

ARTICLE OPEN

Large scale hybrid Monte Carlo simulations for structure and property prediction

Sergei Prokhorenko¹, Kruz Kalke¹, Yousra Nahas¹ and Laurent Bellaïche¹

The Monte Carlo method is one of the first and most widely used algorithms in modern computational physics. In condensed matter physics, the particularly popular flavor of this technique is the Metropolis Monte Carlo scheme. While being incredibly robust and easy to implement, the Metropolis sampling is not well-suited for situations where energy and force evaluations are computationally demanding. In search for a more efficient technique, we here explore the performance of Hybrid Monte Carlo sampling, an algorithm widely used in quantum electrodynamics, as a structure prediction scheme for systems with long-range interactions. Our results show that the Hybrid Monte Carlo algorithm stands out as an excellent computational scheme that can not only significantly outperform the Metropolis sampling but also complement molecular dynamics in materials science applications, while allowing ultra-large-scale simulations of systems containing millions of particles.

npj Computational Materials (2018)4:80; <https://doi.org/10.1038/s41524-018-0137-0>

INTRODUCTION

Following the pioneering computational experiments¹ of Enrico Fermi, Nick Metropolis, Stanislaw Ulam and John von Neumann, the family of Monte Carlo algorithms has been constantly growing and gained an incredible popularity in computational physics. Of particular note is the publication in 1953 of the paper by Nick Metropolis, Marshall and Arianna Rosenbluth, and Edward and Mici Teller, describing for the first time the algorithm that has come to be known as the *Metropolis* algorithm² (we will further refer to it as MMC scheme). This algorithm was the first example of a thermal “importance sampling” method, and it is to this day easily the most widely used Monte Carlo method. The Swendsen-Wang algorithm³ (later refined by Ulli Wolff⁴) established an important advancement of the Metropolis method via the introduction of the non-local, “cluster”, updates of the system’s state. However, while improving the simulation performance near the phase transition¹ where critical slowing down becomes important, the cluster update relies on the short-range nature of interactions and, to the best of our knowledge, has not been generalized beyond nearest neighbors models. For systems exhibiting long-range interactions, the Metropolis algorithm still remains the method of choice mostly due to robustness and simplicity of implementation despite its poor scaling with the system size. Specifically, for a system of N particles, all pairwise interacting, a single Metropolis step requires performing $O(N)$ floating-point operations yielding an $O(N^2)$ scaling for a single sweep, a prohibitively expensive demand for $N \gtrsim 10^4$. A possible remedy to this scaling problem resides in making use of the gradient of energy (forces), the computation of which can be efficiently parallelized on modern computer architectures. One of the Monte Carlo algorithms that would allow to adopt such solution is the *Hybrid* or *Hamiltonian* Monte Carlo (HMC).^{5,6}

Although the main domain of application of the HMC method is computational quantum electrodynamics, the method has been previously applied and found wanted to treat liquid phase

problems^{7,8} (especially the path integral version of the algorithm for quantum treatment of protons). Nonetheless, the HMC algorithm has not gained traction in the solid state community (rare examples of its application can be found in refs 9–11 and references therein) and to the best of our knowledge has never been envisioned as high-performance alternative to MMC scheme for solid state systems with long-range interactions.

In this study we explore the performance of the HMC algorithm in comparison with Metropolis and thermalized molecular dynamics (MD) methods on the example of effective Hamiltonian models that describe properties of ferroelectric, relaxor and multiferroic materials at finite temperatures. Our results reveal selected model cases for which the HMC scheme significantly outperforms the MMC and MD methods and show that the GPU-oriented implementation of the HMC algorithm for effective Hamiltonian models can allow for performance on par with the best general purpose molecular dynamics programs such as NAMD¹² and LAMMPS.¹³

In order to understand the specific cases that are particularly suited for HMC simulations, it is important to recall some basic properties of this algorithm. The HMC algorithm^{5,6} is a Markov chain Monte Carlo sampling¹ that essentially generates a chain of microscopic states s_i

$$S_{N_t} = \{s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_t \rightarrow s_{t+1} \dots \rightarrow s_{N_t}\}, \quad (1)$$

where N_t denotes the total number of HMC iterations and for which the rate of occurrence of any given microscopic state s converges to the Boltzmann probability $\rho_B(s) = e^{-\beta E(s)}/Z$, at sufficiently large iteration t . Here β denotes the inverse temperature in energy units, $E(s)$ is the total energy of the microscopic state s and Z is the canonical partition function. Each iteration of the scheme consists of a suggestion of a trial, “candidate”, state s'_t , followed by an acceptance decision based on the difference of the total internal energies $\Delta E = E(s'_t) - E(s_{t-1})$ (in contrast to the MMC scheme for which only the difference of

¹Physics Department and Institute for Nanoscience and Engineering, University of Arkansas, Fayetteville, AR 72701, USA

Correspondence: Sergei Prokhorenko (prokhorenko.s@gmail.com)

Received: 21 May 2018 Accepted: 22 November 2018

Published online: 21 December 2018

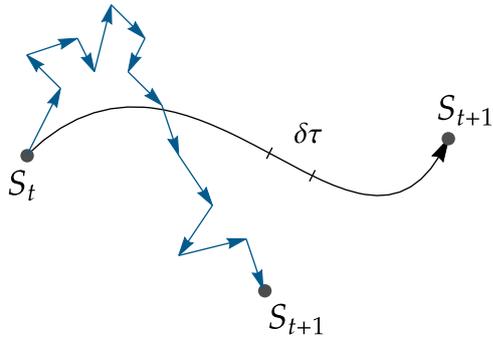


Fig. 1 Schematic representation of MMC and HMC trial state generation procedures. In the MMC algorithm each degree of freedom is updated individually. Such updates are called Monte Carlo steps and the acceptance decision is made at each step. Once an update attempt has been performed consecutively on all degrees of freedom a Monte Carlo sweep is considered to be accomplished and the new state S_{t+1} is generated. A Monte Carlo sweep in the MMC algorithm can be thus represented as a sequence of small random moves (Monte Carlo steps) in the configuration space (blue arrows in the figure). In contrast, within the HMC scheme a trial state S_{t+1} is generated by evaluating a random Hamiltonian trajectory starting from an initial state S_t (black arrow). Therefore, an HMC iteration corresponds to a smooth motion in the phase space. Moreover, all degrees of freedom are updated simultaneously and an acceptance decision is made only once, at the final point of the trial trajectory

