

This article was downloaded by:

On: 21 January 2011

Access details: *Access Details: Free Access*

Publisher *Taylor & Francis*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



International Reviews in Physical Chemistry

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713724383>

Quantum control landscapes

Raj Chakrabarti^a; Herschel Rabitz^a

^a Department of Chemistry, Princeton University, Princeton, New Jersey 08544, USA

To cite this Article Chakrabarti, Raj and Rabitz, Herschel(2007) 'Quantum control landscapes', International Reviews in Physical Chemistry, 26: 4, 671 – 735

To link to this Article: DOI: 10.1080/01442350701633300

URL: <http://dx.doi.org/10.1080/01442350701633300>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Quantum control landscapes

RAJ CHAKRABARTI* and HERSCHEL RABITZ

Department of Chemistry, Princeton University, Princeton, New Jersey 08544, USA

(Received 1 August 2007; in final form 14 August 2007)

Numerous lines of experimental, numerical and analytical evidence indicate that it is surprisingly easy to locate optimal controls steering quantum dynamical systems to desired objectives. This has enabled the control of complex quantum systems despite the expense of solving the Schrödinger equation in simulations and the complicating effects of environmental decoherence in the laboratory. Recent work indicates that this simplicity originates in universal properties of the solution sets to quantum control problems that are fundamentally different from their classical counterparts. Here, we review studies that aim to systematically characterize these properties, enabling the classification of quantum control mechanisms and the design of globally efficient quantum control algorithms.

Keywords: Quantum control theory; Quantum control experiments; Optimal control; Open quantum systems; Laser chemistry

Contents	PAGE
1. Introduction	672
2. Hamiltonian-independent properties of solution sets to quantum control problems	675
2.1. Observable maximization	678
2.2. Quantum gate control	682
2.3. Continuous variable quantum control	683
3. Analytical features of quantum control landscape geometry	685
3.1. The role of analysis in exploring control landscape geometry	685
3.2. Analytical solutions to state control problems	688
3.3. Analytical solutions to gate control problems	690
4. Numerical exploration of quantum control landscape level sets	693
4.1. Algorithms for level set exploration	693
4.2. Quantum control mechanisms and robustness	695
4.3. Hamiltonian-dependence of landscape level set geometry	697

*Corresponding author. Email: rajchak@Princeton.Edu

5. Experimental exploration of quantum control landscapes	699
5.1. Level sets	699
5.2. Landscape topology	701
6. Quantum system controllability and landscape structure	704
6.1. Exact-time controllability of discrete quantum systems	705
6.2. Controllability of classical and continuous variable quantum systems	706
7. Computational complexity of quantum control landscapes	707
7.1. Gradient flows and search complexity	709
7.2. Relation between dynamic and kinematic gradient flows	712
8. Global search algorithms for quantum control	713
8.1. Scalar and matrix tracking algorithms	713
8.2. Extremals of the input-state map	718
9. Open quantum systems	721
10. Conclusion and future challenges	723
Appendix: Mathematical appendices	724
A.1. Critical topology	724
A.1.1. Landscape mapping	724
A.1.2. Hessian quadratic form: observable maximization	724
A.2. Maximum principle and adjoint control systems	725
A.3. Rotating wave approximation	726
A.4. Analytical solutions to state and gate control problems	727
A.4.1. Low-dimensional gate control problems	727
A.4.2. Low-dimensional state control problems	729
A.5. Diffeomorphic homotopy on control landscapes	730
A.6. Controllability on compact Lie groups	733
A.7. Kraus superoperator formalism	733
A.8. Symplectic propagators	734
Acknowledgements	734
References	734

1. Introduction

The notion of controlling quantum systems seems inherently problematic on several counts. First, the extreme sensitivity of quantum systems to environmental interactions would appear to place limits on the maximal achievable control fidelity. Second, from a numerical perspective, given the considerable cost of propagating the Schrödinger equation, unless the search space for optimal controls has particularly simple properties, it would appear impossible to locate controls for all but the smallest quantum systems in reasonable time. However, once the methods of optimal control

began to be applied to molecular systems (thanks to remarkable advances in laser pulse shaping technology) it rapidly became clear that quantum control was not an ill-fated concept, but rather, that controlling quantum systems was surprisingly easy. In the laboratory, this conclusion was particularly apparent in the case of so-called adaptive learning control of quantum dynamics, wherein learning (i.e., typically genetic) algorithms are applied to search the space of laser control parameters for the maximization of the expectation value of a quantum observable. This search space is high-dimensional, normally suggesting that it should be replete with local optima and other unfavourable features that would trap unsophisticated search algorithms, especially in the presence of environmental decoherence. The repeated successes of quantum optimal control experiments and simulations indicated that the so-called 'curse of dimensionality', common in the theory of optimization in high-dimensional spaces, was not prohibitive here.

