

Size-consistent variational approaches to nonlocal pseudopotentials: Standard and lattice regularized diffusion Monte Carlo methods revisited

Cite as: J. Chem. Phys. **132**, 154113 (2010); <https://doi.org/10.1063/1.3380831>

Submitted: 01 February 2010 . Accepted: 13 March 2010 . Published Online: 21 April 2010

Michele Casula, Saverio Moroni, Sandro Sorella, and Claudia Filippi



View Online



Export Citation

ARTICLES YOU MAY BE INTERESTED IN

[Nonlocal pseudopotentials and diffusion Monte Carlo](#)

The Journal of Chemical Physics **95**, 3467 (1991); <https://doi.org/10.1063/1.460849>

[Energy-consistent pseudopotentials for quantum Monte Carlo calculations](#)

The Journal of Chemical Physics **126**, 234105 (2007); <https://doi.org/10.1063/1.2741534>

[Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space](#)

The Journal of Chemical Physics **131**, 054106 (2009); <https://doi.org/10.1063/1.3193710>



Size-consistent variational approaches to nonlocal pseudopotentials: Standard and lattice regularized diffusion Monte Carlo methods revisited

Michele Casula,^{1,a)} Saverio Moroni,^{2,3} Sandro Sorella,^{2,3} and Claudia Filippi⁴

¹*Centre de Physique Théorique, CNRS, École Polytechnique, 91128 Palaiseau Cedex, France*

²*International School for Advanced Studies (SISSA), Via Beirut 2-4, 34014 Trieste, Italy*

³*CNR-IOM Democritos National Simulation Center, Via Beirut 2-4, 34014 Trieste, Italy*

⁴*Faculty of Science and Technology and MESA+Research Institute, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands*

(Received 1 February 2010; accepted 13 March 2010; published online 21 April 2010)

We propose improved versions of the standard diffusion Monte Carlo (DMC) and the lattice regularized diffusion Monte Carlo (LRDMC) algorithms. For the DMC method, we refine a scheme recently devised to treat nonlocal pseudopotential in a variational way. We show that such scheme—when applied to large enough systems—maintains its effectiveness only at correspondingly small enough time-steps, and we present two simple upgrades of the method which guarantee the variational property in a size-consistent manner. For the LRDMC method, which is size-consistent and variational by construction, we enhance the computational efficiency by introducing: (i) an improved definition of the effective lattice Hamiltonian which remains size-consistent and entails a small lattice-space error with a known leading term and (ii) a new randomization method for the positions of the lattice knots which requires a single lattice-space. © 2010 American Institute of Physics. [doi:10.1063/1.3380831]

I. INTRODUCTION

The fixed-node (FN) diffusion Monte Carlo (DMC) method is often the method of choice for accurate computations of many-body systems.¹ Since the scaling of DMC with the number of electrons N is a modest N^4 , the method has been employed in recent years to accurately compute electronic properties of large molecular and solid systems where conventional highly correlated quantum chemistry approaches are very difficult to apply. Unfortunately, for full-core atoms the computational cost of DMC increases approximately^{2,3} as $Z^{5.5-6.5}$ with the atomic number Z . Therefore, the use of pseudopotentials is an essential ingredient in the application of DMC to complex systems to reduce the effective value of Z and significantly improve the efficiency of the method.

The use of pseudopotentials in DMC poses however a problem since pseudopotentials are usually nonlocal and the nonlocality introduces a fermionic sign problem additional to the one due to the antisymmetry of the electronic wave function. The commonly adopted solution is to “localize” the nonlocal pseudopotential on the trial wave function and use this local potential in the DMC simulation.^{4,5} Unfortunately, the so-called locality approximation (LA) does not ensure variationality and alternative schemes employing a different effective Hamiltonian were recently introduced to overcome this difficulty.^{6,7}

In the lattice regularized DMC (LRDMC) algorithm,⁶ both the Laplacian and the nonlocal pseudopotentials are discretized such that the corresponding imaginary time propagator $\langle \exp[-\tau\mathcal{H}] \rangle$ assumes nonzero values only on a finite set

of points, and the lattice Green function Monte Carlo algorithm can be employed that ensures variationality and stability all along the simulation.⁸ Alternatively, another scheme which is based on the standard DMC algorithm was developed.⁷ The latter exploits the discretization of the propagator only in the part depending on the nonlocal pseudopotentials, and a nonlocal effective Hamiltonian is defined in order to fulfill the FN constraint. Here, we show that this variational DMC scheme is however not size-consistent at finite time-steps. Indeed, the time-step error strongly depends on the system size and, upon increasing the number of particles at fixed time-step, the corresponding energies approach those given by DMC with LA. In this paper, we explain how to cure this problem and present a simple formulation of the algorithm which is size-consistent and suffers at the same time from a smaller time-step error. Moreover, we define a better discretization rule for the LRDMC effective Hamiltonian which reduces the lattice-space bias, remains size-consistent as in the original formulation, and improves the efficiency of the method.

In Sec. II, we briefly summarize the problems introduced by the use of nonlocal pseudopotentials in the standard DMC and describe in detail the variational DMC algorithm of Ref. 7 to treat pseudopotentials beyond the commonly used LA. In Sec. III, we present our size-consistent variational (SV) approach to nonlocal pseudopotentials in DMC and demonstrate its effectiveness on a series of oxygen systems of increasing size. In Sec. IV, we briefly describe the LRDMC method which is variational by its own nature, and give a better prescription for the lattice regularization of the continuous Hamiltonian to always guarantee a well defined and faster zero lattice-space extrapolation. Finally, in Sec. V, we discuss the behavior of the discretization error of the differ-

^{a)}Electronic mail: casula@cph.t.polytechnique.fr.

ent DMC algorithms (in time or space as appropriate) and comment on the relative efficiency of the methods presented here.

II. DMC AND NONLOCAL PSEUDOPOTENTIALS

In DMC, the projection to the ground state wave function of an Hamiltonian \mathcal{H} is performed by stochastically applying the operator $\exp[-\tau\mathcal{H}]$ to a trial wave function Ψ_T . If the projection is formulated in real space and importance sampling introduced, the mixed distribution $f(\mathbf{R}, t) = \Psi_T(\mathbf{R})\Psi(\mathbf{R}, t)$ is then propagated as

$$f(\mathbf{R}', t + \tau) = \int d\mathbf{R} G(\mathbf{R}', \mathbf{R}, \tau) f(\mathbf{R}, t), \quad (1)$$

where the importance sampling Green's function is defined as

$$G(\mathbf{R}', \mathbf{R}, \tau) = \frac{\Psi_T(\mathbf{R}')}{\Psi_T(\mathbf{R})} \langle \mathbf{R}' | \exp[-\tau\mathcal{H}] | \mathbf{R} \rangle. \quad (2)$$

The FN approximation is usually employed for fermionic systems to avoid the collapse to the bosonic ground state. In continuous systems, it is implemented by constraining the diffusion process within the nodal pockets of the trial wave function. For long times, the distribution $f(\mathbf{R}, t)$ approaches $\Psi_T(\mathbf{R})\Psi_{FN}(\mathbf{R})$ where $\Psi_{FN}(\mathbf{R})$ is the ground state wave function consistent with the boundary condition that it vanishes at the nodes of Ψ_T . The FN energy is an upper bound to the true fermionic ground state energy.

