

Quantum optimal control in a chopped basis: Applications in control of Bose-Einstein condensates

J. J. W. H. Sørensen, M. O. Aranburu, T. Heinzl, and J. F. Sherson*

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

(Received 26 March 2018; published 14 August 2018)

We discuss quantum optimal control of Bose-Einstein condensates trapped in magnetic microtraps. The objective is to transfer a condensate from the ground state to the first-excited state. This type of control problem is typically solved using derivative-based methods in a high-dimensional control space such as gradient-ascent pulse engineering (GRAPE) and Krotov's method or derivative-free methods in a reduced control space such as Nelder-Mead with a chopped random basis (CRAB). We discuss how these methods can be combined in gradient optimization using parametrization (GROUP) including the finite bandwidth of the control electronics. We compare these methods and find that GROUP converges much faster than Nelder-Mead with CRAB and achieves better results than GRAPE and Krotov's method on the control problem presented here.

DOI: [10.1103/PhysRevA.98.022119](https://doi.org/10.1103/PhysRevA.98.022119)**I. INTRODUCTION**

Technological advances in the experimental toolbox in physical chemistry and atomic, molecular, and optical physics currently enable exciting new developments in the manipulation of complex quantum systems. Gradually, the focus is shifting from verifying the validity of theoretical models towards controlling and manipulating quantum systems for specific technological applications [1,2]. Some examples of this trend are quantum state preparation [3,4], atomic clocks [5], quantum based sensors [6–8], quantum simulators [9], and quantum computers [10,11].

These applications require the ability to steer the quantum dynamics precisely using external control fields. The quality of a control is measured by a cost function, which can describe, e.g., distance to some target state [12], similarity with a unitary gate operator [13], or the amount of experimental signal [14]. Quantum optimal control (QOC) is a framework that enables the design of control strategies that achieve the desired dynamics [15–17]. QOC has been studied in a wide range of physical systems [12,18–21]. Central in QOC are local optimization algorithms that maximize or minimize the cost function. These local algorithms can broadly be divided along two axes [22]. The first axis is derivative-based versus derivative-free algorithms and the second is optimization in a full or reduced basis for admissible controls. It is also possible to extend the local algorithms with global optimization [22].

Derivative-free algorithms operate by directly evaluating the cost functional in carefully selected points. A particularly prevalent example in QOC is the Nelder-Mead algorithm [21,23–25]. More recently, it has been proposed to use gradient-free methods in QOC such as Brent's principal axis [26]. On the other hand, derivative-based methods use both the functional and derivative information, which can speed up the convergence rate. Important examples from QOC are the

gradient-ascent pulse engineering (GRAPE) [27] and Krotov's method [28,29].

In a reduced-basis method, the control space is parametrized by a set of smooth functions [21]. This is motivated by the fact that the high-quality solutions often lie in some lower-dimensional subspace of the full control space [30], so a proper low-dimensional parametrization of the control space can *a priori* steer the optimization in the correct direction. A prominent example is chopped random basis (CRAB) [21,23]. The introduction of a random chopped basis is based on the idea that randomness allows one to explore a range of different bases, which can lead to improved results. A too-low-dimensional parametrization can introduce artificial traps since the parametrization is no longer able to adequately describe the optimal solutions [31]. Within QOC, reduced-basis methods are typically only used in combination with derivative-free algorithms such as Nelder-Mead [21,23–25,32].

Ideally, one should combine both of these two approaches and perform derivative-based optimization in a reduced basis. In the past years, steps have been taken to implement this in a number of different ways [33–37]. We will collectively refer to such methods as gradient optimization using parametrization (GROUP). In part due to the frequent success of derivative-free search in a reduced basis, GROUP-type methods have not been widely adopted in the QOC community. However, in recent years, there has been a growing awareness that the search for time-optimal solutions in quantum engineering leads to an exponentially growing computational complexity [22,38–41], which will necessitate algorithmic improvements. In this work, we present the GROUP formalism and show how to combine GRAPE- and CRAB-type methods. In addition, we clearly demonstrate the efficiency of the GROUP methodology to the field of quantum engineering by applying it to a high-profile challenge in the field and presenting a detailed comparison of performance with conventional algorithms. In this work, we also demonstrate how the GROUP formalism can include a filter function on the controls.

This paper discusses quantum control of Bose-Einstein condensates manipulated in a magnetic microtrap. These systems

*sherson@phys.au.dk




FIG. 1. Comparison of gradient-based and gradient-free optimization methods in an artificial landscape. The shaded blue triangles show the gradient-free method, Nelder-Mead. The solid orange line and the dotted red line are gradient-based algorithms steepest descent and BFGS, respectively. The steepest descent exhibits the characteristic zigzag-type behavior, which BFGS avoids due to the inverse Hessian approximation. Nelder-Mead, steepest descent, and BFGS, respectively, use 45, 34, and 15 iterations for convergence.

are particularly challenging to control due to the nonlinearity in the equations of motion. We will investigate fast excitation from the ground state to the first-excited state in a single well. This system has attracted much attention in QOC and it has been investigated both experimentally and theoretically [12,24,42,43]. Due to the interest in this system, we see it as an ideal test bed for quantum control algorithms.

Here we present a direct comparison of derivative-based search in a reduced basis (GROUP) with the other state-of-the-art algorithms, i.e., Krotov's method, Nelder-Mead using CRAB, and GRAPE. We find that GROUP not only converges faster, but also achieves better end fidelity. This work supplements and extends previous efforts in comparing QOC algorithms [43,44].

Derivative-based methods generally converge faster than derivative-free [44,45]. Derivative-based methods typically employ either steepest descent or the quasi-Newton method of Broyden-Fletcher-Goldfarb-Shanno (BFGS) [45]. BFGS is typically faster than steepest descent since it avoids steepest descent's characteristic inefficient zigzag motion close to the optimum [44,45]. BFGS achieves this by gradually building an approximation of the Hessian using the past gradients [45]. In this paper, we only use BFGS-type descent with GRAPE and GROUP. A graphical illustration of the difference in convergence behavior between a derivative-free Nelder-Mead and the two derivative-based methods is given in Fig. 1.

This paper is organized as follows. In Sec. II, we introduce the control problem of transferring a condensate from the ground state to the first-excited state. In Sec. III, we present the GROUP methodology and the other different quantum control algorithms. In Sec. IV, we present how to include the finite bandwidth of the control electronics. In Sec. V, we compare the different algorithms. Finally, we conclude the paper in Sec. VI.

II. THE CONTROL PROBLEM

We will discuss the manipulation of a condensate trapped on an atom chip [4]. In the experiment described in Refs. [4,42], a source for stimulated emission of matter waves in twin beams is created by transferring a condensate into the collective first-excited state. The typical decay rate of the system is 3 ms, so it is very important to find optimal controls that can transfer the condensate into the excited state faster than this decay rate and still allow time for subsequent experiments [4,24]. This condensate driving control problem has been investigated using a number of different QOC algorithms in Refs. [12,24,43]. We give a brief account of the condensate driving control problem and, for more details on the experimental setup, we refer to Refs. [4,24,42,46]. In a mean-field treatment, the dynamics of the condensate are well described by an effective one-dimensional Gross-Pitaevskii equation (GPE),

$$i \frac{\partial \psi}{\partial t} = -\frac{1}{2m} \frac{\partial^2 \psi}{\partial x^2} + V(x, u) \psi + \beta |\psi|^2 \psi = (\hat{H} + \beta |\psi|^2) \psi, \quad (1)$$

where $\hbar = 1$, β is the effective nonlinear self-interaction, and \hat{H} is the Hamiltonian. The atom-chip experiment is tightly confined along two transverse directions and weakly confined along the axial direction. One of the transverse directions is so strongly confining that the state is frozen into the ground state. The dynamics along the axial direction is slow compared with the other transverse direction, which allows for a description with an effective one-dimensional GPE with a nonlinear dependence of β on the atom number [24,47]. For 700 atoms, one finds $\beta = 1830 \hbar \text{ Hz } \mu\text{m}$ [24].

The potential is an anisotropic Ioffe-Pritchard trap dressed by a radio-frequency potential [24,42,46]. In Ref. [24], this potential is approximated by a polynomial,

$$V(x, u(t)) = p_2 [x - u(t)]^2 + p_4 [x - u(t)]^4 + p_6 [x - u(t)]^6, \quad (2)$$

where the control $u(t)$ is the displacement of the trap. The coefficients are given by $p_2 = 2\pi \hbar \times 310 r_0^{-2} \text{ Hz}$, $p_4 = 2\pi \hbar \times 13.6 r_0^{-4} \text{ Hz}$, and $p_6 = -2\pi \hbar \times 0.0634 r_0^{-6} \text{ Hz}$, with $r_0 = 172 \text{ nm}$ [24]. For an in-depth discussion of the experimental setup, we refer the reader to Refs. [4,24]. The goal is to transfer the initial state ψ_0 into the target state ψ_t after a duration of T . The initial state is the ground state for $u(t=0)=0$ and the target state is the first-excited state for $u(t=T)=0$.

In order to find a control that transfers the initial state into the target state, the problem is expressed as a minimization of the cost functional $J(u, \psi)$,

$$J(u(t), \psi) = \frac{1}{2} [1 - |\langle \psi_t | \psi(T) \rangle|^2] + \frac{\gamma}{2} \int_0^T \dot{u}(t)^2 dt. \quad (3)$$

In the first term, $F = |\langle \psi_t | \psi(T) \rangle|^2$ is the fidelity and $1 - F$ is the infidelity, which quantifies the difference between the final state and the target state [12,34,43,48]. The second term in Eq. (3) is the regularization that penalizes strong fluctuations in the control, which accounts for the fact that very fast changes cannot be realized experimentally. We found that $\gamma = 1 \times 10^{-6}$ gives acceptably smooth solutions.

