

Accelerating Hybrid Monte Carlo simulations of the Hubbard model on the hexagonal lattice

Stefan Krieg^a, Thomas Luu^{b,c}, Johann Ostmeyer^c, Philippos Papaphilippou^d,
Carsten Urbach^c

^a*Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Jülich, Germany*

^b*Institute for Advanced Simulation (IAS-4), Institut für Kernphysik (IKP-3), and Jülich Center for Hadron Physics, Forschungszentrum Jülich GmbH, Jülich, Germany*

^c*Rheinische Friedrich-Wilhelms-Universität Bonn, D-53012 Bonn, Germany*

^d*Dept. of Computing, Dept. of Electrical and Electronic Engineering, Imperial College London, London, United Kingdom*

Abstract

We present different methods to increase the performance of Hybrid Monte Carlo simulations of the Hubbard model in two-dimensions. Our simulations concentrate on a hexagonal lattice, though can be easily generalized to other lattices. It is found that best results can be achieved using a flexible GMRES solver for matrix inversions and the second order Omelyan integrator with Hasenbusch acceleration on different time scales for molecular dynamics. We demonstrate how an arbitrary number of Hasenbusch mass terms can be included into this geometry and find that the optimal speed depends weakly on the choice of the number of Hasenbusch masses and their values. As such, the tuning of these masses is amenable to automatization and we present an algorithm for this tuning that is based on the knowledge of the dependence of solver time and forces on the Hasenbusch masses. We benchmark our algorithms to systems where direct numerical diagonalization is feasible and find excellent agreement. We also simulate systems with hexagonal lattice dimensions up to 102×102 and $N_t = 64$. We find that the Hasenbusch algorithm leads to a speed up of more than an order of magnitude.

Keywords: Hybrid Monte Carlo, Hubbard model, Graphene, Auto-tuner

1. Introduction

The Hubbard model has been used to describe many interesting phenomena relevant to low-dimensional solid-state materials, such as Mott insulating behaviour, anti-ferromagnetic order, etc. For the one-dimensional system, under

Email addresses: s.krieg@fz-juelich.de (Stefan Krieg), t.luu@fz-juelich.de (Thomas Luu), ostmeyer@hiskp.uni-bonn.de (Johann Ostmeyer), pp616@ic.ac.uk (Philippos Papaphilippou), urbach@hiskp.uni-bonn.de (Carsten Urbach)

the Bethe Ansatz [1], there is no first-order transition between an insulating and conducting state. For two-dimensions (and higher), such a transition is expected to occur at some critical coupling, but to date the exact value of this coupling is unknown. For weak coupling below the critical coupling, standard perturbative many-body techniques have been applied to this system [2]. But for coupling near or above the critical coupling, very little is known *analytically* about the Hubbard model in these higher dimensions (see, e.g. [3]). What is known has been deduced from numerical simulations, such as via Random Phase Approximation [4], mean field theories [5], cluster theories [6], or functional renormalization group techniques [7, 8]. Recent calculations on modest system sizes have utilized Quantum Monte Carlo (QMC) techniques [4, 9, 10, 11, 12]. These QMC calculations employ either the Hybrid Monte Carlo (HMC) algorithm [13], a method used extensively in lattice quantum chromodynamics (QCD) [14, 15, 16, 17], or the BSS algorithm [18] which is more prevalent in the condensed matter community [19, 20]. In [20] lattices as large as 36×36 were simulated, providing the most precise value of the critical coupling in which a Mott insulating phase occurs. The largest hexagonal lattice sizes simulated to date are 48×48 with $N_t = 80$ timesteps [21]¹.

Many of these studies have been further motivated by the discovery of graphene, first isolated in 2004 by Andre Geim and Konstantin Novoselov. Consisting of carbon ions within an hexagonal lattice, graphene's unique electrical, thermal, and mechanical properties make it a prime candidate (as well as its derivatives, e.g. carbon nanotubes, ribbons, bi-layers) for next generation electrical devices. Numerical simulations of this system (via the Hubbard model and Hamiltonians with more sophisticated electron-electron interactions [22]) will play an essential role in quantifying graphene's (and its allotropes') electrical properties within macroscopic volumes. The ability to simulate these systems at these scales is, therefore, of prime importance.

The aim of this work is to describe our investigations for speeding up HMC simulations of the Hubbard model on an hexagonal lattice, as well as systems based off this geometry, such as graphene and carbon nanotubes. By utilising techniques originally developed in the lattice QCD community, we are able to obtain a speedup of more than an order of magnitude in some cases. This has allowed us to simulate on system sizes as large as 102×102 using single-node runs. In the following section we give a cursory description of the system and its underlying Hamiltonian. Section 3 and section 4 describe our solvers² and integrators used when integrating our equations of motions, respectively. In

¹The total number of lattice sites for an $n \times m$ hexagonal lattice is $2nm$. The total dimension of a simulated system includes the number of timesteps N_t and is $2nmN_t$.

²We note the recent development of an efficient direct solver via Schur complement [23]. However our work targets system sizes that are much larger than the lattices for which the Schur complement solver is feasible. Ref. [23] argues that a standard CG solver becomes preferable to Schur complement for systems larger than approximately (70×70) unit cells. As our algorithm is much more efficient than standard CG solver we expect it to be faster than Schur complement at even smaller lattices.




Figure 1: Honeycomb lattice of graphene. One red and one black adjacent points form a basis of the triangular Bravais lattice.

section 5 we describe our multi-shift solvers, where we provide the scaling and dependence of our algorithms with the Hasenbusch³ masses. The following section provides our algorithm for the autotuning of these masses, as well as other relevant parameters to our simulations. We demonstrate the speedup of our simulations due to the Hasenbusch solvers in section 8. Here we also benchmark our simulation to the exact solutions from a 4-site problem, ensuring that our solver implementations simulate correct physics. We end with a recapitulation of our findings in section 9.

2. The system

The honeycomb lattice, i.e. a lattice of rhombus-shaped unit cells with a basis of two atoms, consists of two underlying triangle lattices as depicted in figure 1. This bipartite structure means that any site of one sublattice has three nearest neighbours from the other sublattice. The minimal unit cell, from which the hexagonal lattice is tessellated, consists of two sites A and B, each coming from one of the sublattices. For this work, the orientation of the unit cell is chosen in such a way that the distance between A- and B-site is exactly $(a, 0)^T$ where a is the lattice spacing.

In order to investigate the basic properties of the Hubbard model in this geometry, Hybrid Monte Carlo⁴ (HMC) [13] simulations are used. The underlying Hamiltonian, after Hubbard-Stratonovich transformation and introduction of pseudofermions[17], is

$$\mathcal{H} = \frac{1}{2}\phi^T \frac{1}{\delta U} \phi + \chi^\dagger (MM^\dagger)^{-1} \chi + \frac{1}{2}\pi^T \pi \quad (1)$$

where π is the real momentum field, ϕ is the real Hubbard field, χ is a complex pseudofermionic vector field, δ is the step size in euclidean time dimension

³Hasenbusch acceleration is an exact preconditioning scheme first introduced in [24].

⁴Sometimes referred to as Hamiltonian Monte Carlo

weighted by the inverse temperature β of the system, U is the Hubbard onsite interaction, and M is the fermion operator (see ref. [17]) with

$$\begin{aligned} M_{(x,t)(y,t')}^{AA} &= \delta_{xy} (-\delta_{tt'} + (e^{i\phi_{x,t}} - \delta m_s) \delta_{t-1,t'}) \\ M_{(x,t)(y,t')}^{BB} &= \delta_{xy} (\delta_{tt'} - (e^{-i\phi_{x,t}} - \delta m_s) \delta_{t+1,t'}) \\ M_{(x,t)(y,t')}^{AB} &= M_{(x,t)(y,t')}^{BA} = -\delta\kappa\delta_{\langle x,y\rangle}\delta_{t,t'} . \end{aligned} \quad (2)$$

Here κ is the hopping parameter, $\delta_{\langle x,y\rangle}$ denotes nearest neighbour connections, and m_s is the staggered mass. A nonzero m_s corresponds to a broken symmetry between A- and B-sites, as it introduces a magnetic bias where the spin-up-electrons favour A-sites and spin-down-electrons favour B-sites or the other way around depending on the sign of m_s . All our calculations here are performed with $m_s = 0$. (In [25] it was pointed out that certain fermion operator discretizations with $m_s = 0$ suffer from ergodicity issues. We stress that our discretization in eq. (2) does not suffer from these same ergodicity issues.) Details related to the derivation of this Hamiltonian can be found in ref. [17] and [26], where also an implementation of the standard HMC algorithm to this system was discussed. For an $n \times m$ hexagonal lattice, we apply periodic boundary conditions along the oblique axes, as depicted in fig. 4 of [27]. We note that for graphene and carbon nanotube systems, $\kappa \sim 2.7$ eV and $a = 1.42$ Å [28].

