

Tutorial

Path integral Monte Carlo ground state approach: formalism, implementation, and applications

Yangqian Yan^{1,2} and D Blume^{2,3}

¹Department of Physics, Indiana University Purdue University Indianapolis (IUPUI), Indianapolis, IN 46202, United States of America

²Department of Physics and Astronomy, Washington State University, Pullman, WA 99164-2814, United States of America

³Homer L Dodge Department of Physics and Astronomy, The University of Oklahoma, 440 W Brooks Street, Norman, OK 73019, United States of America

E-mail: doerte.blume-1@ou.edu

Received 6 October 2016, revised 22 August 2017

Accepted for publication 19 September 2017

Published 31 October 2017



CrossMark

Abstract

Monte Carlo techniques have played an important role in understanding strongly correlated systems across many areas of physics, covering a wide range of energy and length scales. Among the many Monte Carlo methods applicable to quantum mechanical systems, the path integral Monte Carlo approach with its variants has been employed widely. Since semi-classical or classical approaches will not be discussed in this review, path integral based approaches can for our purposes be divided into two categories: approaches applicable to quantum mechanical systems at zero temperature and approaches applicable to quantum mechanical systems at finite temperature. While these two approaches are related to each other, the underlying formulation and aspects of the algorithm differ. This paper reviews the path integral Monte Carlo ground state (PIGS) approach, which solves the time-independent Schrödinger equation. Specifically, the PIGS approach allows for the determination of expectation values with respect to eigen states of the few- or many-body Schrödinger equation provided the system Hamiltonian is known. The theoretical framework behind the PIGS algorithm, implementation details, and sample applications for fermionic systems are presented.

Keywords: Monte Carlo methods, cold atoms, PIGS

(Some figures may appear in colour only in the online journal)

1. Introduction

Monte Carlo techniques have found many applications, ranging from the modeling of the stock market to the simulation of classical and quantum spin models [1]. This review introduces the path integral Monte Carlo ground state (PIGS) method [2–5], which allows for the treatment of quantum mechanical systems with continuous spatial degrees of freedom at zero temperature. The PIGS method is a variant of the finite temperature path integral Monte Carlo method [6]. The key quantity in the finite temperature path integral Monte

Carlo approach is the density matrix $\hat{\rho}$, $\hat{\rho} = \exp[-\hat{H}/(k_B T)]$, where \hat{H} denotes the quantum mechanical system Hamiltonian, k_B the Boltzmann constant, and T the temperature. Making the formal replacement $(k_B T)^{-1} = t/\hbar$, where t is the time, the density matrix turns into the time evolution operator. Introducing the imaginary time τ , $\tau = t/\hbar$, and repeatedly acting with the ‘imaginary time evolution operator’ $\hat{G}(\Delta\tau)$, $\hat{G}(\Delta\tau) = \exp(-\Delta\tau\hat{H})$ (assuming $\Delta\tau$ is, using a metric to be defined later, small) onto an initial state, the ground state wave function, or more precisely the lowest energy state that has finite overlap with the initial state, is

projected out. This projection idea is the key concept behind the PIGS approach as well as many other imaginary time propagation schemes [7, 8]. Unlike grid based or basis set expansion based approaches, the PIGS approach is applicable to systems with varying degrees of freedom, i.e. few- and many-body systems. This versatility of the PIGS approach stems from the fact that the action of the imaginary time evolution operator on the initial (or propagated) state is evaluated stochastically, i.e. by means of Monte Carlo Metropolis sampling.

While the finite temperature path integral Monte Carlo algorithm, which—as has already been alluded to—has many features in common with the PIGS algorithm, has been reviewed quite extensively [6, 9], the PIGS algorithm has not, despite its generality, been reviewed in detail in the literature. The present paper thus serves three main purposes: (i) it develops, starting from equations that should be familiar to an advanced undergraduate student, the theoretical concepts behind the PIGS algorithm. (ii) It details how the relevant equations can be evaluated numerically, provides a good number of implementation details, and discusses various aspects regarding the algorithm performance. (iii) It presents applications of the PIGS algorithm to fermionic systems.

The PIGS approach allows one to solve the time-independent non-relativistic Schrödinger equation. Since the PIGS algorithm does not provide the full wave function in numerical or analytical form, the type of expectation values that one would like to determine needs to be specified *a priori* rather than *a posteriori*. In particular, an estimator has to be derived and implemented for each observable. The PIGS algorithm works, as already mentioned, through imaginary time propagation. It is imperative to clarify upfront that the imaginary time propagation is a numerical tool that facilitates projecting out unwanted excited state contributions. An extension to the real time evolution is, in general, not feasible, at least not for systems with a large number of degrees of freedom (for small systems, grid based real time propagation schemes do, of course, exist). A key ingredient of the PIGS algorithm is the stochastic evaluation of high-dimensional integrals, which arise from the discretization of the imaginary time and from the intrinsic degrees of freedom (particle coordinates) of the system under study. The stochastic Monte Carlo based approach to evaluating these integrals makes the PIGS method applicable to large systems containing as many as hundreds of particles. However, as in many Monte Carlo techniques, the treatment of identical fermions leads to the infamous Fermi sign problem. This tutorial applies the PIGS algorithm to small fermionic systems with zero-range interactions. It is shown that the sign problem can be ‘postponed’ but not be avoided, i.e. application of the PIGS algorithm to larger fermionic systems with zero-range interactions will necessarily fail.

The PIGS algorithm has been applied to systems relevant to physics and chemistry. For example, our understanding of pristine and doped bosonic helium clusters of varying size has been informed by PIGS calculations [10]. Unlike alternative zero-temperature methods such as the variational Monte Carlo method [7, 11] and the diffusion quantum Monte Carlo method

with mixed estimators [12, 13], the PIGS approach is known to yield unbiased results for structural properties like the radial density and pair distribution function. Here the term ‘unbiased’ refers to the fact that the resulting structural properties are independent of the initial state, provided the initial state has finite overlap with the state of interest and provided the state of interest is the lowest energy state with a particular symmetry. Moreover, the condensate fraction [6, 14] and Renyi entropy [15, 16] are observables that can be calculated relatively straightforwardly within the PIGS algorithm (or at least more straightforwardly than within a number of other approaches). The PIGS algorithm has also been, among others, applied to bulk helium in varying spatial dimensions [14, 17–24], liquid helium in nano-pores [25, 26], molecular para-hydrogen in nano-pores [27, 28], molecular para-hydrogen clusters [29–31], hard-sphere bosons [32, 33], dipolar systems [34, 35], and cold atoms loaded into optical lattices [36].

The applications presented in this review deal with cold atom systems with infinitely large two-body *s*-wave scattering length a_s [37–39], which are—like helium droplets—strongly interacting. However, the average interparticle spacing in cold atom systems tends to be significantly larger than that in helium droplets. This implies that the sample applications presented in this review deal with Hamiltonian that are characterized by vastly different length scales. To describe systems for which the average interparticle distance is large compared to the two-body interaction range, we employ two-body zero-range interactions. The use of two-body zero-range interactions removes the two-body range from the problem. If a_s is sent to infinity, as in the applications presented in this review, then two-component fermions are characterized by the same number of length scales as the corresponding non-interacting system [37, 38, 40, 41].

For bosons, in contrast, a three-body parameter, which can be defined in terms of the size of one of the extremely weakly bound Efimov trimers, sets a length scale of the interacting system even if the range of the two-body interactions is zero. Since the use of two-body zero-range interactions in continuum Monte Carlo calculations is a fairly novel development [42–48], the associated implementation details are discussed in detail. Simulation results are presented for fermionic systems. Application of the algorithm to bosons requires only a few changes in the code; however, due to the existence of a three-body parameter, the number of time slices, e.g., is much larger than for fermionic systems with two-body zero-range interactions.

The remainder of this article is organized as follows. Section 2 introduces, starting from the non-relativistic Schrödinger equation, the key quantum mechanical equations behind the PIGS algorithm. Section 3 discusses a number of theoretical concepts that are needed to reformulate the basic quantum mechanical equations in a form amenable to computer simulations; many considerations in this section do not only apply to the PIGS algorithm but also to other Monte Carlo algorithms. Section 3.1 introduces some basic ideas. Sections 3.2 and 3.3 discuss two different approaches for approximating the short-time propagator, namely a Trotter formula based approach and an approach that utilizes the so-called pair product approximation; these two approaches are

compared in section 3.4. The use of two-body zero-range interactions within the pair product approximation is discussed in section 3.5. Section 4 ‘translates’ the formalism introduced in sections 2 and 3 into an algorithm. Section 4.1 introduces the basics of Monte Carlo sampling of high-dimensional integrals while section 4.2 reviews formal aspects of the Monte Carlo Metropolis sampling. Section 4.3 discusses the generation of new configurations, i.e. the moves employed in the PIGS algorithm; as applicable, differences to the path integral Monte Carlo algorithm are pointed out. Sections 4.4 and 4.5 discuss the determination of expectation values and associated error bars, respectively. Last, section 4.6 discusses how to treat permutations in the PIGS algorithm; this discussion is particularly relevant if the system contains two or more identical fermions.

Section 5 presents a number of applications to harmonically trapped equal-mass two-component Fermi gases. The simulation results are discussed from two different angles. On the one hand, ‘technical aspects’ such as convergence with respect to the propagation time and the time step are discussed. On the other hand, the physical relevance of the simulation results presented is highlighted. Spin-balanced systems with up to $N = 10$ particles and a non-interacting Fermi gas with a single impurity with up to $N = 5$ particles are considered. In both cases, interspecies two-body zero-range interactions with infinitely large s -wave scattering length are employed. The construction of different types of trial functions ψ_T is discussed and the dependence of the simulation results on ψ_T is elucidated. PIGS results for the energy, pair distribution function, and contact are presented and compared to literature results where available. Lastly, section 6 provides a summary and an outlook.

2. Quantum mechanical foundation

We consider N non-relativistic particles described by the time-independent Hamiltonian \hat{H} at zero temperature. The Hamiltonian may contain single-particle potentials, two-body potentials, and higher-body potentials. We work in position space, where the potentials are local, i.e. we consider potentials that only depend on the position vectors and not on the momentum vectors as would be the case if, e.g., spin-orbit coupling terms were present [49–51]. The position vector for the j th particle with mass m_j is denoted by \mathbf{r}_j and we collectively denote the position vectors of all the particles by \mathbf{R} , $\mathbf{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$. The stationary eigen states and corresponding eigen energies are denoted by $\psi_j(\mathbf{R})$ and E_j , where $j = 0, 1, 2, \dots$. The $\psi_j(\mathbf{R})$ form a complete set and we are, throughout this article, interested in systems that support at least one N -body bound state. The treatment of scattering states by means of quantum Monte Carlo approaches is, in general, a challenging task [52–55] that is beyond the scope of this paper. The PIGS algorithm allows one to calculate a subset of the bound state energies as well as expectation values such as the pair distribution functions associated with the corresponding eigen states.

The PIGS algorithm is rooted in imaginary time propagation, a concept that is used widely to find the ground state or selected excited states of linear and non-linear Schrödinger equations [8, 56]. The concept of imaginary time propagation is also used to solve non-quantum mechanical wave equations. In what follows, we restrict ourselves, for concreteness, to the linear Schrödinger equation. To illustrate the key idea behind imaginary time propagation algorithms, we assume that the ground state is non-degenerate, i.e. that $E_0 < E_j$ for $j = 1, 2, \dots$. We consider an initial trial function $\psi_T(\mathbf{R})$, which does not have to be normalized, that has finite overlap with the ground state wave function $\psi_0(\mathbf{R})$. To analyze what happens when this trial function is propagated in imaginary time, we decompose the trial function into the eigen states $\psi_j(\mathbf{R})$ of the Hamiltonian \hat{H} ,

$$\psi_T(\mathbf{R}) = \sum_{j=0}^{\infty} c_j \psi_j(\mathbf{R}), \quad (1)$$

where c_0 is non-zero by assumption. Using equation (1), $\psi_T(\mathbf{R})$,

$$\psi_\tau(\mathbf{R}) = \exp(-\tau\hat{H})\psi_T(\mathbf{R}), \quad (2)$$

can be written as

$$\psi_\tau(\mathbf{R}) = \exp(-\tau E_0) \times \left\{ c_0 \psi_0(\mathbf{R}) + \sum_{j=1}^{\infty} c_j \exp[-\tau(E_j - E_0)] \psi_j(\mathbf{R}) \right\}. \quad (3)$$

Since E_j is, by assumption, greater than E_0 , the excited states contained in $\psi_\tau(\mathbf{R})$ decay out during the imaginary time propagation. In the $\tau \rightarrow \infty$ limit, $\psi_\tau(\mathbf{R})$ approaches, except for an overall factor, the eigen state $\psi_0(\mathbf{R})$. Correspondingly, the energy E_τ ,

$$E_\tau = \frac{\langle \psi_\tau | \hat{H} | \psi_\tau \rangle}{\langle \psi_\tau | \psi_\tau \rangle}, \quad (4)$$

approaches the exact ground state energy E_0 exponentially in the $\tau \rightarrow \infty$ limit. For finite τ , E_τ provides an upper bound to the exact eigen energy. This suggests that one can obtain a reliable estimate of E_0 by extrapolating the E_τ for various finite τ to the $\tau \rightarrow \infty$ limit. Expectation values of an arbitrary operator \hat{O} can be written analogously,

$$O_\tau = \frac{\langle \psi_\tau | \hat{O} | \psi_\tau \rangle}{\langle \psi_\tau | \psi_\tau \rangle}, \quad (5)$$

where O_τ denotes the τ -dependent expectation value. The convergence of O_τ toward the exact expectation value with respect to $\psi_0(\mathbf{R})$ may not be simply exponential and needs to be analyzed carefully for each operator \hat{O} (see section 5 for examples).

Equations (2), (4), and (5) constitute the starting point of the PIGS algorithm (see section 4). Based on these equations, two ingredients or components of the PIGS algorithm can already be identified. (i) An initial trial function $\psi_T(\mathbf{R})$ needs to be supplied by the ‘simulator’. From equation (3) it is clear that the efficiency of the PIGS algorithm depends on the overlap between $\psi_T(\mathbf{R})$ and $\psi_0(\mathbf{R})$: if all c_j with $j > 0$ vanish, then the imaginary time propagation is not needed at all. If the c_j for the states that lie energetically close to E_0 vanish, then

the decay of the excited states is fast, i.e. small τ should suffice. The construction of $\psi_T(\mathbf{R})$ is, of course, strongly dependent on the Hamiltonian under study. Examples are discussed in section 5. (ii) Given an initial trial function $\psi_T(\mathbf{R})$, the action of $\exp(-\tau\hat{H})$ onto $\psi_T(\mathbf{R})$ needs to be evaluated. The PIGS algorithm as well as many other algorithms accomplish this by dividing τ into multiple smaller imaginary time steps $\Delta\tau$. While non-Monte Carlo based approaches are, typically, restricted to relatively small system sizes, the PIGS algorithm as well as other Monte Carlo algorithms are designed to treat systems for which \mathbf{R} can be a high-dimensional vector.

The discussion thus far focused on determining the absolute ground state of the system. The outlined formalism can be readily adopted to the determination of the energetically lowest-lying state with a given symmetry. For concreteness, let us assume that the total angular momentum L and the total parity Π are good quantum numbers and that the absolute ground state has vanishing angular momentum ($L = 0$) and positive parity ($\Pi = +1$). If $\psi_T(\mathbf{R})$ is chosen to have a symmetry other than $(L, \Pi) = (0, +1)$, say (L', Π') symmetry, then the imaginary time propagation projects out the eigen state with (L', Π') symmetry that has the lowest energy. Said differently, the imaginary time propagation preserves the symmetry of $\psi_T(\mathbf{R})$ and acts in a subspace of the full Hilbert space.

It is instructive to compare the PIGS formalism discussed above with another imaginary time propagation based Monte Carlo technique, namely the diffusion Monte Carlo technique (for the purpose of the discussion that follows, the Green's function Monte Carlo technique behaves identically) [7]. The diffusion Monte Carlo approach utilizes, in addition to a trial function, a reference energy E_{ref} that is adjusted continually during the simulation. If equation (3) is multiplied by $\exp(E_{\text{ref}}\tau)$, then the right hand side is, except for an overall

permutation operators; the identical particle characteristics (bosons and/or fermions) of the N -particle wave function are instead encoded via the trial function, combined with the fixed- or released-node approach in the case of identical fermions [57, 58]. The PIGS algorithm, in contrast, explicitly anti-symmetrizes the paths if the system contains identical fermions. If the system contains identical bosons, explicit symmetrization operations are not needed provided the ground state of the system where the bosons are replaced by 'Boltzmann particles' is the same as that of the system containing bosons.

3. PIGS algorithm: general considerations

3.1. Basic concepts

This section rewrites equation (4) in a form amenable to evaluation by Monte Carlo techniques. The actual Monte Carlo implementation is discussed in section 4. Using $|\psi_\tau\rangle = \exp(-\tau\hat{H})|\psi_T\rangle$ and $\langle\psi_\tau| = \langle\psi_T|\exp(-\tau\hat{H})$, equation (4) reads

$$E_\tau = \frac{\langle\psi_T|\exp(-\tau\hat{H})\hat{H}\exp(-\tau\hat{H})|\psi_T\rangle}{\langle\psi_T|\exp(-\tau\hat{H})\exp(-\tau\hat{H})|\psi_T\rangle}. \quad (6)$$

The denominator is commonly denoted by $Z(\tau)$,

$$Z(\tau) = \langle\psi_T|\exp(-2\tau\hat{H})|\psi_T\rangle. \quad (7)$$

To obtain a prescription for evaluating the operators in the integrands, we project \hat{H} and $\exp(-\tau\hat{H})$ onto the position basis. Formally, this amounts to inserting the identity

$$\int_{\mathbf{R}} |\mathbf{R}\rangle \langle\mathbf{R}| d\mathbf{R} = \hat{1}, \quad (8)$$

where $\hat{1}$ denotes the unit operator, multiple times into equation (6),

$$E_\tau = \frac{\int_{\mathbf{R}} \int_{\mathbf{R}'} \int_{\mathbf{R}''} \int_{\mathbf{R}'''} F_{\text{aux}} d\mathbf{R} d\mathbf{R}' d\mathbf{R}'' d\mathbf{R}'''}{\int_{\mathbf{R}} \int_{\mathbf{R}'} \int_{\mathbf{R}''} \langle\psi_T|\mathbf{R}\rangle \langle\mathbf{R}|\exp(-\tau\hat{H})|\mathbf{R}'\rangle \langle\mathbf{R}'|\exp(-\tau\hat{H})|\mathbf{R}''\rangle \langle\mathbf{R}''|\psi_T\rangle d\mathbf{R} d\mathbf{R}' d\mathbf{R}''}, \quad (9)$$

\mathbf{R} -independent factor, independent of τ for sufficiently large τ and $E_{\text{ref}} = E_0$. This is the key idea behind the diffusion Monte Carlo approach. The accumulation of expectation values is started after the excited state contributions have decayed out and after E_{ref} has been adjusted such that $E_{\text{ref}} \approx E_0$. While the diffusion Monte Carlo and PIGS approaches share, as just discussed, similarities, the treatment of particle permutations differs notably if the system contains two or more identical fermions. The diffusion Monte Carlo algorithm does not explicitly apply sequences of two-particle

where

$$F_{\text{aux}} = \langle\psi_T|\mathbf{R}\rangle \langle\mathbf{R}|\exp(-\tau\hat{H})|\mathbf{R}'\rangle \langle\mathbf{R}'|\hat{H}|\mathbf{R}''\rangle \times \langle\mathbf{R}''|\exp(-\tau\hat{H})|\mathbf{R}'''\rangle \langle\mathbf{R}'''\psi_T\rangle. \quad (10)$$

We refer to $G(\mathbf{R}, \mathbf{R}'; \tau)$,

$$G(\mathbf{R}, \mathbf{R}'; \tau) = \langle\mathbf{R}|\exp(-\tau\hat{H})|\mathbf{R}'\rangle, \quad (11)$$

as the imaginary time evolution operator projected onto the position basis or, in short, as the imaginary time evolution operator or propagator. Using equation (11), we obtain

$$E_\tau = \frac{\int_{\mathbf{R}} \int_{\mathbf{R}'} \int_{\mathbf{R}''} \int_{\mathbf{R}'''} \psi_T(\mathbf{R}) G(\mathbf{R}, \mathbf{R}'; \tau) \langle \mathbf{R}' | \hat{H} | \mathbf{R}''' \rangle G(\mathbf{R}''', \mathbf{R}''; \tau) \psi_T(\mathbf{R}'') d\mathbf{R} d\mathbf{R}' d\mathbf{R}'' d\mathbf{R}'''}{\int_{\mathbf{R}} \int_{\mathbf{R}'} \int_{\mathbf{R}''} \psi_T(\mathbf{R}) G(\mathbf{R}, \mathbf{R}'; \tau) G(\mathbf{R}', \mathbf{R}''; \tau) \psi_T(\mathbf{R}'') d\mathbf{R} d\mathbf{R}' d\mathbf{R}''} \quad (12)$$

The normalization factor $Z(\tau)$, equation (7), plays a key role in the simulations. For example, if $Z(\tau)$ is known, instead of evaluating equation (12), one can calculate the energy expectation value E_τ directly,

$$E_\tau = -\frac{1}{Z(\tau)} \frac{\partial Z(\tau)}{\partial(2\tau)}. \quad (13)$$

Equations (12) and (13) generate two distinct energy estimators (see section 4.4 for details).

In the zero propagation time limit, i.e. for $\tau = 0$, $\hat{G} = \exp(-\tau\hat{H})$ becomes the identity operator. This implies that the propagator is simply a δ -function in position space,

$$G(\mathbf{R}, \mathbf{R}'; 0) = \delta(\mathbf{R} - \mathbf{R}'). \quad (14)$$

To propagate to finite imaginary time, one can solve the Bloch equation [6]

$$\frac{\partial \hat{G}}{\partial \tau} = -\hat{H}\hat{G}, \quad (15)$$

which is obtained by taking the partial derivative of the propagator with respect to τ . Equation (15) can be interpreted as a diffusion equation in the imaginary time τ . For the remainder of this section, we write the Hamiltonian \hat{H} as a sum of the kinetic energy operator \hat{K} and the potential energy operator \hat{V} . Moreover, we assume that all particles have the same mass m ; this assumption, which can be readily relaxed, simplifies the notation. If the kinetic energy operator \hat{K} is zero, the propagator can be readily solved for. Similarly, if the potential energy operator \hat{V} is zero, the propagator can also be solved for. In this case, the solution \hat{G}_0 corresponds to free particles diffusing in space (the subscript ‘0’ is used to indicate that the corresponding Hamiltonian contains only kinetic energy terms), i.e. \hat{G}_0 is a product of single-particle Gaussians,

$$G_0(\mathbf{R}, \mathbf{R}'; \tau) = (4\pi\lambda_m\tau)^{-3N/2} \exp\left(-\frac{(\mathbf{R} - \mathbf{R}')^2}{4\lambda_m\tau}\right), \quad (16)$$

where λ_m is equal to $\hbar^2/(2m)$. Equation (16) shows that the off-diagonal terms (terms for which $\mathbf{R} \neq \mathbf{R}'$) of G_0 , expressed in the position basis, are non-zero. This shows explicitly that the kinetic energy operator is non-local in position space. If \hat{V} and \hat{K} are both non-zero, then the propagator at finite τ is known only for a few selected problems such as non-interacting particles in a harmonic trap [59] and two particles with zero-range interactions [45, 60–63]. In general, the N -particle imaginary time evolution operator or propagator is unknown. If it was known, the problem would be ‘trivial’.

The PIGS algorithm is based on the idea of writing the imaginary time evolution operator for large τ as a product over imaginary time evolution operators for small imaginary

time steps,

$$\exp(-\tau\hat{H}) = [\exp(-\tau\hat{H}/n)]^n. \quad (17)$$

Using equation (17) in equation (11) and inserting the unit operator (equation (8)) $n - 1$ times, we obtain

$$\begin{aligned} G(\mathbf{R}_0, \mathbf{R}_n; \tau) &= \langle \mathbf{R}_0 | e^{-\tau\hat{H}/n} \underbrace{\int_{\mathbf{R}_1} |\mathbf{R}_1\rangle \langle \mathbf{R}_1| d\mathbf{R}_1}_{=\hat{1}} e^{-\tau\hat{H}/n} \\ &\times \underbrace{\int_{\mathbf{R}_2} |\mathbf{R}_2\rangle \langle \mathbf{R}_2| d\mathbf{R}_2}_{=\hat{1}} \times \cdots \\ &\times \underbrace{\int_{\mathbf{R}_{n-1}} |\mathbf{R}_{n-1}\rangle \langle \mathbf{R}_{n-1}| d\mathbf{R}_{n-1}}_{=\hat{1}} e^{-\tau\hat{H}/n} |\mathbf{R}_n\rangle \end{aligned} \quad (18)$$

or

$$\begin{aligned} G(\mathbf{R}_0, \mathbf{R}_n; \tau) &= \int_{\mathbf{R}_1} \int_{\mathbf{R}_2} \cdots \int_{\mathbf{R}_{n-1}} \\ &\times G(\mathbf{R}_0, \mathbf{R}_1; \tau/n) G(\mathbf{R}_1, \mathbf{R}_2; \tau/n) \times \cdots \\ &\times G(\mathbf{R}_{n-1}, \mathbf{R}_n; \tau/n) d\mathbf{R}_1 d\mathbf{R}_2 \cdots d\mathbf{R}_{n-1}. \end{aligned} \quad (19)$$

The problem of evaluating the propagator at the desired imaginary time τ has been converted to evaluating n propagators at τ/n and integrating over $n - 1$ (potentially high-dimensional) auxiliary coordinates $\mathbf{R}_1, \dots, \mathbf{R}_{n-1}$. The key points are that one can typically find a fairly accurate but approximate short-time propagator for finite n (see sections 3.2–3.5) and that the $n - 1$ associated ‘auxiliary’ integrations can be performed efficiently by Monte Carlo techniques (see section 4).

