



HAL
open science

Improved Modularity and New Features in ipie: Towards Even Larger AFQMC Calculations on CPUs and GPUs at Zero and Finite Temperatures

Tong Jiang, Moritz Baumgarten, Pierre-François Loos, Ankit Mahajan,
Anthony Scemama, Shu Fay Ung, Jinghong Zhang, Fionn Malone, Joonho Lee

► **To cite this version:**

Tong Jiang, Moritz Baumgarten, Pierre-François Loos, Ankit Mahajan, Anthony Scemama, et al.. Improved Modularity and New Features in ipie: Towards Even Larger AFQMC Calculations on CPUs and GPUs at Zero and Finite Temperatures. *The Journal of Chemical Physics*, 2024, 161 (16), 10.1063/5.0225596 . hal-04624474

HAL Id: hal-04624474

<https://hal.science/hal-04624474v1>

Submitted on 25 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improved Modularity and New Features in `ipie`: Towards Even Larger AFQMC Calculations on CPUs and GPUs at Zero and Finite Temperatures

Tong Jiang,¹ Moritz K. A. Baumgarten,¹ Pierre-François Loos,² Ankit Mahajan,³ Anthony Scemama,² Shu Fay Ung,³ Jinghong Zhang,¹ Fionn D Malone,⁴ and Joonho Lee^{1, a)}

¹*Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA, USA*

²*Laboratoire de Chimie et Physique Quantiques (UMR 5626), Université de Toulouse, CNRS, UPS, France*

³*Department of Chemistry, Columbia University, New York, NY, USA*

⁴*Google Research, Venice, CA 90291, United States*

`ipie` is a Python-based auxiliary-field quantum Monte Carlo (AFQMC) package that has undergone substantial improvements since its initial release [*J. Chem. Theory Comput.*, 2022, 19(1): 109-121]. This paper outlines the improved modularity and new capabilities implemented in `ipie`. We highlight the ease of incorporating different trial and walker types and the seamless integration of `ipie` with external libraries. We enable distributed Hamiltonian simulations, allowing for multi-GPU simulations of large systems. This development enabled us to compute the interaction energy of a benzene dimer with 84 electrons and 1512 orbitals, which otherwise would not have fit on a single GPU. We also support GPU-accelerated multi-slater determinant trial wavefunctions [arXiv:2406.08314] to enable efficient and highly accurate simulations of large-scale systems. This allows for near-exact ground state energies of multi-reference clusters, $[\text{Cu}_2\text{O}_2]^{2+}$ and $[\text{Fe}_2\text{S}_2(\text{SCH}_3)]^{2-}$. We also describe implementations of free projection AFQMC, finite temperature AFQMC, AFQMC for electron-phonon systems, and automatic differentiation in AFQMC for calculating physical properties. These advancements position `ipie` as a leading platform for AFQMC research in quantum chemistry, facilitating more complex and ambitious computational method development and their applications.

I. INTRODUCTION

Auxiliary-field quantum Monte Carlo (AFQMC)^{1,2} has become increasingly popular in quantum chemistry^{3,4} and is furthermore being recognized as a useful approach in the context of quantum algorithms.⁵⁻⁹ A well-maintained AFQMC program with flexibility and robust performance will play a pivotal role at the intersection of many disciplines, including chemistry, physics, materials science, and quantum information science.

Given the initial PAUXY development effort for rapid prototyping, while lacking performance,¹⁰⁻¹⁴ the Python-based AFQMC program `ipie` was designed from scratch for high performance and ease of development. It was officially introduced as a production-level package in Ref. 15. `ipie` was optimized for high-performance computing architectures with both central and graphical processing units (CPUs and GPUs). High performance was largely achieved by integrating Numba's JIT compilation¹⁶ to fine-tune the computational efficiency of specific kernels and MPI parallelism for effective distributed computing. The utility of `ipie` was showcased through the resolution of intricate quantum chemical challenges, notably the $[\text{Cu}_2\text{O}_2]^{2+}$ torture track.^{15,17} Rigorous benchmarks on CPU and GPU platforms position `ipie` competitively, displaying speed on par with—or surpassing—existing Python and C++ codes.^{15,18,19} Since its release, `ipie` has gained widespread adoption in both quantum chemistry and quantum computing studies.^{3,7,8,20-23}

In this article, we describe the recent development and current status of `ipie`, introducing enhanced modularity,

a suite of new features, interfaces to external packages, and associated numerical examples. Below are the key highlights:

a. High degree of modularity and customizability. The AFQMC driver has been restructured to be fully modular, allowing for a straightforward combination of features. This provides greater flexibility in adapting to a wide range of user demands. Advancements in AFQMC algorithms often focus on developing new trial wavefunctions to control better the fermionic sign/phase problem.^{5,23-26} The key routines in an AFQMC calculation compute intermediates using walker and trial wavefunctions, including overlap, force bias, Green's function, and energy estimators. Without altering the internal core code of `ipie`, users and developers can customize all components making up an AFQMC simulation, including trial wavefunctions, walkers, Hamiltonians, propagators, and estimators. The key objects are all structured using object-oriented programming (OOP) principles, facilitating straightforward customization through inheritance.

b. Development-friendly design. This improvement standardizes the component interfaces and workflow processes while offering a flexible system that adapts to various data types and user requirements. The abstract base classes serve as a foundational blueprint, ensuring all components adhere to a uniform structure and interact seamlessly. Complementing this, the factory methods for common workflows simplify the instantiation process, allowing users to set up standard calculations with minimal effort and reduced potential for errors. Integrating type-based dispatch through Plum²⁷ brings increased precision and efficiency in method handling, ensuring every component downstream dynamically adjusts its operations based on the specific trials and walkers requested.

^{a)}Electronic mail: joonholee@chem.harvard.edu

Moreover, we have simplified the integration of `ipie` with external quantum chemistry packages. While the necessary integrals and orbitals for running AFQMC are most commonly obtained through an interface with PySCF,²⁸ our simplified file format also ensures that other packages can be easily used. For sophisticated trial wavefunctions, such as multiple Slater determinant (MSD) trials derived from selected configuration interaction, interfaces with PySCF,²⁸ Dice,^{19,29} and TREXIO³⁰ are available. An additional interface with the Fermionic Quantum Emulator (FQE)³¹ has also been introduced, facilitating the conversion between `ipie`'s MSD wavefunction and quantum circuit wavefunctions, which facilitates AFQMC applications in the quantum information science (QIS) community.

c. New improvements and features. Several features have been added to `ipie` to improve memory management, handle new problems, enhance integration testing, *etc.* To address the significant memory demand associated with storing Cholesky vectors for large systems, `ipie` offers options for shared memory across CPU processes or distributing among CPU processes or GPU cards. Support for complex-valued Hamiltonians has also been added.

While `ipie` is primarily developed for phaseless AFQMC (ph-AFQMC) targeting *ab initio* systems, it supports other AFQMC methods as well, including free-projection AFQMC,²⁴ finite temperature AFQMC,¹² AFQMC for coupled electron-phonon systems,¹³ and automatic differentiation within AFQMC for calculating observables which do not commute with the Hamiltonian.³² Furthermore, `ipie` is equipped with improved integration testing, boosting the package's robustness and adaptability and significantly enhancing its reliability and usability for end-users and developers alike.

The organization of this paper is as follows: Section II overviews the theory of AFQMC; Section III details the components, software architecture, and workflow of `ipie`, including examples to illustrate the framework's adaptability for AFQMC development; Section IV introduces new features in `ipie` and provides corresponding examples; Section V outlines the interfaces to external packages; and Section VI concludes with a summary and outlook.

II. THEORY OF AFQMC

AFQMC is based on the following imaginary time evolution:

$$|\Psi_0\rangle \propto \lim_{\tau \rightarrow \infty} \exp(-\tau \hat{H}) |\Phi_0\rangle = \lim_{n \rightarrow \infty} (\exp(-\Delta\tau \hat{H}))^n |\Phi_0\rangle, \quad (1)$$

where $\Delta\tau$ is an infinitesimal time step, $|\Psi_0\rangle$ is the ground state wavefunction, and $|\Phi_0\rangle$ is an initial state satisfying $\langle \Phi_0 | \Psi_0 \rangle \neq 0$. While `ipie` supports some of the prototypical model Hamiltonians, its development has focused on the simulation of the *ab initio* Hamiltonian, which in

second quantization is given by

$$\hat{H} = \sum_{p,q=1}^N h_{pq} \hat{a}_p^\dagger \hat{a}_q + \frac{1}{2} \sum_{p,q,r,s=1}^N g_{psqr} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s, \quad (2)$$

where the two-electron repulsion integral (ERI) is factorized with the Cholesky decomposition

$$g_{psqr} = (ps|qr) = \sum_{\gamma=1}^{N_\gamma} L_{ps}^\gamma L_{qr}^\gamma. \quad (3)$$

With this factorization, we have

$$\hat{H} = \hat{v}_0 - \frac{1}{2} \sum_{\gamma=1}^{N_\gamma} \hat{v}_\gamma^2, \quad (4)$$

where

$$\hat{v}_0 = \sum_{pq} \left[h_{pq} - \frac{1}{2} \sum_r (pr|rq) \right] a_p^\dagger a_q \quad (5)$$

$$\hat{v}_\gamma = i \sum_{pq} L_{pq}^\gamma \hat{a}_p^\dagger \hat{a}_q. \quad (6)$$

The short-time propagator with Trotter decomposition is written as

$$e^{-\Delta\tau \hat{H}} = e^{-\frac{\Delta\tau}{2} \hat{v}_0} e^{\frac{\Delta\tau}{2} \sum \hat{v}_\gamma^2} e^{-\frac{\Delta\tau}{2} \hat{v}_0} + \mathcal{O}(\Delta\tau^3). \quad (7)$$

Upon applying the Hubbard–Stratonovich transformation,^{33,34} our effective propagator contains only one-body operators,

$$e^{-\Delta\tau \hat{H}} = \int d\mathbf{x} p(\mathbf{x}) \hat{B}(\mathbf{x}, \Delta\tau) + \mathcal{O}(\Delta\tau^2), \quad (8)$$

where $p(\mathbf{x})$ is the standard Gaussian distribution, $\mathbf{x} = (x_1, x_2, \dots, x_{N_\gamma})$ are the auxiliary fields, and the one-body propagator \hat{B} is

$$\hat{B}(\mathbf{x}, \Delta\tau) = e^{-\frac{\Delta\tau}{2} \hat{v}_0} e^{-\sqrt{\Delta\tau} \mathbf{x} \cdot \hat{\mathbf{v}}} e^{-\frac{\Delta\tau}{2} \hat{v}_0}. \quad (9)$$

where $\hat{\mathbf{v}} = (\hat{v}_1, \hat{v}_2, \dots, \hat{v}_{N_\gamma})$.