The Hamiltonian (1) describes two flavours of spinless fermions, which can be interpreted either as a single electron species with spin one half or as “particles” and “holes”. In our formalism the latter interpretation has been chosen.

We outline the basic principles of the standard HMC algorithm. The ‘conjugate momenta’ field π and auxiliary complex field ρ are sampled according to Gaussian distributions $e^{-\pi^2/2}$ and $e^{-\rho^\dagger\rho}$, respectively. Then the pseudofermionic field is obtained as $\chi = M\rho$. With these starting parameters and an initial Hubbard field ϕ a molecular dynamics (MD) trajectory is calculated and the result is accepted with the probability $\min(1, e^{-\Delta\mathcal{H}})$. $\Delta\mathcal{H}$ is the difference in energy resulting from the molecular dynamics.

3. Solver

Note that equation (1) involves a matrix “inversion” via the matrix equation $(MM^\dagger)\eta = \chi$. This equation must be solved for η very often, making it beneficial to optimize the linear solver in this case. Observe that MM^\dagger is hermitian and positive definite. This allows us to use the Conjugate Gradient (CG) method (see e.g. ref. [29]) to solve the linear system involved. The CG “solver”, with relative tolerance set to 10^{-8} (i.e. squared tolerance 10^{-16}), and implemented in C++ with double precision was used initially to calculate η .

As a first step to increase the performance another CG solver was added with single precision arithmetic. This single precision CG (spCG) has been used as a preconditioner for the double precision CG. As the spCG with a precision target is a non stationary solver, we have to use the flexibly preconditioned CG (fCG) method[30]. We have tested this “mixed precision” algorithm on systems

of different sizes. On larger systems it led to a speed up of about factor 1.9, thus using nearly the full advantage of factor two between the bandwidths (and peak flop rates) of single and double precision arithmetics. As a next step we replaced the (“outer”) fCG solver by a flexible Generalized Minimal Residual (fGMRES) solver from ref. [31]. After some fine tuning described in Appendix A, a speed up of factor 2.2 compared to the original solver could be achieved for systems with a total dimension $2L^2N_t$ larger than $\sim 100\,000$.

On very large systems one encounters an additional advantage of the spCG preconditioner. The size of the L3-cache is limited and eventually will be too small to store all the vectors required for the solver in double precision. Using single precision arithmetics gives an additional factor of two in the system size until the memory requirements of the spCG preconditioner exceeds the size L3-cache as well. On the compute nodes we used for testing ($2 \times$ Intel(R) Xeon(R) CPU E5-2680, 40 MB Cache per node), this leads to a speed up of more than factor 4 instead of the expected factor two. For more details on the performance see section 8.

4. Integrators

In order to improve the calculation of the MD trajectory in the HMC simulations, we tested several different reversible, symplectic integrators⁵. An integrator evolves a system of coordinates $\phi(t)$ and momenta $\pi(t)$ depending on the force $F(\phi)$ by a time step Δt . Any integrator which is linear in Δt follows the rules:

$$\phi_i = \phi_{i-1} + c_i \Delta t \pi_{i-1} \quad (3)$$

$$\pi_i = \pi_{i-1} - d_i \Delta t F(\phi_i) \quad (4)$$

Here $(\phi_0, \pi_0) = (\phi(0), \pi(0))$ and $(\phi_n, \pi_n) = (\phi(\Delta t), \pi(\Delta t))$ with $i \in \{0, \dots, n\}$. In addition the condition $\sum_i c_i = \sum_i d_i = 1$ has to hold.

The most simple symmetric, thus reversible, symplectic integrator is the “leapfrog” integrator. The leapfrog is a second order method ($c_1 = c_2 = \frac{1}{2}$, $d_1 = 1$, $d_2 = 0$ in the velocity version).

4.1. Second order Omelyan integrator

In ref. [32] it was shown that the second order integrator using $c_1 = c_3 = \zeta$, $c_2 = 1 - 2\zeta$, $d_1 = d_2 = \frac{1}{2}$, $d_3 = 0$ with $\zeta \approx 0.193$, even though it requires two force calculations, is more effective as it minimizes the leading order error term significantly allowing for a coarser time step. The error can be estimated by expanding the evolution of two non commutative operators a and b (in our case they correspond to the evolution of ϕ and π respectively) in the form

$$e^{(a+b)h} = e^{a\zeta h} e^{bh/2} e^{a(1-2\zeta)h} e^{bh/2} e^{a\zeta h} + Ch^3 + \mathcal{O}(h^4) \quad (5)$$

⁵We confined ourselves to reversible, symplectic integrators to ensure detailed balance in our HMC algorithm[13]. Reversibility is ensured by symmetry which excludes odd order integrators.

and minimizing C . In general, when no assumptions are made about the operators a and b , then the commutator $[a, b]$ is not known. In this case Omelyan used an Euclidean norm spanned by the commutators $[a, [b, a]]$ and $[b, [b, a]]$:

$$C = c_1[a, [b, a]] + c_2[b, [b, a]] \quad (6)$$

$$\Rightarrow \|C\| = \sqrt{c_1^2 + c_2^2}. \quad (7)$$

The minimum has the numerical value $\|C\| \approx 0.009$ when $\zeta = 0.193$ (when $\zeta = 0$ or $1/2$ one obtains the original leapfrog integrator with $\|C_0\| \approx 0.09$). As the Omelyan integrator uses two force calculations per step the error has to be compared at double step size leading to the error ratio

$$\frac{\epsilon_{\text{Omelyan}}}{\epsilon_{\text{leapfrog}}} \approx \frac{0.009}{0.09} 2^2 = 0.4 \quad (8)$$

where ϵ denotes the different errors.

We tested different values of ζ in our simulations with the Omelyan integrator and found the theoretically derived value $\zeta = 0.193$ to be optimal, resulting in a factor of two to three speed-up compared to the generic method.

4.2. Higher order integrators

A detailed list of symmetric symplectic integrators can be found in ref. [33]. Several 4th order integrators were analysed in our molecular dynamics calculations. As our goal was not to obtain extremely precise solutions⁶ but rather to achieve a good acceptance rate with minimal cost, we did not investigate yet higher order integrators. The acceptance is weighted by the Boltzmann factor $e^{-\Delta\mathcal{H}}$ where $\Delta\mathcal{H}$ is the dimensionless deviation in energy accumulated over one trajectory. Targeting an acceptance of about 0.66 translates to $\Delta\mathcal{H} \approx 0.42$.

We found the force gradient integrator as described in ref. [34] most promising. The algorithm reads

$$\begin{aligned} \pi_1 &= \pi_0 - \frac{1}{6}\Delta t F(\phi_0) \\ \phi_1 &= \phi_0 + \frac{1}{2}\Delta t \pi_1 \\ \varphi &= \phi_1 - \frac{1}{24}\Delta t^2 F(\phi_1) \\ \pi_2 &= \pi_1 - \frac{2}{3}\Delta t F(\varphi) \\ \phi_3 &= \phi_2 = \phi_1 + \frac{1}{2}\Delta t \pi_2 \\ \pi_3 &= \pi_2 - \frac{1}{6}\Delta t F(\phi_2) \end{aligned} \quad (9)$$

and has an error scaling in $\mathcal{O}(\Delta t^4)$ with three force calculations per iteration. Figure 2 shows the average absolute deviation $\langle |\Delta\mathcal{H}| \rangle$ made by different inte-

⁶The accept/reject step renders the algorithm exact even for finite precision integrators[13].




Figure 2: Average absolute deviation of the energy after one trajectory $\langle |\Delta \mathcal{H}| \rangle$ for different integrators depending on the cost. The Omelyan integrator requires two force calculations per step, both 4th order integrators require three. The blue line shows a fit of the error of the force gradient integrator.

grators for an arbitrarily chosen system. In the region of interest with the error approximately 0.42 the force gradient integrator is comparable to the Omelyan integrator. However, the problem here is the well known divergence of the error with slight increase of Δt for both algorithms. The fit alone thus does not tell which of the integrators is preferable. In this case the Omelyan integrator seems to be slightly more stable in the considered region but the advantage gained by using one or the other integrator is negligible.