To simplify the notation, the product $G(\mathbf{R}, \mathbf{R}'; \tau) G(\mathbf{R}', \mathbf{R}''; \tau)$ in the denominator of equation (12) is rewritten as $G(\mathbf{R}_0, \mathbf{R}_n; \tau) G(\mathbf{R}_n, \mathbf{R}_{2n}; \tau)$. Each set of coordinates \mathbf{R}_j inserted in equation (18) is referred to as a ‘time slice’. There are three ‘special’ time slices: the initial time slice \mathbf{R}_0 , the middle time slice \mathbf{R}_n , and the final time slice \mathbf{R}_{2n} . The initial and middle time slices are connected by the propagator $G(\mathbf{R}_0, \mathbf{R}_n; \tau)$ and the middle and final time slices are connected by the propagator $G(\mathbf{R}_n, \mathbf{R}_{2n}; \tau)$. Since both propagators are rewritten by inserting $n - 1$ auxiliary time slices, the ‘expanded’ partition function contains a total of $2n + 1$ time slices. The position vector $\mathbf{r}_{k,j}$ of the k th particle in the set of coordinates \mathbf{R}_j is referred to as a ‘bead’. Thus, a single particle is represented by $2n + 1$ beads. The propagator that ‘connects’ two consecutive time slices is referred to as a ‘link’. The inverse temperature corresponding to this link is $\Delta\tau$, where $\Delta\tau = \tau/n$. The propagator that ‘connects’ two consecutive beads is referred to as a ‘single-particle link’. In addition, the set of all time slices $\{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$ is referred to as a configuration. The definitions are summarized in table 1.




Figure 1. Schematic world-line representation for a single particle with mass m in a one-dimensional harmonic trap with angular trapping frequency ω . Panels (a)–(d) show paths for 3, 5, 9, and 17 time slices, respectively.

Table 1. PIGS terminology used in this article. Columns 1–3 show the term, symbol, and explanation, respectively.

Bead	$\mathbf{r}_{k,j}$	A single coordinate of the k th particle at the j th imaginary time index
Time slice	\mathbf{R}_j	A set of beads at the j th imaginary time index
Configuration	$\{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$	The set of all time slices
Link	$G(\mathbf{R}_j, \mathbf{R}_{j+1}; \Delta\tau)$	The propagator connecting two consecutive time slices
Single-particle link	$G(\mathbf{r}_{k,j}, \mathbf{r}_{k,j+1}; \Delta\tau)$	The propagator connecting two consecutive beads

Figure 1 shows the world-line representation of a single particle in a one-dimensional harmonic trap. World lines move in position space (x -axis) and imaginary time (y -axis). Figures 1(a)–(d) show paths for $n = 3, 5, 9,$ and 17 beads, respectively. As n increases, the path is resolved in more detail (each link is more accurate) and observables calculated based on the sampled paths become more accurate.

Figure 2 depicts a single particle (the Hamiltonian only contains the kinetic energy term) in two-dimensional space [64]. Two consecutive beads (circles in figure 2) are connected by a single-particle link (wiggly line in figure 2). The kinetic energy is ‘carried’ by the propagators represented by the links. The expression for the propagator in free space reads (equation (16) for a single particle with position vector $\mathbf{r}_{1,j} = (x_{1,j}, y_{1,j})$)

$$G_0(\mathbf{r}_{1,j}, \mathbf{r}_{1,j+1}; \tau) = (4\pi\lambda_m\tau)^{-1} \exp\left(-\frac{(\mathbf{r}_{1,j} - \mathbf{r}_{1,j+1})^2}{4\lambda_m\tau}\right). \quad (20)$$




Figure 2. Isomorphism between path integrals for a single free particle in two-dimensional space and classical particles connected by springs. In the path integral interpretation, the circles and wiggly lines depict the beads and (single-particle) links of a single particle, respectively. In the classical mechanics formulation, the circles and wiggly lines depict particles and springs, respectively. The position vector of the particle is denoted by $\mathbf{r}_{1,j}$, where j indicates the imaginary time index ($j = 1, 2,$ and 3).

The action S [6],

$$S = -\ln[G_0(\mathbf{r}_{1,j}, \mathbf{r}_{1,j+1}; \tau)], \quad (21)$$

of the single-particle link that connects the beads labeled $\mathbf{r}_{1,j}$ and $\mathbf{r}_{1,j+1}$ reads

$$S = \ln(4\pi\lambda_m\tau) + \frac{(\mathbf{r}_{1,j} - \mathbf{r}_{1,j+1})^2}{4\lambda_m\tau}. \quad (22)$$

It can be seen that the action S has the same form as that of a ‘spring potential’ $V_s(\mathbf{r}_{1,j} - \mathbf{r}_{1,j+1})$ for two classical particles with position vectors $\mathbf{r}_{1,j}$ and $\mathbf{r}_{1,j+1}$ connected via Hooke’s law. The propagator can thus be interpreted as being proportional to the Boltzmann factor $\exp(-\tau V_s)$ of a classical system of springs. Note that $\mathbf{r}_{1,j}$ and $\mathbf{r}_{1,j+1}$ in equations (20) and (22) correspond to position vectors of consecutive beads for one single particle while $\mathbf{r}_{1,j}$ and $\mathbf{r}_{1,j+1}$ in the classical isomorphism correspond to position vectors of two different particles.

In addition to the propagators $G(\mathbf{R}_0, \mathbf{R}_n; \tau)$ and $G(\mathbf{R}_n, \mathbf{R}_{2n}; \tau)$, the partition function contains the trial functions (or ‘weights’) $\psi_T(\mathbf{R}_0)$ and $\psi_T(\mathbf{R}_{2n})$. For the single two-dimensional particle in free space, $\psi_T(\mathbf{r}_0)$ and $\psi_T(\mathbf{r}_{2n})$ can be interpreted as potentials that are felt by the first and last particle of the chain of classical particles. Thus, we can interpret the PIGS simulation of a single particle as a simulation of a chain of classical particles connected by springs (or a polymer with nearest neighbor interactions) and two external forces that act on the particles at the ends of the chain. The stiffness of the springs is determined by $1/\Delta\tau$, i.e. by the inverse of the imaginary time step associated with the links.

The following two sections introduce two different approaches for approximating the short-time propagator, namely the Trotter formula and the pair product approximation.

3.2. Trotter formula

One way to approximate the short-time propagator is to use the Trotter formula [65]. For sufficiently small time steps $\Delta\tau$, the kinetic energy contribution \hat{K} and the potential energy contribution \hat{V} to the propagator can be split,

$$\exp[-\Delta\tau(\hat{K} + \hat{V})] = \exp(-\Delta\tau\hat{K})\exp(-\Delta\tau\hat{V}) + \mathcal{O}(\Delta\tau^2), \quad (23)$$

where the notation $\mathcal{O}(\Delta\tau^2)$ indicates that the leading-order error scales, in general, as $\Delta\tau$ to the power of 2. More specifically, by Taylor expanding the exponentials, one can

prove that the leading-order error is $\Delta\tau^2[\hat{K}, \hat{V}]/2$, where $[\hat{K}, \hat{V}]$ is the commutator between \hat{K} and \hat{V} , $[\hat{K}, \hat{V}] = \hat{K}\hat{V} - \hat{V}\hat{K}$. In the $\Delta\tau \rightarrow 0$ limit (this corresponds to the insertion of infinitely many time slices, i.e. the $n \rightarrow \infty$ limit), the Trotter formula becomes exact. Since n cannot be infinitely large in practice, one performs calculations for different n and extrapolates the observables of interest to the infinite n limit.

Importantly, the Trotter formula can be extended to higher orders. We can readily adopt a $\mathcal{O}(\Delta\tau^3)$ scheme by further splitting the kinetic energy term or the potential energy term,

$$\begin{aligned} \exp[-\Delta\tau(\hat{K} + \hat{V})] &= \exp(-\Delta\tau\hat{K}/2)\exp(-\Delta\tau\hat{V}) \\ &\times \exp(-\Delta\tau\hat{K}/2) + \mathcal{O}(\Delta\tau^3) \end{aligned} \quad (24)$$

or

$$\begin{aligned} \exp[-\Delta\tau(\hat{K} + \hat{V})] &= \exp(-\Delta\tau\hat{V}/2)\exp(-\Delta\tau\hat{K}) \\ &\times \exp(-\Delta\tau\hat{V}/2) + \mathcal{O}(\Delta\tau^3). \end{aligned} \quad (25)$$

In practice, equation (25), which is accurate up to second order [the error is $\mathcal{O}(\Delta\tau^3)$], is easier to use than equation (24). In position space, equation (25) reads

$$\begin{aligned} G(\mathbf{R}, \mathbf{R}'; \Delta\tau) &= \exp[-\Delta\tau V(\mathbf{R})/2]\exp[-\Delta\tau V(\mathbf{R}')/2] \\ &\times G_0(\mathbf{R}, \mathbf{R}'; \Delta\tau) + \mathcal{O}(\Delta\tau^3), \end{aligned} \quad (26)$$

where G_0 (see equation (16)) is the propagator that accounts for the kinetic energy term.

One can reach successively higher accuracy by the repeated use of the Baker–Campbell–Hausdorff formula (see, e.g., [66])

$$e^{\epsilon\hat{A}}e^{\epsilon\hat{B}} = e^{\hat{C}}, \quad (27)$$

where

$$\begin{aligned} \hat{C} &= (\hat{A} + \hat{B})\epsilon + \frac{1}{2}[\hat{A}, \hat{B}]\epsilon^2 + \frac{1}{12}([\hat{A}, [\hat{A}, \hat{B}]] \\ &+ [\hat{B}, [\hat{B}, \hat{A}]])\epsilon^3 \\ &- \frac{1}{24}[\hat{B}, [\hat{A}, [\hat{A}, \hat{B}]]]\epsilon^4 + \mathcal{O}(\epsilon^5). \end{aligned} \quad (28)$$

Using equations (27) and (28) twice, we obtain [67]

$$e^{\epsilon\hat{B}}e^{\epsilon\hat{A}}e^{\epsilon\hat{B}} = e^{\hat{D}}, \quad (29)$$

where

$$\begin{aligned} \hat{D} &= (\hat{A} + 2\hat{B})\epsilon - \frac{1}{6}\epsilon^3[\hat{A}, [\hat{B}, \hat{A}]] \\ &+ \frac{1}{6}\epsilon^3[\hat{B}, [\hat{A}, \hat{B}]] + \mathcal{O}(\epsilon^5). \end{aligned} \quad (30)$$

Applying equations (29) and (30) twice to

$$\begin{aligned} &\exp\left(-\Delta\tau\frac{\hat{V}}{6}\right)\exp\left(-\Delta\tau\frac{\hat{K}}{2}\right)\exp\left(-\Delta\tau\frac{2\tilde{V}}{3}\right) \\ &\times \exp\left(-\Delta\tau\frac{\hat{K}}{2}\right)\exp\left(-\Delta\tau\frac{\hat{V}}{6}\right), \end{aligned} \quad (31)$$




Figure 3. Isomorphism between path integrals for two interacting particles in two-dimensional space and classical particles connected by springs. In the path integral interpretation, the circles and wiggly lines depict the beads and links of a single particle, respectively, and the dotted lines depict the two-body interaction between beads with the same imaginary time index. In the classical mechanics interpretation, the circles and wiggly lines depict particles and springs, respectively, and the dotted lines depict the two-body interaction between selected classical particles. The position vector of the k th particle at the j th imaginary time index j is denoted by $\mathbf{r}_{k,j}$.

we can check that the fourth-order factorization [67]

$$\begin{aligned} \exp[-\Delta\tau(\hat{K} + \hat{V})] &= \exp\left(-\Delta\tau\frac{\hat{V}}{6}\right)\exp\left(-\Delta\tau\frac{\hat{K}}{2}\right) \\ &\times \exp\left(-\Delta\tau\frac{2\tilde{V}}{3}\right)\exp\left(-\Delta\tau\frac{\hat{K}}{2}\right)\exp\left(-\Delta\tau\frac{\hat{V}}{6}\right) \\ &+ \mathcal{O}(\Delta\tau^5), \end{aligned} \quad (32)$$

where \tilde{V} is given by $\hat{V} + \Delta\tau^2[\hat{V}, [\hat{K}, \hat{V}]]/48$, holds. The term $[\hat{V}, [\hat{K}, \hat{V}]]$, in position space, can be simplified to $(\hbar^2/m)\sum_{i=1}^N|\nabla_i V|^2$, where the gradient ∇_i in three spatial dimensions is given by

$$\nabla_i = \hat{x}_i \frac{\partial}{\partial x_i} + \hat{y}_i \frac{\partial}{\partial y_i} + \hat{z}_i \frac{\partial}{\partial z_i}, \quad (33)$$

with \hat{x}_i , \hat{y}_i , and \hat{z}_i denoting unit vectors that point in the x_i , y_i , and z_i directions, respectively. The term $|\nabla_i V|^2$ corresponds to the square of the force on the i th particle. Care needs to be taken in evaluating the derivatives, since V usually contains a double sum over two-body potentials or even a triple sum over three-body potentials. In most cases, the evaluation of the force terms cannot be simplified analytically, implying that the evaluation of double commutators involves double or triple sums over the total number of particles. This makes the numerical evaluation comparatively expensive. Note that the exponentials in equation (32) that contain the potential energy can have different numerical factors. In addition, equation (32) is not the only form of the fourth-order factorization [67, 68].

Using the Trotter formula, the isomorphism between a single quantum particle in free space and the classical spring system can be extended to multiple quantum particles with interactions. Figure 3 depicts two interacting particles in two-dimensional space (it is assumed that the particles do not feel a single-particle potential).

In this case, the propagator for the link that connects the beads labeled by $\mathbf{r}_{1,j}$, $\mathbf{r}_{1,j+1}$, $\mathbf{r}_{2,j}$, and $\mathbf{r}_{2,j+1}$ reads (see

equation (25))

$$\begin{aligned} G(\{\mathbf{r}_{1,j}, \mathbf{r}_{2,j}\}, \{\mathbf{r}_{1,j+1}, \mathbf{r}_{2,j+1}\}; \Delta\tau) \\ = e^{-\Delta\tau V_{2b}(\mathbf{r}_{1,j}, \mathbf{r}_{2,j})/2} e^{-\Delta\tau V_{2b}(\mathbf{r}_{1,j+1}, \mathbf{r}_{2,j+1})/2} \\ \times G_0(\mathbf{r}_{1,j}, \mathbf{r}_{1,j+1}; \Delta\tau) G_0(\mathbf{r}_{2,j}, \mathbf{r}_{2,j+1}; \Delta\tau), \end{aligned} \quad (34)$$

where V_{2b} denotes the two-body interaction potential between particles 1 and 2 and $G_0(\mathbf{r}_{k,j}, \mathbf{r}_{k,j+1}; \Delta\tau)$ the single-particle propagator of the k th particle (see equation (20)). As in figure 2, two consecutive beads for the same particle (e.g., the circles labeled by $\mathbf{r}_{1,j}$ and $\mathbf{r}_{1,j+1}$ in figure 3) are connected by single-particle links (wiggly lines in figure 3) that represent the propagators G_0 . Since the two-body interaction potential (dotted lines in figure 3) is diagonal in position space (see the exponentials on the right hand side of equation (34)), it connects beads of different particles with the same index, i.e. it connects $\mathbf{r}_{1,1}$ and $\mathbf{r}_{2,1}$, $\mathbf{r}_{1,2}$ and $\mathbf{r}_{2,2}$, and $\mathbf{r}_{1,3}$ and $\mathbf{r}_{2,3}$ (or, in general, it connects $\mathbf{r}_{k,j}$ and $\mathbf{r}_{k+1,j}$). Each particle in the PIGS simulation corresponds to $2n + 1$ classical particles connected by springs. Classical particles associated with different chains interact only if they have the same bead index.

3.3. Pair product approximation

To introduce the pair product approximation, we assume for simplicity that the potential energy operator \hat{V} can be written as a sum of two-body terms,

$$V(\mathbf{R}) = \sum_{k=1}^{N-1} \sum_{l=k+1}^N V_{2b}(\mathbf{r}_k - \mathbf{r}_l), \quad (35)$$

i.e. we assume for now that single-particle and three- and higher-body forces are absent. Under these assumptions, the short-time propagator can be evaluated using the pair product approximation [6]. It is convenient to define the two-body kinetic energy operator \hat{K}_{kl} for the k th and l th particle in position space,

$$\hat{K}_{kl} = -\frac{\hbar^2}{m} \nabla_{\mathbf{r}_k - \mathbf{r}_l}^2. \quad (36)$$

The relative non-interacting and interacting two-body Hamiltonian are \hat{K}_{kl} and $\hat{K}_{kl} + \hat{V}_{2b}(\mathbf{r}_k - \mathbf{r}_l)$, respectively. The pair product approximation considers two-body correlations explicitly, but not higher-body correlations, and writes the many-body propagator as a product over single-particle propagators and two-body propagators,

$$\begin{aligned} G(\mathbf{R}, \mathbf{R}'; \Delta\tau) \approx G_0(\mathbf{R}, \mathbf{R}'; \Delta\tau) \\ \left(\prod_{k < l}^N \bar{G}^{\text{rel}}(\mathbf{r}_k - \mathbf{r}_l, \mathbf{r}'_k - \mathbf{r}'_l; \Delta\tau) \right), \end{aligned} \quad (37)$$

where \bar{G}^{rel} ,

$$\begin{aligned} \bar{G}^{\text{rel}}(\mathbf{r}_k - \mathbf{r}_l, \mathbf{r}'_k - \mathbf{r}'_l; \Delta\tau) \\ = \frac{\langle \mathbf{r}_k - \mathbf{r}_l | \exp[-\Delta\tau(\hat{K}_{kl} + \hat{V}_{2b}(\mathbf{r}_k - \mathbf{r}_l))] | \mathbf{r}'_k - \mathbf{r}'_l \rangle}{\langle \mathbf{r}_k - \mathbf{r}_l | \exp(-\Delta\tau \hat{K}_{kl}) | \mathbf{r}'_k - \mathbf{r}'_l \rangle}, \end{aligned} \quad (38)$$

is the reduced pair propagator. The denominator of the reduced pair propagator coincides with the known relative non-interacting two-body propagator. Thus the only ‘non-

trivial input’ is the relative propagator of the interacting two-body system. One can readily see that the pair product approximation is exact for two particles for any propagation time because the center-of-mass and relative degrees of freedom separate in this case. In some cases such as for the two-body zero-range interaction potential, the exact reduced pair propagator is known analytically [60–63]. In other cases such as for the two-body hardcore potential, an approximate reduced pair propagator is known analytically in closed form [69, 70]. If the reduced propagator is not known analytically, one can perform a partial wave decomposition and obtain a numerical representation of the reduced two-body propagator [6].

In dilute gases or weakly bound droplets, the interparticle spacing is typically so large that two-body collisions dominate over three- and higher-body collisions. The leading-order error of the pair product approximation is determined by the importance of three-body correlations. For two-component equal-mass Fermi gases with two-body zero-range interactions, three-body correlations are suppressed by the Pauli exclusion principle. For this system, we found that the pair product approximation provides an extremely good description of the propagator. Specifically, we obtain accurate simulation results for a small number of beads (see section 5 for details). For bosons, in contrast, three-body correlations can be significant. As a consequence, the pair product approximation is not as efficient as for two-component fermions and simulations typically employ a large number of beads (‘large’ in this context means about two orders of magnitude more number of beads as in the simulations for fermions [45]).

To illustrate the pair product approximation, we cannot use the classical isomorphism because the kinetic and potential energy contributions are mixed. One needs to evaluate the single-particle propagator, which can be represented by springs as in figures 2 and 3. However, one also needs to evaluate the reduced two-body propagator, which connects two consecutive beads of one particle’s path with the same consecutive beads of another particle’s path. These ‘four-bead connections’ do not have a simple classical analog.

3.4. Comparison of the two approximations

This section discusses the advantages and disadvantages of approximating the short-time propagator with the help of the Trotter formula and the pair product approximation.

In the Trotter formula based scheme, the kinetic and the potential energy terms are treated separately. Inserting the identity $\int_{\mathbf{R}} |\mathbf{R}\rangle \langle \mathbf{R}| d\mathbf{R} = \hat{1}$, equation (8), multiple times into equations (25) or (32), it can be seen that the potential energy is diagonal in position space. This means that one can directly evaluate the potential energy term at each time slice. The kinetic energy term contains off-diagonal terms and needs to be evaluated at each link instead of at each time slice. Nevertheless, since the kinetic energy term corresponds to a simple Gaussian, the sampling of the kinetic energy piece of the propagator can, in general, be performed efficiently (see section 4.3.2 for details).

Even though the Trotter formula can formally be generalized to expressions that are accurate to order $\Delta\tau^5$, $\Delta\tau^6$, ... [71, 72], many of these expressions are of limited use in practice because they contain either commutators that involve rather complicated expressions or terms that correspond to negative imaginary time, which are not normalizable. There exists a multi-product expansion for the propagator [72]; however, applications thereof are still rare [73]. Thus present-day algorithms mostly employ Trotter formula based decompositions that are accurate to order $\Delta\tau^4$.

In the pair product approximation, the two-body reduced propagator contains kinetic energy and potential energy contributions. This means that the reduced two-body propagator has to be evaluated at each link. Because the reduced two-body propagator is, in general, not a simple Gaussian, the sampling is typically less efficient than in the case where the Trotter formula is used. Furthermore, as discussed in section 4.6, the evaluation of the permutations is computationally more involved.

Our discussion of the pair product approximation assumed that the potential energy can be written as a sum over two-body terms. If the Hamiltonian contains one-, three-, or higher-body potential energy terms, one can include them by combining the Trotter formula and the pair product approximation. To this end, one first splits the propagator into two terms using the Trotter formula. The first term contains the kinetic energy operator and the two-body interactions while the second term contains all other potential terms. One then applies the pair product approximation to the first term. In this approach, it is most convenient to use the second-order Trotter formula for two reasons. First, if a higher-order Trotter formula was used, one would need to evaluate the commutator between the one-, three- and higher-body potential terms and the two-body potential terms. For the two-body zero-range interactions considered in section 5 this is a rather challenging task. Second, both the second-order Trotter formula and the pair product approximation yield errors for the energy that scale quadratically with the time step $\Delta\tau$. While the error in the pair product approximation tends to be smaller than that associated with the pair product approximation, ultimately it is the scaling with $\Delta\tau$ that determines the accuracy and use of the fourth-order Trotter formula typically leads only to a small overall improvement.

From our perspective, the pair product approximation has one key advantage: it can deal with a class of two-body potentials that the Trotter formula based scheme cannot deal with (at least no such treatment is known to us). For example, the two-body hardcore and zero-range potentials contain infinities and can thus not be treated by the Trotter formula based scheme. However, the infinities can, as discussed in the next section exemplarily for the two-body zero-range potential, be dealt with analytically in the pair product approximation.

3.5. Propagator for two-body zero-range interactions

As alluded to in the previous section, two-body zero-range interactions can be incorporated into continuum Monte Carlo

simulations through the pair product approximation [42, 44–46], which employs the relative propagator \bar{G}^{rel} (see equations (37) and (38)). In what follows, we limit our discussion to three spatial dimensions. To determine \bar{G}^{rel} , one considers the relative Hamiltonian H^{rel} for two particles interacting through the regularized zero-range Fermi–Huang pseudopotential $V_{\text{F}}(r)$ [74] in free space,

$$H^{\text{rel}} = -\frac{\hbar^2}{2\mu}\nabla_{\mathbf{r}}^2 + V_{\text{F}}(r), \quad (39)$$

where

$$V_{\text{F}}(r) = \frac{2\pi\hbar^2 a_s}{\mu} \delta^{(3)}(\mathbf{r}) \frac{\partial}{\partial r} r. \quad (40)$$

Here, μ denotes the two-body reduced mass, \mathbf{r} the interparticle distance vector, and a_s the s -wave scattering length. The regularization operator $(\partial/\partial r)r$ in equation (40) ensures that the Hamiltonian is well-behaved. Without this operator, the two-body coupling constant would have to be renormalized. With the regulator, however, the coupling strength is uniquely determined and given by $2\pi\hbar^2 a_s/\mu$.

