

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Reducing the Cost of Energy Differences in Variational Monte Carlo with Spotlight Sampling

### Permalink

<https://escholarship.org/uc/item/4w10v3rk>

### Journal

Journal of Chemical Theory and Computation, 21(22)

### ISSN

1549-9618

### Authors

Bumann, Sonja  
Neuscamman, Eric

### Publication Date

2025-11-25

### DOI

10.1021/acs.jctc.5c01232

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Reducing the Cost of Energy Differences in Variational Monte Carlo with Spotlight Sampling

Sonja Bumann<sup>1,2</sup> and Eric Neuscamman<sup>1,2,a)</sup>

<sup>1)</sup>*Department of Chemistry, University of California, Berkeley, California 94720, USA*

<sup>2)</sup>*Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

(Dated: 28 July 2025)

We investigate an approximate sampling scheme that can significantly reduce the cost scaling of variational Monte Carlo when it is employed to predict the energy differences associated with local chemical changes. Inspired by side-chaining and embedding methods, this spotlight sampling approach adopts an approximate fragmented Hamiltonian and correlated sampling to reduce cost scaling to the point that it is essentially linear with system size, with the potential to go sub-linear if certain conditions are met. In tests on bond stretching energies in alcohols and hydrogen dimer chains, we observe the anticipated linear scaling as well as an explicit cost crossover with standard variational Monte Carlo.

## I. INTRODUCTION

Simultaneously capturing the effects of both weak and strong electron correlation remains challenging for traditional quantum chemistry methods, especially in excited states and large systems.<sup>1–8</sup> In many difficult cases, Quantum Monte Carlo (QMC) methods have proven to be effective alternatives.<sup>9–20</sup> For example, in some excited states, even variational Monte Carlo (VMC), perhaps the most straightforward QMC technique, has proven more accurate than full-triples equation of motion coupled cluster theory.<sup>19,20</sup> However, the significant computational cost required to evaluate large numbers of random samples continues to limit the applicability of both VMC and QMC methods more generally. Indeed, such sampling places an unusually large prefactor in front of VMC’s quartic cost growth with system size,<sup>21</sup> meaning that the advantages of this modest polynomial scaling can be difficult to realize in practice. In this study, we show that a combination of approximate fragmentation and correlated sampling can reduce both cost and scaling while maintaining accuracy when using VMC to predict energy differences associated with local chemical changes, as for example when stretching a chemical bond.

It has long been recognized that, thanks to correlated sampling techniques, it can be advantageous to directly estimate energy differences in VMC rather than attempting to make independent estimates of the energy before and after a chemical change.<sup>22</sup> The basic idea is that, for regions of the system that are only weakly affected by the change, statistical fluctuations will mostly cancel out so long as essentially the sampling of electron positions in those regions is used on either side of the energy difference. The space warp transformation, which we employ in the present study, achieves this goal explicitly.<sup>23,24</sup> What’s more, by trading the task of estimating two extensive energies for the task of estimating an intensive

energy difference, such correlated sampling approaches should in general drop the cost scaling of VMC from quartic to cubic. Here, we seek to combine this advantage with the advantages of local fragmentation approaches, in particular approaches inspired by recent side-chaining methods.<sup>21,25</sup>

Locality has, more generally, long been exploited in attempts to mitigate the cost of VMC.<sup>26–29</sup> Early works focused on leveraging the sparsity of the Slater matrix when using local orbitals in order to reduce the costs associated with Slater determinant ratios,<sup>30–33</sup> which can also be mitigated by complimentary delayed update methods.<sup>34</sup> More recently, Feldt and coworkers have shown that that evaluating local side chains, in which only a few of the electrons near a given nucleus are allowed to move, is an effective way to reduce the cost of dealing with core electrons.<sup>25</sup> This side-chaining approach was then generalized in a study by Bienvenu and coworkers, in which they show that it can reduce the cost of VMC sampling by one power of the system size.<sup>21</sup>

These side-chaining methods have strong parallels to embedding methods,<sup>35–38</sup> in that the side chains sample some electrons well while not sampling the “environment” electrons at all, but they differ in one crucial respect. Unlike embedding methods, in which the interaction energy between the system and its environment is approximated, these two side-chain studies carefully set up their overall energy estimates to be unbiased, and so they do not add any approximation. This juxtaposition begs the question: what might we be able to gain if we instead intentionally introduced simplifications that biased the results? Can the cost growth with system size be reduced further while maintaining reasonable accuracy for relevant chemical quantities like bond stretching energies? In this study, we embark on a preliminary investigation of these questions, concluding that, yes, for some types of predictions, a combination of correlated sampling, side-chaining-inspired fragmentation, and embedding-inspired Hamiltonian approximations allow for VMC cost scaling to be reduced further,

<sup>a)</sup>Electronic mail: eneuscamman@berkeley.edu

in principle to the point that it can become linear or even sub-linear in the system size.

This paper is organized as follows. We begin in Section II A with a theoretical analysis of how the cost scaling of drawing samples in VMC can be reduced by a combination of various techniques, within which we include a review of the fast Slater determinant math used in Markov chains in which only a small number of electrons move. During this analysis, we do not address the question of accuracy, relying instead on the assumption that certain sampling conditions can be met without doing too much damage to VMC’s energy difference predictions. After completing that analysis, we then introduce in Section II B an explicit approximation scheme, spotlight sampling, which seeks to meet these conditions while maintaining accuracy, focusing in particular on how to deal with both erroneous Pauli exclusion effects, and, in Section II C, medium and long range Coulomb interactions. With an approximate Hamiltonian defined, Section II D then completes the cost analysis by showing that its local energy evaluations do not alter the conclusions of the initial sampling cost analysis. We then touch on (Section II E) how this approach intersects with the more general VMC challenge of maintaining numerical stability during long chains of low-rank Slater matrix inverse updates before turning to our results. In these results (Section III), we first verify the necessity of spotlight sampling’s buffering approach, after which we explore proof of principle calculations that demonstrate the cost advantages and test accuracy in a series of alcohols’ O-H bond stretches. We then end in Section IV with some concluding remarks.

## II. THEORY

### A. Reducing the Cost of Sampling

In a straightforward approach to energy estimates in VMC, one approximates the energy as an average of the local energies over samples drawn from the wave function’s probability distribution.

$$\begin{aligned}
 E &= \frac{\langle \Psi | \hat{H} | \Psi \rangle}{\langle \Psi | \Psi \rangle} \\
 &= \frac{\int \Psi(R) \hat{H} \Psi(R) dR}{\int \Psi(R) \Psi(R) dR} \\
 &= \int \rho(R) \frac{\hat{H} \Psi}{\Psi} dR \\
 &= \int \rho(R) E_L(R) dR \approx \frac{1}{K} \sum_m^K E_L(R_m)
 \end{aligned} \tag{1}$$

Here  $E_L(R) = \frac{\hat{H} \Psi(R)}{\Psi(R)}$  is the local energy at the electron positions  $R_m$  of the  $m$ th sample, and  $\rho(R) = |\Psi(R)|^2 / \int |\Psi(R)|^2 dR$  is the 3N-dimensional probability

distribution for the  $N$  electrons. In practice,  $\Psi$  is often chosen as  $\Psi(R) = e^{J(R)} D(R)$ , where  $D(R)$  is a Slater determinant and  $e^{J(R)}$  is a Jastrow factor,<sup>39,40</sup> while the  $K$  samples are drawn from  $\rho(R)$  via a Metropolis-Hastings Markov chain<sup>41,42</sup> in which one electron is moved at a time.<sup>43</sup>