Here we investigate the control problem at $T = 1.09$ ms, which was reported in Ref. [24] to be the value of the quantum speed limit (QSL). The QSL is the lowest-duration T where solutions above $F \geq 0.99$ can be found. In recent work, it has been shown that constrained control problems such as condensate driving become NP-hard close to the QSL [39–41]. Hence, comparing the performance of the quantum control algorithms close to the QSL provides a stringent test of their individual performance. A solution at the QSL is also a solution at all longer durations since the target state is an eigenstate of the GPE. Attempting to solve the problem at shorter durations can give improvements to the estimate of the QSL, but it also requires systematic global exploration, which is discussed for the condensate driving problem in Ref. [22].

III. QUANTUM OPTIMAL CONTROL

In QOC, the goal is to minimize the cost functional $J(u, \psi)$ while satisfying the constraints from the GPE. The constraint can be handled using a Lagrange multiplier,

$$\mathcal{L}(\psi, u, \chi) = J + \text{Re} \int_0^T \langle \chi | i\dot{\psi} - \hat{H}\psi - \beta|\psi|^2\psi \rangle dt, \quad (4)$$

where χ is a Lagrange multiplier [17]. Note that the usual linear Schrödinger equation is a special case for $\beta = 0$, so the methodology we present here can be directly extended to such systems. At a local minimum, all three variational derivatives $D_{\delta\chi}\mathcal{L}$, $D_{\delta\psi}\mathcal{L}$, and $D_{\delta u}\mathcal{L}$ are zero. This gives the following optimality equations (see Appendix 1 for details):

$$i\dot{\psi} = \hat{H}\psi + \beta|\psi|^2\psi, \quad (5)$$

$$i\dot{\chi} = (\hat{H} + 2\beta|\psi|^2)\chi + \beta\psi^2\chi^*, \quad (6)$$

$$\gamma\ddot{u} = -\text{Re}\left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle, \quad (7)$$

and the associated boundary conditions,

$$\psi(0) = \psi_0, \quad (8)$$

$$i\chi(T) = -\langle \psi_T | \psi(T) \rangle \psi_T, \quad (9)$$

$$u(0) = u_0, \quad u(T) = u_T. \quad (10)$$

Ideally, these equations would be solved analytically, which would directly give the optimal solutions. Unfortunately, such analytic solutions are, in general, infeasible so it is necessary to use iterative numerical algorithms [12,48].

When performing the numerical optimization of the cost functional, it is necessary to discretize the control in steps Δt , where Δt is set by the required accuracy when numerically solving the GPE. This means that the control becomes represented by a vector of length $N = \lfloor T/\Delta t \rfloor + 1$. Typically, in GRAPE and Krotov's method, the dimension for the space of admissible control (M) is the same as for the simulation $N = M$. However, with a proper change of basis, the optimal controls could adequately be described by another basis with much smaller dimension [30,37]. In this sense, setting $N = M$ often introduces too many degrees of freedom for the control problem. In the simulations performed here, we have

$N \simeq 3500$, whereas by a proper choice of basis, we have only $M \simeq 50$. These potential large reductions in the dimension of the control problem are the motivation for reduced or chopped basis methods. In recent years, CRAB has emerged as an attractive alternative to the conventional QOC methods since it parametrizes the control in a reduced basis [21,23–25].

As mentioned, the main purpose of this paper is to introduce and numerically test GROUP. We first give a brief review of GRAPE and CRAB since GROUP builds on these methods. When discussing these methods, one typically works with the reduced cost functional $\hat{J}(u) = J(u, \psi_u)$. Here, ψ_u is the unique solution to the GPE, which is found by solving Eq. (1) with $u(t)$ [34]. Our objective is to find a minimum, and preferably the global minimum, of $\hat{J}(u)$.

A. GRAPE

A standard approach to minimizing $\hat{J}(u)$ is the GRAPE algorithm [27,43,48]. The simplest version of this method is to update the control along the gradient,

$$u^{(i+1)} = u^{(i)} - \alpha^{(i)} \nabla \hat{J}(u^{(i)}), \quad i = 0, 1, 2, \dots \quad (11)$$

Here the index i refers to the current iteration. This is the steepest descent method. An appropriate value for α is found using a step-size algorithm, which finds a satisfactory solution to the one-dimensional optimization problem,

$$\alpha^{(i)} = \arg \min_{\alpha} \hat{J}[u^{(i)} - \alpha \nabla \hat{J}(u^{(i)})]. \quad (12)$$

Note that in Eq. (11), the control is updated for all times $0 \leq t \leq T$ at once. This makes GRAPE a concurrent method, which is different from Krotov's method presented later where the control is updated sequentially for each time slice.

An important but subtle point is the use of the gradient $\nabla \hat{J}(u^{(i)})$ in Eq. (11). The complication arises from the fact the gradient is defined in a function space X . In this space, the gradient is the unique element such that $(\nabla \hat{J}, \delta u)_X = D_{\delta u} \hat{J}$ for all possible variations δu . Hence, the gradient depends on the choice of the inner product for the function space X . This has already been discussed in Refs. [34,48]. Specifically, all variations δu should satisfy the boundary condition in Eq. (10) giving $\delta u(0) = \delta u(T) = 0$. A choice that will fulfill this requirement is the H^1 -space with the inner product $(u, v)_{H^1} = \int_0^T \dot{u} \dot{v} dt$. As shown in Appendix 2, this choice gives the result

$$\frac{d^2}{dt^2} [\nabla \hat{J}(u)] = \text{Re}\left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u}, \quad (13)$$

where ψ and χ are the solutions of Eqs. (5) and (6). This is a Poisson equation for the control in time so we can choose the Dirichlet boundary conditions $[\nabla \hat{J}(u)](0) = \nabla[\hat{J}(u)](T) = 0$. These conditions imply that GRAPE preserves the boundary condition in Eq. (10) in any iteration. As pointed out in Refs. [34,48], if we had made the canonical choice of $X = L^2$, then the gradient would not vanish at the boundary. In this case, the boundary condition must be enforced numerically by simply setting the gradient to be zero at the boundary, which can greatly decrease the performance of the algorithm.

The update written in Eq. (11) is the steepest descent algorithm, which is also illustrated in Fig. 1. As shown in this

figure, the optimization can be improved by using a quasi-Newton method such as BFGS instead of steepest descent. Again, special care has to be taken when working in H^1 space, and we use the matrix-free version of L-BFGS described in Ref. [34].

B. Chopped basis and CRAB

As discussed above, the control is represented by a vector of length $N = \lceil T/\Delta t \rceil + 1$ due to the temporal discretization. In GRAPE, the dimensionality of the space of admissible controls (M) is $N = M$. However, by expanding the control in a proper basis, it is possible to substantially reduce the dimension of the control problem,

$$u(t) = u_0(t) + S(t) \sum_{n=1}^M c_n f_n(t), \quad (14)$$

where the f_n 's are the basis functions. $S(t)$ is a shape function that enforces the boundary condition in Eq. (10), which gives $S(0) = S(T) = 0$. In this chopped basis, we optimize the coefficients c_n instead of full control $u(t)$, which implies that the cost function is $\hat{J}(\mathbf{c})$ where $\mathbf{c} = (c_1, c_2, \dots, c_M)$. This method is known as a chopped basis (CB). The functions f_n must be chosen sensibly based on physical insight, which would typically be sinusoidal functions around characteristic frequencies. The reduced dimension of $\hat{J}(\mathbf{c})$ enables the use of the gradient-free Nelder-Mead algorithm. Gradient-free methods have the advantage that there is no need to implement code that calculates the gradient by solving, e.g., Eq. (13), which also requires solving the equation for the Lagrange multiplier given by Eq. (6). This is particularly an advantage whenever calculation of the gradient is infeasible or too resource consuming [23].

For the control problem discussed here, we found the following expansion useful:

$$u(t) = u_0(t) + S(t) \sum_{n=1}^M c_n \sin\left(\frac{\omega_n t}{T}\right), \quad (15)$$

where $\omega_n = n\pi$ is a set of frequencies. This type of chopped basis was extended in Ref. [21] by introducing the chopped random basis or CRAB. In CRAB, the frequencies ω_n are randomly shifted as $\omega_n = (n + r_n)\pi$, where $-0.5 \leq r_n \leq 0.5$ are initially chosen random numbers. The optimization is repeated a number of times with different values of r_n 's. This is a central idea in CRAB since it allows the algorithm to explore different basis functions with slightly similar frequencies starting from the same u_0 . An optimization within a CRAB can principally be done using any method, but it is typically done using the gradient-free method Nelder-Mead [21,23,24].

C. GROUP

In GROUP, the best features of the two previous methods are combined. We parametrize the control in some basis as in Eq. (14). However, instead of using a gradient-free method to search in the chopped basis, we use the gradient-descent methods from GRAPE. Here the gradient is with respect to the expansion coefficients $[\nabla \hat{J}(\mathbf{c})]$. Given this gradient, an iterative update analogous to Eq. (11) can be directly applied.