Our findings suggest that it really does not matter whether one uses the Omelyan integrator or the force gradient integrator. From all inspected integrators these two are the best (compare exemplary the “standard” 4th order integrator by Forest and Ruth [35] in figure 2) and we decided to use the Omelyan integrator because it is more flexible for modifications as described in the following sections. In addition one should mention that due to the improvements stated in this article the calculation of the molecular dynamics is not the bottleneck any more. We spend at least the same amount of time on the measurements after having calculated the trajectories.

5. Multi-scale integration

So far, we have discussed improvements to the linear solvers and to the integrators used in the computation of the MD trajectory. The integration

follows Hamilton's equations of motion

$$\dot{\phi}^T = \frac{\partial \mathcal{H}}{\partial \pi} = \pi^T \quad (10)$$

$$\dot{\pi}^T = -\frac{\partial \mathcal{H}}{\partial \phi} = -\phi^T \frac{1}{\delta U} + 2 \operatorname{Re} \left(\eta^\dagger \frac{\partial M}{\partial \phi} \xi \right) \quad (11)$$

with the two complex vector fields

$$\eta = (MM^\dagger)^{-1} \chi, \quad \xi = M^\dagger \eta. \quad (12)$$

The computational costs are dominated by the linear solver. Further improvements to the algorithm are possible by realizing that the rate of convergence of the linear solver and thus the speed of the MD integration is given by the condition number of the fermion matrix, i.e. it is dominated by its low eigenmodes. The algorithmic performance of the solver has already been improved and, while it is believed that there is still room for optimization, here a second approach is investigated.

5.1. Splitting off the small eigenvalues of the fermion matrix

In ref. [24] Hasenbusch proposed to replace the pseudofermionic term in the Hamiltonian (eqn. 1) by two terms:

$$\chi^\dagger (MM^\dagger)^{-1} \chi \mapsto \chi_1^\dagger (MM^\dagger + \mu^2)^{-1} \chi_1 + \chi_2^\dagger \left(\frac{MM^\dagger}{MM^\dagger + \mu^2} \right)^{-1} \chi_2 \quad (13)$$

This is possible because the determinant

$$\det(MM^\dagger) = \int \mathcal{D}\chi \mathcal{D}\chi^\dagger e^{-\chi^\dagger (MM^\dagger)^{-1} \chi} \quad (14)$$

$$= \int \mathcal{D}\chi_1 \mathcal{D}\chi_1^\dagger \mathcal{D}\chi_2 \mathcal{D}\chi_2^\dagger e^{-\chi_1^\dagger (MM^\dagger + \mu^2)^{-1} \chi_1 - \chi_2^\dagger \left(\frac{MM^\dagger}{MM^\dagger + \mu^2} \right)^{-1} \chi_2} \quad (15)$$

does not change under the given transformation. In the case of lattice QCD, this procedure replaces the original ill-conditioned matrix with one that has its low modes shifted away from zero by the factor μ^2 without changing the high modes significantly, while correcting for the change with an additional term in the Hamiltonian. The relation is exact and if, as it turns out to be, the second term is sub-dominant, i.e. leads to a smaller MD force than the first one, it can be included in the MD integration using a larger time-step, hence speeding up the overall simulation.

With, e.g. Wilson type fermions, adding the shift μ^2 to the matrix MM^\dagger in the determinant is easily achieved by adding a “twisted” mass μ to the fermion operator M . In our case this approach does not work because our fermion matrix lacks any “spinor” structure. However, we are free to add a non-zero staggered mass m_s , as shown in eq. (2). This leads to

$$\chi^\dagger (MM^\dagger)^{-1} \chi \mapsto \chi_1^\dagger (M_\mu M_\mu^\dagger)^{-1} \chi_1 + \chi_0^\dagger \left(\frac{MM^\dagger}{M_\mu M_\mu^\dagger} \right)^{-1} \chi_0 \quad (16)$$

with the notation

$$M = M(m_s = 0), \quad M_\mu = M(m_s = \mu). \quad (17)$$

The generalization to n masses μ_0, \dots, μ_{n-1} reads

$$\begin{aligned} \chi^\dagger (MM^\dagger)^{-1} \chi &\mapsto \\ \chi_n^\dagger (M_{\mu_{n-1}} M_{\mu_{n-1}}^\dagger)^{-1} \chi_n &+ \sum_{i=1}^{n-1} \chi_i^\dagger \left(\frac{M_{\mu_{i-1}} M_{\mu_{i-1}}^\dagger}{M_{\mu_i} M_{\mu_i}^\dagger} \right)^{-1} \chi_i + \chi_0^\dagger \left(\frac{MM^\dagger}{M_{\mu_0} M_{\mu_0}^\dagger} \right)^{-1} \chi_0. \end{aligned} \quad (18)$$

The χ_n -term can now be treated exactly as the standard pseudofermionic χ -term in the original algorithm by replacing M by $M_{\mu_{n-1}}$. In contrast the χ_i -terms with $i < n$ deserve additional treatment.

First of all the creation of χ_i now involves an inversion. After generating ρ_i from $e^{-\rho_i^\dagger \rho_i}$ one obtains χ_i from

$$\chi_i = M_{\mu_i}^{-1} M_{\mu_{i-1}} \rho = M_{\mu_i}^\dagger (M_{\mu_i} M_{\mu_i}^\dagger)^{-1} M_{\mu_{i-1}} \rho \quad \text{for } 0 < i < n, \quad (19)$$

$$\chi_0 = M_{\mu_0}^{-1} M \rho = M_{\mu_0}^\dagger (M_{\mu_0} M_{\mu_0}^\dagger)^{-1} M \rho. \quad (20)$$

In addition one has to consider the changes in eq. (11) where using $\frac{\partial M_{\mu_i}}{\partial \phi} = \frac{\partial M}{\partial \phi}$ leads to

$$\dot{\pi}^T = F_\phi + \sum_{i=0}^n F_{\chi_i} \quad (21)$$

$$F_\phi = -\delta\phi^T V^{-1} \quad (22)$$

$$F_{\chi_n} = 2 \operatorname{Re} \left(\eta_n^\dagger \frac{\partial M}{\partial \phi} \xi_n \right) \quad (23)$$

$$F_{\chi_i} = 2 \operatorname{Re} \left(\eta_i^\dagger \frac{\partial M}{\partial \phi} (\xi_i - \chi_i) \right) \quad \text{for } i < n \quad (24)$$

with η_n and ξ_n obtained like in eq. (12) with $M_{\mu_{n-1}}$ instead of M . The terms in eq. 24 are defined as follows:

$$\eta_i = (M_{\mu_{i-1}} M_{\mu_{i-1}}^\dagger)^{-1} M_{\mu_i} \chi_i \quad (25)$$

$$\xi_i = M_{\mu_{i-1}}^\dagger \eta_i = M_{\mu_{i-1}}^{-1} M_{\mu_i} \chi_i \quad (26)$$

Here $M_{\mu_{-1}} \equiv M$ has been used for simplicity. Note that in the case of $\mu_i = \mu_{i-1}$ this leads to $\xi_i = \chi_i$ and thus $F_{\chi_i} = 0$, leading to the reduced set of equations with one less pseudofermion field. It can be assumed that F_{χ_i} changes continuously in μ_i and μ_{i-1} (see below for more details). Therefore, this part of the force will be weak for small $|\mu_i - \mu_{i-1}|$. On the other hand M_{μ_i} is better conditioned than M because the eigenvalues are bound from below by μ_i . This leads to a splitting of the original force into $n+1$ terms of which some are easy to calculate and the other ones can be chosen small.