In this straightforward approach to estimating  $E$ , the cost of drawing samples is driven by the expense of evaluating the ratio between the new and old determinant values when one electron is moved. After such a move, one row will have changed in the matrix  $M = XC$  of molecular orbital (MO) values, in which  $X$  and  $C$  are the atomic orbital value matrix and the matrix of occupied MO coefficients, respectively. The ratio of the determinant values needed by Metropolis Hastings when moving the  $i$ th electron is then found (via the matrix determinant lemma) by taking the dot product of the  $i$ th row of the new  $M$  with the  $i$ th column of the inverse of the old  $M$ .

$$\frac{D(R_{\text{new}})}{D(R_{\text{old}})} = \sum_j [M_{\text{new}}]_{ij} [(M_{\text{old}})^{-1}]_{ji} \tag{2}$$

To reduce the cost of working with  $M^{-1}$ , it is updated at  $O(N^2)$  cost via the Sherman-Morrison formula after each accepted move.<sup>44–46</sup> Generating a new and independent sample for use in Eq. (1) requires  $O(N)$  of these one-electron moves, and, at least when we are attempting to estimate the total energy  $E$ , the overall number of samples  $K$  needs to grow linearly with  $N$  to counteract the linear growth in the variance.<sup>21</sup> All told, the computational effort comes out as  $O(N^4)$ , which breaks down as taking  $O(N)$  samples, each of which requires  $O(N)$  moves, each of which bears an  $O(N^2)$  cost for the Sherman-Morrison update.

In this study, we will explore a combination of correlated sampling and fragmentation in order to approximate energy differences for local chemical changes at reduced cost. To start, let us focus first on the effort required to draw samples, and in particular how a fragmentation of the system can reduce this effort. Imagine if, as in a PMC side chain,<sup>21</sup> we only ever had to consider situations in which a small group of  $O(1)$  electrons had moved away from their initial positions, while all other electrons remained frozen where they started. In that case, only  $O(1)$  of the rows of  $M$  would ever change, allowing for an efficient evaluation of a key determinant ratio.

$$\frac{D(R_{\text{new}})}{D(R_{\text{init}})} = \frac{\det(M_{\text{new}})}{\det(M_{\text{init}})} = \det(YZ), \tag{3}$$

Here  $Y$  is the rectangular matrix formed by the rows of  $M_{\text{new}}$  that have actually changed and  $Z$  is the matrix formed from the corresponding columns of  $(M_{\text{init}})^{-1}$ . The ratio  $D(R_{\text{new}})/D(R_{\text{old}})$  needed for Metropolis is just the ratio of two evaluations of Eq. (3), and so the cost of a single-electron move is thus dominated by the cost of forming the  $O(1)$  by  $O(1)$  matrix  $YZ$ . If each MO

is localized in space and only has nonzero contributions from  $O(1)$  atomic orbitals,  $YZ$  can be formed and its determinant taken in  $O(1)$  time. If the MOs are not localized, then forming  $YZ$  will require  $O(N)$  time. Either way, this is a reduction compared to the  $O(N^2)$  time per move required by the straightforward approach above.

Although at first glance it sounds like a drastic approximation, our approach will put us in this computationally desirable situation by breaking the overall system into  $O(N)$  local fragments (e.g.,  $\text{CH}_2$  groups), after which we will approximate each fragment’s energy using samples in which only that fragment’s and  $O(1)$  neighboring fragments’ electrons move. Because other electrons are never able to relax to the wave function’s stationary distribution when evaluating a given fragment’s energy, this approach will be biased, and we will discuss below how we minimize the effects of this bias on our overall energy estimate. For now, we focus on the resulting computational cost. If each fragment contains  $O(1)$  atoms and we were to take  $O(1)$  samples per fragment, we could expect the uncertainty in each fragment’s energy contribution to be independent of system size. However, adding these independent fragment energies together would produce an overall energy whose uncertainty grows as  $\sqrt{N}$ , and so, if the goal was to estimate that overall energy with a fixed uncertainty, we would instead need to take  $O(N)$  samples per fragment. Thus, if all we did was apply our fragmentation scheme and then attempt to estimate the total energy with fixed uncertainty, the cost would be  $O(N^3)$  or  $O(N^2)$  depending on whether the orbitals (and Jastrow) were nonlocal or local, which breaks down as  $O(N^2)$  samples ( $O(N)$  per fragment) that each cost either  $O(N)$  or  $O(1)$  each.

In this study, our goal is more specific than total energies: by also leveraging correlated sampling, we wish to predict energy changes due to local chemical changes at an even lower cost. To do so, we make use of Filippi and Umrigar’s space warp correlated sampling technique.<sup>24</sup> In this approach, whose mathematical details are in the Appendix, electron positions are sampled at one molecular geometry and “warped” to the other geometry in a way that leaves unchanged the positions of electrons near atoms that did not move. We would therefore anticipate that the energy difference  $\Delta E^{(l)}$  and energy difference uncertainty  $\sigma^{(l)}$  for fragment  $l$  will be relatively small if the  $l$ th fragment’s atoms do not move during the local change, and relatively large if its atoms do move. If we were to sort the fragments in decreasing order of  $\sigma^{(l)}$  and could bound these uncertainties by  $a/l$  when using some fixed and  $O(1)$  fragment sample size  $K_f$ , where  $a$  is some constant, then estimating the overall energy difference for the local change would no longer require us to take  $O(N)$  samples per fragment. Indeed,  $O(1)$  samples per fragment would now suffice. For example, with our bound,  $K_f$  samples per fragment would give an overall

uncertainty of

$$\sigma = \sqrt{\sum_l (\sigma^{(l)})^2} \leq a \sqrt{\sum_l \frac{1}{l^2}} \leq \frac{\pi a}{\sqrt{6}} \quad (4)$$

where we have used  $\sum_{l=1}^{\infty} l^{-2} = \pi^2/6$ . With this bound, we would thus need  $O(1)$  samples per fragment, or  $O(N)$  samples overall, bringing our overall cost down to  $O(N^2)$  for nonlocal orbitals or  $O(N)$  for local orbitals.

If the falloff in energy difference uncertainty for fragments far from the change were even faster, such as dropping off as  $\sigma_l \leq a/l^2$  when using  $K_f$  samples per fragment, then we could afford to actually decrease the number of samples used for fragments with smaller uncertainties and lower the overall cost still further. If, instead of taking  $K_f$  samples on each fragment, we instead took  $K_f/l^2$  samples, then our overall uncertainty would again be bounded by the  $O(1)$  value  $\pi a/\sqrt{6}$ .

$$\sigma = \sqrt{\sum_l (l\sigma^{(l)})^2} \leq \sqrt{\sum_l \frac{a^2 l^2}{l^4}} \leq \frac{\pi a}{\sqrt{6}} \quad (5)$$

As we will see in an example below, falloff at least as rapid as  $a/l^2$  can be observed in practice, implying that, at least in some applications, decaying the fragment sample size as  $1/l^2$  with fragment number  $l$  should still achieve an  $O(1)$  uncertainty in the overall energy difference. In such a case, the total number of samples is now just  $O(1)$ , again thanks to the fact that  $\sum_{l=1}^{\infty} l^{-2}$  is finite, meaning that the overall cost of drawing our samples should be  $O(1)$  when the orbitals and Jastrow are local and  $O(N)$  then they are nonlocal. This result, if it can be achieved while retaining accuracy, would mark a dramatic reduction in scaling compared to the  $O(N^4)$  cost of the straightforward approach to drawing samples in VMC. Towards that end, we will now turn to the details of how to efficiently evaluate accurate approximations to each fragment’s local energy contribution under this aggressive sampling approach.

