



HAL
open science

Systematic lowering of the scaling of Monte Carlo calculations by partitioning and subsampling

Antoine Bienvenu, Jonas Feldt, Julien Toulouse, Roland Assaraf

► **To cite this version:**

Antoine Bienvenu, Jonas Feldt, Julien Toulouse, Roland Assaraf. Systematic lowering of the scaling of Monte Carlo calculations by partitioning and subsampling. 2022. hal-03656059v1

HAL Id: hal-03656059

<https://hal.science/hal-03656059v1>

Preprint submitted on 30 Apr 2022 (v1), last revised 28 Jul 2022 (v2)

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

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

Systematic lowering of the scaling of Monte Carlo calculations by partitioning and subsampling

Antoine Bienvenu,¹ Jonas Feldt,¹ Julien Toulouse,^{1,2} and Roland Assaraf^{1,*}

¹*Laboratoire de Chimie Théorique, Sorbonne Université and CNRS, F-75005 Paris, France*

²*Institut Universitaire de France, F-75005 Paris, France*

(Dated: April 15, 2022)

We propose to compute physical properties by Monte Carlo calculations using conditional expectation values. The latter are obtained on top of the usual Monte Carlo sampling by partitioning the physical space in several subspaces or fragments, and subsampling each fragment (i.e., performing side-walks) while freezing the environment. No bias is introduced and a zero-variance principle holds in the limit of separability, i.e. when the fragments are independent. In practice, the usual bottleneck of Monte Carlo calculations – the scaling of the statistical fluctuations as a function of the number of particles N – is relieved for extensive observables. We illustrate the method in variational Monte Carlo on the 2D Hubbard model and on metallic hydrogen chains using Jastrow-Slater wave functions. A factor $\mathcal{O}(N)$ is gained in numerical efficiency.

Many domains of physics involve large dimensional integrals which can be computed efficiently with Monte Carlo methods, e.g. statistical physics [1], quantum physics applied to molecules and solids [2], or nuclear physics [3]. Monte Carlo methods reinterpret the energy or other properties as the expectation value of a random variable O over a probability distribution π on a configuration space Ω

$$\mathbb{E}(O) = \int_{x \in \Omega} O(x) \pi(x) dx. \quad (1)$$

Typically, the configuration x corresponds to the $3N$ coordinates of the particles in physical space, but it can also correspond to the N trajectories of the particles in the path-integral formulation of quantum mechanics. The probability distribution π depends on the context. For example, in equilibrium statistical physics, π is the Gibbs distribution. In variational Monte Carlo (VMC), $\pi = \Psi^2$ is the probability density of a wave function Ψ , and if $O = (H\Psi)/\Psi$ is the local energy for a given Hamiltonian H then $\mathbb{E}(O)$ is the variational energy. Expectation values are computed using the ergodic theorem which states that the integral can be written as a time average, $\mathbb{E}(O) = \lim_{M \rightarrow \infty} (1/M) \sum_{i=1}^M O(x^i)$, where the sequence of M configurations (x^i) is built from a π -invariant ergodic stochastic process (usually a Markov chain). The sequence (x^i) is called a sample of the distribution π .

The bottleneck of Monte Carlo methods comes from the statistical fluctuations which usually grow with the system size, as measured by the number of particles N . For a sample of sufficiently large size M , the statistical uncertainty σ on the estimation of $\mathbb{E}(O)$ is

$$\sigma = \sqrt{\frac{V(O)c}{M}}, \quad (2)$$

where $V(O) = \mathbb{E}(O^2) - \mathbb{E}(O)^2$ is the variance of O and $c > 1$ is a correlation factor which takes into account that the configurations are not fully independent. According to

Eq. (2), reaching a given precision σ requires a CPU time $t_M = Mt_1$ proportional to both the time t_1 of performing one step of the sampling and to the variance $V(O)$. The numerical efficiency of the method can then be measured by the asymptotically M -independent quantity