5.2. Different time scales

One can use this result in the same way as proposed in ref. [36]. The force F_{χ_0} can be treated with less accuracy by putting it on a different time scale. One can do this for each successive F_{χ_1} , F_{χ_2} , and so on. Let us for simplicity first consider the case with only one staggered mass term. We use the Omelyan integrator with the parameter $\zeta = 0.193$ for the calculation of the trajectory. Let $T(t)$ denote the evolution of the system by one trajectory of length t , $T_\pi(\Delta t)$ denote the evolution of ϕ over the time Δt , and $T_{\phi, \chi_0, \chi_1}(\Delta t)$ denote the evolution of π corresponding to the forces F_{ϕ, χ_0, χ_1} respectively. Then the naïve Omelyan integrator reads

$$T(\Delta t) = T_{\phi+\chi_0+\chi_1}(\zeta\Delta t) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_0+\chi_1}((1-2\zeta)\Delta t) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_0+\chi_1}(\zeta\Delta t) \quad (27)$$

and $T(t) = T(\Delta t)^{N_{\text{MD}}}$ where $N_{\text{MD}} = \frac{t}{\Delta t}$ is the number of time steps in the molecular dynamics. As we want to weight T_{χ_0} less we can choose a number N_0 that divides N_{MD} and write

$$T(N_0\Delta t) = T_{\chi_0}(N_0\zeta\Delta t) \cdot T_0(\Delta t)^{N_0/2} \cdot T_{\chi_0}(N_0(1-2\zeta)\Delta t) \cdot T_0(\Delta t)^{N_0/2} \cdot T_{\chi_0}(N_0\zeta\Delta t) \quad (28)$$

$$T_0(\Delta t) = T_{\phi+\chi_1}(\zeta\Delta t) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_1}((1-2\zeta)\Delta t) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_1}(\zeta\Delta t) \quad (29)$$

where in case of N_0 odd we define

$$T_0(\Delta t)^{1/2} = T_{\phi+\chi_1}(\zeta\Delta t) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_1}((1-2\zeta)\Delta t/2) \quad (30)$$

if the term stands before $T_{\chi_0}(N_0(1-2\zeta)\Delta t)$ and

$$T_0(\Delta t)^{1/2} = T_{\phi+\chi_1}((1-2\zeta)\Delta t/2) \cdot T_\pi(\Delta t/2) \cdot T_{\phi+\chi_1}(\zeta\Delta t) \quad (31)$$

if it comes later. This manipulation maintains the basic properties of the integrator: It remains symplectic and symmetric.

Now μ_0 has to be chosen in such a way that the calculation of F_{χ_1} is much cheaper than the calculation of F_{χ_0} but still the magnitude of F_{χ_0} is much less than of F_{χ_1} . In this case N_0 should be chosen as high as possible without significantly reducing the accuracy.

For an arbitrary number of Hasenbusch masses and time scales this scheme can be defined recursively in the same way as shown above. $T_{\phi+\chi_n}$ is the innermost part. It is all together repeated N_{MD} times per trajectory, this being N_{n-1} times as often as $T_{\chi_{n-1}}$, N_{n-2} times as often as $T_{\chi_{n-2}}$, and so on, up to N_0 times as often as T_{χ_0} . Here obviously N_{n-1} has to divide N_{n-2} , N_{n-2} has to divide N_{n-3} , ..., N_0 has to divide N_{MD} . Of course it is possible to choose $N_i = N_{i+1}$.

6. Scaling behaviour of the combined method

Using the methods described in the previous sections, we have performed simulations for different systems, one of which is shown here in greater detail: a (15,15)-chirality nanotube with 15 unit lengths and periodic boundary conditions applied at the ends, divided into $N_t = 256$ time slices. Such boundary conditions are defined using an orthogonal basis as depicted in fig. 3 of [27]. The coupling is $U/\kappa = 3.5$ and the inverse temperature $\beta\kappa = 8$.

6.1. Dependence of the solver time on m_s

First of all it is important to understand the scaling of the solver time depending on the staggered mass parameter m_s . It is known (e.g. from ref. [37]) that for the flexible Generalized Minimal Residual (GMRES) solver the time scales as

$$t = t_0 \frac{\lambda_{\max}}{\lambda_{\min}} \quad (32)$$

where λ_{\max} is the largest and λ_{\min} the smallest eigenvalue of the matrix that has to be inverted. As m_s shifts the complete eigenspectrum of the matrix this becomes

$$t(m_s) = t_0 \frac{\lambda_{\max} + m_s}{\lambda_{\min} + m_s}. \quad (33)$$

It turns out that even for our nested solvers this is a good approximation for $\lambda_{\min} < m_s < \lambda_{\max}$ as can be seen in figure 3. The depicted fit corresponds to $\lambda_{\max} = 4.1(2)$ and $\lambda_{\min} = 0.054(2)$. This functional form describes the data




Figure 3: Runtime of one inversion of $M_{m_s} M_{m_s}^\dagger$ depending on the staggered mass m_s .

sufficiently well. It has the advantage that it directly gives the highest and the lowest value of the spectrum of the original matrix MM^\dagger . These values can directly be used to find a good set of Hasenbusch masses.

The largest eigenvalue λ_{\max} can be approximated with the non-interacting limit for special cases (e.g. for graphene without chemical potential using the Hubbard model the largest eigenvalue known from the non-interacting limit shows very small deviations from the values with interaction). However it is not known in general and the purpose of the auto-tuner is much more general. We find that it works extremely well for graphene and carbon nanotubes without having to provide any additional information. We expect it to be applicable even when we add e.g. non-onsite interactions.

6.2. Dependence of F_{χ_i} on μ_i and μ_{i-1}

Let us, for simplicity, assume that we simulate with two staggered masses μ_0 and μ_1 . As any force term F_{χ_i} depends only on at most two masses, this consideration is general enough to understand all force terms except for $F_{\phi+\chi_n}$. From eqs. (24) and (26) follows

$$F_{\chi_1} \propto (M_{\mu_0}^{-1} M_{\mu_1} - 1) \chi_1 \quad (34)$$

$$= (M_{\mu_0}^{-1} M_{\mu_0+(\mu_1-\mu_0)} - 1) \chi_1 \quad (35)$$

$$\approx (M_{\mu_0}^{-1} (M_{\mu_0} + (\mu_1 - \mu_0)) - 1) \chi_1 \quad (36)$$

$$= M_{\mu_0}^{-1} (\mu_1 - \mu_0) \chi_1 \quad (37)$$

where, between eqs. (35) and (36), we assumed that the staggered mass term can be placed on the time diagonal. This approximation is valid as it produces the correct continuum limit. Thus for small differences $\mu_1 - \mu_0$ the norm of the force should be proportional to the absolute value of this difference:

$$|F_{\chi_1}| \propto |\mu_1 - \mu_0| \quad (38)$$

Higher order influences coming from the matrix products have not been considered here. To take them into account one first has to pay attention to the fact that the force should be symmetric under exchange of μ_0 and μ_1 . Empirically we have found that

$$|F_{\chi_1}| = a \frac{|\mu_1 - \mu_0|}{((b + \mu_0)(b + \mu_1))^c} \quad (39)$$

is a symmetric form obeying the proportionality for small differences and fitting the data sufficiently well (see fig. 4). It has been found that, independent of the system size and geometry, one always has $c \approx 0.4$. The parameter b is in the same order of magnitude as λ_{\min} found in the previous section. Thus it can be assumed that the factors $b + \mu_0$ and $b + \mu_1$ in the denominator correspond to the minimal eigenvalues of the matrix at staggered masses μ_0 and μ_1 respectively.




Figure 4: Force $|F_{\chi_1}|$ depending on the staggered mass $m_s = \mu_1$ for several fixed values of μ_0 . We find the fit parameters $a = 0.413(13)$, $b = 0.078(26)$ and $c = 0.380(38)$.

6.3. Dependence of $F_{\phi+\chi_n}$ on μ_{n-1}

As the remaining large force term depends only on the last mass μ_{n-1} we can again consider for simplicity only one staggered mass term m_s . Then one can make the ansatz

$$|F_{\phi+\chi_n}| \approx F_0 - |F_{\chi_{n-1}}| \quad (40)$$

for the functional form with some constant F_0 . However the numerical values of the coefficients have to be fitted independently of the results in the previous section:

$$|F_{\phi+\chi_n}| = F_0 - \tilde{a} \frac{m_s}{\left(\tilde{b} (\tilde{b} + m_s)\right)^{\tilde{c}}} \quad (41)$$

Figure 5 shows that again for $m_s > \lambda_{\min}$ the scaling of the force is described very well by the given functional form. One also sees as expected that the values are about an order of magnitude larger than the ones of $|F_{\chi_i}|$. For the parameters one gets again \tilde{b} (here $\tilde{b} = 0.0992(56)$) in the same order of magnitude as λ_{\min} . $\tilde{c} = 0.789(3)$ in this case.

7. Automatic tuning of the simulation parameters

In order to explore the whole parameter space, we need to run a large number of independent simulations, at least one for each combination of coupling U , temporal lattice spacing δ , number of time steps N_t , and system size L . Such a large set of independent simulations is hard to tune by hand, and therefore the ability to auto-tune these parameters is essential.




Figure 5: Force $|F_{\phi+\chi_n}|$ depending on the staggered mass m_s and fitted to (40).