## B. Spotlight Sampling

When approximating a given fragment’s contribution to an overall energy difference, there are two key issues that must be addressed. First, in the space near frozen electrons, the marginal probability distribution for electrons that are moving will be erroneously distorted due to the over-localized Pauli exclusion holes of the frozen electrons. Fortunately, Pauli exclusion effects decay exponentially with distance, which suggests that the distribution of a moving electron should be only minimally distorted so long as it is not too close to any of the frozen electrons. Second, the standard VMC local energy’s sum of electron-electron repulsions will be incorrect due to the frozen electrons. Unlike the Pauli exclusion effect, this problem decays much more slowly with distance. To see

why, imagine modeling a long alkane chain, each of whose  $\text{CH}_2$  groups is essentially nonpolar. In our approximate approach,  $\text{CH}_2$  groups with electrons frozen at random positions would instead have large dipoles, the erroneous interactions of which with our active fragment would decay only polynomially with distance. In order to produce accurate estimates of the energy differences we seek, both of these issues will need to be tackled.

Let’s start with the Pauli exclusion problem, which we address by adopting a “spotlight” sampling approach. If we imagine shining a spotlight at the fragment whose contribution we want to evaluate, that fragment would be brightly lit, while fragments adjacent to it would be partially lit, fragments adjacent to those would be dimly lit, and the remaining fragments would be in the dark. Grouping fragments into these four regions — which we will label as A, B, C, and D respectively, as shown in Figure 1 — will form the foundation of our approach. The basic idea is to hold frozen all region D electrons while performing standard Metropolis sampling on the electrons in regions A, B, and C. So long as regions B and C are thick enough, the electron positions sampled in the fragment that defines region A should be nearly free of the Pauli exclusion errors we are trying to mitigate. We will therefore refer to regions B and C as our buffer regions, and for now focus on the fact that they protect the electrons in fragment A from Pauli exclusion errors (we will elaborate on the division between B and C when discussing long-range electrostatics below).


Although we now expect that an electron in region A will be mostly free of Pauli exclusion errors, we will need a precise definition of what it means for an electron to be “in” region A in order to construct that fragment’s energy difference contribution. Towards this purpose, we define nuclear ownership factors that smoothly partition space between the different nuclei. Specifically, the factor  $C_I(\vec{r})$  gives the ownership fraction that nucleus  $I$  has over the point  $\vec{r}$  in three-dimensional space.

$$C_I(\vec{r}) = \frac{e^{-d|\vec{r}-\vec{R}_I|^2}}{\sum_{J=1}^{N_{\text{nuc}}} e^{-d|\vec{r}-\vec{R}_J|^2}} \quad (6)$$

Here  $d = 0.3$  Bohr,  $\vec{R}_I$  is the position of the  $I$ th nucleus, and we note that, at any given position  $\vec{r}$ , the ownership fractions across all  $N_{\text{nuc}}$  nuclei sum to one. Thus, if we want to know to what degree an electron moved to a new position is owned by region A, we can evaluate region A’s ownership fraction  $w_A$  of that position as the sum of its nuclei’s ownership fractions.

$$w_A(\vec{r}) = \sum_{I \in A} C_I(\vec{r}) \quad (7)$$

Using this definition of ownership, we can construct an approximate Hamiltonian operator for the fragment that



**Figure 1.** An illustration of spotlight sampling showing three examples of how centering the spotlight on one fragment (region A) creates partially lit (region B), dimly lit (region C), and unlit (region D) portions of the overall system. The system’s total energy is the sum of fragment energies, each of which is evaluated in its own Markov chain in which the spotlight centers on that fragment and in which region D electrons are frozen. See text for details.

defines region A.

$$\hat{H}_A = \sum_k \frac{w_A(\vec{r}_k)}{2} \left( -\xi_k \nabla_k^2 + \sum_{p \neq k} \sum_X^{ABCD} w_X(\vec{r}_p) V_{AX}(\vec{r}_k, \vec{r}_p) \right), \quad (8)$$

Here  $k$  sums over both the nuclei and the unfrozen electrons,  $\xi_k$  is one for electrons and zero for nuclei,  $p$  sums over electrons, nuclei, and multipoles (see next section),  $X$  sums over A, B, C, and D, and, finally,  $V_{AX}(\vec{r}_k, \vec{r}_p)$  is the electrostatic interaction between the particles and/or multipoles at positions  $\vec{r}_k$  and  $\vec{r}_p$ . It is important to note that, in the limit that region B is extended to encompass everything other than fragment A, pointing the spotlight at each fragment in turn and summing up the corresponding versions of Eq. (8) would yield the standard electronic Hamiltonian  $\hat{H}$  (note that multipoles will only be present in C and D and so can be ignored in this limit). In other words, in the large buffer limit, the sum of the fragment energy differences would have the same expectation value as the energy difference one would get from standard VMC. The key questions we arrive at are, therefore, (a) how large do we need to make the buffer regions in practice to address the Pauli exclusion issue and (b) how can we address the issue of repulsions against

frozen electrons? We will need data to address the buffer question, so let us focus for now on the frozen repulsions.

### C. Long Range Electrostatics

Naively defining  $V_{AX}(\vec{r}_k, \vec{r}_p)$  as the standard two-particle Coulomb interaction will create serious errors in Eq. (8) due to the terms in which  $p$  is a frozen electron in region D whose position is stuck at its randomly drawn initial position. Similar errors will also occur for electrons in region C, as their distributions will be artificially skewed by the too-localized Pauli exclusion effects of the nearby frozen electrons. To address these errors, we remove the region C and region D electrons and nuclei from the sum over  $p$  in Eq. (8) and replace them with multipole expansions, placing one expansion at the nuclear center of charge of each fragment in regions C and D. We then define  $V_{AA}$  and  $V_{AB}$  as the standard two-particle Coulomb interaction, while  $V_{AC}$  and  $V_{AD}$  are defined as the interaction energy of particle  $k$  with the multipole expansion centered at  $\vec{r}_p$ . This way, we retain standard VMC interactions between particles whose distributions we trust to be sufficiently buffered by the presence of region C, while resorting to multipole approximations for the remaining interactions. We note that, in the large buffer limit, this approximation also gets better and better, both because fewer interactions are replaced by multipoles and because the multipole approximations themselves improve with distance.

To construct our multipole expansions, we perform one round of spotlight sampling before we take the samples that will be used for the energy difference evaluation. During this round, when it is fragment Y's turn in the spotlight, we construct a point-charge distribution  $\phi_Y$  for that fragment by aggregating its nuclei and the ownership-weighted sampled positions of the unfrozen electrons.

$$\phi_Y(\vec{r}) = \sum_{I \in Y} Z_I \delta(\vec{r} - \vec{R}_I) - \sum_a \sum_k^{L_Y N_u} \frac{w_Y(\vec{r}_k^{(a)})}{L_Y} \delta(\vec{r} - \vec{r}_k^{(a)}) \quad (9)$$

Here  $Z_I$  is the charge on nucleus  $I$ ,  $L_Y$  is the number of samples used for fragment  $Y$ , and  $\vec{r}_k^{(a)}$  is the position of the  $k$ th unfrozen electron in the  $a$ th sample. We then evaluate the finite multipole expansion of  $\phi_X$  (we truncate after the octopole) for use in Eq. (8).