This attractive circumstance for quantum control rests on the mathematical underpinnings of quantum theory being surprisingly simple, owing to the linearity of quantum mechanics and the unitarity of the accompanying transformations. Although quantum systems can be highly sensitive to environmental perturbations, the rules governing their dynamics are in many ways simpler than those governing classical dynamics. Furthermore, the presence of an environment, rather than being an impediment, may under the right conditions aid the control process. Recent work has aimed to understand the precise mathematical properties of quantum mechanics responsible for the surprising simplicity with which quantum phenomena can be controlled. Analytical, numerical and experimental treatments of the problem have been explored. In order to enable the systematic study of these features, the notion of a quantum control landscape, defined as the map between the space of time-dependent controls and associated values of the objective functional, was introduced (figures 1 and 2).

From an analytical perspective, it was recently found that for several classes of low-dimensional problems, it is possible to exactly solve for quantum optimal controls, without any need for numerical search. By contrast, for the analogous classical problems, analytical solutions do not exist. Of course, analytical solutions are still only possible for special small systems; however, beyond this, it has become clear that the numerical or experimental search for optimal controls is often easier for quantum systems than for classical systems. In this regard, the topology of the search space is of fundamental importance. Evidence suggests that the landscapes for both observable maximization and control of dynamical transformations have simpler topological properties for quantum versus classical systems, contributing to rapid convergence of numerical or experimental searches for effective controls. Moreover, besides the simplicity of locating quantum controls, it has been observed that the controls themselves have remarkably simple functional properties, in some cases enabling a direct interpretation of the mechanisms involved in steering about the dynamics.

The ease of locating optimal quantum controls, and the comparatively simple structure of these controls, have pervasive implications for a wide range of quantum technologies. The study of quantum control landscapes is motivated by the practical goals of achieving higher objective function yields and designing control fields with desired properties, but in order to attain these goals, it is necessary to embrace the

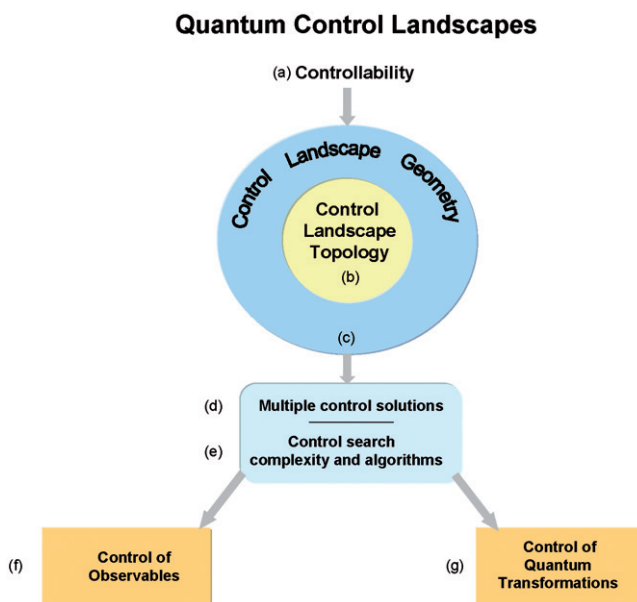


Figure 1. [Colour online] A control landscape is defined as the map between the time-dependent control Hamiltonian and associated values of the control cost functional. The entry point into their study is (a) the controllability of the quantum system, which allows search algorithms to freely traverse the landscape. Quantum control landscape features can be conveniently subdivided into those pertaining to (b) critical topology, i.e., the maxima, minima and saddle points of the landscape, and (c) landscape geometry, namely the characteristic local structures encountered while climbing toward the global optimum. Study of the geometry of quantum control landscapes reveals (d) the existence of multiple control solutions corresponding to any given objective function value. The topology and geometry of quantum control landscapes together determine (e) the search complexity of the control problems, i.e., the scaling of the effort required to locate optimal controls. An ultimate goal in the study of quantum control landscapes is the design of global search algorithms that attain lower bounds on this search complexity. Such algorithms may be applied to either of the two major classes of quantum system manipulation problems, (f) control of quantum observables and (g) control of quantum dynamical transformations (i.e., unitary propagators). An overarching conclusion pertaining to both these types of landscapes is that they contain no suboptimal traps, which has broad-scale implications for both the experimental and computational feasibility of quantum control.

mathematical framework that underlies the remarkable properties of these landscapes. The origin of these counterintuitive properties, and their differences with respect to classical control, constitute the primary subject of this review.

The review is organized as follows. In section 2, we examine the topology of solution sets to quantum observable and unitary transformation control problems. Section 3 reviews the analytical solutions obtained for low-dimensional quantum control problems. In section 4, we consider numerical studies of the solution sets to higher dimensional quantum control problems without analytical solution, focusing on the degeneracy in the set of controls that reach the same objective. Then, section 5 reviews experimental work that has probed the structure of these level sets of multiple solutions, as well as the topology of quantum control landscapes. Section 6 examines how the controllability of quantum systems, compared to that of classical systems, impacts control landscape properties, in particular with respect to search efficiency. In sections 7 and 8, we review approaches to the quantification of quantum control search effort