$$\sigma^2 t_M = V(O)ct_1, \quad (3)$$

which should be as small as possible for maximal efficiency. In the present work, we will not be concerned about the correlation factor c which sometimes diverges with N (e.g. near criticality). A large corpus of work is devoted to reducing its scaling as a function of N , such as parallel tempering based methods (see, e.g., Refs. [4, 5]). Equation (3) indicates a more crucial double penalty of Monte Carlo methods for large systems: both t_1 and $V(O)$ grow with system size N . This double penalty is for example at the origin of the main bottleneck in computing the VMC energy of a fermionic system in real space [2, 6]. Evaluating the wave function involves indeed calculating a Slater determinant of order $\mathcal{O}(N \times N)$ which costs $t_1 = \mathcal{O}(N^3)$ while the variance is typically extensive, $V(O) \propto N$, thus rising the scaling of the overall cost to $\mathcal{O}(N^4)$. This scaling is still larger than some deterministic methods like the celebrated Kohn-Sham density-functional theory which scales as $\mathcal{O}(N^3)$ for a spatially delocalized (i.e., metallic) system [7].

The extensivity of the variance has a physical origin. A large system can in general be approximated by a collection of independent fragments. This ideal case corresponds to the separability limit where the random variable O is the sum of independent variables O_k on each fragment indexed by k , i.e. $O = \sum_k O_k$, and the variance is then $V(O) = \sum_k V(O_k) \propto N$. It is possible to reduce considerably the variance using an improved estimator \tilde{O} built from the approximate solution of a partial differential equation [8–10]. But this type of improved estimator is still a sum of independent random variables in the separability limit, i.e. $\tilde{O} = \sum_k \tilde{O}_k$, and thus does

not change the scaling with respect to N but only reduces the prefactor [11].

To reduce the global computational scaling, a common and obvious strategy is to reduce the cost of the sampling. Some distributions π can be sampled with a linear-scaling algorithm, i.e. $t_1 = \mathcal{O}(N)$, reducing the overall cost to an ideal scaling $\mathcal{O}(N^2)$. One can for example try to use the sparsity of the Slater matrix when localized Wannier functions are used [12]. But such sparsity is highly dependent on the physics of the system, and does not hold for a metallic system. Besides, this linear scaling is only theoretical because of memory-access slow down as N increases. Another strategy consists in using a stable-versus-chaos stochastic dynamics [13] but finding such a stochastic dynamic is not straightforward [14].

Here we propose to reduce the global computational scaling by using the locality of physical observables. The idea of using the locality of information to reduce the variance is not new: the strong locality in time of the Schrödinger equation (a first-order partial differential equation in time) has for example been exploited to remove the dynamical sign problem for bosonic systems [15]. Recently, a method was proposed [16] to exploit the low correlation between different core regions in a molecule, resulting in a reduced scaling as a function of the atomic charge Z but not as a function of N . The present work exploits the fact that in an extended physical system (including a metallic system) correlations between large fragments are small. We construct an improved estimator \tilde{O} with a variance having a reduced scaling with respect to N , without changing the scaling of t_1 , therefore achieving a reduction of the overall computational scaling. The present work shares the same general philosophy as other fragment-based methods (see, e.g., Refs. 17–19). However, while the latter methods are systematic techniques to find a good compromise between a smaller computational time and a larger systematic error, in the present method the reduction of the computational scaling is done without introducing any systematic error.

Theory — A configuration of particles is written as $x = (x_j)_{j \in J}$ where x_j is the j^{th} coordinate and J is the list of coordinate indexes. For a given configuration $x = (x_j)_{j \in J}$, we define a partition of J as p disjoint subsets $J_k(x) \subset J$ such that $\bigcup_{k=1}^p J_k(x) = J$. We also define p subsets $\Omega_k(x)$ of the configuration space Ω such that for any element $x' = (x'_j)_{j \in J}$ in $\Omega_k(x)$, $j \notin J_k(x) \implies x'_j = x_j$. In other words, only the coordinates of x' indexed in $J_k(x)$ can be different from their counterpart in x . These subsets of the configuration space will correspond to moves of some particles in a given region of space (or a fragment of the system), while freezing the other particles. Note that according to this definition $x \in \Omega_k(x)$. We then introduce the following

improved estimator

$$\tilde{O} \equiv O + \sum_{k=1}^p \lambda_k (\mathbb{E}(O|\Omega_k) - O), \quad (4)$$

where λ_k are constants and $\mathbb{E}(O|\Omega_k)$ is the conditional expectation value of the random variable O with respect to Ω_k , defined as the random variable obtained by averaging O over only configurations $x' \in \Omega_k$

$$\mathbb{E}(O|\Omega_k) \equiv \frac{\int_{x' \in \Omega_k} O(x') \pi(x') dx'}{\int_{x' \in \Omega_k} \pi(x') dx'}. \quad (5)$$