potential energies is employed). The probability of accepting S'_t is taken to be equal to $w = \min(1, e^{-\beta\Delta E})$. In the case where the trial is accepted, S_t is set to S'_t , while otherwise the default state S_{t-1} is duplicated, or in other words S_t is set to S_{t-1} . The structure of the HMC scheme is therefore identical to that of the Metropolis algorithm and the main difference between the two methods resides in the recipe for choosing the trial states. While MMC relies on sequentially accumulating small random changes to each individual degree of freedom, the HMC scheme incorporates collective variable updates by generating Hamiltonian trajectories in the phase space of the system (see Fig. 1).

The correct canonical distribution of internal energies is ensured by choosing random values for generalized momenta p_i drawn from the Maxwell-Boltzmann distribution ρ as an initial condition for each trial trajectory

$$\rho(p_i) = \sqrt{\frac{2\pi m_i}{\beta}} e^{-\beta p_i^2 / 2m_i}, \quad (2)$$

where m_i denotes an effective mass associated with each microscopic degree of freedom. Interestingly, the values of the masses m_i can be chosen arbitrarily, since m_i can be eliminated from the equations of motion by a proper choice of the time units.

RESULTS

As a first step, we test the performance of the HMC scheme on the example of the effective Hamiltonian model¹⁴ describing the sequence of ferroelectric phase transitions of BaTiO₃ crystals. For comparison, the same simulations are also performed using MMC and thermalized MD schemes. Barium titanate is known to exhibit three structural phase transitions, all of which are successfully reproduced using the aforementioned set of algorithms (see Fig. 2). However, the results presented in Fig. 2 indicate that the transition temperatures obtained using MMC and MD schemes are slightly lower than the corresponding estimates obtained from HMC simulations. Moreover, the temperature mismatch increases with decreasing temperature—for the paraelectric to tetragonal transition temperature (HMC gives an estimate of $T_{\text{HMC}} \sim 380$ K)

the estimates differ by 3 K, while for the tetragonal to orthorhombic ($T_{\text{HMC}} \sim 285$ K) and orthorhombic to rhombohedral ($T_{\text{HMC}} \sim 230$ K) transition the mismatch reaches 7 and 15 K, respectively. Such differences can be attributed to the first order character of the phase transitions resulting in the temperature hysteresis. Indeed, the discrepancy of transition temperatures estimates obtained from cooling down and heating up cycles are well-expected in MD simulations since the relaxation time of metastable states can extend beyond the time scale reachable by MD algorithms. Therefore, it is expected that the transition temperature estimate would depend on the cooling rate, or equivalently on the number of MD steps performed at each temperature. Similarly, although the memory effects for MC sampling are less pronounced, the transition temperature would still depend on the number of sweeps performed at each temperature. In fact, the hysteresis width shall necessarily grow when decreasing the simulation-to-autocorrelation time ratio. Panels (g)–(i) of Fig. 2 present the temperature evolution of the autocorrelation times of polarization components obtained using the three considered algorithms. As it can be readily seen, in the vicinity of the phase transitions, the HMC scheme yields lower sample correlations than the MMC algorithm. Therefore, at a fixed number of sweeps, the HMC scheme should yield higher accuracy estimates of phase transition temperatures. Similar arguments hold for relative performance of MD and HMC algorithms, thus allowing us to explain the observed mismatch of estimated transition temperature values.

Noting good performance of the HMC scheme in reproducing ferroelectric transitions of bulk BaTiO₃, we now consider a more challenging simulation—the relaxation of the 180° domain wall in the tetragonal phase of bulk BaTiO₃ at $T = 300$ K. The initial state of the system is taken to be a supercell divided into two domains of equal volume and opposite orientations of polarization. Specifically, we choose the polarization in the two domains to be oriented along [010] and [0 $\bar{1}$ 0] pseudo-cubic directions, while the domain wall normal is taken to be aligned with the [100] axis (see panel (a) of Fig. 3). The simulation is performed for two different supercell sizes—in order to compare the performance of HMC scheme and thermalized molecular dynamics (MD) with that of MMC algorithm, we chose a $24 \times 12 \times 12$ supercell, while a more challenging relaxation case with the supercell size of $128 \times 32 \times 32$ will be used as a probe of the HMC scheme capabilities. For MD and HMC simulations we use a time step of 0.5 fs, and a single HMC trial trajectory (sweep) corresponds to 25 fs evolution, or, equivalently, 50 integration steps. To make a fair comparison of performances of HMC and MD algorithms we also define one MD sweep to consist of 50 MD steps.

