

**Towards an Accurate Description of Strongly
Correlated Chemical Systems with Phaseless
Auxiliary-Field Quantum Monte Carlo -
Methodological Advances and Applications**

James Shee

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2019

ABSTRACT

Towards an Accurate Description of Strongly Correlated Chemical Systems with Phaseless Auxiliary-Field Quantum Monte Carlo - Methodological Advances and Applications

James Shee

The exact and phaseless variants of auxiliary-field quantum Monte Carlo (AFQMC) have been shown to be capable of producing accurate ground-state energies for a wide variety of systems including those which exhibit substantial electron correlation effects. The first chapter of this thesis will provide an overview of the relevant electronic structure problem, and the phaseless AFQMC (ph-AFQMC) methodology.

The computational cost of performing these calculations has to date been relatively high, impeding many important applications of these approaches. In Chapter 2 we present a correlated sampling methodology for AFQMC which relies on error cancellation to dramatically accelerate the calculation of energy differences of relevance to chemical transformations. In particular, we show that our correlated sampling-based ph-AFQMC approach is capable of calculating redox properties, deprotonation free energies, and hydrogen abstraction energies in an efficient manner without sacrificing accuracy. We validate the computational protocol by calculating the ionization potentials and electron affinities of the atoms contained in the G2 test set and then proceed to utilize a composite method, which treats fixed-geometry processes with correlated sampling-based AFQMC and relaxation energies via MP2, to compute the ionization potential, deprotonation free energy, and the O-H bond dissociation energy of methanol, all to within chemical accuracy. We show that the efficiency of correlated sampling relative to uncorrelated calculations increases with system and basis set size and that correlated sampling greatly reduces the required number of random walkers to achieve a target statistical error. This translates to reductions in wall-times by

factors of 55, 25, and 24 for the ionization potential of the K atom, the deprotonation of methanol, and hydrogen abstraction from the O-H bond of methanol, respectively.

In Chapter 3 we present an implementation of ph-AFQMC utilizing graphical processing units (GPUs). The AFQMC method is recast in terms of matrix operations which are spread across thousands of processing cores and are executed in batches using custom Compute Unified Device Architecture kernels and the hardware-optimized cuBLAS matrix library. Algorithmic advances include a batched Sherman-Morrison-Woodbury algorithm to quickly update matrix determinants and inverses, density-fitting of the two-electron integrals, an energy algorithm involving a high-dimensional precomputed tensor, and the use of single-precision floating point arithmetic. These strategies result in dramatic reductions in wall-times for both single- and multi-determinant trial wavefunctions. For typical calculations we find speed-ups of roughly two orders of magnitude using just a single GPU card. Furthermore, we achieve near-unity parallel efficiency using 8 GPU cards on a single node, and can reach moderate system sizes via a local memory-slicing approach. We illustrate the robustness of our implementation on hydrogen chains of increasing length, and through the calculation of all-electron ionization potentials of the first-row transition metal atoms. We compare long imaginary-time calculations utilizing a population control algorithm with our previously published correlated sampling approach, and show that the latter improves not only the efficiency but also the accuracy of the computed ionization potentials. Taken together, the GPU implementation combined with correlated sampling provides a compelling computational method that will broaden the application of ph-AFQMC to the description of realistic correlated electronic systems.

In Chapter 4 the bond dissociation energies of a set of 44 *3d* transition metal-containing diatomics are computed with ph-AFQMC utilizing the correlated sampling technique. We investigate molecules with H, N, O, F, Cl, and S ligands, including those in the 3dMLBE20 database first compiled by Truhlar and co-workers with calculated and experimental values that have since been revised by various groups. In order to make a direct comparison of the accuracy of our ph-AFQMC calculations with previously published results from 10 DFT functionals, CCSD(T), and icMR-CCSD(T), we establish an objective selection protocol which utilizes the most recent exper-

imental results except for a few cases with well-specified discrepancies. With the remaining set of 41 molecules, we find that ph-AFQMC gives robust agreement with experiment superior to that of all other methods, with a mean absolute error (MAE) of 1.4(4) kcal/mol and maximum error of 3(3) kcal/mol (parenthesis account for reported experimental uncertainties and the statistical errors of our ph-AFQMC calculations). In comparison, CCSD(T) and B97, the best performing DFT functional considered here, have MAEs of 2.8 and 3.7 kcal/mol, respectively, and maximum errors in excess of 17 kcal/mol (for the CoS diatomic). While a larger and more diverse data set would be required to demonstrate that ph-AFQMC is truly a benchmark method for transition metal systems, our results indicate that the method has tremendous potential, exhibiting unprecedented consistency and accuracy compared to other approximate quantum chemical approaches.

The energy gap between the lowest-lying singlet and triplet states is an important quantity in chemical photocatalysis, with relevant applications ranging from triplet fusion in optical upconversion to the design of organic light-emitting devices. The *ab initio* prediction of singlet-triplet (ST) gaps is challenging due to the potentially biradical nature of the involved states, combined with the potentially large size of relevant molecules. In Chapter 5, we show that ph-AFQMC can accurately predict ST gaps for chemical systems with singlet states of highly biradical nature, including a set of 13 small molecules and the ortho-, meta-, and para- isomers of benzyne. With respect to gas-phase experiments, ph-AFQMC using CASSCF trial wavefunctions achieves a mean averaged error of ~ 1 kcal/mol. Furthermore, we find that in the context of a spin-projection technique, ph-AFQMC using unrestricted single-determinant trial wavefunctions, which can be readily obtained for even very large systems, produces equivalently high accuracy. We proceed to show that this scalable methodology is capable of yielding accurate ST gaps for all linear polyacenes for which experimental measurements exist, i.e. naphthalene, anthracene, tetracene, and pentacene. Our results suggest a protocol for selecting either unrestricted Hartree-Fock or Kohn-Sham orbitals for the single-determinant trial wavefunction, based on the extent of spin-contamination. These findings provide a reliable computational tool with which to investigate specific photochemical processes involving large molecules that may have substantial biradical character. We compute the ST gaps for a set of anthracene derivatives which are potential triplet-triplet annihilators for op-

tical upconversion, and compare our ph-AFQMC predictions with those from DFT and CCSD(T) methods.

We conclude with a discussion of ongoing projects, further methodological improvements on the horizon, and future applications of ph-AFQMC to chemical systems of interest in the fields of biology, drug-discovery, catalysis, and condensed matter physics.

Table of Contents

List of Figures	iv
List of Tables	ix
Chapter 1: Introduction	1
1.1 Overview of AFQMC - Exact Method and the Phaseless Constraint	10
Chapter 2: Chemical transformations approaching chemical accuracy via correlated sampling in auxiliary-field quantum Monte Carlo²	14
2.1 Introduction	14
2.2 Methods	16
2.2.1 Correlated Sampling Methods for AFQMC	16
2.2.2 Utilizing Optimal Correlation in Molecular Applications	20
2.2.3 Computational Details	22
2.3 Results	25
2.3.1 Atomic IPs and EAs	25
2.3.2 The Case of Methanol	27
2.3.3 Basis Set Size, Number of Random Walkers, and CPU-time Reduction	28
2.4 Conclusions and Outlook	34
Chapter 3: Phaseless Auxiliary-Field Quantum Monte Carlo on Graphical Processing Units⁴	36
3.1 GPU Implementation	38

²Based on work published in *J. Chem. Theory Comput.* 2017, 13, 6, 2667-2680.

⁴Based on work published in *J. Chem. Theory Comput.* 2018, 14, 8, 4109-4121.

3.2	Illustration with Hydrogen Chains	43
3.2.1	Computational Details	44
3.2.2	Timings	44
3.2.3	Accuracy	47
3.3	IPs of Transition Metal Atoms	49
3.3.1	Computational Details	49
3.3.2	Results and Discussion	51
3.4	Conclusions and Outlook	56
Chapter 4: On Achieving High Accuracy in Quantum Chemical Calculations of		
3d Transition Metal-containing Diatomics with Auxiliary-Field Quantum Monte		
Carlo⁶		59
4.1	Introduction	59
4.2	Computational Details	65
4.3	Correlated Sampling for BDEs	70
4.4	Experimental BDEs	72
4.5	BDEs of Transition metal-containing Diatomics	74
4.5.1	Selection of Best Values and Comparisons of ph-AFQMC with DFT and CC	
	methods	82
4.6	Conclusions and Outlook	86
Chapter 5: Singlet-Triplet Energy Gaps of Organic Biradicals and Polyacenes		
with Auxiliary-Field Quantum Monte Carlo		89
5.1	Introduction	89
5.2	Computational Details	92
5.3	Results and Discussion	94
5.3.1	Small Molecule Biradical Set	94
5.3.2	Benzyne Isomers	99
5.3.3	Polyacenes	104

⁶Based on work published in *J. Chem. Theory Comput.* 2019, 15, 42346-2358.

5.3.4	Potential Annihilators for Optical Upconversion - Anthracene Derivatives . .	107
5.4	Conclusions	110
5.A	Appendix - Ph-AFQMC data from various trial wavefunctions	111
5.A.1	Small Molecule Biradical Set	111
5.A.2	Benzyne Isomers	113
5.A.3	Polyacenes	114
Chapter 6:	Conclusions	116
6.1	Current Status	116
6.2	Outlook	118
6.2.1	Additional Methodological Improvements	118
6.2.2	Future Applications	119
Bibliography		121

List of Figures

- Figure 2.1 Comparison of correlated and uncorrelated sampling for the IP of the K atom in a 6-31+G* basis, with $\Delta\tau = 0.01$, 12 walkers per repeat, and a HF reference state. The error bars give the standard errors (the standard deviation times $\frac{1}{\sqrt{N_r}}$, where N_r is the number of repeats) of the mean values among the repeats at each τ point in (a), and of the cumulative averages in (b). These standard errors are plotted in the insets, along with the scaled standard error resulting from an uncorrelated run in which PC was used with 360 walkers per repeat (dotted green). 18
- Figure 2.2 Schematic of the PES version of correlated sampling. Preliminary equilibration of the primary system is shown in red. This phase can be as long as necessary due to the use of PC. The resulting walkers are used to initialize the secondary system, after which the two systems are propagated with their own Hamiltonians using correlated sampling. Equilibration of the secondary system is relatively rapid, allowing measurements to be taken before the noise growth becomes prohibitive. . . 19
- Figure 2.3 Standard errors resulting from uncorrelated (red) and correlated (blue) propagation of the repeats for methanol in the cc-pVTZ basis. $\Delta\tau = 0.01$ and 24 walkers per repeat were used. 21
- Figure 2.4 Comparison of correlated and uncorrelated sampling for the fixed-geometry deprotonation of CH₃OH in the cc-pVTZ basis with “ghost” basis functions. We use $\Delta\tau = 0.01$, 12 walkers per repeat, and a HF reference state. The error bars, plotted in the insets, in (a) represent the standard error among the repeats at each τ ; those in (b) give the standard error of the cumulative averages. 22

Figure 2.5	Comparison of correlated and uncorrelated sampling for the energy difference corresponding to the removal of H \cdot from the O-H bond of CH $_3$ OH, with the rest of the geometry fixed. The cc-pVTZ basis is used, with $\Delta\tau = 0.01$, 12 walkers per repeat, and a CASSCF reference state. The error bars, plotted in the insets, in (a) represent the standard error among the repeats at each τ ; those in (b) give the standard error of the cumulative averages.	23
Figure 2.6	Comparison of the standard errors of the cumulative averages resulting from correlated and uncorrelated sampling in computing the IP of the K atom in three different basis sets with 5040 walkers per repeat. The inset zooms in on the lower region, for clarity.	30
Figure 2.7	Comparison of the errors resulting from correlated and uncorrelated sampling in computing the fixed-geometry deprotonation energies of H $_2$ O (M=58), CH $_3$ OH (M=116), and C $_2$ H $_5$ OH (M=174) with 192 walkers per repeat.	30
Figure 2.8	Dependence of standard error on the number of random walkers per repeat for the IP of K in the 6-31+G* basis. The inset highlights the errors of the uncorrelated run with 1056 walkers and the correlated run with 24 walkers, compared with the 0.5 mHa error target.	31
Figure 2.9	Dependence of standard error on the number of random walkers per repeat for the fixed-geometry deprotonation of methanol in the cc-pVTZ basis. The inset highlights the errors of the uncorrelated run with 2400 walkers and the correlated run with 96 walkers, compared with the 1 mHa error target.	32
Figure 2.10	Comparison of the standard errors resulting from the use of correlated sampling with 12 walkers per repeat and uncorrelated sampling with 288 walkers per repeat to calculate the energy difference associated with the fixed-geometry removal of H \cdot from the O-H bond of methanol in the cc-pVTZ basis. The 2 mHa error target is shown in black.	33

Figure 3.1	Flowchart of our AFQMC implementation, for 9 walkers ($k = 0, \dots, 8$) and 3 CPU-GPU pairs (CPU/GPU 0, 1, 2). If the size of the precomputed tensor used in the energy evaluation exceeds the memory capacity of each device, each GPU precomputes and stores only a slice of the tensor, and the energy of a walker is computed by circulating the walker to all other GPUs and tallying the partial energies obtained from the resident slices.	39
Figure 3.2	Propagation time vs the number of configurations in the CASSCF trial wavefunction for the Mn atom in the aug-cc-pwCVQZ-DK basis. Calculations use sp, a CD threshold of $10^{-4}Ha$, 20 walkers and imaginary-time trajectories of length $1Ha^{-1}$ with a time step of $\Delta\tau = 0.005Ha^{-1}$. Walker orthonormalization and local energy measurements were performed every 2 and 20 steps, respectively.	41
Figure 3.3	Propagation time using 1 GPU for hydrogen chains of varying lengths, comparing two types of two-electron integral decompositions, DF vs CD with a 10^{-5} cutoff, within both sp and dp. UHF trial functions are used, and 24 walkers are propagated for an imaginary-time segment of length $1Ha^{-1}$	45
Figure 3.4	Parallel efficiency of our ph-AFQMC code illustrated on H_{50} . We use a CASSCF trial wavefunction with 44 determinants, and 800 walkers propagated for $0.5Ha^{-1}$ with $\Delta\tau = 0.01Ha^{-1}$	47
Figure 3.5	Comparison of the IP of Cu as a function of imaginary-time produced from ph-AFQMC/CS in the aug-cc-pwCVQZ-DK basis compared to the regular ph-AFQMC/PC result in the same basis. The i-FCIQMC result in the aug-cc-pVQZ-DK basis is indicated by a dashed line.	54
Figure 4.1	For a set of node counts on the Summit computing cluster, we plot the parallel efficiency, defined as the speed-up over a 1 node calculation divided by the number of nodes. Each node utilizes 6 NVIDIA V100 GPU cards.	69

Figure 4.2	Ratio of the standard errors, as a function of imaginary time, resulting from ph-AFQMC calculations with correlated vs uncorrelated sampling, for the five Mn-containing diatomic species.	71
Figure 4.3	Deviations [kcal/mol] of various calculations and alternate experiments (when relevant) from the experimental values used by de Oliveira-Filho (GSOF) and co-workers in Ref. 35. For calculations and experiments, error bars represent statistical error and quoted experimental uncertainties, respectively.	75
Figure 4.4	Same as Fig. 4.3, but for Ti-containing diatomics.	76
Figure 4.5	Same as Fig. 4.3, but for V-containing diatomics. VH was not considered in Ref. 35, and we show the experimental result selected by Truhlar and co-workers in Ref. 229.	76
Figure 4.6	Same as Fig. 4.3, but for Cr-containing diatomics. For CrO we also show the experiment selected by Dixon and co-workers in Ref. 230, since it is not consistent with that chosen by de Oliveira-Filho and co-workers[35], given the reported uncertainties.	77
Figure 4.7	Same as Fig. 4.3, but for Mn-containing diatomics.	78
Figure 4.8	Same as Fig. 4.3, but for Fe-containing diatomics.	78
Figure 4.9	Same as Fig. 4.3, but for Co-containing diatomics.	79
Figure 4.10	Same as Fig. 4.3, but for Ni-containing diatomics.	79
Figure 4.11	Same as Fig. 4.3, but for Cu-containing diatomics.	80
Figure 4.12	Same as Fig. 4.3, but for Zn-containing diatomics.	80
Figure 4.13	Statistical summary of the accuracy of bond-dissociation energy predictions, for a representative set of computational methods.	87
Figure 5.1	Deviations [kcal/mol] of various calculations from experimentally-derived reference values. Error bars show the statistical error of the QMC measurements. . . .	95
Figure 5.2	Deviations [kcal/mol] of computational methods selected from Ref. 24 from experimentally-derived reference values, for a subset of 8 biradicals.	98

Figure 5.3	Benzyne isomers and polyacenes studied in this work.	100
Figure 5.4	Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements. Dotted black lines represent the reported uncertainty of the experimental measurements.	103
Figure 5.5	Structures of the anthracene derivatives. Top row, from left to right: DCA, DPA, TIPS. Middle row: CF ₃ , CN-UpDnMethyl, CN-DiagMethyl. Bottom row: DMA, OMe	108
Figure 5.6	Deviations [kcal/mol] of ph-AFQMC predictions with single-determinant trial wavefunctions from experimentally-derived reference values. Error bars show the statistical error of the QMC measurements.	111
Figure 5.7	Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements. Dotted black lines represent the reported uncertainty of the experimental measurements.	114
Figure 5.8	Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements.	115

List of Tables

Table 2.1	Experimental IPs and the deviations of various calculated results (theory - experiment) for atoms in the G2 Test Set in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. QMC calculations using single-determinant trial functions (phaseless and FP) use 5040 walkers per repeat, while those using multi-determinant trial functions use 1056 walkers per repeat.	25
Table 2.2	Experimental EAs and the deviations of various calculated results (theory - experiment) for atoms in the G2 Test Set in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. QMC calculations using single-determinant trial functions (phaseless and FP) use 5040 walkers per repeat, while those using multi-determinant trial functions use 1056 walkers per repeat.	26
Table 2.3	Adiabatic Reaction Energies for Methanol. Experimental values and the deviations of the correlated sampling-based ph-AFQMC and G2 results (theory - experiment) in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. All QMC calculations use 192 walkers per repeat.	29
Table 3.1	Efficiency metrics of the GPU kernels involved in our energy algorithm. For each kernel we show the percentage of total run-time spent in that kernel, compute utilization as a percentage of peak compute performance, the number of registers per thread, memory utilization as a percentage of peak bandwidth (shown only for the memory type exhibiting the highest utilization), and the occupancy, i.e. the percentage of available warps (a group of 32 threads) that are active.	42

Table 3.2	Propagation times (in seconds) for an H_{50} chain with a varying number of determinants that comprise the trial wavefunction. We use CD with a 10^{-5} cutoff, and show the speed-up of a single GPU in sp over a single CPU in dp.	46
Table 3.3	Total electronic energies [Ha] of H_{50} at $R = 1.8$ Bohr in the cc-pVDZ basis. Propagation times [hours] are presented using 8 CPU/GPU pairs. We use 1000 walkers propagated for a length of $200 H a^{-1}$ (including equilibration).	48
Table 3.4	Accuracy of DF and CDs with various cutoffs for the deprotonation energy of methanol. Sp is used, and long imaginary-time trajectories are stabilized with PC. N_{AFs} denotes the resulting number of auxiliary-fields.	49
Table 3.5	Target electron configurations and spin-multiplicities ($2S + 1$), from Refs. 65 and 204.	49
Table 3.6	Comparison of CBS IPs [eV] for Co, Ni, Cu, and Zn with $\Delta\tau \rightarrow 0$ vs $\Delta\tau = 0.005 H a^{-1}$ computed in sp with ph-AFQMC/PC.	50
Table 3.7	Calculated ph-AFQMC IPs [eV] in the CBS limit computed with sp and $\Delta\tau = 0.005 H a^{-1}$, compared with experimental and CCSD(T) values. Experimental IPs have spin-orbit contributions removed.	51
Table 3.8	Same as Table 3.7, but for atoms in the right-half of the row.	51
Table 3.9	Number of active electrons and orbitals in the CASSCF trial wavefunctions for the cation/neutral species, and the number of determinants kept in the ph-AFQMC trial function accounting for 99.5% of the CI weight. For all species in this table the 3p electrons are active.	51
Table 3.10	Same as Table 3.9, but for Cu and Zn with 99.0% of the CI weight retained.	52
Table 3.11	Comparison of ph-AFQMC IPs for Ni, as obtained with CS and regular ph-AFQMC/PC, with CCSD(T) in triple- and quadruple- zeta basis sets and in the CBS limit[204]. The CCSD(T) values were obtained with a composite method, namely cc-pVxZ-DK results plus a core-valence correction, which is the difference in the cc-pwCVxZ-DK basis of CCSD(T) calculations with active spaces defined by $3s3p3d4s$ and $3d4s$ orbitals (x=T,Q). QMC results used the aug-cc-pwCVxZ-DK basis sets.	55

Table 3.12	Total propagation times [hours] required to produce the final all-electron ph-AFQMC IPs in this work.	55
Table 3.13	Speed-ups corresponding to the timings in Table 3.12. All CPU/GPU calculations use dp/sp respectively.	56
Table 4.1	Electronic States, Active Spaces, and Bond Distances used in our ph-AFQMC calculations for diatomics containing Sc through Mn. X/Y means that both active space configurations produced statistically equivalent results. The number in parenthesis is the experimental bond length.	66
Table 4.2	Electronic States, Active Spaces, and Bond Distances used in our ph-AFQMC calculations for diatomics containing Fe through Zn. X/Y means that both active space configurations produced statistically equivalent results. The number in parenthesis is the experimental bond length.	67
Table 4.3	Mean Absolute Error and Maximum Error on D_e shown for AFQMC, CCSD(T), icMRCCSD(T), and DFT methods vs. the experiments selected in Ref. 35 and, when possible, Ref. 237. For reasons justified in the text, we omit VH, CrO, and ZnS from the comparative statistical analysis. In all, our test set contains 41 diatomics. All DFT calculations are in the aug-cc-pVQZ basis with DKH corrections. DFT and CC values taken from Ref. 35. All values are in kcal/mol.	83
Table 5.1	$\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for the 13 small molecule biradical set.	96
Table 5.2	Mean signed, absolute, and maximum errors [kcal/mol] of the theoretical methods shown in Fig. 5.1 for the 13 small molecule biradical set. Sorted by MAE value. Parenthesis indicate statistical errors.	97
Table 5.3	Mean signed, absolute, and maximum errors [kcal/mol] of ph-AFQMC results and other methods for the 8 molecule biradical subset shown in Fig. 5.2. Sorted by MAE value. Parenthesis indicate statistical errors.	99

Table 5.4	$\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for ortho- meta- and para- benzyne molecules.	101
Table 5.5	ST gaps [kcal/mol] for the ortho- meta- and para- benzyne isomers. Parenthesis indicate statistical errors.	102
Table 5.6	$\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for polyacenes $n = 2 - 5$	105
Table 5.7	ST gaps [kcal/mol] for the polyacenes $n=2-5$. Square brackets indicate alternate experiments, and parenthesis indicate statistical errors.	106
Table 5.8	ST gaps for selected anthracene derivatives	109
Table 5.9	Mean signed, absolute, and maximum errors [kcal/mol] of the theoretical methods shown in Fig. 5.6 for the 13 small molecule biradical set. Sorted by MAE value.	111
Table 5.10	Mean signed, absolute, and maximum errors [kcal/mol] of ph-AFQMC results and other methods for the 8 molecule biradical subset shown in Fig. 2 of the main manuscript. Sorted by MAE value.	112
Table 5.11	ST gaps [kcal/mol] for the ortho- meta- and para- benzyne isomers.	113
Table 5.12	Total energies [Ha] for naphthalene with various trial wavefunctions, in the cc-pVTZ basis.	114
Table 5.13	ST gaps [kcal/mol] for the polyacenes $n=2-5$	115

Acknowledgments

I am sincerely grateful for my advisors, Profs. Richard Friesner and David Reichman, for their generous guidance and support throughout my graduate studies. Through their mentorship and example, I have learned how to think critically and creatively about scientific endeavors, the importance of both collaboration and a skeptical spirit of self-reliance, that ripe scientific opportunities can be found at the intersection of different fields, to make good estimates and to always proceed with both audacity and rationality.

I thank Prof. Shiwei Zhang for generously sharing his advice and expertise in AFQMC from the beginning of my graduate studies, and acting essentially as a third advisor to me. I would also like to thank the other members of my committee, Profs. Bruce Berne and Timothy Berkelbach.

The Simons Foundation has provided, through annual summer schools and conferences, many opportunities to forge collaborations and broaden my knowledge of computational approaches to solving the quantum many-body problem.

I thank my collaborators - Profs. Brenda Rubenstein, Garnet Chan, Kenneth Jordan, and Joseph Francisco - for broadening my knowledge of interesting problems in chemistry, and for many insightful discussions.

Scientific interactions with experimental chemists at Columbia have led to a number of fruitful collaborations, and moreover have deepened my understanding (and awe) for the intuition and resourcefulness that abound in the field of chemistry. I would, in particular, like to thank Zack Strater, Andrew Pun, Michael Rauch, and Profs. Tristan Lambert and Tomislav Rovis.

I am indebted to Fengjie Ma for his generous help in the early stages of my graduate studies. My own implementation of AFQMC would not have been possible without side-by-side comparisons

with results from the code of Prof. Zhang and co-workers, and also the many pedagogical discussions I had with Hao Shi.

I thank Qiming Sun for his help with PySCF, Joonho Lee for his perspective and deep understanding of electronic structure theory, especially for discussions on the singlet-triplet gaps of biradicals.

I am grateful to the Reichman group for providing a stimulating intellectual environment, and especially for teaching me about the importance of giving good talks. I thank Roel Tempelaar for the many lunches and espressos, for providing an example of successful experimental collaborations, and introducing me to the field of optical upconversion.

I appreciate my labmates in the Friesner group - Ben Rudsteyn, Jack Weber, and Dilek Coskun - who have, in the mid to late acts of my time at Columbia, really become part of the day-to-day AFQMC team, which together has accomplished so much more than I ever could have alone.

I am thankful for the Bunsen Burners - our practices, games (IM champions, Spring 2019!), and banquets; for Tom Gold Dance, the Columbia Ballet Collaborative, and Silas Farley who have given me artistic opportunities that I will treasure for a lifetime; for my friends - Alyssa, Chelsy, Christine, Chuck, E-Dean, Elinor, Greta, Joanna, Jyoti, Lilia, Liz, Mary-Liz, Natalie, Pierre, Shoshana, Stephanie, Tianchuan - with whom I've shared so many wonderful moments.

Mario Motta has been a cherished friend and academic colleague, to whom I could always turn for support. His flowering generosity in both scientific and artistic endeavors is greatly appreciated.

I owe much of this work, and many good laughs, to Evan Arthur. I am grateful for his kind friendship, GPU expertise, and trusty company in our (numerous) bug-squashing adventures.

And to my dear friend Ian, who has been there from day one, with a benevolent earnestness, contagious wonder and curiosity. I could ask him anything, and he'd entertain the idea with unwaivering sincerity and thoughtfulness. Ours is a friendship in which science and smiles, bike rides and fish-bagels, the profound and wonderfully ridiculous coexist... and I am so glad to have experienced so much of what is good in life together.

Finally, words cannot describe the gratitude in my heart for my love, Mariana; my brother, Kevin, and my mom and dad. I love you, and am so thankful for your endless love and support.

Chapter 1

Introduction

A central goal in the field of quantum chemistry is to solve the Schrödinger equation

$$i\hbar\frac{\partial}{\partial t}\Psi = \hat{H}\Psi \tag{1.1}$$

for a realistic chemical system. Typically the Born-Oppenheimer approximation[1, 2] must be invoked, in which the wavefunction is split into an electronic and nuclear part, with the former depending parametrically on the nuclear coordinates. As the nuclei are much more massive than the electrons, the kinetic energy of the nuclei is ignored (and the nuclear-nuclear repulsion is simply a constant). The resulting Hamiltonian contains the electronic kinetic energy and the electron-nuclear and electron-electron Coulomb interactions:

$$\hat{H}_e = \hat{T}_e + \hat{V}_{en} + \hat{V}_{ee}. \tag{1.2}$$

The electronic structure community seeks an *ab initio* solution of the corresponding electronic Schrödinger equation

$$\hat{H}_e\psi = E_e\psi. \tag{1.3}$$

Of particular interest, and arguably most relevant to chemical reactivity, are the ground and low-lying eigenvalues and eigenstates of \hat{H}_e . Ground-states are especially important due to the broad

utility of predicting thermochemical properties, e.g. reaction equilibria, oxidation and reduction potentials, bond dissociation energies, and pKa’s. In addition, the free-energy at the nuclear geometry corresponding to a transition state yields information about the kinetics of a chemical reaction. Low-lying excited states, in turn, can be readily populated via excitations due to light or thermal energy, and therefore the determination of their relative energetics plays a key role in the mechanistic elucidation of a host of dynamical processes.

However, exact solutions of the so-called “electronic structure problem” can be feasibly obtained only for a very small number of electrons, even after having narrowed the quantum mechanical scope in light of the Born-Oppenheimer approximation. The scientific community has for decades been in pursuit of an *ab initio* quantum chemical method that is both accurate for a wide variety of chemical phenomena and scalable to systems with thousands of electrons. However, such a general method has proven elusive (though much progress has been made). The difficulty stems from the exponential complexity of the quantum many-body problem. The space of Slater determinants, i.e. antisymmetrized products of occupied single-particle orbitals, grows exponentially with system size, and thus, e.g., diagonalization of H_e in such a space quickly becomes intractable.

Indeed, if a single Slater determinant provides a good description of the total wavefunction, the above difficulty can be avoided, and large chemical systems can be investigated with reasonable accuracy via independent particle approximations. Hartree-Fock (HF) theory[3] variationally minimizes the energy with respect to the orbital coefficients of a single Slater determinant. Density functional theory (DFT) [4, 5], while formally exact, in practice yields unsystematic accuracy due to the fact that the exchange-correlation energy functional is unknown and therefore must be approximated. These types of theories have formed the cornerstone of computational quantum chemistry, and have been utilized to great effect, with both astonishing successes and failures.

There are myriad instances when a single determinant is insufficient to properly describe the interacting electrons in realistic chemical systems. In these cases, the energy that is not captured by HF theory, known as the “correlation energy,” is a large percentage of the exact total energy. “Electron correlation” is frequently used rather imprecisely in the electronic structure community; in this work, it will refer simply to a quantum mechanical phenomenon which must be described

by a superposition of many-particle states with *more than one* significant coefficient. Electron correlation is often classified into static and dynamic correlation. While formally indistinguishable from a physical point of view, in conventional usage the former is said to describe situations in which there are multiple degenerate or nearly-degenerate states that contribute to the wavefunction, as encountered, e.g., in transition metal atoms with partial occupancy of five degenerate *d* orbitals, or biradicals with two relevant electronic configurations. Dynamic correlation is encountered when the inclusion of electronic configurations which represent excitations from occupied orbitals of a reference state to (often high-lying) virtual orbitals is necessary to recover substantial amounts of correlation energy. The accurate description of inter-electron repulsion in accordance with the Pauli exclusion principle, most severely encountered in regions near the electron cusps, can require configurations with orbitals of high energy and angular momentum.

Strongly correlated electrons abound in realistic molecules and materials, and in fact allow for the possibility of very rich and perhaps exotic chemical reactivity. In this thesis, we will focus on the electronic properties of transition metals and biradicals, for which an accurate description of strong correlation is essential in predicting properties such as ionization and bond-dissociation energies, and the energy gaps separating the lowest-lying singlet and triplet states.

A predictive understanding of the fundamental electronic properties of transition metal containing systems would constitute a huge advance in fields such as drug discovery, metalloenzymes, energy conversion, and chemical catalysis. Consider, e.g., the case of cisplatin - a small molecule with a platinum center, FDA-approved in 1978 as a prominent antitumor agent whose binding to guanine bases results in cellular apoptosis[6, 7]. The ability to understand its binding specificity and mechanism of action, quantum-mechanically, would enable the rational design of even more effective drug molecules. Also, roughly half of all enzymes rely on the presence of a metal ion to properly function[8], and thus a predictive tool to accurately model metal-ligand interactions would be a critical advance in drug design[9]. Molecular dynamics simulations, which are increasingly being used to predict relative binding free-energies of drug-like molecules[10], rely on classical force fields to reproduce the quantum mechanical interactions involving electrons and nuclei. However, bond-breaking events and especially interactions involving transition metals remain a challenging research

field[11], and stand to benefit from larger data sets of accurate quantum mechanical data. Metalloproteins are another important class of systems for which accurate computational insights would be greatly beneficial, given the limitations and scarcity of reliable experimental techniques. Much remains to be understood, e.g., about the biological mechanisms of metalloporphyrins[12, 13, 14], the redox properties of iron-sulfur clusters[15] in nitrogenase enzymes and the electron transport chain, and the structure and water-splitting mechanism of the Oxygen-Evolving Complex (OEC) in Photosystem II[16, 17, 18]. Finally, catalysts for, e.g., sensitization in optical upconversion[19] and CO₂ reduction typically involve transition metals[20]; and synthetic metal clusters[21], with the cubane motif in common with the OEC, are used as building blocks for superatomic constructions which exhibit tunable redox and magnetic properties[22].

Another domain of strong correlation that we will focus on in this thesis involves molecules with biradical character. We define a biradical as a system in which the frontier electronic structure is characterized by two electrons in two degenerate or nearly-degenerate orbitals. Note this is different than the typical definition known to most chemists, i.e. when two spatially distinct orbitals are singly occupied (as in all triplets). Interestingly, the present definition of biradical character encompasses homolytic single-bond breaking products, and closed-shell systems with a low-lying excited state, which is the case for most carbenes. In general, the singlet states of such systems are poorly described by a single Slater determinant, since one can show in the two electron two orbital model[23, 24] that all three possible singlet states, after configuration interaction, are necessarily a superposition of two electronic configurations. This is also the case for one of the three triplet states, however this state is typically not encountered due to the fact that a specified spin multiplicity puts a constraint on the value of S_z in most quantum chemical methods (generalized HF theory is a notable exception). An accurate description of the strong correlation in singlet biradicals is a prerequisite for the prediction of, e.g., singlet-triplet gaps in long polyacene molecules and other potential photocatalysts, and also in understanding the important roles of, e.g., singlet oxygen in biology[25].

The pursuit of chemical accuracy (as defined by errors not exceeding 1 kcal/mol) in the *ab initio* computation of the energetic properties of *generic* chemical systems, including the strongly

correlated ones introduced above, is a long-standing goal that has yet to be reached[26, 27, 28, 29, 30]. As mentioned above, the formally exact procedure involving direct diagonalization of the Hamiltonian in the complete space of determinants, known as Full Configuration Interaction (FCI)[3], scales exponentially with system size. Coupled cluster theory with single, double, and perturbative triple excitations (CCSD(T))[31, 32], commonly known as the “gold-standard” of quantum chemistry, scales with the seventh power of the system size. Despite many advantageous properties, chief among them size-extensivity and a wavefunction ansatz that efficiently describes dynamic correlations, CCSD(T) can break down for systems which exhibit strong multi-reference character[33, 34, 35], as electronic configurations which cannot be obtained via a truncated set of excitations from a single reference state are not represented in the wavefunction. This limitation is shared, too, by methods such as second-order Møller-Plesset Perturbation Theory (MP2)[36] and truncated CI expansions.

In contrast, there are a number of approaches designed to target the multi-reference nature of strongly correlated systems. Methods such as Complete Active Space Self-Consistent Field (CASSCF) [37, 38, 39, 40] supplemented with second order perturbation theory[41, 42], and stochastic variants such as selected CI approaches[43], have been shown to yield benchmark-quality results even for strongly correlated electronic systems, including those containing transition metals[44, 45]. However, while a judicious choice of the active space can make such calculations tractable for small systems, exponential scaling will inevitably prohibit the use of these methods to studying most realistic systems of interest in biology, materials science, and chemical catalysis. As a result, with the exception of isolated studies using specialized, albeit still approximate, wavefunction methods[46, 15, 47], these systems can only be investigated feasibly with less accurate but more economical approaches such as DFT, and the importance of having a systematically improvable *ab initio* method which is both accurate and feasible becomes readily apparent.

Quantum Monte Carlo (QMC) methods present an alternative approach[48] to obtaining approximate solutions to the Schrödinger equation, and in general rely on a stochastic sampling of electronic configurations from a probability distribution function to average observables such as the ground-state energy. Various QMC methods are reviewed in detail in Refs. 49, 50, 51, 52, and

53. To set the stage for the work presented in this thesis, we first illustrate key concepts by briefly describing the predominant QMC approaches utilized in the chemistry community.

We will refer to the nonrelativistic Born-Oppenheimer Hamiltonian

$$\hat{H} = -\frac{1}{2} \sum_i \nabla_i^2 - \sum_i \sum_\alpha \frac{Z_\alpha}{|R_i - d_\alpha|} + \frac{1}{2} \sum_i \sum_{j \neq i} \frac{1}{|R_i - R_j|}, \quad (1.4)$$

where R_i and d_α are electronic and nuclear positions, respectively, and Z_α are nuclear charges. The Variational Monte Carlo (VMC) method is the simplest QMC approach[54], which computes the following energy expression

$$E_T = \frac{\int \phi_T^*(R) \hat{H} \phi_T(R) dR}{\int \phi_T^*(R) \phi_T(R) dR} = \int \frac{\hat{H} \phi_T(R)}{\phi_T(R)} P(R) dR > E_0, \quad (1.5)$$

where $P(R) = \frac{|\phi_T(R)|^2}{\int |\phi_T(R)|^2 dR}$. Sampling Q configurations from $P(R)$, Eq. (1.5) can be evaluated via Monte Carlo as

$$\frac{1}{Q} \sum_i^Q \frac{\hat{H} \phi_T(R)}{\phi_T(R)} \Big|_{R=R_i}. \quad (1.6)$$

The resulting variational parameterization of ϕ_T is typically utilized as the guiding, or trial, wavefunction for subsequent projector approaches.

In such approaches, the substitution $\tau = -it$ (and assuming atomic units, i.e. $\hbar = 1$) transforms Eq. (1.1) to

$$-\frac{\partial}{\partial \tau} \Psi = \hat{H} \Psi. \quad (1.7)$$

When the Hamiltonian is shifted by a constant, E_0 , equal to the exact ground-state energy, the general solution to Eq. (1.7) is

$$\Psi = \sum_n a_n \psi_n e^{\tau(E_0 - E_n)}, \quad (1.8)$$

where $\hat{H} \psi_n = E_n \psi_n$. At large values of τ , the contributions to Ψ from the excited eigenstates, ψ_n ($n > 0$), will decay exponentially, and thus only the ground-state of the Hamiltonian persists. That is, given $\langle \Psi_0 | \Psi \rangle \neq 0$,

$$\lim_{\tau \rightarrow \infty} e^{\tau(E_0 - \hat{H})} |\Psi\rangle \rightarrow |\psi_0\rangle. \quad (1.9)$$

In Green's Function Monte Carlo (GFMC), Eq. (1.9) is realized in coordinate space by iterating the following integral equation

$$\Psi^{n+1}(R) = E_0 \int G(R, R') \Psi^n(R') dR', \quad (1.10)$$

where $\hat{H}G(R, R') = \delta(R - R')[55, 56, 57]$.

Alternatively, in the most widely-used variant of QMC, Diffusion Monte Carlo (DMC), the Hamiltonian and wavefunction are also represented in real-space, but the imaginary-time propagator is discretized such that the ground-state projection takes the form

$$\lim_{N \rightarrow \infty} (e^{\Delta\tau(E_0 - \hat{H})})^N |\Psi\rangle \rightarrow |\psi_0\rangle. \quad (1.11)$$

This is equivalent to iterating

$$\Psi(R, \tau + \Delta\tau) = \int G(R, R', \Delta\tau) \Psi(R', \tau) dR'. \quad (1.12)$$

The choice of a small imaginary-time step enables the use of the Trotter-Suzuki decomposition[58, 59], which yields an approximate form of the propagator

$$G(R, R', \Delta\tau) \sim \exp\left[-\frac{(R - R')^2}{2\Delta\tau}\right] \exp[\Delta\tau(2E_0 - V(R) - V(R'))/2] \quad (1.13)$$

with an error that grows as $O(\Delta\tau^3)$. Thus, in this basic formulation of DMC,

$$\Psi(R', \tau) \rightarrow \sum_Q \delta(R' - R_Q), \quad (1.14)$$

and Eq. (1.12) becomes

$$\Psi(R, \tau + \Delta\tau) \sim \sum_Q \int G(R, R', \Delta\tau) \delta(R' - R_Q) dR' \quad (1.15)$$

$$= \sum_Q G(R, R_Q, \Delta\tau) \quad (1.16)$$

$$= \sum_Q \exp\left[-\frac{(R - R_Q)^2}{2\Delta\tau}\right] \exp\left[\Delta\tau(2E_0 - V(R) - V(R_Q))/2\right]. \quad (1.17)$$

For each Q , a new position is sampled from the Gaussian, and the value of the second exponential determines the weight of the walker, and whether or not it is replicated or eliminated via a so-called birth/death algorithm[53].

Ref. 53 also details an importance sampling algorithm, in which the introduction of a trial wavefunction defines a new function $f(R, \tau) = \Psi(R, \tau)\phi_T(R)$, such that

$$f(R, \tau + \Delta\tau) = \int \tilde{G}(R, R', \Delta\tau) f(R', \tau) dR', \quad (1.18)$$

where $\tilde{G}(R, R', \Delta\tau) \sim \exp\left[\frac{-(R - R' - \Delta\tau v_D(R'))^2}{2\Delta\tau}\right] \exp\left[\Delta\tau(2E_0 - E_L(R) - E_L(R'))/2\right]$. The drift velocity is defined as $v_D(R) = \phi_T(R)^{-1} \nabla \phi_T(R)$, which guides the walkers toward increasing $|\phi_T|$. The walker weights are now expressed in terms of a quantity called the “local energy,” $E_L(R) = \phi_T(R)^{-1} \hat{H} \phi_T(R)$. This greatly reduces the statistical fluctuations in the random walks assuming the use of a trial wavefunction that is close to the true ground-state wavefunction, since the resulting E_L are typically much more stable than sampled values of the potential energy surface.

DMC calculations typically employ the fixed-node constraint[60], in which the random walks are restricted to the space defined by $\langle \phi_T | \Psi \rangle > 0$. Sampling thus in coordinate-space, in the presence of complicated nodal surfaces, can result in large fluctuations especially due to the core electrons, which must be treated with an uncontrolled approximation, e.g. most commonly with pseudopotentials[53]. The fixed-node constraint in DMC is expected to result in a bias more severe than that incurred in alternative QMC approaches that sample in the space of determinants, which by construction respect fermionic antisymmetry.

One such method is known as FCI-QMC, which has undergone rapid development in recent years[61, 62, 63, 64, 65]. In this method, the propagator is approximated as $1 - \tau \hat{H}$. Despite its high level of accuracy, the method is limited by the initiator bias; also, despite a much-reduced prefactor (relative to, e.g., FCI) FCI-QMC still inevitably suffers from exponential scaling with system size.

Auxiliary-field QMC (AFQMC) is another projector method that involves Monte Carlo sampling in determinant space. As will be detailed in the next section, the imaginary-time propagator is transformed such that the interacting system is exactly mapped onto a sum of non-interacting systems in auxiliary-fields. In the “phaseless” variant (ph-AFQMC)[66, 67, 68] a constraint is introduced to control statistical noise and to preserve fourth power (or less[69]) computational scaling with system size for molecular systems, yet a salient feature is that the resulting bias can, in principle, be systematically reduced. First introduced and developed in the context of model systems for strongly correlated condensed matter systems, ph-AFQMC has produced state-of-the-art results in the context of electronic structure theory[70, 71, 72, 73, 74, 75, 76, 77, 78, 79]. In particular, the methodology shows great promise in the accurate description of a variety of strongly correlated electronic systems[80, 81, 82], including transition metal-containing species[83, 84, 85, 86, 33, 87, 88, 78, 89].

This thesis will present two major methodological developments aimed at reducing the wall-time prefactor of ph-AFQMC calculations - a correlated sampling approach to accelerate (and improve the accuracy of) calculations of energy differences, and an efficient implementation on graphical processing units. These advances have enabled us to obtain robust results and predictions for a host of chemically relevant molecular systems, with particular focus on the strongly correlated regime.

1.1 Overview of AFQMC - Exact Method and the Phaseless Constraint

In AFQMC the general electronic Hamiltonian is written in second-quantized form:

$$\hat{H} = \sum_{ij}^M T_{ij} \sum_{\sigma} c_{i\sigma}^{\dagger} c_{j\sigma} + \frac{1}{2} \sum_{ijkl}^M V_{ijkl} \sum_{\sigma, \tau} c_{i\sigma}^{\dagger} c_{j\tau}^{\dagger} c_{l\tau} c_{k\sigma}, \quad (1.19)$$

where M is the size of the orthonormal one-particle basis, and $c_{i\sigma}^{\dagger}$ and $c_{i\sigma}$ are the fermionic creation and annihilation operators with particle and spin labels. The two-body matrix elements, V_{ijkl} , can be expressed in terms of Cholesky vectors as $V_{ijkl} = \sum_{\alpha} L_{ik}^{\alpha} L_{jl}^{\alpha}$. [77] Defining the one-body operator $\hat{v}_{\alpha} \equiv i \sum_{ik} L_{ik}^{\alpha} \sum_{\sigma} c_{i\sigma}^{\dagger} c_{k\sigma}$, and subtracting the expectation value with respect to the trial wavefunction $\langle \hat{v}_{\alpha} \rangle$ from \hat{v}_{α} , the Hamiltonian can be written as the sum of all one-body operators, \hat{H}_1 , plus the following two-body operator

$$\hat{H}_2 = -\frac{1}{2} \sum_{\alpha} (\hat{v}_{\alpha} - \langle \hat{v}_{\alpha} \rangle)^2. \quad (1.20)$$

Use of the Trotter-Suzuki decomposition gives

$$e^{-\Delta\tau\hat{H}} = e^{-\Delta\tau\hat{H}_1/2} e^{-\Delta\tau\hat{H}_2} e^{-\Delta\tau\hat{H}_1/2} + O(\Delta\tau^3). \quad (1.21)$$

The exponential terms involving \hat{H}_2 may be decomposed using a Hubbard-Stratonovich (HS) transformation [90, 91]

$$e^{\frac{1}{2}\Delta\tau(\hat{v}_{\alpha}-\langle\hat{v}_{\alpha}\rangle)^2} = \int_{-\infty}^{\infty} dx_{\alpha} \left(\frac{e^{-\frac{1}{2}x_{\alpha}^2}}{\sqrt{2\pi}} \right) e^{\sqrt{\Delta\tau}x_{\alpha}(\hat{v}_{\alpha}-\langle\hat{v}_{\alpha}\rangle)}, \quad (1.22)$$

which expresses the exponential of a two-body operator as the exponential of a one-body operator integrated over auxiliary-fields (AFs). This transformation allows for practical propagation in terms of the Thouless theorem, which states that the application of an exponential of a one-body operator on a Slater determinant produces another Slater determinant [92], which can be implemented via a

simple matrix multiplication[93, 94].