In AFQMC, each walker samples auxiliary fields and represents a statistical sample of the global wavefunction at imaginary time τ , written as

$$|\Psi(\tau)\rangle = \sum_i^{N_w} w_i(\tau) \frac{|\psi_i(\tau)\rangle}{\langle \Psi_T | \psi_i(\tau) \rangle}, \quad (10)$$

where $|\psi_i(\tau)\rangle$ is the wavefunction of the i -th walker at time τ and $|\Psi_T\rangle$ is the trial wavefunction used for importance sampling. The energy estimator is then

$$E(\tau) = \sum_i^{N_w} w_i(\tau) E_{\text{loc},i}(\tau) = \sum_i^{N_w} w_i(\tau) \frac{\langle \Psi_T | \hat{H} | \psi_i(\tau) \rangle}{\langle \Psi_T | \psi_i(\tau) \rangle}. \quad (11)$$

The walker state $|\psi_i(\tau)\rangle$ is updated by applying the discrete-time propagator, and the weight $w_i(\tau)$ is updated following the phaseless approximation:²

$$|\psi_i(\tau + \Delta\tau)\rangle = \hat{B}(\mathbf{x}_i - \bar{\mathbf{x}}_i, \Delta\tau) |\psi_i(\tau)\rangle \quad (12)$$

$$w_i(\tau + \Delta\tau) = I_{\text{ph}}(\mathbf{x}_i, \bar{\mathbf{x}}_i, \tau, \Delta\tau) \times w_i(\tau). \quad (13)$$

We dynamically shift the distribution of auxiliary fields using the force bias $\bar{\mathbf{x}}_i$ defined by

$$\bar{\mathbf{x}}_i(\Delta\tau, \tau) = -\sqrt{\Delta\tau} \frac{\langle \Psi_T | \hat{\mathbf{v}} - \langle \hat{\mathbf{v}} \rangle_T | \psi_i(\tau) \rangle}{\langle \Psi_T | \psi_i(\tau) \rangle}, \quad (14)$$

where $\langle \hat{\mathbf{v}} \rangle_T$ is usually called the mean-field shift in the AFQMC literature.⁴ Furthermore, the phaseless importance function^{1,2} used in the weight update (13) is given by

$$I_{\text{ph}}(\mathbf{x}_i, \bar{\mathbf{x}}_i, \tau, \Delta\tau) = |I(\mathbf{x}_i, \bar{\mathbf{x}}_i, \tau, \Delta\tau)| \times \max[0, \cos(\theta_i(\tau))], \quad (15)$$

which remains real and positive throughout the propagation. The so-called hybrid importance function is given by

$$I(\mathbf{x}_i, \bar{\mathbf{x}}_i, \tau, \Delta\tau) = S_i(\tau, \Delta\tau) e^{\mathbf{x}_i \cdot \bar{\mathbf{x}}_i - \bar{\mathbf{x}}_i \cdot \bar{\mathbf{x}}_i / 2}, \quad (16)$$

and the overlap ratio of the i -th walker is

$$S_i(\tau, \Delta\tau) = \frac{\langle \Psi_T | \hat{B}(\Delta\tau, \mathbf{x}_i - \bar{\mathbf{x}}_i) | \psi_i(\tau) \rangle}{\langle \Psi_T | \psi_i(\tau) \rangle}. \quad (17)$$

We define the phase of the overlap as

$$\theta_i(\tau) = \arg[S_i(\tau, \Delta\tau)]. \quad (18)$$

The computation of local energies in Eq. (11), propagation of wavefunctions in Eq. (12), and evaluation of force biases in Eq. (14) are the primary computational hotspots in AFQMC calculations. We optimize these routines by employing optimized algorithms and adopting high-performance computing techniques. We provide an overview of the phaseless AFQMC algorithm described above in Fig. 1.

III. SOFTWARE ARCHITECTURE AND DESIGN PRINCIPLES

A. AFQMC driver

We overview the software architecture of `ipie` with emphasis on improved modularity, as shown in Fig. 2. `ipie` generates an AFQMC simulation by assembling the generic AFQMC driver from problem-specific components and then running it to execute the simulation. The driver requires several inputs to define the QMC simulation, including the trial wavefunction, walkers, and the Hamiltonian, as well as more generic parameters such as the number of blocks and time step. The following

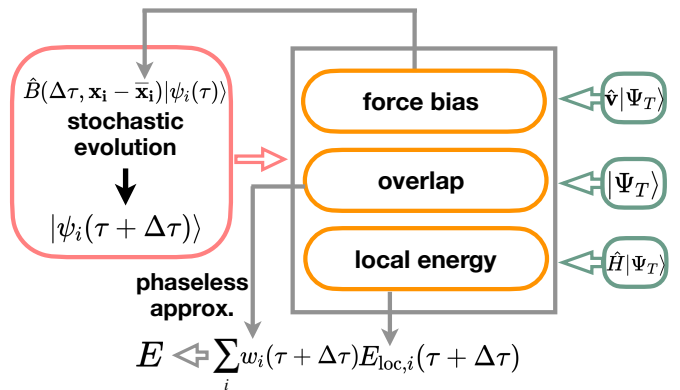


FIG. 1. Overview of phaseless AFQMC in terms of computational hotspots.

Python Code Snippet shows how the driver can be instantiated:

Code Snippet 1 The AFQMC class within `ipie`.

```
class AFQMC(object):
    def __init__(self, system, hamiltonian, trial,
                 walkers, propagator, mpi_handler, params:
                 QMCPParams):
        .....
```

The code takes several objects needed for a typical AFQMC driver:

- (1) `system` contains information that defines the problem, including the number of spin-up and spin-down electrons.
- (2) `hamiltonian` stores molecular integrals as detailed in Sec. III B.
- (3) `trial` represents a trial wavefunction object, for which `ipie` provides many options as detailed in Sec. III C.
- (4) `walkers` manages all information about walkers during the imaginary time propagation. This includes the initial walker wavefunctions and the number of walkers, among other details. `ipie` allows walkers to be explicitly passed into the AFQMC driver or dynamically dispatched with default settings, as detailed in Sec. III D.
- (5) `propagator` handles the update of walker weights and wavefunctions during the imaginary time evolution. `ipie`'s AFQMC driver can take the propagator class as an input or dispatch it internally with default settings.
- (6) `mpi_handler` manages information related to the Message Passing Interface (MPI) for parallel computing, including rank details, chunking, and




FIG. 2. The workflow of `ipie`.

groups (i.e., a group is a collection of MPI processes). The parameter `shared_comm` in `mpi_handler` refers to the MPI rank within a group where the chunking integrals are distributed. The specifics of these are elaborated in Section IV A.

- (7) `params` encompasses other fundamental parameters for the QMC simulation, such as the time step.

While users can directly provide all of the aforementioned inputs to construct the AFQMC driver object, especially for development purposes, a factory method is provided to simplify the process greatly. This is especially useful for standard AFQMC calculations using single or multiple Slater determinant trials. This method significantly reduces the number of inputs required to construct the AFQMC object, as demonstrated in Code Snippet 2:

Code Snippet 2 Factory method for AFQMC driver.

```
def build(
    num_elec: Tuple[int, int],
    hamiltonian,
    trial_wavefunction
) -> "AFQMC":
    .....
```

As illustrated in the Code Snippet 2 and Fig. 2, the construction of the AFQMC driver ultimately relies on two key inputs: the Hamiltonian and the trial wavefunction. The Hamiltonian object is assembled with one-electron integrals and the Cholesky decomposition (or density fitting) of two-electron integrals, detailed in Sec. III B. As for trial wavefunctions, `ipie` offers various options that utilize interfaces with external packages such as `PySCF`,²⁸ `Dice`,¹⁹ `TREXIO`³⁰ to facilitate standardized workflows. `ipie` can also accommodate user-

customized trials, which offers flexibility in future developments. Benefited from the flexibility of `ipie`, we are also able to incorporate more complicated trial wavefunctions, such as density matrix renormalization group (DMRG), and in Sec. III C we show the AFQMC calculation with matrix product states (MPS) trials by incorporating `ipie` with the DMRG package `Renormalizer`.³⁵

The `build` method in the AFQMC class is designed to streamline the setup of an AFQMC calculation by requiring only three inputs: the number of electrons, the Hamiltonian, and the trial wavefunction. Its main purpose is automating the construction of key components such as walkers, propagators, and estimators, *etc.* Users need only specify the `hamiltonian` and trial inputs `trial_wavefunction`, with the `build` method handling downstream instantiation of the rest.

Once the driver is built, wavefunction propagation, stabilization, population control, and estimator calculation are performed with the provided QMC parameters. Simulation outputs in both text file and `hdf5` file formats are saved and updated after each block, which enables real-time reblocking analysis³⁶ via built-in tools.

B. Hamiltonian

The construction of a Hamiltonian object can be achieved in several ways. One can construct the object directly by inputting the one-electron integrals and the Cholesky decomposition (Eq. (3)) of either 8-fold (for real symmetric integrals) or 4-fold (for complex hermitian integrals) symmetric two-body integrals:

Code Snippet 3 Constructing the Hamiltonian object with provided electron integrals

```
from ipie.hamiltonians.generic import Generic as
↳ HamGeneric
ham = HamGeneric(h1e, chol, ecore)
```

where `h1e`, `chol` and `ecore` are the one-electron integrals, the Cholesky decomposition (Eq. (6)) of the two-electron integrals within the desired orbitals in Eq. (2), and the nuclear repulsion energy, respectively.