Since we assume periodic boundary conditions along all of the three Cartesian directions to mimic a bulk crystal, the depolarizing fields that usually provoke breaking of the system into ferroelectric domains¹⁵ are absent and the equilibrium state at 300 K corresponds to a monodomain configuration with homogeneous distribution of polarization. The bi-domain state taken as the initial configuration is in fact unstable and we expect all the algorithms to converge to the equilibrium monodomain configuration. Panel (b) of Fig. 3 shows the evolution with sweeps of the total supercell polarization magnitude P obtained using the HMC, MD and MMC algorithms for the $24 \times 12 \times 12$ supercell size. While for the bi-domain state the polarization magnitude is zero $P = 0$, the monodomain state yields an equilibrium value of $P \sim 0.35$ C/m² and hence P can be used as a reaction coordinate characterizing the convergence. It can be readily seen (panel (b) of Fig. 3) that the HMC algorithm is able to arrive to an equilibrium state within ~ 800 sweeps ($\sim 2 \times 10^{11}$ floating-point operations, or 0.2 Tflop), while the MMC algorithm convergence is achieved within $\sim 10,000$ sweeps (4.8 Tflop). Furthermore, we find that the MD scheme is unable to converge within 10000 sweeps (500,000 MD steps or 2.8 Tflop)—the MD simulation convergence is achieved

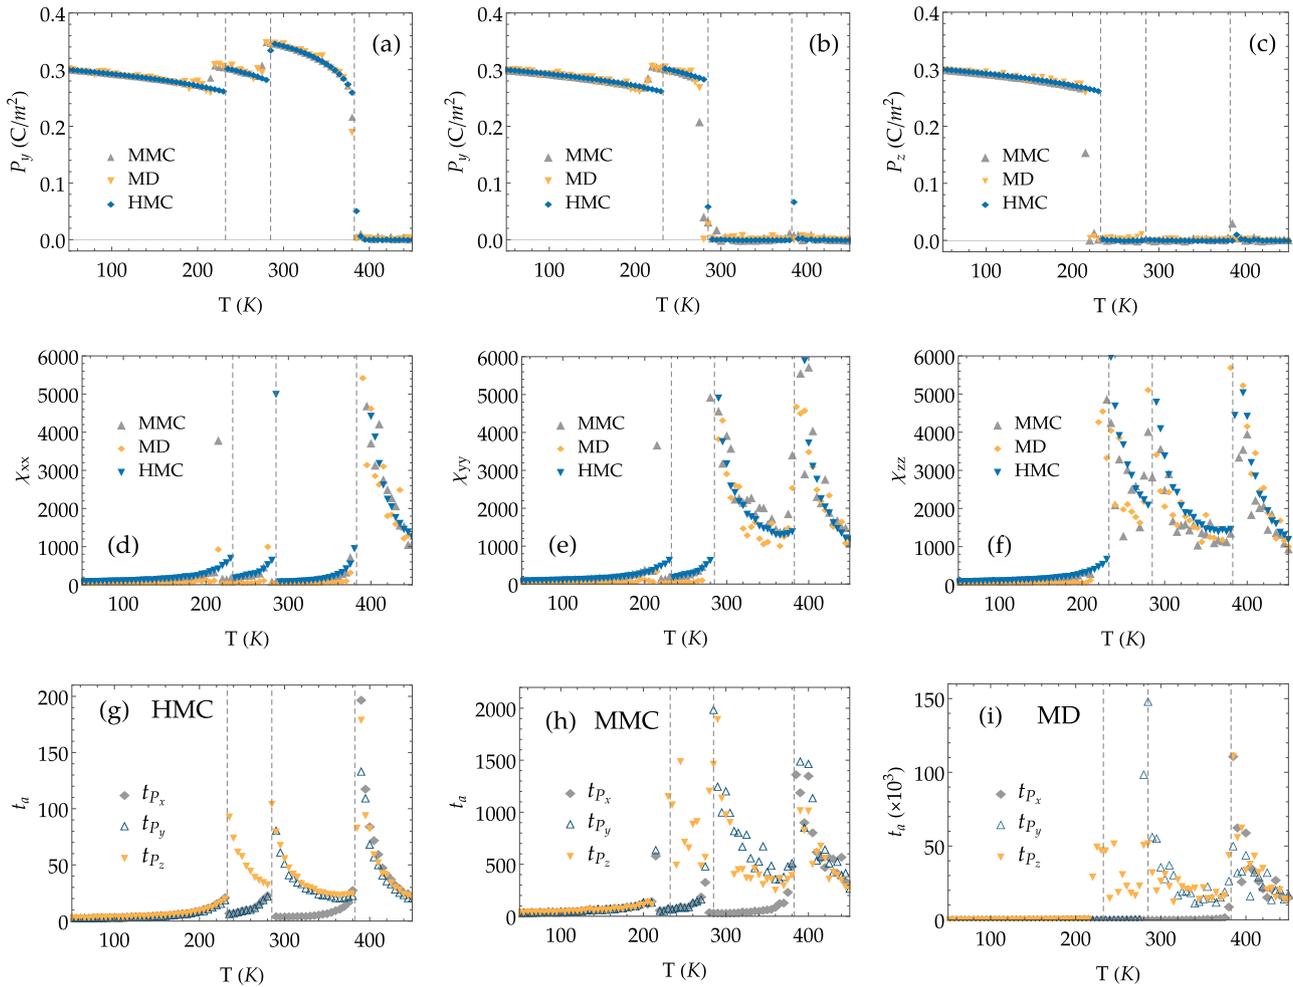


Fig. 2 Dependence of the polarization components (a–c) and dielectric susceptibility (d–f) of BaTiO₃ crystal on temperature calculated using HMC, MMC and MD algorithms. Simulation were performed using a $16 \times 16 \times 16$ supercell. At each temperature, 40,000 HMC (MMC) sweeps were performed, out of which the first 10,000 were considered as thermalization sweeps and the thermodynamic averages were computed over the remaining 30,000 sweeps. In MD simulations, we have 0.5 millions of thermalization and 1.0 million of averaging steps. The sequence of structural phase transitions (corresponding transition temperatures are indicated by dashed lines) is successfully reproduced by all algorithms. However, HMC scheme yields higher accuracy of thermodynamic averages as can be seen from (g–i) that show temperature evolution of autocorrelation times of polarization components obtained for HMC, MMC and MD simulations respectively. Note that while the autocorrelation time in (g, h) is in the units of MC sweeps, for (i) the autocorrelation time units are 10^3 MD steps. For each temperature, the total simulation cpu time for MMC, MD and HMC simulations was of 35.4, 19.5, and 19.6 min, respectively