7.1. Tuning the HMC step size

Our auto-tuner executes a series of successive simulations, each time continuing the simulation for a single MD trajectory, until some termination conditions (which the user supplies) are met. The length of the trajectory is fixed and has to be chosen by the user. It has been set to $t = 1$ here for reasons of simplicity yielding $\Delta t = \frac{1}{N_{\text{MD}}}$. There is one main loop consisting of two phases. In the first phase, the auto-tuner runs the simulator multiple times with the same parameters until the mean acceptance rate has a small confidence interval, making sure that the acceptance rate observation for the current N_{MD} is sufficiently accurate. In the second phase, the auto-tuner collects the observations from all previous iterations and tries to create a new model for the relationship between acceptance rate and N_{MD} . After the model is created, the N_{MD} for the next iteration is selected based on the models prediction for yielding a preset acceptance rate, e.g. $p_{\text{acc}} = 0.66$. The loop stops when the parameters have converged or the number of trajectories has reached a specified maximum limit.

First phase – minimizing the confidence interval. In the first phase, the auto-tuner tries to get an accurate measurement of the acceptance rate for the current N_{MD} . In order to do this, a number of samples of trajectories is simulated. Initially, a minimum number of trajectories is required. Next, the confidence interval is updated after each additional trajectory and the phase stops after this measurement gets smaller than a specified limit. The confidence interval is based on the variance of a normally distributed population and then the bounds are minimized if a portion of it lies outside the range between 0 and 1.

It is important to note that here we make the assumption that the acceptance rates for the samples of trajectories are normally distributed across the samples and, therefore, it may be subject for further refinement. In our experience the number of trajectories needed in this phase ranges from five to 30.

Second phase – creating a model. The design of the second phase was based on visual observations from the shapes of the relationship of p_{acc} versus N_{MD} . The resulting figures are reminiscent of the Sigmoid function, such as a logistic function, not only in shape, but also in the range of the dependent variable, which is between 0 and 1 in both cases. These are important pieces of information that we have provided to the model in order to make it more accurate and representative, even with a very small number of training data. For our purposes it seemed that the most appropriate solution was to fit the Cumulative Distribution Function (CDF) of the skew-normal distribution [38], which is not always symmetric as a Sigmoid function. The skew-normal distribution is a generalization of the normal distribution that allows non-zero skewness. This impacts the shape of its CDF accordingly, giving our model the ability to change the amount of steepness in the two ‘bend points’ of its ‘S’-like shape (see fig. 6 for an example for modelling real data).

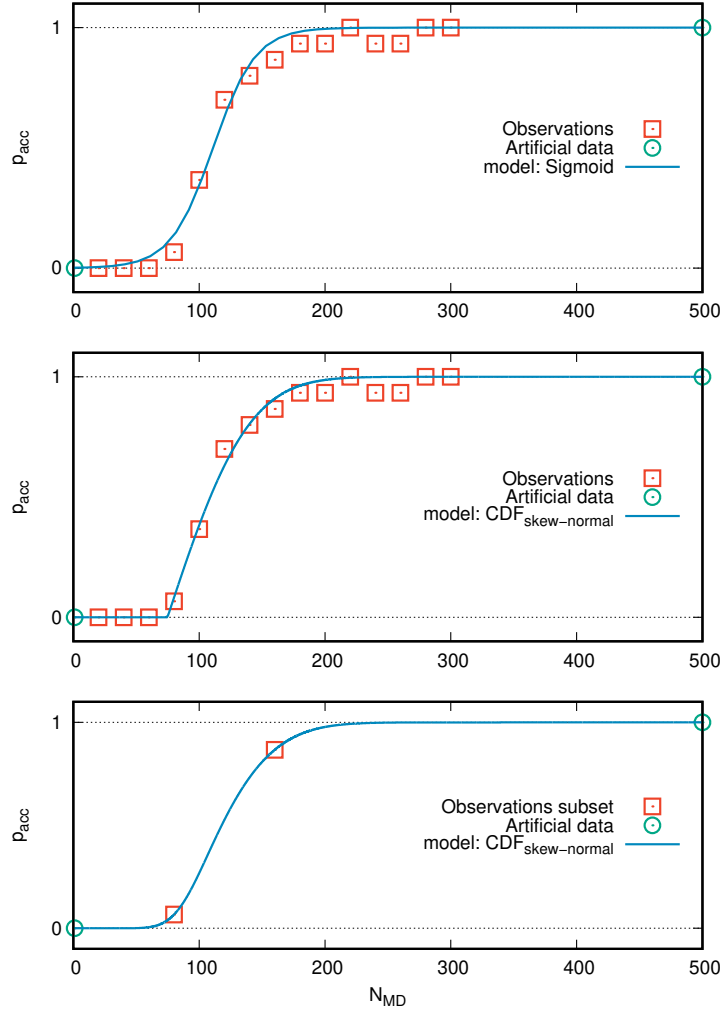


Figure 6: The model based on the CDF of the skew-normal distribution (middle) seems to fit the real observations better than a sigmoid function (top). By providing two extra artificial points, the model gives a reasonable shape even if a small subset of the measurements is available (bottom).

The auto-tuner is written in python and uses the numpy and scipy libraries [39], as well as the python version of the ‘sn’ package [40] for calculating the CDF of the skew-normal distribution for different parameters. The fitting of the model is done by calling the `scipy.optimize.least_squares` function of scipy, that implements the Levenberg–Marquardt algorithm [41] for minimizing a loss function, ‘soft_l1’ in this case, which is a smooth approximation of the L_1 norm condition for robustness [42] (see alg. B.1 for the pseudocode in Appendix B).

7.2. Tuning the Hasenbusch parameters

We now shift our aim to find a routine that optimizes the Hasenbusch parameters automatically. Using the auto-tuner of the previous section, we first thermalize the system with or without some fixed Hasenbusch parameters. After this we find appropriate staggered masses and time scales and, in the end, optimize N_{MD} for this set of parameters using the HMC auto-tuner once more.

Empirical data suggests a very broad minimum in parameter space (see fig. 7 for one staggered mass term). This means that the tuning does not have to be very accurate to achieve a speed that is nearly optimal. This is why we do not attempt to find the exact minimum of the CPU-time per trajectory. Instead we provide a simple procedure that can be automatized to find a set of parameters good enough for any practical purpose.




Figure 7: Runtime of one trajectory depending on the staggered mass m_s and the ratio between the time scales N_0 normalized by the runtime without Hasenbusch acceleration and weighted by the square root of the error $\Delta\mathcal{H}$ made by the calculation. The weighting compensates that all calculations have been performed with the same $N_{\text{MD}} = 540$. Values greater than 1 have been set to one 1 in this plot.

For the systems considered until now we did not see any significant improvements in using more than two Hasenbusch terms. Because of this presently our auto-tuner uses exactly two staggered masses, but it can be easily generalized to more masses should the need arise. Since the number of mass terms is two, the masses themselves have to be calculated. It seems reasonable to distribute them equidistantly in the spectrum, where equidistantly has to be understood

in a geometrical sense. For n masses this leads to

$$\mu_i = \lambda_{\min} \left(\frac{\lambda_{\max}}{\lambda_{\min}} \right)^{\frac{i+1}{n+1}} \quad (42)$$

with $\lambda_{\min, \max}$ gained from a fit of the solver time. The time scales then can be found knowing that we use a second order integrator and thus the error of the Hamiltonian $\Delta\mathcal{H}$ made in one trajectory scales approximately with

$$\Delta\mathcal{H} \propto \sum_i |F_i| N_i^2 \quad (43)$$

where i counts all forces introduced above. Because of symmetry this term is minimized near the point where all addends have the same value. Thus we get a first approximation for the time scale ratios via

$$N'_i = \sqrt{\frac{F_{\phi+\chi_n}(\mu_{n-1})}{F_{\chi_i}(\mu_{i-1}, \mu_i)}}. \quad (44)$$

For two masses one finds that $N'_0 \approx N'_1$ independently of the system as long as $c \approx 0.4$. As N_1 has to divide N_0 the only appropriate choice is $N_0 = N_1$. One finds that in most cases the results are best if one chooses

$$N_0 = N_1 = \left\lfloor \sqrt{N'_0 \cdot N'_1} \right\rfloor \quad (45)$$

instead of rounding.

In summary the combined auto-tuner works as follows:

1. Thermalize the system with some fixed Hasenbusch parameters (one staggered mass with $m_s = 0.5$, $N_0 = 1$ works quite good) and find a good value for N_{MD} .
2. Calculate solver time and forces for some different mass terms.
We calculate with $N_0 = N_1 = 1$ for the tuples $(m_0, m_1) \in \{(0.05, 0.2), (0.1, 0.3), (0.2, 1.5), (0.4, 0.8), (1.2, 2)\}$.
3. Fit the results according to the eqs. 33, 39, and 41.
4. Calculate the two Hasenbusch masses according to eq. 42 and the time scales according to 44 and 45.
5. Find the new optimum for N_{MD} .