`ipie` also offers built-in functions for assembling the Hamiltonian object from PySCF calculations. One can provide: (i) a PySCF check file; (ii) a PySCF mol object together with the molecular orbital (MO) coefficient matrix `mo_coeff`; or (iii) a mol object, `mo_coeff`, and a basis transformation matrix `X` that transforms MOs to orthogonal atomic orbitals (OAO), natural orbitals and so on, depending on the single particle basis used in the AFQMC calculation:

Code Snippet 4 Constructing the Hamiltonian object from PySCF.

```
from ipie.utils.from_pyscf import generate_hamiltonian,
↳ generate_hamiltonian_from_chk

# OAO basis
ham = generate_hamiltonian_from_chk(
    'scf.chk', use_mcsf=False,
    chol_cut=1e-5, num_frozen_core=0, ortho_ao=True
)

# MO basis
ham = generate_hamiltonian(mol, mo_coeff, h1e, mo_coeff,
    chol_cut=1e-5, num_frozen_core=0)
# basis with transform matrix X
ham = generate_hamiltonian(mol, mo_coeff, h1e, X,
    chol_cut=1e-5)
```

The argument `chol_cut` specifies the cutoff for the Cholesky decomposition; `use_mcscf` specifies whether to use the multi-configurational self-consistent field (MC-SCF) MO coefficients; and `num_frozen_core` specifies the number of frozen cores for the subsequent AFQMC calculations.

Given that the size of Cholesky vectors scales as $\mathcal{O}(N^3)$, with N as the number of orbitals, large systems with large basis sets can generate electron integrals of considerable size. For such systems, one cannot store copies of integrals for each MPI process and must resort to other strategies. Utilizing shared memory across different MPI processes, `ipie` offers an approach that allows processes on the same node to access Cholesky vectors in the same memory block. The following code snippet can be adopted to ensure the use of shared memory Hamiltonian objects:

Code Snippet 5 Constructing the Hamiltonian object with shared memory.

```
from ipie.utils.mpi import get_shared_comm
from ipie.hamiltonians.utils import get_hamiltonian
from mpi4py import MPI
shared_comm = get_shared_comm(MPI.COMM_WORLD)
ham = get_hamiltonian("ham.h5", comm=shared_comm,
    pack_chol=True)
```

This reads the integrals from the `ham.h5` file and generates the Hamiltonian object with shared integrals across all MPI processes. `pack_chol` refers to another strategy used by `ipie` to reduce memory usage by a factor of two and speed up the propagation. Namely, the permutation symmetry of Cholesky vectors is exploited to utilize only the upper-triangular part of the Cholesky vectors. Propagation using the packed matrix contracts only the symmetric part in Eq. (12) and is thus more efficient.

However, the shared memory strategy often creates a memory access overhead and reduces the MPI parallel efficiency when many cores are used. We also enable a chunking (or distributed-memory) strategy. The Cholesky vectors are divided into segments, or ‘chunked,’ and distributed across different MPI processes to optimize memory usage and computational efficiency. The Cholesky vectors are divided into several chunks, each of which will be allocated to separate MPI processes. The specifics of this chunking process and the associated MPI communication strategy will be thoroughly explained in Sec. IV A. Given enough nodes, our approach completely removes the memory bottleneck in AFQMC and can work on both CPU and GPU nodes. We especially recommend our distributed-memory implementation for GPUs, where memory capacity is much more limited than that of CPUs.

C. Trial wavefunctions

Any new trial wavefunction may inherit from the `TrialWavefunctionBase` class, which contains methods to calculate essential quantities including the overlap (Eq. (17)), force bias (Eq. (14)), *etc.*, as shown in Fig. 3. Some trials, such as single/multiple Slater determinants, also require computation of Green’s function and half rotation of electron repulsion integrals.¹¹ There is flexibility to either inherit these functions directly or overwrite them with customized implementations, ensuring both consistency in fundamental operations and adaptability for specialized needs.

Single Slater determinant (SD) trials can be generated with




FIG. 3. The functions to be defined for a trial object. *: The function is optional depending on the trial type, used in certain trials to computations like overlap, force bias, or local energy.

Code Snippet 6 The single determinant trial object.

```

from ipie.trial_wavefunction.single_det import SingleDet
trial = SingleDet(np.hstack([orbs_a, orbs_b]), nelecs,
                  num_basis)

```

where `orbs_a` and `orbs_b` are the spin-up and spin-down coefficient matrices for the SD trials. Similarly, MSD trials can be generated with:

Code Snippet 7 The multiple Slater determinant trial object.

```

from ipie.trial_wavefunction.particle_hole import
  ParticleHole
trial = ParticleHole(wfn, nelecs, num_basis)

```

with `wfn` containing the occupation indices in the basis used in the configuration interaction (CI) expansion of the MSD and their coefficients.

A new customized trial can be defined straightforwardly in `ipie`. For instance, we consider a “noisy” SD trial that can be inherited from our existing SD trial class. The noisy trial adds Gaussian noise to the overlap evaluation and keeps everything else the same as for regular SD trials. The following code snippet achieves this:

Code Snippet 8 Example of constructing a custom noisy trial class.

```

class NoisySingleDet(SingleDet):
    def __init__(self, wavefunction, num_elec,
                 num_basis, noise_level=1e-12):
        super().__init__(wavefunction, num_elec,
                         num_basis)
        self._noise_level = noise_level
    def calc_overlap(self, walkers) -> np.ndarray:
        ovlp = super().calc_overlap(walkers)
        noise = np.random.normal(
            scale=self._noise_level,
            size=ovlp.size
        )
        return ovlp * (1 + noise)

```

Recently, we have also implemented a more sophisti-




FIG. 4. Application of MPS-AFQMC to two-dimensional hydrogen lattice with $r = 4.2a_0$. Energy deviation from exact results using MPS-AFQMC, UHF-AFQMC, and DMRG with different bond dimensions.²³

cated extension using a matrix product state (MPS) trial wavefunction.²³ MPS is the variational ansatz used in the density matrix renormalization group (DMRG) algorithm, and we obtain the MPS solution from external DMRG packages such as `Renormalizer`.³⁵ The theory explaining the structure of the MPS trial and associated functions is detailed in Ref. 23. A representative calculation is shown in Fig. 4.

D. Dispatchers of walkers, propagators and estimators

As the program grows, additional complexity arises from the need to dispatch different features for arbitrary combinations of smaller building blocks. While there are numerous ways to achieve this, `ipie` previously relied on conditional `if` statements at the computational subroutines, which—despite its simplicity—will result in many intractable conditions as the number of features increases. The latest release of `ipie` thus utilizes the multiple dispatch library `Plum`²⁷ to dispatch features for different combinations of objects in a compact and flexible manner. This can be especially useful in

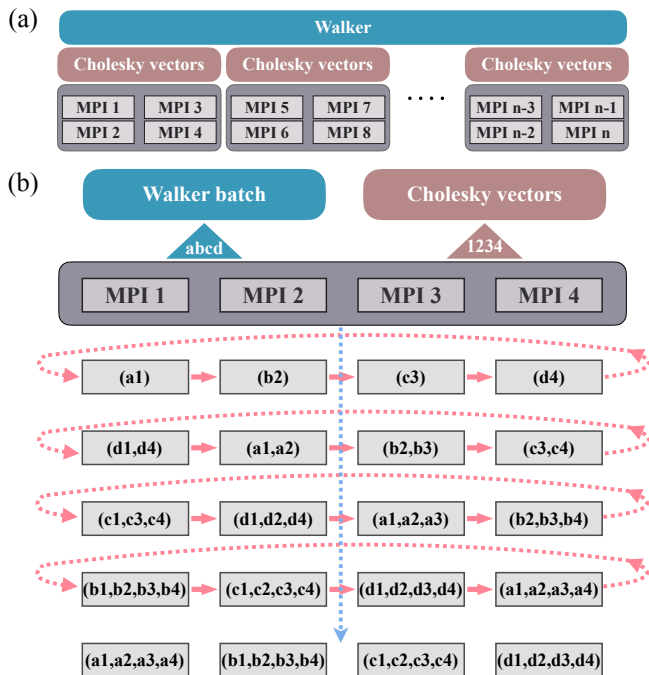


FIG. 6. (a) Distribution of walkers and the chunked Hamiltonian over multiple groups. (b) MPI communication between MPI processes within a group with the distributed Hamiltonian and walker batches.

the tensor is required by the computational subroutines, for example, in computing the force bias or local energy, the walkers in each member of the group are communicated to other members in the same group, as illustrated in Fig. 6(b). We employ a cyclic data passing scheme of walkers among MPI processes to apply all existing Cholesky chunks to given walkers. With this strategy, one can accumulate locally contracted quantities before moving to the next QMC time step.

We performed timing benchmarks as an application by computing the parallel-displaced benzene dimer, a Van der Waals complex taken from the S22 dataset.³⁸ Disk usage for the benzene dimer calculations varies with the choice of basis set: 11GB for the aug-cc-pVTZ basis and 67GB for the aug-cc-pVQZ basis with a Cholesky decomposition threshold of 10^{-3} . The threshold is set based on the Hartree-Fock energy error around typical density fitting errors³⁹ ($< 50\mu E_h$ per atom here). We used our shared memory framework for CPU calculations and distributed Cholesky vectors and half-rotated integrals¹¹ across multiple cards were used for GPU computations. For the benzene dimer analysis employing an aug-cc-pVTZ basis set, a single A100 GPU card achieves performance around tenfold faster than 36 CPU cores (see Fig. 7). Increasing the GPU count for non-chunked Hamiltonian configurations (`nmembers=1`) results in a nearly linear acceleration proportional to the number of GPUs deployed, as shown in Fig. 7(a). Distributing the Hamiltonian across additional GPUs slightly slows




FIG. 7. (a) Timing benchmarks for AFQMC calculations on the benzene dimer with the aug-cc-pVTZ basis employing 36 CPUs with shared memory and multiple A100 GPUs with the distributed Hamiltonian. The data were obtained by averaging 100 blocks, each containing 25 steps with a time step of $0.005 E_h^{-1}$. The interaction energy of the dimer computed from AFQMC is compared against other methods using the aug-cc-pVTZ basis in (b) and in the complete basis set limit obtained from a two-point basis extrapolation⁴⁰ from aug-cc-pVTZ and aug-cc-pVQZ calculations in (c). The MP2, CCSD, and CCSD(T) data using the aug-cc-pVTZ basis in (b) were obtained from Refs. 41 and 42. The CCSD(T) and FN-DMC data in (c) were extracted from Ref. 43. All AFQMC results were obtained with 1224 total walkers.

down the performance, predominantly during the VHS step, which involves the construction of two-body propagators (where HS stands for Hubbard-Stratonovich) during propagation.

