

©Copyright 2026

Agam Shayit

Numerically Exact Configuration Interaction at
Quadrillion-Determinant Scale

Agam Shayit

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Xiaosong Li, Chair

Silas Beane

Mark Rudner

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Numerically Exact Configuration Interaction at Quadrillion-Determinant Scale

Agam Shayit

Chair of the Supervisory Committee:

Xiaosong Li

Department of Chemistry

Quantum chemistry is the field of study that uses quantum mechanics to make chemical predictions. The chemistry of the systems of interest (atoms and molecules) is faithfully described by the electronic Dirac equation. The equation describes the relativistic dynamics of the electrons in a molecule subject to electric forces from the nuclei and other electrons. For virtually all systems of practical interest, the equation is impossible to solve by means of exact analytical techniques. Thus, the use of numerical methods becomes inevitable. The most accurate such method is configuration interaction (CI), which yields the exact formal solution of the Dirac equation in the complete-basis-set limit. Due to the factorial growth in the resource requirements of CI calculations, they have historically been applied to the smallest of chemical systems. In this dissertation, we introduce categorical compression of CI vectors within the small tensor product distributed active space (STP-DAS) CI framework. We demonstrate its capabilities by conducting a CI calculation consisting of over one quadrillion determinants, the largest CI calculation to date. We then explore strategies to accelerate STP-DAS CI calculations by adapting the STP-DAS σ -build step to graphics-processing units (GPUs). Finally, we use the developed framework to conduct previously untenable benchmark studies of two state-of-the-art quantum chemical methods, namely coupled cluster (CC) and density matrix renormalization group (DMRG).

TABLE OF CONTENTS

	Page
List of Figures	iii
Glossary	vi
Chapter 1: Introduction	1
1.1 Quantum Mechanics	1
1.2 Relativity	4
1.3 Relativistic Quantum Mechanics	6
1.4 Relativistic Electronic Structure Theory	9
1.5 Spectral Methods: Electronic Structure Theory in a Basis	11
1.6 The Mean-Field Solution	15
1.7 Configuration Interaction and Electronic Correlation	17
Chapter 2: Review of the STP-DAS Framework	22
2.1 Introduction	22
2.2 Distributed Active Space Configuration Interaction	25
2.3 Computational Details	37
2.4 Results and Discussion	39
2.5 Conclusions	50
Chapter 3: Categorical Compression within the STP-DAS Framework	52
3.1 Introduction	52
3.2 Results	56
3.3 Methods	65
3.4 Conclusions	70
Chapter 4: Accelerating Small Tensor Product Distributed Active Space (STP-DAS) Configuration Interaction using Graphics Processing Units (GPUs)	73
4.1 Introduction	73
4.2 Methods	76

4.3	Results	78
4.4	Conclusions	80
Chapter 5:	Benchmarking Coupled Cluster and Density Matrix Renormalization Group: Accuracy and Convergence Trends Against Large Scale Nu- merically Exact Configuration Interaction	83
5.1	Introduction	83
5.2	Results	89
5.3	Conclusions	91
Chapter 6:	Conclusions and Outlook	93
Bibliography	95
Appendix A:	Supplementary information to Categorical Compression within the STP-DAS Framework	122
A.1	Convergence Behavior of Compression-Compatible Preconditioning	122
A.2	Cost Analysis	124
A.3	Exploiting Symmetry	124

LIST OF FIGURES

Figure Number	Page
2.1	Space Partitioning in STP-DAS. Boxes symbolize active regions. Within these boxes, shaded portions indicate electron occupations, while white areas represent empty orbitals. A CAS problem is mapped onto multiple DASs. Distinct categories of configurations are produced by categorical excitations, depicted by green arrows. In every category, each DAS is treated as a complete active space, including all possible excitations, illustrated by blue arrows. 28
2.2	A comparison between a traditional indexing scheme and the indexing scheme used in STP-DAS. In DAS, the global address of a determinant is found by combining a global offset of the category of a given determinant with the local index of that determinant within the category. Boxes symbolize determinant categories. 32
2.3	Examples of excitations between categories. Boxes symbolize active regions. Within these boxes, shaded portions indicate electron occupations, while white areas represent empty orbitals. All partially filled DASs have an identical number of electrons and holes. 35
2.4	Distribution of determinants per node for a $\binom{44}{29}$ X2C-DASCI calculation using 600 nodes with (top) 8 DAS spaces and (bottom) 9 DAS spaces. The theoretical ideal distribution of determinants is denoted by the red line in both plots, and the solid line denotes the median of the observed distributions. 42
2.5	Execution time of a single σ build using the STP-DAS framework for the TIH test case with increasing CAS space sizes with respect to the number of nodes (1 MPI per node, 20 SMP threads per MPI). 45
2.6	Relative speedup of a single σ build for TIH (top) $\binom{37}{24}$ and (bottom) $\binom{38}{24}$ with respect to the number of nodes (1 MPI per node, 20 SMP threads per MPI). Relative speed-up is defined with respect to the performance on the smallest node count capable of solving the problem. 46
2.7	Execution time of a single σ build using the STP-DAS framework for the TIH test case with increasing CAS space sizes (Perlmutter 256 Nodes, 1 MPI per node, 64 SMP threads per MPI). 48
2.8	(top) σ build time (seconds) and (bottom) the distribution of the MPI communication idle time of each process during the σ build process of the TIH test case with a increasing numbers of correlated orbitals (Perlmutter 256 Nodes, 1 MPI per node, 64 SMP threads per MPI). 49

3.1	The evolution of the state-of-the-art for CI calculations over time [8, 93, 95, 122, 211, 302]. Each historical point is colored according to the nature of the key development of the respective work, which we classify as either intrinsic algorithmic developments (in purple) or optimizations on the HPC hardware of the time (in blue). * Only one CI iteration was performed.	53
3.2	Categorical Compression within the small-tensor-product distributed active space (STP-DAS) framework. (A), The STP-DAS framework decomposition of a complete active space configuration interaction (CASCI) calculation into a direct sum of categorical excitations. The large excitation lists can be factored into much smaller categorical excitation lists. Purple sections within active spaces represent electron-occupied orbitals. (B), The exact two-component full configuration interaction (X2C-FCI) ground state energy of the Mg^{2+} ion within the cc-pVNZ-DK [53, 229] ($N = 2, 3, 4$) basis sets, along with the extrapolated complete basis set limit. (C), Average execution time (in seconds) of the compression-compatible STP-DAS algorithm per Davidson iteration of a thallium hydride (TIH) test case versus the node count (5 iterations, 1 message passing interface (MPI) process per node, 40 symmetric multiprocessing (SMP) threads per MPI process). Here, \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, and $\boldsymbol{\sigma}$ is their product. The dashed lines illustrate the ideal strong scaling behavior of each CASCI calculation. (D), The representation of the subspace expansion vector in a traditional configuration interaction (CI) picture, the decomposed subspace vector in the STP-DAS framework, and the numerically exact compression of the subspace expansion vector in the categorically compression-compatible STP-DAS representation. The color of CI coefficients indicates their configuration category, while their brightness symbolizes their magnitude. White indicates a magnitude of zero. (E), A schematic representation of the lossless, compression-compatible, STP-DAS $\boldsymbol{\sigma}$ -build algorithm. The Hamiltonian matrix is represented as a heatmap, where brighter elements have larger magnitudes. The color of the vector elements indicates the configuration category of the corresponding CI coefficients. Note that the $\boldsymbol{\sigma}$ -build preserves categorical compression. (F), An illustration comparing the traditional Davidson preconditioner with the compression-compatible preconditioner to generate successive subspace expansion vectors. The compression-compatible preconditioner appends the subspace with the same effective search direction as the traditional Davidson preconditioner without compromising its compression.	72

4.1	(Red: 8 CPU cores + 8 GPUs) Marginal runtime of the GPU portion of Algorithm 2 (matrix multiplications and data transfers) for various batch sizes of the fourth σ -build iteration for an Au ₂ test case using an active space of 48 orbitals and 8 electrons within the x2c-TZVPall-2c [227] basis set. (Blue: 8 CPU cores) The reference runtime of an identical calculation with a CPU-only version of Algorithm 2 [122]. A preconditioning dropping threshold of $\varepsilon = 10^{-16}$ was used in all calculations. Both axes are log-scaled for convenience.	80
4.2	(Red) Marginal speedup of the GPU portion of Algorithm 2 (matrix multiplications and data transfers) as a function of the number of GPUs used. (Blue) The corresponding ideal strong-scaling behavior. The data used the fourth σ -build iteration of an Xe ₂ test in an active space of 48 orbitals and 8 electrons within the x2c-TZVPall-2c [227] basis set. A preconditioning dropping threshold of $\varepsilon = 10^{-16}$ was used in all calculations.	81

GLOSSARY

- : **AO:** atomic orbitals
- : **BSE:** Bethe–Salpeter equation
- : **CAS:** complete active space
- : **CASCI:** complete active space configuration interaction
- : **CASPT2:** complete active space second-order perturbation theory
- : **CASSCF:** complete active space self-consistent field
- : **CC:** coupled cluster
- : **CCSD:** coupled cluster singles and doubles
- : **CCSD(T):** coupled cluster singles and doubles with perturbative triples
- : **CCSDT:** coupled cluster singles, doubles and triples
- : **CI:** configuration interaction
- : **CISD:** configuration interaction - singles and doubles
- : **CISDT:** configuration interaction - singles, doubles and triples
- : **CPU:** central processing unit
- : **DAS:** distributed active space

- : **DASCI:** distributed active space configuration interaction
- : **DFT:** density-functional theory
- : **DKH:** Douglas-Kroll-Hess
- : **DMRG:** density matrix renormalization group
- : **FLOP:** floating-point operation
- : **GAS:** generalized active space
- : **GPU:** graphics-processing unit
- : **GTO:** Gaussian-type orbital
- : **MCPDFT:** multiconfiguration pair-density functional theory
- : **MCSCF:** multi-configurational self-consistent field
- : **MO:** molecular orbital
- : **MPI:** message passing interface
- : **MRCI:** multi-reference configuration interaction
- : **MRPT2:** multi-reference second-order perturbation theory
- : **NEVPT2:** *N*-electron valence state second-order perturbation theory
- : **OpenMP:** open multi-processing
- : **ORMAS:** occupation restricted multiple active space

- : **QMC:** quantum Monte Carlo
- : **RAM:** random-access memory
- : **RAS:** restricted active space
- : **RMA:** remote memory access
- : **SIMD:** single instruction, multiple data
- : **STO:** Slater-type orbital
- : **VRAM:** video random-access memory
- : **X2C:** exact two-component
- : **1eX2C:** one-electron exact two-component

ACKNOWLEDGMENTS

The accomplishments reported in this dissertation would not be possible without the immense courage of my mother to uproot her and her family's lives and move to the United States, the unshakable world leader in scientific discovery and research. I owe tremendous gratitude to her and the sacrifices she made along the way.

I owe my success to Prof. Xiaosong Li, my Ph.D. supervisor, who took me into his research group in the second year of my Ph.D. studies. Through his guidance, I was introduced to the realm of electronic structure theory and method development, which I ultimately fell in love with. Throughout my time in the Li group, Xiaosong materially supported me with a research associateship, thereby allowing me to dedicate my utmost attention and productivity to research projects. Frankly, I would not have reached the milestones reported here without his supervisory and financial support.

I am grateful to fellow members of the Li group for teaching me almost everything I know about electronic structure theory. Coming from a computer science and physics background, I lacked some of the chemistry domain knowledge needed to penetrate the field. Special thanks go to Shiv Upadhyay for his close mentorship throughout my Ph.D. studies and to Mikael Kovtun, a fellow physics Ph.D. student, for sharing his understanding of the field from a physical perspective.

Most importantly, I am extremely grateful to my wife, Krista Anderson, for supporting me and us throughout my time in graduate school. The financial burden of life overwhelmingly fell on her for years as I tinkered with equations, implemented them efficiently in C++ code, and wrote manuscripts. My success would simply not be possible without her presence in my life, and as my graduate studies come to an end, it is now her rightful turn to be on the other side of the financial seesaw.

DEDICATION

To my wonderful wife and best friend, Krista Anderson, for supporting me in every possible aspect throughout my graduate studies.

To my father, Ran Shayit, who has always guided, encouraged, and materially facilitated my academic curiosity.

And to my mother, Sharon Bach, who, in 2016, made the life-changing decision to uproot and relocate our family to the United States of America in pursuit of a better life.

ACKNOWLEDGMENTS TO PREVIOUS WORK

Chapter 2 and parts of Section 1.7 are adapted with permission from H. Hu, S. Upadhyay, L. Lu, A. J. Jenkins, T. Zhang, A. Shayit, S. Knecht, and X. Li. Small Tensor Product Distributed Active Space (STP-DAS) Framework for Relativistic and Non-relativistic Multiconfiguration Calculations: Scaling from 10^9 on a Laptop to 10^{12} Determinants on a Supercomputer. *Chem. Phys. Rev.*, 5(4):041404, 2024. DOI: 10.1063/5.0227122. Copyright 2024 AIP Publishing LLC.

Chapter 3 is adapted with permission from A. Shayit, C. Liao, S. Upadhyay, H. Hu, T. Zhang, A. E. DePrince III, C. Yang, and X. Li. Numerically Exact Configuration Interaction at Quadrillion-Determinant Scale. *Nat. Commun.*, 16(1):11016, 2025. DOI: 10.1038/s41467-025-65967-7. Copyright 2025 Springer Nature.

Chapter 4 is adapted with permission from A. Shayit, M. Kovtun, S. Upadhyay, H. Hu, and X. Li. Accelerating Small Tensor Product Distributed Active Space (STP-DAS) Configuration Interaction using Graphics Processing Units (GPUs). Manuscript in preparation, 2026.

Chapter 5 is adapted with permission from S. Upadhyay, A. Shayit, T. Zhang, and X. Li. Definitive Assessment of the Accuracy, Variationality, and Convergence of Relativistic Coupled Cluster and Density Matrix Renormalization Group in 100-Orbital Space. Submitted to *J. Chem. Theory Comput.*, 2026.

Chapter 1

INTRODUCTION

1.1 Quantum Mechanics

The very scientific understanding of the universe we inhabit underwent a tumultuous revolution in the eighteenth and nineteenth centuries. Before then, the physical sciences were simply a collection of seemingly unrelated empirical laws that were deduced by meticulous observationalists. A canonical example is the formulation of Kepler's laws [139], which quantify the motion of stars in orbits around the sun. After extensive analysis of precise astronomical data observed by Tycho Brahe [140], astronomer Johannes Kepler was able to effectively fit the data into a mathematical pattern, thereby arriving at the correct expressions for the elliptical orbits of planets and the time-independence of the area they sweep.

In 1687, Sir Issac Newton revolutionized the relationship between physical models and experimental observations. In his landmark work *Principia Mathematica* [206], he formulated the mathematical laws of classical mechanics, from which one can *derive* the equations of motion governing the dynamics of arbitrary objects from the net force applied on them. Classical mechanics marks a complete paradigm shift: instead of fitting specific observational data into a post-hoc model to explain specialized phenomena, it constructs a universal model out of first principles rooted in empirical observations. These axioms, Newton's three laws of motion, were then used to *derive* theoretical predictions such as Kepler's laws. This paradigm shift crowned Newton as the father of classical physics.

The remarkable success of axiomatic physical models continued well into the nineteenth century with Maxwell's theory of electrodynamics [192] and the theory of thermodynamics pioneered by Carnot [34], Clausius [40], Gibbs [98], and others. These theories exhibited remarkable predictive successes and arguably ushered in the industrial revolution with the invention of the steam engine, electrical grids, and telephone communication lines [33, 123,

198,276]. For a brief moment in time, it seemed as though the physical laws governing the universe were completely understood.

Despite its unprecedented successes, the foundations of classical physics began showing cracks by the end of the nineteenth century with the discoveries of the photoelectric effect [112], blackbody radiation [142], and Einstein's landmark explanation of Brownian motion [74]. The latter implied the existence of atoms and molecules which contradicts the theoretical foundations of classical physics. In a classical model of an atom, in which a negatively charged electron orbits a positively charged nucleus, the electron experiences a net force and is therefore accelerating, per Newton's second law of motion. One can then use Larmor's formula [164] to calculate the radiation power of the electron as a function of its acceleration as

$$P(t) = \frac{2}{3} \frac{e^2}{c^3} a(t)^2, \quad (1.1)$$

where e is the charge of the electron and c is the speed of light. This expression can then be used to obtain the equation of motion of the electron within the Newtonian framework. The prediction is the complete *collapse* of the electron into the nucleus in about 10^{-11} seconds [214]. Simply put, atoms, molecules, and therefore chemistry, do not exist in the realm of classical physics. This necessitated a complete rejection of the axioms of classical physics which had previously been accepted as absolute truths.

The first 27 years of the twentieth century saw some of the brightest minds, such as Planck [226], Einstein [73], Debye [54], Bohr [24], and de Broglie [51], try to dig physics out of this catastrophic hole. While these efforts certainly inched us closer to the understanding of the microscopic world, they did not constitute a *theory* of the same, and more closely resembled Kepler's empirical model rather than an axiomatic one. The first successful such theory was Schödinger's theory of quantum mechanics [253]. This theory describes the possible configurations of all physical systems as a complex *Hilbert space* \mathcal{H} of states. In this theory, physical observables are Hermitian *operators* on \mathcal{H} . The eigenvalues of the observables are the possible outcomes of measurements, which are made real by the Hermiticity requirement. For a given observable, \mathcal{H} includes its eigenstates, as well as arbitrary linear combinations of the same. Such linear combinations do not have definite values with

respect to the observable, and the predictive extent of the theory is the expectation value and variance of the measurement.

The dynamics of the system are described by the trajectory its quantum state takes in the Hilbert space. In accordance with classical physics, the time evolution of states is generated by the Hamiltonian \hat{H} , which is an operator representing the total energy of the system. Given an initial quantum state $|\psi\rangle \in \mathcal{H}$, its dynamics obey the Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle. \quad (1.2)$$

Besides its axiomatic approach to the dynamics of systems, Schrödinger's theory bears little resemblance to classical mechanics. Newton's theory is *deterministic*, that is, once the initial conditions (position and velocity) of the system are known, the second-order differential equation of motion can be solved to predict its exact trajectory. This way, a set of initial conditions corresponds to exactly *one* trajectory, and vice versa. This is one of the core axioms completely abandoned by quantum mechanics: a given set of initial conditions admits a Hilbert space of possible solutions, containing a (possibly infinite) set of basis states and all of their linear combinations. A full description of the trajectories can never be determined, and instead, the theory predicts statistical expectation values and moments of observables.

In the quantum theory, the role of the initial conditions is served by the initial quantum state of the system. For closed systems, those for which \hat{H} is time independent, a set of such steady states is provided by the eigenstates of the Hamiltonian (energy eigenstates). This is because the time dependence of energy eigenstates is realized as a simple multiplicative phase:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = E |\psi(t)\rangle \implies |\psi(t)\rangle = e^{-iEt/\hbar} |\psi\rangle. \quad (1.3)$$

These are steady states because for any observable A ,

$$\langle A^n \rangle (t) = \langle \psi(t) | A^n | \psi(t) \rangle = \langle \psi | A^n | \psi \rangle e^{+iEt/\hbar} e^{-iEt/\hbar} = \langle \psi | A^n | \psi \rangle, \quad (1.4)$$

rendering all expectation values, variances, and higher-order statistical moments time-

independent.

Concrete predictions can therefore be made by calculating the spectrum of the Hamiltonian, that is, the set of energy eigenstates and eigenvalues. This is done by solving the time-independent Schrödinger equation

$$\hat{H} |\psi\rangle = E |\psi\rangle. \quad (1.5)$$

The solution is commonly found by expressing Eq. (1.5) as a differential equation in position space and adopting the non-relativistic form of the Hamiltonian:

$$\begin{aligned} \langle \mathbf{x} | \hat{H} |\psi\rangle &= \int d\mathbf{x}' \langle \mathbf{x} | \hat{H} | \mathbf{x}' \rangle \langle \mathbf{x}' | \psi \rangle = \int d\mathbf{x}' \delta(\mathbf{x} - \mathbf{x}') \left(-\frac{\hbar^2}{2m} \nabla'^2 + V(\mathbf{x}') \right) \langle \mathbf{x}' | \psi \rangle \\ &= \left(-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{x}) \right) \langle \mathbf{x} | \psi \rangle = E \langle \mathbf{x} | \psi \rangle. \end{aligned} \quad (1.6)$$

A quantum theory of atoms and molecules was then obtained by substituting Coulomb's electrostatic potential [52] for $V(\mathbf{x})$. Due to its simple two-body nature, the first atomic system to be modeled and solved in this framework was the hydrogen atom [253]. The Hamiltonian was shown to admit a discrete spectrum of energy eigenvalues, meaning the steady state solutions have a restricted set of energy *quanta* they can occupy. Consequently, states cannot continuously dissipate energy, and the hydrogen atom is stable. The stability of atoms is in stark contrast with the classical theory and is one of quantum mechanics's most celebrated achievements to date.

1.2 Relativity

Around the time the quantum nature of reality was being discovered by the pioneering work of Planck, Einstein, de Broglie, and others, an orthogonal set of challenges was launched at classical mechanics from both theoretical and experimental directions.

On the theoretical front, the two main classical theories, classical mechanics and electrodynamics, were shown to be inconsistent with one another. The equation of motion of classical mechanics,

$$F(\vec{x}(t)) = m\ddot{\vec{x}}(t), \quad (1.7)$$

is invariant under the following change of coordinates, called a Galilean boost [92]:

$$\begin{cases} \vec{x} & \rightarrow \vec{x}' = \vec{x} - \vec{v}_0 t, \\ t & \rightarrow t' = t + t_0, \end{cases} \quad (1.8)$$

for any constant velocity vector \vec{v} and time shift t_0 . Together with spatial translations and rotations, these transformations form a mathematical *group* of coordinate changes that leave the equations of motion invariant. The existence of such a group has profound consequences: by Noether's theorem [207], each of the aforementioned transformations corresponds to a conserved quantity. In the case of classical mechanics, these correspond to the conservation laws of energy, momentum, and angular momentum. The conserved quantity associated with Galilean boosts can be shown to be the uniform motion of a system's center of mass, a result known as the center of mass theorem.

The equations governing electrodynamics, namely Maxwell's equations, are not invariant under Eq. (1.8). Instead, they are invariant under the more cumbersome Lorentz boosts [182]:

$$\begin{cases} x & \rightarrow x' = \gamma(x - v_0 t), \\ t & \rightarrow t' = \gamma(t - vx/c^2), \\ y & \rightarrow y', \\ z & \rightarrow z', \end{cases} \quad (1.9)$$

where $\gamma = \frac{1}{\sqrt{1-(v/c)^2}}$ is the *Lorentz factor*, and a boost in the x -direction was chosen without loss of generality.

The inconsistency between the symmetry groups of both theories cannot be reconciled. Eq. (1.8) implies a universally agreed-upon time coordinate t for all observers and the invariance of physical laws under the addition of arbitrarily high velocities, while Eq. (1.9) does away with a universal notion of time, allows transformations that “rotate” time into spatial coordinates, and conserves the speed of light c in all reference frames, ensuring that the trajectories of objects cannot exceed the speed of light in any frame.

The need to abandon the principles of classical mechanics became even more apparent

after the Michelson-Morley experiment [196] disproved the existence of the luminiferous aether in 1887. This result served as empirical evidence for the inadequacy of classical mechanics and ultimately led to the adoption of Einstein's theory of special relativity [75]. In special relativity, just as in electrodynamics, equations of motion are invariant under the Lorentz boosts depicted in Eq. (1.9).

1.3 Relativistic Quantum Mechanics

Despite the remarkable success of the Schrödinger equation in predicting properties of atoms and molecules [253], its manifestly non-relativistic form Eq. (1.6) cannot constitute a complete quantum theory of the microscopic world. The failure to account for relativistic effects ultimately diminishes its domain of validity, making it unsuitable for the study of heavy late-row chemical elements. This is because the speed of the electrons is proportional to the charge of the nucleus they orbit. Upon solving Eq. (1.6) for a hydrogen-like atom consisting of one electron and one nucleus of charge Z , one can compute the root mean square (RMS) value of the speed of the electron in the ground state $|n = 1, l = 0, m = 0\rangle$ as [209]

$$\sqrt{\langle v^2 \rangle} = \frac{1}{m} \sqrt{\langle p^2 \rangle} = Z\alpha c, \quad (1.10)$$

where $\alpha \approx 1/137$ is the fine structure constant. Therefore, the non-relativistic prediction for the average speed clearly breaks down when $Z \gg 1$, as is the case with heavy chemical elements. For instance, the RMS speed of a ground-state electron in Tellurium is over a third of the speed of light, exhibiting strong relativistic effects and the inadequacy of Eq. (1.6). Furthermore, although relativistic effects are strongest in the electronic core, the valence electrons do not escape unscathed: since the valence electronic states are orthogonal to the core electronic states, significant relativistic corrections to the core states necessarily affect the valence electronic states. Put another way, the core and valence electronic states are not independent. Therefore, relativistic effects have profound consequences for the valence electrons, and may even alter the very nature of chemical bonds [68]. Other implications include the liquidity of mercury and the yellow color of gold [209, 231, 232, 288]. All of these suggest that the incorporation of relativity into quantum theory is crucial for the correct

qualitative and quantitative description of heavy-element chemical systems.

To account for relativistic effects, one must somehow equalize the footing of the spatial and the time coordinates in Eq. (1.6). Specifically, the goal is to replace the non-relativistic Hamiltonian in Eq. (1.6), in which the kinetic energy takes the form $\frac{p^2}{2m}$, by a Lorentz-covariant counterpart. In accordance with the classical Hamiltonian, the corresponding quantum operator should obey the operator-valued relativistic energy-momentum relation

$$p^\mu p_\mu = E^2/c^2 - P^2 = m^2 c^2 \implies \hat{H}^2 - c^2 \hat{P}^2 = m^2 c^4 \mathbf{1}. \quad (1.11)$$

Klein and Gordon proposed a relativistic quantum theory [143] which implemented Eq. (1.11) indirectly by squaring both sides of Eq. (1.2):

$$-\hbar^2 \frac{\partial^2}{\partial t^2} |\psi\rangle = \hat{H}^2 |\psi\rangle = \left(c^2 \hat{P}^2 + m^2 c^4 \right) |\psi\rangle = \left(-\hbar^2 c^2 \nabla^2 + m^2 c^4 \right) |\psi\rangle. \quad (1.12)$$

By using the relativistic energy-momentum relation, the resulting second-order partial differential equation is manifestly Lorentz invariant. However, this equation is not a suitable quantum wave equation because it is not norm-preserving: it can be shown [277] that the Klein-Gordon wavefunction dissipates probability over time. Simply said, while the Klein-Gordon equation obeys relativity, it does not obey quantum mechanics!

A different approach was pursued by Dirac [57], in which the energy-momentum relation of Eq. (1.11) is explicitly implemented into the construction of the Hamiltonian operator. To do so, a general form of the Hamiltonian is guessed by dimensional analysis:

$$\hat{H} = c\vec{\alpha} \cdot \vec{p} + \beta mc^2, \quad (1.13)$$

where $\langle \mathbf{x}' | \vec{p} | \phi \rangle = -i\hbar \vec{\nabla} \langle \mathbf{x}' | \phi \rangle$ is the usual momentum operator and $\vec{\alpha}$ and β are some constant factors. The factors are to be chosen such that Eq. (1.11) is satisfied:

$$\hat{H}^2 = (c\vec{\alpha} \cdot \vec{p} + \beta mc^2)^2 = c^2 (\vec{\alpha} \cdot \vec{p})^2 + \beta^2 m^2 c^4 + mc^3 (\vec{\alpha} \cdot \vec{p} \beta + \beta \vec{\alpha} \cdot \vec{p}). \quad (1.14)$$

Comparing Eqs. (1.11) and (1.14), the constant factors must satisfy:

$$\{\alpha_i, \alpha_j\} = 2\delta_{ij}, \{\alpha_i, \beta\} = 0, \beta^2 = \mathbf{1}, \quad (1.15)$$

where $\{X, Y\} = XY + YX$ is the *anticommutator*. Trivially, the coefficients $\vec{\alpha}$ and β cannot be c-numbers, because $\{\alpha_i, \beta\} = 2\alpha_i\beta = 0$ would imply that $\vec{\alpha} = 0$ or $\beta = 0$, making the implementation of the other two conditions impossible. Furthermore, it can be shown that the smallest objects satisfying Eq. (1.15), known as the Clifford algebra [41], admit a four-dimensional matrix representation [57]. It is trivial to verify that Eq. (1.15) is satisfied by

$$\alpha_i = \begin{pmatrix} 0 & \sigma_i \\ \sigma_i & 0 \end{pmatrix}, \beta = \begin{pmatrix} \mathbf{1} & 0 \\ 0 & -\mathbf{1} \end{pmatrix}. \quad (1.16)$$

Finally, the Dirac equation is obtained by substituting the Hamiltonian defined by Eqs. (1.13) and (1.16) into the abstract Schrödinger in Eq. (1.2):

$$i\hbar \frac{\partial}{\partial t} \psi = -i\hbar c \vec{\alpha} \cdot \vec{\nabla} + \beta mc^2 \psi. \quad (1.17)$$

Unlike the Klein-Gordon equation in Eq. (1.12), the Dirac equation successfully meshes relativity and quantum mechanics together. The former is satisfied because the Dirac matrices, as composed out of Pauli matrices, form a representation of the group $SU(2) \oplus SU(2)$ with the spin quantum numbers $(1/2, 0) \oplus (0, 1/2)$. As the algebra of the Lorentz group $SO(3, 1)$ can be “diagonalized” into the algebras of two non-overlapping copies of $SU(2)$ [277], the Lagrangian corresponding to the Dirac equation turns out to be Lorentz invariant. The latter can be seen by taking the time derivative of the probability density and applying Eq. (1.17):

$$\begin{aligned} \partial_t (\psi^\dagger \psi) &= \psi^\dagger (\partial_t \psi) + \text{c.c} = \psi^\dagger \left(-c \vec{\alpha} \cdot \vec{\nabla} \psi - i/\hbar \beta mc^2 \psi \right) + \text{c.c} \\ &= 2\text{Re} \left(-c \psi^\dagger \vec{\alpha} \cdot \vec{\nabla} \psi - i/\hbar mc^2 |\psi|^2 \right) = 2\text{Re} \left(-c \psi^\dagger \vec{\alpha} \cdot \vec{\nabla} \psi \right) \\ &= -c \left(\psi^\dagger \vec{\alpha} \cdot \vec{\nabla} \psi + \psi \vec{\alpha} \cdot \vec{\nabla} \psi^\dagger \right) = -c \vec{\nabla} \cdot \left(\psi^\dagger \vec{\alpha} \psi \right), \end{aligned} \quad (1.18)$$

where we used the fact that $\|\psi\|^2 \in \mathbb{R}$ in the fourth step. This is the continuity equation [78], and the total conserved charge is the wavefunction norm:

$$\partial_t \int_{\mathbb{R}^3} d^3x \psi^\dagger \psi = \int_{\mathbb{R}^3} d^3x \partial_t (\psi^\dagger \psi) = -c \int_{\mathbb{R}^3} d^3x \vec{\nabla} \cdot (\psi^\dagger \vec{\alpha} \psi) = -c \int_{\partial\mathbb{R}^3} d\vec{a} \cdot (\psi^\dagger \vec{\alpha} \psi) = 0. \quad (1.19)$$

By satisfying both Lorentz covariance and probability conservation, the Dirac equation became the first successful relativistic quantum theory. Further investigation of its properties [57] reveals that it describes a spin-one-half fermionic particle. It admits both negative and positive energy solutions, today interpreted as the existence of both electrons and positrons in the spectrum of the theory [3, 58]. Its four-component nature, due to the dimensionality of the Dirac matrices of Eq. (1.16), corresponds to the simultaneous and coupled description of electronic and positronic degrees of freedom as a single quantum state.

1.4 Relativistic Electronic Structure Theory

The ability of the Dirac equation to successfully describe a relativistic spin-one-half fermion makes it a natural choice in the study of the electronic relativistic effects in atomic and molecular systems mentioned in Section 1.3. Steady-state solutions can in principle be obtained by solving the time-independent Schrödinger equation with the molecular Dirac Hamiltonian, the many-body equivalent of Eq. (1.13), containing the kinetic energies and electric interactions of all electrons and nuclei in the system. This approach, while formally exact, is impractical: no exact solutions exist to the molecular Dirac equation beyond the hydrogen atom [290]. The need to employ approximations in solving the molecular Dirac equation is apparent.

One such approximation is the *Born-Oppenheimer* approximation [26], in which the nuclei are treated as classical degrees of freedom. Because the mass of the nuclei is significantly larger than that of the electrons, it is assumed that the effect of electronic motion on the nuclei is negligible, and the kinetic nuclear terms are simply omitted from the equation. By doing so, one arrives at the *electronic* Dirac equation, in which the nuclear degrees of

freedom are external *parameters*:

$$\begin{pmatrix} V(\{\vec{R}_i\}) + m_e c^2 & c\vec{\sigma} \cdot \vec{p} \\ c\vec{\sigma} \cdot \vec{p} & V(\{\vec{R}_i\}) - m_e c^2 \end{pmatrix} \begin{pmatrix} \psi_e^L(\{\vec{R}_i\}; \{\vec{r}_j\}) \\ \psi_e^S(\{\vec{R}_i\}; \{\vec{r}_j\}) \end{pmatrix} = E \begin{pmatrix} \psi_e^L(\{\vec{R}_i\}; \{\vec{r}_j\}) \\ \psi_e^S(\{\vec{R}_i\}; \{\vec{r}_j\}) \end{pmatrix}. \quad (1.20)$$

Here, $\{\vec{R}_i\}$ are the parametric nuclear positions, $\{\vec{r}_j\}$ are the electronic positions, and V is the net electric potential energy of the system, consisting of Coulomb (and possibly Breit [28]) interactions. Crucially, the momentum operator \vec{p} in Eq. (1.20) is only with respect to the electronic positions $\{\vec{r}_j\}$. This considerably reduces the dimensionality of the configuration space, and is the ubiquitous starting point of electronic structure theory [68, 283].

Notably, the Born-Oppenheimer approximation is insufficient for the analysis of many-body electronic systems at low temperatures [141]. The Einstein-Debye model [54] demonstrates that at temperatures below the Debye temperature, the heat capacity of solids deviates from the classical Dulong–Petit law [62]. This is because at such temperatures, nuclear quantum effects become significant, and nuclei may access vibrational degrees of freedom. A possible treatment of such low-temperature regimes relies on the nuclear-electronic orbital (NEO) framework [105] and is beyond the scope of this manuscript.

A further reduction in dimensionality can be achieved by decoupling the electronic and protonic degrees of freedom in the electronic Dirac equation. Because of its Hermiticity, this can be achieved by unitarily transforming the four-component Dirac Hamiltonian:

$$U^\dagger \hat{H}_{4C} U = U^\dagger \begin{pmatrix} \hat{H}_{LL} & \hat{H}_{LS} \\ \hat{H}_{SL} & \hat{H}_{SS} \end{pmatrix} U = \begin{pmatrix} \hat{H}_+ & 0 \\ 0 & \hat{H}_- \end{pmatrix}. \quad (1.21)$$

Such a transformation expresses Eq. (1.20) as two decoupled sets of equations for the negative and positive sectors of the system. With the exception of positronic scattering in molecular systems, the typical chemical energy scale (1 Å) renders the positronic degrees of freedom ignorable. Consequently, for our purposes, the electronic Dirac equation may be solved in the positive energy sector alone, resulting in a two-fold reduction in the dimension of the Hilbert space.

One approach to apply such a unitary transformation is the Douglas-Kroll-Hess (DKH) method [60, 113, 200, 220, 234, 235, 294, 309, 310]. The method formally decouples \hat{H}_{4C} by applying a sequence of unitary transformations based on the free-particle Foldy-Wouthuysen (FW) transformation [88]. The procedure is perturbative, because successive transformations decouple the Hamiltonian to an additional order in the electric interaction V .

Another method, which we employ throughout this manuscript, is the exact two-component (X2C) method [64, 69, 70, 100, 121, 126, 131, 155, 163, 171, 178, 180, 181, 221, 222, 250, 286]. The X2C method seeks to explicitly construct the unitary matrix U by forming two independent sets of basis vectors for the large and small blocks (see Eq. (1.21)). Since the exact construction of the decoupling unitary transformation is computationally intractable, one often uses various approximations as proxies for U . One such approximation is the one-electron X2C (1eX2C) method which only uses the one-electron portion of \hat{H}_{4C} to approximately decouple the Hamiltonian into positive and negative energy sectors [64–67, 69, 70, 100, 119, 126, 155, 163, 171, 177, 178, 180, 181, 183, 221, 222, 260]. The one-electron portion of the Hamiltonian contains the kinetic energies and the electron-nuclear interactions, and is computationally easier to decouple than the two-body electric interaction present in V .

1.5 Spectral Methods: Electronic Structure Theory in a Basis

Even after employing the Born-Oppenheimer and 1eX2C approximations, solving the many-body electronic Dirac equation is still computationally challenging. As exact solutions do not exist beyond the hydrogen atom [290], one often resorts to numerical methods to solve the eigenvalue equation Eq. (1.20). A naive approach would be to discretize the N_e -body electronic position space as a grid and express all derivative and multiplicative operators as matrices, thereby reducing the problem into an ordinary matrix eigenproblem. Besides the inability to enforce boundary conditions at spatial infinity, which are critical to ensure the normalizability of the solutions, the main problem with this approach is its explosive dimensionality. In this paradigm, the N_e -electron Dirac Hamiltonian is represented by a $D^{N_e} \times D^{N_e}$ -dimensional matrix, where D is the size of the grid. Such exponential growth makes this approach incapable of handling even modestly-sized chemical systems.