The estimator \tilde{O} in Eq. (4) is never biased, i.e. $\mathbb{E}(\tilde{O}) = \mathbb{E}(O)$, because of the well-known law of total expectation $\mathbb{E}(\mathbb{E}(O|\Omega_k)) = \mathbb{E}(O)$. We note that λ_k might more generally be constant in Ω_k while depending in principle on the environment of the fragment. Let us prove now that the estimator \tilde{O} has a zero-variance property in the separability limit when we choose $\forall k, \lambda_k = 1$. In this limit, O is a sum of p independent contributions on each fragment, $O = \sum_{k=1}^p O_k((x_j)_{j \in J_k})$. Independence then implies that $\mathbb{E}(O|\Omega_k) - O = \mathbb{E}(O_k) - O_k$ and thus $\tilde{O} = \sum_{k=1}^p \mathbb{E}(O_k) = \mathbb{E}(O)$, which proves that the estimator \tilde{O} has indeed a variance of zero in the separability limit for $\lambda_k = 1$. In this limit, only one parent configuration x is sufficient for sampling \tilde{O} , the algorithm becomes equivalent to p independent Monte Carlo simulations of the p subsystems.

Of course, we do not know $\mathbb{E}(O|\Omega_k)$, but we can sample it from the marginal distribution $\pi(\cdot|\Omega_k)$. This is done through a side-walk which samples only Ω_k , i.e. moving the coordinates indexed by J_k in a given fragment while the other coordinates are frozen. In practice, we thus use the following improved estimator

$$\tilde{O} \equiv O + \sum_{k=1}^p \frac{\lambda_k}{m_k} \sum_{i=1}^{m_k} (O_k^i - O), \quad (6)$$

where O_k^i is the value of the random variable O at the i^{th} step of the k^{th} side-walk (moving only the coordinates indexed by J_k) of length m_k . We call this scheme the partition Monte Carlo (PMC) method. We expect this scheme to reduce the variance with a low numerical cost because the p subsamplings correspond to handling $p = \mathcal{O}(N)$ low-dimensional problems.

In practice, the parameters λ_k and m_k have to be adjusted to lower the variance of \tilde{O} for a given CPU time. Also, for optimal efficiency, we can generalize the estimator \tilde{O} in Eq. (4) using different control variates $\mathbb{E}(G_k|\Omega_k) - G_k$ which converge to $\mathbb{E}(O|\Omega_k) - O$ in the separability limit. The random variables G_k can be obtained from O by neglecting terms outside of the fragment k , reducing the computational cost while retaining the unbiasedness and the zero-variance property in the separability limit. For example, for a fragment k of a

molecule containing n_k electrons we can take G_k as the local energy obtained from the truncated Hamiltonian

$$H_k = \sum_{i=1}^{n_k} \left(-\frac{1}{2} \nabla_i^2 - \sum_A \frac{Z_A}{r_{iA}} + \sum_j \frac{1}{r_{ij}} \right), \quad (7)$$

where i runs over the electrons in the fragment. The first term is the kinetic-energy operator and the last two terms are the Coulomb interactions of the electrons of the fragment with the nuclei A (charges Z_A) and electrons j lying in a given neighborhood of the fragment.

Let us see now how the PMC method relieves the variance bottleneck. As an example, we consider VMC calculations using Jastrow-Slater wave functions

$$\Psi(x) = e^{J(x)} \Phi(x), \quad (8)$$

where $J(x)$ is any real symmetric function of the electron configuration x , and $\Phi(x) = \det(A)$ with the Slater matrix $A = XC$ where X is a rectangular matrix of localized atomic orbitals (Kronecker functions in the case of a lattice model) and C is the rectangular matrix of the orbital coefficients. For one fragment of the system we introduce now the matrix P which selects the lines corresponding to the electrons of that fragment. For a side-walk in that fragment, X takes different values X' such that only the lines PX might differ from the lines PX' . The new determinant is [20, 21]

$$\begin{aligned} \Phi(x') &= \det(X'C) \\ &= \det(A) \det(X'CA^{-1}) \\ &= \det(A) \det(PX'Q^T QCA^{-1}P^T), \end{aligned} \quad (9)$$