### D. Overall Cost Scaling

With both the Pauli exclusion errors and electron-electron repulsion errors that arise from freezing electrons now addressed, let us consider the cost of evaluating our energy difference. At each individual sample of the unfrozen electron positions, the local energy  $\frac{\hat{H}_A \Psi}{\Psi}$  contains  $O(1)$  kinetic energy terms. Each costs  $O(1)$  time to evaluate if both the orbitals are local and the locality of the

Jastrow factor is exploited. Otherwise each costs  $O(N)$  time (see Appendix). Summing the  $V_{AX}$  interactions requires  $O(1)$  time for the standard two-particle Coulomb interactions (note that, for a given spotlight position,  $w_A$  and  $w_B$  are only nonzero for  $O(1)$  nuclei) and  $O(N)$  time for the charge-multipole interactions, although this could be reduced via the fast multipole method in sufficiently large systems.<sup>47,48</sup> Thus, if as discussed above the uncertainty contributions from the ordered fragments decay rapidly enough that the total number of samples required across all spotlight positions is  $O(1)$ , then the overall cost of evaluating the energy difference for our local chemical change will add up to  $O(N)$  when nonlocal orbitals are used, and it could become sublinear in large systems if local orbitals, Jastrow locality, and the fast multipole method are employed. However, the cost of first evaluating the multipoles would still be  $O(N)$ , as we can expect to need  $O(1)$  samples for each of the  $O(N)$  multipoles.

### E. Numerical Stability

In this spotlight approach, as in regular VMC, care is required to mitigate the buildup of finite precision arithmetic errors when repeatedly executing low-rank matrix inverse updates via Sherman-Morrison. In fact, there are three contributing factors that make this issue particularly important in the present approach. First, the inverse matrix  $(YZ)^{-1}$  on which we would like to perform rank one updates after each one-electron move already contains a matrix inverse inside of it:  $Z$  is built from columns of  $(M_{\text{init}})^{-1}$ . Although we do not update  $Z$  as the unfrozen electrons move around, we do update  $Y$ , and in numerical testing we find that the condition number of  $YZ$  can become more extreme than that of the standard  $M$  that would be used in regular VMC, leading to faster round-off buildup. Second, in this study at least, we are employing explicitly cusped core orbitals<sup>49</sup> whose values can range by multiple orders of magnitude depending on how close to the nucleus an electron is. All else being equal, a larger range of element magnitudes tends to worsen matrix condition numbers (consider, for example, what happens in the case of a diagonal matrix). Third, the space warp transformation we employ in correlated sampling can place the warped electrons near a node even when, pre-warp, the guiding Markov chain is not close to a node. The larger the number of electrons, the more nodes there are to contend with, and we have observed that poor conditioning of the post-warp  $YZ$  matrix is more common in larger systems.

We address the first two issues by limiting the number of steps we allow the Markov chain to take before explicitly re-inverting  $YZ$  and the third issue by imposing an approximate modification of our probability distribution. Note that, we could, if we wanted, explicitly re-invert  $YZ$  at every step while still maintaining the same cost scaling, as  $YZ$ 's matrix dimension is  $O(1)$ . However, especially when employing buffer regions, the number of

unfrozen electrons is large enough to make a Sherman-Morrison approach worthwhile. The question is where to strike the balance between practical efficiency and numerical stability, and, at least in the systems tested here, we find that re-inverting  $YZ$  after each unfrozen electron has attempted 100 moves to be an effective compromise. To address our third issue, we note that the conditioning of the pre-warp and post-warp matrices is likely to differ the most when the numerator of the correlated sampling weight factor  $W_i$  (see Appendix) is significantly different than one. This numerator is connected to the relative condition number, because it contains the ratio of the post-warp to pre-warp determinants inside it. To avoid samples that are likely to be ill conditioned, we simply set the probability distribution equal to zero when this numerator is less than  $10^{-5}$  or greater than 10. This is of course an approximation, but we note that it leads to increasing the Metropolis rejection rate by just 0.01%, and, as we will see in our results, it does not prevent the method from predicting correct energy differences. It does, however, make for a significant improvement in the round-off error of our rank one updates, as desired.


## F. Initialization

We begin by placing electrons according to a Lewis structure so that the molecule is roughly spatially and spin neutral. The position of each electron (except for those in the carbon and oxygen cores) is then randomized within a 1 Bohr box centered on its Lewis placement. After separating the nuclei into fragments — in our tests, each carbon and any hydrogens or O-H groups bonded to it form a fragment, as do individual  $H_2$  molecules — we then group the initial electron positions by fragment so that each fragment gets the  $n_e$  nearest electrons, where  $n_e$  is the fragment’s nuclear charge. We burn in the electron positions using a “round-robin” approach in which we loop over each fragment and, for each one, unfreeze and sample its electrons and those of its two nearest neighbors in each direction. At the beginning of each fragment’s burn in, and once after they have all been burned in, we re-invert  $M$ , for a total of  $n_{\text{fragments}} + 1$  large matrix inversions. In the systems we test, these inversions together cost much less than the initial Hartree-Fock calculation. After the round-robin burn in, we reassign electrons to fragments, again taking the closest so many so that the fragment is neutral. At this point, we have our initial positions from which we separately launch the Markov chains that evaluate each fragment’s multipole and then continue on to evaluate each fragment’s energy using the approximate Hamiltonian in Eq. (8).

## III. RESULTS

### A. Computational Details

We performed all quantum chemistry calculations for this project using the PYSCF package.<sup>50–52</sup> We optimized all geometries at the MP2/6-31G level of theory unless otherwise specified. For our Hartree-Fock references, we used restricted Slater determinants with localized Foster-Boys orbitals<sup>53</sup> using the 6-31G(d) basis set. To ensure electron-nuclear cusps<sup>39</sup> were satisfied for the VMC sampling, we employed a recently developed cusping algorithm<sup>49</sup> in which all AOs are explicitly cusped at each nucleus. We used a simple two-body Jastrow factor<sup>40,44,54</sup> to satisfy the electron-electron cusps (A parameter = 0.03). Regular and spotlight sampling VMC calculations were performed using our own pilot code.



**Figure 2.** Predictions for the energy required to stretch an O-H bond by 0.2 Bohr from standard VMC and from spotlight sampling with different buffer region choices.

### B. Impact of the Buffer Region

We begin by investigating the effect that different buffer region choices have on predicting the energy difference associated with an O-H bond stretch, as shown in Figure 2. We adopt a 4-letter naming convention for the different buffer choices we test, in which the first letter is always “A” to specify that it refers to the active fragment. The second letter refers to that fragment’s nearest neighbors, the third to next-nearest neighbors, and the fourth to all remaining fragments. Thus, “ADDD” corresponds to using no buffer region at all, with only one fragment’s electrons unfrozen in each Markov chain. “ABDD” also unfreezes nearest-neighbor fragment’s electrons and uses that fragment’s particles for explicit Coulomb interac-

tions, while using multipole for next-nearest and farther fragments. “ACCC” unfreezes all the electrons, but employs the multipole approximation for all non-active fragments. Finally, “ABCD” unfreezes electrons in nearest and next-nearest neighbor fragments, uses explicit Coulomb interactions for the nearest neighbor fragment, and uses multipole for next-nearest and farther fragments.


As seen in the figure, only the ABCD approach reproduces the predictions of standard VMC within statistical error. Freezing all other fragments’ electrons (ABDD) produces both a significant error and a significantly larger uncertainty, which we suspect is driven by the effect of the randomly frozen Pauli exclusion holes in the nearest-neighbor fragments. The uncertainty is smaller for ACCC, in which all electrons move, but its prediction is incorrect due to errors associated with using a multipole approximation at short, nearest-neighbor range. Although such errors should be reduced in ABDD, in which the closest multipole use is next-nearest neighbor, the Coulomb interactions with the nearest neighbor are still inaccurate due to the bias in the nearest neighbor fragment’s electron positions coming from the frozen Pauli exclusion holes in the next-nearest neighbor fragment. It is only when we adopt the ABCD strategy depicted in Figure 1 that we see the spotlight sampling prediction match the standard VMC prediction. In this approach, we mitigate multipole errors by keeping the closest multipole use at next-nearest neighbor range and explicit Coulomb errors by using next-nearest neighbors as buffers in between the nearest neighbors and the frozen electrons. We adopt this approach for all remaining cal-

culations.


