

Achieving fast high-fidelity optimal control of many-body quantum dynamics

Jesper Hasseriis Mohr Jensen ^{1,*}, Frederik Skovbo Møller,^{1,2} Jens Jakob Sørensen,¹ and Jacob Friis Sherson^{1,†}

¹*Department of Physics and Astronomy, Aarhus University, Ny Munkegade 120, 8000 Aarhus C, Denmark*

²*Vienna Center for Quantum Science and Technology, Atominstytut, TU Wien, Stadionallee 2, 1020 Vienna, Austria*



(Received 4 September 2020; accepted 27 October 2021; published 16 November 2021)

We demonstrate the efficiency of a recent exact-gradient optimal control methodology by applying it to a challenging many-body problem, crossing the superfluid to Mott-insulator phase transition in the Bose-Hubbard model. The system size necessitates a matrix product state representation and this seamlessly integrates with the requirements of the algorithm. We observe fidelities in the range 0.99–0.9999 with associated minimal process duration estimates displaying an exponential fidelity-duration trade-off across several orders of magnitude. The corresponding optimal solutions are characterized in terms of a predominantly linear sweep across the critical point followed by bang-bang-like structure. This is quite different from the smooth and monotonic solutions identified by earlier gradient-free optimizations which are hampered in locating the higher complexity protocols in the regime of high fidelities at low process durations. Overall, the comparison suggests significant methodological improvements also for many-body systems in the ideal open-loop setting. Acknowledging that idealized open-loop control may deteriorate in actual experiments, we discuss the merits of using such an approach in combination with closed-loop control—in particular, high-fidelity physical insights extracted with the former can be used to formulate practical, low-dimensional search spaces for the latter.

DOI: [10.1103/PhysRevA.104.052210](https://doi.org/10.1103/PhysRevA.104.052210)

I. INTRODUCTION

With experimental and theoretical advances in the preparation and engineering of quantum mechanical systems, precise manipulation of fragile quantum systems has become increasingly important [1]. To this end, *quantum optimal control* [2] is a particularly successful tool for designing controls that implement desired physical transformations with wide applications in numerous research areas, such as superconducting qubits [3–7], nuclear magnetic resonance systems [8–12], nitrogen vacancy centers [13–16], cold molecules [17–20], and cold atoms [21–28], to name a few. On the theoretical side, open-loop optimal control design can be considered the union of numerical simulation and optimization methodologies. Therefore, with growing Hilbert space sizes such as in many-body contexts [29–39], the performance capacity of both these components must be streamlined to succeed in finite time within this paradigm.

Simulating very high-dimensional many-body systems exactly requires an exponential amount of memory and computation time [40]. Current techniques for breaking this *curse of dimensionality* finds recourse in tensor networks ansätze and, in the case of 1D systems, the appropriate structures are *matrix product states* (MPS). The entanglement entropy for these systems exhibits a constant area scaling law with the number of particles and can therefore be effectively simulated classically with subexponential resources [6,40–43].

Among the significant number of nonlinear optimization methodologies [9,12,21,22,44–55], it is conventional understanding [49,50,56] that derivative-based local search

techniques is the best way to identify local optima for any given optimization objective J . Coupled with an appropriate multistarting scheme [57], global optima may be uncovered. In addition to the local gradient ∇J , it is well-known that also incorporating curvature information through the local Hessian $\nabla^2 J$ is essential for the convergence rates of local optimization. The theoretical ideal is to include both exactly, but in practice the full Hessian is often prohibitively expensive to calculate. It is instead standard practice to apply, e.g., Broyden-Fletcher-Goldfarb-Shanno (BFGS) schemes to progressively construct a Hessian approximation from successive gradients [56]. If the individual gradients are inexact, these errors will accumulate in each iteration and make the Hessian approximation unreliable as the optimization progresses. This will in turn manifest as increasingly poor update directions that ultimately lead the optimization irreparably astray and significantly hamper the discovery of local optima as discussed in, e.g., Refs. [44,58]. Hence, exact gradients are central to local search paradigms. Applying these in very high-dimensional Hilbert spaces, however, has been a major obstacle because the currently known exact gradient element calculations become “extremely resource consuming, if not impossible” [21] due to their unfavorable scaling with the Hilbert space dimensionality $D_{\mathcal{H}}$. Specifically, their evaluation requires either full diagonalization [59] or exponentiation [48] of matrices that are at least of size $D_{\mathcal{H}} \times D_{\mathcal{H}}$ potentially followed by a recursive commutator series summation [44,58]. Therefore, control problems in, e.g., the many-body limit have either been approached using derivative-free methods [21,22] or, contemporarily to this work, inexact first-order gradients [60] which are prone to the reduced overall optimization capabilities described above.

In this paper, we apply a new exact derivative methodology, introduced in our recent parallel work [58] and reviewed in

*jhasseriis@phys.au.dk

†sherson@mgmt.au.dk

Sec. II, to a paradigmatic problem in the complex many-body regime. Briefly, the methodology circumvents the computational bottlenecks mentioned above. This is achieved by (1) applying Trotterization schemes and (2) representing the problem in a diagonal basis for the control Hamiltonian. Explicitly including these effects in the analytical derivations leads to very simple forms for the gradient depending on the specific Trotter scheme. The expressions are exact with only a first-order term as all higher-order correction terms vanish and the only computational effort lies in time evolving an auxiliary state in addition to the usual quantum state. That is, these exact gradient calculations scale only with the time it takes to solve the quantum dynamics which is the fundamental operation of any numerical optimal control algorithm. It is of particular convenience that both (1) and (2) are by themselves common and independently implemented in the general context of numerical simulations. For example, a large family of algorithms for performing time evolution of the MPS considered here rely on Trotterization [61]. Our exact derivative methodology thus meshes quite naturally with established standard numerical techniques.

The central ideas and derivations in Ref. [58] are valid for all unitary control tasks. For concreteness, the special case of maximizing state transfer fidelity in Hilbert spaces of dimension $D_{\mathcal{H}} = 2 - 100$ was considered there. For fixed $D_{\mathcal{H}}$, we demonstrated orders of magnitude improvement in terms of computational speedup relative to similar exact gradients and that the gained relative speedup is exponential as a function of $D_{\mathcal{H}}$. It was also demonstrated that inexact gradients lead to the expected poor performance in terms of achievable fidelity.

Here we demonstrate that the techniques remain viable for fidelity requirements above 0.99 for many-body systems with $D_{\mathcal{H}} \sim 10^{11}$ where other currently known exact gradient approaches are computationally prohibited. This evidences the possibility for significantly enhanced optimal open-loop control also over very high-dimensional Hilbert spaces operating in the fast high-fidelity regime.

Having showcased an advance in open-loop control capabilities, we finally turn to the broader context of quantum optimal control. Outlining first the benefits of open- and closed-loop methodologies separately, we give our perspective on how the role played by the open-loop component in a unified-loop may remain useful in the future.

II. MANY-BODY STATE TRANSFER

A. Bose-Hubbard model

As a challenging representative example from the class of many-body problems, we examine the superfluid-Mott insulator phase transition in the one-dimensional Bose-Hubbard model. This model describes the physics of N_p interacting spinless bosons in a lattice with N_s sites by the Hamiltonian [29,62,63]

$$\hat{H}_{\text{SI}} = \hat{H}^{J_x} + \hat{H}^U = J_x \sum_{i=1}^{N_s-1} \hat{h}_{[i,i+1]}^{J_x} + \frac{U}{2} \sum_{i=1}^{N_s} \hat{h}_{[i]}^U, \quad (1)$$

$$\hat{h}_{[i,i+1]}^{J_x} = -(\hat{a}_{i+1}^\dagger \hat{a}_i + \text{H.c.}), \quad \hat{h}_{[i]}^U = \hat{n}_i(\hat{n}_i - 1). \quad (2)$$

The operators \hat{a}_i^\dagger and \hat{a}_i are the bosonic creation and annihilation operators for site i , respectively, while

$\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ counts the number of particles occupying the site, $\hat{n}_i |n_1, \dots, n_i, \dots, n_{N_s}\rangle = n_i |n_1, \dots, n_i, \dots, n_{N_s}\rangle$. We assume a fixed number of particles $N_p = \sum_{i=1}^{N_s} n_i$ and unit filling $N_s/N_p = 1$.

The energy J_x is associated with the hopping (tunneling) operator $\hat{h}_{[i,i+1]}^{J_x}$ along the x direction, and U is the energy associated with the on-site interaction operator $\hat{h}_{[i]}^U$. The ratio U/J_x characterizes the quantum phase of the system. We associate with $J_x \gg U$ the superfluid phase in which the ground state |SF) is a delocalized particle distribution across the lattice with sizable site occupation variance. The Mott insulator phase is conversely associated with $U \gg J_x$ and the ground state |Mott) = $|1, 1, \dots, 1\rangle$ is a single Fock state component with unit occupancy on each site in the thermodynamic limit.

We seek to dynamically connect the ground states, |SF) \rightarrow |Mott), on opposing sides of the critical point by controlling the time-dependent ratio of on-site interaction- and tunneling energies $u(t) = U(t)/J_x(t)$. In particular, we are interested in estimating the lowest possible transfer duration T consistent with high-fidelity requirements, e.g., $F(T) \geq 0.99$, yielding empirical estimates for the minimal process duration T_{min}^F . To this end, we turn to numerical quantum optimal control techniques as described in Sec. II B. A particular difficulty in driving the transition is that the energy excitation spectrum becomes approximately gapless (exactly gapless in the thermodynamic limit) in the superfluid limit. This implies very long (diverging) adiabatic timescales for crossing the critical point of the phase transition. Thus, residual population defects may become pinned, which can be estimated either by Kibble-Zurek theory or as a cascade of independent Landau-Zener avoided crossing transitions [64–67].

One possible physical context for desiring such a fast and precise transfer lies in neutral atoms trapped with a cubic optical lattice. There the |Mott) state is a candidate for quantum information processing [68] tasks and quantum simulation of spin systems [69], but experimental protocols for the initial lattice loading leaves the system in the |SF) state. Within this physical implementation of the Bose-Hubbard model, the occupation number n_i is associated with the lowest band Wannier function maximally localized on site i , see Appendix A for a more detailed discussion. The characteristic energies $J_x(v_x)$ and $U(v_x, v_y, v_z)$ are then implicitly related to the lattice trapping depths v_x, v_y, v_z as denoted and for fixed v_y, v_z this translates into a functional dependence $v_x(U/J_x)$. These monotonic mappings are calculated numerically for a set of experimentally relevant lattice parameters in Appendix A (J_x is exponentially decreasing with v_x and v_x is exponentially increasing with U/J_x).

Matrix product states turn out to be an effective description for one-dimensional many-body systems. The price paid for such extended simulatory treatments is a significant increase in analytical and numerical code logistics. See Appendix B for a more detailed discussion of MPS.

B. Exact gradient optimization

Here we briefly review derivative-based optimal control and the main results of Ref. [58] in the context of the present paper.

In quantum optimal control, we seek to dynamically steer a quantum mechanical process in a controlled way such as to maximize a desired physical yield. For unitary evolution, any such task can be encoded as a minimization over an appropriate cost functional $J[\hat{U}(T; 0)]$, where

$$\hat{U}(T; 0) = \mathcal{T} \exp \left(-i \int_0^T \hat{H}(t) dt \right), \quad (3)$$

is the time evolution operator in units where $\hbar = 1$ from time $t = 0 \rightarrow T$, \mathcal{T} denotes time ordering, and \hat{H} is the system Hamiltonian carrying some generic time dependence. The Hamiltonian can be generically decomposed as

$$\hat{H} = \hat{H}^d(t, u(t)) = \hat{H}^d(t) + \hat{H}^c(t, u(t)), \quad (4)$$

where \hat{H}^d is the uncontrollable drift Hamiltonian and \hat{H}^c is the control Hamiltonian parametrized by the control $u(t)$ which allows manipulation of the unitary time evolution. The extension to more than one control parameter is discussed in Ref. [58]. Discretizing time on a regular grid of length N_t spaced by δt , we obtain

$$t \in [t_1, t_2, \dots, t_{N_t}] = [0, \delta t, \dots, T], \quad t_j = (j-1)\delta t, \quad (5a)$$

$$\hat{U}(T; 0) \approx \prod_{j=1}^{N_t-1} \hat{U}_j = \hat{U}_{N_t-1} \dots \hat{U}_2 \hat{U}_1, \quad (5b)$$

with time indices denoted as subscripts and \hat{U}_n the propagator across the time interval $[t_n, t_{n+1}] = [t_n, t_n + \delta t]$. The propagators \hat{U}_n depend on the discretized control vector $\mathbf{u} = (u_1, \dots, u_{N_t})$ in some manner depending on the chosen discretization scheme. We can then numerically minimize $J(\mathbf{u}) = J(\hat{U}_{N_t-1} \dots \hat{U}_2 \hat{U}_1)$ through iterative local updates of the control vector

$$\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \alpha^{(k)} \mathbf{p}^{(k)}, \quad (6)$$

such that $J(\mathbf{u}^{(k+1)}) \leq J(\mathbf{u}^{(k)})$. The step size $\alpha^{(k)} > 0$ at iteration k is found by line searching and there are several choices for the search direction $\mathbf{p}^{(k)}$ which all depend on the local gradient $\nabla J(\mathbf{u}^{(k)})$. This includes, e.g., the steepest descent, BFGS, and Newton direction [56].