Just as in GRAPE, we need to find an analytic expression for the gradient similar to Eq. (13). The partial derivative of $\hat{J}(\mathbf{c})$ with respect to c_n can be found using the chain rule for variational derivatives (see Appendix 3 for more details). The result is

$$\begin{aligned} \frac{\partial \hat{J}(\mathbf{c})}{\partial c_n} &= D_{S(t)f_n(t)} \hat{J}(u) \\ &= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) S(t) f_n(t) dt. \end{aligned} \quad (16)$$

Note that these partial derivatives only differ in the $f_n(t)$ function in the integrand. This expression is valid for any CB and CRAB. The quantity in the bracket only needs to be computed once, which contains ψ and χ that are found by the numerically expensive solution of Eqs. (5) and (6). The cost of calculating all the partial derivatives and hence the full gradient $[\nabla \hat{J}(\mathbf{c})]$ is dominated by the solution of Eqs. (5) and (6), which implies that the time needed for calculating the gradient in GROUP is comparable to GRAPE.

In the comparative numerical studies presented below, we performed the GROUP optimization using the quasi-Newton method BFGS. Note that when optimizing $\hat{J}(\mathbf{c})$, the optimization is done in the usual l^2 -space and not H^1 -space, so all the standard methods for BFGS can be directly applied.

All numerical GROUP results in this paper were obtained using the formalism from Eq. (16) extended with a filter function. In Ref. [33], it was proposed to compute the gradient in a reduced basis $\partial \hat{J}(\mathbf{c})/\partial c_n$, but using another method named gradient optimization of analytic controls (GOAT). We give a brief account of how this method can be extended to the condensate driving control problem in Appendix 4. Briefly, the GOAT algorithm is more numerically expensive than GROUP, but could potentially offer advantages if ultralow infidelities are required.

D. Krotov's method

Krotov's method is an alternative to the standard Lagrange multiplier method used in GRAPE and GROUP [49]. In Krotov's method, the cost functional given by Eq. (3) is rewritten so the GPE appears explicitly and conditions for a guaranteed decrease in the cost are directly built in [28,29,43,49]. This allows Krotov's method to give an optimal control algorithm that directly ensures a monotonic decrease in the cost and it is expected to provide fast convergence. The resulting update for the control is (see Appendix 5 for more details)

$$u^{(i+1)}(t) = u^{(i)}(t) + \alpha S(t) \text{Re} \left\langle \chi^{(i)}(t) \left| \frac{\partial \hat{H}}{\partial u} \right|_{u^{(i)}} \psi^{(i+1)}(t) \right\rangle. \quad (17)$$

Here, $\psi^{(i)}$ and $\chi^{(i)}$ are the solutions to the GPE (6) and the Lagrange multiplier (5) from GRAPE. $S(t)$ is again a shape function $0 \leq S(t) \leq 1$ that turns the control off at $t = 0$ and $t = T$, which ensures that the boundary condition for the control is always satisfied in Eq. (10). α is the step-size parameter that must be selected for proper convergence. Note that both the current iteration index i and the next index $i + 1$ appear in the equation. Specifically, the control at the next iteration $u^{(i+1)}$ depends on the states in the next iteration $\psi^{(i+1)}$. This makes Krotov's method a sequential algorithm where

the next control $u^{(i+1)}$ is being calculated while the equations of motion are being solved along that control. This is very different from the other methods presented here, where the control is updated concurrently for all values of $0 < t < T$. Note that since the GPE is nonlinear in the states, a monotonic decrease in the cost can only be guaranteed if Eq. (17) includes an additional term that is proportional to the difference in the states between iterations [28,29,43]. However, for the small values of β discussed here, neglecting this term does not notably affect the monotonic decrease in the cost [43]. The derivative is typically with respect to the next iteration ($u^{(i+1)}$), but for the small values of α used here it is acceptable to use the current iteration ($u^{(i)}$) [29,43].

IV. FILTER FUNCTION

In order to obtain a close match to the experimental conditions in the condensate driving control problem, it is also necessary to include the finite bandwidth of the control electronics. This effect causes the control $u(t)$ to become distorted into $v(t)$ and the atoms experience the potential from $V(x, v)$ [24]. The distortion is large enough to cause the fidelity to drop by a couple of percent, so it must be included in the modeling. It has previously been discussed how to include this type of effect into GRAPE [50] and Krotov-type methods [51,52]. The distorted control is given by a convolution with the filter $h(\tau)$,

$$v(t) = (h * u)(t) = \int_0^t h(\tau)u(t - \tau)d\tau. \quad (18)$$

The presence of the filter changes the expressions for the gradients (see Appendices 2 and 3 for more details). The GRAPE gradient [Eq. (13)] becomes

$$\begin{aligned} \frac{d^2}{dt^2}[\nabla \hat{J}(u)] = & -\gamma \dot{v}(T)h(T - t) \\ & + \int_t^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle + \gamma \ddot{v} \right) h(t' - t) dt'. \end{aligned} \quad (19)$$

Similarly, the expression for the GROUP gradient becomes

$$\begin{aligned} \frac{\partial \hat{J}(\mathbf{c})}{\partial c_n} = & - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle + \gamma \ddot{v} \right) (h * S f_n)(t) dt \\ & + \gamma \dot{v}(T)(h * S f_n)(T). \end{aligned} \quad (20)$$

It is not straightforward to include such a filter function in Krotov's method [51]. For the comparison below, the simulations for Krotov are performed without the filter [51]. The filter gives rise to another potential complication, which is that although $u(T) = 0$, it can occur that $v(T) \neq 0$. This could slightly perturb the final state since the control cannot be instantaneously quenched to zero due to the filter. For the optimal solutions found here, this effect did not notably affect the fidelity.

V. NUMERICAL RESULTS

In the last section, we gave an introduction to four different quantum control algorithms. In this section, we will compare

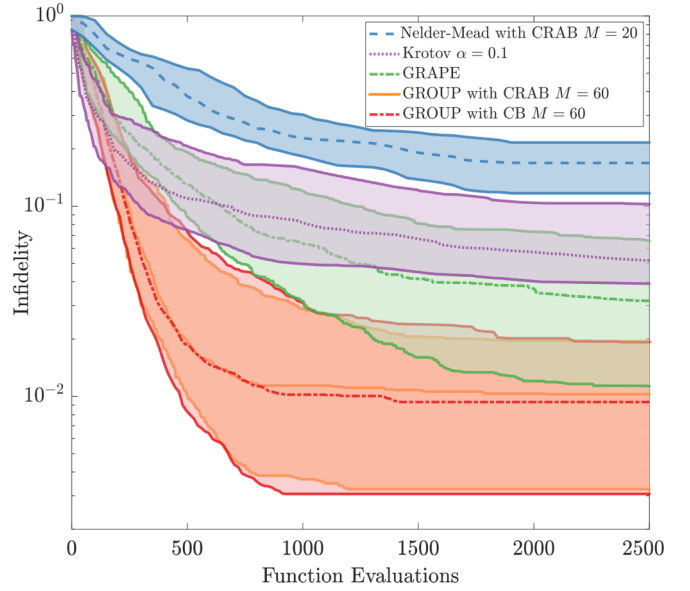


FIG. 2. The infidelity as a function of the number of evaluations for each algorithm. One function evaluation is a solution of the GPE or Lagrange multiplier, given by Eq. (6). The different algorithms are shown at the basis size or steps size where they performed the best (see legend). The dotted line shows the median and the shaded area indicates the 25% and 75% quartiles found from 100 different random initial controls. The quasi-Newton method BFGS was used together with GRAPE and GROUP.

them numerically when applied on the condensate driving control problem.

A. Convergence behavior

We applied the following QOC algorithms: GRAPE, Krotov, Nelder-Mead with CRAB, and GROUP with CB and CRAB, to the condensate driving control problem. We applied the algorithms to the same 100 initial controls, which were randomly generated using Eq. (14). The convergence behavior of the different methods is illustrated in Fig. 2. Here the median and 25% and 75% quartiles are shown for the different algorithms, which gives an impression of the expected behavior for each method on this problem. One function evaluation is a solution of the GPE or the equation for the Lagrange multiplier [Eq. (6)].

Throughout the optimization, GROUP achieves the lowest infidelities. At the end of the optimization, GROUP has the best infidelity, followed by GRAPE and Krotov. There is no particular difference between GROUP using CB or CRAB so we will refer to them both as GROUP. Nelder-Mead using CRAB has the slowest convergence rate of the four methods since it does not utilize derivative information. This is in accordance with the picture presented in Fig. 1, which shows that derivative-based methods are typically faster than derivative-free.

The optimization curves in Fig. 2 can be split into two regimes: one at a high number of function evaluations after 600 and one below. Below 600 evaluations, the three derivative-based methods have similar rates of convergence, but GROUP and Krotov perform better than GRAPE. In the high number of function evaluations regime, GROUP performs better than




FIG. 3. (a) The best-case controls found using GROUP with CB (red) and Nelder-Mead with CRAB (blue) after 2500 iterations. The control is held constant after the vertical dashed line. (b) The density for the condensate ($|\psi(x, t)|^2$) found when using the solution from GROUP. After the vertical line, the density is constant since the state has converged with $F = 0.999$ to the first-excited state. (c) The same is shown as in (b) for Nelder-Mead with CRAB, which has residual oscillations since the final fidelity is $F = 0.95$.

GRAPE and Krotov. We attribute this to the fact that the basis gradually steers GROUP towards a more profitable part of the optimization landscape. This shows the complexity of the optimization landscape since, although the algorithms start at the same point and perform local greedy optimization, they converge towards different points with different infidelities. This is different from the situation in Fig. 1 where all the algorithms converge to the same point since the two-dimensional landscape is much simpler than the control problem, which has dimension $M \simeq 3500$ or $M \simeq 50$.