only at $\sim 25,000$ MD sweeps (1,250,000 MD steps or 7 Tflop). Panel (c) shows the evolution with sweeps of the potential energy of the system obtained using HMC algorithm. During the first hundreds of sweeps, the evolution of the state can be described as the motion of the wall along its normal. At this stage, despite the growing polarization, the potential energy does not significantly change. An abrupt reduction of the potential energy happens only when the volume of one of the domains becomes small enough to allow the destruction of the domain wall that is triggered at \sim the 500th HMC sweep. Plotting the values of potential energy at each HMC sweep with respect to the polarization (see panel (d) of Fig. 3) allows estimation of the energy profile which results in a flat energy plateau in the vicinity of the initial state followed by a steep well at the monodomain minimum happening upon the collapse of the wall. In section 3 we provide arguments explaining higher performance of HMC scheme for such types of energy profiles. Furthermore, we find that for the case of $128 \times 32 \times 32$, the HMC algorithm appears to be the only scheme out of the three considered algorithms that allows efficient relaxation towards an equilibrium monodomain state. The MMC simulation at this supercell size is practically impossible due to the poor

scaling of the algorithm with the system size, while MD simulation yields at least a ~ 30 times larger relaxation time as established in the simulation of the $24 \times 12 \times 12$ supercell. We find that for the $128 \times 32 \times 32$ supercell test case, the HMC algorithm converges at $\sim 500,000$ sweeps (see Fig. S1 of supplemental material). Naturally, the increased relaxation time can be explained by a significantly larger distance in the configuration space between bi-domain initial and the monodomain states.

In order to achieve high computational performance of the HMC algorithm implementation for effective Hamiltonian simulations, we have employed the approach described in refs. ^{16,17} Specifically, computation of all energies and forces stemming from the non-local harmonic interactions are carried out in the reciprocal space using fast fourier transformed¹⁸ local mode and strain fields, while the single unit-cell quantities are computed using the corresponding lattice variables in real space. Such an approach allows for separate diagonalization of all parts of the Hamiltonian and results in $O(N/\log N)$ computational complexity. Such methodology proves useful for MD algorithms too, since it allows for efficient computation of the long-range dipolar forces at all lattice sites once the update of all lattice fields has been

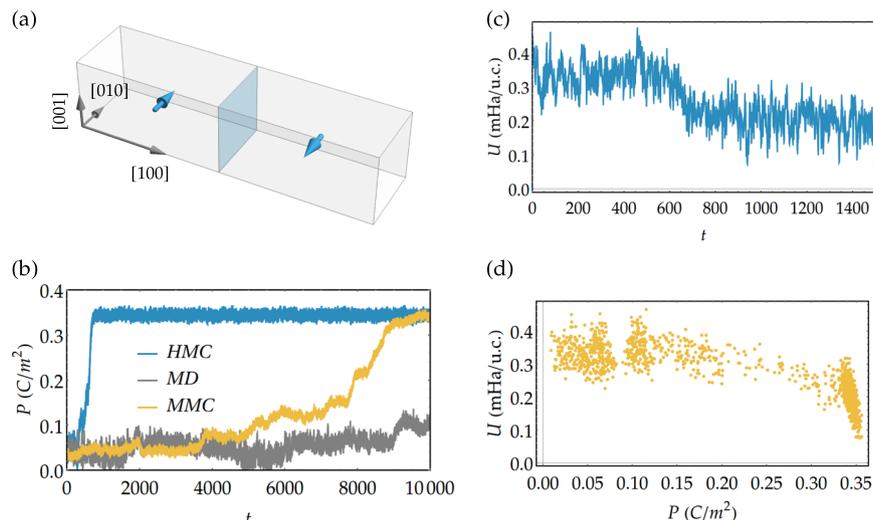


Fig. 3 Relaxation of 180° domain wall in the tetragonal phase of BaTiO_3 . **a** shows the schematization of the initial state with equal volume domains having opposite polarization orientations—along $[010]$ in the domain on the left and along $[0\bar{1}0]$ in the domain on the right. **b** shows the dependence of the absolute value of polarization on sweeps for HMC, MMC and MD algorithms. For the thermalized MD, we assume that a sweep is finalized after 50 molecular dynamics steps are performed. Such definition allows fair comparison of HMC and MD algorithms. **c** shows the evolution with sweeps of the potential energy obtained using HMC scheme. The plot of potential energy values versus the polarization magnitude at each HMC sweep is shown in **d**. The total cpu time of MMC, MD and HMC simulation was of 5.1, 5.5, and 5.5 min, respectively

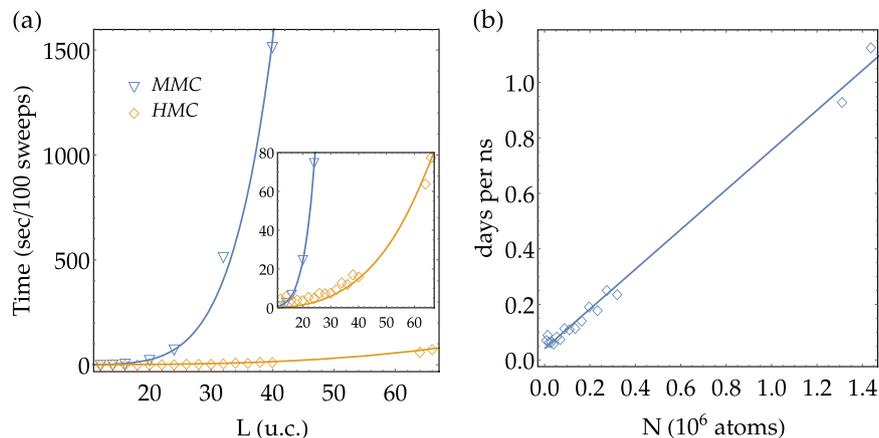


Fig. 4 Performance of the GPU-oriented implementation of the effective Hamiltonian model. **a** shows comparison of the times needed to execute 100 MC sweeps using Metropolis and HMC algorithms depending on the lateral size of the supercell. **b** presents the number of days needed to perform the HMC simulation of the time evolution of the system for 1 nanosecond (ns) time interval