The reduced (or normalized) relative propagator corresponding to the Hamiltonian given in equation (39) reads [61, 63]

$$\begin{aligned} \bar{G}^{\text{rel}}(\mathbf{r}, \mathbf{r}'; \tau) &= 1 + \frac{\hbar^2\tau}{\mu rr'} \exp\left(-\frac{\mu rr'(1 + \cos\theta)}{\hbar^2\tau}\right) \\ &\times \left(1 + \frac{\hbar}{a_s} \sqrt{\frac{\pi\tau}{2\mu}} \text{erfc}(v) \exp(v^2)\right), \end{aligned} \quad (41)$$

where $\cos\theta = \mathbf{r} \cdot \mathbf{r}'/(rr')$ and $v = [r + r' - \tau\hbar^2/(\mu a_s)]/\sqrt{2\tau\hbar^2/\mu}$. For $|a_s| = \infty$, the length scale a_s drops out of the expression for the propagator and equation (41) simplifies to

$$\bar{G}^{\text{rel}}(\mathbf{r}, \mathbf{r}'; \tau) = 1 + \frac{\hbar^2\tau}{\mu rr'} \exp\left(-\frac{\mu rr'(1 + \cos\theta)}{\hbar^2\tau}\right). \quad (42)$$

Importantly, the reduced relative propagator diverges when r or r' approach zero. These divergencies have implications for the Monte Carlo sampling of the paths. As discussed in detail in section 4.3.3, moves have to be designed carefully such that detailed balance is fulfilled. For example, while \bar{G}^{rel} diverges for r and $r' \rightarrow 0$, the probability to find two particles at vanishing interparticle distance does not diverge. The treatment of systems with two-body hardcore interactions is similar in spirit to that detailed here for two-body zero-range interactions.

Adding the spherically symmetric harmonic confining potential $V_{\text{trap}}(r) = \mu\omega^2 r^2/2$ for the relative degrees of freedom to the Hamiltonian H^{rel} given in equation (39) and assuming that a_s is infinitely large, the reduced relative propagator $\bar{G}^{\text{rel}}(\mathbf{r}, \mathbf{r}'; \tau)$ reads

$$\begin{aligned} \bar{G}^{\text{rel}}(\mathbf{r}, \mathbf{r}'; \tau) &= 1 \\ &+ \frac{2a_{\text{ho}}^2}{rr'} \sinh(\tau\hbar\omega) \exp\left(-\frac{rr'(1 + \cos\theta)}{2a_{\text{ho}}^2 \sinh(\tau\hbar\omega)}\right), \end{aligned} \quad (43)$$

where $a_{\text{ho}} = \sqrt{\hbar/(m\omega)}$. In the limit that the angular trapping frequency ω goes to zero, equation (43) reduces to

Table 2. Definitions of Monte Carlo sampling terminology used in this article. Columns 1–3 show the symbol, name, and related equation number, respectively. The configuration \mathbf{x} is defined as $\mathbf{x} = \{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$.

$\pi(\mathbf{x})$	Probability distribution	Equation (45)
$p(\mathbf{x})$	Probability density function	Equation (48)
$w(\mathbf{x})$	Weight function (observable specific)	Equations (46) and (47); section 4.4
$P(\mathbf{x} \rightarrow \mathbf{x}')$	Transition probability	Equation (49)
$\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$	Proposal distribution (selected by simulator)	Around equations (49) and (50)
$\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$	Acceptance distribution	Equation (50)
$\psi_T(\mathbf{R})$	Trial function (selected by simulator)	Equation (1)

equation (42). Expression (43) is used in section 5 to treat harmonically trapped two-component Fermi gases with two-body zero-range interactions at unitarity using the pair product approximation.

4. Monte Carlo techniques and the PIGS algorithm

Throughout this section we assume that the trial function $\psi_T(\mathbf{R})$ is given and that its value can be determined for any set of coordinates \mathbf{R} . The functional form of $\psi_T(\mathbf{R})$ depends sensitively on the system under study. The choice of $\psi_T(\mathbf{R})$ and the dependence of the PIGS results on $\psi_T(\mathbf{R})$ will be discussed in section 5 for harmonically trapped two-component Fermi gases.

4.1. General sampling scheme: importance sampling

Equation (19) writes the long-time propagator as a high-dimensional integral over a product of short-time propagators. This implies that the evaluation of the normalization factor $Z(\tau)$ is a high-dimensional integral. This section discusses the Monte Carlo sampling of this high-dimensional integral over $\{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$ (there are $3 \times (2n + 1) \times N$ independent coordinates if we are considering three spatial dimensions). To proceed, we write $Z(\tau)$ explicitly in terms of the short-time propagator,

$$Z(\tau) = \int_{\mathbf{R}_0} \dots \int_{\mathbf{R}_{2n}} \pi(\mathbf{R}_0, \dots, \mathbf{R}_{2n}) d\mathbf{R}_0 \dots d\mathbf{R}_{2n}, \quad (44)$$

where

$$\begin{aligned} \pi(\mathbf{R}_0, \dots, \mathbf{R}_{2n}) &= \psi_T(\mathbf{R}_0) G(\mathbf{R}_0, \mathbf{R}_1; \Delta\tau) G(\mathbf{R}_1, \mathbf{R}_2; \Delta\tau) \\ &\times \dots \times G(\mathbf{R}_{2n-1}, \mathbf{R}_{2n}; \Delta\tau) \psi_T(\mathbf{R}_{2n}). \end{aligned} \quad (45)$$

To simplify the notation, we denote the configuration $\{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$ by \mathbf{x} and the probability distribution $\pi(\mathbf{R}_0, \dots, \mathbf{R}_{2n})$ by $\pi(\mathbf{x})$. The notation of these and other quantities is summarized in tables 1 and 2. The expectation value

$\langle O \rangle$ of an arbitrary observable O can be written as

$$\langle O \rangle = \frac{\int_{\mathbf{x}} w(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}}{\int_{\mathbf{x}} \pi(\mathbf{x}) d\mathbf{x}}, \quad (46)$$

where the integration goes over $3 \times (2n + 1) \times N$ coordinates and the weight function $w(\mathbf{x})$ needs to be determined, as will be discussed in section 4.4, for each observable. To see the structure of $\langle O \rangle$ more clearly, we rewrite equation (46) as

$$\langle O \rangle = \int_{\mathbf{x}} w(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}, \quad (47)$$

where the probability density function $p(\mathbf{x})$ is defined as

$$p(\mathbf{x}) = \frac{\pi(\mathbf{x})}{\int_{\mathbf{x}'} \pi(\mathbf{x}') d\mathbf{x}'}. \quad (48)$$

In contrast to the probability distribution $\pi(\mathbf{x})$, the probability density function $p(\mathbf{x})$ is normalized; $w(\mathbf{x})$ and $p(\mathbf{x})$ represent the weight contributed to the observable by the configuration \mathbf{x} and the normalized probability to be in the configuration \mathbf{x} , respectively. Equation (47) provides the basis of importance sampling: configurations are not blindly distributed uniformly in space but instead are distributed according to $p(\mathbf{x})$. The advantage of importance sampling is that most computer time is used to sample configurations that are physically relevant and little time to sample configurations that do not contribute significantly to $\langle O \rangle$.

The general idea of the PIGS algorithm is to generate configurations \mathbf{x} according to $p(\mathbf{x})$ and to use the generated configurations to accumulate the weight functions $w(\mathbf{x})$ for a set of observables. Thus, it is crucial to have correct and efficient sampling schemes that explore the full configuration space with a relatively high acceptance ratio and without getting stuck around a local minimum. Section 4.2 reviews the basics of selected Monte Carlo methods, which are then used in the subsequent sections.

4.2. Some background on Monte Carlo methods

This section discusses how to update or generate configurations using the Metropolis algorithm. A Markov process is uniquely defined by the transition probability $P(\mathbf{x} \rightarrow \mathbf{x}')$ to go from configuration \mathbf{x} to configuration \mathbf{x}' . The Metropolis algorithm satisfies the detailed balance condition [7]

$$\pi(\mathbf{x}) P(\mathbf{x} \rightarrow \mathbf{x}') = \pi(\mathbf{x}') P(\mathbf{x}' \rightarrow \mathbf{x}), \quad (49)$$

which states that the flow of probability from \mathbf{x} to \mathbf{x}' is equal to that from \mathbf{x}' to \mathbf{x} . This means that there is no net flow of probability. The Metropolis algorithm needs to ensure ergodicity of the Markov process. If the process is ergodic, the Markov chain (i) returns to any previously generated configuration \mathbf{x} after a sufficiently long simulation time and (ii) is not periodic (a Markov chain of $\{\mathbf{x}, \mathbf{x}', \mathbf{x}, \mathbf{x}', \dots\}$, e.g., is periodic). The ergodicity ensures that the probability distribution $\pi(\mathbf{x})$ gets sampled fully. For example, as discussed in [75], if we use the traditional scheme of treating the permutations [6], for a two-component Fermi gas with zero-range interactions, the Markov process ends up with a configuration in which all particles sit on top of each other and

the configuration never returns to the original configuration \mathbf{x} . This means that ergodicity is violated and that the Markov process does not generate samples according to $p(\mathbf{x})$. This renders the sampled configurations meaningless. We note, however, that while the detailed balance condition together with the ergodicity guarantees that the equilibrium distribution coincides with the desired probability distribution $\pi(\mathbf{x})$, there exist other Monte Carlo methods that do not satisfy the detailed balance condition but yield an equilibrium distribution that coincides with the desired probability distribution $\pi(\mathbf{x})$ [76].

The Metropolis algorithm consists of two steps [7]: (i) the generation of a proposed configuration (or move) and (ii) the acceptance or rejection of the proposed configuration (or move). The combination of (i) and (ii) leads to a new configuration. Starting from the configuration \mathbf{x} , we propose a new configuration \mathbf{x}' according to a proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ and accept (the new configuration would be \mathbf{x}') or reject the new configuration (the new configuration would be \mathbf{x}) according to the acceptance distribution $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$. This implies that $P(\mathbf{x} \rightarrow \mathbf{x}')$ is given by $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$. The Metropolis algorithm chooses $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ such that [7]

$$\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}') = \min\left(1, \frac{\pi(\mathbf{x}')\mathcal{G}(\mathbf{x}' \rightarrow \mathbf{x})}{\pi(\mathbf{x})\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')}\right). \quad (50)$$

We verify that the detailed balance condition (equation (49)) is satisfied in the following. If $\pi(\mathbf{x}')\mathcal{G}(\mathbf{x}' \rightarrow \mathbf{x})$ is smaller than $\pi(\mathbf{x})\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$, we obtain from equation (50) that

$$P(\mathbf{x} \rightarrow \mathbf{x}') = \frac{\pi(\mathbf{x}')\mathcal{G}(\mathbf{x}' \rightarrow \mathbf{x})}{\pi(\mathbf{x})} \quad (51)$$

and

$$P(\mathbf{x}' \rightarrow \mathbf{x}) = \mathcal{G}(\mathbf{x}' \rightarrow \mathbf{x}). \quad (52)$$

Plugging the right-hand sides of equations (51) and (52) into equation (49), we confirm that equation (49) holds. If $\pi(\mathbf{x}')\mathcal{G}(\mathbf{x}' \rightarrow \mathbf{x})$ is larger or equal to $\pi(\mathbf{x})\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$, it can be checked similarly that equation (49) holds. Thus, we have shown that detailed balance is fulfilled.

A key task is to design proposal distributions $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ that ensure the complete and, ideally, efficient exploration of the entire configuration space. In most cases, efficient simulation schemes are achieved if more than one proposal distribution (and hence type of move) is utilized. As discussed more in the next section, the proposal distribution might be designed based on the knowledge of the non-interacting system (see, e.g., section 4.3.2) or based on knowledge of certain limiting behaviors of the interacting system (see, e.g., section 4.3.3).

In practice, the acceptance ratio A , i.e. one minus the fraction of rejected moves, should be monitored (note, A is a number and not an \mathbf{x} - and \mathbf{x}' -dependent function). The acceptance ratio A for Metropolis sampling is different from the acceptance ratio encountered in the rejection sampling. In the rejection sampling, a rejected configuration does not lead to a new configuration. In the Metropolis sampling, in contrast, a rejected configuration does lead to a new configuration. When a configuration is rejected, the old configuration becomes the new configuration. For most of the updates (i.e. the generation of

proposed new configurations), the acceptance ratio should not be too large and not be too small. A high acceptance ratio typically implies that the deviation between the old and new configurations is, on average, small. This means that the configuration space is explored comparatively slowly. A low acceptance ratio, in contrast, means that the Markov chain contains many identical configurations; again, typically this means that the configuration space is explored comparatively slowly. Both cases can result in large correlations of the sample and should be avoided. As a rule of thumb, the acceptance ratio should lie roughly between 30% and 50% [77].

4.3. Moves

The previous section outlined the basics of the Metropolis algorithm. This section discusses the PIGS moves that are used to update the configurations. For all moves, the proposed new configuration \mathbf{x}' is chosen based on the proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ and accepted/rejected based on the acceptance distribution $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$. Once $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ is specified, $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ follows from equation (50). This section discusses three different moves. The ‘naive move’ and the ‘wiggle move’ are ‘all purpose’ moves, which have proven to be useful for nearly all systems. In some cases, the use of these two types of moves alone does not lead to an efficient (or even correct) exploration of the configuration space. For systems with two-body zero-range interactions, e.g., the ‘pair distance move’ is needed. In general, the simulator decides on the frequency with which the individual moves are used. The optimal ratio can be found empirically from the performance of the simulation itself or through the implementation of some sort of learning algorithm.

The list of moves discussed below does not include a ‘permutation move’, i.e. the stochastic sampling of the permutations is not discussed. The reason for this is twofold. If the system contains identical bosons, the ground state wave function is typically identical to that of Boltzmann particles, eliminating the need for an explicit symmetrization. If the system contains identical fermions, we employ the on-the-fly anti-symmetrization scheme discussed in section 4.6, which is particularly useful if two-body zero-range interactions are employed.

The moves employed in the PIGS algorithm have much in common with the moves employed in the finite-temperature path integral Monte Carlo algorithm. As already discussed, one difference is that the PIGS algorithm contains the trial function ψ_T while the finite-temperature path integral Monte Carlo algorithm does not. Quite generally, whether a move depends on the trial function ψ_T or not depends on whether or not the beginning time slice \mathbf{R}_0 or the ending time slice \mathbf{R}_{2n} are being updated. In the implementations of the moves discussed below, the wiggle move does not depend on ψ_T , and the naive move and the pair move may depend on ψ_T (it depends on whether or not the randomly selected time slice to be updated is the 0th or 2nth time slice). It should be noted, though, that the moves can, in principle, be implemented in a variety of ways, i.e. slightly different proposal distributions $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ might be used in different implementations and be referred to by the same move name.




Figure 4. Illustration of the naive move for a single particle in a one-dimensional harmonic trap for $n = 32$ beads. The black circles depict the old bead positions. The red square shows the proposed bead position for the 16th time slice index. It can be seen that only two links (namely the link involving the 15th and 16th beads and that involving the 16th and 17th beads) are changed.

4.3.1. Naive move. The simplest move (the naive move) consists of shifting the position vector \mathbf{r}_{old} of a single bead by $\delta\mathbf{r}$, where $\delta\mathbf{r}$ is drawn uniformly from the interval $[-\Delta\mathbf{r}, \Delta\mathbf{r}]$. The basic idea behind this move is that the propagator is a smooth function of \mathbf{x} and that a small change in \mathbf{x} does not introduce a huge change in the probability distribution $\pi(\mathbf{x})$. The size $2\Delta\mathbf{r}$ of the interval (if we are simulating a three-dimensional system, then the interval corresponds to a cube) can be adjusted such that the acceptance ratio of the proposed new position vector is around 50%. The proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ is equal to a constant if the new bead lies in the interval $[\mathbf{r}_{\text{old}} - \Delta\mathbf{r}, \mathbf{r}_{\text{old}} + \Delta\mathbf{r}]$; otherwise, $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ is equal to 0. Following equation (50), the move is accepted according to

$$\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}') = \min\left(1, \frac{\pi(\mathbf{x}')}{\pi(\mathbf{x})}\right). \quad (53)$$

Importantly, one cannot choose an ‘unbalanced’ interval like $[-\epsilon\Delta\mathbf{r}, \Delta\mathbf{r}]$, where $\epsilon < 1$, since the detailed balance condition, equation (49), is not satisfied in this case. The reason is that it is possible to go from \mathbf{r} to $\mathbf{r} + \Delta\mathbf{r}$ in one move but that it is impossible to go from $\mathbf{r} + \Delta\mathbf{r}$ to \mathbf{r} in one move.

The algorithm for the naive move can be summarized as follows. (i) Let the current configuration be $\mathbf{x} = \{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\}$. Randomly select a particle index k and a time slice index j , where j can take any value from 0 to $2n$. Set $\mathbf{r}_{\text{old}} = \mathbf{r}_{k,j}$ and calculate the old probability distribution $\pi_{\text{old}} = \pi(\mathbf{R}_0, \dots, \mathbf{R}_{2n})$. (ii) Generate a new position $\mathbf{r}_{\text{new}} = \mathbf{r}_{\text{old}} + \delta\mathbf{r}$, where $\delta\mathbf{r}$ is drawn uniformly from the interval $[-\Delta\mathbf{r}, \Delta\mathbf{r}]$. Define $\mathbf{R}_j^{\text{new}} = \{\mathbf{r}_{1,j}, \dots, \mathbf{r}_{k-1,j}, \mathbf{r}_{\text{new}}, \mathbf{r}_{k+1,j}, \dots, \mathbf{r}_{N,j}\}$ and calculate the new probability distribution $\pi_{\text{new}} = \pi(\mathbf{R}_0, \dots, \mathbf{R}_{j-1}, \mathbf{R}_j^{\text{new}}, \mathbf{R}_{j+1}, \dots, \mathbf{R}_{2n})$. (iii) Calculate the ratio $\pi_{\text{new}}/\pi_{\text{old}}$. If this ratio is larger than a random number drawn uniformly from 0 to 1, accept the move and set $\mathbf{r}_{k,j} = \mathbf{r}_{\text{new}}$; otherwise, reject the move

and set $\mathbf{r}_{k,j} = \mathbf{r}_{\text{old}}$ (i.e. do not change $\mathbf{r}_{k,j}$). Figure 4 illustrates the naive move for a single particle in a one-dimensional harmonic trap. It can be seen that the proposed move involves only one bead.

Although the naive move attempts to change only one bead at a time, whether the proposed move gets accepted or rejected depends, in principle, on all the beads, i.e. the coordinates of all particles at all time slices, since the acceptance/rejection depends on the ratio $\pi_{\text{new}}/\pi_{\text{old}}$. However, because the probability distribution $\pi(\mathbf{x})$ is a product over propagators and the trial function ψ_T evaluated at the two ends (equation (45)), the only terms that contribute to the ratio $\pi_{\text{new}}/\pi_{\text{old}}$ are the propagators $G(\mathbf{R}_{j-1}, \mathbf{R}_j; \Delta\tau)$ and $G(\mathbf{R}_j, \mathbf{R}_{j+1}; \Delta\tau)$ and, if j equals 0 or $2n$, the trial function $\psi_T(\mathbf{R}_0)$ or $\psi_T(\mathbf{R}_{2n})$. If one uses the second-order Trotter formula (equation (26)), which treats the potential and kinetic energy terms separately, the terms that contribute to $G(\mathbf{R}_{j-1}, \mathbf{R}_j; \Delta\tau)$ and $G(\mathbf{R}_j, \mathbf{R}_{j+1}; \Delta\tau)$ are the potential energy term $\exp[-\Delta\tau V(\mathbf{R}_j)]$ and the kinetic energy terms $G_0(\mathbf{r}_{k,j-1}, \mathbf{r}_{k,j}; \Delta\tau)$ and $G_0(\mathbf{r}_{k,j}, \mathbf{r}_{k,j+1}; \Delta\tau)$.

The caveat of the naive move is that the correlation length is typically large. In the best case scenario (i.e. in the case where all beads of all particles are considered exactly once and all proposed moves are accepted), $(2n + 1) \times N$ moves are needed to generate a configuration in which every bead differs from the starting configuration. Thus, we calculate observables for every $(\alpha \times (2n + 1) \times N)$ th configuration, where α is a constant greater than 1 that is adjusted to ensure that the observables are calculated from configurations with small correlations. In practice, we find that α lies between 2 and 20 for the applications considered in this review.

4.3.2. Wiggle move. The wiggle move uses the non-interacting propagator to design the proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$. Since the non-interacting propagator is a product of simple Gaussians for particles in free space (equation (16)) and for particles confined in a harmonic trap [45, 59], one can generate configurations efficiently with 100% acceptance ratio using the Box–Muller transformation [78] or with finite acceptance ratio using the Marsaglia polar method (the acceptance ratio is around 80%) [79] or the Zigurat algorithm (the acceptance ratio is around 98%) [80, 81]. If the difference between the propagator of the system to be simulated and the non-interacting propagator is small, the acceptance ratio for a move generated based on the propagator of the non-interacting system is high. Despite the large acceptance ratio, the correlation between consecutive configurations is, in general, small. In the non-interacting limit, the acceptance ratio is exactly 1.

Depending on the number of beads changed simultaneously, the wiggle move is a single-bead move or a multi-bead move. The single-bead version of the wiggle move randomly selects a particle index k and a time slice index j ($j > 0$ and $j < 2n$). Since the beads to be moved exclude the time slices 0 and $2n$, the wiggle move does not involve the trial function ψ_T . We denote the new proposed position vector by $\mathbf{r}_{k,j}^{\text{new}}$ (how to choose $\mathbf{r}_{k,j}^{\text{new}}$ is discussed below) and define

$\mathbf{R}_j^{\text{new}} = (\mathbf{r}_{1,j}, \dots, \mathbf{r}_{k-1,j}, \mathbf{r}_{k,j}^{\text{new}}, \mathbf{r}_{k+1,j}, \dots, \mathbf{r}_{N,j})$. The old and proposed configurations read

$$\mathbf{x} = \{\mathbf{R}_0, \dots, \mathbf{R}_{2n}\} \quad (54)$$

and

$$\mathbf{x}' = \{\mathbf{R}_0, \dots, \mathbf{R}_{j-1}, \mathbf{R}_j^{\text{new}}, \mathbf{R}_{j+1}, \dots, \mathbf{R}_{2n}\}, \quad (55)$$

respectively. We choose $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ according to the propagator G_0 (equation (16)) of the non-interacting system without confinement (the proposal distribution based on the propagator of the non-interacting harmonic oscillator Hamiltonian can be treated similarly),

$$\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}') = \exp\left(-\frac{(\mathbf{r}_{k,j}^{\text{new}} - \mathbf{r}_{k,j-1})^2 + (\mathbf{r}_{k,j}^{\text{new}} - \mathbf{r}_{k,j+1})^2}{4\lambda_m \Delta\tau}\right) \quad (56)$$

or, rearranging the exponent,

$$\begin{aligned} \mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}') &= \exp\left(-\frac{(\mathbf{r}_{k,j-1} - \mathbf{r}_{k,j+1})^2}{8\lambda_m \Delta\tau}\right) \\ &\times \exp\left(-\frac{[\mathbf{r}_{k,j}^{\text{new}} - (\mathbf{r}_{k,j-1} + \mathbf{r}_{k,j+1})/2]^2}{2\lambda_m \Delta\tau}\right). \end{aligned} \quad (57)$$

The first exponential on the right-hand side of equation (57) is independent of $\mathbf{r}_{k,j}^{\text{new}}$ and can be interpreted as a normalization constant. The second exponential on the right-hand side of equation (57) is equal to a Gaussian whose mean value is given by the midpoint of the $(j-1)$ th and the $(j+1)$ th bead of the k th particle and whose variance is $\lambda_m \Delta\tau$. Thus, $\mathbf{r}_{k,j}^{\text{new}}$ can be generated using the Box–Muller transformation, the Marsaglia polar method, or the Ziggurat algorithm discussed above. If we use the second-order Trotter formula, the acceptance probability $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ (equation (50)) takes a fairly simple form since a large number of terms (those not involving time slice j and those not involving particle k) in the ratio $\pi(\mathbf{x}')/\pi(\mathbf{x})$ can be cancelled. If the pair product approximation is used, the evaluation of $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ is more involved since fewer terms in the ratio $\pi(\mathbf{x}')/\pi(\mathbf{x})$ can be cancelled due to the fact that the kinetic energy and the potential energy terms are ‘linked’ in the pair product approximation.