The propagator (1.21) now takes the form of a multi-dimensional integral

$$e^{-\Delta\tau\hat{H}} = \int d\mathbf{x} P(\mathbf{x}) \hat{B}(\mathbf{x}), \quad (1.23)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_\alpha)$, $P(\mathbf{x})$ is a normal distribution with unit variance, and

$$\hat{B}(\mathbf{x}) = e^{-\Delta\tau\hat{H}_1} e^{\sqrt{\Delta\tau}\mathbf{x} \cdot (\hat{\mathbf{v}} - \langle \hat{\mathbf{v}} \rangle)} e^{-\Delta\tau\hat{H}_1}. \quad (1.24)$$

The integral in (1.23) may be approximated using a Monte Carlo scheme, with walkers whose propagation in the space of Slater determinants is guided by the complex importance function $\langle \phi_T | \phi \rangle$ which is proportional to the walker weights. The representation of the total wavefunction is thus a weighted sum over walker determinants

$$|\Phi\rangle = \sum_k \frac{w_k |\phi_k\rangle}{\langle \phi_T | \phi_k \rangle}, \quad (1.25)$$

yielding essentially a multi-reference description. The energy is calculated at intervals using the mixed-estimator

$$\frac{\langle \phi_T | \hat{H} | \Phi \rangle}{\langle \phi_T | \Phi \rangle} = \frac{\sum_k w_k E_L(\phi_k)}{\sum_k w_k}, \quad (1.26)$$

where the “local energy” is given by $E_L(\phi_k) = \frac{\langle \phi_T | \hat{H} | \phi_k \rangle}{\langle \phi_T | \phi_k \rangle}$.

The method as described above is called the “Free Projection” (FP) approach[95]. While it is formally exact and can yield excellent results for small system sizes, this method suffers from the “phase problem,” which is a generalization of the Fermionic “sign problem” to the complex plane[96, 97, 98, 99]. For the standard Coulomb interaction the \hat{v}_α operators are purely imaginary, and each application of $e^{\sqrt{\Delta\tau}\mathbf{x} \cdot (\hat{\mathbf{v}} - \langle \hat{\mathbf{v}} \rangle)}$ can be thought of as a rotation of the Slater determinant $|\phi\rangle$, causing an evolution of the overlap $\langle \phi_T | \phi \rangle$ in the complex plane. Over the course of the random walk, a determinant accumulates a phase $e^{i\theta}$, and the infinitely many possible values of $\theta \in [0, 2\pi)$ result in the possibility of infinitely many indistinguishable determinants. Furthermore, over the

course of the propagation, walkers will populate the origin where $\langle \phi_T | \phi \rangle = 0$, and subsequent propagation yields only noise from signal cancellation effects and divergences in the weights and local energies.

The ph-AFQMC employs a multifaceted strategy to control this problem. The weights are initialized to a positive real constant, and after each propagation step are projected back onto the real axis, *i.e.* the rotated weights are multiplied by $\max\{0, \cos(\Delta\theta)\}$, where we have defined the phase of the overlap ratio $\Delta\theta \equiv \text{Im}\{\ln \frac{\langle \phi_T | \phi^{(\tau+1)} \rangle}{\langle \phi_T | \phi^{(\tau)} \rangle}\}$ [100]. For this phase projection to work, the AFs are shifted by a force bias (FB) $\bar{\mathbf{x}}$ [67, 101], the optimal choice of which is obtained by minimizing the fluctuations of the weights with respect to the AFs at their average value:

$$\frac{\partial}{\partial x_\alpha} \left[\frac{\langle \phi_T | e^{\sqrt{\Delta\tau}(x_\alpha - \bar{x}_\alpha)(\hat{v}_\alpha - \langle \hat{v}_\alpha \rangle)} | \phi_k \rangle}{\langle \phi_T | \phi_k \rangle} e^{-\bar{x}_\alpha^2/2 + x_\alpha \bar{x}} \right]_{x_\alpha=0} = 0. \quad (1.27)$$

This is, in essence, a stationary-phase approximation. Expanding the expression inside the brackets to $O(\sqrt{\Delta\tau})$ and taking the derivative gives

$$\bar{x}_\alpha = -\sqrt{\Delta\tau} [\bar{v}_\alpha - \langle \hat{v}_\alpha \rangle], \quad (1.28)$$

where $\bar{v}_\alpha \equiv \frac{\langle \phi_T | \hat{v}_\alpha | \phi \rangle}{\langle \phi_T | \phi \rangle}$. The introduction of the FB does *not* add any additional approximations, as the integration variable in (1.22) is merely shifted by a constant, yet it is crucial for two reasons. First, it diverges when the “nodal surface” (as defined in the complex plane of overlaps) is approached, pushing the walker away from the origin, reminiscent of the drift velocity in DMC. Second, since \hat{v}_α is complex, $\text{Im}[\bar{\mathbf{x}}]$ reduces the amount of physical information discarded in the phase projection. Similarly, the subtraction of $\langle \hat{v} \rangle$ from \hat{v} in the propagator also greatly reduces the severity of this projection, as the smaller diagonal matrix elements of the resulting propagator cause milder rotations of the phases of the orbitals[71].

The choice of FB allows the weight factor which multiplies the previous weight after a propagation step to be written as $W(\phi) = e^{-\Delta\tau E_L(\phi)}$, in analogy to DMC. The second approximation in ph-AFQMC, which is much milder than the first, takes the real part of the local energy in the

weight above:

$$E_L(\phi) \equiv \text{Re}\left\{\frac{\langle\phi_T|\hat{H}|\phi\rangle}{\langle\phi_T|\phi\rangle}\right\}. \quad (1.29)$$

The severity of the phaseless constraint in ph-AFQMC can be reduced with the use of more accurate trial wavefunctions, especially those with the correct symmetry properties.[95] This can be seen from (1.29), for when $|\phi_T\rangle = |\Phi_0\rangle$, the local energy which determines the weights and energy measurements is a real constant (equal to the exact ground state energy).

Chapter 2

Chemical transformations approaching chemical accuracy via correlated sampling in auxiliary-field quantum Monte Carlo¹

2.1 Introduction

ph-AFQMC for quantum chemical applications has traditionally required a relatively high computational cost, as the favorable wall-time scaling is masked by a large prefactor. In this chapter we present an approach to greatly reducing this prefactor which involves the use of correlated sampling for a particular class of important processes. The general idea is that for sufficiently similar systems, energy *differences* are expected to converge more rapidly, *i.e.* with smaller error bars, than total energies when the errors or statistical fluctuations in the calculations are biased in the same direction. Indeed, error cancellation is largely responsible for the success of many approximate methods such as DFT in the computation of energy differences. Correlated sampling has previously been adapted to reduce the statistical errors in QMC approaches via the use of the same

¹Based on work published in *J. Chem. Theory Comput.* 2017, 13, 6, 2667-2680.

set of configurations sampled for both the primary and secondary systems[52]. This technique is often referred to in the literature as “Differential” QMC, and the details of its implementation can vary depending on the type of QMC being used. The potential energy curves of H_2 and BH have been calculated using correlated sampling with VMC[102], and similarly that of the H_3 cation with Differential GFMC[103]. The latter method has been used to compute the dipole moment of LiH, [104] and to calculate infinitesimal energy differences from which forces and various polarizabilities have been obtained[105]. This idea has also been extended to DMC, which has been used to compute forces and potential energy surfaces for the first row diatomics[106]. Correlated sampling has also been used to calculate energy differences between ground and excited states of the same Hamiltonian, as illustrated by a VMC study of particle-hole excitations in the two-dimensional electron gas[107]. In addition, the concept has been extended to enable concerted propagation of a system with different time steps, in order to extrapolate the Trotter error in Differential DMC[108].

Correlated sampling is, in fact, also well-suited to model the energetics of myriad chemical reactions, since *only* energy differences, as opposed to total energies, are relevant. In this paper we present a novel correlated sampling-based AFQMC approach, and show that it is capable of computing ground-state energy differences corresponding to redox, deprotonation, and hydrogen abstraction reactions to a given statistical error in a fraction of the time previously required, without any loss of accuracy. Redox reactions (often involving TMs) abound, for example, in metabolic and photosynthetic processes[109, 18, 110], battery chemistry[111], and catalysis (*e.g.* CO_2 and O_2 reduction)[112, 113]. A reliable *ab initio* method to calculate deprotonation free energies would provide an improvement upon existing computational approaches to determining pKa’s and protonation states[114, 115, 116], which would have significant ramifications for drug discovery[117, 118], materials science[119, 120, 121], and the structural determination of biological complexes[17, 122]. In the context of chemical catalysis, proton removal is known to be the rate-limiting step in many important reactions, *e.g.* oxygen reduction on the surface of TiO_2 [123]. Hydrogen abstraction reactions are ubiquitous and play a major role in combustion and the oxidation of hydrocarbons[124, 125], diamond growth via chemical vapor deposition[126], biochemical processes involving *e.g.* the antioxidant vitamin E [127] and various metalloenzymes[128, 129],

industrial processes[130], and organic synthesis[131]. AFQMC would be a useful benchmark for previous *ab initio* studies[129, 132, 133, 134, 135, 136, 137, 138, 139] in predicting thermodynamic properties of this difficult class of chemical reactions. Thus the class of applications we target is large and important. We highlight the fact that the cost of our correlated sampling approach, relative to the uncorrelated method, decreases with increasing system and basis set size, opening the door to the treatment of realistic large chemical systems with correlated sampling-based AFQMC in the near future.

This chapter is organized as follows: Section 2.2.1 will present our correlated sampling approach to AFQMC. We justify this approach for modeling molecular systems in Section 2.2.2, and Section 2.2.3 will disclose further computational details. Section 2.3.1 will present calculations of the ionization potentials (IPs) and electron affinities (EAs) of the 1st row atoms in the G2 Ion Test Set[140], while Section 2.3.2 will consider adiabatic molecular properties, taking the IP, deprotonation energy, and O-H bond dissociation energy of methanol as examples. The efficacy of correlated sampling as a function of both basis set size and number of random walkers will be explored in Sections 2.3.3. The latter subsection will provide an assessment of the reduction in compute time afforded by the use of correlated sampling. In Section 2.4 we conclude, emphasizing opportunities for further gains in computational efficiency that may be possible and future targets of investigation.

2.2 Methods

2.2.1 Correlated Sampling Methods for AFQMC

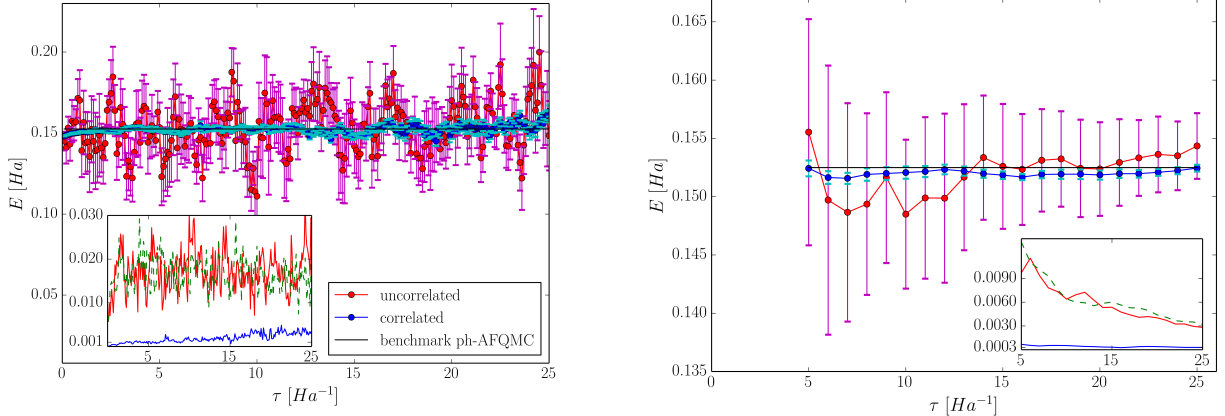
Given that the statistical error in an AFQMC calculation arises solely from the MC evaluation of the integral over AFs, a natural way to implement correlated sampling in the calculation of an energy difference is to pair the walkers of the two systems, and use the same set of AFs to propagate each pair of walkers. To be precise, the correlation is established on the level of each term in (1.20), for each value of α . This, in essence, matches Cholesky vectors of the same iteration, and becomes more effective the more similar the interactions are in the two systems.

Sampling in this way, however, requires one to relinquish optimal importance sampling, which is usually implemented via a population control (PC) scheme in which walkers with large (small) weights are replicated (stochastically purged), since performing independent PC for each system separately would quickly destroy the walker-pair correlation. Alternatively, one can implement PC with respect to the weights of the primary system *for both systems*. This, however, will only be effective if the two systems are essentially identical. In the absence of a PC scheme, the noise from the accumulation and persistence of divergent walkers inevitably grows with propagation time. We find that the immediate reduction in statistical error following the correlation of the AFs, augmented as needed by what we call the “preliminary equilibration scheme” below, allows converged averages to be obtained at short and intermediate projection times. In light of the fact that the stability of the random walks at long times, as afforded by a branching scheme, appears to be only marginally relevant when our correlated sampling approach is used, we simply choose *not* to implement PC when the AFs are correlated. In the event that a walker’s weight becomes zero or negative (in the latter case the phase projection sets the weight to zero), the walker is no longer propagated and the random number stream is updated if necessary such that the correlation of the other walker-pairs is unaffected.

The data shown in blue in Fig. 2.1 illustrates features associated with the typical correlated sampling protocol used in this work. In this example, the propagation in imaginary time (measured in Ha^{-1}) is performed using correlated AFs and repeated 11 times using different random number seeds. The mean energy difference and associated standard error at each τ point is computed among the repeats (see left plot). We then choose the imaginary-time at which the energy difference is seen to stabilize (in this case at $\tau \sim 4$), and for each repeat calculate the *cumulative* average at each τ , which represents the running average of the energy measurements taken after the end of the equilibration period up to the given value of τ . To obtain the final result, we compute the mean and standard error of the cumulative averages among the repeats (see right plot), and choose the value corresponding to the τ at which the standard error reaches a minimum (if there are multiple equivalent minima the one occurring earliest is chosen) or falls below a target error level. We note that our choice of 11 repeats is arbitrary; a larger number could be used to reduce the standard

error as necessary. Clearly in this example correlating the AFs for $\tau \geq 0$ drastically reduces the error relative to the uncorrelated runs shown in red, and the resulting IP agrees with that obtained from independent $\tau = 80$ ph-AFQMC runs of the neutral and cationic species. This benchmark, indicated by the solid black line, was found to have a negligible standard error of 0.2 mHa after employing a reblocking analysis which corrects for auto-correlation[141].

As it is often the case that a very small population size is sufficient to achieve a desired statistical error via the correlated sampling approach, we choose not to employ PC in the uncorrelated comparisons since this would result in a bias that typically goes as $1/N_{wlk}$ [142]. A second reason is that over the relatively short imaginary-time scales relevant for the correlated sampling method, PC is expected to have little effect on the uncorrelated comparisons as the walker weights usually do not stray far from unity. This is confirmed by the similarity of the error curves plotted in the insets of Fig. 2.1, corresponding to uncorrelated runs with (dotted green) and without (red) PC. The former is obtained by using a large enough population size such that the bias from the PC algorithm is negligible (360 walkers per repeat), and rescaling the resulting standard error by $\sqrt{\frac{360}{12}}$.



(a) Averaged IPs (circles) among the repeats at each τ along the imaginary-time propagation.

(b) Mean values (circles) of the cumulative averages taken for $\tau > 4$.

Figure 2.1: Comparison of correlated and uncorrelated sampling for the IP of the K atom in a 6-31+G* basis, with $\Delta\tau = 0.01$, 12 walkers per repeat, and a HF reference state. The error bars give the standard errors (the standard deviation times $\frac{1}{\sqrt{N_r}}$, where N_r is the number of repeats) of the mean values among the repeats at each τ point in (a), and of the cumulative averages in (b). These standard errors are plotted in the insets, along with the scaled standard error resulting from an uncorrelated run in which PC was used with 360 walkers per repeat (dotted green).

If one or both of the comparative systems requires a long equilibration time, which can be the

case for, *e.g.*, initial populations which are severely spin-contaminated or for strongly correlated systems in which the “guiding” trial function poorly describes the true ground state, then walker pair correlation can be lost prior to convergence, and the associated noise growth can make measurements impossible. The simplest way to overcome this problem is to use a better trial function, *e.g.* a multi-determinant CASSCF wavefunction instead of the single Hartree-Fock (HF) determinant. Alternatively, we have devised what we will refer to as the “preliminary equilibration scheme” (PES). First, one of the two systems is equilibrated using PC for the required interval, then the walkers of the secondary system are initialized with the resulting determinants of the equilibrated primary system, with weights scaled by $\frac{\langle \phi_T^{\text{secondary}} | \phi \rangle}{\langle \phi_T^{\text{primary}} | \phi \rangle}$ (and any resulting phases projected) to reflect the appropriate importance sampling. Finally, both systems are propagated using correlated sampling with their respective Hamiltonians without PC for a short period of time after which energy measurements are collected. A schematic of this procedure is shown in Fig. 2.2. This protocol is repeated with different random number seeds to obtain satisfactory statistics on the energy difference between the ground states. We note that a similar scheme was published many years ago by Traynor et al[103].

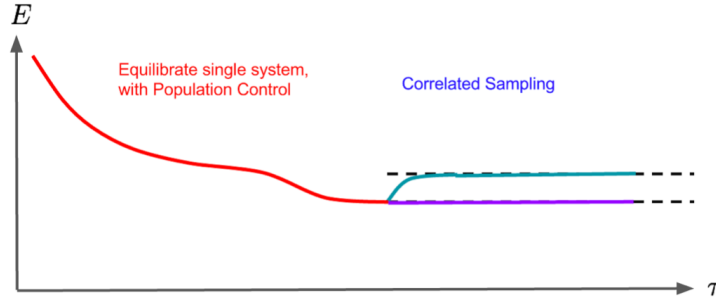


Figure 2.2: Schematic of the PES version of correlated sampling. Preliminary equilibration of the primary system is shown in red. This phase can be as long as necessary due to the use of PC. The resulting walkers are used to initialize the secondary system, after which the two systems are propagated with their own Hamiltonians using correlated sampling. Equilibration of the secondary system is relatively rapid, allowing measurements to be taken before the noise growth becomes prohibitive.

Inspection of the ph-AFQMC propagator

$$\hat{B}(\mathbf{x} - \bar{\mathbf{x}}) = e^{-\Delta\tau\hat{H}_1/2} e^{\sqrt{\Delta\tau}(\mathbf{x}-\bar{\mathbf{x}})(\hat{\mathbf{v}}-\langle\hat{\mathbf{v}}\rangle)} e^{-\Delta\tau\hat{H}_1/2}, \quad (2.1)$$

demonstrates that stricter approaches to correlated sampling exist. In one such alternative to correlating only the AFs, each pair of walkers is propagated using both the same AFs and FBs. However taking the simple average of the FBs of the primary and secondary systems was found to cause an early onset of noise even for marginally different systems. This issue likely arises from divergences in the local energies and weights which can occur more frequently due to the use of sub-optimal FBs, a fact which follows from the discussion centered around Eq. (1.28). Separately, we choose *not* to correlate the $\langle \hat{\mathbf{v}} \rangle$ since any gain in sampling efficiency would come at the cost of an increased bias from the phaseless constraint[71].

Finally, we note that our method of correlated sampling is not limited to the phaseless version of AFQMC, and can be used in exactly the same manner for FP calculations. In fact, using the same AFs in the latter case results in a relatively stricter form of correlated sampling, since the FB is not present in the post-HS propagator used in FP.

2.2.2 Utilizing Optimal Correlation in Molecular Applications

Having justified our choice to correlate only the AFs, we claim that maximal correlation between a pair of walkers is achieved when the primary and secondary systems use the same set of basis functions. To see this, recall the definition of the two-electron matrix elements in the Hamiltonian (1.19):

$$V_{ijkl} = \int d\mathbf{r}_1 d\mathbf{r}_2 \phi_i(\mathbf{r}_1) \phi_j(\mathbf{r}_2) \frac{1}{r_{12}} \phi_k(\mathbf{r}_1) \phi_l(\mathbf{r}_2). \quad (2.2)$$

When $\{\phi\}_{primary} = \{\phi\}_{secondary}$, the V_{ijkl} and, in turn, the one-body operators $\hat{\mathbf{v}}$ in the propagators (2.1) for both systems will be identical. This ideal condition is obviously satisfied when the two systems are an atom or molecule with the same geometry but *e.g.* different charges, since atomic basis functions such as Gaussian-type orbitals are usually centered on the positions of the nuclei. In fact, for the calculation of vertical IPs and EAs, exactly the same T_{ij} and V_{ijkl} elements are used in the imaginary-time propagation of both systems. For adiabatic redox processes, however, the ground-state geometries of the neutral and ionized species are, in general, different. In most cases the differences in the V_{ijkl} are slight, and correlated sampling is still found to be effective,

albeit with reduced efficiency. For example, Fig. 2.3 compares the correlated and uncorrelated standard errors corresponding to the vertical and adiabatic IPs of methanol. While the errors from the uncorrelated runs are roughly similar in magnitude, that from the correlated vertical case is significantly smaller and more constant compared to that from the correlated adiabatic case. This suggests the use of a two-step process to compute adiabatic energy differences, in which the fixed-geometry transition is calculated with correlated sampling-based AFQMC, and the geometry relaxation energy is obtained from a lower level of theory. Further details will be presented in Section III.B.

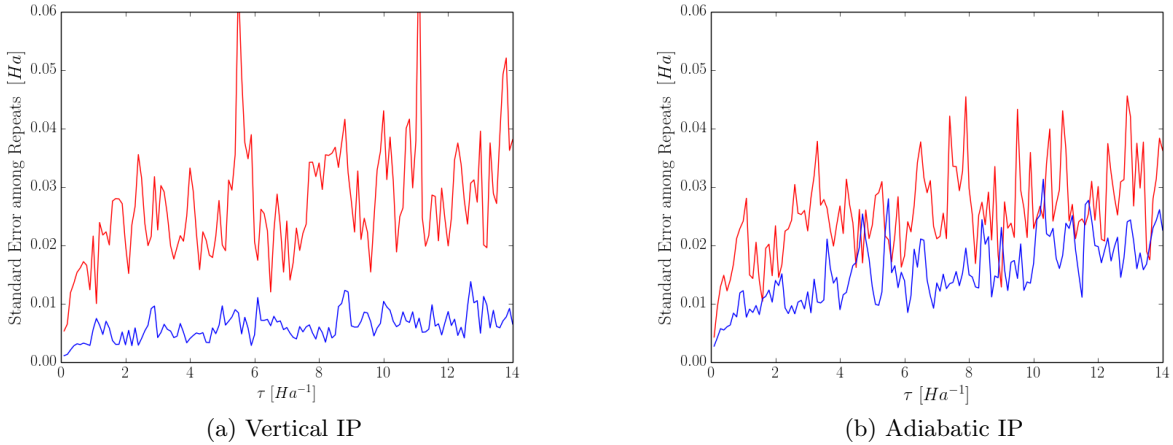


Figure 2.3: Standard errors resulting from uncorrelated (red) and correlated (blue) propagation of the repeats for methanol in the cc-pVTZ basis. $\Delta\tau = 0.01$ and 24 walkers per repeat were used.

In the same spirit, we have devised a protocol which enables the calculation of the change in energy corresponding to the removal of a proton while the rest of the geometry remains fixed, in which the optimal condition mentioned above is realized with the use of so-called “ghost” basis functions. In this scheme, the deprotonated species uses the same set of basis functions as the acid, *i.e.* the position of the removed proton still serves as a center for hydrogen basis functions but *not* as a center of nuclear charge. As a result, the number of basis functions remains the same for the primary and secondary systems as required by our correlated sampling method. Moreover the V_{ijkl} and thus the $\hat{\mathbf{v}}$ are, by construction, identical. We note that even though the T_{ij} now differ due to the altered electron-nucleus attraction terms, the statistical error in the calculated energy difference

is *not* exacerbated since it results only from the MC evaluation of the integral over AFs. Fig. 2.4 illustrates the efficacy of this procedure in a calculation of the deprotonation energy for methanol, showing a drastic reduction in the statistical errors when the AFs are correlated. The “ghost” basis

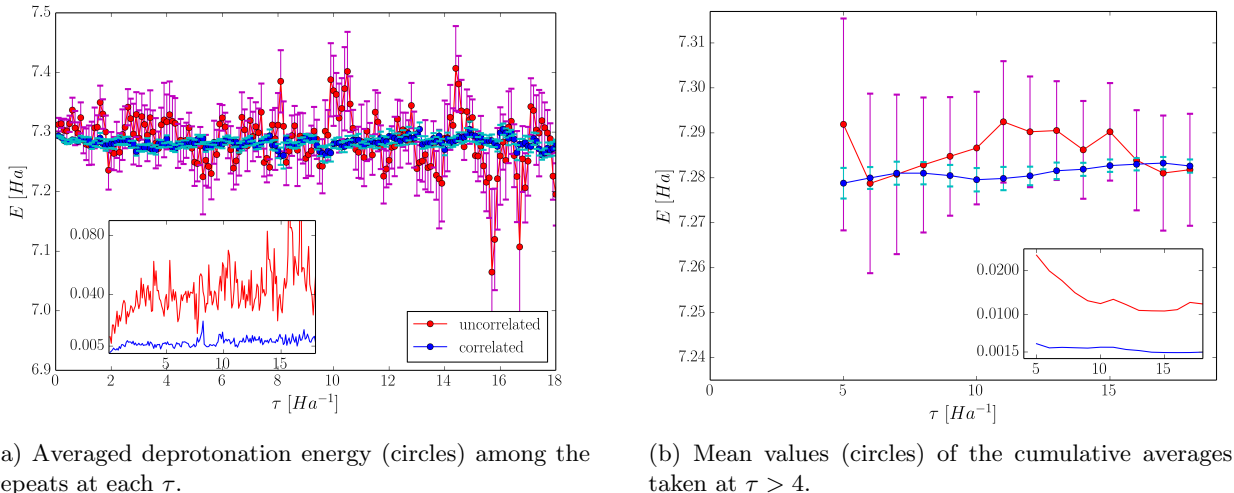


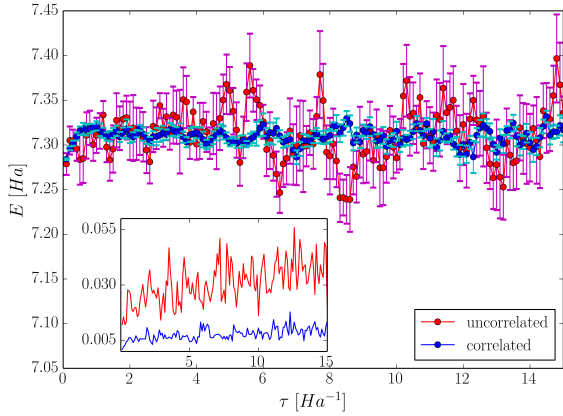
Figure 2.4: Comparison of correlated and uncorrelated sampling for the fixed-geometry deprotonation of CH₃OH in the cc-pVTZ basis with “ghost” basis functions. We use $\Delta\tau = 0.01$, 12 walkers per repeat, and a HF reference state. The error bars, plotted in the insets, in (a) represent the standard error among the repeats at each τ ; those in (b) give the standard error of the cumulative averages.

function strategy can also be directly applied effectively to hydrogen abstraction reactions. Fig. 2.5 demonstrates the error reduction afforded in a calculation of the MeOH \rightarrow MeO \cdot reaction energy.

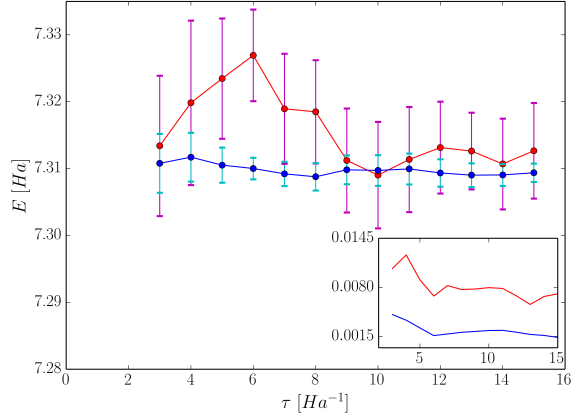
Thus we have presented a correlated sampling approach which has the potential to reduce the statistical error for redox, deprotonation, and hydrogen abstraction reactions. The extension of this protocol to obtain adiabatic energy differences will be described in Section III.B.

2.2.3 Computational Details

One-electron and overlap integrals, Cholesky vectors, restricted open-shell and unrestricted HF trial wavefunctions were all obtained from a modified version of NWChem[143, 77]. The maximum residual error in the Cholesky decomposition was chosen to be 1×10^{-6} Ha. This cutoff has a negligible effect on the QMC energies (*i.e.* orders of magnitude smaller than the statistical error bars), and is utilized to decrease the number of AFs required, which in turn decreases both the error from the random noise and the computational cost of the calculations.



(a) Averaged H abstraction energy (circles) among the repeats at each τ .



(b) Mean values (circles) of the cumulative averages taken at $\tau > 2$.

Figure 2.5: Comparison of correlated and uncorrelated sampling for the energy difference corresponding to the removal of H \cdot from the O-H bond of CH₃OH, with the rest of the geometry fixed. The cc-pVTZ basis is used, with $\Delta\tau = 0.01$, 12 walkers per repeat, and a CASSCF reference state. The error bars, plotted in the insets, in (a) represent the standard error among the repeats at each τ ; those in (b) give the standard error of the cumulative averages.

For calculations using a single-determinant trial wavefunction with unrestricted reference, spin-contamination can occur due to the presence of higher multiplicity spin configurations in the trial function, which can prolong equilibration times and reduce accuracy. We use a spin-projection technique[75] in AFQMC to minimize such effects. In this scheme, the walkers are initialized with a restricted open-shell HF determinant, which is an eigenfunction of \hat{S}^2 , such that propagation with $e^{\sqrt{\Delta\tau}\mathbf{x}\cdot(\hat{\mathbf{v}}-\langle\hat{\mathbf{v}}\rangle)}$ preserves spin symmetry. The unrestricted HF determinant, while *not* an eigenfunction of the total spin operator, is known in most cases to provide a better description of the ground-state energy, and is therefore used to implement the phaseless constraint and to estimate the energy.

Multi-determinant trials were obtained from CASSCF calculations performed with PySCF[144]. The resulting expansions of the wavefunctions in CI space are truncated such that determinants associated with coefficients below a specified threshold are discarded. In what follows, CASSCF(X,Y) will denote an active space with X electrons and Y orbitals; for cases in which an energy difference is calculated we use the notation $X \equiv N_{neutral}/N_{ion}$ and $Y \equiv M_{neutral}/M_{ion}$ to denote the numbers of electrons and orbitals in the neutral and charged systems, respectively.

In all AFQMC calculations the spin -up and -down sectors of the walker determinants are sep-

arately orthonormalized via a modified Gram-Schmidt procedure[145] every 0.05 Ha^{-1} to preserve the anti-symmetry of the walker determinants. When required, we use a PC algorithm in which the number of walkers is fixed throughout the entire calculation[142]. PC (when used) and energy measurements are performed every 0.1 Ha^{-1} .

Use of the “hybrid” formulation[76] of ph-AFQMC allows for the calculation of the local energy, the bottleneck in the algorithm’s scaling, at intervals rather than at every step since E_L is not required to compute the weight factors, which are now calculated explicitly as $\frac{\langle \phi_T | \phi^{(\tau+1)} \rangle}{\langle \phi_T | \phi^{(\tau)} \rangle} e^{\mathbf{x} \cdot \bar{\mathbf{x}} - \bar{\mathbf{x}}^2/2}$. We chose to implement the hybrid method primarily because it is much faster than the local energy method, and offers more flexibility when devising correlated sampling strategies.

The comparison of QMC energies with experimental results in general requires extrapolations to mitigate errors associated with the use of a finite time step and basis set. The Trotter error, due to the decomposition in (1.21) and three bounding conditions[76], can be estimated from a linear extrapolation of the energy differences from three independent simulations at $\Delta\tau = 0.02, 0.01$, and 0.005 Ha^{-1} . For the IPs and EAs of the G2 atomic test set, the energies resulting from the 0.01 Ha^{-1} time step were found in all cases to be well within 1 mHa of the $\Delta\tau$ -extrapolated value. Hence, we use only the $\Delta\tau = 0.01$ value to produce the results in Section III. Extrapolation to the complete basis set (CBS) limit was performed following the MP2-assisted protocol detailed in Ref. [77]. AFQMC calculations were performed at $x = 3$ (where x is the cardinal number of the basis set); UHF calculations were run at $x = 2-5$ and MP2 calculations at $x = 3, 4$ using GAMESS[146] or, for the reactions which utilize “ghost” basis functions, NWChem. IPs were computed using the cc-pVxZ basis[147], while EAs and the deprotonation energies to be compared with experimental data were computed using the aug-cc-pVxZ basis sets[148].

QMC statistical error bars were propagated through the data analysis procedure. In assessing the statistical error on the deviation of QMC from experiment, we assume that there are negligible uncertainties associated with 1) the experimental measurements, 2) the exponential (HF) and linear (MP2) fits in the CBS extrapolation procedure where we found that the former error is an order of magnitude or more smaller than the QMC error in representative cases, and 3) the scaling factor for the zero-point energies. With regard to 1), experimental uncertainties can be found in the

NIST database[149], and are ignored in order to isolate the statistical error associated with the QMC measurements, since the quantities presented in this work are differences between calculated energies and experimental measurements.

Calculations in this chapter utilized our first implementation of ph-AFQMC, in Fortran 90 linked with OpenBLAS[150, 151] and Expokit[152]. Random numbers were generated with the 48 bit Linear Congruential Generator with Prime Addend, as implemented in SPRNG5[153].

2.3 Results

2.3.1 Atomic IPs and EAs

In this section, we use our correlated sampling-based AFQMC approach to calculate the IPs and EAs of the 1st row atoms included in the G2 Ion Test set, which have experimental uncertainties of less than 0.05 eV[140]. The deviations from experiment within calculated AFQMC results are presented in Tables 2.1 and 2.2 alongside the deviations resulting from G2 theory[154], and DFT with the B3LYP exchange-correlation functional[155, 156] in the 6-311+G(3df,2p) basis[140].

Table 2.1: Experimental IPs and the deviations of various calculated results (theory - experiment) for atoms in the G2 Test Set in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. QMC calculations using single-determinant trial functions (phaseless and FP) use 5040 walkers per repeat, while those using multi-determinant trial functions use 1056 walkers per repeat.

Atom	Expt.	Δ ph-HF/QMC	Δ FP	Δ ph-CAS/QMC	Δ G2	Δ B3LYP
B	8.2980	-0.156(10) [*]	-0.0012(50)	-0.0162(26) ^a	0.10	-0.44
C	11.2603	0.0342(90) [*]		0.0045(43) ^b	0.08	-0.29
N	14.5341	0.1214(61) [*]	0.0050(84)	0.0100(38) ^c	0.06	-0.14
O	13.6181	-0.0830(22) [*]		-0.0360(24) ^d	0.08	-0.55
F	17.4228	0.0010(35)		0.0015(46) ^e	0.03	-0.34
Ne	21.5645	0.0775(51)		0.0159(34) ^f	-0.05	-0.21

^{*} PES

^a QMC trial from CASSCF(5/4,8) with minimum CI coefficient of 0.034 (20/14 determinants)

^b QMC trial from CASSCF(4/3,8) with minimum CI coefficient of 0.034 (14/7 determinants)

^c QMC trial from CASSCF(5/4,8) with minimum CI coefficient of 0.01 (29/32 determinants)

^d QMC trial from CASSCF(6/5,13) with minimum CI coefficient of 0.01 (67/62 determinants)

^e QMC trial from CASSCF(7/6,8) with minimum CI coefficient of 0.033 (7/2 determinants)

^f QMC trial from CASSCF(8/7,16) with minimum CI coefficient of 0.0085 (101/104 determinants)

For the calculations which used the unrestricted HF determinant as the trial function, PES was

Table 2.2: Experimental EAs and the deviations of various calculated results (theory - experiment) for atoms in the G2 Test Set in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. QMC calculations using single-determinant trial functions (phaseless and FP) use 5040 walkers per repeat, while those using multi-determinant trial functions use 1056 walkers per repeat.

Atom	Expt.	Δ ph-HF/QMC	Δ FP	Δ ph-CAS/QMC	Δ G2	Δ B3LYP
B	0.2797	-0.0422(31)		-0.0090(30) ^a	0.09	-0.18
C	1.2621	0.0762(84) [*]		0.0031(43) ^b	0.07	-0.11
O	1.4620	0.0505(86)		0.0327(53) ^c	0.06	-0.22
F	3.4013	0.222(18) [*]	-0.033(20)	0.0474(49) ^d	-0.08	-0.13

^{*} PES

^a QMC trial from CASSCF(3/4,8) with minimum CI coefficient of 0.01 (34/46 determinants)

^b QMC trial from CASSCF(4/5,8) with minimum CI coefficient of 0.01 (30/53 determinants)

^c QMC trial from CASSCF(6/7,8) with minimum CI coefficient of 0.01 (39/77 determinants)

^d QMC trial from CASSCF(7/8,16) with minimum CI coefficient of 0.0075 (145/243 determinants)

used for the IPs of B, C, N, and O with equilibration times of 35, 20, 15, and 15 Ha^{-1} , respectively, and similarly for the EAs of C and F with 15 and 10 Ha^{-1} , respectively. For the remaining species, equilibration was facile and thus the AFs were correlated from the beginning of the imaginary-time propagation. Notable deviations from the experimental values are found for the IPs of B and N, and the EA of F when the UHF state is used as the trial function (nearly identical errors have been previously reported within ph-AFQMC[71]). These discrepancies are resolved in the correlated sampling-based FP results. Since the energies from FP are not biased by the phaseless constraint and thus insensitive to the trial function used, (any) small residual errors can be attributed to the MP2-assisted CBS extrapolation scheme.

The long equilibration times and inaccuracies encountered in the above cases are manifestations of the fact that single-determinant trial functions obtained from mean-field calculations are generally not well-suited to describe the open-shell systems involved in redox reactions. For instance in the IP of B, comparing FP and ph-AFQMC calculations for both the neutral and cationic species exposes the fact that the error in the IP stems from inaccuracies in the computation of the total energy of the neutral B atom, which has a single unpaired electron in the triply-degenerate p orbital manifold. The use of trial wavefunctions with proper symmetry properties has previously been shown to lead to improved accuracy within the phaseless approximation[95], so we now consider trial functions which are eigenfunctions of \hat{S}^2 for use with the phaseless constraint. While

employing a restricted open-shell HF trial affords no appreciable gain in accuracy, the use of a truncated CASSCF trial with a very modest active space size and a small number of determinants is sufficient to eliminate most of the error for the IPs of both B and N.

The accurate description of the EA of F proves to be more demanding. It is possible that the phase problem is particularly problematic for the case of F- due to the small atomic radius which may lead to very strong electronic correlations. Such sizable correlations manifest themselves in the algorithm as a large imaginary component of the propagator, which in turn leads to the elimination of crucial physical information during the projection to the real axis if the trial function does not adequately provide the gauge information on the Slater determinants of the ground state. In such cases, excitations into a large number of virtual orbitals will make significant contributions to the correlation energy, and therefore a large active space is required to generate the CASSCF trial function so that the resulting phase projections are sufficiently benign.

In general, the implementation and efficacy of the correlated sampling scheme presented in this work remain unaltered in the case of a multi-determinant trial function since walker determinants are propagated in exactly the same manner. We do find a reduction in the required equilibration times for all multi-determinant ph-AFQMC calculations that we have performed, relative to single-determinant calculations of the same systems, rendering the use of PES unnecessary. In addition, the use of multi-determinant trial functions results in drastically smaller error bars, even when more than 4x *fewer* walkers are used. Increasing the quality of the CASSCF trial function is a promising way to systematically reduce the error from the phaseless constraint, and we find that in all cases the use of a reasonable number of determinants produces redox energies within the maximum experimental error of 0.05 eV for the G2 Test Set.

2.3.2 The Case of Methanol

Motivated by the arguments set forth in Section II.C, we now describe the details of a composite method for calculating the *adiabatic* IPs, deprotonation free-energies, and hydrogen-dissociation energies of molecular systems, and illustrate the accuracy of our approach on the case of methanol. We utilize a stepwise process consisting of: 1) the fixed-geometry process calculated with correlated

sampling-based ph-AFQMC, and 2) a geometry-relaxation step calculated within MP2. Optimal correlation can be achieved in 1) since the same set of basis functions is used, while in 2) we expect large error cancellation resulting from the fact that the initial and final ground-state geometries are typically very similar in redox, deprotonation, and hydrogen abstraction reactions. The calculations are performed in a triple-zeta basis with additional diffuse basis functions for all species involved in the deprotonation reaction. The resulting energy differences from these two steps are added together, and the endpoints are extrapolated to the CBS limit.

For all reaction types, zero-point energies are calculated at the level of HF/6-31G* and scaled by a factor of 0.899 to account for anharmonicity and the known shortcomings of HF theory, following the G2 protocol. Deprotonation free-energy results incorporate the value of -6.28 kcal/mol as the free-energy of a proton at 298 K[157], and we use the exact ground-state energy of H \cdot (-0.5 Ha) to calculate the bond dissociation enthalpy of the O-H bond.

Table 2.3 illustrates the accuracy of our correlated sampling protocol with respect to experimental results for all three reaction types. The quality of our IP and deprotonation free-energy results surpass that of the more costly G2 method, while the O-H bond dissociation result has an error of comparable magnitude. We note that the single-determinant trial wavefunction is sufficient to produce a near-exact deprotonation free-energy, which we attribute to the fact that in this reaction type both the protonated and deprotonated species are closed-shell. Moreover, our correlated sampling-based ph-AFQMC calculations, using an extremely modest number of walkers, predicts the energy differences corresponding to all of these reaction types to within chemical accuracy, which is not the case for the G2 method.

2.3.3 Basis Set Size, Number of Random Walkers, and CPU-time Reduction

2.3.3.1 Basis Set Size

The statistical noise of individual AFQMC runs is expected to increase with the number of AFs, α , yet the precise scaling is subtle. Each AF contributes additional noise. On the other hand, the magnitude of the contribution is moderated by a degree which depends on the form of the interaction and the decomposition procedure leading to Eq. (1.23). For example, the error bar

Table 2.3: Adiabatic Reaction Energies for Methanol. Experimental values and the deviations of the correlated sampling-based ph-AFQMC and G2 results (theory - experiment) in eV. QMC statistical errors in the two right-most digits are shown in parenthesis. All QMC calculations use 192 walkers per repeat.

Reaction	Expt.	Δ ph-QMC	Δ G2
Ionization Potential	10.84	0.034(27) [†]	-0.11
Deprotonation Free-Energy	16.2695	-0.005(21) ^{††}	0.0484 ^a
O-H Bond Dissociation Energy	4.5359	0.039(14) [†]	0.0166 ^b

[†] QMC trial from CASSCF(10/11, 10/11) with minimum CI coefficient of 0.02 (27/32 determinants)

^{††} HF trial

^a Ref. [158]

^b Ref. [159]

is seen to change little beyond a modest cutoff in plane-wave AFQMC[73]. In principle α grows as M^2 , but our (rather conservative) truncation of the Cholesky decomposition via the cutoff mentioned in Section II.D results in $\alpha \sim 10M$. As a result, we expect the statistical error in an uncorrelated calculation to increase with M before saturating toward the CBS limit. This could be problematic given that large basis sets containing functions associated with high angular momenta are frequently required to accurately describe the correlation energy[147], and also given the fact that most interesting applications involve large systems.

Here we show that the magnitude of the reduction in statistical error enabled by the use of correlated sampling grows with M , such that the error is nearly independent of M . Fig. 2.6 shows the standard errors resulting from calculations of the IP of the K in the 6-31G*, 6-31+G*, and 6-311+G* basis sets (which consist of 23, 35, and 45 basis functions, respectively); Fig. 2.7 illustrates the same effect for fixed-geometry deprotonation reactions of water, methanol, and ethanol (58, 116, and 174 basis functions). For both reaction types, while the errors from the uncorrelated calculations increase significantly with M , we find that those resulting from correlated sampling remained roughly constant. We anticipate that this finding will hold in general, and will be of crucial importance in future applications of correlated sampling-based AFQMC to larger molecules.

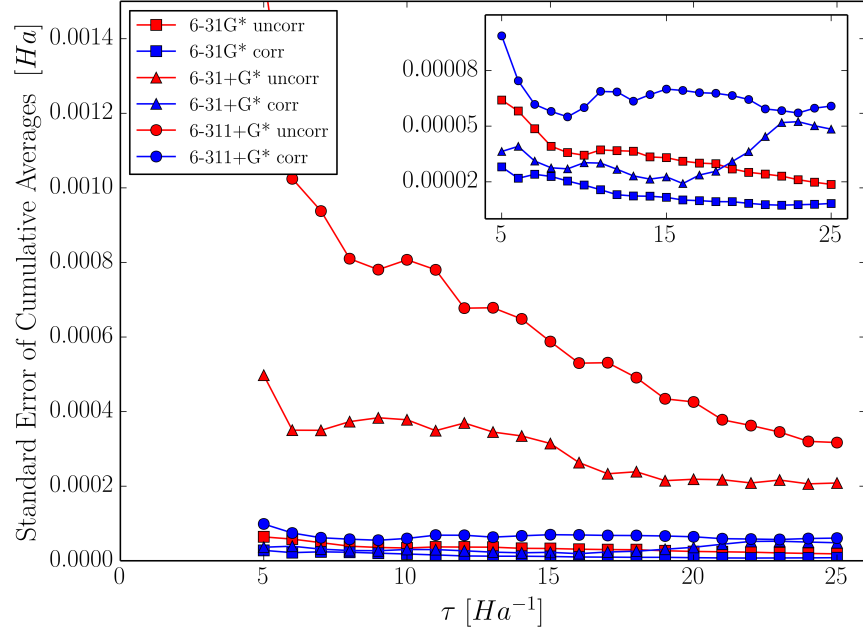


Figure 2.6: Comparison of the standard errors of the cumulative averages resulting from correlated and uncorrelated sampling in computing the IP of the K atom in three different basis sets with 5040 walkers per repeat. The inset zooms in on the lower region, for clarity.

