Amaya, W. Hu, N. Nakatani, S. Sharma, J. Yang, and G. K.-L. Chan. The ab-initio density matrix renormalization group in practice. *The Journal of Chemical Physics*, 142(3):034102, 2015.
- [153] G. Maier, R. Wolf, and H.-O. Kalinowski. How High is the Barrier for the Valence Isomerization of Cyclobutadiene? *Angewandte Chemie International Edition in English*, 31(6):738–740, 1992.
- [154] M. Eckert-Maksić, M. Vazdar, M. Barbatti, H. Lischka, and Z. B. Maksić. Automerization reaction of cyclobutadiene and its barrier height: An ab initio benchmark multireference average-quadratic coupled cluster study. *The Journal of Chemical Physics*, 125(6):064310, 2006.
- [155] O. Demel and J. Pittner. Multireference Brillouin-Wigner coupled clusters method with noniterative perturbative connected triples. *The Journal of Chemical Physics*, 124(14):144112, 2006.
- [156] O. Demel, K. R. Shamasundar, L. Kong, and M. Nooijen. Application of Double Ionization State-Specific Equation of Motion Coupled Cluster Method to Organic Diradicals. *The Journal of Physical Chemistry A*, 112(46):11895–11902, 2008.
- [157] P. B. Karadakov. Ground- and Excited-State Aromaticity and Antiaromaticity in Benzene and Cyclobutadiene. *The Journal of Physical Chemistry A*, 112(31):7303–7309, 2008.
- [158] T. Saito, S. Nishihara, Y. Kitagawa, T. Kawakami, S. Yamanaka, M. Okumura, and K. Yamaguchi. A broken-symmetry study on the automerization of cyclobutadiene. Comparison with UNO- and DNO-MRCC methods. *Chemical Physics Letters*, 498(4): 253–258, 2010. ISSN 0009-2614.
- [159] D. I. Lyakh, V. F. Lotrich, and R. J. Bartlett. The ‘tailored’ CCSD(T) description of the automerization of cyclobutadiene. *Chemical Physics Letters*, 501(4):166–171, 2011. ISSN 0009-2614.

- [160] F. Fantuzzi, T. M. Cardozo, and M. A. C. Nascimento. The Nature of the Singlet and Triplet States of Cyclobutadiene as Revealed by Quantum Interference. *ChemPhysChem*, 17(2):288–295, 2016.
- [161] P. C. Varras and P. S. Gritzapis. The transition state of the automerization reaction of cyclobutadiene: A theoretical approach using the Restricted Active Space Self Consistent Field method. *Chemical Physics Letters*, 711:166–172, 2018. ISSN 0009-2614.
- [162] J. E. T. Smith, B. Mussard, A. A. Holmes, and S. Sharma. Cheap and Near Exact CASSCF with Large Active Spaces. *Journal of Chemical Theory and Computation*, 13(11):5468–5478, 2017.
- [163] D. A. Mazziotti. Two-electron reduced density matrices from the anti-Hermitian contracted Schrödinger equation: Enhanced energies and properties with larger basis sets. *The Journal of Chemical Physics*, 126(18):184101, 2007.
- [164] A. E. DePrince, E. Kamarchik, and D. A. Mazziotti. Parametric two-electron reduced-density-matrix method applied to computing molecular energies and properties at nonequilibrium geometries. *The Journal of Chemical Physics*, 128(23):234103, 2008.
- [165] D. A. Mazziotti. Parametrization of the two-electron reduced density matrix for its direct calculation without the many-electron wave function: Generalizations and applications. *Phys. Rev. A*, 81:062515, Jun 2010.
- [166] J. W. Mullinax, E. Epifanovsky, G. Gidofalvi, and A. E. DePrince. Analytic Energy Gradients for Variational Two-Electron Reduced-Density Matrix Methods within the Density Fitting Approximation. *Journal of Chemical Theory and Computation*, 15(1):276–289, 2019.
- [167] R. M. Richard, K. U. Lao, and J. M. Herbert. Understanding the many-body expansion for large systems. I. Precision considerations. *The Journal of Chemical Physics*, 141(1):014108, 2014.
- [168] K. U. Lao, K.-Y. Liu, R. M. Richard, and J. M. Herbert. Understanding the many-body expansion for large systems. II. Accuracy considerations. *The Journal of Chemical Physics*, 144(16):164105, 2016.