In the high function evaluations regime, Krotov and GRAPE switch places and GRAPE finds better infidelities than Krotov. These results show that Krotov achieves fast initial reductions in the infidelity, but it slows down considerably when approaching the optimum. GRAPE does not exhibit this behavior, which we attribute to the Hessian approximation from BFGS since the cost function can be well described by a second-order expansion close to the optimum [45]. Similar results were also reported in Ref. [43]. In principle, Krotov's method can also be combined with a BFGS-type method [53], but it has been reported in Refs. [43,53] that this does not significantly improve the convergence.

This comparison focuses on the expected behavior. Each of the methods have a few optimization runs that perform significantly better, which reflects that each algorithm has some specific seeds where it just happens to search the optimization landscape in the most favorable manner. GROUP had the individual optimization runs with the lowest infidelity.

The best controls found for GROUP with CB and Nelder-Mead with CRAB are shown in Fig. 3 with, respectively, $F = 0.999$ and $F = 0.95$. This figure also shows the density for the




FIG. 4. (a) The final infidelity after 2500 function evaluations for GROUP with CB (red) and Nelder-Mead with CRAB (blue) for different basis sizes. The dotted line is the median and the shaded area displays the 25% and 75% quartiles found from the 100 different random initial controls. (b) The same as (a) for Nelder-Mead and GROUP with dressed CRAB. The gray plot shows the final infidelity for Nelder-Mead with dressed CRAB after 25 000 iterations rather than 2500 iterations.

condensate when propagated along the optimal controls. After T , the control is held constant at $u = 0$. The Nelder-Mead with CRAB solution has a residual oscillation after T due to residual excited-states' components in the solution.

B. Reduced-basis size and dressed methods

A limitation of the reduced-basis methods is that the optimization might get caught in an artificial trap introduced by limited bandwidth of the parametrization. As an example, consider the expansion in Eq. (15) at some fixed M . If the optimal solution requires frequencies above M , then the optimization can never converge to this solution and it will get caught in an artificial trap. This effect would favor larger values of M . On the other hand, increasing M also increases the dimension of the control problem, which is exactly what the reduced-basis methods attempt to avoid.

In Fig. 4(a), we compare the trade-off between the risk of artificial traps and retaining a low dimension for Nelder-Mead with CRAB and GROUP with CB. As expected, a too small basis $M \leq 10$ gives poor results for both algorithms since the optimal controls cannot be adequately described with these small basis sizes. Nelder-Mead with CRAB clearly has an optimal basis size around $M = 20$ and it becomes worse with a larger dimension. This is due to the fact that the Nelder-Mead algorithm cannot effectively search within a large dimension. Surprisingly, the performance in GROUP is very robust with respect to the basis size, and it even seems to prefer large basis sizes. We attribute this result to the fact that the large frequencies have small weights in the gradient, meaning that they contribute little to the search. This type of behavior is

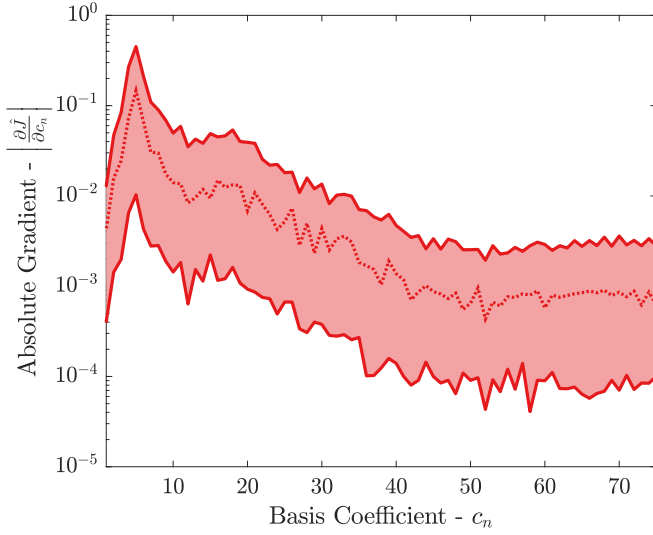


FIG. 5. The absolute value of the partial derivative with respect to basis coefficients $|\partial \hat{J}(\mathbf{c})/\partial c_n|$ found using Eq. (16). The dotted line shows the median of the 100 gradients taken at the start of the optimization. The shaded area indicates the 25% and 75% quartiles.

illustrated in Fig. 5, which shows the median initial gradient for the 100 initial points. Above $n = 50$, the partial derivatives are much smaller than for $n = 5$, and hence it is the low frequencies that dominate the search.

A solution to the artificial trap problem was proposed in Ref. [31] where a so-called dressed CRAB (dCRAB) was introduced. In dCRAB, the optimization is restarted with new basis functions from the last optimum ($u^{(j-1)}$),

$$u^{(j)} = u^{(j-1)} + S(t) \sum_{n=1}^M c_n^{(j)} f_n^{(j)}(t, r_n^{(j)}). \quad (21)$$

These iterations are known as superiterations [31]. The new basis functions $f_n^{(j)}(t, r_n^{(j)})$ are found by reselecting the random frequency shifts (r_n) in CRAB, giving new basis functions in each superiteration. The initial coefficients are $c_n^{(j)} = 0$ for all n so the optimization starts from the previous control. The new basis functions give the algorithm the possibility to escape the artificial traps [31]. Clearly, this formalism cannot be used in a normal chopped basis (CB). We apply the same methodology to GROUP by simply following the same procedure for the superiterations and calculate the gradients using Eq. (20).

In Fig. 4(b), we compare the effect of dressing for different basis sizes. GROUP in combination with dCRAB (dGROUP) only does slightly better than GROUP with CB, which indicates that the solutions were already highly optimal after 2500 evaluations. We observe the same behavior for Nelder-Mead with dCRAB after 2500 evaluations. However, if Nelder-Mead with dCRAB is allowed to perform 25 000 evaluations, then it finds much better results especially for basis sizes around $M = 20$. This shows that for low basis sizes, Nelder-Mead with dCRAB can escape the artificial traps and greatly improve the infidelity with enough function evaluations. The Nelder-Mead with dCRAB using 25 000 evaluations finds similar infidelities to GROUP with dCRAB using 2500 evaluations, which shows that if Nelder-Mead runs for longer times, it can find results

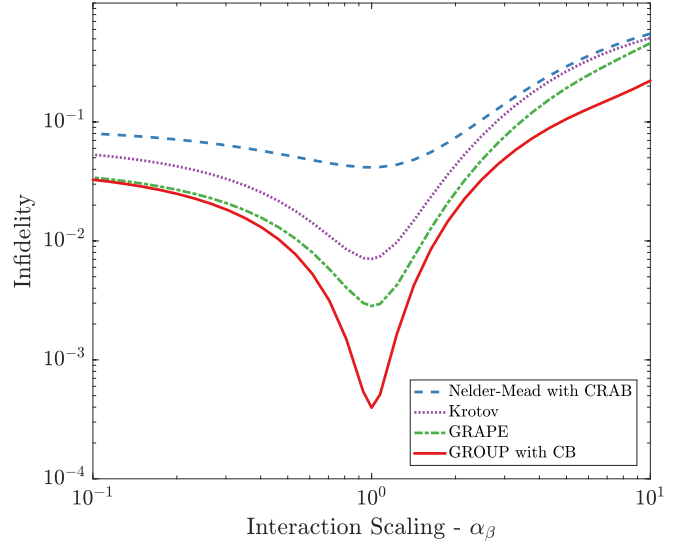


FIG. 6. The robustness of the optimal solutions from each algorithm to rescaling of the self-interaction in Eq. (1): $\tilde{\beta} = \alpha_\beta \beta$, where α_β is the scaling.

similar to gradient descent. At large basis sizes, dCRAB does not notably improve the infidelity, which we interpret as the optimization being blocked by Nelder-Mead's inability to efficiently search in high-dimensional landscapes rather than artificial traps.

C. Robustness analysis

The solutions presented in this paper can only be directly applied in an experiment if they are stable against unavoidable experimental fluctuations in the system parameters. As reported in Ref. [24], the primary experimental fluctuations occur in the self-interaction (β) in Eq. (1) due to variations in the atom number. This gives variations in β below 14%. For this reason, we have investigated the stability of the solutions against scaling of the self-interaction. In Fig. 6, the infidelity is calculated with the best solution from each algorithm using $\tilde{\beta} = \alpha_\beta \beta$, where α_β is the scaling factor. The solutions are generally stable with respect to fluctuations in the self-interaction since atom interactions have little influence over the relative short duration investigated here. The GROUP solution clearly exhibits superior behavior over the entire range of parameter variation. The GROUP solution remains above $F = 0.99$ in the range $0.5 \leq \alpha_\beta \leq 1.5$, which is above the experimental fluctuations reported in Ref. [24].

We have also investigated the robustness of the solutions against perturbations in the potential. In Fig. 7, the infidelity is calculated for each algorithm's best solution with rescaled potential coefficients ($\tilde{p}_i = \alpha_p p_i$) in Eq. (2). The solutions are much more sensitive to perturbations in the potential compared with the self-interaction.

We note that it would most likely be possible to find much more stable solutions by including the stability directly in the optimization. This could, for example, be achieved by optimizing over an ensemble of different scaled interactions and potentials and defining the average cost over the ensemble.




FIG. 7. The robustness of the optimal solutions from each algorithm to rescaling of the potential coefficients in Eq. (2): $\tilde{p}_i = \alpha_p p_i$, $i = 2, 4, 6$, where α_p is the scaling.