performed.¹⁷ In contrast, the MMC scheme could hardly benefit from such Hamiltonian diagonalization—the lattice variables are updated sequentially, one after the other and each accepted trial move calls for an update of long-range fields at all sites yielding an $O(N) < O(N \log N)$ complexity for a single MMC step. However, since all N variables ought to be updated, the MMC sweep complexity increases to $O(N^2)$, while complexity of the HMC sweep stays on the order of $O(N \log N)$ since the algorithm relies on the MD-based generation of trial states. In other words, it is the sequential nature of MMC steps that represents a significant performance bottleneck. The comparison of real-life performance of MMC and HMC schemes is shown in panel (a) of Fig. 4, which makes the described scaling gap evident. The performance difference is all the more pronounced since in contrast to ref.¹⁷ we here adopt the GPU-oriented parallelization strategy.¹⁹ The use of such massively parallel architectures allows for MD and HMC simulation performance of ~ 1 ns per day for $N \sim O(10^6)$ using a single GPU as attested by our benchmark results shown in panel (b) of Fig. 4.

DISCUSSION

It can be readily noticed that the HMC scheme is more advantageous than the Metropolis algorithm since, theoretically, in the former, all generated trial states should be accepted irrespectively of the length of the trial trajectory. This follows from the total energy conservation property of Hamiltonian dynamics employed in HMC. In contrast, increasing the acceptance ratio within an MMC simulation comes at the cost of constraining the magnitude of random variations introduced to individual degrees of freedom during each MMC step. Indeed, in order to obtain higher acceptance probability w , the difference of energies between initial and trial states has to be reduced, which can only be achieved by making “shorter” steps in the configuration space. In other words, an attempt to reduce the amount of redundant states within the MMC chain by increasing the acceptance ratio ineluctably leads to an increase of redundancy due to generation of states that are very close to each other. Therefore, both “long” as well as “short” random steps yield a reduced sampling efficiency since more MMC sweeps will be required to obtain accurate

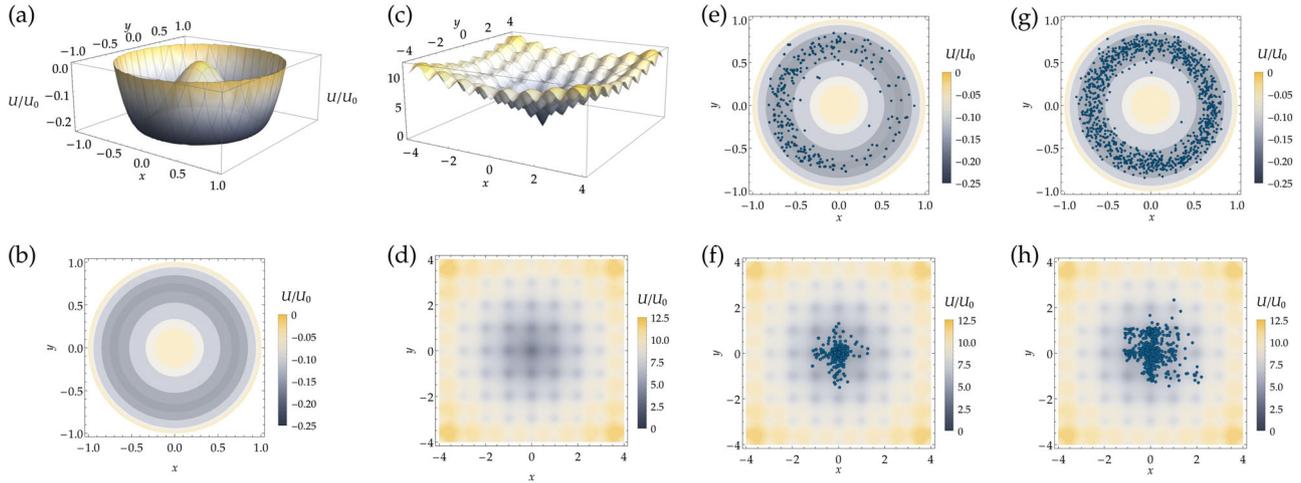


Fig. 5 Model *mexican-hat* **a, b** and *Ackley's function* **c, d** potentials used to test the relative performance of MMC and HMC algorithms. The generated sequences of 10^3 equilibrium states (shown as blue dots) for the model potential **a** are presented in **e, g**, while generated states for potential **(c)** are shown in **f, h**. **e, f** correspond to the results obtained using MMC sampling while **g, h** show the results obtained using the HMC algorithm. The temperature is chosen such that $k_B T = 0.02U_0$ for the potential **a** and $k_B T = 0.7U_0$ for potential **c**

expectation values. The optimal performance for the MMC algorithm was estimated²⁰ to be achieved when the acceptance ratio is between 20–40% meaning that at best 60 to 80% of the computational effort is wasted when using the MMC algorithm. In contrast, the HMC algorithm removes such a trade off scenario—the autocorrelation time can be decreased by simply increasing the trial trajectory length while conserving the acceptance ratio at its maximum. Note that in practice when the numerical integration is used to estimate trial Hamiltonian trajectories, the total energy conservation can be only approximately achieved and it is very important to opt for symplectic integration schemes⁶ to avoid the energy drift problems (See supplemental material for more information). Nonetheless, lowering the integration step size can always allow for acceptance ratios close to 100% in the HMC scheme.