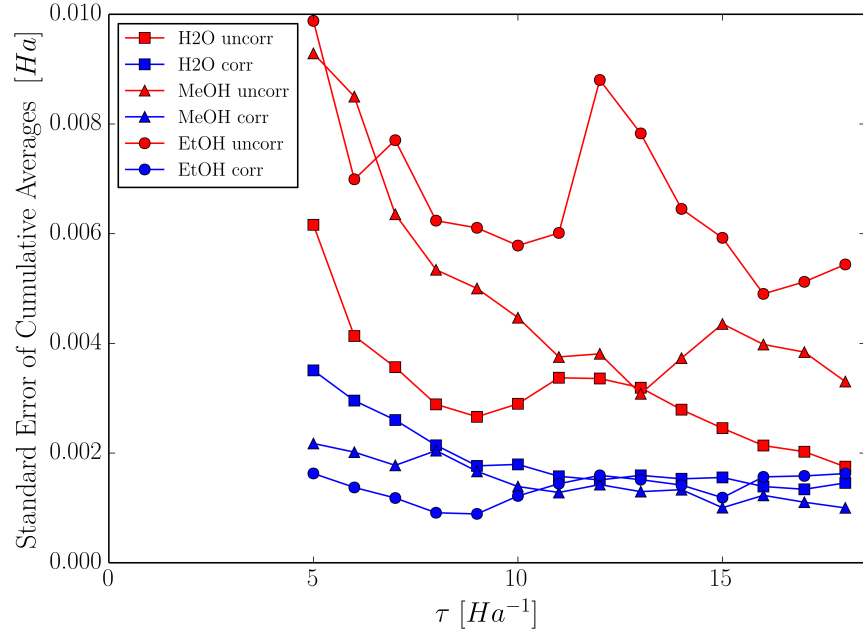


Figure 2.7: Comparison of the errors resulting from correlated and uncorrelated sampling in computing the fixed-geometry deprotonation energies of H_2O ($M=58$), CH_3OH ($M=116$), and $\text{C}_2\text{H}_5\text{OH}$ ($M=174$) with 192 walkers per repeat.

2.3.3.2 Reduction in the Number of Random Walkers and CPU-Time

In an uncorrelated QMC calculation for a fixed length of propagation time, the resulting standard error can be reduced by increasing the number of random walkers, N_{wlk} , used in the MC evaluation of the HS integral in Eq. (1.23). However, given that the required computational expense increases linearly with N_{wlk} , using a brute-force approach that simply increases N_{wlk} is less practical for many systems. For energy differences, correlated sampling provides a much cheaper alternative as it allows for a dramatic reduction in the N_{wlk} required to achieve a given statistical error.

The errors of the cumulative averages of the IP of K in the 6-31+G* basis (M=35) and the fixed-geometry deprotonation of methanol in the cc-pVTZ basis (M=116) are shown for different values of N_{wlk} in Figs. 2.8 and 2.9, respectively. For both reaction types, we make the following observations: First, for a given N_{wlk} , the standard error is significantly lower in the correlated sampling case. Second, the magnitude of this reduction is greater for smaller N_{wlk} . Finally, while the standard errors of the uncorrelated runs increase as N_{wlk} is reduced, in the correlated sampling runs the error is relatively insensitive to N_{wlk} .

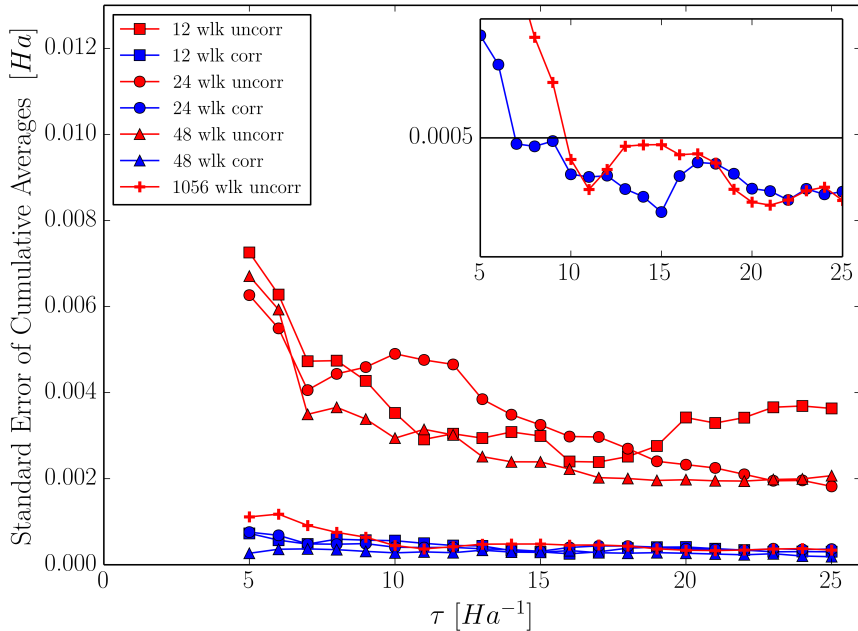


Figure 2.8: Dependence of standard error on the number of random walkers per repeat for the IP of K in the 6-31+G* basis. The inset highlights the errors of the uncorrelated run with 1056 walkers and the correlated run with 24 walkers, compared with the 0.5 mHa error target.

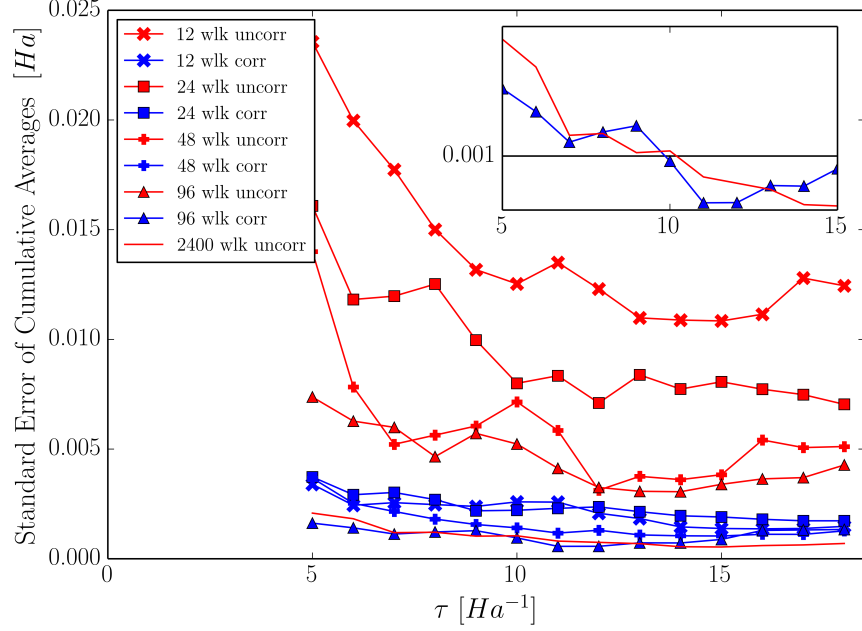


Figure 2.9: Dependence of standard error on the number of random walkers per repeat for the fixed-geometry deprotonation of methanol in the cc-pVTZ basis. The inset highlights the errors of the uncorrelated run with 2400 walkers and the correlated run with 96 walkers, compared with the 1 mHa error target.

In light of these findings, we are in a position to understand how a reduction in the standard error due to correlating the AFs translates into a reduction in CPU-time. Considering first the IP of K in the 6-31+G* basis, we choose a target standard error of 0.5 mHa on the QMC energy (which defines the 99% confidence interval as the cumulative average of the energy ± 1 kcal/mol), and compare the total CPU-time required to achieve this via correlated and uncorrelated approaches. In the former case, we find that using only 24 walkers in each of the 11 repeat calculations is sufficient to achieve the target error and a resulting energy in agreement (*i.e.* within the 99% confidence interval) with a benchmark result obtained with 5040 walkers. In fact, as few as 6 walkers produced the same level of accuracy in some cases. The inset of Fig. 2.8 shows that the statistical error falls below the target at $\tau \sim 7$. The total CPU-time required to propagate 11 repeats for this length of imaginary-time is 41.2 minutes on a single 2.60 GHz Intel Xeon processor. Using the same number of walkers *without* correlating the AFs, we find that the target error is not reached even after 200 Ha^{-1} . Using 1056 walkers gives rise to a standard error that falls below 0.5 mHa after 10 Ha^{-1} , as shown in the inset of Fig. 2.8, and a resulting QMC energy that is in agreement with the benchmark

result. This calculation takes 2262.6 minutes on a single processor, and we thus conclude that our correlated sampling approach reduces the CPU-time by a factor of approximately 55.

For the deprotonation of methanol we use a target error of 1 mHa. As shown in the inset of Fig. 2.9, both the uncorrelated run with 2400 walkers and the correlated run with 96 walkers yield results that fall below our target error at $\tau \sim 10$. We use the fact that the total CPU-time is proportional to the product of the number of walkers and the propagation time to estimate that correlating the AFs reduces the CPU-time by a factor of approximately 25. Currently the total calculation, including all 11 repeats, requires ~ 154 hours on a single CPU core. We perform a similar analysis for the dissociation of $\text{H}\cdot$ from the O-H bond of methanol. Due to the relatively large computational cost of using a CASSCF trial function (with the same active space and truncation scheme described in Table III), we increase our target error to 2 mHa. Figure 2.10 shows that the errors on the cumulative averages from both an uncorrelated run with 288 walkers and a correlated run with 12 walkers fall below our target error at $\tau \sim 6$. Thus we deduce a speed-up factor of 24 for this H abstraction reaction. On a single CPU core this requires ~ 157 hours.

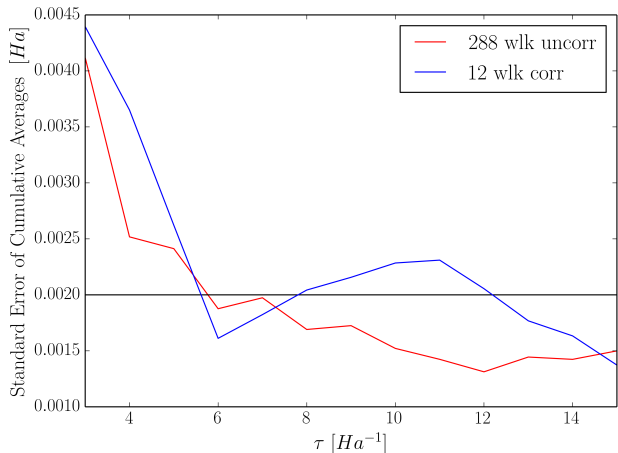


Figure 2.10: Comparison of the standard errors resulting from the use of correlated sampling with 12 walkers per repeat and uncorrelated sampling with 288 walkers per repeat to calculate the energy difference associated with the fixed-geometry removal of $\text{H}\cdot$ from the O-H bond of methanol in the cc-pVTZ basis. The 2 mHa error target is shown in black.

As the previous sections have shown, the magnitude of the reduction in the statistical error, and consequently the CPU-time, as compared with an uncorrelated calculation depends on the number of walkers used and the size of the basis set. In addition, one intuitively expects our

correlated sampling approach to work better for systems that more closely resemble each other. This is crudely the case, yet the subtleties involved in such a claim warrant further discussion. Indeed, referring to the propagator in (2.1), while our correlated sampling method ensures that the $\hat{\mathbf{v}}$ operators and the AFs, \mathbf{x} , are the same for both the primary and secondary systems, the FBs, $\bar{\mathbf{x}}$, as defined in (1.28) and the expectation values with respect to the trial functions $\langle \hat{\mathbf{v}} \rangle$ will in general be different. In the limit that the primary and secondary systems are identical, the entire propagator in (2.1) is identical for both systems, and the statistical error in the energy difference will be exactly and trivially zero as a result of perfect walker-pair correlation. Otherwise, the reduction in statistical error afforded by our correlated sampling approach becomes less pronounced the more the trial wavefunctions of the primary and secondary systems differ, since the the FBs and $\langle \hat{\mathbf{v}} \rangle$ are expectation values that depend explicitly on the trial wavefunctions. It is encouraging to note that despite the differences in trial functions correlating only the AFs yields such large speed-ups in CPU-time. While additionally correlating the $\langle \hat{\mathbf{v}} \rangle$, possibly by using some combination of the trial functions for both systems, would compromise the accuracy of the phaseless approximations, future studies will explore optimal ways to pair walkers such that the similarity in the FBs of walker pairs is maximized. It is encouraging that even for systems for which the *electronic* energies of the primary and secondary systems differ by some 200 eV, as is the case in the deprotonation of methanol, our correlated sampling approach still yields dramatic efficiency gains.

2.4 Conclusions and Outlook

In this chapter we have devised a correlated sampling protocol for the calculation of chemically relevant energy differences within the exact and phaseless variants of AFQMC. For molecules we utilize a two step strategy in which optimal walker-pair correlation is achieved in the ph-AFQMC description of the fixed-geometry process, while the geometry relaxation energy is calculated with the confines of MP2. Together with an MP2-assisted CBS extrapolation method we obtain calculated IPs, EAs, deprotonation free-energies, and bond dissociation energies that are in excellent agreement with experiments. Moreover, our correlated sampling approach yields large reductions

in the statistical errors relative to those obtained from uncorrelated approaches. In contrast to uncorrelated AFQMC, where the error bars are found to increase with system size and/or the number of basis functions, correlating the AFs keeps the statistical error relatively constant as chemical complexity increases. In addition, our approach drastically reduces the number of walkers required to achieve a given statistical error target, which translates into large reductions in CPU-time.

We utilize a “ghost” basis function strategy that enables the application of a correlated sampling-based approach to processes involving large energetic changes. Indeed, given that the correlated sampling scheme outlined here is successful for electron, proton, and H \cdot transfer reactions, we are optimistic that other chemical changes are within reach. In the future, we plan to systematically investigate chemical reactions which involve substantial changes in geometry, including the addition/removal of larger, more complex functional groups. Along the same lines we are optimistic about the savings that our correlated sampling approach may yield when basis functions that are independent of the nuclear coordinates, such as plane-waves, are used. We anticipate that the insensitivity of the statistical error to basis set size that we observe in this work may partially or totally offset the relatively large number of plane-waves typically required for convergence.

Chapter 3

Phaseless Auxiliary-Field Quantum Monte Carlo on Graphical Processing Units¹

In the previous Chapter we introduced a correlated sampling (CS) approach for quantities involving energy differences which is capable of reducing computational prefactors by approximately an order of magnitude. In this chapter we present a different but complementary strategy involving hardware optimization on graphical processing units (GPUs) which can drastically reduce the prefactors in calculations of general ground-state properties.

GPUs have several distinct advantages over traditional Central Processing Units (CPUs), including the ability to perform efficiently parallelized matrix operations both in serial and in “batches,” and the use of single-precision (sp) floating-point arithmetic with significant gains in computational speed. We refer the reader to Ref. 160 for a lucid exposition of many general properties of GPU hardware. In recent years the use of GPUs has been extended well beyond traditional image visualization tasks into many fields such as machine learning[161] and molecular mechanics[162]. Of particular relevance to our work presented here is the progress in performing electronic structure calculations on GPUs. This hardware has been utilized to ef-

¹Based on work published in *J. Chem. Theory Comput.* 2018, 14, 8, 4109-4121.

ficiently evaluate the integrals required in ab-initio calculations[163, 164, 165, 166], to perform HF[167, 168] and DFT calculations[169, 170, 171], and to study model systems such as the Hubbard Model within the dynamical cluster approximation[172] and the Ising model[173, 174]. In addition there have been recent GPU implementations of MP2[175, 176, 177, 178], CC methods[179], TDDFT[180], CI[181, 182], and CASSCF approaches[183, 184]. Efficient algorithms to compute energy gradients[185, 186] and tensor contractions[187] have also been developed.

With respect to QMC methods, GPU implementations have been formulated primarily for real-space approaches. For example, DMC with sp arithmetic has been accelerated by a factor of $\sim 6x$ on a GPU versus a quad-core CPU[188]. A recent study employing a multi-GPU implementation has reported speed-ups of a factor of 10-15x relative to a quad-core CPU for VMC and DMC for real materials[189]. Very recently an open-source QMC suite, QMCPACK[190], has released scalable implementations of real-space QMC methods. An implementation of AFQMC is mentioned in Ref. 190, although data illustrating its efficiency and accuracy is not yet available.

In this chapter we detail our GPU implementation of the phaseless variant of AFQMC (ph-AFQMC), and illustrate its performance and accuracy via calculations of the total energies of linear chains of hydrogen atoms and the all-electron IPs of the first-row transition metal (TM) atoms. We explicitly compare our GPU wall-times with CPU timings from a code of equivalent algorithmic sophistication. Speed-ups from the GPU port of two orders of magnitude are seen in large systems, with the potential for even greater reductions of the scaling prefactor depending on the system-size. The robustness and accuracy of our implementation are shown by comparing our calculated values to either exact numerical techniques or experiment.

This chapter is organized as follows: In Sec. 3.1 we detail our GPU implementation and highlight significant algorithmic additions. In Sec. 3.2 and 2.3 we present timing and accuracy results for the hydrogen chains and TM IPs, respectively, and comment on the advantages of the correlated sampling approach. In Sec. 3.4 we conclude with a summary of our results and a discussion of future work.

3.1 GPU Implementation

In contrast to traditional computing paradigms which utilize CPUs to execute all computing tasks, we employ a strategy in which CPUs offload a majority of the computational effort to one or more GPU cards. A typical GPU device has 4-12 GB of memory which is separate from that accessible by the CPU; therefore, data is usually allocated on both the host and device, and inter-communication between these different memory spaces requires the explicit copying of data back and forth. In this work great care is taken to minimize such transfers, and we create custom memory structures that organize memory addresses and facilitate switching between sp and double-precision (dp).

A flowchart outlining the ph-AFQMC algorithm with 3 CPU/GPU pairs is presented in Fig. 3.1. First the root CPU reads in relevant quantities such as matrix elements of \hat{H} , overlap integrals, and the initial and trial wavefunctions, and then completes a preliminary setup which includes transformations to an orthonormal basis, walker and operator initializations. These quantities are sent to all devices, after which the tensor (or slices of it) used in the energy evaluation is precomputed directly on the GPUs. Throughout the entire sequence of functions involved in propagating a walker, all operations are performed on the devices, i.e. without any data transfers or operations involving the CPUs. We utilize NVIDIA’s Compute Unified Device Architecture (CUDA) Basic Linear Algebra Subprograms (cuBLAS) library to execute, e.g., the matrix multiplications that propagate walker determinants by a one-body operator, and have supplemented this library with custom C/CUDA functions which can be classified roughly into two types. Those in the first constitute a matrix library of kernels which carry out, most notably, element-wise matrix additions and matrix sums (i.e. matrix \rightarrow scalar). These are used frequently to compute the trace of a matrix product, which is utilized in the calculation of expectation values such as the force-bias and local energy. In addition to such library-type functions, we also wrote GPU kernels to sample auxiliary-fields, compute the force-bias, assemble and exponentiate one-body operators, carry out the Sherman-Morrison-Woodbury (SMW) updates, orthonormalize the orbitals of walker determinants, and measure the local energy. With the exception of the SMW and energy measurement functions, which we will subsequently detail, the GPU port of the above functions did not involve

notable algorithmic improvements over our CPU implementation.

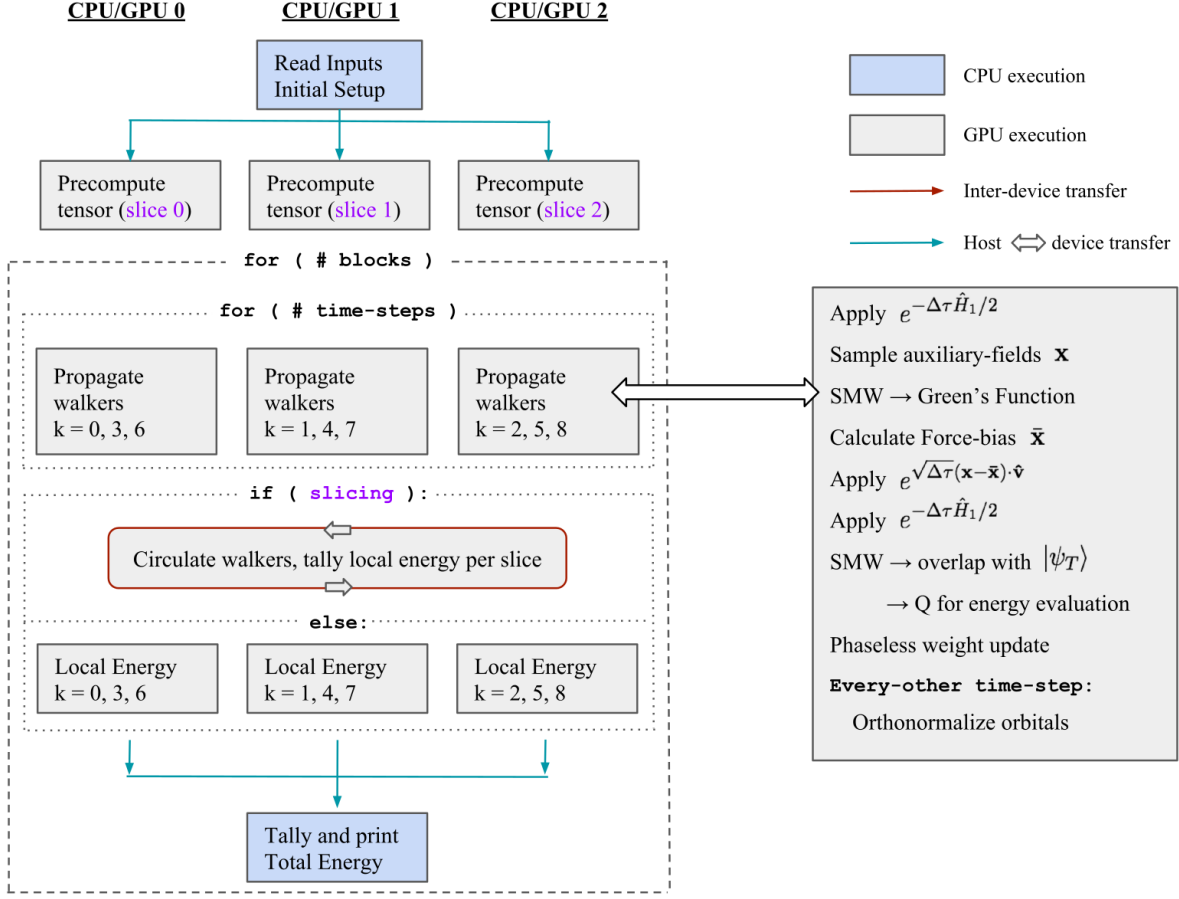


Figure 3.1: Flowchart of our AFQMC implementation, for 9 walkers ($k = 0, \dots, 8$) and 3 CPU-GPU pairs (CPU/GPU 0, 1, 2). If the size of the precomputed tensor used in the energy evaluation exceeds the memory capacity of each device, each GPU precomputes and stores only a slice of the tensor, and the energy of a walker is computed by circulating the walker to all other GPUs and tallying the partial energies obtained from the resident slices.

Once the code enters the loops shown in Fig. 3.1, data need only be transferred from the current device when the energy is measured, which for our typical time-step choice of $\Delta\tau = 0.005 \text{ Ha}^{-1}$ happens once in 20 propagation steps. An explanation of the slicing variant of the energy algorithm will be presented later in Sec. 3.2. For now we simply wish to emphasize that throughout the majority of an AFQMC calculation data does not need to leave the devices. This is in large part why the current implementation leads to such pronounced speed-ups compared to our initial attempts to simply offload the matrix multiplications.

A relatively new addition to the cuBLAS library are so-called “batched” functions which perform

many smaller operations simultaneously, e.g. a set of matrix-matrix multiplies or lower-upper (LU) decompositions. These batched functions are well-suited for operations that are individually too small to parallelize effectively across thousands of cores. We utilize this feature heavily in our implementation of SMW updates to quickly compute equal-time Green's functions when multi-determinant trial functions are used. We note that previous Diffusion MC studies have utilized similar SMW updates[191, 192, 189]. Given a reference matrix A , the following formulas are used to compute the determinants and inverses of matrices which differ from A by one or more row or column:

$$\begin{aligned} \det(A + U_i V_i^T) &= \det(I + V_i^T A^{-1} U_i) \det(A) \\ (A + U_i V_i^T)^{-1} &= A^{-1} - A^{-1} U_i (I + V_i^T A^{-1} U_i)^{-1} V_i^T A^{-1}. \end{aligned} \tag{3.1}$$

In the context of ph-AFQMC, suppose we use a multi-determinant trial wavefunction, $|\Psi_T\rangle = \sum_{i=0} c_i |\psi_{T,i}\rangle$, where $\langle \psi_{T,i} | \psi_{T,j} \rangle = \delta_{ij}$. Then, for the k^{th} walker determinant $|\phi_k\rangle$, $A = [\psi_{T,i=0}]^\dagger [\phi_k]$, where the square brackets denote a matrix representation, U_i, V_i are of dimension $N_\sigma \times E_i$, where N_σ is either the number of spin-up or spin-down electrons, and E_i is the number of excitations required to form the i^{th} configuration of the multi-determinant expansion from the reference configuration. The determinant and inverse of the reference matrix corresponding to zero excitations ($i = 0$) is computed first for spin-up and spin-down configurations, followed by batched SMW updates for all $i \neq 0$. Sub-cubic scaling with respect to particle number is achieved since $E_i \ll N_\sigma$. Fig. 3.2 highlights the efficiency of our batched implementation of the SMW algorithm, for the Mn atom in the aug-cc-pwCVQZ-DK basis (185 basis functions, 25 electrons). “Propagation time” denotes the total wall-time minus the time spent on initial setup, e.g., memory allocation, input/output, and precomputation of the operators and required intermediates. Previously, going from, e.g., 10 to 1200 determinants would multiply the propagation time by a factor of 120. In contrast, our SMW algorithm reduces this to a mere factor of 3.9.

We have developed a GPU-optimized algorithm for evaluating the local two-electron energy of a walker. 4-index tensors are precomputed once at the start of a simulation, which, in the spin-free

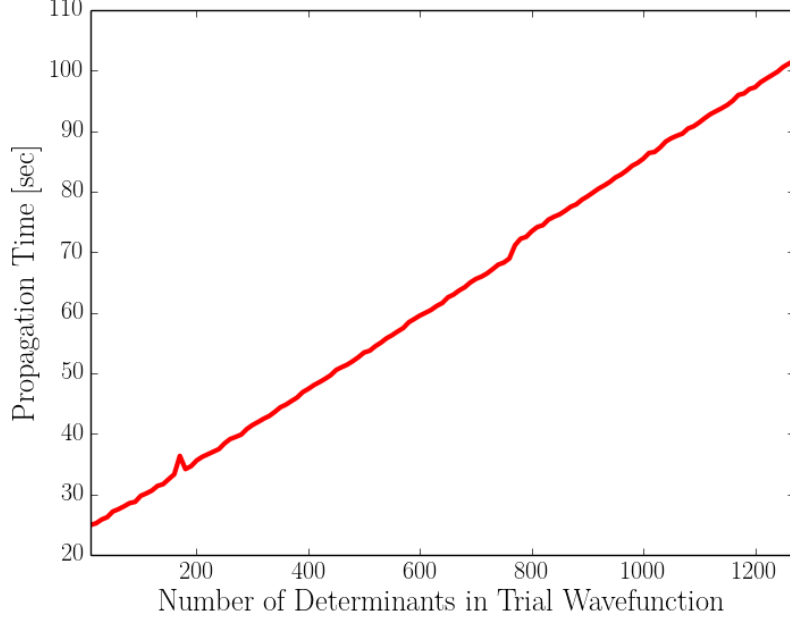


Figure 3.2: Propagation time vs the number of configurations in the CASSCF trial wavefunction for the Mn atom in the aug-cc-pwCVQZ-DK basis. Calculations use sp, a CD threshold of $10^{-4} Ha$, 20 walkers and imaginary-time trajectories of length $1 Ha^{-1}$ with a time step of $\Delta\tau = 0.005 Ha^{-1}$. Walker orthonormalization and local energy measurements were performed every 2 and 20 steps, respectively.

and single-determinant trial case (for simplicity), are of the form:

$$Y_{ijab} = \sum_{kl} \sum_{\alpha} L_{ik}^{\alpha} L_{jl}^{\alpha} [\psi_{T,ak}^{\dagger} \psi_{T,bl}^{\dagger} - \psi_{T,al}^{\dagger} \psi_{T,bk}^{\dagger}], \quad (3.2)$$

where the L^{α} arise from decomposing the two-electron integrals via CD or DF ($V_{ijkl} = \sum_{\alpha} L_{ik}^{\alpha} L_{jl}^{\alpha}$), indices i, j, k, l run from 1 to M (basis size) while indices a, b run from 1 to N (number of electrons). ψ_T is a matrix with columns composed of the orbitals in the trial function. Importantly, the sums over k, l and over auxiliary-fields α , the number of which typically scales as $2-10M$, need only be computed once at the start of the simulation. The energy is evaluated by pairing indices $i, a \rightarrow \gamma$ and $j, b \rightarrow \delta$, thereby flattening the 4-index tensor in (3.2) to a 2-index tensor, and then performing the following contraction:

$$E_{2e}[\phi] = Tr(QY) = \sum_{\gamma\delta} Q_{\gamma\delta} Y_{\gamma\delta}, \quad (3.3)$$

where $Q_{\gamma\delta} = [\phi(\psi_T^\dagger\phi)^{-1}]_\gamma[\phi(\psi_T^\dagger\phi)^{-1}]_\delta$. The Q and Y matrices are of size $MN \times MN$, and hence the energy evaluation dominates the scaling, with respect to system size, of both the required memory and run-time of an AFQMC calculation. In Sec. 3.2 we describe a strategy to split the memory burden among multiple GPU cards on a single node, and suggest approaches to attain additional scalability in Sec. 3.4. In what follows we illustrate the efficiency of our current implementation.

Table 3.1 shows select performance metrics, from an analysis using NVIDIA’s nvprof code profiler, for the GPU kernels involved in our energy algorithm. While these metrics can, in general, vary widely depending on the particulars of both the device architecture and the description of the chemical system under study (e.g. choices of basis and trial function), we chose to optimize our code’s performance for ph-AFQMC calculations using large multi-determinant trials, in light of our interest in strongly correlated systems. As before, we show data for the Mn atom in the aug-cc-pwCVQZ-DK basis using sp, with 1200 determinants in the trial function (as is used to calculate the IP reported in Sec. 3.3). Nearly 90% of the wall-time is spent in CGEMM from the cuBLAS library, which we use to compute a quantity analogous to Q in Eq. (3.3) but generalized to the case of a multi-determinant trial function. It appears that our custom element-wise matrix multiplication and matrix sum kernels, while at peak warp utilization, are limited by the device memory (DRAM) bandwidth. Additional fine-tuning of the latter kernels’ usage of the memory hierarchy will at best result in a small improvement in the overall performance of the energy function for this system ($\sim 1.1\times$, from Amdahl’s law), given that the majority of the time is spent executing the highly optimized CGEMM kernel.

Table 3.1: Efficiency metrics of the GPU kernels involved in our energy algorithm. For each kernel we show the percentage of total run-time spent in that kernel, compute utilization as a percentage of peak compute performance, the number of registers per thread, memory utilization as a percentage of peak bandwidth (shown only for the memory type exhibiting the highest utilization), and the occupancy, i.e. the percentage of available warps (a group of 32 threads) that are active.

Kernel	% Run-time	% Comput. Util.	Reg/Thread	% Mem. Util. (type)	% Occ.
CGEMM	88.2	95-100	84	50-60 (shared)	25
El. MatMul	6.8	< 10	12	80-90 (device)	91.6
Matrix Sum	2.3	< 10	16	80-90 (device)	96.1

Finally we introduce the use of DF[193] in AFQMC calculations, where effective densities $\bar{\rho}_{ij}(\mathbf{r})$

are fit to auxiliary basis functions, $\chi(\mathbf{r})$:

$$V_{ijkl} = \int d\mathbf{r}_1 d\mathbf{r}_2 \phi_i(\mathbf{r}_1) \phi_j(\mathbf{r}_1) \frac{1}{r_{12}} \phi_k(\mathbf{r}_2) \phi_l(\mathbf{r}_2) \quad (3.4)$$

$$\sim \int d\mathbf{r}_1 d\mathbf{r}_2 \bar{\rho}_{ij}(\mathbf{r}_1) \frac{1}{r_{12}} \phi_k(\mathbf{r}_2) \phi_l(\mathbf{r}_2) \quad (3.5)$$

$$= \sum_{\nu} d_{\nu}^{ij}(\nu|kl) \quad (3.6)$$

The last equality follows from inserting $\bar{\rho}_{ij}(\mathbf{r}_1) = \sum_{\nu} d_{\nu}^{ij} \chi_{\nu}(\mathbf{r}_1)$, and defining the three-center integrals $(\nu|kl) = \int d\mathbf{r}_1 d\mathbf{r}_2 \chi_{\nu}(\mathbf{r}_1) \frac{1}{r_{12}} \phi_k(\mathbf{r}_2) \phi_l(\mathbf{r}_2)$. The expansion coefficients can be chosen such that $d_{\nu}^{ij} = \sum_{\mu} (ij|\mu) \mathbf{J}_{\mu\nu}^{-1}$, where $\mathbf{J}_{\mu\nu} = \int d\mathbf{r}_1 d\mathbf{r}_2 \chi_{\mu}(\mathbf{r}_1) \frac{1}{r_{12}} \chi_{\nu}(\mathbf{r}_2)$. Expressing $\mathbf{J}_{\mu\nu}^{-1}$ as a contraction over a third index allows the two-electron integrals to be written in a form suitable for AFQMC:

$$V_{ijkl} = \sum_{\alpha} \left(\sum_{\mu} (ij|\mu) \mathbf{J}_{\mu\alpha}^{-1/2} \right) \left(\sum_{\nu} \mathbf{J}_{\alpha\nu}^{-1/2} (\nu|kl) \right) = \sum_{\alpha} L_{ij}^{\alpha} L_{kl}^{\alpha}. \quad (3.7)$$

The number of terms in the sum over α is equal to the number of auxiliary-fields sampled by each walker in AFQMC, which via DF is typically reduced to $\sim 2M$. As a result the calculation of the force-bias and the assembly of the one-body operator in Eq. (1.22) can be done faster, and fewer L matrices (each with M^2 elements) need to be stored in memory relative to when CD is used. In addition, fewer auxiliary-fields generally leads to a reduction in statistical noise. The accuracy of the DF approximation will be assessed in Sec. 3.2.

3.2 Illustration with Hydrogen Chains

In this section we explore the effects on both computational efficiency and accuracy due to the use of sp vs dp, and DF vs CD for linear chains of hydrogen atoms. These systems have played an important role in benchmarking new theories of correlated electronic materials[194, 195, 80, 196, 197, 198, 81]. While these systems do not capture many nuances of more realistic molecular systems, they are nevertheless a useful prototype capable of (1) yielding wall-time and scaling insights due to the ability to systematically increase the system size, (2) providing an atomistic

analogue of well-studied model systems such as the Heisenberg and Hubbard models albeit with a more realistic description of long-range Coulomb interactions, while (3) exhibiting strong static correlation at large bond lengths.

3.2.1 Computational Details

For all hydrogen chain calculations we use the cc-pVDZ basis, for which there is abundant benchmark data[81]. In this basis there are 5 basis functions per electron, a notably smaller number than used in typical molecular calculations. The Weigend Coulomb-fitting basis set[199] is employed as the auxiliary basis for DF, and CDs in this section employ a threshold of $10^{-5} Ha$ (as chosen in Ref. 81).

We use PySCF[144] to compute all inputs required of our ph-AFQMC code. Unless otherwise specified we use an imaginary-time step of $0.005 Ha^{-1}$. Walker orbitals are orthonormalized after every two propagation steps, to preserve the anti-symmetry of the walker configurations and also to keep the magnitude of orbital coefficients and associated quantities as small as possible (thus extending the accuracy of sp). We employ the hybrid method of ph-AFQMC[76] to minimize evaluations of the local energy, which is measured every $0.1 Ha^{-1}$. The total number of walkers is fixed throughout each simulation, and when required we use a population control (PC) algorithm at intervals of $0.1 Ha^{-1}$. Long imaginary-time runs utilizing PC use a reblocking analysis[141] to obtain statistical errors uncontaminated by autocorrelation. All calculations are run on NVIDIA GeForce GTX 1080 GPUs, with Intel Xeon E5-2620 v4 CPUs running at a maximum of 2.10GHz.

3.2.2 Timings

Employing an unrestricted HF trial for ph-AFQMC has been shown to produce very accurate energies for hydrogen chains near their equilibrium bond lengths[81]. Using a bond length of $1.880(2)$ Bohr as given by Density Matrix Renormalization Group (DMRG) in the cc-pVDZ basis, we compare propagation times using a single GPU card for an increasing number of hydrogen atoms. Sample propagation times for several variants of precision and means of decomposing the two-electron terms are shown in Fig. 3.3. For H_{60} DF is 2.0x faster than CD in sp and 1.5x faster

in dp. Sp is 2.1x faster than dp when DF is used, and 1.6x faster using CD. Generally we find that the relative speed-ups afforded by sp over dp, and DF over CD, increase with system size. The non-monotonicity of the propagation times vs system size is a unique and rather unexpected artifact of the GPU architecture, and we observe that the sp (dotted) and dp (lines) trajectories move together, suggesting a different treatment of sp and dp at the hardware level. The GPUs used in this work can perform sp and dp floating point operations at a maximum of 8876 and 277.36 GigaFLOP/s, respectively. Our observed speed-up going from dp to sp is significantly less than what these peak metrics would imply. This is because for the H chain sizes investigated here with single-determinant trials the GPU performance is not compute limited but rather bound by the device memory bandwidth. This suggests additional speed-ups can be expected for calculations of this type, and we plan to pursue further memory optimization in the near future.

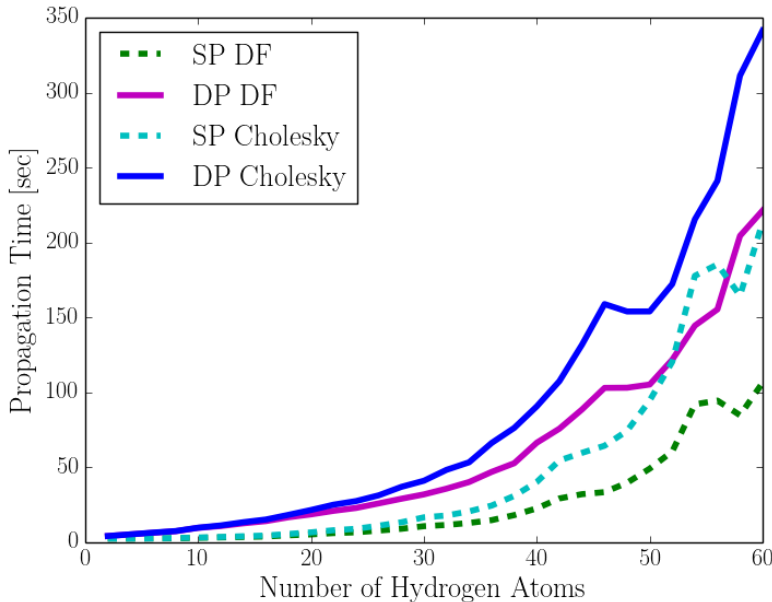


Figure 3.3: Propagation time using 1 GPU for hydrogen chains of varying lengths, comparing two types of two-electron integral decompositions, DF vs CD with a 10^{-5} cutoff, within both sp and dp. UHF trial functions are used, and 24 walkers are propagated for an imaginary-time segment of length 1 Ha^{-1} .

In Table 3.2 we benchmark the performance of our GPU implementation for multi-determinant trial functions with CD. We compare against our latest CPU code, which utilizes the same SMW algorithm but without the batching scheme and with cuBLAS kernels replaced by calls to equivalent

functions in the Intel Math Kernel Library (MKL). The energy algorithm implemented in the CPU code also utilizes the precomputed tensor shown in Eq. (3.2). Furthermore the CPU code defines the same C structure types for matrices in sp and dp, and uses analogous algorithms to, e.g., copy and exponentiate these matrix structures. We believe that for these reasons a fair comparison between our CPU and GPU codes can be made. The GPU-accelerated code in sp achieves large speed-ups ranging from 87.2x with two determinants to 670.1x with 1000 determinants, compared to our CPU code in dp. Importantly we find that the relative speed-up increases with the number of determinants present in the trial function. This is due to efficient batched processing in the evaluation of mixed-expectation values involving the trial wavefunction.

Table 3.2: Propagation times (in seconds) for an H_{50} chain with a varying number of determinants that comprise the trial wavefunction. We use CD with a 10^{-5} cutoff, and show the speed-up of a single GPU in sp over a single CPU in dp.

	$N_{det} = 2$	$N_{det} = 50$	$N_{det} = 100$	$N_{det} = 500$	$N_{det} = 1000$
GPU sp	99.9	105.3	111.8	158.0	218.7
CPU dp	8775.2	15019.7	22993.7	79713.3	148544.2
Speed-up	87.8x	142.7x	205.6x	504.6x	679.1x

To parallelize across GPU cards on a single node, we divide the total number of walkers into subsets which are independently propagated and measured on different GPU cards. We use Open Multi-Processing (OpenMP) to achieve shared-memory parallelization of the CPU threads, and to each CPU thread we associate a partner GPU device. Fig. 4.1 highlights the near-unity parallel efficiency of our implementation, defined as the multi-GPU speed-up over 1 GPU divided by the number of GPUs utilized.

To treat larger system sizes we have implemented a local memory strategy which spreads slices of the 4-dimensional tensors in Eq. (3.2) across 8 cards for the entire simulation. At the intervals where the energy is measured, the random walkers propagated on, e.g., GPU 0 are sent to GPUs 1-7 to compute the components of the two-electron energy derived from the elements stored locally on GPUs 1-7. This is done simultaneously for walkers on all GPUs, after which the energy components are gathered and tallied for each walker. The nodes utilized have 8 GPU cards each with 8 GB of RAM. Using DF and sp, this local memory-slicing algorithm enables us to treat systems as large

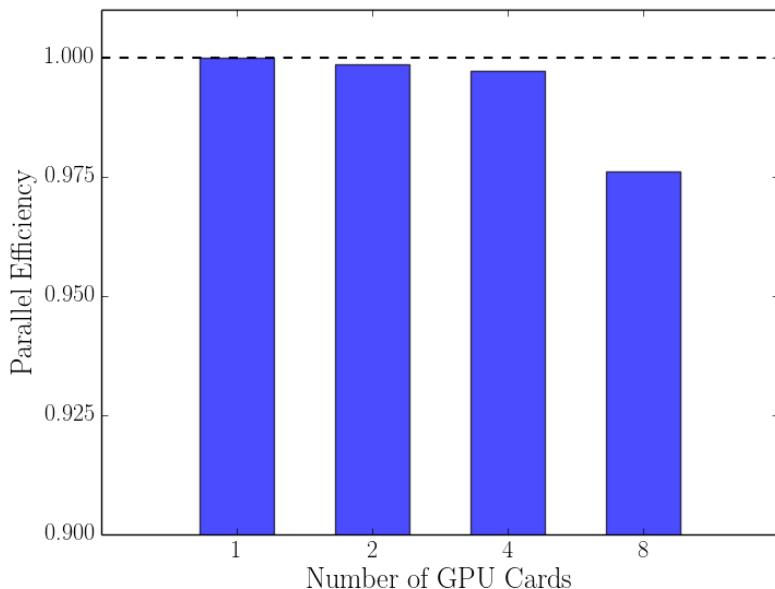


Figure 3.4: Parallel efficiency of our ph-AFQMC code illustrated on H_{50} . We use a CASSCF trial wavefunction with 44 determinants, and 800 walkers propagated for 0.5 Ha^{-1} with $\Delta\tau = 0.01 \text{ Ha}^{-1}$.

as H_{100} in the double-zeta basis ($M = 500$, $N = 100$).

3.2.3 Accuracy

To benchmark the accuracy of our algorithm when sp and DF are used, we compute the total energy of H_{50} in the cc-pVDZ basis with a bond length of 1.8 Bohr, and compare with results from a recent study[81] presenting data from state-of-the-art methods including ph-AFQMC, DMRG, and restricted CCSD(T), among others. DMRG is essentially exact for one-dimensional systems[194], and RCCSD(T) is expected to provide a high level of accuracy as the bond-length is near its equilibrium value[31].

Our GPU results are shown in Table 3.3 along with the previously published data. The differences in the total energies of sp vs dp for our GPU calculations are 1.2(9) mHa for DF and 0.7(8) mHa for CD. Importantly, both of these are smaller than the resolution required for chemical accuracy (1.6 mHa), confirming that for this system size we can take advantage of the hardware-optimized sp arithmetic on the GPU without incurring a significant loss of accuracy.

Table 3.3: Total electronic energies [Ha] of H_{50} at $R = 1.8$ Bohr in the cc-pVDZ basis. Propagation times [hours] are presented using 8 CPU/GPU pairs. We use 1000 walkers propagated for a length of 200 Ha^{-1} (including equilibration).

	Electronic Energy	Propagation Time
GPU sp DF	-125.2107(7)	14.6
GPU dp DF	-125.2119(6)	31.4
GPU sp Chol	-125.2239(6)	27.8
GPU dp Chol	-125.2246(5)	45.2
CPU dp Ref. 81	-125.2242(8)	
RCCSD(T) Ref. 81	-125.2067	
DMRG Ref. 81	-125.2210(1)	

DF produces about half the number of auxiliary-fields compared with CD (550 vs 1105), reducing propagation times by a factor of 1.9 for sp and 1.4 for dp. In terms of the resulting accuracy, it is well known that while the DF decomposition may not be sufficient to produce total energies within chemical accuracy, it can recover sub- mHa accuracy in the calculation of relative energies[200, 201]. Indeed, we find that for the total energy of H_{50} DF differs from CD with a 10^{-5} cutoff by 13.2(9) mHa in sp and 12.7(8) mHa in dp, respectively. Yet to put these errors into context we note in passing that DF ph-AFQMC in both sp and dp produces total energies for this system that are closer to the DMRG reference by $\sim 4 \text{ mHa}$ than RCCSD(T), known to many as the “gold standard” of quantum chemistry[202, 203].

To conclude this section we illustrate the capacity of DF and more aggressive CD truncation thresholds to recover chemically accurate energy differences for the deprotonation of methanol. Table 3.4 shows errors of $\sim 3 \text{ mHa}$ for the total energies of the neutral and deprotonated species; however the deviation of the energy difference from that of the most stringent CD cutoff is negligible, taking statistical errors into account. In this molecular case, compared to H_{50} , we find a more pronounced reduction in the number of auxiliary-fields, implying a $\sim 4x$ speed-up (vs $\sim 2x$ for the hydrogen chain) over CD with a 10^{-5} threshold.

Table 3.4: Accuracy of DF and CDs with various cutoffs for the deprotonation energy of methanol. Sp is used, and long imaginary-time trajectories are stabilized with PC. N_{AFs} denotes the resulting number of auxiliary-fields.

	N_{AFs}	MeOH	MeO ⁻	ΔE	$\Delta E - \Delta E_{CD10^{-6}}$
DF	142	-155.8503(4)	-148.5340(7)	0.6777(8)	-0.0004(10)
CD 10^{-2}	187	-155.8144(6)	-148.4985(4)	0.6772(7)	-0.0008(10)
CD 10^{-3}	385	-155.8531(4)	-148.5357(4)	0.6787(6)	0.0007(9)
CD 10^{-4}	471	-155.8542(4)	-148.5366(4)	0.6790(6)	0.0009(9)
CD 10^{-5}	617	-155.8544(4)	-148.5371(3)	0.6787(6)	0.0007(8)
CD 10^{-6}	855	-155.8533(5)	-148.5366(4)	0.6780(6)	0

3.3 IPs of Transition Metal Atoms

In this section we compute the IPs of the first-row TM atoms correlating *all* electrons, and compare the calculated ph-AFQMC results to experiment and previous electronic structure calculations.