- [169] K.-Y. Liu and J. M. Herbert. Understanding the many-body expansion for large systems. III. Critical role of four-body terms, counterpoise corrections, and cutoffs. *The Journal of Chemical Physics*, 147(16):161729, 2017.
- [170] J. M. Herbert. Fantasy versus reality in fragment-based quantum chemistry. *The Journal of Chemical Physics*, 151(17):170901, 2019.
- [171] K.-Y. Liu and J. M. Herbert. Energy-Screened Many-Body Expansion: A Practical Yet Accurate Fragmentation Method for Quantum Chemistry. *Journal of Chemical Theory and Computation*, 16(1):475–487, 2020.

- [172] K. Kitaura, S.-I. Sugiki, T. Nakano, Y. Komeiji, and M. Uebayasi. Fragment molecular orbital method: analytical energy gradients. *Chemical Physics Letters*, 336(1):163–170, 2001. ISSN 0009-2614.
- [173] D. G. Fedorov, T. Ishida, M. Uebayasi, and K. Kitaura. The Fragment Molecular Orbital Method for Geometry Optimizations of Polypeptides and Proteins. *The Journal of Physical Chemistry A*, 111(14):2722–2732, 2007.
- [174] T. Nagata, D. G. Fedorov, and K. Kitaura. Derivatives of the approximated electrostatic potentials in the fragment molecular orbital method. *Chemical Physics Letters*, 475(1):124–131, 2009. ISSN 0009-2614.
- [175] T. Nagata, D. G. Fedorov, and K. Kitaura. Importance of the hybrid orbital operator derivative term for the energy gradient in the fragment molecular orbital method. *Chemical Physics Letters*, 492(4):302–308, 2010. ISSN 0009-2614.
- [176] T. Nagata, K. Brorsen, D. G. Fedorov, K. Kitaura, and M. S. Gordon. Fully analytic energy gradient in the fragment molecular orbital method. *The Journal of Chemical Physics*, 134(12):124115, 2011.
- [177] P. Pulay. Analytical derivatives, forces, force constants, molecular geometries, and related response properties in electronic structure theory. *WIREs Computational Molecular Science*, 4(3):169–181, 2014.
- [178] S. Keller, K. Boguslawski, T. Janowski, M. Reiher, and P. Pulay. Selection of active spaces for multiconfigurational wavefunctions. *The Journal of Chemical Physics*, 142(24):244104, 2015.
- [179] V. Guner, K. S. Khuong, A. G. Leach, P. S. Lee, M. D. Bartberger, and K. N. Houk. A Standard Set of Pericyclic Reactions of Hydrocarbons for the Benchmarking of Computational Methods: The Performance of ab Initio, Density Functional, CASSCF, CASPT2, and CBS-QB3 Methods for the Prediction of Activation Barriers, Reaction Energetics, and Transition State Geometries. *The Journal of Physical Chemistry A*, 107(51):11445–11459, 2003.
- [180] F. Fracchia, R. Cimiraglia, and C. Angeli. Assessment of Multireference Perturbation Methods for Chemical Reaction Barrier Heights. *The Journal of Physical Chemistry A*, 119(21):5490–5495, 2015.
- [181] C. W. Bauschlicher and S. R. Langhoff. Full configuration-interaction study of the ionic–neutral curve crossing in LiF. *The Journal of Chemical Physics*, 89(7):4246–4254, 1988.
- [182] A. Sanchez de Meras, M.-B. Lepetit, and J.-P. Malrieu. Discontinuity of valence CASSCF wave functions around weakly avoided crossing between valence configurations. *Chemical Physics Letters*, 172(2):163–168, 1990. ISSN 0009-2614.

- [183] W. J. Glover. Communication: Smoothing out excited-state dynamics: Analytical gradients for dynamically weighted complete active space self-consistent field. *The Journal of Chemical Physics*, 141(17):171102, 2014.
- [184] J. Cullen. Generalized valence bond solutions from a constrained coupled cluster method. *Chemical Physics*, 202(2):217–229, 1996. ISSN 0301-0104.
- [185] J. Gerratt, D. L. Cooper, P. B. Karadakov, and M. Raimondi. Modern valence bond theory. *Chem. Soc. Rev.*, 26:87–100, 1997.