To illustrate these arguments, we have conducted both MMC and HMC samplings for two model potentials – the so-called “*mexican-hat*” function²¹ (see panels **(a, b)** of Fig. 5)

$$U(x, y) = U_0 \left(-\sqrt{(x^2 + y^2)} + (x^4 + y^4) \right),$$

and *Ackley's function*²² (see panels **(c, d)** of Fig. 5)

$$U(x, y) = U_0 \left(-20e^{-0.2\sqrt{(x^2 + y^2)/2}} \right. \\ \left. - e^{-0.5(\cos 2\pi x + \cos 2\pi y)} \right. \\ \left. - e + 20 \right).$$

As can be readily seen, the chosen model potentials have different topologies and therefore can be used to test sampling efficiency in qualitatively different model situations. Indeed, the “*mexican-hat*” potential has a $U(1)$ degenerate ground state, while *Ackley's function* possesses a rugged energy landscape with a single and sharp global minimum at the origin of the coordinate system. Panels **(e)** and **(f)** of Fig. 5 show a sample of 10^3 equilibrium states (blue dots) obtained using MMC sampling scheme while the corresponding results obtained using HMC algorithm are presented in panels **(g)** and **(h)**. A superior efficiency of HMC sampling is evident for both test cases. Indeed, the HMC sampling of the degenerate ground state in the case of the “*mexican-hat*” potential is much more homogeneous (see panel **(g)** of Fig. 5), while for the case of *Ackley's function* the sampling of both the global minimum as well as higher energy, “metastable” states is significantly denser (see panel **(h)** of Fig. 5). The metastable states become more reachable due to

high acceptance ratio accessible at longer separations between initial and trial states, which in fact allows to explain better performance of HMC as compared to MMC algorithm in the BaTiO_3 annealing test.

An interesting feature of the HMC scheme that can be revealed in the test involving the *Mexican-hat* potential is the ability of the algorithm to efficiently sample energy plateaus and shallow valleys. In the *Mexican-hat* example, the shallow valley corresponds to the hat's brim—a continuously degenerate minimum of the potential energy. Within the HMC scheme, generating a random initial momentum tangential (or close-to-tangential) to the energy isolines will generate a quasi-circular trial trajectory that progresses along the brim. The corresponding generated trial state s'_i can easily advance far away from the initial point s_i . In contrast, sampling shallow valleys can be rather challenging for MMC and thermalized MD algorithms since on flat energy profiles, both schemes generate motion through the configuration space equivalent to a random walk. Panels **(a)** and **(b)** of Fig. 6 provide an illustration of this argument by showing 20 trial states generated using MMC (panel **(a)**) and HMC algorithms (panel **(b)**). In case of thermalized MD, the flatness of the energy landscape translates into smallness of intrinsic forces ($-\partial U/\partial x$) in comparison with the random force (also friction in case of, e.g., Langevin thermostat) stemming from the interaction of the system with the thermostat. For the *Mexican-hat* example, the positive and negative values of the projection of such random force on the curvilinear axis aligned with potential energy minimum isoline would be equiprobable. Therefore, since the value of the random force changes at each MD step the propagation along the brim of the hat would have only diffusive character. In contrast, in case of the HMC sampling, the random shuffling of momenta does not happen at each Hamiltonian dynamics step but rather at the beginning of each HMC sweep which allows the system to propagate much further along the potential energy isoline.

The argument discussed above allows to easily understand the superior efficiency of the HMC scheme, not only in the case of the described toy-model example (see panels **(g)** and **(e)** of Fig. 5), but also in the bulk BaTiO_3 domain-wall relaxation simulation discussed in section 2. In the latter case, the initial state $s_{\uparrow\downarrow}$ (corresponding to supercell comprising two domains of equal volume) is equidistant in configuration space from the mono-domain ground states s_{\uparrow} and s_{\downarrow} with polarization oriented along $[010]$ and $[0\bar{1}0]$, respectively. The potential energy excess $\delta\varepsilon > 0$ of