When a nonlocal potential \mathcal{V}^{NL} is employed to remove the core electrons, the off-diagonal elements of the Hamiltonian in real space are generally nonzero and the standard DMC approach cannot be applied. If we analyze the behavior of the propagator at short time-steps

$$\langle \mathbf{R}' | \exp[-\tau\mathcal{H}] | \mathbf{R} \rangle \approx \delta_{\mathbf{R}', \mathbf{R}} - \tau \langle \mathbf{R}' | \mathcal{H} | \mathbf{R} \rangle, \quad (3)$$

we note that, while the diagonal elements can always be made positive by choosing τ small enough, the off-diagonal elements are non-negative if and only if the off-diagonal elements of the Hamiltonian are nonpositive. This condition is not always met in the presence of nonlocal pseudopotentials, so the fermionic sign problem reappears even if one works in the FN approximation. Consequently, in addition to the FN approximation, the LA is commonly introduced where the nonlocal potential \mathcal{V}^{NL} is replaced by a local quantity \mathcal{V}^{LA} obtained by ‘‘localizing’’ the potential on the trial wave function

$$\mathcal{V}^{LA}(\mathbf{R}) = \frac{\langle \mathbf{R} | \mathcal{V}^{NL} | \Psi_T \rangle}{\langle \mathbf{R} | \Psi_T \rangle}. \quad (4)$$

The DMC algorithm in the LA yields the FN ground state of the effective Hamiltonian \mathcal{H}^{LA} with the local potential \mathcal{V}^{LA} instead of the original nonlocal \mathcal{V}^{NL} . The FN energy in the LA is equal to

$$E_{FN}^{LA} = \frac{\langle \Psi_{FN}^{LA} | \mathcal{H}^{LA} | \Psi_{FN}^{LA} \rangle}{\langle \Psi_{FN}^{LA} | \Psi_{FN}^{LA} \rangle}, \quad (5)$$

and estimated by sampling the mixed distribution $\Psi_T \Psi_{FN}^{LA}$ as

$$E_{FN}^{LA} = \frac{\langle \Psi_{FN}^{LA} | \mathcal{H}^{LA} | \Psi_T \rangle}{\langle \Psi_{FN}^{LA} | \Psi_T \rangle} = \frac{\langle \Psi_{FN}^{LA} | \mathcal{H} | \Psi_T \rangle}{\langle \Psi_{FN}^{LA} | \Psi_T \rangle}. \quad (6)$$

Since Ψ_{FN}^{LA} is the FN ground state of \mathcal{H}^{LA} and not of the original Hamiltonian \mathcal{H} , the mixed average energy of \mathcal{H} is not equal to its expectation value on the wave function Ψ_{FN}^{LA}

$$\frac{\langle \Psi_{FN}^{LA} | \mathcal{H} | \Psi_{FN}^{LA} \rangle}{\langle \Psi_{FN}^{LA} | \Psi_{FN}^{LA} \rangle}. \quad (7)$$

Therefore, E_{FN}^{LA} is in general not an upper bound to the ground state of \mathcal{H} and the variational principle may not apply.

A. Beyond the LA

The LRDMC algorithm was recently developed to overcome this difficulty⁶ and then extended to the continuum formulation of DMC.⁷ Both algorithms provide a variational scheme to treat nonlocal pseudopotentials in DMC by introducing an effective Hamiltonian, different from the one used in the LA approximation, which provides an upper bound to the ground state of the original Hamiltonian. We briefly describe here the algorithm in the framework of continuum DMC.

We first apply a Trotter expansion for small time-steps to the importance sampling Green's function

$$G(\mathbf{R}', \mathbf{R}, \tau) \approx \int d\mathbf{R}'' T^{NL}(\mathbf{R}', \mathbf{R}'', \tau) G^{loc}(\mathbf{R}'', \mathbf{R}, \tau), \quad (8)$$

where we have split the Hamiltonian into a local and a non-local operator. The propagator $G^{loc}(\mathbf{R}'', \mathbf{R}, \tau)$ is equal to the drift-diffusion-branching Green's function for the local component of the Hamiltonian

$$\frac{1}{(2\pi\tau)^{3N/2}} e^{-[\mathbf{R}' - \mathbf{R} - \tau\mathbf{V}(\mathbf{R})]^2/2\tau} e^{-\tau E_L^{loc}(\mathbf{R}')}, \quad (9)$$

where the velocity is defined as $\mathbf{V}(\mathbf{R}) = \nabla \Psi_T(\mathbf{R}) / \Psi_T(\mathbf{R})$ and $E_L^{loc}(\mathbf{R}) = \mathcal{H}^{loc} \Psi_T(\mathbf{R}) / \Psi_T(\mathbf{R})$ is the local energy of the local part of the Hamiltonian (kinetic \mathcal{K} plus local potential \mathcal{V}^{loc}). The transition T^{NL} contains the nonlocal potential

$$\begin{aligned} T^{NL}(\mathbf{R}', \mathbf{R}, \tau) &= \frac{\Psi_T(\mathbf{R}')}{\Psi_T(\mathbf{R})} \langle \mathbf{R}' | \exp[-\tau\mathcal{V}^{NL}] | \mathbf{R} \rangle \\ &\approx \delta_{\mathbf{R}', \mathbf{R}} - \tau V_{\mathbf{R}', \mathbf{R}}. \end{aligned} \quad (10)$$

where $V_{\mathbf{R}', \mathbf{R}} = \Psi_T(\mathbf{R}') / \Psi_T(\mathbf{R}) \langle \mathbf{R}' | \mathcal{V}^{NL} | \mathbf{R} \rangle$. In both the variational Monte Carlo and the standard DMC method with the LA approximation, one adopts a quadrature rule with a discrete mesh of points, belonging to a regular polyhedron used to evaluate the projection of the nonlocal component on a given trial wave function.

Consequently, the number of elements $V_{\mathbf{R}', \mathbf{R}}$ is finite and the transition T^{NL} corresponds to the move of one electron on the grid obtained by considering the union of the quadrature points generated for each electron and pseudoatom (center of a nonlocal pseudopotential). Moreover, in order to work with a small quadrature mesh, the vertices of the polyhedron are

defined in a frame rotated by θ and ϕ , the azimuthal and planar angle, respectively, which are taken randomly for each electron.

As discussed above, performing a transition based on T^{NL} poses however a problem since T^{NL} can be negative given that both $\Psi_{\text{T}}(\mathbf{R}')/\Psi_{\text{T}}(\mathbf{R})$ and $\langle \mathbf{R}' | \mathcal{V}^{\text{NL}} | \mathbf{R} \rangle$ can change sign. A solution is to apply the FN approximation not only to G^{loc} but also to T^{NL} by keeping only the transition elements which are positive

$$T_{\text{FN}}^{\text{NL}}(\mathbf{R}', \mathbf{R}, \tau) = \delta_{\mathbf{R}', \mathbf{R}} - \tau V_{\mathbf{R}', \mathbf{R}}^- \quad (11)$$

where $V_{\mathbf{R}', \mathbf{R}}^{\pm} = [V_{\mathbf{R}', \mathbf{R}} \pm |V_{\mathbf{R}', \mathbf{R}}|]/2$. The discarded elements are included in the so-called sign-flip term, which is then added to the diagonal local potential as

$$\mathcal{V}_{\text{eff}}^{\text{loc}}(\mathbf{R}) = \mathcal{V}^{\text{loc}}(\mathbf{R}) + \sum_{\mathbf{R}'} V_{\mathbf{R}', \mathbf{R}}^+ \quad (12)$$