Using our GPU implementation, we calculated the interaction energy of the benzene dimer with counterpoise correction.⁴⁴ We compared AFQMC results with those from other methodologies, as shown in Fig. 7(b)(c). While coupled cluster with singles and doubles (CCSD) underestimates the binding energy, second-order Møller-Plesset perturbation theory overestimates. Crucially, we observed that RHF-AFQMC is on par with CCSD with perturbative triples (CCSD(T)).

B. Support of complex Cholesky vectors

Most calculations in quantum chemistry deal with real-valued integrals. However, in some cases, especially in the presence of external magnetic fields, when considering spin-orbit coupling or employing Bloch orbitals in solid-state calculations, the wavefunctions (and orbitals) can become complex-valued.⁴⁵ This adds a layer of mathematical and computational complexity. To handle these cases, `ipie` supports complex Cholesky vectors; the num-

ber of auxiliary fields is hence twice as many as the number of Cholesky vectors. The implementation details are available in Ref. 46.

C. GPU accelerated MSD-AFQMC and timing benchmarks

ph-AFQMC with MSD trials is a powerful tool for systems with multi-reference characteristics, especially when the phaseless bias can be converged away. Building on top of our successful CPU implementation¹⁵ based on Wick’s theorem,³⁷ As detailed in Ref. 47, `ipie` has now enabled GPU-accelerated MSD-AFQMC implementations. Here, we summarize some of the important features of our implementation and discuss applications of this feature. Interested readers are referred to Ref. 47 for more details.

Similar to our single determinant GPU implementation, `ipie` employs `cupy.einsum` via the `cuTENSOR` library to enhance matrix multiplication performance. Additionally, Wick’s algorithm is adapted for GPU execution using custom CUDA kernels, effectively speeding up the computational hotspots in MSD-AFQMC, such as Green’s function, overlap, and local energy computations.

The GPU implementation of MSD-AFQMC is benchmarked on the example $[\text{Cu}_2\text{O}_2]^{2+}$ from `ipie`’s first release paper.¹⁵ In addition to employing a total of 10 walkers to directly compare with our previous results in Ref. 15, calculations using a total of 640 walkers were also performed to obtain a more realistic comparison. With 10 walkers, the GPU implementation on a single NVIDIA A100 card is six times faster than using a single CPU core when the number of determinants in the trial is less than 10^3 . As the number of determinants increases, the performance improvement is even more pronounced: tenfold with 10^4 determinants and approximately 100-fold with 10^6 determinants, as illustrated in Fig. 8(a). With 640 walkers, the GPU code on a single A100 card is compared against with 32 CPU cores. As shown in Fig. 8(b), a fourfold speedup is observed with less than 10^2 determinants that increase to around sixfold with more determinants. The determinants were divided into chunks and computed sequentially to handle a larger number of determinants within the memory limits. This approach introduces an overhead that mostly affects the energy estimator, as reported in Ref. 47.

The timing and absolute energy benchmarks are further benchmarked on the (20o, 30e) active space of the $[\text{Fe}_2\text{S}_2(\text{SCH}_3)]^{2-}$ cluster,⁴⁷ as shown in Fig. 9. In contrast to the 108-orbital calculation of $[\text{Cu}_2\text{O}_2]^{2+}$, a merely two-fold speedup is observed using a single A100 card compared to 32 CPUs. This speedup increases to almost tenfold when the number of determinants exceeds 10^4 . As shown in Fig. 9(b), another observation is using natural orbitals for AFQMC calculation leads to chemical accuracy with around 3×10^5 determinants in the trial.




FIG. 8. (a) Comparisons of the time cost per block for AFQMC calculations on $[\text{Cu}_2\text{O}_2]^{2+}$ using the BS1 basis¹⁵ with 10 walkers on a single CPU core and a single A100 GPU. (b) Time per block with 640 walkers on a single A100 against 32 CPU cores. Figure adapted from Ref. 47.

In contrast, reaching chemical accuracy is more challenging when using localized orbitals employed for DMRG calculations.⁴⁸ These observations highlight the critical need to select an appropriate set of orbitals for AFQMC calculations involving strongly correlated systems. Additionally, it is noteworthy that the energies obtained using localized orbitals with a limited number of determinants are initially lower than the FCI energy, and they tend to converge upward toward the FCI reference as the number of determinants increases.

D. Free projection AFQMC

Free projection (fp-) AFQMC is a numerically exact method for calculating the eigenvalues of a Hamiltonian. While the more commonly used ph-AFQMC variant employs the phaseless constraint to control the sign problem, fp-AFQMC attempts to sample the ground state energy brute-force with exponential-scaling sample complexity. Despite this scaling, it is possible to perform relatively large active space calculations in practice using accurate trial states.

In fp-AFQMC, the ground-state energy is estimated by

$$E(\tau) = \frac{\langle \psi_l | \hat{H} e^{-\tau \hat{H}} | \psi_r \rangle}{\langle \psi_l | e^{-\tau \hat{H}} | \psi_r \rangle} = \frac{\int d\mathbf{x} p(\mathbf{x}) \langle \psi_l | \hat{H} \hat{B}(\mathbf{x}, \Delta\tau) | \psi_r \rangle}{\int d\mathbf{x} p(\mathbf{x}) \langle \psi_l | \hat{B}(\mathbf{x}, \Delta\tau) | \psi_r \rangle}, \quad (19)$$

where $|\psi_l\rangle$ and $|\psi_r\rangle$ are left and right trial states, respectively; $\hat{B}(\mathbf{x}, \Delta\tau)$ is defined in Eq. 9; and τ is the imaginary time. Analogously to the ph-AFQMC propagator sampling, the auxiliary fields \mathbf{x} are sampled from the Gaussian distribution $p(\mathbf{x})$. Unlike ph-AFQMC, we




FIG. 9. Time cost and absolute energy benchmarks on the $[\text{Fe}_2\text{S}_2(\text{SCH}_3)]^{2-}$ cluster. (a) The time cost per block with 640 walkers on 32 CPU cores and a single A100 GPU. (b) Absolute energies are derived using localized atomic orbitals and natural orbitals. Figure adapted from Ref. 47.

do not use the force bias to perform importance sampling. Instead, only mean-field subtraction is used. We also do not employ population control.

A judicious choice of $|\psi_l\rangle$ and $|\psi_r\rangle$ offers two advantages. First, more accurate trial states (*eg.*, selected CI) reduce the imaginary time required to project the ground state energy to a given accuracy, thereby reducing the noise in the energy estimate.^{24,37} Integrating the GPU-accelerated MSD code⁴⁷ will also significantly accelerate the fp-AFQMC calculations with MSD trials. Second, since energies are measured using the state $|\psi_l\rangle$, the closer this state is to the ground state, the smaller the variance in the energy estimate due to the zero variance principle.⁴⁹ Furthermore, one can use a CCSD wavefunction as the initial state $|\psi_r\rangle$, which reduces the projection time. CCSD wavefunction is employed by performing a Hubbard–Stratonovich transformation on the exponential of the cluster operator.²⁴ Using trial states belonging to specific symmetry sectors also allows one to target the lowest energy states in the corresponding sectors.

As an illustrative example, we consider the D_{4h} symmetric transition state of cyclobutadiene. This state has a biradical character, which makes it a challenging problem for single reference electronic structure methods, in-

cluding spin-restricted and spin-unrestricted CCSD(T) (RCCSD(T) and UCCSD(T)). We performed fp-AFQMC calculations on this system using an MSD trial, obtained via heat-bath configuration interaction (HCI), as $|\psi_l\rangle$ and the spin-restricted CCSD (RCCSD) state as $|\psi_r\rangle$. The HCI state was obtained from an HCI calculation with a crude $\epsilon_1 = 10^{-4}$ in the full space except for four frozen HF orbitals. These orbitals were kept frozen in all correlated calculations. The geometry was taken from Ref. 24, which also reported fp-AFQMC energies for this system. As additional validation of our results, we converged the ph-AFQMC energy with respect to the number of determinants in the trial HCI state; we expect this energy to be nearly exact. Results are shown in Fig. 10.

The fp-AFQMC energy converges to the converged ph-AFQMC energy within statistical error bars. Our results suggest that hybrid (H)-AFQMC energies reported in Ref. 50 are likely biased as they are too much lower than our converged ph-AFQMC and fp-AFQMC energies. RCCSD(T) is about 13.5 mE_h higher than ph-AFQMC, while UCCSD(T) is closer but still about 4.5 mE_h higher. The differences in the barrier height largely come from transition state energies because all these methods work well for the equilibrium D_{2h} geometry.




FIG. 10. Convergence of the fp-AFQMC energy for the cyclobutadiene transition state in the cc-pVDZ basis set with projection time. The inset shows the convergence of ph-AFQMC energy with respect to the number of determinants in the trial.

E. AFQMC beyond the ground-state electronic structure energy

The preceding discussions concentrate on the ground-state *ab initio* electronic structure energy calculations. Apart from this, `ipie` now also accommodates other types of calculations within the framework of AFQMC, such as property calculations,³² finite temperature cal-

culations^{12,51} and ground-state calculations of electron-phonon coupled model systems.¹³ The code structures for these different features mirror those of the zero temperature *ab initio* electronic structure implementation.