A local minimizer \mathbf{u}^* of J is called an optimal control. The cost typically contains several contributions, $J = \sum_i J_i$, where each term encodes a desired feature that prospective optimal controls strives to fulfill. This always includes a term explicitly related to the quantum dynamics, $J_{\text{dyn}}(\hat{U}_{N_t-1} \dots \hat{U}_1)$, and potentially various control constraints, $J_{\text{con}}(\mathbf{u})$, that do not (see Appendix C). The total gradient is then given by the sum of the individual gradients, $\nabla J = \sum_i \nabla J_i$. The functional form of the cost J_{dyn} depends on the particular unitary task [53], yet the associated gradient elements can always be reduced to evaluating expressions on the form

$$\frac{\partial J_{\text{dyn}}}{\partial u_n} \sim \frac{\partial}{\partial u_n} (\hat{U}_{N_t-1} \dots \hat{U}_2 \hat{U}_1). \quad (7)$$

Evaluating the right-hand side of (7) constitutes the most numerically expensive computation in this optimization paradigm and is what leads to the typical computational bottlenecks mentioned in the Introduction. Once calculated, however, it is straightforward to assemble the gradient for

any $J[\hat{U}]$. A main result of Ref. [58] is that Eq. (7) can be efficiently evaluated by

$$\frac{\partial}{\partial u_n} \left(\prod_{j=1}^{N_t-1} \hat{U}_j \right) = \left(\prod_{j=n}^{N_t-1} \hat{U}_j \right) \left(-i \delta t \frac{\partial \hat{H}_n^c}{\partial u_n} \right) \left(\prod_{j=1}^{n-1} \hat{U}_j \right), \quad (8a)$$

given that the system is described in a basis where $\hat{H}_n^c = \hat{H}^c(t_n, u_n)$ is diagonal and simultaneously employing a specific [70] Trotterized propagator,

$$\hat{U}_n = \hat{U}_n^{\text{ST}} \equiv \hat{U}_{n+1}^{c/2} \hat{U}_n^d \hat{U}_n^{c/2}, \quad (8b)$$

where we defined the control and drift propagators

$$\hat{U}_n^{c/2} \equiv \exp(-i \hat{H}_n^c \delta t / 2), \quad \hat{U}_n^d \equiv \exp(-i \hat{H}_n^d \delta t), \quad (8c)$$

and $\hat{H}_n^d = \frac{1}{2}(\hat{H}_{n+1}^d + \hat{H}_n^d)$, where $\hat{H}_n^d = \hat{H}^d(t_n)$.

The $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ state transfer described in Sec. II A is conveniently encoded as a maximization of the fidelity or minimization of an associated J_{dyn} cost, respectively, given by [58]

$$F = |\langle \psi_{\text{tgt}} | \psi(T) \rangle|^2 = |\langle \psi_{\text{tgt}} | \hat{U}(T; 0) | \psi_{\text{ini}} \rangle|^2, \quad (9a)$$

$$J_F = \frac{1}{2}(1 - F), \quad (9b)$$

where $|\psi_{\text{tgt}}\rangle = |\text{Mott}\rangle$ is the target state, $|\psi_{\text{ini}}\rangle = |\text{SF}\rangle$ is the initial state, and $|\psi_{N_t}\rangle = |\psi(T)\rangle = \hat{U}(T; 0) |\psi_{\text{ini}}\rangle$ is the time-evolved state at final time T . Using Eqs. (8), the analytically exact gradient elements [71] of the fidelity cost can then be written as

$$\frac{\partial J_F}{\partial u_n} = \text{Re} \left(i o^* \left\langle \chi_n \left| \frac{\partial \hat{H}_n^c}{\partial u_n} \right| \psi_n \right\rangle \right) \delta t, \quad (10a)$$

where $|\psi_{n+1}\rangle = \hat{U}_n |\psi_n\rangle$ is the forward propagated initial state $|\psi_1\rangle = |\psi_{\text{ini}}\rangle$ while $|\chi_n\rangle = \hat{U}_n^\dagger |\chi_{n+1}\rangle$ is the backward propagated target state $|\chi_{N_t}\rangle = |\psi_{\text{tgt}}\rangle$ with the transfer amplitude $o = \langle \chi_{N_t} | \psi_{N_t} \rangle$. The exact gradient computation is only dominated by the time it takes to solve the dynamics—once forward and once backward. Any optimal control algorithm, including derivative-free ones, is in each iteration bounded from below by this propagation time. This gradient thus saturates this bound up to a small constant scaling (as opposed to other known exact gradients).

It is noted that a variational treatment of the cost functional prior to temporal discretization leads to a gradient expression that is superficially similar to Eq. (10a), but such an approach assumes an infinitesimal time step. Time steps are, however, always finite in numerical practice. This subtlety means that the details of the chosen propagator are important and that the well-known variationally obtained gradient is only a first-order approximation. It is then only under the conditions surrounding Eqs. (8) that the corrections vanish and Eq. (10a) is exact, which by happenstance coincides with the approximate variational calculation. See Ref. [58] for an expanded discussion.

The control may also be decomposed on a set of parametrized functions $\{f_l(t; \boldsymbol{\theta})\}_{l=1}^L$,

$$u(t; \boldsymbol{\theta}) = f_0(t) + \sum_{l=1}^L f_l(t; \boldsymbol{\theta}), \quad (10b)$$

where $f_0(t)$ is a fixed reference function and $\theta \in \mathbb{R}^M$ are now the optimization parameters. This is widely referred to [72] as the chopped random basis (CRAB) technique [54,73], and a typical choice of functional basis consists of randomized Fourier components. By choosing an appropriate set of functions, the search space can be both significantly reduced and focused on certain realistic control shapes, depending on the application. Exact gradient optimization in a CRAB parametrization is also admitted by Eqs. (8) through the chain rule:

$$\frac{\partial J(\mathbf{u}(\theta))}{\partial \theta_i} = \sum_{n=1}^{N_t} \frac{\partial J}{\partial u_n} \frac{\partial u_n}{\partial \theta_i} = \sum_{j=1}^{N_t} \sum_{l=1}^L \frac{\partial J}{\partial u_n} \frac{\partial f_l(t_n; \theta)}{\partial \theta_i}. \quad (10c)$$

The gradient element with respect to the parameter $(\theta)_i = \theta_i$ therefore still depends on the nonparametrized gradient elements $\partial J / \partial u_n$ described above with negligible further computational overhead.

In applying the exact gradient methodology here, it is convenient to rescale Eq. (1) at time t_n by the instantaneous tunneling rate:

$$\hat{H}_n = \frac{\hat{H}_{\text{SI}}(t_n)}{J_x(t_n)} = \sum_{i=1}^{N_s-1} \hat{h}_{[i,i+1]}^J + \frac{1}{2} \left(\frac{U(t_n)}{J_x(t_n)} \right) \sum_{i=1}^{N_s} \hat{h}_{[i]}^U. \quad (11a)$$

We identify the drift and control Hamiltonian as, respectively,

$$\hat{H}^d = \sum_{i=1}^{N_s-1} \hat{h}_{[i,i+1]}^J, \quad \hat{H}^c = \frac{u_n}{2} \sum_{i=1}^{N_s} \hat{h}_{[i]}^U, \quad (11b)$$

where $u_n = U(t_n)/J_x(t_n)$ is the control parameter. The control Hamiltonian is thus diagonal, as required for the validity of Eqs. (8a) and (10a), when represented in the site-occupation basis. As an additional numerical benefit, the drift Hamiltonian is time independent, $\hat{H}_n^d = \hat{H}^d$, and so is the control Hamiltonian derivative needed in Eq. (10a):

$$\frac{\partial \hat{H}_n^c}{\partial u_n} = \frac{1}{2} \sum_{i=1}^{N_s} \hat{h}_{[i]}^U. \quad (11c)$$

This structure of Eq. (11a) can be exploited to significantly accelerate the time evolution of the MPS as described in Appendix B.

The control vector $\mathbf{u} = (\dots, u_n, \dots) = (\dots, U(t_n)/J_x(t_n), \dots)$ is in our simulations independent of the physical implementation of the Bose-Hubbard model. It can be translated into a corresponding time sequence of laboratory parameters such as, e.g., optical lattice depth $\mathbf{v}_x(\mathbf{u}) = (\dots, v_x(u_n), \dots)$ by specifying the relationship between v_x and U/J_x as discussed in the Appendixes. The value $v_x(u_n)$ must be applied for a duration proportional to $\hbar/J_x(u_n)$, and the total duration measured in SI units therefore depends on the control $T_{\text{SI}}(\mathbf{u})$. The translated \mathbf{v}_x is therefore qualitatively speaking a scaled version of the corresponding \mathbf{u} .

In the following, we denote durations by the format $T = T_{\text{sim}}(T_{\text{SI}})$, where $T = T_{\text{sim}}$ is the process duration entering in Eqs. (5). The presented SI numbers are based on the set of

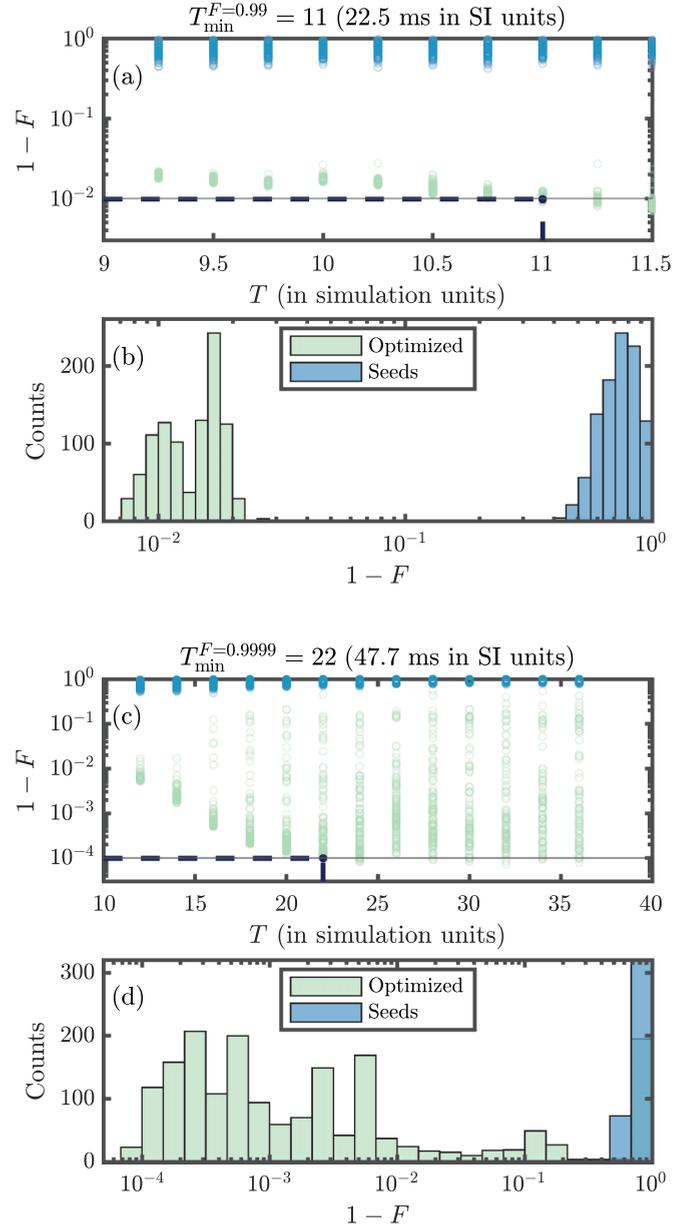


FIG. 1. Optimization results for $N_s = N_p = 20$ showing infidelity $1 - F$ (lower is better) for each seed before (blue) and after (green) optimization. (a) Each dot represents a control and the transparency informs about the distribution density. In this batch, we find $T_{\min}^{F=0.99} = 11$ (22.5 ms) indicated with the dashed line. (b) Histogram (left is better) prior to and after optimization with the same color scheme. (c), (d) As (a) and (b), but at higher durations and finding $T_{\min}^{F=0.9999} = 22$ (47.7 ms).

experimental lattice parameters described in Ref. [22], see Appendix A.

C. Results

We now present exact gradient optimization results for the $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ many-body transfer at unit filling $N_p = N_s = 20$ using MPS. A numerical exact diagonalization treatment [74] is prohibitively resource intensive at these numbers as the Hilbert space has dimensionality $D_{\mathcal{H}} \sim 10^{11}$ and

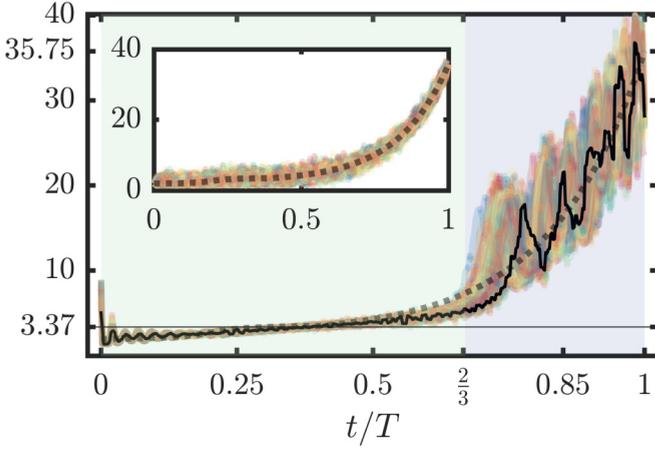


FIG. 2. Sample duration-normalized control ramps (multicolored faded lines) $u(t/T)$ for solutions with $F \geq 0.99$. The corresponding seeds are shown in the inset, the black dotted line indicates the adiabatically inspired reference control, the horizontal line denotes the critical point value for the phase transition (see Appendixes), and the black solid line highlights the optimal control with duration $T_{\min}^{F=0.99} = 11$ (22.5 ms). The optimized controls are characterized by two distinct segments denoted by the shaded areas.

requires ~ 1 TB memory just to store a single generic state. The initial starting points (seeds) for the optimization were generated by adding a sum of randomized Fourier components to an adiabatically inspired reference control, i.e., an exponential function that slowly crosses the critical point before ramping up. After generating a seed, it is optimized using Eq. (10a)—it is not bound by a parametrization and each $u_n = u(t_n)$ is independently adjustable. Appendix C gives an overview of further optimization details.