- [186] T. Van Voorhis and M. Head-Gordon. Connections between coupled cluster and generalized valence bond theories. *The Journal of Chemical Physics*, 115(17):7814–7821, 2001.
- [187] G. J. O. Beran, B. Austin, A. Sodt, and M. Head-Gordon. Unrestricted Perfect Pairing: The Simplest Wave-Function-Based Model Chemistry beyond Mean Field. *The Journal of Physical Chemistry A*, 109(40):9183–9192, 2005.
- [188] K. V. Lawler, D. W. Small, and M. Head-Gordon. Orbitals That Are Unrestricted in Active Pairs for Generalized Valence Bond Coupled Cluster Methods. *The Journal of Physical Chemistry A*, 114(8):2930–2938, 2010.
- [189] A. I. Krylov, C. D. Sherrill, E. F. C. Byrd, and M. Head-Gordon. Size-consistent wave functions for nondynamical correlation energy: The valence active space optimized orbital coupled-cluster doubles model. *The Journal of Chemical Physics*, 109(24):10669–10678, 1998.
- [190] A. Köhn and J. Olsen. Orbital-optimized coupled-cluster theory does not reproduce the full configuration-interaction limit. *The Journal of Chemical Physics*, 122(8):084116, 2005.
- [191] J. A. Parkhill, K. Lawler, and M. Head-Gordon. The perfect quadruples model for electron correlation in a valence active space. *The Journal of Chemical Physics*, 130(8):084101, 2009.
- [192] Y. Shao, Z. Gan, E. Epifanovsky, A. T. B. Gilbert, M. Wormit, J. Kussmann, A. W. Lange, A. Behn, J. Deng, X. Feng, D. Ghosh, M. Goldey, P. R. Horn, L. D. Jacobson, I. Kaliman, R. Z. Khaliullin, T. Kúš, A. Landau, J. Liu, E. I. Proynov, Y. M. Rhee, R. M. Richard, M. A. Rohrdanz, R. P. Steele, E. J. Sundstrom, H. L. Woodcock III, P. M. Zimmerman, D. Zuev, B. Albrecht, E. Alguire, B. Austin, G. J. O. Beran, Y. A. Bernard, E. Berquist, K. Brandhorst, K. B. Bravaya, S. T. Brown, D. Casanova, C.-M. Chang, Y. Chen, S. H. Chien, K. D. Closser, D. L. Crittenden, M. Diedenhofen, R. A. DiStasio Jr., H. Dop, A. D. Dutoi, R. G. Edgar, S. Fatehi, L. Fusti-Molnar, A. Ghysels, A. Golubeva-Zadorozhnaya, J. Gomes, M. W. D. Hanson-Heine, P. H. P. Harbach, A. W. Hauser, E. G. Hohenstein, Z. C. Holden, T.-C. Jagau, H. Ji, B. Kaduk, K. Khistyayev, J. Kim, J. Kim, R. A. King, P. Klunzinger, D. Kosenkov, T. Kowalczyk, C. M. Krauter, K. U. Lao, A. Laurent, K. V. Lawler, S. V. Levchenko, C. Y. Lin, F. Liu, E. Livshits, R. C. Lochan, A. Luenser, P. Manohar, S. F. Manzer, S.-P. Mao, N. Mardirossian, A. V.

- Marenich, S. A. Maurer, N. J. Mayhall, C. M. Oana, R. Olivares-Amaya, D. P. O'Neill, J. A. Parkhill, T. M. Perrine, R. Peverati, P. A. Pieniazek, A. Prociuk, D. R. Rehn, E. Rosta, N. J. Russ, N. Sergueev, S. M. Sharada, S. Sharma, D. W. Small, A. Sodt, T. Stein, D. Stück, Y.-C. Su, A. J. W. Thom, T. Tsuchimochi, L. Vogt, O. Vydrov, T. Wang, M. A. Watson, J. Wenzel, A. White, C. F. Williams, V. Vanovschi, S. Yeganeh, S. R. Yost, Z.-Q. You, I. Y. Zhang, X. Zhang, Y. Zhou, B. R. Brooks, G. K. L. Chan, D. M. Chipman, C. J. Cramer, W. A. Goddard III, M. S. Gordon, W. J. Hehre, A. Klamt, H. F. Schaefer III, M. W. Schmidt, C. D. Sherrill, D. G. Truhlar, A. Warshel, X. Xua, A. Aspuru-Guzik, R. Baer, A. T. Bell, N. A. Besley, J.-D. Chai, A. Dreuw, B. D. Dunietz, T. R. Furlani, S. R. Gwaltney, C.-P. Hsu, Y. Jung, J. Kong, D. S. Lambrecht, W. Liang, C. Ochsenfeld, V. A. Rassolov, L. V. Slipchenko, J. E. Subotnik, T. Van Voorhis, J. M. Herbert, A. I. Krylov, P. M. W. Gill, and M. Head-Gordon. Advances in molecular quantum chemistry contained in the Q-Chem 4 program package. *Mol. Phys.*, 113:184–215, 2015.