An alternative approach is to use *spectral* methods, in which one-electron wavefunctions are represented as a linear combination of pre-defined, analytical spinor *orbitals*, $\{\phi_{P_i}\}_{i=1}^{N_o}$. Here, $N_o > N_e$, allowing the set of one-electron orbitals to describe electrons in ground and excited states. Throughout this dissertation, we will focus on the methods suitable for wavefunctions which decay at spatial infinity. Such methods yield wavefunctions which faithfully describe localized atoms and molecules. This is in contrast with the family of methods suitable to treat infinite systems, which employ periodic boundary conditions.

The one-electron orbitals can be used to construct N_e -electron wavefunctions as *Slater determinants*:

$$\langle \vec{r}_1 \dots \vec{r}_{N_e} | a_{P_1}^\dagger \dots a_{P_{N_e}}^\dagger | \emptyset \rangle \equiv |\phi_{P_1} \dots \phi_{P_{N_e}}| = \frac{1}{\sqrt{N_e!}} \begin{vmatrix} \phi_{P_1}(\vec{r}_1) & \dots & \phi_{P_{N_e}}(\vec{r}_1) \\ \vdots & \ddots & \vdots \\ \phi_{P_1}(\vec{r}_{N_e}) & \dots & \phi_{P_{N_e}}(\vec{r}_{N_e}) \end{vmatrix}. \quad (1.22)$$

Note that Slater determinants of this form are totally antisymmetric in their arguments by construction, and are therefore suitable N_e -electron wavefunctions obeying the Pauli exclusion principle. Expressing the determinant states via creation operators acting on a vacuum, one can verify that the total antisymmetry yields the anti-commutation relations [68]

$$\{a_{P_i}, a_{P_j}\} = 0; \quad \{a_{P_i}^\dagger, a_{P_j}\} = \delta_{P_i P_j}, \quad (1.23)$$

thereby defining a fermionic Fock space. The vacuum state $|0\rangle$ contains no electrons, and the action of $a_{P_i}^\dagger$ is to occupy the orbital $\phi_{P_i}(\vec{r})$ with an electron. Using the Slater-Condon rules [42, 273], the electronic Dirac Hamiltonian in Eq. (1.20) can be expressed as

$$\hat{H} = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} g_{pqrs} a_p^\dagger a_r^\dagger a_s a_q. \quad (1.24)$$

The coefficients h_{pq} and g_{pqrs} can be determined by the appropriate one- and two-electron

integrals

$$\begin{aligned}
 h_{pq} &\propto \int d^3x \phi_{P_p}^*(\vec{x}) \left(\vec{\sigma} \cdot \vec{\nabla} \text{ or } \frac{1}{\|\vec{x} - \vec{R}\|} \right) \phi_{P_q}(\vec{x}), \\
 g_{pqrs} &\propto \int d^3x_1 d^3x_2 \phi_{P_p}^*(\vec{x}_1) \phi_{P_r}^*(\vec{x}_2) \left(\frac{1}{\|\vec{x}_1 - \vec{x}_2\|} \right) \phi_{P_q}(\vec{x}_1) \phi_{P_s}(\vec{x}_2), \quad (1.25)
 \end{aligned}$$

where \vec{R} is a nuclear position.

Notably, the number of N_e -electron Slater determinants constructed from the set $\{\phi_{P_i}\}_{i=1}^{N_o}$ is equal to the number of ways to construct a totally-antisymmetric N_e -variable function out of N_o basis functions, which is $\binom{N_o}{N_e}$. Consequently, the set of all Slater determinants forms a *complete basis* for all totally-antisymmetric N_e -variable wavefunctions constructed from the given orbitals. Therefore, the spectral solution of the N_e -electron Dirac equation can be expressed as

$$\langle \vec{r}_1, \dots, \vec{r}_{N_e} | \psi \rangle = \sum_{P_1 \dots P_{N_e}} C^{P_1 \dots P_{N_e}} |\phi_{P_1} \dots \phi_{P_{N_e}}| \quad (1.26)$$

for some complex coefficients $C^{P_1 \dots P_{N_e}}$. By projecting Eq. (1.20) onto the basis of Slater determinants, the task of solving the N_e -electron Dirac equation reduces to an $\binom{N_o}{N_e}$ -dimensional matrix eigenproblem. Because the low-lying spectrum is most pertinent to chemistry, this dimension is often significantly smaller than that of the grid-based approach (D^{N_e}) in practical applications.

Evidently, the successful use of spectral methods hinges on the construction of a set of orbitals $\{\phi_{P_i}\}_{i=1}^{N_o}$ capable of accurately expressing the low-lying spectrum of the system via Slater determinants. It is convenient to construct such a set as linear combinations of *atomic orbitals* (AOs), which describe the wavefunctions of a hydrogen-like system (in which a single electron orbits a single nucleus). Because such a system is one of the few example of exactly-solvable quantum systems, exact expressions for AOs are obtainable [48]. The solutions can be expressed as

$$\chi_{nlm}(r, \theta, \phi) = R_{nl}(r) Y_{lm}(\theta, \phi), \quad (1.27)$$

where $r = 0$ is the nuclear position, R is the radial function, and Y_{lm} are the usual spherical harmonics. Similarly to the non-relativistic hydrogen atom [253], the leading-order behavior of the solutions near the nucleus is [48]

$$\chi_{nlm} \approx r^l e^{-r/(2a_0)} Y_{lm}(\theta, \phi). \quad (1.28)$$

While formally exact, such one-electron functions (called Slater-type orbitals or STOs [274]) are seldom used in computational spectral methods. This is because projecting Eq. (1.20) onto STOs entails the evaluation of Eq. (1.25) with exponential functions. Because there are no closed-form expressions of such integrals, the adoption of STOs requires the (generally undesirable) numerical evaluation of one- and two-electron integrals.

An alternative set of AOs can be constructed by replacing the exponential in Eq. (1.28) with a Gaussian centered around the nucleus, resulting in a Gaussian-type orbital (GTO) [27]:

$$\xi_{\alpha lm} \approx r^l e^{-\alpha r^2} Y_{lm}(\theta, \phi). \quad (1.29)$$

A major advantage of GTOs is that the one- and two-electron integrals become simple Gaussian integrals, which have closed-form expressions obtainable from known recursion relations [107, 210]. Therefore, in contrast with STOs, no numerical integration is required. This saves computational costs and avoids numerical accuracy and stability issues. As GTOs are not theoretically motivated and are used merely for convenience, we expect a suitable AO to be expressed as a linear combination of GTOs,

$$\chi_{nlm} = \sum_{i'l'm'} A_{nlm}^{i'l'm'} \xi_{\alpha_i l' m'}. \quad (1.30)$$

The choice of contraction coefficients $A_{nlm}^{i'l'm'}$ and exponents α_i defines a basis set of AOs.

GTO basis sets are often sufficient for practical computations, but are not without pathologies. An infamous example is its failure to satisfy Kato's cusp condition [136]. As evident in Eq. (1.28), the s -orbital STO (for which $l = 0$) exhibits a cusp at $r = 0$ stemming from the singularity of the Coulomb potential at the nucleus. The Gaussians used in GTOs, on the other hand, are smooth at $r = 0$, and no linear combination of them is capable of

reproducing the cusp. More severely, purely electronic cusps also arise from the singularity inherent in the electron-electron interaction. Thus, the computational advantages of GTOs come at the cost of a complete description of the sought-after electronic wavefunction.

1.6 The Mean-Field Solution

Equipped with a set of AOs describing the wavefunction of a single electron orbiting a single nucleus, we seek a set of *molecular orbitals* (MOs) that describe the wavefunction of a single electron orbiting *multiple* nuclei. A common choice is to construct a set $\{\phi'_{P_i}\}_{i=1}^{N_o}$ of MOs that minimizes the energy expectation value of a single Slater determinant. Expressed in terms of the AOs, we seek a unitary (inner-product-preserving) transformation $\phi'_{P_i} = \sum_j U_{ij} \phi_{P_j}$ such that

$$\langle 0 | \hat{H} | 0 \rangle \equiv \langle \emptyset | a'_{P_{N_e}} \cdots a'_{P_1} \hat{H} a'_{P_1}^\dagger \cdots a'_{P_{N_e}}^\dagger | \emptyset \rangle \quad (1.31)$$

is minimized, where the identification between the orbitals ϕ'_{P_j} and the Fock-space creation operators $a'_{P_j}^\dagger$ of Eq. (1.22) was used.

To this end, we parameterize the sought-after unitary transformation in Fock space as $U = e^{\hat{\kappa}}$, where $\hat{\kappa}$ is an anti-Hermitian generator of orbital rotations (ensuring the unitarity of \hat{U}). Without loss of generality, we take $\hat{\kappa}$ to be of first order:

$$\hat{\kappa} = \sum_{pq} \kappa_{pq} a_p^\dagger a_q, \quad (1.32)$$

where the anti-Hermiticity of $\hat{\kappa}$ and the anticommutation relations of Eq. (1.23) imply that

$$\kappa_{pq} = -\kappa_{qp}^*. \quad (1.33)$$

We seek rotation parameters that minimize Eq. (1.31). Plugging in Eq. (1.32) and

applying the Baker–Campbell–Hausdorff (BCH) formula [11, 32, 106], we obtain [68]:

$$\begin{aligned}
E' &= \langle 0 | e^{-\hat{\kappa}} \hat{H} e^{\hat{\kappa}} | 0 \rangle = \langle 0 | \left(\hat{H} + [\hat{H}, \hat{\kappa}] + \frac{1}{2!} [[\hat{H}, \hat{\kappa}], \hat{\kappa}] + \dots \right) | 0 \rangle \\
&= \langle 0 | \hat{H} | 0 \rangle + \langle 0 | [\hat{H}, \hat{\kappa}] | 0 \rangle + \langle 0 | [[\hat{H}, \hat{\kappa}], \hat{\kappa}] | 0 \rangle \\
&= E + \sum_{pq} \kappa_{pq} \langle 0 | [\hat{H}, a_p^\dagger a_q] | 0 \rangle + \sum_{pqrs} \kappa_{pq} \kappa_{rs} \langle 0 | [[\hat{H}, a_p^\dagger a_q], a_r^\dagger a_s] + \dots \\
&\equiv E + \sum_{pq} \kappa_{pq} \hat{g}_{pq} + \sum_{pqrs} \kappa_{pq} \kappa_{rs} \hat{G}_{pq;rs} + \dots
\end{aligned} \tag{1.34}$$

We thus obtain an expression for the energy E' of the Slater determinant of the rotated MOs as an infinite expansion in powers of the rotation parameters. The first term in the expansion is the energy of the original Slater determinant. The factors multiplying the first- and second-order terms, \hat{g}_{pq} and $\hat{G}_{pq;rs}$, are therefore called the orbital *gradients* and *Hessians*.

The expression for the energy of the transformed Slater determinant can then be minimized by setting its derivatives with respect to all rotation parameters κ_{pq} equal to zero:

$$0 = \frac{\partial E'}{\partial \kappa_{pq}} = \hat{g}_{pq} + \frac{1}{2} \sum_{rs} \kappa_{rs} \left(\hat{G}_{pq;rs} + \hat{G}_{rs;pq} \right) + \dots \tag{1.35}$$

Truncated at the order presented here, Eq. (1.35) is a linear equation for the rotation parameters κ_{pq} . These linear equations can be solved iteratively until the solution to the untruncated equation is obtained, at which point the orbital gradients will vanish.

To interpret the solution prescribed by the procedure above, we use Eq. (1.24) to obtain the explicit orbital gradient as [68]

$$g_{pq} = \langle 0 | [\hat{H}, a_p^\dagger a_q] | 0 \rangle = \dots = \left(h_{qp} + \sum_r g_{rrqp} - g_{qrrp} \right) - \left(h_{pq} + \sum_r g_{rrpq} - g_{prrq} \right)^* = f_{qp} - f_{pq}^*, \tag{1.36}$$

where

$$f_{qp} \equiv h_{qp} + \sum_r g_{rrqp} - g_{qrrp} \tag{1.37}$$

is the *Fock* matrix. Exact solutions of Eq. (1.35) correspond to $g_{pq} = f_{qp} - f_{pq}^* = 0$, satisfied

by a set of MOs that *diagonalize* the Fock matrix. As Eq. (1.37) contains contributions from all orbitals $r \neq p, q$, it is readily understood as capturing the *average* electric interaction an electron experiences due to the electric field generated by all other electrons. For this reason, this solution is called the mean-field (or Hartree-Fock) solution to Eq. (1.20).

1.7 Configuration Interaction and Electronic Correlation

Armed with the mean-field MOs, we can now solve Eq. (1.20) by projecting it onto the corresponding $\binom{N_o}{N_e}$ -dimensional basis of Slater determinants. In the configuration interaction (CI) method, all possible Slater determinants (configurations) are generated from the reference mean-field determinant as

$$|\psi\rangle_{\text{CI}} = \sum_K C_K |K\rangle, \quad (1.38)$$

$$|K\rangle \in \{|0\rangle, |0_i^a\rangle, |0_{ij}^{ab}\rangle, \dots\},$$

where $\{|0_i^a\rangle, |0_{ij}^{ab}\rangle, \dots\}$ are Slater determinants generated from the reference mean-field determinant $|0\rangle$ by applying the excitation operators:

$$|0_i^a\rangle = \hat{E}_{ai} |0\rangle \equiv a_a^\dagger a_i |0\rangle; \quad |0_{ij}^{ab}\rangle = \hat{e}_{aibj} |0\rangle \equiv a_a^\dagger a_b^\dagger a_j a_i |0\rangle. \quad (1.39)$$

The coefficients C_K are then obtained by diagonalizing the $\binom{N_o}{N_e} \times \binom{N_o}{N_e}$ -dimensional Hamiltonian matrix $H \equiv \langle L | \hat{H} | K \rangle$. The inclusion of all possible excited configurations in Eq. (1.38) leads to the full CI (FCI) solution of Eq. (1.20).

FCI offers the most complete and accurate description of the electronic structure of chemical systems within a given basis set, providing the exact spectral solution to the non-relativistic electronic Schrödinger equation [134, 263, 264, 268, 284, 295]. Due to its variational nature, FCI is particularly well-suited for treating strong relativistic effects, such as spin-orbit and spin-spin couplings [2, 17, 18, 120, 121, 130, 132, 146–149, 183, 238, 260, 270, 289, 300].

At its core, solving the FCI problem reduces to diagonalizing a large many-electron Hamiltonian matrix. This matrix is Hermitian, sparse, and typically diagonally dominant, making it well-suited for iterative diagonalization techniques that can efficiently converge

a few eigenstates without requiring the full storage or construction of the entire matrix. However, as the number of determinants grows factorially with system size (reflecting the combinatorial nature of Slater determinant enumeration within the full Hilbert space), even iterative methods become intractable beyond a certain threshold.

One approximate method that aims to approximate the FCI solution is complete active space CI (CASCI). In the CASCI method [85–87, 95, 145–147, 150, 213, 243, 302], it is assumed that only a subspace of the full Hilbert space of the system contains meaningful correlation, and the FCI wavefunction is approximated as the CI wavefunction in the truncated space (the so-called active space). The truncation often leads to an underestimation of dynamic correlation. Applying more computationally demanding methods such as multiconfigurational self-consistent field (MCSCF) [17, 94, 132, 144, 174, 185, 187, 188, 238, 270, 284, 289, 301, 305, 306, 319], multireference configuration interaction (MRCI) [115, 121, 174–176, 189, 246, 284, 304], and many-body perturbation theory (MRPT2, CASPT2, NEVPT2, MC-PDFT) [5–7, 35, 172, 183, 186, 244, 270, 300] is typically required to achieve qualitative and quantitative agreement with experiment.

In both the FCI and CASCI cases, the Hamiltonian matrix is usually diagonalized by Krylov subspace methods [159] suitable for large sparse matrices. One such method is the iterative Davidson method [49], which diagonalizes the Hamiltonian in a linear subspace spanned by CI coefficient vectors $\{\{\vec{C}^1\}, \dots, \{\vec{C}^N\}\}$. The method requires the evaluation of the computationally-intensive matrix-vector product

$$\begin{aligned}
\sigma_K &= \sum_{pq} \sum_L h_{pq} \langle K | \hat{E}_{pq} | L \rangle C_L \\
&\quad + \frac{1}{2} \sum_{pqrs} \sum_L g_{pqrs} \langle K | \hat{e}_{pqrs} | L \rangle C_L \\
&= \sum_{pq} \sum_L h'_{pq} \langle K | \hat{E}_{pq} | L \rangle C_L \\
&\quad + \frac{1}{2} \sum_{pqrs} \sum_{L,J} g_{pqrs} \langle K | \hat{E}_{pq} | J \rangle \langle J | \hat{E}_{rs} | L \rangle C_L,
\end{aligned} \tag{1.40}$$

where we used [150, 271]

$$\hat{e}_{pqrs} = \hat{E}_{pq} \hat{E}_{rs} - \delta_{qr} \hat{E}_{ps} \tag{1.41}$$

and

$$h'_{pq} = h_{pq} - \frac{1}{2} \sum_r g_{prrq}. \quad (1.42)$$

The result of this contraction, called the σ -vector, is then used to project the configuration-space Hamiltonian matrix onto the selected subspace. Upon diagonalizing the subspace Hamiltonian matrix, one computes the Ritz residual [239] of the subspace eigenvector \vec{C} as

$$\vec{r} \equiv H\vec{C} - E\vec{C} = H\vec{C} - (\vec{C}^\dagger H\vec{C})\vec{C}. \quad (1.43)$$

The residual measures the vectorial deviation between the subspace eigenvector \vec{C} and its true configuration-space counterpart. Because the residual is orthogonal to \vec{C} , it is often used to expand the subspace. Thus, Krylov methods iteratively expand the subspace, thereby converging to the true configuration-space solution.

The factorial growth of the configuration space with the number of correlated MOs and electrons renders the explicit construction of the Hamiltonian matrix practically impossible. Therefore, within the Davidson procedure, Eq. (1.40) is evaluated using one-electron excitation lists. These lists map each Slater determinant K to a list of all configurations L *connected* to K by exciting a single electron. Because K is connected to exactly $N_e \times (N_o - N_e + 1)$ configurations, the total size of the one-electron excitation lists is $\binom{N_o}{N_e} \times N_e \times (N_o - N_e + 1)$. Thus, the use of excitation lists drastically reduces the memory footprint of the iterative diagonalization, which otherwise would require the explicit storage of the $\binom{N_o}{N_e} \times \binom{N_o}{N_e}$ Hamiltonian matrix.

Software implementations of Eq. (1.40) based on excitation lists must assign a binary code, namely an *address*, to each determinant K in the configuration space. Under a generic addressing scheme, the addresses of all relevant determinants would need to be computed and stored before the diagonalization procedure, and the evaluation of Eq. (1.40) would require the list of all determinant addresses to be sorted. To circumvent these computationally-expensive steps, Knowles and Handy [150, 152] introduced their string-based paradigm based on a graphical representation of the orbital space [61].

The strings, which correspond to walks in the graph, form a contiguously-ordered one-

to-one map between configuration determinants and integer-valued addresses. To generate the addresses, the graph vertices (labeled by the number of electrons e and the number of orbitals o) are assigned a weight $w(e, o)$. Edges connecting two vertices are assigned edge weights $w((e, o), (e', o'))$. Under the usual lexical ordering, the tail vertex (N_e, N_o) is assigned a weight of 1. The weights of all other vertices are then computed recursively as [150, 152, 267]

$$w(e, o) = w(e + 1, o + 1) + w(e, o + 1). \quad (1.44)$$

Edges of the form $((e, o), (e, o + 1))$ correspond to a configuration in which MO $o + 1$ is unoccupied and are assigned zero weight. The other type of edge occupies MO $o + 1$ with electron $e + 1$ and is weighted as

$$w((e, o), (e + 1, o + 1)) = w(e + 1, o + 1) + \cdots + w(e + 1, e + 1). \quad (1.45)$$

The address strings are then generated from a walk in the graph $\{e_o\}_{o=0}^{N_o-1}$ as

$$|K\rangle \mapsto \sum_{o=0}^{N_o-1} w((e_o, o), (e_{o+1}, o + 1)). \quad (1.46)$$

The generation of the excitation lists and the corresponding determinant addressing scheme, which will be discussed in detail throughout this dissertation, are at the core of any software implementation of the CI procedure. The purpose of implementing this computationally-nontrivial procedure is ultimately to recover electronic correlation energy into the mean-field solution discussed in Section 1.6, defined as the difference between the true and mean-field energies,

$$E_{\text{corr}} = E_{\text{exact}} - E_{\text{MF}}. \quad (1.47)$$

Other methods to recover electronic correlation energy back into the mean-field solution include density-functional theory (DFT) [19, 20, 55, 56, 71, 103, 116, 127, 135, 154, 166, 167, 173, 219, 223, 225, 257, 278, 324], multireference second-order perturbation theory (MRPT2, CASPT2) [5, 6, 35, 183, 186, 244, 270, 300], coupled cluster (CC) [14, 15, 47, 208, 230, 256], the GW and Bethe-Salpeter equation (BSE) methods [23, 29, 109, 125, 215, 241, 279], as

well as stochastic methods such as quantum Monte Carlo (QMC) [25, 36, 37, 89, 193, 195, 237]. Nevertheless, CI serves as the most accurate post-mean-field method, recovering the formally-exact spectral solution to Eq. (1.20) in the limit of a complete basis set. Therefore, CI is highly useful in the benchmarking of the aforementioned methods. As will be shown in Chapter 5, such benchmark studies can reveal the domain of validity and general convergence behavior of more approximate state-of-the-art electronic structure methods.

Chapter 2

REVIEW OF THE STP-DAS FRAMEWORK

2.1 Introduction

Full configuration interaction (CI) provides the most accurate electronic structure description of a molecular system in a given basis [134, 263, 264, 268, 284, 295]. However, its computational cost grows factorially with the system's size due to the combinatorial increase in the number of configurations, as it includes all possible excited states within the complete orbital space. Significant work has been invested in developing approximate CI techniques, as exemplified by the complete active space (CAS) approach [213, 243] along with its numerous variants and enhancements. The CAS approach specifically focuses on configurations that arise from a limited active space, which consists of a selection of orbitals and electrons deemed chemically significant. While the CAS framework effectively addresses static correlation, it often fails to capture a substantial portion of the dynamic correlation arising from excitations beyond the active space. This limitation can be overcome by incorporating multireference configuration interaction (MRCI) to account for the missing dynamic correlation. [115, 121, 174–176, 284, 304] and multireference second-order perturbation theory (MRPT2, CASPT2) [5, 6, 35, 183, 186, 244, 270, 300].

To minimize the number of configurations in the CI expansion further, implementing limitations on excitation operators is effective, resulting in the development of the restricted active space (RAS) [144, 190, 213] and occupation restricted multiple active space (ORMAS) methods [128, 129]. A more generalized strategy involves dividing the total correlation space into smaller generalized active spaces (GASs) [85–87, 146, 147, 185, 213, 301]. The primary benefit of GAS lies in its ability to apply a broad framework of excitation constraints, providing a flexible framework for truncated CI types of electronic structure calculations.

Recent advancements in CI-based relativistic methods, which variationally incorporate both scalar relativity and spin-orbit couplings at the molecular orbital level, have spurred a

new era in multi-reference electronic structure theory [2, 17, 18, 120, 121, 130, 132, 146–149, 183, 238, 260, 270, 289, 300]. This progress is largely motivated by improvements in core-electron spectroscopies, research in magnetic and spintronic materials, as well as the chemistry of rare-earth and heavy elements. Employing two- and four-component complex-valued orbitals within Kramers’ unrestricted framework [68, 236] inherently results in increased computational requirements for CI-based relativistic methods, as manifested by a significantly larger configurational space, a less sparse CI vector, and a higher count of floating-point operations compared to their non-relativistic counterparts.

In computations using CI or its various approximations, constructing the Hamiltonian matrix demands an algorithm capable of efficiently navigating through all determinants, with constraints on excitation operators and the active space. This process is commonly referred to as the CI addressing scheme. Since the CI addressing scheme is the core engine of any high-performance CI calculation, considerable efforts have been directed toward developing highly vectorized CI algorithms optimized to take advantage of the capabilities of modern computational infrastructures. Handy’s string-based addressing scheme [150] stands out as one of the most effective methodologies and is implemented in almost all high-performance CI programs. The string-based addressing scheme can efficiently produce a unique address for each configuration, leading to a precomputed list of non-zero excitations that facilitates a highly vectorized construction of the CI Hamiltonian. However, the storage of the excitation list from calculations involving large active spaces presents a significant computational challenge. For a CI problem with N configurations, the size of each CI vector grows linearly with respect to N while the size of the one-electron excitation list expands quadratically. For example, in a CAS calculation involving 42 Kramers’ unrestricted orbitals and 24 electrons, the combinatorial factor $\binom{42}{24}$ implies that the storage requirement for each CI vector is 1 TB (using complex-valued double precision). However, to store the non-zero elements of the one-electron excitation list, an estimated 1.8 petabytes (PB) of memory would be needed even with bit-wise compression. The size of this list is proportional to $n_e \times (n_h + 1) \times N$, where n_e is the number of electrons and n_h is the number of holes (unoccupied orbitals) in the complete active space.

Simply put, the bottleneck for CI methods is the memory requirement arising from the

storage of the excitation list, which practical CI-based methods must work to circumvent. The seminal works for direct CI calculations aimed to optimize memory storage and locality [150, 212, 271, 318] sometimes at the cost of additional computation. In the modern era of heterogeneous computation, this bottleneck is particularly critical for performance on accelerators such as graphical processing units (GPUs), as the memory on accelerators is limited and data transfer with the accelerator might be constrained by the bandwidth and latency [83, 84]. A straightforward solution is to distribute the excitation list among multiple computing nodes [59]. Indeed, several recent efforts have focused on developing a distributed CI framework, albeit at the cost of utilizing several hundreds of high-memory computing nodes [95, 146, 147, 302]. Nonetheless, this approach does not alleviate the demand for large-scale CI applications that can be executed on conventional workstations or small computing clusters.

Minimizing the memory needs for extensive CI calculations on a single workstation or laptop, while simultaneously enhancing scalability for massively parallel computing, may appear as divergent goals. Yet, they form the central objective of this work, and we demonstrate that optimizations for both small-scale and large-scale resources can complement each other. Additionally, we aim to establish a versatile framework capable of supporting various CI methods such as CAS, RAS, GAS, MRCI, and MRPT2, applicable to both relativistic and non-relativistic theories.

In this work, we advance CI methods by introducing a novel small-tensor-product (STP) addressing scheme that dramatically reduces storage requirements while being fully compatible with distributed computing environments. This advancement leverages the proposed distributed active space (DAS) framework, which separates inter-space and intra-space excitations in the loop structure, as well as global and local phase factors in the symbolic matrix elements. Through a carefully designed three-tiered addressing structure, the STP-DAS framework eliminates the necessity of storing a global excitation list. Instead, it adopts a small-tensor-product algorithm using a localized addressing strategy and is tailored for efficient vectorization and parallel processing. The hallmark of this research is the STP-DAS framework’s seamless support for large-scale CI calculations, applicable across both single-node systems with constrained memory resources and expansive, distributed super-

computing infrastructures.

2.2 Distributed Active Space Configuration Interaction

2.2.1 Statement of the Problem

As stated above the σ vector formation is the most expensive part of a CI calculation. Eq. (1.40) can be readily partitioned into one-electron and two-electron contributions,

$$\begin{aligned}\sigma_K &= {}^{1e}\sigma_K + {}^{2e}\sigma_K \\ {}^{1e}\sigma_K &= \sum_{pq} \sum_L h'_{pq} \langle K | \hat{E}_{pq} | L \rangle C_L \\ {}^{2e}\sigma_K &= \frac{1}{2} \sum_{pqrs} \sum_{L,J} g_{pqrs} \langle K | \hat{E}_{pq} | J \rangle \langle J | \hat{E}_{rs} | L \rangle C_L\end{aligned}\quad (2.1)$$

The bulk of the computational effort comes from the two-electron part of Eq. (2.1). As such the efficient construction of the two-electron contribution to the σ vector formation has been the focus of several studies over the years primarily for nonrelativistic Hamiltonians [150,212,271,318]. The emergent state of the art methods for large-scale CI problems from the decades of effort can be characterized as direct methods which generate matrix elements on-the-fly from one-electron excitation lists $\langle K | \hat{E}_{pq} | J \rangle$ with differing orders of contraction in Eq. (2.1). However, for extremely large CI calculations distributed over many nodes the storage of the excitation lists becomes unmanageable.

In this work, we aim to reformulate the structure of the loops in Eq. (2.1) by factorization into many small tensor products and considering only nonzero combinations – which is henceforth referred to as a tensor-looping algorithm – facilitated by a distributed active space partitioning scheme. The goal is to identify a partitioning of the orbital space which allows us to circumvent the storage of extremely large excitation lists by writing the two-electron contribution solely in terms of small tensor products local to the partitioned space. If such a partitioning is possible, one would be able to reuse excitation lists between spaces and additionally store much smaller excitation lists. Crucially, we seek a partitioning which is *exact* and does not introduce any approximations relative to the full space when all excitations between the partitioned space are allowed.

In the CI community, ORMAS [128, 129] and GAS are common orbital partitioning schemes [85–87, 146, 147, 185, 213, 301]. These partitioning schemes differ based on their occupation schemes. ORMAS bounds the minimum and maximum occupation for each space, and GAS places bounds on the accumulated electron occupation number for each successive orbital space. These methods have been successful when forming approximate full CI spaces, however neither partitioning scheme provides a computational advantage when exactly partitioning a full CI space, *i.e.*, when there is no limit on the interspace excitations.

As we will demonstrate in the following sections, the STP-DAS framework can provide a computational advantage even in the limit of full interspace excitations, by breaking the sigma build into small tensor products. The following sections are technical but lay out important foundations of the STP-DAS framework for CI-based calculations. Readers who are primarily interested in the final STP-DAS σ build expressions may skip directly to Section 2.2.7.

2.2.2 The Distributed Active Space Framework

In the DAS framework, the total correlation space is defined with a collection of active orbitals $\{\phi\}$ whose cardinality is $M = |\{\phi\}|$ (*i.e.*, the total number of orbitals), a total number of active electrons n_e , and the determinants $|K\rangle$ generated either by the CAS or MRCI method.

Analogous to the generalized active space (GAS) approach, the total correlation space can be partitioned into an arbitrary number of distributed active spaces (DASs). In the absence of excitation constraints between these spaces, DAS becomes equivalent to CAS. On the other hand, imposing limits on the number of electrons or holes within each DAS creates a situation reminiscent of RAS, ORMAS, and GAS.

In the CI σ build, the global address of each K is uniquely defined. Addressing each determinant in CAS can be efficiently done with a string-based method, leading to a highly vectorized algorithm. A similar string-based addressing scheme has been extended to RAS with excitation restrictions [121]. However, for large CAS calculations, it is not feasible to

save the complete global address with the excitation list. An efficient addressing scheme is critical for the large-scale parallelization of CI codes, and reducing the memory demands for the excitation list is essential for their broad practical use. Our strategy to accomplish these objectives involves fully separating global and local addresses to enable small tensor products in the σ build.

Figure 2.1 illustrates the STP-DAS mapping scheme, using CAS as example. In this illustration, a CAS problem is mapped onto multiple DASs. In contrast to the RAS addressing scheme where excitation restrictions are enforced in the map that generates strings [121], the first step in the STP-DAS framework involves differentiating between excitations within a space (intra-space) and those between different spaces (inter-space). This separation is achieved through a process known as categorical excitations, which creates various configuration categories by promoting electrons from one DAS to another, following specific excitation constraints. This process is shown schematically in Figure 2.1.

For CAS, there are no limitations on inter-space excitations. Imposing constraints on these excitations transforms the approach into RAS, ORMAS, GAS, or MRCI methods, thus rendering STP-DAS a versatile structure suitable for a broad spectrum of multiconfigurational techniques. Following the separation of inter-space and intra-space excitations, each DAS is treated as an intra-space complete active space. This approach incorporates every possible excitation, making it compatible with the efficient string-based addressing scheme for computational processing, allowing for small tensor products embedded in a tensor looping algorithm.

The STP-DAS framework does not result in decreased accuracy as the total number of excitations, and thus the total number of determinants, remains unchanged. However, the advantage of STP-DAS comes from the fact that it is only necessary to store strings for intra-space excitations used in small tensor products and tensor looping. This means that when two DASs share the same configuration, even if they belong to different categories or consist of different orbitals, they can utilize the same intra-space excitation list and small tensor products. For instance, in CAS calculations involving 40 spinor orbitals and 20 electrons, $\binom{40}{20}$, 637 TB of memory would be required to store the excitation list. By adopting the DAS approach illustrated in Figure 2.1, where the correlation space is segmented into four DASs

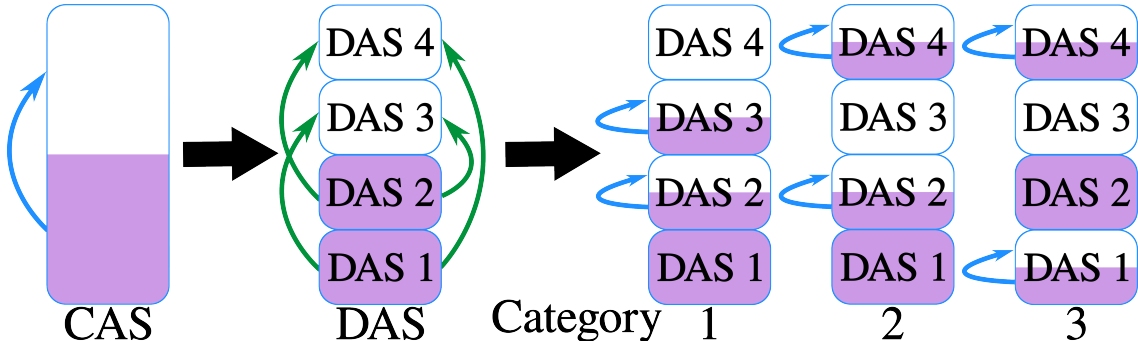


Figure 2.1. Space Partitioning in STP-DAS. Boxes symbolize active regions. Within these boxes, shaded portions indicate electron occupations, while white areas represent empty orbitals. A CAS problem is mapped onto multiple DASs. Distinct categories of configurations are produced by categorical excitations, depicted by green arrows. In every category, each DAS is treated as a complete active space, including all possible excitations, illustrated by blue arrows.

each containing 10 orbitals, the memory requirement is dramatically reduced. In this case, one only needs to store intra- and inter-space excitation lists for $\binom{10}{1}$, $\binom{10}{2}$, \dots , $\binom{10}{9}$ CAS strings, which collectively only require 200 MB of memory. Note that any type of space partition will reduce the total number of strings generated from the smaller combinatorics. However, we advocate for equal space partitioning to maximize the reduction and take advantage of the reuse of the excitation list (see Section 3.4 for a detailed discussion).

The discussion above highlights the efficiency of the STP-DAS framework in managing computational resources through the reuse of intra-space excitation lists in small tensor products. Nonetheless, the process of identifying and reutilizing these lists is a challenging task. The complexity arises because, within the excitation list, the symbolic matrix elements vary across different DASs, even when they utilize identical local addressing strings. This complexity presents a significant challenge and restricts the application of the RAS/GAS approach to merely imposing excitation constraints, rather than achieving a reduction in computational demands.

2.2.3 Space Partitioning and Organization of Excitations

In this section, the following notations are used, unless otherwise specified:

- $\mu, \nu, \kappa, \lambda$: DAS indices

- $\mathcal{A}, \mathcal{B}, \mathcal{C}$: configuration category
- $|K\rangle, |L\rangle$: determinants in the full space
- $|\mathbb{K}\rangle, |\mathbb{L}\rangle$: sub-determinants in DAS
- p, q, r, s : general MOs

Distributed Active Space

A DAS, $\mathbb{X}_\mu \equiv \{\{\phi\}_\mu, n_\mu, \mathbb{K}_\mu\}$, is defined by a set of orbitals $\{\phi\}_\mu$, an electron occupation number n_μ in the space, and sub-determinants \mathbb{K}_μ . Each DAS is a complete active space where the total number of determinants is given by

$$N_\mu = \binom{M_\mu}{n_\mu} \quad (2.2)$$

where $M_\mu = |\{\phi\}_\mu|$ is the cardinality of the orbital set in \mathbb{X}_μ , *i.e.*, the number of orbitals. Here, we assume the Kramers' unrestricted two-component or four-component condition where each orbital is a singly occupied spinor.

It's important to emphasize that within the STP-DAS framework, intra-space excitations are not subject to any constraints, as the excitation types typically associated with RAS or MRCI, known as inter-space excitations, are separated and utilized for the creation of configuration categories.

Configuration Category

Applying inter-space excitations, also referred to as categorical excitations, on the reference electron configurations leads to the formation of different electron configuration categories (see Figure 2.1 for an example). These categories will from now on be referred to as either configuration categories or simply as categories.

A configuration category $\mathcal{A} = \{\mathbb{X}_\mu^{\mathcal{A}}, \mathbb{X}_\nu^{\mathcal{A}}, \dots\}$ is constructed from a set of DASs with unique electron occupation numbers. Note that the DAS orbital partitioning is unchanged. Therefore, it is important to associate each DAS with its parent category through the

notation, $\mathbb{X}_\mu^{\mathcal{A}}$. Categorical excitation operators are defined as:

$$\begin{aligned}\hat{\mathcal{E}}_{pq}^{\mathbb{X}_\mu^{\mathcal{A}}\mathbb{X}_\nu^{\mathcal{A}}} &= \{\hat{E}_{pq} : p \in \mathbb{X}_\mu, q \in \mathbb{X}_\nu\} \\ \hat{\mathcal{E}}_{pqrs}^{\mathbb{X}_\mu^{\mathcal{A}}\mathbb{X}_\nu^{\mathcal{A}}\mathbb{X}_\lambda^{\mathcal{A}}\mathbb{X}_\kappa^{\mathcal{A}}} &= \{\hat{e}_{pqrs} : p \in \mathbb{X}_\mu, q \in \mathbb{X}_\nu, r \in \mathbb{X}_\lambda, s \in \mathbb{X}_\kappa\} \\ &\dots\end{aligned}$$

Figure 2.1 shows three examples of different categories generated from categorical excitations.