Figure 1 shows the empirically achieved fidelity through numerical optimization as a function of process duration T . From these, we find minimal duration estimates $T_{\min}^{F=0.99} = 11$ (22.5 ms) and $T_{\min}^{F=0.9999} = 22$ (47.7 ms). The highest fidelities as a function of T are found to approximately follow an exponential law. The effect of the optimization is seen to be significant as solutions are typically improved by several orders of magnitude. Sample optimization trajectories are shown in Appendix C.

This successful application of exact-gradient optimal control techniques in the many-body limit constitutes the first main result of this paper.

Figure 2 shows a subset of the optimized control ramps from Fig. 1(a) and we observe two control segments with very distinct behaviors.

Segment 1—“linear”: $t/T \in [0, \frac{2}{3}]$.

With nearly vanishing variance and following an initial perturbation, this predominantly linear segment slowly crosses the critical point with high-frequency, low-amplitude oscillations.

Segment 2—“bang-bang”: $t/T \in [\frac{2}{3}, 1]$.

With large variance, this segment ramps to the final control value and consists of several low-frequency, high-amplitude oscillations. These oscillations turn out to be smoothed approximations to a bang-bang structure.

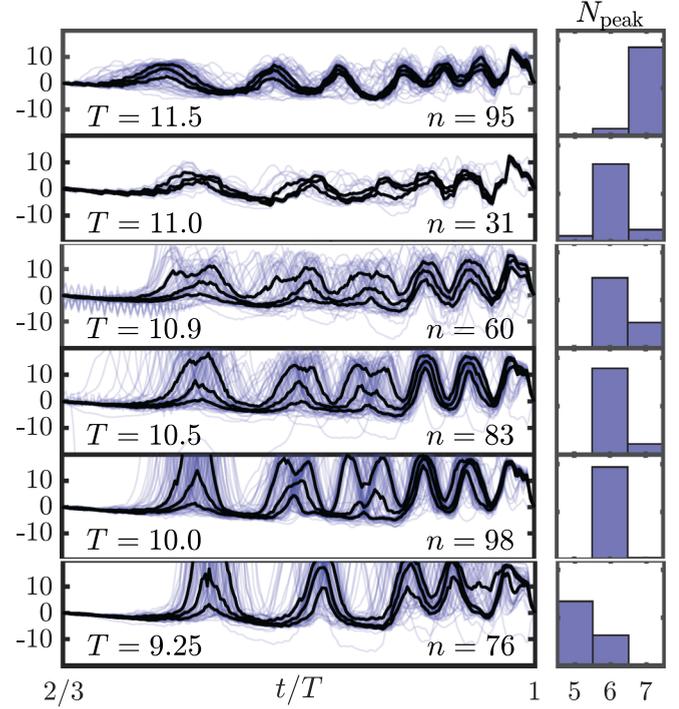


FIG. 3. Distinct optimal control strategies characterized by a different number of peaks (bangs) in segment 2 that are active at different T (rows). The left panels show n found high-fidelity controls (translucent lines) and their (25,50,75)% quantiles (black lines) (the linear contribution has been subtracted). The right panels show the relative distribution of the n solutions into the three identified strategies at these T .

Focusing on segment 2, subtracting the linear contribution there and resolving the optimized solutions according to their duration T , we obtain a clearer picture in Fig. 3, namely, there are several optimal control strategies [75] that are individually characterized by an integer number of peaks (“bangs”) $N_{\text{peak}} = 5, 6, 7$. The best and most prominent strategy depends on T where higher T leads to a higher optimal number of peaks. This trend is verified to continue outside of this particular range of durations where different N_{peak} are prevalent. At lower T the peaks are increasingly bang-bang-like while at higher T the peak shapes are smoother and less extreme [76]. Since the effect on the dynamics is largely the same, however, we will for simplicity refer to these peaks collectively as bangs although they may not always meet the strictest definition. Translating $u(t) = U(t)/J_x(t)$ into, e.g., optical lattice depths $v_x(t)$ scale the shape but leave the overall structure intact—the peaks then appear to have roughly the same width ≈ 1.3 ms and center distance $\approx 2 - 2.5$ ms independently of T . The recoil energy time scale is roughly $\hbar/E_R \approx 0.5$ ms for comparison.

The discovery of this family of optimal solutions for the $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ transfer constitutes the second main result of this paper. The existence of other types of optimal solution shapes cannot be excluded and we leave these types of searches for future work.

To understand the physical mechanisms at play in the problem, we inspect the quantum dynamics induced by these

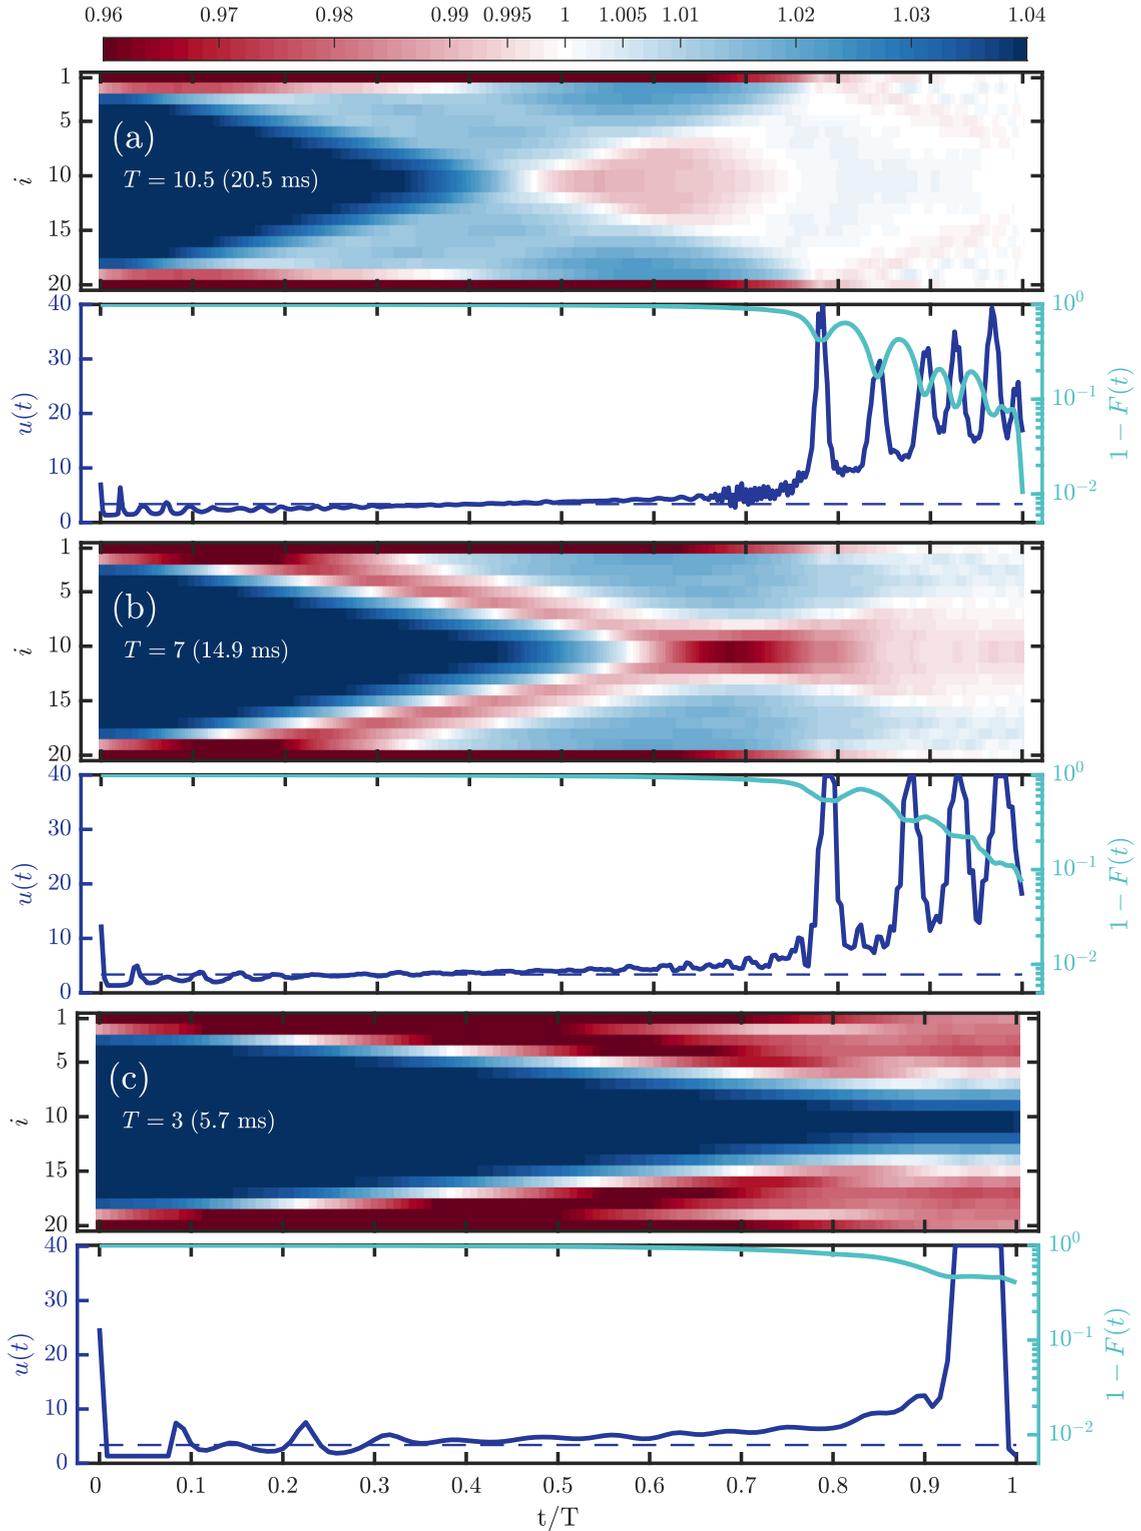


FIG. 4. (a)–(c) Time evolution along sample optimized controls at three different durations. Each upper panel shows the site occupations $\langle \hat{n}_i \rangle_{\psi(t)}$ (sensitive near unit occupancy and note $\langle \hat{n}_1 \rangle_{\psi(0)} = 0.67$) while each lower panel shows the corresponding control (solid purple line, left axis) and critical value $(U/J)_{\text{crit}}$ (dashed purple line, left axis) and infidelity (solid teal line, right axis). The t axis is slightly extended to better display the behavior at $t/T = 0, 1$.

fidelity-optimized controls at different durations in Fig. 4. From $\langle \hat{n}_i \rangle_{\psi(t)} = \langle \psi(t) | \hat{n}_i | \psi(t) \rangle$ in Fig. 4, we find that a light-cone-like homogenization of site population takes place when crossing the critical point in segment 1. The population is oth-

erwise initially concentrated in the bulk (roughly sites 4–17) due to finite edge effects. This process is clearly limited by the propagation velocity $\partial \langle \hat{n}_i \rangle_{\psi(t)} / \partial t$ and thus also influences the minimal duration T_{min}^F for high-fidelity transfers. The tardiness

TABLE I. Summary of problem specifications in this paper and Refs. [21,22]. The optimization in this paper does not rely on explicit lattice parameters as discussed in Appendix A. References [21,22] do not report the concrete details of the control parametrization and certain other parameters for the optimization.

	This paper	Ref. [21]	Ref. [22]
Physical system			
Sites and particles	$N_s = N_p = 20$	$N_s = N_p = 20$ (and others)	$N_s = 32, N_p = 16$
Harmonic trapping	No	No and yes (weak)	Yes (weak)
Lattice parameters	$\lambda = 1064$ nm, $v_y = v_z = 20 E_R$	$\lambda = 826$ nm, $v_y = v_z = 30 E_R$	$\lambda = 1064$ nm, $v_y = v_z = 20 E_R$
Optimization details			
Objective	$J_F + J_\alpha + J_Y$ Eq. (C1)	Energy minimization halting below $\rho = 10^{-3}$ Eq. (13)	η Eq. (14)
Update rule	Exact gradients (BFGS)	Gradient-free (Nelder-Mead)	Gradient-free (Nelder-Mead)
Control	U/J_x	v_x	v_x
Parametrization	None	Fourier CRAB Eq. (10b)	Fourier CRAB Eq. (10b)
Bandwidth limited by	J_Y cost	Parametrization	Parametrization
Timescale	\hbar/J_x	\hbar/E_R	\hbar/E_R
Initial and target depths	$3E_R \rightarrow 13E_R$	$2E_R \rightarrow 22E_R$	$3E_R \rightarrow 14E_R$
MPS parameters	$D = 200, s_{\max} = 10^{-12}, d = 5$	$D \leq 100$	$D \leq 24, s_{\max} = 10^{-5}$
δt for simulation	0.025	0.01 – 0.001	0.01

hereof is somewhat expected because we are crossing a phase transition and the adiabatic timescales approach infinity, as discussed earlier. Even upon reaching a largely homogenized site population after segment 1, however, the fidelity remains very small. The following bang-bang process in segment 2 corresponds to alternating between tunneling events and site-locking—during the latter, individual c phases are imprinted on the individual Fock components

$$\hat{U}_n |n_1, \dots, n_{N_s}\rangle \approx |n_1, \dots, n_{N_s}\rangle \times \exp\left(-\frac{i u \delta t}{2} \sum_{i=1}^{N_s} n_i (n_i - 1)\right), \quad (12)$$

since tunneling is approximately negligible in the deep lattice limit, $u \lesssim u^{\max}$, see Appendix A. This leads to a nontrivial interplay [$F(t)$ is nonmonotonic] between the canonically conjugate population and phase variables that evidently produces the correct interferences finally leading to the target |Mott) state.