where we have used the determinant lemma. We inserted the projector $Q^T Q$ where Q^T selects on the right of PX' only the few columns which may differ from zero for this fragment. These columns are very few because the atomic orbitals are localized. In conclusion updating the determinant along the side-walk is equivalent to multiplying it by a low-order effective Slater determinant

$$\Phi(x') = \det(A) \det(\bar{X}\bar{C}), \quad (10)$$

where $\bar{X} = PX'Q^T$ and $\bar{C} = QCA^{-1}P^T$. The matrix \bar{C} represents effective orbitals for the fragment and is computed only once at each step of the usual main walk, at a $\mathcal{O}(N^3)$ numerical cost. Once \bar{C} has been built and stored the side-walk costs only $\mathcal{O}(n^3)$ where n is the number of electrons in the fragment. The local energy of the subsystem involves a truncated Hamiltonian and can be computed with the same cost $\mathcal{O}(n^3)$ [20, 21]. The cost of subsampling $\mathcal{O}(N)$ fragments is thus $\mathcal{O}(N)$ for an extended system with a finite correlation length. This allows us to perform up to $\sum_k m_k = \mathcal{O}(N^3)$ total steps in the side-walks without modifying the scaling of the main walk. Therefore, we can perform $m_k = \mathcal{O}(N^2)$ steps in each fragment and the improved estimator in Eq. (6) will

have consequently a variance reduced by a factor up to $\mathcal{O}(N^2)$, which is achieved in the separability limit.

Results — We now illustrate the PMC method on the calculation of the ground-state energy of the 2D Hubbard model and of metallic hydrogen chains.

The Hubbard systems that we employ consist in 2D square grids of $L \times L$ sites with periodic boundary conditions, filled to half-capacity with $N \approx L^2$ electrons evenly distributed between the spins. Designating by $c_{i\sigma}^\dagger$ and $c_{i\sigma}$ the creation and annihilation operators of site i with spin $\sigma \in \{\uparrow, \downarrow\}$, and by $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ the corresponding number operators, the Hamiltonian takes the form [22]

$$H = - \sum_{i \neq j, \sigma} t_{ij} c_{i\sigma}^\dagger c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (11)$$

where $t_{ij} = 1$ if i and j are adjacent, and $t_{ij} = 0$ otherwise, and $U = 1$ is the on-site interaction parameter. We have chosen the trial ground-state wave function to be a Slater determinant of plane waves without any Jastrow factor. We choose the subsystems as adjacent squares of $l \times l$ sites. The number of iterations of the main walk is kept constant at $M = 500$.

As an example of a simple system with a continuum configuration space, we consider metallic hydrogen chains with a regular interatomic distance of $1.4 a_0$. The Hamiltonian is given by Eq. (7) except of course that there is no restriction in the sums for the full system. For the trial ground-state wave function, we use a simple Jastrow function [16] multiplied by the Hartree-Fock Slater determinant obtained from a basis made of the exact hydrogen 1s orbital on each atom. We choose the subsystems as consisting in n adjacent hydrogen atoms.

The first parameter of the PMC method whose impact is to be explored is the side-walk length m (chosen to be the same for all subsystems). Figure 1 reports the variance gain $V_{\text{VMC}}/V_{\text{PMC}}$, the CPU time increase $t_{\text{PMC}}/t_{\text{VMC}}$, and the uncorrelated efficiency gain $(V_{\text{VMC}} t_{\text{VMC}})/(V_{\text{PMC}} t_{\text{PMC}})$ [efficiency gain assuming a correlation factor $c = 1$] of the PMC method over the standard variational Monte Carlo (VMC) method. The efficiency gain is plotted as a function of the side-walk length m for the 2D Hubbard model with total size $L = 20$ and subsystem size $l = 5$, and for hydrogen chains with $N = 320$ total atoms and $n = 12$ atoms in the subsystems. Two regimes are clearly visible. For small m , the variance gain increases linearly with m while the CPU time is almost constant (the cost of a side-walk step is very small compared to that of a main walk step). This leads to a linear increase of the uncorrelated efficiency gain. For large m , the variance gain saturates while the CPU time ratio increases linearly, driving the uncorrelated efficiency gain down. Between these two regimes, there is a plateau corresponding to optimal values of the side-walk length m . The saturation of the variance gain originates from the correlation between subsystems. In-




FIG. 1. Variance gain $V_{\text{VMC}}/V_{\text{PMC}}$, time increase $t_{\text{PMC}}/t_{\text{VMC}}$, uncorrelated efficiency gain $(V_{\text{VMC}}t_{\text{VMC}})/(V_{\text{PMC}}t_{\text{PMC}})$ of the PMC method over the standard VMC method as a function of the side-walk length m for (a) the 20×20 square Hubbard model (half filling) and (b) the H_{320} metallic hydrogen chain. The real efficiency gain $(\sigma_{\text{VMC}}^2 t_{\text{VMC}})/(\sigma_{\text{PMC}}^2 t_{\text{PMC}})$ differs from the uncorrelated efficiency gain only for hydrogen chains.