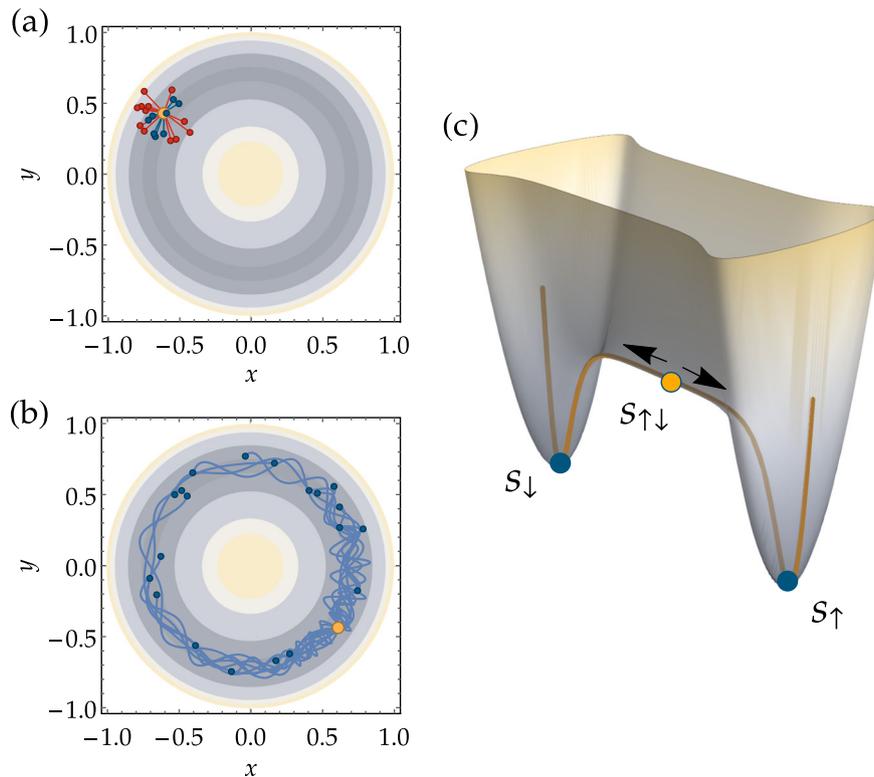


Fig. 6 Sets of 20 randomly generated trial states for the model Mexican-hat potential using MMC **a** and HMC **b** algorithms. Temperature is chosen so that $k_B T = 0.02 U_0$. Trial states that would have been rejected and accepted are indicated with red and blue dots, respectively. The initial state is indicated by a yellow dot. **c** shows a schematic representation of the energy landscape for the case of domain wall relaxation in bulk BaTiO₃ discussed in section 2. The yellow circle indicates the $s_{\uparrow\downarrow}$ microscopic state corresponding to equal-volume domains with opposite polarization orientations, while s_{\uparrow} and s_{\downarrow} states (blue circles) correspond to two symmetry-related monodomain ground states with polarization pointing along [001] and [00 $\bar{1}$] p.c. directions, respectively. Motion through the shallow valley connecting $s_{\uparrow\downarrow}$ and s_{\uparrow} (s_{\downarrow}) microscopic states is achieved by a shifting the domain wall along its normal

the $s_{\uparrow\downarrow}$ state with respect to either s_{\uparrow} or s_{\downarrow} is due to the energy of the domain wall.²³ Furthermore, shifting the domain wall along its normal leads to a small ($\ll \delta\epsilon$) reduction of the potential energy, while other transformations of the dipolar configuration lead to an energy increase. For example, the rotation of the domain wall under periodic boundary conditions necessarily creates additional boundaries between the two domains. All these considerations lead to the following conclusions: (i) the $s_{\uparrow\downarrow}$ state corresponds to a saddle point; (ii) the steepest descent from $s_{\uparrow\downarrow}$ to either one of the monodomain states is achieved by the parallel translation of the wall along its normal and (iii) the one dimensional passage between s_{\uparrow} and s_{\downarrow} via $s_{\uparrow\downarrow}$ is narrow and rather flat in the vicinity of the saddle. A schematization of such energy landscape is presented in Fig. 6c. In the vicinity of the initial state $s_{\uparrow\downarrow}$, the potential energy profile topologically resembles the brim of the Mexican-hat—a situation for which the HMC scheme performance is expected to be higher than that of MMC and thermalized MD algorithms.

It is equally important to note the trade offs that come along with the described advantages of the HMC scheme. Firstly, in terms of simplicity of implementation, the MMC scheme remains unmatched since an HMC code, especially its parallel implementation, requires much more effort to develop and debug. Moreover, parallelization strategies used in this study are not as efficient for problems involving small number of degrees of freedom (N of the order of 10^4 or less) and can be very challenging to optimize for systems with only short-range interactions. Therefore, in some situations it might be more practical to resort to implementation of Metropolis algorithm even though the number of sweeps required to achieve the same accuracy as with HMC simulation

might be higher. Finally, the HMC scheme inherits from molecular dynamics its inapplicability to models with discrete degrees of freedom, e.g., Ising model and its extensions, lattice gas, etc.

The performed analysis therefore reveals that Hybrid Monte Carlo scheme can prove to be useful as an algorithm for structure and property predictions in the computationally demanding case of systems with long-range interactions. Specifically, the tests we performed using the effective Hamiltonian model of a prototypical ferroelectric material BaTiO₃ show that the HMC scheme not only inherits efficient parallelization strategies that allow simulation of systems consisting of $N \sim 10^6$ particles, but can also offer significant performance gains when compared to Metropolis Monte Carlo and thermalized MD simulations. Cases particularly suitable for HMC simulations include systems with energy landscapes exhibiting plateaux and shallow valleys. Landscapes with multiple metastable states can be also efficiently simulated using HMC algorithm, although for this particular case more specialized algorithms will most likely exhibit better performance. Based on the presented arguments, we strongly believe that HMC simulations would allow to tackle current challenging problems, such as simulations of new functional materials at a new level of accuracy and length scale. Moreover, the use of this scheme can extend the reach of MC methods not only to systems exhibiting long-range interactions, but more generally to structure prediction cases where the energy calculations are computationally demanding, such as ab initio simulations. Although the test of the HMC performance for on-the-fly DFT calculations lies beyond the scope of this study, we have implemented the algorithm in the Abinit software suite²⁴. Based on several model tests used to validate the implementation (single unit cell PbTiO₃ structural relaxation and

72 atom amorphous SiO₂ structural relaxation) we did not find significant difference in performance of thermalized MD and HMC algorithms. However, these tests revealed several practical advantages of the HMC algorithm. Firstly, the Metropolis decision test present in the HMC algorithm allows to automatically reject configurations for which the self consistent cycle did not converge to a required accuracy in contrast to thermalized MD simulation. This can allow to save some time required for simulation setup. Another advantage of the HMC scheme resides in that a hybrid Monte Carlo update of reduced coordinates can be easily combined with MMC updates of lattice vectors needed to optimize the unit cell geometry. Such combination of algorithms would therefore remove the need for implementing barostats, introduce additional auxiliary input parameters (e.g. mass of the barostat) and ease the implementation of geometrical constraints. Furthermore, based on the test cases presented in the manuscript we believe that HMC scheme can prove useful for large-scale on-the-fly DFT simulations and that this study would encourage further tests of our implementation of the HMC algorithm in Abinit. Our open-source GPU oriented implementation of the effective Hamiltonian code will be shortly made publicly available at www.lattiscope.com.