The single-bead version of the wiggle move can be generalized to multiple consecutive beads. Since the multi-bead move leads to a deformation of a segment of the path, the move is called ‘wiggle move’. In what follows, our discussion is guided by [9]. Instead of a single bead of the path we propose to change a path segment consisting of multiple beads according to a proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$ that generalizes the expression given in equation (57). We denote the time slice indices of the two ends that are unchanged by j and $j+s$, where s is an integer power of 2; the condition for s allows one, as will become clear

below, to organize the move into ‘levels’. The corresponding position vectors are $\mathbf{r}_{k,j}$ and $\mathbf{r}_{k,j+s}$, where $s > 0$. Note that the wiggle move does not explicitly involve the trial function ψ_T since the path segment to be changed has to be continuous. This means that the zeroth bead can be at the beginning of the segment but nowhere else and that the $2n$ th bead can be at the end of the segment but nowhere else. In the finite-temperature path integral Monte Carlo approach, in contrast, any path segment can be considered, provided the path is closed.

We now outline the multi-bead move, both without and with ‘staging’. The algorithm without staging is less efficient but can be employed in connection with the Trotter formula and the pair product approximation. The staging version can only be used in connection with the Trotter formula. Both multi-bead move versions generate a proposed new path segment $\{\mathbf{r}_{k,j+1}^{\text{new}}, \dots, \mathbf{r}_{k,j+s-1}^{\text{new}}\}$ that is completely independent of the old path segment $\{\mathbf{r}_{k,j+1}, \dots, \mathbf{r}_{k,j+s-1}\}$. Here, $j+s$ has to be smaller than or equal to $2n$.

To motivate the strategy of the multi-bead wiggle move, we write the pieces of the non-interacting propagator G_0 that depend on the particle index k and the time slice indices j to $j+s$ out explicitly,

$$\begin{aligned} G_0 &\sim \exp\left[-\frac{(\mathbf{r}_{k,j} - \mathbf{r}_{k,j+s})^2}{4s\lambda_m \Delta\tau}\right] \\ &\quad \underbrace{\text{constant}} \\ &\times \exp\left[-\frac{(\mathbf{r}_{k,j+s/2} - \bar{\mathbf{r}}_{k,j,j+s})^2}{s\lambda_m \Delta\tau}\right] \\ &\quad \underbrace{\text{zeroth level}} \\ &\times \exp\left[-\frac{(\mathbf{r}_{k,j+s/4} - \bar{\mathbf{r}}_{k,j,j+s/2})^2}{s\lambda_m \Delta\tau/2}\right] \exp\left[-\frac{(\mathbf{r}_{k,j+3s/4} - \bar{\mathbf{r}}_{k,j+s/2,j+s})^2}{s\lambda_m \Delta\tau/2}\right] \\ &\quad \underbrace{\text{first level}} \\ &\times \dots, \end{aligned} \quad (58)$$

where $\bar{\mathbf{r}}_{k,\alpha,\beta} = (\mathbf{r}_{k,\alpha} + \mathbf{r}_{k,\beta})/2$. If s is equal to 2^l , equation (58) contains l levels (the zeroth level is counted as one level but the constant term is not). Equation (58) suggests that the sampling can be done level by level. For example, for a path segment consisting of three time slices ($s = 4$ and $j = 0$), the beginning bead is $\mathbf{r}_{k,0}$ and the ending bead is $\mathbf{r}_{k,4}$. Thus, there exist two levels in total. First, the new midpoint bead $\mathbf{r}_{k,2}^{\text{new}}$ is proposed according to the 0th level (partial) proposal distribution $\mathcal{G}_{0\text{th}}(\mathbf{x} \rightarrow \mathbf{x}')$, i.e. $\mathbf{r}_{k,2}^{\text{new}}$ is generated by sampling a three-dimensional Gaussian distribution with variance $s\lambda_m \Delta\tau/2$. Second, the new midpoint beads $\mathbf{r}_{k,1}^{\text{new}}$ and $\mathbf{r}_{k,3}^{\text{new}}$ are proposed according to the 1st level (partial) proposal distribution $\mathcal{G}_{1\text{st}}(\mathbf{x} \rightarrow \mathbf{x}')$, i.e. $\mathbf{r}_{k,1}^{\text{new}}$ and $\mathbf{r}_{k,3}^{\text{new}}$ are generated by sampling three-dimensional Gaussian distributions with variance $s\lambda_m \Delta\tau/4$.

In general, the u th level (partial) proposal distribution $\mathcal{G}_{u\text{th}}(\mathbf{x} \rightarrow \mathbf{x}')$ reads

$$\mathcal{G}_{u\text{th}}(\mathbf{x} \rightarrow \mathbf{x}') = \prod_{v=1}^{2^u} \exp\left(-\frac{[\mathbf{r}_{k,j+s(2^v-1)/2^{u+1}}^{\text{new}} - (\mathbf{r}_{k,j+s(v-1)/2^u} + \mathbf{r}_{k,j+s v/2^u})/2]^2}{\lambda_m s \Delta\tau/2^u}\right), \quad (59)$$




Figure 5. Illustration of the wiggle move for a single particle in a one-dimensional harmonic trap. The black circles and red squares depict the old and proposed new configurations, respectively. It is assumed that the construction of the new path segment was continued after the construction of the first midpoint bead, the next two midpoint beads, and so on.

which implies that $\mathbf{r}_{k,j+s(2v-1)/2^{u+1}}^{\text{new}}$ can be generated by sampling a three-dimensional Gaussian with variance $\lambda_m s \Delta\tau / 2^{u+1}$. Since the u th level proposal distribution $\mathcal{G}_{u\text{th}}(\mathbf{x} \rightarrow \mathbf{x}')$ depends only on the position vectors of the $(u-1)$ th level, the new path segment $\{\mathbf{r}_{k,j+1}^{\text{new}}, \dots, \mathbf{r}_{k,j+s-1}^{\text{new}}\}$ can, indeed, be generated level by level. The product of $\mathcal{G}_{u\text{th}}(\mathbf{x} \rightarrow \mathbf{x}')$ over all levels (i.e. $\prod_{u=0}^{l-1} \mathcal{G}_{u\text{th}}$) yields the ‘full’ proposal distribution $\mathcal{G}(\mathbf{x} \rightarrow \mathbf{x}')$. Denoting the time slices that involve the newly proposed beads by $\mathbf{R}_v^{\text{new}}$, where v ranges from $j+1$ to $j+s-1$, and using the second-order Trotter formula, the acceptance distribution $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ becomes

$$\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}') = \min \left(1, \prod_{v=j+1}^{j+s-1} \frac{\exp[-\Delta\tau V(\mathbf{R}_v^{\text{new}})]}{\exp[-\Delta\tau V(\mathbf{R}_v)]} \right). \quad (60)$$

The staging algorithm allows one to reject the multi-bead wiggle move in advance, i.e. before the entire new path segment has been generated, if ‘bad bead positions’ are drawn [9]. The ‘in advance rejection’ is checked for at each level u . Let us assume that we are considering level u with the new midpoint beads $\mathbf{r}_{k,j+s(2v-1)/2^{u+1}}^{\text{new}}$, where v ranges from 1 to 2^u . Using the second-order Trotter formula, the move is accepted or rejected based on

$$\mathcal{A}_{u\text{th}}(\mathbf{x} \rightarrow \mathbf{x}') = \min \left(1, \prod_{v=1}^{2^u} \frac{\exp[-\Delta\tau V(\mathbf{R}_{j+s(2v-1)/2^{u+1}}^{\text{new}})]}{\exp[-\Delta\tau V(\mathbf{R}_{j+s(2v-1)/2^{u+1}})]} \right). \quad (61)$$

If $\mathcal{A}_{u\text{th}}$ is smaller than a random number drawn uniformly from the interval $[0,1]$, the move is rejected at the u th level and the new configuration is set equal to the old configuration; otherwise, the move is accepted. If the move is accepted at the u th level, we go to the $(u+1)$ th level and repeat the procedure. If the final level is reached and the new proposed

beads are accepted, then the entire path segment consisting of the proposed new beads $\mathbf{r}_{k,j+1}^{\text{new}}, \dots, \mathbf{r}_{k,j+s-1}^{\text{new}}$ is accepted and a configuration with a new path segment has been generated.

The outcome of the ‘multi-bead sampling + staging’ algorithm is equivalent to that of the multi-bead algorithm without staging, which proposes all the beads of the path segment considered first and then accepts or rejects at the very end. The staged (or in-advance) rejection speeds up the algorithm. Importantly, the staging algorithm only works if the Trotter formula is used. If the pair product approximation is used, the rejection needs to be done at the very end because the propagator for consecutive time slices cannot be reorganized into different levels.

Figure 5 illustrates the wiggle move for a single particle in a one-dimensional harmonic trap ($j=14$ and $s=16$). The proposed new path is constructed as follows: a new midpoint bead with index 22 is proposed and tested according to equation (61); if rejected (i.e. if the random number generated is smaller than $\mathcal{A}_{0\text{th}}$), the move is aborted in advance and the new configuration is set to the old configuration; if not rejected, the construction of the new path segment is continued (this is what is assumed in making figure 5). In the latter case, two new midpoint beads with index 18 and 26 are proposed and tested simultaneously according to equation (61). If rejected, the move is aborted in advance and the new configuration is set to the old configuration; if not rejected, four new midpoint beads with index 16, 20, 24, and 28 are proposed and tested simultaneously according to equation (61). If rejected, the move is aborted in advance and the new configuration is set to the old configuration; if not rejected, eight new midpoint beads with index 15, 17, 19, 21, 23, 25, and 29 are tested simultaneously according to equation (61). If rejected, the move is aborted with the new configuration being the old configuration; if not rejected, the move is accepted in its entirety and the path segment that involves the beads with index 15–29 is changed to the new position vectors.

4.3.3. Pair distance move. The pair distance move is employed in systems with two-body zero-range interactions. As discussed earlier, two-body zero-range interactions can be treated using the pair product approximation but not, to the best of our knowledge, using the Trotter formula based decomposition of the propagator. The use of the pair distance move is especially important if the two-body s -wave scattering length diverges. The key motivation is that two particles can, if zero-range interactions are present, be close to each other or even on top of each other. Traditional moves such as the wiggle move and the naive move, however, do not generate configurations in which particles sit on top of each other. The reason is that the scaled pair distribution function $4\pi r^2 P_{12}(r)$ for non-interacting particles or for uniformly distributed particles is zero at $r=0$, implying that configurations with vanishing pair distance are not generated by the traditional moves. The pair distance move involves two particles with the same time slice index j , where $j=0, \dots, 2n$. The proposed move keeps the center of mass

of the selected pair unchanged but modifies its relative distance vector.

The pair distance move can be implemented as follows. (i) Randomly choose the indices k , l and j of the single-particle beads $\mathbf{r}_{k,j}$ and $\mathbf{r}_{l,j}$ involved in the move and set $\mathbf{r}_{\text{old}} = \mathbf{r}_{k,j} - \mathbf{r}_{l,j}$ and $r_{\text{old}} = |\mathbf{r}_{\text{old}}|$. Store the old center-of-mass vector \mathbf{b} of the selected pair, $\mathbf{b} = (\mathbf{r}_{k,j} + \mathbf{r}_{l,j})/2$. (ii) Generate a new relative distance vector $\mathbf{r}_{\text{new}} = \mathbf{r}_{\text{old}} + \delta r \hat{\mathbf{r}}_{\text{old}}$, where δr is obtained by choosing a value uniformly from the pre-set interval $[-\Delta r, \Delta r]$. This prescription implies that δr can be negative and that \mathbf{r}_{new} , in turn, lies along the directions \mathbf{r}_{old} or $-\mathbf{r}_{\text{old}}$. (iii) Calculate the ratio w , $w = (r_{\text{old}} + \delta r)^2 \pi_{\text{new}} / [(r_{\text{old}})^2 \pi_{\text{old}}]$. If this ratio is larger than a random number drawn uniformly from 0 to 1, accept the move and set $\mathbf{r}_{k,j} = \mathbf{b} + \mathbf{r}_{\text{new}}/2$ and $\mathbf{r}_{l,j} = \mathbf{b} - \mathbf{r}_{\text{new}}/2$; otherwise, reject the move (in this case, $\mathbf{r}_{k,j}$ and $\mathbf{r}_{l,j}$ remain unchanged). The value of Δr is adjusted such that approximately 50% of the proposed moves are accepted.

The acceptance/rejection step involves the quantity w . If j is not equal to 0 or $2n$, w reduces to

$$w = \frac{(r_{\text{old}} + \delta r)^2 G(\mathbf{R}_{j-1}, \mathbf{R}_j^{\text{new}}; \Delta\tau) G(\mathbf{R}_j^{\text{new}}, \mathbf{R}_{j+1}; \Delta\tau)}{(r_{\text{old}})^2 G(\mathbf{R}_{j-1}, \mathbf{R}_j^{\text{old}}; \Delta\tau) G(\mathbf{R}_j^{\text{old}}, \mathbf{R}_{j+1}; \Delta\tau)}. \quad (62)$$

If j is equal to 0, one finds

$$w = \frac{(r_{\text{old}} + \delta r)^2 \psi_T(\mathbf{R}_0^{\text{new}}) G(\mathbf{R}_0^{\text{new}}, \mathbf{R}_1; \Delta\tau)}{(r_{\text{old}})^2 \psi_T(\mathbf{R}_0^{\text{old}}) G(\mathbf{R}_0^{\text{old}}, \mathbf{R}_1; \Delta\tau)}. \quad (63)$$

The expression for $j = 2n$ is similar to that for $j = 0$. In equations (62) and (63), $\mathbf{R}_j^{\text{new}}$ is defined as

$$\mathbf{R}_j^{\text{new}} = \{\mathbf{r}_1, \dots, \mathbf{r}_{k-1,j}, \mathbf{b} + \mathbf{r}_{\text{new}}/2, \mathbf{r}_{k+1,j}, \dots, \mathbf{r}_{l-1,j}, \mathbf{b} - \mathbf{r}_{\text{new}}/2, \mathbf{r}_{l+1,j}, \dots, \mathbf{r}_{N,j}\}. \quad (64)$$

As already alluded to, the propagator G entering in the expressions for w is expressed using the pair product approximation, i.e. each of the G 's is replaced by the expression given in equation (37). Careful inspection of the resulting expression for w shows that a number of terms in the numerator can be canceled by corresponding terms in the denominator. However, since the reduced pair propagator \bar{G}_{rel} 'connects' the k th particle with all other particles and the l th particle with all other particles, two products over a particle dummy index survive for each of the G 's, making the pair distance move computationally more expensive than the naive move implemented using the Trotter decomposition (in fact, the argument just given explains why the pair product approximation is, generally speaking, computationally more demanding than Trotter formula based schemes). The ratio $(r_{\text{old}} + \delta r)^2 / (r_{\text{old}})^2$, which is included in the acceptance step, ensures that the small interparticle distance behavior is described properly. The reader is referred to [45] for more details.

4.4. Expectation values

The previous section outlined how to generate new configurations. Assuming that no symmetrization or anti-symmetrization is needed (see section 4.6 for details) and that a suitable trial function ψ_T is known, the missing piece for completing the PIGS algorithm is the determination of the weight function $w(\mathbf{x})$ (see equation (47)). This section discusses the derivation of the form of the weight function $w(\mathbf{x})$ for selected observables; the steps outlined below can be generalized to other observables. Explicit expressions for $w(\mathbf{x})$ can be derived for many observables using either quantum estimator relations such as equation (12) or thermodynamic type relations such as equation (13). The determination of the superfluid or condensate fractions is more involved and not considered in this review.

4.4.1. Example: energy estimator. Using the thermodynamic type relation, equation (13), and plugging in one of the approximate expressions for the propagator, an explicit expression for the weight function $w(\mathbf{x})$ can be derived. As an example, we consider the thermodynamic energy estimator for the second-order Trotter formula for particles without permutations. Using equation (19) and equation (25) without the $\mathcal{O}(\Delta\tau^3)$ term in equation (7), the normalization factor $Z(\tau)$ reads

$$\begin{aligned} Z(\tau) &= \int_{\mathbf{R}_0} \dots \int_{\mathbf{R}_{2n}} \psi_T(\mathbf{R}_0) \psi_T(\mathbf{R}_{2n}) G_0(\mathbf{R}_0, \mathbf{R}_1; \Delta\tau) \\ &\quad \times \dots \times G_0(\mathbf{R}_{2n-1}, \mathbf{R}_{2n}; \Delta\tau) \\ &\quad \times \exp \left[-\Delta\tau \left(\frac{1}{2} V(\mathbf{R}_0) + \frac{1}{2} V(\mathbf{R}_{2n}) + \sum_{j=1}^{2n-1} V(\mathbf{R}_j) \right) \right] \\ &\quad \times d\mathbf{R}_0 \dots d\mathbf{R}_{2n}. \end{aligned} \quad (65)$$

Using equation (65) in equation (13), recalling that τ is equal to $n\Delta\tau$, and denoting the energy estimator E_τ —calculated for a finite number of time slices by $\langle E_\tau \rangle$ —we obtain

$$\begin{aligned} \langle E_\tau \rangle &= -\frac{1}{Z(\tau)} \int_{\mathbf{R}_0} \dots \int_{\mathbf{R}_{2n}} \frac{1}{2n} \psi_T(\mathbf{R}_0) \psi_T(\mathbf{R}_{2n}) \\ &\quad \times \frac{\partial}{\partial \Delta\tau} \left\{ G_0(\mathbf{R}_0, \mathbf{R}_1; \Delta\tau) \times \dots \times G_0(\mathbf{R}_{2n-1}, \mathbf{R}_{2n}; \Delta\tau) \right. \\ &\quad \times \exp \left[-\Delta\tau \left(\frac{1}{2} V(\mathbf{R}_0) + \frac{1}{2} V(\mathbf{R}_{2n}) + \sum_{j=1}^{2n-1} V(\mathbf{R}_j) \right) \right] \left. \right\} \\ &\quad d\mathbf{R}_0 \dots d\mathbf{R}_{2n}. \end{aligned} \quad (66)$$

The goal is now to rewrite the right-hand side of equation (66) such that we can read off $w(\mathbf{x})$ by comparing with equation (47). Combining equation (26) (without the $\mathcal{O}(\Delta\tau^3)$ term) and equation (45), we recognize that equation (66) can be rewritten in terms of $\pi(\mathbf{x})$,

$$\langle E_\tau \rangle = -\frac{1}{Z(\tau)} \int_{\mathbf{x}} \frac{1}{2n} \frac{\partial \pi(\mathbf{x})}{\partial \Delta\tau} d\mathbf{x}. \quad (67)$$

The probability distribution $\pi(\mathbf{x})$ depends on $\Delta\tau$ through the $2n$ propagators G_0 . Applying the chain rule to evaluate the derivative with respect to $\Delta\tau$, we obtain

$$\begin{aligned} \langle E_T \rangle &= \frac{1}{Z(\tau)} \\ &\times \int_{\mathbf{x}} \frac{1}{2n} \left[\sum_{j=0}^{2n-1} \left(\frac{3N}{2\Delta\tau} - \frac{(\mathbf{R}_j - \mathbf{R}_{j+1})^2}{4\lambda_m \Delta\tau^2} \right) \right. \\ &\quad \left. + \frac{1}{2} V(\mathbf{R}_0) + \frac{1}{2} V(\mathbf{R}_{2n}) + \sum_{j=1}^{2n-1} V(\mathbf{R}_j) \right] \\ &\times \pi(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (68)$$

Comparing equation (68) with equation (47), one reads off

$$\begin{aligned} w(\mathbf{x}) &= \frac{1}{2n} \left[\sum_{j=0}^{2n-1} \left(\frac{3N}{2\Delta\tau} - \frac{(\mathbf{R}_j - \mathbf{R}_{j+1})^2}{4\lambda_m \Delta\tau^2} \right) \right. \\ &\quad \left. + \frac{1}{2} V(\mathbf{R}_0) + \frac{1}{2} V(\mathbf{R}_{2n}) + \sum_{j=1}^{2n-1} V(\mathbf{R}_j) \right]. \end{aligned} \quad (69)$$

The right-hand side of equation (69) can be evaluated straightforwardly, provided the configuration \mathbf{x} is known.

Using the quantum estimator relation, equation (12), an alternative energy estimator can be derived. Here, we derive the quantum energy estimator for particles without permutations using the second-order Trotter formula as an example. Our goal is to rewrite equation (12) such that we can read off the form of $w(\mathbf{x})$ by comparing with equation (47). To this end, we derive an auxiliary identity (see equation (74)) that we use below to rewrite the integrand of the numerator of equation (12).

Using the position representation of the Hamiltonian \hat{H} [82],

$$\langle \mathbf{R} | \hat{H} | \mathbf{R}' \rangle = H_{\mathbf{R}} \delta(\mathbf{R} - \mathbf{R}'), \quad (70)$$

where

$$H_{\mathbf{R}} = -\lambda_m \nabla_{\mathbf{R}}^2 + V(\mathbf{R}), \quad (71)$$

one finds

$$\begin{aligned} &\int_{\mathbf{R}'} \langle \mathbf{R} | \hat{H} | \mathbf{R}' \rangle G(\mathbf{R}', \mathbf{R}''; \tau) d\mathbf{R}' \\ &= \int_{\mathbf{R}'} \{ [-\lambda_m \nabla_{\mathbf{R}}^2 + V(\mathbf{R})] \delta(\mathbf{R} - \mathbf{R}') \} G(\mathbf{R}', \mathbf{R}''; \tau) d\mathbf{R}' \end{aligned} \quad (72)$$

or, integrating by parts twice,

$$\begin{aligned} &\int_{\mathbf{R}'} \langle \mathbf{R} | \hat{H} | \mathbf{R}' \rangle G(\mathbf{R}', \mathbf{R}''; \tau) d\mathbf{R}' \\ &= \int_{\mathbf{R}'} \delta(\mathbf{R} - \mathbf{R}') \{ [-\lambda_m \nabla_{\mathbf{R}}^2 + V(\mathbf{R})] G(\mathbf{R}', \mathbf{R}''; \tau) \} d\mathbf{R}'. \end{aligned} \quad (73)$$

Performing the integration over \mathbf{R}' on the right-hand side of equation (73), we have

$$\begin{aligned} &\int_{\mathbf{R}'} \langle \mathbf{R} | \hat{H} | \mathbf{R}' \rangle G(\mathbf{R}', \mathbf{R}''; \tau) d\mathbf{R}' \\ &= [-\lambda_m \nabla_{\mathbf{R}}^2 + V(\mathbf{R})] G(\mathbf{R}, \mathbf{R}''; \tau). \end{aligned} \quad (74)$$

We denote the quantum estimator E_τ (equation (12)), evaluated using a finite number of time slices, by $\langle E_H \rangle$. Inserting the closure relation (equation (8)) $2n - 2$ times into equation (12) and using equation (74) with τ replaced by $\Delta\tau$, we obtain

$$\begin{aligned} \langle E_H \rangle &= \frac{1}{Z(\tau)} \\ &\times \int_{\mathbf{R}_0} \cdots \int_{\mathbf{R}_{2n}} \psi_T(\mathbf{R}_0) G(\mathbf{R}_0, \mathbf{R}_1; \Delta\tau) G(\mathbf{R}_1, \mathbf{R}_2; \Delta\tau) \\ &\times \cdots \times G(\mathbf{R}_{n-2}, \mathbf{R}_{n-1}; \Delta\tau) \\ &\times \{ [-\lambda_m \nabla_{\mathbf{R}_n}^2 + V(\mathbf{R}_n)] G(\mathbf{R}_{n-1}, \mathbf{R}_n; \Delta\tau) \} \\ &\times G(\mathbf{R}_n, \mathbf{R}_{n+1}; \Delta\tau) \times G(\mathbf{R}_{n+1}, \mathbf{R}_{n+2}; \Delta\tau) \times \cdots \\ &\times G(\mathbf{R}_{2n-1}, \mathbf{R}_{2n}; \Delta\tau) \psi_T(\mathbf{R}_{2n}) d\mathbf{R}_0 \cdots d\mathbf{R}_{2n}. \end{aligned} \quad (75)$$

Applying the second-order Trotter formula (equation (26) without the $\mathcal{O}(\Delta\tau^3)$ term) to equation (75), we obtain

$$\begin{aligned} \langle E_H \rangle &= \frac{1}{Z(\tau)} \int_{\mathbf{x}} \left[\frac{3N}{2\Delta\tau} - \frac{(\mathbf{R}_n - \mathbf{R}_{n-1})^2}{4\lambda_m \Delta\tau^2} + V(\mathbf{R}_n) \right. \\ &\quad \left. - (\mathbf{R}_n - \mathbf{R}_{n-1}) \cdot \nabla_{\mathbf{R}_n} V(\mathbf{R}_n) \right. \\ &\quad \left. + \lambda_m \Delta\tau \nabla_{\mathbf{R}_n}^2 V(\mathbf{R}_n) - \lambda_m \Delta\tau^2 (\nabla_{\mathbf{R}_n} V(\mathbf{R}_n))^2 \right] \pi(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (76)$$

Because \hat{H} commutes with the propagator, \hat{H} can be applied to any time slice (in the derivation above, \hat{H} was applied to the n th time slice). This implies that one can average over all time slices to improve the accuracy (i.e. to take more ‘measurements’ for each configuration). Averaging over all possible time slice indices, we obtain

$$\begin{aligned} \langle E_H \rangle &= \frac{1}{Z(\tau)} \int_{\mathbf{x}} \frac{1}{2n} \left\{ \sum_{j=1}^{2n} \left[\frac{3N}{2\Delta\tau} - \frac{(\mathbf{R}_j - \mathbf{R}_{j-1})^2}{4\lambda_m \Delta\tau^2} + V(\mathbf{R}_j) \right. \right. \\ &\quad \left. \left. - (\mathbf{R}_j - \mathbf{R}_{j-1}) \cdot \nabla_{\mathbf{R}_j} V(\mathbf{R}_j) \right. \right. \\ &\quad \left. \left. + \lambda_m \Delta\tau \nabla_{\mathbf{R}_j}^2 V(\mathbf{R}_j) - \lambda_m \Delta\tau^2 (\nabla_{\mathbf{R}_j} V(\mathbf{R}_j))^2 \right] \right\} \pi(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (77)$$

Comparing equation (77) with equation (47), we obtain

$$\begin{aligned} w(\mathbf{x}) &= \frac{1}{2n} \left\{ \sum_{j=1}^{2n} \left[\frac{3N}{2\Delta\tau} - \frac{(\mathbf{R}_j - \mathbf{R}_{j-1})^2}{4\lambda_m \Delta\tau^2} + V(\mathbf{R}_j) \right. \right. \\ &\quad \left. \left. - (\mathbf{R}_j - \mathbf{R}_{j-1}) \cdot \nabla_{\mathbf{R}_j} V(\mathbf{R}_j) \right. \right. \\ &\quad \left. \left. + \lambda_m \Delta\tau \nabla_{\mathbf{R}_j}^2 V(\mathbf{R}_j) - \lambda_m \Delta\tau^2 (\nabla_{\mathbf{R}_j} V(\mathbf{R}_j))^2 \right] \right\}. \end{aligned} \quad (78)$$

In equation (78), the head (0th time slice) and tail (2nth time slice) are not treated on equal footing because of the partial derivative. The expression can be made ‘symmetric’ by

averaging over additional terms for which the derivative yields a term that contains the factor $\mathbf{R}_{j+1} - \mathbf{R}_j$. Compared to equation (69), equation (78) contains three extra terms in the sum. In the $n \rightarrow \infty$ limit, both estimators approach the true expectation value. However, for finite n , $\langle E_T \rangle$ and $\langle E_H \rangle$ generally give different estimates of the energy. To obtain accurate results, one needs to extrapolate the finite $\Delta\tau$ calculations to the zero time step limit. The difference between the two estimators for a single $\Delta\tau$ may be used as a rough estimate of the systematic error [6].