3.3.1 Computational Details

Our computational protocol begins with a restricted (open-shell) HF calculation. We visually inspect the occupied orbitals of this solution to ensure that the electron configurations shown in Table 3.5 are obtained. For some atomic species, HF provides a qualitatively incorrect description of the single-particle orbital occupancies, requiring initialization from custom density matrices to converge subsequent HF calculations to the target ground-state configurations. We note that for the V⁺ cation the initial density matrix guess was constructed with the $L = 2$ orbital unoccupied. In all cases the canonical HF orbitals are used to initialize a restricted CASSCF calculation.

Table 3.5: Target electron configurations and spin-multiplicities ($2S + 1$), from Refs. 65 and 204.

System	Sc	Ti	V	Cr	Mn	Fe	Co	Ni	Cu	Zn
Neutral	$4s^2 3d^1$	$4s^2 3d^2$	$4s^2 3d^3$	$4s^1 3d^5$	$4s^2 3d^5$	$4s^2 3d^6$	$4s^2 3d^7$	$4s^2 3d^8$	$4s^1 3d^{10}$	$4s^2 3d^{10}$
Spin Mult.	2	3	4	7	6	5	4	3	2	1
Cation	$4s^1 3d^1$	$4s^1 3d^2$	$3d^4$	$3d^5$	$4s^1 3d^5$	$4s^1 3d^6$	$3d^8$	$3d^9$	$3d^{10}$	$4s^1 3d^{10}$
Spin Mult.	3	4	5	6	7	6	3	2	1	2

All ph-AFQMC calculations in this section use a CD cutoff of 10^{-4} . We utilize basis sets that have been optimized to account for scalar relativistic effects[204], and use the spin-free exact two-component approach[205, 206] to decouple the electronic degrees of freedom from the Dirac

equation. This approximation produces one-body terms which we simply add to the non-relativistic Hamiltonian in Eq. (1.19).

To compare calculations in finite basis sets to experiments we extrapolate the correlation energies to the CBS limit using two data points fit to $1/x^3$ ($x = 3, 4$ for TZ, QZ)[85, 207, 204]. We confirmed for a subset of the atoms that the inclusion of the aug-cc-pwCV5Z-DK energies did not significantly change the extrapolated results, consistent with Ref. 204. Following Ref. 65 and our own observation that the HF energies converge relatively quickly in this sequence of basis sets, we use the 5Z value for the CBS HF energies.

The Trotter error due to finite imaginary-time discretization can be extrapolated to 0 using progressively smaller time steps. Here we use $\Delta\tau = 0.005, 0.01$, and 0.02 Ha^{-1} . For Co through Zn we compared the CBS estimate from such an extrapolation with values from the smallest time step only, $\Delta\tau = 0.005 \text{ Ha}^{-1}$. In the latter approach we observe a substantial yet systematic cancellation of error, and CBS estimates of equivalent accuracy compared to the 3-point extrapolation approach are shown in Table 3.6. In light of this data we use only $\Delta\tau = 0.005 \text{ Ha}^{-1}$ for all calculations.

Table 3.6: Comparison of CBS IPs [eV] for Co, Ni, Cu, and Zn with $\Delta\tau \rightarrow 0$ vs $\Delta\tau = 0.005 \text{ Ha}^{-1}$ computed in sp with ph-AFQMC/PC.

	Co	Ni	Cu	Zn
Expt.	7.87	7.59	7.73	9.39
$\Delta\tau \rightarrow 0$	7.87(3)	7.61(3)	7.54(3)	9.33(4)
$\Delta\tau = 0.005 \text{ Ha}^{-1}$	7.89(3)	7.59(3)	7.55(3)	9.37(3)

Details of the CS procedure can be found in Chapter 2. In short, we run a set of independent calculations called repeats, each of which uses a distinct random number seed to propagate both the neutral and cationic species such that pairs of walkers sample the same auxiliary-fields. After an initial equilibration period, cumulative averages of the energy difference are computed along each of the imaginary-time trajectories, and are averaged among the set of repeats to obtain an estimate of statistical error. In the present case, stochastic error cancellation leads to a pronounced reduction in the variance of the IPs, and convergence at very short imaginary times can be achieved when the standard error drops below the target error tolerance and upon visual observation of a plateau in the measured quantity. We will show in the next section that this CS approach leads

not only to significant reductions in computational cost, relative to the uncorrelated approach, but also to systematically improved accuracy.

3.3.2 Results and Discussion

Tables 3.7 and 3.8 summarize our results for the all-electron IPs of the first-row TM atoms. We show values obtained from both PC and CS ph-AFQMC approaches, and compare with experimental and CCSD(T) values.

Table 3.7: Calculated ph-AFQMC IPs [eV] in the CBS limit computed with sp and $\Delta\tau = 0.005 \text{ Ha}^{-1}$, compared with experimental and CCSD(T) values. Experimental IPs have spin-orbit contributions removed.

	Sc	Ti	V	Cr	Mn
ph-AFQMC/PC	6.51(1)	6.71(2)	6.74(1)	6.75(2)	7.41(2)
ph-AFQMC/CS	6.52(3)	6.80(3)	6.74(3)	6.74(3)	7.45(3)
Expt.	6.56	6.83	6.73	6.77	7.43
CCSD(T)*	6.54	6.81	6.73	6.79	7.42

* Ref. 204

Table 3.8: Same as Table 3.7, but for atoms in the right-half of the row.

	Fe	Co	Ni	Cu	Zn
ph-AFQMC/PC	7.86(2)	7.89(3)	7.59(3)	7.55(3)	9.37(3)
ph-AFQMC/CS	7.89(2)	7.87(3)	7.61(2)	7.68(3)	9.37(3)
Expt.	7.90	7.87	7.59	7.73	9.39
CCSD(T)	7.89	7.88	7.59	7.72	9.37

Table 3.9: Number of active electrons and orbitals in the CASSCF trial wavefunctions for the cation/neutral species, and the number of determinants kept in the ph-AFQMC trial function accounting for 99.5% of the CI weight. For all species in this table the 3p electrons are active.

	Sc	Ti	V	Cr	Mn
Active Space	8/9e,16o	9/10e,16o	10/11e,19o	11/12e,16o	12/13e,18o [†]
N_{dets} TZ	146/224	240/442	366/751	303/271	423/584
N_{dets} QZ	143/439	293/388	300/903	92/262	852/1266

[†] Three 5p orbitals replaced by five 4d orbitals in the active space.

The active spaces employed for the neutral and cationic species in the TZ and QZ basis sets are described in detail in Tables 3.9 and 3.10. In general, the use of truncated CASSCF trial wavefunctions in ph-AFQMC involves subtleties that require careful consideration, since the CASSCF

Table 3.10: Same as Table 3.9, but for Cu and Zn with 99.0% of the CI weight retained.

	Fe	Co	Ni	Cu	Zn
Active Space	7/8e,18o	8/9e,13o	9/10e,13o	10/11e,18o	11/12e,13o
N_{dets} TZ	23/227	210/85	138/156	374/322	299/518
N_{dets} QZ	23/121	237/66	159/161	504/507	277/526

calculation itself can become an expensive pre-processing step when large active spaces are required, and the truncation breaks size extensivity. This approach is viable if the ph-AFQMC result converges quickly with trial wavefunctions generated from active spaces much smaller than the full Hilbert space. For atoms and molecules this is typically the case, and an internal validation procedure within ph-AFQMC can be employed involving a series of calculations using various active space and truncation cutoffs. In particular, for Fe-Zn we started by including the $4s$ and $3d$ electrons in an active space composed of 13 active orbitals. While the resulting truncated CASSCF trial wavefunctions produced sufficiently accurate ph-AFQMC/PC results in the CBS and $\Delta\tau \rightarrow 0$ limits for Co, Ni, and Zn, 18 orbitals were required in the case of Fe. The improvement in the IP resulting from the inclusion of a second shell of d orbitals in the CASSCF active space is a manifestation of the so-called “double-shell” effect[208, 209]. We find that this effect is less pronounced in the case of all-electron ph-AFQMC since the application of $e^{-\tau\hat{H}}$ to walker configurations can explore the space of excitations into virtual d orbitals even if such excitations are not represented in the trial function.

For the left half of the 1st row of transition metals in the periodic table, Sc-Mn, we designate the $3p$ electrons as active in addition to the $4s$ and $3d$ electrons. In an effort to maintain consistency (i.e. to include HF virtuals of similar character in the initial guesses for the CASSCF procedure) among all atoms in the row, for those in the left-half we start with 16 active orbitals. This produced accurate ph-AFQMC/PC results for Sc and Cr. For V we noticed a sharp drop in energy in both the neutral and cationic species going from 16 to 19 active orbitals; for Mn an accurate IP required the replacement of three $5p$ orbitals with five $4d$ in the CASSCF active space to accommodate the double-shell effect.

The case of Cu proves to be particularly challenging, and illustrates an additional merit of the

CS approach. A trial function with 18 active orbitals approaches the memory limit of traditional CASSCF solvers, but is still insufficient to produce results of the desired accuracy within ph-AFQMC/PC. With additional active orbitals, approximate CASSCF solvers utilizing DMRG[210] did converge, but only a subset of the resulting configurations and CI coefficients could be accessed with the current implementation of selected CI in PySCF. Even with moderate selection cutoffs, when such a wavefunction was used as a trial function in ph-AFQMC we found a significant increase in statistical error, in addition to larger deviations of the resulting IP from experiment.

In contrast to regular ph-AFQMC/PC, which stabilizes long imaginary-time trajectories, a key advantage of the CS approach is that averaging among independent repeats at short times allows for not only a vast variance reduction when the auxiliary fields are correlated, but also the ability to converge measurements of the energy difference *before* the full onset of the bias that results from the phaseless constraint. Even though the phaseless approximation is made after each time step, the walker weights at early times stay relatively closer to their true unconstrained values than at long times when the phaseless constraint has fully equilibrated. To illustrate this we plot the IP of Cu in the TZ basis at short imaginary-times in Fig. 3.5. At longer imaginary times (not shown) the CS IP appears to approach the ph-AFQMC/PC result (albeit with substantial noise due to the absence of PC), yet from 2-7 $H a^{-1}$ ph-AFQMC/CS unambiguously converges on an answer consistent with iFCI-QMC, which is expected to be very accurate here[65]. Moreover, this value after CBS extrapolation is within range of chemical accuracy with respect to experiment.

The case of Ni is also quite remarkable. Both CS and PC methods produce IPs consistent with the experimental value and each other in the CBS limit, however a detailed comparison with CCSD(T) values in each basis set, shown in Table 3.11, reveals that this agreement is due to fortuitous cancellations of error. While the CCSD(T) values approach the CBS limit from above, the ph-AFQMC/PC values approach the same value from below. ph-AFQMC/CS calculations, on the other hand, produce statistically consistent results with CCSD(T) in each basis and in the CBS limit.

For the case of Ti, having observed a quick equilibration time in the PC run with 16 active orbitals we chose to use CS as a much cheaper alternative to further increasing the size of the active

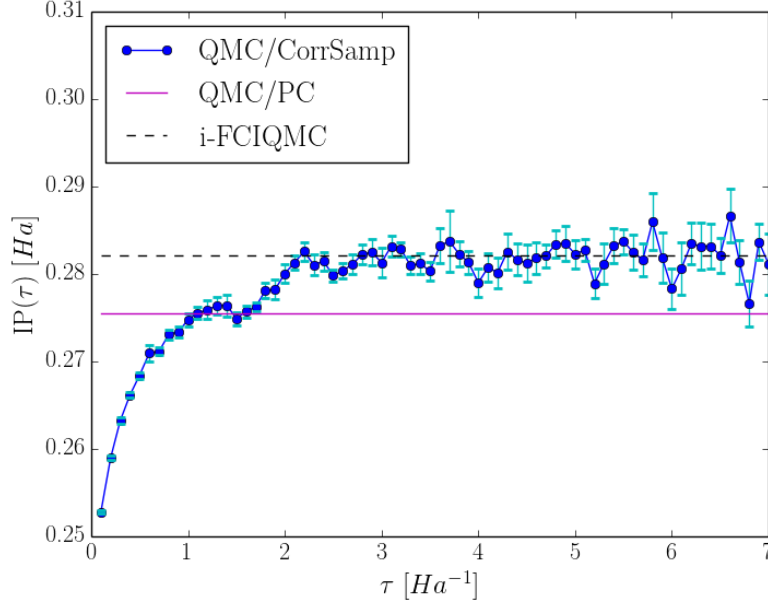


Figure 3.5: Comparison of the IP of Cu as a function of imaginary-time produced from ph-AFQMC/CS in the aug-cc-pwCVQZ-DK basis compared to the regular ph-AFQMC/PC result in the same basis. The i-FCIQMC result in the aug-cc-pVQZ-DK basis is indicated by a dashed line.

space. We note, however, that this alternative may not always be feasible, e.g. when a poor trial function results in long equilibration times. Generally, for all atoms in this work CS results exhibit equivalent or better accuracy compared to the conventional method of running ph-AFQMC calculations with PC. Moreover, the ability to consistently produce chemically accurate results while using sp is reassuring, given that the total energies involved in these calculations are on the order of $\sim -1000 Ha$. This implies that mHa energy scales require precision out to at least 7 significant figures, which would be stretching the typical capabilities of sp arithmetic in deterministic algorithms. In the future any differences in the sensitivity of stochastic vs deterministic algorithms to numerical precision could be explored further using half-precision, however on our current GPUs the peak performance of half-precision in GigaFLOP/s is 128x slower than sp and 4x slower than dp.

Table 3.12 shows the total propagation-times required to produce the final IPs in Tables 3.7 and 3.8, which account for calculations of the total energies of the neutral and cation in the TZ and QZ basis sets. 8-core CPU times are estimated by scaling the propagation time of a small

Table 3.11: Comparison of ph-AFQMC IPs for Ni, as obtained with CS and regular ph-AFQMC/PC, with CCSD(T) in triple- and quadruple- zeta basis sets and in the CBS limit[204]. The CCSD(T) values were obtained with a composite method, namely cc-pVxZ-DK results plus a core-valence correction, which is the difference in the cc-pwCVxZ-DK basis of CCSD(T) calculations with active spaces defined by $3s3p3d4s$ and $3d4s$ orbitals (x=T,Q). QMC results used the aug-cc-pwCVxZ-DK basis sets.

	CCSD(T)	Δ QMC/CS	Δ QMC/PC
TZ	7.68	7.68(1)	7.56(2)
QZ	7.63	7.64(1)	7.57(1)
CBS	7.59	7.61(2)	7.59(3)

20 walker system propagated for 1 Ha^{-1} by the required factors to reproduce the parameters of the GPU/PC calculations, i.e. 2000 walkers propagated in the TZ/QZ bases for 120/130 Ha^{-1} for Cr and Fe, and for 200/230 Ha^{-1} for Cu. We assume perfect parallel efficiency in projecting our single-core CPU estimates to 8-cores. Obtaining comparable error bars using CS with our GPU code required the propagation of 200 walkers in the TZ/QZ bases for 5/3 Ha^{-1} for Cr, 4/3 Ha^{-1} for Cr, and 10/6 Ha^{-1} for Cu. For Cr and Cu we use 16 repeats in both basis sets, and for Fe we found that only 5/8 repeats in TZ/QZ are needed. We note that the larger wall-times for Cu are due to 1) the larger number of both particles and determinants in the trial functions employed, and 2) the relatively poor trial function (compared to the exact ground state) which leads to longer propagation and equilibration times in the PC and CS methods, respectively. It may be the case that the relatively small atomic radius of Cu results in larger dynamical correlations compared to the rest of the atoms in the row, which are unaccounted for in the CASSCF trial wavefunctions (this explanation is consistent with the relative difficulty we encountered previously in calculating the electron affinity of the fluorine atom [211]). The corresponding speed-ups for these selected atoms are shown in Table 3.13.

Table 3.12: Total propagation times [hours] required to produce the final all-electron ph-AFQMC IPs in this work.

	Cr	Fe	Cu
Est. 8-core CPU PC dp	9780	6970	34200
8-card GPU PC sp	50.5	47.5	136.4
8-card GPU CS sp	2.4	0.9	6.0

We conclude this section with a few remarks. Currently we use a simple combining method[142] to implement PC. More sophisticated schemes are possible which may improve the statistical accuracy

Table 3.13: Speed-ups corresponding to the timings in Table 3.12. All CPU/GPU calculations use dp/sp respectively.

	Cr	Fe	Cu
GPU PC vs CPU PC	194x	138x	250x
GPU CS vs GPU PC	21x	53x	23x
GPU CS vs CPU PC	4100x	7700x	5700x

of the calculation. While this would slightly reduce the wall-times for the ph-AFQMC/PC method, the accuracy of the results with respect to experiments will be unchanged, since any bias due to PC vanishes when a large population (~ 2000 walkers) is used. At the time of writing, auxiliary basis sets optimized for the scalar relativistic Hamiltonian and DK basis sets used in this work are not publicly available. The ability to use a DF decomposition would certainly provide additional speed-ups, although its effect on accuracy remains to be tested for these TM systems. Finally, we note that the capacity of our GPU code to treat $O(1000)$ determinants in the trial wavefunction will likely enable the accurate study of many strongly-correlated systems. We anticipate that fewer determinants will be needed for metal-ligand complexes (in which the ligand is a non-metal), as TM atoms typically exhibit larger static correlation effects than most coordinated complexes. In addition, the use of symmetry constraints in the CASSCF calculations will greatly reduce the number of configurations in the CI expansions.

3.4 Conclusions and Outlook

We have designed a GPU implementation of ph-AFQMC for single- and multi-determinant trial wavefunctions which can drastically reduce the scaling prefactor in realistic electronic structure calculations with near-unity parallel efficiency. Our strategy utilizes new batched SMW and energy algorithms, along with the ability to use sp and the DF decomposition. We validate performance enhancements with ph-AFQMC calculations of linear chains of hydrogen atoms and the atomic IPs of Sc through Zn, finding speed-ups relative to the CPU in dp of two orders of magnitude and which increase with the number of determinants in the trial wavefunction. For H_{50} and TM IPs, sp is sufficient to produce accuracy on the scale of 1 kcal/mol with respect to exact methods and experiment, respectively. In this work we also demonstrate that our previously outlined CS

approach to ph-AFQMC enables both additional speed-ups of an order of magnitude, as well as the ability to converge measurements before the full onset of the bias due to the phaseless constraint. For all TM atoms, CS produces equivalent and often more accurate IPs, and in a fraction of the wall-time.

While we have shown that the code segment which scales most steeply with system size, i.e. the energy algorithm, has been implemented with a very high level of device utilization, we anticipate that calculations on small systems, and especially those employing single-determinant trial functions, can still be substantially accelerated by additional tuning iterations in which the various utilization metrics from nvprof are prioritized. However, it must be stressed that the optimal choice of parameters (e.g. grid and block sizes) and memory strategy to most efficiently utilize the device architecture will depend on the particular choice of hardware. For this reason we postpone these fine-tuning optimizations until a target application on a large-scale computing cluster is ascertained.

We are optimistic that in the near future the investigation of many large, realistic systems will be feasible with ph-AFQMC. In what follows we anticipate issues of scalability and describe the possible solutions we envision. While our current implementation exclusively uses NVIDIA hardware with CUDA and cuBLAS, it would be straightforward to adapt it to a more universal standard, e.g., Open Computing Language (OpenCL) and associated BLAS packages. To enable large-scale calculations that efficiently utilize available High-Performance Computing clusters, we have designed a simple and scalable scheme to parallelize *across* GPU nodes for cases in which all required data for a ph-AFQMC calculation can be stored on a single node. Once a small population is equilibrated, walker data can be copied to all available nodes and used to initialize independent trajectories on separate nodes, each with a different random number seed. These sub-trajectories can later be combined into a single trajectory from which averages and error bars can be obtained.

We note that our current memory limitation is rather artificial in the sense that GPU architectures and computing capabilities are improving at a rapid pace, suggesting that the memory capacity of GPU cards will continue to increase. Also, in later Chapters we report an implementation that uses Message Passing Interface (MPI) to extend our local memory-slice scheme to multiple nodes. At the time of writing, NVIDIA’s NVLink boasts transfer speeds of ~ 300 GB/sec between

Tesla V100 GPUs, and we expect that future improvements in device-to-device and host-to-device transfer speeds will further reduce the overhead associated with MPI communication or possibly other strategies utilizing CPU memory to store the high-dimensional tensors.

Combining the speed-ups due to the GPU and CS, we now have a robust and efficient computational protocol that is approximately three orders of magnitude faster than previous AFQMC procedures. This will enable routine ph-AFQMC calculations of a variety of chemically relevant properties with an unprecedented level of throughput and systematically improvable accuracy.

Chapter 4

On Achieving High Accuracy in Quantum Chemical Calculations of $3d$ Transition Metal-containing Diatomics with Auxiliary-Field Quantum Monte Carlo¹

4.1 Introduction

Transition metals play a vital role in a wide range of important processes in biology [212] and materials science [213]. Many redox and catalytic reactions, such as the water splitting reaction in Photosystem II[214], are dependent upon the electronic structure of specific transition metal-containing clusters. A precise understanding of the chemistry and physics of these processes at an atomic level of detail can only be elucidated by accurate quantum chemical calculations in conjunction with extensive experimental data. However, quantum chemical methods have had

¹Based on work published in *J. Chem. Theory Comput.* 2019, 15, 42346-2358.

great difficulty in the treatment of transition metal-containing systems[30, 215]. Even for small molecules, the accuracy of high level *ab initio* approaches for these systems has been far from clear. For larger systems, DFT has been the only viable alternative. Much has been learned from applying DFT to complex systems[216], but while in many cases surprisingly good quantitative results have been obtained, there are also cases where errors as large as 40 kcal/mol can be observed[35]. A benchmark quality quantum chemical methodology which can be scaled up efficiently to treat systems 30-100 atoms in size would be a transformative advance.

Validation of benchmark accuracy must start with molecules containing only a few atoms, as was the case for organic systems, where coupled cluster (CC) based approaches, predominantly CCSD(T), have been able to demonstrate accuracy to better than 1 kcal/mol, with steady, systematic improvement over the past 20 years [32]. For transition metals, the challenge is compounded by uncertainties in many of the experimental measurements used as relevant test cases, as is apparent in recent investigations using a variety of computational methods on small molecules[217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 227]. Focusing on CC methods, electronic excitations for atoms are well described by CCSD(T) calculations using large basis sets and correcting for relativistic effects [32]. However, even for problems involving simple diatomic molecules, such as the dissociation energy of NiH, there is considerable uncertainty as to the degree of accuracy that CCSD(T) methods can achieve [35]. Error bars in the experimental gas phase measurements of dissociation energies of transition metal-containing diatomics are reported to be as large as ~ 5 -10 kcal/mol in unfavorable cases (and for a few experiments may exceed that threshold) [35]. With this level of possible error, it is very challenging to carry out robust statistical assessments of various approaches, as was done successfully for organic systems using the G2 [140] and G3 [228] databases of Pople and co-workers.

Over the past decade, there have been a number of efforts to evaluate the accuracy of CC approaches for small transition metal-containing molecules. The most recent work over the past 5 years has focused principally on diatomic species. The electronic structure problem is still qualitatively more difficult than it is for atoms, but the minimal size of the system enables very high level theoretical methods to be applied on relatively large data sets, and the experimental errors are in

general more well controlled than for more diverse test cases (although severe individual problematic cases remain). In addition to the experimental uncertainty, a key issue that has emerged is that the CC numbers can vary considerably depending upon the details of the calculations. The treatment of relativistic effects, spin-orbit coupling, and basis set extrapolation can have large effects on the accuracy of predicted bond dissociation energies. Early work from this period did not necessarily utilize a complete treatment of such aspects. For example, Ref. 229 employed single point calculations only in the triple zeta basis set, without any basis set extrapolation. Subsequent work has established standard protocols (which we discuss in more detail below) which appear to be sufficient to handle these particular aspects of the problem to near-chemical accuracy[230, 231]. Nevertheless, significant discrepancies between theory and experiment remain, and have been challenging to analyze definitively.

The current state-of-the-art is well reflected in the recent work of de Oliveira-Filho and coworkers[35]. They consider the bond dissociation energies of 60 diatomic species, each consisting of one transition metal atom and one hydrogen or second or third row acceptor. Of these systems, 42 contain a first row transition metal, to which we will limit our consideration in the present work (we plan to consider higher row transition metals in subsequent work). This data set of diatomics is expanded in size as compared to earlier efforts along the same lines, e.g. the 3dMLBE20 data set of Truhlar and coworkers, which contains 20 molecules, 19 of which are included in Ref. 229. All of the test cases have available experimental results that are at least plausible, although the issues with uncertainty noted above remain. We adopt the data set of de Oliveira-Filho and co-workers as a starting point for our analysis in what follows, adding and subtracting a few cases based on consideration of the experimental results, as will be discussed in detail below. A larger and more diverse data set enables more robust conclusions to be drawn concerning the performance of quantum chemical approaches in thermochemical calculations. Calculated errors can vary dramatically among molecules that are apparently very similar, as can be seen by examining the performance of DFT methods in calculating atomization energies for molecules in the G3 database (222 molecules)[232]. While the present data set is in our view not sufficiently large or diverse to draw rigorous conclusions concerning benchmark quality (on the order of 1 kcal/mole mean absolute error (MAE) across the

entire range of first row transition metal chemistry), it does represent a reasonable place to start an assessment of whether a given method is a candidate for such performance, assuming that the experimental errors can be sufficiently well understood.

CC-based calculations are carried out in reference 35 at the state of the art level, carefully converging results to the CBS limit and incorporating core-valence and relativistic effects. In addition to single-reference (SR) CCSD(T) calculations, multi-reference (MR) CCSD(T) calculations are also reported. Such computations require nontrivial approximations, due to the potentially large computational expense incurred by the use of MR wavefunctions. Nevertheless, it is of great interest to observe the effects of attempting to employ a methodology that, in principle, represents a systematic improvement over CCSD(T), addressing the well known presence of multiple relevant low-lying states in the electronic structure of transition metals. The results presented in that work provide a qualitative picture of the accuracy of CC based approaches for transition metal-containing systems. In many cases, both the SR and MR approaches are within a few kcal/mole of the experimental value of the dissociation energy. In others, the MR calculation provides a dramatic correction to SR results that were in considerable disagreement with experiment, by as much as 14.6 kcal/mole. In still other cases, the MR results continue to exhibit large disagreements with experiment, up to 11.6 kcal/mole. For these remaining outliers, even at the best (MR-CCSD(T)) level of theory employed, the question remains as to the relative contribution of computational and experimental errors to the discrepancies. A reasonable conclusion to be drawn is that SR-CCSD(T) is not capable of benchmark quality results for transition metal-containing systems (in contrast to non-metal systems, where MAEs < 1 kcal/mole have been reported for a subset of the G2 database[233]).

QMC approaches are an alternative to the CC methodology, and have shown encouraging accuracy in the prediction of transition metal properties. DMC within the fixed-node constraint has been utilized to compute the dissociation energies of the 20 diatomics in the 3dMLBE20 data set[234]. Deviations as large as 10 kcal/mol between calculated and experimental values were encountered, due to uncontrolled biases arising from the use of pseudopotentials and single-determinant trial wavefunctions, and possibly to the use of erroneous experimental values. We note that the use of

multi-determinant trial wavefunctions, e.g. from selected CI, can produce sub kcal/mol accuracy for the dissociation energy of the FeS diatomic[235], however the high computational cost of obtaining the coefficients for the millions of required determinants in the trial wavefunction would make such calculations on a large set of molecules highly challenging, and systematic accuracy has not been demonstrated.

Another QMC approach that is, in principle, capable of achieving systematically improvable and benchmark-quality accuracy for transition metal-containing systems is the ph-AFQMC methodology [66, 67, 71, 70, 68]. In previous chapters we have described a number of technical advances which have demonstrated dramatic reductions in the computational requirements for ph-AFQMC calculations, while in some cases actually improving their accuracy and robustness. The first of these is the use of correlated sampling[211]. With correlated sampling, energy differences between two states are computed by sampling both states with the same set of auxiliary fields, leading to significant cancellation of error. This enables energy differences to be computed in a much shorter amount of propagation time and with fewer samples than would normally be required to obtain a given statistical error[211]. Furthermore, these measurements at short propagation times are often converged before the full accumulation of the errors associated with the phaseless approximation, thus yielding results that are closer to the unbiased, exact value[211]. The second advance is the development of an efficient implementation of ph-AFQMC on graphical processing units (GPUs), including the use of the Sherman-Morrison-Woodbury (SMW) algorithm to accelerate calculations using multideterminantal trial wavefunctions[236]. For problems where correlated sampling is applicable, the combination of these two techniques can reduce the computational effort by more than two orders of magnitude, enabling the method to be applied to larger systems, and also to substantially larger data sets. Further efficiency improvements are feasible (reducing both the scaling and the prefactor), leading to the possibility that ph-AFQMC will emerge as a scalable benchmark methodology for transition metal-containing systems.

In the present chapter, we apply our ph-AFQMC methodology to a subset of the diatomics considered in Ref. 35, specifically all those containing first row transition metals (44 test cases in all). In Chapter 3 we showed that ph-AFQMC yields excellent accuracy for the IPs of first row transition

metal atoms. This finding is a good starting point, but it is clear from previous efforts in the literature that diatomic dissociation energies are much harder to compute with kcal/mol accuracy[35], and that the validation problem is more challenging given the issues with the experimental data.

The first objective of this chapter is to address key methodological issues that are critical to achieving robust and accurate results with ph-AFQMC for diatomic dissociation energies. Firstly, we demonstrate that correlated sampling can be made to work well for heavy atom dissociation, building on previous work which only considered removal of a hydrogen atom[211]. We find that correlated sampling not only provides substantial reductions in computational effort, but is essential in obtaining accurate energetics for these systems. The ability to treat heavy atom dissociation substantially expands the domain of applicability of correlated sampling to a wide range of chemical and biological problems.

Secondly, in ph-AFQMC calculations it is essential to utilize a sufficiently “good” trial function. We explore CASSCF type wavefunctions[40] for this purpose, and take advantage of the fact that for these small systems the dissociation energies can be converged with respect to active space size, making our calculations effectively size-consistent. The ph-AFQMC calculations for the diatomic molecules in our test set used between 100 and 5700 determinants. Our efficient GPU implementation of the SMW approach is essential for the utilization of large multideterminant trial functions of this form while keeping the increase in computer time at only a small factor.

Thirdly, we investigate three different approaches to estimating the CBS limit. All strategies employ a ph-AFQMC calculation in the triple zeta basis, and two-point extrapolations based on MP2, CCSD(T), and entirely based on ph-AFQMC. MP2 extrapolation suffices for many, but not all cases. CCSD(T) extrapolation usually does better, if not similarly to MP2. For a subset of the cases which we found to be exceptionally difficult, we show that direct AFQMC extrapolation is consistently able to improve the MP2 and CCSD(T) results.

Fourthly, we include new experimental values published in Ref. 237, and we identify one case (ZnS) where we believe that the experimental result is problematic, i.e. outside the error bars reported in the experimental papers. The very large discrepancies of experiment with both state-of-the-art CC and QMC results, along with a detailed analysis of the experiments, lead us to believe

that the experiment is in error. Theory cannot evolve to benchmark status without such conclusions being drawn along the way. With an optimized methodology defined, and with an objectively chosen set of reference values, we find remarkably good agreement between the ph-AFQMC results and the experimental data (taking into account the experimental error bars). We compare the MR-CCSD(T) and CCSD(T) values reported in Ref. 35 to the reference values, and find that the CC methods display a number of large outliers (fewer for the MR-corrected version). We also analyze the performance of 10 DFT functionals, reported in Ref. 35. Assessment of DFT results has been a feature of many of the papers cited above; however, the accuracy of the assessment has been problematic due to the uncertain nature of the reference values.

Finally, we discuss computational efficiency and the feasibility of scaling up to larger systems. It is possible to parallelize AFQMC efficiently across a large farm of GPUs (we plan to report the results of such an implementation in the near future), so with sufficient computational hardware resources, AFQMC calculations with a large number of basis functions can be carried out in a reasonable wall clock time. Furthermore, significant improvements in the AFQMC algorithm are still possible, and likely will be necessary to handle grand challenge problems with the goal of achieving true benchmark status. As noted above, the generation of sufficiently good trial functions may turn out to be the leading challenge to be faced in this scale up effort.

This chapter is organized as follows. In Sec. 4.2, we provide computational details. In Sec. 4.3, we describe the extension of our correlated sampling approach to the computation of bond dissociation energies. Sec. 4.4 includes a discussion of the landscape of experimental methods. In Sec. 4.5, we present our results for the D_e 's of the $3d$ transition metal diatomics, and justify our selection of the reference values used in the comparative statistical analysis of the various computational methods. In Sec. 4.6, we offer concluding remarks.

4.2 Computational Details

We use PySCF[144] to obtain all inputs required by our ph-AFQMC calculations. To compute the trial wavefunctions used in this work, we first perform restricted (open-shell) HF calculations,

ensuring that the electronic configurations are consistent with the term symbols published in Refs. 35 and 65. Canonical HF orbitals are used to initialize restricted CASSCF calculations. The resulting wavefunctions are truncated such that the sum of the squares of the CI coefficients kept is $> 98\%$ in all cases, resulting in ~ 800 determinants on average.

We stress that the spin *and* orbital symmetries, which we enforce at the RHF level, cannot be overlooked[238, 239]. For example, the latter can change the computed D_e in TiH by a staggering 15 kcal/mol. The diatomic term symbols, active space specifications, and bond lengths are shown in Tables 4.1 and 4.2.

Table 4.1: Electronic States, Active Spaces, and Bond Distances used in our ph-AFQMC calculations for diatomics containing Sc through Mn. X/Y means that both active space configurations produced statistically equivalent results. The number in parenthesis is the experimental bond length.

	electronic state	CASSCF Active Space	R_e [Å] CC (expt)
ScH	$^1\Sigma$	10e18o	1.762 (1.7754)
ScO	$^2\Sigma$	13e15o	1.664 (1.6661)
ScF	$^1\Sigma$	14e15o	1.787 (1.787)
ScS	$^2\Sigma$	13e15o	2.132 (2.1353)
TiH	$^4\Phi$	10e18o	1.768 (1.777)
TiN	$^2\Sigma$	13e15o/7e18o	1.57 (1.5802)
TiO	$^3\Delta$	14e15o	1.617 (1.6203)
TiF	$^4\Phi$	15e15o	1.8311 (1.8311)
TiS	$^3\Delta$	10e18o	2.0827 (2.0827)
TiCl	$^4\Phi$	15e15o	2.2642 (2.2697)
VH	$^5\Delta$	12e13o	1.684 (1.730)
VN	$^3\Delta$	14e15o/10e17o	1.544 (1.5703)
VO	$^4\Sigma$	15e15o	1.5839 (1.5893)
VCl	$^5\Delta$	16e15o/10e16o	2.2273 (2.2145)
CrH	$^6\Sigma$	13e18o	1.6293 (1.6554)
CrO	$^5\Pi$	10e16o	1.6116 (1.615)
CrF	$^6\Sigma$	11e17o	1.776 (1.7839)
CrCl	$^6\Sigma$	17e15o/11e17o	2.1688 (2.194)
MnH	$^7\Sigma$	14e18o	1.727 (1.7309)
MnO	$^6\Sigma$	17e15o/11e18o	1.638 (1.6446)
MnF	$^7\Sigma$	18e15o	1.834 (1.836)
MnS	$^6\Sigma$	17e15o/11e18o	2.0633 (2.0663)
MnCl	$^7\Sigma$	18e15o/12e18o	2.2355 (2.2352)

Table 4.2: Electronic States, Active Spaces, and Bond Distances used in our ph-AFQMC calculations for diatomics containing Fe through Zn. X/Y means that both active space configurations produced statistically equivalent results. The number in parenthesis is the experimental bond length.

	electronic state	CASSCF Active Space	R_e [Å] CC (expt)
FeH	$^4\Delta$	9e18o	1.5478 (1.606)
FeO	$^5\Delta$	12e17o	1.612 (1.6164)
FeS	$^5\Delta$	12e17o	2.009 (2.0140)
FeCl	$^6\Delta$	13e17o	2.1751 (2.1742)
CoH	$^3\Phi$	10e15o/10e18o	1.5049 (1.5327)
CoO	$^4\Delta$	13e17o	1.6286 (1.5286)
CoS	$^4\Delta$	13e17o	1.9786 (1.9786)
CoCl	$^3\Phi$	14e17o	2.0749 (2.0656)
NiH	$^2\Delta$	11e15o/11e19o	1.4538 (1.4538)
NiO	$^3\Sigma$	14e17o	1.626 (1.6271)
NiF	$^2\Pi$	15e17o	1.733 (1.7387)
NiCl	$^2\Pi$	15e17o	2.0539 (2.0615)
CuH	$^1\Sigma$	12e15o/12e19o	1.4593 (1.4626)
CuO	$^2\Pi$	15e17o	1.709 (1.7246)
CuF	$^1\Sigma$	16e17o/10e19o	1.745 (1.7449)
CuS	$^2\Pi$	11e19o	2.051 (2.0499)
CuCl	$^1\Sigma$	16e17o	2.0498 (2.0512)
ZnH	$^2\Sigma$	13e15o	1.5899 (1.5935)
ZnO	$^1\Sigma$	16e12o	1.6989 (1.7047)
ZnS	$^1\Sigma$	16e17o	2.0427 (2.0464)
ZnCl	$^2\Sigma$	17e16o	2.1274 (2.1300)

Our ph-AFQMC calculations correlate *all* electrons (i.e. no frozen-core), and utilize the “hybrid” formulation of the algorithm[76]. With an imaginary-time step of 0.005 Ha^{-1} , walker orbitals are orthonormalized every other propagation step, and energy measurements are taken every 0.1 Ha^{-1} . We employ a cutoff of 10^{-4} for the Cholesky decomposition of the two-electron integrals. We have verified that these parameter choices result in biases smaller than the statistical error bar[236]. We use the aug-cc-pwCVxZ-DKH basis sets[204] and the spin-free exact two-component approach[205, 206] to account for scalar relativistic effects. For the 3dMLBE20 molecules, the combination of this level of theory and basis sets has produced good results for CC calculations[231].

The MP2-assisted CBS extrapolation protocol is detailed in Refs. 211 and 77. After a ph-AFQMC calculation in the triplet-zeta (TZ) basis, a CBS correction is obtained by extrapolating the correlation energies as computed with MP2 using the $\frac{1}{x^3}$ form, with $x = 3, 4$ [85, 207, 204]. We

employ a scaling factor, which is the ratio of the MP2 and QMC values in the TZ basis. For all diatomics we performed both restricted and unrestricted HF calculations to compute correlation energies, and choose the method which leads to a scaling factor closest to 1. Finally, following Ref. 65 and our own observation that the HF energies converge relatively quickly in this sequence of basis sets, we use the 5Z ($x = 5$) value for the CBS HF energies, and add this to the extrapolated correlation energy to arrive at our final result.

For the small molecules considered here, CCSD(T) calculations can be performed in large basis sets, and results have been made available in the Supporting Information of Ref. 35. To evaluate the reliability of the MP2-assisted protocol for transition metal-containing systems, we use the published CCSD(T) data to extrapolate our ph-AFQMC results to the CBS limit as follows: We take the difference between de Oliveira-Filho’s CCSD(T)(CV)/CBS estimate of D_e , as was obtained via $1/x^3$ extrapolation of the correlation energy at $x = \text{Q}, 5$ with respect to the restricted open-shell HF reference, and their value in the aug-cc-pwCVTZ basis. We then add this term to our ph-AFQMC result in the aug-cc-pwCVTZ-DKH basis. We estimate the statistical error in the CBS limit using that in the TZ basis combined with their $x = \text{T}, \text{Q}$ CCSD(T) values. We note that this procedure assumes that the optimal bond lengths and scalar relativistic contribution to the BDEs are independent of basis size, as is done in Ref. 35, among other works. The spin-orbit term, ΔSO , in Eq. (4.1) is taken from Ref. 35, in which values are computed using CASSCF in a QZ basis.

We emphasize that in both of these extrapolation approaches, AFQMC calculations are only performed in the TZ basis. In our view, this is a significant source of computational expedience, as the convergence of observables with the size of the CASSCF active space used in the trial function is expected to be slower in basis sets of increasing size. For large chemical systems with substantial multireference character we note that other methods such as CASPT2 or even ph-AFQMC with a single-determinant trial wavefunction can be used to compute a CBS correction.

For a select number of cases where we encountered significant discrepancies among our calculated methods and with respect to experiment, we perform ph-AFQMC calculations in both the TZ and QZ basis sets, and extrapolate to the CBS limit. We view this extrapolation protocol to be of the highest quality, and for the purposes of this paper we employ it as required.

All ph-AFQMC calculations use single precision floating point arithmetic (which we have verified to give consistent results within statistics as double precision calculations[236]) and were run on NVIDIA GeForce GTX 1080, Tesla P100 and V100 graphical processing units. Our code is parallelized with Message Passing Interface (MPI), and we observe excellent strong-scaling parallel efficiency, shown in Fig. 4.1 for the CoO diatomic in the QZ basis. Using 360 GPUs on 60 nodes of the Summit supercomputer, the parallel efficiency of our implementation is still 90%. This allows us to run large calculations in minutes, a capability not possible for traditional, non-stochastic quantum chemical methods.

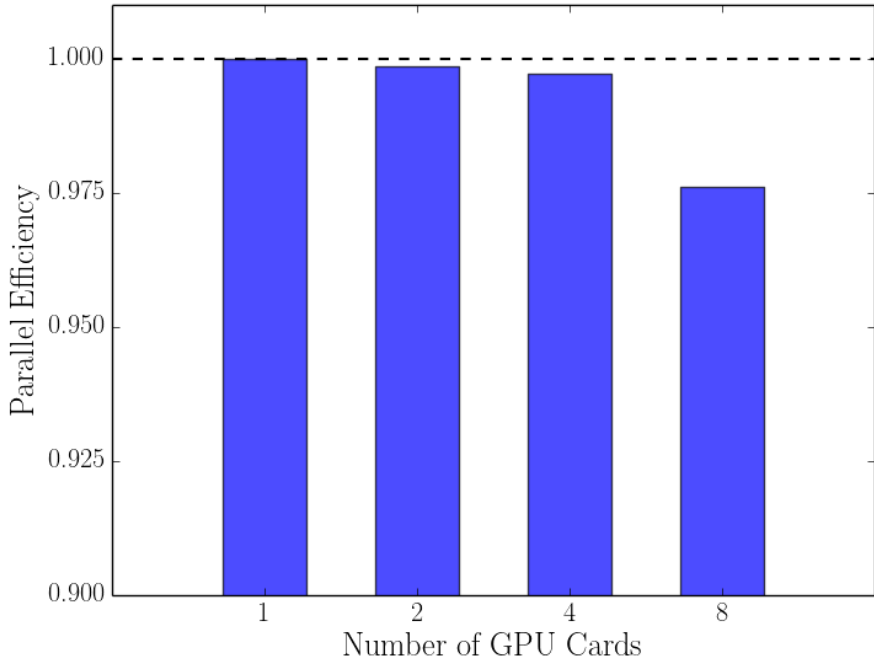


Figure 4.1: For a set of node counts on the Summit computing cluster, we plot the parallel efficiency, defined as the speed-up over a 1 node calculation divided by the number of nodes. Each node utilizes 6 NVIDIA V100 GPU cards.

Not all of the “experimental” bond lengths in Ref. 229 were actually obtained from experiments, so we use calculated bond lengths for the 3dMLBE20 set at the CCSD(T)/CBS level of theory, taken from Ref. 230. When the experimental bond lengths, as given in Ref. 35, differ by more than 0.01\AA from the CCSD(T) values, we ran ph-AFQMC calculations at both bond lengths. In the future we will consider using geometries optimized within ph-AFQMC[240].

The active spaces utilized to generate the trial wavefunctions were initially chosen with the intention of realizing a maximal cancellation in the systematic error associated with finite active spaces. That is, the number of active electrons (orbitals) of the isolated metal and ligand should sum to the number of electrons (orbitals) in the metal-ligand dimer. However, for these diatomic systems we found it possible to converge the BDE with respect to increasing active space sizes, and prioritized this convergence at times over the balanced protocol described above. The size of the active spaces is limited by the current CI module in PySCF, yet we were able to employ active spaces with up to 19 orbitals, allowing for satisfactory convergence throughout.

For the isolated ligands, we confirmed the convergence of the energy from ph-AFQMC/PC calculations with increasingly large active space sizes, and found that in all cases except for the sulfur and fluorine atoms, CASSCF did not lower the energy by more than a milliHartree with respect to unrestricted Hartree Fock (UHF). Hence, we use UHF for these cases, CASSCF(6e,8o) for S, and CASSCF(7e,16o) for F. The latter is consistent with our previous work[211], in which we found that an active space of this size was necessary to obtain chemically accurate electron affinities for the F atom.

4.3 Correlated Sampling for BDEs

Recently we have introduced a correlated sampling (CS) approach for quantities involving energy differences which is capable of reducing computational prefactors[211] and in some cases the severity of the phaseless approximation[236]. In this section, we show that significant reductions in statistical errors are obtained not only for hydrogen abstraction reactions, as shown previously, but also for bond breaking events between a transition metal and a heavier ligand atom.

For diatomic molecules consisting of a metal (M) and ligand (L), the following equation for the bond dissociation energy is employed:

$$D_e = E(M) + E(L) - E(ML) + \Delta SO. \tag{4.1}$$

We use CS to compute $E(M) - E(ML)$, where in the former term so-called “ghost” basis functions

centered at the coordinates of L are added, but without the nuclear charge or electrons from the ligand. We note that the basis set superposition imbalance[241], if any, that is introduced at the TZ level vanishes in the CBS limit. $E(L)$ is computed using the population control (PC) method detailed in Refs. 211 and 142, and ΔSO is the calculated energy difference due to spin-orbit coupling taken from Ref. 35.

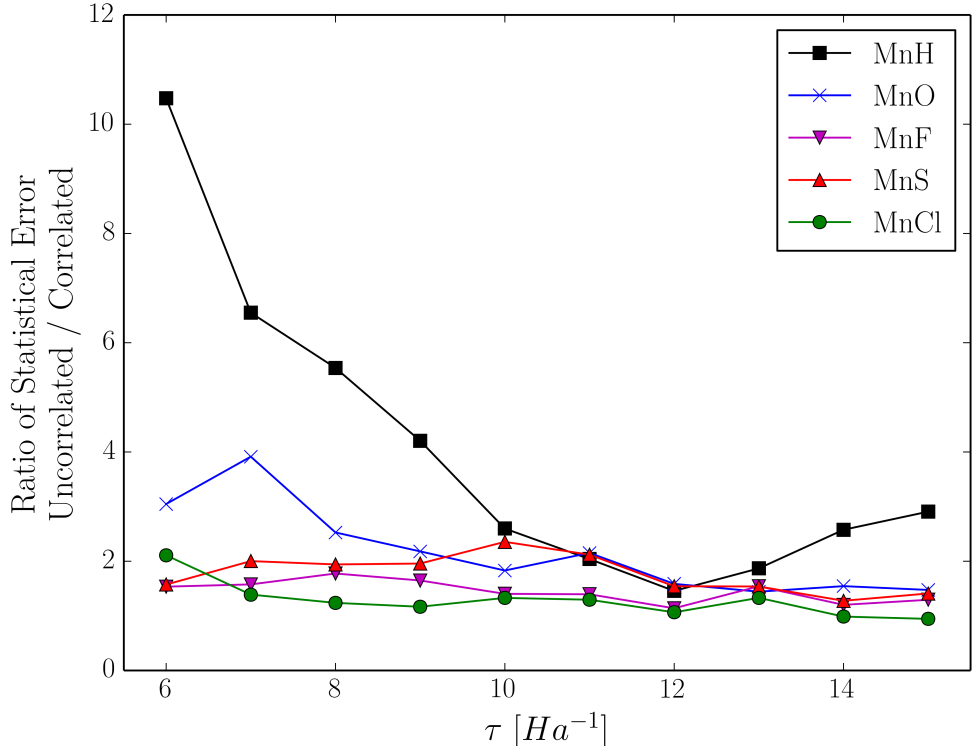


Figure 4.2: Ratio of the standard errors, as a function of imaginary time, resulting from ph-AFQMC calculations with correlated vs uncorrelated sampling, for the five Mn-containing diatomic species.