The resulting effective Hamiltonian \mathcal{H}_{eff} is therefore given by

$$\langle \mathbf{R} | \mathcal{H}_{\text{eff}} | \mathbf{R} \rangle = \langle \mathbf{R} | \mathcal{K} | \mathbf{R} \rangle + \mathcal{V}_{\text{eff}}^{\text{loc}}(\mathbf{R}), \quad (13)$$

$$\langle \mathbf{R}' | \mathcal{H}_{\text{eff}} | \mathbf{R} \rangle = \langle \mathbf{R}' | \mathcal{V}^{\text{NL}} | \mathbf{R} \rangle \quad \text{if } V_{\mathbf{R}', \mathbf{R}} < 0,$$

and yields the same local energy as the original Hamiltonian \mathcal{H} . In contrast to the LA Hamiltonian, its ground state energy is an upper bound to the ground state energy of the true Hamiltonian. Therefore, the variational principle is recovered and, in addition, the use of \mathcal{H}_{eff} in combination with the $T_{\text{FN}}^{\text{NL}}$ transition cures the instabilities which are commonly observed in a DMC run with the LA Hamiltonian and are due to the negative divergences of the localized potential on the nodes of Ψ_{T} .

In the branching term of G^{loc} [Eq. (9)], the local potential \mathcal{V}^{loc} is replaced by $\mathcal{V}_{\text{eff}}^{\text{loc}}$ [Eq. (12)] and the weights of the walkers are multiplied by an additional factor which enters in the normalization of the transition T_{FN} and to order τ is equal to

$$\sum_{\mathbf{R}'} T_{\text{FN}}^{\text{NL}}(\mathbf{R}', \mathbf{R}) \approx \exp\left(-\tau \sum_{\mathbf{R}'} V_{\mathbf{R}', \mathbf{R}}\right). \quad (14)$$

The weights are therefore given by

$$w = w_{\text{eff}}^{\text{loc}} \sum_{\mathbf{R}'} T_{\text{FN}}^{\text{NL}}(\mathbf{R}', \mathbf{R}) = \exp[-\tau E_{\text{L}}(\mathbf{R})], \quad (15)$$

where $E_{\text{L}}(\mathbf{R}) = \mathcal{H}_{\text{eff}} \Psi_{\text{T}}(\mathbf{R}) / \Psi_{\text{T}}(\mathbf{R}) = \mathcal{H} \Psi_{\text{T}}(\mathbf{R}) / \Psi_{\text{T}}(\mathbf{R})$.

The basic algorithm as proposed in Ref. 7 therefore consists of the following steps:

1. The walker drifts and diffuses from \mathbf{R} to \mathbf{R}' . The move is followed by an accept/reject step as in standard DMC.
2. The weight of the walker is multiplied by the branching factor $\exp[-\tau(E_{\text{L}}(\mathbf{R}') - E_{\text{T}})]$ where the trial energy E_{T} has been introduced.
3. The walker moves to \mathbf{R}'' according to the transition probability $T_{\text{FN}}^{\text{NL}}(\mathbf{R}'', \mathbf{R}') / \sum_{\mathbf{R}'''} T_{\text{FN}}^{\text{NL}}(\mathbf{R}''', \mathbf{R}')$.

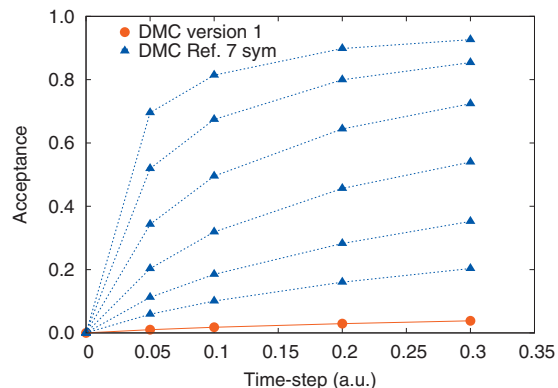


FIG. 1. Acceptance of the $T_{\text{FN}}^{\text{NL}}$ move as a function of time-step. The increasing dotted curves correspond to systems with 1, 2, 4, 8, 16, and 32 oxygen atoms aligned at a distance of 30 Å from each other. The dotted curves are obtained with the algorithm of Ref. 7 while the lowest continuum curve is obtained with the size-consistent DMC algorithm (version 1) we proposed. We only show the size-consistent curve with one atom as it is indistinguishable from the ones obtained for the larger systems.

For large systems, the first step is implemented not by moving all the electron together but by sequentially drifting and diffusing each electron and applying the accept/reject step after each single-electron (SE) move.

III. SIZE-CONSISTENCY

In the move governed by the transition $T_{\text{FN}}^{\text{NL}}$, only one electron is displaced on the grid of the quadrature points generated by considering all the pseudoatoms and all the electrons. Therefore, for given time-step, the probability of a successful move will increase with the system size (i.e., the number of electrons) and saturate to one for sufficiently large systems. In this limit, the effect of the move will become independent of the system size and lead to one electron being displaced at each step. Therefore, for sufficiently large systems, the overall impact of the nonlocal move will decrease and the algorithm will effectively behaves more and more like in the LA procedure.

To demonstrate the size-consistency problem of the algorithm as originally formulated in Ref. 7, we consider a series of systems consisting of an increasing number M of oxygen atoms aligned 30 Å apart. The oxygen atom is described by an s -nonlocal energy-consistent Hartree–Fock pseudopotential.⁹ The trial wave function is of the Jastrow–Slater type with a single determinant expressed on a cc-pVDZ basis⁹ and a Jastrow factor which includes electron–electron and electron–nucleus terms.¹⁰ All Jastrow and orbital parameters are optimized in energy minimization¹¹ for a single atom and the wave function of a system with more than one oxygen is obtained by replicating the wave function of one atom on the other centers. In Fig. 1, we plot the acceptance of the $T_{\text{FN}}^{\text{NL}}$ move as a function of time-step for systems containing 1, 2, 4, 8, 16, and 32 oxygen atoms. For each system size, the probability goes to 0 at small time-steps and increases for larger values of τ as expected from the expression of $T_{\text{FN}}^{\text{NL}}$ [Eq. (11)]. The acceptance increases

with the size of the system; as a function of the time-step, it approaches its asymptotic value of one more quickly for the larger systems.

To better understand the overall behavior of the algorithm with increasing system size, we also analyze the FN energy as a function of time-step. We are interested in comparing the results obtained with the conventional LA approach and with the algorithm of Ref. 7 described in the previous section. For a more meaningful and clear comparison with conventional DMC with the LA which employs a symmetrized branching factor, we modify the original algorithm as described in the previous section to also use a symmetrized branching factor

$$\exp[-\tau(E_L(\mathbf{R}) + E_L(\mathbf{R}'))/2] \quad (16)$$

where \mathbf{R} and \mathbf{R}' are the coordinates before the drift-diffusion move of the first electron and after the drift-diffusion move of the last electron, respectively, if the electrons are displaced subsequently. Such a simple modification is allowed as it only entails a different time-step error which we actually find to be significantly smaller than the one obtained with the algorithm of Ref. 7 as we will detail in the Sec. V.

From the results reported in Fig. 2, we observe that the FN energies obtained with the algorithm of Ref. 7 significantly increase with M and approach the energies obtained with the LA algorithm. Already with a system with 32 oxygen atoms, the FN energies at $\tau=0.1$ obtained with these two approaches become equivalent within the error bars. The energies given by the two algorithms must however extrapolate to different values as the time-step goes to 0.⁷ The lower panel of Fig. 2 shows that for $M=32$, in particular, they have to depart from each other within a tiny time-step interval near the origin. Because of this behavior, the algorithm of Ref. 7 is bound to have a problematic extrapolation to the 0 time-step limit for large enough systems.