4.4.2. Example: structural properties. The energy estimator is special in that the information carried by all $2n + 1$ time slices can be used (see the sum over j in equations (69) and (78)). The reason is that the Hamiltonian operator commutes with the propagator. For other estimators, only the information carried by the middle or n th time slice and the associated propagators can, in general, be used. This section exemplarily discusses the determination of structural properties within the PIGS framework.

Quite generally, the operator \hat{D} corresponding to a structural observable can be written as a function $f(\mathbf{R})$ times a δ -function. For example, for the scaled pair distribution function $4\pi r^2 P_{12}(r)$ for particles one and two, $f(\mathbf{R})$ is equal to one and the δ -function is equal to $\delta(r_{\text{ref}} - |\mathbf{r}_{12}|)$, where \mathbf{r}_{12} is the distance vector between particles one and two. Replacing $\langle \mathbf{R} | \hat{H} | \mathbf{R}'' \rangle$ on the right-hand side of equation (12) by $\delta(r_{\text{ref}} - |\mathbf{r}_{12}|) \delta(\mathbf{R}' - \mathbf{R}'')$, one finds $w(\mathbf{x}) = \delta(r_{\text{ref}} - |\mathbf{r}_{1,n} - \mathbf{r}_{2,n}|)$. Similarly to the stochastic evaluation of structural properties for a given many-body (zero-temperature) wave function, the δ -function in the operator yielding the pair distribution function amounts to sorting the configurations into small intervals or bins and counting the number of configurations that fall into each of the intervals.

In practice, to obtain the scaled pair distribution function $4\pi r^2 P_{12}(r)$ for particles one and two, we discretize the pair distance r_{12} , $r_{12} = |\mathbf{r}_1 - \mathbf{r}_2|$, into a series of k_{max} bins $[k\delta r, (k+1)\delta r]$, where k range from 0 to $k_{\text{max}} - 1$. During the simulation, the pair distance is calculated for many configurations and sorted into the bins, i.e. a histogram of the pair distances is collected. For each configuration considered (note, we may skip configurations to ensure that the samples collected have negligible correlations), the pair distance r_{12} is calculated for the middle time slice \mathbf{R}_n . The bin number l of the histogram is calculated by evaluating $l = \text{Floor}(r_{12}/\delta r)$, where $\text{Floor}(x)$ gives the largest integer smaller or equal to x , and the histogram value v_l of the l th bin is increased by one. At the end, the histogram defined by the v_l is normalized by dividing by the total number B_l of pair distances considered and the bin size δr . The histogram created is a discretized version of the scaled pair distribution function $4\pi r^2 P_{12}(r)$. The approach outlined yields the correct normalization even if some pair distances generated during the simulation are larger than $k_{\text{max}}\delta r$. We typically monitor how many distances cannot be sorted into the histogram by comparing $\sum_l v_l$ with B_l . If the fraction is too large, then the ‘cutoff’ $k_{\text{max}}\delta r$ needs to be increased.

Because the process involved in calculating different structural properties such as the pair distribution function and triple distribution function is the same, the data structure used to accumulate different distribution functions and the accumulation process can be described by a single class in object oriented programming languages. This avoids duplication of the code. In the code, the desired estimator (the ‘object’) such as the pair distribution function estimator and the triple distribution function estimator can be constructed according to the same class and can be initialized with observable specific parameters such as the bin size, the bin number, and the number of particles. To accumulate the weight and finalize the results, the same virtual methods can be called for different estimators. The actual implementation of these virtual methods may or may not be the same for different estimators. For example, the scaled pair and scaled triple distribution functions can share the same implementation since both are described by an operator of the form $\delta(r_{\text{ref}} - r)$, where r_{ref} is either the pair distance or the three-body hyperradius (see, e.g., [83, 84] for the resulting distribution function). The (unscaled) pair distribution function, in contrast, is described by an operator of the form $\delta(r_{\text{ref}} - r)/(4\pi r^2)$ and has to be implemented separately.

4.5. Error analysis

The expectation value $\langle f(\mathbf{x}) \rangle$ of a function $f(\mathbf{x})$ with respect to the probability density function $p(\mathbf{x})$ is defined as

$$\langle f(\mathbf{x}) \rangle = \int_{\mathbf{x}} f(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}. \quad (79)$$

In the PIGS algorithm, we generate a finite series X of configurations \mathbf{x}_j ,

$$X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M\}, \quad (80)$$

according to the probability distribution $\pi(\mathbf{x})$. The expectation value $\langle O \rangle$ of an operator can then be estimated by the mean value \bar{O} of the series X ,

$$\bar{O} = \frac{1}{M} \sum_{j=1}^M w(\mathbf{x}_j). \quad (81)$$

We refer to \bar{O} as $\langle O \rangle_X$. In the limit $M \rightarrow \infty$, the mean value \bar{O} approaches the expectation value $\langle O \rangle$ (see equation (47)).

According to the central limit theorem, the mean value \bar{O} of the series X approaches the expectation value $\langle O \rangle$ in a predictive manner. The ‘standard’ central limit theorem states that the mean of a sufficiently large number of random samples, drawn from a distribution with a well-defined mean value and variance, is approximately normally distributed [59, 85]. Since the Markov chain generates a series of data that is correlated for small ‘lag’ and uncorrelated for large ‘lag’, the standard central limit theorem cannot be applied directly. However, it has been shown that the central limit theorem can be extended to Markov-chain generated data [86]. Thus, we divide the series X , obtained from the PIGS

samples $w(\mathbf{x}_j)$, into L blocks, each with $l = M/L$ configurations. Defining the block averages S_k ,

$$S_k = \frac{1}{l} \sum_{j=(k-1) \times l + 1}^{k \times l} w(\mathbf{x}_j), \quad (82)$$

we construct the series $\{S_1, \dots, S_L\}$. Provided l is sufficiently large (in our applications, the value of l ranges from 10 to 10^4), the block averages S_k are normally distributed and the variance σ^2 of the block averages can be estimated from the sample variance $\langle \sigma^2 \rangle_X$,

$$\langle \sigma^2 \rangle_X = \frac{1}{L-1} \sum_{j=1}^L (S_j - \bar{S}), \quad (83)$$

where

$$\bar{S} = \frac{1}{L} \sum_{j=1}^L S_j. \quad (84)$$

Note that \bar{S} is equal to \bar{O} (this can be seen by comparing equations (81) and (82)). In our simulations, we estimate the error of the expectation value $\langle O \rangle$ using $\langle \sigma_{\bar{O}} \rangle_X$,

$$\langle \sigma_{\bar{O}} \rangle_X = \sqrt{\frac{\langle \sigma^2 \rangle_X}{L}}, \quad (85)$$

i.e. we report the mean \bar{O} with error $\langle \sigma_{\bar{O}} \rangle_X$.

In our simulations, the number l of configurations per block is determined such that the block averages S_k are normally distributed. Alternatively, the value of l (and, assuming M is fixed, that of L) can be determined by calculating the autocorrelation length [1, 77]. The latter approach is more commonly used and is somewhat simpler to implement. The two approaches should yield comparable results.

Considering Q simulations, each yielding a series X_j and correspondingly $\sqrt{\langle \sigma^2 \rangle_{X_j}}$ (assuming finite L), the estimate of the standard deviation is biased because the mean value of a square root function is not equal to the square root of the mean, i.e.

$$\frac{\sum_{j=1}^Q \sqrt{\langle \sigma^2 \rangle_{X_j}}}{Q} \neq \sqrt{\frac{\sum_{j=1}^Q \langle \sigma^2 \rangle_{X_j}}{Q}}. \quad (86)$$

Since the bias becomes negligible for sufficiently large L , there is no need to correct for the bias. In our simulations, L is typically 80 or larger. Because the elements S_j in $\{S_1, \dots, S_L\}$ are normally distributed, the variance $\langle \sigma^2 \rangle_X$ is approximately a constant for sufficiently large L and the error $\langle \sigma_{\bar{O}} \rangle_X$ scales, according to equation (85), as $1/\sqrt{L}$. Thus, to improve the accuracy of an observable by an order of magnitude, the computational time needs to be increased by two orders of magnitude.

To check whether the final distribution is approximately normal, one can make a histogram of the observable under study. Figure 6 shows the normalized histogram for the energy of the (3, 3) system at unitarity. The notation (3, 3) refers to three spin-up fermions and three spin-down fermions under external harmonic confinement with angular frequency ω (the oscillator energy of $\hbar\omega$ is denoted by E_{ho}). The physics of this small fermionic system is discussed in more detail in section 5. The example considered here uses $\tau = 0.5(E_{\text{ho}})^{-1}$,




Figure 6. Normalized histogram of the energy E/E_{ho} , calculated using the thermal estimator, for the harmonically trapped (3, 3) system at unitarity. The histogram is constructed using 38 400 block averages generated using 480 processors (see the text for details). The red solid line shows the normal distribution with mean value of $E = 8.273E_{\text{ho}}$ and standard deviation of $1.124E_{\text{ho}}$ (the error of the mean is $1.124E_{\text{ho}}/\sqrt{38\,400} = 0.006E_{\text{ho}}$).

$\Delta\tau = 0.125(E_{\text{ho}})^{-1}$, and ψ_T given in equation (120); the resulting extrapolated $\Delta\tau = 0$ energy is reported in table 3. The simulation is done on 480 processors with each processor producing 80 block averages. This yields a total of 38 400 block averages. Even though these block averages are not obtained from a single Markov chain but from 480 independent Markov chains, we calculate the mean and error of the mean using, respectively, equations (84) and (85) with $L = 38\,400$. The resulting sample mean is $8.273E_{\text{ho}}$ with an error or uncertainty of $0.006E_{\text{ho}}$. Using the calculated mean and standard deviation, the solid line in figure 6 shows the corresponding normal distribution. It can be seen that the solid line provides a faithful description of the histogram, indicating that the underlying samples are indeed normally distributed.

The presented analysis requires a sufficiently large number of block averages. In practice, it may not be feasible or advisable to calculate many block averages. In such a case, one can check if the error scales as $\sqrt{1/L}$ with the number of blocks L . Reducing the number of blocks by a factor of two, one should observe that, if the block averages are normally distributed, the error increases roughly by a factor of $\sqrt{2}$. This check can be performed for as few as 5 or 10 blocks and provides, in many cases, enough information to reliably assign error bars.

To check explicitly whether the samples are independent, one needs to perform autocorrelation (or serial correlation) tests [7]. Given a series of numbers $\{x_1, \dots, x_L\}$, the lag k correlation coefficient r_k , which measures the correlation of the series of numbers $\{x_1, \dots, x_{L-k}\}$ and $\{x_{1+k}, \dots, x_L\}$, is given by [89]

$$r_k = \frac{\sum_{j=1}^{L-k} (x_j - \bar{x})(x_{j+k} - \bar{x})}{\sum_{j=1}^L (x_j - \bar{x})^2}, \quad (87)$$

where \bar{x} denotes the average of the numbers $\{x_1, \dots, x_L\}$. If the samples are truly uncorrelated, the r_k approximately follow a

Table 3. Spin-balanced two-component Fermi gas ($N/2 \leq 5$) with zero-range interactions at unitarity. Column 1 lists the $(N/2, N/2)$ system considered. Column 2 reports the value of α , obtained from equation (127) using the energies reported in [87]. Columns 3–6 show the propagation time τ , the scheme used to extrapolate the energy to $\Delta\tau = 0$, the n used (the number of time slices is $2n + 1$), and the resulting extrapolated $\Delta\tau = 0$ PIGS energy E_{PIGS} with error bars, respectively. The abbreviation ‘extrap.’ in the header of column 4 stands for ‘extrapolation’ and ‘4th’ and ‘2nd’ are to be read as ‘4th order’ and ‘2nd order’, respectively. For comparison, columns 7 and 8 show energies from the literature, obtained using the explicitly correlated Gaussian (ECG) approach [87] (the energies are denoted by E_{ECG}) and the diffusion Monte Carlo method [88] (the energies are denoted by E_{DMC}), respectively; the energies E_{ECG} and E_{DMC} are obtained by extrapolating a series of finite-range energies to the zero-range limit.

	α	τE_{ho}	extrap.	n used	$E_{\text{PIGS}}/E_{\text{ho}}$	$E_{\text{ECG}}/E_{\text{ho}}$	$E_{\text{DMC}}/E_{\text{ho}}$
(2, 2)	0.505	1	4th	4, 5, 6, 8	5.0069(29)	5.0091(4)	5.028(2)
(3, 3)	0.390	0.5	2nd	3, 4	8.353(14)	8.337(4)	8.377(3)
(4, 4)	0.335	0.5	2nd	2, 4	11.99(7)	12.03(3)	12.04(1)
(5, 5)	0.306	0.25	2nd	1, 2	16.12(6)	16.12(6)	16.10(1)




Figure 7. The crosses show the correlation coefficient r_k as a function of the lag k ($k = 1-40$) for the sample and observable considered in figure 6. A series of 80 block averages, obtained from a single Markov chain (and processor), is analyzed. The upper and lower solid lines show the 95% confidence interval defined in equation (88).

normal distribution for sufficiently large $L - k$ and the variance of r_k is approximately equal to $1/L$. Furthermore, the probability that r_k falls into the interval

$$\left[\frac{-1 - 1.96\sqrt{L - k - 1}}{L - k}, \frac{-1 + 1.96\sqrt{L - k - 1}}{L - k} \right] \quad (88)$$

is 95% [89]. Based on hypothesis testing theory [90], it is claimed, with 95% confidence, that a sample is correlated if r_k (based on a single test for one k) falls outside the interval given in equation (88).

Figure 7 shows the correlation coefficient r_k (equation (87)) with $L = 80$) for 80 block averages generated on a single processor (i.e. obtained from a single Markov chain) as a function of the lag k for the system and observable considered in figure 6. Since the correlation coefficients for $k \geq 1$ all lie within the confidence band, it is said that the data pass the autocorrelation test. For the data shown in figure 6, similar correlation coefficient plots are obtained for each of the 80 block averages generated by the other 479 processors. This verifies that the samples are truly independent.

4.6. Permutations: on-the-fly anti-symmetrization scheme

To account for the particle statistics, one needs to ensure the proper behavior of the propagator under particle permutations.

The Hilbert space for identical bosons or identical fermions is restricted compared to that of Boltzmann particles described by the same Hamiltonian. The discussion so far, including the short-time approximations for the propagator introduced in section 3.1, applies to Boltzmann particles.

In the ‘standard’ approach of generating paths for systems containing identical particles, the symmetrizer and anti-symmetrizer are evaluated stochastically (the corresponding move is referred to as ‘permutation move’) [6, 9]. This implies that Bose and Fermi systems are simulated by the same paths. Expectation values, in contrast, are accumulated by including ‘weight factors’ (plus and minus signs) that account for the particle statistics. This standard approach can be thought of as an analog of a post-symmetrization scheme, where one first generates configurations that represent the entire Hilbert space and then projects out those configurations that have the proper symmetry.

Here, we introduce an alternative ‘on-the-fly’ symmetrization/anti-symmetrization scheme that explicitly enforces the proper symmetry at each imaginary time index. This scheme is particularly useful for fermionic systems with zero-range interactions. Without a three-body regulator and without this on-the-fly anti-symmetrization scheme, two-component Fermi gases would undergo Thomas collapse [91]. A downside of the scheme discussed below is that the computational effort scales factorially with the number of identical particles; as a consequence, the scheme becomes prohibitively expensive with increasing number of particles. While the on-the-fly symmetrization scheme can be applied to systems that contain identical bosons, the explicit symmetrization is typically not needed in this case since the ground state wave function of the system in which the bosons are replaced by Boltzmann particles is identical to that of the system with bosons. Thus, the discussion in this section is most useful for fermions.

We start with a general discussion that will be useful for our on-the-fly anti-symmetrization scheme and then discuss on-the-fly anti-symmetrization scheme and PIGS specific aspects. To this end, we introduce the symmetrizer \hat{P} [6]. For single-component Bose and Fermi systems (N identical particles), \hat{P} can be written as [92]

$$\hat{P} = \frac{1}{N!} \sum_{\sigma} (\pm 1)^{N_I(\sigma)} \hat{P}_{\sigma}, \quad (89)$$

where σ denotes the permutation of particle indices, $N_1(\sigma)$ the number of inversions in σ ⁴, and \hat{P}_σ the corresponding permutation operator. The plus (minus) sign in equation (89) applies to identical bosons (fermions). For example, the symmetrizers (to unify the notation, we use the term symmetrizer for bosons and fermions) for two and three identical fermions are $\hat{P} = \hat{A}_2$ and $\hat{P} = \hat{A}_3$, $\hat{A}_2 = (1 - \hat{P}_{12})/2$ and $\hat{A}_3 = (1 - \hat{P}_{12} - \hat{P}_{13} - \hat{P}_{23} + \hat{P}_{123} + \hat{P}_{132})/6$, respectively. Here, $\hat{P}_{ijk\dots l}$ replaces the identity of particle i (i.e. its entire ‘information’ including spatial coordinates, spin degrees of freedom, etc) with that of particle j , that of particle j with that of particle k , \dots , and that of particle l with that of particle i . The symmetrizer \hat{P} commutes with \hat{P}_i if the i th and j th particles are identical. In the previous examples, \hat{A}_2 commutes with \hat{P}_{12} and \hat{A}_3 commutes with \hat{P}_{12} , \hat{P}_{13} , and \hat{P}_{23} . The definition of the symmetrizer \hat{P} can be generalized to multi-component Bose and Fermi systems as well as Bose–Fermi mixtures. In these cases, the total symmetrizer is written as a product of symmetrizers for each component. For example, the symmetrizer for the mixture of two identical bosons (particles 1 and 2) and two identical fermions (particles 3 and 4) reads $(1 + \hat{P}_{12})(1 - \hat{P}_{34})/4$.

The symmetrizer \hat{P} also commutes with the Hamiltonian \hat{H} and the propagator \hat{G} . \hat{P} serves the purpose of projecting out the wave functions that satisfy the proper exchange symmetry, i.e. it divides the Hilbert space into two parts: (i) if ψ_s is an eigen state with the proper symmetry, then one has $\hat{P}\psi_s = \psi_s$. (ii) If, in contrast, ψ_{ns} is an eigen state that does not have the proper exchange symmetry, then we have $\hat{P}\psi_{ns} = 0$. We note that the eigen values of the symmetrizer \hat{P} are 0 and 1 while those of the two-particle permute operator \hat{P}_{12} are -1 and 1 . As we will show in the following, the fact that the eigen values of \hat{P} are either 0 or 1 implies

$$\hat{P}^2 = \hat{P}. \quad (90)$$

To prove equation (90), we introduce a unitary matrix \mathcal{U} that diagonalizes the Hermitian symmetrizer \hat{P} , i.e. \mathcal{U} is constructed such that $\hat{D} = \mathcal{U}\hat{P}\mathcal{U}^{-1}$ is diagonal. Because \hat{P} and \hat{D} are related through a unitary transformation, \hat{D} and \hat{P} share the same eigen values. Since the eigen values of \hat{D} are either 0 or 1, \hat{D} is diagonal with diagonal elements 0 or 1. This implies that \hat{D}^2 is equal to \hat{D} . We now rewrite \hat{P}^2 using $\hat{P} = \mathcal{U}^{-1}\hat{D}\mathcal{U}$,

$$\hat{P}^2 = (\mathcal{U}^{-1}\hat{D}\mathcal{U})(\mathcal{U}^{-1}\hat{D}\mathcal{U}). \quad (91)$$

Using $\mathcal{U}^{-1}\mathcal{U} = 1$, we have

$$\hat{P}^2 = \mathcal{U}^{-1}\hat{D}^2\mathcal{U}. \quad (92)$$

Since \hat{D}^2 is equal to \hat{D} (see above), one finds $\hat{P}^2 = \mathcal{U}^{-1}\hat{D}\mathcal{U}$ and thus $\hat{P}^2 = \hat{P}$, which is what we set out to prove.

The symmetrized propagator can, as we prove below, be written as $\hat{G}_{\text{unsymm}}\hat{P}$, where \hat{G}_{unsymm} is the unsymmetrized propagator, i.e. the propagator for the corresponding system

with Boltzmann particles. In position space, the symmetrized propagator $G(\mathbf{R}, \mathbf{R}', \hat{P}; \tau)$ can be rewritten as [6]

$$G(\mathbf{R}, \mathbf{R}', \hat{P}; \tau) = \langle \mathbf{R} | \exp(-\tau\hat{H})\hat{P} | \mathbf{R}' \rangle \quad (93)$$

or as a sum over the unsymmetrized propagators,

$$G(\mathbf{R}, \mathbf{R}', \hat{P}; \tau) \propto \sum_{\sigma} \text{sgn}(\sigma) G(\mathbf{R}, \hat{P}_{\sigma}\mathbf{R}'; \tau), \quad (94)$$

where $\text{sgn}(\sigma)$ is the sign for the permutation σ (for single-component fermions, $\text{sgn}(\sigma) = (-1)^{N_1(\sigma)}$). In equation (94), we use the proportionality symbol since the ‘normalization factor’ depends on the number of identical particles in the system (for N identical particles, the proportionality symbol becomes an equal sign if the right-hand side is multiplied by $(N!)^{-1}$).