The reduction in statistical error, compared to the uncorrelated sampling approach, is shown in Fig. 4.2 for the Mn-containing diatomics in our set (the saving in computational time is, as usual, given by the square of the error ratio). We find that the effect of correlated sampling is largest for the hydride ligand, with decreasing noise reduction efficiency as the ligand atomic number increases.

As shown in Chapter 3, CS results exhibit equivalent or *better* accuracy compared to the conventional method of running ph-AFQMC calculations with PC employing the same trial wavefunctions, for the ionization potentials of first row transition metal atoms. We find the same behavior in the

calculation of the BDEs in this study. For ph-AFQMC calculations of the D_0 of MnCl in the aug-cc-pwCVTZ-DKH basis, PC and CS yield values of 86(1) and 82(2) kcal/mol, compared with the experimental value of 80(2). For MnS PC, CS, and experimental D_0 values are 78(1), 71(1), and 70(3). Thus, in light of significantly improved computational efficiency *and* accuracy, we use CS for the $E(M) - E(ML)$ part of all the BDE calculations in this chapter.

4.4 Experimental BDEs

The experimental BDEs given by Truhlar and co-workers for the 3dMLBE20 set were determined either from experimental enthalpies of formation (TiCl, VH, VO, VCl, CrO, CrCl, MnS, MnCl, FeCl, CoCl, NiCl, CuCl, ZnH, ZnO, ZnS, and ZnCl), which have their own error bars, or from direct measurements of D_0 at 0 K (CrO, FeH, CoH, and CuH)[229]. Both were converted to D_e via scaled DFT calculations of zero-point energies.

In the follow-up work by Dixon and co-workers, the experimental D_e ’s for the hydride diatomics were replaced by values derived from hydride transfer experiments[230]. These experiments involve the following reaction



By combining the energy of this reaction (referred to in Ref. 230 as $E_{threshold}$) with the ionization potential (IP) of the metal (M), the electron affinity (EA) of the hydrogen atom, and the heterolytic bond dissociation energy of a C-H bond in an organic molecule (typically a hydrocarbon or amine) ($BDE_{heterolytic}(R-H)$) the BDE of the metal hydride is obtained:

$$BDE(MH) = BDE_{heterolytic}(R-H) - IP(M) - E_{threshold} - EA(H). \quad (4.3)$$

In these measurements, $IP(M)$, $EA(H)$, and $E_{threshold}$ are known relatively accurately. The $BDE_{heterolytic}(R-H)$ values have more uncertainty, but Dixon and co-workers confirmed the experimental quantities with G3MP2 calculations[230]. However, the quantity $BDE_{heterolytic}(R-H)$ -

$E_{threshold}$ is not constant for various R, often varying up to 10 kcal/mol. Therefore, Dixon and co-workers give two values: one that is the average of all the measurements with different R's and one measurement that is the closest to their calculated CCSD(T)-level value. [230] In the present work, when referring to the values from Dixon and co-workers we only consider the measurements derived from the former method (averaged values). Stanton and co-workers use similar experimental values for VH and CrH using hydride transfer reactions and also confirm the validity of the $BDE_{heterolytic}(R-H)$ using their own HEAT345-Q protocol[231]. These experimental values for the D_e may be an underestimate of the true D_e as the $E_{threshold}$ may be affected by competition with side reactions, which may explain some of persistent disagreement between theory and experiment[231].

Dixon and co-workers also replaced the values for the chlorides with direct mass spectrometric measurements using Ag-M-Cl vapors, the value for VO with a direct measurement using a Eu-V-O system, and the value for ZnO with a mass spectrometry experiment, and the value for other compounds, particularly ZnS using different, fully experimental, heats of formation using more accurate Joint Army-Navy-NASA-Air Force (JANAF) thermochemical values[230]. Their selection of best experimental values for the 3dMLBE20 set are listed in Table 5 of Ref. [230].

Recently, Morse has reviewed his group's progress in obtaining highly precise measurements using resonant two-photon ionization spectroscopy to obtain predissociation thresholds that are equivalent to the BDE's of those diatomics with a very high density of states[237]. This experiment works by increasing the frequency of the incoming laser pulse until the excited state cation can no longer be detected (the predissociation threshold), because it has dissociated from the excited state's rovibrational state to the ground-state separated atom limit via other unstable excited states. Thus, this technique requires there to be a high density of states to ensure the method is precise and accurate, which precludes study of diatomics containing Cr, Mn, Cu, or Zn. For molecules where this technique is appropriate, it is more precise than many high-temperature Knudsen effusion measurements of gas-phase equilibria and guided ion beam mass spectrometry. Morse also shows that the measurements are also amenable to testing via a thermodynamic cycle with other precise measurements[237]. In our study we convert D_0 to D_e using the ZPE data in Ref. 35.

In the study by de Oliveira-Filho and co-workers only spectroscopically-derived data is referenced[35], which may explain their omission of the VH molecule.

4.5 BDEs of Transition metal-containing Diatomics

In this section, we show our computed values of D_e for the set of diatomic molecules containing first row transition metal atoms, and compare the calculated ph-AFQMC results to experiments, and to the MR-CCSD(T) calculations performed in Ref. 35.

As we expect the finite basis set error to be less sensitive to the method used to calculate the correlation energy, we examine various strategies to minimize the compute time required to reach the CBS limit. The MP2-assisted and CC protocols produced CBS results that are very similar in the majority of cases. Importantly, an extreme value (with respect to unity) of the scaling factor, which is the ratio of the correlation energies at the MP2 and ph-AFQMC levels in the TZ basis, can serve to flag an unreliable MP2 extrapolation. For instance, the scaling factors for CrH and NiH are both larger than 2.5. The resulting MP2-assisted predictions for D_e in the CBS limit for NiH was the furthest from experiment that we observed.

The CC method of CBS extrapolation is more reliable, and we checked that all scaling factors are between 0.8 and 1.2. When the resulting CBS BDE value still differed substantially from experiment, i.e. for ScH, TiS, CrO, CrF, CrCl, CoH, NiH, NiO, NiCl, CuO, and ZnS, we performed ph-AFQMC calculations in both the TZ and QZ basis sets, and extrapolated to the CBS limit. In all of these cases, except for CrO and ZnS, this procedure produced results consistent (within experimental uncertainty and statistical error) with experimental values. These two cases will be discussed in detail below. As control cases, we utilized this procedure for two cases, CoO and CrO, for which the CC CBS estimate is already accurate. As expected, the pure ph-AFQMC CBS results produced essentially the same values.

For all species in which the CC and experimental bond lengths, shown in Tables 4.1 and 4.2, differed by more than 0.1 Å, we performed ph-AFQMC calculations at both bond lengths. No significant differences in D_e resulted, so the results corresponding to the CC values are shown.

In Figs. 4.3 - 4.12 we present our results by metal species. We show the selection of experimental values and the MR-CCSD(T) results from Ref. 35 by default. We also show experimental values from the predissociation technique of Ref. 237 when available, and note their extremely small uncertainties. We considered the experimental selections in Ref. 230 as well, and plot their choice only when it is not within error bars of the corresponding value from Ref. 35.

We would like to highlight that the measurements in Ref. 237 were published *after* the ph-AFQMC calculations in this work were completed. It is rather remarkable that in all six relevant cases - ScS, TiN, TiS, VN, and FeS - our best QMC results (QMC/CCcbs and, when available, QMCCbs) are consistent with the newly available, and presumably of higher quality, experimental data. While demonstrating consistency with past results is obviously a necessary phase in the development of any new method, we are certainly encouraged by the predictive capability already shown by our methodology.

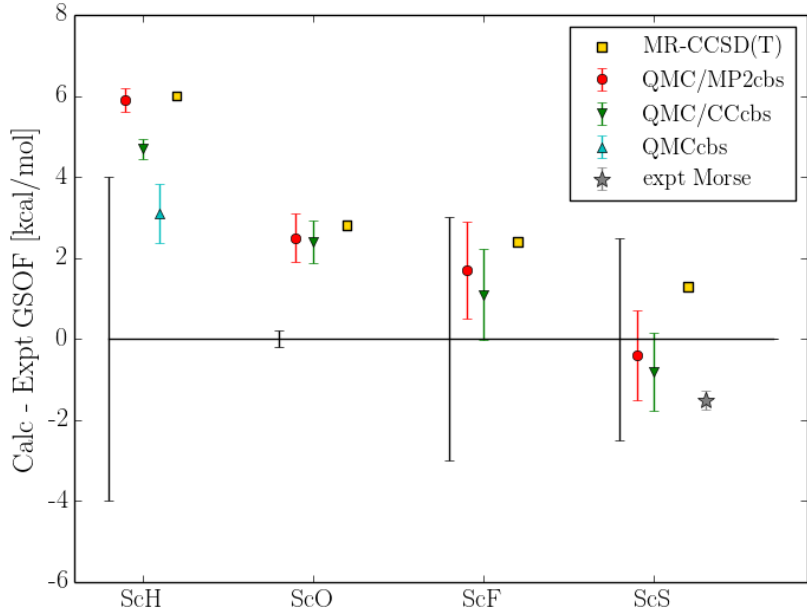


Figure 4.3: Deviations [kcal/mol] of various calculations and alternate experiments (when relevant) from the experimental values used by de Oliveira-Filho (GSO) and co-workers in Ref. 35. For calculations and experiments, error bars represent statistical error and quoted experimental uncertainties, respectively.

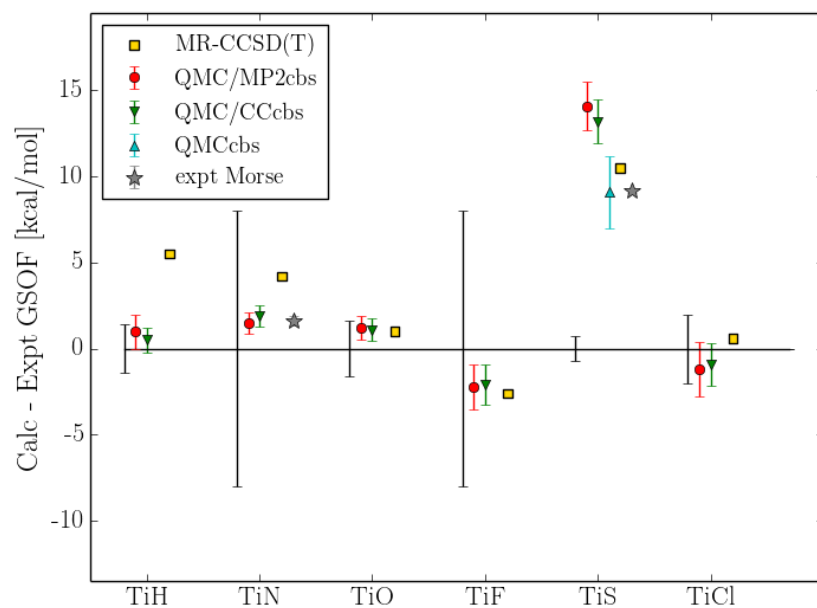


Figure 4.4: Same as Fig. 4.3, but for Ti-containing diatomics.

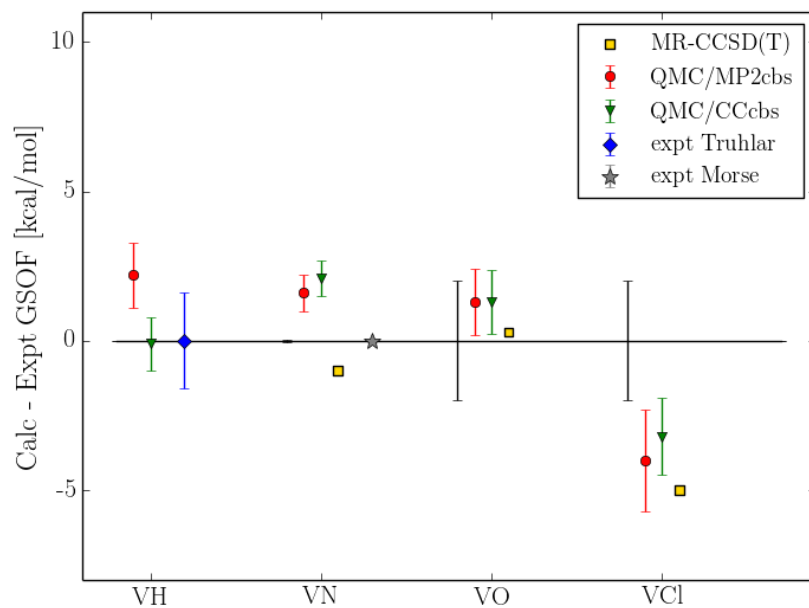


Figure 4.5: Same as Fig. 4.3, but for V-containing diatomics. VH was not considered in Ref. 35, and we show the experimental result selected by Truhlar and co-workers in Ref. 229.

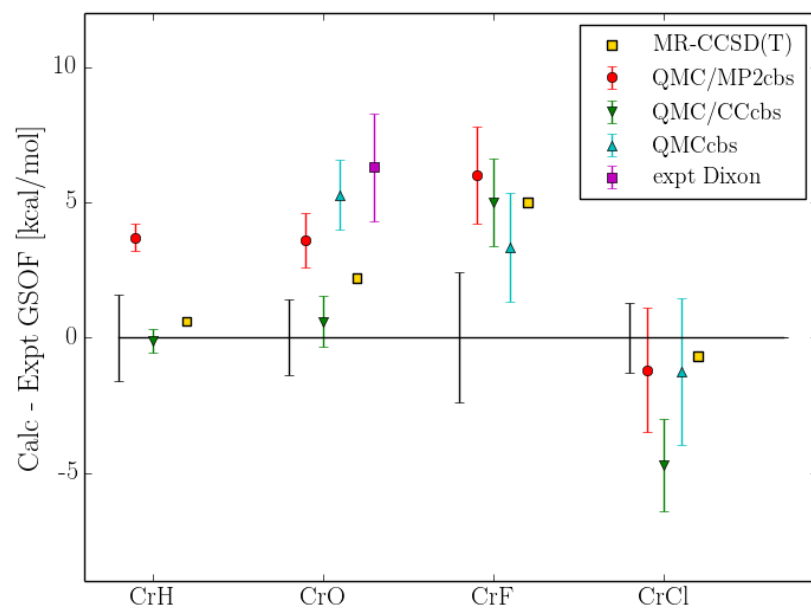


Figure 4.6: Same as Fig. 4.3, but for Cr-containing diatomics. For CrO we also show the experiment selected by Dixon and co-workers in Ref. 230, since it is not consistent with that chosen by de Oliveira-Filho and co-workers[35], given the reported uncertainties.

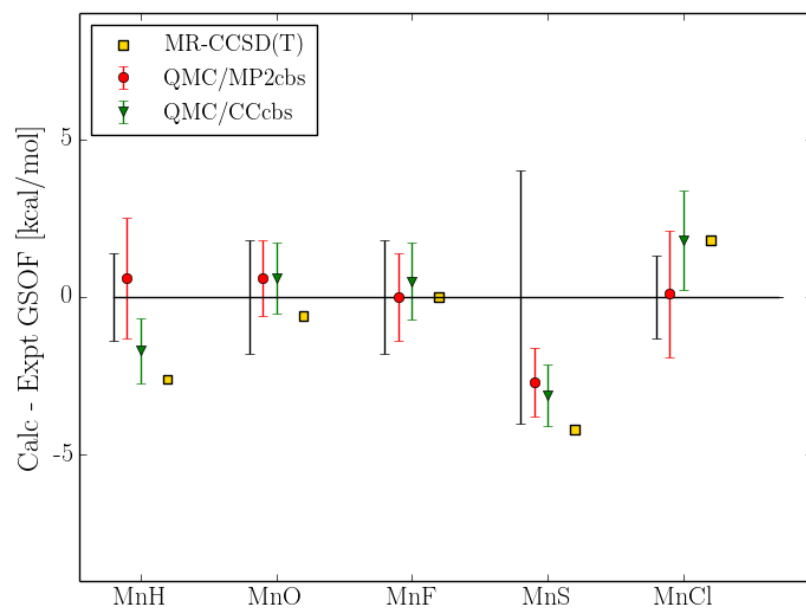


Figure 4.7: Same as Fig. 4.3, but for Mn-containing diatomics.

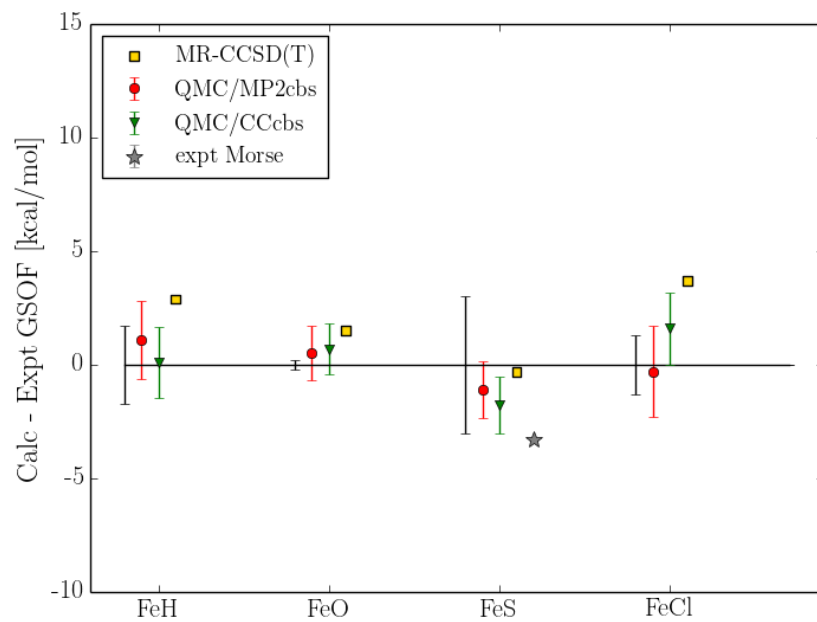


Figure 4.8: Same as Fig. 4.3, but for Fe-containing diatomics.

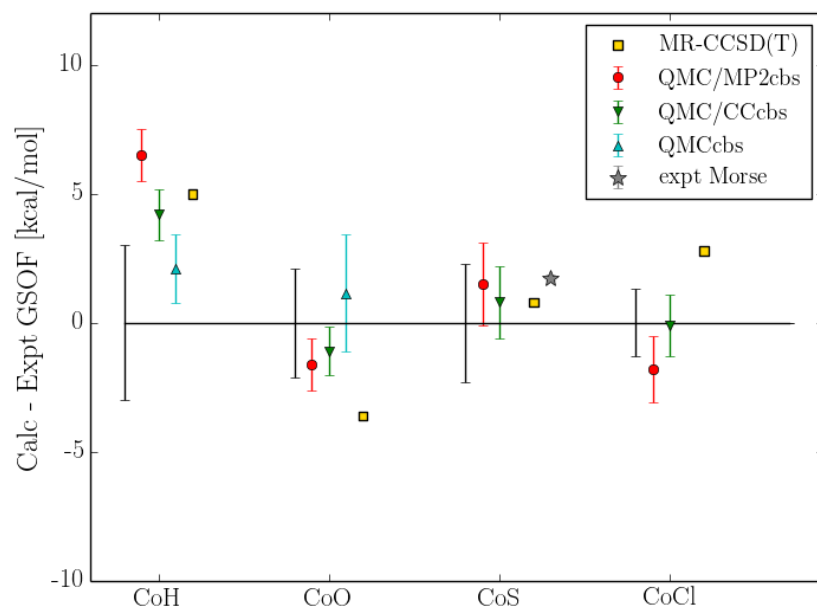


Figure 4.9: Same as Fig. 4.3, but for Co-containing diatomics.

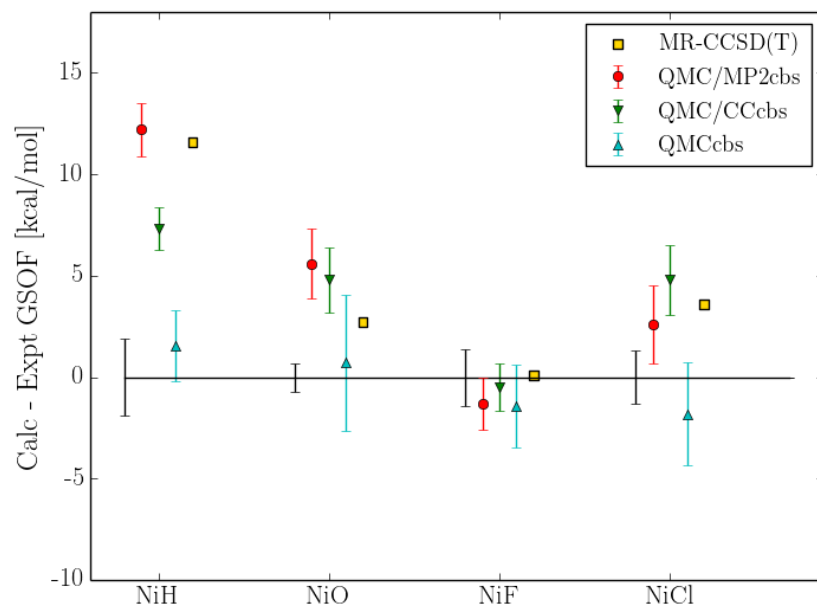


Figure 4.10: Same as Fig. 4.3, but for Ni-containing diatomics.

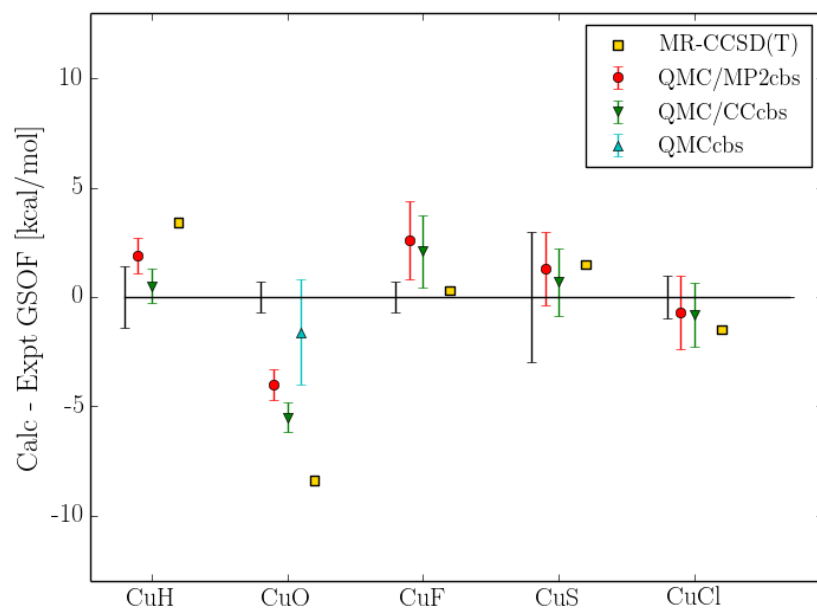


Figure 4.11: Same as Fig. 4.3, but for Cu-containing diatomics.

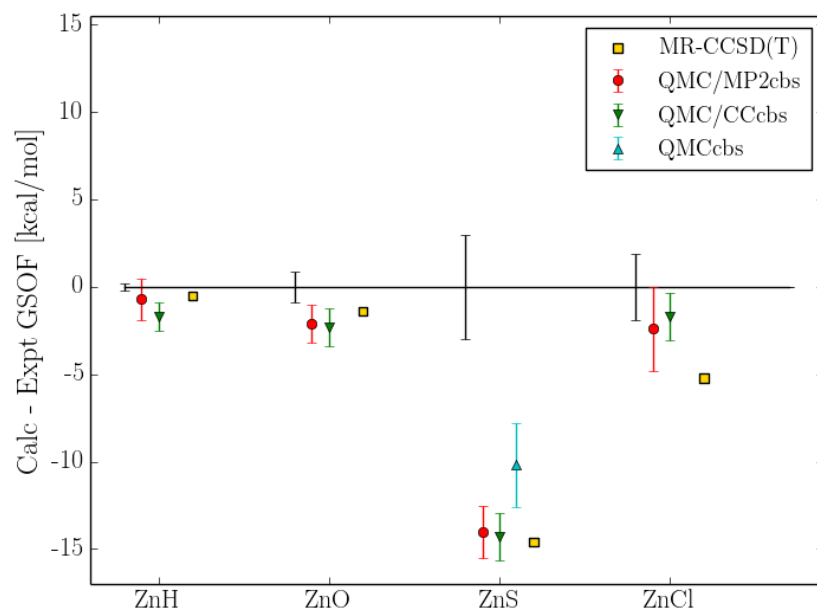


Figure 4.12: Same as Fig. 4.3, but for Zn-containing diatomics.

We now proceed to highlight a number of notable cases:

de Oliveira-Filho and coworkers do not consider VH, presumably because the experimental dissociation energy has not been measured spectroscopically. In Fig. 4.5 we show the experimental value selected by Truhlar and coworkers[229] which is derived from enthalpies of formation. We prefer this number over those proposed in Ref. 230, since the disparity of values shown in Ref. 230 (52.5 ± 4.4 and 56.1 ± 1.5 kcal/mol) illustrates the sensitivity of D_e to the R-dependent quantity $BDE_{heterolytic}(R-H)$, as discussed in Sec. IV.

For CrO, even though our QMC/CCcbs result agrees with the value chosen by de Oliveira-Filho and co-workers, our QMC/MP2cbs and QMCcbs value deviate from that value by 4(2) and 5(2) kcal/mol, respectively. However, we note two alternative experimental values: 111.1 ± 2 reported by Dixon and co-workers, and 110 ± 2 reported using ion molecule reactions in Ref. 242. Both of these experiments are consistent with our best method, QMCcbs, which gives 110.1 ± 1.3 kcal/mol.

For CoS, single-reference CCSD(T), the “gold standard” to many, is in error by a sizable 14.6 kcal/mol. The large MR correction, in the TZ basis, brought the CCSD(T) value within error bars of experiment. ph-AFQMC calculations, which also show excellent agreement with experiment, strengthen our confidence in the reliability of the experimental measurement.

For CoO, CCSD(T) again makes a large error of 10.9 kcal/mol. The MR correction, however, is not sufficient this time, as the MR-CCSD(T) result is still off by 3.6 kcal/mol. de Oliveira-Filho and coworkers suggest that an MR correction in a larger basis may fix this. We also note that the scalar relativistic correction is rather large for this molecule, at -4.5 kcal/mol. All CBS extrapolation variants of our ph-AFQMC calculations produce values in agreement with the experimental value.

For NiH, CCSD(T) is 8.5 kcal above experiment, and MR-CCSD(T) brings the BDE further away from experiment by an additional 3 kcal/mol. Our QMC/MP2cbs result has a deviation comparable to that from MR-CCSD(T), and while QMC/CCcbs provides a substantial improvement, it is still off by 7.3 kcal/mol. Only QMCcbs brings NiH within the error bars of experiment. We note that our CS calculation in the QZ basis, with the CAS(11e,19o) trial, resolved two measurable plateaus. Since it is not feasible to include more orbitals in the active space in this case, we measure the first plateau, as done (and justified) previously in Ref. 236.

CuO is another difficult case for all but our best full QMC treatment. CCSD(T) is off by 9.3 kcal/mol, and MR-CCSD(T) by 8.4 kcal/mol. QMC/MP2cbs, QMC/CCcbs, and QMCCbs differ from experiment by 4, 5.5, and 1.6 kcal/mol.

4.5.1 Selection of Best Values and Comparisons of ph-AFQMC with DFT and CC methods

With the goal of obtaining a robust, objective statistical comparison among ph-AFQMC, CCSD(T), MR-CCSD(T), and various DFT functionals, we now lay out a protocol to construct a reliable test set of reference experimental values.

We begin with a baseline set of experiments selected by de Oliveira-Filho and co-workers[35]. When possible, we substitute the high quality experimental results obtained via the predissociation technique of Ref. 237. We then consider the independently selected best experimental values in the work of Dixon and co-workers[230], and remove from the test set any case for which there is disagreement, considering error bars, with the original experiment from either de Oliveira-Filho or Morse. This situation arises only once, for CrO, and we therefore exclude it from our comparison.

For ZnS, the BDEs computed with QMC/MP2cbs, QMC/CCcbs, and MR-CCSD(T) are all in excellent agreement, with a value roughly 15 kcal/mol below experiment, which is 49.1 ± 3 kcal/mol as suggested by de Oliveira-Filho and co-workers[35] and derived from the work of Marquat and Berkowitz [243] and de Maria et al[244]. This experimental value is in stark contrast to the value of 34.3 ± 1.0 kcal/mol suggested by Truhlar and co-workers[229]. The latter value, as pointed out by de Oliveira-Filho and co-workers[35], is derived from a thermochemical analysis by von Szentpaly[245], which used theoretical, not experimental, values provided by Peterson et al[246]. Papakondylis, who used various theoretical methods to calculate the D_e of ZnS[247], pointed out that the aforementioned experimental papers used outdated values for the equilibrium bond lengths and frequencies as compared to more recent measurements[248, 249], bringing the older experimental measurements into doubt. Therefore, more experimental investigation into ZnS should be done to see if indeed the calculations are correct. For these reasons, we omit ZnS from our analysis, and simply report to the community our prediction of 38.8 ± 2.4 kcal/mol, as obtained with

our QMCcbs method. We believe that the quantitatively consistent results provided by completely independent high-level electronic structure approaches will eventually prove to predict the correct experimental value, but this speculation will have to wait for further experiments to confirm or refute.

The MAE of all QMC methods, the CC approaches, and 10 DFT functionals, with respect to the experimental set of 40 molecules as selected above, are given in Table 4.3.

Table 4.3: Mean Absolute Error and Maximum Error on D_e shown for AFQMC, CCSD(T), icMR-CCSD(T), and DFT methods vs. the experiments selected in Ref. 35 and, when possible, Ref. 237. For reasons justified in the text, we omit VH, CrO, and ZnS from the comparative statistical analysis. In all, our test set contains 41 diatomics. All DFT calculations are in the aug-cc-pVQZ basis with DKH corrections. DFT and CC values taken from Ref. 35. All values are in kcal/mol.

Method	MAE	Max Error
PBE	15.66	40.90
BP86	14.78	38.17
TPSS	12.83	31.00
M06-2X	12.05	37.95
BLYP	11.64	37.10
M06-L	8.44	21.85
M06	7.06	22.25
PBE0	4.73	21.85
B3LYP	4.45	23.45
B97	3.70	17.25
CCSD(T)	2.84	17.35
icMRCCSD(T)	2.76	11.60
QMC/MP2cbs	2.3(4)	12(2)
QMC/CCcbs	2.1(4)	7(2)
QMC/MP2+QZcbs [†]	1.5(4)	4(3)
QMC/CC+QZcbs [†]	1.4(4)	3(3)

[†] Includes QMC TZ/QZ CBS extrapolations when available

We note that the treatment of spin-orbit effects at the CASSCF level does not include dynamic correlation, which can in some cases be very large (e.g., -3 kcal/mol for NiCl, -3.1 for NiO, -2.4 for NiF). For reasons such as this, DeYonkers and co-workers have suggested that “chemical accuracy”

for transition metal species is ± 3 kcal/mol[250]. None of the DFT functionals considered in this study meets this criterion. As is well known, DFT results are highly dependent on the exchange-correlation functional used, with MAEs ranging from 3.7 to 15.7 kcal/mol. The highest level of accuracy is obtained with the B97 functional, and our data suggests that it should be chosen in DFT studies of similar transition metal chemistries.

The results of this work would argue that Truhlar’s original claim, that CCSD(T) and DFT produce comparable accuracy, must be qualified. A head to head comparison of the B97 functional in a large basis set with state-of-the-art CCSD(T) in the CBS limit shows that while both exhibit equally large maximum errors (~ 17 kcal/mol), the MAE of the latter is slightly, but significantly, lower. That said, the accuracy of DFT depends entirely, and perhaps unsystematically, on the functional employed, and we note that in comparison to the majority of functionals, CCSD(T) should be preferred assuming one has adequate computing capacity.

For the CCSD(T) approach and its MR-corrected variant, our analysis gives MAEs of 2.84 and 2.76 kcal/mol, and maximum errors of 17.35 and 11.6 kcal/mol, respectively. A robust benchmark method for transition metal chemistry, in our view, cannot make such large errors for individual cases. That is *not* to say that CCSD(T) is necessarily unfit for benchmark applications. Specifically, both the CCSD(T) protocol and the MR-CCSD(T) results we present, as performed in Ref. 35, involve a number of assumptions that may lead to suboptimal accuracy. Chief among them are the additivity assumptions involving the core-valence and scalar relativistic corrections. In the CC protocol, the CBS limit is estimated with the non-relativistic Hamiltonian. Then, a relativistic correction using the DKH Hamiltonian in a TZ-level basis is added. Similarly, the MR-CCSD(T) values shown simply add a multi-reference corrections computed in the TZ basis without core-valence effects treated explicitly, and without the DKH Hamiltonian and corresponding basis sets. Indeed, the relativistic corrections can be quite large, e.g. -8.1, 6.4, and -7.8 kcal/mol for NiCl, CoCl, and NiF, respectively. In such cases, among others, it is plausible that the additivity assumptions mentioned above break down.

Another potential source of error is that the MR-CCSD(T) calculations have not been converged with respect to active space size, likely due to the high computational expense associated with such

a procedure. Indeed, the full-valence (and sometimes smaller) active spaces may be insufficient for cases in which excitations into high-lying virtual and/or from low-lying occupied orbitals contribute significantly to the correlation energy.

Thus, these points suggest that more accurate CC results are possible, in principle, but only if one is willing to bear the high computational expense required to carry out a more rigorous computational protocol.

Turning to our ph-AFQMC methods, we first notice that QMC/MP2cbs, with an MAE of 2.3(4) kcal/mol and maximum error of 12(2) kcal/mol, is of comparable quality to, or arguably slightly more robust than, MR-CCSD(T). This is remarkable given that the latter involves CC calculations, which scale as the seventh power with system size, in QZ and 5Z basis sets, while the former requires a ph-AFQMC calculation in the TZ basis only, followed by a relatively inexpensive two-point MP2 extrapolation. The near-perfect parallel efficiency of the QMC calculation, and its acceleration on graphical processing units, are advantages that are not enjoyed by traditional CC implementations.

QMC/CCcbs achieves notable reductions in both the MAE and maximum error, at 2.1(4) and 7(2) kcal/mol, respectively. We note that for larger systems, using localized orbital implementations of CCSD(T) can drastically reduce the computational cost, while preserving systematic improvability with regard to localization errors. For systems with substantial MR character, methods such as CASSCF of selected CI supplemented with perturbation theory can be used to replace CCSD(T) to perform the CBS extrapolation.

When the 13 cases for which we performed CBS extrapolations entirely with ph-AFQMC are taken into account, the maximum error of our best method, QMC/CC+QZcbs is reduced to 3(3) kcal/mol, with an MAE of 1.4(4) kcal/mol. For larger systems, the QZ extrapolation option becomes relatively more advantageous, given its accuracy at low-polynomial scaling.

4.6 Conclusions and Outlook

In summary, we have computed the D_e of 44 $3d$ transition metal-containing diatomic molecules with ph-AFQMC. We describe the extension of a recently developed correlated sampling approach to the calculation of bond dissociation energies, and report improvements in both efficiency and accuracy compared to uncorrelated calculations. In order to assess the robustness of various CBS extrapolation techniques, and moreover to compare our ph-AFQMC results to the DFT, CCSD(T), and MR-CCSD(T) calculations performed in Ref. 35, we carefully assemble a set of reference experimental values via the following unbiased protocol. We use the experimental values selected by de Oliveira-Filho and co-workers, and, when available, predissociation measurements from a recently published work by Morse and co-workers. VH was omitted in Ref. [35], thereby depriving us not only of a consistently-chosen experimental reference but also of consistently-computed DFT and CC values, so it is not included in our statistical analysis. We omit cases where the experimental values selected by de Oliveira-Filho and Dixon are significantly different (with non-overlapping error bars), which necessitates the removal of CrO from the test set. Finally, we remove ZnS from the analysis on the grounds of concerns regarding the validity of the reported experimental number, which has been voiced previously in the literature, and emboldened substantially by the observation that ph-AFQMC, CCSD(T), and MR-CCSD(T) all disagree with the experimental value, and roughly agree with each other given statistical error bars.

Using this set of reference values, we assess the accuracy of our ph-AFQMC calculations alongside previously published results from 10 DFT functionals, CCSD(T), and MR-CCSD(T). The results are plotted in Fig. 4.13.

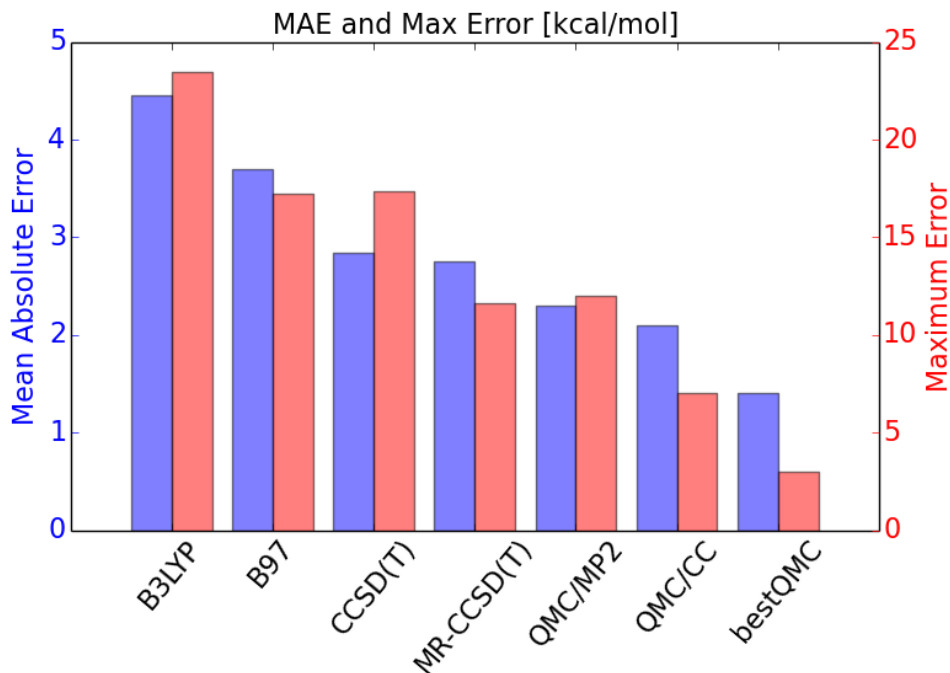


Figure 4.13: Statistical summary of the accuracy of bond-dissociation energy predictions, for a representative set of computational methods.

We find that of the DFT functionals, B97 performs best, and suggest its use for future DFT studies of transition metal-containing systems. We find that CC methods, while more accurate on average than DFT approaches, are not suitable benchmark methods for these systems due to the persistence of outliers with errors in excess of 10 kcal/mol. We take advantage of the systematic improvability of the ph-AFQMC method to attain high-quality predictions for these diatomic systems, and experiment with various cost-effective CBS extrapolation methods utilizing MP2 or CC. The final MAE of our best calculations is 1.4(4) kcal/mol, with maximum error of 3(3) kcal/mol. We would like to draw particular attention to the need for more robust experimental determinations of the dissociation energy of ZnS (as discussed before) and CrO, as, with regard to the latter, our most reliable ph-AFQMC calculation predicts 110.1 ± 1.3 kcal/mol, which is in agreement with two other published experiments[230, 242], but not the one put forth by de Oliveira-Filho and co-workers.

Although they are composed only of two atoms, these transition metal systems exhibit very

complex electronic structures, with a wide range of both static and dynamic correlation, core-valence and relativistic phenomena. The presence of many competing, low-lying states is common in transition metal-containing systems, e.g. in the Ni atom[251] and FeS[235], and we have shown that our ph-AFQMC protocol is capable of constraining calculations to targeted, experimentally-observed angular momenta and spatial symmetries that characterize the ground states.

That our QMC calculations can achieve such high accuracy is even more remarkable given that the trial wavefunctions used to implement the phaseless constraint utilize between 100 and 5700 determinants (~ 800 on average). However, obtaining CASSCF wavefunctions with sufficiently many active electrons and orbitals will become a challenge when larger systems are considered. We stress the need to experiment with alternative trial wavefunctions, the choice of which will likely depend on the target application. We are optimistic that explorations into more efficient descriptions of dynamic correlation and ways to exploit the locality of entanglements will lead the way toward scalable trials for accurate ph-AFQMC calculations.

Chapter 5

Singlet-Triplet Energy Gaps of Organic Biradicals and Polyacenes with Auxiliary-Field Quantum Monte Carlo

5.1 Introduction

The energy gap separating the lowest-lying singlet and triplet states of a molecule is an important property relevant to many chemical processes. For example, light absorption by chlorophyll in Photosystem II can produce triplet states which in turn react with triplet oxygen to produce short-lived and highly reactive singlet oxygen[252]. Additionally, the relative energetics of first-excited singlet and triplet states in dopants utilized in organic light-emitting diodes governs the efficiency of such devices, and is a useful parameter for the design of light-emitting electronics[253]. In the field of photocatalysis the singlet-triplet (ST) gap is directly relevant to a variety of redox reactions[254]. In addition the ST gap is a quantity of crucial importance for determining the energetic feasibility of the optical processes known as singlet fission[255] and upconversion[19]. In the former process, a

single photon produces a high energy singlet state which eventually splits into two triplet excitons; in the latter, two triplet excitons annihilate or fuse to form a high energy emissive singlet. In certain cases the ST gap can be probed under cryogenic conditions via phosphorescence measurements, however for a large number of relevant molecules direct measurement is not possible.

The electronic structures relevant to these types of applications can be complicated by the presence of biradical character in one or more of the involved spin-states. Biradicals are molecules in which two valence electrons can occupy two degenerate but spatially distinct molecular orbitals[256] (the species are referred to as biradicaloids if these two orbitals are nearly-degenerate, but we do not make this distinction here). The many possible electronic configurations give rise to their capacity to exhibit remarkably specific chemical reactivity[257, 258, 259, 260, 261]. Singlet states can be characterized as either closed-shell or open-shell. In the former, one of the two valence states is doubly occupied, which is typically the case, e.g., in carbenes. These species can simultaneously display Lewis base and Lewis acid character, and thereby undergo concerted addition reactions, e.g. with alkenes, to produce stereospecific products. Open-shell singlet states[262] are characterized by single-occupancy of each of the two valence states, as in all triplet states due to Pauli exclusion. Such open-shell molecules yield products of mixed stereochemistry.

The electronic structure community has witnessed the development of promising theoretical methods for computing the ST gap in biradical molecules. This is a challenging problem for traditional single-reference *ab initio* methods, e.g. Hartree-Fock (HF) and Density Functional Theory (DFT), as a minimal quantum-mechanical description of the wavefunctions corresponding to all singlet states and one triplet state ($M_s = 0$) in the two electron two orbital model is necessarily a superposition of two electronic configurations[23, 24], thus requiring more than one Slater determinant. Even for simple chemical species such as NH and O₂, which have triplet ground states, the singlet spin-configuration is notoriously difficult to describe. Accurate predictions are further complicated by the requirement that static and dynamic electron correlation be well balanced, with spin-states of both multiplicities treated on equal theoretical footing. Biradical systems have been studied with DFT[263, 264] and fractional spin variants[265, 266], Generalized Valence Bond theory[267, 268, 269, 270], spin-flip (SF) methods[23, 271, 272, 273], Complete Active Space

Self-Consistent Field (CASSCF) with 2nd order perturbation theory[274], multiconfiguration pair-density functional theory (pDFT)[275, 276, 277, 278, 279], Coupled Cluster (CC) methods[280, 281, 282], doubly electron-attached equation-of-motion CC theory[283, 284], spin-extended Configuration Interaction (CI) with singles and doubles[285], incremental Full CI[286], Difference Dedicated CI[287, 288, 289], the particle-particle random phase approximation (pp-RPA)[24, 290], the Density Matrix Renormalization Group (DMRG)[291], Yamaguchi spin projection[292] and its recent combination with orbital-optimized MP2[293].

We will use ph-AFQMC to compute ST gaps in this Chapter. While imaginary-time projection is most frequently used to yield ground-state properties of a system, the formalism has also been used to accurately compute low-lying excited states of materials[294], molecular diatomics[295] and dipole-bound species[296]. These calculations rely on the fact that eigenfunctions of the Hamiltonian are orthogonal. In practice, when the exact eigenfunctions are unknown, ph-AFQMC calculations use a so-called “trial wavefunction” to project out orthogonal components that may be sampled along the imaginary-time random walks. For example, for the molecules with triplet ground-states relevant to this work, the energy of the lowest-lying excited singlet state can be sampled using a trial wavefunction with $\langle S^2 \rangle = 0$, due to its near-orthogonality to the true triplet ground-state. In the limit of an exact trial wavefunction, this symmetry-constraining approach would be exact.

In what follows we will show that an unrestricted single-determinant trial wavefunction is capable of accurately describing multi-reference biradical species. This is a significant result because the computational cost of CI and CASSCF calculations, despite many recent advances[297, 298, 299, 300], scales exponentially with system size and thus such methods are infeasible as trial wavefunctions for ph-AFQMC. As an example, the number of π electrons in the polyacene series is $4n+2$ ($n=1,2,\dots$ for benzene, naphthalene,...), which typically must be included in the active space for accurate MCSCF-based predictions. Typical CI solvers can handle up to ~ 16 active electrons and orbitals, which would be insufficient for $n > 3$.

We view the ability of an electronic structure theory to accurately describe biradicals as a prerequisite for future studies of large, typically conjugated systems that catalyze photochemical pro-

cesses such as upconversion. After showing that a spin-projected approach to ph-AFQMC[75] with unrestricted single-determinant trial wavefunctions produces accurate results for a set of strongly biradical small molecules and three benzyne isomers, we then illustrate the scalability of our approach by taking a first step toward relevant and relatively large photocatalytic molecules, namely the polyacene series including naphthalene, anthracene, tetracene, and pentacene. These molecules are well-studied both experimentally (with measurements available in the literature for $n=1-5$, though not beyond) and theoretically, as they have myriad applications in organic electronics (see references in the first paragraph of Ref. 290). While biradical and polyradical character is predicted to be responsible for the instability of hexacene and longer acenes[291, 301], we focus on the $n=2-5$ molecules since they are representative of the majority of molecules in our target class of photocatalysts[254]. In fact, $n=3,4$ are known to perform upconversion[19], and $n=3-5$ are known singlet-fission catalysts[255].