D. Convergence behavior at different parameter values

The results presented in this section have so far only discussed a single choice of the system parameter values, meaning that the underlying optimization problem has been the same. The performance of an optimization algorithm may depend on the choice of problem, so in this section we investigate the performance of each algorithm for different problems by rescaling the system parameters. In Fig. 2, we compared the convergence behavior for a number of different seeds, but fixed system parameters. In Figs. 8 and 9, we hold the seed fixed and vary the system parameters examined in the previous section. The seeds are those that gave the best result for each algorithm in the previous analysis. As in Fig. 2, we have performed the analysis for different basis sizes and step sizes and we present the best case behavior in Figs. 8 and 9.

In Fig. 8, we compare the convergence behavior of each algorithm when optimizing with 35 different values of $\tilde{\beta} = \alpha_\beta \beta$ ranging from $\alpha_\beta = 0.1$ to $\alpha_\beta = 10$ as in the previous section. Although GRAPE and GROUP have similar medians for their final infidelities, GROUP has a better overall convergence behavior. Both GROUP and GRAPE have faster rates of convergence than Nelder-Mead with CRAB and Krotov's method. In Fig. 2, Krotov's method has a better final infidelity than Nelder-Mead with CRAB. However, in Fig. 8, Nelder-Mead with CRAB catches up with Krotov after a high number of function evaluations. In general, Krotov's method is struggling for the high values of β , which we attribute to the lack of including the term proportional to the difference in Eq. (17). We believe that correctly including this term could substantially improve the convergence rate of Krotov's method. Note that here we present the results with the best basis for Nelder-Mead after optimizing over a range of basis sizes. The performance is significantly worse for nearby basis sizes 15 and 30.

In Fig. 9, we perform a similar analysis for each algorithm just optimizing 35 different values of the potential with rescaled coefficients $\tilde{p}_i = \alpha_p p_i$, with $i = 2, 4, 6$ ranging from

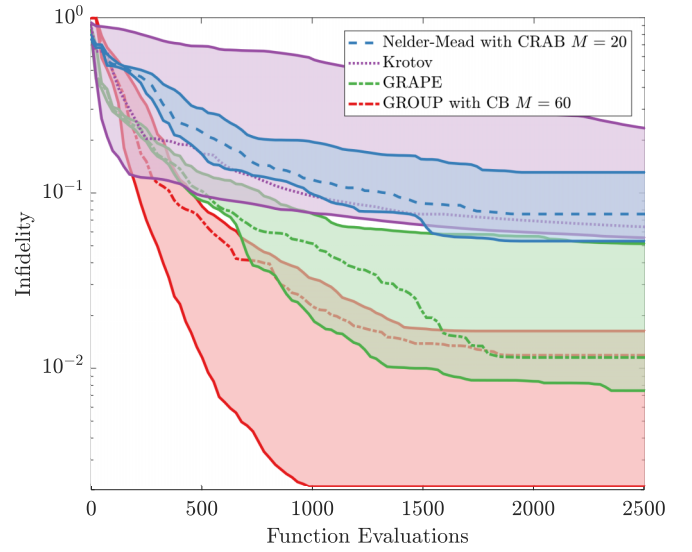


FIG. 8. Comparison of the convergence behavior for different values of the self-interaction for each algorithm. The optimizations are performed for 35 different scaled values of the self-interaction $\tilde{\beta} = \alpha_\beta \beta$, where α_β is the scaling. The dotted line shows the median and the shaded area indicates the 25% and 75% quartiles. The quasi-Newton method BFGS was used together with GRAPE and GROUP.

$\alpha_p = 0.91$ to $\alpha_p = 1.09$ as in the previous section. Again, we observe that GROUP followed by GRAPE finds the solutions with the lowest infidelities. Krotov's method has fast initial reductions in the infidelity below 600 function evaluations, but it slows down above 1000 function evaluations. In this case, we also observe that Nelder-Mead with CRAB catches up with

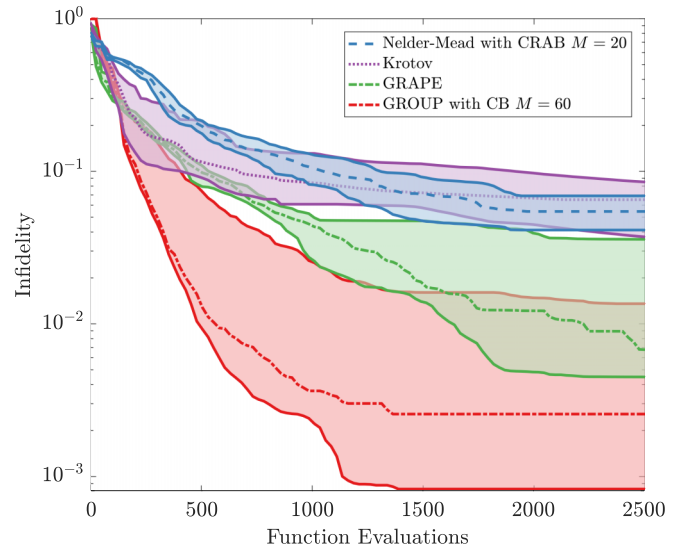


FIG. 9. Comparison of the convergence behavior for different values of the potential coefficients. The optimizations are performed for 35 different scaled values of the potential coefficients $\tilde{p}_i = \alpha_p p_i$ with $i = 2, 4, 6$, where α_p is the scaling. The dotted line shows the median and the shaded area indicates the 25% and 75% quartiles. The quasi-Newton method BFGS was used together with GRAPE and GROUP.

Krotov's method after a high number of function evaluations. Similarly to before, we observe in our other optimizations that Nelder-Mead has significantly worse performance for the nearby basis sizes 15 and 30.

We have also performed a similar analysis for rescaled values of the regularization factor (γ) in Eq. (3), which showed that the regularization factor only has a minor impact on the results for all algorithms.

VI. CONCLUSION AND OUTLOOK

We have presented in detail how to perform gradient-based optimal control in a reduced basis in GROUP, which combines the advantages from the CRAB and GRAPE methods. We have presented a benchmark with other quantum optimal control methods close to the quantum speed limit in the condensate driving control problem. Here, GROUP-type methods are competitive with other quantum control algorithms. We have also presented how to extend GROUP-type methods with experimentally motivated filter functions. It would be very interesting to compare these methods on other control problems in order to better understand the advantages and disadvantages of each method. The formalism for GROUP-type methods presented here can also be applied to systems with a linear Schrödinger equation such as control of many-body systems [23,24]. We believe our method and analysis is a relevant addition to the repository of quantum control algorithms.

ACKNOWLEDGMENTS

This work has been supported by the European Research Council and the Lundbeck Foundation. We would also like to thank Jesper H. M. Jensen for illuminating discussions and help with adjusting the manuscript. Antonio Negretti and Marie Bonneau helped with supplying the filter function and clarifying details in the modeling of the experiment.

APPENDIX: VARIATIONS AND GRADIENTS

1. Equations of optimality

Here we give a brief derivation of the optimality equations presented in the main text [Eqs. (5)–(10)]. The presentation here is based on Refs. [12,34,43,48]. In order to find equations for an optimum that also satisfies the GPE, we need to calculate the variations with respect to the Lagrangian

$$\mathcal{L}(\psi, u, \chi) = J(u, \psi) + \text{Re} \int_0^T \langle \chi | Z(u, \psi) \rangle dt, \quad (\text{A1})$$

where $Z(u, \psi)$ is the constraint given by

$$Z(u, \psi) = i\dot{\psi} - \hat{H}\psi - \beta|\psi|^2\psi. \quad (\text{A2})$$

Clearly this constraint is zero for any $\psi(t)$ that also satisfies the GPE equation. The optimality system is found by requiring that all first-order variations vanish. This gives us the equations

$$D_{\delta\chi}\mathcal{L} = D_{\delta\psi}\mathcal{L} = D_{\delta u}\mathcal{L} = 0, \quad (\text{A3})$$

for all admissible variations. We now calculate these three variations one by one. The variation with respect to χ

gives

$$\begin{aligned} D_{\delta\chi}\mathcal{L} &= D_{\delta\chi}\text{Re} \int_0^T \langle \chi | i\dot{\psi} - \hat{H}\psi - \beta|\psi|^2\psi \rangle dt \\ &= \text{Re} \int_0^T \langle \delta\chi | i\dot{\psi} - \hat{H}\psi - \beta|\psi|^2\psi \rangle dt. \end{aligned} \quad (\text{A4})$$

This variation must be zero for all variations $\delta\chi(t)$, which gives Eq. (5). Next consider the variation with respect to ψ . A variation of the constraint gives

$$\begin{aligned} D_{\delta\psi}\text{Re} \int_0^T \langle \chi | Z \rangle dt &= D_{\delta\psi} \left(\text{Re} \langle \chi | i\psi \rangle \Big|_0^T - \text{Re} \int_0^T \langle \dot{\chi} | i\psi \rangle dt \right) \\ &\quad - \text{Re} \int_0^T \langle \chi | \hat{H}\delta\psi + \beta\psi^2\delta\psi^* + 2\beta|\psi|^2\delta\psi \rangle dt \\ &= \text{Re} \langle i\chi(0) | \delta\psi(0) \rangle - \text{Re} \langle i\chi(T) | \delta\psi(T) \rangle \\ &\quad + \text{Re} \int_0^T \langle i\dot{\chi} - \hat{H}\chi - 2\beta|\psi|^2\chi - \beta\psi^2\chi^* | \delta\psi \rangle dt. \end{aligned} \quad (\text{A5})$$