It is interesting to note that a recent work [77] argues that optimal controls should exhibit a bang-smooth-bang structure for a certain class of problems. Said problems are equivalent to special cases of the state transfer formulation where the total Hamiltonian, $\hat{H}(u(t)) = u(t)\hat{H}_1 + (1 - u(t))\hat{H}_2$, is a linear interpolation with $0 \leq u(t) \leq 1$ between Hamiltonians \hat{H}_1 and \hat{H}_2 , and $|\psi_{\text{ini}}\rangle = |\text{GS}; u = 1\rangle$ and $|\psi_{\text{tgt}}\rangle = |\text{GS}; u = 0\rangle$ are the ground states for \hat{H}_1 and \hat{H}_2 . The problem studied here is approximately on this form and it appears plausible that our optimized controls are related to this bang-smooth-bang prediction.

D. Comparison to previous open-loop efforts

We now discuss our two main results in relation to previous open-loop efforts [21,22] of optimizing the |SF) \rightarrow |Mott) transition. These explicitly assume an optical lattice implementation and employed derivative-free Nelder-Mead

methodologies to optimize the expansion coefficients of a Fourier component CRAB. The respective optimization results and methodologies are discussed in turn.

Our first result was to show that exact gradients are now a viable option in the many-body limit, which was previously considered potentially infeasible [21]. Direct quantitative comparisons to Refs. [21,22] in terms of fidelity and process durations are obstructed by the differences in the exact problem formulations summarized in Table I. Nevertheless, the conservative quantitative inferences discussed below indicate that the optimization technique is indeed promising. Strong comparative conclusions between optimization methodologies can only be drawn if they are applied in numerical environments where everything else is equal. We consider such extended numerical comparisons outside the scope of this demonstration of exact gradients in the many-body limit.

The optimized observables reported in Refs. [21,22] are not the |Mott) fidelity. Instead they are related only to the site population statistics and are thus insensitive to the relative phases between Fock state components. They are thus correlated with the fidelity but not in one-to-one correspondence.

Reference [21] reports on the final density of defects:

$$\rho = \frac{1}{N_s} \sum_{i=1}^{N_s} |\langle \hat{n}_i \rangle_{\psi(T)} - 1|. \quad (13)$$

Vanishing ρ values [78] are indicative of a |Mott) state and optimization of a solution in Ref. [21] was halted if it achieved $\rho \leq \rho_c = 10^{-3}$. The density of defects should not be affected by the differences in Table I, and this allows a point of comparison when evaluating ρ for all our fidelity-optimized solutions as shown in Fig. 5(a). We are able to verify that a fixed value of ρ can display a large variance in F and vice versa. For example, $\rho \approx 3 \times 10^{-3} \geq \rho_c$ can correspond to both $F \approx 0.75$ and $F \approx 0.93$ while $F \approx 0.93$ can also correspond to $\rho = 6 \times 10^{-4} \leq \rho_c$. Thus ρ_c is not an ideal stopping condition for the |Mott) state in terms of fidelity and it is not

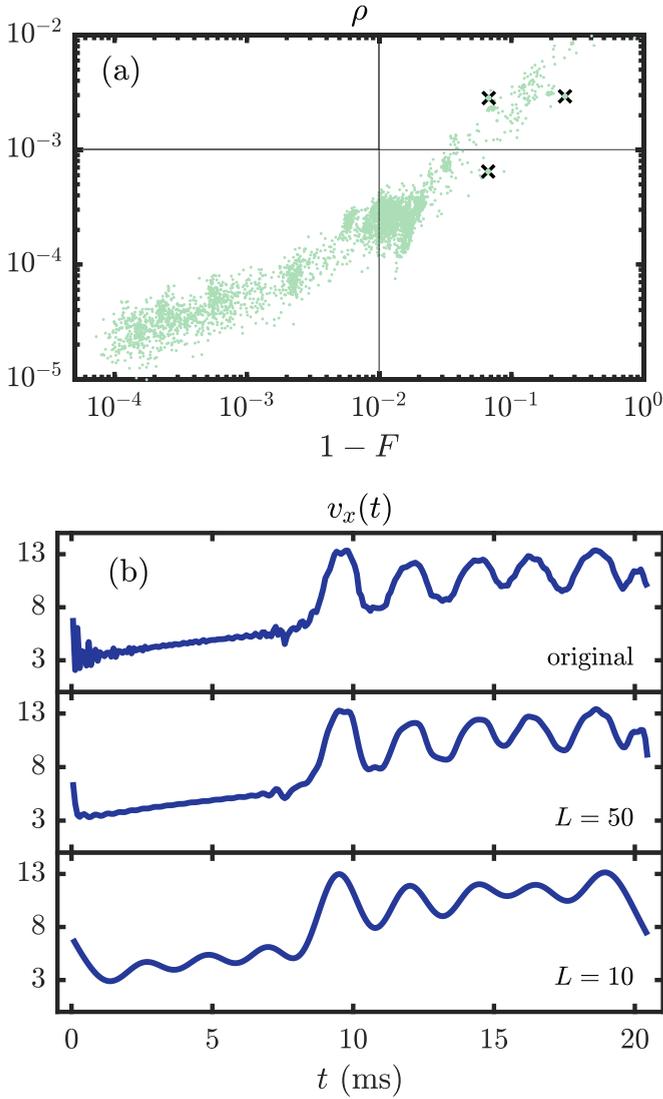


FIG. 5. (a) Density of defects Eq. (13) for all the fidelity-optimized solutions (green dots). The horizontal line denotes $\rho_c = 10^{-3}$ from Ref. [21] and the vertical line denotes $F = 0.99$. The black crosses highlight solution pairs that have nearly the same ρ but very different F and vice versa. (b) Optical lattice depths $v_x(t)$ mapped from $u(t)$ in Fig. 4(a) and the corresponding signal when removing Fourier components with a frequency larger than $\nu_L = 2\pi L/T$. The y axis is in units of recoil energies using lattice parameters from Ref. [22].

possible to know what the fidelity distribution in Ref. [21] would look like. However, all our solutions with $F \geq 0.99$ have less than $\rho \approx 4 \times 10^{-4}$ and solutions with $F \approx 0.9999$ are distributed near $\rho \approx 2 \times 10^{-5}$.

Our numerical calculation of the relationship $v_x(U/J_x)$ in Fig. 6—which is used to calculate SI times for a given set of lattice parameters described in Appendix A—yields values that are consistent with those stated in Ref. [22]. When we instead use the tighter lattice parameters in Ref. [21] the SI durations are reduced by roughly a factor of 2–3. This places our process durations on the order of ms, which is similar to the 3.09 ms used in Ref. [21]. There is, however, not a basis for direct comparison between these since we observe

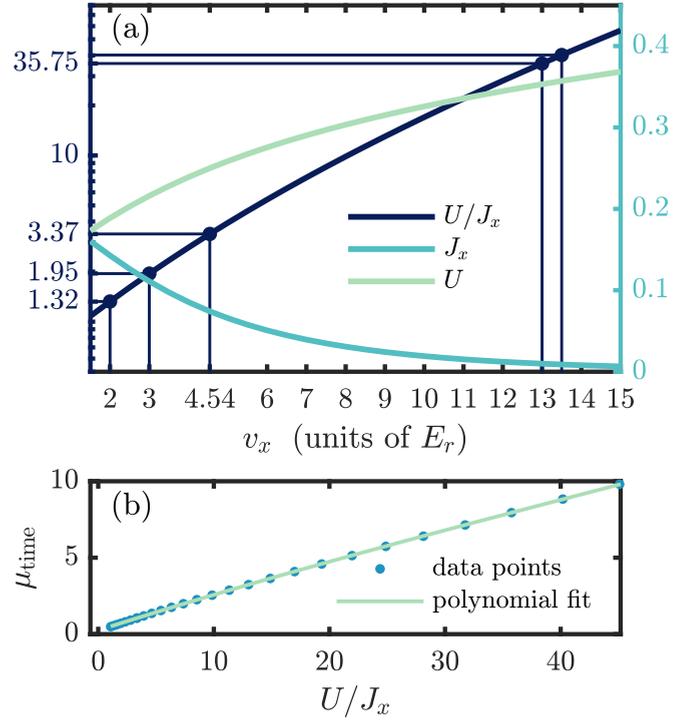


FIG. 6. (a) Constitutive relations $U(v_x)$, $J_x(v_x)$ for the concrete optical lattice described in the Appendix (right axis in units of E_R). The dots denoted on the U/J_x curve (left axis) indicate values of note. Left to right, they are the minimally allowed $v_x = 2E_R$, the v_x defining |SF), the critical point of the |SF) \rightarrow |Mott) phase transition, the v_x defining |Mott), and the maximally allowed $v_x = 13.5E_R$ (corresponding to $U/J = 40.18$). (b) The timescaling is $\mu_{\text{time}} = \hbar/J_x(v_x)$ in units of ms and the transfer duration in SI units Eq. (A6) thus depends on the control vector $\mathbf{u} = (\dots, (U/J_x)_n, \dots)$.

a discrepancy in $v_x(U/J_x)$ at the values reported in Ref. [21]. (The variance between values of F and ρ above suggests that $F \geq 0.99$ requires more than 3.09 ms in the geometry of Ref. [21]).

The work in Ref. [22] assumes $N_p = 16$, $N_s = 32$, and includes a harmonic potential $\Omega \times \sum_{i=1}^{N_s} (i - i_0)^2 \hat{n}_i$, where $i_0 = (N_s - 1)/2$ and $\Omega = 2.4 \times 10^{-3} E_R$, which is weak compared to the on-site interaction U at all lattice depths (see Fig. 6). Note that Ref. [21] shows optimization results both with and without a similarly weak harmonic potential and in that instance there does not seem to be a significant difference between the quantitative outcomes.

The observable reported and minimized in Ref. [22] is the rescaled average variance in the center of the trap

$$\eta = \frac{1}{8} \sum_{i=N_s/2-3}^{N_s/2+4} \frac{\Delta \hat{n}_i(T)}{\Delta \hat{n}_i(0)}, \quad (14)$$

where $\Delta \hat{n}_i(t) = \langle \hat{n}_i^2 \rangle_{\psi(t)} - \langle \hat{n}_i \rangle_{\psi(t)}^2$. The smallest value reported is $\eta \approx 1 \times 10^{-1}$ at roughly 15 ms. Similarly to ρ , we verify in our data that a fixed value of η (where we extend the sum to all lattice sites) can exhibit a large variances in F . We observe $\eta = 1 \times 10^{-1}$ at $F = 0.75$ while all fidelities $F \geq 0.99$ have $\eta \leq 1.4 \times 10^{-2}$ and at most reached $\eta \approx 1.2 \times 10^{-2}$. Since the harmonic potential is quite weak and

that the reported η more leniently covers only eight sites, these numbers suggest that it is quite unlikely that the corresponding states in Ref. [22] are close to $F = 0.99$ with respect to the ground state. In terms of process durations, the $F = 0.75$ control in Fig. 4(c) attains $\eta \approx 1 \times 10^{-1}$ at about a third of the duration in Ref. [22].

These assessments indicate that our exact gradient methodology is able to achieve results that are competitive with the methodologies employed by Refs. [21,22] in terms of solution quality and transfer times.

Our second main result was that the $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ transition has a family of solutions consisting of distinct segments consisting of a predominantly linear sweep followed by bangs. This overall control structure is preserved when mapped to optical lattice depths as seen in Fig. 5(b). From a qualitative point of view, the optimized controls reported in Refs. [21,22] are in contrast very smooth and are generally similar to each other. A more quantitative statement is that our obtained optimized controls have a higher information-theoretic *control complexity* as initially introduced in Ref. [79] where an operational definition is given by the number of Fourier components needed to solve the problem to a given fidelity threshold (these and associated notions were shortly after treated in more generality in Ref. [43]). A partial reason for the observed disparity can be understood as follows.

As mentioned, Refs. [21,22] are based on the CRAB technique in Eq. (10b) with derivative-free Nelder-Mead optimization. They utilize Fourier components centered on low-lying harmonic frequencies, e.g., $\nu_l = 2\pi l(1 + r_l)/T$ where $r_l \in [0 : 1]$ are randomized frequency offsets [21]. That is, the basis functions are on the form $f_l(t, \theta) \propto \theta_l^{\sin} \sin \nu_l t + \theta_l^{\cos} \cos \nu_l t$, where the control $\theta = (\theta^{\sin}, \theta^{\cos}) = (\theta_1^{\sin}, \dots, \theta_L^{\sin}, \theta_1^{\cos}, \dots, \theta_L^{\cos})$ represents the $M = 2L$ expansion coefficients. The number of components used in Refs. [21,22] is stated as small but is not explicitly specified. A low-dimensional Fourier decomposition effectively means that the control signal is bandwidth, or rate of change, limited. Our presented optimizations instead incorporate rate of change limitations through a cost functional, see the Appendixes. The restrictions imposed in this way are somewhat more difficult to characterize, but it is observed that the maximal rate of change due to the approximate bangs in, e.g., Fig. 4 is similar to the maximal rate of change near $t = T$ in Refs. [21,22].