8. Results

The data presented here was produced at $\beta\kappa = 8$ and $U/\kappa = 2.5$ on single nodes $2 \times \text{Intel(R) Xeon(R) CPU E5-2680}$, 40 MB Cache per node. The exact numbers deviate for other choices of parameters⁷ and on other computers but the qualitative behaviour is representative.

⁷Smaller values of U/κ lead to significantly faster calculations and reduce the advantage of the presented algorithms compared to a standard CG solver. Yet even at $U/\kappa = 0.5$ our optimized algorithm reaches a speedup of factor three to four. For $U/\kappa > 2.5$ we find only small deviations from the presented data. Explicitly $U/\kappa = 5$ has runtime and speedup deviations of less than 25 % from the $U/\kappa = 2.5$ case.

The auto-tuner as explained above has been implemented and was stable for all systems of carbon nanotubes as well as graphene tested up to this point. All the results (with and without Hasenbusch acceleration) presented in this section have been tuned in this way. Figures 8 and 9 visualize the speedup depending on the system size, achieved for one trajectory by the introduction of the fGMRES solver and additional Hasenbusch acceleration respectively. The lengths to which the trajectories have been tuned can be found in table 1.




Figure 8: Top: Time needed per trajectory by the different algorithms for graphene sheets with $(L \times L)$ unit cells and $N_t = 64$. Bottom: Speedup over the standard HMC algorithm (using double precision CG) for graphene sheets with $(L \times L)$ unit cells and $N_t = 64$. The double precision CG solver exceeds L3 capacity at $L = 51$, whereas the spCG exceeds L3 capacity at $L = 72$. Our best algorithm is highlighted in bold.




Figure 9: Top: Time needed per trajectory by the different algorithms for graphene sheets with (15×15) unit cells depending on N_t . Bottom: Speedup over the standard HMC algorithm (using double precision CG) for graphene sheets with (15×15) unit cells depending on N_t . Our best algorithm is highlighted in bold.

We analysed the runtime scaling with the system size. One finds a monomial dependence of the time on L and N_t in some regions but not globally. The dependence in the spatial size L (see fig. 8) is comparable for all algorithms but the algorithm including Hasenbusch acceleration and the new fGMRES solver is

| L | N_t | CG-solver | fGMRES-solver | fGMRES-solver, Hasenbusch | | | |
|-----|-------|-----------------|-----------------|---------------------------|---------------------|---------|---------|
| | | N_{MD} | N_{MD} | N_{MD} | N_{MD}/N_0 | μ_0 | μ_1 |
| 12 | 64 | 34 | 35 | 18 | 9 | 0.51 | 1.62 |
| 15 | 32 | 12 | 12 | 8 | 4 | 0.42 | 0.73 |
| 15 | 64 | 37 | 35 | 18 | 9 | 0.39 | 1.18 |
| 15 | 96 | 72 | 73 | 26 | 13 | 0.33 | 1.37 |
| 15 | 128 | 126 | 124 | 44 | 22 | 0.51 | 2.36 |
| 15 | 192 | 269 | 263 | 70 | 35 | 0.50 | 2.27 |
| 15 | 256 | 464 | 474 | 98 | 49 | 0.41 | 2.09 |
| 21 | 64 | 38 | 39 | 22 | 11 | 0.54 | 1.84 |
| 30 | 64 | 41 | 43 | 20 | 10 | 0.41 | 1.12 |
| 39 | 64 | 48 | 49 | 30 | 10 | 0.33 | 0.74 |
| 48 | 64 | 45 | 47 | 30 | 10 | 0.33 | 0.70 |
| 66 | 64 | 54 | 61 | 24 | 12 | 0.41 | 0.94 |
| 90 | 64 | 64 | 68 | 26 | 13 | 0.39 | 0.99 |
| 102 | 64 | 66 | 63 | 30 | 15 | 0.41 | 1.10 |

Table 1: Lengths of the auto-tuned trajectories for the different algorithms. In the case of the Hasenbusch-algorithm in addition the two staggered masses and their time scale are given. The third column from the right gives the length of the trajectory on the most coarse time scale.

much faster than the algorithm without Hasenbusch acceleration but with mixed precision solves, which in turn is significantly faster than the algorithm without Hasenbusch acceleration using only the double precision CG solver. In addition one encounters a jump in the runtime between $L = 48$ and $L = 66$ for the original algorithm and between $L = 66$ and $L = 90$ for the algorithms using the spCG preconditioner. At these points the memory needed by the solver exceeds the L3-cache of 40 MB. After this threshold the scaling of the runtime becomes stronger and the ratio between the single precision and the double precision algorithms increases dramatically. It is expected that this advantage is going to decrease as the system size grows even larger because L3-cache is not going to play an important role any more neither for double nor single precision solver. We see this behaviour already at $L = 102$, so we are not going to simulate yet larger systems on single node in recent future.

The L3-threshold is not reached in figure 9, but here one can see that the scaling of the runtime with N_t is much stronger without Hasenbusch acceleration. Not only the runtime but also the tuning takes much longer without Hasenbusch acceleration regardless of the much lower amount of trajectories that have to be calculated. This is why we can present results for systems as large as $N_t = 512$ with Hasenbusch acceleration but only $N_t = 256$ without. Our experience up to this point leads us to conclude that the continuum limit ($N_t \rightarrow \infty$) is not only more expensive but also much more crucial for physical results than the infinite volume limit ($L \rightarrow \infty$), making this advantage of the Hasenbusch algorithm rather important.




Figure 10: The two correlators (labelled M and Γ) calculated from a 4-site model with $U/\kappa = 4$ and $\beta\kappa = 8$. The number of timesteps N_t is given in the figure. Uncertainties are too small to be seen in many cases. The black lines are the exact results.

8.1. Benchmarking to the 4-site problem

To ensure that our implementation of the auto-tuner and the Hasenbusch mass shifts produce correct physics, we compute the allowed correlators for the 4-site model with nearest neighbours drawn from a cubic geometry (i.e. hexagonal lattice), and compare to the corresponding correlators obtained from direct diagonalization of the Hubbard Hamiltonian. We simulate with a Hubbard ratio $U/\kappa = 4$, $\beta\kappa = 8$, and $N_t = 64, 96, 128$, and 160 . Figure 10 shows our calculations, excluding the $N_t = 160$ results to reduce clutter. Here the expected convergence of the correlators as we increase N_t is observed.

We can extract the interacting quasi-particle eigenenergies of this system by fitting an exponential to these correlators, as the expected behaviour of these correlators is $C(\tau) \sim e^{-E\tau}$ for large τ . In practice, we fitted exponentials within a given time window. In particular, E_Γ was fitted to an exponential in the time window $\tau\kappa \in [.2, 1]$, whereas E_M was fitted in $\tau\kappa \in [.6, 2]$. In Table 2 we show these extracted energies as well as the exact results from direct diagonalization. As these results were only for benchmarking, we did not optimize our fit windows when performing our fits. Still, we find excellent agreement with the exact results for $N_t = 128$ and higher.

9. Conclusion

In this paper we investigated different attempts at accelerating lattice Monte Carlo simulations of the Hubbard model on a hexagonal lattice. We implemented higher order symplectic integrators in our MD integrations, and introduced mixed-precision CG solvers as a ‘preconditioner’. In the former case we

Table 2: Extracted energies from correlators of the 4-site model, as a function of N_t . E_Γ was extracted by fitting an exponential to the correlator in the time window $\tau\kappa \in [.2, 1]$. E_M was extracted from exponential fits in window $\tau\kappa \in [.6, 2]$. Uncertainties were obtained from a bootstrap ensemble of fits. The last column gives the exact results for direct diagonalization of the Hubbard Hamiltonian.

| N_t | 64 | 96 | 128 | 160 | exact |
|-------------------|-----------|-----------|-----------|-----------|-------|
| E_Γ/κ | 3.555(16) | 3.501(11) | 3.445(12) | 3.447(14) | 3.429 |
| E_M/κ | 1.647(23) | 1.629(21) | 1.542(18) | 1.567(22) | 1.565 |

found the Omelyan integrator to be competitive with the force gradient method, while in the latter the mixed precision preconditioner gave a speedup of approximately 1.9. In large systems, such as 66×66 or larger, the introduction of the fGMRES solver also provided a factor of ~ 2.2 speedup.

Our greatest speedup came with the implementation of Hasenbusch acceleration via mass shifts. With just two Hasenbusch masses, and in combination with the fGMRES solver, we were able to get a speedup of well over an order of magnitude for lattices of size 90×90 and higher.