A full-space determinant in category \mathcal{A} can be constructed using sub-determinants from each DAS space

$$|K^{\mathcal{A}}\rangle = |\mathbb{K}_\mu^{\mathcal{A}}\rangle \oplus |\mathbb{K}_\nu^{\mathcal{A}}\rangle \oplus \dots \quad (2.3)$$

The total number of determinants in category \mathcal{A} is:

$$N^{\mathcal{A}} = N_\mu^{\mathcal{A}} \cdot N_\nu^{\mathcal{A}} \dots \quad (2.4)$$

Considering the definition of a configuration category and its constituent DASs, it becomes apparent how they relate to the overall number of determinants (N), the total electron count (n_e), and the complete set of orbitals ($\{\phi\}$) in the entire correlation space:

$$N = \sum_{\mathcal{A}} N^{\mathcal{A}} \quad (2.5)$$

$$n_e = \sum_{\mu} n_\mu^{\mathcal{A}} \quad (2.6)$$

$$\{\phi\} = \{\phi\}_\mu \cup \{\phi\}_\nu \cup \dots \quad (2.7)$$

It's important to note that the notation for categories is not applied in the DAS orbital representation, ϕ_μ , since every category employs the same scheme for orbital partitioning.

2.2.4 Small-Tensor Addressing in String-Based DAS

The objective is to develop an STP addressing algorithm that exclusively utilizes and reuses local address strings, thereby avoiding the need for explicitly constructing and storing the complete global list of excitations. In the STP-DAS framework, a three-level addressing protocol is adapted from Kozłowski and Pulay's RAS addressing scheme [157].

For a determinant K that belongs to category \mathcal{B} , its global address can be defined as

$$A(K) = A(\mathcal{B}) + A(K^{\mathcal{B}}) \quad (2.8)$$

Here, we define A as an address function. $A(\mathcal{B})$ returns the global offset for all determinants in category \mathcal{B} . It can be easily calculated as

$$A(\mathcal{B}) = \sum_{\mathcal{A} < \mathcal{B}} N^{\mathcal{A}} \quad (2.9)$$

which naturally arises from Eq. (2.5). Figure 2.2 illustrates the address mapping strategy utilized in the STP-DAS framework. In order to compute the global offset for each category, it is necessary to systematically organize all categories for efficient record-keeping. $A(K^{\mathcal{B}})$ returns the address of the determinant K in category \mathcal{B} .

Eq. (2.8) serves as the outer loop in constructing the STP-DAS σ build, where it is also possible to achieve effective load balancing across categories in a parallel computing environment. Although Eq. (2.8) successfully divides global addresses into more manageable categorical address sections, it does not reduce the overall size of the excitation list. This is due to the uniqueness of each category, requiring that each $A(K^{\mathcal{B}})$ be stored individually.

From Eq. (2.3), we can effectively write the categorical address $A(K^{\mathcal{B}})$ via separations of active and inactive DASs, based on its usage in the σ build. Since only the one-electron excitation list is saved, we only need to consider local addresses associated with

$$\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle, \quad p \in \mathbb{X}_{\mu}^{\mathcal{A}}, \quad q \in \mathbb{X}_{\nu}^{\mathcal{B}} \quad (2.10)$$

Only two scenarios are possible:

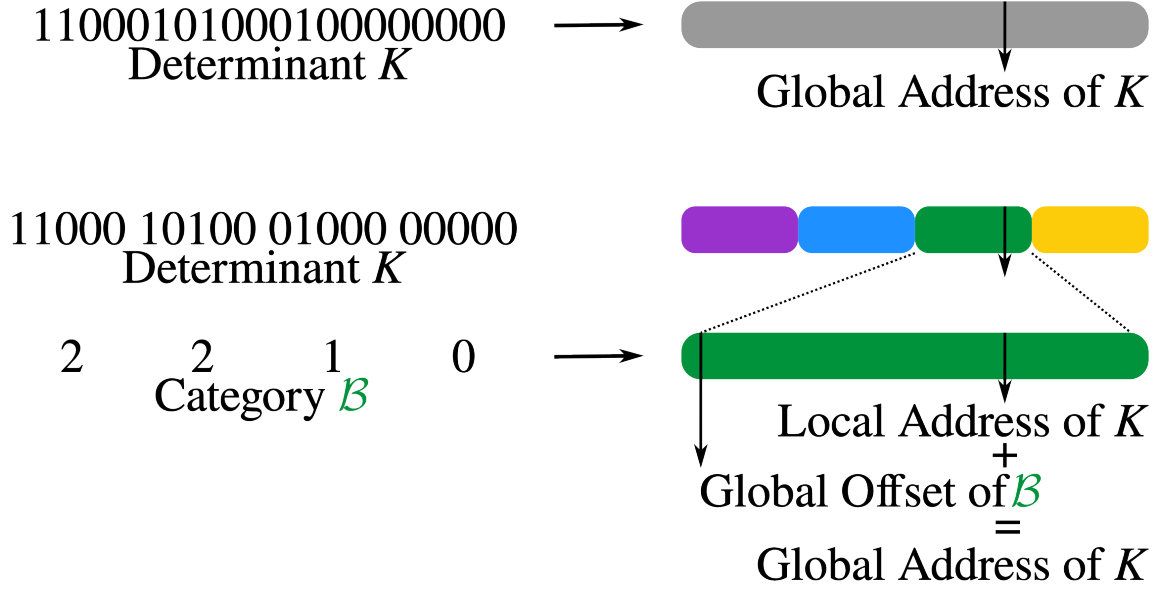


Figure 2.2. A comparison between a traditional indexing scheme and the indexing scheme used in STP-DAS. In DAS, the global address of a determinant is found by combining a global offset of the category of a given determinant with the local index of that determinant within the category. Boxes symbolize determinant categories.

- $\mu = \nu$, *i.e.*, bracket excitations between an identical DAS in different categories.

$$\begin{aligned}
 A(L^{\mathcal{A}}) &= \dots + A(\mathbb{L}_{\mu}^{\mathcal{A}}) \cdot \prod_{\kappa < \mu} N_{\kappa}^{\mathcal{A}} + \dots \\
 A(K^{\mathcal{B}}) &= \dots + A(\mathbb{K}_{\mu}^{\mathcal{B}}) \cdot \prod_{\kappa < \mu} N_{\kappa}^{\mathcal{B}} + \dots
 \end{aligned} \tag{2.11}$$

where “...” refers to concurrent looping over addresses in DASs that are not associated with the excitations, *i.e.* inactive DASs, for both $A(L^{\mathcal{A}})$ and $A(K^{\mathcal{B}})$.

- $\mu \neq \nu$, *i.e.*, bracket excitations between two different DASs in different categories.

$$\begin{aligned}
A(L^{\mathcal{A}}) &= \cdots + A(\mathbb{L}_{\mu}^{\mathcal{A}}) \cdot \prod_{\kappa < \mu} N_{\kappa}^{\mathcal{A}} + \cdots \\
&\quad + A(\mathbb{L}_{\nu}^{\mathcal{A}}) \cdot \prod_{\kappa < \nu} N_{\kappa}^{\mathcal{A}} + \cdots \\
A(K^{\mathcal{B}}) &= \cdots + A(\mathbb{K}_{\mu}^{\mathcal{B}}) \cdot \prod_{\kappa < \mu} N_{\kappa}^{\mathcal{B}} + \cdots \\
&\quad + A(\mathbb{K}_{\nu}^{\mathcal{B}}) \cdot \prod_{\kappa < \nu} N_{\kappa}^{\mathcal{B}} + \cdots
\end{aligned} \tag{2.12}$$

Eq. (2.11) and Eq. (2.12) efficiently separate out the inactive DASs linked to the excitation list in $\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle$. For single-electron excitations, the local addresses within the inactive DASs must be identical for both bra and ket. This requirement facilitates an efficient tensor looping design, where a single outer loop iterates over the addresses of determinants within the inactive DASs, enabling concurrent tensor looping for both bra and ket. The address structure of Eq. (2.11) and Eq. (2.12) allows the inner tensor loop to use only the local addresses of sub-determinants $A(\mathbb{L}_{\mu}^{\mathcal{A}})$ and $A(\mathbb{K}_{\mu}^{\mathcal{B}})$, leading to the small-tensor-product algorithm.

It is readily seen that when two DASs share the same configuration, characterized by an identical number of orbitals and electrons, they can utilize the same set of small-tensor address strings. This is true even if the orbitals and electrons between the two DASs differ in physical character. This principle is fundamental to the shift from a single, large global excitation list to numerous, smaller, shared local excitation lists.

In adopting this approach to streamline determinant addressing, a new challenge emerges. The CI excitation list features a symbolic matrix $\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle$, where each element is determined by the global positions of the orbitals involved in the excitation operator. This global dependency poses a conflict with the local addressing framework. To address this, the subsequent section will present an algorithm designed to differentiate between global and local phase factors during the construction of the symbolic matrix $\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle$, enabling the storage of solely local excitation lists in small tensor products.

2.2.5 Local One-Electron Excitation List

To account for the sparsity of the Hamiltonian, a one-electron excitation list that contains the information of non-zero elements of $\langle L | \hat{E}_{tu} | K \rangle$ is pre-computed and stored in memory. Only elements that are non-zero in the excitation list are summed in Eq. (1.40). Based on the addressing scheme shown in Eqs. (2.10) to (2.12), the local excitation list can be defined, and the symbolic matrix element $\langle \mathbb{L}_\mu^A | \hat{E}_{pq} | \mathbb{K}_\nu^B \rangle$ can be computed as

- $\mu = \nu$:

$$\left\{ \begin{array}{l} A(\mathbb{L}_\mu^A) \\ p \in \mathbb{X}_\mu^A \\ q \in \mathbb{X}_\nu^B \\ A(\mathbb{K}_\nu^B) \\ \langle \mathbb{L}_\mu^A | \hat{E}_{pq} | \mathbb{K}_\nu^B \rangle \end{array} \right\} \quad (2.13)$$

In this intra-space one-electron excitation scenario, for each $|\mathbb{K}_\nu^B\rangle$, there are $n_\mu^- (n_\mu^+ + 1)$ number of non-zero elements, where n_μ^- and n_μ^+ are the number of electrons (occupied orbitals) and holes (unoccupied orbitals), respectively, in DAS $\mathbb{X}_\mu = \mathbb{X}_\nu$. The non-zero element is computed as

$$\begin{aligned} & \forall p > q : \langle \mathbb{L}_\mu^A | \hat{E}_{pq} | \mathbb{K}_\nu^B \rangle \\ & = \begin{cases} +1, & \text{if } \left(\sum_{i=q+1}^{p-1} b_i \right) \text{ is even} \\ -1, & \text{if } \left(\sum_{i=q+1}^{p-1} b_i \right) \text{ is odd} \end{cases} \end{aligned} \quad (2.14)$$

where $b_i = 0$ or 1 is the electron occupancy.

- $\mu \neq \nu$: This inter-space one-electron excitation list involves four sub-determinants in different DASs,

$$\left\{ \begin{array}{l} A(\mathbb{L}_\mu^A), A(\mathbb{L}_\nu^A) \\ p \in \mathbb{X}_\mu^A \\ q \in \mathbb{X}_\nu^B \\ A(\mathbb{K}_\mu^B), A(\mathbb{K}_\nu^B) \\ \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B \rangle \end{array} \right\} \quad (2.15)$$

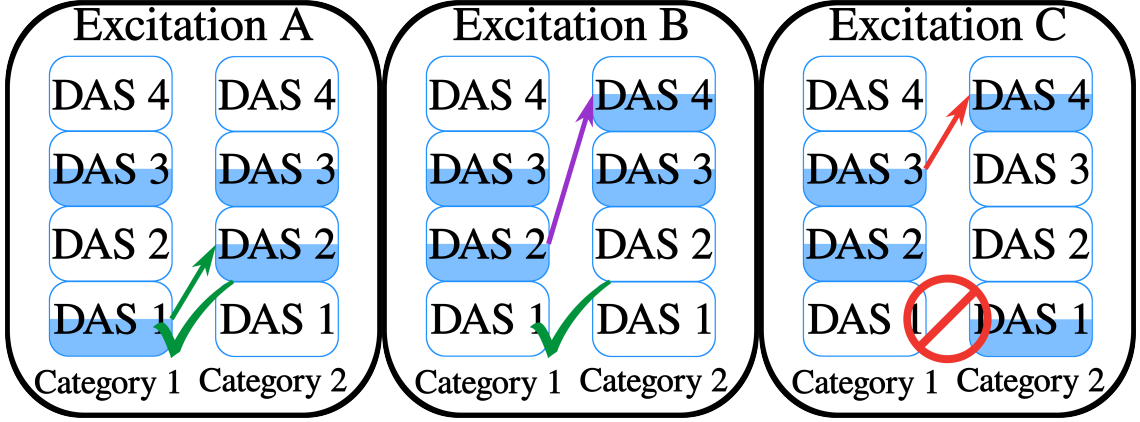


Figure 2.3. Examples of excitations between categories. Boxes symbolize active regions. Within these boxes, shaded portions indicate electron occupations, while white areas represent empty orbitals. All partially filled DASs have an identical number of electrons and holes.

The non-zero one-electron symbolic matrix elements can be computed as:

$$\begin{aligned}
 \forall \mu < \nu : & \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B \rangle \\
 = & \begin{cases} +1, & \text{if } \left(\sum_{i=p+1}^{M_\mu} b_i + \sum_{j=0}^{q-1} b_j \right) \text{ is even} \\ -1, & \text{if } \left(\sum_{i=p+1}^{M_\mu} b_i + \sum_{j=0}^{q-1} b_j \right) \text{ is odd} \end{cases} \quad (2.16)
 \end{aligned}$$

Eq. (2.13) and Eq. (2.15) define the one-electron excitation lists in the context of STP-DAS using local addresses. These lists are free of global addresses, enabling their reuse across excitations with identical DAS configurations. As illustrated in Figure 2.3, the green and purple excitations (excitations A and B), although associated with different orbitals and electrons, have identical DAS configurations, allowing them to share the same local excitation lists used in the small tensor product. This commonality highlights that dividing the correlation space evenly among numerous DASs can significantly reduce the number of unique local excitation lists. Achieving this efficiency is a principal objective in the conceptualization of STP-DAS.

2.2.6 Global Phase Factor in the Small Tensor Product algorithm

Equipped with determinantal addresses and one-electron excitation lists completely defined in STP-DAS, the CI σ build can be ideally performed with small tensor products using local excitation lists. Nevertheless, two critical aspects need to be considered for enabling the σ matrix assembly with small tensor products using these local lists.

One issue is distinguishing between different types of excitations, such as excitations A and B, as shown Figure 2.3, during the σ matrix construction. This process requires global symbolic matrix elements, which are derived from the global positioning of the orbitals. Another challenge is identifying invalid excitations, such as excitation C in Figure 2.3. The excitation lists mainly provide local DAS specifics, leading to a loss of global context. However, STP-DAS's architecture facilitates the straightforward retrieval of global address-dependent information during the CI σ matrix construction, resolving both issues.

The distinction between excitations A and B in Figure 2.3 stems from a difference in the phase factor between the global $\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle$ and local $\langle \mathbb{L}_{\mu}^{\mathcal{A}} \oplus \mathbb{L}_{\nu}^{\mathcal{A}} | \hat{E}_{pq} | \mathbb{K}_{\mu}^{\mathcal{B}} \oplus \mathbb{K}_{\nu}^{\mathcal{B}} \rangle$ symbolic matrix element. To address this, we define a global phase factor $P_{\mu\nu}$ as:

$$p \in \mathbb{X}_{\mu}^{\mathcal{A}}, q \in \mathbb{X}_{\nu}^{\mathcal{B}}, \forall \mu \leq \nu :$$

$$P_{\mu\nu} = \begin{cases} +1, & \text{if } \left(\sum_{\mathbb{X}_{\kappa}=\mathbb{X}_{\nu+1}}^{\mathbb{X}_{\nu-1}} n_{\kappa}^{-} \right) \text{ is even} \\ -1, & \text{if } \left(\sum_{\mathbb{X}_{\kappa}=\mathbb{X}_{\nu+1}}^{\mathbb{X}_{\nu-1}} n_{\kappa}^{-} \right) \text{ is odd} \end{cases} \quad (2.17)$$

where n_{κ}^{-} is the number of electrons in DAS \mathbb{X}_{κ} . $P_{\mu\nu}$ takes into account the contributions of DASs that are not directly involved in the excitation to the symbolic matrix element.

In Figure 2.3, excitation C is an invalid excitation because it cannot be defined with one-electron excitation lists. To determine the validity of an excitation, $\langle L^{\mathcal{A}} | \hat{E}_{pq} | K^{\mathcal{B}} \rangle$, we need to check if the non-excitation DASs between categories are identical. Here, we introduce a Kronecker δ function $\delta_{\bar{\mathbb{X}}_{\mu}^{\mathcal{A}} \bar{\mathbb{X}}_{\nu}^{\mathcal{B}}}$ where $\bar{\mathbb{X}}_{\mu}^{\mathcal{A}}$ refers to all but \mathbb{X}_{μ} DASs in category \mathcal{A} . The δ function is non-zero only when non-excitation DASs between categories are identical.

2.2.7 σ Build using Small Tensor Products

Taking everything discussed above into consideration, we now have the final working expressions for the CI σ build using the STP algorithm with only local excitation lists:

$$\sigma_{L^A} = {}^{1e}\sigma_{L^A} + {}^{2e}\sigma_{L^A} \quad (2.18)$$

$${}^{1e}\sigma_{L^A} = \sum_{\mathcal{B}} \sum_{\mathbb{K}_\mu^{\mathcal{B}} \oplus \mathbb{K}_\nu^{\mathcal{B}}} \sum_{pq} P_{\mu\nu} \delta_{\mathbb{X}_{\mu\nu}^A, \mathbb{X}_{\mu\nu}^{\mathcal{B}}} h'_{pq} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{K}_\mu^{\mathcal{B}} \oplus \mathbb{K}_\nu^{\mathcal{B}} \rangle C_{K^{\mathcal{B}}} \quad (2.19)$$

$${}^{2e}\sigma_{L^A} = \frac{1}{2} \sum_{\mathcal{BC}} \sum_{\mathbb{J}_\mu^{\mathcal{B}} \oplus \mathbb{J}_\nu^{\mathcal{B}}} \sum_{\mathbb{J}_\kappa^{\mathcal{B}} \oplus \mathbb{J}_\lambda^{\mathcal{B}}} \sum_{\mathbb{K}_\kappa^{\mathcal{C}} \oplus \mathbb{K}_\lambda^{\mathcal{C}}} \sum_{pqrs} P_{\mu\nu} P_{\kappa\lambda} \delta_{\mathbb{X}_{\mu\nu}^A, \mathbb{X}_{\mu\nu}^{\mathcal{B}}} \delta_{\mathbb{X}_{\kappa\lambda}^{\mathcal{B}}, \mathbb{X}_{\kappa\lambda}^{\mathcal{C}}} g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^{\mathcal{B}} \oplus \mathbb{J}_\nu^{\mathcal{B}} \rangle \langle \mathbb{J}_\kappa^{\mathcal{B}} \oplus \mathbb{J}_\lambda^{\mathcal{B}} | \hat{E}_{rs} | \mathbb{K}_\kappa^{\mathcal{C}} \oplus \mathbb{K}_\lambda^{\mathcal{C}} \rangle C_{K^{\mathcal{C}}} \quad (2.20)$$

where $p \in \mathbb{X}_\mu^A$, $q \in \mathbb{X}_\nu^{\mathcal{B}}$, $r \in \mathbb{X}_\kappa^{\mathcal{B}}$, $s \in \mathbb{X}_\lambda^{\mathcal{C}}$.

Eqs. (2.18) to (2.20) are completely formulated in terms of small tensor products using local DAS one-electron excitation lists. These expressions can be adapted to any σ build algorithm in any CI-based method. In this chapter, we used the Knowles-Handy algorithm [150] for the categorical σ build. The outer loop is over $\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{B}}$, and maximum sparsity is utilized to avoid large intermediates.

2.3 Computational Details

All calculations in this chapter are performed with a development version of the Chronus Quantum software package [308] with the STP-DAS framework. The speed of light utilized in this study is 137.035999074 a.u. All calculations utilized the standard Gaussian nuclear model [299]. Relativistic calculations are done in the Kramers' unrestricted exact-two-component (X2C) framework where all spinor orbitals are singly occupied. The one-electron X2C transformation is a one-electron-Hamiltonian-based one-step procedure that "folds" small component wave function into a pseudo-large component so that the four-component Dirac equation becomes an effective two-component eigenfunction prob-

lem. [64–67, 69, 70, 100, 119, 126, 155, 163, 171, 177, 178, 180, 181, 183, 221, 222, 260] The one-electron X2C approach makes use of the effective one-electron spin–orbit Hamiltonian and avoids the four-component self-consistent-field procedure. In this work, we use the new Dirac–Coulomb–Breit-parameterized effective one-electron spin–orbit Hamiltonian in the X2C approach [72].

STP-DAS stands as a scalable, high-performance CI framework supporting various electronic structure theories – including CAS, RAS, GAS, ORMAS, MRCI, and MRPT2 – in both relativistic and non-relativistic domains. The fundamental characteristics and dimensionalities of these electronic structure methods remain unchanged within this framework. In the current work, we benchmark the performance and analyze the memory requirement of the DAS framework applied to X2C-CASCI calculations.

To benchmark the STP-DAS framework, we use molecular thallium monohydride (TlH), the heaviest stable monohydride species observed experimentally [292, 303]. An accurate electronic structure characterization of TlH requires the use of a relativistic many-body approach [81, 254]. In this study, we concentrate on evaluating the algorithm’s performance as the correlation space expands. Furthermore, we investigate the load balance among nodes within a high-performance computing environment, as well as the reduction in memory demands when computing resources are constrained, such as on a laptop. In the following benchmark calculations, the Tl atom used the x2c-TZVPall-2c basis set [227] and the H atom used the aug-cc-pVTZ basis set [63, 138].

When forming Slater determinants in the Kramers’ unrestricted two-component or four-component no-virtual-pair CASCI method, the number of possible determinants is given by

$$N_{\text{CAS}} = \binom{M_{\text{spinor}}}{n_e} \quad (2.21)$$

where M_{spinor} is the number of active spinor orbitals. Since the spin-symmetry is no longer enforced, for a same number of active electrons and molecular orbitals, the CI dimension in the Kramers’ unrestricted relativistic CAS is much bigger than that in the non-relativistic calculation. In relativistic computations, the floating point operations (FLOP) count for

constructing the σ matrix experiences a sixfold increase, arising from complex-valued arithmetic, compared to a non-relativistic (NR) calculation.

The calculation of the memory requirement in all test cases takes into account both the sparsity of the excitation list, *i.e.*, non-zero elements only, and bit-wise compression of the determinant address, representing the minimum requirement in a CAS calculation. For X2C-CASCI calculations, the size of this list (in Bytes) can be calculated as $n_{\text{save}} \times n_e \times (n_h + 1) \times N_{\text{CAS}}$, where n_e is the number of electrons and n_h is the number of holes in the complete active space. Here, n_{save} represents the amount of information saved in the excitation list, including the address of small tensors, p , p , and the phase factor. Depending on the level of bit-wise compression, n_{save} varies from 4 to 11 Bytes for most applications.

2.4 Results and Discussion

2.4.1 A Large-Scale X2C-CASCI Calculation on a Laptop

A significant benefit of the STP-DAS framework lies in its capacity to facilitate large-scale CI computations using constrained computational resources, like a laptop. This was illustrated through the performance of X2C-CAS calculations on the TIH system, utilizing the STP-DAS framework on an Apple M3 Max laptop equipped with 14 compute cores and 128 GB of RAM.

Table 2.1. The storage needs for the one-electron excitation list in a $\binom{44}{36}$ X2C-CASCI calculation (Apple M3 Max laptop with 14 compute cores and 128 GB of RAM). The total number of X2C determinants is 177×10^6 , which is comparable to the computational cost of 1.1×10^9 determinants in non-relativistic calculations when measured by the number of FLOPs.

# of DASs (Orbital Partitions)	Excitation List Storage ^a	σ Build Time
1 (44)	402. GB	–
2 (22,22)	210. GB	–
6 (7,7,7,7,9)	7.91×10^{-3} GB	2981 s
9 (5,5,5,5,5,5,4)	9.06×10^{-5} GB	5788 s

^a Only non-zero elements are considered in the one-electron excitation list. Bit-wise compression of the determinant address is used.

Table 2.1 illustrates the STP-DAS framework’s capability to reduce storage demands, thereby enabling large-scale CI calculations on a laptop. For a $\binom{44}{36}$ X2C-CASCI calcula-

tion with STP-DAS, the resulting 177×10^6 X2C determinants (1.1×10^9 non-relativistic determinants equivalent in terms of FLOP count) would require 402 GB of RAM to hold the excitation list with non-zero elements. This requirement makes the task impractical for a personal laptop and many smaller workstations.

As shown in Table 2.1, partitioning the correlation space into multiple DASs significantly reduces the storage needs. Once the number of DASs hits 6, the storage demands for one-electron excitations become minimal, making it feasible to conduct the X2C-CASCI calculations on a laptop with ease. Each σ build only takes 0.8 hours on the latest Apple laptop for a 177×10^6 determinant X2C-CASCI calculation.

With the introduction of additional STP-DAS for space partitioning using small tensor products, the storage requirement continues to decrease. However, this leads to an increase in the time required to build the σ vector. The primary reason for this is the increased overhead associated with extensive tensor looping necessary to locate each local address within the global framework, as indicated in Eq. (2.11) and Eq. (2.12). This analysis reveals that while STP-DAS is effective in reducing memory requirements, excessively fine partitioning might result in added costs associated with small tensor mapping from local to global addresses.

2.4.2 A Large-Scale X2C-CASCI Calculation on a Supercomputer

In high-performance computing, especially for large CI calculations on a supercomputer, it's crucial to ensure that the workload is evenly distributed across all compute nodes. This is where the STP-DAS framework is particularly effective. The STP-DAS framework is capable of decreasing memory demands while concurrently leveraging distributed memory on a massively parallel high-performance computing system.

As illustrated in Table 2.2 with a $\binom{44}{29}$ X2C-CASCI calculation example, the STP-DAS approach achieves memory reduction by employing localized excitation lists within the distributed active space. Expanding the X2C-CASCI of $\binom{44}{29}$ results in over 230 billion determinants (1.4×10^{12} non-relativistic determinants equivalent in terms of FLOP count), requiring 1,173 terabytes (TB) of memory to maintain the one-electron excitation list within

the conventional CAS setup. Assuming that each compute node allocates 0.5 TB of memory for this purpose, it would require over 2,346 compute nodes with distributed memory systems to perform a CI calculation. A configuration with 4 DASs can lower the memory needs to a manageable 1.7 GB for a small distributed computing system.

# of DASs (Orbital Partitions)	Excitation List Storage ^a	# of Categories
1 (44)	1,173 TB	1
2 (22,22)	563 TB	16
4 (11,11,11,11)	2 GB	736
6 (7,7,7,7,7,9)	1×10^{-2} GB	11,292
8 (6,6,6,6,6,6,2)	4×10^{-4} GB	80,823
9 (5,5,5,5,5,5,5,4)	9×10^{-5} GB	260,656

Table 2.2. The storage needs for the one-electron excitation list in a $\binom{44}{29}$ X2C-CASCI calculation. The total number of X2C determinants is 230×10^9 , which is comparable to the computational cost of 1.4×10^{12} determinants in non-relativistic calculations when measured by the number of FLOPs.

^a Only non-zero elements are considered in the one-electron excitation list. Bit-wise compression of the determinant address is used.

In the STP-DAS framework, we allocate workloads among the compute nodes according to different configuration categories. As we introduce more categories, we can achieve a more balanced distribution of the total workload for constructing the σ matrix. The effectiveness of the STP-DAS framework is illustrated in Table 2.2, which demonstrates a rapid increase in the number of configuration categories as additional DASs are introduced into the system for space and tensor partitioning. This expansion significantly enhances the efficiency of the σ matrix construction. A figure of merit for computational work distribution is the percent difference of the median number of determinants per node relative to the theoretical ideal mean distribution of determinants. The observed distribution of determinants to each node in a 600-node calculation is illustrated in Figure 2.4. The 8 DAS distribution is markedly more extended with a much longer tail. This results in a few nodes with many more determinants than all of the other nodes shifting the median number of determinants further from the ideal distribution. When increasing the number of DASs from 8 to 9, the percent difference in the work distribution drops from 4% to 0.15% resulting in a nearly ideal division of determinants among the nodes. With the resulting balanced workload, each σ build for the $\binom{44}{29}$ X2C-CASCI calculation with 9 DASs containing 230×10^9 determinants

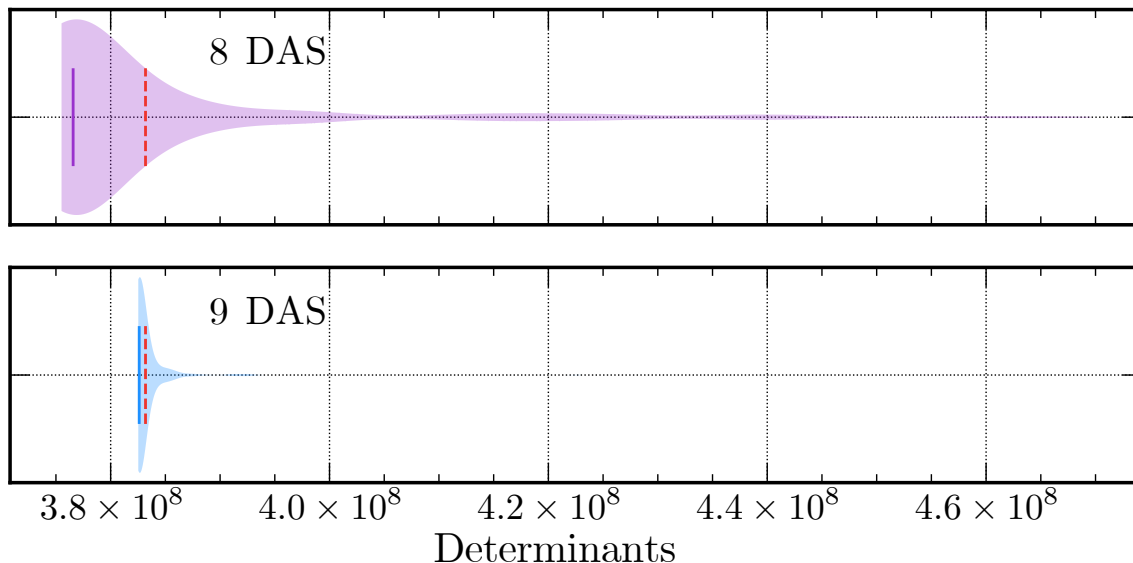


Figure 2.4. Distribution of determinants per node for a $\binom{44}{29}$ X2C-DASCI calculation using 600 nodes with (top) 8 DAS spaces and (bottom) 9 DAS spaces. The theoretical ideal distribution of determinants is denoted by the red line in both plots, and the solid line denotes the median of the observed distributions.

(1.4×10^{12} NR equivalent determinants) only took 7 hours. This calculation was run on the Department of Energy’s Perlmutter high-performance super computer with a total of 16000 compute cores (AMD EPYC 7763 Milan, 200 GB/s NIC, 1 MPI per node and 16 SMP threads per MPI process). This analysis shows the scalability and efficiency improvements in computational performance that can be realized through optimized workload distribution in the STP-DAS framework.

2.4.3 Massively Parallel Performance

In this section, we examine how the STP-DAS framework functions within extensive supercomputing environments. In order to demonstrate the massively parallel performance of the STP-DAS algorithm, we performed X2C-CASCI calculations on the TIH molecule with increasing size active spaces of 35 to 40 spin orbitals using 24 correlated electrons.

Table 2.3 lists the total number of determinants in a X2C-CASCI calculation and the equivalent number of non-relativistic determinants measured by the number of FLOPs. The minimal memory storage requirement for each case is also computed should all active

Table 2.3. Number of determinants in an X2C-CASCI calculation using the STP-DAS framework of TIH with 24 correlated electrons and a varying number of active spinor orbitals.

# of Orbitals	# of X2C Determinants (NR Equivalent ^a)	Excitation List Storage ^b
35	417,225,900 (2.5×10^9)	841 GB
36	1,251,677,700 (7.5×10^9)	2,733 GB
37	3,562,467,300 (2.1×10^{10})	8,378 GB
38	9,669,554,100 (5.8×10^{10})	38 TB
39	25,140,840,660 (1.5×10^{11})	106 TB
40	62,852,101,650 (3.8×10^{11})	282 TB

^a Non-relativistic (NR) equivalent number of determinants in non-relativistic calculations when measured by the number of FLOPs.

^b Only non-zero elements are considered in the one-electron excitation list. Bit-wise compression of the determinant address is used.

# of Orbitals (Orbital Partitions)	# of Categories
35 (5,5,5,5,5,5,5)	9,142
36 (5,5,5,5,5,5,5,1)	21,259
37 (5,5,5,5,5,5,5,2)	36,526
38 (5,5,5,5,5,5,5,3)	54,853
39 (5,5,5,5,5,5,5,4)	75,846
40 (5,5,5,5,5,5,5,5)	98,813

Table 2.4. STP-DAS space partition schemes and the resulting numbers of configuration categories used in the benchmark for the massively parallel performance of the STP-DAS framework.

orbitals and electrons be included in a single active space, *e.g.*, in a conventional CASCI calculation. Table 2.3 shows that as the active space expands, the memory needed to store non-zero elements of the excitation list increases rapidly, becoming exceedingly demanding.

Table 2.4 lists the STP-DAS partition schemes employed in the following HPC benchmark study. The excitation list storage requirement is not presented because it is reduced to less than 100 KB for all test cases. Two high-performance computing (HPC) systems were used in this benchmark study. The first is a medium size HPC system, Hyak, managed by the University of Washington (UW). Each Hyak node has two Intel Xeon 6230 Gold CPUs with a single 100 GB/s network interface card. A maximum of 250 nodes on this medium sized HPC system were available to the authors. The second system is the Department of Energy’s Perlmutter high-performance super computer with up to 512 nodes (AMD EPYC 7763 (Milan), 200 GB/s NIC).

Strong scaling

We first study the strong scaling of the X2C-CASCI calculation using the STP-DAS framework on a medium size HPC system, the UW Hyak. Figure 2.5 shows the strong scaling performance of the STP-DAS σ build on varying active space sizes for the TIH test case. For the log-log presentation of runtime versus number of nodes, an ideally scaling algorithm would be represented by straight, decreasing line. As such, one can observe excellent strong scaling of the STP-DAS implementation for a variety of problem sizes.

For real world applications, an important feature that can be extracted from a strong scaling plot is a stagnation point where one can observe that speedups have ceased despite an increase in computational power. This is usually a sign that the amount of work is insufficient for the number of nodes. Focusing on the $\binom{35}{24}$ problem, one can observe this stagnation point at around 75 nodes. As one increases the problem size, this stagnation point shifts to a greater node count as more work is available to divide amongst the nodes. This can be observed for the $\binom{39}{24}$ X2C-CASCI problem, which still benefits from extra computational resources and scales up to 250 nodes.




Figure 2.5. Execution time of a single σ build using the STP-DAS framework for the TIH test case with increasing CAS space sizes with respect to the number of nodes (1 MPI per node, 20 SMP threads per MPI).

Relative speed-up

A strong scaling analysis is helpful at understand the overall behavior of an algorithm, but this does not yield an understanding of the origins of the stagnation point. For a detailed understanding of the STP-DAS framework, we turn to a detailed analysis of the speed up of a certain problem size with respect to the number of nodes. Ideally the concept of speedup would require a definition of a serial run time, however since the size of the CI calculations described in Table 2.3 would not fit on a single node, we instead present a strong scaling as a relative speed-up with respect to the performance on 16 nodes. We study this speedup in detail for two problem sizes $\binom{37}{24}$ and $\binom{38}{24}$ to extract the salient features of the STP-DAS algorithm.

Figure 2.6 (top) shows the relative speed up of an X2C-CASCI calculation using the STP-DAS framework of TIH with an active space of $(37o,24e)$. We additionally plot the relative speedup of the computation of the σ build and the MPI communication idle time as the total runtime is a combination of these two times. It is immediately apparent that for this problem size the strong scaling stagnation onset occurs around 125 nodes. Additionally, one can observe that the MPI communication idle time does not scale with the number of nodes and that the actual computation time scales extremely well with the number of nodes.




Figure 2.6. Relative speedup of a single σ build for TIH (top) $\binom{37}{24}$ and (bottom) $\binom{38}{24}$ with respect to the number of nodes (1 MPI per node, 20 SMP threads per MPI). Relative speed-up is defined with respect to the performance on the smallest node count capable of solving the problem.

It is important to recognize that a superlinear scaling of the computation, the σ build time, is not an indication of a superlinear scaling of the algorithm as this does not include the MPI communication idle time. Additionally, it is important to stress that the onset of the strong scaling stagnation is not an indication that the calculation takes the same amount of time irrespective of the number of nodes beyond that point. Beyond the stagnation point, one observes deviation from ideal scaling, but there is still a reduction in total runtime as the computation scales with the number of nodes while the MPI communication idle time does not.