Regardless of bandwidth considerations, the restriction to a small M in Refs. [21,22] is necessary because the heuristic Nelder-Mead methodology does not generally perform well as M becomes large [56]. Reference [50] reports, for example, on a similar parametrization for a different control problem, and the statistically best performance for Nelder-Mead was limited to $F \approx 0.9$ at only $M \approx 20$ (the performance deteriorated for higher M). Results in the same parametrization obtained by gradient-based means showed statistically, however, that roughly a factor of 2–3 more components were at minimum required to resolve $F = 0.99$ controls. A similar number of components was needed in Ref. [12]. Assuming that $M \approx 20$ is a reasonable general guideline for Nelder-Mead's maximum basis size leaves only the lowest $L = M/2 \approx 10$ harmonic frequencies for the expansion in Refs. [21,22]. The spectral cutoffs [80] in Fig. 5(b) show that this is insufficient to represent the control structure yielding $F = 0.99$ from Sec. II C,

and roughly a factor of 5 more frequency components are needed.

The initial practicality of restricting the search space to a limited set of basis functions such as Fourier components may therefore not be ideal because there is no guarantee that the ultimately best solutions are captured by the parametrization. A typical extension known as dressed-CRAB (dCRAB) [81] attempts to remedy this phenomenon by drawing a new set of basis functions, e.g., by sampling different random frequency offsets for the Fourier components. This can be expected to work reasonably well only if the overall parametric structure can capture the solutions. In this case, however, a significant increase in the spectral content is needed to produce the structure in Sec. II C. This cannot be accounted for by small perturbations around a limited number of low harmonic frequencies. The Fourier CRAB can only attain the required control complexity by increasing the number of components L , but this is at odds with Nelder-Mead requiring a small number of optimization parameters. A derivative-free Nelder-Mead with a standard Fourier CRAB is thus likely prohibited in identifying the high-fidelity optimal controls in Sec. II C.

The discussion above highlights the potential drawbacks of preemptively choosing a parametrization. It does not imply that parametrization are always detrimental. For a given parametrization, however, we expect the gradient-based update rule in Eq. (10c) to statistically yield better results than a gradient-free update rule in the ideal open-loop context [50] (the opposite may be true in the closed-loop context and this leads into their potential unification discussed in Sec. III). It is also straightforward to calculate exact gradients with and without parametrizations for ρ and η observables in Eqs. (13) and (14). Since these observables are much more lenient than fidelity, it is expected that $\rho = 10^{-3}$ and $\eta = 10^{-1}$ solutions can be found at significantly reduced process durations, but that these will not correspond to high-fidelity solutions.

Despite the potential issues encountered by Nelder-Mead, it is quite clear that the optimized solutions in Refs. [21,22] improve significantly on both the observable metrics and transfer times compared to adiabatic solutions. This fact can be interpreted based on the present analyses as follows. Recall that the control strategies were characterized by a predominantly linear sweep followed by a number of bangs where a lower number corresponds to lower durations and fidelity. Fewer bangs suggest that the full control requires fewer spectral components for a faithful representation. This implies a more convenient situation for Nelder-Mead with a low-dimensional Fourier CRAB. Indeed, the smooth and monotonic controls in Refs. [21,22] also consist roughly of a linear sweep followed by a very steep rise that may be interpreted as the beginning of a bang. The results of Refs. [21,22] thus appear consistent with the single bang strategy which is associated with fidelities on the order of mid tens of percent. Given the spread in Fig. 5(a), such fidelities could reasonably well correspond to the optimized observable values ρ and η reported in Refs. [21,22].

In conclusion, the presented numerical studies indicate that the exact gradient methodology of Ref. [58] is a promising tool in the open-loop setting also for many-body systems.

III. UNIFIED OPTIMAL CONTROL

The results presented in Sec. II focused solely on ideal open-loop conditions. We now discuss the strengths and weaknesses of open- and closed-loop optimization, and how the former can be best utilized in prospective unified frameworks.

A. Open loop

Model-based open-loop methodologies can draw on very efficient derivative-based nonlinear optimization machinery [56] and allows a large degree of numerical parallelization. The produced controls are optimal with respect to the chosen mathematical optimization objective $J_{\text{theory}}(\mathbf{u})$ which typically contains contributions from both the quantum process and the experimental constraints. An optimal control \mathbf{u}^* numerically attained through this procedure can therefore only be expected to be as good as the underlying model of the experimental reality, $J_{\text{experiment}}(\mathbf{u}^*) \gtrsim J_{\text{theory}}(\mathbf{u}^*)$, and the acceptance of this potential performance degradation depends on the given context. These modeling errors could originate from many sources, for example, imperfect equipment fabrication, drifting or fluctuating noisy signals, or nondeterministic run-to-run system preparation. If these errors can be characterized as statistical uncertainties, one can endow the optimal controls with robustness towards these by, e.g., *ensemble optimization*. This places an increased emphasis on both gradient computation speed and exactness. The exact gradient optimization applied here addresses both these points and is therefore particularly well suited for such extensions; see Ref. [58] for a more detailed discussion.

Nevertheless, if the error sources are unknown and the model cannot be sufficiently refined to an acceptable degree then a \mathbf{u}^* found by open-loop control is inadequate.

B. Closed loop

The primary strength of closed-loop control lies in integrating the experiment itself in the optimization objective. This model-free approach allows $J_{\text{experiment}}(\mathbf{u})$ to be optimized directly and robustness toward fluctuations is achieved more naturally, given that they are not too large. Within this paradigm, gradients can in principle be calculated by finite differences but this is typically too impractical [82]. (This is true regardless of how experimentally accessible the optimized observable is, e.g., fidelity, density of defects or the rescaled average variance in Sec. II). By instead employing derivative-free direct search approaches, requiring only evaluation of $J_{\text{experiment}}$ itself, the numerical programming effort is significantly reduced since it can be outsourced to off-the-shelf black-box optimization implementations. These approaches are enabled in large by the widely adopted CRAB technique [54,73] discussed in Sec. II. By choosing an appropriate set of functions, e.g., Fourier components where the expansion coefficients play the role of optimization parameters, the search space is both significantly reduced and focused on certain realistic control shapes.

However, as also discussed in Sec. II, derivative-free methods are practically limited in the number of optimization parameters, and this potentially obfuscates the ultimately

best implementable solutions. Additionally, parallelization is greatly prohibited because the optimization relies on evaluating the cost on (possibly singular) experimental hardware.

C. Unified loop

Open- and closed-loop control thus have distinct and individual advantages and disadvantages. The fundamental goal of either approach is the same and combining these in a complementary unified-loop approach appears sensible [83]. Such metaiterative techniques can be strengthened further by also including model calibration [84].

We argue that an open-loop component is valuable even in the simplest constellation. Derivative-based algorithms and their parallelizable nature allow large-scale studies and, given that they are sufficiently efficient and accurate, can lead to the discovery of optimal control strategies [75]. Identification of such general solution structures and associated physical processes reveals a more complete characterization of a given problem than any singular solution does, and these insights can be used for choosing appropriate seeds (starting points for the optimization) and parametrizations for Eq. (10b). Thus, if the modeling is sufficient for correctly representing the optimal *strategy*, it is of much less importance if any *individual* control degrade in practice because the gap between theory and experiment can be minimized in subsequent closed-loop optimizations and calibrations.

As an example, the control strategies found in the Sec. II structure are not *a priori* obvious. Like in the open-loop case, an isolated derivative-free closed-loop effort using a standard [73] low-dimensional chopped Fourier parametrization could not simultaneously capture both the linear and bang-bang-like behavior. However, based on the strategy insights achieved by the presented gradient open-loop efforts, one could now, e.g., choose an appropriate reference ramp $u_{\text{ref}}(t)$ and $\{f_i(t; \boldsymbol{\theta})\}_{i=1}^L$ to be smoothed bangs where the optimization parameters $\boldsymbol{\theta}$ are the widths, positions, and heights of these. It was also found that the optimal number of bangs depended on T and this gives an idea of what L should be. In addition to subsequent closed-loop optimization in this refined parametrization, the gradient-based open-loop version in Eq. (10c) may also prove useful in this paradigm.

This broadens the possibility of alternating between open- and closed-loops as, e.g., proposed in Refs. [83,84].

IV. CONCLUSION

We applied a recently developed quantum optimal control methodology [58] to a representative example from the class of very high-dimensional, many-body state transfer problems described by MPS where exact diagonalization is prohibited. Our methodology provides exact gradients by explicitly including in their analytical derivation the effect of (1) Trotterizing the dynamics and (2) choosing a basis where the control Hamiltonian is diagonal. This circumvents detrimental scalings with the Hilbert space dimensionality which prohibit the use of other currently known exact gradients in high-dimensional, e.g., many-body, systems. The gradient computation time scales instead only with the time it takes to solve the quantum dynamics, i.e., the fundamental operation of any quantum control algorithm.

The exact gradient optimizations lead to very high-fidelity results for this class of problems (0.99–0.9999). An alternative type of solution for the $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ transition was also uncovered, a predominantly linear sweep across the critical point of the phase transition followed by a variable number of bangs (depending on the process duration). These solutions remain hidden for the earlier gradient-free approaches [21,22] since the employed CRAB parametrization cannot practically include a sufficient number of Fourier components, i.e., have a high enough bandwidth, to resolve these overall shapes. A (small) CRAB parametrization is a practical computational requirement for gradient-free methods, but for gradient-based methods—once unencumbered from the aforementioned detrimental scalings—it is a matter of choice which is free of principal basis size limitations.

Direct quantitative comparisons to Refs. [21,22] were somewhat hampered by the differences in the specific problem formulations, e.g., choice of figure of merit, but the afforded inferences are very encouraging. When the gradient is readily available, gradient-based approaches are generally considered favorable over gradient-free in the idealized open-loop setting when all else is equal (figure of merit, control parametrization, bandwidth limitation, etc.), and we expect that this would also apply here given a more extensive investigation in a collated numerical environment.

The documented efficiency over wide spans of Hilbert space dimensionality (in Ref. [58] and here) suggests that the methodology could be useful for future quantum optimization tasks.

In discussing the role of open-loop methodologies as a whole, we argued that they are particularly useful both for initially identifying appropriate control subspaces for a CRAB parametrization and, in conjunction with subsequent closed-loop methods, may remain a relevant optimization paradigm in the future.

ACKNOWLEDGMENTS

We thank T. Calarco, S. Montangero, K. Mølmer, M. Dalgaard, and C. A. Weidner for feedback. This work was funded by the ERC, H2020 Grant No. 639560 (MECTRL), and the John Templeton and Carlsberg Foundations. The numerical results presented in this paper were obtained at the Centre for Scientific Computing, Aarhus.

APPENDIX A: BOSE-HUBBARD MODEL IN OPTICAL LATTICE

The Bose-Hubbard model in Eq. (1) can be realized in a variety of physical platforms [31] and here we review the standard treatment for a cubic optical lattice. We first specify how the energetic parameters in Eq. (1) relate to the optical trapping depths and the conversion between simulation and laboratory timescales. These are then numerically calculated for a set of experimental parameters.

We assume a simple cubic periodic lattice approximated near the trap center by

$$V(\mathbf{r}) \approx \sum_{q=x,y,z} v_q \sin^2 k_l q = V_x + V_y + V_z, \quad (\text{A1})$$

with $V_q \equiv v_q \sin^2 k_l q$ being the potential in q direction with depth v_q . Here, $k_l = 2\pi/\lambda_l = \pi/a_{\text{lat}}$ is the laser wave number, λ_l is the laser wavelength, and a_{lat} is the lattice site separation. Thus, the potential is separable in all directions and we may decompose arbitrary wave functions as $\psi(\mathbf{r}) = \psi_x \psi_y \psi_z$. In particular, we may focus on a single direction for the single-particle stationary states. Choosing the x direction, we write $\hat{H}_x^{\text{1p}} \phi_k^n(x) = E_k^n \phi_k^n(x)$ where \hat{H}_x^{1p} is the single-particle operator in the x direction, n is the band index, and $k_l \leq k \leq k_l$ defines the first Brillouin zone of quasimomentum with intraquidistant spacing $\Delta k = 2\pi/L = 2\pi/(N_s a_{\text{lat}})$, where L is the length of the chain. The Bloch wave expansion reads $\phi_k^n(x) = e^{ikx} u_k^n(x)$, where u_k^n inherits the V periodicity. The Fourier series for both quantities contain only a few terms and substitution into the eigenproblem yields a particularly small, simple system of equations for the Fourier expansion coefficients of u_k^n [85]. After numerically obtaining $\phi_k^n(x)$ for a given value of v_x , the n th band Wannier state centered on site i is defined by

$$w_{n,x}(x - x_i) = \frac{1}{\sqrt{\mathcal{N}}} \sum_k^{\text{first Brillouin}} e^{-ikx_i} \phi_k^n(x), \quad (\text{A2})$$

where $x_i = i a_{\text{lat}}$ and \mathcal{N} is a normalization constant. To progress, we make the standard assumptions that the lattice has been loaded in the tube of sites defined by $\mathbf{r}_i = (x_i, 0, 0)$ and that v_y and v_z are sufficiently deep to suppress all tunneling events in their respective directions and that tunneling along x is nearest neighbor only. Additionally assuming that only the $n = 0$ band is occupied in each direction ($v_i \gtrsim 2E_R$) and dropping the index, the bosonic field operator can be expanded as $\hat{\Psi}(\mathbf{r}) \approx \sum_{i=1}^{N_s} \hat{a}_{x_i,0,0} \times w_x(x - x_i) w_y(y) w_z(z)$. Inserting this expansion in the many-body Hamiltonian for a dilute bosonic system [86], one obtains Eq. (1) by letting $\hat{a}_i \equiv \hat{a}_{x_i,0,0}$ and defining the constitutive relations

$$J_x(v_x) = - \int_{-\infty}^{\infty} w_x(x - x_i) \hat{H}_x^{\text{1p}} w_x(x - x_{i+1}) dx, \quad (\text{A3})$$

$$U(v_x, v_y, v_z) = g_{3D} \int |w_x(x - x_i) w_y(y) w_z(z)|^4 d\mathbf{r}, \quad (\text{A4})$$

where $g_{3D} = 4\pi \hbar^2 a_s/m$ is the two-body collision coupling strength, $a_s = 101a_0$ the s -wave scattering length of Rubidium 87, and a_0 is the Bohr radius. Thus, for a cubic optical lattice loaded with ultracold atoms, the energies are related to the trapping depths v_x, v_y, v_z through the Wannier states. Note the integrals over y and z in J_x are equal to one due to their normalization, and U factors into three independent one-dimensional integrals. Both J_x and U nontrivially depend on v_x , and although we assume only 1D dynamics along the x direction, the frozen out transverse y and z directions still implicitly enter in U via the associated Wannier functions w_y and w_z . With the constitutive equations in Eqs. (A3) and (A4), we can map any $u_n = U(t_n)/J_x(t_n)$ ramp into the corresponding trapping depth $v_x(t_n)$, as shown below.