### C. Cost Scaling

In investigating whether observed computational cost scaling matches the expectations from our analysis above, let us begin with a simple test in which we take a fixed number of samples for each fragment. As our implementation does not exploit orbital locality, we expect each sample to carry an  $O(N)$  cost, which, when combined with the fact that there are  $O(N)$  fragments, leads us to expect  $O(N^2)$  cost growth overall. We test this theory in a chain of hydrogen dimers (bond length = 1.393 Bohr) arranged in a line, in which each  $H_2$  molecule’s center is placed 4.03 Bohr away from the previous center, and all bond axes are arranged at an angle of 7.12 degrees relative to the line passing through the centers. Using each  $H_2$  as its own fragment and taking 110 samples per fragment, we see in Figure 3 that the spotlight sampling cost does indeed grow as  $O(N^2)$ . While our current code cannot test the reduction to  $O(N)$  that would result from local orbitals, that reduction is well established,<sup>21,25</sup> and so we instead turn to the question of how effectively costs can be reduced by taking fewer samples in fragments farther away from the local chemical change in question.


To address this question, we first perform tests of how much uncertainty each fragment contributes to a bond stretching energy when we take a fixed number of samples per fragment. We perform our first test in 1-octanol, where we perform 150 runs on each fragment, in which each run uses 720,000 samples. For each run in each fragment, we use a blocking analysis to evaluate the un-



**Figure 3.** Overall spotlight sampling cost for  $(H_2)_n$  when employing a fixed number of samples per fragment. Each calculation used 1 core on an Intel Xeon Gold 6230 processor.



**Figure 4.** Each fragment’s average uncertainty contribution to the O-H bond stretching energy in 1-octanol. Each fragment contains one C atom and any H or O-H bonded to it.



**Figure 5.** Each fragment’s average uncertainty contribution to the O-H bond stretching energy in methanol- $(\text{H}_2)_{16}$ . Each molecule is one fragment.


certainty of that fragment’s contribution to the bond stretching energy in that run. The averages of these uncertainties over all 150 runs are plotted in Figure 4, where we see that the uncertainties decrease as one moves away from the  $\text{CH}_2\text{OH}$  fragment (fragment 1) that contains the bond stretch. In fact, these uncertainties decay fast enough that they are bounded by an  $a/l^2$  function, with  $a = 0.01$  Hartree and  $l$  the fragment number. As discussed in Section II A, this is a highly desirable bound, as it suggests that decreasing the sample size for farther fragments should be highly effective.

We repeat this per-fragment uncertainty quantification in a system in which a methanol’s O-H bond is stretched when the methanol is placed next to a line of 16  $\text{H}_2$  molecules, this time taking 480,000 samples per fragment. In Figure 5, we see that the uncertainty contributions again become small for fragments far away from the bond stretch, and indeed we find that they can also be bounded by  $a/l^2$ , this time with  $a = 0.004$  Hartree. Note that, in both cases, fragment 3 has a larger uncertainty contribution than one might at first expect, but this result makes sense once we note that the fragment 3 Markov chain contains the largest number of unfrozen electrons and that these electrons include those of the stretched bond.

Seeing that uncertainties decay rapidly as one gets farther from the local chemical change, we now exploit this behavior to demonstrate that, even without making use of local orbital accelerations, spotlight sampling can achieve lower costs than traditional VMC. To do so, we again calculate the O-H bond stretching energy for methanol- $(\text{H}_2)_n$ , but this time we employ blocks of 960,000 samples for the first three fragments and much

smaller blocks of 4,000 samples for the remaining fragments, versus whole-system blocks of 960,000 samples for standard VMC. In each system, we compute enough sampling blocks so that the overall uncertainty in the stretching energy is below 0.0005 Hartree, always using at least 15 blocks to ensure a reliable uncertainty estimate.

In Figure 6, we see that this test of spotlight sampling achieves near-linear scaling with system size, with spotlight’s cost lower than standard correlated sampling once there are about 100 electrons. The linear scaling is expected, as the vast majority of the samples are in the first three fragments, making the total number of samples nearly independent of system size. Indeed, the linear cost growth comes instead from our code’s  $O(N)$  cost per sample. The small quadratic term in the cost growth that begins to be visible in the spotlight timing data in the largest system is due to the fact that we used fixed-size 4,000 sample blocks even for the farthest fragments, and could be eliminated by tapering this sample size further. It is important to re-emphasize that this crossover with standard sampling has been achieved *without* exploiting local orbitals. Were one to do so, the crossover would occur even earlier, and the cost growth of spotlight sampling could drop to sub-linear, assuming local Jastrows and fast multipole methods were also employed. Although each spotlight prediction agreed with standard sampling within error bars, this system is somewhat artificial, and so we now turn to a series of alcohols as a




**Figure 6.** Wall time comparison of spotlight versus standard correlated sampling for the O-H bond stretch in  $\text{MeOH}-(\text{H}_2)_n$  when employing reduced spotlight sampling efforts for farther fragments. In all cases, 15 sampling blocks were sufficient for spotlight sampling to reach the 0.0005 Hartree uncertainty threshold, whereas the number needed by standard sampling varied between 15 and 27. Each calculation used five Intel Xeon Gold 6230 processors.

more realistic test of the method’s accuracy.

#### D. Accuracy

To test whether our approximations to the Hamiltonian and probability distribution produce meaningful errors in predicted energy differences, we apply the spotlight sampling approach to predict the energetic cost of a 0.2 Bohr O-H bond stretch in a series of alcohols: 1-butanol through 1-octanol. As seen in Figure 7, spotlight sampling reproduces the predictions of standard correlated sampling in all cases. We note that this accuracy comes despite the fact that, in octanol, the Markov chain that estimates the contribution from the fragment containing the stretched bond froze more than half of the electrons.



**Figure 7.** Predictions for the energy needed to stretch an O-H bond in the alcohols 1-butanol through 1-octanol.

#### IV. CONCLUSION

This study has explored the efficacy of a spotlight sampling approach to predicting the energy differences associated with local chemical changes in VMC. In this approach, standard correlated sampling is approximated by using a series of separate Markov chains, one for each fragment in the molecule, in which only the electrons near that fragment move. Our theoretical analysis suggests that the cost of the approach can be made linear with system size when employing nonlocal orbitals, and that this cost should become sub-linear when local orbitals, local Jastrows, and fast multipole methods are used. Timing calculations show that the overall cost

grows quadratically when we take a fixed number of samples per spotlighted fragment, and that this cost growth becomes essentially linear, as expected, once we exploit the fact that per-fragment uncertainty decays rapidly for fragments farther from the local change. To maintain accuracy despite the biases inherent to freezing most electrons’ positions, we have developed an approximate fragment Hamiltonian that employs multipoles for long range interactions. Tests on a series of alcohols show that this approach can reproduce the predictions of standard sampling within statistical uncertainty.