1. Electronic structure at finite temperatures

An extension of the phaseless AFQMC method at zero temperature, finite-temperature AFQMC (FT-AFQMC), was developed to study systems at finite temperatures.¹² It is customary (though not necessary⁵²) to work in the grand canonical ensemble described by the temperature T , volume V , and chemical potential μ . The central quantities of interest are thermal expectation values computed from the partition function

$$\Xi = \text{Tr} \left[e^{-\beta(\hat{H} - \mu\hat{N})} \right], \quad (20)$$

where \hat{N} is the total number operator and μ is the chemical potential. The Boltzmann factor in Eq. (20) can be interpreted as a propagator in imaginary time $\tau = \beta$.

Analogously to the zero temperature case, τ is first discretized into l intervals of length $\Delta\tau = \tau/l$,

$$\Xi = \text{Tr} \left[e^{-\beta(\hat{H} - \mu\hat{N})} \right] = \text{Tr} \left[\lim_{l \rightarrow \infty} \prod_{k=1}^{k=l} e^{-\Delta\tau(\hat{H} - \mu\hat{N})} \right], \quad (21)$$

the short-time propagator is then Trotter decomposed as

$$e^{-\Delta\tau(\hat{H} - \mu\hat{N})} \simeq e^{-\frac{\Delta\tau}{2}(\hat{v}_0 - \mu\hat{N})} e^{\frac{\Delta\tau}{2} \sum_{\gamma} \hat{v}_{\gamma}^2} e^{-\frac{\Delta\tau}{2}(\hat{v}_0 - \mu\hat{N})}, \quad (22)$$

and the application of the Hubbard–Stratonovich transformation gives

$$e^{-\Delta\tau(\hat{H} - \mu\hat{N})} \simeq \int d\mathbf{x} p(\mathbf{x}) \hat{B}(\mathbf{x}, \Delta\tau, \mu), \quad (23)$$

where the one-body propagator \hat{B} is now also function of μ :

$$\hat{B}(\mathbf{x}, \Delta\tau, \mu) = e^{-\frac{\Delta\tau}{2}(\hat{v}_0 - \mu\hat{N})} e^{-\sqrt{\Delta\tau} \mathbf{x} \cdot \hat{\mathbf{v}}} e^{-\frac{\Delta\tau}{2}(\hat{v}_0 - \mu\hat{N})}. \quad (24)$$

The grand canonical partition function is thus evaluated as

$$\begin{aligned} \Xi &= \text{Tr} \left[e^{-\beta(\hat{H} - \mu\hat{N})} \right] \\ &= \int d\mathbf{x}_1 \cdots d\mathbf{x}_l \underbrace{p(\mathbf{x}_1) \cdots p(\mathbf{x}_l)}_{p(\mathbf{x}_1, \dots, \mathbf{x}_l)} \text{Tr} \left[\prod_{k=1}^l \hat{B}(\mathbf{x}_k, \Delta\tau, \mu) \right], \end{aligned} \quad (25)$$

with $p(\mathbf{x}_1, \dots, \mathbf{x}_l)$ being the probability of sampling a specific path designated by auxiliary fields $\mathbf{x}_1, \dots, \mathbf{x}_l$.

Furthermore, the trace in Eq. (25) can be written in terms of a determinant,^{53–55} which finally yields

$$\Xi = \int d\mathbf{x}_1 \cdots d\mathbf{x}_l p(\mathbf{x}_1, \dots, \mathbf{x}_l) \det \left[\mathbf{I} + \prod_{k=1}^l \mathbf{B}(\mathbf{x}_k, \Delta\tau, \mu) \right], \quad (26)$$

where \mathbf{I} is the identity matrix and \mathbf{B} is a matrix representation of \hat{B} in a single-particle basis. It should be noted that the partition function itself is not explicitly calculated—expectation values derived from it are. For some generic observable \hat{O} , its expectation value is

$$\langle \hat{O} \rangle = \frac{1}{\Xi} \text{Tr} \left[e^{-\beta(\hat{H} - \mu\hat{N})} \hat{O} \right] \quad (27)$$

$$= \int d\mathbf{X} \frac{p(\mathbf{X}) \text{Tr} [A(\mathbf{X}, \Delta\tau, \mu)]}{\int d\mathbf{Y} p(\mathbf{Y}) \text{Tr} [A(\mathbf{Y}, \Delta\tau, \mu)]} \times \quad (28)$$

$$\frac{\text{Tr} [A(\mathbf{X}, \Delta\tau, \mu) \hat{O}]}{\text{Tr} [A(\mathbf{X}, \Delta\tau, \mu)]}, \quad (29)$$

where we introduced the shorthand \mathbf{X} for the set of auxiliary fields along an imaginary path $\{\mathbf{x}_1, \dots, \mathbf{x}_l\}$, and

$$A(\mathbf{X}, \Delta\tau, \mu) = \prod_{k=1}^l \hat{B}(\mathbf{x}_k, \Delta\tau, \mu). \quad (30)$$

Rewriting Eq. (27) in the form Eq. (29) allows $\langle \hat{O} \rangle$ to be estimated through an importance sampling procedure. The field configurations \mathbf{X} are obtained via Monte Carlo sampling from the modified probability distribution $\tilde{p}(\mathbf{X}) = p(\mathbf{X}) \text{Tr} [A(\mathbf{X}, \Delta\tau, \mu)]$, while the computed random variables are the *local* expectation values

$$O_L(\mathbf{X}, \Delta\tau, \mu) = \frac{\text{Tr} [A(\mathbf{X}, \Delta\tau, \mu) \hat{O}]}{\text{Tr} [A(\mathbf{X}, \Delta\tau, \mu)]}. \quad (31)$$

Similar to the zero temperature case, the importance sampling is implemented by initializing a set of walkers in the space of auxiliary fields with weights w_i and propagating them in imaginary time. We therefore evaluate Eq. (29) in practice as

$$\langle \hat{O} \rangle = \frac{\sum_i w_i O_L(\mathbf{X}_i)}{\sum_i w_i}. \quad (32)$$

An example of the imaginary time trace for observables is provided in Fig. 11, which reproduces Fig. 3 in Ref. 12.




FIG. 11. Imaginary time trace from FT-AFQMC of the uniform electron gas model with 2 electrons. The top panel depicts the average electron number at imaginary time τ , $\langle \hat{N} \rangle_\tau$, while the bottom panel shows the error in the total energy compared to FCI, *i.e.* $E_{\text{FT-AFQMC}}(\tau) - E_{\text{FCI}}$. These results reproduce Fig. 3 in Ref. 12.

2. Electrons coupled to phonons

Utilizing the tools provided in ipie, one may write a new projector Monte Carlo method to calculate the ground state of electrons coupled to phonons (*i.e.*, lattice vibration).¹³

$$\hat{H} = \underbrace{\sum_{pq} h_{pq} a_p^\dagger a_q}_{\hat{H}_{\text{el}}} + \underbrace{\sum_{\nu} \omega_{\nu} b_{\nu}^\dagger b_{\nu}}_{\hat{H}_{\text{ph}}} + \underbrace{\sum_{pq\nu} g_{pq\nu} a_p^\dagger a_q (b_{\nu} + b_{\nu}^\dagger)}_{\hat{H}_{\text{el-ph}}} \quad (33)$$

where $b^{(\dagger)}$ represents the bosonic annihilation (and creation) operator, the first term represents the electronic band term, the second term represents the phonon band term, and the third term represents the coupling between electrons and phonons.

The corresponding ground-state projector is obtained from trotterizing the imaginary time propagator corresponding to Eq. (33),

$$e^{-\Delta\tau\hat{H}} = e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{el}}} e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{ph}}} e^{-\Delta\tau\hat{H}_{\text{el-ph}}} e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{ph}}} e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{el}}} + \mathcal{O}(\Delta\tau^3) \quad (34)$$

In our method, we choose walkers of the form $|\psi_w\rangle \otimes |\mathbf{X}_w\rangle \equiv |\psi_w(\tau), \mathbf{X}_w(\tau)\rangle$, with $|\psi_w\rangle$ being a single determinant and $|\mathbf{X}_w\rangle$ being the coordinates for phonon displacements. Using the Monte Carlo sampling, we can work in the position space of phonons without invoking any boson number truncation. This strategy differs from other standard Monte Carlo approaches in this area.^{56,57}

With importance sampling via a trial wavefunction, $|\Psi_T\rangle$, our global wavefunction is a weighted linear com-




FIG. 12. Reproduction of Fig. 3 in Ref. 13. Results are shown for a 20-site Holstein model at half-filling. We used a coherent state trial to obtain the AFQMC/S results. CSMP2 data points correspond to perturbation theory-based computations based on molecular orbitals obtained from SCF calculations for a Holstein Fock operator.¹³ Panels (a)-(d) show scans along phonon frequencies ω for λ equals (a) 0.1, (b) 0.3, (c) 0.8, and (d) 2.

bination of walker vibronic wavefunctions,

$$|\Psi(\tau)\rangle = \sum_{w=1}^{N_{\text{walkers}}} w_w(\tau) \frac{|\psi_w(\tau), \mathbf{X}_w(\tau)\rangle}{\langle \Psi_T | \psi_w(\tau), \mathbf{X}_w(\tau) \rangle}. \quad (35)$$

The propagation of the bosonic degrees of freedom under \hat{H}_{ph} is performed via a diffusion Monte Carlo algorithm.⁵⁸ With the walker representation, the propagation with \hat{H}_{ph} and $\hat{H}_{\text{el-ph}}$ is straightforward by exploiting the Thouless theorem. More details can be found in Ref. 13. As an example calculation, we picked the one-dimensional Holstein model under a periodic boundary condition,⁵⁹ which involves $h_{pq} = -t\delta_{p+1,q}$, $\omega_{\nu} = \omega$, and $g_{pq\nu} = g\delta_{pq}\delta_{p\nu}$. For 20-site at half-filling, we considered four different unitless electron-phonon coupling strengths $\lambda = \frac{g^2}{2t\omega}$, which are displayed in Fig. 12.