- [193] M. Feyereisen, G. Fitzgerald, and A. Komornicki. Use of approximate integrals in ab initio theory. An application in MP2 energy calculations. *Chemical Physics Letters*, 208(5):359–363, 1993. ISSN 0009-2614.
- [194] F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs. RI-MP2: optimized auxiliary basis sets and demonstration of efficiency. *Chemical Physics Letters*, 294(1):143–152, 1998. ISSN 0009-2614.
- [195] H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, M. Schütz, et al. MOLPRO, version 2012.1, a package of ab initio programs, 2012.
- [196] M. Palucki, N. S. Finney, P. J. Pospisil, M. L. Güler, T. Ishida, and E. N. Jacobsen. The Mechanistic Basis for Electronic Effects on Enantioselectivity in the (salen)Mn(III)-Catalyzed Epoxidation Reaction. *Journal of the American Chemical Society*, 120(5): 948–954, 1998.
- [197] E. M. McGarrigle and D. G. Gilheany. Chromium and Manganese-salen Promoted Epoxidation of Alkenes. *Chemical Reviews*, 105(5):1563–1602, 2005.
- [198] C. Linde, B. Åkermark, P.-O. Norrby, and M. Svensson. Timing Is Critical: Effect of Spin Changes on the Diastereoselectivity in Mn(salen)-Catalyzed Epoxidation. *Journal of the American Chemical Society*, 121(21):5083–5084, 1999.
- [199] Y. G. Abashkin, J. R. Collins, and S. K. Burt. (Salen)Mn(III)-Catalyzed Epoxidation Reaction as a Multichannel Process with Different Spin States. Electronic Tuning of Asymmetric Catalysis: A Theoretical Study. *Inorganic Chemistry*, 40(16):4040–4048, 2001.
- [200] J. Ivanic, J. R. Collins, and S. K. Burt. Theoretical Study of the Low Lying Electronic States of oxoX(salen) (X = Mn, Mn-, Fe, and Cr-) Complexes. *The Journal of Physical Chemistry A*, 108(12):2314–2323, 2004.

- [201] J. S. Sears and C. D. Sherrill. The electronic structure of oxo-Mn(salen): Single-reference and multireference approaches. *The Journal of Chemical Physics*, 124(14):144314, 2006.
- [202] S. Wouters, T. Bogaerts, P. Van Der Voort, V. Van Speybroeck, and D. Van Neck. Communication: DMRG-SCF study of the singlet, triplet, and quintet states of oxo-Mn(Salen). *The Journal of Chemical Physics*, 140(24):241103, 2014.
- [203] C. J. Stein and M. Reiher. Automated Selection of Active Orbital Spaces. *Journal of Chemical Theory and Computation*, 12(4):1760–1771, 2016.
- [204] S. Sharma, G. Knizia, S. Guo, and A. Alavi. Combining Internally Contracted States and Matrix Product States To Perform Multireference Perturbation Theory. *Journal of Chemical Theory and Computation*, 13(2):488–498, 2017.
- [205] P. Zimmerman. Reliable Transition State Searches Integrated with the Growing String Method. *Journal of Chemical Theory and Computation*, 9(7):3043–3050, 2013.
- [206] P. M. Zimmerman. Growing string method with interpolation and optimization in internal coordinates: Method and examples. *The Journal of Chemical Physics*, 138(18):184102, 2013.
- [207] P. M. Zimmerman. Single-ended transition state finding with the growing string method. *Journal of Computational Chemistry*, 36(9):601–611, 2015.