First consider $0 < t < T$. Since all variations must vanish, we find the optimality condition in Eq. (6). For $t = T$, the cost functional $J(u, \psi)$ also contributes with the term

$$\begin{aligned} D_{\delta\psi}J &= -\frac{1}{2}D_{\delta\psi}\langle \psi_T | \psi(T) \rangle \langle \psi(T) | \psi_T \rangle \\ &= -\text{Re}[\langle \psi(T) | \psi_T \rangle \langle \psi_T | \delta\psi(T) \rangle]. \end{aligned} \quad (\text{A6})$$

Combining this with Eq. (A5) gives

$$\begin{aligned} 0 &= -\text{Re}[\langle i\chi(T) | + \langle \psi(T) | \psi_T \rangle \langle \psi_T |] \delta\psi(T) \rangle] \\ &\quad + \text{Re} \langle i\chi(0) | \delta\psi(0) \rangle. \end{aligned} \quad (\text{A7})$$

The first term in this equation gives us the boundary condition for the Lagrange multiplier in Eq. (9). There is also the initial condition that $\psi(0) = \psi_0$ [Eq. (8)], which gives $\delta\psi(0) = 0$, implying that the second term in the equation above is zero. Finally, consider the variation with respect to the control u . First, consider the variation with respect to the cost functional,

$$D_{\delta u}J = \gamma[\dot{u}(T)\delta u(T) - \dot{u}(0)\delta u(0)] - \gamma \int_0^T \ddot{u}\delta u dt. \quad (\text{A8})$$

The terms at the boundary vanish since the control must be fixed at the boundary [Eq. (10)]. With this expression, we find the variation of the Lagrangian to be

$$D_{\delta u}\mathcal{L} = - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) \delta u dt. \quad (\text{A9})$$

Since this must vanish for all variations, we find the last optimality condition in Eq. (10).

2. Calculation of GRAPE gradients

In the main text, we introduce the gradients for GRAPE. Here we give a brief derivation of these equations. Our considerations follow Refs. [34,48]. In numerical simulations, we solve the GPE for a given u , so it is more natural to consider the reduced cost functional $\hat{J}(u, \psi) = J(u, \psi_u)$, where ψ_u is the solution to the GPE for a given u . Note that $\hat{J}(u) = \mathcal{L}(\psi_u, u, \chi)$

since $Z(u, \psi_u) = 0$. This is correct for any χ , hence χ is a free variable that we can choose conveniently. In order to find the gradient for GRAPE, we need the variation of the reduced cost functional,

$$D_{\delta u} \hat{J}(u) = D_{\delta u} \mathcal{L}(\psi, u, \chi) + D_{\delta \psi_u} \mathcal{L}(\psi, u, \chi), \quad (\text{A10})$$

where we have used the total derivative. The derivative $D_{\delta \psi_u} \mathcal{L}$ is the induced variation in ψ_u from the variation of u . This extra term appears since ψ depends implicitly on u through the GPE equation. This was not the case when discussing the Lagrangian [Eq. (A1)] since ψ was taken to be a free variable and the GPE is a constraint. Notice that performing this variation is formally the same as the $D_{\delta \psi}$, just with the induced variation instead. Hence, we would get the same equations as above for $D_{\delta \psi} \mathcal{L}$, just with $\delta \psi$ replaced with $\delta \psi_u$. Specifically, we find Eq. (A5) again for the induced variation. Recall that χ is now a free variable that we can select. If we pick χ to satisfy Eq. (6), then the induced variation $D_{\delta \psi_u}$ will vanish, leading to the conclusion that

$$\begin{aligned} D_{\delta u} \hat{J}(u) &= D_{\delta u} \mathcal{L}(\psi, u, \chi) \\ &= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) \delta u dt, \end{aligned} \quad (\text{A11})$$

where we have used the result from Eq. (A9). Recall from the discussion in the main text that the GRAPE gradient is defined as the unique element such that $(\nabla \hat{J}, \delta u)_X = D_{\delta u} \hat{J}$. The gradient depends on the choice of the inner product. The typical choice of the L^2 inner product gives

$$\begin{aligned} (\nabla \hat{J}, \delta u)_{L^2} &= \int_0^T \nabla \hat{J} \delta u dt \\ &= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) \delta u dt. \end{aligned} \quad (\text{A12})$$

From this equation, we can immediately recognize that

$$\nabla \hat{J} = -\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle - \gamma \ddot{u} \text{ for } L^2. \quad (\text{A13})$$

Note that there is no reason for this expression to vanish at the boundaries ($t = 0$ and $t = T$). As discussed in the main text, failing to satisfy the boundary conditions for the control [Eq. (10)] can cause instabilities in the algorithm. It turns out that a more suitable choice is the H^1 inner product. Here the gradient is

$$(\nabla \hat{J}, \delta u)_{H^1} = \delta u \frac{d}{dt} \nabla \hat{J} \Big|_0^T - \int_0^T \delta u \frac{d^2}{dt^2} \nabla \hat{J} dt \quad (\text{A14})$$

$$= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) \delta u dt. \quad (\text{A15})$$

The first term in Eq. (A14) vanishes since $\delta u(0) = \delta u(T) = 0$. From this expression, we can directly read off the H^1 gradient,

$$\frac{d^2}{dt^2} [\nabla \hat{J}(u)] = \gamma \ddot{u} + \text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle \text{ for } H^1, \quad (\text{A16})$$

which is the result given in Eq. (13). This is a Poisson equation, so the Dirichlet boundary condition that the gradients vanish at the boundaries can be chosen. This is the motivation for using the H^1 gradient over the L^2 gradient.

As discussed in the main text, it is necessary in condensate driving to take the finite bandwidth of the electronics into account. This effect distorts the control $u(t)$ into $v(t)$, which enters into the GPE equation. This can be modeled using a filter function $h(t)$,

$$v(t) = (h * u)(t) = \int_0^t h(\tau) u(t - \tau) d\tau, \quad (\text{A17})$$

where $v(t)$ is the distorted control. Again, we can calculate the variation. This can be done with the chain rule for variations,

$$\begin{aligned} D_{\delta u} \hat{J}(v) &= D_{D_{\delta u} v} \hat{J}(v) \\ &= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle + \gamma \ddot{v} \right) (h * \delta u)(t) dt \\ &\quad + \gamma \dot{v}(T) (h * \delta u)(T). \end{aligned} \quad (\text{A18})$$

The first term is directly found from Eq. (A9) where the expression inside the bracket is evaluated along the distorted control $v(t)$. Note that when calculating the variation of the regularization with respect to u , two boundary terms were zero due to the boundary conditions for δu . However, the first of these terms is not zero when the filter is included. This gives the second term in the equation above. This equation can be rewritten as

$$\begin{aligned} D_{\delta u} \hat{J}(v) &= \int_0^T \delta u(t) \left[\gamma \dot{v}(T) h(T-t) - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle \right. \right. \\ &\quad \left. \left. + \gamma \ddot{v} \right) h(t' - t) \Theta(t' - t) dt' \right] dt, \end{aligned} \quad (\text{A19})$$

where $\Theta(t' - t)$ is the Heaviside-step function. From this expression, we can identify the gradients for both L^2 and H^1 . A similar expression for the filter gradient was reported in Ref. [50]. With the same arguments as in Eq. (A14), the H^1 gradient is given as

$$\begin{aligned} \frac{d^2}{dt^2} [\nabla \hat{J}(u)] &= - \gamma \dot{v}(T) h(T-t) \\ &\quad + \int_t^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle + \gamma \ddot{v} \right) h(t' - t) dt'. \end{aligned} \quad (\text{A20})$$

This is the expression given in the main text.

3. Calculation of GROUP gradients

Here we give a brief deviation of the GROUP gradient expression presented in the main text. As discussed in the main text, in GROUP, the control is expressed as the linear combination

$$u(t) = u_0(t) + S(t) \sum_{n=1}^M c_n f_n(t), \quad (\text{A21})$$

where the $f_n(t)$ are smooth functions. Here the optimization is over the expansion coefficients c_n 's. The partial derivative with respect to one of these coefficients can be found using the chain rule for variational derivatives,

$$\begin{aligned} \frac{\partial \hat{J}(\mathbf{c})}{\partial c_n} &= D_{S(t)f_n(t)} \hat{J}(u) \\ &= - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle + \gamma \ddot{u} \right) S(t) f_n(t) dt, \end{aligned} \quad (\text{A22})$$

where we have simply reused the result from Eq. (A9) and replaced $\delta u(t)$ with $S(t)f_n(t)$. There is no contribution from the boundary terms in the regularization since $S(0) = S(T) = 0$. The chain rule also allows us to take the filter function into account since we can replace $\delta u(t)$ with $S(t)f_n(t)$ in Eq. (A18),

$$\frac{\partial \hat{J}(v)}{\partial c_n} = - \int_0^T \left(\text{Re} \left\langle \chi \left| \frac{\partial \hat{H}}{\partial v} \right| \psi \right\rangle + \gamma \ddot{v} \right) (h * S f_n)(t) dt + \gamma \dot{v}(T) (h * S f_n)(T). \quad (\text{A23})$$

The expression inside the bracket is evaluated along the distorted control $v(t)$.