By increasing the problem size, one shifts the location of the strong scaling stagnation point. Figure 2.6 (bottom) shows the relative speedup of an X2C-CASCI calculation using the STP-DAS framework of TIH with an active space of (38o, 24e). The behavior now differs from the smaller (37o, 24e) case. For the range of calculations presented, ideal scaling is observed for the full range of the number of nodes. One may often see scaling plots such as this one rather than the previous case, which does not show the stagnation point. However, it is important to note that this does not imply that the stagnation point does not exist only that it has shifted to larger node counts.

Computation time, MPI communication idle time, and load balancing

To further understand the communication versus computation time of the STP-DAS algorithm, we plot the raw execution time as a function of problem size. These calculations were run on the large Department of Energy’s Perlmutter HPC system. From Figure 2.7, it is evident that the total execution time exhibits a linear relationship with the number of determinants. This linear scaling is due to the fact that, with a fixed number of electrons, the number of non-zero elements in the one-electron excitation list increases linearly as the number of unoccupied orbitals grows. Additionally, it is observed that the MPI communication idle time grows very slowly with problem size. In the largest test case $\binom{40}{24}$, the MPI communication idle time is only 18% of the total σ build time.

In the STP-DAS algorithm the σ build computation time and the MPI communication idle time vary slightly per MPI process as each MPI process receives a different set of




Figure 2.7. Execution time of a single σ build using the STP-DAS framework for the TIH test case with increasing CAS space sizes (Perlmutter 256 Nodes, 1 MPI per node, 64 SMP threads per MPI).

categories to process. To study the load balancing performance of the STP-DAS algorithm, Figure 2.8 (top) shows the build time for each of the X2C-CASCI calculations. The solid line represents the mean of the σ build times from each process, and the shaded region represents the histogram of the σ build times from each MPI process. The small range of times over which the histogram is spread, irrespective of the calculation size, is an indicator of the strong load balancing in the current iteration of the STP-DAS algorithm. This shows that breaking the problem into smaller DAS spaces allows for a balanced distribution of categories of determinants.

Once each MPI process has completed the assigned work for the σ build, it must idle and communicate its result with the other processes. The distribution of the MPI communication idle times is represented in Figure 2.8 (bottom). The MPI communication idle time tends to be closely grouped around the average for most problem sizes. However, for the largest calculation, $\binom{40}{24}$, the distribution exhibits a minor tail. It is anticipated that the incorporation of Remote Memory Access into MPI will enhance both the efficiency of MPI communication idle time and the load balancing, especially in cases of such substantial computational magnitude.

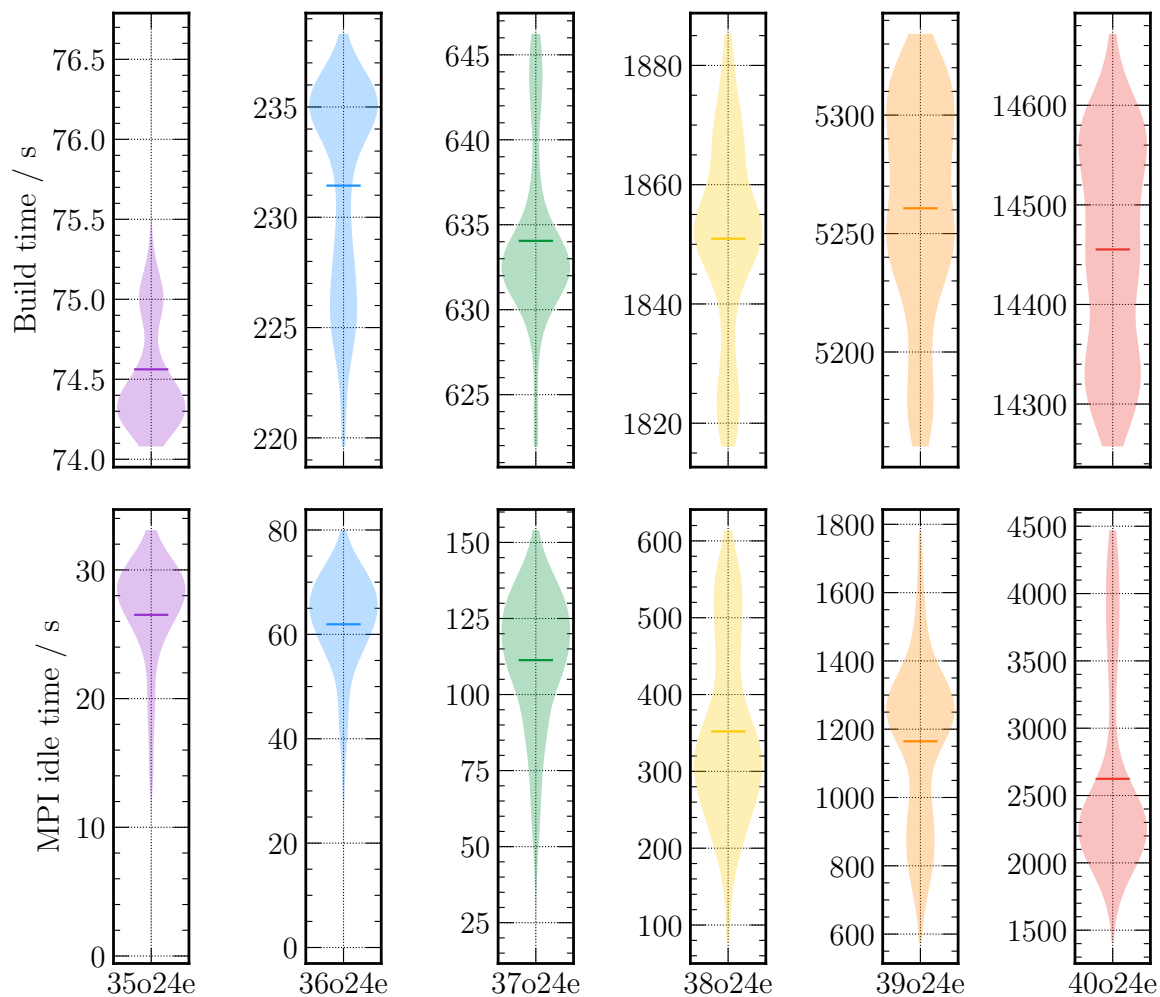


Figure 2.8. (top) σ build time (seconds) and (bottom) the distribution of the MPI communication idle time of each process during the σ build process of the T1H test case with a increasing numbers of correlated orbitals (Perlmutter 256 Nodes, 1 MPI per node, 64 SMP threads per MPI).

Effect of the active space partitioning

To test the effect of active space partitioning on the performance of the STP-DAS algorithm, we used four different DAS partitionings for a small test case $\binom{28}{18}$ (13 million X2C determinants) to run the calculations on a single node. The runtimes for one σ -build are collected in Table 2.5.

We observed a decrease in the σ -build time from one to four DAS partitionings. This reduction is due to shifting work from internal loops (over excitations within DAS spaces) to the outer loop over categories. However, when the space is partitioned into seven DASs, the runtime increases significantly due to the extra work required to locate each local address, as seen previously in Section 5.1. This observation suggests that there is an optimal condition for STP-DAS, but it is strongly dependent on the system size and the nature of the computing architecture.

Table 2.5. The storage needs for the one-electron excitation list and runtimes for one σ build in a $\binom{28}{18}$ X2C-CASCI calculation.

# of DASs (Orbital Partitions)	Excitation List Storage ^a	σ Build Time
1 (28)	18.12 GB	876 s
2 (14,14)	0.86 GB	404 s
4 (7,7,7,7)	1.98×10^{-3} GB	86.7 s
7 (4,4,4,4,4,4,4)	1.06×10^{-5} GB	365.4 s

^a Only non-zero elements are considered in the one-electron excitation list. Bit-wise compression of the determinant address is used. 2 Intel Xeon Gold 6148 processors with 20 physical cores each and 250 GB of memory.

2.5 Conclusions

In this work, we introduced a small tensor product distributed active space (STP-DAS) framework, characterized by adjustable space partitioning and many small tensor products with an efficient tensor loop used in the σ build. This framework is designed to support a variety of configuration interaction (CI) methodologies. It is also compatible with both relativistic (2-component and 4-component) and non-relativistic electronic structure methods.

The CI engine within the STP-DAS framework leverages this adjustable partitioning to significantly reduce the memory requirements for large-scale multiconfigurational calcula-

tions, offering scalability from single workstations to massively parallel computing environments.

Our benchmark tests, conducted on two different supercomputers using realistic X2C-CASCI calculations using the STP-DAS framework with determinant numbers ranging from 10^9 (billion) to 10^{12} (trillion), consistently demonstrated robust parallel scalability and excellent load balancing.

A standout feature of the STP-DAS framework is its capacity to facilitate extensive CI calculations with limited computational resources. Illustrating this, we performed a relativistic CI calculation involving 177 million X2C determinants, a task computationally equivalent to 1.1 billion non-relativistic determinants based on FLOP count, on a laptop. This capability showcases the STP-DAS framework’s potential to broaden the applicability of CI methods in computational science research.

Although determinant-based algorithms vary in their contraction order and the size of the intermediates they form, they all benefit from the STP-DAS framework through significantly reduced excitation list sizes and the ability to express contractions solely in terms of local sub-determinants. Benchmarking these different algorithms requires fine-tuning each one within the STP-DAS framework, an ongoing research endeavor that will be presented in a future publication.

The STP-DAS algorithm provides an opportunity to optimize CI-based methods for specific hardware architectures. Although the optimal STP-DAS scheme cannot be determined a priori, benchmarking various space partitioning schemes enables users to tailor the algorithm to a particular large-scale high-performance computing facility. This approach helps to leverage hardware configurations effectively and minimize the impact of communication latency on overall computational efficiency.

Chapter 3

**CATEGORICAL COMPRESSION WITHIN THE STP-DAS
FRAMEWORK****3.1 Introduction**

Full configuration interaction (FCI) offers the most complete and accurate description of a molecule’s electronic structure within a given basis set, providing the exact spectral solution to the non-relativistic electronic Schrödinger equation [134, 263, 264, 268, 284, 295]. Due to its variational nature, FCI is particularly well-suited for treating relativistic effects, such as spin–orbit and spin–spin couplings beyond perturbation theory [2, 17, 18, 120, 121, 130, 132, 146–149, 183, 238, 260, 270, 289, 300], which are fundamentally rooted in the electronic Dirac equation.

At its core, solving the FCI problem reduces to diagonalizing a large many-electron Hamiltonian matrix. This matrix is Hermitian, sparse, and typically diagonally dominant, making it well-suited to iterative diagonalization techniques that can efficiently converge on a few eigenstates without requiring full storage or construction of the entire matrix. However, as the number of determinants grows factorially with system size, reflecting the combinatorial nature of Slater determinant enumeration within the full Hilbert space, even iterative methods become intractable beyond a certain threshold.

The CI wavefunction is often expressed as a linear combination of Slater determinants, typically generated by excitations from a mean-field self-consistent field (SCF) ground-state reference. In the relativistic regime, this framework must be reformulated using complex-valued 2- or 4-spinor wavefunctions, as required by the Dirac formalism [68, 236].

Figure 3.1 illustrates the historical progression of CI implementations, highlighting major breakthroughs in the achievable scale of determinants. Prior to this work, over a span of 35 years (1990–2024), the field advanced from handling billions to trillions of determinants, driven largely by advances in computer hardware technologies. Although CI is amenable to

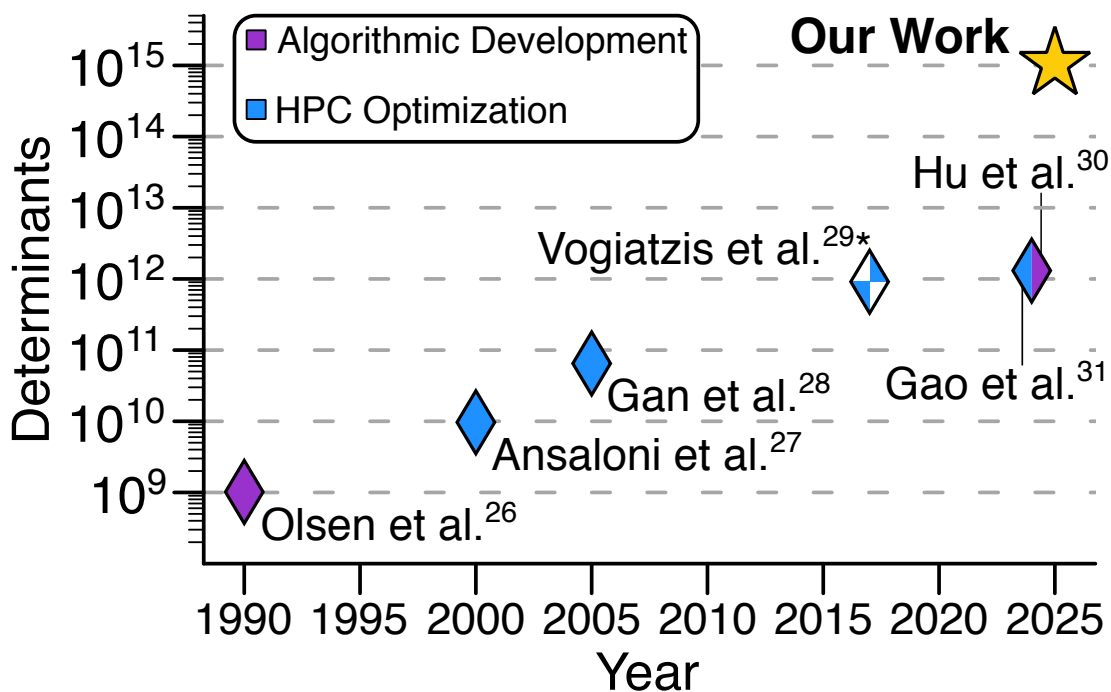


Figure 3.1. The evolution of the state-of-the-art for CI calculations over time [8, 93, 95, 122, 211, 302]. Each historical point is colored according to the nature of the key development of the respective work, which we classify as either intrinsic algorithmic developments (in purple) or optimizations on the HPC hardware of the time (in blue). * Only one CI iteration was performed.

large-scale parallel processing schemes [8, 93, 95, 122, 211, 302], the explosive growth in memory requirements has historically restricted its applicability to only the smallest chemical systems. Relativistic CI is even more limited, due to the intrinsically larger spinor configuration space associated with complex-valued 2- or 4-component wavefunctions. Simply put, enabling CI for practical quantum chemistry applications demands alternative theoretical frameworks and data representations that can circumvent the brute-force enumeration of the CI space.

Many CI-based wavefunction methods aim to approximate the FCI solution. These methods use different types of approximations, which affect the accuracy of the resulting wavefunction. The two main approaches are complete active space CI (CASCI) and selected CI (SCI). While both CASCI and SCI methods effectively truncate the Hilbert space of the system to a subspace of significant determinants, this significance is determined differently

and at different stages of the computation.

In the CASCI method [85–87, 95, 145–147, 150, 213, 243, 302], it is assumed that only a subspace of the full Hilbert space of the system contains meaningful correlation, and the FCI wavefunction is approximated as the CI wavefunction in the truncated space (the so-called active space). The truncation often leads to an underestimation of dynamic correlation. Applying more computationally demanding methods such as multiconfigurational self-consistent field (MCSCF) [17, 94, 132, 144, 174, 185, 187, 188, 238, 270, 284, 289, 301, 305, 306, 319], multireference configuration interaction (MRCI) [115, 121, 174–176, 189, 246, 284, 304], and many-body perturbation theory (MRPT2, CASPT2, NEVPT2, MC-PDFT) [5–7, 35, 172, 183, 186, 244, 270, 300] is typically required to achieve qualitative and quantitative agreement with experiment.

SCI-based methods estimate the importance of each configuration in the total wavefunction based on a predefined significance criterion, which depends on the chemical problem of interest [4, 21, 31, 80, 99, 102, 117, 118, 124, 168, 179, 252, 262, 291, 307, 320, 321]. Once the most significant determinants are identified, the Hamiltonian is constructed and diagonalized within this reduced space to approximate the FCI wavefunction. As in CASCI, SCI approaches are often combined with perturbation theory to recover contributions from neglected configurations and improve quantitative accuracy [80, 117, 124, 252, 291].

Even with advances in dimensionality reduction, CI remains fundamentally constrained by memory limitations. State-of-the-art implementations still require explicit storage of either Hamiltonian matrix elements or excitation lists to support on-the-fly matrix–vector operations, commonly referred to as the σ -build [150, 212, 271, 318]. For large CI spaces, storing the full Hamiltonian matrix and performing direct diagonalization is clearly impractical. In on-the-fly CI algorithms, the one-electron excitation list, which encodes the allowed excitations between determinants for efficient Hamiltonian construction, scales as $n_e \times (n_h + 1) \times N$, where N is the number of determinants, and n_e and n_h denote the numbers of electrons and holes (unoccupied orbitals), respectively. Since N increases factorially with system size, the associated memory requirements grow rapidly, making conventional CI calculations infeasible for anything beyond the smallest systems.

For a CI problem involving N determinants, the size of each CI expansion vector scales

linearly with N . For instance, a relativistic CI problem with one quadrillion (10^{15} , 100 orbitals, 88 electrons) determinants would require approximately 16 petabytes (PB) of memory just to store a single CI vector composed of complex-valued double-precision coefficients. While this memory footprint alone makes such problems challenging to tackle, the memory required to store the excitation list can easily scale to the exabyte (EB) regime. This poses a fundamental barrier to scalability, even before considering computational cost.

A recently introduced CI matrix–vector product algorithm [122] leverages the exact factorization of the active space into small tensor products of distributed active spaces—an approach known as the small-tensor product distributed active space (STP-DAS) framework, illustrated in Figure 3.2(A). The STP-DAS algorithm reformulates the large CI matrix–vector product into a sequence of small tensor products, each embedded within a distributed active space, computed on-the-fly using string-based methods. This advance exploits the mathematical condition governing the phase relationship between the global address and the local DAS address of any CI matrix element, enabling the use of only small local determinant address strings in the CI matrix–vector build and overcoming the memory bottleneck associated with storing the full excitation list. This formulation enables extensive reuse of Hamiltonian excitation lists, leading to a dramatic reduction in memory demands. For a CI problem involving one quadrillion (10^{15}) determinants, the STP-DAS framework reduces the excitation list memory requirement from 12 exabytes (EB) to 25 gigabytes (GB), an 8-orders-of-magnitude reduction! In addition, by evenly distributing the computation of small tensor products, the STP-DAS algorithm achieves excellent load balance with minimal node-to-node communication overhead, ensuring strong scalability across both single-node and large-scale parallel architectures.

Since the excitation list typically dominates the storage requirements in CI calculations, the STP-DAS framework overcomes a longstanding memory bottleneck and enables CI computations that were previously deemed intractable. By effectively eliminating the memory footprint of the excitation lists, the storage of CI vector coefficients now emerges as the bottleneck in many large-scale CI problems.

Revisiting the earlier CI example of 10^{15} determinants: even after eliminating the memory bottleneck associated with storing excitation lists, the same calculation still demands

Table 3.1. Details of the convergence of the X2C-CASCI calculation (100 2-spinor orbitals, 88 electrons, 10^{15} 2-spinor determinants) of the ground state of HBrTe (x2c-TZVPall [227]). Each row corresponds to a Davidson iteration. The fourth and fifth columns track the changes in the computed Rayleigh–Ritz eigenpair across iterations, and the sixth column tracks the residual norm, $\|\mathbf{r}\| = \|\mathbf{H}\mathbf{C} - E\mathbf{C}\|$, where \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\varepsilon = 1.0 \times 10^{-5}$.

Iteration	Duration (s)	Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ \mathbf{r}\ $ (E_h)
1	122	-9812.080102321897	—	—	2.62×10^{-1}
2	111	-9812.129587025540	4.95×10^{-2}	5.45×10^{-2}	8.73×10^{-2}
3	262	-9812.133341658482	3.75×10^{-3}	1.36×10^{-2}	3.23×10^{-2}
4	710	-9812.133735698319	3.94×10^{-4}	2.26×10^{-3}	1.40×10^{-2}
5	1536	-9812.133784486688	4.88×10^{-5}	1.08×10^{-3}	7.14×10^{-3}
6	3293	-9812.133790409140	5.92×10^{-6}	4.97×10^{-4}	4.09×10^{-3}
7	6432	-9812.133791245860	8.37×10^{-7}	1.31×10^{-4}	2.54×10^{-3}
8	11565	-9812.133791408300	1.62×10^{-7}	9.17×10^{-5}	1.61×10^{-3}
9	20384	-9812.133791451150	4.28×10^{-8}	4.25×10^{-5}	9.71×10^{-4}

16 PB of memory to store the numerically exact coefficients of a single CI expansion vector in an iterative solver. Given that practical CI calculations typically require multiple subspace vectors for convergence, it becomes clear that storing these expansion vectors now represents the dominant memory bottleneck in large CI calculations. To address this challenge, we leverage the locally compressed nature of the STP-DAS framework to efficiently compute ultra-large-scale CI problems involving up to a quadrillion (10^{15}) determinants. This approach yields deterministic, numerically exact solutions and effectively shifts CI calculations from being memory-bound to compute-bound.

3.2 Results

Before presenting methodological details and performance benchmarks, we highlight the largest CASCI calculation to date for the ground state of HBrTe, enabled by the compression-compatible STP-DAS framework to be introduced herein. HBrTe is a substituted form of a hydrogen chalcogenide where one of the hydrogens was substituted with a bromine atom to decrease the symmetry to the molecule. We performed a relativistic exact two-component [64, 65, 70, 119, 126, 155, 163, 177, 180, 181, 183, 221, 222, 260] CASCI (X2C-CASCI) calculation (100 2-spinor orbitals, 88 electrons, complex-valued 1.05×10^{15} 2-spinor deter-

minants) of the ground state of HBrTe using the compression-compatible STP-DAS framework. We also performed a calculation with the same number of determinants for the ground state of a magnesium atom (see Section S3 of the Supplementary Information). The calculation ran on the National Energy Research Scientific Computing Center’s Perlmutter high-performance supercomputer with a total of 1000 nodes (AMD EPYC 7763 Milan, 128,000 compute cores, 512 GB of RAM per node, 200 GB/s NIC, 2 message passing interface (MPI) processes per node, and 64 symmetric multiprocessing (SMP) threads per MPI process).

The excitation list is generated without assuming any symmetry of the target state. Consequently, the calculation is formally performed in a quadrillion-determinant space, with all determinants explicitly included. Because the CI coefficients are complex, the memory footprint is twice that of an analogous nonrelativistic calculation. Moreover, the complex arithmetic makes the computational cost (in FLOP count) equivalent to that of a nonrelativistic calculation with more than twice as many determinants.

Table 3.1 summarizes the results of the HBrTe calculation. The ground-state energy converged in 9 iterations to microhartree precision ($< 10^{-6}$ a.u.) with a total runtime of 34.5 hours. In each iteration, an additional CI expansion vector was introduced to accelerate convergence, while the compression algorithm dynamically adapted to the expanding vector space, leading to a gradual increase in computational cost. A total of 9 expansion vectors were involved in the σ -build. Despite the enormous configuration space of over one quadrillion (1.05×10^{15}) complex-valued 2-spinor determinants, the average σ -build time per vector remained just 3.8 hours.

The ground-state energy of the HBrTe molecule from this CI calculation is -9395.0280045 a.u. Leveraging the gap theorem [50, 217], we determine that our $\binom{100}{88}$ X2C-CASCI result lies within 10×10^{-6} a.u. of the true x2c-TZVPall [227] ground state energy within that active space, well below any chemically meaningful threshold. A detailed analysis is provided in Methods.

This work represents a 3-orders-of-magnitude increase in CI space and a 6-orders-of-magnitude increase in FLOP count, which is estimated using $\mathcal{O}(N^2)$, compared to the previous state-of-the-art in CI calculations [95, 122]. Compared to previous state-of-the-

art CI calculations, this work also achieves a 6-orders-of-magnitude speedup in time-to-completion, as measured in core seconds per exaFLOP ($\frac{\text{core-second}}{\text{exaFLOP}}$, see Section S2 in the Supplementary Information for analysis). This ultra-large-scale CI calculation is enabled by the STP-DAS-based numerically exact categorical compression scheme, which reduces the memory required to store the 9 CI expansion vectors from 134 PB to less than 500 TB while maintaining a good load balance [122], making the computation feasible on most existing supercomputing infrastructures.

We now describe the algorithmic developments that enable such CI calculations to be performed on existing supercomputing resources. Detailed algorithms, parallel implementation strategies, and error bound analyses are provided in the Supplementary Information. The central concept of the STP-DAS framework is the systematic partitioning of the full CI orbital space into a collection of distributed active spaces. Within each active space, configurations are further classified into categorical subspaces, rigorously defined by distinct electron occupation patterns, as illustrated in Figure 3.2(A) [122]. The STP-DAS framework reformulates the CI σ -build as a sum of small-tensor products, each uniquely addressed via a global tensor looping structure. The major memory bottleneck associated with storing the excitation list is eliminated by allowing categorical subspaces to share compact, local excitation lists.

In the largest CASCI calculation (10^{15} determinants) presented here, employing 13 distributed active spaces, the STP-DAS approach reduces the excitation list memory requirement from 12×10^9 GB to just 25 GB. However, storing all nine CI expansion vectors for a system with 10^{15} determinants would require approximately 134 PB of memory. On a high-performance computing system such as Perlmutter, this translates to more than 275,000 nodes, each equipped with 512 GB of memory. A straightforward element-wise sparsity treatment, however, does not meet the STP-DAS condition. To overcome this limitation, the following section introduces a categorical compression scheme that achieves the necessary memory reduction, making it possible to carry out large relativistic CI calculations on a medium-sized computing cluster.

With the STP-DAS framework, the memory bottleneck associated with storing CI expansion vectors is effectively eliminated by applying numerically exact categorical compression.

The STP-DAS CI expansion vectors take the form

$$\mathbf{C} = \bigoplus_{\mathcal{B}} \mathbf{C}^{\mathcal{B}}, \quad (3.1)$$

where \mathcal{B} is a category, defined by a unique electron occupation pattern within the distributed active spaces [122].

The compression scheme, categorical compression (see Figure 3.2(D)), stores the $\mathbf{C}^{\mathcal{B}}$ vectors as compressed sparse column (CSC) vectors. In this format, each categorical expansion vector $\mathbf{C}^{\mathcal{B}}$ is represented by two equally sized arrays: one, $V_{\text{CSC}}^{\mathbf{C}^{\mathcal{B}}} \equiv \{C_{K^{\mathcal{B}}} : C_{K^{\mathcal{B}}} \neq 0\}$, stores the values of the nonzero coefficients, while the other, $A_{\text{CSC}}^{\mathbf{C}^{\mathcal{B}}} \equiv \{K^{\mathcal{B}} : C_{K^{\mathcal{B}}} \neq 0\}$, stores their corresponding local addresses. The tensor-loop structure in the STP-DAS σ -build algorithm is reformulated in terms of categorically compressed local addresses together with their corresponding global phase factors (see Figure 3.2(E)).

In contrast to element-wise compression, categorical compression can eliminate all configurations within a category, i.e., skip an entire category at once. This approach is better able to preserve the vectorized structure compared to element-wise compression. The major difficulty of any sparse matrix-vector product algorithm is the lack of a priori knowledge of the location of the nonzero elements. This difficulty leads to a major bottleneck rooted in the nonuniform accesses to memory. The categorical compression localizes nonuniform memory accesses within a category, which is always orders-of-magnitudes smaller than the size of the full CI space. This localized memory access pattern is the central strength of the categorical compression scheme.

Within each category, element-wise compression is still applicable to maximize sparsity. Most importantly, a category-based compression scheme naturally supports distributed small tensor products, taking advantage of the reduced memory footprint and improved parallel load balance of the STP-DAS framework.

In summary, the numerically exact categorical compression introduced here allows the STP-DAS σ -build algorithm to bypass entire categories of determinants while preserving both the vectorized structure and the local addressing scheme of STP-DAS. Because the compression is fully lossless, the omitted determinants have no impact on the resulting Ritz

eigenvalue–eigenvector pair.

The final hurdle in reducing CI memory demands lies in the iterative solver. In CI, the Hamiltonian operator is diagonalized iteratively within the full Hilbert space of the system. As a result, the error in the computed Ritz value directly reflects the missing correlation energy in the associated approximate wavefunction. This enables the Ritz residual, defined as $\mathbf{r} \equiv \|\mathbf{H}\mathbf{C} - E\mathbf{C}\|$, to be computed and appended as an additional CI expansion vector. The norm of the residual provides rigorous bounds on the missing correlation energy relative to the true eigenpair (eigenvector and eigenvalue) of the Hamiltonian in the chosen basis [46, 50, 151, 152, 197, 217, 239, 242, 325]. A widely used approach that leverages this principle is the Davidson iterative solver [49]. Other related methods that use residual norm as the convergence criterion include the locally optimal block preconditioned conjugate gradient (LOBPCG) method [153], the Jacobi-Davidson [91] method and generalized preconditioned locally harmonic residual (GPLHR) method [135, 297].

By applying numerical or convergence thresholds at various stages of the Davidson method [49], one can exploit the sparsity of newly generated CI expansion vectors. With sufficiently tight thresholds, these vectors can span the part of the Hilbert space required to accurately represent the desired wavefunction and drive the iterative diagonalization to any desired level of precision. The Davidson method [49] utilized the Davidson preconditioner, which generates the i th component of the next trial expansion vector according to

$$t_i \leftarrow \begin{cases} 0, & \text{if } |\lambda - H_{ii}| \text{ is small} \\ \frac{r_i}{\lambda - H_{ii}}, & \text{else} \end{cases} \quad (3.2)$$

Here, λ is the Ritz value of the current iteration, H_{ii} is the i th element of the diagonal of the Hamiltonian \mathbf{H} , and r_i is the i th component of the current residual. Among the various preconditioners employed in the Davidson method [46, 151, 152, 173, 218, 326], compression-compatible preconditioners, which discard terms t_i below a numerical threshold ε , have been shown to achieve convergence to the exact same results as the traditional Davidson preconditioner [111, 151, 152, 197, 242].

In this work, we apply the compression-compatible categorical preconditioner

$$\begin{aligned} s_i &\leftarrow \begin{cases} \frac{r_i}{\lambda - H_{ii}}, & \text{if } |\lambda - H_{ii}| \geq 10^{-12} \\ 0, & \text{else} \end{cases} \\ t_i &\leftarrow \begin{cases} s_i, & \text{if } |s_i| \geq \varepsilon \|\mathbf{s}\| \\ 0, & \text{else} \end{cases} \end{aligned} \quad (3.3)$$

using the nonzero residual elements r_i . We do this on-the-fly to avoid explicitly storing the prohibitively large diagonal of \mathbf{H} (a dense vector of size N_{dets}). Figure 3.2(F) illustrates the expansion of the CI vector space enabled by the compression-compatible categorical preconditioner used in the Davidson method implemented here. Eq. (3.3) closely resembles the preconditioner proposed in Ref. 111, with the key distinction being the inclusion of the Davidson-preconditioned residual norm $\|\mathbf{s}\|$ in the dropping criterion. This facilitates dynamic threshold adjustment: as the iterations progress and $\|\mathbf{s}\|$ decreases, the criterion becomes more stringent. More importantly, the factor of $\|\mathbf{s}\|$ ensures that numerical thresholding is applied to the generated expansion vectors relative to the total norm of their exact (traditional Davidson) counterparts, rather than some fixed cutoff on the absolute values of their entries. This results in a very accurate CI space expansion scheme at the cost of computing and contracting a significant number of Hamiltonian matrix elements (to evaluate s exactly), which would be intractable without the STP-DAS framework.

Algorithms and pseudocodes of the compression-compatible STP-DAS method are presented in Methods, along with discussions on load balancing and parallel implementation. The convergence behavior of the STP-DAS framework equipped with the compression-compatible preconditioner defined in Eq. (3.3) was evaluated across three systems with varying degrees of electron correlation: the magnesium atom, diatomic nitrogen, and a model carbon nanotube. The results are provided in Section S1 of the Supplementary Information.

The analysis reveals that overly aggressive thresholding can cause the Davidson procedure to stagnate, thereby preventing convergence to the correct electronic wavefunction.

When thresholds are too loose, newly generated trial vectors quickly become linearly dependent on the existing subspace vectors, signaling that the span of the modified subspace has saturated before achieving convergence. However, when the preconditioning threshold satisfies

$$\varepsilon \lesssim \frac{10}{\sqrt{N_{\text{dets}}}}, \tag{3.4}$$

the resulting energies agree with their exact values to better than $10^{-7} E_h$, and the residual norms become correspondingly small, demonstrating successful and reliable convergence.

The CI wavefunction of highly correlated systems is comprised of a large number of determinants with small CI coefficients. Because the Hilbert space of the problem is never truncated, and no determinants are discarded from the wavefunction itself. As a result, the compression-compatible preconditioner easily facilitates convergence to the exact wavefunction, provided that the preconditioning threshold ε is sufficiently small for the true eigenvector to be accurately represented in the subspace spanned by the modified expansion vectors.

With the capability to perform large CI calculations, energetic extrapolation to the correlation limit becomes feasible for many-electron systems. Figure 3.2(B) illustrates the correlation-consistent extrapolation for the Mg^{2+} ion using double-zeta (DZ, 36 orbitals), triple-zeta (TZ, 68 orbitals), and quadruple-zeta (QZ, 118 orbitals) basis sets, involving 2.54×10^8 , 2.91×10^{11} , and 9.75×10^{13} 2-spinor determinants, respectively. The complete basis set limit of -199.16704295 a.u. was obtained using a mixed Gaussian extrapolation scheme tailored for correlation-consistent basis sets [224].

To demonstrate the strong scaling behavior of the compression-compatible STP-DAS σ -build algorithm, we performed relativistic X2C-CASCI calculations on the thallium hydride (TIH) molecule using active spaces of 40 and 41 2-spinor orbitals and 24 electrons, $\binom{40}{24}$ and $\binom{41}{24}$, corresponding to 63 and 152 billion 2-spinor determinants, respectively. Five distributed active spaces (DASs) were employed. These calculations were executed on the University of Washington’s Hyak HPC system, a small-sized cluster where each node is equipped with two Intel Xeon 6230 Gold CPUs and a single 100 GB/s network interface card. As shown in Figure 3.2(C), even with just 5 compute nodes, relativistic X2C-CASCI

calculations involving tens to hundreds of billions of determinants require only 4-6 minutes per σ -build on average. Increasing to 30 nodes further reduces the cost to just over one minute. Past 30 nodes, the communication time dominates the runtime and the calculations no longer scale. This benchmark demonstrates that billion- and even trillion-determinant CI calculations are now feasible on a small-scale computing cluster.

Additionally, we performed X2C-CASCI calculations for the ground states of two highly correlated systems, showcasing the applicability of the compression-compatible STP-DAS framework to highly correlated systems. These systems were chosen from opposite ends of the correlation spectrum from strongly statically correlated to strongly dynamically correlated. Square Rb_4 , a relativistic analogue of H_4 [9, 43, 97, 110, 137, 201, 202, 249, 251, 255, 258, 272, 293], displays strong static correlation and Xe_2 is a dynamically correlated noble gas dimer [30, 101, 133, 160, 165, 275, 287, 323, 328].

Tables 3.2 and 3.3 summarize the results of the Rb_4 and Xe_2 calculations. The Rb_4 calculation (50 2-spinor orbitals, 28 electrons, 8.9×10^{13} 2-spinor determinants) ran on the National Energy Research Scientific Computing Center’s Perlmutter high-performance supercomputer with a total of 100 nodes (AMD EPYC 7763 Milan, 12,800 compute cores, 512 GB of RAM per node, 200 GB/s NIC, 2 MPI processes per node, and 64 SMP threads per MPI process) and took 6 iterations and 11.8 hours to converge. The Xe_2 calculation (60 2-spinor orbitals, 12 electrons, 1.4×10^{12} 2-spinor determinants) ran on the same platform with 256 nodes and took 7 iterations and 36.1 hours to converge.

The ground-state energies of the Rb_4 and Xe_2 molecules from these calculations are -11916.1527249 a.u. and -14889.6466956 a.u., accordingly. The gap theorem [50, 217] guarantees that these X2C-CASCI results lie within 0.53×10^{-6} a.u. and 17.56×10^{-6} a.u. of the true ground state energies within the corresponding active spaces and basis sets (see Methods).

In summary, by combining compression-compatible preconditioners with compression-compatible categorical CI vectors, the STP-DAS framework drastically reduces the memory footprint of both the excitation lists and the CI expansion vectors. In the largest relativistic CASCI calculation presented here, spanning 10^{15} determinants across 13 distributed active spaces, the STP-DAS approach reduces the memory required for the excitation list from

Table 3.2. Details of the convergence of the X2C-CASCI calculation (50 2-spinor orbitals, 28 electrons, 8.9×10^{13} 2-spinor determinants) of the ground state of Rb_4 (cc-pvtz-x2c [114]). Each row corresponds to a Davidson iteration. The third and fourth columns track the changes in the computed Rayleigh–Ritz eigenpair across iterations, and the fifth column tracks the residual norm, $\|\mathbf{r}\| = \|\mathbf{HC} - E\mathbf{C}\|$, where \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\varepsilon = 1.0 \times 10^{-6}$.

Iteration	Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ \mathbf{r}\ $ (E_h)
1	-90.038089795466	–	–	4.48×10^{-2}
2	-90.048919589524	1.08×10^{-2}	1.21×10^{-1}	1.74×10^{-2}
3	-90.049940164869	1.02×10^{-3}	1.64×10^{-2}	4.95×10^{-3}
4	-90.050002638482	6.25×10^{-5}	2.53×10^{-3}	1.20×10^{-3}
5	-90.050006376495	3.74×10^{-6}	1.37×10^{-3}	2.96×10^{-4}
6	-90.050006593144	2.17×10^{-7}	1.55×10^{-4}	7.36×10^{-5}

Table 3.3. Details of the convergence of the X2C-CASCI calculation (60 2-spinor orbitals, 12 electrons, 1.4×10^{12} 2-spinor determinants) of the ground state of Xe_2 (x2c-TZVPall-2c [227]). Each row corresponds to a Davidson iteration. The third and fourth columns track the changes in the computed Rayleigh–Ritz eigenpair across iterations, and the fifth column tracks the residual norm, $\|\mathbf{r}\| = \|\mathbf{HC} - E\mathbf{C}\|$, where \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\varepsilon = 1.0 \times 10^{-5}$.