METHODS

For simulations reproducing the BaTiO₃ phase transition sequence, we use a 16 × 16 × 16 supercell. In case of HMC and MMC schemes 30,000 sweeps are used at each temperature to compute thermodynamic averages after a 10,000 sweep thermalization. For MD simulations we use the Evans-Hoover thermostat that allows for sampling the NPT ensemble. At each temperature 1.5 × 10⁶ MD steps of 1 fs are performed, out of which the first 200 ps are considered as a thermalization period. Within HMC simulations, each trial trajectory corresponds to 40 steps of 1 fs. For domain wall relaxation simulations we use 0.5 fs integration steps for both MD and HMC simulations. For the HMC test, the trial trajectory consists of 50 integration steps for both considered supercell sizes (24 × 12 × 12 and 128 × 32 × 32). The employed effective Hamiltonian model¹⁴ includes on-site, short-range and long-range (dipolar) local mode interactions, the strain elastic energy, as well as electrostrictive interactions of homogeneous and inhomogeneous strains with local modes.

DATA AVAILABILITY

The authors declare that all data supporting the findings of this study are available within the paper and its supplementary information file.

ACKNOWLEDGEMENTS

S.P. and L.B. thank the DARPA Grant HR0011-15-2-0038 (MATRIX program). K.K. acknowledges a SURF grant from the state of Arkansas, Y.N. and L.B. thank the DARPA Grant No. HR0011727183-D18AP00010 (TEE Program). All authors are grateful for support provided by NVIDIA via the NVIDIA GPU Grant. Computations were made possible thanks to the use of the Arkansas High Performance Computing Center and the Arkansas Economic Development Commission. DARPA Grant HR0011-15-2-0038 (MATRIX program), DARPA Grant No. HR0011727183-D18AP00010 (TEE Program), SURF grant from the state of Arkansas, NVIDIA GPU Grant.

AUTHOR CONTRIBUTIONS

S.P. initiated the study. S.P. and K.K. implemented the GPU oriented code for effective Hamiltonian simulations. S.P. and Y.N. performed numerical simulations. L.B. supervised the study. All authors participated in discussing the results and manuscript preparation.

ADDITIONAL INFORMATION

Supplementary information accompanies the paper on the *npj Computational Materials* website (<https://doi.org/10.1038/s41524-018-0137-0>).

Competing interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

- Newman, M. E. J. & Barkema, G. T. *Monte Carlo Methods in Statistical Physics*. (Oxford University Press Inc., New York, 2001).
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N. & Teller, A. H. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).
- Swendsen, R. H. & Wang, J.-S. Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.* **58**, 86–88 (1987).
- Wolff, U. Collective Monte Carlo updating for spin systems. *Phys. Rev. Lett.* **62**, 361 (1989).
- Duane, S., Kennedy, A. D., Pendleton, B. J. & Roweth, D. Hybrid Monte Carlo. *Phys. Lett. B* **195**, 216 (1987).
- Betancourt, M. A conceptual introduction to Hamiltonian Monte Carlo. Preprint at <https://arxiv.org/abs/1701.02434> (2017).
- Tagawa, T., Kaneko, T. & Miura, Sh On computational efficiency of the hybrid Monte Carlo method applied to the multicanonical ensemble. *Mol. Simul.* **43**, 1291 (2017).
- Knott, B. C. et al. Homogeneous nucleation of methane hydrates: unrealistic under realistic conditions. *J. Am. Chem. Soc.* **134**, 19544 (2012).
- Mehlig, B., Heermann, D. W. & Forrest, B. M. Hybrid Monte Carlo method for condensed-matter systems. *Phys. Rev. B* **45**, 679 (1992).
- Drut, J. E. & Porte, W. J. Hybrid Monte Carlo approach to the entanglement entropy of interacting fermions. *Phys. Rev. B* **92**, 125126 (2015).
- Körner, M., Smith, D., Buividovich, P., Ulybyshev, M. & Smekal, L. Hybrid Monte Carlo study of monolayer graphene with partially screened Coulomb interactions at finite spin density. *Phys. Rev. B* **96**, 195408 (2017).
- Phillips, J. C., Braun, R., Wang, W., Gumbart, J., Tajkhorshid, E., Villa, E., Chipot, Ch, Skeel, R. D., Kale, L. & Schulten, K. Scalable molecular dynamics with NAMD. *J. Comput. Chem.* **26**, 1781 (2005).
- Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comp. Phys.* **117**, 1 (1995).
- Walizer, L., Lisenkov, S. & Bellaiche, L. Finite-temperature properties of (Ba,Sr)TiO₃ systems from atomistic simulations. *Phys. Rev. B* **73**, 144105 (2006).
- Kittel, C. *Introduction To Solid State Physics* 8th edn, (Wiley, Hoboken, NJ, USA, 2004).
- Waghmare, U., Cockayne, E. J. & Burton, B. P. Ferroelectric phase transitions in nano-scale chemically ordered PbSc_{0.5}Nb_{0.5}O₃ using a first-principles model hamiltonian. *Ferroelectrics* **291**, 187 (2003).
- Nishimatsu, T., Waghmare, U. V., Kawazoe, Y. & Vanderbilt, D. Fast molecular-dynamics simulation for ferroelectric thin-film capacitors using a first-principles effective Hamiltonian. *Phys. Rev. B* **78**, 104104 (2008).
- Brigham, E. O. *The Fast Fourier Transform*. (Prentice-Hall, New York, 2002).
- John Nickolls, J., Buck, I., Garland, M. & Skadron, K. Scalable parallel programming with CUDA. *ACM Queue* **6**, 40 (2008).
- Landau, D. P. & Binder, K. *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, Cambridge, UK, 2014).
- Altland, A. & Simons B. D. *Condensed Matter Field Theory* (Cambridge University Press, Cambridge, UK, 2010).
- Ackley, D. H. *A Connectionist Machine for Genetic Hillclimbing*. (Kluwer Academic Publishers, Boston MA, 1987).
- Chaikin P. M. & Lubensky T. C. *Principles of Condensed Matter Physics* (Cambridge University Press, Cambridge, UK, 2012).
- Gonze, X. et al. Recent developments in the ABINIT software package. *Comput. Phys. Commun.* **205**, 106 (2016).



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018