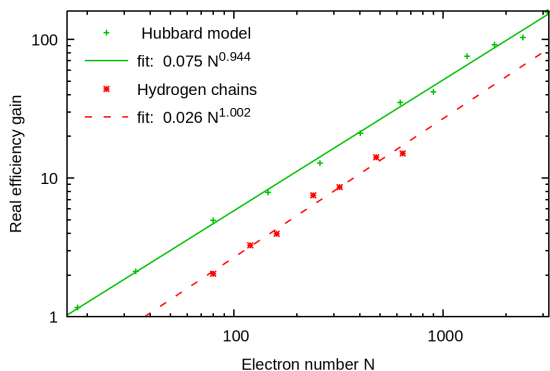


FIG. 2. Optimal real efficiency gain $(\sigma_{\text{VMC}}^2 t_{\text{VMC}})/(\sigma_{\text{PMC}}^2 t_{\text{PMC}})$ as a function of electron number N for the Hubbard model and metallic hydrogen chains.

deed, if the subsystems were independent, the variance would converge to zero as m increases (zero-variance principle in the separability limit) and the variance gain to infinity.

One may ask the role of the correlation factor c in Eq. (2). For the Hubbard model, c has been found to be very close to 1 leading to a real efficiency gain almost identical to the uncorrelated efficiency gain. For the hydrogen chains $c \simeq 2.5$ for $m = 0$ (VMC) and c is reduced for small m (about 40% less for H_{320} and $m \in [5, 40]$) before increasing slowly for larger values of m . This explains the difference between the uncorrelated efficiency gain and the real efficiency gain $(\sigma_{\text{VMC}}^2 t_{\text{VMC}})/(\sigma_{\text{PMC}}^2 t_{\text{PMC}})$ in Fig. 1. In particular, the optimal real efficiency gain is 40% higher than the optimal uncorrelated efficiency gain.

We now consider systems of increasing sizes. For the Hubbard model, the optimal subsystem size has been found to be $l \approx \sqrt{L}$, and similarly for the metallic hydrogen chains we find $n \approx \sqrt{N}/2$. The fact that the op-

timal subsystem size does not saturate to a finite value as the system size increases is an indication of the non-separability of the system. The optimal side-walk length m also increases with system size since larger systems result in more decorrelated subsystems and cheaper side-walks compared to the main walk. Figure 2 reports the real efficiency gain as a function of the electron number N for the Hubbard model and the hydrogen chains up to N of the order of 10^3 . Both metallic systems present a real efficiency gain scaling linearly with N , which hovers around $0.075N$ for the Hubbard model and $0.025N$ for the hydrogen chains. This real efficiency gain is entirely achieved by decreasing the variance of the local energy from $\mathcal{O}(N)$ to a behavior close to $\mathcal{O}(1)$, as shown in Fig. 3.

Conclusions — We introduced a general and simple method to reduce the scaling of Monte Carlo calculations of extensive properties. It only requires to have an explicit formula [Eq. (1)] for the integral to be computed, and therefore can be used in any Markov Chain Monte Carlo application. The method was illustrated on VMC calculations of metallic systems of N particles, providing an efficiency gain of order $\mathcal{O}(N)$. The present idea can be applied in many contexts, including fixed-node path-integral Monte Carlo approaches [23, 24] since these schemes sample explicit probability distributions. Finally, the method can in principle be extended to derivatives of extensive properties to reduce the scaling for calculating response properties or optimizing variational wave functions.

* assaraf@lct.jussieu.fr

[1] K. Binder and D. W. Heermann, *Monte Carlo Simulation in Statistical Physics: An Introduction; 5th ed.*, Graduate Texts in Physics (Springer, Berlin, Heidelberg, 2010).