We now prove that $\hat{G}_{\text{unsymm}}\hat{P}$ is, indeed, the symmetrized propagator. In Schrödinger quantum mechanics, the symmetrized propagator in position space reads

$$G(\mathbf{R}, \mathbf{R}', \hat{P}; \tau) = \langle \mathbf{R} | \exp(-\tau\hat{H}) \sum_j |\psi_{\text{symm},j}\rangle \langle \psi_{\text{symm},j} | \mathbf{R}' \rangle, \quad (95)$$

where $\{\psi_{\text{symm},j}\}$ is the complete set of symmetrized eigen states, e.g., for N identical bosons or N identical fermions. The complete set of unsymmetrized eigen states of \hat{H} , i.e. the set of eigen states for Boltzmann particles is denoted by $\{\psi_{\text{unsymm},j}\}$. Recall that the \hat{P} operator can be diagonalized using the unitary matrix \mathcal{U} . \mathcal{U} ‘reorganizes’ the eigen states $\psi_{\text{unsymm},j}$ such that the new eigen states are also eigen states of \hat{P} . The resulting eigen states $\psi_{r,j}$,

$$\psi_{r,j} = \sum_l \mathcal{U}_{jl} \psi_{\text{unsymm},l}, \quad (96)$$

either have the proper symmetry, i.e. $\hat{P}\psi_{r,j}$ is equal to $\psi_{r,j}$ (in this case, the eigen value of \hat{P} is 1) or $\hat{P}\psi_{r,j}$ gives zero (in this case, the eigen value of \hat{P} is 0). The subset of eigen states $\{\psi_{r,j}\}$, for which $\hat{P}\psi_{r,j}$ is equal to $\psi_{r,j}$, coincides with the complete set of symmetrized eigen states. This process of constructing a set of properly symmetrized eigen states from a complete set of unsymmetrized eigen states is known as post-symmetrization. For later reference, we write down the auxiliary identity

$$\mathcal{U} \sum_l |\psi_{\text{unsymm},l}\rangle \langle \psi_{\text{unsymm},l}| \mathcal{U}^{-1} = \sum_l |\psi_{r,l}\rangle \langle \psi_{r,l}|, \quad (97)$$

which can be obtained using the matrix form of equation (96),

$$\{\psi_{r,1}, \psi_{r,2}, \dots\}^T = \mathcal{U} \{\psi_{\text{unsymm},1}, \psi_{\text{unsymm},2}, \dots\}^T. \quad (98)$$

We now manipulate the right-hand side of equation (93) so that it can be readily related to equation (95). Starting with the right-hand side of equation (93) and using equation (90), we find

$$\langle \mathbf{R} | \exp(-\tau\hat{H})\hat{P} | \mathbf{R}' \rangle = \langle \mathbf{R} | \exp(-\tau\hat{H})\hat{P}\hat{P} | \mathbf{R}' \rangle. \quad (99)$$

Inserting $\mathcal{U}\mathcal{U}^{-1} = \hat{1}$ and then

$$\sum_l |\psi_{\text{unsymm},l}\rangle \langle \psi_{\text{unsymm},l}| = \hat{1}, \quad (100)$$

⁴ In Mathematica, the number of inversions $N_1(\sigma)$ is represented by the function `Inversions`.

we find

$$\begin{aligned} \langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} | \mathbf{R}' \rangle &= \langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} \mathcal{U} \\ &\times \sum_l |\psi_{\text{unsymm},l}\rangle \langle \psi_{\text{unsymm},l} | \mathcal{U}^{-1} \hat{\mathcal{P}} | \mathbf{R}' \rangle. \end{aligned} \quad (101)$$

Using equation (97) in equation (101), we find

$$\begin{aligned} \langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} | \mathbf{R}' \rangle &= \langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} \\ &\times \sum_l |\psi_{r,l}\rangle \langle \psi_{r,l} | \hat{\mathcal{P}} | \mathbf{R}' \rangle. \end{aligned} \quad (102)$$

Finally, noting that $\hat{\mathcal{P}} \sum_l |\psi_{r,l}\rangle \langle \psi_{r,l} | \hat{\mathcal{P}}$ is equal to $\sum_l |\psi_{\text{symm},l}\rangle \langle \psi_{\text{symm},l} |$, we arrive at

$$\begin{aligned} \langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} | \mathbf{R}' \rangle &= \langle \mathbf{R} | \exp(-\tau \hat{H}) \\ &\times \sum_l |\psi_{\text{symm},l}\rangle \langle \psi_{\text{symm},l} | \mathbf{R}' \rangle, \end{aligned} \quad (103)$$

i.e. we have proven that $\hat{G}_{\text{unsymm}} \hat{\mathcal{P}}$ (the left-hand side of equation (103) in position space) is identical to the symmetrized propagator in position space in Schrödinger quantum mechanics (equation (95)).

Replacing the propagator $\langle \mathbf{R} | \exp(-\tau \hat{H}) | \mathbf{R}' \rangle$ by $\langle \mathbf{R} | \exp(-\tau \hat{H}) \hat{\mathcal{P}} | \mathbf{R}' \rangle$ in all expressions involving the propagator (such as the normalization factor $Z(\tau)$, the probability distribution $\pi(\mathbf{x})$, and the weight function $w(\mathbf{x})$), we have all elements of the PIGS algorithm for bosons and fermions; as indicated at the beginning of this section, the trial function ψ_T will be discussed in section 5 for specific examples. The symmetrized probability distribution $\pi_{\text{symm}}(\mathbf{x})$, e.g., reads

$$\begin{aligned} \pi_{\text{symm}}(\mathbf{R}_0, \dots, \mathbf{R}_{2n}) &= \psi_T(\mathbf{R}_0) G(\mathbf{R}_0, \mathbf{R}_1, \hat{\mathcal{P}}; \Delta\tau) \\ &\times G(\mathbf{R}_1, \mathbf{R}_2, \hat{\mathcal{P}}; \Delta\tau) \times \dots \times G(\mathbf{R}_{2n-1}, \mathbf{R}_{2n}, \hat{\mathcal{P}}; \Delta\tau) \psi_T(\mathbf{R}_{2n}), \end{aligned} \quad (104)$$

where each symmetrized propagator is a sum over unsymmetrized propagators with permuted configurations. This implies that the complexity of the symmetrized PIGS algorithm is increased by up to a factor of $N!$ compared to the non-symmetrized PIGS algorithm.

Since the symmetrizer is applied at each link (see equation (104)), the number of operations needed to evaluate the symmetrized propagator scales as $2nN!$ for N identical particles. This implies that the on-the-fly symmetrization scheme becomes inefficient if the number of time slices is too large; section 5 demonstrates that reliable results for two-component Fermi gases can be obtained for a series of n as small as 1, 2 and 4. If the Trotter formula is used, some terms can typically be pulled out of the sum over the permutations, reducing the computational effort somewhat. If the pair product approximation is used, analogous simplifications are, in general, not possible. However, as mentioned at the beginning of this section, the on-the-fly symmetrization scheme is particularly useful for fermionic systems with zero-range interactions, which cannot be treated using the Trotter formula.

The change of the probability distribution from $\pi(\mathbf{x})$ to $\pi_{\text{symm}}(\mathbf{x})$ needs to be accounted for in the implementation of the moves and the collection of the expectation values. For example, the acceptance function for the naive move is equal to $\pi_{\text{symm}}(\mathbf{x})$, implying that $\mathcal{A}(\mathbf{x} \rightarrow \mathbf{x}')$ in equation (53) is

given by $\min(1, \pi_{\text{symm}}(\mathbf{x}')/\pi_{\text{symm}}(\mathbf{x}))$. The pair distance move has to be modified analogously. To obtain the acceptance function for the wiggle move, the derivation outlined in section 4.3.2 needs to be carried out using the symmetrized propagator. Similarly, the estimators need to be changed accordingly. For example, to account for the permutations in the thermodynamic energy estimator $\langle E_T \rangle$, $\pi(\mathbf{x})$ and $Z(\tau)$ in equation (66) have to be replaced by the corresponding symmetrized quantities. If we wrote out, in analogy to equation (69), the fully symmetrized expression for the weight function $w(\mathbf{x})$ using the second-order Trotter formula, it would be rather lengthy since the symmetrized propagator contains a sum over permutations at each time slice.

The ground state of fermionic systems usually corresponds to an excited state of the corresponding system with Boltzmann statistics. This implies that explicit anti-symmetrization is necessary to propagate the trial function to the ground state with the correct particle statistics. The anti-symmetrization introduces a ‘sign problem’ since the probability distribution can be positive or negative. The fact that the probability distribution can take either sign can be dealt with using the ideas of [83, 93].

The probability distribution for a given configuration can take either sign. Integrating all the positive (negative) portions, we obtain Z_+ (Z_-). The normalization factor $Z(\tau)$ is the sum of Z_+ and Z_- , $Z(\tau) = |Z_+| - |Z_-|$. When accumulating observables, the sign needs to be kept track of. In general, the expectation value of an observable O can be written as

$$\langle O \rangle = \frac{|Z_+|}{|Z_+| - |Z_-|} \langle O_+ \rangle - \frac{|Z_-|}{|Z_+| - |Z_-|} \langle O_- \rangle. \quad (105)$$

For the thermal energy estimator, e.g., this can be worked out explicitly. Using $Z(\tau) = |Z_+| - |Z_-|$, we find

$$\langle E_T \rangle = -(|Z_+| - |Z_-|)^{-1} \frac{\partial(|Z_+| - |Z_-|)}{\partial\tau} \quad (106)$$

or

$$\langle E_T \rangle = (|Z_+| - |Z_-|)^{-1} (|Z_+| \langle E_{T,+} \rangle - |Z_-| \langle E_{T,-} \rangle), \quad (107)$$

where

$$\langle E_{T,\pm} \rangle = -(|Z_{\pm}|)^{-1} \frac{\partial(|Z_{\pm}|)}{\partial\tau}. \quad (108)$$

In an actual calculation, the contributions $\langle O_+ \rangle$ and $\langle O_- \rangle$ to the estimator are first calculated separately and then weighted according to their relative magnitudes. Alternatively, equation (105) can be written as

$$\langle O \rangle = S^{-1} \left(\frac{|Z_+|}{|Z_+| + |Z_-|} \langle O_+ \rangle + \frac{|Z_-|}{|Z_+| + |Z_-|} \langle -O_- \rangle \right), \quad (109)$$

where

$$S = \frac{|Z_+| - |Z_-|}{|Z_+| + |Z_-|}. \quad (110)$$

This suggests that one can think of the simulation as yielding a purely positive normalization factor $|Z_+| + |Z_-|$; however, if one does so, a minus sign needs to be included in the

observable if the probability distribution for the chosen configuration is negative. The final result is obtained if the expression is multiplied by the factor S .

In equations (105) and (109), the term $|Z_+| - |Z_-|$ appears in the denominator. If $|Z_+|$ becomes closer to $|Z_-|$ with increasing propagation time, then the simulation becomes increasingly more challenging since the statistical noise needs to be smaller than the difference between $|Z_+|$ and $|Z_-|$.

Finally, one may wonder if there are any constraints on the trial function ψ_T . For the fixed-node diffusion Monte Carlo simulations, the trial function ψ_T needs to be an eigen state of the symmetrizer. This is not the case for PIGS simulations because the symmetrized propagator projects out the corresponding wave function.

5. Application to fermionic systems

5.1. General considerations

This section discusses applications of the PIGS algorithm to harmonically trapped equal-mass two-component Fermi gases consisting of n_1 spin-up and n_2 spin-down particles ($N = n_1 + n_2$) in three-dimensional space. We refer to these systems as (n_1, n_2) . The model Hamiltonian \hat{H} reads

$$\hat{H} = \hat{H}_{\text{free-space}} + \hat{V}_{\text{trap}}, \quad (111)$$

where

$$\hat{H}_{\text{free-space}} = \sum_{j=1}^N \frac{-\hbar^2}{2m} \nabla_j^2 + \sum_{j=1}^{N-1} \sum_{k>j}^N V_F(r_{jk}). \quad (112)$$

The interspecies two-body zero-range potential $V_F(r_{jk})$ is given in equation (40) and the confining potential \hat{V}_{trap} with angular trapping frequency ω reads

$$\hat{V}_{\text{trap}} = \frac{1}{2} m \omega^2 \sum_{j=1}^N \mathbf{r}_j^2. \quad (113)$$

Throughout this section, we assume that the interspecies two-body interaction is characterized by an infinitely large s -wave scattering length a_s , i.e. we consider systems at unitarity. Like fermions are assumed to be non-interacting, i.e. no intraspecies interactions are considered. This assumption is realized in cold atom systems provided one operates at magnetic field strengths away from p - and higher-partial-wave resonances. As discussed in sections 3.3–3.5, two-body zero-range interactions are most conveniently treated using the pair product approximation. In the calculations presented below, the reduced relative propagator $\bar{G}^{\text{rel}}(\mathbf{r}, \mathbf{r}'; \tau)$ (equation (43)), which accounts for the two-body zero-range interactions, the kinetic energy, and the relative two-body confining potential, is being used.

The Hamiltonian \hat{H} is, for infinitely large a_s , characterized by one (meaningful) length scale, the harmonic oscillator length a_{ho} , $a_{\text{ho}} = \sqrt{\hbar/(m\omega)}$. The harmonic oscillator length also characterizes the non-interacting system. The range of the interaction potential, which is zero, and the s -wave

scattering length, which is infinitely large, do not define meaningful length scales. Moreover, for two-component fermions with equal masses, the three-body system in free space is unbound, implying that the three-body system does not introduce a new (finite) length scale; in particular, Efimov physics is absent [94, 95]. In what follows, we express lengths in units of a_{ho} and energies in units of the harmonic oscillator energy E_{ho} , $E_{\text{ho}} = \hbar\omega$.

Two-component Fermi gases with vanishing interaction range and infinitely large interspecies s -wave scattering length have been and continue to be a paradigmatic strongly correlated system, for which few analytical results are known and which are challenging to treat numerically. The PIGS applications presented in this section have not been published before. The examples are chosen for their pedagogical value and for their relevance with regards to obtaining a more complete understanding of small two-component Fermi gases. Two types of systems are considered, spin-balanced systems ($n_1 = n_2 = N/2$; see section 5.2) and spin-imbalanced systems ($n_1 = N - 1$ and $n_2 = 1$; see section 5.3).

5.2. Spin-balanced Fermi gas

The ground state of spin-balanced two-component Fermi gases has $(L, \Pi) = (0, +1)$ symmetry, i.e. vanishing total orbital angular momentum L and positive parity Π . Intuitively, this can be understood by realizing that each spin-up fermion is paired with a spin-down fermion. In reality, the pairing respects the identical particle characteristics, i.e. each spin-up fermion is paired with $(1/n_2)$ th of each spin-down fermion and each spin-down fermion is paired with $(1/n_1)$ th of each spin-up fermion. The particle statistics is enforced along the paths by explicitly applying the symmetrizer to each of the $2n$ propagators (see, e.g., equation (104) for the symmetrized probability distribution). The applications below use trial functions ψ_T that have the proper particle symmetry build in. In general, one could employ any trial function that has finite overlap with the eigen state to be determined. In practice, however, it seems best to build as much ‘prior knowledge’ as possible into the trial function.

In the following, we discuss the construction of the trial function $\psi_T(\mathbf{R})$ and the dependence of the energy on the variational parameters entering into ψ_T . In addition, the convergence of the energy with respect to the total imaginary propagation time τ is analyzed. As can be seen from equation (3) and the surrounding discussion, τ should, in principle, be taken to infinity to allow for the excited state contributions to fully die out. In practice, this is not feasible since the noise or error that arises due to the anti-symmetrization (the sign error) increases with increasing τ . Thus, the task is to find a regime of τ values, for which the excited state contributions can be neglected and the sign error is sufficiently small. For each fixed τ simulation, the convergence of the results with respect to $\Delta\tau$ or, equivalently, the number of time slices needs to be checked. Typically, for each fixed τ , the results for several n are, in a first step, extrapolated to the

infinite n limit. In a second step, the $n \rightarrow \infty$ results for several τ are considered to determine for which τ excited state contributions can be neglected. In considering larger τ , it has to be checked that the sign error is sufficiently small. Last, the calculations should, ideally, be repeated for different ψ_T to ensure that the trial function does not introduce a bias. Having an overview of the general PIGS procedure, we now discuss the construction of the trial function ψ_T .

Quite generally, the construction of the trial function is guided by physical considerations. For example, one may parameterize the trial function ψ_T in terms of a set of variational parameters α , which are optimized by minimizing the expectation value of the Hamiltonian, calculated using ψ_T , with respect to the variational parameters α . A good trial function ψ_T is associated with a small energy variance. In fact, if the variance is zero, the trial function coincides with one of the exact eigen states of the model Hamiltonian. While the outlined optimization strategy has been applied quite fruitfully to a number of systems, it cannot—in general—be used for the model Hamiltonian and trial functions considered in this article since both the kinetic energy and the potential energy expectation values diverge for Hamiltonian with zero-range interactions. For an exact eigen state, the divergencies cancel, yielding a finite energy expectation value. For a trial function that does not fulfill the boundary conditions imposed by the two-body zero-range interactions (see below), the infinities prevent one from estimating the energy expectation value reliably. As a consequence, we optimize the ‘variational parameters’ contained in ψ_T using the PIGS approach itself or using analytical arguments.

To motivate the functional form of our trial function ψ_T , we consider the behavior of the many-body eigen function ψ when a spin-up fermion (particle j) and a spin-down fermion (particle k) approach each other while the other $3N - 3$ coordinates, collectively denoted by \mathbf{Y} , where $\mathbf{Y} = \{(\mathbf{r}_j + \mathbf{r}_k)/2, \mathbf{r}_1, \dots, \mathbf{r}_{j-1}, \mathbf{r}_{j+1}, \dots, \mathbf{r}_{k-1}, \mathbf{r}_{k+1}, \dots, \mathbf{r}_N\}$, are kept fixed [37],

$$\psi|_{r_{jk} \rightarrow 0} \propto \left(\frac{1}{r_{jk}} - \frac{1}{a_s} \right) B(\mathbf{Y}). \quad (114)$$

Here, B is a function that is independent of the distance vector \mathbf{r}_{jk} . The Bethe–Peierls boundary condition, equation (114), is a direct consequence of the interspecies two-body zero-range interactions and holds for any up–down pair distance r_{jk} . If the scattering length a_s diverges, as assumed throughout this section, the boundary condition reduces to

$$\psi|_{r_{jk} \rightarrow 0} \propto \frac{1}{r_{jk}} B(\mathbf{Y}). \quad (115)$$

The trial function is constructed such that ψ_T (i) approximately fulfills the boundary condition, equation (115); (ii) approximately describes a state containing $N/2$ up–down pairs; (iii) changes sign under the exchange of any two identical fermions; (iv) approximately accounts for the external harmonic confinement; and (v) has $(L, \Pi) = (0, +1)$ symmetry.

Guideline (ii) suggests a term of the form $\prod_{j=1}^{N/2} (r_{j,N/2+j})^{-1}$. This term fulfills the boundary condition for the pairs containing

the first and $(N/2 + 1)$ st particle, the second and $(N/2 + 2)$ nd particle, and so on but not the boundary condition for pairs containing, e.g., the first and $(N/2 + 2)$ nd particle. Application of the symmetrizer $\hat{\mathcal{P}}$ makes the short-distance behavior ‘less ideal’. To see this, let us exemplarily consider the $(n_1, n_2) = (2, 2)$ system and act with $\hat{\mathcal{P}}$ onto $(r_{13}r_{24})^{-1}$,

$$\hat{\mathcal{P}} \left(\frac{1}{r_{13}r_{24}} \right) = \frac{1}{2r_{13}r_{24}} - \frac{1}{2r_{23}r_{14}}. \quad (116)$$

Rewriting the second term on the right-hand side of equation (116) in terms of the independent Jacobi vectors \mathbf{r}_{13} , \mathbf{r}_{24} , and $\mathbf{r}_{13,24}$, where $\mathbf{r}_{13,24} = \mathbf{r}_2 + \mathbf{r}_4 - (\mathbf{r}_1 + \mathbf{r}_3)$ (these coordinates correspond to one of the so-called H-trees [96]), we obtain

$$\hat{\mathcal{P}} \left(\frac{1}{r_{13}r_{24}} \right) = \frac{1}{2r_{13}r_{24}} - \frac{2}{|\mathbf{r}_{13,24} + \mathbf{r}_{13} - \mathbf{r}_{24}| |\mathbf{r}_{13,24} - \mathbf{r}_{13} + \mathbf{r}_{24}|}. \quad (117)$$

It can now be seen that the right-hand side of equation (117) cannot be brought into the form of equation (115), implying that application of the symmetrizer leads to a functional form for which neither the ‘paired’ nor the ‘unpaired’ interspecies distances obey the Bethe–Peierls boundary condition.

Attempting to find a compromise between guidelines (i) and (ii), we write $\psi_T(\mathbf{R})$ as

$$\psi_T(\mathbf{R}) = f_{\text{trap}}(\mathbf{R}) \Phi_\alpha(\mathbf{R}), \quad (118)$$

where $f_{\text{trap}}(\mathbf{R})$ accounts for the external confinement,

$$f_{\text{trap}}(\mathbf{R}) = \exp \left(- \sum_{j=1}^N \frac{r_j^2}{2a_{\text{ho}}^2} \right), \quad (119)$$

and $\Phi_\alpha(\mathbf{R})$ for the correlations,

$$\Phi_\alpha(\mathbf{R}) = (a_{\text{ho}})^{-\alpha N/2(N/2-1)-N} \hat{\mathcal{P}} \left(\frac{\prod_{j=1}^{N/2} \prod_{k=1, k \neq j}^{N/2} (r_{j,N/2+k})^\alpha}{\prod_{j=1}^{N/2} r_{j,N/2+j}} \right). \quad (120)$$

It can be readily checked that the right-hand side of equation (120) has units of $\text{length}^{-3N/2}$, as required for a $3N$ -particle wave function. Since $\Phi_\alpha(\mathbf{R})$ depends only on scalars, i.e. interparticle distances, $\psi_T(\mathbf{R})$ has the desired $(L, \Pi) = (0, +1)$ symmetry. In equation (118), α is an adjustable parameter. For $\alpha = 0$, $\Phi_\alpha(\mathbf{R})$ fulfills the Bethe–Peierls boundary condition for all up–down pairs, provided the symmetrizer $\hat{\mathcal{P}}$ is dropped. A finite value of α reduces the probability of spin up–spin down particles that are ‘not paired’ via the product in the denominator to be close to each other. We pursue two avenues to determine the optimal α . We use results from the literature to fix α , and we determine the optimal α by analyzing our PIGS results.

To determine α using literature results, we rewrite the time-independent Schrödinger equation in terms of the hyperspherical

coordinates R and Ω , where R denotes the hyperradius,

$$R^2 = \sum_{k=1}^N \mathbf{r}_k^2, \quad (121)$$

and Ω the $3N - 1$ hyperangles. Note that the hyperradius R , which is simply given by $|\mathbf{R}|$, is defined without separating off the center-of-mass degrees of freedom. For our purposes, the exact definition of the hyperangles is not important. The key ingredient for our train of thought is that the hyperradial and hyperangular degrees of freedom decouple when the s -wave scattering length is infinitely large [97]. The eigen value of the hyperangular equation [we denote the hyperangular function by $\phi_\nu(\Omega)$] is typically written in terms of s_ν , which—in turn—determines the total energy of the system, $E_{\nu q} = (2q + s_\nu + 1)E_{\text{ho}}$, where $E_{\nu q}$ includes the center-of-mass energy of $3E_{\text{ho}}/2, 5E_{\text{ho}}/2, \dots$, and q is the hyperradial quantum number, which takes the values $q = 0, 1, \dots$. Writing the total wave function ψ as $R^{-(3N-1)/2} F_{\nu q}(R) \phi_\nu(\Omega)$, the hyperradial Schrödinger-like equation reads

$$\left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial R^2} + V_{\text{eff}}(R) + \frac{1}{2} m \omega R^2 \right] F_{\nu q}(R) = E_{\nu q} F_{\nu q}(R), \quad (122)$$

where

$$V_{\text{eff}}(R) = \frac{\hbar^2 (s_\nu^2 - 1/4)}{2mR^2}. \quad (123)$$

Solving the differential equation, one finds

$$F_{\nu q}(R) = \exp\left(-\frac{R^2}{2(a_{\text{ho}})^2}\right) \frac{R^{s_\nu+1/2}}{(a_{\text{ho}})^{s_\nu+1}} L_q^{(s_\nu)}((R^2/(a_{\text{ho}})^2)), \quad (124)$$

where $L_q^{(s_\nu)}$ denotes the associated Laguerre polynomial and $F_{\nu q}(R)$ is normalized according to $\int_0^\infty |F_{\nu q}(R)|^2 dR = 1$. For the ground state ($q = 0, \nu = 0$, and no center-of-mass excitations), the hyperradial solution becomes

$$F_{00}(R) = \exp\left(-\frac{R^2}{2(a_{\text{ho}})^2}\right) \frac{R^{s_0+1/2}}{(a_{\text{ho}})^{s_0+1}}. \quad (125)$$

Comparing the power of a_{ho} on the right-hand side of equation (120) with the power of a_{ho} in equation (125), we deduce

$$\alpha N/2(N/2 - 1) + N = s_0 + 1 \quad (126)$$

or

$$\alpha = \frac{E_{00}/E_{\text{ho}} - N}{N/2(N/2 - 1)}. \quad (127)$$

Using this reasoning, α can be estimated using the ground state energy of the $(N/2, N/2)$ Fermi system. Selected α values, obtained using the energies reported in [87], are shown in column 2 of table 3. Our first set of PIGS calculations use the value of α reported in table 3. In a second set of calculations, α is varied and optimized using the PIGS results themselves. We emphasize that the trial function constructed above provides a fairly descent description of the hyperradial degree of freedom. The hyperangular degrees of freedom are, however, less well described; we return to this point below.




Figure 8. Illustration of the dependence of the PIGS energy on $\Delta\tau$ for the $(2, 2)$ system. The symbols with error bar show the PIGS energy for $\tau = (E_{\text{ho}})^{-1}$, obtained using the trial function given in equations (118)–(120) with $\alpha = 0.505$. The solid and dotted lines show second- and fourth-order fits to the energies (see text for details).