A. Size-consistent variational formulations: SVDMC Version 1

To address this problem, the original algorithm of Ref. 7 can be easily reformulated in a size-consistent manner by observing that the transition T^{NL} [Eq. (10)] can be factorized as

$$\begin{aligned} T^{\text{NL}}(\mathbf{R}', \mathbf{R}, \tau) &= \frac{\Psi_T(\mathbf{R}')}{\Psi_T(\mathbf{R})} \prod_{i=1}^N \langle \mathbf{r}'_i | e^{-\tau \nu^{\text{NL}}} | \mathbf{r}_i \rangle \\ &= \prod_{i=1}^N \frac{\Psi_T(\mathbf{r}'_1, \dots, \mathbf{r}'_i, \mathbf{r}_{i+1}, \dots)}{\Psi_T(\mathbf{r}'_1, \dots, \mathbf{r}'_{i-1}, \mathbf{r}_i, \dots)} \langle \mathbf{r}'_i | e^{-\tau \nu^{\text{NL}}} | \mathbf{r}_i \rangle \\ &= \prod_{i=1}^N (\delta_{\mathbf{r}'_i \mathbf{r}_i} - \tau v_{\mathbf{r}'_i \mathbf{r}_i}) = \prod_{i=1}^N t^{\text{NL}}(\mathbf{r}'_i, \mathbf{r}_i, \tau), \quad (17) \end{aligned}$$

where ν^{NL} is the nonlocal potential acting only on one electron due to all the atomic centers so that the total nonlocal potential is given by the sum over the electrons $\langle \mathbf{R}' | \nu^{\text{NL}} | \mathbf{R} \rangle = \sum_i \langle \mathbf{r}'_i | \nu^{\text{NL}} | \mathbf{r}_i \rangle$. We defined the matrix element $v_{\mathbf{r}'_i \mathbf{r}_i}$ as

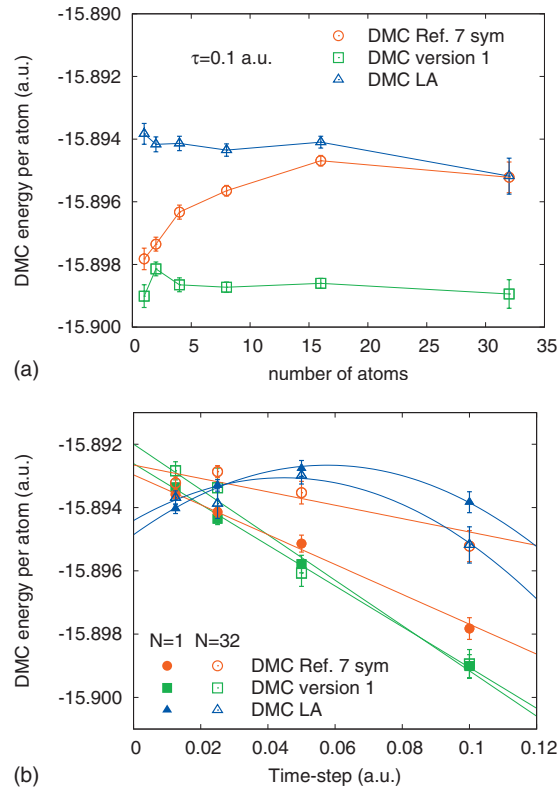


FIG. 2. Upper panel: DMC FN energy per atom for systems of M isolated oxygen atoms for $\tau=0.1$. For the given time-step, the results of the algorithm of Ref. 7 (red circles) approach those of the LA (blue triangles) upon increasing the system size, whereas the present algorithm (DMC version 1, green squares) gives values independent of M . In the three algorithms, the branching factor is updated after having moved all the electrons (AE branching). Lower panel: time-step dependence of the DMC FN energy for the same three algorithms for $M=1$ (filled symbols) and $M=32$ (open symbols). The algorithm of Ref. 7 will have a problematic extrapolation to 0 time-step for large enough system size. For $M=32$, the linear extrapolation (red curve with open symbols) is consistent, as expected, with the corresponding result for $M=1$ (red curve with filled circles). However a better χ^2 would be obtained with a quadratic extrapolation which in turn would require a sudden upturn at very small τ to recover the correct zero time-step limit.

$$v_{\mathbf{r}'_i \mathbf{r}_i} = \frac{\Psi_T(\mathbf{r}'_1, \dots, \mathbf{r}'_i, \mathbf{r}_{i+1}, \dots)}{\Psi_T(\mathbf{r}'_1, \dots, \mathbf{r}'_{i-1}, \mathbf{r}_i, \dots)} \langle \mathbf{r}'_i | \nu^{\text{NL}} | \mathbf{r}_i \rangle. \quad (18)$$

The transition $t^{\text{NL}}(\mathbf{r}'_i, \mathbf{r}_i, \tau)$ displaces the i th electron over the grid of quadrature points generated by considering only the i th electron and all pseudoatoms which host the i th electron in their core region.

The FN approximation is applied separately to each SE transition as

$$t_{\text{FN}}^{\text{NL}}(\mathbf{r}', \mathbf{r}, \tau) = \delta_{\mathbf{r}' \mathbf{r}} - \tau v_{\mathbf{r}' \mathbf{r}}, \quad (19)$$

where $v_{\mathbf{r}' \mathbf{r}}^{\pm}$ is defined in analogy to the case of the total nonlocal potential so that only the positive transition elements are kept in the transition matrix. In this formulation, the third step of the DMC algorithm detailed above consists of a loop over the electrons where each electron is subsequently moved according to the SE transition, $t_{\text{FN}}^{\text{NL}}$. Therefore, while in the original algorithm of Ref. 7 the configuration generated in the $T_{\text{FN}}^{\text{NL}}$ step differs from the starting configuration only in the coordinate of one electron, the configuration resulting from this size-consistent move will generally change

in more than one electronic coordinate and the number of electrons being moved will increase with the size of the system.

To understand that the drift-diffusion-branching steps in the original algorithm do not need to be modified, we observe that the expression of the effective Hamiltonian \mathcal{H}_{eff} we are working with is the same as in Eq. (13) in the limit of τ going to 0. In particular, the sign-flip term obtained by summing all discarded terms $v_{\mathbf{r}'_i, \mathbf{r}_i}^+$ over all the electrons is equal to the sign-flip term in $\mathcal{V}_{\text{eff}}^{\text{loc}}$ [Eq. (12)] to zero-order in τ . Similarly, we have that, to order τ

$$\prod_{i=1}^N \sum_{\mathbf{r}'_i} t_{\text{FN}}^{\text{NL}}(\mathbf{r}'_i, \mathbf{r}_i, \tau) \approx \sum_{\mathbf{R}'} T_{\text{FN}}^{\text{NL}}(\mathbf{R}', \mathbf{R}, \tau), \quad (20)$$

and we recover the same branching factor as in the original algorithm. Therefore, both algorithms extrapolate to the same limit at 0 time-step. We will refer to this improved algorithm as “DMC version 1” to distinguish it from another size-consistent version we will define later in this section. We stress here that by “version 1” we do not only mean the use of the product of single-particle $t_{\text{FN}}^{\text{NL}}$ in step 3, but also the symmetrization of the weights in step 2, as described in Eq. (16), where the initial configuration is taken before the diffusion process (step 1), and the final is the one after step 1.