3. Automatic differentiable AFQMC

The computation of observables that do not commute with \hat{H} poses additional challenges in projector Monte Carlo, such as AFQMC. A recently proposed approach by Mahajan *et. al.*³² aims to compute the response estimator (Eq. (36)), which, though still not exact due to the discontinuity in the distribution, is accurate enough to give reliable results:

$$\langle \hat{O} \rangle_{\text{response}} = \left. \frac{dE_{\text{AFQMC}}(\lambda)}{d\lambda} \right|_{\lambda=0}. \quad (36)$$

Here λ is a parameter for the perturbed Hamiltonian: $\hat{H}(\lambda) = \hat{H} + \lambda\hat{O}$, and AFQMC energy $E_{\text{AFQMC}}(\lambda)$ is given by the Monte Carlo estimator

$$E_{\text{AFQMC}}(\lambda) = \frac{\sum_i w_i(\lambda) \frac{\langle \Psi_T(\lambda) | \hat{H}(\lambda) | \psi_i(\lambda) \rangle}{\langle \Psi_T(\lambda) | \psi_i(\lambda) \rangle}}{\sum_i w_i(\lambda)}, \quad (37)$$

where $w_i(\lambda)$ is the weight for the i -th walker, and $|\psi_i(\lambda)\rangle$ is the i -th walker state at coupling strength λ .

We implemented this scheme for computing observables using the Automatic Differentiation (AD) functionality of PyTorch⁶⁰ as an add-on within `ipie`. It is noteworthy that storing the computation graph for an entire AFQMC run requires substantial memory. In practice, we use the concept of AD blocks.³² We only track the computation graph within an AD block; there is no connection between different AD blocks. Using this approach, we manage memory costs by adjusting block size. We also use the gradient checkpointing technique⁶¹ to reduce the memory cost further.

For relatively large systems, the AFQMC calculation is parallelized using MPI to distribute walkers over MPI tasks. The differentiation of the AFQMC global energy estimator is not embarrassingly parallel since walkers will mix between different MPI ranks via population control. The AD implementation in `ipie` does not support communication between MPI tasks. Therefore, we perform ‘‘local’’ population control within an MPI rank. Because the local population control does not mix the walkers between different ranks, it is valid to regard those estimates as independent samples. We thus perform block analysis on those samples to obtain the final result. Effectively, this amounts to running very low walker population simulations (50 walkers per task here), which is only practically possible for small system sizes where population control is not a concern. For intermediate system sizes, one could imagine using large memory nodes and OpenMP threading or using single GPUs with a sufficient amount of memory.

Here, we present the results of AD-AFQMC calculations on various molecular systems. The molecular integrals are obtained by PySCF,²⁸ and the modified Cholesky decomposition is performed with `ipie`, with a threshold of 10^{-5} . We used a time step of 0.01 a.u., and periodic reorthogonalization of walkers is performed every 5 steps for all calculations. We used the restricted Hartree Fock (RHF) trial for all systems. For accurate statistical analysis, we perform block analysis for ≥ 200 gradient samples for each calculation. All AFQMC calculations are performed using the frozen-core approximation.

We benchmarked our implementation on various small molecules using aug-cc-pVTZ⁶² basis set and compared the results to Ref. 32 in Table I. All dipole moments align strongly with the AD-AFQMC results reported in Ref. 32. Furthermore, except for CO, AD-AFQMC matches experimental values better than RCCSD and exhibits accuracy comparable to RCCSD(T). This supports the widely held view that AFQMC’s accuracy falls

between CCSD and CCSD(T).

F. Enhanced integration testing and no-MPI mode

As `ipie` expands its functionalities, robust testing workflows become crucial. `ipie` supports a new integration testing framework that enhances the package’s robustness and adaptability and significantly contributes to its reliability and ease of use for the end-users. The GitHub continuous integration (CI) workflow automatically tests new pull requests, ensuring that every change to the codebase does not break existing functionalities. The workflow encompasses a comprehensive test suite and executes various linting and code formatting checks before running serial and parallel unit tests and integration tests.

Recognizing the diverse computational environments in which `ipie` might be deployed, we introduced the no-MPI mode. This mode is designed for situations where MPI is unavailable, or its use is not desired, offering greater flexibility for users and developers. The CI workflow includes a job that tests `ipie`’s functionality without the MPI dependency. This mode is particularly beneficial for users who wish to perform quantum Monte Carlo simulations on personal computing setups or in environments where setting up MPI is challenging.

V. INTERFACES TO EXTERNAL PACKAGES

A. Dice interface and SHCI-AFQMC

`ipie` includes utilities for converting the output from Dice,¹⁹ a package that employs semistochastic heat bath configuration interaction (SHCI)^{19,29} as the complete active space (CAS) solver.

Code Snippet 11 Convert Dice output to the MSD trial in `ipie`.

```
python -u /path-to-ipie/tools/extract_dice.py --dice-wfn
↪ /path-to-Dice-output/dets.bin --sort --verbose
```

which produces the `wfn.h5` file containing the coefficients and the indices of occupied orthogonal orbitals for spin-up and spin-down sectors that follows the block-formatted ($\alpha\alpha\cdots\beta\beta\cdots$) orbital convention in `ipie`. The SHCI calculation generates a good MSD trial for challenging systems.^{15,37} However, it is not a black-box approach and often requires careful handling, such as selecting CAS and using natural orbitals, as discussed in Section IV C.

These detailed discussions are beyond the scope of this article, and we briefly mention the procedures in the semi-blackbox example provided within `ipie`:¹⁵

TABLE I. Comparison of AD-AFQMC dipole moment (in a.u.) of various molecules at equilibrium geometry with the implementation in Ref. 32 and other quantum chemistry methods. The data using RCCSD and RCCSD(T) are also extracted from Ref. 32.

Molecule	This work	AD-AFQMC in Ref. 32	RCCSD	RCCSD(T)	Experiment
H ₂ O	0.723(2)	0.720(2)	0.7335	0.7247	0.730 ⁶³
NH ₃	0.592(2)	0.592(2)	0.6015	0.5938	0.581(1) ⁶⁴
CO	0.022(4)	0.019(4)	0.0199	0.0429	0.048(1) ⁶⁵
HCl	0.428(1)	0.429(1)	0.4318	0.4273	0.430 ⁶⁶
HBr	0.332(2)	0.329(2)	0.3289	0.3245	0.325 ⁶⁶

1. Initiate a preliminary rough SHCI calculation in an extensive active space for the system under study.
2. Derive the SHCI one-electron reduced density matrix (1-RDM), extracting the resultant natural orbital occupation number (NOON) and natural orbitals.
3. Define an active space criterion based on a predetermined NOON threshold.
4. Adjust the orbitals through rotation, aligning them with the unitary transformation specified by the natural orbitals.
5. Execute a refined SHCI self-consistent field calculation within the new active space determined in the previous step.
6. Use this MSD trial in AFQMC.

This strategy ensures the trial wavefunction encapsulates static correlations within the active space via the MSD trial. At the same time, AFQMC incorporates the residual dynamic correlations. This procedure is folded into the factory utility method `ipie.utils.from_dice.build_driver_from_shcisf`.

B. TREXIO support and CIPSI-AFQMC

The TREXIO library and file format have been developed to offer a robust and efficient solution for storing and exchanging wavefunction parameters and matrix elements.³⁰ This library supports bindings in several programming languages, including Python, and can be conveniently installed via the `pip` package manager.

The compatibility of `ipie` with the TREXIO format facilitates its integration with various software packages. Specifically, it allows `ipie` to utilize trial wavefunctions produced by Quantum Package,⁶⁷ along with the associated one-electron integrals and Cholesky-decomposed electron repulsion integrals. This interface allowed us to check if AFQMC could complement configuration interaction using perturbative selection done iteratively (CIPSI) calculations to improve full configuration interaction (FCI) energy estimates of large systems.

Typically, to estimate the FCI energy from a CIPSI calculation, one extrapolates to zero the variational energy,




FIG. 13. Energy of the trial wavefunction (E_{trial}) as a function of ΔE , where $\Delta E = E_{\text{rPT2}}$ for CIPSI calculations and $\Delta E = E_{\text{AFQMC}} - E_{\text{trial}}$ for AFQMC. exFCI is the extrapolated FCI energy obtained from CIPSI calculations.

E_{trial} , as a function of the renormalized second-order perturbative correction, E_{rPT2} .⁶⁷ However, the AFQMC energy, E_{AFQMC} , is anticipated to provide a closer approximation to the FCI energy than the sum of E_{trial} and E_{rPT2} . This expectation is particularly relevant for systems with large E_{rPT2} corrections.

We conducted AFQMC calculations on the nitroxyl and benzene molecules using CIPSI trial wavefunctions of increasing sizes. The results of these calculations are detailed in Table II. Supporting our expectation, Fig. 13 demonstrates that the AFQMC corrections align the data points along a straight line,⁶⁸ validating the hypothesis that AFQMC energies are more reliable for the extrapolation towards the FCI value, especially in cases with large E_{rPT2} corrections. Employing calculations with comparatively small wavefunctions for benzene, as shown in Table II, a three-point linear extrapolation based on the $r\text{PT2}$ correction yields a correlation energy of $-858.6 mE_h$. In contrast, extrapolation using AFQMC energies results in a correlation energy of $-862.2 mE_h$. This latter value is significantly closer to the correlation energy of $-863.4 mE_h$ achieved through CIPSI with 167 million determinants.⁶⁹

These preliminary calculations enabled by the TREXIO interface illustrate that the integration of AFQMC with CIPSI emerges as a promising methodol-

TABLE II. Renormalized PT2 correction E_{rPT2} , Variational energy E_{var} , and AFQMC energy E_{AFQMC} for the nitroxyl and benzene molecules as functions of the number of determinants in the variational space.

System, Basis	N_{det}	E_{trial}	$E_{\text{trial}} + E_{rPT2}$	E_{AFQMC}
HNO, 6-31G	1	-129.712554	-129.987361	-129.971(1)
	2	-129.752113	-129.987158	-129.974(1)
	41	-129.781149	-129.981219	-129.9727(6)
	748	-129.894775	-129.973405	-129.9724(2)
	2838	-129.942391	-129.975057	-129.9726(1)
	14201	-129.955368	-129.973620	-129.9727(1)
C ₆ H ₆ , cc-pVDZ	1	-230.7204904	-231.39330	-231.5887(7)
	23	-230.7615816	-231.42167	-231.5866(7)
	828	-230.8687072	-231.44602	-231.5864(6)
	15690	-231.1300087	-231.49330	-231.5848(4)
	109869	-231.3152399	-231.53022	-231.5842(4)

ogy for estimating the FCI energy of systems larger than those currently feasible.