This chapter is organized as follows. Details of our computational approach are described in Sec. 5.2. In Sec. 5.3 we compute ST gaps for a set of 13 small organic molecules which have singlet states of highly biradical nature, and compare with experimental measurements. Next we examine ortho- meta- and para- isomers of benzyne, and show that very high accuracy can be obtained with both CASSCF and single-determinant trial wavefunctions using a basis set of moderate size. Having shown that a single-determinant trial wavefunction combined with a spin projection technique is capable of accurately describing multi-reference biradical species, we proceed to compute ST gaps with ph-AFQMC for the increasingly large (but not necessarily biradical) systems naphthalene, anthracene, tetracene, and pentacene, and compare our results with state-of-the-art electronic structure theories and experimental measurements. Finally, we predict the ST gaps for a set of anthracene derivatives which are potential annihilators for optical upconversion.

5.2 Computational Details

All calculations utilize an imaginary time step of $\Delta\tau = 0.005Ha^{-1}$. Walker orthonormalization, population control, and local energy measurements are carried out every 2, 20, and 20 steps,

respectively. We utilize a modified Cholesky decomposition of the electron repulsion integrals (ERIs)[77], with cutoffs of 10^{-5} for the small molecule biradicals, and 10^{-4} for polyacenes, $n=2-4$. For pentacene ($n=5$) we use density fitting with the Weigend Coulomb-fitting basis set[199] to reduce the memory requirements of the calculation, while preserving high accuracy in energy differences[236].

All calculations utilize single-precision floating point arithmetic, though we note that inputs such as the trial wavefunction, one-electron integrals, and decomposed ERIs are obtained using double-precision. For ionization and bond-dissociation energies of transition metal atoms and diatomics, respectively, this yielded very high accuracy while reducing the computational cost compared to double-precision calculations[236, 302]. For pentacene, the largest molecule considered in this work, we verified with separate calculations in the STO-3G basis that single- and double-precision calculations gave statistically indistinguishable results.

For the small molecule biradicals, we use the aug-cc-pV x Z basis sets[148], with $x=T,Q$. Unrestricted HF (UHF), restricted HF (RHF) and its open-shell variant (ROHF), unrestricted Kohn-Sham DFT with the B3LYP functional (UB3LYP), and CASSCF trial wavefunctions are obtained using PySCF[144]. ST gaps are extrapolated to the complete basis set (CBS) limit using exponential and $1/x^3$ forms for the mean-field and correlation energies, respectively, as detailed in, e.g., Ref. 236.

For the benzyne isomers and polyacenes we report ST gaps in the cc-pVTZ basis[147], primarily for computational expedience, since pentacene has nearly 900 basis functions. For hydrocarbon systems, full CBS extrapolation typically alters the triple-zeta results by ~ 1 kcal/mol or less. This empirical finding is consistent with previous studies of polyacenes using other wavefunction methods which show little basis set dependence[280, 303, 290].

Our ph-AFQMC calculations utilize the spin-projection technique detailed in Ref. 75. The walkers are initialized with RHF for singlets and ROHF for triplets, which have $\langle S^2 \rangle$ of exactly 0 and 2, respectively. With an appropriate form of the Hubbard-Stratonovich transformation this can ensure that the single-particle imaginary-time propagator preserves spin-symmetry despite the use of a (possibly) spin-contaminated unrestricted trial wavefunction.

In what follows, we propose a simple empirical protocol, “AFQMC/U”, to guide the selection of an optimal unrestricted orbital set, among UHF and UB3LYP orbitals, to be used as the trial wavefunction in AFQMC calculations of potentially biradical systems. UHF orbitals are used by default, unless $\langle S^2 \rangle_{singlet}^{UHF} > 1.1$ or $\langle S^2 \rangle_{triplet}^{UHF} > 2.1$, in which case UB3LYP orbitals are utilized. Our rationale is based on the observation that wavefunctions constructed from unrestricted Kohn-Sham orbitals are known to exhibit less spin-contamination compared to UHF solutions[304]. To justify the selected thresholds, we note that for a perfect singlet biradical, as exemplified by the dissociated H_2 system, UHF provides a correct description of the dissociation limit, and $\langle S^2 \rangle_{singlet}^{UHF} = 1$ [305]. This would suggest that spin-contamination in significant excess of this value may represent distortions that are irrelevant to the essential physics of singlet biradicals, and thus UB3LYP orbitals, which are less severely contaminated, are a preferable alternative. Biradical triplet states, in general, are *not* expected to exhibit spin-contamination at the UHF level, as they are typically well-described by a single determinant (the two-determinant $M_s = 0$ triplet state is not encountered due to the constraint on $\langle S_z \rangle$). Thus, any spin-contamination in significant excess of the spin-pure value of $\langle S^2 \rangle_{triplet} = 2$ is likely undesirable, and in such situations AFQMC/U utilizes a trial wavefunction with UB3LYP orbitals.

5.3 Results and Discussion

5.3.1 Small Molecule Biradical Set

ST gaps are computed for the set of 13 small molecules with singlet states that exhibit substantial biradical character, recently examined in Ref. 275. Molecular geometries for the biradical species are taken from the Supporting Information of Ref. 275, which used QCISD/MG3S for CF_2 and QCISD(T)/aug-cc-pVQZ for the rest, with unrestricted (restricted) HF references for triplet (singlet) multiplicities, respectively. ph-AFQMC results utilizing CASSCF trial wavefunctions are plotted in Fig. 5.1, relative to experimental reference values, along with data provided in Ref. 275 for broken-symmetry DFT with the BLYP functional (UKS/BLYP), its spin-projected variant (WA-KS/BLYP), and a composite method which has been shown to produce comparable accuracy

to CCSD(T)/CBS[306] (W2X). We note that we use a different reference value for NH_2^+ , which is from a genuine experimental measurement[307], following Ref. 308.

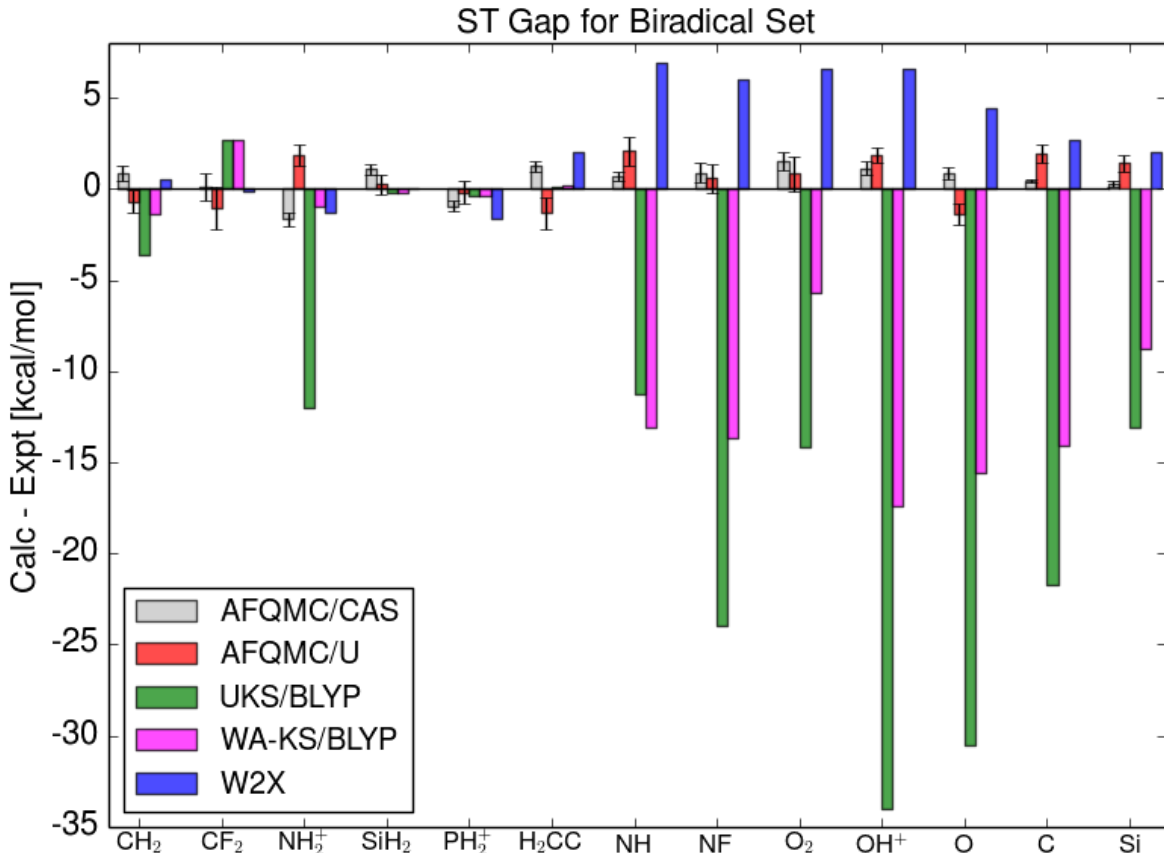


Figure 5.1: Deviations [kcal/mol] of various calculations from experimentally-derived reference values. Error bars show the statistical error of the QMC measurements.

The ph-AFQMC ST gaps have been converged with respect to active space sizes of the CASSCF trial wavefunctions, an approach which has been shown to produce very high accuracy even for strongly correlated systems[302, 236]. However the generation of such trial wavefunctions is, in practice, limited to moderate system sizes due to the procedure’s exponential scaling. In light of applications to large systems such as those found in organic electronics which we will focus on in future investigations, we explore the simplest scalable alternatives, namely UHF and RHF wavefunctions, and single determinants constructed from UB3LYP orbitals. $\langle S^2 \rangle$ values with respect to the trial wavefunctions are shown in Table 5.1.

Table 5.1: $\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for the 13 small molecule biradical set.

	UHF singlet	UHF triplet	UB3LYP singlet	UB3LYP triplet
CH ₂	0	2.02	0.57	2.01
CF ₂	0.14	2.01	0	2.00
NH ₂ ⁺	0.82	2.02	0.75	2.01
SiH ₂	0.43	2.01	0	2.00
PH ₂ ⁺	0.44	2.01	0	2.00
H ₂ CC	0.16	2.35	0	2.03
NH	1.01	2.02	1.00	2.01
NF	1.01	2.02	1.00	2.01
O ₂	1.02	2.05	1.01	2.01
OH ⁺	1.01	2.01	1.00	2.01
O	1.01	2.01	1.00	2.00
C	1.02	2.01	1.01	2.00
Si	1.05	2.02	1.01	2.00

Table 5.9 provides a rudimentary statistical representation of the accuracy of selected theoretical approaches with respect to experiment, comparing the mean signed error (MSE), mean absolute error (MAE), and maximum error (MaxE). As expected, utilizing CASSCF trial wavefunctions can obtain excellent accuracy, with an MAE of less than a kcal/mol and MaxE of 1.7(3) kcal/mol for O₂. For this molecule, we have tried active spaces of 12e8o, 8e12o, and 10e15o. The latter two active spaces produce statistically indistinguishable ST gaps in the CBS limit. Additionally, we verified that the total energies of both the singlet and triplet states in the aug-cc-pvtz basis differ by less than $1mHa$ going from one active space to the next. This suggests that the AFQMC/CAS result is converged with respect to active space size, and the remaining deviation from experiment is likely due to inaccuracies of the optimized geometries and/or the zero-point energies used to correct the experimental result.

Table 5.2: Mean signed, absolute, and maximum errors [kcal/mol] of the theoretical methods shown in Fig. 5.1 for the 13 small molecule biradical set. Sorted by MAE value. Parenthesis indicate statistical errors.

	MSE	MAE	MaxE
AFQMC/CAS	0.5(4)	0.9(4)	1.7(3)
AFQMC/U	0.5(7)	1.2(7)	2.1(8)
W2X	2.7	3.1	6.9
WA-KS/BLYP	-6.8	7.3	17.4
UKS/BLYP	-12.5	12.9	34

The performance of UKS/BLYP is unsurprisingly poor, given the high level of spin-contamination in the singlet states revealed in Table 5.1. The notable onset of larger errors in Fig. 5.1, i.e. for NH, NF, O₂, OH⁺, O, C, and Si, is found to roughly correlate with the presence of spin-contamination in the unrestricted wavefunctions. The Yamaguchi correction clearly improves upon the UKS/BLYP results, however the MAE of 7.3 kcal/mol and MaxE of 17.4 are still very large. W2X is relatively more robust, with an MAE of 3.1 kcal/mol; however, the MaxE of 6.9 kcal/mol illustrates the difficulty in describing biradical systems even with “gold-standard” single-reference methods.

In contrast, the AFQMC/U approach shows a significant improvement in accuracy (all ph-AFQMC results using UHF, UB3LYP, and RHF trials are shown in the Appendix). We note that the MaxE for AFQMC/UHF is for the H₂CC molecule. The UHF wavefunction for the triplet state, which should largely be of single-reference nature, still exhibits significant spin-contamination, as seen in Table 5.1. The Slater determinant derived from UB3LYP orbitals appears to be relatively uncontaminated, with an $\langle S^2 \rangle$ value of 2.03 while still benefiting from the additional variational freedom due to the use of unrestricted orbitals. When this is used as the trial wavefunction for ph-AFQMC, the resulting ST gap in the CBS limit is -49.9(9) kcal/mol, which significantly reduces the deviation from experiment from 3.1(7) to -1.3(9) kcal/mol. While more data points are needed to validate a more general claim, this case suggests that when a single-reference spin-state exhibits spin-contamination, AFQMC/UB3LYP can improve the accuracy of ST predictions over AFQMC/UHF. This is reflected in the protocol specified earlier for AFQMC/U, which utilizes UHF trial wavefunctions for triplet states when $\langle S^2 \rangle \leq 2.1$, and UB3LYP otherwise.

To make a broader comparison regarding the ability of a variety of electronic structure methods (with similar computational scaling with respect to system size) to predict ST gaps in biradicals, we include calculated values from methods highlighted in Ref. 24 for a subset of 8 biradicals. We plot this data in Fig. 5.2, and provide a comparative statistical summary in Table 5.10.

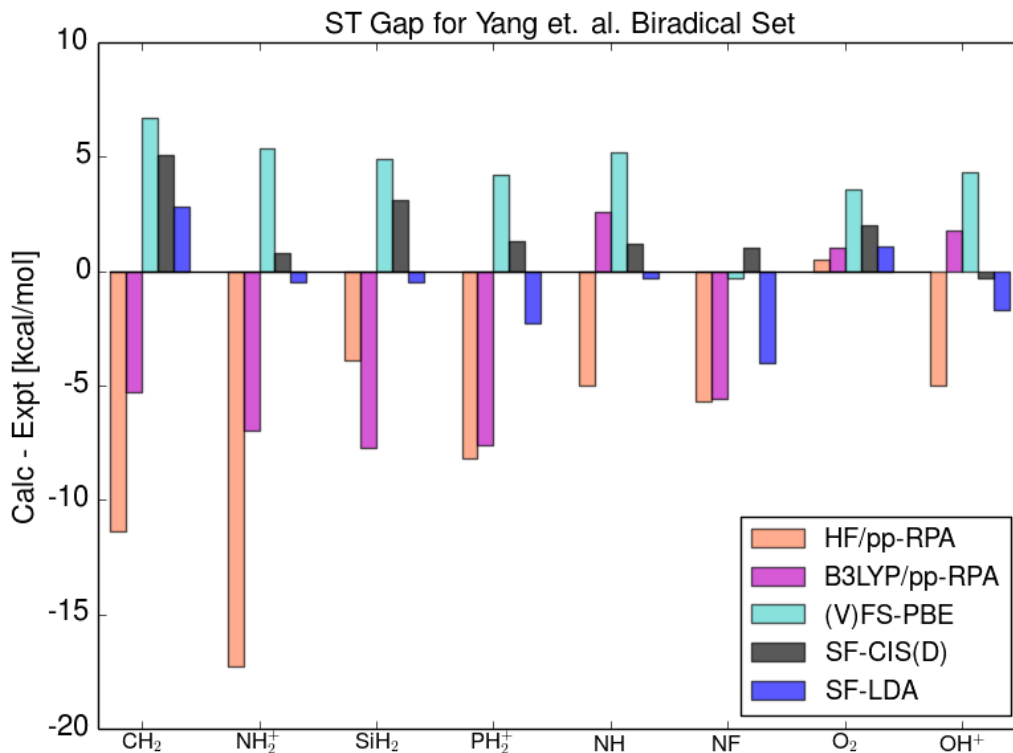


Figure 5.2: Deviations [kcal/mol] of computational methods selected from Ref. 24 from experimentally-derived reference values, for a subset of 8 biradicals.

Table 5.3: Mean signed, absolute, and maximum errors [kcal/mol] of ph-AFQMC results and other methods for the 8 molecule biradical subset shown in Fig. 5.2. Sorted by MAE value. Parenthesis indicate statistical errors.

	MSE	MAE	MaxE
AFQMC/U [†]	0.83(7)	1.0(7)	2.1(7)
AFQMC/CAS [†]	0.5(4)	1.1(4)	1.7(3)
SF-LDA ^a	-0.7	1.7	4.0
SF-CIS(D) ^b	1.8	1.9	5.1
W2X ^c	3.0	3.7	6.9
(V)FS-PBE ^d	4.2	4.3	6.7
B3LYP/pp-RPA ^e	-3.5	4.8	7.7
WA-KS/BLYP ^c	-6.6	6.6	17.5
HF/pp-RPA ^e	-7.0	7.1	17.3
UKS/BLYP ^c	-12.5	12.5	34.0

[†] This work.

^a SF-TDDFT with LDA functional and non-collinear kernel, TZ2P basis, from Ref. 309.

^b cc-pVQZ basis for NH, OH⁺, NF, O₂; TZ2P basis for CH₂, NH₂⁺, SiH₂, PH₂⁺; from Ref. 23

^c from Ref. 275.

^d (Variational) Fractional-Spin method, 6-311++G(2d,2p) basis, from Ref. 266.

^e aug-cc-pVDZ basis, from Ref. 24.

For this subset of cases, there is no distinction between the AFQMC/UHF and AFQMC/U procedures, and AFQMC/U and AFQMC/CAS yield equivalent accuracy, considering statistical error bars. Both produce MAEs of ~ 1 kcal/mol and MaxEs of ~ 2 kcal/mol, comparing favorably to all other methods. As shown in the Appendix, the accuracy of AFQMC/UB3LYP and AFQMC/RHF is similar to that obtained via spin-flip methods for these systems.

5.3.2 Benzyne Isomers

In this section we consider the ortho-, meta-, and para-benzyne isomers, shown in Fig. 5.3, and compare predicted ST gaps with precise, gas-phase experimental measurements. The ground state for all isomers is a singlet, and biradical character correlates with the distance between the unpaired

electrons (ortho < meta < para)[23]. These systems are of scientific interest in their own right, e.g. singlet para-benzyne is a biradical that can abstract hydrogen atoms from specific positions in DNA, potentially enabling antitumor activity[310, 311].

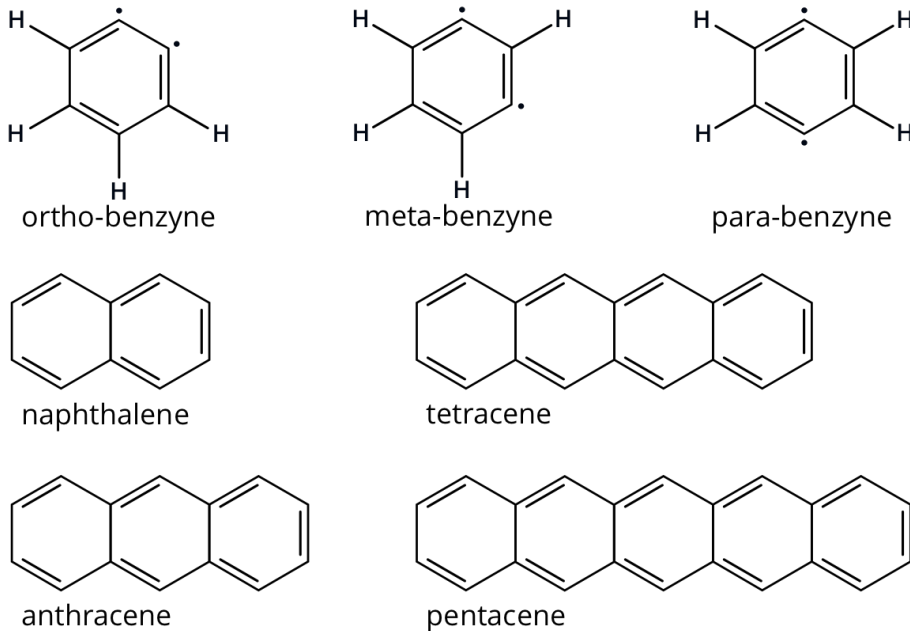


Figure 5.3: Benzyne isomers and polyacenes studied in this work.

Ab initio calculations are difficult due to the strongly correlated biradical electrons. In particular, singlet para-benzyne exhibits orbital instabilities at the RHF level, and subsequent RHF-based correlation methods produce poor results. Better results are obtained with the broken-symmetry UHF reference, with complete spin symmetry-restoration via, e.g., subsequent CCSD calculation[311]. With this in mind, we explore UHF and UB3LYP trial wavefunctions for ph-AFQMC, in addition to multi-determinant trial wavefunctions from CASSCF.

Our ph-AFQMC calculations utilize geometries obtained from SF-DFT/6-311G* using the so-called 50/50 functional[23]. The ortho- and meta- geometries are provided in Ref. 24, the para- geometry in Ref. 273. In Ref. 24, ST gaps of para-benzyne are based on SF-CCSD/cc-pVTZ optimized geometries[273]. However we have performed ph-AFQMC calculations on this SF-CCSD geometry with CASSCF trial wavefunctions utilizing increasingly large active spaces (up to 8 electrons in 16 orbitals), and still find a residual error from experiment of 2.9(7) kcal/mol in the

cc-pVTZ basis (which is altered by only 1 kcal/mol in the estimated CBS limit). For this reason, and for consistency with the ortho- and meta- isomers, we use the SF-DFT/6-311G* geometry for the para-isomer as well.

It is noteworthy that Refs. 24 and 271 do not account for ZPE contributions in comparing calculated electronic ST energies to experimentally measured quantities. For meta-benzynes this correction is ~ 1 kcal/mol, so we choose here to subtract out ZPE values, taken from Ref. 23, from the experimental data when comparing with purely electronic predictions.

$\langle S^2 \rangle$ values of the candidate unrestricted trial wavefunctions are shown in Table 5.4 for the benzyne isomers. The UHF singlet and triplet states are both significantly contaminated in all isomers, with the singlet state of para-benzynes is severe case ($S^2=1.68$). UB3LYP reduces the amount of spin-contamination in all cases, however it does not always eliminate it, e.g. $\langle S^2 \rangle$ of singlet para-benzynes is reduced to 0.92. The AFQMC/U method will use UB3LYP trial wavefunctions for all isomers, since $\langle S^2 \rangle_{singlet}^{UHF} > 1.1$ for ortho- and para- benzyne, and $\langle S^2 \rangle_{triplet}^{UHF} > 2.1$ for all.

Table 5.4: $\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for ortho- meta- and para-benzynes molecules.

	UHF singlet	UHF triplet	UB3LYP singlet	UB3LYP triplet
ortho	1.26	2.31	0.00	2.01
meta	0.97	2.68	0.10	2.02
para	1.68	2.31	0.92	2.01

The resulting ST gaps are shown in Table 5.11, and the deviations of the predicted values from ZPE-corrected experimental results are shown in Fig. 5.7. AFQMC/UHF and AFQMC/UB3LYP values are shown separately in the Appendix.

Table 5.5: ST gaps [kcal/mol] for the ortho- meta- and para- benzyne isomers. Parenthesis indicate statistical errors.

	ortho	meta	para
expt [*]	37.5 \pm 0.3	21.0 \pm 0.3	3.8 \pm 0.3
ZPE ^{**}	-0.6	1.0	0.5
ZPE-corr'd expt	38.1	20.0	3.3
AFQMC/CAS [†]	37.4(6)	20.7(8)	4.5(5)
AFQMC/U [†]	37.6(7)	18.9(9)	2.2(9)
UB3LYP ^a	29.4	14.2	2.4
HF/pp-RPA ^b	45.6	35.5	4.0
B3LYP/pp-RPA ^b	37.4	22.1	0.6
SF-CIS(D) ^c	35.7	19.4	2.1
SF-B3LYP ^d	46.9	26.1	6.9
SF-CCSD(T) ^e	37.3	20.6	4.0
SF-oo-CCD ^c	37.6	19.3	3.9

^{*} Ref. 312

^{**} SF-DFT/6-311G*, Ref. 23.

[†] SF-DFT/6-311G* geometries from Ref. 23, cc-pVTZ basis. This work.

^a 6-31G* basis, Ref. 271

^b Ref. 24. pp-RPA calculations in aug-cc-pVDZ basis. o- and m- geometries from SF-DFT with the 50/50 functional (Ref.271). p- geometry from SF-CCSD, Ref. 273.

^c SF-DFT/6-311G* geometries, cc-pVTZ basis, Ref. 23.

^d cc-pVTZ basis, Ref. 273.

^e SF-DFT/6-311G* geometries, cc-pVTZ basis, Ref. 313.

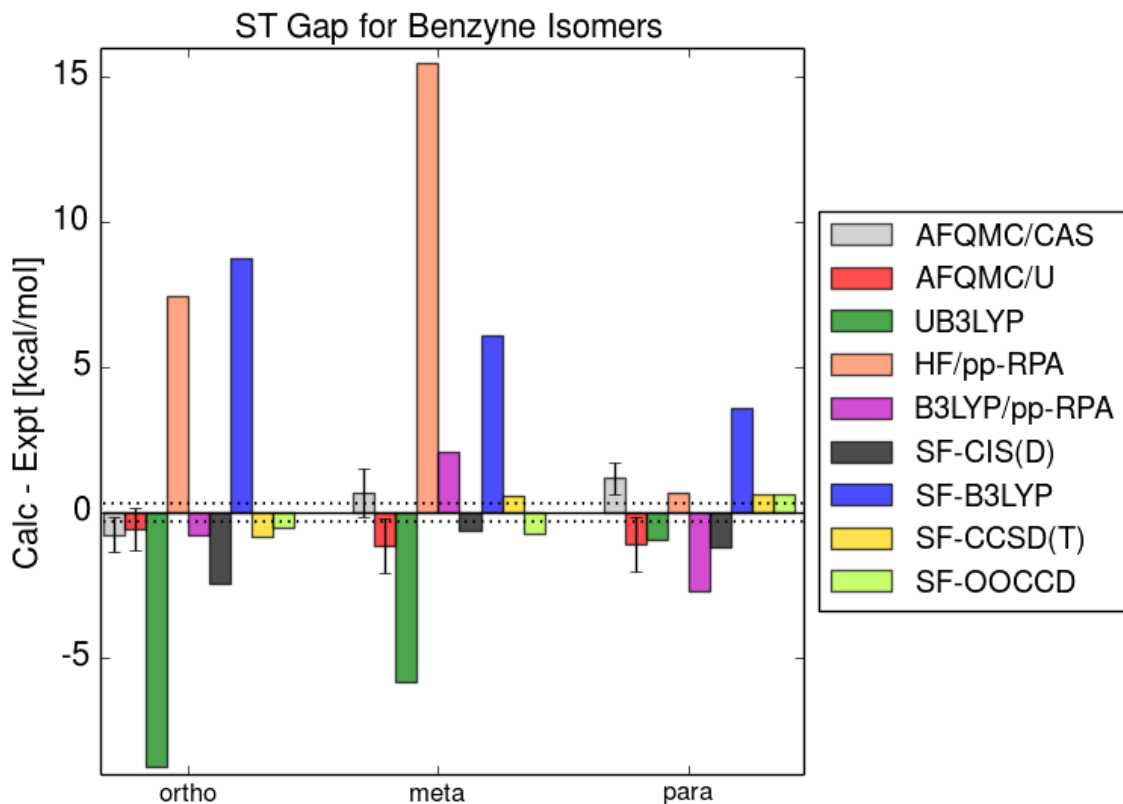


Figure 5.4: Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements. Dotted black lines represent the reported uncertainty of the experimental measurements.

SF-oo-CCD and SF-CCSD(T) both consistently obtain sub kcal/mol accuracy, however they are computationally infeasible for larger systems due to the respective $O(M^6)$ and $O(M^7)$ scaling. We thus focus presently on methods with lower scaling.

AFQMC/CAS and AFQMC/U produce predictions of comparable accuracy, with maximum errors just outside 1 kcal/mol. AFQMC/UHF achieves good accuracy for the ortho- and meta-isomers, however the severely spin-contaminated singlet state of para-benzyne results in a relatively large overestimation of the ST gap. For this latter system AFQMC/UB3LYP substantially reduces the deviation from experiment from 5.7(8) to 1(1) kcal/mol, and we note that similarly pronounced corrections are observed in pp-RPA calculations when the B3LYP reference for the $(N-2)$ -electron

system is used instead of HF[24]. AFQMC/U and AFQMC/UB3LYP are equivalent for this set of molecules, and thus we again find that AFQMC/U, utilizing only unrestricted single-determinant trials, can produce results of comparable accuracy to both AFQMC/CAS and experiment.

We must emphasize the importance of explicitly breaking the spin-symmetry when obtaining UHF and UB3LYP trial wavefunctions, and choosing the solution with lowest-energy. For instance, the calculated ST gap using a spin-pure ($\langle S^2 \rangle = 0$) B3LYP trial for para-benzyne is -15 kcal/mol, vs +2.2 kcal/mol when the lower-energy $\langle S^2 \rangle = 0.92$ unrestricted singlet state is used.

As a final remark in this section, we note that while indeed the errors from UB3LYP (i.e. without subsequent QMC) are generally reduced in comparison with those from the small molecule biradical set, we still find significant errors of -8.7, -5.8, and -0.9 kcal/mol for the ortho-, meta-, and para-benzynes. UB3LYP systematically underestimates the ST gaps, due to the unrealistically low energy of the broken-symmetry singlet state. Rather unexpectedly, however, the magnitude of the errors here are inversely correlated with diradical character.

These results suggest that for the benzyne isomers, which exhibit strong biradical character while sharing features similar to the planar aromatic ring systems relevant to chemical photocatalysis, the ST gaps are accurately predicted by ph-AFQMC with both CASSCF and single-determinant trials, and in the cc-pVTZ basis. The accuracy of AFQMC/UHF, even in its spin-projected form, is compromised by the heavily spin-contaminated singlet state in para-benzyne, though AFQMC/UB3LYP provides an improved prediction, and is utilized in our AFQMC/U method.

5.3.3 Polyacenes

Having shown that AFQMC/U can accurately describe molecules with strong biradical nature, we now show that this computational approach can scale to larger molecules, focusing on polyacenes from naphthalene to pentacene, shown in Fig. 5.3. These molecules all have singlet ground states. We use geometries from Ref. 314 for the acenes, which were computed at the unrestricted B3LYP/6-31G(d) level of theory.

Table 5.6 shows that the extent of spin-contamination in the singlet UHF states increases with

the number of fused rings. In contrast, wavefunctions constructed from UB3LYP orbitals are spin-pure, consistent with previous computational studies[315, 281, 280].

Table 5.6: $\langle S^2 \rangle$ of the singlet and triplet unrestricted solutions for polyacenes $n = 2 - 5$.

n	UHF singlet	UHF triplet	UB3LYP singlet	UB3LYP triplet
2	1.10	2.30	0	2.02
3	1.78	2.68	0	2.02
4	2.43	2.91	0	2.03
5	3.06	3.44	0	2.03

ST gaps calculated with ph-AFQMC in the cc-pVTZ basis are shown in Table 5.13, alongside predictions from state-of-the-art *ab initio* methods and available experimental data. When more than one experimental value is given in Ref. 303, we choose the value that is closest to that shown in Ref. 290. Zero point energy corrections, which are subtracted out of the experimental values, are required in order to compare calculated electronic energies with experiment, and we utilize numbers from Ref. 303 derived from B3LYP/6-31G(d) geometry optimizations and frequency calculations.

It is important to recognize that all the experimental values, except in the case of anthracene, *cannot* be fairly compared directly with gas-phase calculations. Ref. 281 conveniently provides details of many of the experimental measurements, which are reproduced here. While the adiabatic ST gap of anthracene was obtained via gas-phase photoelectron spectroscopy, the reported experimental measurement for naphthalene was done in ether-isopentane-alcohol (solid) solvents at 77K. The tetracene measurement was done in poly(methyl methacrylate) matrix at 298K, and pentacene was measured in a tetracene matrix at 298K. Clearly, the gas-phase 0K conditions assumed in our calculations are not consistent with the realistic experimental conditions for most of the polyacenes studied here.

Table 5.7: ST gaps [kcal/mol] for the polyacenes $n=2-5$. Square brackets indicate alternate experiments, and parenthesis indicate statistical errors.

	naphthalene	anthracene	tetracene	pentacene
expt [*]	[60.9] 61.0	[42.6] 43.1	29.4	19.8±0.7
ZPE [*]	-3.4	-2.3	-1.8	-1.5
ZPE-corr'd expt	64.4	45.4	31.2	21.3
AFQMC/U ^a	68.0(1.2)	46.2(1.2)	34.0(1.6)	25.2(1.6)
UB3LYP ^b	62.6	41.8	27.7	17.9
CCSD(T)/FPA ^c	65.8	48.2	33.5	25.3
B3LYP/pp-RPA ^d	66.2	45.7	32.1	22.6
GAS-pDFT (FP-1) ^e	70.6	45.5	33.6	25.4
GAS-pDFT (WFP-3) ^e	64.7	43.1	28.8	20.5
ACI-DSRG-MRPT2 ^f	62.2	43.2	28.3	18.0
DMRG-pDFT ^g	67.1	46.1	31.6	22.6

^{*} Taken from Ref. 303

^a this work

^b UB3LYP/6-31G(d) geometries and energies. Ref. 291

^c B3LYP/cc-pVTZ geometries. Ref. 281

^d UB3LYP/6-31G* geometries, B3LYP reference, cc-pVDZ basis. Ref. 290

^e UB3LYP/6-31G(d,p) geometries, tPBE/6-31G(p,d). Active spaces defined in Ref. 277

^f UB3LYP/6-31G(d) geometries. Ref. 303

^g UB3LYP/6-31G(d,p) geometries. tPBE/6-31+G(p,d). Ref. 279

We observe that the ph-AFQMC predictions are insensitive to the trial wavefunction used for these polyacene systems, and show results from UHF, RHF, and UB3LYP trial wavefunctions in the Appendix. Given that naphthalene through pentacene exhibit little biradical character[290], the extreme spin-contamination of the UHF solutions shown in Table 5.6 is likely *not* representative of strong electron correlation, and hence restricted trial wavefunctions can also be expected to yield accurate results. Moreover, all ph-AFQMC variants unambiguously achieve high accuracy with respect to the gas-phase measurement for anthracene. The ph-AFQMC predictions for the other polyacenes, which were experimentally probed in (solid) solvent matrix, appear to systematically overestimate the experimental values by a few kcal/mol. For naphthalene, we performed ph-AFQMC calculations with CASSCF(10e,10o) trial wavefunctions, which gave a ST gap of 67(1) kcal/mol. This value is in agreement with that from AFQMC/U, given statistical error bars, and lies above the reported ZPE-corrected experiment by some 3 kcal/mol. Very similar overestimations by

AFQMC/U are seen for $n=2,4,5$ and are corroborated by CCSD(T)/FPA, B3LYP/pp-RPA, and DMRG-pDFT (though not by UB3LYP and ACI-DSRG-MRPT2). Thus it seems reasonable to hypothesize that the calculations’ neglect of molecular environment may be responsible, though admittedly there are a number of other factors that might be expected to contribute. For instance, adiabatic ST gaps are known to be sensitive to the optimized geometries, which can result in variations on the order of 1-3 kcal/mol[303].

Overall, given the uncertainties due to the treatment of temperature, solvent, and molecular geometries in these acene calculations, ph-AFQMC with single-determinant trial wavefunctions gives satisfactory agreement with both experiments and other highly-accurate electronic structure methods, and moreover can scale with near-perfect parallel efficiency to systems as large as pentacene in a triple-zeta basis, which has 146 electrons and 856 basis functions.

5.3.4 Potential Annihilators for Optical Upconversion - Anthracene Derivatives

In this section, we predict the ST gaps of a set of anthracene derivatives with AFQMC/U, and compare with results from UB3LYP and a localized-orbital variant of CCSD(T) theory, DLPNO-CCSD(T)[316, 317]. The structures are shown in Fig. 5.5. The anthracene with two methyl substituents, known as DMA, is a known triplet-triplet annihilator in optical upconversion schemes[318], i.e. it catalyzes the fusion of two triplet excitons into a higher energy singlet state. In an effort to design new annihilators, we generate candidate compounds by replacing the methyl substituents with various functional groups that are synthetically feasible, and probe the effects, if any, on the ST gaps.

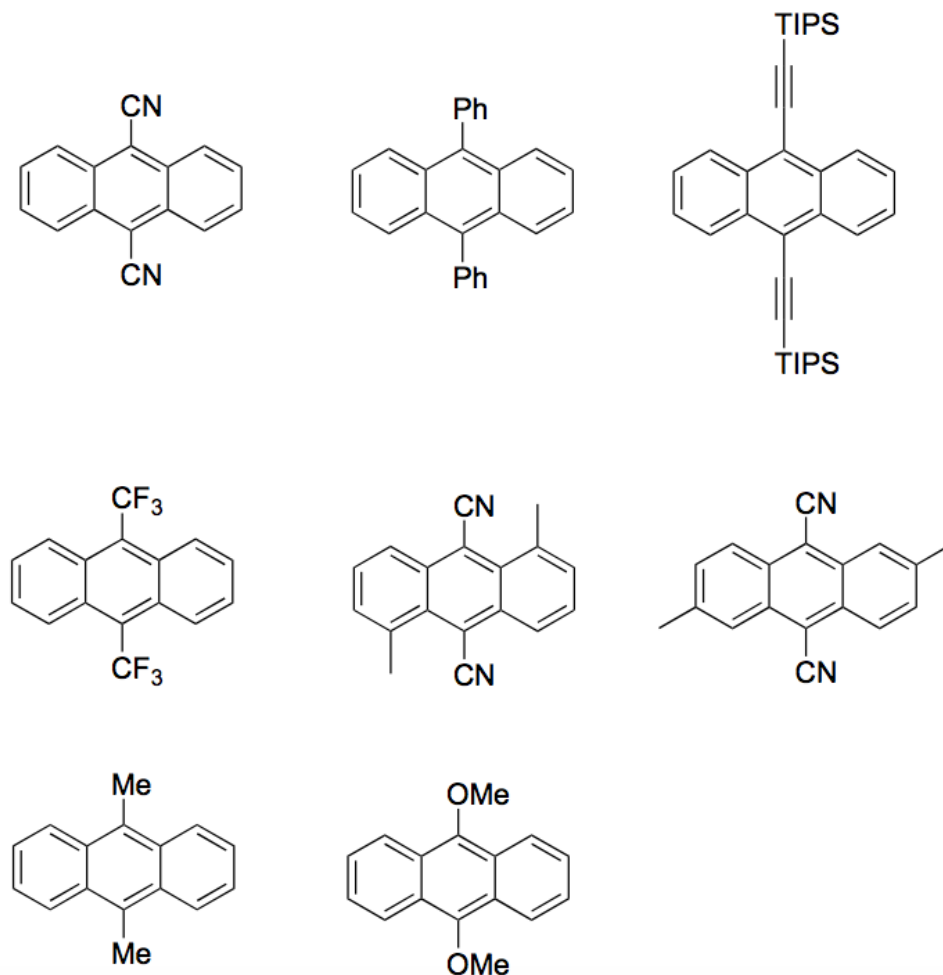


Figure 5.5: Structures of the anthracene derivatives. Top row, from left to right: DCA, DPA, TIPS. Middle row: CF₃, CN-UpDnMethyl, CN-DiagMethyl. Bottom row: DMA, OMe

UB3LYP and DLPNO-CCSD(T) calculations were performed with the ORCA quantum chemistry program[319]. Geometries were optimized at the UB3LYP/cc-pVTZ level of theory. Vibrational frequencies were computed numerically with an increment of 0.001 Bohr. Zero-point energies and thermal vibrational corrections are scaled by a factor obtained from the NIST computational chemistry database (0.965 for B3LYP/cc-pVTZ)[149] and added to the resulting energies.

The DFT results reported in Table 5.8 are obtained from geometries and vibrational frequencies computed within the conductor-like polarizable continuum model (CPCM) to estimate the effects of toluene solvent. The DLPNO-CCSD(T) and AFQMC/U values are obtained by first calculating

the ST energy difference in the gas-phase; energetic corrections to this gap due to ZPEs and thermal occupation of vibrational states, computed in the gas-phase, are scaled and then added. Finally, a solvent correction is obtained via DFT as the difference between the vibrationally corrected ST gaps in toluene, as represented by the CPCM model, and in gas-phase.

Table 5.8: ST gaps for selected anthracene derivatives

	UB3LYP	DLPNO-CCSD(T)	AFQMC/U
DCA	1.50	1.75	1.75(07)
DPA	1.66	1.91	1.82(11)
TIPS	1.39	1.66	*
CF3	1.66	1.95	2.05(11)
CN-UpDnMethyl	1.28	1.58	1.56(12)
CN-DiagMethyl	1.42	1.70	1.73(11)
DMA	1.54	1.77	1.69(07)
OMe	1.60	1.85	1.81(10)

* Given the general agreement for these systems between AFQMC/U and DLPNO-CCSD(T) we omit the TIPS calculation, as it requires ~ 1000 basis functions and thus significant computational effort.

In every case, the DLPNO-CCSD(T) and AFQMC/U results agree to within the statistical error bars of the latter, though both are significantly different than the DFT values, which appear to systematically underestimate the gaps. We are aware of only one experimental measurement of the ST gap, which for DCA is 1.8 eV[320]. UB3LYP performs poorly for this system, underestimating the gap by 0.3 eV, whereas both DLPNO-CCSD(T) and AFQMC/U are in good agreement with the experimental value.

For all molecules except DPA and TIPS, we investigated the effects of CBS extrapolation on the ST gaps, and find that with DLPNO-CCSD(T) the difference between the cc-pVTZ result and our extrapolated estimate of the CBS limit is never larger than 0.06 eV. Given that our QMC error bars are larger than this, further extrapolation of the AFQMC/U results appears to be unnecessary.

5.4 Conclusions

With CASSCF trial wavefunctions, ph-AFQMC can predict ST gaps with sub kcal/mol accuracy with respect to gas-phase experimental measurements for a set of 13 small molecules with singlet states of strong biradical character. However, for large systems the generation of such a trial wavefunction quickly becomes impractical. The main result of this work is that near-chemical accuracy for gas-phase ST gaps can also be obtained, even for strongly correlated biradical systems, with a spin-projected ph-AFQMC technique, which initializes walkers with restricted determinants while using an unrestricted single-determinant trial wavefunction to implement the phaseless constraint. We establish a quantitative criteria for choosing UHF or UB3LYP orbitals based on the spin-contamination of the UHF wavefunction, and the resulting AFQMC/U methodology is validated on the small molecule test set, the ortho-, meta-, and para-benzyne isomers, and of all the polyacenes for which experimental results are reported (though a true gas-phase experiment is only available for anthracene). Having shown that the ph-AFQMC method can provide a balanced and robust approach to accurately predict ST gaps, we use AFQMC/U to compute this quantity for a set of anthracene derivatives, in an effort to discover novel triplet-triplet annihilators for optical upconversion.

5.A Appendix - Ph-AFQMC data from various trial wavefunctions

5.A.1 Small Molecule Biradical Set

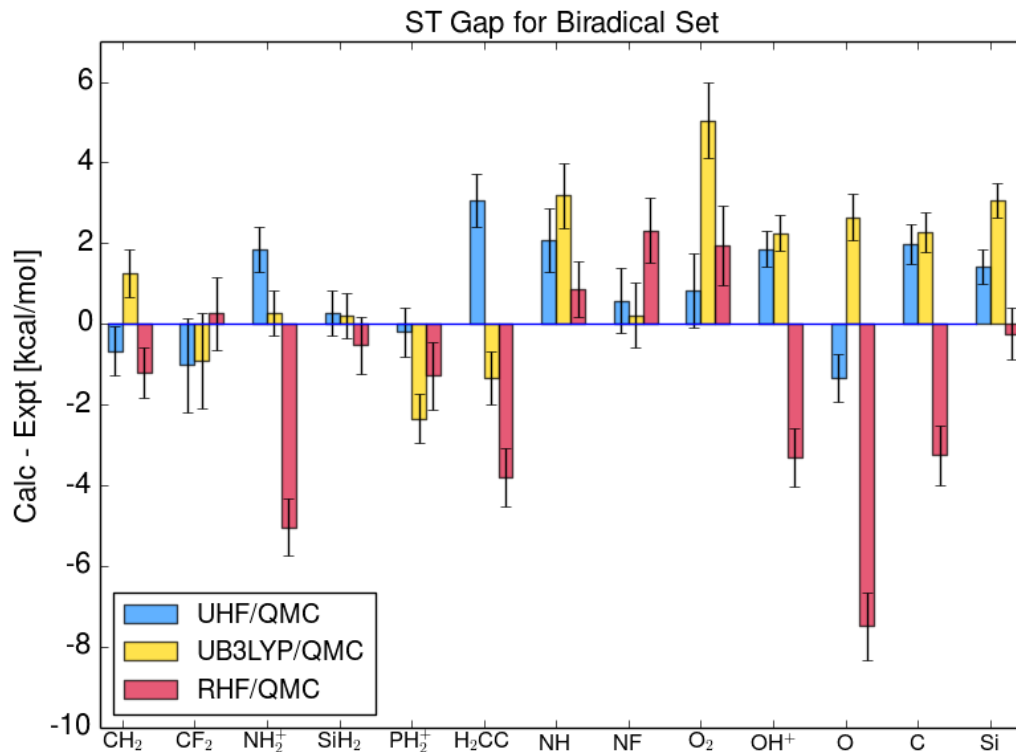


Figure 5.6: Deviations [kcal/mol] of ph-AFQMC predictions with single-determinant trial wavefunctions from experimentally-derived reference values. Error bars show the statistical error of the QMC measurements.

Table 5.9: Mean signed, absolute, and maximum errors [kcal/mol] of the theoretical methods shown in Fig. 5.6 for the 13 small molecule biradical set. Sorted by MAE value.