Using $\alpha = 0.505$ and $\tau = (E_{\text{ho}})^{-1}$, the symbols with error bars in figure 8 show the PIGS energies for the $(N/2, N/2) = (2, 2)$ system for four different $\Delta\tau$. The solid line shows a second-order fit of the form $a + b\Delta\tau^2$. The extrapolated $\Delta\tau = 0$ energy is $5.0038(12)E_{\text{ho}}$, where the error bar in brackets represents the fit uncertainty, which takes the error bars of the finite $\Delta\tau$ PIGS energies into account. For comparison, the dotted line shows a fourth-order fit of the form $a + b\Delta\tau^2 + c\Delta\tau^4$. The extrapolated $\Delta\tau = 0$ energy is $5.0069(29)E_{\text{ho}}$, where the error bar in brackets represents—as before—the fit uncertainty, which takes the error bars of the finite $\Delta\tau$ PIGS energies into account. The fact that the extrapolated second- and fourth-order energies agree within error bars suggests that the extrapolated energies are reasonably good. This is confirmed by comparing with the highly accurate energy E_{ECG} obtained via a basis set expansion approach, which employs explicitly correlated Gaussians (see column 6 of table 3). Assuming, for a moment, that the $\tau = (E_{\text{ho}})^{-1}$ result is identical to that for the $\tau \rightarrow \infty$ limit, we can estimate the systematic uncertainties of the second- and fourth-order extrapolations. The fourth-order energy agrees with E_{ECG} (see table 3) within error bars while the second-order energy deviates by about four standard deviations. We thus estimate that the systematic error that originates from the second-order fit is of the order of 0.1%. This suggests that one needs to use the fourth- or an even higher-order extrapolation scheme or perform additional calculations for smaller $\Delta\tau$ if the statistical error is of the order of 0.1% or smaller.

To analyze how close the $\tau = (E_{\text{ho}})^{-1}$ energy is to the $\tau \rightarrow \infty$ limit, circles with error bars in figure 9 show the extrapolated energy, using the fourth-order scheme, of the $(2, 2)$ system for various τ ; as before, α is set to 0.505. As expected, the energy decreases with increasing τ and flattens out for large τ . To better show how the energy behaves with increasing τ , the inset replots the extrapolated $\Delta\tau = 0$ energies with error bars as a function of $1/\tau$. It can be seen that the energies for the two smallest $1/\tau$ (two largest τ) do not agree within error bars. This means that, strictly speaking, the large τ limit has not yet been reached. However, as discussed further below, going to larger τ is rather challenging because




Figure 9. Illustration of the dependence of the PIGS energy on the propagation time τ for the (2, 2) system. The circles and crosses show the extrapolated $\Delta\tau = 0$ energies with error bars (the fourth-order scheme is used), obtained using the trial function given in equations (118)–(120) with $\alpha = 0.505$ and 1, respectively. The inset replots the energies as a function of $1/\tau$.




Figure 10. Illustration of the dependence of the PIGS energy for the (2, 2) system on the trial function parameter α . The symbols show the extrapolated $\Delta\tau = 0$ energies with error bars (the fourth-order scheme is used) for $\tau = 0.125(E_{ho})^{-1}$ as a function of $1/\alpha$, obtained using the trial function given in equations (118)–(120).

of the Fermi sign problem. This implies that, ultimately, the accuracy of the PIGS energy is limited, as already alluded to above, by the systematic error that originates from not going to the $\tau \rightarrow \infty$ limit and not by the statistical error bars. For the $N \geq 6$ systems, the computational time is chosen such that the statistical error is of the order of the estimated systematic error; the reasoning behind this is that a smaller statistical error would not allow one to gain more insight into the exact value of the energy.

For comparison, the crosses with error bars in the main part of figure 9 show the extrapolated (2, 2) PIGS energies for $\alpha = 1$. These energies lie above those for $\alpha = 0.505$ for all τ , reflecting the fact that the trial function with $\alpha = 0.505$ provides a better description of the (2, 2) system than the trial function with $\alpha = 1$. The difference between the PIGS energies for the calculations with the two different α values decreases with increasing τ , reflecting the fact that, in principle, any trial function that has finite overlap with the exact ground state wave function could be used. However, the better ψ_T , the smaller the resulting error bars.

To more systematically investigate the dependence of the PIGS energy, and correspondingly the speed of the




Figure 11. Symbols show the quantity S , equation (110), as a function of τ for the (2, 2) system. The error bars are smaller than the symbol size. The time step $\Delta\tau$ is fixed at $0.125(E_{ho})^{-1}$. As a guide to the eye, the solid line connects consecutive data points.

convergence with increasing τ , on α , we fix τ at $0.125(E_{ho})^{-1}$. These small τ calculations are computationally comparatively inexpensive and hence allow one to survey the α dependence more exhaustively. Ultimately, one needs, of course, to go to larger τ . However, to find the best α (or more generally, the best trial function), it is often times sufficient to consider a relatively small τ . Figure 10 shows the PIGS energy with error bars as a function of α^{-1} for $\tau = 0.125(E_{ho})^{-1}$. The lowest energy is obtained for α around 0.5, confirming our choice of α based on the hyperspherical coordinate approach. This suggests that the optimal α could alternatively be determined iteratively. To this end, let us assume that the ground state energy is unknown. One would then choose an initial value of α to obtain a first PIGS energy estimate for small τ . Using this (non-converged) PIGS energy, one would obtain an improved α value using equation (127) and perform another PIGS calculation. After a few iterations, the optimal α value would be found.

As already alluded to, the Fermi sign problem limits the maximum propagation time τ that can be reached with a finite amount of computational resources. The symbols in figure 11 show the quantity S (see equation (110)), which appears in the denominator of the expression for all observables, as a function of τ for the (2, 2) system for the trial function with $\alpha = 0.505$. For this series of calculations, $\Delta\tau$ is fixed at $\Delta\tau = 0.125(E_{ho})^{-1}$, i.e. the number of time slices increases with increasing τ . The solid line shows a fit to the data, demonstrating that S decreases exponentially with increasing τ or, equivalently, increasing number of time slices. An S value close to 1 indicates that the sign problem is irrelevant. The smaller S , the more severe the sign problem becomes. As a consequence, for a given τ , there exists a maximum $\Delta\tau$ for which the calculation is feasible. For smaller $\Delta\tau$, the errors that originate from the sign problem are too large to be useful. For large τ , the smallest $\Delta\tau$ that can be treated reliably might not be sufficiently small to allow for a reliable extrapolation to $\Delta\tau = 0$. For the (2, 2) system, e.g., $\tau = (E_{ho})^{-1}$ is a good compromise. The excited state contributions have, essentially, decayed and the extrapolation to the $\Delta\tau = 0$ limit is reliable.

In addition to the (2, 2) system, we treat the (3, 3), (4, 4), and (5, 5) systems using the α values determined from the known ground state energies via equation (127). The value of τ (see column 3 of table 3) is chosen such that the estimated systematic error, due to the use of a finite τ , is comparable to or smaller than the error of the extrapolated $\Delta\tau = 0$ energy for this τ (see columns 4–6 of table 3). As N goes up, the propagation time τ is chosen to be smaller and smaller (see table 3). The reason is that the simulations for larger N are more computationally demanding since the Fermi sign problem becomes more severe with increasing N . This means that we are limited by the number of time slices and, correspondingly, the largest τ we can use. The number of time slices $2n + 1$ considered in the $\Delta\tau \rightarrow 0$ extrapolation are chosen based on our detailed analysis of the (2, 2) system. Because of the relatively small τ considered, the energies reported for the $N/2 = 4$ and 5 systems should be regarded, within error bars, as variational upper bounds. Our (3, 3) energy agrees, within error bars, with the basis set expansion energy E_{ECG} but lies slightly below the diffusion Monte Carlo energy E_{DMC} . For the (4, 4) and (5, 5) systems, the PIGS energies agree, within error bars, with E_{ECG} and E_{DMC} . For the (5, 5) system, we performed an additional calculation using—as before— $\tau = 0.125(E_{\text{ho}})^{-1}$ but using an α value that is larger than that listed in table 3, namely $\alpha = 0.335$. The extrapolated $\Delta\tau = 0$ PIGS energy is $16.22(14)E_{\text{ho}}$, which agrees within error bars with our result listed in table 3. The larger error bar reflects the fact that the larger α value provides a less good trial function.

The trial function used so far (see equations (118)–(120)) contains a single adjustable parameter, namely α , that primarily determines the correlations in the hyperangular degree of freedom. Our goal is now to design a trial function that provides an improved description of the hyperangular degrees of freedom. In doing so, we are guided by the analytically known wave function of the harmonically trapped (2, 1) system with $(L, \Pi) = (0, +1)$ symmetry at unitarity [98]. The hyperangular part of the wave function that yields the lowest energy with $(0, +1)$ symmetry is proportional to

$$(1 - \hat{P}_{12})\sin(\bar{s}_0(\theta_1 - \pi/2))/\sin(2\theta_1), \quad (128)$$

where $\theta_1 = \arcsin(r_{13}/(2\sqrt{\bar{R}}))$, $\bar{s}_0 = 2.166$, and $\bar{R}^2 = \sum_{j < k} r_{jk}^2/N$. Application of \hat{P}_{12} changes θ_1 into θ_2 , where $\theta_2 = \arcsin(r_{23}/(2\sqrt{\bar{R}}))$. The factor $\sin(\bar{s}_0(\theta_1 - \pi/2))$ enhances the probability to find two particles at vanishing hyperangle θ_1 , i.e. at vanishing distance between the unlike particles 1 and 3. In the non-interacting limit, \bar{s}_0 is equal to 4, which implies that the probability to find two particles at vanishing hyperangle θ_1 vanishes. The quantities \bar{R} and \bar{s}_0 , which are defined by excluding the center-of-mass degrees of freedom, are related to R and s_0 , $\bar{R}^2 = R^2 - N(R_{\text{cm}})^2$ and $\bar{s}_0 = s_0 - 3/2$.

Motivated by the hyperangular wave function of the (2, 1) system with $(0, +1)$ symmetry at unitarity, we consider the following alternative form of the trial function,

$$\psi_{\tau} = f_{\text{trap}}(\mathbf{R})\Phi_{\alpha,\beta,\gamma}(\mathbf{R}), \quad (129)$$

where

$$\Phi_{\alpha,\beta,\gamma}(\mathbf{R}) = \frac{\bar{R}^{\alpha}}{(a_{\text{ho}})^{\alpha+3N/2}} \hat{P} \prod_{j=1}^{N/2} \frac{\sin[\beta(\theta_j - \pi/2)](\cos\theta_j)^{\gamma+1}}{\sin(2\theta_j)} \quad (130)$$

and $\theta_j = \arcsin(r_{j,N/2+j}/(2\sqrt{\bar{R}}))$. The factor $(\cos\theta_j)^{\gamma+1}$ is introduced to increase the tunability of the trial function. As before, the optimal value of α is obtained by matching to the known hyperradial solution. This yields $\alpha = 0.166$. The values of β and γ , in contrast, are determined by performing PIGS simulations for small τ .

Considering the (2, 2) system and using $\alpha = 2.5091$, $\beta = 1.1$, and $\gamma = 2$, we obtain the extrapolated $\Delta\tau = 0$ PIGS energy of $5.022(3)E_{\text{ho}}$ for $\tau = 0.0625(E_{\text{ho}})^{-1}$. This energy is significantly lower than the PIGS energy of $5.193(4)E_{\text{ho}}$ that we obtained using the correlation factor $\Phi_{\alpha}(\mathbf{R})$ for the same τ . For $\tau = 0.25(E_{\text{ho}})^{-1}$, we obtain $5.011(8)E_{\text{ho}}$, which agrees within error bars with the ground state energy E_{ECG} (see table 3). Considering the (3, 3) system and using the trial parameters $\alpha = 5.837$, $\beta = 1.5$, and $\gamma = 8$, we obtain the extrapolated $\Delta\tau = 0$ PIGS energy of $8.357(8)E_{\text{ho}}$ for $\tau = 0.25(E_{\text{ho}})^{-1}$. For comparison, the (3, 3) PIGS energy reported in table 3 is for a larger τ , namely $\tau = 0.5(E_{\text{ho}})^{-1}$. For the (4, 4) and (5, 5) systems, the trial function with the correlation factor $\Phi_{\alpha,\beta,\gamma}$ did not yield an improved energy compared to that for Φ_{α} . The reason could be that the trial parameters were not fully optimized or that the degrees of freedom of the larger systems are less well described by the trial function (e.g., that three- and higher-body correlations are needed).

In addition to the energies, we use the PIGS approach to calculate structural properties. The scaled pair distribution functions reported below are obtained for a finite $\Delta\tau$; no extrapolation to the $\Delta\tau = 0$ limit was performed. To determine a suitable $\Delta\tau$, we consider the (2, 2) system and perform calculations for $\tau = 0.25(E_{\text{ho}})^{-1}$ using three different $\Delta\tau$, i.e. $\Delta\tau = 0.25(E_{\text{ho}})^{-1}$, $0.125(E_{\text{ho}})^{-1}$, and $0.0625(E_{\text{ho}})^{-1}$. We find that the scaled pair distribution function for $\Delta\tau = 0.25(E_{\text{ho}})^{-1}$ differs slightly from those for $\Delta\tau = 0.125(E_{\text{ho}})^{-1}$ and $\Delta\tau = 0.0625(E_{\text{ho}})^{-1}$. However, no visual difference is observed between the scaled pair distribution functions for $\Delta\tau = 0.125(E_{\text{ho}})^{-1}$ and $\Delta\tau = 0.0625(E_{\text{ho}})^{-1}$. Motivated by this observation, we calculate the scaled pair distribution functions for the spin-balanced systems with $N/2 \leq 5$ using $\Delta\tau = 0.125(E_{\text{ho}})^{-1}$. The propagation times and trial functions are the same as those used to obtain the energies reported in table 3.

Solid lines in figures 12(a)–(d) show the resulting scaled pair distribution function for the spin-balanced systems with $N = 4$ –10. For comparison, the dashed lines show the scaled pair distribution function obtained from basis set calculations for an attractive two-body Gaussian potential with infinitely large s -wave scattering length and effective range of approximately $0.12a_{\text{ho}}$ (the range is $0.06a_{\text{ho}}$) [87]. The Gaussian potential used supports exactly one zero-energy two-body bound state in free space. If the effective range were taken to zero, the two different approaches should yield the




Figure 12. Scaled pair distribution functions $4\pi r^2 P_{12}(r)$ for the (a) (2, 2), (b) (3, 3), (c) (4, 4), and (d) (5, 5) systems at unitarity. The solid lines show our PIGS results for the model Hamiltonian with two-body zero-range interactions. The results are obtained using the propagation times and trial function parameters listed in table 3. The imaginary time step is set to $\Delta\tau = 0.125(E_{\text{ho}})^{-1}$. For comparison, the dashed lines show basis set expansion based results [87] for an attractive two-body Gaussian interaction with finite effective range (see text for details).

Table 4. Contact for spin-balanced two-component Fermi gas ($N/2 \leq 5$) with zero-range interactions at unitarity. Column 2 reports the value of the contact C_{PIGS} obtained from our PIGS simulations with $\Delta\tau = 0.125(E_{\text{ho}})^{-1}$. The error reported does not include the systematic error introduced by not extrapolating to zero $\Delta\tau$. Column 3 shows the contact C_{ECG} from the literature, obtained by extrapolating the contact obtained by the ECG approach for finite-range interactions to the zero-range limit [87].

(n_1, n_2)	$C_{\text{PIGS}} a_{\text{ho}}$	$C_{\text{ECG}} a_{\text{ho}}$
(2, 2)	25.91(6)	25.74(1)
(3, 3)	42.26(5)	40.39(8)
(4, 4)	59.5(2)	55.4(5)
(5, 5)	79.3(3)	72.3(8)

same result. As can be seen from figure 12, the zero-range results deviate a bit from the finite-range results at small interparticle distances; in particular, the scaled pair distribution functions take a finite value for vanishing interparticle distance if the zero-range interaction model is used and go to zero if the finite-range interaction model is used. At larger interparticle distances, the agreement between the dashed and solid lines is quite good, suggesting that the PIGS scaled pair distribution functions are, indeed, quite well converged for $r \gtrsim 0.5a_{\text{ho}}$.

To check the convergence at small interparticle distances, we report the contact C obtained from the $r = 0$ value of the scaled pair distribution function [99],

$$C = 4\pi n_1 n_2 \lim_{r \rightarrow 0} 4\pi r^2 P_{12}(r). \quad (131)$$

The resulting contacts are summarized in table 4. The contact calculated by the PIGS approach is, for all N considered, larger than the contact calculated by extrapolating basis set expansion results (namely, the slope of the energy with respect to the inverse of the s -wave scattering length) for finite-range interactions to the zero-range limit. As $N/2$ increases from 2 to 5, the difference between the two sets of results increases from about two to about six standard deviations. The numerical

determination of the contact through the scaled pair distribution function is, in general, quite challenging since the contact probes—if determined in this manner—a small portion of the Hilbert space. As a consequence, the convergence of the contact can be slow. It is presently unclear which set of results is more reliable and if, possibly, the errorbars of the results obtained by either of the two methods was underestimated: the basis set expansion results may be ‘contaminated’ by basis set extrapolation and zero-range extrapolation errors while the PIGS results may be ‘contaminated’ by finite propagation time and finite time step errors.

5.3. Non-interacting Fermi gas with a single impurity

This section considers a non-interacting Fermi gas with a single impurity, i.e. the $(n_1, n_2) = (N - 1, 1)$ system with $N = 3-5$. In a zeroth-order approximation, one can think of this system as consisting of one up-down pair and $N - 2$ unpaired spin-up atoms. Of course, this picture needs to be refined to account for the fact that the system contains $N - 1$ identical fermions and for the fact that the presence of additional spin-up fermions ‘disturbs’ the single pair. Nevertheless, the simple picture correctly suggests that the ground state of the $(N - 1, 1)$ system does not have $(L, \Pi) = (0, +1)$ symmetry. Rather, the ground state of the (2, 1), (3, 1) and (4, 1) systems has $(L, \Pi) = (1, -1)$, $(1, +1)$ and $(0, -1)$ symmetry [98, 100, 101]. Roughly, this can be understood by realizing that the (2, 1), (3, 1) and (4, 1) systems contain one, two and three unpaired spin-up atoms, each of which carry one quantum of angular momentum (the p -shell is being filled). Thus, due to the angular momentum carried by the single unpaired spin-up atom the (2, 1) system has $(1, -1)$ symmetry. In the (3, 1) system, the angular momenta of each of the two unpaired atoms can couple to an angular momentum 0, 1, or 2, with the parity being even. The (1, +1) channel turns out to have the lowest energy [101]. Last, in the (4, 1) system, the angular momenta of each of the three unpaired atoms can couple to an angular momentum 0, 1, 2, or 3, with the parity being odd. Consistent with the idea of a closed shell configuration, the $(0, -1)$ channel turns out to have the lowest energy.

As our first application of the PIGS approach to spin-imbalanced systems, we treat the (2, 1) system with $(L, \Pi) = (0, +1)$ symmetry. This application illustrates that the PIGS approach can be used to describe the energetically lowest-lying state (not the ‘absolute ground state’) of a given symmetry. Following the logic that motivated the trial function given in equations (118)–(120), we write

$$\psi_{\mathcal{T}}(\mathbf{R}) = f_{\text{trap}}(\mathbf{R}) \Phi_{\alpha'}(\mathbf{R}), \quad (132)$$

where

$$\Phi_{\alpha'}(\mathbf{R}) = (a_{\text{ho}})^{-\alpha'(N-2)-3N/2+1} \hat{p} \left(\frac{\prod_{j=2}^{N-1} (r_{j,N})^{\alpha'}}{r_{1,N}} \right). \quad (133)$$

Using a relatively small τ , namely $\tau = 0.25(E_{\text{ho}})^{-1}$, table 5 shows the extrapolated $\Delta\tau = 0$ energy. It lies about six sigma above the exact energy obtained within the hyperspherical

Table 5. Spin-imbalanced fermionic (2, 1) system with zero-range interactions at unitarity. Columns 2–4 report the propagation time, the n used (the number of time slices is $2n+1$), and the trial function employed for two different symmetries (see column 1). Column 5 shows the resulting PIGS energy with error bars. For the (0, +1) channel, no extrapolation to the $\Delta\tau = 0$ limit was performed. For the (1, -1) channel, a second-order extrapolation was used. For comparison, column 6 shows the exact zero-range energies obtained using the formalism developed in [98].

(L, Π)	τE_{ho}	n used	Trial function	$E_{\text{PIGS}}/E_{\text{ho}}$	$E_{\text{exact}}/E_{\text{ho}}$
(0, +1)	0.25	4	Equation (132), $\alpha = 1.66622$	4.687(4)	4.66622
(0, +1)	0.25	4	Equation (128)	4.676(10)	4.66622
(1, -1)	0.5	2,4	Equation (134), $\alpha = 0.772724$	4.276(9)	4.27272

Table 6. Spin-imbalanced fermionic ($N - 1, 1$) system ($N = 4$ and 5) with zero-range interactions at unitarity. Column 1 lists the ($N - 1, 1$) system considered. The symmetry of the ground state is reported in column 2. Columns 3, 4, 5, and 6 report the propagation time, the n used ($2n+1$ is the number of time slices included in the extrapolation of the energy to $\Delta\tau = 0$), the value of α , and the equation number of the trial function used. Column 7 shows the resulting extrapolated $\Delta\tau = 0$ PIGS energy with error bars, obtained using a second-order extrapolation. For comparison, column 8 reports energies obtained using the explicitly correlated Gaussian approach [101, 102]; the energies, which are denoted by E_{ECG} , are obtained by extrapolating a series of finite-range energies to the zero-range limit. Column 9 reports the (4, 1) energy E_{DMC} for a square-well potential with range $r_0 = 0.01a_{\text{ho}}$ obtained using the diffusion Monte Carlo method [103]; no extrapolation to the zero-range limit was done. The uncertainties for the ECG and DMC calculations are, according to [101–103], in the last digit reported.

(L, Π)	τE_{ho}	n	α	ψ_T	$E_{\text{PIGS}}/E_{\text{ho}}$	$E_{\text{ECG}}/E_{\text{ho}}$	$E_{\text{DMC}}/E_{\text{ho}}$
(3, 1)	(1, +1)	0.5	4, 8	0.791 (135)	6.60(1)	6.5819	
(4, 1)	(0, -1)	0.5	2, 4	0.667 (136)	8.93(7)	8.95	8.93
(4, 1)	(0, -1)	1	4, 8	0 (137)	8.92(4)	8.95	8.93

coordinate approach [98]. If we repeated the calculation for larger τ , we would expect to obtain a PIGS energy closer to the exact energy. To prove that the lowest (0, +1) energy can be obtained exactly within the PIGS approach, i.e. to prove that the PIGS approach does, indeed, preserve the symmetry of the trial function, we use the exact (analytically known) eigen state [98] as the trial function. The resulting extrapolated $\Delta\tau = 0$ energy (see table 5) agrees to within error bars with the exact energy.

As a proof-of-principle, we apply the PIGS approach to the (2, 1) system with $(L, \Pi) = (1, -1)$ symmetry at unitarity. Knowing that the orbital angular momentum is carried by the Jacobi vector $\mathbf{r}_{13,2}$ when particles 1 and 3 form a pair and by the Jacobi vector $\mathbf{r}_{23,1}$ when particles 2 and 3 form a pair [98, 100], we write

$$\psi_T(\mathbf{R}) = (a_{\text{ho}})^{-9/2} f_{\text{trap}}(\mathbf{R}) \left(\frac{\bar{R}}{a_{\text{ho}}} \right)^\alpha \hat{\mathcal{P}} \left(\frac{x_{13,2}}{r_{13,2}} \right). \quad (134)$$

Owing to the three-fold degeneracy of $L = 1$ states, alternatively one can use $y_{13,2}$ or $z_{13,2}$ instead of $x_{13,2}$. In an equivalent formulation, the term $x_{13,2}/r_{13,2}$ in equation (134) is replaced by $Y_{1,m_l}(\hat{\mathbf{r}}_{13,2})$, where Y_{1,m_l} denotes the spherical harmonic and m_l can take the values ± 1 and 0. Since the spherical harmonics with $m_l \neq 0$ are complex, spherical harmonics are less convenient from a numerical/implementation perspective than the real version used in equation (134). Using equation (134) with $\tau = 0.5(E_{\text{ho}})^{-1}$, we find the extrapolated $\Delta\tau = 0$ energy $E_{\text{PIGS}} = 4.276(9)E_{\text{ho}}$, which agrees, within error bars, with the exact zero-range energy from [98] (see row 3 of table 5).