The natural simulation timescales in Eq. (11a), $\hat{H}_{\text{SI}}/J_x = \hat{H}$, depend on the control. This can be seen by considering the

nondimensionalized propagator,

$$\begin{aligned}\hat{U} &= \exp\left(-i\frac{\hat{H}_{\text{SI}}\delta t_{\text{SI}}}{\hbar}\right) = \exp\left(-i\frac{\hat{H}_{\text{SI}}}{J_x}\left\{\frac{J_x\mu_{\text{time}}}{\hbar}\right\}\delta t_{\text{sim}}\right) \\ &= \exp(-i\hat{H}\delta t_{\text{sim}}) \Rightarrow \mu_{\text{time}} = \hbar/J_x,\end{aligned}\quad (\text{A5})$$

where time dependences have been omitted for clarity. These are working equations corresponding to $\hbar = 1$ and where $\delta t_{\text{SI}} = \mu_{\text{time}} \times \delta t_{\text{sim}}$. Time steps expressed in SI units, δt_{SI} , are related to (constant) dimensionless simulation numbers δt_{sim} , through the timescale $\mu_{\text{time}} = \hbar/J_x(u_n)$ which depends on the control value. In particular, the total duration of the transfer process given in SI time is

$$T_{\text{SI}}(\mathbf{u}) = \hbar\delta t_{\text{sim}} \sum_{n=1}^{N_t-1} J_x^{-1}(u_n), \quad (\text{A6})$$

i.e., the relevant timescales are given by the specific realization of the physical platform and depends on the control vector. Elsewhere in the paper, subscripts are dropped and we write $\delta t = \delta t_{\text{sim}}$ and all quantities of time are implicitly given in nondimensional simulation values unless followed by a unit.

Experimental parameters

To enable a degree of quantitative comparison in the main text, we consider the experimental lattice parameters given in Ref. [22] as follows. A summary of our problem specification compared to Refs. [21,22] is shown in Table I.

The lattice recoil energy is $E_R = \hbar^2\pi^2/(2a_{\text{lat}}^2m) \approx 2.03 \text{ kHz} \times \hbar$, where $m = 87 \text{ amu}$ is the mass of rubidium 87 and a_{lat} is the lattice site separation. Additionally, h is the Planck's constant, and \hbar the reduced Planck's constant. We assume a lattice of wavelength $\lambda = 1064 \text{ nm}$ with lattice spacing $a_{\text{lat}} = \lambda/2 = 532 \text{ nm}$. The transverse trapping depths are fixed at $v_y = v_z = 20 E_R$ and the $|\text{SF}\rangle \rightarrow |\text{Mott}\rangle$ transition is defined for concreteness by

$$|\psi_{\text{ini}}\rangle = |\text{SF}\rangle \equiv |\text{GS}; v_x = 3E_R\rangle, \quad (\text{A7})$$

$$|\psi_{\text{tgt}}\rangle = |\text{Mott}\rangle \equiv |\text{GS}; v_x = 13E_R\rangle, \quad (\text{A8})$$

where GS refers to the ground state at the specified longitudinal depth v_x . The phase transition is driven by varying v_x with the requirement that $v_x \geq 2 E_R$ at all times to satisfy the modeling assumptions in Eqs. (A3) and (A4). With this choice of parameters, the constitutive equations between v_x and the energies U and J_x are calculated numerically and shown in Fig. 6. The conversion to U/J_x for several relevant depths v_x are shown, e.g., at $(v_x)_{\text{crit}} \approx 4.5 E_R$ that in Ref. [22] corresponds to the critical point for the phase transition. At this depth, we obtain $(U/J_x)_{\text{crit}} \approx 3.4$, which agrees with the number stated in Ref. [22], and we consider this a verification for our numerical calculation of $J_x(v_x)$ and $U(v_x, v_y, v_z)$.

APPENDIX B: MATRIX PRODUCT STATES

We briefly present the many-body ansatz of MPS, see Ref. [87] for an excellent and more detailed introduction. We then discuss a t-DMRG (time-dependent density matrix renor-

malization group) algorithm tailored to the necessary problem representation to significantly accelerate computations.

The explosive growth of Hilbert space with the number of constituents is well-known. This exponential scaling, however, is in a sense a convenient illusion since the majority of physically relevant states occupy only a small corner of the full Hilbert space [40]. These are usually characterized by low entanglement, as measured, e.g., by entanglement entropy, and includes ground states and reachable states from these in finite time. Tensor networks and their bespoke algorithms are capable of targeting this much reduced subspace with subexponential resources [6,42,43]. The success of such approaches is owed to the fact that the size of the corner is governed by favorable so-called area scaling laws [88] for the entanglement entropy. Matrix product states, also known as tensor trains, are the appropriate types of tensor networks for 1D systems and their area scaling law is constant with the number of constituents.

The general form of an MPS for a finite unclosed chain of N_s constituents (sites) is

$$|\psi\rangle = \sum_{j_1, j_2, \dots, j_{N_s}} \mathbf{A}^{j_1} \mathbf{A}^{j_2} \dots \mathbf{A}^{j_{N_s}} |j_1, \dots, j_{N_s}\rangle, \quad (\text{B1})$$

where $j_i \in \{1, 2, \dots, d\}$ is the *physical index* (degree of freedom) for the i th constituent, d is the size of the local Fock space, and $\mathbf{A}^{j_i} \in \mathbb{C}^{a_{i-1} \times a_i}$, where a_i is the *bond index* with the only requirement that the product of all the matrices yields a scalar. The ansatz Eq. (B1) is simply a decomposition of the expansion coefficient $c_{j_1, j_2, \dots, j_{N_s}}$ tensor of rank N_s into N_s rank-three tensors $\{A_{a_{i-1}, a_i}^{j_i}\}_{i=1}^{N_s}$, which is always possible by repeated singular value decomposition (SVD) or similar and is in principle an *exact* representation. Properties of, e.g., the SVD procedure, however, allows significant truncation of the matrix \mathbf{A}^{j_i} dimensions associated with the bond indices for low-entanglement states: singular values of the SVD correspond to the expansion coefficients in the Schmidt decomposition across a given bipartitioning of the system, many of which are close to (or exactly) zero for states in the small corner. Thus, we can choose to keep only singular values larger than a given threshold s_{max} and/or impose a maximum number values D to keep, depending on the desired accuracy. Even though the MPS is in practice not constructed directly from the coefficients since it requires exponential amounts of storage, virtually all basic MPS algorithms such as diagonalization (DMRG) and time evolution (t-DMRG) similarly employ SVD (or QR) decompositions. This enables a natural way of keeping resource consumption in check, typically by specifying a given s_{max} and/or D in advance.

In the present case of the Bose-Hubbard model Eq. (11a) with unit filling, j_i is the site occupation number and $d = N_p = N_s$ where N_p is the number of particles. The Hilbert space dimension scales exponentially,

$$D_{\mathcal{H}} = \frac{(N_s + N_p - 1)!}{N_p!(N_s - 1)!}, \quad (\text{B2})$$

which limits the computational feasibility of exact diagonalization approaches to roughly $N_p = 10 - 13$ with increasing layers of analytical and numerical sophistication needed for relatively small gains [74]. At such a low number of sites, the

bulk of the system is constituted by only a relatively small fraction of sites. Matrix product states, on the other hand, are associated with polynomial scaling [40,42,87] and can comfortably extend this range into the low-to-mid tens of particles in a time-dependent setting [21,22,89] or low hundreds in a static setting [90].

t-DMRG for Bose-Hubbard Model

Time evolution is the fundamental operation for quantum optimal control. For this reason, we present here a t-DMRG variant similar to that in Ref. [91] tailored to the structure of Eq. (11a) to speed up our computations.

We start by considering the Suzuki-Trotter expansion [58] in Eqs. (10). Each term in the diagonal control Hamiltonian commutes and we may write exactly

$$\hat{U}_n^{c/2} = \exp\left(-i\left(\frac{u_n}{2} \sum_{i=1}^{N_s} \hat{h}_{[i]}^U\right) \frac{\delta t}{2}\right) = \prod_i \hat{U}_{n,[i]}^U, \quad (\text{B3})$$

where $\hat{U}_{n,[i]}^U = \exp(-iu_n \hat{h}_{[i]}^U \delta t / 4)$. For the drift Hamiltonian, we can apply the same technique as in standard t-DMRG [61,87] for nearest-neighbor Hamiltonians to obtain a first-order Suzuki-Trotter expansion with associated error $O(\delta t^2)$,

$$\begin{aligned} \hat{U}^d &= e^{-i(\hat{H}_{\text{even}}^d + \hat{H}_{\text{odd}}^d)\delta t} \approx e^{-i\hat{H}_{\text{even}}^d \delta t} e^{-i\hat{H}_{\text{odd}}^d \delta t} \\ &= \left(\prod_{i \text{ even}}^{N_s-1} \hat{U}_{[i,i+1]}^{J_x}\right) \left(\prod_{i \text{ odd}}^{N_s-1} \hat{U}_{[i,i+1]}^{J_x}\right), \end{aligned} \quad (\text{B4})$$

where $\hat{U}_{[i,i+1]}^{J_x} = \exp(-i\hat{h}_{[i,i+1]}^{J_x} \delta t)$. The enabling step in this expansion is to group even and odd terms:

$$\hat{H}^d = \hat{H}_{\text{even}}^d + \hat{H}_{\text{odd}}^d = \sum_{i \text{ even}}^{N_s-1} \hat{h}_{[i,i+1]}^{J_x} + \sum_{i \text{ odd}}^{N_s-1} \hat{h}_{[i,i+1]}^{J_x}. \quad (\text{B5})$$

Although $[\hat{H}_{\text{even}}^d, \hat{H}_{\text{odd}}^d] \neq 0$ causes the $O(\delta t^2)$ error, each term has total internal-commutativity, allowing the subsequent exact product form Eq. (B4). Combining the above expressions and moving each individual even (odd) $\hat{h}_{[i,i+1]}^{J_x}$ to the left (right) until they meet a noncommutative operator, we obtain for even N_s

$$\begin{aligned} \hat{U}_n^{\text{ST}} &\approx \prod_i \hat{U}_{n+1,[i]}^U \prod_{i \text{ even}}^{N_s-1} \hat{U}_{[i,i+1]}^{J_x} \prod_{i \text{ odd}}^{N_s-1} \hat{U}_{[i,i+1]}^{J_x} \prod_i \hat{U}_{n,[i]}^U \\ &= \hat{U}_{n+1,[1]}^U \left(\prod_{N_s-1}^{i \text{ even}} \hat{U}_{n+1,[i]}^U \hat{U}_{n+1,[i+1]}^U \hat{U}_{[i,i+1]}^{J_x}\right) \\ &\quad \times \hat{U}_{n+1,[N_s]}^U \left(\prod_{i \text{ odd}}^{N_s-1} \hat{U}_{[i,i+1]}^{J_x} \hat{U}_{n,[i+1]}^U \hat{U}_{n,[i]}^U\right) \\ &\equiv \hat{U}_{n+1,[1]}^U \left(\prod_{N_s-1}^{i \text{ even}} \hat{U}_{n+1,[i,i+1]}^{UUJ_x}\right) \\ &\quad \times \hat{U}_{n+1,[N_s]}^U \left(\prod_{i \text{ odd}}^{N_s-1} \hat{U}_{n,[i,i+1]}^{J_xUU}\right) \end{aligned} \quad (\text{B6})$$

$$\times \hat{U}_{n+1,[N_s]}^U \left(\prod_{i \text{ odd}}^{N_s-1} \hat{U}_{n,[i,i+1]}^{J_xUU}\right) \quad (\text{B7})$$

$$\equiv \hat{U}_{n+1,[1]}^U \hat{U}_{\text{back-sweep}} \hat{U}_{n+1,[N_s]}^U \hat{U}_{\text{forwardsweep}}. \quad (\text{B8})$$

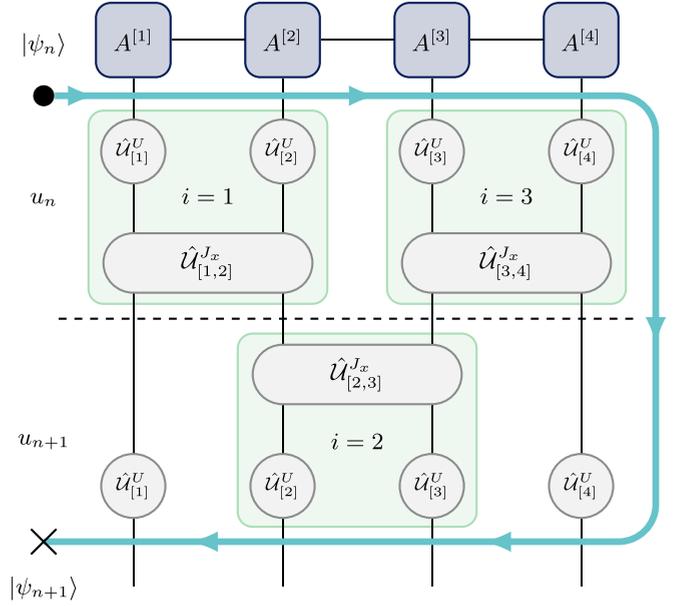


FIG. 7. Example tensor network diagram [87] showing the full calculation $|\psi_{n+1}\rangle = \hat{U}_n^{\text{ST}} |\psi_n\rangle = \hat{U}_{n+1}^c \hat{U}^d \hat{U}_n^c |\psi_n\rangle$ (B8) due to the Hamiltonian in Eq. (11a) for $N_s = 4$.