	MSE	MAE	MaxE
AFQMC/CAS	0.5(4)	0.9(4)	1.7(3)
AFQMC/U	0.5(7)	1.2(7)	2.1(8)
AFQMC/UHF	0.8(7)	1.3(7)	3.1(7)
AFQMC/UB3LYP	1.2(7)	1.9(7)	5.0(1.2)
AFQMC/RHF	-1.6(8)	2.4(8)	7.5(8)

Table 5.10: Mean signed, absolute, and maximum errors [kcal/mol] of ph-AFQMC results and other methods for the 8 molecule biradical subset shown in Fig. 2 of the main manuscript. Sorted by MAE value.

	MSE	MAE	MaxE
AFQMC/UHF ^{†, ††}	0.83(7)	1.0(7)	2.1(7)
AFQMC/CAS [†]	0.5(4)	1.1(4)	1.7(3)
SF-LDA ^a	-0.7	1.7	4.0
AFQMC/UB3LYP [†]	1.26(7)	1.8(7)	5.0(1.2)
SF-CIS(D) ^b	1.8	1.9	5.1
AFQMC/RHF [†]	-0.8(8)	2.1(8)	5.0(7)
W2X ^c	3.0	3.7	6.9
(V)FS-PBE ^d	4.2	4.3	6.7
B3LYP/pp-RPA ^e	-3.5	4.8	7.7
WA-KS/BLYP ^c	-6.6	6.6	17.5
HF/pp-RPA ^e	-7.0	7.1	17.3
UKS/BLYP ^c	-12.5	12.5	34.0

[†] This work.

^{††} Equivalent to AFQMC/U

^a SF-TDDFT with LDA functional and non-collinear kernel, TZ2P basis, from Ref. 309.

^b cc-pVQZ basis for NH, OH⁺, NF, O₂; TZ2P basis for CH₂, NH₂⁺, SiH₂, PH₂⁺; from Ref. 23

^c from Ref. 275.

^d (Variational) Fractional-Spin method, 6-311++G(2d,2p) basis, from Ref. 266.

^e aug-cc-pVDZ basis, from Ref. 24.

5.A.2 Benzyne Isomers

Table 5.11: ST gaps [kcal/mol] for the ortho- meta- and para- benzyne isomers.

	ortho	meta	para
expt	37.5 ± 0.3	21.0 ± 0.3	3.8 ± 0.3
ZPE*	-0.6	1.0	0.5
ZPE-corr'd expt	38.1	20.0	3.3
AFQMC/CAS [†]	37.4(6)	20.7(8)	4.5(5)
AFQMC/UHF [†]	40.0(9)	20.7(8)	9.0(8)
AFQMC/UB3LYP ^{†, ††}	37.6(7)	18.9(9)	2.2(9)

* SF-DFT/6-311G*, Ref. 23.

[†] SF-DFT/6-311G* geometries from Ref. 23, cc-pVTZ basis. This work.

^{††} Equivalent to AFQMC/U

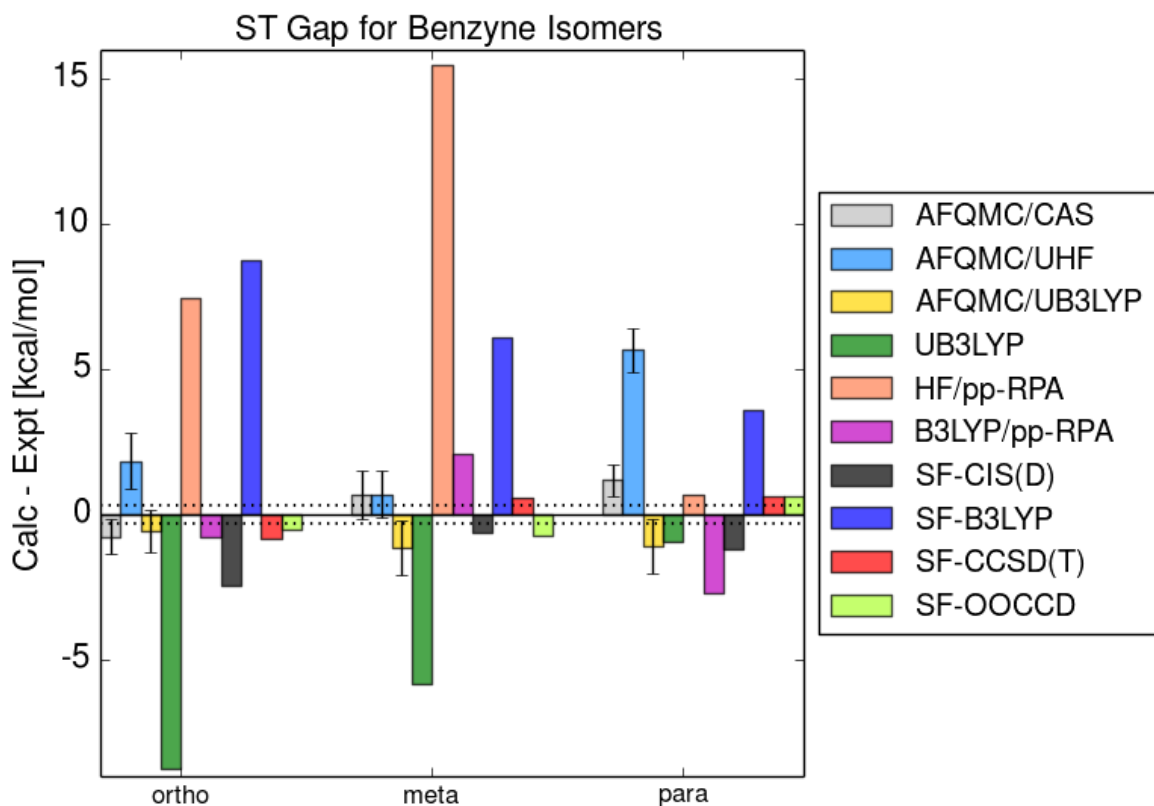


Figure 5.7: Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements. Dotted black lines represent the reported uncertainty of the experimental measurements.

5.A.3 Polyacenes

Table 5.12: Total energies [Ha] for naphthalene with various trial wavefunctions, in the cc-pVTZ basis.

Spin-state	Trial Wavefunction	ph-AFQMC Energy [Ha]
singlet	CASSCF(10e10o)	-385.334(1)
	UHF	-385.334(1)
	UB3LYP	-385.331(2)
	RHF	-385.334(2)
triplet	CASSCF(10e10o)	-385.227(1)
	UHF	-385.225(1)
	UB3LYP	-385.223(1)
	ROHF	-385.225(1)

Table 5.13: ST gaps [kcal/mol] for the polyacenes $n=2-5$.

	naphthalene	anthracene	tetracene	pentacene
expt	[60.9] 61.0	[42.6] 43.1	29.4	19.8 \pm 0.7
ZPE	-3.4	-2.3	-1.8	-1.5
ZPE-corr'd expt	64.4	45.4	31.2	21.3
AFQMC/UHF ^a	68.1(1.2)	44.7(1.1)	33.7(1.9)	24.4(1.5)
AFQMC/RHF ^a	68.1(1.2)	45.2(1.4)	35.1(1.7)	23.4(1.6)
AFQMC/UB3LYP ^{a, ††}	68.0(1.2)	46.2(1.2)	34.0(1.6)	25.2(1.6)

^{††} Equivalent to AFQMC/U

^a this work

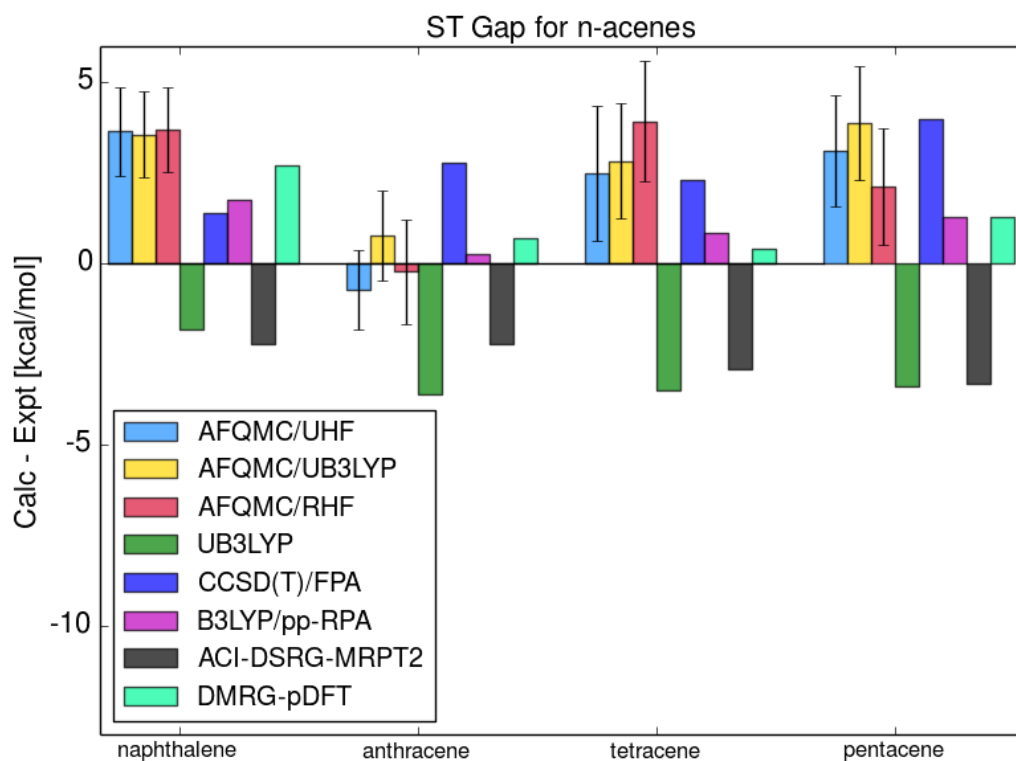


Figure 5.8: Deviations [kcal/mol] of various calculations from ZPE-corrected experimental measurements. Error bars show the statistical error of the QMC measurements.

Chapter 6

Conclusions

6.1 Current Status

The methodological improvements detailed in this work, together with the increasing availability of large-scale GPU-accelerated supercomputers, have enabled the broad application of ph-AFQMC to explore a wide variety of chemical applications, and we are continuing this push with a number of ongoing projects.

The ability to compute and compare barrier heights corresponding to competing reaction mechanisms is a useful tool in predicting the rate of a chemical reaction. In fact, the prevalence of one or more stretched bond at the transition state of a bond-breaking reaction coordinate demands the use of an *ab initio* method which can accurately describe significant electron correlation effects. To explore the performance of ph-AFQMC, we have computed reaction profiles for a number of isoprene-ozone reaction mechanisms that may possibly occur in the atmosphere. The results will be reported elsewhere, and will be followed by future explorations into the reactivity of transition metal complexes and clusters.

The Friesner group has, for a number of years, been developing an empirical correction scheme for DFT energy differences, called DFT-LOC[321, 322, 323, 14, 115]. We are utilizing ph-AFQMC to supplement a collection of experimental IPs (and, quite often, to select among many values reported with large uncertainties). This reference data is being used to fit a set of physically

relevant descriptors, which are parameters that depend, e.g., on the molecular geometry and other qualitative properties of the relevant chemical reaction. Having produced many hundreds of data points, the resulting model to correct DFT IPs is shown to yield notable improvements in accuracy compared to uncorrected DFT results. The success of this procedure paves the way for extensions to other types of chemical reactions, e.g. spin-splittings and pKa’s of transition metal complexes. In the near future, we intend to produce much larger data sets which will enable the investigation of various machine learning-based alternatives.

Producing benchmark-quality thermochemical predictions for realistic transition metal containing species is a momentous goal, which we are intent on pursuing, since accurate and precise experimental measurements in the gas phase especially are extremely rare. As discussed in Chapter 4, new experimental techniques from Morse and co-workers have recently been reported, which appear to yield reliable bond-dissociation measurements for small transition metal containing molecules with *sub* kcal/mol precision[237]. Having utilized such measurements for the 3*d*-containing diatomics, we are keen to compare ph-AFQMC predictions with the measurements for compounds containing heavier transition metals, e.g., Pt, Ag, and Au. In the near future we plan to validate a ph-AFQMC protocol for such elements, which will involve an examination of various scalar relativistic approximations and pseudopotentials. In addition, we have identified a modest number of ligand dissociation measurements for select transition metal complexes with reported experimental uncertainties of ~ 1 kcal/mol or less. We are currently using this reference data to explore how the bias due to the phaseless constraint depends upon the choice of trial wavefunction; in other words, we are asking the question, “What properties of a trial wavefunction are necessary to produce accurate energy differences for transition metal complexes?” The results in Chapter 5 suggest that an unrestricted single-determinant trial wavefunction is sufficient to describe some organic molecules, even those with significant biradical character; however, the results in Chapters 3 and 4 suggest that such trials are insufficient for transition metal atoms and diatomics (though we do expect, on average, that the extent of static electron correlation may be slightly ameliorated in response to ligand coordination). We are developing a systematic protocol to select active spaces, informed by ligand field theory and molecular orbital diagrams, for trial wavefunctions, and are comparing

the resulting ph-AFQMC results with experiment and high-level quantum chemical predictions.

6.2 Outlook

6.2.1 Additional Methodological Improvements

The projects mentioned above are being pursued alongside concurrent efforts to improve the ph-AFQMC methodology and its computational efficiency.

While, in principle, we have eliminated the memory bottleneck, as both the decomposed two-electron integrals and precomputed tensors required to efficiently evaluate the local energy can be split among an arbitrarily large number of GPU cards across multiple nodes, under such circumstances the parallel efficiency drops rather severely as the number of nodes is increased. Thus, further methodological advances, to increase the computational efficiency while at least maintaining current levels of accuracy, are desired. Our primary focus is to explore the utility of the DLPNO approximations in the context of AFQMC, which exploit the short-range nature of dynamic correlations to restrict the vast number of electronic excitations from localized, occupied orbitals into virtual orbitals which do not contribute significantly to the correlation energy. This approach, in the context of MP2[324], CC methods[316, 325, 317], and MCSCF methods[326], has been shown to drastically reduce both the scaling and memory requirements, and we have reason to believe the same will hold true in the case of AFQMC. Indeed, a related strategy known as “downfolding,” in which terms in the Hamiltonian involving virtual orbitals above a specified cutoff are excluded, has produced promising results[85], and may even be combined with DLPNO approaches to further accelerate ph-AFQMC.

A second methodological advance, which can potentially lead to large reductions in computational cost, involves a more efficient way to estimate the CBS limit. While our current MP2- or CCSD(T)-assisted extrapolation protocol appears to produce robust results in the vast majority of cases, we did find, e.g., in Chapter 4 that ph-AFQMC calculations in both triple- *and* quadruple- ζ basis sets are necessary to produce accurate predictions for highly multi-reference systems. An alternative approach would be to utilize an auxiliary set of “explicitly correlated” basis functions, as

is done in, e.g., Canonical Transcorrelation (CT) theory[327]. An effective Hamiltonian is produced with the electron cusps analytically removed. Such a transformation is analogous to including Jastrow factors in DMC[53]. We anticipate that this will reduce the number of basis functions required to reach the CBS limit[328, 329], and that the high scaling of the original CT procedure might be reduced by the use of localized orbital approximations.

Regarding code development, we note that the parallel efficiency of our implementation for large systems can be further improved by incorporating asynchronous MPI communications. This approach will enable more efficient utilization of the GPUs, as the code can be organized such that the devices can simultaneously perform computations while (unrelated) data is being transferred elsewhere. Finally, we anticipate that AFQMC calculations on large systems may require the incorporation of mixed-precision and sparse linear algebra routines[330, 175, 189].

6.2.2 Future Applications

Computing the relative energetics of low-lying excited-states is possible within the current ph-AFQMC framework[295], and is highly relevant especially to fields such as photocatalysis. Chapter 5 describes an accurate methodology to obtain the lowest excited-state of a different spin-multiplicity, which fits naturally into the ground-state projection scheme. We have also shown in Ref. 296 that an excited dipole-bound anionic state can be described given careful choice of trial wavefunction. An exciting frontier involves computing a manifold of excited-state wavefunctions with, e.g., CASSCF, TD-DFT, or CI methods, and utilizing these states as trials for subsequent ph-AFQMC calculations. This could be a powerful technology, especially if correlated sampling is used effectively.

Indeed, a similar approach utilizing an orthogonalization constraint with DFT virtual orbitals has been introduced to compute excited-state properties of solids, e.g. band structures and optical gaps[294]. As a new research direction, we envision that our efficient ph-AFQMC implementation will open the door to the accurate and fully *ab initio* investigation of strongly correlated solids. In addition to the recent advances mentioned above, we have outlined a way to compute quasiparticle

gaps, Δ , with the addition-removal expression[331]

$$2\Delta = [E(N + 1) - E(N)] + [E(N - 1) - E(N)], \quad (6.1)$$

where correlated sampling is utilized to converge the quantities in square brackets. While we anticipate that large finite-size effects will make such a calculation, in practice, far from trivial, we believe that ph-AFQMC has the potential to be a robust *ab initio* tool for band-gap prediction of semiconducting materials. Finally, we envision that atomistic ph-AFQMC simulations can accurately probe the charge- and spin- stripe orders known to be present in the ground-state of high-temperature superconducting materials such as the cuprates[332, 333]. We intend to corroborate the results from a recent study[82] in which the ground-state of the doped two-dimensional Hubbard model is characterized by state-of-the-art quantum-mechanical simulations. Another possible extension of the results in Chapter 5 is the computation of spin gaps in solids[334] utilizing the correlated sampling technique.

The far-reaching significance of the developments presented in this thesis can be illustrated by the potential use of ph-AFQMC to not only study systems such as the OEC of Photosystem II with quantum mechanical resolution, but also to improve the accuracy of molecular mechanics simulations for transition metal containing species. The parameterization of accurate classical force fields for transition metals would be highly beneficial to the fields of drug discovery, metalloenzymes, catalysis, and materials science. The required parameterization would likely involve the calculation of potential energy surfaces of a metal atom coordinated by a wide variety of ligands or amino acid residues. With accurate interaction potentials, it would be possible to use, e.g., the free-energy perturbation technique to investigate relative binding affinities of candidate drug molecules with metal cofactors in an active site. Furthermore, with such potentials one could rigorously compute solvation free-energies of metal containing complexes, which would properly account for important factors such as entropic effects at physiological temperatures. These results might then be used to improve the parameterization of implicit solvation models for transition metal ions and complexes, which could be used routinely and reliably to model metal chemistry in solution.

Bibliography

- [1] Max Born and Robert Oppenheimer. Zur quantentheorie der molekeln. *Annalen der physik*, 389(20):457–484, 1927.
- [2] Max Born and Kun Huang. *Dynamical theory of crystal lattices*. Clarendon press, 1954.
- [3] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 1989.
- [4] Robert G Parr. Density functional theory of atoms and molecules. In *Horizons of Quantum Chemistry*, pages 5–15. Springer, 1980.
- [5] Walter Kohn, Axel D Becke, and Robert G Parr. Density functional theory of electronic structure. *J. Phys. Chem.*, 100(31):12974–12980, 1996.
- [6] Shaloam Dasari and Paul Bernard Tchounwou. Cisplatin in cancer therapy: molecular mechanisms of action. *Eur. J. Pharmacol.*, 740:364–378, 2014.
- [7] Lloyd Kelland. The resurgence of platinum-based cancer chemotherapy. *Nat. Rev. Cancer*, 7(8):573, 2007.
- [8] Kevin J Waldron, Julian C Rutherford, Dianne Ford, and Nigel J Robinson. Metalloproteins and metal sensing. *Nature*, 460(7257):823, 2009.
- [9] Laura Riccardi, Vito Genna, and Marco De Vivo. Metal–ligand interactions in drug design. *Nat. Rev. Chem.*, page 1, 2018.

- [10] Lingle Wang, Yujie Wu, Yuqing Deng, Byungchan Kim, Levi Pierce, Goran Krilov, Dmitry Lupyan, Shaughnessy Robinson, Markus K Dahlgren, Jeremy Greenwood, et al. Accurate and reliable prediction of relative ligand binding potency in prospective drug discovery by way of a modern free-energy calculation protocol and force field. *J. Am. Chem. Soc.*, 137(7):2695–2703, 2015.
- [11] Thomas P Senftle, Sungwook Hong, Md Mahbubul Islam, Sudhir B Kylasa, Yuanxia Zheng, Yun Kyung Shin, Chad Junkermeier, Roman Engel-Herbert, Michael J Janik, Hasan Metin Aktulga, et al. The reaxff reactive force-field: development, applications and future directions. *Npj Comput. Mater.*, 2:15011, 2016.
- [12] Dmytro Bykov and Frank Neese. Six-electron reduction of nitrite to ammonia by cytochrome c nitrite reductase: insights from density functional theory studies. *Inorg. Chem.*, 54(19):9303–9316, 2015.
- [13] Jing Shen, Manuel J Kolb, Adrien J Gottle, and Marc TM Koper. Dft study on the mechanism of the electrochemical reduction of co2 catalyzed by cobalt porphyrins. *J. Phys. Chem. C*, 120(29):15714–15721, 2016.
- [14] Thomas F Hughes and Richard A Friesner. Development of accurate dft methods for computing redox potentials of transition metal complexes: results for model complexes and application to cytochrome p450. *J. Chem. Theory Comput.*, 8(2):442–459, 2012.
- [15] Sandeep Sharma, Kantharuban Sivalingam, Frank Neese, and Garnet Kin-Lic Chan. Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics. *Nat. Chem.*, 6(10):927–933, 2014.
- [16] Mikhail Askerka, Gary W Brudvig, and Victor S Batista. The o2-evolving complex of photosystem ii: recent insights from quantum mechanics/molecular mechanics (qm/mm), extended x-ray absorption fine structure (exafs), and femtosecond x-ray crystallography data. *Acc. Chem. Res.*, 50(1):41–48, 2016.

- [17] William Ames, Dimitrios A Pantazis, Vera Krewald, Nicholas Cox, Johannes Messinger, Wolfgang Lubitz, and Frank Neese. Theoretical evaluation of structural models of the s2 state in the oxygen evolving complex of photosystem ii: protonation states and magnetic interactions. *J. Am. Chem. Soc.*, 133(49):19743–19757, 2011.
- [18] Per EM Siegbahn and Margareta RA Blomberg. Transition-metal systems in biochemistry studied by high-accuracy quantum chemical methods. *Chem. Rev.*, 100(2):421–438, 2000.
- [19] Jing Zhou, Qian Liu, Wei Feng, Yun Sun, and Fuyou Li. Upconversion luminescent materials: Advances and applications. *Chem. Rev.*, 115(1):395–465, 2015. PMID: 25492128.
- [20] Jinli Qiao, Yuyu Liu, Feng Hong, and JiuJun Zhang. A review of catalysts for the electroreduction of carbon dioxide to produce low-carbon fuels. *Chem. Soc. Rev.*, 43(2):631–675, 2014.
- [21] Peter Baran, Roman Boca, Indranil Chakraborty, John Giapintzakis, Radovan Herchel, Qing Huang, John E McGrady, Raphael G Raptis, Yiannis Sanakis, and Athanasios Simopoulos. Synthesis, characterization, and study of octanuclear iron-oxo clusters containing a redox-active fe₄o₄-cubane core. *Inorg. Chem.*, 47(2):645–655, 2008.
- [22] Andrew Pinkard, Anouck M Champsaur, and Xavier Roy. Molecular clusters: Nanoscale building blocks for solid-state materials. *Acc. Chem. Res.*, 51(4):919–929, 2018.
- [23] Lyudmila V. Slipchenko and Anna I. Krylov. Singlet-triplet gaps in diradicals by the spin-flip approach: A benchmark study. *J. Chem. Phys.*, 117(10):4694–4708, 2002.
- [24] Yang Yang, Degao Peng, Ernest R Davidson, and Weitao Yang. Singlet–triplet energy gaps for diradicals from particle–particle random phase approximation. *J. Phys. Chem. A*, 119(20):4923–4932, 2015.
- [25] Thomas Devasagayam and Jayashree P Kamat. Biological significance of singlet oxygen. 2002.

- [26] Kirk A Peterson, David Feller, and David A Dixon. Chemical accuracy in ab initio thermochemistry and spectroscopy: current strategies and future challenges. *Theor. Chem. Acc.*, 131(1):1–20, 2012.
- [27] Richard A Friesner. Ab initio quantum chemistry: methodology and applications. *Proc. Natl. Acad. Sci. USA*, 102(19):6648–6653, 2005.
- [28] S Langhoff. *Quantum mechanical electronic structure calculations with chemical accuracy*, volume 13. Springer Science & Business Media, 2012.
- [29] C David Sherrill. Frontiers in electronic structure theory. *J. Chem. Phys.*, 132(11):110902, 2010.
- [30] Jeremy N Harvey. On the accuracy of density functional theory in transition metal chemistry. *Annu. Rep. Prog. Chem., Sect. C: Phys. Chem.*, 102:203–226, 2006.
- [31] George D Purvis III and Rodney J Bartlett. A full coupled-cluster singles and doubles model: The inclusion of disconnected triples. *J. Chem. Phys.*, 76(4):1910–1918, 1982.
- [32] Rodney J Bartlett and Monika Musiał. Coupled-cluster theory in quantum chemistry. *Rev. Mod. Phys.*, 79(1):291, 2007.
- [33] Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. An auxiliary-field quantum monte carlo study of the chromium dimer. *J. Chem. Phys.*, 142(6):064302, 2015.
- [34] Antara Dutta and C David Sherrill. Full configuration interaction potential energy curves for breaking bonds to hydrogen: an assessment of single-reference correlation methods. *J. Chem. Phys.*, 118(4):1610–1619, 2003.
- [35] Yuri A. Aoto, Ana Paula de Lima Batista, Andreas Kohn, and Antonio GS de Oliveira-Filho. How to arrive at accurate benchmark values for transition metal compounds: Computation or experiment? *J. Chem. Theor. Comput.*, 13(11):5291–5316, 2017.
- [36] Chr Møller and Milton S Plesset. Note on an approximation treatment for many-electron systems. *Phys. Rev.*, 46(7):618, 1934.

- [37] Björn O Roos. The complete active space self-consistent field method and its applications in electronic structure calculations. *Advances in Chemical Physics: Ab Initio Methods in Quantum Chemistry Part 2, Volume 69*, pages 399–445, 1987.
- [38] Jeppe Olsen. The casscf method: A perspective and commentary. *Int. J. Quantum Chem.*, 111(13):3267–3272, 2011.
- [39] Michael W Schmidt and Mark S Gordon. The construction and interpretation of mcscf wavefunctions. *Annu. Rev. Phys. Chem.*, 49(1):233–266, 1998.
- [40] Björn O Roos. The complete active space self-consistent field method and its applications in electronic structure calculations. *Adv. Chem. Phys.*, pages 399–445, 2007.
- [41] Kerstin Andersson, Per Aake Malmqvist, Björn O Roos, Andrzej J Sadlej, and Krzysztof Wolinski. Second-order perturbation theory with a casscf reference function. *J. Phys. Chem.*, 94(14):5483–5488, 1990.
- [42] Kerstin Andersson, Per-Åke Malmqvist, and Björn O Roos. Second-order perturbation theory with a complete active space self-consistent field reference function. *J. Chem. Phys.*, 96(2):1218–1226, 1992.
- [43] Adam A Holmes, Norm M Tubman, and CJ Umrigar. Heat-bath configuration interaction: An efficient selected configuration interaction algorithm inspired by heat-bath sampling. *J. Chem. Theory Comput.*, 12(8):3674–3680, 2016.
- [44] Kristine Pierloot. Transition metals compounds: Outstanding challenges for multiconfigurational methods. *Int. J. Quantum Chem.*, 111(13):3291–3301, 2011.
- [45] Sandeep Sharma, Adam A Holmes, Guillaume Jeanmairet, Ali Alavi, and Cyrus J Umrigar. Semistochastic heat-bath configuration interaction method: selected configuration interaction with semistochastic perturbation theory. *J. Chem. Theory Comput.*, 13(4):1595–1604, 2017.
- [46] Garnet Kin-Lic Chan and Sandeep Sharma. The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.*, 62:465–481, 2011.

- [47] Yuki Kurashige, Garnet Kin-Lic Chan, and Takeshi Yanai. Entangled quantum electronic wavefunctions of the mn4cao5 cluster in photosystem ii. *Nat. Chem.*, 5(8):660–666, 2013.
- [48] Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. *J. Am. Stat. Assoc.*, 44(247):335–341, 1949.
- [49] M Peter Nightingale and Cyrus J Umrigar. *Quantum Monte Carlo methods in physics and chemistry*. Number 525. Springer Science & Business Media, 1998.
- [50] James Gubernatis, Naoki Kawashima, and Philipp Werner. *Quantum Monte Carlo Methods*. Cambridge University Press, 2016.
- [51] Brian L Hammond, William A Lester Jr, and Peter James Reynolds. *Monte Carlo methods in ab initio quantum chemistry*, volume 1. World Scientific, 1994.
- [52] William A Lester Jr and Brian L Hammond. Quantum monte carlo for the electronic structure of atoms and molecules. *Annu. Rev. Phys. Chem.*, 41(1):283–311, 1990.
- [53] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum monte carlo simulations of solids. *Rev. Mod. Phys.*, 73:33–83, Jan 2001.
- [54] Brenda Rubenstein. Introduction to the variational monte carlo method in quantum chemistry and physics. In *Variational Methods in Molecular Modeling*, pages 285–313. Springer, 2017.
- [55] M. H. Kalos. Monte carlo calculations of the ground state of three- and four-body nuclei. *Phys. Rev.*, 128:1791–1795, Nov 1962.
- [56] Michael A. Lee, K. E. Schmidt, M. H. Kalos, and G. V. Chester. Green’s function monte carlo method for liquid ^3He . *Phys. Rev. Lett.*, 46:728–731, Mar 1981.
- [57] Michael A. Lee and Kevin E. Schmidt. Green’s function monte carlo. *Computers in Physics*, 6(2):192–197, 1992.
- [58] Hale F Trotter. On the product of semi-groups of operators. *Proc. Am. Math. Soc.*, 10(4):545–551, 1959.

- [59] Masuo Suzuki. Generalized trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.*, 51(2):183–190, 1976.
- [60] James B Anderson. Quantum chemistry by random walk. *J. Chem. Phys.*, 65(10):4121–4127, 1976.
- [61] George H. Booth, Alex J. W. Thom, and Ali 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.
- [62] George H. Booth, Deidre Cleland, Alex J. W. Thom, and Ali Alavi. Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum monte carlo methods. *J. Chem. Phys.*, 135(8):084104, 2011.
- [63] DM Cleland, George H Booth, and Ali Alavi. A study of electron affinities using the initiator approach to full configuration interaction quantum monte carlo. *J. Chem. Phys.*, 134(2):024112, 2011.
- [64] Deidre Cleland, George H Booth, Catherine Overy, and Ali Alavi. Taming the first-row diatomics: a full configuration interaction quantum monte carlo study. *J. Chem. Theory Comput.*, 8(11):4138–4152, 2012.
- [65] Robert E. Thomas, George H. Booth, and Ali Alavi. Accurate ab initio calculation of ionization potentials of the first-row transition metals with the configuration-interaction quantum monte carlo technique. *Phys. Rev. Lett.*, 114:033001, Jan 2015.
- [66] Shiwei Zhang, J Carlson, and James E Gubernatis. Constrained path quantum monte carlo method for fermion ground states. *Phys. Rev. Lett.*, 74(18):3652, 1995.
- [67] Shiwei Zhang and Henry Krakauer. Quantum monte carlo method using phase-free random walks with slater determinants. *Phys. Rev. Lett.*, 90:136401, Apr 2003.

- [68] Shiwei Zhang. Ab initio electronic structure calculations by auxiliary-field quantum monte carlo. *Handbook of Materials Modeling: Methods: Theory and Modeling*, pages 1–27, 2018.
- [69] Mario Motta, James Shee, Shiwei Zhang, and Garnet Kin-Lic Chan. Efficient ab initio auxiliary-field quantum monte carlo calculations in gaussian bases via low-rank tensor decomposition. *J. Chem. Theory Comput.*, 2019. PMID: 31091103.
- [70] Mario Motta and Shiwei Zhang. Ab initio computations of molecular systems by the auxiliary-field quantum monte carlo method. *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, 8(5):e1364, 2018.
- [71] WA Al-Saidi, Shiwei Zhang, and Henry Krakauer. Auxiliary-field quantum monte carlo calculations of molecular systems with a gaussian basis. *J. Chem. Phys.*, 124(22):224101, 2006.
- [72] WA Al-Saidi, Henry Krakauer, and Shiwei Zhang. Auxiliary-field quantum monte carlo study of first-and second-row post-d elements. *J. Chem. Phys.*, 125(15):154110, 2006.
- [73] Malliga Suewattana, Wirawan Purwanto, Shiwei Zhang, Henry Krakauer, and Eric J. Walter. Phaseless auxiliary-field quantum monte carlo calculations with plane waves and pseudopotentials: Applications to atoms and molecules. *Phys. Rev. B*, 75:245123, Jun 2007.
- [74] WA Al-Saidi, Henry Krakauer, and Shiwei Zhang. A study of $\text{h} + \text{h}_2$ and several h-bonded molecules by phaseless auxiliary-field quantum monte carlo with plane wave and gaussian basis sets. *J. Chem. Phys.*, 126(19):194105, 2007.
- [75] Wirawan Purwanto, WA Al-Saidi, Henry Krakauer, and Shiwei Zhang. Eliminating spin contamination in auxiliary-field quantum monte carlo: Realistic potential energy curve of f_2 . *J. Chem. Phys.*, 128(11):114309, 2008.
- [76] Wirawan Purwanto, Henry Krakauer, and Shiwei Zhang. Pressure-induced diamond to β -tin transition in bulk silicon: A quantum monte carlo study. *Phys. Rev. B*, 80:214116, Dec 2009.

- [77] Wirawan Purwanto, Henry Krakauer, Yudistira Virgus, and Shiwei Zhang. Assessing weak hydrogen binding on Ca^+ centers: An accurate many-body study with large basis sets. *J. Chem. Phys.*, 135(16):164105, 2011.
- [78] Fengjie Ma, Shiwei Zhang, and Henry Krakauer. Auxiliary-field quantum monte carlo calculations with multiple-projector pseudopotentials. *Phys. Rev. B*, 95(16):165103, 2017.
- [79] Mario Motta and Shiwei Zhang. Computation of ground-state properties in molecular systems: back-propagation with auxiliary-field quantum monte carlo. *J. Chem. Theory Comput.*, 13(11):5367–5378, 2017.
- [80] Wissam A Al-Saidi, Shiwei Zhang, and Henry Krakauer. Bond breaking with auxiliary-field quantum monte carlo. *J. Chem. Phys.*, 127(14):144101, 2007.
- [81] Mario Motta, David M. Ceperley, Garnet Kin-Lic Chan, John A. Gomez, Emanuel Gull, Sheng Guo, Carlos A. Jiménez-Hoyos, Tran Nguyen Lan, Jia Li, Fengjie Ma, Andrew J. Millis, Nikolay V. Prokof'ev, Ushnish Ray, Gustavo E. Scuseria, Sandro Sorella, Edwin M. Stoudenmire, Qiming Sun, Igor S. Tupitsyn, Steven R. White, Dominika Zgid, and Shiwei Zhang. Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods. *Phys. Rev. X*, 7:031059, Sep 2017.
- [82] Bo-Xiao Zheng, Chia-Min Chung, Philippe Corboz, Georg Ehlers, Ming-Pu Qin, Reinhard M Noack, Hao Shi, Steven R White, Shiwei Zhang, and Garnet Kin-Lic Chan. Stripe order in the underdoped region of the two-dimensional hubbard model. *Science*, 358(6367):1155–1160, 2017.
- [83] WA Al-Saidi, Henry Krakauer, and Shiwei Zhang. Auxiliary-field quantum Monte Carlo study of TiO and MnO molecules. *Phys. Rev. B*, 73(7):075103, 2006.
- [84] Yudistira Virgus, Wirawan Purwanto, Henry Krakauer, and Shiwei Zhang. Ab initio many-body study of cobalt adatoms adsorbed on graphene. *Phys. Rev. B*, 86(24):241406, 2012.

- [85] Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. Frozen-orbital and downfolding calculations with auxiliary-field quantum monte carlo. *J. Chem. Theory Comput.*, 9(11):4825–4833, 2013.
- [86] Yudistira Virgus, Wirawan Purwanto, Henry Krakauer, and Shiwei Zhang. Stability, energetics, and magnetic states of cobalt adatoms on graphene. *Phys. Rev. Lett.*, 113(17):175502, 2014.
- [87] Fengjie Ma, Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. Quantum monte carlo calculations in solids with downfolded hamiltonians. *Phys. Rev. Lett.*, 114(22):226401, 2015.
- [88] Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. Auxiliary-field quantum monte carlo calculations of the molybdenum dimer. *J. Chem. Phys.*, 144(24):244306, 2016.
- [89] Shuai Zhang, Fionn D Malone, and Miguel A Morales. Auxiliary-field quantum monte carlo calculations of the structural properties of nickel oxide. *J. Chem. Phys.*, 149(16):164102, 2018.
- [90] RL Stratonovich. A method for the computation of quantum distribution functions. *Dokl. Akad. Nauk SSSR*, 115(6):1097–1100, 1957.
- [91] J. Hubbard. Calculation of partition functions. *Phys. Rev. Lett.*, 3:77–78, Jul 1959.
- [92] D. R. Hamann and S. B. Fahy. Energy measurement in auxiliary-field many-electron calculations. *Phys. Rev. B*, 41:11352–11363, Jun 1990.
- [93] Huy Nguyen. Quantum monte carlo calculation of the imaginary-time greens function in the hubbard model. Undergraduate thesis, Reed College, 2014.
- [94] Brenda M. Rubenstein, Shiwei Zhang, and David R. Reichman. Finite-temperature auxiliary-field quantum monte carlo technique for bose-fermi mixtures. *Phys. Rev. A*, 86:053606, Nov 2012.
- [95] Hao Shi and Shiwei Zhang. Symmetry in auxiliary-field quantum monte carlo calculations. *Phys. Rev. B*, 88:125132, Sep 2013.

- [96] Shiwei Zhang, J. Carlson, and J. E. Gubernatis. Constrained path monte carlo method for fermion ground states. *Phys. Rev. B*, 55:7464–7477, Mar 1997.
- [97] Shiwei Zhang. Auxiliary-field quantum monte carlo for correlated electron systems. *Emergent Phenomena in Correlated Matter*, 3, 2013.
- [98] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B*, 41:9301–9307, May 1990.
- [99] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Phys. Rev. Lett.*, 94:170201, May 2005.
- [100] M Motta, Davide Emilio Galli, Saverio Moroni, and Eleonora Vitali. Imaginary time correlations and the phaseless auxiliary field quantum monte carlo. *J. Chem. Phys.*, 140(2):024107, 2014.
- [101] Wirawan Purwanto and Shiwei Zhang. Quantum monte carlo method for the ground state of many-boson systems. *Phys. Rev. E*, 70:056702, Nov 2004.
- [102] CJ Umrigar. Two aspects of quantum monte carlo: determination of accurate wavefunctions and determination of potential energy surfaces of molecules. *Int. J. Quantum Chem.*, 36(S23):217–230, 1989.
- [103] CA Traynor and James B Anderson. Parallel monte carlo calculations to determine energy differences among similar molecular structures. *Chem. Phys. Lett.*, 147(4):389–394, 1988.
- [104] Bryan H Wells. The differential green’s function monte carlo method. the dipole moment of lih. *Chem. Phys. Lett.*, 115(1):89–94, 1985.
- [105] Jan Vrbik, Daniel A Legare, and Stuart M Rothstein. Infinitesimal differential diffusion quantum monte carlo: Diatomic molecular properties. *J. Chem. Phys.*, 92(2):1221–1227, 1990.

- [106] Claudia Filippi and C. J. Umrigar. Correlated sampling in quantum monte carlo: A route to forces. *Phys. Rev. B*, 61:R16291–R16294, Jun 2000.
- [107] Yongkyung Kwon, D. M. Ceperley, and Richard M. Martin. Quantum monte carlo calculation of the fermi-liquid parameters in the two-dimensional electron gas. *Phys. Rev. B*, 50:1684–1694, Jul 1994.
- [108] David R Garmer. Extrapolation of the time-step bias in diffusion quantum monte carlo by a differential sampling technique. *J. Comput. Chem.*, 10(2):176–185, 1989.
- [109] David L Nelson, Albert L Lehninger, and Michael M Cox. *Lehninger principles of biochemistry*. Macmillan, 2008.
- [110] Bernard Meunier, Samuel P De Visser, and Sason Shaik. Mechanism of oxidation reactions catalyzed by cytochrome p450 enzymes. *Chem. Rev.*, 104(9):3947–3980, 2004.
- [111] M Rosa Palacin. Recent advances in rechargeable battery materials: a chemist’s perspective. *Chem. Soc. Rev.*, 38(9):2565–2575, 2009.
- [112] Michael Gratzel. *Energy resources through photochemistry and catalysis*. Elsevier, 2012.
- [113] Cyrille Costentin, Marc Robert, and Jean-Michel Savéant. Catalysis of the electrochemical reduction of carbon dioxide. *Chem. Soc. Rev.*, 42(6):2423–2436, 2013.
- [114] Emil Alexov, Ernest L Mehler, Nathan Baker, António M Baptista, Yong Huang, Francesca Milletti, Jens Erik Nielsen, Damien Farrell, Tommy Carstensen, Mats HM Olsson, Jana K. Shen, Jim Warwicker, Sarah Williams, and Michael J. Word. Progress in the prediction of pka values in proteins. *Proteins: Struct., Funct., Bioinf.*, 79(12):3260–3275, 2011.
- [115] Steven V Jerome, Thomas F Hughes, and Richard A Friesner. Accurate pka prediction in first-row hexaaqua transition metal complexes using the b3lyp-dbloc method. *J. Phys. Chem. B*, 118(28):8008–8016, 2014.
- [116] Art D. Bochevarov, Mark A. Watson, Jeremy R. Greenwood, and Dean M. Philipp. Multi-conformation, density functional theory-based pka prediction in application to large, flexible

- organic molecules with diverse functional groups. *J. Chem. Theory Comput.*, 12(12):6001–6019, 2016. PMID: 27951674.
- [117] Hemaka A Rajapakse, Philippe G Nantermet, Harold G Selnick, James C Barrow, Georgia B McGaughey, Sanjeev Munshi, Stacey R Lindsley, Mary Beth Young, Phung L Ngo, M Katherine Holloway, Ming-Tain Laid, Amy S. Espesethd, Xiao-Ping Shid, Dennis Colussid, Beth Pietrakd, Ming-Chih Crouthameld, Katherine Tugushevad, Qian Huangd, Min Xud, Adam J. Simond, Lawrence Kuoc, Daria J. Hazudad, Samuel Grahama, and Joseph P. Vaccaa. Sar of tertiary carbinamine derived bace1 inhibitors: Role of aspartate ligand amine pka in enzyme inhibition. *Bioorg. Med. Chem. Lett.*, 20(6):1885–1889, 2010.
- [118] DG Sprous, RK Palmer, JT Swanson, and M Lawless. Qsar in the pharmaceutical research setting: Qsar models for broad, large problems. *Curr. Top. Med. Chem.*, 10(6):619–637, 2010.
- [119] Jun Cheng and Michiel Sprik. Acidity of the aqueous rutile tio2 (110) surface from density functional theory based molecular dynamics. *J. Chem. Theory Comput.*, 6(3):880–889, 2010.
- [120] Dennis Roman Gallus, Ralf Wagner, Simon Wiemers-Meyer, Martin Winter, and Isidora Cekic-Laskovic. New insights into the structure-property relationship of high-voltage electrolyte components for lithium-ion batteries using the pka value. *Electrochim. Acta*, 184:410–416, 2015.
- [121] Vyacheslav S Bryantsev. Predicting the stability of aprotic solvents in li-air batteries: pka calculations of aliphatic c–h acids in dimethyl sulfoxide. *Chem. Phys. Lett.*, 558:42–47, 2013.
- [122] Jamal El Yazal, Franklyn G Prendergast, David Elliot Shaw, and Yuan-Ping Pang. Protonation states of the chromophore of denatured green fluorescent proteins predicted by ab initio calculations. *J. Am. Chem. Soc.*, 122(46):11411–11415, 2000.
- [123] Jia Chen, Ye-Fei Li, Patrick Sit, and Annabella Selloni. Chemical dynamics of the first proton-coupled electron transfer of water oxidation on tio2 anatase. *J. Am. Chem. Soc.*, 135(50):18774–18777, 2013.

- [124] George A Olah and Arpad Molnar. *Hydrocarbon chemistry*. John Wiley & Sons, 2003.
- [125] Hai Wang and Michael Frenklach. Calculations of rate coefficients for the chemically activated reactions of acetylene with vinylic and aromatic radicals. *J. Phys. Chem.*, 98(44):11465–11489, 1994.
- [126] Michael Page and Donald W Brenner. Hydrogen abstraction from a diamond surface. ab initio quantum chemical study with constrained isobutane as a model. *J. Am. Chem. Soc.*, 113(9):3270–3274, 1991.
- [127] Graham W Burton and Keith U Ingold. Vitamin e as an in vitro and in vivo antioxidant. *Ann. N. Y. Acad. Sci.*, 570(1):7–22, 1989.
- [128] James M Mayer. Hydrogen atom abstraction by metal-oxo complexes: Understanding the analogy with organic radical reactions. *Acc. Chem. Res.*, 31(8):441–450, 1998.
- [129] Margareta RA Blomberg, Per EM Siegbahn, Stenbjörn Styring, Gerald T Babcock, Björn Åkermark, and Peter Korall. A quantum chemical study of hydrogen abstraction from manganese-coordinated water by a tyrosyl radical: A model for water oxidation in photosystem ii. *J. Am. Chem. Soc.*, 119(35):8285–8292, 1997.
- [130] Robert K Grasselli. Advances and future trends in selective oxidation and ammoxidation catalysis. *Catal. Today*, 49(1):141–153, 1999.
- [131] Barry B Snider. Manganese (iii)-based oxidative free-radical cyclizations. *Chem. Rev.*, 96(1):339–364, 1996.
- [132] Harold Basch and Shmaryahu Hoz. Ab initio study of hydrogen abstraction reactions. *J. Phys. Chem. A*, 101(24):4416–4431, 1997.
- [133] Michelle L Coote. Reliable theoretical procedures for the calculation of electronic-structure information in hydrogen abstraction reactions. *J. Phys. Chem. A*, 108(17):3865–3872, 2004.