The acceptance as a function of time-step using the size-consistent DMC algorithm (version 1) is shown in Fig. 1. We only report the result obtained with $M=1$ as the curves for the other system sizes are exactly equivalent within statistical error. This finding can be easily understood since the probability of moving a given electron on the grid generated by considering all centers will be practically the same as the probability computed using only the atom close to the electron as all other centers are at least 30 Å far apart. Therefore, in the new size-consistent algorithm when more atoms are added to increase the size of the system, the loop over the electrons will ensure that each electron attempts a move around its closest center. The acceptance remains therefore constant as more oxygen atoms are added. In a more realistic systems (e.g., with closer oxygen atoms) there will be a weak dependence on the size of the system but, after enough atoms have been added for most electrons to experience an equivalent environment, the acceptance will become independent on the system size given the short-range nature of the nonlocal components of the pseudopotentials.

The FN energies obtained for the oxygen systems with this size-consistent algorithm are compared in Fig. 2 with the results of the LA and of the original algorithm. We observe that as expected the FN energies of the size-consistent algorithm extrapolate to the same value as the original algorithm as τ goes to zero. On the other hand, while the size-consistent FN energies are close to the values obtained with the original method for the smallest, one atom system, the FN results obtained by the two methods depart from each other as the system size increases. Importantly, the FN energies of the size-consistent scheme do not approach the LA results for large systems at finite τ and their extrapolation to zero time-step is therefore as smooth for large as for small systems.

B. Size-consistent variational formulations: SVMC Version 2

An alternative scheme to address the size-consistent problem of the original algorithm of Ref. 7 can be obtained through a different route by starting from Eq. (10) and breaking it up in N terms with time-step of τ/N such that

$$\sum_{\mathbf{R}'} T_{\text{FN}}^{\text{NL}}(\mathbf{R}', \mathbf{R}, \tau) = \sum_{\mathbf{R}_1, \dots, \mathbf{R}_N} \prod_{i=1}^N T_{\text{FN}}^{\text{NL}}(\mathbf{R}_i, \mathbf{R}_{i-1}, \tau/N), \quad (21)$$

with $\mathbf{R}_N = \mathbf{R}'$, $\mathbf{R}_0 = \mathbf{R}$, and the sum over the quadrature points sampled by the chain $\{\mathbf{R}_0, \dots, \mathbf{R}_i, \dots, \mathbf{R}_N\}$ generated during the random walk. This is another way to evaluate the quantity in Eq. (20). The difference is that Eq. (21) involves a product of N all-electron factors, while Eq. (20) is a factorization of N SE terms. Both will avoid the saturation of the acceptance probability of the nonlocal Green’s function $T_{\text{FN}}^{\text{NL}}$, and therefore they will ensure a size-consistent time-step error. Since Eq. (21) requires the calculation of *all* matrix elements $V_{\mathbf{R}', \mathbf{R}}$ each time, it is more convenient to split the N factors in such a way that the drift-diffusion move involving the i th electron could be placed between the $i-1$ th and i th factor, and the corresponding branching weight updated as a product of subsequent SE components:

$$\prod_{i=1}^N \exp \left[-\frac{\tau}{N} E_L(\mathbf{r}'_1, \dots, \mathbf{r}'_i, \mathbf{r}_{i+1}, \dots, \mathbf{r}_N) \right], \quad (22)$$

where we can exploit the knowledge of $V_{\mathbf{R}', \mathbf{R}}$ to compute also E_L for every SE move. We will call this algorithm “DMC version 2.”¹² It consists of the following steps:

1. Drift-diffusion move of the i th electron.
2. The weight of the walker is multiplied by the branching factor $\exp[-\frac{\tau}{N} E_L(\mathbf{r}'_1 \dots \mathbf{r}'_i, \mathbf{r}_{i+1}, \dots, \mathbf{r}_N)]$.
3. The walker moves to \mathbf{R}'' according to the transition probability $T_{\text{FN}}^{\text{NL}}(\mathbf{R}'', \mathbf{R}', \tau/N) / \sum_{\mathbf{R}'''} T_{\text{FN}}^{\text{NL}}(\mathbf{R}''', \mathbf{R}', \tau/N)$ which involves *all* the electrons.

In contrast to the original algorithm, and the “DMC version 1,” these three steps need to be performed inside a loop over the electrons. In the “version 2” formulation of the DMC size-consistent algorithm, each electron drifts and diffuses in a time τ and the branching factor is updated at each SE move with the total local energy E_L and time τ/N in the exponent. After each SE branching update, a nonlocal transition is performed with $T_{\text{FN}}^{\text{NL}}(\mathbf{R}'', \mathbf{R}', \tau/N)$ where one electron among *all* electrons is displaced over the grid of quadrature points obtained by considering all electrons and all pseudopotentials. Therefore, the electron displaced in the nonlocal move may differ from the electron which is currently being moved in the drift-diffusion step.

IV. LRDMC AND NONLOCAL PSEUDOPOTENTIALS

The main difference between the effective DMC Hamiltonian reported in Eq. (13) and the LRDMC one is the kinetic operator \mathcal{K} . In the LRDMC approach, \mathcal{K} is replaced by a discretized Laplacian and treated on the same footing as \mathcal{V}^{NL} . In the original formulation, the discretized Laplacian is

a linear combination of two discrete operators with incommensurate lattice spaces a and a' , introduced to sample densely the continuous space by performing discrete moves whose length is either a or a' . This method can be simplified by noticing that all the continuous space can be visited using only a single displacement length a , provided we randomize the direction of the Cartesian coordinates each time the electron positions are updated. The randomization of the lattice mesh is similar to the well established approach used to perform the angular integration in the nonlocal part of the pseudopotential.⁴ Therefore, in the LRDMC approach, we can extend the definition of the kinetic part by including both the discretized Laplacian and the nonlocal part of the pseudopotentials. The total nonlocal operator reads

$$\mathcal{K}^a = - \sum_{i=1}^N \Delta_i^a(\theta_i, \phi_i)/2 + \mathcal{V}^{\text{NL}}, \quad (23)$$

where $\Delta_i^a(\theta_i, \phi_i)$ is the Laplacian acting on the i th electron and discretized to second order so that $\Delta_i^a(\theta_i, \phi_i) = \Delta_i + O(a^2)$. The discretized Laplacian is computed in a frame rotated by the angles θ_i and ϕ_i which are chosen randomly and independently of the ones used to compute \mathcal{V}^{NL} . In this formulation, we need to evaluate only six off-diagonal elements of the Green's function instead of 12 as in the original algorithm, gaining a speedup of a factor of 1.6 in full-core calculations and of $(16+n_{\text{quad}})/(10+n_{\text{quad}})$ with pseudopotentials, where n_{quad} is the number of quadrature points per electron, and the cost of computing both the Laplacian and the gradient has been taken into account.¹³ We notice that, with this simplification, the LRDMC error in the extrapolation to the continuous limit depends on a single parameter a , and the method can therefore be compared fairly with the DMC approach where the discretization of the diffusion process also depends on a single scale, i.e., the time-step τ .