Iteration	Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ \mathbf{r}\ $ (E_h)
1	-21.059579374232	–	–	4.76×10^{-1}
2	-21.194173478231	1.35×10^{-1}	3.90×10^{-2}	1.86×10^{-1}
3	-21.206728805229	1.26×10^{-2}	8.89×10^{-3}	5.62×10^{-2}
4	-21.207569776081	8.41×10^{-4}	8.74×10^{-4}	1.75×10^{-2}
5	-21.207634517148	6.47×10^{-5}	4.51×10^{-4}	7.88×10^{-3}
6	-21.207644283591	9.77×10^{-6}	1.19×10^{-4}	4.48×10^{-3}
7	-21.207646647756	2.36×10^{-6}	3.90×10^{-5}	2.93×10^{-3}

12×10^9 GB to just 25 GB, and for 9 CI expansion vectors from 134 PB to less than 500 TB. These reductions make quadrillion-determinant calculations tractable on current supercomputing architectures. While most of the community may not have access to the hundreds of compute nodes required for such runs, this work also demonstrates the practical feasibility of trillion-determinant calculations on just a few nodes and even on a laptop.

3.3 Methods

3.3.1 Lossless σ -Build using the Categorical Compression of Small Tensor Products

The categorical σ -build algorithm within STP-DAS [122] can be reformulated to exploit the categorical compression of the expansion vectors. The compact nature of the categorical representation enables fast and memory-efficient computation of σ -vectors. The categorical σ -build algorithm implements the evaluation of

$$\sigma_{L^A} = {}^{1e}\sigma_{L^A} + {}^{2e}\sigma_{L^A}, \quad (3.5)$$

$${}^{1e}\sigma_{L^A} = \sum_{\mathcal{B}} \sum_{\mathbb{K}_\mu^{\mathcal{B}} \oplus \mathbb{K}_\nu^{\mathcal{B}}} \sum_{pq} P_{\mu\nu} \delta_{\bar{\mathbb{X}}_{\mu\nu}^A \bar{\mathbb{X}}_{\mu\nu}^{\mathcal{B}}}$$

$$h'_{pq} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{K}_\mu^{\mathcal{B}} \oplus \mathbb{K}_\nu^{\mathcal{B}} \rangle C_{K^{\mathcal{B}}}, \quad (3.6)$$

$${}^{2e}\sigma_{L^A} = \frac{1}{2} \sum_{\mathcal{CB}} \sum_{\mathbb{J}_\mu^{\mathcal{C}} \oplus \mathbb{J}_\nu^{\mathcal{C}}} \sum_{\mathbb{J}_\kappa^{\mathcal{C}} \oplus \mathbb{J}_\lambda^{\mathcal{C}}} \sum_{\mathbb{K}_\kappa^{\mathcal{B}} \oplus \mathbb{K}_\lambda^{\mathcal{B}}} \sum_{pqrs} P_{\mu\nu} P_{\kappa\lambda}$$

$$\delta_{\bar{\mathbb{X}}_{\mu\nu}^A \bar{\mathbb{X}}_{\mu\nu}^{\mathcal{C}}} \delta_{\bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{C}} \bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{B}}} g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^{\mathcal{C}} \oplus \mathbb{J}_\nu^{\mathcal{C}} \rangle$$

$$\langle \mathbb{J}_\kappa^{\mathcal{C}} \oplus \mathbb{J}_\lambda^{\mathcal{C}} | \hat{E}_{rs} | \mathbb{K}_\kappa^{\mathcal{B}} \oplus \mathbb{K}_\lambda^{\mathcal{B}} \rangle C_{K^{\mathcal{B}}}, \quad (3.7)$$

where $p \in \mathbb{X}_\mu^A$, $q \in \mathbb{X}_\nu^{\mathcal{C}}$, $r \in \mathbb{X}_\kappa^{\mathcal{C}}$, $s \in \mathbb{X}_\lambda^{\mathcal{B}}$, $\langle \mathbb{X}_\mu^A \oplus \mathbb{X}_\nu^A | \hat{E}_{pq} | \mathbb{X}_\mu^{\mathcal{B}} \oplus \mathbb{X}_\nu^{\mathcal{B}} \rangle$ are categorical one electron excitation lists, $P_{\mu\nu}$ are global phase factors, and h'_{pq} (g_{pqrs}) are one (two) body Hamiltonian elements. See Ref. 122 for algorithmic details.

Eqs. (3.5) to (3.7) define the categorical σ -vector in terms of local STP-DAS one-electron excitation lists and the categorical CI expansion vector. Notably, when the expansion coefficients $C_{K^{\mathcal{B}}}$ are categorically compressed, the resulting σ -vector coefficients σ_{L^A} are also categorically compressed. In such cases, the categorical σ -build reduces to a contraction between a categorically compressed expansion vector and categorically compressed STP-DAS one-electron excitation lists. Thus, Eqs. (3.5) to (3.7) yield a categorically compressed σ -vector, in a manner directly analogous to the compression-preserving behavior of sparse matrix–sparse vector products (SpMSPV). Importantly, this compression preservation is general and independent of the specific storage format used for the categorically compressed representations.

The categorically compressed representation can be implemented in various forms, with the choice of storage format guided primarily by computational efficiency. Since the categorical σ -build algorithm often involves reading numerous expansion coefficients with increasing local addresses during contraction, it is natural to adopt a compressed sparse column (CSC) format for storing the categorical expansion coefficients.

3.3.2 Eigenvalue Bound Analysis

We wish to apply the gap theorem [50, 217, 327] to bound the error in the computed X2C-CASCI ground state energy:

$$|\delta E| \leq \frac{\|\mathbf{r}\|^2}{\gamma_0}, \quad (3.8)$$

where \mathbf{r} is the Ritz residual of the computed ground state and the gap $\gamma_0 \equiv E_1 - \tilde{E}_0$ is the difference between E_1 , the (unknown) exact energy of the first excited state, and \tilde{E}_0 , the computed Ritz value of the ground state. Because γ_0 is unknown, one can estimate its order-of-magnitude using approximate methods or use experimental values to compute a surrogate for the true gap. One can also obtain an exact lower bound on the gap by including the posterior error bound of the first excited state [247, 313, 327] in the Davidson calculation [325]:

$$\gamma_0 = E_1 - \tilde{E}_0 \geq \left(\tilde{E}_1 - \|\mathbf{r}_1\| \right) - \tilde{E}_0 \equiv \gamma_0^-, \quad (3.9)$$

where \mathbf{r}_1 is the residual associated with the Ritz value \tilde{E}_1 .

Using the gap theorem [50, 217, 327], we can place an exact bound on the error in the computed X2C-CASCI ground-state energy. This requires an estimate of the energy gap between the ground and first excited states of the X2C-CASCI Hamiltonian. To obtain this, we performed an X2C-CISD calculation for the two lowest-lying states and determined a gap of approximately 0.095 a.u., in good agreement with experimental values. Based on this estimate, the gap theorem bounds the error in our X2C-CASCI ground-state energy to within 10 microhartree, which is well below any chemically meaningful threshold.

3.3.3 Compression-Compatible STP-DAS Algorithm

Algorithm 1 shows the categorically compressed STP-DAS algorithm, in which only nonzero elements contribute to the categorical σ -vector. Its advantage over the traditional STP σ -build algorithm [122] is twofold: in the outermost loop, where we skip entire categories whose expansion vector vanishes (see line 2-3), and in the inner loop of line 7, where we only process excitations $\langle \mathbb{J}_\kappa^{\mathcal{C}} \oplus \mathbb{J}_\lambda^{\mathcal{C}} | \hat{E}_{rs} | \mathbb{K}_\kappa^{\mathcal{B}} \oplus \mathbb{K}_\lambda^{\mathcal{B}} \rangle$ for which $C_{K^{\mathcal{B}}} \neq 0$.

There is a potential workload imbalance associated with Algorithm 1: because the collection of categorical expansion vectors is distributed among computing nodes, the contraction workload of a given node is proportional to the number of nonzero categorical expansion coefficients it has. We alleviated some of the resulting computational delay by implementing passive one-sided MPI communication of categorical expansion vectors using remote memory access (RMA). This allows idle nodes to contract more categorical expansion vectors with their excitation lists without waiting for the corresponding busy nodes to broadcast them. Ultimately, overcoming this load-balancing issue requires dynamically redistributing categorical expansion vectors according to their sparsity, which changes during the iterative diagonalization.

As illustrated in Algorithm 1, the computational cost, both in memory and runtime, of the compression-compatible STP-DAS σ -build procedure increases with the density of the expansion vectors. Therefore, it is essential to maintain maximal compression in these vectors. To achieve this, we replace the traditional Davidson preconditioner with a compression-compatible alternative for generating new trial expansion vectors. This modification alters only the subspace expansion strategy in the Davidson algorithm, while preserving exact treatment of the full determinantal space. As a result, the computed matrix-vector product \mathbf{HC} and the corresponding residual norm $\|\mathbf{r}\| = \|\mathbf{HC} - \lambda\mathbf{C}\|$ remain exact, unlike in selected CI and other truncated approaches, where both the Hamiltonian and CI vectors are explicitly approximated.

The compression-compatible STP-DAS σ -build algorithm significantly reduces the overall workload associated with the σ build. The reduction in workload can be nonuniform: the contraction workload associated with a determinant $\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{C}}$ is proportional to the num-

Algorithm 1: Two-electron σ -build using categorical compression. Bold text represents algorithmic logic and typewritten text represents comments.

Data: the categorically compressed expansion vector $\mathbf{C} = \bigoplus_{\mathcal{B}} \left(V_{\text{CSC}}^{\mathcal{C}\mathcal{B}}, A_{\text{CSC}}^{\mathcal{C}\mathcal{B}} \right)$ and integrals g_{pqrs}

Result: the categorically compressed contraction ${}^{2e}\sigma$

1 **for** categories $\mathcal{A}, \mathcal{B}, \mathcal{C}$ and DAS indices $\mu, \nu, \kappa, \lambda$ for which $\delta_{\bar{\mathbb{X}}_{\mu\nu}^{\mathcal{A}} \bar{\mathbb{X}}_{\mu\nu}^{\mathcal{C}}} \delta_{\bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{C}} \bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{B}}} = 1$

parallel do

2 **if** $|A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}| = 0$ **then**

3 continue to next iteration

4 **for** $\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{C}}$ **do**

 // Obtain the sorted DAS one-electron excitation list of $\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{C}}$

5 $A \leftarrow \{ \mathbb{K}_{\kappa}^{\mathcal{B}} \oplus \mathbb{K}_{\lambda}^{\mathcal{B}} : \exists r \in \mathbb{X}_{\kappa}^{\mathcal{C}}, \exists s \in \mathbb{X}_{\lambda}^{\mathcal{B}} \text{ S.T. } \langle \mathbb{J}_{\kappa}^{\mathcal{C}} \oplus \mathbb{J}_{\lambda}^{\mathcal{C}} | \hat{E}_{rs} | \mathbb{K}_{\kappa}^{\mathcal{B}} \oplus \mathbb{K}_{\lambda}^{\mathcal{B}} \rangle \neq 0 \}$

 // Intersect $A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}$ with A to contract nonzero expansion coefficients

6 $i, j \leftarrow 0$

7 **while** $i < |A|$ and $j < |A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}|$ **do**

8 **if** $A[i] < A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j]$ **then**

 // Advance i to the lowest i' for which $A_{i'} \geq A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j]$

9 exponentialSearch($A, i, A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j]$)

10 **else if** $A[i] > A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j]$ **then**

 // Advance j to the lowest j' for which $A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j'] \geq A[i]$

11 exponentialSearch($A_{\text{CSC}}^{\mathcal{C}\mathcal{B}}, j, A[i]$)

12 **else**

 // Process the σ contribution at the address $L^{\mathcal{A}}$ into a thread-safe hash table belonging to category \mathcal{A}

13 **for** $p \in \mathbb{X}_{\mu}^{\mathcal{A}}, q \in \mathbb{X}_{\nu}^{\mathcal{C}}$ **do**

14 ${}^{2e}\sigma_{L^{\mathcal{A}}} \leftarrow {}^{2e}\sigma_{L^{\mathcal{A}}} + P_{\mu\nu} P_{\kappa\lambda} g_{pqrs} \langle \mathbb{L}_{\mu}^{\mathcal{A}} \oplus \mathbb{L}_{\nu}^{\mathcal{A}} | \hat{E}_{pq} | \mathbb{J}_{\mu}^{\mathcal{C}} \oplus \mathbb{J}_{\nu}^{\mathcal{C}} \rangle$

15 $\times \langle \mathbb{J}_{\kappa}^{\mathcal{C}} \oplus \mathbb{J}_{\lambda}^{\mathcal{C}} | \hat{E}_{rs} | \mathbb{K}_{\kappa}^{\mathcal{B}} \oplus \mathbb{K}_{\lambda}^{\mathcal{B}} \rangle V_{\text{CSC}}^{\mathcal{C}\mathcal{B}}[j]$

16 $i \leftarrow i + 1$

17 $j \leftarrow j + 1$

18 **Subroutine:** exponentialSearch(A, i, x)

Result: An index $i < j < |A|$ such that $A[j] \geq x$ or $j = |A|$ if nonexistent

19 $\delta \leftarrow 1$

20 **while** $j < |A|$ and $A[j] < x$ **do**

21 $j \leftarrow j + \delta$

22 $\delta \leftarrow 2\delta$

23 $j_{\min} \leftarrow j - \delta/2$

24 $j_{\max} \leftarrow \min\{j, |A| - 1\}$

 // Return the lower bound of x in the range $A[j_{\min} \dots j_{\max}]$

25 **return** binarySearch($A[j_{\min} \dots j_{\max}], x$)

ber of local addresses containing both nonzero categorical expansion vectors elements and Hamiltonian matrix elements. To improve the load-balance, we implemented dynamic SMP thread-level parallelism in the outermost loop (line 1 in Algorithm 1) instead of in the loop over determinants $\{\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{C}}\}$ (line 4 in Algorithm 1). Under dynamic parallelism, some threads execute many light contractions, while others execute fewer heavy contractions, resulting in a more uniform distribution of contraction workload. Such dynamic parallelism is ineffective in the loop over determinants $\{\mathbb{J}_{\mu\nu\kappa\lambda}^{\mathcal{C}}\}$ due to the small tensor product nature of the STP-DAS framework.

3.4 Conclusions

In this work, we conducted a relativistic configuration interaction (CI) calculation for the ground state of HBrTe in a quadrillion-determinantal space. This calculation was enabled by numerically exact categorical compression within the STP-DAS framework, which effectively eliminates the memory bottlenecks associated with storing both excitation lists and CI expansion vectors. Compared to previous state-of-the-art CI calculations, this work represents a 3-orders-of-magnitude increase in CI space, a 6-orders-of-magnitude increase in FLOP count, and a 6-orders-of-magnitude improvement in computational speed.

We introduced a categorically compressed representation of the CI expansion vectors and reformulated the STP-DAS σ -build algorithm to take advantage of this structure. By expressing the global expansion vector as a direct sum of compressed local components, the algorithm efficiently skips all coefficients that do not contribute to the categorical σ -vector. This approach is further enabled by a compression-compatible preconditioner, which generates compressed expansion directions within the Davidson procedure.

The resulting categorically compressed STP-DAS σ -build algorithm demonstrates excellent strong scaling behavior and yields dramatic reductions in both runtime and memory footprint. These benefits extend seamlessly to both relativistic (two- and four-component) and non-relativistic CI calculations. To highlight this capability, we computed the $\binom{100}{88}$ X2C-CASCI ground-state energy of HBrTe using over one quadrillion (10^{15}) complex-valued 2-spinor determinants. The categorically compressed STP-DAS approach spans 10^{15} determinants across 13 distributed active spaces, reducing the memory required for the excitation list from 12×10^9 GB to only 25 GB, and for nine CI expansion vectors from 134 PB to under 500 TB. It converges the ground-state wave function of HBrTe in just nine iterations over a 34.5-hour runtime. This achievement represents the largest CI calculation reported to date. Additionally, we achieved σ -build times of just 5 minutes for systems with approximately 150 billion complex-valued 2-spinor determinants using only a few compute nodes. The capability to perform large CI calculations makes basis set extrapolations to the complete basis set limit and computations on highly correlated molecular systems readily achievable with CI.

The integration of categorical compression with STP-DAS marks a paradigm shift in tackling large-scale CI problems. As quantum chemistry continues to push the limits of system complexity, the ability to carry out quadrillion-determinant calculations within tractable resource bounds establishes a powerful foundation for studying highly correlated, multireference, relativistic systems. While access to hundreds of compute nodes for quadrillion-determinant calculations may remain out of reach for most of the community, this work demonstrates the practical feasibility of trillion-determinant calculations on a small cluster.

For transition-metal, rare-earth, and heavy-element complexes, such large-scale CI calculations enable predictive simulations of electronic structure properties (bond order, covalency, polarization, etc.), spectroscopic observables (UV/Vis, X-ray, etc.), and reaction pathways, with the full orbital space consisting of both metal and ligand orbitals, treated on an equal footing.

The ability to simulate a full CI space of 100 orbitals on a classical computer not only challenges current notions of quantum supremacy, but also establishes a robust platform for developing and benchmarking quantum algorithms aimed at achieving chemical accuracy.

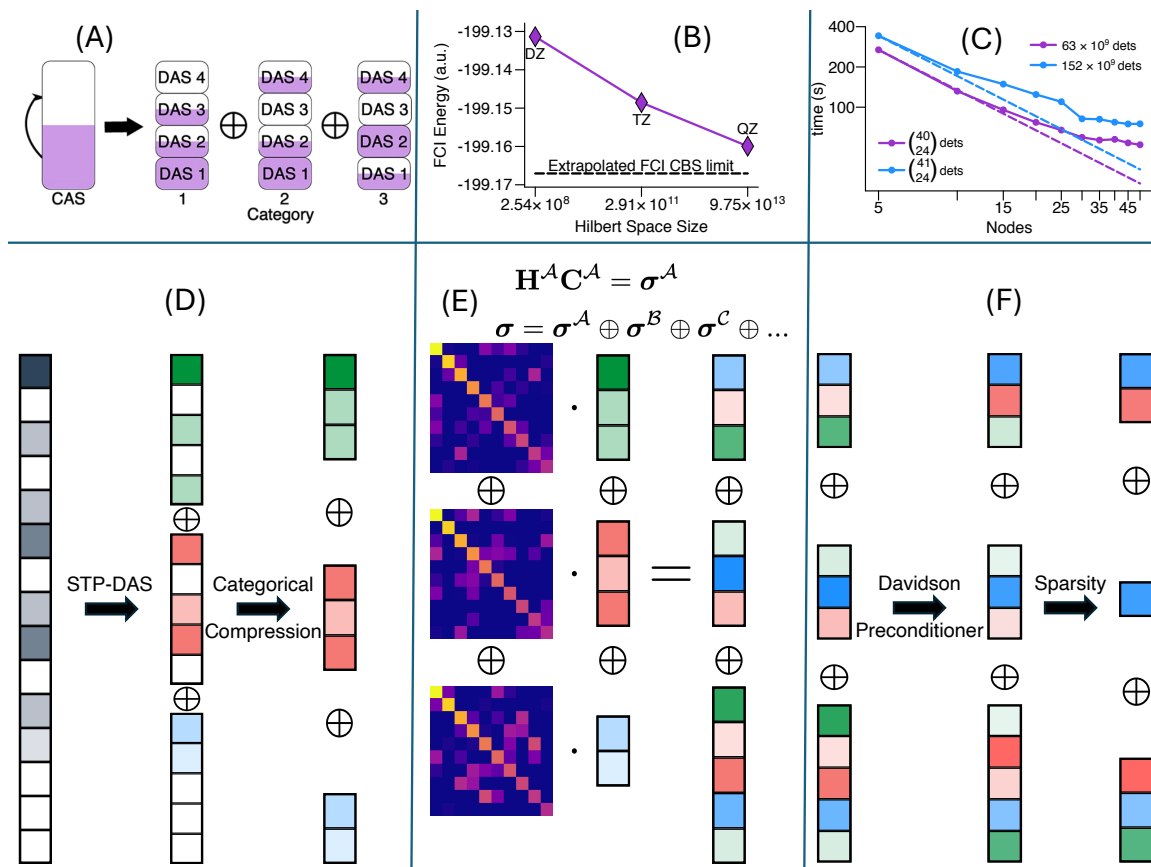


Figure 3.2. Categorical Compression within the small-tensor-product distributed active space (STP-DAS) framework. (A), The STP-DAS framework decomposition of a complete active space configuration interaction (CASCI) calculation into a direct sum of categorical excitations. The large excitation lists can be factored into much smaller categorical excitation lists. Purple sections within active spaces represent electron-occupied orbitals. (B), The exact two-component full configuration interaction (X2-FCI) ground state energy of the Mg^{2+} ion within the cc-pVNZ-DK [53, 229] ($N = 2, 3, 4$) basis sets, along with the extrapolated complete basis set limit. (C), Average execution time (in seconds) of the compression-compatible STP-DAS algorithm per Davidson iteration of a thallium hydride (TIH) test case versus the node count (5 iterations, 1 message passing interface (MPI) process per node, 40 symmetric multiprocessing (SMP) threads per MPI process). Here, \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, and σ is their product. The dashed lines illustrate the ideal strong scaling behavior of each CASCI calculation. (D), The representation of the subspace expansion vector in a traditional configuration interaction (CI) picture, the decomposed subspace vector in the STP-DAS framework, and the numerically exact compression of the subspace expansion vector in the categorically compression-compatible STP-DAS representation. The color of CI coefficients indicates their configuration category, while their brightness symbolizes their magnitude. White indicates a magnitude of zero. (E), A schematic representation of the lossless, compression-compatible, STP-DAS σ -build algorithm. The Hamiltonian matrix is represented as a heatmap, where brighter elements have larger magnitudes. The color of the vector elements indicates the configuration category of the corresponding CI coefficients. Note that the σ -build preserves categorical compression. (F), An illustration comparing the traditional Davidson preconditioner with the compression-compatible preconditioner to generate successive subspace expansion vectors. The compression-compatible preconditioner appends the subspace with the same effective search direction as the traditional Davidson preconditioner without compromising its compression.

Chapter 4

**ACCELERATING SMALL TENSOR PRODUCT DISTRIBUTED
ACTIVE SPACE (STP-DAS) CONFIGURATION INTERACTION
USING GRAPHICS PROCESSING UNITS (GPUS)****4.1 Introduction**

The formal description of all quantum mechanical systems is governed by the abstract Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \hat{H} |\psi\rangle, \quad (4.1)$$

where $|\psi\rangle$ is the state of the system and \hat{H} is the Hamiltonian governing it. In the context of the electronic structure theory of molecules in finite basis sets, the solution to Eq. (4.1) is given by the full configuration interaction (FCI) method [134,263,264,268,284,295]. The many-body state $|\psi\rangle$ is computed in a given basis set by explicitly diagonalizing the corresponding electronic Hamiltonian matrix. Its straightforward nature makes FCI applicable to a wide range of many-body problems. For instance, its adaptation to the two- and four-component Dirac formalisms [68,236] enables the non-perturbative description of relativistic quantum effects [2,17,18,120,121,130,132,146–149,183,238,260,270,289,300].

The explicit diagonalization of the Hamiltonian matrix has historically rendered its applicability to small chemical systems. This is because the number of possible configurations increases factorially with the system size, causing the dimension of the Hilbert space to explode. As such, even with its special properties (sparse, Hermitian, and diagonally-dominant), the diagonalization of the Hamiltonian matrix quickly becomes intractable.

Due to the explosive growth of the dimension of the space of configurations, practical calculations often rely on approximate configuration interaction (CI) wavefunction methods. The two main families of methods are complete active space CI (CASCI) [85–87,95,145–147,150,213,243,302] and selected CI (SCI) [76,80,99,117,118,124,168,179,252,262,291,320,321]. Both methods work by effectively truncating the Hilbert space of the system to a subspace

of significant determinants.

The CASCI method approximates the FCI wavefunction as the CI solution within the configuration space generated by a subset of MOs (called the active space). The outright exclusion of occupied and unoccupied orbitals often leads to an underestimation of dynamic correlation. Therefore, to attain quantitative agreement with experimental results, it is often required to apply additional methods on top of CASCI to recover the missing correlation. Such methods include multiconfigurational self-consistent field (MCSCF) [17, 94, 132, 144, 174, 185, 187, 188, 238, 270, 284, 289, 301, 305, 306, 319], multireference configuration interaction (MRCI) [115, 121, 174–176, 189, 246, 284, 304], and many-body perturbation theory (MRPT2, CASPT2, NEVPT2, MC-PDFT) [5–7, 35, 172, 183, 186, 244, 270, 300].

A recent CI framework [122] remedies the explosive memory footprint of the Hamiltonian matrix by factorizing the active space into small tensor products of distributed active spaces (DASs). In the small tensor product distributed active space (STP-DAS) approach, the computation of the large matrix-vector product associated with the iterative diagonalization of the Hamiltonian matrix is decomposed into the evaluation of a sequence of small tensor products. By employing an addressing scheme compatible with the small-tensor-product approach, the on-the-fly evaluation of matrix-vector products can significantly reuse the address strings local to each DAS to emulate the full excitation list. This reuse leads to dramatic reductions in the memory footprint of the Hamiltonian excitation list. This framework was shown to successfully facilitate CI calculations consisting of over one quadrillion (10^{15}) determinants [265].

The engine that powers CI is the matrix-vector product algorithm within the iterative diagonalization procedure, also known as the σ -build. As such, CI is amenable to large-scale parallel processing paradigms [8, 93, 95, 122, 211, 302], so long as subtle issues such as load balancing are properly addressed. Beyond the traditional host-centered paradigms of multi-threading (OpenMP, OpenACC) and message passing interface (MPI), electronic structure method development is increasingly adopting the use of graphics processing units (GPUs) and other accelerators to maximize throughput and parallelism [13, 77, 82, 84, 162, 191, 194, 228, 259, 266, 314]. Such accelerators consist of (tens of) thousands of cores purposefully designed to perform vectorized operations with peak efficiency. Their hardware design

implements the “single instruction, multiple data” (SIMD) principle, in which all cores execute the exact same operation on different data elements concurrently. While this limits the functionality of GPU cores compared to traditional central processing unit (CPU) cores, the design drastically reduces the per-core cost of vectorized operations such as matrix-vector products. Despite this remarkable capability, the effective exploitation of the SIMD cores of the accelerator requires careful consideration of the data flow and processing patterns of the underlying algorithm. This includes the clever distribution of the computational workload between the host and the accelerator and the implementation of pipelining. The latter is a technique that hides the overhead associated with data transfers to and from the accelerator by overlapping them with host and accelerator computation. In addition, the adoption of batched processing and coalesced data transfers is often needed to maximize GPU throughput and performance.

Despite their vast computational power, the video random-access memory (VRAM) capacity of GPUs is rather limited, hindering their ability to accelerate certain applications. This general restriction does not apply in the case of STP-DAS: the framework decomposes the large matrix-vector product into a sequence of small tensor products, each of which can comfortably fit into the limited memory of GPUs!

In this work, we accelerate the STP-DAS framework by adapting the two-electron portion of the small-tensor-product σ -build algorithm to GPUs. The resulting algorithm implements batched processing and coalesced data transfers within the Knowles-Handy approach [150, 152] to contract the Hamiltonian excitation list with CI expansion vectors. It supports multi-GPU and cooperative multi-process execution via MPI and Multi-Process Service (MPS). To showcase its features and performance, we use the accelerated small-tensor-product σ -build algorithm to conduct exact two-component CASCI (X2C-CASCI) calculations for the relativistic ground states of Au_2 and Xe_2 .

4.2 Methods

4.2.1 Accelerating STP-DAS σ -Build using GPUs

The σ -build algorithm within STP-DAS [122] implements the evaluation of

$$\sigma_{L^A} = {}^{1e}\sigma_{L^A} + {}^{2e}\sigma_{L^A}, \quad (4.2)$$

$${}^{1e}\sigma_{L^A} = \sum_{\mathcal{B}} \sum_{\mathbb{K}_{\mu}^{\mathcal{B}} \oplus \mathbb{K}_{\nu}^{\mathcal{B}}} \sum_{pq} P_{\mu\nu} \delta_{\bar{\mathbb{X}}_{\mu\nu}^A, \bar{\mathbb{X}}_{\mu\nu}^{\mathcal{B}}},$$

$$h'_{pq} \langle \mathbb{L}_{\mu}^A \oplus \mathbb{L}_{\nu}^A | \hat{E}_{pq} | \mathbb{K}_{\mu}^{\mathcal{B}} \oplus \mathbb{K}_{\nu}^{\mathcal{B}} \rangle C_{K^{\mathcal{B}}}, \quad (4.3)$$

$${}^{2e}\sigma_{L^A} = \frac{1}{2} \sum_{\mathcal{CB}} \sum_{\mathbb{J}_{\mu}^{\mathcal{C}} \oplus \mathbb{J}_{\nu}^{\mathcal{C}}} \sum_{\mathbb{J}_{\kappa}^{\mathcal{C}} \oplus \mathbb{J}_{\lambda}^{\mathcal{C}}} \sum_{\mathbb{K}_{\kappa}^{\mathcal{B}} \oplus \mathbb{K}_{\lambda}^{\mathcal{B}}} \sum_{pqrs} P_{\mu\nu} P_{\kappa\lambda}$$

$$\delta_{\bar{\mathbb{X}}_{\mu\nu}^A, \bar{\mathbb{X}}_{\mu\nu}^{\mathcal{C}}} \delta_{\bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{C}}, \bar{\mathbb{X}}_{\kappa\lambda}^{\mathcal{B}}} g_{pqrs} \langle \mathbb{L}_{\mu}^A \oplus \mathbb{L}_{\nu}^A | \hat{E}_{pq} | \mathbb{J}_{\mu}^{\mathcal{C}} \oplus \mathbb{J}_{\nu}^{\mathcal{C}} \rangle$$

$$\langle \mathbb{J}_{\kappa}^{\mathcal{C}} \oplus \mathbb{J}_{\lambda}^{\mathcal{C}} | \hat{E}_{rs} | \mathbb{K}_{\kappa}^{\mathcal{B}} \oplus \mathbb{K}_{\lambda}^{\mathcal{B}} \rangle C_{K^{\mathcal{B}}}, \quad (4.4)$$

where $p \in \mathbb{X}_{\mu}^A$, $q \in \mathbb{X}_{\nu}^{\mathcal{C}}$, $r \in \mathbb{X}_{\kappa}^{\mathcal{C}}$, $s \in \mathbb{X}_{\lambda}^{\mathcal{B}}$, $\langle \mathbb{X}_{\mu}^A \oplus \mathbb{X}_{\nu}^A | \hat{E}_{pq} | \mathbb{X}_{\mu}^{\mathcal{B}} \oplus \mathbb{X}_{\nu}^{\mathcal{B}} \rangle$ are categorical one electron excitation lists, $P_{\mu\nu}$ are global phase factors, h'_{pq} and g_{pqrs} are one- and two-body Hamiltonian elements, and $C_{K^{\mathcal{B}}}$ are CI coefficients of determinants K belonging to configuration category \mathcal{B} (see Ref. 122 for definitions and algorithmic details). From the summation structure of Eq. (4.4), it is evident that the computational bottleneck is the evaluation of the two-electron portion of the σ -build. Consequently, it has historically been the focus of algorithmic optimization efforts [150, 152, 213].

One such optimization is inspired by the Knowles-Handy approach [150, 152], which evaluates Eq. (4.4) using vectorized operations, thereby enabling the full exploitation of state-of-the-art cache optimization techniques. Vectorization, which comes at the cost of additional memory requirements, is achieved by organizing the Hamiltonian elements and the CI coefficients into *intermediates* and matrix-multiplying them:

$$X_{rs,pq}^{L^A, J^C} \leftarrow g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle, \quad (4.5)$$

$$\Omega_{rs, \bar{\mathbb{K}}_{\kappa\lambda}^B}^{J^C, K^B} \leftarrow C_{K^B} \text{ such that } \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle \neq 0, \quad (4.6)$$

$$\Lambda_{pq, \bar{\mathbb{K}}_{\kappa\lambda}^A}^{L^A} \leftarrow \left(X_{rs,pq}^{L^A, J^C} \right)^\top \times \Omega_{rs, \bar{\mathbb{K}}_{\kappa\lambda}^B}^{J^C, K^B}, \quad (4.7)$$

where $\bar{\mathbb{K}}$ and $\bar{\mathbb{L}}$ are the *local* addresses of the portion of the determinant K not belonging to the DASs numbered κ or λ . The locality of the addresses renders the column dimensions of the intermediates in Eqs. (4.6) and (4.7) small, making the evaluation of Eq. (4.4) on GPUs possible despite VRAM limitations. This small dimensionality, however, presents challenges in maximizing GPU computational throughput and minimizing frequent data transfers between the CPU and the accelerator.

Computational throughput can be maximized by the incorporation of *batching* in the evaluation of Eq. (4.4). Instead of successively forming Eqs. (4.5) and (4.6) and computing Eq. (4.7) for every connecting determinant J^C , one can form a batch $\left\{ X_{rs,pq}^{L^A, J^C}, \Omega_{rs, \bar{\mathbb{K}}_{\kappa\lambda}^B}^{J^C, K^B} \right\}_{J^C}$ of intermediates and evaluate Eq. (4.7) for the entire batch concurrently. While this approach has a larger memory footprint than naive successive evaluation, its computational advantages are twofold: the increased utilization of the compute cores of the accelerator, and the coalescence of data transfers between the CPU and the accelerator [169].

Because the batching technique organizes the evaluation of Eqs. (4.5) and (4.6) into large computational tasks, it also enables the implementation of GPU *pipelining*, which hides the cost of large data transfers by overlapping them with significant CPU and GPU computation. To this end, we adopted the use of asynchronous operations in our implementation, in which the CPU submits a task to be executed by the GPU without needing to wait for the GPU to complete it. Consequently, upon forming Eqs. (4.5) and (4.6) for batch n , the CPU can form intermediates for batch $(n+1)$ while, *concurrently*, the intermediates of batch n are transferred to the GPU, the GPU evaluates Eq. (4.7) for batch n , and the computed contractions for batch n are transferred back to the CPU for processing into the σ -vector, as prescribed in Eq. (4.4). This pipelined processing of batches of intermediates not only effec-

tively hides the latency associated with large and frequent data transfers between the CPU and the accelerator, but also dramatically shrinks the idle time the CPU spends waiting for the GPU to complete the contractions of intermediates. The precise execution pattern that yields these optimizations is depicted in Algorithm 2.

The computational load distribution structure within the STP-DAS algorithm [122] naturally enables the multi-GPU acceleration of the evaluation of Eq. (4.4). By assigning each MPI rank to a unique GPU, Algorithm 2 can be executed on multiple accelerators to compute and populate distinct sectors of the σ -vector concurrently. Cooperative multi-process execution can also be realized by assigning multiple MPI ranks to the same GPU and partitioning it by utilizing technologies such as NVIDIA’s MPS or multi-instance GPU (MIG). This can maximize the performance of Algorithm 2 on multi-node, single-accelerator systems. Notably, this approach to multi-GPU multi-process execution does not require communication among the accelerators. Therefore, in contrast with other implementations [1], Algorithm 2 does not require the use of GPU-aware MPI libraries.

4.3 Results

To showcase the dependence of the intermediates batch size on the runtime of Algorithm 2, we conducted a sequence of calculations for the X2C-CASCI ground state of Au₂ with varying batch sizes within the compression-compatible STP-DAS framework [265]. The calculations were executed on the University of Washington’s AI-accelerated research computing platform Tillicum, a small-sized cluster where each node is equipped with an Intel Emerald Rapids CPU, eight NVIDIA Hopper-H200 SXM GPUs, a 400 GB/s network interface card, and an NVLink™ 4.0 GPU interconnect providing 900 GB/s GPU-to-GPU bandwidth.