C. FQE interfaces

The Fermionic quantum emulator (FQE)³¹ is a lightweight fermionic circuit simulator, which is particularly useful in quantum computing where it aids in the development and testing of quantum algorithms tailored for fermionic systems. `ipie` provides the conversion between the `ipie`'s MSD wavefunction and the FQE wavefunction.

VI. CONCLUSIONS AND OUTLOOKS

This paper summarized the improvements and new features added in `ipie` since its original release.¹⁵ These improvements enhance modularity and computational efficiency and offer intuitive user-end APIs. New features and interfaces aim to expand a broader spectrum of AFQMC calculations in quantum chemistry.

We summarize the key features we highlighted in this manuscript:

1. *Distributed Hamiltonians to remove the memory bottleneck.* We demonstrated `ipie`'s capacity for studying large systems deploying GPUs with significantly higher efficiency than CPU-based implementations, exemplified in our case study assessing the interaction energies in a benzene dimer.
2. *GPU support for MSD trial wavefunctions.* With customized CUDA kernels, we enabled an efficient realization of Wick's theorem. Timing benchmarks for $[\text{Cu}_2\text{O}_2]^{2+}$ and $[\text{Fe}_2\text{S}_2(\text{SCH}_3)]^{2-}$ were shown to achieve more than an order of magnitude speedup compared to our CPU implementation for a large MSD trial.⁴⁷
3. *Support for complex-valued Cholesky vectors.* `ipie` can handle complex-valued Cholesky vectors that

may arise when the underlying basis functions are complex-valued.

4. *Free-projection AFQMC.* A numerically exact AFQMC approach can be used to study small strongly correlated systems.
5. *Finite-temperature AFQMC.* A finite-temperature AFQMC algorithm based on the grand canonical ensemble was added.
6. *Electron-phonon QMC.* A QMC algorithm that computes the ground state of electron-phonon problems was added.
7. *Automatic differentiable AFQMC.* We offer AFQMC property calculations via automatic differentiation.
8. *External package interfaces.* `ipie` is now interfaced with PySCF, Dice, TREXIO, and FQE.

We hope that `ipie` will serve as a community code base for developing *ab initio* AFQMC methods and their applications. Furthermore, as `ipie` is written mainly in Python, we anticipate its use in machine learning and quantum computing communities will also grow.

VII. ACKNOWLEDGEMENTS

Work by T.J. and J.L. were supported by Harvard University's startup funds, the quantum information science project from the U.S. Department of Energy (DOE), and Wellcome Leap as part of the Quantum for Bio Program. Computations were carried out partly on the FASRC cluster supported by the FAS Division of Science Research Computing Group at Harvard University. This work also used the Delta system at the National Center for Supercomputing Applications through allocation CHE230032, CHE230088, and PHY230192 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation

grants #2138259, #2138286, #2138307, #2137603, and #2138296. P.F.L. and A.S. have received financial support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement no. 863481). A.S. was also supported by the European Centre of Excellence in Exascale Computing (TRES) which has received funding from the European Union’s Horizon 2020 — Research and Innovation program — under grant agreement no. 952165. S.F.U. thanks David Reichman for his support. We thank Yifei Huang and Dingshun Lv for their GPU-MSD code contributions.⁴⁷

VIII. DATA AVAILABILITY

The data that support the findings of this study are openly available in our Zenodo repository at <https://doi.org/10.5281/zenodo.10939113>.

- ¹S. Zhang, J. Carlson, and J. E. Gubernatis, “Constrained Path Quantum Monte Carlo Method for Fermion Ground States,” *Phys. Rev. Lett.* **74**, 3652–3655 (1995).
- ²S. Zhang and H. Krakauer, “Quantum Monte Carlo Method using Phase-Free Random Walks with Slater Determinants,” *Phys. Rev. Lett.* **90**, 136401 (2003).
- ³J. Lee, H. Q. Pham, and D. R. Reichman, “Twenty Years of Auxiliary-Field Quantum Monte Carlo in Quantum Chemistry: An Overview and Assessment on Main Group Chemistry and Bond-Breaking,” *J. Chem. Theory Comput.* **18**, 7024–7042 (2022).
- ⁴M. Motta and S. Zhang, “Ab initio computations of molecular systems by the auxiliary-field quantum Monte Carlo method,” *WIREs Computational Molecular Science* **8**, e1364 (2018).
- ⁵W. J. Huggins, B. A. O’Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, “Unbiasing fermionic quantum Monte Carlo with a quantum computer,” *Nature* **603**, 416–420 (2022).
- ⁶K. Wan, W. J. Huggins, J. Lee, and R. Babbush, “Matchgate shadows for fermionic quantum simulation,” *Commun. Math. Phys.* **404**, 629–700 (2023).
- ⁷M. Amsler, P. Deglmann, M. Degroote, M. P. Kaicher, M. Kiser, M. Kühn, C. Kumar, A. Maier, G. Samsonidze, A. Schroeder, et al., “Classical and quantum trial wave functions in auxiliary-field quantum monte carlo applied to oxygen allotropes and a cubr2 model system,” *J. Chem. Phys.* **159** (2023).
- ⁸M. Kiser, A. Schroeder, G.-L. R. Anselmetti, C. Kumar, N. Moll, M. Streif, and D. Vodola, “Classical and quantum cost of measurement strategies for quantum-enhanced auxiliary field quantum monte carlo,” *New J. Phys.* (2023).
- ⁹B. Huang, Y.-T. Chen, B. Gupt, M. Suchara, A. Tran, S. McArdle, and G. Galli, “Evaluating a quantum-classical quantum Monte Carlo algorithm with Matchgate shadows,” *arXiv* (2024), 2404.18303.
- ¹⁰J. Lee, F. D. Malone, and M. A. Morales, “An auxiliary-field quantum monte carlo perspective on the ground state of the dense uniform electron gas: An investigation with hartree-fock trial wavefunctions,” *J. Chem. Phys.* **151** (2019).
- ¹¹J. Lee and D. R. Reichman, “Stochastic resolution-of-the-identity auxiliary-field quantum monte carlo: Scaling reduction without overhead,” *J. Chem. Phys.* **153** (2020), 10.1063/5.0015077.
- ¹²J. Lee, M. A. Morales, and F. D. Malone, “A phaseless auxiliary-field quantum monte carlo perspective on the uniform electron gas at finite temperatures: Issues, observations, and benchmark study,” *J. Chem. Phys.* **154** (2021).
- ¹³J. Lee, S. Zhang, and D. R. Reichman, “Constrained-path auxiliary-field quantum monte carlo for coupled electrons and phonons,” *Phys. Rev. B* **103**, 115123 (2021).
- ¹⁴J. Lee, F. D. Malone, M. A. Morales, and D. R. Reichman, “Spectral functions from auxiliary-field quantum monte carlo without analytic continuation: The extended koopmans’ theorem approach,” *J. Chem. Theory Comput.* **17**, 3372–3387 (2021).
- ¹⁵F. D. Malone, A. Mahajan, J. S. Spencer, and J. Lee, “Ipie: A python-based auxiliary-field quantum monte carlo program with flexibility and efficiency on cpus and gpus,” *J. Chem. Theory Comput.* **19**, 109–121 (2022).
- ¹⁶S. K. Lam, A. Pitrou, and S. Seibert, “Proceedings of the second workshop on the llvm compiler infrastructure in hpc. llvm’15,” (2015).
- ¹⁷C. J. Cramer, M. Wloch, P. Piecuch, C. Puzzarini, and L. Gagliardi, “Theoretical models on the cu2o2 torture track: Mechanistic implications for oxytyrosinase and small-molecule analogues,” *J. Phys. Chem. A* **110**, 1991–2004 (2006).
- ¹⁸P. R. Kent, A. Annaberdiyev, A. Benali, M. C. Bennett, E. J. Landinez Borda, P. Doak, H. Hao, K. D. Jordan, J. T. Krogel, I. Kylänpää, et al., “Qmcpack: Advances in the development, efficiency, and application of auxiliary field and real-space variational and diffusion quantum monte carlo,” *J. Chem. Phys.* **152** (2020).
- ¹⁹S. Sharma, A. A. Holmes, G. Jeanmairet, A. Alavi, and C. J. Umrigar, “Semistochastic heat-bath configuration interaction method: Selected configuration interaction with semistochastic perturbation theory,” *J. Chem. Theory Comput.* **13**, 1595–1604 (2017).
- ²⁰L. Hehn, P. Deglmann, and M. Kühn, “Chelate complexes of 3d transition metal ions—a challenge for electronic-structure methods?” *J. Chem. Theory Comput.* (2024).
- ²¹V. P. Vysotskiy, C. Filippi, and U. Ryde, “Scalar relativistic all-electron and pseudopotential ab initio study of a minimal nitrogenase [fe (sh) 4h]- model employing coupled-cluster and auxiliary-field quantum monte carlo many-body methods,” *J. Phys. Chem. A* (2024).
- ²²M. S. Chen, J. Lee, H.-Z. Ye, T. C. Berkelbach, D. R. Reichman, and T. E. Markland, “Data-efficient machine learning potentials from transfer learning of periodic correlated electronic structure methods: Liquid water at afqmc, ccSD, and ccSD (t) accuracy,” *J. Chem. Theory Comput.* **19**, 4510–4519 (2023).
- ²³T. Jiang, B. O’Gorman, A. Mahajan, and J. Lee, “Unbiasing fermionic auxiliary-field quantum monte carlo with matrix product state trial wavefunctions,” *arXiv* (2024), 2405.05440.
- ²⁴A. Mahajan and S. Sharma, “Taming the sign problem in auxiliary-field quantum monte carlo using accurate wave functions,” *J. Chem. Theory Comput.* **17**, 4786–4798 (2021).
- ²⁵E. Vitali, P. Rosenberg, and S. Zhang, “Calculating ground state properties of correlated fermionic systems with BCS trial wave functions in Slater determinant path-integral approaches,” *Phys. Rev. A* **100**, 023621 (2019).
- ²⁶C.-C. Chang, B. M. Rubenstein, and M. A. Morales, “Auxiliary-field-based trial wave functions in quantum Monte Carlo calculations,” *Phys. Rev. B* **94**, 235144 (2016).
- ²⁷“Plum: Multiple dispatch in python,” <https://github.com/beartype/plum>.
- ²⁸Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, et al., “Pyscf: the python-based simulations of chemistry framework,” *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **8**, e1340 (2018).
- ²⁹A. A. Holmes, N. M. Tubman, and C. Umrigar, “Heat-bath configuration interaction: An efficient selected configuration interaction algorithm inspired by heat-bath sampling,” *J. Chem. Theory Comput.* **12**, 3674–3680 (2016).
- ³⁰E. Posenitskiy, V. G. Chilkuri, A. Ammar, M. Hapka, K. Pernal, R. Shinde, E. J. Landinez Borda, C. Filippi, K. Nakano, O. Kohulák, S. Sorella, P. de Oliveira Castro, W. Jalby, P. L. Ríos, A. Alavi, and A. Scemama, “TRESIO: A file format and library for quantum chemistry,” *J. Chem. Phys.* **158** (2023), 10.1063/5.0148161, 174801.
- ³¹N. C. Rubin, K. Gunst, A. White, L. Freitag, K. Throssell, G. K.-L. Chan, R. Babbush, and T. Shiozaki, “The fermionic quantum