In the LRDMC choice of the Hamiltonian, we further regularize the single-particle operator ν , defined as the electron-ion Coulomb interaction in full-core atoms or the local part ν^{loc} of the pseudopotential, so that $\nu(\mathbf{r}_i) \rightarrow \mathcal{V}_i^x(\mathbf{R})$ as

$$\mathcal{V}_i^x(\mathbf{R}) = \nu(\mathbf{r}_i) - \frac{(\Delta_i - \Delta_i^a)\Psi_T(\mathbf{R})}{2\Psi_T(\mathbf{R})}, \quad (24)$$

when acting on the i th electron. The single-particle operator ν acquires therefore a many-body term and $\mathcal{V}_i^x(\mathbf{R})$ depends on the all-electron configuration \mathbf{R} . The total potential term is then given by

$$\mathcal{V}^x = \sum_{i=1}^N \mathcal{V}_i^x + \mathcal{V}_{\text{ee}} + \mathcal{V}_{\text{nn}}, \quad (25)$$

where no regularization is employed in the electron-electron \mathcal{V}_{ee} and ion-ion \mathcal{V}_{nn} Coulomb terms. This lattice regularization leads to an approximate Hamiltonian $\mathcal{H}^a = \mathcal{K}^a + \mathcal{V}^x$ which converges to the exact Hamiltonian as $\mathcal{H}^a = \mathcal{H} + a^2 \Delta \mathcal{H}$ for $a \rightarrow 0$ where we denote with $a^2 \Delta \mathcal{H}$ the $O(a^2)$ LRDMC error on \mathcal{H} .

The lattice Green's function Monte Carlo algorithm can then be employed to sample exactly the lattice regularized Green's function $\Lambda - \mathcal{H}^a$ and project the trial wave function

Ψ_T to the approximate ground state Ψ_a^{LRDMC} which fulfills the FN constraint based on Ψ_T , in complete analogy to the DMC framework.⁶ Note that, since the spectrum of \mathcal{H}^a is not bounded from above, we need to take the limit $\Lambda \rightarrow \infty$ which can be handled with no loss of efficiency as described in Ref. 14. The usual DMC Trotter breakup results in a time-step error, while the LRDMC formulation yields a lattice-space error, but both approaches share the same upper bound property and converge to the same projected FN energy in the limit of 0 time-step and lattice-space, respectively.

Since the discretized Laplacian and the nonlocal potential are treated on the same footing, and the sampling of the Green's function is based on a sequence of single-particle moves generated both from the Laplacian and the nonlocal part, the LRDMC is intrinsically size-consistent (in the sense previously discussed for the DMC algorithm), and no modification is necessary to make the lattice-space bias independent of the system size. It will depend however on the quality of the trial wave function in the way detailed below.

A. Small a^2 correction for good trial function

The regularization of the potential [Eq. (24)] in the definition of the lattice Hamiltonian \mathcal{H}^a implies that the correction $\Delta \mathcal{H}$ satisfies

$$\Delta \mathcal{H} |\Psi_T\rangle = 0. \quad (26)$$

Using this property, we can estimate the leading-order error of the lattice regularization by simple perturbation theory as

$$\begin{aligned} E^a &= E^0 + a^2 \langle \Psi_0^{\text{LRDMC}} | \Delta \mathcal{H} | \Psi_0^{\text{LRDMC}} \rangle \\ &= E^0 + a^2 \langle \Psi_0^{\text{LRDMC}} - \Psi_T | \Delta \mathcal{H} | \Psi_0^{\text{LRDMC}} - \Psi_T \rangle \\ &= E^0 + O(a^2 |\Psi_0^{\text{LRDMC}} - \Psi_T|^2), \end{aligned} \quad (27)$$

where E^a is the expectation value of the Hamiltonian \mathcal{H}^a on the approximate FN ground state Ψ_a^{LRDMC} and E^0 the extrapolated value as $a \rightarrow 0$. Thus, the approach to the continuous limit is particularly fast for good trial functions, namely, for Ψ_T close to the ground state solution, since Ψ_0^{LRDMC} is a state with lower energy than Ψ_T and has to approach the ground state at least as Ψ_T does. The leading corrections to the continuous limit are quadratic in the wave function error. This property is not easily generalized to the usual DMC method and, to our knowledge, has not been established so far.

B. Well defined lattice regularization

As in any lattice model, the Hamiltonian \mathcal{H}^a has a finite ground state energy only if the potential \mathcal{V}^x is always limited from below. If $\mathcal{V}^x(\mathbf{R}_0) = -\infty$ for some configuration \mathbf{R}_0 , the variational state $\Psi(\mathbf{R}) = \delta_{\mathbf{R}, \mathbf{R}_0}$ will have unbounded negative energy expectation value and the ground state energy of \mathcal{H}^a is not defined. Unfortunately, the regularized potential $\mathcal{V}^x(\mathbf{R})$ in Eq. (24) is not bounded from below when \mathbf{R} belongs to the $(3N-1)$ -dimensional nodal surface \mathcal{N} defined by the equation $\Psi_T(\mathbf{R}) = 0$. To cure these divergences, we need to be able to establish when a configuration is close to the nodal surface. In the lattice regularized formulation, we can assign an electron position \mathbf{r}_i to the nodal surface, i.e., $\mathbf{r}_i \in \mathcal{N}_a$, if

$\Psi_T(\mathbf{r}_i + a\vec{\mu})$ has the opposite sign of $\Psi_T(\mathbf{r}_i)$ for at least one of the six points used to evaluate the finite difference Laplacian (i.e., $\vec{\mu}$ is one of the six unit vectors $\pm\hat{x}$, $\pm\hat{y}$, or $\pm\hat{z}$ of the reference frame randomly oriented according to the angles θ_i and ϕ_i). \mathcal{N}_a correctly defines the nodal surface \mathcal{N} in the limit $a \rightarrow 0$.

With this definition of nodal surface, we can modify \mathcal{V}_i^a so that it remains finite when $\mathbf{r}_i \in \mathcal{N}_a$

$$\tilde{\mathcal{V}}_i^a(\mathbf{R}) = \begin{cases} \text{Max}[\nu(\mathbf{r}_i), \mathcal{V}_i^a(\mathbf{R})] & \text{if } \mathbf{r}_i \in \mathcal{N}_a \\ \mathcal{V}_i^a(\mathbf{R}) & \text{otherwise.} \end{cases} \quad (28)$$

If $\mathbf{r}_i \notin \mathcal{N}_a$, we use the original LRDMC definition of \mathcal{V}_i^a since \mathcal{V}_i^a remains finite even when an electron approaches a nucleus for trial functions which satisfy the electron-ion cusp conditions. If $\mathbf{r}_i \in \mathcal{N}_a$, we need to distinguish two cases. If the electron is not close to a nucleus, the regularized \mathcal{V}_i^a can diverge negatively while $\nu(\mathbf{r}_i)$ remains finite and, according to Eq. (28), the potential $\tilde{\mathcal{V}}_i^a$ coincides with ν . If the electron is close to a nucleus in a full-core calculation, both \mathcal{V}_i^a and $\nu(\mathbf{r}_i)$ diverge, so we need to further regularize $\nu(\mathbf{r}_i)$ in the right hand side of Eq. (28) and use an expression bounded from below. In this particular case, we choose to replace the divergent electron-ion contribution $-Z/|\mathbf{r}_{in}|$ in $\nu(\mathbf{r}_i)$ with $-Z/a$ whenever the electron-ion distance $|\mathbf{r}_{in}| < a$.