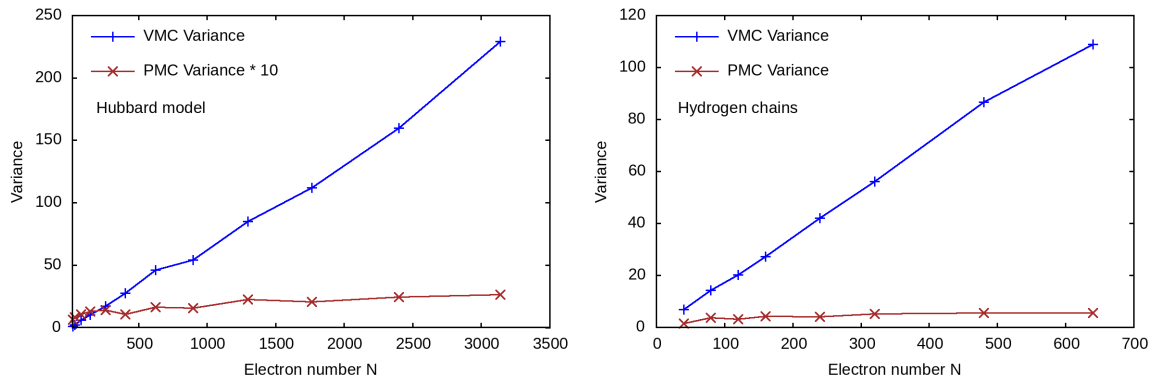


FIG. 3. Variance of the local energy in standard VMC and in PMC (for optimal m) as a function of electron number N for (a) the Hubbard model (PMC variance multiplied by 10 on the plot) and (b) metallic hydrogen chains.

- [2] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- [3] J. Lynn, I. Tews, S. Gandolfi, and A. Lovato, *Annu. Rev. Nucl. Part. Sci.* **69**, 279 (2019).
- [4] J. Goodman and A. D. Sokal, *Phys. Rev. D* **40**, 2035 (1989).
- [5] J. Weare, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 12657 (2007).
- [6] J. Toulouse, R. Assaraf, and C. J. Umrigar, *Adv. Quantum Chem.* **73**, 285 (2016).
- [7] S. Mohr, M. Eixarch, M. Amsler, M. J. Mantsinen, and L. Genovese, *Nucl. Mater. Energy* **15**, 64 (2018).
- [8] R. Assaraf and M. Caffarel, *Phys. Rev. Lett.* **83**, 4682 (1999).
- [9] A. Mira, R. Solgi, and D. Imperato, *Stat. Comput.* **23**, 653 (2012).
- [10] D. Borgis, R. Assaraf, B. Rotenberg, and R. Vuilleumier, *Mol. Phys.* **111**, 3486 (2013).
- [11] R. Assaraf and D. Domin, *Phys. Rev. E* **89**, 033304 (2014).
- [12] A. J. Williamson, R. Q. Hood, and J. C. Grossman, *Phys. Rev. Lett.* **87**, 246406 (2001).
- [13] R. Assaraf, *Phys. Rev. E* **90**, 063317 (2014).
- [14] R. Assaraf, B. Jourdain, T. Lelièvre, and R. Roux, *Stoch. Partial Differ. Equ.: Anal. Comput.* **6**, 125 (2017).
- [15] G. Cohen, E. Gull, D. R. Reichman, and A. J. Millis, *Phys. Rev. Lett.* **115**, 266802 (2015).
- [16] J. Feldt and R. Assaraf, *J. Chem. Theory Comput.* **17**, 1380 (2021).
- [17] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [18] G. Knizia and G. K.-L. Chan, *Phys. Rev. Lett.* **109**, 186404 (2012).
- [19] F. Zahariev and M. S. Gordon, *Phys. Chem. Chem. Phys.* **23**, 14308 (2021).
- [20] C. Filippi, R. Assaraf, and S. Moroni, *J. Chem. Phys.* **144**, 194105 (2016).
- [21] R. Assaraf, S. Moroni, and C. Filippi, *J. Chem. Theory Comput.* **13**, 5273 (2017).
- [22] M. Cyrot, *Physica B+C* **91**, 141 (1977).
- [23] S. Baroni and S. Moroni, *Phys. Rev. Lett.* **82**, 4745 (1999).
- [24] J. Shumway and M. Gilbert, *J. Phys. Conf. Ser.* **35**, 190 (2006).