- emulator, version 0.2.0," (2021); "The fermionic quantum emulator," *Quantum* **5**, 568 (2021).
- ³²A. Mahajan, J. S. Kurian, J. Lee, D. R. Reichman, and S. Sharma, "Response properties in phaseless auxiliary field quantum Monte Carlo," *J. Chem. Phys.* **159**, 184101 (2023).
- ³³J. Hubbard, "Calculation of partition functions," *Phys. Rev. Lett.* **3**, 77 (1959).
- ³⁴R. Stratonovich, *Sov. Phys. Dokl.* **2**, 416 (1958).
- ³⁵J. Ren, W. Li, T. Jiang, Y. Wang, and Z. Shuai, "The renormalizer package," <https://github.com/shuaigroup/renormalizer> (2021); "Time-dependent density matrix renormalization group method for quantum dynamics in complex systems," *Wiley Interdiscip Rev Comput Mol Sci* **12**, e1614 (2022).
- ³⁶H. Flyvbjerg and H. G. Petersen, "Error estimates on averages of correlated data," *J. Chem. Phys.* **91**, 461–466 (1989).
- ³⁷A. Mahajan, J. Lee, and S. Sharma, "Selected configuration interaction wave functions in phaseless auxiliary field quantum monte carlo," *J. Chem. Phys.* **156** (2022), <https://doi.org/10.1063/5.0087047>.
- ³⁸P. Jurečka, J. Šponer, J. Černý, and P. Hobza, "Benchmark database of accurate (mp2 and ccsd (t) complete basis set limit) interaction energies of small model complexes, dna base pairs, and amino acid pairs," *Phys. Chem. Chem. Phys.* **8**, 1985–1993 (2006).
- ³⁹F. Weigend, M. Kattannek, and R. Ahlrichs, "Approximated electron repulsion integrals: Cholesky decomposition versus resolution of the identity methods," *J. Chem. Phys.* **130** (2009).
- ⁴⁰A. Halkier, T. Helgaker, P. Jørgensen, W. Klopper, H. Koch, J. Olsen, and A. K. Wilson, "Basis-set convergence in correlated calculations on ne, n2, and h2o," *Chem. Phys. Lett.* **286**, 243–252 (1998).
- ⁴¹"Benchmark data for S22 reaction set," (2024).
- ⁴²L. A. Burns, M. S. Marshall, and C. D. Sherrill, "Appointing silver and bronze standards for noncovalent interactions: A comparison of spin-component-scaled (scs), explicitly correlated (f12), and specialized wavefunction approaches," *J. Chem. Phys.* **141** (2014), <https://doi.org/10.1063/1.4903765>.
- ⁴³Y. S. Al-Hamdani, P. R. Nagy, A. Zen, D. Barton, M. Kállay, J. G. Brandenburg, and A. Tkatchenko, "Interactions between large molecules pose a puzzle for reference quantum mechanical methods," *Nat. Commun.* **12**, 3927 (2021).
- ⁴⁴S. F. Boys and F. Bernardi, "The calculation of small molecular interactions by the differences of separate total energies. some procedures with reduced errors," *Mol. Phys.* **19**, 553–566 (1970).
- ⁴⁵S. Blaschke and S. Stopkowicz, "Cholesky decomposition of complex two-electron integrals over gao: Efficient mp2 computations for large molecules in strong magnetic fields," *J. Chem. Phys.* **156** (2022), <https://doi.org/10.1063/5.0076588>.
- ⁴⁶M. Siewattana, W. Purwanto, S. Zhang, H. Krakauer, and E. J. Walter, "Phaseless auxiliary-field quantum monte carlo calculations with plane waves and pseudopotentials: Applications to atoms and molecules," *Phys. Rev. B* **75**, 245123 (2007).
- ⁴⁷Y. Huang, Z. Guo, H. Q. Pham, and D. Lv, "Gpu-accelerated auxiliary-field quantum monte carlo with multi-slater determinant trial states," *arXiv* (2024), 2406.08314.
- ⁴⁸Z. Li and G. K.-L. Chan, "Spin-projected matrix product states: Versatile tool for strongly correlated systems," *J. Chem. Theory Comput.* **13**, 2681–2695 (2017).
- ⁴⁹R. Assaraf and M. Caffarel, "Zero-Variance Principle for Monte Carlo Algorithms," *Phys. Rev. Lett.* **83**, 4682–4685 (1999).
- ⁵⁰Y. Chen, L. Zhang, W. E, and R. Car, "Hybrid auxiliary field quantum monte carlo for molecular systems," *J. Chem. Theory Comput.* **19**, 4484–4493 (2023).
- ⁵¹Y.-Y. He, M. Qin, H. Shi, Z.-Y. Lu, and S. Zhang, "Finite-temperature auxiliary-field quantum monte carlo: Self-consistent constraint and systematic approach to low temperatures," *Phys. Rev. B* **99**, 045108 (2019).
- ⁵²T. Shen, Y. Liu, Y. Yu, and B. M. Rubenstein, "Finite temperature auxiliary field quantum Monte Carlo in the canonical ensemble," *J. Chem. Phys.* **153** (2020), 10.1063/5.0026606.
- ⁵³R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, "Monte carlo calculations of coupled boson-fermion systems. i," *Phys. Rev. D* **24**, 2278–2286 (1981).
- ⁵⁴J. E. Hirsch, "Two-dimensional hubbard model: Numerical simulation study," *Phys. Rev. B* **31**, 4403–4419 (1985).
- ⁵⁵R. R. d. Santos, "Introduction to quantum monte carlo simulations for fermionic systems," *Braz. J. Phys.* **33**, 36–54 (2003).
- ⁵⁶A. Macridin, "Phonons, Charge and Spin in Correlated Systems," PhD Thesis (2003).
- ⁵⁷D. J. J. Marchand, P. C. E. Stamp, and M. Berciu, "Dual coupling effective band model for polarons," *Phys. Rev. B* **95**, 035117 (2017).
- ⁵⁸M. H. Kalos, D. Levesque, and L. Verlet, "Helium at zero temperature with hard-sphere and other forces," *Phys. Rev. A* **9**, 2178–2195 (1974).
- ⁵⁹T. Holstein, "Studies of polaron motion : Part I. The molecular-crystal model," *Ann. Phys.* **8**, 325–342 (1959).
- ⁶⁰A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc., 2019) pp. 8024–8035.
- ⁶¹T. Chen, B. Xu, C. Zhang, and C. Guestrin, "Training deep nets with sublinear memory cost," *arXiv:1604.06174* (2016), <https://doi.org/10.48550/arXiv.1604.06174>.
- ⁶²T. H. Dunning Jr, "Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen," *J. Chem. Phys.* **90**, 1007–1023 (1989).
- ⁶³S. L. Shostak, W. L. Eberstein, and J. S. Muentzer, "The dipole moment of water. i. dipole moments and hyperfine properties of h2o and hdo in the ground and excited vibrational states," *J. Chem. Phys.* **94**, 5875–5882 (1991).
- ⁶⁴F. Shimizu, "Stark spectroscopy of nh3 ν_2 band by 10- μ co2 and n2o lasers," *J. Chem. Phys.* **52**, 3572–3576 (1970).
- ⁶⁵J. Muentzer, "Electric dipole moment of carbon monoxide," *J. Mol. Spectrosc.* **55**, 490–491 (1975).
- ⁶⁶F. Lovas, E. Tiemann, J. Coursey, S. Kotochigova, J. Chang, K. Olsen, and R. Dragoset, "Diatomic spectral database," (2003).
- ⁶⁷Y. Garniron, T. Applencourt, K. Gasperich, A. Benali, A. Ferté, J. Paquier, B. Pradines, R. Assaraf, P. Reinhardt, J. Toulouse, P. Barbaresco, N. Renon, G. David, J.-P. Malrieu, M. Véril, M. Caffarel, P.-F. Loos, E. Giner, and A. Scemama, "Quantum package 2.0: An open-source determinant-driven suite of programs," *J. Chem. Theory Comput.* **15**, 3591–3609 (2019).
- ⁶⁸H. G. A. Burton and P.-F. Loos, "Rationale for the extrapolation procedure in selected configuration interaction," *J. Chem. Phys.* **160**, 104102 (2024).
- ⁶⁹P.-F. Loos, Y. Damour, and A. Scemama, "The performance of CIPSI on the ground state electronic energy of benzene," *J. Chem. Phys.* **153** (2020), 10.1063/5.0027617.