If we employ the regularized potential $\tilde{\mathcal{V}}^a$ in the Hamiltonian \mathcal{H}^a , we no longer satisfy Eq. (26) and, in principle, it is not possible to compute E^a by averaging the local energy $\mathcal{H}\Psi_T/\Psi_T$. However, the use of $\tilde{\mathcal{V}}^a$ introduces only negligible errors in the computation of E^a because the regularization is adopted only in a region of volume $S \times a$, where S is the area of the nodal surface \mathcal{N} . Since both the trial and the LRDMC wave function vanish $\approx a$ close to the nodal surface, the finite lattice error corresponds to averaging $(\mathcal{H}^a - \mathcal{H})\Psi_T/\Psi_T$ ($\propto a$) over $\Psi_T\Psi_{\text{FN}}(\propto a^2)$ in a nodal region of extension $\propto a$. If we collect these contributions, we find that the present regularization introduces a bias in the nodal region which vanishes as a^4 for $a \rightarrow 0$ and is always negligible compared to the dominant contribution $O(a^2|\Psi_0 - \Psi_T|^2)$. Moreover, since the regularization in Eq. (28) acts independently on each electron, it does not affect the size-consistent character of the algorithm, and the energy of N independent atoms at large distances is equal to N times the energy of a single atom. Therefore, we did not perform any LRDMC calculations for the oxygen systems since the energy per atom as a function of a is exactly independent of N .

V. PERFORMANCE OF THE PROPOSED METHODS

An important point to address is the efficiency of our revised techniques. This involves not only the computational cost per Monte Carlo step, but also the elimination of the discretization error (in time or space, as appropriate) by extrapolation to the continuum limit. Indeed, a smaller and smoother bias enhances the overall efficiency, as does the knowledge of the leading term in the discretization parameter.




FIG. 3. FN energies as a function of time-step for the oxirane ($\text{C}_2\text{H}_4\text{O}$) molecule, obtained using a simple (top) and a more sophisticated (bottom) trial wave function. We employ different schemes, i.e., the original algorithm as in Ref. 7 and with an improved symmetrized branching factor (“sym”), the two size-consistent approaches we proposed (“DMC version 1” and “DMC version 2”) are the LA approach and the LRDMC method. The lattice-space has been mapped into the time-step via the relation $\tau=0.6a^2$ which guarantees the same autocorrelation time between the “DMC version 1” and the “LRDMC” method in this particular case.

A. Time-step error

We study the time-step error on the FN energy computed with the various algorithms discussed above, using the oxirane molecule ($\text{C}_2\text{H}_4\text{O}$) as a test case. Our aim here is in particular to assess the reduction of the time-step error with respect to the original algorithm.⁷ In the DMC “version 1” this reduction is due to the symmetrization of the weights, while in the DMC “version 2” it is due to the update of the branching factor after SE moves. We employ nonlocal energy-consistent Hartree-Fock pseudopotentials⁹ for the oxygen and the carbon atoms in combination with the corresponding cc-pVDZ basis sets, and construct two single determinant Jastrow-Slater wave functions of different quality. The first wave function is built from B3LYP orbitals and a very simple electron-electron Jastrow factor of the form $b[1 - \exp(-\kappa r_{ij})]/\kappa$, where $b=1/2$ or $1/4$ for antiparallel-spin and parallel-spin electrons, respectively. The parameter κ is optimized in energy minimization and is equal to 1.91. The second wave function is characterized by a more sophisticated Jastrow factor comprising of electron-electron and electron-nucleus terms, and all orbital and Jastrow parameters in the wave function are optimized in energy minimization.

The top panel of Fig. 3 shows results obtained with the

simple wave function. Consistently with previous studies on the water molecule,^{15,16} the LA energies extrapolate to a lower value (not necessarily variational) than the original algorithm of Ref. 7, with a smaller time-step error; symmetrization of the branching factor in the original algorithm is already sufficient to reduce the time-step error down to a value similar to that found in the LA. As expected, given the small size of the system considered, the original and the size-consistent algorithm give nearly identical results, as shown here for its version 1 with AE branching.

The main result shown in the top panel of Fig. 3 is the remarkable reduction of the time-step error obtained with a SE branching factor. The data shown in the figure refer to version 2 of the size-consistent algorithm. We also mention, without reporting the data, that when the branching factor is updated after SE moves, the symmetrization of the local energy in the exponent does not improve the time-step error significantly.

The improvement obtained with a SE branching factor, however, is strongly dependent on the quality of the trial function. The lower panel of Fig. 3 shows results obtained with the more sophisticated wave function. We still find a lower energy with the LA (with its possibly problematic behavior at very small time-step), and a large time-step error with the original algorithm of Ref. 7. All the other cases, however, display similar behavior, or at least comparable quality, in terms of the time-step error.

Also the LRDMC energy values are reported in Fig. 3 where the lattice-space has been converted into time-step based on the equal autocorrelation time between Monte Carlo generations in the DMC and LRDMC algorithms. This is the fairest mapping since it keeps the final statistical error equivalent for the same sample length. In this case, it gives $\tau \sim 0.6a^2$. One can see that the LRDMC energies are always converging from below in a monotonic way, usually easier to extrapolate than the corresponding DMC energies.

In order to make a more quantitative analysis of the predictions reported in Sec. IV for the lattice-space error, we studied the lattice-space extrapolation of the oxirane molecule with the DFT-B3LYP Slater determinant, and Jastrow factors going from the simple two-body one to the most complicated comprising of one-, two-, and three-body terms. The results are reported in Fig. 4. For good trial wave functions, a reliable extrapolation can be obtained even by using very large values of a , where small statistical errors can be obtained with much less computational effort. Also, the FN energies are basically independent of the shape of the trial wave function already for a rather simple Jastrow with one-body and two-body terms, implying that the “locality error” becomes negligible in the variational formulation even for not-so-accurate trial wave functions. This consideration applies also to the DMC variational energies, since the zero-lattice-space zero-time-step limits are equivalent.

B. Relative efficiency

In all the methods presented here, there is an extra computational cost per Monte Carlo step with respect to the stan-




FIG. 4. FN LRDMC energies as a function of lattice-space a for the oxirane (C_2H_4O) molecule, obtained using three types of Jastrow factors. The fitting curves include a quadratic and quartic term, namely, $f(a) = E_0 + ba^2 + ca^4$. The finite lattice-space error improves dramatically with a better wave function. For the simple two-body Jastrow $b \approx -0.15$, while for the most accurate Jastrow factor $b' \approx -0.03$. The ratio of the variances of the two trial wave functions is roughly equal to b/b' , in agreement with Eq. (27).