The above trial function can be extended to the (3, 1) system. Assuming the formation of a pair consisting of atoms 1 and 4, one quantum of orbital angular momentum each is assumed to be carried by the vectors $\mathbf{r}_{14,2}$ and $\mathbf{r}_{14,3}$. Coupling $Y_{1,m_1}(\mathbf{r}_{14,2})$ and $Y_{1,m_2}(\mathbf{r}_{14,3})$ such that the resulting function has $(L, \Pi) = (1, -1)$ symmetry, we obtain the desired correlation factor. Since we prefer to work with real quantities, we write

$$\psi_T(\mathbf{R}) = (a_{\text{ho}})^{-6} f_{\text{trap}}(\mathbf{R}) \left(\frac{\bar{R}}{a_{\text{ho}}} \right)^\alpha \hat{\mathcal{P}} \left(\frac{(\mathbf{r}_{14,2} \times \mathbf{r}_{14,3}) \cdot \hat{\mathbf{z}}}{r_{14,2} r_{14,3}} \right), \quad (135)$$

where the dot product serves to select the z -component of the vector that results when taking the cross product. Instead of the z -component, the x - or y -components can be used. For the (4, 1) system, we use

$$\psi_T(\mathbf{R}) = (a_{\text{ho}})^{-15/2} f_{\text{trap}}(\mathbf{R}) \left(\frac{\bar{R}}{a_{\text{ho}}} \right)^\alpha \hat{\mathcal{P}} \left(\frac{(\mathbf{r}_{15,2} \times \mathbf{r}_{15,3}) \cdot \mathbf{r}_{15,4}}{r_{15,2} r_{15,3} r_{15,4}} \right), \quad (136)$$

which has the desired $(L, \Pi) = (0, -1)$ symmetry. Alternatively, one could use

$$\psi_T(\mathbf{R}) = (a_{\text{ho}})^{-21/2} f_{\text{trap}}(\mathbf{R}) \left(\frac{\bar{R}}{a_{\text{ho}}} \right)^\alpha \hat{\mathcal{P}}((\mathbf{r}_{15,2} \times \mathbf{r}_{15,3}) \cdot \mathbf{r}_{15,4}). \quad (137)$$

The resulting extrapolated $\Delta\tau = 0$ energies for the (3, 1) and (4, 1) systems are reported in table 6. The (3, 1) PIGS energy deviates by two sigma from the highly accurate basis set expansion energy E_{ECG} . The small disagreement may be due to the fact that τ is not quite large enough or that the error bar of




Figure 13. Symbols show the contact C of the (3, 1) system as a function of τ^{-1} . The time step $\Delta\tau$ is fixed at $0.125(E_{ho})^{-1}$. As a guide to the eye, the symbols are connected by a solid line. For comparison, the dashed horizontal line shows the contact $C = 10.84(2)(a_{ho})^{-1}$, obtained by extrapolating basis set expansion results for finite-range interactions to the zero-range limit.

the extrapolated energy is, in fact, slightly larger than what is reported in table 6. The (4, 1) PIGS energies for $\tau = 0.5(E_{ho})^{-1}$ and $\tau = (E_{ho})^{-1}$ agree, within error bars, with the energy E_{DMC} . Since the diffusion Monte Carlo energy was not extrapolated to the zero-range limit, the true zero-range energy is probably somewhat smaller than E_{DMC} .

In addition to the energy, we determine the contact C from the $r = 0$ behavior of the scaled pair distribution function (see equation (131)). In general, the convergence rate of the energy and that of other observables can be different. For the case at hand, namely the contact, this can be understood by realizing that only a small fraction of the wave function amplitude is located at small r . Thus, while the energies shown in table 6 appear to be converged, the contact may not be. Indeed, this is what we find. Figure 13 shows the contact of the (3, 1) system as a function of τ^{-1} . The calculations are performed for $\Delta\tau = 0.125(E_{ho})^{-1}$. The contacts for the two largest τ values considered differ by roughly 2%. For even larger τ , the curve should flatten out. Thus, we can interpret our calculations as providing a lower bound on C , $C \geq 10.7(a_{ho})^{-1}$. Indeed, calculations that employ a correlated Gaussian basis set yield a contact of $C = 10.84(2)(a_{ho})^{-1}$ (see the dashed horizontal line in figure 13), which is close to the PIGS contact for the largest τ considered. For $\tau = 0.5(E_{ho})^{-1}$, we performed an additional calculation for a smaller $\Delta\tau$, i.e. for $\Delta\tau = 0.0625(E_{ho})^{-1}$. The results for $\Delta\tau = 0.125(E_{ho})^{-1}$ and $\Delta\tau = 0.0625(E_{ho})^{-1}$ differ by 0.4%, which is small compared to the error introduced by not extrapolating to the $\tau = \infty$ limit.

6. Summary and outlook

This article provided a detailed introduction to the path integral ground state Monte Carlo (PIGS) algorithm. Implementation details and convergence properties were discussed using general arguments and subsequently illustrated for selected systems and observables. While the primary focus was on the PIGS approach, several aspects are applicable more broadly. For example, the discussion of the error analysis is relevant to essentially all Monte Carlo algorithms and

many features of the path generation apply also to the finite-temperature path integral Monte Carlo approach.

The PIGS algorithm takes a trial function, which is selected by the simulator, and propagates it in imaginary time. For sufficiently large imaginary time τ , the lowest eigen state of the system Hamiltonian, which has finite overlap with the trial function, is being projected out. A crucial question is how to assess whether the resulting eigen energy and other observables are truly converged. This question is particularly pressing for fermionic systems, for which the largest τ considered is restricted by numerical instabilities due to the Fermi sign problem. In the absence of identical fermions, the convergence analysis is relatively simple since there are essentially no restrictions on the τ that can be considered. The applications to fermions considered in this tutorial employed a multi-faceted approach to the convergence analysis. For small systems, comparisons with established literature results were used as a benchmark. For larger systems, an analysis of the error bars was used to establish where the Fermi sign problem sets in. For τ not noticeably impacted by the Fermi sign problem, the resulting energies, extrapolated to the infinite time slice limit, provided variational upper bounds. Additionally, the calculations were performed for different trial functions and the runs were checked for consistency. Ultimately, there is no guarantee that the resulting observables are not biased by the trial function. However, the various checks provide one with tools for (roughly) estimating and minimizing the variational bias.

The sample applications presented concern strongly correlated Fermi gases. In cold atom experiments, the two-body van der Waals length is typically much smaller than the average interparticle spacing and the two-body s -wave scattering length (for two-component Fermi systems, this is the interspecies (and not the intraspecies) scattering length). Thus, cold atom systems realize, to a very good approximation, idealized systems in which the two-body interaction range is zero. From a theoretical point of view, a vanishing two-body range is particularly interesting as this implies that the range drops out of the problem. For infinitely large two-body s -wave scattering length, e.g., the system exhibits a scale invariance, reflecting underlying symmetries of the Hamiltonian. While scale-invariance based arguments and formulations have led to a great deal of insight into these paradigmatic, strongly correlated systems, few analytical or numerical techniques exist that can reliably predict the energy, Tan contact, superfluid fraction, or other observables. The PIGS approach treats two-body zero-range interactions by building the exact two-body Bethe–Peierls boundary condition into the propagator using the pair product approximation. For fermions, the treatment is limited to small number of particles since the Fermi sign problem becomes exponentially more severe with increasing number of identical fermions. The system sizes considered in this article are the same as those that have been treated by the explicitly correlated Gaussian basis set expansion approach [87]. For both (the PIGS and basis set expansion approach), the computational effort increases tremendously as N is increased beyond what is considered in this article.

Future applications of the PIGS approach to systems with zero-range interactions may include unequal-mass two-component Fermi gases, Fermi gases in non-spherically symmetric external traps (including effectively low-dimensional systems), or Bose droplets without and with an impurity. These applications can be tackled with the technology already developed. An interesting and challenging future development is the treatment of spin-orbit coupled systems, where the spatial degrees of freedom are coupled to the spin degrees of freedom. It will be interesting to marry the treatment of spin degrees of freedom with the use of two-body zero-range interactions.

Acknowledgments

We are grateful to Xiangyu Yin for providing the scaled pair distribution functions from [87] (dashed lines in figure 12) in tabular form. Support by the National Science Foundation (NSF) through Grant No. PHY-1415112 is gratefully acknowledged. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF Grant No. OCI-1053575, and the WSU HPC.

References

- [1] Binder K 1986 *Monte Carlo Methods in Statistical Physics* (Berlin: Springer)
- [2] Hetényi B, Rabani E and Berne B J 1999 Path-integral diffusion Monte Carlo: calculation of observables of many-body systems in the ground state *J. Chem. Phys.* **110** 6143
- [3] Sarsa A, Schmidt K E and Magro W R 2000 A path integral ground state method *J. Chem. Phys.* **113** 1366
- [4] Cuervo J E, Roy P-N and Boninsegni M 2005 Path integral ground state with a fourth-order propagator: application to condensed helium *J. Chem. Phys.* **122** 114504
- [5] Rota R, Casulleras J, Mazzanti F and Boronat J 2010 High-order time expansion path integral ground state *Phys. Rev. E* **81** 016707
- [6] Ceperley D M 1995 Path integrals in the theory of condensed helium *Rev. Mod. Phys.* **67** 279
- [7] Hammond B L, Lester W A Jr and Reynolds P J 1994 *Monte Carlo Methods in Ab Initio Quantum Chemistry (Lecture and Course Notes In Chemistry Series)* vol 1 (Singapore: World Scientific)
- [8] Press W 2007 *Numerical Recipes: The Art of Scientific Computing* 3rd edn (Cambridge: Cambridge University Press)
- [9] Boninsegni M 2005 Permutation sampling in path integral Monte Carlo *J. Low Temp. Phys.* **141** 27
- [10] Whitley H D, DuBois J L and Whaley K B 2009 Spectral shifts and helium configurations in $^4\text{He}_N$ -tetracene clusters *J. Chem. Phys.* **131** 124514
- [11] Guardiola R 1988 Monte Carlo techniques in the many body problem *First International Course on Condensed Matter* ed D P S Rosati and G Violini (Singapore: World Scientific)
- [12] Casulleras J and Boronat J 1995 Unbiased estimators in quantum Monte Carlo methods: application to liquid ^4He *Phys. Rev. B* **52** 3654
- [13] Reynolds P J, Barnett R N, Hammond B L and Lester W A Jr 1986 Molecular physics and chemistry applications of quantum Monte Carlo *J. Stat. Phys.* **43** 1017
- [14] Moroni S and Boninsegni M 2004 Condensate fraction in liquid ^4He *J. Low Temp. Phys.* **136** 129
- [15] Herdman C M, Inglis S, Roy P-N, Melko R G and Maestro A Del 2014 Path-integral Monte Carlo method for Rényi entanglement entropies *Phys. Rev. E* **90** 013308
- [16] Hastings M B, González I, Kallin A B and Melko R G 2010 Measuring renyi entanglement entropy in quantum Monte Carlo simulations *Phys. Rev. Lett.* **104** 157201
- [17] Galli D E and Reatto L 2004 Disorder phenomena in quantum solids with vacancies *J. Low Temp. Phys.* **134** 121
- [18] Vitali E, Rossi M, Tramonto F, Galli D E and Reatto L 2008 Path-integral ground-state Monte Carlo study of two-dimensional solid ^4He *Phys. Rev. B* **77** 180505
- [19] Rossi M, Vitali E, Reatto L and Galli D E 2012 Microscopic characterization of overpressurized superfluid ^4He *Phys. Rev. B* **85** 014525
- [20] Galli D E, Reatto L and Saslow W M 2007 Bounds for the superfluid fraction from exact quantum Monte Carlo local densities *Phys. Rev. B* **76** 052503
- [21] Rossi M, Vitali E, Galli D E and Reatto L 2010 Quantum dislocations: the fate of multiple vacancies in two-dimensional solid ^4He *J. Phys. Condens. Matter* **22** 145401
- [22] Nava M, Galli D E, Moroni S and Vitali E 2013 Dynamic structure factor for ^3He in two dimensions *Phys. Rev. B* **87** 144506
- [23] Arrigoni F, Vitali E, Galli D E and Reatto L 2013 Excitation spectrum in two-dimensional superfluid ^4He *Low Temp. Phys.* **39** 793
- [24] Carleo G, Moroni S and Baroni S 2009 Zero-temperature dynamics of solid ^4He from quantum Monte Carlo simulations *Phys. Rev. B* **80** 094301
- [25] Galli D E and Reatto L 2004 The shadow path integral ground state method: study of confined solid ^4He *J. Low Temp. Phys.* **136** 343
- [26] Nava M, Galli D E, Cole M W and Reatto L 2012 Adsorption of He isotopes on fluorographene and graphane: Fluid and superfluid phases from quantum Monte Carlo calculations *Phys. Rev. B* **86** 174509
- [27] Turnbull J D and Boninsegni M 2005 Adsorption of parahydrogen on fullerenes *Phys. Rev. B* **71** 205421
- [28] Turnbull J and Boninsegni M 2007 Molecular hydrogen isotopes adsorbed on krypton-preplated graphite *Phys. Rev. B* **76** 104524
- [29] Cuervo J E and Roy P-N 2006 Path integral ground state study of finite-size systems: application to small (parahydrogen) $_N$ ($N = 2-20$) clusters *J. Chem. Phys.* **125** 124314
- [30] Cuervo J E and Roy P-N 2008 On the solid- and liquidlike nature of quantum clusters in their ground state *J. Chem. Phys.* **128** 224509
- [31] Cuervo J E and Roy P-N 2009 Weakly bound complexes trapped in quantum matrices: structure, energetics, and isomer coexistence in (para- H_2) $_N$ (ortho- D_2) $_3$ clusters *J. Chem. Phys.* **131** 114302
- [32] Rossi M and Salasnich L 2013 Path-integral ground state and superfluid hydrodynamics of a bosonic gas of hard spheres *Phys. Rev. A* **88** 053617
- [33] Rota R, Tramonto F, Galli D E and Giorgini S 2013 Quantum Monte Carlo study of the dynamic structure factor in the gas and crystal phase of hard-sphere bosons *Phys. Rev. B* **88** 214505
- [34] Macia A, Hufnagl D, Mazzanti F, Boronat J and Zillich R E 2012 Excitations and stripe phase formation in a two-dimensional dipolar Bose Gas with tilted polarization *Phys. Rev. Lett.* **109** 235307

- [35] Moroni S and Boninsegni M 2014 Coexistence, interfacial energy, and the fate of microemulsions of 2D dipolar Bosons *Phys. Rev. Lett.* **113** 240407
- [36] Pilati S and Troyer M 2012 Bosonic superfluid-insulator transition in continuous space *Phys. Rev. Lett.* **108** 155301
- [37] Giorgini S, Pitaevskii L P and Stringari S 2008 Theory of ultracold atomic Fermi gases *Rev. Mod. Phys.* **80** 1215
- [38] Blume D 2012 Few-body physics with ultracold atomic and molecular systems in traps *Rep. Prog. Phys.* **75** 046401
- [39] Levinsen J, Massignan P, Endo S and Parish M M 2017 Universality of the unitary Fermi gas: a few-body perspective *J. Phys. B: At. Mol. Opt. Phys.* **50** 072001
- [40] Heiselberg H 2001 Fermi systems with long scattering lengths *Phys. Rev. A* **63** 043606
- [41] Werner F and Castin Y 2012 General relations for quantum gases in two and three dimensions: two-component fermions *Phys. Rev. A* **86** 013626
- [42] Casula M, Ceperley D M and Mueller E J 2008 Quantum Monte Carlo study of one-dimensional trapped fermions with attractive contact interactions *Phys. Rev. A* **78** 033607
- [43] Piątek S and Krauth W 2014 Efimov-driven phase transitions of the unitary Bose gas *Nat. Commun.* **5** 3503
- [44] Pessoa R, Vitiello S A and Schmidt K E 2015 Monte Carlo calculations for Fermi gases in the unitary limit with a zero-range interaction *J. Low Temp. Phys.* **180** 168
- [45] Yan Y and Blume D 2015 Incorporating exact two-body propagators for zero-range interactions into N-body Monte Carlo simulations *Phys. Rev. A* **91** 043607
- [46] Pessoa R, Gandolfi S, Vitiello S A and Schmidt K E 2015 Contact interaction in a unitary ultracold Fermi gas *Phys. Rev. A* **92** 063625
- [47] Peña Ardila L A and Giorgini S 2016 Bose polaron problem: effect of mass imbalance on binding energy *Phys. Rev. A* **94** 063640
- [48] Parisi L and Giorgini S 2017 Quantum Monte Carlo study of the Bose-polaron problem in a one-dimensional gas with contact interactions *Phys. Rev. A* **95** 023619
- [49] Lin Y-J, Jiménez-García K and Spielman I B 2011 Spin-orbit-coupled Bose-Einstein condensates *Nature* **471** 83
- [50] Dalibard J, Gerberier F, Juzeliūnas G and Öhberg P 2011 Colloquium: Artificial gauge potentials for neutral atoms *Rev. Mod. Phys.* **83** 1523
- [51] Zhai H 2015 Degenerate quantum gases with spin-orbit coupling: a review *Rep. Prog. Phys.* **78** 026001
- [52] Carlson J, Schmidt K E and Kalos M H 1987 Microscopic calculations of ^3He with realistic interactions *Phys. Rev. C* **36** 27
- [53] Shumway J and Ceperley D M 2001 Quantum Monte Carlo treatment of elastic exciton-exciton scattering *Phys. Rev. B* **63** 165209
- [54] Shumway J and Ceperley D M 2005 Quantum Monte Carlo simulations of exciton condensates *Solid State Commun.* **134** 19
- [55] Nollett K M, Pieper S C, Wiringa R B, Carlson J and Hale G M 2007 Quantum Monte Carlo calculations of neutron- α scattering *Phys. Rev. Lett.* **99** 022502
- [56] Magnus W 1954 On the exponential solution of differential equations for a linear operator *Commun. Pure Appl. Math.* **7** 649
- [57] Reynolds P J, Ceperley D M, Alder B J and Lester W A Jr 1982 Fixed-node quantum Monte Carlo for molecules *J. Chem. Phys.* **77** 5593
- [58] Ceperley D M and Alder B J 1980 Ground state of the electron gas by a stochastic method *Phys. Rev. Lett.* **45** 566
- [59] Krauth W 2006 *Statistical Mechanics: Algorithms and Computations (Oxford Master Series in Physics)* (Oxford: Oxford University Press)
- [60] Gaveau B and Schulman L S 1986 Explicit time-dependent Schrödinger propagators *J. Phys. A: Math. Gen.* **19** 1833
- [61] Lawande S V and Bhagwat K V 1988 Feynman propagator for the δ -function potential *Phys. Lett. A* **131** 8
- [62] Blinder S M 1988 Green's function and propagator for the one-dimensional δ -function potential *Phys. Rev. A* **37** 973
- [63] Wódkiewicz K 1991 Fermi pseudopotential in arbitrary dimensions *Phys. Rev. A* **43** 68
- [64] Chandler D and Wolynes P G 1981 Exploiting the isomorphism between quantum theory and classical statistical mechanics of polyatomic fluids *J. Chem. Phys.* **74** 4078
- [65] Trotter H F 1959 On the product of semi-groups of operators *Proc. Am. Math. Soc.* **10** 545
- [66] Rossmann W 2006 *Lie Groups: An Introduction Through Linear Groups (Oxford Graduate Texts in Mathematics)* (Oxford: Oxford University Press)
- [67] Chin S A 1997 Symplectic integrators from composite operator factorizations *Phys. Lett. A* **226** 344
- [68] Suzuki M 1995 Hybrid exponential product formulas for unbounded operators with possible applications to Monte Carlo simulations *Phys. Lett. A* **201** 425
- [69] Barker J A 1979 A quantum-statistical Monte Carlo method; path integrals with boundary conditions *J. Chem. Phys.* **70** 2914
- [70] Cao J and Berne B J 1992 A new quantum propagator for hard sphere and cavity systems *J. Chem. Phys.* **97** 2382
- [71] Chin S A 2005 Structure of positive decompositions of exponential operators *Phys. Rev. E* **71** 016703
- [72] Zillich R E, Mayrhofer J M and Chin S A 2010 Extrapolated high-order propagators for path integral Monte Carlo simulations *J. Chem. Phys.* **132** 044103
- [73] Nava M, Galli D E, Cole M W and Reatto L 2012 Adsorption of He isotopes on fluorographene and graphane: Fluid and superfluid phases from quantum Monte Carlo calculations *Phys. Rev. B* **86** 174509
- [74] Huang K and Yang C N 1957 Quantum-mechanical many-body problem with hard-sphere interaction *Phys. Rev.* **105** 767
- [75] Yan Y and Blume D 2016 Path-Integral Monte Carlo determination of the fourth-order virial coefficient for a unitary two-component Fermi gas with zero-range interactions *Phys. Rev. Lett.* **116** 230401
- [76] Bernard E P, Krauth W and Wilson D B 2009 Event-chain Monte Carlo algorithms for hard-sphere systems *Phys. Rev. E* **80** 056704
- [77] Berg B 2004 *Markov Chain Monte Carlo Simulations and Their Statistical Analysis: With Web-based Fortran Code* (Singapore: World Scientific)
- [78] Box G E P and Muller M E 1958 A note on the generation of random normal deviates *Ann. Math. Stat.* **29** 610
- [79] Marsaglia G and Bray T A 1964 A convenient method for generating normal variables *SIAM Rev.* **6** 260
- [80] Marsaglia G and Tsang W W 2000 The ziggurat method for generating random variables *J. Stat. Softw.* **5** 1
- [81] Nadler B 2006 Design flaws in the implementation of the Ziggurat and Monty Python methods (and some remarks on Matlab randn) arXiv:math/0603058
- [82] Sakurai J J and Napolitano J 2011 *Modern Quantum Mechanics* 2nd edn (New York: Addison-Wesley)
- [83] Yan Y and Blume D 2014 Temperature dependence of small harmonically trapped atom systems with Bose, Fermi, and Boltzmann statistics *Phys. Rev. A* **90** 013620
- [84] Yan Y and Blume D 2015 Energy and structural properties of N-boson clusters attached to three-body Efimov states: Two-body zero-range interactions and the role of the three-body regulator *Phys. Rev. A* **92** 033626
- [85] Feller W 1971 *An Introduction to Probability Theory and its Applications (Wiley Series in Probability and Mathematical Statistics)* (New York: Wiley)

- [86] Jones G L 2004 On the Markov chain central limit theorem *Probab. Surv.* **1** 299
- [87] Yin X Y and Blume D 2015 Trapped unitary two-component Fermi gases with up to ten particles *Phys. Rev. A* **92** 013608
- [88] Carlson J and Gandolfi S 2014 Predicting energies of small clusters from the inhomogeneous unitary Fermi gas *Phys. Rev. A* **90** 011601(R)
- [89] Salas J 1980 *Applied Modeling of Hydrologic Time Series* (Highlands Ranch, CO: Water Resources Publications)
- [90] Gonick L 1993 *Cartoon Guide to Statistics (Cartoon Guide Series)* (New York: HarperCollins)
- [91] Thomas L H 1935 The interaction between a neutron and a proton and the structure of H^3 *Phys. Rev.* **47** 903
- [92] Baym G 1990 *Lectures on Quantum Mechanics (Advanced Book Program)* (New York: Westview Press)
- [93] Loh E Y, Gubernatis J E, Scalettar R T, White S R, Scalapino D J and Sugar R L 1990 Sign problem in the numerical simulation of many-electron systems *Phys. Rev. B* **41** 9301
- [94] Petrov D S 2003 Three-body problem in Fermi gases with short-range interparticle interaction *Phys. Rev. A* **67** 010703(R)
- [95] Skorniakov G V and Ter-Martirosian K A 1956 Three body problem for short range forces: I. Scattering of low energy neutrons by deuterons *Zh. Eksp. Teor. Fiz.* **31** 775
- Skorniakov G V and Ter-Martirosian K A 1957 *Sov. Phys.—JETP* **4** 648
- [96] Avery J 1989 *Hyperspherical Harmonics: Applications in Quantum Theory (Reidel Texts in the Mathematical Sciences)* (Norwell, MA: Kluwer Academic)
- [97] Werner F and Castin Y 2006 Unitary gas in an isotropic harmonic trap: symmetry properties and applications *Phys. Rev. A* **74** 053604
- [98] Werner F and Castin Y 2006 Unitary quantum three-body problem in a harmonic trap *Phys. Rev. Lett.* **97** 150401
- [99] Tan S 2008 Energetics of a strongly correlated Fermi gas *Ann. Phys., NY* **323** 2952
- [100] Daily K M and Blume D 2010 Energy spectrum of harmonically trapped two-component Fermi gases: three- and four-particle problem *Phys. Rev. A* **81** 053615
- [101] Rakshit D, Daily K M and Blume D 2012 Natural and unnatural parity states of small trapped equal-mass two-component Fermi gases at unitarity and fourth-order virial coefficient *Phys. Rev. A* **85** 033634
- [102] Blume D and Daily K M 2010 Few-body resonances of unequal-mass systems with infinite interspecies two-body s -wave scattering length *Phys. Rev. A* **82** 063612
- [103] Blume D 2008 Trapped polarized Fermi gas at unitarity *Phys. Rev. A* **78** 013635