- [208] C. Aldaz, J. A. Kammeraad, and P. M. Zimmerman. Discovery of conical intersection mediated photochemistry with growing string methods. *Phys. Chem. Chem. Phys.*, 20:27394–27405, 2018.
- [209] P. Pinski and F. Neese. Analytical gradient for the domain-based local pair natural orbital second order Møller-Plesset perturbation theory method (DLPNO-MP2). *The Journal of Chemical Physics*, 150(16):164102, 2019.
- [210] B. Q. Pham and M. S. Gordon. Development of the FMO/RI-MP2 Fully Analytic Gradient Using a Hybrid-Distributed/Shared Memory Programming Model. *Journal of Chemical Theory and Computation*, 16(2):1039–1054, 2020.
- [211] M. Izakovič, J. Šima, and M. Žitňanský. Effects of wavelength and substituents on iron(III) photoreduction in trans-[Fe(R-sal-R-en)(CH₃OH)(NCS)] complexes. *Journal of Coordination Chemistry*, 58(12):1039–1046, 2005.
- [212] L. Rigamonti, F. Demartin, A. Forni, S. Righetto, and A. Pasini. Copper(II) Complexes of salen Analogues with Two Differently Substituted (PushPull) Salicylaldehyde Moieties. A Study on the Modulation of Electronic Asymmetry and Nonlinear Optical Properties. *Inorganic Chemistry*, 45(26):10976–10989, 2006.
- [213] A. Trujillo, M. Fuentealba, D. Carrillo, C. Manzur, I. Ledoux-Rak, J.-R. Hamon, and J.-Y. Saillard. Synthesis, Spectral, Structural, Second-Order Nonlinear Optical Properties and Theoretical Studies On New Organometallic DonorAcceptor Substituted

- Nickel(II) and Copper(II) Unsymmetrical Schiff-Base Complexes. *Inorganic Chemistry*, 49(6):2750–2764, 2010.
- [214] S. Aono, M. Nakagaki, and S. Sakaki. Theoretical study of one-electron-oxidized salen complexes of group 7 (Mn(iii), Tc(iii), and Re(iii)) and group 10 metals (Ni(ii), Pd(ii), and Pt(ii)) with the 3D-RISM-GMC-QDPT method: localized vs. delocalized ground and excited states in solution. *Phys. Chem. Chem. Phys.*, 19:16831–16849, 2017.
- [215] I. B. Bersuker. Modern Aspects of the Jahn-Teller Effect Theory and Applications To Molecular Problems. *Chemical Reviews*, 101(4):1067–1114, 2001.
- [216] A. J. Johansson. Teaching the Jahn–Teller Theorem: A Simple Exercise That Illustrates How the Magnitude of Distortion Depends on the Number of Electrons and Their Occupation of the Degenerate Energy Level. *Journal of Chemical Education*, 90(1): 63–69, 2013.
- [217] D. W. Whitman and B. K. Carpenter. Limits on the activation parameters for automerization of cyclobutadiene-1,2-d2. *Journal of the American Chemical Society*, 104(23): 6473–6474, 1982.
- [218] D.-K. Dang, L. W. Wilson, and P. M. Zimmerman. The numerical evaluation of Slater integrals on graphics processing units. *Journal of Computational Chemistry*, 43(25): 1680–1689, 2022.
- [219] T. Kato. On the eigenfunctions of many-particle systems in quantum mechanics. *Communications on Pure and Applied Mathematics*, 10(2):151–177, 1957.
- [220] P. Reinhardt and P. E. Hoggan. Cusps and derivatives for wave-functions expanded in Slater orbitals: A density study. *International Journal of Quantum Chemistry*, 109(14):3191–3198, 2009.
- [221] L. E. McMurchie and E. R. Davidson. One- and two-electron integrals over cartesian gaussian functions. *Journal of Computational Physics*, 26(2):218–231, 1978. ISSN 0021-9991.
- [222] S. Obara and A. Saika. Efficient recursive computation of molecular integrals over Cartesian Gaussian functions. *The Journal of Chemical Physics*, 84(7):3963–3974, 1986.
- [223] S. Obara and A. Saika. General recurrence formulas for molecular integrals over Cartesian Gaussian functions. *The Journal of Chemical Physics*, 89(3):1540–1559, 1988.