dard DMC with LA since an extra step is needed in order to sample correctly the Green’s function related to the nonlocal pseudopotentials. However, we have seen that, in *all* the variational methods, the nonlocal pseudopotential operator will displace only one electron at a time since the nonlocal pseudopotential gives a one-body contribution to the Hamiltonian. This means that, in order to update all the quantities needed by the simulation as the wave function ratios $V_{\mathbf{R}',\mathbf{R}}$, the gradients and Laplacian terms, one can exploit the Sherman–Morrison algebra which scales as N^2 . For instance, to update the non-local term $V_{\mathbf{R}',\mathbf{R}}$, as well as the \mathcal{K}^a in the LRDMC, one employs the same algebra as the one used to update the gradient (i.e., the drift term) in the standard DMC with importance sampling. After a single-particle move, the cost to *fully* update $V_{\mathbf{R}',\mathbf{R}}$ scales as $n_{\text{off}}N^2$ where n_{off} is the number of nonlocal mesh points per electron ($n_{\text{off}} = n_{\text{quad}}$ in DMC with nonlocal pseudopotentials and $n_{\text{off}} = n_{\text{quad}} + 6$ in LRDMC with nonlocal pseudopotentials and a single lattice-space in the Laplacian). In the size-consistent DMC (both “version 1” and “version 2”), the pseudopotential move has to be performed N times in a single time-step τ , so the overall cost per time-step coming from the pseudopotential operator is $\eta n_{\text{off}}N^3$, where η is the acceptance ratio of the nonlocal part. Since $n_{\text{off}} \approx 20$, and $\eta \approx 0.1$ at convergence (see Fig. 1), it is clear that the DMC “version 1” will be only a prefactor ≈ 2 slower than the standard DMC with LA. The “version 2” might be slightly slower than the “version 1” since it requires the calculation of the local energy after each single-particle move, but again the difference will be just a prefactor. The LRDMC approach is the slowest because a cycle with N SE updates of the local energy takes $(10 + n_{\text{quad}})/(4 + n_{\text{quad}}) \leq 2.5$ more operations (the worst case is for full-core calculations when $n_{\text{quad}} = 0$). Moreover, there is also an additional slowing down compared to the DMC “version 1” approach because, in the latter case, all operations involving the local energy can be done at the end of a cycle and cast in a very efficient form using matrix-matrix multiplications of size $\sim N$. These operations, for large $N > \approx 1000$, can be much more efficient than SE matrix updates (by a factor ranging from 2 to 20, depending on the computer

hardware and software). At present, it is difficult to estimate how much slower LRDMC will be on a particular machine,¹⁷ also considering that further algorithmic and software developments are expected in the near future, which should allow faster updates. However, even though LRDMC is certainly slower, it has the advantage of a much smoother lattice-space extrapolation as discussed above.

VI. CONCLUSIONS

In conclusion, we have introduced important developments in the DMC and LRDMC methods in the context of electronic structure simulations with nonlocal pseudopotentials. We have explained how to modify the DMC variational formulation for nonlocal potentials of Ref. 7 in order to make it size-consistent. We have shown that, for large systems, the original algorithm⁷ will depart from the usual localization approximation only for small time-steps, making the zero time-step extrapolation possibly problematic. Instead, the two DMC algorithms presented here, based on a more accurate Trotter breakup for the nonlocal operator and a better branching factor, have a smaller and size-consistent time-step error. The DMC version 1, which features a single-particle representation of the nonlocal operator and a branching factor symmetric with respect to the application of the diffusion operator, is straightforward to implement in the existing codes. The DMC version 2 is closer to the LRDMC spirit, since the nonlocal part is further split in τ/N factor always acting on the all-electron configuration, and the branching factor is accumulated after every single-particle move. The latter version can give an even better time-step error [order $O(\tau/N)$ in the nonlocal part], particularly for relatively poor wave functions. In general, it is slightly more time consuming than the version 1, since it requires the evaluation of the full nonlocal matrix after every single-particle move.

We have made significant progress also in the LRDMC approach. In the present formulation, it is no longer necessary to use two lattice meshes to randomize the electron position, but a single lattice-space a is sufficient, provided the orientation of the Cartesian coordinates of the discretized Laplacian is changed randomly during the diffusion process. We have defined a better lattice regularization of the Hamiltonian in order to have always a potential bounded from below, with a cutoff depending on a . This leads to a well defined and size-consistent lattice-space extrapolation since, in the $a \rightarrow 0$ limit, we recover the variational expectation value of the continuous Hamiltonian with a lattice-space error whose leading term is quadratic in a . Moreover, we showed that the prefactor of the a^2 term vanishes quadratically in $|\Psi_0 - \Psi_T|$. Therefore, for good wave functions, the extrapolation to the $a \rightarrow 0$ limit is particularly rapid and

smooth with a computational effort $\propto 1/a^2$. The DMC error appears instead to be less correlated with the quality of the guiding function and may display a turn-down behavior for small time-steps (observed here and elsewhere¹⁵), which makes the time-step extrapolation much harder than in the LRDMC lattice-space approach. Regarding the computational cost, the LRDMC approach is slower but the overall efficiency is comparable to the two variational and size-consistent DMC formulations presented here since LRDMC allows one to work with large values of a due to the robust extrapolation to the zero lattice-space limit.

ACKNOWLEDGMENTS

C.F. acknowledges the support from the Stichting Nationale Computerfaciliteiten (NCF-NWO) for the use of the SARA supercomputer facilities. M.C. thanks the computational support provided by the NCSA of the University of Illinois at Urbana-Champaign. S.S. and S.M. acknowledge the support from CINECA and COFIN'07.

- ¹W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- ²D. M. Ceperley, *J. Stat. Phys.* **43**, 815 (1986).
- ³B. L. Hammond, P. J. Reynolds, and W. A. Lester, Jr., *J. Chem. Phys.* **87**, 1130 (1987).
- ⁴S. Fahy, X. W. Wang, and S. G. Louie, *Phys. Rev. B* **42**, 3503 (1990).
- ⁵L. Mitáš, E. L. Shirley, and D. M. Ceperley, *J. Chem. Phys.* **95**, 3467 (1991).
- ⁶M. Casula, C. Filippi, and S. Sorella, *Phys. Rev. Lett.* **95**, 100201 (2005).
- ⁷M. Casula, *Phys. Rev. B* **74**, 161102(R) (2006).
- ⁸D. F. B. ten Haaf, H. J. M. van Bommel, J. M. J. van Leeuwen, W. van Saarloos, and D. M. Ceperley, *Phys. Rev. B* **51**, 13039 (1995).
- ⁹M. Burkatzki, C. Filippi, and M. Dolg, *J. Chem. Phys.* **126**, 234105 (2007).
- ¹⁰C. Filippi and C. J. Umrigar, *J. Chem. Phys.* **105**, 213 (1996). As Jastrow correlation factor, we use the exponential of the sum of two fifth-order polynomials of the electron-nuclear (e-n) and the electron-electron (e-e) distances, respectively. The Jastrow factor is adapted to deal with pseudoatoms, and the scaling factor κ is set to 0.8 a.u. for the oxygen systems and 0.5 a.u. for oxirane.
- ¹¹C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Henning, *Phys. Rev. Lett.* **98**, 110201 (2007).
- ¹²The short time propagator of DMC version 2 lends itself to an implementation of the Reptation quantum Monte Carlo method [S. Baroni and S. Moroni, *Phys. Rev. Lett.* **82**, 4745 (1999)] with single-particle moves.
- ¹³For heavy elements treated either as full-core atoms or with small-core pseudopotentials, the two-mesh discretization might remain a better choice since it allows a more efficient sampling of the different length scales in the core and valence regions.
- ¹⁴S. Sorella and L. Capriotti, *Phys. Rev. B* **61**, 2599 (2000).
- ¹⁵I. G. Gurtubay and R. J. Needs, *J. Chem. Phys.* **127**, 124306 (2007).
- ¹⁶M. W. Schmidt, K. K. Baldrige, J. A. Boatz, S. T. Elbert, M. S. Gordon, J. H. Jensen, S. Koseki, N. Matsunaga, K. A. Nguyen, S. Su, T. L. Windus, M. Dupuis, and J. A. Montgomery, *J. Comput. Chem.* **14**, 1347 (1993).
- ¹⁷The relative efficiency between LRDMC and “DMC version 1” within the same machine (Intel quad-core) and the same code (Ref. 18) is less than a factor of 2 for the cases shown in this work.
- ¹⁸S. Sorella, <https://qe-forge.org/projects/turborvb/>.