As seen in Figure 4.1, the contraction speedup greatly depends on the batch size used in Algorithm 2. This is expected: the larger the batch size, the more likely the calculation is to fully saturate all compute cores of the GPU, resulting in ideal throughput and performance. This way, the GPU computation overlaps with the intermediate batch formation on the hosting CPU, resulting in the complete hiding of GPU computation and data transfers latency. Past a certain point, further increasing the size of the batch has at most marginal

Algorithm 2: The GPU-accelerated two-electron STP-DAS σ -build. Bold text represents algorithmic logic and typewritten text represents comments. Curly brackets represent code-interfacing operations.

```

// Process first batch  $B_1$  of  $J^C$ 's
1  $B_1 \leftarrow \{0, \dots, \min(S_{\text{batch}}, |\{J^C\}|)\}$ 
// Form Eqs. (4.5) and (4.6) for all  $J^C \in B_1$ 
2 for  $J^C \in B_1$  do
3  $\left[ \begin{array}{l} X^{B_1} \leftarrow X^{B_1} \cup \left\{ \left\{ g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle \right\} \right\} \\ \Omega^{B_1} \leftarrow \Omega^{B_1} \cup \left\{ \left\{ C_{KB} : \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle \neq 0 \right\} \right\} \end{array} \right.$ 
4 {Asynchronously task GPU with contracting intermediates in  $B_1$ }
// Continue processing next batch  $B_n$  without waiting for the GPU
5 for  $n = 2, \dots, |\{J^C\}|/S_{\text{batch}}$  do
6  $B_n \leftarrow \{(n-1) \times S_{\text{batch}}, \dots, n \times S_{\text{batch}}\}$ 
7  $X^{B_n}, \Omega^{B_n}, \Lambda^{B_n} \leftarrow \emptyset$ 
// Form Eqs. (4.5) and (4.6) for all  $J^C \in B_n$ 
8 for  $J^C \in B_n$  do
9  $\left[ \begin{array}{l} X^{B_n} \leftarrow X^{B_n} \cup \left\{ \left\{ g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle \right\} \right\} \\ \Omega^{B_n} \leftarrow \Omega^{B_n} \cup \left\{ \left\{ C_{KB} : \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\kappa^B \oplus \mathbb{K}_\lambda^B \rangle \neq 0 \right\} \right\} \end{array} \right.$ 
10 {Asynchronously task GPU with contracting intermediates in  $B_n$ }
11 {Wait for the GPU to finish processing batch  $B_{n-1}$ }
// Process completed batch  $\Lambda^{B_{n-1}}$  into  ${}^{2e}\sigma$ 
12 for  $J^C \in B_{n-1}$  do
13  $\left[ \begin{array}{l} {}^{2e}\sigma_{LA+} = \Lambda_{pq, \mathbb{K}_{\kappa\lambda}^A}^{L^A} \text{ such that } \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \end{array} \right.$ 
14  $B_{\text{leftover}} \leftarrow \{(|\{J^C\}|/S_{\text{batch}} - 1) \times S_{\text{batch}}, \dots, |\{J^C\}|\}$ 
15 {Wait for the GPU to finish processing batch  $B_{\text{leftover}}$ }
// Process the leftover batch into  ${}^{2e}\sigma$ 
16 for  $J^C \in B_{\text{leftover}}$  do
17  $\left[ \begin{array}{l} {}^{2e}\sigma_{LA+} = \Lambda_{pq, \mathbb{K}_{\kappa\lambda}^A}^{L^A} \text{ such that } \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \end{array} \right.$ 

```

benefits, as the compute capabilities of the GPU are already fully saturated. The modest increase in speedup at this regime can be attributed to the coalescence of data transfers, since at 900 GB/s, the NVLink™ 4.0 GPU interconnect may still be unsaturated.

To study the strong-scaling behavior of Algorithm 2, we conducted compression-compatible STP-DASCI [265] calculations for the X2C-CASCI ground state of Xe₂ with varying number of MPI ranks, each assigned a unique GPU. As seen in Figure 4.2, the excellent strong-

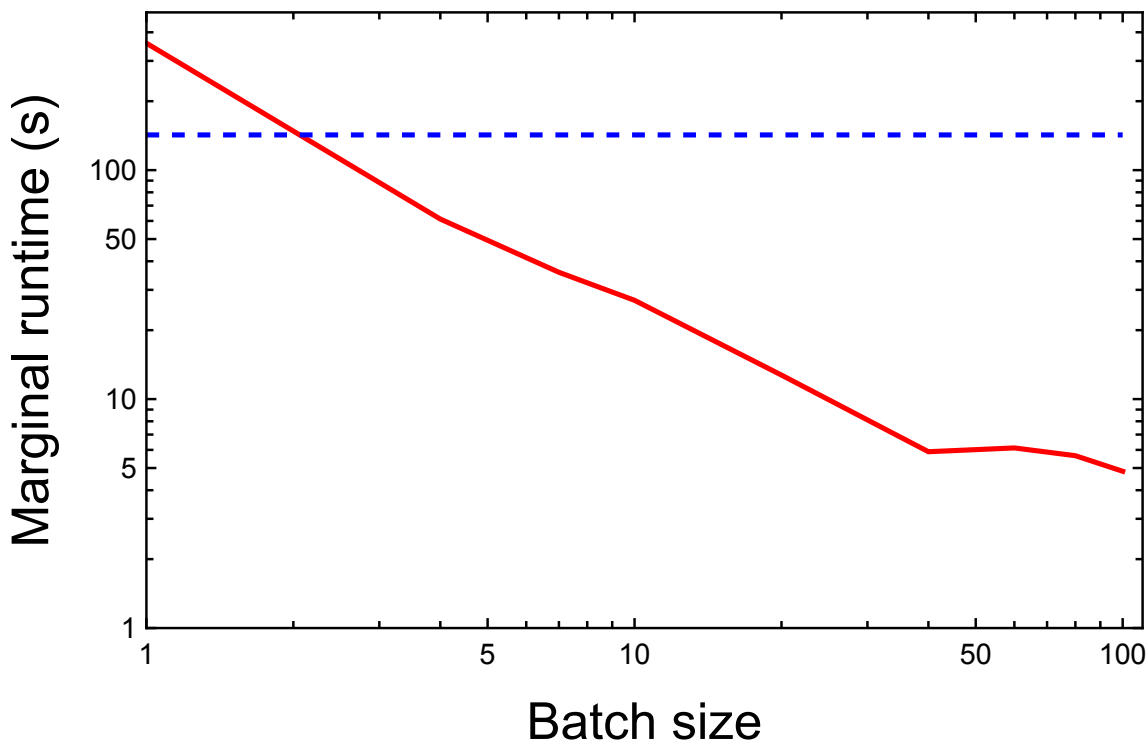


Figure 4.1. (Red: 8 CPU cores + 8 GPUs) Marginal runtime of the GPU portion of Algorithm 2 (matrix multiplications and data transfers) for various batch sizes of the fourth σ -build iteration for an Au_2 test case using an active space of 48 orbitals and 8 electrons within the x2c-TZVPall-2c [227] basis set. (Blue: 8 CPU cores) The reference runtime of an identical calculation with a CPU-only version of Algorithm 2 [122]. A preconditioning dropping threshold of $\varepsilon = 10^{-16}$ was used in all calculations. Both axes are log-scaled for convenience.

scaling behavior inherent in the STP-DAS framework [122] directly transfers to the GPU-accelerated version of the same. Moreover, the depicted strong-scaling behavior clearly illustrates that Algorithm 2 does not require inter-GPU communication, resulting in excellent load balancing and minimal communication latency.

4.4 Conclusions

In this work, we showcased the acceleration of the STP-DAS σ -build algorithm using GPUs. By adopting traditional GPU optimization techniques such as batching and pipelined execution, the algorithm significantly speeds up the evaluation of matrix-vector contractions and enables rapid CI calculations. The algorithm naturally supports multi-GPU processing

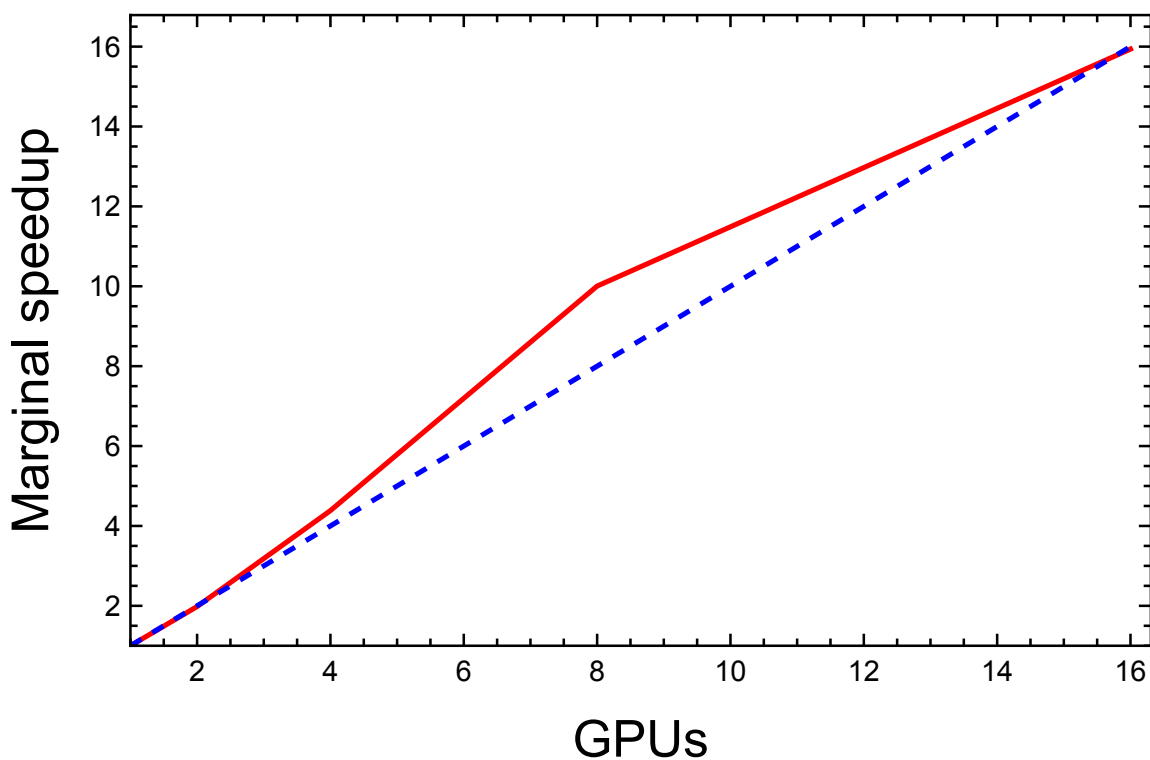


Figure 4.2. (Red) Marginal speedup of the GPU portion of Algorithm 2 (matrix multiplications and data transfers) as a function of the number of GPUs used. (Blue) The corresponding ideal strong-scaling behavior. The data used the fourth σ -build iteration of an Xe_2 test in an active space of 48 orbitals and 8 electrons within the x2c-TZVPall-2c [227] basis set. A preconditioning dropping threshold of $\epsilon = 10^{-16}$ was used in all calculations.

and exhibits excellent strong scaling behavior.

Our benchmark studies demonstrate the ability of the GPU-accelerated STP-DAS σ -build algorithm to massively accelerate the execution of compressed CI calculations. The batched processing of intermediates within the algorithm implementation delivers excellent throughput and maximizes GPU utilization. Such a workload distribution is demonstrated to effectively hide the latency associated with data transfers, making the use of accelerators highly efficient. The accelerated algorithm is also shown to exhibit excellent strong-scaling behavior, similarly to the original STP-DAS σ -build procedure [122]. This enables multi-GPU processing of compressed CI calculations with minimal communication overhead, capable of facilitating large-scale calculations at previously-unattainable speed.

Together with categorical compression within the STP-DAS framework [265], the ability to carry out accelerated large-scale CI calculations in an inherently relativistic level-of-theory facilitates the accurate simulation of the electronic structure properties and observables of strongly-relativistic systems such as transition metals, rare-earth complexes, and heavy elements.

Chapter 5

**BENCHMARKING COUPLED CLUSTER AND DENSITY MATRIX
RENORMALIZATION GROUP: ACCURACY AND CONVERGENCE
TRENDS AGAINST LARGE SCALE NUMERICALLY EXACT
CONFIGURATION INTERACTION****5.1 Introduction**

Full configuration interaction (FCI) provides the formally exact description of the electronic structure of a chemical system within a given basis set by explicitly solving the eigenvalue problem that arises from the quantum equation of motion [134, 263, 264, 268, 284, 295],

$$\hat{H}|\psi\rangle = E|\psi\rangle, \quad (5.1)$$

where \hat{H} is the Hamiltonian governing the dynamics of the system, $|\psi\rangle$ is the electronic quantum state, and E is its energy. The non-perturbative nature of FCI enables the treatment of strong relativistic effects, such as spin-orbit and spin-spin couplings [2, 17, 18, 120, 121, 130, 132, 146–149, 183, 238, 260, 270, 289, 300].

The FCI solution, therefore, is obtained by diagonalizing a sparse many-electron Hamiltonian matrix. Due to its size, the Hamiltonian matrix is almost never stored explicitly, and the low-lying spectrum is computed iteratively. However, because the dimension of the Hilbert space grows factorially with the system size, even iterative methods become intractable for all but the smallest chemical systems. The relativistic regime is even more restrictive because of the larger configuration space associated with complex-valued two- and four-component wavefunctions. The dimensionality problem of FCI can be partially alleviated using parallel processing schemes, algorithmic optimization, and data compression techniques [8, 93, 95, 122, 211, 265, 302]. Even after these optimizations, however, the diagonalization procedure requires enormous amounts of computational resources and time, pushing practitioners and users towards alternative computational methods.

Many alternative wavefunction-based methods were developed to address the inherent memory limitations of CI. These methods seek approximate solutions to the FCI problem by utilizing various approximations to truncate the dimensionality of the system. Within the CI framework, a straightforward strategy is to take the Hartree-Fock determinant as the reference and to truncate the determinant excitation level up to double or triple excitations. This leads to the CISD and CISDT methods with polynomial scaling at $\mathcal{O}(n_o^2 n_v^4)$ and $\mathcal{O}(n_o^3 n_v^5)$, where n_o and n_v denote the number of spin-orbitals that are occupied and unoccupied in the reference determinant, respectively. However, truncated CI methods suffer from slow convergence to FCI and size-consistency problems [14]. Two other ubiquitous classes of methods approximating FCI are coupled cluster (CC) and density matrix renormalization group (DMRG). Both methods are post-Hartree-Fock methods, generating an approximated wavefunction solution from a mean-field reference Slater determinant. We refer the readers to Ref. 15 for a thorough review of the CC family of methods.

At the same scaling as truncated CI methods, the coupled cluster series of methods provides more rapid convergence towards FCI [14]. In CC methods, electron correlation is recovered into the mean-field reference determinant by including excited configurations in the form of an exponential ansatz [296]:

$$|\psi\rangle = e^{\hat{T}} |\Phi\rangle = e^{\hat{T}_1 + \hat{T}_2 + \dots} |\Phi\rangle, \quad (5.2)$$

where Φ is the Hartree-Fock determinant and T_n is the n -body cluster operator, defined as

$$\hat{T}_n \equiv \left(\frac{1}{n!}\right)^2 \sum_{ij\dots ab\dots}^n t_{ij\dots}^{ab\dots} \hat{a}^\dagger \hat{b}^\dagger \dots \hat{j} \hat{i}, \quad (5.3)$$

where $t_{ij\dots}^{ab\dots}$ are cluster amplitudes and $\hat{a}^\dagger \hat{b}^\dagger \dots$ and $\hat{j} \hat{i} \dots$ are creation and annihilation operators of unoccupied $ab\dots$ and occupied $ij\dots$ orbitals, respectively. Often, the expansion is truncated at a particular excitation level, with the most common choices being $\hat{T} = \hat{T}_1 + \hat{T}_2$ and $\hat{T} = \hat{T}_1 + \hat{T}_2 + \hat{T}_3$, resulting in the coupled cluster singles and doubles (CCSD) [47, 230] and coupled cluster singles, doubles, and triples (CCSDT) [208, 256] methods. The cluster

amplitudes $t_{ij\dots}^{ab\dots}$ are usually solved iteratively in a projectional manner by requiring

$$\langle \Phi_{ij\dots}^{ab\dots} | \bar{H} | \Phi \rangle = 0 \quad \forall |\Phi_{ij\dots}^{ab\dots}\rangle, \quad (5.4)$$

in which $\bar{H} = e^{-\hat{T}} \hat{H} e^{\hat{T}}$ is the similarity-transformed Hamiltonian and $\Phi_{ij\dots}^{ab\dots}$ are excited determinants where electrons are excited from occupied orbitals $ij\dots$ to unoccupied orbitals $ab\dots$. The CC correlation energy is computed as

$$E^{(\text{CC})} = \langle \Phi | \bar{H} | \Phi \rangle. \quad (5.5)$$

While the CCSD method is widely used in medium-size quantum chemistry calculations with hundreds of electrons and orbitals, the CCSDT method is much less commonly utilized due to its iterative $\mathcal{O}(n_o^3 n_v^5)$ FLOPs scaling and $\mathcal{M}(n_o^3 n_v^3)$ memory requirement to store the t_{ijk}^{abc} amplitudes. Methods of greater accuracy than CCSD and smaller computational footprint than CCSDT have been developed, among which is the CCSD(T) method [15]. In the CCSD(T) method, after a converged CCSD calculation, the correlation energy is improved by a one-shot $\mathcal{O}(n_o^3 n_v^4)$ calculation without explicit storage of the t_{ijk}^{abc} amplitudes [16, 233]. We refer readers to Ref. 233 for the details of this non-iterative approximation.

Correlation energy is usually artificially separated into two parts [15]:

1. Dynamic correlation, which keeps electrons apart, and is composed of small contributions from a large number of determinants;
2. Static correlation (or non-dynamic correlation), which describes near-degenerate configurations, and is composed of large contribution from a small fraction of determinants.

All CC methods mentioned above are in the category of single-reference methods, in which dynamic correlation is described well by the exponential ansatz. However, their single-reference nature is not efficient in describing static correlation. Many multi-reference CC methods have been reported to describe static and dynamic correlation equally well [79, 184, 216], but they are beyond the scope of this chapter.

A common alternative to the truncated CI and coupled cluster approaches is the density matrix renormalization group (DMRG) method, in which the wavefunction is represented by a tensor network [10, 38, 311]. Tensor-network wavefunctions are factorizations of the exponentially-large FCI coefficient tensor that come in various flavors, including the matrix product state (MPS) [10, 38, 311], the tree tensor network (TTN) state [45, 199, 203, 269, 285], and the projected entangled pair states (PEPS) [298]. The DMRG algorithm allows for the efficient variational optimization of MPS wavefunctions. The MPS ansatz is given by

$$|\Psi\rangle = \sum_{\sigma} \sum_{a_1}^m \cdots \sum_{a_{N-1}}^m M_{1,a_1}^{\sigma_1} M_{a_1,a_2}^{\sigma_2} \cdots M_{a_{N-1},1}^{\sigma_N} |\sigma_1, \sigma_2, \cdots, \sigma_N\rangle, \quad (5.6)$$

where σ are sites with $|\sigma_1, \sigma_2, \cdots, \sigma_N\rangle$ representing an occupation vector, N is the number of sites, M are site tensors with auxiliary indices a that are truncated at a given bond dimension m . For quantum chemical problems, the orbitals are mapped directly to sites.

Because electronic-structure Hamiltonians are weakly-coupled (rendering perturbative methods [5–7, 35, 172, 183, 186, 244, 270, 300] applicable), the MPS ansatz scales polynomially in the number of parameters and the DMRG optimization scales polynomially in the system size. This combination makes DMRG an attractive method for electronic structure problems, as it allows for a polynomially-scaling algorithm with a compact ansatz and a controllable convergence to the FCI result (by increasing the bond dimension m). Notably, the DMRG approach approximately solves a corresponding CI problem, and therefore may suffer the same lack of dynamic correlation if the underlying CI space is truncated. Therefore, methods to recover dynamic correlation into truncated CI solutions may be used to recover the same into the DMRG tensor-network wavefunction [22, 39, 104, 108, 161, 204, 205, 240, 248, 261, 312].

The practical validity of approximate methods such as CC and DMRG can ultimately be assessed by their predictive accuracy relative to experiments [322]. Such experimental benchmarks, however, cannot resolve the error due to approximations underlying a given method, because the discrepancy with experimental results will also include errors sourced in the finite atomic basis set and the level of theory used in the studies. Alternatively, more controlled benchmark studies can be conducted by comparing the chemical observables

computed using various computational methods with those of FCI. Because FCI yields the formal wavefunction solution of the chemical system within a given atomic basis and level of theory, such comparisons trivially control for the effects these truncations have on the computed observables relative to experiments. Because of the explosive growth in memory requirements associated with FCI, such studies have been restricted to only the smallest chemical systems. To overcome this limitation, one could in principle obtain benchmark values from approximated CI methods, such as the various flavors of selected CI (SCI) [76, 80, 99, 117, 118, 124, 168, 179, 252, 262, 291, 320, 321]. While such studies can provide useful insights into the convergence behavior of the benchmarked methods, their conclusiveness is limited. This is because approximated CI methods have an unknown domain of validity at large scale, as they themselves rely on approximations within the CI expansion to reduce memory and computational costs.

Unlike benchmarks against approximated CI methods, the results presented here definitively classify the error in the computed coupled-cluster and DMRG states. This is due to the lossless nature of the compression-compatible STP-DAS method [265]. The small-tensor-product decomposition of the Hamiltonian excitation list and the categorical compression of CI expansion vectors allow for the *exact* computation of σ -vectors, even in extremely large Hilbert spaces. This then facilitates the exact computation of the Ritz residual of the CI vector, which can be used to obtain *exact* error bounds on the true eigenvalue of the Hamiltonian matrix within the given basis [325]. Of these bounds, the gap theorem [50, 217, 327] often offers the tightest lower bound, which is computed using the Ritz eigenvector and eigenvalue as

$$|\delta E| \leq \frac{\|\mathbf{r}\|^2}{\gamma_0}, \quad (5.7)$$

where \mathbf{r} is the Ritz residual of the desired state and the gap $\gamma_0 \equiv E_1 - \tilde{E}_0$ is the difference between E_1 , the (unknown) exact energy of the first excited state, and \tilde{E}_0 , the computed Ritz value of the ground state. Because γ_0 is unknown, one can estimate its order of magnitude using approximate methods or use experimental values to compute a surrogate for the true gap. This is often sufficient for gapped systems, for which γ_0 is much larger than the values $\|\mathbf{r}\|^2$ reasonably attains after a few Davidson iterations [49].

In more pathological cases, such as ground-state near-degeneracy, one can instead obtain an exact lower bound on the gap γ_0 by including the posterior error bound of the first excited state [247, 313, 327] in the Davidson calculation [325]:

$$\gamma_0 = E_1 - \tilde{E}_0 \geq \left(\tilde{E}_1 - \|\mathbf{r}_1\| \right) - \tilde{E}_0 \equiv \gamma_0^-, \quad (5.8)$$

where \mathbf{r}_1 is the residual associated with the Ritz value \tilde{E}_1 . Combined with the variational bound associated with all Krylov subspace methods [325],

$$E_0 \leq \tilde{E}_0, \quad (5.9)$$

one can place tight upper and lower bounds on the true eigenvalue:

$$\tilde{E}_0 - \frac{\|\mathbf{r}\|^2}{\gamma_0} \leq E_0 \leq \tilde{E}_0. \quad (5.10)$$

In this chapter we use the exact eigenvalue bounds provided by (5.10) within the compression-compatible STP-DAS framework [265] to analyze the convergence properties and absolute accuracy of relativistic CC and DMRG. To this end, we used the methods to compute the ground states of three systems of varying degrees of symmetry and correlation: HBrTe, Rb₄, and Xe₂. Relativistic effects were incorporated variationally within the Dirac formalism using the one-electron exact two-component (1eX2C) variant of the exact two-component (X2C) method along with a Dirac-Coulomb-Breit-parameterized Boettger factor [72]. The 1eX2C method uses the one-electron portion of the Hamiltonian to approximately reduce the four-component Dirac equation into a two-component eigenvalue problem for the positive energy (electronic) sector of the Hamiltonian [64–67, 69, 70, 100, 119, 126, 155, 163, 171, 177, 178, 180, 181, 183, 221, 222, 260].

The ability to obtain exact error bounds for the results of large-scale CI calculations opens the door to previously intractable precise benchmarking of various computational chemistry methods beyond small systems. These bounds reveal the range in which the true energy eigenvalue falls, and thus allow the benchmarked methods to be faithfully judged at various correlation regimes.

5.2 Results

The relativistic coupled-cluster ground state energies of HBrTe (100 2-spinor orbitals, 88 electrons, x2c-TZVPall [227]), Rb₄ (50 2-spinor orbitals, 28 electrons, cc-pVTZ-x2c [114]), and Xe₂ (60 2-spinor orbitals, 12 electrons, x2c-TZVPall-2c [227]) were computed using the X2C-CC method at various excitation-truncation levels. These results were obtained from the coupled-cluster code [12, 44, 156, 170, 315–317, 322] within the Chronus Quantum software package [308]. The relativistic DMRG energies were calculated using the X2C-DMRG method within PySCF [280–282]. All calculations were conducted on the National Energy Research Scientific Computing Center’s Perlmutter high-performance supercomputer (AMD EPYC 7763 Milan, 128 compute cores per node, 512 GB of RAM per node, and 200 GB/s NIC). The largest X2C-DMRG calculation (HBrTe, 100 2-spinor orbitals, 88 electrons, bond dimension $m = 1000$) was conducted on 30 such nodes, while the corresponding X2C-CC calculation was conducted on 4 nodes.

Table 5.1 shows the discrepancy between the computed X2C-CC energies and the corresponding X2C-CASCI energies, which were previously obtained using the compression-compatible STP-DAS method [265]. It is evident that while all benchmarked excitation-restriction levels attain near-chemical accuracy, none attain microhartree accuracy. The latter is crucial when the resolution of individual wavenumbers is needed, such as in spectroscopy simulations.

Unsurprisingly, the most accurate coupled-cluster energies are obtained by CCSDT, which incorporates triple-excitation amplitudes converged iteratively rather than one-step perturbatively (as is the case with CCSD(T)). Notably, the error in the X2C-CCSDT ground state energy of Rb₄ is an order of magnitude larger than in the other two systems. This behavior is expected because coupled cluster is inherently a single-reference method. Because the ground state of Rb₄ is a relativistic analogue of the H₄ system [96], it is highly multireference, and coupled cluster will not accurately capture its static correlation.

Additionally, these benchmarks showcase the non-variational nature of coupled-cluster-based methods. This is evident in the negative entry in Table 5.1 corresponding to the X2C-CCSDT ground state energy of Xe₂. At first glance, because of Eq. (5.9), one might assume

that it is plausible that the X2C-CCSDT energy is closer to the true energy eigenvalue of the exact two-component Hamiltonian of Xe_2 (60 2-spinor orbitals, 12 electrons, x2c-TZVPall-2c [227]) than the benchmark X2C-CASCI value [265]. However, that is not the case: applying the result of Eq. (5.10) reveals the following bounds on the error between the computed X2C-CASCI energy and the true eigenvalue of the Hamiltonian matrix [265]:

$$0 \leq \tilde{E}_0 - E_0 \leq 1.756 \times 10^{-5}. \quad (5.11)$$

This result definitively shows that the X2C-CCSDT ground state energy of the Xe_2 system demonstrates the nonvariational nature of coupled-cluster-based methods: the X2C-CCSDT value is guaranteed to be *at least* $29.7 \mu E_h$ below the *exact* energy eigenvalue.

Table 5.2 shows the discrepancy between the computed X2C-DMRG energies and the corresponding X2C-CASCI benchmark values [265]. The results demonstrate the variationality of DMRG energies with increasing bond dimension. Furthermore, the DMRG energies converge monotonically to the benchmark values. This is because the formal solution of the eigenvalue problem can be obtained by DMRG at infinite bond dimension. In contrast with the coupled cluster results of Table 5.1, it is evident that DMRG performs significantly better on Rb_4 , a multireference system dominated by static correlation, than on Xe_2 , which is dominated by dynamic correlation. This is because the finite bond dimension truncation associated with DMRG limits the entanglement between sites (orbitals), thus failing to accurately capture dynamic correlation.

Table 5.1. The discrepancy (in Hartree), $E_{\text{CC}} - E_{\text{CASCI}}$, between X2C-CC ground state energies and the corresponding X2C-CASCI values for HBrTe (100 2-spinor orbitals, 88 electrons, x2c-TZVPall [227]), Rb_4 (50 2-spinor orbitals, 28 electrons, cc-pVTZ-x2c [114]), and Xe_2 (60 2-spinor orbitals, 12 electrons, x2c-TZVPall-2c [227]).

	HBrTe	Rb_4	Xe_2
CCSD	1.26×10^{-3}	7.06×10^{-4}	2.52×10^{-3}
CCSD(T)	1.97×10^{-4}	2.27×10^{-4}	3.41×10^{-4}
CCSDT	1.62×10^{-5}	1.61×10^{-4}	-4.73×10^{-5}

Table 5.2. The discrepancy (in Hartree), $E_{\text{DMRG}} - E_{\text{CASCI}}$, between X2C-DMRG ground state energies and the corresponding X2C-CASCI values at various bond dimensions. The systems are HBrTe (100 2-spinor orbitals, 88 electrons, x2c-TZVPall [227]), Rb₄ (50 2-spinor orbitals, 28 electrons, cc-pVTZ-x2c [114]), and Xe₂ (60 2-spinor orbitals, 12 electrons, x2c-TZVPall-2c [227]).

Bond dimension	HBrTe	Rb ₄	Xe ₂
$m = 50$	4.01×10^{-3}	1.44×10^{-4}	3.82×10^{-2}
$m = 100$	1.71×10^{-3}	4.99×10^{-5}	1.22×10^{-2}
$m = 200$	6.31×10^{-4}	1.63×10^{-5}	9.93×10^{-3}
$m = 500$	1.50×10^{-4}	2.08×10^{-6}	5.58×10^{-3}
$m = 1000$	5.24×10^{-5}	5.32×10^{-7}	2.40×10^{-3}

5.3 Conclusions

The ability to perform extremely large-scale numerically exact CI calculations, enabled by categorical compression within the STP-DAS CI framework [265], facilitates previously intractable precise benchmarking of the convergence behavior and accuracy of all other computational quantum chemistry methods. In this work, we performed such benchmark studies on two common (and complementary) computational chemistry methods: coupled cluster and DMRG. The former relies on an exponential ansatz to capture dynamic correlation with unparalleled efficiency, and the latter approximates the underlying eigenvalue problem by a tensor network of a finite bond dimension, easily capturing multireference effects. The benchmark studies were conducted using the HBrTe, Rb₄, and Xe₂ systems. These systems exhibit relativistic effects such as spin-orbit coupling and various degrees of point-group symmetry and correlation. Because these systems have a significant number of correlated orbitals and electrons, obtaining accurate benchmark energies (along with exact error bounds) is only made possible by the compression-compatible STP-DAS framework.

The benchmark study of coupled cluster showcases its ability to capture the dynamic correlation of single-reference systems. At the highest order tested (CCSDT), coupled cluster can reproduce the CI energies of single-reference systems (HBrTe and Xe₂) to within a few wavenumbers. The same benchmark study, however, also reveals an order-of-magnitude worse performance on multireference systems (Rb₄), as expected. Additionally, the non-variational nature of coupled cluster is definitively demonstrated by proving that the X2C-

CCSDT energy of Xe_2 is at least $29.7 \mu E_h$ lower than the true energy eigenvalue of the Hamiltonian in the given basis.

The benchmark study of DMRG reveals an expected, complementary convergence behavior to that of coupled cluster. As shown in Table 5.2, DMRG can easily recover the CI energy of Rb_4 even at modest bond dimensions. At the same time, even large bond dimensions are insufficient to recover the CI energy of the dynamically-correlated systems (HBrTe and Xe_2) to the level of accuracy often needed in quantum chemical applications. While expected, these results help quantify the error made by the tensor network approximation within the DMRG method.

Benchmark studies such as those reported here allow for the accurate assessment of the plethora of computational methods in quantum chemistry on systems of varying levels of molecular symmetry and correlation. This allows for the determination of the domain of validity of these methods, which ensures their predictive capabilities within simulations of electronic structure properties (bond order, covalency, polarization, etc.), spectroscopic observables (UV/Vis, X-ray, etc.), and reaction pathways.

Chapter 6

CONCLUSIONS AND OUTLOOK

In Chapters 2 and 3, we presented the STP-DAS framework and its ability to drastically shrink the memory footprint of the Hamiltonian matrix and state vectors required to obtain the solution of the electronic wave equation via configuration interaction. The reduction in memory footprint, rooted in the small-tensor-product factorization of the active orbital space and the corresponding categorical compression of CI expansion vectors, results in a nontrivial indexing and processing patterns when computing the σ -vector defined in Eq. (1.40). This shifts some of the overall computational burden away from memory requirements and toward sheer processing time, ultimately facilitating CI calculations consisting of over one quadrillion determinants. As shown in Chapter 5, the CI results computed using categorical compression within STP-DAS yield better results than other state-of-the-art quantum chemistry methods. Moreover, equipped with the ability to compute exact error bounds on the CI energy by means of the Ritz residual, the categorically-compressed STP-DAS calculations can reveal important trends in the behavior of other quantum chemistry methods. This is showcased in Section 5.2 by the proof that the CCSDT energy of Xe_2 falls below the true FCI value of the same, thereby demonstrating the non-variational nature of the coupled cluster method.

As computing power emerges as a new bottleneck in quantum chemistry calculations, the adoption of heterogeneous computing paradigms presents ample opportunity to explore previously unthinkable frontiers. Such platforms are capable of accelerating data processing by asynchronously offloading workloads to designated accelerator chips such as GPUs and application-specific integrated circuits (ASICs). These accelerators specialize in particular data-processing patterns, such as vectorization, making the workload-sharing between them and the hosting CPU ideal for tasks consisting of heterogeneous processing patterns, such as iterative diagonalization. Therefore, the rapid development of new data centers and

acceleration technologies due to the artificial intelligence economic boom also presents many promises for the future of quantum chemistry and other sciences.

BIBLIOGRAPHY

- [1] M. Alkan, B. Q. Pham, D. Del Angel Cruz, J. R. Hammond, T. A. Barnes, and M. S. Gordon. LibERI—A Portable and Performant Multi-GPU Accelerated Library for Electron Repulsion Integrals via OpenMP Offloading and Standard Language Parallelism. *J. Chem. Phys.*, 161(8):082501, 08 2024.
- [2] A. Almoukhalalati, S. Knecht, H. J. A. Jensen, K. G. Dyall, and T. Saue. Electron Correlation within the Relativistic No-pair Approximation. *J. Chem. Phys.*, 145(7):074104, 2016.
- [3] C. D. Anderson. The Positive Electron. *Phys. Rev.*, 43(6):491–494, 1933.
- [4] J. S. Anderson, F. Heidar-Zadeh, and P. W. Ayers. Breaking the Curse of Dimension for the Electronic Schrödinger Equation with Functional Analysis. *Comput. Theor. Chem*, 1142:66–77, 2018.
- [5] K. Andersson, P. Malmqvist, and B. O. Roos. Second-Order Perturbation Theory with a Complete Active Space Self-Consistent Field Reference Function. *J. Chem. Phys.*, 96(2):1218–1226, 01 1992.
- [6] K. Andersson, P. A. Malmqvist, B. O. Roos, A. J. Sadlej, and K. Wolinski. Second-Order Perturbation Theory with a CASSCF Reference Function. *J. Chem. Phys.*, 94(14):5483–5488, 1990.
- [7] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J. P. Malrieu. Introduction of n-Electron Valence States for Multireference Perturbation Theory. *J. Chem. Phys.*, 114:10252–10264, 2001.
- [8] R. Ansaloni, G. L. Bendazzoli, S. Evangelisti, and E. Rossi. A Parallel Full-CI Algorithm. *Comp. Phys. Comm.*, 128(1):496–515, 2000.
- [9] U. Baek, D. Hait, J. Shee, O. Leimkuhler, W. J. Huggins, T. F. Stetina, M. Head-Gordon, and K. B. Whaley. Say NO to Optimization: A Nonorthogonal Quantum Eigensolver. *Phys. Rev. X Quantum*, 4(3):030307, July 2023.
- [10] A. Baiardi and M. Reiher. The Density Matrix Renormalization Group in Chemistry and Molecular Physics: Recent Developments and New Challenges. *J. Chem. Phys.*, 152(4):040903, Jan. 2020.

- [11] H. F. Baker. Alternants and Continuous Groups. *Proc. London Math. Soc.*, s2-3(1):24–47, 1905.
- [12] S. Banerjee, R. R. Li, B. C. Cooper, T. Zhang, E. F. Valeev, X. Li, and I. DePrince, A. Eugene. Relativistic Core-Valence-Separated Molecular Mean-Field Exact-Two-Component Equation-of-Motion Coupled Cluster Theory: Applications to L-edge X-ray Absorption Spectroscopy. *APL Comput. Phys.*, page (In Press), 2026.
- [13] G. M. J. Barca, M. Alkan, J. L. Galvez-Vallejo, D. L. Poole, A. P. Rendell, and M. S. Gordon. Faster Self-Consistent Field (SCF) Calculations on GPU Clusters. *J. Chem. Theory Comput.*, 17(12):7486–7503, 2021.
- [14] R. J. Bartlett. Coupled-Cluster Theory and Its Equation-of-Motion Extensions. *WIREs Comput. Mol. Sci.*, 2(1):126–138, 2012.
- [15] R. J. Bartlett and M. Musiał. Coupled-Cluster Theory in Quantum Chemistry. *Rev. Mod. Phys.*, 79:291–352, 2007.
- [16] R. J. Bartlett, J. Watts, S. Kucharski, and J. Noga. Non-Iterative Fifth-Order Triple and Quadruple Excitation Energy Corrections in Correlated Methods. *Chem. Phys. Lett.*, 165(6):513–522, 1990.
- [17] J. E. Bates and T. Shiozaki. Fully Relativistic Complete Active Space Self-Consistent Field for Large Molecules: Quasi-Second-Order Minimax Optimization. *J. Chem. Phys.*, 142:044112, 2015.
- [18] S. Battaglia, S. Keller, and S. Knecht. Efficient Relativistic Density-Matrix Renormalization Group Implementation in a Matrix-Product Formulation. *J. Chem. Theory Comput.*, 14(5):2353–2369, 2018.
- [19] A. D. Becke. Density-Functional Exchange-Energy Approximation with Correct Asymptotic Behavior. *Phys. Rev. A*, 38(6):3098–3100, 1988.
- [20] A. D. Becke. Density-Functional Thermochemistry. III. The Role of Exact Exchange. *J. Chem. Phys.*, 98(7):5648–5652, 1993.
- [21] C. F. Bender and E. R. Davidson. Studies in Configuration Interaction: The First-Row Diatomic Hydrides. *Phys. Rev.*, 183:23–30, Jul 1969.
- [22] P. Beran, M. Matoušek, M. Hapka, K. Pernal, and L. Veis. Density Matrix Renormalization Group with Dynamical Correlation via Adiabatic Connection. *J. Chem. Theory Comput.*, 17(12):7575–7585, Dec. 2021.