#### V. APPENDIX

##### A. Correlated Sampling with the Space-Warp Coordinate Transformation

We employ the correlated sampling<sup>55</sup> technique known as the space-warp coordinate transformation<sup>22</sup> to reduce the statistical uncertainty of simple bond stretches (e.g., O-H stretch in 1-n-ol,  $n$  = butan-, pentan-, hexan-,...). The relative energy difference for such a stretch,  $E_s - E$ , is rewritten to be sampled from the unstretched  $\Psi$  reference:

$$E_s - E = \langle \Psi_s | \hat{H}_s | \Psi_s \rangle - \langle \Psi | \hat{H} | \Psi \rangle \approx \frac{1}{N} \sum_{R_i \in |\Psi|^2} \left( \frac{\hat{H}_s \Psi_s(R_i)}{\Psi_s(R_i)} W_i - \frac{\hat{H} \Psi(R_i)}{\Psi(R_i)} \right) \quad (10)$$

where

$$W_i = \frac{|\Psi_s(R_i)|^2 / |\Psi(R_i)|^2 J(R_i)}{\frac{1}{N} \sum_j |\Psi_s(R_j)|^2 / |\Psi(R_j)|^2 J(R_j)} \quad (11)$$

$J(R_i)$  is the Jacobian of the space-warp coordinate transformation,<sup>22,24</sup> and

$$r_i^s = r_i + \sum_{A=1}^{n_{nuc}} (r_A^s - r_A) \omega_A(r_i) \quad (12)$$

$$\omega_A(r_i) = \frac{|r_i - r_A|^{-4}}{\sum_{B=1}^{n_{nuc}} |r_i - r_B|^{-4}} \quad \sum_{A=1}^{n_{nuc}} \omega_A(r_i) = 1 \quad (13)$$

When sampling electron positions from the unstretched reference geometry, Eq. 12 shifts the electron positions for the stretched geometry to account for the displacement of the nucleus  $R_A$ . The weights  $W_i$  are updated using the same techniques as the determinant ratios discussed in the main text, and therefore this correlated sampling technique shares the same per-sample cost scaling as the spotlight algorithm would if used on only the unstretched geometry.

## B. Kinetic Energy Evaluation

As in many QMC contexts, evaluating the kinetic energy for the  $i$ th active electron can be simplified via the following identity.

$$\frac{\nabla_i^2 \Psi}{\Psi} = \nabla_i^2 \ln \Psi + (\nabla_i \ln \Psi) \cdot (\nabla_i \ln \Psi) \quad (14)$$

Start with the  $\nabla_i^2$  part and focus for now on the determinant part of the wave function. Summing the contributions over the active electrons, we have

$$\sum_i \frac{\nabla_i^2 D(R_{\text{new}})}{D(R_{\text{new}})} = \sum_i \frac{D(R_{\text{init}})}{D(R_{\text{new}})} \frac{\nabla_i^2 D(R_{\text{new}})}{D(R_{\text{init}})} \quad (15)$$

$$= \sum_i \frac{\det(Y^{(i)}Z)}{\det(YZ)}, \quad (16)$$

where we have used Eq. (3), and where  $Y^{(i)}$  is the matrix formed by replacing the orbital values in the row of  $Y$  belonging to the  $i$ th electron with the corresponding orbital Laplacians. As the matrices in the numerator and denominator of Eq. (16) differ by just one row, the matrix determinant lemma can be used to evaluate the ratio in  $O(N)$  time (or  $O(1)$  time with local orbitals). In particular, the matrix  $(YZ)^{-1}$  that the lemma uses is already available, as we are already keeping it updated via Sherman-Morrison as discussed in the main text. Similar arguments show that the  $\nabla_i \ln D$  part has the same cost scaling, and since there are only  $O(1)$  active electrons, the overall cost scaling for the determinant part of the kinetic energy for one sample in one fragment is  $O(N)$  with nonlocal orbitals and  $O(1)$  with local orbitals.

The kinetic energy contribution of the Jastrow factor,  $e^{J(R)}$ , for electron  $i$  requires calculating the terms

$$\nabla_i J(R_{ij}) = \frac{A}{r_{ij}} \begin{pmatrix} x_i - x_j \\ y_i - y_j \\ z_i - z_j \end{pmatrix} \left( \frac{1 - e^{-r_{ij}/F}}{r_{ij}^2} - \frac{e^{-r_{ij}/F}}{r_{ij}F} \right) \quad (17)$$

and

$$\begin{aligned} \nabla_i^2 J(R_{ij}) = & \sum_{\alpha \in x,y,z} \frac{A}{r_{ij}} \left( \frac{1 - e^{-r_{ij}/F}}{r_{ij}^2} - \frac{e^{-r_{ij}/F}}{r_{ij}F} \right) \\ & + (\alpha_i - \alpha_j)^2 \left[ e^{-r_{ij}/F} \left( \frac{3}{r_{ij}^3 F} + \frac{1}{r_{ij}^2 F^2} \right) - \frac{3(1 - e^{-r_{ij}/F})}{r_{ij}^4} \right] \end{aligned} \quad (18)$$

where  $j$  is another electron. To ensure linear scaling for the kinetic energy evaluation for each fragment, we only evaluate Eqs. 17 and 18 for terms involving active electrons. We find this approximation to be accurate thanks to our buffer region and the fact that the Jastrow factor derivatives decay quickly with distance.

## VI. ACKNOWLEDGEMENTS

The authors thank Scott Garner, Leon Otis, Trine Quady, and Connie Robinson for helpful discussions. This work was supported by the Office of Science, Office of Basic Energy Sciences, the U.S. Department of Energy, Contract no. DE-AC02-05CH11231, through the Gas Phase Chemical Physics program. Calculations were performed using the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley and the Lawrencium computational cluster resource provided by the IT Division at the Lawrence Berkeley National Laboratory.

## REFERENCES

- <sup>1</sup>J. Q. Benjamin G. Levine, Chaehyuk Ko and T. J. Martínez, “Conical intersections and double excitations in time-dependent density functional theory,” *Molecular Physics* **104**, 1039–1051 (2006).
- <sup>2</sup>C. Sherrill, “Chapter 4 Bond Breaking in Quantum Chemistry,” (Elsevier, 2005) pp. 45–56.
- <sup>3</sup>D. Hait, N. M. Tubman, D. S. Levine, K. B. Whaley, and M. Head-Gordon, “What Levels of Coupled Cluster Theory Are Appropriate for Transition Metal Systems? A Study Using Near-Exact Quantum Chemical Values for 3d Transition Metal Binary Compounds,” *Journal of Chemical Theory and Computation* **15**, 5370–5385 (2019).
- <sup>4</sup>Y. A. Aoto, A. P. de Lima Batista, A. Köhn, and A. G. S. de Oliveira-Filho, “How To Arrive at Accurate Benchmark Values for Transition Metal Compounds: Computation or Experiment?” *Journal of Chemical Theory and Computation* **13**, 5291–5316 (2017).
- <sup>5</sup>T. Van Voorhis and M. Head-Gordon, “Benchmark variational coupled cluster doubles results,” *The Journal of Chemical Physics* **113**, 8873–8879 (2000).
- <sup>6</sup>J. N. Harvey, “On the accuracy of density functional theory in transition metal chemistry,” *Annu. Rep. Prog. Chem., Sect. C: Phys. Chem.* **102**, 203–226 (2006).
- <sup>7</sup>K. D. Vogiatzis, D. Ma, J. Olsen, L. Gagliardi, and W. A. de Jong, “Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations,” *The Journal of Chemical Physics* **147**, 184111 (2017).
- <sup>8</sup>J. Olsen, “The CASSCF method: A perspective and commentary,” *International Journal of Quantum Chemistry* **111**, 3267–3272 (2011).
- <sup>9</sup>J. Shee, B. Rudsteyn, E. J. Arthur, S. Zhang, D. R. Reichman, and R. A. Friesner, “On Achieving High Accuracy in Quantum Chemical Calculations of 3d Transition Metal-Containing Systems: A Comparison of Auxiliary-Field Quantum Monte Carlo with Coupled Cluster, Density Functional Theory, and Experiment for Diatomic Molecules,” *Journal of Chemical Theory and Computation* **15**, 2346–2358 (2019).
- <sup>10</sup>R. J. Needs, M. D. Towler, N. D. Drummond, P. López Ríos, and J. R. Trail, “Variational and diffusion quantum Monte Carlo calculations with the CASINO code,” *The Journal of Chemical Physics* **152**, 154106 (2020).
- <sup>11</sup>P. R. C. Kent, A. Annaberdiyev, A. Benali, M. C. Bennett, E. J. Landinez Borda, P. Doak, H. Hao, K. D. Jordan, J. T. Krogel, I. Kylänpää, J. Lee, Y. Luo, F. D. Malone, C. A. Melton, L. Mitas, M. A. Morales, E. Neuscamman, F. A. Reboredo, B. Rubenstein, K. Saritas, S. Upadhyay, G. Wang, S. Zhang, and L. Zhao, “QMCPACK: Advances in the development, efficiency, and application of auxiliary field and real-space variational and diffusion