4. GOAT

In Refs. [33,54], it was proposed to calculate the gradient in a reduced basis $\partial \hat{J}(\mathbf{c})/\partial c_n$ using another method named gradient optimization of analytic controls (GOAT), which we briefly discuss here. Ignoring the regularization in Eq. (3), the derivative of the cost is

$$\frac{\partial \hat{J}}{\partial c_n} = -\text{Re} \left\langle \psi_i \left| \frac{\partial \psi}{\partial c_i}(T) \right\rangle. \quad (\text{A24})$$

In Ref. [33], the derivative of the state is directly computed from the equations of motion, which gives

$$i \partial_t \begin{pmatrix} \psi \\ \partial_{c_n} \psi \end{pmatrix} = \begin{pmatrix} H + \beta |\psi|^2 & 0 \\ \partial_{c_n} H + \beta \psi \partial_{c_n} \psi^* & H + 2\beta |\psi|^2 \end{pmatrix} \begin{pmatrix} \psi \\ \partial_{c_n} \psi \end{pmatrix}. \quad (\text{A25})$$

This is the straightforward extension of the method from Ref. [33] to the nonlinear dynamics of the GPE, and for $\beta = 0$ we obtain the original result. In order to compute the derivative in a basis with size $M = 60$, it would then be necessary to compute $\partial_{c_1} \psi, \partial_{c_2} \psi, \dots, \partial_{c_{60}} \psi$ using the relation above. As we discussed in the main text, $M = 60$ is numerically found to be the optimal basis size. This can be done by constructing one huge matrix of size $[(M+1)d] \times (2d)$, where d is the number of grid points used when discretizing the GPE. An alternative is to calculate 60 independent solutions of Eq. (A25), which could be done in parallel. Compared with the expression in Eq. (16) that only requires integrating two equations of motion, calculating the gradient using GOAT is much more demanding. Nevertheless, the GOAT approach potentially offers a higher numerical accuracy in the gradients since the accuracy of Eq. (16) is limited by the error from discretizing the time evolution. This is an advantage when optimizing for very low errors ($1 - F \leq 10^{-9}$) as needed for some error-correction protocols in quantum computing [33]. Very low infidelities below 10^{-4} are not required in the condensate driving control problem, so calculating the derivative using the more expensive GOAT method would most likely not be an advantage. A fair comparison between these two methods for calculating the gradient would require studying a wider class of problems and discussing other alternatives for calculating high-accuracy gradients, which is beyond the scope of this paper.

5. Krotov's method

As discussed in the main text, Krotov's method is an alternative to the Lagrange multiplier approach for deriving

optimal control algorithms [29,49]. Here we briefly present Krotov's method as described in Refs. [28,29,53]. We do not include the regularization term in Krotov so the cost functional becomes $J_T = [1 - |\langle \psi_i | \psi(T) \rangle|^2]/2$. In Krotov's method, the cost functional is rewritten so the GPE appears explicitly. This is done by adding the vanishing quantity $0 = \phi(T) - \phi(0) - \int_0^T \frac{d\phi}{dt} dt$. Here, ϕ is some arbitrary field, which we can select freely. In order to present the derivation more clearly, we write the variational derivatives differently here and adopt the standard notation that $\dot{\psi} = f(\psi, u) = -i \hat{H}_{\text{NL}} \psi$, where \hat{H}_{NL} is the GPE Hamiltonian from Eq. (1). This allows the cost functional to be rewritten as

$$L = G(\psi(T)) - \phi(0) - \int_0^T R(\psi, u) dt, \quad (\text{A26})$$

where

$$G(\psi(T)) = J_T + \phi(T), \quad (\text{A27})$$

$$R(\psi, u) = \frac{\partial \phi}{\partial t} + \frac{\delta \phi}{\delta \psi} f + f^* \frac{\delta \phi}{\delta \psi^*}. \quad (\text{A28})$$

From the arguments above and the definition of R , it is seen that $L = J$, as discussed in Ref. [28]. So minimizing J is equivalent to minimizing L . In Krotov's method, it is directly required that the cost decreases at each iteration $J^{(i+1)} \leq J^{(i)}$, which is equivalent to $0 \leq L^{(i)} - L^{(i+1)} = \Delta_1 + \Delta_2 + \Delta_3$, where

$$\Delta_1 = G(\psi^{(i)}(T)) - G(\psi^{(i+1)}(T)), \quad (\text{A29})$$

$$\Delta_2 = \int_0^T R(\psi^{(i+1)}, u^{(i+1)}) - R(\psi^{(i+1)}, u^{(i)}) dt, \quad (\text{A30})$$

$$\Delta_3 = \int_0^T R(\psi^{(i+1)}, u^{(i)}) - R(\psi^{(i)}, u^{(i)}) dt. \quad (\text{A31})$$

A sufficient condition for a decrease in the cost is that the Δ 's are positive. Central in Krotov's method is the unintuitive notion that L is maximized with respect to the states at the current iteration. This implies that any change in the states caused by selecting a new value of the control for the next iteration would decrease the value of L [28,29]. If L is maximal with respect to the states, then the first-order derivatives with respect to ψ^* must vanish. These derivatives are

$$\begin{aligned} \frac{\delta R}{\delta \psi^*} &= \left(\frac{\partial}{\partial t} + f \frac{\delta}{\delta \psi} + f^* \frac{\delta}{\delta \psi^*} \right) \frac{\delta \phi}{\delta \psi^*} + \frac{\delta \phi}{\delta \psi} \frac{\delta f}{\delta \psi^*} + \frac{\delta f^*}{\delta \psi^*} \frac{\delta \phi}{\delta \psi^*} \\ &= \frac{d}{dt} \frac{\delta \phi}{\delta \psi^*} + \frac{\delta \phi}{\delta \psi} \frac{\delta f}{\delta \psi^*} + \frac{\delta f^*}{\delta \psi^*} \frac{\delta \phi}{\delta \psi^*}, \end{aligned} \quad (\text{A32})$$

and

$$\frac{\delta G(\psi(T))}{\delta \psi^*} = \frac{\delta J_T}{\delta \psi^*} + \frac{\delta \phi(T)}{\delta \psi^*}. \quad (\text{A33})$$

From the discussion above, we require that $\delta R/\delta \psi^*|_{\psi^{(i)}, u^{(i)}} = 0$ and $\delta G/\delta \psi^*|_{\psi^{(i)}, u^{(i)}} = 0$. This condition gives for R that

$$\begin{aligned} \frac{d}{dt} \frac{\delta \phi}{\delta \psi^*} \Big|_{\psi^{(i)}, u^{(i)}} &= \left(i \psi \frac{\delta \hat{H}_{\text{NL}}}{\delta \psi^*} \frac{\delta \phi}{\delta \psi} - i \psi^* \frac{\delta \hat{H}_{\text{NL}}}{\delta \psi^*} \frac{\delta \phi}{\delta \psi^*} \right. \\ &\quad \left. - i \hat{H}_{\text{NL}} \frac{\delta \phi}{\delta \psi^*} \right) \Big|_{\psi^{(i)}, u^{(i)}}. \end{aligned} \quad (\text{A34})$$

The requirement that L is maximal with respect to ψ also puts requirements on the second-order derivatives of R and G . A good ansatz for ϕ is a second-order expansion in the states

$$\phi = \frac{1}{2}(\langle \xi | \psi \rangle + \langle \psi | \xi \rangle) + \frac{1}{4} \langle \Delta \psi | \sigma(t) | \Delta \psi \rangle, \quad (\text{A35})$$

where ξ are some expansion coefficients, $\sigma(t)$ is an arbitrary time-dependent function, and $\Delta \psi = \psi - \psi^{(i)}$ is the difference from the next iteration to the current. If this expression is inserted into Eqs. (A33) and (A34), one obtains

$$i\dot{\xi}^{(i)} = [\hat{H}(u^{(i)}) + 2\beta|\psi^{(i)}|^2]\xi^{(i)} - (\psi^{(i)})^2\beta\xi^{(i)*}, \quad (\text{A36})$$

$$\xi^{(i)}(T) = -|\psi_i\rangle\langle\psi_i|\psi^{(i)}(T)\rangle. \quad (\text{A37})$$

Incidentally, these equations are the same as those for χ given in Eqs. (6) and (9) with an extra phase factor, so we have $\xi = i\chi$. We now discuss the additional conditions that ensure that each Δ is positive.

The boundary condition for ξ [Eq. (A37)] can be combined with the definition of G [Eq. (A27)], which gives

$$G(\psi(T)) = \frac{1}{2}[1 - \langle \psi(T) | \hat{P} | \psi(T) \rangle + \langle \psi^{(i)}(T) | \hat{P} | \psi(T) \rangle + \langle \psi(T) | \hat{P} | \psi^{(i)}(T) \rangle], \quad (\text{A38})$$

where $\hat{P} = |\psi_i\rangle\langle\psi_i|$ is the projection operator for the target state. From this expression, Δ_1 can be directly rewritten as $2\Delta_1 = \langle \Delta \psi | \hat{P} | \Delta \psi \rangle$, which is always non-negative due to the positivity of \hat{P} .

Generally, it is more difficult to ensure that Δ_3 is positive. Additional conditions on the second-order derivatives with respect to the states on R are also required, e.g., $\delta^2 R / \delta \psi \delta \psi^* |_{\psi^{(i)}} > 0$. This can be ensured with a proper choice of the $\sigma(t)$ function. A number of different strategies for choosing $\sigma(t)$ have been discussed in the literature [28,29,43]. For the moderate values of β discussed here, a good strategy is simply to select $\sigma(t) = 0$ and forfeit the strict guarantee that Δ_3 is positive [43]. Note that if $\beta = 0$, then R is independent of ψ and $\Delta_3 = 0$. If $\Delta_3 = 0$, we can just pick $\sigma(t) = 0$ and the algorithm is guaranteed to decrease the cost at every iteration.