- [134] Jingzhi Pu and Donald G Truhlar. Benchmark calculations of reaction energies, barrier heights, and transition-state geometries for hydrogen abstraction from methanol by a hydrogen atom. *J. Phys. Chem. A*, 109(5):773–778, 2005.
- [135] Edson FV Carvalho, Alessandra N Barauna, Francisco BC Machado, and Orlando Roberto-Neto. Dft study for the reactions of h atoms with ch3oh and c2h5oh. *Int. J. Quantum Chem.*, 108(13):2476–2485, 2008.
- [136] EFV Carvalho, Alessandra N Barauna, Francisco BC Machado, and Orlando Roberto-Neto. Theoretical calculations of energetics, structures, and rate constants for the h + ch3oh hydrogen abstraction reactions. *Chem. Phys. Lett.*, 463(1):33–37, 2008.
- [137] Francesco Fracchia, Claudia Filippi, and Claudio Amovilli. Barrier heights in quantum monte carlo with linear-scaling generalized-valence-bond wave functions. *J. Chem. Theory Comput.*, 9(8):3453–3462, 2013.
- [138] Yosuke Kanai and Noboru Takeuchi. Toward accurate reaction energetics for molecular line growth at surface: Quantum monte carlo and density functional theory calculations. *J. Chem. Phys.*, 131(21):214708, 2009.
- [139] AC Kollias, O Couronne, and WA Lester Jr. Quantum monte carlo study of the reaction: Cl + ch3oh → ch2oh + hcl. *J. Chem. Phys.*, 121(3):1357–1363, 2004.
- [140] Larry A Curtiss, Paul C Redfern, Krishnan Raghavachari, and John A Pople. Assessment of gaussian-2 and density functional theories for the computation of ionization potentials and electron affinities. *J. Chem. Phys.*, 109(1):42–55, 1998.
- [141] H. Flyvbjerg and H. G. Petersen. Error estimates on averages of correlated data. *J. Chem. Phys.*, 91(1):461–466, 1989.
- [142] Huy Nguyen, Hao Shi, Jie Xu, and Shiwei Zhang. Cpmc-lab: A matlab package for constrained path monte carlo calculations. *Comput. Phys. Commun.*, 185(12):3344–3357, 2014.

- [143] Marat Valiev, Eric J Bylaska, Niranjana Govind, Karol Kowalski, Tjerk P Straatsma, Hubertus JJ Van Dam, Dunyong Wang, Jarek Nieplocha, Edoardo Apra, Theresa L Windus, and Wibe A. de Jong. Nwchem: a comprehensive and scalable open-source solution for large scale molecular simulations. *Comput. Phys. Commun.*, 181(9):1477–1489, 2010.
- [144] Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhen-dong Li, Junzi Liu, James D. McClain, Elvira R. Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. Pyscf: the pythonbased simulations of chemistry framework. *WIREs Comput. Mol. Sci.*, 8(1):e1340, 2018.
- [145] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar. Numerical study of the two-dimensional hubbard model. *Phys. Rev. B*, 40:506–516, Jul 1989.
- [146] Michael W Schmidt, Kim K Baldridge, Jerry A Boatz, Steven T Elbert, Mark S Gordon, Jan H Jensen, Shiro Koseki, Nikita Matsunaga, Kiet A Nguyen, Shujun Su, Theresa L. Windus, Michel Dupuis, and John A. Montgomery Jr. General atomic and molecular electronic structure system. *J. Comput. Chem.*, 14(11):1347–1363, 1993.
- [147] Thom H Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen. *J. Chem. Phys.*, 90(2):1007–1023, 1989.
- [148] Rick A Kendall, Thom H Dunning Jr, and Robert J Harrison. Electron affinities of the first-row atoms revisited. systematic basis sets and wave functions. *J. Chem. Phys.*, 96(9):6796–6806, 1992.
- [149] Russell J. Johnson III. NIST Computational Chemistry Comparison and Benchmark Database. <http://cccbdb.nist.gov/>. NIST Standard Reference Database Number 101, Release 17b, September 2015.
- [150] Qian Wang, Xianyi Zhang, Yunquan Zhang, and Qing Yi. Augem: Automatically generate high performance dense linear algebra kernels on x86 cpus. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC ’13, pages 25:1–25:12, New York, NY, USA, 2013. ACM.

- [151] Z. Xianyi, W. Qian, and Z. Yunquan. Model-driven level 3 blas performance optimization on loongson 3a processor. In *2012 IEEE 18th International Conference on Parallel and Distributed Systems*, pages 684–691, Dec 2012.
- [152] Roger B Sidje. Expokit: a software package for computing matrix exponentials. *ACM Trans. Math. Softw.*, 24(1):130–156, 1998.
- [153] Michael Mascagni and Ashok Srinivasan. Algorithm 806: Sprng: A scalable library for pseudorandom number generation. *ACM Trans. Math. Softw.*, 26(3):436–461, 2000.
- [154] Larry A Curtiss, Krishnan Raghavachari, Gary W Trucks, and John A Pople. Gaussian-2 theory for molecular energies of first-and second-row compounds. *J. Chem. Phys.*, 94(11):7221–7230, 1991.
- [155] Axel D Becke. Density-functional thermochemistry. iii. the role of exact exchange. *J. Chem. Phys.*, 98(7):5648–5652, 1993.
- [156] Chengteh Lee, Weitao Yang, and Robert G. Parr. Development of the colle-salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B*, 37:785–789, Jan 1988.
- [157] GJ Tawa, IA Topol, SK Burt, RA Caldwell, and AA Rashin. Calculation of the aqueous solvation free energy of the proton. *J. Chem. Phys.*, 109:4852–4863, 1998.
- [158] Grant N Merrill and Steven R Kass. Calculated gas-phase acidities using density functional theory: is it reliable? *J. Phys. Chem.*, 100(44):17465–17471, 1996.
- [159] Larry A Curtiss, L David Kock, and John A Pople. Energies of ch₂oh, ch₃o, and related compounds. *J. Chem. Phys.*, 95(6):4040–4043, 1991.
- [160] Ivan S Ufimtsev and Todd J Martinez. Quantum chemistry on graphical processing units. 2. direct self-consistent-field implementation. *J. Chem. Theory Comput.*, 5(4):1004–1015, 2009.
- [161] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015.

- [162] Joshua A Anderson, Chris D Lorenz, and Alex Traveset. General purpose molecular dynamics simulations fully implemented on graphics processing units. *J. Comput. Phys.*, 227(10):5342–5359, 2008.
- [163] Ivan S Ufimtsev and Todd J Martinez. Quantum chemistry on graphical processing units. 1. strategies for two-electron integral evaluation. *J. Chem. Theory Comput.*, 4(2):222–231, 2008.
- [164] Nathan Luehr, Ivan S Ufimtsev, and Todd J Martínez. Dynamic precision for electron repulsion integral evaluation on graphical processing units (gpu). *J. Chem. Theory Comput.*, 7(4):949–954, 2011.
- [165] Jaroslaw Kalinowski, Frank Wennmohs, and Frank Neese. Arbitrary angular momentum electron repulsion integrals with graphical processing units: Application to the resolution of identity hartree-fock method. *J. Chem. Theory Comput.*, 2017.
- [166] Jorg Kussmann and Christian Ochsenfeld. Hybrid cpu/gpu integral engine for strong-scaling ab initio methods. *J. Chem. Theory Comput.*, 13(7):3153–3159, 2017.
- [167] Ivan S Ufimtsev and Todd J Martinez. Graphical processing units for quantum chemistry. *Comput. Sci. Eng.*, 10(6):26–34, 2008.
- [168] Takeshi Yoshikawa and Hiromi Nakai. Linear-scaling self-consistent field calculations based on divide-and-conquer method using resolution-of-identity approximation on graphical processing units. *J. Comput. Chem.*, 36(3):164–170, 2015.
- [169] Koji Yasuda. Accelerating density functional calculations with graphics processing unit. *J. Chem. Theory Comput.*, 4(8):1230–1236, 2008.
- [170] Luigi Genovese, Matthieu Ospici, Thierry Deutsch, Jean-François Méhaut, Alexey Neelov, and Stefan Goedecker. Density functional theory calculation on many-cores hybrid central processing unit-graphic processing unit architectures. *J. Chem. Phys.*, 131(3):034103, 2009.

- [171] Mohamed Hacene, Ani Anciaux-Sedrakian, Xavier Rozanska, Diego Klahr, Thomas Guignon, and Paul Fleurat-Lessard. Accelerating vasp electronic structure calculations using graphic processing units. *J. Comput. Chem.*, 33(32):2581–2589, 2012.
- [172] Jeremy S Meredith, Gonzalo Alvarez, Thomas A Maier, Thomas C Schulthess, and Jeffrey S Vetter. Accuracy and performance of graphics processors: A quantum monte carlo application case study. *Parallel Comput.*, 35(3):151–163, 2009.
- [173] Tal Levy, Guy Cohen, and Eran Rabani. Simulating lattice spin models on graphics processing units. *J. Chem. Theory Comput.*, 6(11):3293–3301, 2010.
- [174] Benjamin Block, Peter Virnau, and Tobias Preis. Multi-gpu accelerated multi-spin monte carlo simulations of the 2d ising model. *Comput. Phys. Commun.*, 181(9):1549–1556, 2010.
- [175] Roberto Olivares-Amaya, Mark A Watson, Richard G Edgar, Leslie Vogt, Yihan Shao, and Alán Aspuru-Guzik. Accelerating correlated quantum chemistry calculations using graphical processing units and a mixed precision matrix multiplication library. *J. Chem. Theory Comput.*, 6(1):135–144, 2009.
- [176] Alexander E Doran and So Hirata. Monte carlo mp2 on many graphical processing units. *J. Chem. Theory Comput.*, 12(10):4821–4832, 2016.
- [177] Chenchen Song and Todd J Martínez. Atomic orbital-based sos-mp2 with tensor hypercontraction. i. gpu-based tensor construction and exploiting sparsity. *J. Chem. Phys.*, 144(17):174111, 2016.
- [178] Leslie Vogt, Roberto Olivares-Amaya, Sean Kermes, Yihan Shao, Carlos Amador-Bedolla, and Alán Aspuru-Guzik. Accelerating resolution-of-the-identity second-order møller- ples-set quantum chemistry calculations with graphical processing units. *J. Phys. Chem. A*, 112(10):2049–2057, 2008.
- [179] A Eugene DePrince III and Jeff R Hammond. Coupled cluster theory on graphics processing units i. the coupled cluster doubles method. *J. Chem. Theory Comput.*, 7(5):1287–1295, 2011.

- [180] Christine M Isborn, Nathan Luehr, Ivan S Ufimtsev, and Todd J Martínez. Excited-state electronic structure with configuration interaction singles and tamm-dancoff time-dependent density functional theory on graphical processing units. *J. Chem. Theory Comput.*, 7(6):1814–1823, 2011.
- [181] B Scott Fales and Benjamin G Levine. Nanoscale multireference quantum chemistry: Full configuration interaction on graphical processing units. *J. Chem. Theory Comput.*, 11(10):4708–4716, 2015.
- [182] B Scott Fales, Yinan Shu, Benjamin G Levine, and Edward G Hohenstein. Complete active space configuration interaction from state-averaged configuration interaction singles natural orbitals: Analytic first derivatives and derivative coupling vectors. *J. Chem. Phys.*, 147(9):094104, 2017.
- [183] Edward G Hohenstein, Nathan Luehr, Ivan S Ufimtsev, and Todd J Martínez. An atomic orbital-based formulation of the complete active space self-consistent field method on graphical processing units. *J. Chem. Phys.*, 142(22):224103, 2015.
- [184] James W Snyder Jr, B Scott Fales, Edward G Hohenstein, Benjamin G Levine, and Todd J Martínez. A direct-compatible formulation of the coupled perturbed complete active space self-consistent field equations on graphical processing units. *J. Chem. Phys.*, 146(17):174113, 2017.
- [185] Ivan S Ufimtsev and Todd J Martinez. Quantum chemistry on graphical processing units. 3. analytical energy gradients, geometry optimization, and first principles molecular dynamics. *J. Chem. Theory Comput.*, 5(10):2619–2628, 2009.
- [186] Chenchen Song and Todd J Martínez. Analytical gradients for tensor hyper-contracted mp2 and sos-mp2 on graphical processing units. *J. Chem. Phys.*, 147(16):161723, 2017.
- [187] Ilya A Kaliman and Anna I Krylov. New algorithm for tensor contractions on multi-core cpus, gpus, and accelerators enables ccsd and eom-ccsd calculations with over 1000 basis functions on a single compute node. *J. Comput. Chem.*, 38(11):842–853, 2017.

- [188] Amos G Anderson, William A Goddard, and Peter Schröder. Quantum monte carlo on graphical processing units. *Comput. Phys. Commun.*, 177(3):298–306, 2007.
- [189] Kenneth Esler, Jeongnim Kim, David Ceperley, and Luke Shulenburger. Accelerating quantum monte carlo simulations of real materials on gpu clusters. *Comput. Sci. Eng.*, 14(1):40–51, 2012.
- [190] Jeongnim Kim, Andrew Baczewski, Todd Beaudet, Anouar Benali, Chandler Bennett, Mark Berrill, Nick Blunt, Edgar Josue Landinez Borda, Michele Casula, David Ceperley, Simone Chiesa, bryan K clark, Raymond Clay, Kris Delaney, Mark Dewing, Ken Esler, Hongxia Hao, Olle Heinonen, Paul R C Kent, Jaron T. Krogel, Ilkka Kylanpaa, Ying Wai Li, M. Graham Lopez, Ye Luo, Fionn Malone, Richard Martin, Amrita Mathuriya, Jeremy McMinis, Cody Melton, Lubos Mitas, Miguel A. Morales, Eric Neuscamman, William Parker, Sergio Flores, Nichols A Romero, Brenda Rubenstein, Jacqueline Shea, Hyeondeok Shin, Luke Shulenburger, Andreas Tillack, Joshua Townsend, Norman Tubman, Brett van der Goetz, Jordan Vincent, D. ChangMo Yang, Yubo Yang, Shuai Zhang, and Luning Zhao. Qmcpack : An open source ab initio quantum monte carlo package for the electronic structure of atoms, molecules, and solids. *J. Phys. Condens. Mat.*, 2018.
- [191] Bryan K Clark, Miguel A Morales, Jeremy McMinis, Jeongnim Kim, and Gustavo E Scuseria. Computing the energy of a water molecule using multideterminants: A simple, efficient algorithm. *J. Chem. Phys.*, 135(24):244105, 2011.
- [192] T McDaniel, EF DAzevedo, YW Li, K Wong, and PRC Kent. Delayed slater determinant update algorithms for high efficiency quantum monte carlo. *J. Chem. Phys.*, 147(17):174107, 2017.
- [193] Hans-Joachim Werner, Frederick R Manby, and Peter J Knowles. Fast linear scaling second-order møller-plesset perturbation theory (mp2) using local and density fitting approximations. *The Journal of chemical physics*, 118(18):8149–8160, 2003.
- [194] Johannes Hachmann, Wim Cardoen, and Garnet Kin-Lic Chan. Multireference correlation

- in long molecules with the quadratic scaling density matrix renormalization group. *J. Chem. Phys.*, 125(14):144101, 2006.
- [195] Takashi Tsuchimochi and Gustavo E. Scuseria. Strong correlations via constrained-pairing mean-field theory. *J. Chem. Phys.*, 131(12):121102, 2009.
- [196] Anton V Sinitskiy, Loren Greenman, and David A Mazziotti. Strong correlation in hydrogen chains and lattices using the variational two-electron reduced density matrix method. *J. Chem. Phys.*, 133(1):014104, 2010.
- [197] Nan Lin, CA Marianetti, Andrew J Millis, and David R Reichman. Dynamical mean-field theory for quantum chemistry. *Phys. Rev. Lett.*, 106(9):096402, 2011.
- [198] Lorenzo Stella, Claudio Attaccalite, Sandro Sorella, and Angel Rubio. Strong electronic correlation in the hydrogen chain: A variational monte carlo study. *Phys. Rev. B*, 84(24):245117, 2011.
- [199] Florian Weigend. Accurate coulomb-fitting basis sets for h to rn. *Phys. Chem. Chem. Phys.*, 8(9):1057–1065, 2006.
- [200] David B Krisiloff, Caroline M Krauter, Francis J Ricci, and Emily A Carter. Density fitting and cholesky decomposition of the two-electron integrals in local multireference configuration interaction theory. *J. Chem. Theory Comput.*, 11(11):5242–5251, 2015.
- [201] Francesco Aquilante, Roland Lindh, and Thomas Bondo Pedersen. Unbiased auxiliary basis sets for accurate two-electron integral approximations. *J. Chem. Phys.*, 127(11):114107, 2007.
- [202] Raghunath O Ramabhadran and Krishnan Raghavachari. Extrapolation to the gold-standard in quantum chemistry: Computationally efficient and accurate ccsd(t) energies for large molecules using an automated thermochemical hierarchy. *J. Chem. Theor. Comput.*, 9(9):3986–3994, 2013.
- [203] John A Pople. Nobel lecture: Quantum chemical models. *Rev. Mod. Phys.*, 71(5):1267, 1999.

- [204] Nikolai B Balabanov and Kirk A Peterson. Systematically convergent basis sets for transition metals. i. all-electron correlation consistent basis sets for the 3d elements sc–zn. *J. Chem. Phys.*, 123(6):064107, 2005.
- [205] Werner Kutzelnigg and Wenjian Liu. Quasirelativistic theory equivalent to fully relativistic theory. *J. Chem. Phys.*, 123(24):241102, 2005.
- [206] Daoling Peng and Markus Reiher. Exact decoupling of the relativistic fock operator. *Theor. Chem. Acc.*, 131(1):1081, 2012.
- [207] Trygve Helgaker, Wim Klopper, Henrik Koch, and Jozef Noga. Basis-set convergence of correlated calculations on water. *J. Chem. Phys.*, 106(23):9639–9646, 1997.
- [208] Charles W Bauschlicher, Per Siegbahn, and Lars GM Pettersson. The atomic states of nickel. *Theor. Chim. Acta*, 74(6):479–491, 1988.
- [209] Kerstin Andersson and Björn O Roos. Excitation energies in the nickel atom studied with the complete active space scf method and second-order perturbation theory. *Chem. Phys. Lett.*, 191(6):507–514, 1992.
- [210] Qiming Sun, Jun Yang, and Garnet Kin-Lic Chan. A general second order complete active space self-consistent-field solver for large-scale systems. *Chem. Phys. Lett.*, 683:291–299, 2017.
- [211] James Shee, Shiwei Zhang, David R Reichman, and Richard A Friesner. Chemical transformations approaching chemical accuracy via correlated sampling in auxiliary-field quantum monte carlo. *J. Chem. Theory Comput.*, 13(6):2667–2680, 2017.
- [212] Alfred Trautwein. *Bioinorganic Chemistry: Transition Metals in Biology and Their Coordination Chemistry*. Wiley-VCH, 1997.
- [213] D.I. Khomskii. *Transition Metal Compounds*. Cambridge University Press, 2014.
- [214] Mikhail Askerka, Gary W. Brudvig, and Victor S. Batista. The O₂-Evolving Complex of Photosystem II: Recent Insights from Quantum Mechanics/Molecular Mechanics (QM/MM),

- Extended X-ray Absorption Fine Structure (EXAFS), and Femtosecond X-ray Crystallography Data. *Acc. Chem. Res.*, 50:41–48, December 2016.
- [215] Richard A Friesner and Steven V Jerome. Localized orbital corrections for density functional calculations on transition metal containing systems. *Coord. Chem. Rev.*, 344:205–213, 2017.
 - [216] Tom Ziegler. Approximate density functional theory as a practical tool in molecular energetics and dynamics. *Chem. Rev.*, 91(5):651–667, 1991.
 - [217] David H Bross, J Grant Hill, H-J Werner, and Kirk A Peterson. Explicitly correlated composite thermochemistry of transition metal species. *J. Chem Phys.*, 139(9):094302, 2013.
 - [218] Sivabalan Manivasagam, Marie L Laury, and Angela K Wilson. Pseudopotential-based correlation consistent composite approach (rp-ccca) for first-and second-row transition metal thermochemistry. *J. Phys. Chem. A*, 119(26):6867–6874, 2015.
 - [219] Wanyi Jiang, Marie L Laury, Mitchell Powell, and Angela K Wilson. Comparative study of single and double hybrid density functionals for the prediction of 3d transition metal thermochemistry. *J. Comput. Theor. Comput.*, 8(11):4102–4111, 2012.
 - [220] Wenjing Zhang, Donald G Truhlar, and Mingsheng Tang. Tests of exchange-correlation functional approximations against reliable experimental data for average bond energies of 3d transition metal compounds. *J. Chem. Theor. Comput.*, 9(9):3965–3977, 2013.
 - [221] Klaus August Moltved and Kasper P Kepp. Chemical bond energies of 3d transition metals studied by density functional theory. *J. Chem. Theor. Comput.*, 2018.
 - [222] Wanyi Jiang, Nathan J DeYonker, John J Determan, and Angela K Wilson. Toward accurate theoretical thermochemistry of first row transition metal complexes. *J. Phys. Chem. A*, 116(2):870–885, 2011.
 - [223] Rebecca K Carlson, Giovanni Li Manni, Andrew L Sonnenberger, Donald G Truhlar, and Laura Gagliardi. Multiconfiguration pair-density functional theory: Barrier heights and main group and transition metal energetics. *J. Chem. Theor. Comp.*, 11(1):82–90, 2014.

- [224] Junwei Lucas Bao, Samuel O Odoh, Laura Gagliardi, and Donald G Truhlar. Predicting bond dissociation energies of transition-metal compounds by multiconfiguration pair-density functional theory and second-order perturbation theory based on correlated participating orbitals and separated pairs. *J. Chem. Theor. Comput.*, 13(2):616–626, 2017.
- [225] Junwei Lucas Bao, Xin Zhang, Xuefei Xu, and Donald G Truhlar. Predicting bond dissociation energy and bond length for bimetallic diatomic molecules: A challenge for electronic structure theory. *Phys. Chem. Chem. Phys.*, 19(8):5839–5854, 2017.
- [226] Kamal Sharkas, Laura Gagliardi, and Donald G Truhlar. Multiconfiguration pair-density functional theory and complete active space second order perturbation theory. bond dissociation energies of fec, nic, fes, nis, fese, and nise. *J. Phys. Chem. A*, 121(48):9392–9400, 2017.
- [227] Lan Nguyen Tran, Sergei Iskakov, and Dominika Zgid. Spin-unrestricted self-energy embedding theory. *J. Phys. Chem. Lett.*, 2018.
- [228] Larry A Curtiss, Krishnan Raghavachari, Paul C Redfern, and John A Pople. Assessment of gaussian-3 and density functional theories for a larger experimental test set. *J. Chem. Phys.*, 112(17):7374–7383, 2000.
- [229] Xuefei Xu, Wenjing Zhang, Mingsheng Tang, and Donald G Truhlar. Do practical standard coupled cluster calculations agree better than kohn–sham calculations with currently available functionals when compared to the best available experimental data for dissociation energies of bonds to 3d transition metals? *J. Chem. Theor. Comput.*, 11(5):2036–2052, 2015.
- [230] Zongtang Fang, Monica Vasiliu, Kirk A Peterson, and David A Dixon. Prediction of bond dissociation energies/heats of formation for diatomic transition metal compounds: Ccsd(t) works. *J. Chem. Theor. Comput.*, 13(3):1057–1066, 2017.
- [231] Lan Cheng, Jurgen Gauss, Branko Ruscic, Peter B Armentrout, and John F Stanton. Bond dissociation energies for diatomic molecules containing 3d transition metals: Benchmark

- scalar-relativistic coupled-cluster calculations for 20 molecules. *J. Chem. Theor. Comput.*, 13(3):1044–1056, 2017.
- [232] Richard A Friesner, Eric H Knoll, and Yixiang Cao. A localized orbital analysis of the thermochemical errors in hybrid density functional theory: Achieving chemical accuracy via a simple empirical correction scheme. *J. Chem. Phys.*, 125(12):124107, 2006.
- [233] David Feller and David A Dixon. Extended benchmark studies of coupled cluster theory through triple excitations. *J. Chem. Phys.*, 115(8):3484–3496, 2001.
- [234] Katharina Doblhoff-Dier, Jorg Meyer, Philip E Hoggan, Geert-Jan Kroes, and Lucas K Wagner. Diffusion monte carlo for accurate dissociation energies of 3d transition metal containing molecules. *J. Chem. Theor. Comput.*, 12(6):2583–2597, 2016.
- [235] Anthony Scemama, Yann Garniron, Michel Caffarel, and Pierre-François Loos. Deterministic construction of nodal surfaces within quantum monte carlo: The case of fes. *J. Chem Theor. Comput.*, 14(3):1395–1402, 2018.
- [236] James Shee, Evan J Arthur, Shiwei Zhang, David R Reichman, and Richard A Friesner. Phaseless auxiliary-field quantum monte carlo on graphical processing units. *J. Chem. Theor. Comput.*, 2018.
- [237] Michael D. Morse. Predissociation measurements of bond dissociation energies. *Acc. Chem. Res.*, 52(1):119–126, 2019.
- [238] Stephen R Langhoff and Charles W Bauschlicher Jr. Ab initio studies of transition metal systems. *Ann. Rev. Phys. Chem.*, 39(1):181–212, 1988.
- [239] J. Anglada, P. J. Bruna, S. D. Peyerimhoff, and R.J. Buenker. Mrd-ci calculations for the low-lying electronic states of scandiumhydride and titaniumhydride. *J. Molec. Struc. THEOCHEM*, 93:299 – 308, 1983.
- [240] Mario Motta and Shiwei Zhang. Communication: Calculation of interatomic forces and

- optimization of molecular geometry with auxiliary-field quantum monte carlo. *J. Chem. Phys.*, 148(18):181101, 2018.
- [241] Sílvia Simon, Miquel Duran, and JJ Dannenberg. How does basis set superposition error change the potential surfaces for hydrogen-bonded dimers? *J. Chem. Phys.*, 105(24):11024–11031, 1996.
- [242] H Kang and JL Beauchamp. Gas-phase studies of alkene oxidation by transition-metal oxides. ion-beam studies of cro+. *J. Am. Chem. Soc.*, 108(19):5663–5668, 1986.
- [243] J R Marquart and J Berkowitz. Dissociation energies of some metal sulfides. *J. Chem. Phys.*, 39(2):283–285, 1963.
- [244] G De Maria, P Goldfinger, L Malaspina, and V Piacente. Mass-spectrometric study of gaseous molecules. zns, znse and znte. *Trans. Faraday Soc.*, 61:2146–2152, 1965.
- [245] László von Szentpály. Atom-based thermochemistry: Predictions of the sublimation enthalpies of group 12 chalcogenides and the formation enthalpies of their polonides. *J. Phys. Chem. A*, 112(49):12695–12701, 2008.
- [246] K. A. Peterson, B. C. Shepler, and JM Singleton. The group 12 metal chalcogenides: An accurate multireference configuration interaction and coupled cluster study. *Mol. Phys.*, 105(9):1139–1155, 2007.
- [247] Aristotle Papakondylis. Ab initio study of the ground and excited states of the zinc sulfide diatomic system, zns. *Chem. Phys. Lett.*, 512(1-3):44–48, 2011.
- [248] Shilin Hou, Mei Long, Ronger Zheng, and Weiguo Sun. Studies on rotational spectra of the ground states of zno, zns, sise and site isotopic molecules using an isotopic error compensation approach. *Mol. Phys.*, 108(11):1491–1496, 2010.
- [249] LN Zack and Lucy M Ziurys. The pure rotational spectrum of zns ($x^1\sigma^+$). *J. Molec. Spec.*, 257(2):213–216, 2009.

- [250] Nathan J DeYonker, Kirk A Peterson, Gideon Steyl, Angela K Wilson, and Thomas R Cundari. Quantitative computational thermochemistry of transition metal species. *J.Phys. Chem. A*, 111(44):11269–11277, 2007.
- [251] Beatrice H Botch, Thom H Dunning Jr, and James F Harrison. Valence correlation in the s^2d^n , sd^{n+1} , and d^{n+2} states of the first-row transition metal atoms. *J. Chem. Phys.*, 75(7):3466–3476, 1981.
- [252] Anja Krieger-Liszkay. Singlet oxygen production in photosynthesis. *J. Exp. Bot.*, 56(411):337–346, 2005.
- [253] H Shaun Kwak, David J Giesen, Thomas F Hughes, Alexander Goldberg, Yixiang Cao, Jacob Gavartin, Steve Dixon, and Mathew D Halls. In silico evaluation of highly efficient organic light-emitting materials. In *Organic Light Emitting Materials and Devices XX*, volume 9941, page 994119. International Society for Optics and Photonics, 2016.
- [254] Nathan A. Romero and David A. Nicewicz. Organic photoredox catalysis. *Chem. Rev.*, 116(17):10075–10166, 2016. PMID: 27285582.
- [255] Millicent B. Smith and Josef Michl. Singlet fission. *Chem. Rev.*, 110(11):6891–6936, 2010. PMID: 21053979.
- [256] Lionel Salem and Colin Rowland. The electronic properties of diradicals. *Angew. Chem. Int. Ed.*, 11(2):92–111, 1972.
- [257] Manabu Abe. Diradicals. *Chem. Rev.*, 113(9):7011–7088, 2013. PMID: 23883325.
- [258] Robert A Moss, Matthew S Platz, and Maitland Jones Jr. *Reactive intermediate chemistry*. John Wiley & Sons, 2004.
- [259] Didier Bourissou, Olivier Guerret, Franois P. Gabba, and Guy Bertrand. Stable carbenes. *Chem. Rev.*, 100(1):39–92, 2000. PMID: 11749234.

- [260] Darrin M. Flanigan, Fedor Romanov-Michailidis, Nicholas A. White, and Tomislav Rovis. Organocatalytic reactions enabled by n-heterocyclic carbenes. *Chem. Rev.*, 115(17):9307–9387, 2015. PMID: 25992594.
- [261] Wolfgang Kirmse. *Carbene chemistry*, volume 1. Elsevier, 2013.
- [262] Tullimilli Y Gopalakrishna, Wangdong Zeng, Xuefeng Lu, and Jishan Wu. From open-shell singlet diradicaloids to polyradicaloids. *Chem. Commun.*, 54(18):2186–2199, 2018.
- [263] Michael Bendikov, Hieu M Duong, Kyle Starkey, KN Houk, Emily A Carter, and Fred Wudl. Oligoacenes: theoretical prediction of open-shell singlet diradical ground states. *J. Am. Chem. Soc.*, 126(24):7416–7417, 2004.
- [264] Robert R Squires and Christopher J Cramer. Electronic interactions in aryne biradicals. ab initio calculations of the structures, thermochemical properties, and singlet- triplet splittings of the didehydronaphthalenes. *J. Phys. Chem. A*, 102(45):9072–9081, 1998.
- [265] Daniel H Ess, Erin R Johnson, Xiangqian Hu, and Weitao Yang. Singlet- triplet energy gaps for diradicals from fractional-spin density-functional theory. *J. Phys. Chem. A*, 115(1):76–83, 2010.
- [266] Degao Peng, Xiangqian Hu, Deepa Devarajan, Daniel H. Ess, Erin R. Johnson, and Weitao Yang. Variational fractional-spin density-functional theory for diradicals. *J. Chem. Phys.*, 137(11):114112, 2012.
- [267] Emily A Carter and William A Goddard III. Electron correlation, basis sets, and the methylene singlet–triplet gap. *J. Chem. Phys.*, 86(2):862–865, 1987.
- [268] Emily A Carter and William A Goddard. New predictions for singlet-triplet gaps of substituted carbenes. *J. Phys. Chem.*, 91(18):4651–4652, 1987.
- [269] Emily A Carter and William A Goddard III. Correlation-consistent singlet–triplet gaps in substituted carbenes. *J. Chem. Phys.*, 88(3):1752–1763, 1988.

- [270] Karl K Irikura, WA Goddard III, and JL Beauchamp. Singlet-triplet gaps in substituted carbenes cxy (x, y= h, fluoro, chloro, bromo, iodo, silyl). *J. Am. Chem. Soc.*, 114(1):48–51, 1992.
- [271] Yihan Shao, Martin Head-Gordon, and Anna I. Krylov. The spinflip approach within time-dependent density functional theory: Theory and applications to diradicals. *J. Chem. Phys.*, 118(11):4807–4818, 2003.
- [272] Anna I. Krylov. Spin-flip equation-of-motion coupled-cluster electronic structure method for a description of excited states, bond breaking, diradicals, and triradicals. *Acc. Chem. Res.*, 39(2):83–91, 2006. PMID: 16489727.
- [273] Yves A Bernard, Yihan Shao, and Anna I Krylov. General formulation of spin-flip time-dependent density functional theory using non-collinear kernels: Theory, implementation, and benchmarks. *J. Chem. Phys.*, 136(20):204103, 2012.
- [274] Samuel J Stoneburner, Jun Shen, Adeayo O Ajala, Piotr Piecuch, Donald G Truhlar, and Laura Gagliardi. Systematic design of active spaces for multi-reference calculations of singlet–triplet gaps of organic diradicals, with benchmarks against doubly electron-attached coupled-cluster data. *J. Chem. Phys.*, 147(16):164120, 2017.
- [275] Junwei Lucas Bao, Andrew Sand, Laura Gagliardi, and Donald G Truhlar. Correlated-participating-orbitals pair-density functional method and application to multiplet energy splittings of main-group divalent radicals. *J. Chem. Theory Comput.*, 12(9):4274–4283, 2016.
- [276] Liam Wilbraham, Pragya Verma, Donald G Truhlar, Laura Gagliardi, and Ilaria Ciofini. Multiconfiguration pair-density functional theory predicts spin-state ordering in iron complexes with the same accuracy as complete active space second-order perturbation theory at a significantly reduced computational cost. *J. Phys. Chem. Lett.*, 8(9):2026–2030, 2017.
- [277] Soumen Ghosh, Christopher J. Cramer, Donald G. Truhlar, and Laura Gagliardi. Generalized-active-space pair-density functional theory: an efficient method to study large, strongly correlated, conjugated systems. *Chem. Sci.*, 8:2741–2750, 2017.

- [278] Samuel J Stoneburner, Donald G Truhlar, and Laura Gagliardi. Mc-pdf can calculate singlet–triplet splittings of organic diradicals. *J. Chem. Phys.*, 148(6):064108, 2018.
- [279] Prachi Sharma, Varinia Bernales, Stefan Knecht, Donald G. Truhlar, and Laura Gagliardi. Density matrix renormalization group pair-density functional theory (dmrg-pdf): singlet–triplet gaps in polyacenes and polyacetylenes. *Chem. Sci.*, 10:1716–1723, 2019.
- [280] B. Hajgat, D. Szieberth, P. Geerlings, F. De Proft, and M. S. Deleuze. A benchmark theoretical study of the electronic ground state and of the singlet-triplet split of benzene and linear acenes. *J. Chem. Phys.*, 131(22):224321, 2009.
- [281] Balzs Hajgat, Matija Huzak, and Michael S. Deleuze. Focal point analysis of the singlettriplet energy gap of octacene and larger acenes. *J. Phys. Chem. A*, 115(33):9282–9293, 2011. PMID: 21755960.
- [282] Marta Włoch, Jeffrey R Gour, and Piotr Piecuch. Extension of the renormalized coupled-cluster methods exploiting left eigenstates of the similarity-transformed hamiltonian to open-shell systems: A benchmark study. *J. Phys. Chem. A*, 111(44):11359–11382, 2007.
- [283] Jun Shen and Piotr Piecuch. Doubly electron-attached and doubly ionized equation-of-motion coupled-cluster methods with 4-particle–2-hole and 4-hole–2-particle excitations and their active-space extensions. *J. Chem. Phys.*, 138(19):194102, 2013.
- [284] Adeayo O Ajala, Jun Shen, and Piotr Piecuch. Economical doubly electron-attached equation-of-motion coupled-cluster methods with an active-space treatment of three-particle–one-hole and four-particle–two-hole excitations. *J. Phys. Chem. A*, 121(18):3469–3485, 2017.
- [285] Takashi Tsuchimochi and Seiichiro Ten-no. Bridging single-and multireference domains for electron correlation: spin-extended coupled electron pair approximation. *J. Chem. Theory Comput.*, 13(4):1667–1681, 2017.
- [286] Paul M Zimmerman. Singlet–triplet gaps through incremental full configuration interaction. *J. Phys. Chem. A*, 121(24):4712–4720, 2017.

- [287] VM Garcia, O Castell, M Reguero, and R Caballol. Singlet-triplet energy gap in halogen-substituted carbenes and silylenes: a difference-dedicated configuration interaction calculation. *Mol. Phys.*, 87(6):1395–1404, 1996.
- [288] J Cabrero, N Ben-Amor, and R Caballol. Singlet- triplet gap in α -n-dehydrotoluene and related biradicals: An ab initio configuration interaction study. *J. Phys. Chem. A*, 103(31):6220–6224, 1999.
- [289] Carmen J Calzado, Celestino Angeli, Coen de Graaf, and Rosa Caballol. Analysis of the magnetic coupling in nitroxide organic biradicals. *Theor. Chem. Acc.*, 128(4-6):505–519, 2011.
- [290] Yang Yang, Ernest R. Davidson, and Weitao Yang. Nature of ground and electronic excited states of higher acenes. *Proc. Natl. Acad. Sci.*, 113(35):E5098–E5107, 2016.
- [291] Johannes Hachmann, Jonathan J Dorando, Michael Avilés, and Garnet Kin-Lic Chan. The radical character of the acenes: A density matrix renormalization group study. *J. Chem. Phys.*, 127(13):134309, 2007.
- [292] K. Yamaguchi, F. Jensen, A. Dorigo, and K.N. Houk. A spin correction procedure for unrestricted hartree-fock and mller-plesset wavefunctions for singlet diradicals and polyradicals. *Chem. Phys. Lett.*, 149(5):537 – 542, 1988.
- [293] Joonho Lee and Martin Head-Gordon. Two single-reference approaches to singlet biradicaloid problems: Complex, restricted orbitals and approximate spin-projection combined with regularized orbital-optimized møller-plesset perturbation theory. *arXiv preprint arXiv:1903.11225*, 2019.
- [294] Fengjie Ma, Shiwei Zhang, and Henry Krakauer. Excited state calculations in solids by auxiliary-field quantum monte carlo. *New J. Phys.*, 15(9):093017, 2013.
- [295] Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. Excited state calculations using

- phaseless auxiliary-field quantum monte carlo: Potential energy curves of low-lying c_2 singlet states. *J. Chem. Phys.*, 130(9):094107, 2009.
- [296] Hongxia Hao, James Shee, Shiv Upadhyay, Can Ataca, Kenneth D Jordan, and Brenda Rubenstein. Accurate predictions of electron binding energies of dipole-bound anions via quantum monte carlo methods. *J. Phys. Chem. Lett.*, 2018.
- [297] Dominika Zgid and Marcel Nooijen. The density matrix renormalization group self-consistent field method: Orbital optimization with the density matrix renormalization group method in the active space. *J. Chem. Phys.*, 128(14):144116, 2008.
- [298] Edward G. Hohenstein, Nathan Luehr, Ivan S. Ufimtsev, and Todd J. Martinez. An atomic orbital-based formulation of the complete active space self-consistent field method on graphical processing units. *J. Chem. Phys.*, 142(22):224103, 2015.
- [299] James E. T. Smith, Bastien Mussard, Adam A. Holmes, and Sandeep Sharma. Cheap and near exact casscf with large active spaces. *J. Chem. Theory Comput.*, 13(11):5468–5478, 2017. PMID: 28968097.
- [300] Konstantinos D. Vogiatzis, Dongxia Ma, Jeppe Olsen, Laura Gagliardi, and Wibe A. de Jong. Pushing configuration-interaction to the limit: Towards massively parallel mscf calculations. *J. Chem. Phys.*, 147(18):184111, 2017.
- [301] Miquel Sol. Forty years of clar’s aromatic π -sextet rule. *Front. Chem.*, 1:22, 2013.
- [302] James Shee, Benjamin Rudsteyn, Evan J Arthur, Shiwei Zhang, David R Reichman, and Richard A Friesner. On achieving high accuracy in quantum chemical calculations of 3d transition metal-containing systems: A comparison of auxiliary-field quantum monte carlo with coupled cluster, density functional theory, and experiment for diatomic molecules. *J. Chem. Theory Comput.*, 2019.
- [303] Jeffrey B. Schriber, Kevin P. Hannon, Chenyang Li, and Francesco A. Evangelista. A combined selected configuration interaction and many-body treatment of static and dynami-

- cal correlation in oligoacenes. *J. Chem. Theory Comput.*, 14(12):6295–6305, 2018. PMID: 30388005.
- [304] Jon Baker, Andrew Scheiner, and Jan Andzelm. Spin contamination in density functional theory. *Chem. Phys. Lett.*, 216(3-6):380–388, 1993.
- [305] Eloy Ramos-Cordoba and Pedro Salvador. Diradical character from the local spin analysis. *Phys. Chem. Chem. Phys.*, 16(20):9565–9571, 2014.
- [306] Bun Chan and Leo Radom. W2x and w3x-l: Cost-effective approximations to w2 and w4 with kj mol⁻¹ accuracy. *J. Chem. Theory Comput.*, 11(5):2109–2119, 2015.
- [307] ST Gibson, JP Greene, and J Berkowitz. Photoionization of the amidogen radical. *J. Chem. Phys.*, 83(9):4319–4328, 1985.
- [308] Jeffrey C. Stephens, Yukio Yamaguchi, C. David Sherrill, and Henry F. Schaefer. \tilde{X}^3b_1 , \tilde{a}^1a_1 , \tilde{b}^1b_1 , and $\tilde{c}^1\sigma_g^+$ electronic states of nh_2^+ . *J. Phys. Chem. A*, 102(22):3999–4006, 1998.
- [309] Fan Wang and Tom Ziegler. The performance of time-dependent density functional theory based on a noncollinear exchange-correlation potential in the calculations of excitation energies. *J. Chem. Phys.*, 122(7):074109, 2005.
- [310] Wolfram Sander. m-benzyne and p-benzyne. *Acc. Chem. Res*, 32(8):669–676, 1999.
- [311] T Daniel Crawford, Elfi Kraka, John F Stanton, and Dieter Cremer. Problematic p-benzyne: Orbital instabilities, biradical character, and broken symmetry. *J. Chem. Phys.*, 114(24):10638–10650, 2001.
- [312] Paul G Wenthold, Robert R Squires, and WC Lineberger. Ultraviolet photoelectron spectroscopy of the o-, m-, and p-benzyne negative ions. electron affinities and singlet- triplet splittings for o-, m-, and p-benzyne. *J. Am. Chem. Soc.*, 120(21):5279–5290, 1998.
- [313] Prashant U Manohar and Anna I Krylov. A noniterative perturbative triples correction for the spin-flipping and spin-conserving equation-of-motion coupled-cluster methods with single and double substitutions. *J. Chem. Phys.*, 129(19):194105, 2008.

- [314] Collins U Ibeji and Debashree Ghosh. Singlet–triplet gaps in polyacenes: a delicate balance between dynamic and static correlations investigated by spin–flip methods. *Phys. Chem. Chem. Phys.*, 17(15):9849–9856, 2015.
- [315] Xiuhui Zhang, Qian-Shu Li, Yaoming Xie, and Henry F Schaefer III. The lowest triplet electronic states of polyacenes and perfluoropolyacenes. *Mol. Phys.*, 105(19-22):2743–2752, 2007.
- [316] Christoph Riplinger, Peter Pinski, Ute Becker, Edward F Valeev, and Frank Neese. Sparse maps? a systematic infrastructure for reduced-scaling electronic structure methods. ii. linear scaling domain based pair natural orbital coupled cluster theory. *The Journal of chemical physics*, 144(2):024109, 2016.
- [317] Yang Guo, Christoph Riplinger, Ute Becker, Dimitrios G Liakos, Yury Minenkov, Luigi Cavallo, and Frank Neese. Communication: An improved linear scaling perturbative triples correction for the domain based local pair-natural orbital based singles and doubles coupled cluster method [dlpno-ccsd (t)]. *J. Chem. Phys.*, 148(1):011101, 2018.
- [318] Tanya N Singh-Rachford, Radiy R Islangulov, and Felix N Castellano. Photochemical upconversion approach to broad-band visible light generation. *The Journal of Physical Chemistry A*, 112(17):3906–3910, 2008.
- [319] Frank Neese. The orca program system. *Wiley Interdiscip. Rev. Comput. Mol. Sci*, 2(1):73–78, 2012.
- [320] AP Darmanyan. Experimental study of singlet-triplet energy transfer in liquid solutions. *Chem. Phys. Lett.*, 110(1):89–94, 1984.
- [321] Eric H Knoll and Richard A Friesner. Localized orbital corrections for the calculation of ionization potentials and electron affinities in density functional theory. *J. Phys. Chem. B*, 110(38):18787–18802, 2006.

- [322] David Rinaldo, Li Tian, Jeremy N Harvey, and Richard A Friesner. Density functional localized orbital corrections for transition metals. *J. Chem. Phys.*, 129(16):164108, 2008.
- [323] Thomas F Hughes and Richard A Friesner. Correcting systematic errors in dft spin-splitting energetics for transition metal complexes. *J. Chem. Theory Comput.*, 7(1):19–32, 2010.
- [324] Peter Pinski, Christoph Riplinger, Edward F Valeev, and Frank Neese. Sparse maps? a systematic infrastructure for reduced-scaling electronic structure methods. i. an efficient and simple linear scaling local mp2 method that uses an intermediate basis of pair natural orbitals. *J. Chem. Phys.*, 143(3):034108, 2015.
- [325] Masaaki Saitow, Ute Becker, Christoph Riplinger, Edward F Valeev, and Frank Neese. A new near-linear scaling, efficient and accurate, open-shell domain-based local pair natural orbital coupled cluster singles and doubles theory. *J. Chem. Phys.*, 146(16):164105, 2017.
- [326] Yang Guo, Kantharuban Sivalingham, Edward F Valeev, and Frank Neese. Sparsemaps? a systematic infrastructure for reduced-scaling electronic structure methods. iii. linear-scaling multireference domain-based pair natural orbital n-electron valence perturbation theory. *J. Chem. Phys.*, 144(9):094111, 2016.
- [327] Takeshi Yanai and Toru Shiozaki. Canonical transcorrelated theory with projected slater-type geminals. *J. Chem. Phys.*, 136(8):084107, 2012.
- [328] Sandeep Sharma, Takeshi Yanai, George H Booth, C J Umrigar, and Garnet Kin-Lic Chan. Spectroscopic accuracy directly from quantum chemistry: Application to ground and excited states of beryllium dimer. *J. Chem. Phys.*, 140(10):104112, 2014.
- [329] Gerald Knizia, Thomas B Adler, and Hans-Joachim Werner. Simplified ccSD (t)-f12 methods: Theory and benchmarks. *J. Chem. Phys.*, 130(5):054104, 2009.
- [330] Dominik Göddeke, Robert Strzodka, and Stefan Turek. Performance and accuracy of hardware-oriented native-, emulated-and mixed-precision solvers in fem simulations. *Int. J. Parallel, Emergent Distrib. Syst.*, 22(4):221–256, 2007.

- [331] Ettore Vitali, Hao Shi, Mingpu Qin, and Shiwei Zhang. Computation of dynamical correlation functions for many-fermion systems with auxiliary-field quantum monte carlo. *Phys. Rev. B*, 94(8):085140, 2016.
- [332] J Orenstein and AJ Millis. Advances in the physics of high-temperature superconductivity. *Science*, 288(5465):468–474, 2000.
- [333] Patrick A Lee, Naoto Nagaosa, and Xiao-Gang Wen. Doping a mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.*, 78(1):17, 2006.
- [334] Brian Busemeyer, Gregory J MacDougall, and Lucas K Wagner. Prediction for the singlet-triplet excitation energy for the spinel mgti 2 o 4 using first-principles diffusion monte carlo. *Phys. Rev. B*, 99(8):081118, 2019.