- quantum Monte Carlo,” *The Journal of Chemical Physics* **152**, 174105 (2020).
- <sup>12</sup>B. K. Clark, M. A. Morales, J. McMinis, J. Kim, and G. E. Scuseria, “Computing the energy of a water molecule using multideterminants: A simple, efficient algorithm,” *The Journal of Chemical Physics* **135**, 244105 (2011).
  - <sup>13</sup>R. Assaraf, S. Moroni, and C. Filippi, “Optimizing the Energy with Quantum Monte Carlo: A Lower Numerical Scaling for Jastrow–Slater Expansions,” *Journal of Chemical Theory and Computation* **13**, 5273–5281 (2017).
  - <sup>14</sup>F. R. Petruzielo, J. Toulouse, and C. J. Umrigar, “Approaching chemical accuracy with quantum Monte Carlo,” *The Journal of Chemical Physics* **136**, 124116 (2012).
  - <sup>15</sup>R. A. Emmanuel Giner and J. Toulouse, “Quantum Monte Carlo with reoptimised perturbatively selected configuration-interaction wave functions,” *Molecular Physics* **114**, 910–920 (2016).
  - <sup>16</sup>M. Dash, J. Feldt, S. Moroni, A. Scemama, and C. Filippi, “Excited States with Selected Configuration Interaction-Quantum Monte Carlo: Chemically Accurate Excitation Energies and Geometries,” *Journal of Chemical Theory and Computation* **15**, 4896–4906 (2019).
  - <sup>17</sup>M. Dash, S. Moroni, C. Filippi, and A. Scemama, “Tailoring CIPSI Expansions for QMC Calculations of Electronic Excitations: The Case Study of Thiophene,” *Journal of Chemical Theory and Computation* **17**, 3426–3434 (2021).
  - <sup>18</sup>S. D. Pineda Flores and E. Neuscamman, “Excited State Specific Multi-Slater Jastrow Wave Functions,” *The Journal of Physical Chemistry A* **123**, 1487–1497 (2019).
  - <sup>19</sup>L. Otis and E. Neuscamman, “A promising intersection of excited-state-specific methods from quantum chemistry and quantum Monte Carlo,” *WIREs Computational Molecular Science* **13**, e1659 (2023).
  - <sup>20</sup>L. Otis, I. M. Craig, and E. Neuscamman, “A hybrid approach to excited-state-specific variational Monte Carlo and doubly excited states,” *The Journal of Chemical Physics* **153**, 234105 (2020).
  - <sup>21</sup>A. Bienvenu, J. Feldt, J. Toulouse, and R. Assaraf, “Systematic lowering of the scaling of Monte Carlo calculations by partitioning and subsampling,” *Phys. Rev. E* **106**, 025301 (2022).
  - <sup>22</sup>C. J. Umrigar, “Two aspects of quantum monte carlo: Determination of accurate wavefunctions and determination of potential energy surfaces of molecules,” *International Journal of Quantum Chemistry* **36**, 217–230 (1989).
  - <sup>23</sup>S. Sorella and L. Capriotti, “Algorithmic differentiation and the calculation of forces by quantum Monte Carlo,” *The Journal of Chemical Physics* **133**, 234111 (2010).
  - <sup>24</sup>C. Filippi and C. J. Umrigar, “Correlated sampling in quantum Monte Carlo: A route to forces,” *Phys. Rev. B* **61**, R16291–R16294 (2000).
  - <sup>25</sup>J. Feldt and R. Assaraf, “Stochastic Effective Core Potentials, toward Efficient Quantum Monte Carlo Simulations of Molecules with Large Atomic Numbers,” *Journal of Chemical Theory and Computation* **17**, 1380–1389 (2021).
  - <sup>26</sup>W. Kohn, “Analytic Properties of Bloch Waves and Wannier Functions,” *Phys. Rev.* **115**, 809–821 (1959).
  - <sup>27</sup>P. E. Maslen, C. Ochsenfeld, C. A. White, M. S. Lee, and M. Head-Gordon, “Locality and Sparsity of Ab Initio One-Particle Density Matrices and Localized Orbitals,” *The Journal of Physical Chemistry A* **102**, 2215–2222 (1998).
  - <sup>28</sup>J. C. Burant, G. E. Scuseria, and M. J. Frisch, “A linear scaling method for Hartree–Fock exchange calculations of large molecules,” *The Journal of Chemical Physics* **105**, 8969–8972 (1996).
  - <sup>29</sup>E. Schwegler and M. Challacombe, “Linear scaling computation of the Hartree–Fock exchange matrix,” *The Journal of Chemical Physics* **105**, 2726–2734 (1996).
  - <sup>30</sup>A. J. Williamson, R. Q. Hood, and J. C. Grossman, “Linear-Scaling Quantum Monte Carlo Calculations,” *Phys. Rev. Lett.* **87**, 246406 (2001).
  - <sup>31</sup>S. Manten and A. Lüchow, “Linear scaling for the local energy in quantum Monte Carlo,” *The Journal of Chemical Physics* **119**, 1307–1312 (2003).
  - <sup>32</sup>D. Alfè and M. J. Gillan, “Linear-scaling quantum Monte Carlo technique with non orthogonal localized orbitals,” *Journal of Physics: Condensed Matter* **16**, L305 (2004).
  - <sup>33</sup>B. Austin, A. Aspuru-Guzik, R. Salomón-Ferrer, and W. A. J. Lester, “Linear-Scaling Evaluation of the Local Energy in Quantum Monte Carlo,” in *Advances in Quantum Monte Carlo* (2006) Chap. 5, pp. 55–68, <https://pubs.acs.org/doi/pdf/10.1021/bk-2007-0953.ch005>.
  - <sup>34</sup>T. McDaniel, E. F. D’Azevedo, Y. W. Li, K. Wong, and P. R. C. Kent, “Delayed Slater determinant update algorithms for high efficiency quantum Monte Carlo,” *The Journal of Chemical Physics* **147**, 174107 (2017).
  - <sup>35</sup>G. Knizia and G. K.-L. Chan, “Density Matrix Embedding: A Simple Alternative to Dynamical Mean-Field Theory,” *Phys. Rev. Lett.* **109**, 186404 (2012).
  - <sup>36</sup>Q. Sun and G. K.-L. Chan, “Quantum Embedding Theories,” *Accounts of Chemical Research* **49**, 2705–2712 (2016).
  - <sup>37</sup>M. Welborn, T. Tsuchimochi, and T. Van Voorhis, “Bootstrap embedding: An internally consistent fragment-based method,” *The Journal of Chemical Physics* **145**, 074102 (2016).
  - <sup>38</sup>H.-Z. Ye, N. D. Ricke, H. K. Tran, and T. Van Voorhis, “Bootstrap Embedding for Molecules,” *Journal of Chemical Theory and Computation* **15**, 4497–4506 (2019).
  - <sup>39</sup>T. Kato, “On the Eigenfunctions of Many-Particle Systems in Quantum Mechanics,” *Communications on Pure and Applied Mathematics* **10**, 151–177 (1957).
  - <sup>40</sup>W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Rev. Mod. Phys.* **73**, 33–83 (2001).
  - <sup>41</sup>N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of State Calculations by Fast Computing Machines,” *The Journal of Chemical Physics* **21**, 1087–1092 (1953).
  - <sup>42</sup>W. K. Hastings, “Monte carlo sampling methods using Markov chains and their applications,” *Biometrika* **57**, 97–109 (1970).
  - <sup>43</sup>J. Toulouse, R. Assaraf, and C. J. Umrigar, “Chapter Fifteen - Introduction to the Variational and Diffusion Monte Carlo Methods,” in *Electron Correlation in Molecules – ab initio Beyond Gaussian Quantum Chemistry*, *Advances in Quantum Chemistry*, Vol. 73, edited by P. E. Hoggan and T. Ozdogan (Academic Press, 2016) pp. 285–314.
  - <sup>44</sup>D. Ceperley, “Ground state of the fermion one-component plasma: A Monte Carlo study in two and three dimensions,” *Phys. Rev. B* **18**, 3126–3138 (1978).
  - <sup>45</sup>C. J. Umrigar, “Accelerated Metropolis method,” *Phys. Rev. Lett.* **71**, 408–411 (1993).
  - <sup>46</sup>R. M. Lee, G. J. Conduit, N. Nemec, P. López Ríos, and N. D. Drummond, “Strategies for improving the efficiency of quantum Monte Carlo calculations,” *Phys. Rev. E* **83**, 066706 (2011).
  - <sup>47</sup>V. Rokhlin, “Rapid solution of integral equations of classical potential theory,” *Journal of Computational Physics* **60**, 187–207 (1985).
  - <sup>48</sup>L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *Journal of Computational Physics* **73**, 325–348 (1987).
  - <sup>49</sup>T. K. Quady, S. Bumann, and E. Neuscamman, “Method-independent cusps for atomic orbitals in quantum Monte Carlo,” *The Journal of Chemical Physics* **162**, 104104 (2025).
  - <sup>50</sup>Q. Sun, “Libcint: An Efficient General Integral Library for Gaussian Basis Functions,” *Journal of Computational Chemistry* **36**, 1664–1671 (2015).
  - <sup>51</sup>Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan, “PySCF: The Python-based Simulations of Chemistry Framework,” *WIREs Computational Molecular Science* **8**, e1340 (2018).
  - <sup>52</sup>Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui, J. J. Eriksen,