Finally, we discuss how to ensure Δ_2 is positive by a proper update for the control. Ideally, we seek a control such that the derivative vanishes ($\partial R / \partial u |_{u^{(i+1)}, \psi^{(i+1)}} = 0$). This derivative is

$$\begin{aligned} \frac{\partial R}{\partial u} \Big|_{u^{(i+1)}, \psi^{(i+1)}} &= \text{Im} \left[\left\langle \xi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle \right. \\ &\quad \left. + \frac{\sigma(t)}{4} \left\langle \Delta \psi \left| \frac{\partial \hat{H}}{\partial u} \right| \psi \right\rangle \right] \Big|_{u^{(i+1)}, \psi^{(i+1)}}. \end{aligned} \quad (\text{A39})$$

It is difficult to ensure that this is zero [29]. Instead of choosing




FIG. 10. The final infidelity after 2500 evaluations for Krotov as a function of the step size (α) used in Eq. (17). The dotted line and shaded area show the median, 25%, and 75% quartiles for the 100 initial controls.

the optimal control, we take a small step in the direction of the gradient,

$$u^{(i+1)} = u^{(i)} + \alpha S(t) \frac{\partial R}{\partial u} \Big|_{\psi^{(i+1)}, u^{(i)}}, \quad (\text{A40})$$

where $\alpha > 0$ and $S(t)$ is the shape function that vanishes for $t = 0$ and $t = T$ so the boundary conditions on the control can be satisfied [Eq. (10)]. If $R(\psi^{(i+1)}, u^{(i+1)})$ in Δ_2 is Taylor expanded around the last control, we find

$$\begin{aligned} \Delta_2 &\approx \int_0^T \frac{\partial R}{\partial u} \Big|_{\psi^{(i+1)}, u^{(i)}} (u^{(i+1)} - u^{(i)}) dt \\ &= \int_0^T \alpha S(t) \left(\frac{\partial R}{\partial u} \Big|_{\psi^{(i+1)}, u^{(i)}} \right)^2 dt \geq 0. \end{aligned} \quad (\text{A41})$$

These arguments show that if the control is chosen as in Eq. (A40), then all three Δ 's are positive and the cost will decrease. If Eq. (A40) is rewritten with χ , then Eq. (17) given in the main text is found.

A proper value of the step size α must also be selected in order to ensure a fast convergence. The performance of the Krotov algorithm in condensate driving is shown for different step sizes in Fig. 10. If the step size is too small, the algorithm converges slowly due to the small steps in the update. On the other hand, if the steps are too large, then there is no longer a guarantee that Δ_3 is positive and the update might not decrease the cost. The optimal step size is around 0.1.

- [1] I. Bloch, J. Dalibard, and W. Zwerger, *Rev. Mod. Phys.* **80**, 885 (2008).
 [2] R. J. Gordon and S. A. Rice, *Annu. Rev. Phys. Chem.* **48**, 601 (1997).

- [3] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune *et al.*, *Nature (London)* **477**, 73 (2011).

- [4] R. Bücker, J. Grond, S. Manz, T. Berrada, T. Betz, C. Koller, U. Hohenester, T. Schumm, A. Perrin, and J. Schmiedmayer, *Nat. Phys.* **7**, 608 (2011).
- [5] B. Bloom, T. Nicholson, J. Williams, S. Campbell, M. Bishof, X. Zhang, W. Zhang, S. Bromley, and J. Ye, *Nature (London)* **506**, 71 (2014).
- [6] B. Lücke, M. Scherer, J. Kruse, L. Pezzé, F. Deuretzbacher, P. Hyllus, J. Peise, W. Ertmer, J. Arlt, L. Santos *et al.*, *Science* **334**, 773 (2011).
- [7] M. F. Riedel, P. Böhi, Y. Li, T. W. Hänsch, A. Sinatra, and P. Treutlein, *Nature (London)* **464**, 1170 (2012).
- [8] C. Gross, T. Zibold, E. Nicklas, J. Esteve, and M. K. Oberthaler, *Nature (London)* **464**, 1165 (2012).
- [9] I. Bloch, J. Dalibard, and S. Nascimbene, *Nat. Phys.* **8**, 267 (2012).
- [10] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. Porto, *Nature (London)* **448**, 452 (2007).
- [11] D. Kielpinski, C. Monroe, and D. J. Wineland, *Nature (London)* **417**, 709 (2002).
- [12] U. Hohenester, P. K. Rekdal, A. Borzi, and J. Schmiedmayer, *Phys. Rev. A* **75**, 023602 (2007).
- [13] J. P. Palao and R. Kosloff, *Phys. Rev. Lett.* **89**, 188301 (2002).
- [14] A. Kaiser and V. May, *J. Chem. Phys.* **121**, 2528 (2004).
- [15] 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.*, *Eur. Phys. J. D* **69**, 279 (2015).
- [16] J. Werschnik and E. Gross, *J. Phys. B* **40**, R175 (2007).
- [17] A. P. Peirce, M. A. Dahleh, and H. Rabitz, *Phys. Rev. A* **37**, 4950 (1988).
- [18] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, *npj Quantum Inf.* **3**, 37 (2017).
- [19] C. P. Koch, J. P. Palao, R. Kosloff, and F. Masnou-Seeuws, *Phys. Rev. A* **70**, 013402 (2004).
- [20] T. Nöbauer, A. Angerer, B. Bartels, M. Trupke, S. Rotter, J. Schmiedmayer, F. Mintert, and J. Majer, *Phys. Rev. Lett.* **115**, 190801 (2015).
- [21] T. Caneva, T. Calarco, and S. Montangero, *Phys. Rev. A* **84**, 022326 (2011).
- [22] J. J. Sørensen, M. Aranburu, T. Heinzel, and J. Sherson, *arXiv:1802.07521*.
- [23] P. Doria, T. Calarco, and S. Montangero, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [24] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco *et al.*, *Sci. Rep.* **6**, 34187 (2016).
- [25] I. Brouzos, A. I. Streltsov, A. Negretti, R. S. Said, T. Caneva, S. Montangero, and T. Calarco, *Phys. Rev. A* **92**, 062110 (2015).
- [26] R. E. Goetz, M. Merkel, A. Karamatskou, R. Santra, and C. P. Koch, *Phys. Rev. A* **94**, 023420 (2016).
- [27] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *J. Magn. Reson.* **172**, 296 (2005).
- [28] S. E. Sklarz and D. J. Tannor, *Phys. Rev. A* **66**, 053619 (2002).
- [29] D. M. Reich, M. Ndong, and C. P. Koch, *J. Chem. Phys.* **136**, 104103 (2012).
- [30] S. Lloyd and S. Montangero, *Phys. Rev. Lett.* **113**, 010502 (2014).
- [31] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, *Phys. Rev. A* **92**, 062343 (2015).
- [32] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey *et al.*, *Phys. Rev. Lett.* **112**, 240504 (2014).
- [33] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [34] G. Von Winckel and A. Borzi, *Inverse Probl.* **24**, 034007 (2008).
- [35] D. G. Lucarelli, *Phys. Rev. A* **97**, 062346 (2018).
- [36] T. E. Skinner and N. I. Gershenzon, *J. Magn. Reson.* **204**, 248 (2010).
- [37] F. Motzoi, J. M. Gambetta, S. T. Merkel, and F. K. Wilhelm, *Phys. Rev. A* **84**, 022307 (2011).
- [38] J. J. W. Sørensen, M. K. Pedersen, M. Munch, P. Haikka, J. H. Jensen, T. Planke, M. G. Andreasen, M. Gajdacz, K. Mølmer, A. Lieberoth *et al.*, *Nature (London)* **532**, 210 (2016).
- [39] M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, *arXiv:1705.00565*.
- [40] A. G. Day, M. Bukov, P. Weinberg, P. Mehta, and D. Sels, *arXiv:1803.10856*.
- [41] M. Bukov, A. G. R. Day, P. Weinberg, A. Polkovnikov, P. Mehta, and D. Sels, *Phys. Rev. A* **97**, 052114 (2018).
- [42] R. Bücker, T. Berrada, S. Van Frank, J.-F. Schaff, T. Schumm, J. Schmiedmayer, G. Jäger, J. Grond, and U. Hohenester, *J. Phys. B* **46**, 104012 (2013).
- [43] G. Jäger, D. M. Reich, M. H. Goerz, C. P. Koch, and U. Hohenester, *Phys. Rev. A* **90**, 033628 (2014).
- [44] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, *Phys. Rev. A* **84**, 022305 (2011).
- [45] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer, Berlin, 2006).
- [46] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer, *Phys. Rev. A* **73**, 033619 (2006).
- [47] F. Gerbier, *Europhys. Lett.* **66**, 771 (2004).
- [48] J.-F. Mennemann, D. Matthes, R.-M. Weishäupl, and T. Langen, *New J. Phys.* **17**, 113027 (2015).
- [49] A. Konnov and V. F. Krotov, *Avtom. Telemek. 10*, 77 (1999).
- [50] G. Jäger and U. Hohenester, *Phys. Rev. A* **88**, 035601 (2013).
- [51] D. M. Reich, J. P. Palao, and C. P. Koch, *J. Mod. Opt.* **61**, 822 (2014).
- [52] C. Gollub, M. Kowalewski, and R. de Vivie-Riedle, *Phys. Rev. Lett.* **101**, 073002 (2008).
- [53] R. Eitan, M. Mundt, and D. J. Tannor, *Phys. Rev. A* **83**, 053426 (2011).
- [54] P. De Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, *J. Magn. Reson.* **212**, 412 (2011).