- [23] X. Blase, C. Attaccalite, and V. Olevano. First-Principles GW Calculations for Fullerenes, Porphyrins, Phtalocyanine, and Other Molecules of Interest for Organic Photovoltaic Applications. *Phys. Rev. B*, 83:115103, Mar 2011.
- [24] N. Bohr. On the Constitution of Atoms and Molecules. *Philos. Mag.*, 26(151):1–25, 1913.
- [25] 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. *J. Chem. Phys.*, 131(5):054106, 2009.
- [26] M. Born and J. R. Oppenheimer. Zur Quantentheorie der Molekeln. *Ann. Phys.*, 389(20):457–484, 1927.
- [27] S. F. Boys. Electronic Wave Functions. I. A General Method of Calculation for the Stationary States of Any Molecular System. *Proc. A*, 200(1063):542–554, 1950.
- [28] G. Breit. The Effect of Retardation on the Interaction of Two Electrons. *Phys. Rev.*, 34(4):553–573, 1929.
- [29] F. Bruneval, T. Rangel, S. M. Hamed, Y. Shao, C. Yang, and J. B. Neaton. molgw 1: Many-Body Perturbation Theory Software for Atoms, Molecules, and Clusters. *Comp. Phys. Comm.*, 208:149–161, 2016.
- [30] J. V. Burda, R. Zahradník, and P. Hobza. Dimers of Rare Gas Atoms: CCSD(T), CCSDT and FCI Calculations on the (He)₂ Dimer, CCSD(T) and CCSDT Calculations on the (Ne)₂ Dimer, and CCSD(T) All-Electron and Pseudopotential Calculations on the Dimers from (Ne)₂ through (Xe)₂. *Mol. Phys.*, 89(2):425–432, 1996.
- [31] L. Bytautas and K. Ruedenberg. A Priori Identification of Configurational Deadwood. *Chem. Phys.*, 356(1):64–75, 2009.
- [32] J. E. Campbell. On a Law of Combination of Operators Bearing on the Theory of Continuous Transformation Groups. *Proc. London Math. Soc.*, s1-28(1):381–390, 1896.
- [33] D. S. L. Cardwell. *From Watt to Clausius: The Rise of Thermodynamics in the Early Industrial Age*. Cornell University Press., 1971.
- [34] S. Carnot. *Réflexions sur la Puissance Motrice du Feu et sur les Machines Propres à Développer cette Puissance*. Bachelier, Paris, 1824.
- [35] P. Celani and H.-J. Werner. Multireference Perturbation Theory for Large Restricted and Selected Active Space Reference Wave Functions. *J. Chem. Phys.*, 112(13):5546–5557, 04 2000.

- [36] D. Ceperley, G. V. Chester, and M. H. Kalos. Monte Carlo Simulation of a Many-Fermion Study. *Phys. Rev. B*, 16(7):3081–3099, 1977.
- [37] D. M. Ceperley and B. J. Alder. Ground State of the Electron Gas by a Stochastic Method. *Phys. Rev. Lett.*, 45(7):566–569, 1980.
- [38] G. K.-L. Chan and S. Sharma. The Density Matrix Renormalization Group in Quantum Chemistry. *Annu. Rev. Phys. Chem.*, 62(Volume 62, 2011):465–481, May 2011.
- [39] Y. Cheng, Z. Xie, and H. Ma. Post-Density Matrix Renormalization Group Methods for Describing Dynamic Electron Correlation with Large Active Spaces. *J. Phys. Chem. Lett.*, 13(3):904–915, Jan. 2022.
- [40] R. Clausius. Über die Bewegende Kraft der Wärme. *Ann. Phys.*, 155(3):368–397, 1850.
- [41] W. K. Clifford. Applications of Grassmann’s Extensive Algebra. *Am. J. Math.*, 1(4):350–358, 1878.
- [42] E. Condon. The Theory of Complex Spectra. *Phys. Rev.*, 36(7):1121–1133, 1930.
- [43] D. Contant, M. Casula, and M. Hellgren. Assessing Many-Body Methods on the Potential Energy Surface of the $(\text{H}_2)_2$ Hydrogen Dimer. *J. Chem. Phys.*, 161(18):074106, Nov. 2024.
- [44] B. C. Cooper, L. N. Koulias, D. R. Nascimento, X. Li, and A. E. DePrince. Short Iterative Lanczos Integration in Time-Dependent Equation-of-Motion Coupled-Cluster Theory. *J. Phys. Chem. A*, 125(24):5438–5447, 2021.
- [45] P. Corboz and G. Vidal. Fermionic Multiscale Entanglement Renormalization Ansatz. *Phys. Rev. B*, 80(16):165129, Oct. 2009.
- [46] S. J. Cotton. A Truncated Davidson Method for the Efficient “Chemically Accurate” Calculation of Full Configuration Interaction Wavefunctions without Any Large Matrix Diagonalization. *J. Chem. Phys.*, 157(22):224105, 2022.
- [47] J. M. Cullen and M. C. Zerner. The Linked Singles and Doubles Model: An Approximate Theory of Electron Correlation Based on the Coupled-Cluster Ansatz. *J. Chem. Phys.*, 77(8):4088–4109, 10 1982.
- [48] C. G. Darwin. The Wave Equations of the Electron. *Proc. A*, 118(780):654–680, 04 1928.

- [49] E. R. Davidson. The Iterative Calculation of a Few of the Lowest Eigenvalues and Corresponding Eigenvectors of Large Real-Symmetric Matrices. *J. Chem. Phys.*, 17:87–94, 1975.
- [50] C. Davis and W. Kahan. The Rotation of Eigenvectors by a Perturbation. III. *SIAM J. Numer. Anal.*, 7:1–46, 1970.
- [51] L. de Broglie. Ondes et Quanta. *C. R. Acad. Sci.*, 177:507–510, 1923.
- [52] C.-A. de Coulomb. Premier Mémoire sur l'Électricité et le Magnétisme. *Hist. Acad. R. Sci.*, 1785:569–577, 1785.
- [53] W. A. de Jong, R. J. Harrison, and D. A. Dixon. Parallel Douglas–Kroll Energy and Gradients in NWChem: Estimating Scalar Relativistic Effects Using Douglas–Kroll Contracted Basis Set. *J. Chem. Phys.*, 114:48–53, 2001.
- [54] P. Debye. Zur Theorie der spezifischen Wärme. *Ann. Phys.*, 344(14):789–839, 1912.
- [55] F. Ding, C. T. Chapman, W. Liang, and X. Li. Mechanisms of Bridge-Mediated Electron Transfer: A TDDFT Electronic Dynamics Study. *J. Chem. Phys.*, 137:22A512, 2012.
- [56] F. Ding, E. B. Guidez, C. M. Aikens, and X. Li. Quantum Coherent Plasmon in Silver Nanowires: A Real-Time TDDFT Study. *J. Chem. Phys.*, 140(24):244705, 2014.
- [57] P. A. M. Dirac. The Quantum Theory of the Electron. *Proc. A*, 117(778):610–624, 1928.
- [58] P. A. M. Dirac. A Theory of Electrons and Protons. *Proc. A*, 126(801):360–365, 1930.
- [59] A. J. Dobbyn, P. J. Knowles, and R. J. Harrison. Parallel Internally Contracted Multireference Configuration Interaction. *J. Comput. Chem.*, 19(11):1215–1228, 1998.
- [60] M. Douglas and N. M. Kroll. Quantum Electrodynamical Corrections to the Fine Structure of Helium. *Ann. Phys.*, 82:89–155, 1974.
- [61] W. Duch. *GRMS or Graphical Representation of Model Spaces: Vol. 1 Basics*, volume 1. Springer Science & Business Media, Berlin, Germany, 2012.
- [62] P. L. Dulong and A. T. Petit. Recherches sur Quelques Points Importants de la Théorie de la chaleur. *Ann. Chim. Phys.*, 10:395–413, 1819.
- [63] 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.

- [64] K. G. Dyall. Interfacing Relativistic and Nonrelativistic Methods. I. Normalized Elimination of the Small Component in the Modified Dirac Equation. *J. Chem. Phys.*, 106(23):9618–9626, 1997.
- [65] K. G. Dyall. Interfacing Relativistic and Nonrelativistic Methods. II. Investigation of a Low-Order Approximation. *J. Chem. Phys.*, 109:4201–4208, 1998.
- [66] K. G. Dyall. Interfacing Relativistic and Nonrelativistic Methods. IV. One- and Two-Electron Scalar Approximations. *J. Chem. Phys.*, 115(20):9136–9143, 11 2001.
- [67] K. G. Dyall and T. Enevoldsen. Interfacing Relativistic and Nonrelativistic Methods. III. Atomic 4-Spinor Expansions and Integral Approximations. *J. Chem. Phys.*, 111(22):10000–10007, 12 1999.
- [68] K. G. Dyall and K. Fægri, Jr. *Introduction to Relativistic Quantum Chemistry*. Oxford University Press, 2007.
- [69] F. Egidi, J. J. Goings, M. J. Frisch, and X. Li. Direct Atomic-Orbital-Based Relativistic Two-Component Linear Response Method for Calculating Excited-State Fine Structures. *J. Chem. Theory Comput.*, 12(8):3711–3718, 2016.
- [70] F. Egidi, S. Sun, J. J. Goings, G. Scalmani, M. J. Frisch, and X. Li. Two-Component Non-Collinear Time-Dependent Spin Density Functional Theory for Excited State Calculations. *J. Chem. Theory Comput.*, 13(6):2591–2603, 2017.
- [71] F. Egidi, D. B. Williams-Young, A. Baiardi, J. Bloino, G. Scalmani, M. J. Frisch, X. Li, and V. Barone. Effective Inclusion of Mechanical and Electrical Anharmonicity in Excited Electronic States: the VPT2-TDDFT Route. *J. Chem. Theory Comput.*, 13(6):2789–2803, 2017.
- [72] J. Ehrman, E. Martinez-Baez, A. J. Jenkins, and X. Li. Improving One-Electron Exact-Two-Component Relativistic Methods with the Dirac–Coulomb–Breit-Parameterized Effective Spin–Orbit Coupling. *J. Chem. Theory Comput.*, 19:5785–5790, 2023.
- [73] A. Einstein. On a Heuristic Viewpoint Concerning the Production and Transformation of Light. *Ann. Phys.*, 17(6):132–148, 1905.
- [74] A. Einstein. Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Ann. Phys.*, 322(8):549–560, 1905.
- [75] A. Einstein. Zur Elektrodynamik bewegter Körper. *Ann. Phys.*, 322(10):891–921, 1905.

- [76] J. J. Eriksen. The Shape of Full Configuration Interaction to Come. *J. Phys. Chem. Lett.*, 12(1):418–432, Jan. 2021.
- [77] K. Esler, J. Kim, D. Ceperley, and L. Shulenburger. Accelerating Quantum Monte Carlo Simulations of Real Materials on GPU Clusters. *Comput. Sci. Eng.*, 14(1):40–51, 2012.
- [78] L. Euler. Principes Généraux du Mouvement des Fluides. *Mém. Acad. Sci. Berlin*, 11:274–315, 1757.
- [79] F. A. Evangelista. Perspective: Multireference Coupled Cluster Theories of Dynamical Electron Correlation. *J. Chem. Phys.*, 149(3):030901, 07 2018.
- [80] S. Evangelisti, J.-P. Daudey, and J.-P. Malrieu. Convergence of an Improved CIPSI Algorithm. *Chem. Phys.*, 75(1):91–102, 1983.
- [81] K. Fægri Jr and L. Visscher. Relativistic Calculations on Thallium Hydride. *Theor. Chem. Acc.*, 105:265–267, 2001.
- [82] B. S. Fales, E. R. Curtis, K. G. Johnson, D. Lahana, S. Seritan, Y. Wang, H. Weir, T. J. Martínez, and E. G. Hohenstein. Performance of Coupled-Cluster Singles and Doubles on Modern Stream Processing Architectures. *J. Chem. Theory Comput.*, 16(7):4021–4028, 2020.
- [83] B. S. Fales and B. G. Levine. Nanoscale Multireference Quantum Chemistry: Full Configuration Interaction on Graphical Processing Units. *J. Chem. Theory Comput.*, 11(10):4708–4716, 2015.
- [84] B. S. Fales and T. J. Martínez. Efficient Treatment of Large Active Spaces through Multi-GPU Parallel Implementation of Direct Configuration Interaction. *J. Chem. Theory Comput.*, 16(3):1586–1596, 2020.
- [85] T. Fleig, H. J. A. Jensen, J. Olsen, and L. Visscher. The Generalized Active Space Concept for the Relativistic Treatment of Electron Correlation. III. Large-Scale Configuration Interaction and Multiconfiguration Self-Consistent-Field Four-Component Methods with Application to UO₂. *J. Chem. Phys.*, 124:104106, 2006.
- [86] T. Fleig, J. Olsen, and C. M. Marian. The Generalized Active Space Concept for the Relativistic Treatment of Electron Correlation. I. Kramers-Restricted Two-Component Configuration Interaction. *J. Chem. Phys.*, 114:4775–4790, 2001.
- [87] T. Fleig, J. Olsen, and L. Visscher. The Generalized Active Space Concept for the Relativistic Treatment of Electron Correlation. II. Large-Scale Configuration Interaction Implementation Based on Relativistic 2- and 4-Spinors and Its Application. *J. Chem. Phys.*, 119:2963–2971, 2003.

- [88] L. L. Foldy and S. A. Wouthuysen. On the Dirac Theory of Spin 1/2 Particles and Its Non-Relativistic Limit. *Phys. Rev.*, 78(1):29–36, 1950.
- [89] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum Monte Carlo Simulations of Solids. *Rev. Mod. Phys.*, 73(1):33–83, 2001.
- [90] J. T. Frey and D. D. J. TubeGen 3.4. <http://turin.nss.udel.edu/research/tubegenonline.html> (accessed 2025-04-03), 2011.
- [91] G. L. G. Sleijpen and H. A. Van der Vorst. A Jacobi–Davidson Iteration Method for Linear Eigenvalue Problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.
- [92] G. Galilei. *Dialogo Sopra i Due Massimi Sistemi del Mondo*. Battista Landini, Florence, 1632.
- [93] Z. Gan and R. Harrison. Calibrating Quantum Chemistry: A Multi-Teraflop, Parallel-Vector, Full-Configuration Interaction Program for the Cray-X1. In *SC '05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, pages 22–22, 2005.
- [94] D. Ganyushin and F. Neese. A Fully Variational Spin-Orbit Coupled Complete Active Space Self-Consistent Field Approach: Application to Electron Paramagnetic Resonance g-Tensors. *J. Chem. Phys.*, 138:104113, 2013.
- [95] H. Gao, S. Imamura, A. Kasagi, and E. Yoshida. Distributed Implementation of Full Configuration Interaction for One Trillion Determinants. *J. Chem. Theory Comput.*, 20(3):1185–1192, 2024.
- [96] K. Gasperich, M. Deible, and K. D. Jordan. H4: A Model System for Assessing the Performance of Diffusion Monte Carlo Calculations Using a Single Slater Determinant Trial Function. *J. Chem. Phys.*, 147(7):074106, 08 2017.
- [97] C. Genovese, A. Meninno, and S. Sorella. Assessing the Accuracy of the Jastrow Antisymmetrized Geminal Power in the H₄ Model System. *J. Chem. Phys.*, 150(8):084102, Feb. 2019.
- [98] J. W. Gibbs. On the Equilibrium of Heterogeneous Substances. *Am. J. Sci.*, s3-16(96):441–458, Dec. 1878.
- [99] J. J. Goings, H. Hu, C. Yang, and X. Li. Reinforcement Learning Configuration Interaction. *J. Chem. Theory Comput.*, 17(9):5482–5491, 2021.
- [100] J. J. Goings, J. M. Kasper, F. Egidi, S. Sun, and X. Li. Real Time Propagation of the Exact Two Component Time-Dependent Density Functional Theory. *J. Chem. Phys.*, 145(10):104107, 2016.

- [101] E. Goll, H.-J. Werner, and H. Stoll. A Short-Range Gradient-Corrected Density Functional in Long-Range Coupled-Cluster Calculations for Rare Gas Dimers. *Phys. Chem. Chem. Phys.*, 7:3917–3923, 2005.
- [102] S. M. Greene, R. J. Webber, J. Weare, and T. C. Berkelbach. Beyond Walkers in Stochastic Quantum Chemistry: Reducing Error Using Fast Randomized Iteration. *J. Chem. Theory Comput.*, 15(9):4834–4850, 2019.
- [103] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg. A Consistent and Accurate *Ab Initio* Parametrization of Density Functional Dispersion Correction (DFT-D) for the 94 Elements H-Pu. *J. Chem. Phys.*, 132(15):154104, 2010.
- [104] S. Guo, M. A. Watson, W. Hu, Q. Sun, and G. K.-L. Chan. N-Electron Valence State Perturbation Theory Based on a Density Matrix Renormalization Group Reference Function, with Applications to the Chromium Dimer and a Trimer Model of Poly(p-Phenylenevinylene). *J. Chem. Theory Comput.*, 12(4):1583–1591, Apr. 2016.
- [105] S. Hammes-Schiffer. Nuclear–Electronic Orbital Methods: Foundations and Prospects. *The Journal of Chemical Physics*, 155(3):030901, 07 2021.
- [106] F. Hausdorff. Die Symbolische Exponentialformel in der Gruppentheorie. *Ber. Verh. Saechs. Akad. Wiss. Leipzig*, 58:19–48, 1906.
- [107] M. Head-Gordon and J. A. Pople. A Method for Two-Electron Gaussian Integral and Integral Derivative Evaluation Using Recurrence Relations. *J. Chem. Phys.*, 89(9):5777–5786, 1988.
- [108] E. D. Hedegård, S. Knecht, J. S. Kielberg, H. J. A. Jensen, and M. Reiher. Density Matrix Renormalization Group with Efficient Dynamical Electron Correlation through Range Separation. *J. Chem. Phys.*, 142(22):224108, June 2015.
- [109] L. Hedin. New Method for Calculating the One-Particle Green’s Function with Application to the Electron-Gas Problem. *Phys. Rev.*, 139(3A):A796–A823, 1965.
- [110] M. I. Hernández and D. C. Clary. Four-Center Reactions: A Quantal Model for H₄. *J. Chem. Phys.*, 104(21):8413–8423, June 1996.
- [111] T. M. Hernandez, R. Van Beeumen, M. A. Caprio, and C. Yang. A Greedy Algorithm for Computing Eigenvalues of a Symmetric Matrix with Localized Eigenvectors. *Numer. Linear Algebra Appl.*, 28(2):e2341, 2021.
- [112] H. Hertz. Ueber einen Einfluss des ultravioletten Lichtes auf die electrische Entladung. *Ann. Phys.*, 267(8):983–1000, 1887.

- [113] B. A. Hess. Relativistic Electronic-Structure Calculations Employing a Two-Component No-Pair Formalism with External-Field Projection Operators. *Phys. Rev. A*, 33:3742–3748, 1986.
- [114] G. J. Hill and K. A. Peterson. Gaussian Basis Sets for Use in Correlated Molecular Calculations. XI. Pseudopotential-Based and All-Electron Relativistic Basis Sets for Alkali Metal (K–Fr) and Alkaline Earth (Ca–Ra) Elements. *J. Chem. Phys.*, 147(24):244106, 12 2017.
- [115] K. Hirao, editor. *Recent Advances in Multireference Methods*. World Scientific, 1999.
- [116] P. Hohenberg and W. Kohn. Inhomogeneous Electron Gas. *Phys. Rev.*, 136(3B):B864–B871, 1964.
- [117] 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. *J. Chem. Theory Comput.*, 12(8):3674–3680, Aug. 2016.
- [118] A. A. Holmes, C. J. Umrigar, and S. Sharma. Excited States Using Semistochastic Heat-Bath Configuration Interaction. *J. Chem. Phys.*, 147(16):164111, Oct. 2017.
- [119] C. E. Hoyer, H. Hu, L. Lu, S. Knecht, and X. Li. Relativistic Kramers-Unrestricted Exact-Two-Component Density Matrix Renormalization Group. *J. Phys. Chem. A*, 126(30):5011–5020, 2022.
- [120] C. E. Hoyer, L. Lu, H. Hu, K. D. Shumilov, S. Sun, S. Knecht, and X. Li. Correlated Dirac–Coulomb–Breit Multiconfigurational Self-Consistent-Field Methods. *J. Chem. Phys.*, 158(4):044101, 2023.
- [121] H. Hu, A. J. Jenkins, H. Liu, J. M. Kasper, M. J. Frisch, and X. Li. Relativistic Two-Component Multireference Configuration Interaction Method with Tunable Correlation Space. *J. Chem. Theory Comput.*, 16(5):2975–2984, 2020.
- [122] H. Hu, S. Upadhyay, L. Lu, A. J. Jenkins, T. Zhang, A. Shayit, S. Knecht, and X. Li. Small Tensor Product Distributed Active Space (STP-DAS) Framework for Relativistic and Non-relativistic Multiconfiguration Calculations: Scaling from 10^9 on a Laptop to 10^{12} Determinants on a Supercomputer. *Chem. Phys. Rev.*, 5(4):041404, 10 2024.
- [123] T. P. Hughes. *Networks of Power: Electrification in Western Society, 1880-1930*. Johns Hopkins University Press, Baltimore, 1993.
- [124] B. Huron, J. P. Malrieu, and P. Rancurel. Iterative Perturbation Calculations of Ground and Excited State Energies from Multiconfigurational Zeroth-Order Wavefunctions. *J. Chem. Phys.*, 58(12):5745–5759, June 1973.

- [125] M. S. Hybertsen and S. G. Louie. Electron Correlation in Semiconductors and Insulators: Band Gaps and Quasiparticle Energies. *Phys. Rev. B*, 34(8):5390–5413, 1986.
- [126] M. Ilias and T. Saue. An Infinite-Order Relativistic Hamiltonian by a Simple One-Step Transformation. *J. Chem. Phys.*, 126:064102, 2007.
- [127] C. M. Isborn, X. Li, and J. C. Tully. TDDFT Ehrenfest Dynamics: Collisions between Atomic Oxygen and Graphite Clusters. *J. Chem. Phys.*, 126:134307, 2007.
- [128] J. Ivanic. Direct Configuration Interaction and Multiconfigurational Self-Consistent-Field Method for Multiple Active Spaces with Variable Occupations. I. Method. *J. Chem. Phys.*, 119(18):9364–9376, 11 2003.
- [129] J. Ivanic. Direct Configuration Interaction and Multiconfigurational Self-Consistent-Field Method for Multiple Active Spaces with Variable Occupations. II. Application to oxoMn(salen) and N₂O₄. *J. Chem. Phys.*, 119(18):9377–9385, 11 2003.
- [130] A. J. Jenkins, H. Hu, L. Lu, M. J. Frisch, and X. Li. Two-Component Multireference Restricted Active Space Configuration Interaction for the Computation of L-Edge X-ray Absorption Spectra. *J. Chem. Theory Comput.*, 18(1):141–150, 2022.
- [131] A. J. Jenkins, H. Liu, J. M. Kasper, M. J. Frisch, and X. Li. Variational Relativistic Complete Active Space Self-Consistent Field Method. *J. Chem. Theory Comput.*, 15:2974–2982, 2019.
- [132] H. Jørgen Aa. Jensen, K. G. Dyall, T. Saue, and K. Fægri. Relativistic Four-Component Multiconfigurational Self-Consistent-Field Theory for Molecules: Formalism. *J. Chem. Phys.*, 104:4083–4097, 1996.
- [133] F. O. Kannemann and A. D. Becke. Van der Waals Interactions in Density-Functional Theory: Rare-Gas Diatomics. *J. Chem. Theory Comput.*, 5(4):719–727, 2009.
- [134] J. A. Karwowski and I. Shavitt. Configuration Interaction. In S. Wilson, editor, *Handbook of Molecular Physics and Quantum Chemistry*. John Wiley & Sons, Ltd, 2003.
- [135] J. M. Kasper, D. B. Williams-Young, E. Vecharynski, C. Yang, and X. Li. A Well-Tempered Hybrid Method for Solving Challenging Time-Dependent Density Functional Theory (TDDFT) Systems. *J. Chem. Theory Comput.*, 14(4):2034–2041, 2018.
- [136] T. Kato. On the Eigenfunctions of Many-Particle Systems in Quantum Mechanics. *Commun. Pure Appl. Math.*, 10(2):151–177, 1957.

- [137] S. Kedzuch, O. Demel, J. Pittner, S. Ten-no, and J. Noga. Multireference F12 Coupled Cluster Theory: The Brillouin-Wigner Approach with Single and Double Excitations. *Chem. Phys. Lett.*, 511(4-6):418–423, Aug. 2011.
- [138] R. A. Kendall, T. H. Dunning, and R. J. Harrison. Electron Affinities of the First-Row Atoms Revisited. Systematic Basis Sets and Wave Functions. *J. Chem. Phys.*, 96:6796–6806, 1992.
- [139] J. Kepler. *Astronomia Nova*. Haeredes Tychonis Brahe, Praga, 1609.
- [140] J. Kepler. *Tabulae Rudolphinae, Quibus Astronomicae Scientiae, Temporum Longinquitate Collapsae Restauratio Continentur*. Typis J. Saurii, Ulmae (Ulm), 1627.
- [141] G. I. Kerley. On Corrections to the Born-Oppenheimer Approximation. arXiv:1306.6574, 2013.
- [142] G. Kirchhoff. Über das verhältniß zwischen dem emissionsvermögen und dem absorptionsvermögen der körper für wärme und licht. *Ann. Phys.*, 185(2):275–301, 1860.
- [143] O. Klein. Quantentheorie und fünfdimensionale relativitätstheorie. *Z. Phys.*, 37(12):895–906, 1926.
- [144] M. Klene, M. A. Robb, L. Blancafort, and M. J. Frisch. A New Efficient Approach to the Direct Restricted Active Space Self-Consistent Field Method. *J. Chem. Phys.*, 119:713–728, 2003.
- [145] M. Klene, M. A. Robb, M. J. Frisch, and P. Celani. Parallel Implementation of the CI-Vector Evaluation in Full CI/CAS-SCF. *J. Chem. Phys.*, 113:5653–5665, 2000.
- [146] S. Knecht, H. J. A. Jensen, and T. Fleig. Large-Scale Parallel Configuration Interaction. I. Nonrelativistic and Scalar-Relativistic General Active Space Implementation with Application to (Rb–Ba)⁺. *J. Chem. Phys.*, 128:014108, 2008.
- [147] S. Knecht, H. J. A. Jensen, and T. Fleig. Large-Scale Parallel Configuration Interaction. II. Two- and Four-Component Double-Group General Active Space Implementation with Application to BiH. *J. Chem. Phys.*, 132:014108, 2010.
- [148] S. Knecht, H. J. A. Jensen, and T. Saue. Relativistic Quantum Chemical Calculations Show that the Uranium Molecule U₂ has a Quadruple Bond. *Nat. Chem.*, 11(1):40–44, 2019.
- [149] S. Knecht, O. Legeza, and M. Reiher. Communication: Four-Component Density Matrix Renormalization Group. *J. Chem. Phys.*, 140(4):041101, 2014.

- [150] P. Knowles and N. Handy. A New Determinant-Based Full Configuration Interaction Method. *Chem. Phys. Lett.*, 111:315–321, 1984.
- [151] P. J. Knowles. Very Large Full Configuration Interaction Calculations. *Chem. Phys. Lett.*, 155(6):513–517, 1989.
- [152] P. J. Knowles and N. C. Handy. A Determinant Based Full Configuration Interaction Program. *Comp. Phys. Comm.*, 54(1):75–83, Apr. 1989.
- [153] A. V. Knyazev. Toward the Optimal Preconditioned Eigensolver: Locally Optimal Block Preconditioned Conjugate Gradient Method. *SIAM J. Sci. Comp.*, 23(2):517–541, 2001.
- [154] W. Kohn and L. J. Sham. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.*, 140(4A):A1133–A1138, 1965.
- [155] L. Konecny, M. Kadek, S. Komorovsky, O. L. Malkina, K. Ruud, and M. Repisky. Acceleration of Relativistic Electron Dynamics by Means of X2C Transformation: Application to the Calculation of Nonlinear Optical Properties. *J. Chem. Theory Comput.*, 12(12):5823–5833, Dec 2016.
- [156] L. N. Koulias, D. B. Williams-Young, D. R. Nascimento, A. E. DePrince, and X. Li. Relativistic Time-Dependent Equation-of-Motion Coupled-Cluster. *J. Chem. Theory Comput.*, 15:6617–6624, 2019.
- [157] P. M. Kozłowski and P. Pulay. The Unrestricted Natural Orbital-Restricted Active Space Method: Methodology and Implementation. *Theor. Chem. Acc.*, 100:12–20, 1998.
- [158] R. Krishnan, J. S. Binkley, R. Seeger, and J. A. Pople. Self-Consistent Molecular Orbital Methods. XX. A Basis Set for Correlated Wave Functions. *J. Chem. Phys.*, 72:650–654, 1980.
- [159] A. N. Krylov. On the Numerical Solution of Equations Whose Solution Determines the Frequency of Small Vibrations of Material Systems (in Russian). *Izv. Akad. Nauk. SSSR*, 1:491–539, 1931.
- [160] O. Kullie and T. Saue. Range-Separated Density Functional Theory: A 4-Component Relativistic Study of the Rare Gas Dimers He₂, Ne₂, Ar₂, Kr₂, Xe₂, Rn₂ and Uuo₂. *Chem. Phys.*, 395:54–62, 2012.
- [161] Y. Kurashige and T. Yanai. Second-Order Perturbation Theory with a Density Matrix Renormalization Group Self-Consistent Field Reference Function: Theory and application to the study of chromium dimer. *J. Chem. Phys.*, 135(9):094104, Sept. 2011.

- [162] J. Kussmann and C. Ochsenfeld. Hybrid CPU/GPU Integral Engine for Strong-Scaling *Ab Initio* Methods. *J. Chem. Theory Comput.*, 13(7):3153–3159, 2017.
- [163] W. Kutzlenigg and W. Liu. Quasirelativistic Theory Equivalent to Fully Relativistic Theory. *J. Chem. Phys.*, 123:241102, 2005.
- [164] J. Larmor. On the Theory of the Magnetic Influence on Spectra; and On the Radiation from Moving Ions. *Philos. Mag.*, 44(271):503–512, 1897.
- [165] E. F. Laschuk, M. M. Martins, and S. Evangelisti. *Ab Initio* Potentials for Weakly Interacting Systems: Homonuclear Rare Gas Dimers. *Int. J. Quant. Chem.*, 95(3):303–312, 2003.
- [166] C. Lee, W. Yang, and R. G. Parr. Development of the Colle-Salvetti Correlation-Energy Formula into a Functional of the Electron Density. *Phys. Rev. B*, 37(2):785–789, 1988.
- [167] P. J. Lestrange, P. D. Nguyen, and X. Li. Calibration of Energy-Specific TDDFT for Modeling K-Edge XAS Spectra of Light Elements. *J. Chem. Theory Comput.*, 11:2994–2999, 2015.
- [168] J. Li, M. Otten, A. A. Holmes, S. Sharma, and C. J. Umrigar. Fast Semistochastic Heat-Bath Configuration Interaction. *J. Chem. Phys.*, 149(21):214110, Dec. 2018.
- [169] R. Li, B. Hanindhito, S. Yadav, Q. Wu, K. Kavi, G. Mehta, N. J. Yadwadkar, and L. K. John. Performance Implications of Pipelining the Data Transfer in CPU-GPU Heterogeneous Systems. *ACM Trans. Archit. Code Optim.*, 22(3), Sept. 2025.
- [170] R. R. Li, S. H. Yuwono, M. D. Liebenthal, T. Zhang, X. Li, and I. DePrince, A. Eugene. Relativistic Two-Component Double Ionization Potential Equation-of-Motion Coupled Cluster with the Dirac–Coulomb–Breit Hamiltonian. *J. Chem. Phys.*, 163(10):104112, 09 2025.
- [171] Z. Li, Y. Xiao, and W. Liu. On the Spin Separation of Algebraic Two-Component Relativistic Hamiltonians. *J. Chem. Phys.*, 137(15):154114, 10 2012.
- [172] G. Li Manni, R. K. Carlson, S. Luo, D. Ma, J. Olsen, D. G. Truhlar, and L. Gagliardi. Multiconfiguration Pair-Density Functional Theory. *J. Chem. Theory Comput.*, 10(9):3669–3680, 2014.
- [173] W. Liang, S. A. Fischer, M. J. Frisch, and X. Li. Energy-Specific Linear Response TDHF/TDDFT for Calculating High-Energy Excited States. *J. Chem. Theory Comput.*, 7(11):3540–3547, 2011.

- [174] H. Lischka, T. Müller, P. G. Szalay, I. Shavitt, R. M. Pitzer, and R. Shepard. Columbus—a Program System for Advanced Multireference Theory Calculations. *WIREs Comput. Mol. Sci.*, 1:191–199, 2011.
- [175] H. Lischka, R. Shepard, F. B. Brown, and I. Shavitt. New Implementation of the Graphical Unitary Group Approach for Multireference Direct Configuration Interaction Calculations. *Int. J. Quant. Chem.*, 20(S15):91–100, 1981.
- [176] B. Liu. *Ab Initio* Potential Energy Surface for Linear H₃. *J. Chem. Phys.*, 58(5):1925–1937, 03 1973.
- [177] J. Liu and L. Cheng. Relativistic Coupled-Cluster and Equation-of-Motion Coupled-Cluster Methods. *WIREs Comput. Mol. Sci.*, 11(6):1536, 2021.
- [178] W. Liu. Ideas of Relativistic Quantum Chemistry. *Mol. Phys.*, 108(13):1679–1706, 2010.
- [179] W. Liu and M. R. Hoffmann. iCI: Iterative CI Toward Full CI. *J. Chem. Theory Comput.*, 12(3):1169–1178, Mar. 2016.
- [180] W. Liu and D. Peng. Infinite-Order Quasirelativistic Density Functional Method Based on the Exact Matrix Quasirelativistic Theory. *J. Chem. Phys.*, 125:044102, 2006.
- [181] W. Liu and D. Peng. Exact Two-Component Hamiltonians Revisited. *J. Chem. Phys.*, 131(3):031104, 2009.
- [182] H. A. Lorentz. Electromagnetic Phenomena in a System Moving with Any Velocity Smaller than That of Light. *Proc. R. Neth. Acad. Arts Sci.*, 6:809–831, 1904.
- [183] L. Lu, H. Hu, A. J. Jenkins, and X. Li. Exact-Two-Component Relativistic Multireference Second-Order Perturbation Theory. *J. Chem. Theory Comput.*, 18(5):2983–2992, 2022.
- [184] D. I. Lyakh, M. Musiał, V. F. Lotrich, and R. J. Bartlett. Multireference Nature of Chemistry: The Coupled-Cluster View. *Chem. Rev.*, 112:182–243, 2012.
- [185] D. Ma, G. Li Manni, and L. Gagliardi. The Generalized Active Space Concept in Multiconfigurational Self-Consistent Field Methods. *J. Chem. Phys.*, 135:044128, 2011.
- [186] D. Ma, G. Li Manni, J. Olsen, and L. Gagliardi. Second-Order Perturbation Theory for Generalized Active Space Self-Consistent-Field Wave Functions. *J. Chem. Theory Comput.*, 12(7):3208–3213, 2016.

- [187] P. Å. Malmqvist. Calculation of Transition Density Matrices by Nonunitary Orbital Transformations. *Int. J. Quant. Chem.*, 30:479–494, 1986.
- [188] P.-Å. Malmqvist and B. O. Roos. The CASSCF State Interaction Method. *Chem. Phys. Lett.*, 155:189–194, 1989.
- [189] P.-Å. Malmqvist and B. O. Roos. The Restricted Active Space Self-Consistent-Field Method, Implemented with a Split Graph Unitary Group Approach. *J. Phys. Chem.*, 94:5477–5482, 1990.
- [190] P. Å. Malmqvist, B. O. Roos, and B. Schimmelpfennig. The Restricted Active Space (RAS) State Interaction Approach with Spin–Orbit Coupling. *Chem. Phys. Lett.*, 357:230–240, 2002.
- [191] F. D. Malone, S. Zhang, and M. A. Morales. Accelerating Auxiliary-Field Quantum Monte Carlo Simulations of Solids with Graphical Processing Units. *J. Chem. Theory Comput.*, 16(7):4286–4297, 2020.
- [192] J. C. Maxwell. A Dynamical Theory of the Electromagnetic Field. *Phil. Trans. R. Soc. Lond.*, 155:459–512, 1865.
- [193] W. L. McMillan. Ground State of Liquid He⁴. *Phys. Rev.*, 138(2A):A442–A451, 1965.
- [194] A. Menczer, M. van Damme, A. Rask, L. Huntington, J. Hammond, S. S. Xantheas, M. Ganahl, and O. Legeza. Parallel Implementation of the Density Matrix Renormalization Group Method Achieving a Quarter petaFLOPS Performance on a Single DGX-H100 GPU Node. *J. Chem. Theory Comput.*, 20(19):8397–8404, 2024.
- [195] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.*, 21(6):1087–1092, 1953.
- [196] A. A. Michelson and E. W. Morley. On the Relative Motion of the Earth and the Luminiferous Ether. *Am. J. Sci.*, s3–34(203):333–345, 1887.
- [197] A. O. Mitrushenkov. Passing the Several Billions Limit in FCI Calculations on a Mini-Computer. *Chem. Phys. Lett.*, 217(5–6):559–565, Jan. 1994.
- [198] J. Mokyr. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press, 1990.
- [199] V. Murg, F. Verstraete, Ö. Legeza, and R. M. Noack. Simulating Strongly Correlated Quantum Systems with Tree Tensor Networks. *Phys. Rev. B*, 82(20):205105, Nov. 2010.