- Y. Gao, S. Guo, J. Hermann, M. R. Hermes, K. Koh, P. Koval, S. Lehtola, Z. Li, J. Liu, N. Mardirossian, J. D. McClain, M. Motta, B. Mussard, H. Q. Pham, A. Pulkin, W. Purwanto, P. J. Robinson, E. Ronca, E. R. Sayfutyarova, M. Scheurer, H. F. Schurkus, J. E. T. Smith, C. Sun, S.-N. Sun, S. Upadhyay, L. K. Wagner, X. Wang, A. White, J. D. Whitfield, M. J. Williamson, S. Wouters, J. Yang, J. M. Yu, T. Zhu, T. C. Berkelbach, S. Sharma, A. Y. Sokolov, and G. K.-L. Chan, “Recent Developments in the PySCF Program Package,” *The Journal of Chemical Physics* **153**, 024109 (2020).
- <sup>53</sup>J. M. Foster and S. F. Boys, “Canonical Configurational Interaction Procedure,” *Reviews of Modern Physics* **32**, 300–302 (1960).
- <sup>54</sup>D. M. Ceperley and B. J. Alder, “Ground State of the Electron Gas by a Stochastic Method,” *Phys. Rev. Lett.* **45**, 566–569 (1980).
- <sup>55</sup>R. E. Lowther and R. L. Coldwell, “Monte Carlo calculation of the Born-Oppenheimer potential between two helium atoms,” *Phys. Rev. A* **22**, 14–21 (1980).

## VII. SUPPORTING INFORMATION

### A. Molecular Geometries

**Table S1.** Coordinates of the methanol and 16 hydrogen dimer chain system. See Figures 5 and 6 in the paper. For shorter and longer H<sub>2</sub> chains, the same placement pattern below is applied. The “stretched” geometry is obtained by stretching the O-H bond by 0.2 Bohr along the original bond axis. Units are in Bohr.

| Z | x             | y             | z             |
|---|---------------|---------------|---------------|
| C | -0.088 441 23 | 1.309 295 04  | 0.000 000 00  |
| O | -0.087 760 49 | -1.470 368 55 | 0.000 000 00  |
| H | -2.078 824 13 | 1.850 932 39  | 0.000 000 00  |
| H | 0.831 955 56  | 2.081 773 80  | 1.694 536 62  |
| H | 0.831 955 56  | 2.081 773 80  | -1.694 536 62 |
| H | 1.648 924 76  | -2.105 958 89 | 0.000 000 00  |
| H | -0.088 441 23 | 5.000 000 00  | 0.000 000 00  |
| H | -0.088 441 23 | 6.393 713 29  | 0.000 000 00  |
| H | -0.588 441 23 | 9.000 000 00  | 0.000 000 00  |
| H | -0.588 441 23 | 10.393 713 29 | 0.000 000 00  |
| H | -1.088 441 23 | 13.000 000 00 | 0.000 000 00  |
| H | -1.088 441 23 | 14.393 713 29 | 0.000 000 00  |
| H | -1.588 441 23 | 17.000 000 00 | 0.000 000 00  |
| H | -1.588 441 23 | 18.393 713 29 | 0.000 000 00  |
| H | -2.088 441 23 | 21.000 000 00 | 0.000 000 00  |
| H | -2.088 441 23 | 22.393 713 29 | 0.000 000 00  |
| H | -2.588 441 23 | 25.000 000 00 | 0.000 000 00  |
| H | -2.588 441 23 | 26.393 713 29 | 0.000 000 00  |
| H | -3.088 441 23 | 29.000 000 00 | 0.000 000 00  |
| H | -3.088 441 23 | 30.393 713 29 | 0.000 000 00  |
| H | -3.588 441 23 | 33.000 000 00 | 0.000 000 00  |
| H | -3.588 441 23 | 34.393 713 29 | 0.000 000 00  |
| H | -4.088 441 23 | 37.000 000 00 | 0.000 000 00  |
| H | -4.088 441 23 | 38.393 713 29 | 0.000 000 00  |
| H | -4.588 441 23 | 41.000 000 00 | 0.000 000 00  |
| H | -4.588 441 23 | 42.393 713 29 | 0.000 000 00  |
| H | -5.088 441 23 | 45.000 000 00 | 0.000 000 00  |
| H | -5.088 441 23 | 46.393 713 29 | 0.000 000 00  |
| H | -5.588 441 23 | 49.000 000 00 | 0.000 000 00  |
| H | -5.588 441 23 | 50.393 713 29 | 0.000 000 00  |
| H | -6.088 441 23 | 53.000 000 00 | 0.000 000 00  |
| H | -6.088 441 23 | 54.393 713 29 | 0.000 000 00  |
| H | -6.588 441 23 | 57.000 000 00 | 0.000 000 00  |
| H | -6.588 441 23 | 58.393 713 29 | 0.000 000 00  |
| H | -7.088 441 23 | 61.000 000 00 | 0.000 000 00  |
| H | -7.088 441 23 | 62.393 713 29 | 0.000 000 00  |
| H | -7.588 441 23 | 65.000 000 00 | 0.000 000 00  |
| H | -7.588 441 23 | 66.393 713 29 | 0.000 000 00  |