The numerous parameters available in our simulations required the use of an auto-tuner to find their optimal values. By establishing scaling relations between the Hasenbusch masses and forces, as well as between the acceptance rate and number of MD steps, we devised an efficient and stable procedure for optimizing the parameter sets in our lattice MC simulations to a target 66% acceptance rate.

To ensure that our implementation of the Hasenbusch acceleration and auto-tuner performed correctly, we simulated the 4-site system and compared our results to direct diagonalization, finding excellent agreement.

Our increased speedup now allows us to simulate hexagonal lattices of unprecedented sizes. With only single-node resources, we were able to simulate up to 102×102 lattices with $N_t = 64$. Results with even larger spatial sizes and N_t are in preparation. If optimized over multi-nodes, we anticipate the ability to simulate on systems an order of magnitude in size larger, thereby approaching system sizes comparable to physical applications. We are currently developing this capability.

Acknowledgements

We thank E. Berkowitz, C. Körber, T. Lähde, and J.-L. Wynen for their insightful discussions related to this work. We also thank M. Ulybyshev for his comments. This work is supported in part by the DFG grants SFB TRR-55 and SFB TRR-110 (Sino German CRC). The simulations were carried out using computing resources which received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under Grant Agreement n° 287530 (DEEP) and 610476 (DEEP-ER), and from the Horizon 2020 Programme (H2020-FETHPC) under Grant Agreement n° 754304 (DEEP-EST). P. Papaphilippou was supported by the Summer of HPC programme of the

PRACE-5IP project (2017-2019, RI-730913) carried under EC Research and Innovation Action of the H2020 Programme. The present publication reflects only the authors' views. The European Commission is not liable for any use that might be made of the information contained therein.

Appendix A. Implementation of the Flexible Generalized Minimal Residual solver

The algorithm of the flexible GMRES from [31] with the spCG as preconditioner in our implementation can be found in algorithm A.1 with some relative

```

input : TOL,  $b$ ,  $A$ 
output:  $x$ 
1 solve  $A \cdot x = b$  for  $x$  using spCG;
2  $r = b - A \cdot x$ ;
3  $ERR = |r|$ ;
4  $w = r$ ;
5 while  $ERR > TOL \cdot |b|$  do
6   for  $j = 0, \dots, m - 1$  do
7     if  $ERR < TOL \cdot |b|$  then
8       break;
9     end
10     $v_j = w / |w|$ ;
11    if  $TOL \cdot a \cdot |b| / ERR < 1$  then
12      solve  $A \cdot z_j = v_j$  for  $z_j$  using spCG with  $TOL = TOL \cdot a \cdot |b| / ERR$ ;
13    end
14    else
15       $z_j = v_j$ ;
16    end
17     $w = A \cdot z_j$ ;
18     $h_{ij} = w \cdot v_i$ ;
19     $w = w - \text{sum}(h_{ij} \cdot v_i, i = 0, \dots, j)$ ;
20     $h_{j+1,j} = |w|$ ;
21     $ERR = ERR \cdot |w|$ ;
22  end
23   $y = \text{argmin}(|r|, 0, \dots, 0) - h \cdot y$ ;
24   $x = x + \text{sum}(y_i \cdot z_i, i = 0, \dots, j)$ ;
25   $r = b - A \cdot x$ ;
26   $ERR = |r|$ ;
27   $w = r$ ;
28 end

```

Algorithm A.1: Flexible GMRES

tolerance TOL, the complex vector $b \in \mathbb{C}^n$ and the matrix $A \in \mathbb{C}^{n \times n}$ as input

parameters and $x, r, w \in \mathbb{C}^n$, $v, z \in \mathbb{C}^{m \times n}$, $h \in \mathbb{R}^{m+1 \times m}$, $y \in \mathbb{R}^m$. Here $|\cdot|$ denotes the euclidean norm. The minimization in line 23 has been performed by solving $h^T h y = |r| \cdot h^T e_1$ for y using a Cholesky decomposition. $m = 10$ has been chosen as the restart parameter and it showed to be good to set $a = 5$ in the tolerance of the spCG preconditioner.

Appendix B. Tuning the HMC step size

In alg. B.1 we can see the pseudocode for the auto-tuner of the HMC step size. It can be seen as a single routine that takes a parameter file for all parameters of the simulator, ignores the provided $N_{MD} = \frac{1}{\Delta t}$, runs a number of simulations with automatically set N_{MD} values and stops when the final N_{MD} value is expected to continue yielding a mean trajectory acceptance rate of 0.66.

```

Result: The  $N_{MD}$  for which the acceptance rate is close to 0.66
1  $N_{MD} \leftarrow 500$ ;
2 initialize or resume from a previous experiment;
3 /* Main loop */
4 while True do
5   /* 1st phase - minimizing the confidence interval */
6   trajectories  $\leftarrow 0$ ;
7   while conf. interval  $> 0.25$  or trajectories  $\leq 5$  do
8     pacc  $\leftarrow$  simulate( $N_{MD}$ );
9     trajectories  $\leftarrow$  trajectories + 1;
10  end
11  Append  $N_{MD}$  to  $N_{MD}list$ ;
12  Append mean pacc to pacclist;
13  /* 2nd phase - building a model */
14  training_set  $\leftarrow$  ( $N_{MD}list$ , pacclist);
15  fit_function  $\leftarrow f(x) = CDF_{skew-normal}(\beta_1 x + \beta_0, \alpha) + \epsilon$ ;
16  model.build(training_set, fit_function);
17   $N_{MD} \leftarrow$  model.predict_ $N_{MD}$ .from.pacc(0.66);
18  /* Stopping conditions - convergence heuristic */
19  if count total trajectories  $\geq 500$  or  $N_{MD}list.count(N_{MD}) \geq 3$  or
    times 0.66 was in conf. interval  $\geq 3$  then
20    break;
21  end
22 end

```

Algorithm B.1: High-level description of the HMC step size auto-tuner

References

- [1] E. H. Lieb, F. Y. Wu, Absence of Mott transition in an exact solution of the short-range, one-band model in one dimension, Phys. Rev. Lett. 20

- (1968) 1445–1448, [Erratum: Phys. Rev. Lett. 21, 192 (1968)]. doi:10.1103/PhysRevLett.21.192.2, 10.1103/PhysRevLett.20.1445.
- [2] A. Giuliani, V. Mastropietro, The two-dimensional hubbard model on the honeycomb lattice, Communications in Mathematical Physics 293 (2) (2009) 301. doi:10.1007/s00220-009-0910-5.
URL <https://doi.org/10.1007/s00220-009-0910-5>
- [3] S. Arya, P. V. Sriluckshmy, S. R. Hassan, A.-M. S. Tremblay, Antiferromagnetism in the hubbard model on the honeycomb lattice: A two-particle self-consistent study, Phys. Rev. B 92 (2015) 045111. doi:10.1103/PhysRevB.92.045111.
URL <https://link.aps.org/doi/10.1103/PhysRevB.92.045111>
- [4] N. Furukawa, Antiferromagnetism of the hubbard model on a layered honeycomb lattice is mgb2 a nearly-antiferromagnetic metal? , Journal of the Physical Society of Japan 70 (6) (2001) 1483–1486. arXiv:<https://doi.org/10.1143/JPSJ.70.1483>, doi:10.1143/JPSJ.70.1483.
URL <https://doi.org/10.1143/JPSJ.70.1483>
- [5] A. Georges, G. Kotliar, W. Krauth, M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, Rev. Mod. Phys. 68 (1996) 13–125. doi:10.1103/RevModPhys.68.13.
URL <https://link.aps.org/doi/10.1103/RevModPhys.68.13>
- [6] K. Fang, G. Fernando, A. Balatsky, A. Kocharian, Possible phase separation in square and honeycomb hubbard model: A variational cluster study, Physics Letters A 379 (37) (2015) 2230 – 2238. doi:<https://doi.org/10.1016/j.physleta.2015.06.059>.
URL <http://www.sciencedirect.com/science/article/pii/S0375960115005782>
- [7] C. Honerkamp, Density waves and cooper pairing on the honeycomb lattice, Phys. Rev. Lett. 100 (2008) 146404. doi:10.1103/PhysRevLett.100.146404.
URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.146404>
- [8] S. Raghu, X.-L. Qi, C. Honerkamp, S.-C. Zhang, Topological mott insulators, Phys. Rev. Lett. 100 (2008) 156401. doi:10.1103/PhysRevLett.100.156401.
URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.156401>
- [9] S. Sorella, Y. Otsuka, S. Yunoki, Absence of a Spin Liquid Phase in the Hubbard Model on the Honeycomb Lattice, Scientific Reports 2 (2012) 992.
URL <http://dx.doi.org/10.1038/srep00992><https://www.nature.com/articles/srep00992#supplementary-information>