- [224] M. Head-Gordon and J. A. Pople. A method for two-electron Gaussian integral and integral derivative evaluation using recurrence relations. *The Journal of Chemical Physics*, 89(9):5777–5786, 1988.
- [225] P. M. W. Gill, M. Head-Gordon, and J. A. Pople. Efficient computation of two-electron - repulsion integrals and their nth-order derivatives using contracted Gaussian basis sets. *The Journal of Physical Chemistry*, 94(14):5564–5572, 1990.

- [226] W. J. Hehre, R. F. Stewart, and J. A. Pople. Self-Consistent Molecular-Orbital Methods. I. Use of Gaussian Expansions of Slater-Type Atomic Orbitals. *The Journal of Chemical Physics*, 51(6):2657–2664, 1969.
- [227] W. J. Hehre, R. Ditchfield, R. F. Stewart, and J. A. Pople. Self-Consistent Molecular Orbital Methods. IV. Use of Gaussian Expansions of Slater-Type Orbitals. Extension to Second-Row Molecules. *The Journal of Chemical Physics*, 52(5):2769–2773, 1970.
- [228] J. Fernández Rico, R. López, A. Aguado, I. Ema, and G. Ramírez. New program for molecular calculations with Slater-type orbitals. *International Journal of Quantum Chemistry*, 81(2):148–153, 2001.
- [229] M. Caffarel. Evaluating two-electron-repulsion integrals over arbitrary orbitals using zero variance Monte Carlo: Application to full configuration interaction calculations with Slater-type orbitals. *The Journal of Chemical Physics*, 151(6):064101, 2019.
- [230] M. Nightingale and C. Umrigar. *Quantum Monte Carlo methods in physics and chemistry*. Kluwer Academic Publishers, 1999.
- [231] S. Zhang and H. Krakauer. Quantum Monte Carlo Method using Phase-Free Random Walks with Slater Determinants. *Phys. Rev. Lett.*, 90:136401, Apr 2003.
- [232] P. M. Zimmerman, J. Toulouse, Z. Zhang, C. B. Musgrave, and C. J. Umrigar. Excited states of methylene from quantum Monte Carlo. *The Journal of Chemical Physics*, 131(12):124103, 2009.
- [233] B. M. Austin, D. Y. Zubarev, and W. A. Lester. Quantum Monte Carlo and Related Approaches. *Chemical Reviews*, 112(1):263–288, 2012.
- [234] W. Koch and M. C. Holthausen. *A Chemist’s Guide to Density Functional Theory*. Wiley - VCH, Weinheim - New York, 2nd edition, 2001.
- [235] M. A. Watson, N. C. Handy, and A. J. Cohen. Density functional calculations, using Slater basis sets, with exact exchange. *The Journal of Chemical Physics*, 119(13):6475–6481, 2003.
- [236] A. Förster, M. Franchini, E. van Lenthe, and L. Visscher. A Quadratic Pair Atomic Resolution of the Identity Based SOS-AO-MP2 Algorithm Using Slater Type Orbitals. *Journal of Chemical Theory and Computation*, 16(2):875–891, 2020.
- [237] A. Förster and L. Visscher. Double hybrid DFT calculations with Slater type orbitals. *Journal of Computational Chemistry*, 41(18):1660–1684, 2020.
- [238] A. Förster and L. Visscher. Low-Order Scaling G0W0 by Pair Atomic Density Fitting. *Journal of Chemical Theory and Computation*, 16(12):7381–7399, 2020.
- [239] A. Förster and L. Visscher. GW100: A Slater-Type Orbital Perspective. *Journal of Chemical Theory and Computation*, 0(0):null, 0.

- [240] J. F. Stanton and R. J. Bartlett. The equation of motion coupled-cluster method. A systematic biorthogonal approach to molecular excitation energies, transition probabilities, and excited state properties. *The Journal of Chemical Physics*, 98(9):7029–7039, 1993.
- [241] F. A. Evangelista, E. Prochnow, J. Gauss, and H. F. Schaefer. Perturbative triples corrections in state-specific multireference coupled cluster theory. *The Journal of Chemical Physics*, 132(7):074107, 2010.
- [242] F. A. Evangelista and J. Gauss. An orbital-invariant internally contracted multireference coupled cluster approach. *The Journal of Chemical Physics*, 134(11):114102, 2011.