- [200] T. Nakajima and K. Hirao. The Higher-Order Douglas-Kroll Transformation. *J. Chem. Phys.*, 113(18):7786–7789, 2000.
- [201] M. Nakano, T. Minami, H. Fukui, R. Kishi, Y. Shigeta, and B. Champagne. Full Configuration Interaction Calculations of the Second Hyperpolarizabilities of the H₄ Model Compound: Summation-over-States Analysis and Interplay with Diradical Characters. *J. Chem. Phys.*, 136(2):024315, Jan. 2012.
- [202] H. Nakashima and H. Nakatsuji. Solving the Schrödinger Equation of a Planar Model H₄ Molecule. *Chem. Phys. Lett.*, 815(140359):140359, Mar. 2023.
- [203] N. Nakatani and G. K.-L. Chan. Efficient Tree Tensor Network States (TTNS) for Quantum Chemistry: Generalizations of the Density Matrix Renormalization Group Algorithm. *J. Chem. Phys.*, 138(13):134113, Apr. 2013.
- [204] N. Nakatani and S. Guo. Density Matrix Renormalization Group (DMRG) Method as a Common Tool for Large Active-Space CASSCF/CASPT2 Calculations. *J. Chem. Phys.*, 146(9):094102, Mar. 2017.
- [205] E. Neuscamman, T. Yanai, and G. K.-L. Chan. Quadratic Canonical Transformation Theory and Higher Order Density Matrices. *J. Chem. Phys.*, 130(12):124102, Mar. 2009.
- [206] I. Newton. *Philosophiæ Naturalis Principia Mathematica*. Jussu Societatis Regiæ ac Typis Josephi Streater, Londini, 1687.
- [207] E. Noether. Invariante Variationsprobleme. *Nachr. Ges. Wiss. Göttingen, Math.-Phys. Kl.*, 1918:235–257, 1918.
- [208] J. Noga and R. J. Bartlett. The Full CCSDT Model for Molecular Electronic Structure. *J. Chem. Phys.*, 86(12):7041–7050, 06 1987.
- [209] L. J. Norrby. Why Is Mercury Liquid? Or, Why Do Relativistic Effects Not Get into Chemistry Textbooks? *J. Chem. Educ.*, 68(2):110, 1991.
- [210] S. Obara and A. Saika. Efficient Recurrence Relations for the Computation of Molecular Integrals over Gaussian Basis Functions. *J. Chem. Phys.*, 84(7):3963—3974, 1986.
- [211] J. Olsen, P. Jørgensen, and J. Simons. Passing the One-Billion Limit in Full Configuration-Interaction (FCI) Calculations. *Chem. Phys. Lett.*, 169(6):463–472, 1990.

- [212] J. Olsen, B. O. Roos, and P. Jørgensen. Determinant Based Configuration Interaction Algorithms for Complete and Restricted Configuration Interaction Spaces. *J. Chem. Phys.*, 89(4):2185–2192, 1988.
- [213] J. Olsen, B. O. Roos, P. Jørgensen, and H. J. A. Jensen. Determinant Based Configuration Interaction Algorithms for Complete and Restricted Configuration Interaction Spaces. *J. Chem. Phys.*, 89(4):2185–2192, Aug. 1988.
- [214] J. D. Olsen and K. T. McDonald. Classical Lifetime of a Bohr Atom. Princeton University Report, Princeton University, 2017.
- [215] G. Onida, L. Reining, and A. Rubio. Electronic Excitations: Density-Functional Versus Many-Body Green’s-Function Approaches. *Rev. Mod. Phys.*, 74:601–659, Jun 2002.
- [216] J. Paldus. Externally and Internally Corrected Coupled Cluster Approaches: An Overview. *J. Math. Chem.*, 55:477–502, 2017.
- [217] B. N. Parlett. *The Symmetric Eigenvalue Problem*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1998.
- [218] B. Peng, P. J. Lestrange, J. J. Goings, M. Caricato, and X. Li. Energy-Specific Equation-of-Motion Coupled-Cluster Methods for High-Energy Excited States: Application to K-Edge X-Ray Absorption Spectroscopy. *J. Chem. Theory Comput.*, 11(9):4146–4153, 2015.
- [219] B. Peng, D. B. Lingerfelt, F. Ding, C. M. Aikens, and X. Li. Real-Time TDDFT Studies of Exciton Decay and Transfer in Silver Nanowire Arrays. *J. Phys. Chem. C*, 119(11):6421–6427, 2015.
- [220] D. Peng and K. Hirao. An Arbitrary Order Douglas-Kroll Method with Polynomial Cost. *J. Chem. Phys.*, 130(4):044102, 2009.
- [221] D. Peng, W. Liu, Y. Xiao, and L. Cheng. Making Four- and Two-Component Relativistic Density Functional Methods Fully Equivalent Based on the Idea of From Atoms to Molecule. *J. Chem. Phys.*, 127:104106, 2007.
- [222] D. Peng, N. Middendorff, F. Weigend, and M. Reiher. An Efficient Implementation of Two-Component Relativistic Exact-Decoupling Methods for Large Molecules. *J. Chem. Phys.*, 138:184105, 2013.
- [223] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.*, 77(18):3865–3868, 1996.

- [224] K. A. Peterson, D. E. Woon, and J. Dunning, Thom H. Benchmark Calculations with Correlated Molecular Wave Functions. IV. The Classical Barrier Height of the $\text{H}+\text{H}_2\rightarrow\text{H}_2+\text{H}$ Reaction. *J. Chem. Phys.*, 100(10):7410–7415, 05 1994.
- [225] A. Petrone, D. B. Lingerfelt, N. Rega, and X. Li. From Charge-Transfer to a Charge-Separated State: A Perspective from the Real-Time TDDFT Excitonic Dynamics. *Phys. Chem. Chem. Phys.*, 16(44):24457–24465, 2014.
- [226] M. Planck. On the Law of Energy Distribution in the Normal Spectrum. *Ann. Phys.*, 4:553–563, 1901.
- [227] P. Pollak and F. Weigend. Segmented Contracted Error-Consistent Basis Sets of Double- and Triple- ζ Valence Quality for One- and Two-Component Relativistic All-Electron Calculations. *J. Chem. Theory Comput.*, 13(8):3696–3705, 2017.
- [228] J. V. Pototschnig, A. Papadopoulos, D. I. Lyakh, M. Repisky, L. Halbert, A. Severo Pereira Gomes, H. J. A. Jensen, and L. Visscher. Implementation of Relativistic Coupled Cluster Theory for Massively Parallel GPU-Accelerated Computing Architectures. *J. Chem. Theory Comput.*, 17(9):5509–5529, 2021.
- [229] B. P. Prascher, D. E. Woon, K. A. Peterson, T. H. Dunning Jr., and A. K. Wilson. Gaussian Basis Sets for Use in Correlated Molecular Calculations. VII. Valence, Core-Valence, and Scalar Relativistic Basis Sets for Li, Be, Na, and Mg. *Theor. Chem. Acc.*, 128:69–82, 2010.
- [230] G. D. Purvis and R. J. Bartlett. A Full Coupled-Cluster Singles and Doubles Model: The Inclusion of Disconnected Triples. *J. Chem. Phys.*, 76:1910–1918, 1982.
- [231] P. Pyykkö. Relativistic Effects in Structural Chemistry. *Chem. Rev.*, 88(3):563–594, 1988.
- [232] P. Pyykkö. Relativistic Effects in Chemistry: More Common Than You Thought. *Annu. Rev. Phys. Chem.*, 63(1):45–64, 2012.
- [233] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon. A Fifth-Order Perturbation Comparison of Electron Correlation Theories. *Chem. Phys. Lett.*, 157(6):479–483, 1989.
- [234] M. Reiher and A. Wolf. Exact Decoupling of the Dirac Hamiltonian. I. General Theory. *J. Chem. Phys.*, 121:2037, 2004.
- [235] M. Reiher and A. Wolf. Exact Decoupling of the Dirac Hamiltonian. II. The Generalized Douglas-Kroll-Hess Transformation Up to Arbitrary Order. *J. Chem. Phys.*, 121:10945, 2004.

- [236] M. Reiher and A. Wolf. *Relativistic Quantum Chemistry*. Wiley-VCH, second edition, 2015.
- [237] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, Jr. Fixed-Node Quantum Monte Carlo for Molecules. *J. Chem. Phys.*, 77(11):5593–5603, 1982.
- [238] R. D. Reynolds, T. Yanai, and T. Shiozaki. Large-Scale Relativistic Complete Active Space Self-Consistent Field with Robust Convergence. *J. Chem. Phys.*, 149:014106, 2018.
- [239] W. Ritz. Über eine neue methode zur lösung gewisser variationsprobleme der mathematischen physik. *J. Reine Angew. Math.*, 135:1–61, 1909.
- [240] M. Roemelt, S. Guo, and G. K.-L. Chan. A Projected Approximation to Strongly Contracted N-electron Valence Perturbation Theory for DMRG Wavefunctions. *J. Chem. Phys.*, 144(20):204113, May 2016.
- [241] M. Rohlfing and S. G. Louie. Electron-Hole Excitations and Optical Spectra from First Principles. *Phys. Rev. B*, 62(8):4927–4944, 2000.
- [242] Z. Rolik, A. Szabados, and P. R. Surján. A Sparse Matrix Based Full-Configuration Interaction Algorithm. *J. Chem. Phys.*, 128(14):144101, Apr. 2008.
- [243] B. O. Roos. The Complete Active Space Self-Consistent Field Method and Its Applications in Electronic Structure Calculations. *Adv. Chem. Phys.*, 69:399–445, 1987.
- [244] B. O. Roos, K. Andersson, M. P. Fülscher, P.-Å. Malmqvist, L. Serrano-Andrés, K. Pierloot, and M. Merchán. Multiconfigurational Perturbation Theory: Applications in Electronic Spectroscopy. *Adv. Chem. Phys.*, 93:219–331, 1996.
- [245] B. O. Roos, R. Lindh, P.-Å. Malmqvist, V. Veryazov, and P.-O. Widmark. Main Group Atoms and Dimers Studied with a New Relativistic ANO Basis Set. *J. Phys. Chem. A*, 108:2851–2858, 2004.
- [246] B. O. Roos, P. R. Taylor, and P. E. M. Sigbahn. A Complete Active Space SCF Method (CASSCF) using a Density Matrix Formulated Super-CI Approach. *Chem. Phys.*, 48:157–173, 1980.
- [247] Y. Saad. *Numerical Methods for Large Eigenvalue Problems*. Society for Industrial and Applied Mathematics, 2011.
- [248] M. Saitow, Y. Kurashige, and T. Yanai. Multireference Configuration Interaction Theory Using Cumulant Reconstruction with Internal Contraction of Density Matrix Renormalization Group Wave Function. *J. Chem. Phys.*, 139(4):044118, July 2013.

- [249] A. M. Sand and D. A. Mazziotti. Parametric Two-Electron Reduced-Density-Matrix Method with Application to Diradical Rectangular H_4 . *Comput. Theor. Chem*, 1003:44–49, Jan. 2013.
- [250] T. Saue. Relativistic Hamiltonians for Chemistry: A Primer. *ChemPhysChem*, 12:3077–3094, 2011.
- [251] Z. Schätzle, J. Hermann, and F. Noé. Convergence to the Fixed-Node Limit in Deep Variational Monte Carlo. *J. Chem. Phys.*, 154(12):124108, Mar. 2021.
- [252] J. B. Schriber and F. A. Evangelista. Communication: An Adaptive Configuration Interaction Approach for Strongly Correlated Electrons with Tunable Accuracy. *J. Chem. Phys.*, 144(16):161106, Apr. 2016.
- [253] E. Schrödinger. Quantisierung als Eigenwertproblem. *Ann. Phys.*, 384(4):361–376, 1926.
- [254] P. Schwerdtfeger. Metal-Metal Bonds in Thallium(I)–Thallium(I) Compounds: Fact or Fiction? *Inorg. Chem.*, 30(7):1660–1663, 1991.
- [255] G. E. Scuseria, C. A. Jiménez-Hoyos, T. M. Henderson, K. Samanta, and J. K. Ellis. Projected Quasiparticle Theory for Molecular Electronic Structure. *J. Chem. Phys.*, 135(12):124108, Sept. 2011.
- [256] G. E. Scuseria and H. F. Schaefer. A New Implementation of the Full CCSDT Model for Molecular Electronic Structure. *Chem. Phys. Lett.*, 152(4):382–386, 1988.
- [257] R. D. Senanayake, D. B. Lingerfelt, G. U. Kuda-Singappulige, X. Li, and C. M. Aikens. Real-Time TDDFT Investigation of Optical Absorption in Gold Nanowires. *J. Phys. Chem. C*, 123(23):14734–14745, 2019.
- [258] B. Senjean, S. Yalouz, N. Nakatani, and E. Fromager. Reduced Density Matrix Functional Theory from an *Ab Initio* Seniority-Zero Wave Function: Exact and Approximate Formulations Along Adiabatic Connection Paths. *Phys. Rev. A*, 106(3):032203, Sept. 2022.
- [259] S. Seritan, C. Bannwarth, B. S. Fales, E. G. Hohenstein, C. M. Isborn, S. I. L. Kokkila-Schumacher, X. Li, F. Liu, N. Luehr, J. W. Snyder Jr., C. Song, A. V. Titov, I. S. Ufimtsev, L.-P. Wang, and T. J. Martinez. TeraChem: A Graphical Processing Unit-Accelerated Electronic Structure Package for Large-Scale *Ab Initio* Molecular Dynamics. *WIREs Comput. Mol. Sci.*, 11(2):e1494, 2021.
- [260] P. Sharma, A. J. Jenkins, G. Scalmani, M. J. Frisch, D. G. Truhlar, L. Gagliardi, and X. Li. Exact-Two-Component Multiconfiguration Pair-Density Functional Theory. *J. Chem. Theory Comput.*, 18(5):2947–2954, 2022.

- [261] S. Sharma and A. Alavi. Multireference Linearized Coupled Cluster Theory for Strongly Correlated Systems Using Matrix Product States. *J. Chem. Phys.*, 143(10):102815, Aug. 2015.
- [262] 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. *J. Chem. Theory Comput.*, 13(4):1595–1604, Apr. 2017.
- [263] I. Shavitt. The Method of Configuration Interaction. In H. F. I. Schaefer, editor, *Methods of Electronic Structure Theory*, pages 189–275. Plenum: New York, 1977.
- [264] I. Shavitt. The History and Evolution of Configuration Interaction. *Mol. Phys.*, 94(1):3–17, 1998.
- [265] A. Shayit, C. Liao, S. Upadhyay, H. Hu, T. Zhang, A. E. DePrince III, C. Yang, and X. Li. Numerically Exact Configuration Interaction at Quadrillion-Determinant Scale. *Nat. Commun.*, 16(1):11016, 2025.
- [266] J. Shee, E. J. Arthur, S. Zhang, D. R. Reichman, and R. A. Friesner. Phaseless Auxiliary-Field Quantum Monte Carlo on Graphical Processing Units. *J. Chem. Theory Comput.*, 14(8):4109–4121, 2018.
- [267] C. D. Sherrill. An Introduction to Configuration Interaction Theory. Technical report, School Of Chemistry And Biochemistry, Georgia Institute Of Technology, Atlanta, GA, 1995.
- [268] C. D. Sherrill and H. F. Schaefer. The Configuration Interaction Method: Advances in Highly Correlated Approaches. *Adv. Quantum Chem.*, 34:143–269, 1999.
- [269] Y.-Y. Shi, L.-M. Duan, and G. Vidal. Classical Simulation of Quantum Many-Body Systems with a Tree Tensor Network. *Phys. Rev. A*, 74(2):022320, Aug. 2006.
- [270] T. Shiozaki and W. Mizukami. Relativistic Internally Contracted Multireference Electron Correlation Methods. *J. Chem. Theory Comput.*, 11:4733–4739, 2015.
- [271] P. E. Siegbahn. A New Direct CI Method for Large CI Expansions in a Small Orbital Space. *Chem. Phys. Lett.*, 109(5):417–423, 1984.
- [272] D. M. Silver and R. M. Stevens. Reaction Paths on the H₄ Potential Energy Surface. *J. Chem. Phys.*, 59(6):3378–3394, Sept. 1973.
- [273] J. C. Slater. The Theory of Complex Spectra. *Phys. Rev.*, 34:1293–1322, Nov 1929.

- [274] J. C. Slater. Atomic Shielding Constants. *Phys. Rev.*, 36(1):57–64, 1930.
- [275] P. Slavíček, R. Kalus, P. Paška, I. Odvárková, P. Hobza, and A. Malijevský. State-of-the-Art Correlated *Ab Initio* Potential Energy Curves for Heavy Rare Gas Dimers: Ar₂, Kr₂, and Xe₂. *J. Chem. Phys.*, 119(4):2102–2119, 07 2003.
- [276] V. Smil. *Energy in World History*. Westview Press, Boulder, 1994.
- [277] M. Srednicki. *Quantum Field Theory*. Cambridge University Press, Cambridge, 2007.
- [278] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch. *Ab Initio* Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields. *J. Phys. Chem.*, 98(45):11623–11627, 1994.
- [279] G. Strinati. Application of the Green’s function method to the study of the optical properties of semiconductors. *Riv. Nuovo Cimento*, 11:1–86, 1988.
- [280] Q. Sun. Libcint: An efficient general integral library for gaussian basis functions. *Journal of Computational Chemistry*, 36(22):1664–1671, 2015.
- [281] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan. PySCF: The Python-Based Simulations of Chemistry Framework. *WIREs Comput. Mol. Sci.*, 8(1):e1340, 2018.
- [282] Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui, J. J. Eriksen, Y. Gao, S. Guo, J. Hermann, M. R. Hermes, K. Koh, P. Koval, S. Lehtola, Z. Li, J. Liu, N. Mardirossian, J. D. McClain, M. Motta, B. Mussard, H. Q. Pham, A. Pulkin, W. Purwanto, P. J. Robinson, E. Ronca, E. R. Sayfutyarova, M. Scheurer, H. F. Schurkus, J. E. T. Smith, C. Sun, S.-N. Sun, S. Upadhyay, L. K. Wagner, X. Wang, A. White, J. D. Whitfield, M. J. Williamson, S. Wouters, J. Yang, J. M. Yu, T. Zhu, T. C. Berkelbach, S. Sharma, A. Y. Sokolov, and G. K.-L. Chan. Recent Developments in the PySCF Program Package. *J. Chem. Phys.*, 153(2):024109, 07 2020.
- [283] A. Szabo and N. S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Publications, Mineola, NY, 1996.
- [284] P. G. Szalay, T. Müller, G. Gidofalvi, H. Lischka, and R. Shepard. Multiconfiguration Self-Consistent Field and Multireference Configuration Interaction Methods and Applications. *Chem. Rev.*, 112:108–181, 2012.
- [285] L. Tagliacozzo, G. Evenbly, and G. Vidal. Simulation of Two-Dimensional Quantum Systems Using a Tree Tensor Network That Exploits the Entropic Area Law. *Phys. Rev. B*, 80(23):235127, Dec. 2009.

- [286] P. K. Tamukong, M. R. Hoffmann, Z. Li, and W. Liu. Relativistic GVVPT2 Multireference Perturbation Theory Description of the Electronic States of Y_2 and Tc_2 . *J. Phys. Chem. A*, 118(8):1489–1501, 2014.
- [287] J. Tao and J. P. Perdew. Test of a Nonempirical Density Functional: Short-Range Part of the Van der Waals Interaction in Rare-Gas Dimers. *J. Chem. Phys.*, 122(11):114102, 03 2005.
- [288] J. S. Thayer. Relativistic Effects and the Chemistry of the Heaviest Main-Group Element. *J. Chem. Educ.*, 82(11):1721, 2005.
- [289] J. Thyssen, T. Fleig, and H. J. A. Jensen. A Direct Relativistic Four-Component Multiconfiguration Self-Consistent-Field Method for Molecules. *J. Chem. Phys.*, 129:034109, 2008.
- [290] I. Toli and S. Zou. Schrödinger Equation with Coulomb Potential Admits No Exact Solutions. *Chem. Phys. Lett.*, 737:100021, 2019. Articles initially published in *Chemical Physics Letters*: X 1-4, 2019.
- [291] 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. *J. Chem. Phys.*, 145(4):044112, July 2016.
- [292] R.-D. Urban, A. H. Bahnmaier, U. Magg, and H. Jones. The Diode Laser Spectrum of Thallium Hydride (^{205}TlH and ^{203}TlH) in Its Ground Electronic State. *Chem. Phys. Lett.*, 158(5):443–446, 1989.
- [293] T. Van Voorhis and M. Head-Gordon. Benchmark Variational Coupled Cluster Doubles Results. *J. Chem. Phys.*, 113(20):8873–8879, Nov. 2000.
- [294] C. van Wüllen. Relation Between Different Variants of the Generalized Douglas-Kroll Transformation through Sixth Order. *J. Chem. Phys.*, 120(16):7307–7313, 2004.
- [295] P. Čársky. Configuration Interaction. In P. v. R. Schleyer, editor, *Encyclopedia of Computational Chemistry*. John Wiley & Sons, Ltd, 2002.
- [296] J. Čížek. On the Correlation Problem in Atomic and Molecular Systems. Calculation of Wavefunction Components in Ursell-Type Expansion Using Quantum-Field Theoretical Methods. *J. Chem. Phys.*, 45(11):4256–4266, 12 1966.
- [297] E. Vecharynski, C. Yang, and F. Xue. Generalized Preconditioned Locally Harmonic Residual Method for Non-Hermitian Eigenproblems. *SIAM J. Sci. Comp.*, 38(1):A500–A527, 2016.

- [298] F. Verstraete, V. Murg, and J. Cirac. Matrix Product States, Projected Entangled Pair States, and Variational Renormalization Group Methods for Quantum Spin Systems. *Adv. Phys.*, 57(2):143–224, Mar. 2008.
- [299] L. Visscher and K. Dyall. Dirac-Fock Atomic Electronic Structure Calculations Using Different Nuclear Charge Distributions. *At. Data Nucl. Data Tables*, 67(2):207 – 224, 1997.
- [300] B. Vlaisavljevich and T. Shiozaki. Nuclear Energy Gradients for Internally Contracted Complex Active Space Second-Order Perturbation Theory: Multistate Extensions. *J. Chem. Theory Comput.*, 12:3781–3787, 2016.
- [301] K. D. Vogiatzis, G. Li Manni, S. J. Stoneburner, D. Ma, and L. Gagliardi. Systematic Expansion of Active Spaces beyond the CASSCF Limit: A GASSCF/SplitGAS Benchmark Study. *J. Chem. Theory Comput.*, 11:3010–3021, 2015.
- [302] 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. *J. Chem. Phys.*, 147:184111, 2017.
- [303] X. Wang, P. F. Souter, and L. Andrews. Infrared Spectra of Antimony and Bismuth Hydrides in Solid Matrixes. *J. Phys. Chem. A*, 107(21):4244–4249, 2003.
- [304] H.-J. Werner and P. J. Knowles. An Efficient Internally Contracted Multiconfiguration-Reference Configuration Interaction Method. *J. Chem. Phys.*, 89:5803–5814, 1988.
- [305] H.-J. Werner and W. Meyer. A Quadratically Convergent MCSCF Method for the Simultaneous Optimization of Several States. *J. Chem. Phys.*, 74:5794–5801, 1981.
- [306] O. Weser, K. Guther, K. Ghanem, and G. Li Manni. Stochastic Generalized Active Space Self-Consistent Field: Theory and Application. *J. Chem. Theory Comput.*, 18:251–272, 2022.
- [307] J. L. Whitten and M. Hackmeyer. Configuration Interaction Studies of Ground and Excited States of Polyatomic Molecules. I. The CI Formulation and Studies of Formaldehyde. *J. Chem. Phys.*, 51(12):5584–5596, 12 1969.
- [308] D. B. Williams-Young, A. Petrone, S. Sun, T. F. Stetina, P. Lestrangle, C. E. Hoyer, D. R. Nascimento, L. Koulias, A. Wildman, J. Kasper, J. J. Goings, F. Ding, A. E. DePrince III, E. F. Valeev, and X. Li. The Chronus Quantum (ChronusQ) Software Package. *WIREs Comput. Mol. Sci.*, 10:e1436, 2020.
- [309] A. Wolf and M. Reiher. Exact Decoupling of the Dirac Hamiltonian. III. Molecular Properties. *J. Chem. Phys.*, 121:064102, 2004.

- [310] A. Wolf, M. Reiher, and B. A. Hess. The Generalized Douglas-Kroll Transformation. *J. Chem. Phys.*, 117(20):9215–9226, 2002.
- [311] S. Wouters and D. Van Neck. The Density Matrix Renormalization Group for *Ab Initio* Quantum Chemistry. *Euro. Phys. J. D*, 68(9):272, Sept. 2014.
- [312] E. Xu, D. Zhao, and S. Li. Multireference Second Order Perturbation Theory with a Simplified Treatment of Dynamical Correlation. *J. Chem. Theory Comput.*, 11(10):4634–4643, Oct. 2015.
- [313] K. Yosida. *Functional Analysis*, volume 123 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, New York, 6th edition, 1980.
- [314] V. W.-z. Yu and M. Govoni. GPU Acceleration of Large-Scale Full-Frequency GW Calculations. *J. Chem. Theory Comput.*, 18(8):4690–4707, 2022.
- [315] S. H. Yuwono, B. C. Cooper, T. Zhang, X. Li, and I. DePrince, A. Eugene. Time-Dependent Equation-of-Motion Coupled-Cluster Simulations with a Defective Hamiltonian. *J. Chem. Phys.*, 159(4):044113, 2023.
- [316] S. H. Yuwono, R. R. Li, T. Zhang, X. Li, and I. DePrince, A. Eugene. Two-Component Relativistic Equation-of-Motion Coupled Cluster for Electron Ionization. *J. Chem. Phys.*, 162(8):084110, 02 2025.
- [317] S. H. Yuwono, R. R. Li, T. Zhang, K. A. Surjuse, E. F. Valeev, X. Li, and A. I. Eugene DePrince. Relativistic Coupled Cluster with Completely Renormalized and Perturbative Triples Corrections. *J. Phys. Chem. A*, 128(31):6521–6539, 2024.
- [318] S. Zarrabian, C. Sarma, and J. Paldus. Vectorizable Approach to Molecular CI Problems Using Determinantal Basis. *Chem. Phys. Lett.*, 155(2):183–188, 1989.
- [319] B. Zhang, J. E. Vandezande, R. D. Reynolds, and H. F. Schaefer. Spin–Orbit Coupling via Four-Component Multireference Methods: Benchmarking on p-Block Elements and Tentative Recommendations. *J. Chem. Theory Comput.*, 14:1235–1246, 2018.
- [320] N. Zhang, W. Liu, and M. R. Hoffmann. Iterative Configuration Interaction with Selection. *J. Chem. Theory Comput.*, 16(4):2296–2316, Apr. 2020.
- [321] N. Zhang, W. Liu, and M. R. Hoffmann. Further Development of iCIPT2 for Strongly Correlated Electrons. *J. Chem. Theory Comput.*, 17(2):949–964, Feb. 2021.
- [322] T. Zhang, S. Banerjee, L. N. Koulias, E. F. Valeev, A. E. I. DePrince, and X. Li. Dirac–Coulomb–Breit Molecular Mean-Field Exact-Two-Component Relativistic Equation-of-Motion Coupled-Cluster Theory. *J. Phys. Chem. A*, 128(17):3408–3418, 2024.

- [323] Y. Zhang, W. Pan, and W. Yang. Describing Van der Waals Interaction in Diatomic Molecules with Generalized Gradient Approximations: The Role of the Exchange Functional. *J. Chem. Phys.*, 107(19):7921–7925, 11 1997.
- [324] Y. Zhao and D. G. Truhlar. The M06 Suite of Density Functionals for Main Group Thermochemistry, Thermochemical Kinetics, Noncovalent Interactions, Excited States, and Transition Elements: Two New Functionals and Systematic Testing of Four M06-Class Functionals and 12 Other Functionals. *Theor. Chem. Acc.*, 120:215–241, 2008.
- [325] Y. Zhou, R. Shepard, and M. Minkoff. Computing Eigenvalue Bounds for Iterative Subspace Matrix Methods. *Comp. Phys. Comm.*, 167(2):90–102, 2005.
- [326] Z. Zhou and S. M. Parker. Converging Time-Dependent Density Functional Theory Calculations in Five Iterations with Minimal Auxiliary Preconditioning. *J. Chem. Theory Comput.*, 20(15):6738–6746, 2024.
- [327] P. Zhu, M. E. Argentati, and A. V. Knyazev. Bounds for the Rayleigh Quotient and the Spectrum of Self-Adjoint Operators. *SIAM J. Matrix Anal. Appl.*, 34(1):244–256, 2013.
- [328] W. Zhu, J. Toulouse, A. Savin, and J. G. Ángyán. Range-Separated Density-Functional Theory with Random Phase Approximation Applied to Noncovalent Intermolecular Interactions. *J. Chem. Phys.*, 132(24):244108, 06 2010.

Appendix A

**SUPPLEMENTARY INFORMATION TO CATEGORICAL
COMPRESSION WITHIN THE STP-DAS FRAMEWORK*****A.1 Convergence Behavior of Compression-Compatible Preconditioning***

Table A.1 presents the convergence behavior of the STP-DAS framework using a compression-compatible preconditioner. The error in energy ($|\delta E|$) and the norm of the residual ($\|\mathbf{r}\|$) with varying truncation level (ϵ) are explored. Three systems with varying degrees of electron correlation, including a magnesium atom, a diatomic nitrogen, and a model carbon nanotube (CNT), were studied using relativistic X2C-CASCI calculations across different active space sizes. Diatomic nitrogen calculations were performed using a bond length of 1.097 Å. The CNT geometry was generated using TubeGen 3.4 with two unit cells and a chirality index of (3, 3) [90]. Results are compared against exact X2C-CASCI calculations performed with the traditional Davidson preconditioner [49].

It is evident in Table A.1, where all three systems, regardless of the level of correlation, attain eigenvalues within $10^{-7} E_h$ of their exact X2C-CASCI counterparts when the residual norms $\|\mathbf{r}\|$ are small. For loose thresholds, new trial expansion vectors quickly become linearly dependent on the existing subspace expansion vectors, indicating that the span of the modified expansion vectors has grown to its limit before convergence. The resulting energies at these thresholds differ significantly from their exact counterparts, as expected. Furthermore, the residual norm remains fairly high, indicating that the Ritz vector is still far from the true eigenvector.

Table A.1. The error $|\delta E|$ (in Hartree) in X2C-CASCI ground state energies attained by the compression-compatible preconditioner (relative to the corresponding energies attained by the conventional Davidson preconditioner) for magnesium (cc-pVTZ-DK [53,229]), diatomic nitrogen (ANO-RCC [245]), and a carbon nanotube (6-311G(d,p) [158]). Active spaces are labeled as (number of orbitals, number of electrons). The truncation threshold is ε and the residual norm is $\|\mathbf{r}\|$. Red indicates errors above $10^{-7} E_h$.

Threshold	Active Space					
Mg	(64, 8)		(38, 8)		(22, 8)	
N_{det}	$\approx 4.4 \times 10^9$		$\approx 4.9 \times 10^7$		$\approx 3.2 \times 10^5$	
ε	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $
1×10^{-1}	1.4×10^{-2}	2.5×10^{-1}	5.8×10^{-3}	1.7×10^{-1}	3.4×10^{-4}	3.9×10^{-2}
1×10^{-2}	6.0×10^{-4}	6.0×10^{-2}	8.3×10^{-5}	2.4×10^{-2}	$< 10^{-7}$	6.0×10^{-5}
1×10^{-3}	1.3×10^{-6}	3.6×10^{-3}	$< 10^{-7}$	9.7×10^{-5}	$< 10^{-7}$	8.10×10^{-6}
1×10^{-4}	$< 10^{-7}$	1.2×10^{-4}	$< 10^{-7}$	4.3×10^{-5}	$< 10^{-7}$	4.0×10^{-6}
1×10^{-5}	$< 10^{-7}$	4.6×10^{-5}	$< 10^{-7}$	6.5×10^{-5}	$< 10^{-7}$	2.9×10^{-6}
N ₂	(44, 10)		(30, 10)		(22, 10)	
N_{det}	$\approx 2.5 \times 10^9$		$\approx 3.0 \times 10^7$		$\approx 6.5 \times 10^5$	
ε	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $
1×10^{-1}	1.7×10^{-2}	1.7×10^{-1}	5.5×10^{-3}	1.0×10^{-1}	3.4×10^{-4}	2.7×10^{-2}
1×10^{-2}	4.3×10^{-3}	3.7×10^{-2}	1.0×10^{-5}	5.9×10^{-3}	$< 10^{-7}$	8.7×10^{-5}
1×10^{-3}	5.9×10^{-7}	1.8×10^{-3}	$< 10^{-7}$	7.5×10^{-5}	$< 10^{-7}$	7.6×10^{-5}
1×10^{-4}	$< 10^{-7}$	5.5×10^{-5}	$< 10^{-7}$	7.7×10^{-5}	$< 10^{-7}$	4.7×10^{-5}
1×10^{-5}	$< 10^{-7}$	5.1×10^{-5}	$< 10^{-7}$	6.5×10^{-5}	$< 10^{-7}$	4.6×10^{-5}
CNT	(40, 12)		(28, 12)		(20, 12)	
N_{det}	$\approx 5.6 \times 10^9$		$\approx 3.0 \times 10^7$		$\approx 1.3 \times 10^5$	
ε	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $	$ \delta E $	$\ \mathbf{r}\ $
1×10^{-1}	6.9×10^{-2}	1.9×10^{-1}	4.9×10^{-2}	1.6×10^{-1}	1.7×10^{-2}	9.2×10^{-2}
1×10^{-2}	8.4×10^{-3}	8.5×10^{-2}	3.4×10^{-3}	5.5×10^{-2}	3.7×10^{-6}	2.0×10^{-3}
1×10^{-3}	4.9×10^{-4}	2.5×10^{-2}	1.1×10^{-7}	4.3×10^{-4}	$< 10^{-7}$	8.3×10^{-5}
1×10^{-4}	$< 10^{-7}$	4.6×10^{-4}	$< 10^{-7}$	7.2×10^{-5}	$< 10^{-7}$	6.9×10^{-5}
1×10^{-5}	$< 10^{-7}$	5.0×10^{-5}	$< 10^{-7}$	5.7×10^{-5}	$< 10^{-7}$	6.9×10^{-5}

A.2 Cost Analysis

Table A.2. Computational Cost Analysis.

Ref.	Time-to-Completion (hours)	# of dets	# of Nodes	# of Cores per Node	Cost ^a
This Work	34.5	1.05×10^{15}	1,000	128	2.4×10^{-3}
This Work	0.8	6.29×10^{10}	25	40	1.2×10^2
This Work	1.1	1.52×10^{11}	25	40	2.9×10^1
95	113.6	1.31×10^{12}	256	80	4.9×10^3
122 ^b	7.0	2.30×10^{11}	1,000	16	–
302 ^c	13.9	9.14×10^{11}	256	32	–

^a The computational cost is estimated using $\frac{\text{Core-Time-to-Completion}}{(\# \text{ of dets})^2}$ and $\frac{\text{Core-Time-to-Completion}}{6(\# \text{ of dets})^2}$ for non-relativistic and relativistic calculations, respectively. The unit is $\frac{\text{core-second}}{\text{exaFLOP}}$.

^b Only a single iteration was performed using the full CI vector without any screening or sparsity applied.

^c Only a single iteration was performed using the full CI vector with sparsity.

A.3 Exploiting Symmetry

To showcase the performance of the compression-compatible STP-DAS framework on highly-symmetric systems, we performed an X2C-FCI calculation for the ground state of a magnesium atom within the aug-cc-pVTZ-DK [53,229] basis set. Table A.3 summarizes the results of this calculation. The ground-state energy converged in 7 iterations to microhartree precision ($< 10^{-6}$ a.u.), with a total σ -build time of 8.6 hours. A total of 7 expansion vectors were involved in the σ -build. Despite the enormous configuration space of over one quadrillion (1.05×10^{15}) complex-valued 2-spinor determinants, the average σ -build time per vector remained just 1.2 hours.

The energy of the magnesium atom computed from this large CI calculation is -199.973455458424 a.u. Leveraging the gap theorem [50,217], we determine that our X2C-FCI result lies within 7×10^{-6} a.u. of the true aug-cc-pVTZ-DK [53, 229] ground state energy, well below any chemically meaningful threshold.

Table A.3. Details of the convergence of the X2C-FCI calculation (100 2-spinor orbitals, 12 electrons, 10^{15} 2-spinor determinants) of the ground state of a Mg atom (aug-cc-pVTZ-DK [53, 229]). Each row corresponds to a Davidson iteration. The fourth and fifth columns track the changes in the computed Rayleigh–Ritz eigenpair across iterations, and the sixth column tracks the residual norm, $\|\mathbf{r}\| = \|\mathbf{H}\mathbf{C} - E\mathbf{C}\|$, where \mathbf{H} is the Hamiltonian, \mathbf{C} is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\varepsilon = 1.0 \times 10^{-5}$.

Iteration	Duration (s)	Energy (E_h)	ΔE (E_h)	$\max\{ \Delta c \}$	$\ \mathbf{r}\ $ (E_h)
1	8	-199.919557510276	–	–	3.83×10^{-1}
2	11	-199.968406808290	4.88×10^{-2}	6.34×10^{-2}	1.25×10^{-1}
3	255	-199.973110240682	4.70×10^{-3}	1.77×10^{-2}	4.12×10^{-2}
4	1429	-199.973426870322	3.17×10^{-4}	3.71×10^{-3}	1.36×10^{-2}
5	3211	-199.973453427878	2.66×10^{-5}	1.29×10^{-3}	4.67×10^{-3}
6	7920	-199.973455300845	1.87×10^{-6}	2.56×10^{-4}	1.74×10^{-3}
7	18202	-199.973455458424	1.58×10^{-7}	7.56×10^{-5}	8.31×10^{-4}