- [10] P. V. Buividovich, M. V. Ulybyshev, Applications of lattice QCD techniques for condensed matter systems, *Int. J. Mod. Phys. A* **31** (22) (2016) 1643008. [arXiv:1602.08431](#), [doi:10.1142/S0217751X16430089](#).
- [11] S. Beyl, F. Goth, F. F. Assaad, Revisiting the Hybrid Quantum Monte Carlo Method for Hubbard and Electron-Phonon Models, *Phys. Rev. B* **97** (8) (2018) 085144. [arXiv:1708.03661](#), [doi:10.1103/PhysRevB.97.085144](#).
- [12] H.-F. Lin, H.-D. Liu, H.-S. Tao, W.-M. Liu, Phase transitions of the ionic Hubbard model on the honeycomb lattice, *Scientific Reports* **5** (2015) 9810.
URL <http://dx.doi.org/10.1038/srep09810><http://10.0.4.14/srep09810>
- [13] S. Duane, A. D. Kennedy, B. J. Pendleton, D. Roweth, Hybrid Monte Carlo, *Phys. Lett. B* **195** (1987) 216–222. [doi:10.1016/0370-2693\(87\)91197-X](#).
- [14] W. Armour, S. Hands, C. Strouthos, Monte Carlo Simulation of the Semimetal-Insulator Phase Transition in Monolayer Graphene, *Phys. Rev. B* **81** (2010) 125105. [arXiv:0910.5646](#), [doi:10.1103/PhysRevB.81.125105](#).
- [15] J. E. Drut, T. A. Lahde, Lattice field theory simulations of graphene, *Phys. Rev. B* **79** (2009) 165425. [arXiv:0901.0584](#), [doi:10.1103/PhysRevB.79.165425](#).
- [16] D. Smith, L. von Smekal, Monte-Carlo simulation of the tight-binding model of graphene with partially screened Coulomb interactions, *Phys. Rev. B* **89** (19) (2014) 195429. [arXiv:1403.3620](#), [doi:10.1103/PhysRevB.89.195429](#).
- [17] T. Luu, T. A. Lähde, Quantum Monte Carlo calculations for carbon nanotubes, *Phys. Rev. B* **93** (2016) 155106. [doi:10.1103/PhysRevB.93.155106](#).
URL <https://link.aps.org/doi/10.1103/PhysRevB.93.155106>
- [18] R. Blankenbecler, D. J. Scalapino, R. L. Sugar, Monte Carlo calculations of coupled boson-fermion systems. I, *Phys. Rev. D* **24** (1981) 2278–2286. [doi:10.1103/PhysRevD.24.2278](#).
URL <https://link.aps.org/doi/10.1103/PhysRevD.24.2278>
- [19] Z. Y. Meng, T. C. Lang, S. Wessel, F. F. Assaad, A. Muramatsu, Quantum spin liquid emerging in two-dimensional correlated Dirac fermions, *Nature* **464** (7290) (2010) 847–851.
URL <http://dx.doi.org/10.1038/nature08942>
- [20] Y. Otsuka, S. Yunoki, S. Sorella, Universal quantum criticality in the metal-insulator transition of two-dimensional interacting Dirac electrons, *Phys.*

- Rev. X6 (1) (2016) 011029. [arXiv:1510.08593](#), [doi:10.1103/PhysRevX.6.011029](#).
- [21] T. Stauber, P. Parida, M. Trushin, M. Ulybyshev, D. Boyda, J. Schliemann, Interacting Electrons in Graphene: Fermi Velocity Renormalization and Optical Response, *Phys. Rev. Lett.* 118 (26) (2017) 266801. [arXiv:1704.03747](#), [doi:10.1103/PhysRevLett.118.266801](#).
 - [22] R. C. Brower, C. Rebbi, D. Schaich, Hybrid Monte Carlo Simulation of Graphene on the Hexagonal Lattice [arXiv:1101.5131](#).
 - [23] M. Ulybyshev, N. Kintscher, K. Kahl, P. Buividovich, Schur complement solver for Quantum Monte-Carlo simulations of strongly interacting fermions [arXiv:1803.05478](#).
 - [24] M. Hasenbusch, Speeding up the hybrid Monte Carlo algorithm for dynamical fermions, *Physics Letters B* 519 (1) (2001) 177 – 182. [doi:https://doi.org/10.1016/S0370-2693\(01\)01102-9](#).
URL <http://www.sciencedirect.com/science/article/pii/S0370269301011029>
 - [25] M. V. Ulybyshev, S. N. Valgushev, Path integral representation for the Hubbard model with reduced number of Lefschetz thimbles [arXiv:1712.02188](#).
 - [26] D. Smith, L. von Smekal, Monte Carlo simulation of the tight-binding model of graphene with partially screened Coulomb interactions, *Phys. Rev. B* 89 (2014) 195429. [doi:10.1103/PhysRevB.89.195429](#).
URL <https://link.aps.org/doi/10.1103/PhysRevB.89.195429>
 - [27] D. Chakrabarti, S. Hands, A. Rago, Topological aspects of fermions on a honeycomb lattice, *Journal of High Energy Physics* 2009 (06) (2009) 060.
URL <http://stacks.iop.org/1126-6708/2009/i=06/a=060>
 - [28] R. Saito, G. Dresselhaus, M. S. Dresselhaus, *Physical Properties of Carbon Nanotubes*, World Scientific Publishing, 1998, ISBN 978-1-86094-093-4 (hb) ISBN 978-1-86094-223-5 (pb).
 - [29] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, H. V. der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM, Philadelphia, PA, 1993.
 - [30] Y. Notay, Flexible conjugate gradients, *SIAM J. Sci. Comput* 22 (2000) 1444–1460.
 - [31] Y. Saad, A flexible Inner-Outer preconditioned GMRES algorithm, *SIAM Journal on Scientific Computing* 14 (1993) 461–469.

- [32] I. P. Omelyan, I. M. Mryglod, R. Folk, Optimized Verlet-like algorithms for molecular dynamics simulations, *Phys. Rev. E* 65 (2002) 056706. doi: 10.1103/PhysRevE.65.056706.
URL <https://link.aps.org/doi/10.1103/PhysRevE.65.056706>
- [33] I. Omelyan, I. Mryglod, R. Folk, Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations, *Computer Physics Communications* 151 (3) (2003) 272 – 314. doi:[https://doi.org/10.1016/S0010-4655\(02\)00754-3](https://doi.org/10.1016/S0010-4655(02)00754-3).
URL <http://www.sciencedirect.com/science/article/pii/S0010465502007543>
- [34] H. Yin, R. D. Mawhinney, Improving DWF Simulations: the Force Gradient Integrator and the Möbius Accelerated DWF Solver, in: *High Energy Physics - Lattice*, 2011.
- [35] E. Forest, R. D. Ruth, Fourth-order symplectic integration, *Physica D Nonlinear Phenomena* 43 (1990) 105–117. doi:10.1016/0167-2789(90)90019-L.
- [36] C. Urbach, K. Jansen, A. Shindler, U. Wenger, HMC algorithm with multiple time scale integration and mass preconditioning, *Computer Physics Communications* 174 (2) (2006) 87 – 98. doi:<https://doi.org/10.1016/j.cpc.2005.08.006>.
URL <http://www.sciencedirect.com/science/article/pii/S0010465505004753>
- [37] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. SCI. STAT. COMPUT.* 7 (3) (1986) 856–869.
- [38] A. Azzalini, A class of distributions which includes the normal ones, *Scandinavian journal of statistics* (1985) 171–178.
- [39] S. v. d. Walt, S. C. Colbert, G. Varoquaux, The NumPy array: a structure for efficient numerical computation, *Computing in Science & Engineering* 13 (2) (2011) 22–30.
- [40] A. Azzalini, J. van Dijk, Python routines – The Skew-Normal Probability Distribution , <http://azzalini.stat.unipd.it/SN>, online; accessed 18-February-2018 (2013).
- [41] J. J. Moré, The Levenberg-Marquardt algorithm: implementation and theory, in: *Numerical analysis*, Springer, 1978, pp. 105–116.
- [42] B. Triggs, P. F. McLauchlan, R. I. Hartley, A. W. Fitzgibbon, Bundle adjustment a modern synthesis, in: *International workshop on vision algorithms*, Springer, 1999, pp. 298–372.