- [243] SlaterGPU. <https://github.com/ZimmermanGroup/SlaterGPU>.
- [244] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin. On some approximations in applications of χ_α theory. *The Journal of Chemical Physics*, 71(8):3396–3402, 1979.
- [245] H.-J. Werner, F. R. Manby, and P. J. Knowles. Fast linear scaling second-order Møller-Plesset perturbation theory (MP2) using local and density fitting approximations. *The Journal of Chemical Physics*, 118(18):8149–8160, 2003.
- [246] R. A. Distasio JR., R. P. Steele, Y. M. Rhee, Y. Shao, and M. Head-Gordon. An improved algorithm for analytical gradient evaluation in resolution-of-the-identity second-order Møller-Plesset perturbation theory: Application to alanine tetrapeptide conformational analysis. *Journal of Computational Chemistry*, 28(5):839–856, 2007.
- [247] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, USA, 3 edition, 2007. ISBN 0521880688.
- [248] P. M. Boerrigter, G. Te Velde, and J. E. Baerends. Three-dimensional numerical integration for electronic structure calculations. *International Journal of Quantum Chemistry*, 33(2):87–113, 1988.
- [249] A. D. Becke. A multicenter numerical integration scheme for polyatomic molecules. *The Journal of Chemical Physics*, 88(4):2547–2553, 1988.
- [250] C. W. Murray, N. C. Handy, and G. J. Laming. Quadrature schemes for integrals of density functional theory. *Molecular Physics*, 78(4):997–1014, 1993.
- [251] O. Treutler and R. Ahlrichs. Efficient molecular numerical integration schemes. *The Journal of Chemical Physics*, 102(1):346–354, 1995.
- [252] M. E. Mura and P. J. Knowles. Improved radial grids for quadrature in molecular density-functional calculations. *The Journal of Chemical Physics*, 104(24):9848–9858, 1996.
- [253] V. Lebedev. Quadratures on a sphere. *USSR Computational Mathematics and Mathematical Physics*, 16(2):10–24, 1976. ISSN 0041-5553.

- [254] Implementations of the gamma functions. <http://www.netlib.org/cephes/>.
- [255] I. S. Ufimtsev and T. J. Martínez. Quantum Chemistry on Graphical Processing Units. 2. Direct Self-Consistent-Field Implementation. *Journal of Chemical Theory and Computation*, 5(4):1004–1015, 2009.
- [256] N. Luehr, I. S. Ufimtsev, and T. J. Martínez. Dynamic Precision for Electron Repulsion Integral Evaluation on Graphical Processing Units (GPUs). *Journal of Chemical Theory and Computation*, 7(4):949–954, 2011.
- [257] P. Pokhilko, E. Epifanovsky, and A. I. Krylov. Double Precision Is Not Needed for Many-Body Calculations: Emergent Conventional Wisdom. *Journal of Chemical Theory and Computation*, 14(8):4088–4096, 2018.
- [258] S. Seritan, C. Bannwarth, B. S. Fales, E. G. Hohenstein, S. I. L. Kokkila-Schumacher, N. Luehr, J. W. Snyder, C. Song, A. V. Titov, I. S. Ufimtsev, and T. J. Martínez. TeraChem: Accelerating electronic structure and ab initio molecular dynamics with graphical processing units. *The Journal of Chemical Physics*, 152(22):224110, 2020.
- [259] E. Van Lenthe and E. J. Baerends. Optimized Slater-type basis sets for the elements 1–118. *Journal of Computational Chemistry*, 24(9):1142–1156, 2003.
- [260] Q. Sun. Libcint: An efficient general integral library for Gaussian basis functions. *Journal of Computational Chemistry*, 36(22):1664–1671, 2015.
- [261] P. M. Gill, B. G. Johnson, and J. A. Pople. A standard grid for density functional calculations. *Chemical Physics Letters*, 209(5):506–512, 1993. ISSN 0009-2614.
- [262] M. Khalilov and A. Timoveev. Performance analysis of CUDA, OpenACC and OpenMP programming models on TESLA V100 GPU. *Journal of Physics: Conference Series*, 1740:012056, jan 2021.
- [263] F. J. Lovas and E. Tiemann. Diatomic Spectral Database. Technical Report NIST Standard Reference Database 114, National Institute of Standards and Technology, 2005.