If N_s is odd, replace $\hat{U}_{n+1,[N_s]}^U \rightarrow \hat{U}_{n,[N_s]}^U$ in the final expressions. In the language of MPS, application of one-site $\hat{U}_{n,[i]}^U$ gates and two-site gates $\hat{U}_{[i,i+1]}^{J_x}$ can be done very efficiently when exploiting left- and right-normalization of the site tensors. The one-site gates are particularly cheap to compute because $\hat{h}_{[i]}^U$ is diagonal. The two-site gates, which would otherwise entail the most expensive operation, are time-independent and can be precomputed, stored on the disk, and be loaded into memory on runtime. Additionally, the grouping of product triples, e.g., $\hat{U}_{n,[i,i+1]}^{UUJ_x}$, acting only on nearest-neighbor pairs of indices $[i, i+1]$ provides a way of reducing overhead in the tensor network contraction $\hat{U}_n^{\text{ST}} |\psi\rangle$ by advancing the central site (gauge) of the MPS: apply the product of triples and contract the site tensors in a forward sweep over odd i (B7) and then in a backward sweep over even i (B6) as illustrated in Fig. 7. A more technical detailing is as follows.

The site tensors (purple nodes) are connected by auxiliary bond indices (black solid horizontal lines). The one- and two-site gates (gray nodes, time index suppressed) come in triples (green boxes) and are applied to physical indices of the site tensors (black solid vertical lines), in the order indicated by the thick teal arrowed line according to the sweeping order. The beginning (end) is marked by a black dot (cross), and the dashed line separates the forward sweep with u_n (above the line) from the backward sweep with u_{n+1} (below the line). Application of product triple i in the forward sweep entails the following:

(1) Contract the two site tensors with physical indices i and $i+1$ over their common bond index into a temporary two-site tensor.

(2) Apply the two one-site gates followed by the two-site gate.

(3) Split the temporary two-site tensor by SVD back into two individual site tensors with the central site (gauge) moved to $i + 1$.

(4) Shift the gauge by SVD an additional site to the right such that the central site is at $i + 2$.

Each arrow tip demarcates the gauge position during the sweeps, where sites to the right (left) are right (left)-normalized, with the exception that the next site intersecting the teal line is the central site. After applying the first leftover one-site gate, the backward sweep is similarly performed with the following modifications: (1) the site indices are $i - 1$ and i , (2) the order of gate application reversed, (3) the central site is placed on $i - 1$, and (4) the central site is gauged to $i - 2$. Finally, the second leftover one-site gate is applied and this completes the time step. The same procedure with new control values u_{n+1} and u_{n+2} can subsequently be applied to obtain $|\psi_{n+2}\rangle = \hat{U}_{n+1}^{\text{ST}} |\psi_{n+1}\rangle$. Implementing the backward propagation $|\psi_n\rangle = \hat{U}_n^{\text{ST}\dagger} |\psi_{n+1}\rangle$ is similar, but with reversed arrow tips and order of application.

APPENDIX C: OPTIMIZATION DETAILS

Our MPS computations are performed using the ITENSOR library [92]. We use an auxiliary dimension of $D = 200$, a singular value threshold of $s_{\text{max}} = 10^{-12}$, and a reduced local Fock space $d = 5$ (higher local occupation numbers do not contribute significantly to the dynamics due to the exponential on-site energy penalty). For reference, the benchmark Refs. [21,22] used $D \leq 100$ and $D \leq 24$, $s_{\text{max}} = 10^{-5}$, respectively, corresponding to less computationally expensive, more approximate, low-entanglement representations of the model. We use the DMRG algorithm implemented in ITENSOR to obtain the initial and target states.

For the given system size of $N_s = N_p = 20$, the durations required to approach the minimal duration for fidelity $F = 0.99$, $T_{\text{min}}^{F=0.99}$, with sufficiently low Trotterization error necessitates about $N_t = 350 - 450$ time steps for time steps of size $\delta t = 0.025$. To accelerate the optimizations, we take δt to be a homotopy, or continuation, parameter [58]: we sequentially optimize on increasingly fine grained time grids, specifically $\delta t = 0.1 \rightarrow 0.05 \rightarrow 0.025$. By halving the values, the new grid points coincide with the old but with doubled resolution as each newly inserted point is set to the value of old point immediately prior corresponding to $\hat{U}_n(\delta t) \approx \hat{U}_n(\delta t/2)\hat{U}_n(\delta t/2)$, where \hat{U}_n is the time evolution operator at time index n . The benefit is that the coarser optimizations can yield relatively rapid fidelity improvements since fine grained resolution is typically not needed for the overall shape of the solution. Care should be taken not to spend too much time on these, since they are not fully coincidental with the final optimization landscape [58]. Note that this technique is enabled by the exactness of the Trotterized gradient not being dependent on δt which is not the case for the exact propagator gradient with finite summation cutoffs, see Ref. [58]. We stress that the exact derivative in Eq. (10a) is the main workhorse whereas the homotopy and time evolution in Appendix B are secondary but effective acceleration techniques.

The SI timescaling in Eq. (A6) depends on J_x^{-1} and thus the control value. Figure 6 shows that larger U/J_x values correspond to longer SI times. Since we desire the

fastest possible optimal controls in real time, we place an upper bound corresponding to $v_x \leq 13.5 E_R$ during optimization to limit this artifact of the nondimensionalization. We also add slight preference toward lower control values by introducing a regularization cost term, J_α , for the control amplitude [Eq. (A28) in Ref. [58]]. Due to the limited bandwidth of experimental electronics we also add a regularization cost term, J_γ , for the temporal derivative of the control, shifting preference toward smoother controls [Eq. (A30) in Ref. [58]]. The strength of these terms are controlled by the parameters $\alpha, \gamma \geq 0$, respectively, and typically $\alpha, \gamma \sim 10^{-7} - 10^{-10}$. The total optimization objective is thus

$$J(\mathbf{u}) = J_F(\mathbf{u}) + J_\alpha(\mathbf{u}) + J_\gamma(\mathbf{u}). \quad (\text{C1})$$

The derivatives for these cost terms are calculated in the Appendix of Ref. [58] and are included in the optimization.

For the optimization [i.e., search direction and step size line searching in Eq. (6)], we employ the nonlinear interior-point algorithm implemented in IPOPT [93] by supplying the exact derivatives. Briefly, interior-point methods can handle control constraints by including them explicitly when solving for the search direction, which in our case is $1.32 \leq u_n \leq 40.18$ for all time indices n . Being a second-order derivative method, the search direction includes the Hessian or a gradient-based approximation thereof (BFGS). We found the exact Hessian calculation [58] (timescale of days per iteration) for the problem under consideration to be outside our time budget even when including the homotopy, and therefore opted for the BFGS approach (timescale of hours per iteration). The seeds (initial points for the optimization) were optimized in parallel on individual cores in a computer cluster. The results reported in the main text were allotted roughly between three to seven

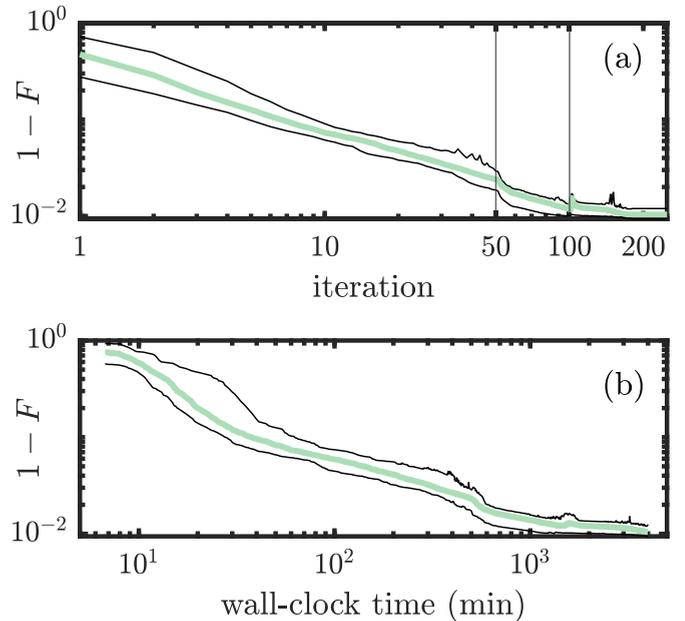


FIG. 8. Optimization trajectories (5%, 50%, 95%) quantiles for $1 - F$ as a function of (a), iteration, and (b), optimization wall-clock time for the 100 seeds at $T = 11$ in Fig. 1(a). The vertical lines mark changes in the δt homotopy parameter.

days of optimization time. Our seeding strategy is based on an exponential reference control overlaid with a sum of random Fourier components. As a verification for our implementation of, e.g., the exact analytical derivatives (gradient and Hessian) and time evolution, we compared the analytical derivatives to their finite difference counterpart and found that they agreed to the same precision as in Ref. [58].

Figure 8 shows optimization trajectory statistics. Changes in the δt homotopy parameter manifest as kinks at 50 and 100 iterations due to δt changes $0.1 \rightarrow 0.05$, and $0.05 \rightarrow 0.025$. A dip in infidelity is seen at the first handover as the increased time resolution of the control allows more complex and fine-tuned dynamics. The homotopy approach accelerates the computations and roughly doubles the number of iterations achieved within the allocated time budget without sacrificing performance since the infidelity iteration trajectories follow roughly the same exponential law for all three homotopy parameter regions.

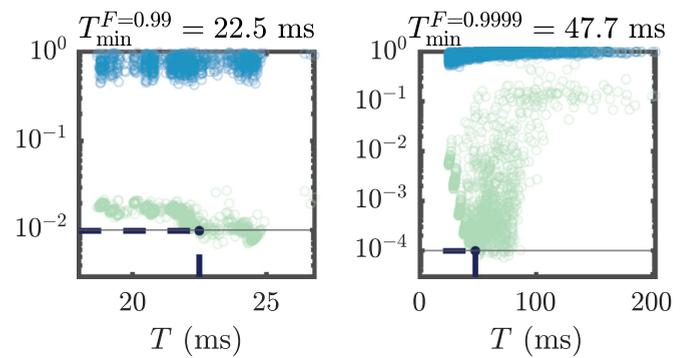


FIG. 9. The $1 - F(T)$ results from Fig. 1 plotted against SI instead of simulation units.

Figure 9 shows the results in Fig. 1 plotted against their SI duration using Eq. (A6).

-
- [1] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko *et al.*, The quantum technologies roadmap: A European community view, *New J. Phys.* **20**, 080201 (2018).
- [2] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen *et al.*, Training Schrödinger's cat: Quantum optimal control, *Eur. Phys. J. D* **69**, 279 (2015).
- [3] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [4] D. J. Egger and F. K. Wilhelm, Optimized controlled-z gates for two superconducting qubits coupled through a resonator, *Supercond. Sci. Technol.* **27**, 014001 (2013).
- [5] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, Charting the circuit QED design landscape using optimal control theory, *npj Quantum Inf.* **3**, 37 (2017).
- [6] S. Montangero, *Introduction to Tensor Network Methods* (Springer International Publishing, Switzerland, Cham, 2018).
- [7] M. Dalgaard, F. Motzoi, J. J. Sørensen, and J. Sherson, Global optimization of quantum dynamics with alphazero deep exploration, *npj Quantum Inf.* **6**, 6 (2020).
- [8] C. T. Kehlet, A. C. Sivertsen, M. Bjerring, T. O. Reiss, N. Khaneja, S. J. Glaser, and N. C. Nielsen, Improving solid-state NRM dipolar recoupling by optimal control, *J. Am. Chem. Soc.* **126**, 10202 (2004).
- [9] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [10] N. C. Nielsen, C. Kehlet, S. J. Glaser, and N. Khaneja, Optimal control methods in NMR spectroscopy, *Encyclopedia of Nuclear Magnetic Resonance* (Wiley, New York, 2010).
- [11] W. Kallies and S. J. Glaser, Cooperative broadband spin echoes through optimal control, *J. Magn. Reson.* **286**, 115 (2018).
- [12] J. J. Sørensen, J. S. Nyemann, F. Motzoi, J. Sherson, and T. Vosegaard, Optimization of pulses with low bandwidth for improved excitation of multiple-quantum coherences in NMR of quadrupolar nuclei, *J. Chem. Phys.* **152**, 054104 (2020).
- [13] J. Scheuer, X. Kong, R. S. Said, J. Chen, A. Kurz, L. Marseglia, J. Du, P. R. Hemmer, S. Montangero, T. Calarco *et al.*, Precise qubit control beyond the rotating wave approximation, *New J. Phys.* **16**, 093022 (2014).
- [14] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen *et al.*, High-fidelity spin entanglement using optimal control, *Nat. Commun.* **5**, 3371 (2014).
- [15] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. Du, P. Neumann *et al.*, Quantum error correction in a solid-state hybrid spin register, *Nature (London)* **506**, 204 (2014).
- [16] Y. Chou, S.-Y. Huang, and H.-S. Goan, Optimal control of fast and high-fidelity quantum gates with electron and nuclear spins of a nitrogen-vacancy center in diamond, *Phys. Rev. A* **91**, 052315 (2015).
- [17] C. P. Koch, J. P. Palao, R. Kosloff, and F. Masnou-Seeuws, Stabilization of ultracold molecules using optimal control theory, *Phys. Rev. A* **70**, 013402 (2004).
- [18] C. P. Koch, E. Luc-Koenig, and F. Masnou-Seeuws, Making ultracold molecules in a two-color pump-dump photoassociation scheme using chirped pulses, *Phys. Rev. A* **73**, 033408 (2006).
- [19] E. F. De Lima, T.-S. Ho, and H. Rabitz, Optimal laser control of molecular photoassociation along with vibrational stabilization, *Chem. Phys. Lett.* **501**, 267 (2011).
- [20] K. M. Tibbetts, X. Xing, and H. Rabitz, Optimal control of molecular fragmentation with homologous families of photonic reagents and chemical substrates, *Phys. Chem. Chem. Phys.* **15**, 18012 (2013).
- [21] P. Doria, T. Calarco, and S. Montangero, Optimal Control Technique for Many-Body Quantum Dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [22] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco *et al.*, Optimal control of complex atomic quantum systems, *Sci. Rep.* **6**, 34187 (2016).

- [23] M. Mundt and D. J. Tannor, Optimal control of interacting particles: A multi-configuration time-dependent Hartree-Fock approach, *New J. Phys.* **11**, 105038 (2009).
- [24] G. Jäger and U. Hohenester, Optimal quantum control of Bose-Einstein condensates in magnetic microtraps: Consideration of filter effects, *Phys. Rev. A* **88**, 035601 (2013).
- [25] J. Cui, R. van Bijnen, T. Pohl, S. Montangero, and T. Calarco, Optimal control of Rydberg lattice gases, *Quantum Sci. Technol.* **2**, 035006 (2017).
- [26] S. Patsch, D. M. Reich, J.-M. Raimond, M. Brune, S. Gleyzes, and C. P. Koch, Fast and accurate circularization of a Rydberg atom, *Phys. Rev. A* **97**, 053418 (2018).
- [27] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi *et al.*, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).
- [28] A. Larrouy, S. Patsch, R. Richaud, J.-M. Raimond, M. Brune, C. P. Koch, and S. Gleyzes, Fast Navigation in a Large Hilbert Space using Quantum Optimal Control, *Phys. Rev. X* **10**, 021058 (2020).
- [29] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153 (2014).
- [30] R. Blatt and C. F. Roos, Quantum simulations with trapped ions, *Nat. Phys.* **8**, 277 (2012).
- [31] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, Cambridge, 2011).
- [32] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Entanglement in many-body systems, *Rev. Mod. Phys.* **80**, 517 (2008).
- [33] C. Kollath, A. M. Läuchli, and E. Altman, Quench Dynamics and Nonequilibrium Phase Diagram of the Bose-Hubbard Model, *Phys. Rev. Lett.* **98**, 180601 (2007).
- [34] J. Eisert, M. Friesdorf, and C. Gogolin, Quantum many-body systems out of equilibrium, *Nat. Phys.* **11**, 124 (2015).
- [35] I. Frérot and T. Roscilde, Quantum Critical Metrology, *Phys. Rev. Lett.* **121**, 020402 (2018).
- [36] L. Garbe, M. Bina, A. Keller, M. G. A. Paris, and S. Felicetti, Critical Quantum Metrology with a Finite-Component Quantum Phase Transition, *Phys. Rev. Lett.* **124**, 120504 (2020).
- [37] A. Osterloh, L. Amico, G. Falci, and R. Fazio, Scaling of entanglement close to a quantum phase transition, *Nature (London)* **416**, 608 (2002).
- [38] G. De Chiara and A. Sanpera, Genuine quantum correlations in quantum many-body systems: A review of recent progress, *Rep. Prog. Phys.* **81**, 074002 (2018).
- [39] M. A. Valdez, D. Jaschke, D. L. Vargas, and L. D. Carr, Quantifying Complexity in Quantum Phase Transitions Via Mutual Information Complex Networks, *Phys. Rev. Lett.* **119**, 225301 (2017).
- [40] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, Quantum Simulation of Time-Dependent Hamiltonians and the Convenient Illusion of Hilbert Space, *Phys. Rev. Lett.* **106**, 170501 (2011).
- [41] G. Vidal, Efficient Simulation of One-Dimensional Quantum Many-Body Systems, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [42] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [43] S. Lloyd and S. Montangero, Information Theoretical Analysis of Quantum Optimal Control, *Phys. Rev. Lett.* **113**, 010502 (2014).
- [44] P. De Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, Second order gradient ascent pulse engineering, *J. Magn. Reson.* **212**, 412 (2011).
- [45] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [46] F. F. Floether, P. De Fouquieres, and S. G. Schirmer, Robust quantum gates for open systems via optimal control: Markovian versus non-markovian dynamics, *New J. Phys.* **14**, 073023 (2012).
- [47] D. Goodwin and I. Kuprov, Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories, *J. Chem. Phys.* **143**, 084113 (2015).
- [48] D. Goodwin and I. Kuprov, Modified Mewton-Raphson grape methods for optimal control of spin systems, *J. Chem. Phys.* **144**, 204107 (2016).
- [49] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [50] J. J. W. H. Sørensen, M. O. Aranburu, T. Heinzl, and J. F. Sherson, Quantum optimal control in a chopped basis: Applications in control of Bose-Einstein condensates, *Phys. Rev. A* **98**, 022119 (2018).
- [51] D. J. Tannor, V. Kazakov, and V. Orlov, Control of photochemical branching: Novel procedures for finding optimal pulses and global upper bounds, in *Time-Dependent Quantum Molecular Dynamics* (Springer, Boston, MA, 1992), pp. 347–360.
- [52] J. P. Palao and R. Kosloff, Quantum Computing by an Optimal Control Algorithm for Unitary Transformations, *Phys. Rev. Lett.* **89**, 188301 (2002).
- [53] S. G. Schirmer and P. de Fouquieres, Efficient algorithms for optimal control of quantum dynamics: The Krotov method unencumbered, *New J. Phys.* **13**, 073029 (2011).
- [54] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [55] X. Li, D. Pęczak, T. Sowiński, J. Sherson, and A. E. B. Nielsen, Global optimization for quantum dynamics of few-fermion systems, *Phys. Rev. A* **97**, 033602 (2018).
- [56] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer Science & Business Media, New York, 2006).
- [57] Z. Ugray, L. Lasdon, J. Plummer, F. Glover, J. Kelly, and R. Martí, Scatter search and local NLP solvers: A multistart framework for global optimization, *INFORMS J. Comput.* **19**, 328 (2007).
- [58] J. H. M. Jensen, F. S. Møller, J. J. Sørensen, and J. F. Sherson, Approximate dynamics leading to more optimal control: Efficient exact derivatives, *Phys. Rev. A* **103**, 062612 (2021).
- [59] M. Dalggaard, F. Motzoi, J. H. M. Jensen, and J. Sherson, Hessian-based optimization of constrained quantum control, *Phys. Rev. A* **102**, 042612 (2020).
- [60] D. Quiñones-Valles, S. Dolgov, and D. Savostyanov, Tensor product approach to quantum control, in *Integral Methods in Science and Engineering: Analytic Treatment and Numerical Approximations*, edited by C. Constanda and P. Harris (Springer International Publishing, Cham, 2019), pp. 367–379.
- [61] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys.* **411**, 167998 (2019).

- [62] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, *Nature (London)* **415**, 39 (2002).
- [63] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen, and U. Sen, Ultracold atomic gases in optical lattices: Mimicking condensed matter physics and beyond, *Adv. Phys.* **56**, 243 (2007).
- [64] B. Damski, The Simplest Quantum Model Supporting the Kibble-Zurek Mechanism of Topological Defect Production: Landau-Zener Transitions from a New Perspective, *Phys. Rev. Lett.* **95**, 035701 (2005).
- [65] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, Theory of quantum annealing of an Ising spin glass, *Science* **295**, 2427 (2002).
- [66] T. Caneva, T. Calarco, R. Fazio, G. E. Santoro, and S. Montangero, Speeding up critical system dynamics through optimized evolution, *Phys. Rev. A* **84**, 012312 (2011).
- [67] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, Optimal Control at the Quantum Speed Limit, *Phys. Rev. Lett.* **103**, 240501 (2009).
- [68] J. V. Porto, S. Rolston, B. Laburthe Tolra, C. J. Williams, and W. D. Phillips, Quantum information with neutral atoms as qubits, *Philos. Trans. R. Soc. A* **361**, 1417 (2003).
- [69] S. Hild, T. Fukuhara, P. Schauß, J. Zeiher, M. Knap, E. Demler, I. Bloch, and C. Gross, Far-From-Equilibrium Spin Transport in Heisenberg Quantum Magnets, *Phys. Rev. Lett.* **113**, 147205 (2014).
- [70] The perhaps more familiar Trotterization $\exp(-i\hat{H}_n^c \delta t/2) \exp(-i\hat{H}_n^d \delta t) \exp(-i\hat{H}_n^c \delta t/2)$ leads to a similar but distinct expression, see Ref. [58].
- [71] An additional factor 1/2 is present for the end points $n = 1, N_t$.
- [72] The name implies that the set of functions becomes a complete basis as $L \rightarrow \infty$, but in numerical practice this is not a necessary requirement.
- [73] M. M. Müller, R. S. Said, F. Jelezko, T. Calarco, and S. Montangero, One decade of quantum optimal control in the chopped random basis, [arXiv:2104.07687](https://arxiv.org/abs/2104.07687).
- [74] J. Zhang and R. Dong, Exact diagonalization: The Bose-Hubbard model as an example, *Eur. J. Phys.* **31**, 591 (2010).
- [75] J. H. M. Jensen, M. Gajdacz, S. Z. Ahmed, J. H. Czarkowski, C. Weidner, J. Rafner, J. J. Sørensen, K. Mølmer, and J. F. Sherson, Crowdsourcing human common sense for quantum control, *Phys. Rev. Research* **3**, 013057 (2021).
- [76] This may be partially due to the regularization discussed in Appendix C not being duration-normalized.
- [77] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal Protocols in Quantum Annealing and Quantum Approximate Optimization Algorithm Problems, *Phys. Rev. Lett.* **126**, 070505 (2021).
- [78] Our chosen target state has a defect density of $\rho \approx 2.8 \times 10^{-6}$.
- [79] T. Caneva, A. Silva, R. Fazio, S. Lloyd, T. Calarco, and S. Montangero, Complexity of controlling quantum many-body dynamics, *Phys. Rev. A* **89**, 042322 (2014).
- [80] The spectral cutoffs were achieved by Fourier transforming the signal, zeroing components with frequency larger than ν_L , and transforming the signal back to the time domain.
- [81] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape, *Phys. Rev. A* **92**, 062343 (2015).
- [82] Z. Leng, P. Mundada, S. Ghadimi, and A. Houck, Robust and efficient algorithms for high-dimensional black-box quantum optimization, [arXiv:1910.03591](https://arxiv.org/abs/1910.03591) (2019).
- [83] D. J. Egger and F. K. Wilhelm, Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems, *Phys. Rev. Lett.* **112**, 240503 (2014).
- [84] N. Wittler, F. Roy, K. Pack, M. Werninghaus, A. S. Roy, D. J. Egger, S. Filipp, F. K. Wilhelm, and S. Machnes, Integrated Tool Set for Control, Calibration, and Characterization of Quantum Devices Applied to Superconducting Qubits, *Phys. Rev. Appl.* **15**, 034080 (2021).
- [85] M. Weidemüller and C. Zimmermann, *Interactions in Ultracold Gases: From Atoms to Molecules* (John Wiley & Sons, Darmstadt, 2011).
- [86] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, Theory of bose-einstein condensation in trapped gases, *Rev. Mod. Phys.* **71**, 463 (1999).
- [87] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [88] D. Pérez-García, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, *Quantum Inf. Comput.* **7**, 401 (2007).
- [89] L. Kohn, P. Silvi, M. Gerster, M. Keck, R. Fazio, G. E. Santoro, and S. Montangero, Superfluid-to-mott transition in a Bose-Hubbard ring: Persistent currents and defect formation, *Phys. Rev. A* **101**, 023617 (2020).
- [90] S. Ejima, H. Fehske, F. Gebhard, K. zu Münster, M. Knap, E. Arrigoni, and W. von der Linden, Characterization of Mott-insulating and superfluid phases in the one-dimensional Bose-Hubbard model, *Phys. Rev. A* **85**, 053644 (2012).
- [91] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces, *J. Stat. Mech.: Theory Exp.* (2004) P04005.
- [92] ITensor Library (version 2), <http://itensor.org>.
- [93] A. Wächter and L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, *Math. Program.* **106**, 25 (2006).