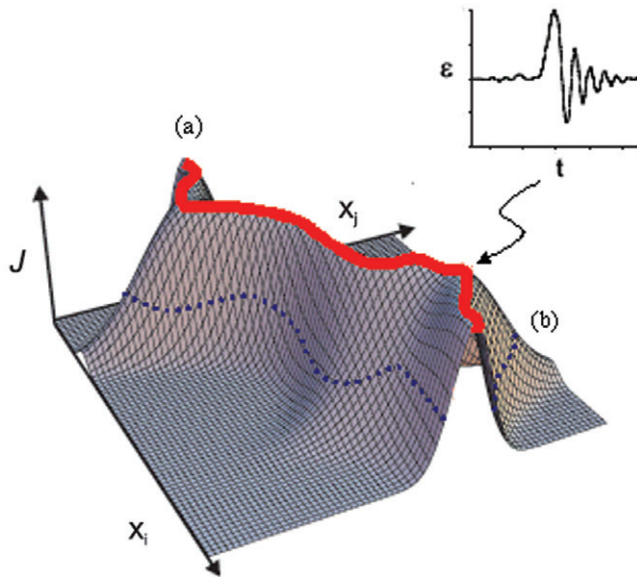


Figure 2. In the most common case of state-to-state population transfer, there are no local optima and the global maximum is a continuous manifold, depicted in the picture as the solid curve. Level sets (b) of the control landscape correspond to control fields that produce the same objective function value at the final dynamical time T (shown as a dotted line).

complexity and the design of global search algorithms that aim to attain lower bounds on complexity scaling. In section 9, we consider the effects of quantum decoherence on the structure of quantum control landscapes derived in previous sections. Finally, in section 10, we draw conclusions and discuss future directions.

2. Hamiltonian-independent properties of solution sets to quantum control problems

The most fundamental property of landscapes of solution sets to variational problems is their critical topology, i.e., the number of solutions, their associated functional values, and their optimality status (figure 1). In the context of optimal control theory, these critical points correspond to (possibly suboptimal) solutions of the optimal control problem, including both global and local minimizers of the objective functional.

An early work [1] explicitly explored the multiplicity of solutions to quantum optimal control problems aimed at the maximization of the expectation value of an observable operator at a final dynamical time T . Through the application of a perturbation theory approach to the nonlinear variational equation, it was shown that in general, a denumerably infinite number of solutions (control fields) exist to such quantum control problems. Moreover, multiple unitary propagators $U(T)$ are typically associated with the various possible local and global optima of the objective. Recent work, discussed throughout this review, has aimed to identify how the infinite number of solutions to quantum optimal control problems are distributed among the global and

local optima of the control landscape. The optimality status of these solutions plays an important role in determining the performance of local search algorithms as they traverse the landscape. A remarkable feature of quantum optimal control landscapes is that these properties can all be determined analytically for many problems of interest, whereas for general variational problems outside of quantum control, such information is exceedingly difficult to acquire. Moreover, many of these properties are independent of the Hamiltonian of the quantum system.

In this section we review work aimed at characterizing the critical topologies of quantum optimal control problems. In the absence of additional information, it is natural to expect that the control landscape will possess multiple local maxima and minima that are capable of trapping the search for optimal controls at suboptimal values of the objective. We will show that in stark contrast to this generically expected situation, the critical topologies of quantum optimal control landscapes are surprisingly simple. In what follows we consider a controllable (see Appendix A.6 for a review of the definition of controllability) quantum system of N discrete levels whose dynamics are driven by the Hamiltonian $H = H[H_0, \{\varepsilon_k\}]$, depending on a free Hamiltonian H_0 describing the uncontrolled evolution of the system and an appropriate set $\{\varepsilon_k\}$ of control variables (e.g., phases and amplitudes in an optimal control experiment (OCE) pulse shaper).

A generic quantum optimal control cost functional can be written:

$$J = \Phi(U(T), T) - \text{Re} \left[\text{Tr} \int_0^T \left\{ \left(\frac{\partial U(t)}{\partial t} + \frac{i}{\hbar} H(t)U(t) \right) \beta(t) \right\} dt \right] - \lambda \int_0^T f^0(\varepsilon(t)) dt \quad (1)$$

where $\beta(t)$ is a Lagrange multiplier operator constraining the quantum system dynamics to obey the Schrödinger equation, $\varepsilon(t)$ is the time-dependent control field, and λ weights the importance of an auxiliary physically motivated penalty term on the field. This latter penalty decreases the degeneracy of solutions to the optimal control problem; a common choice for f^0 in simulations is $1/s(t)|\varepsilon(t)|^2$, where $s(t)$ is the pulse envelope, corresponding to a penalty on the total field fluence. Other, even more general cost functions can easily be generated with alternative choices for f^0 as well as additional terms involving the evolving quantum state.

The two most common types of quantum optimal control problems are the maximization of the expectation value of a Hermitian observable and the maximization of the fidelity of a quantum unitary transformation. Optimizing the expectation value of a Hermitian observable operator describes a broad variety of problems in quantum control, such as performing selective chemical fragmentation and rearrangement [2, 3], redirecting energy transfer in biomolecules [4], creating ultra-fast optical switches and tailoring high harmonic generation [5]. This problem corresponds to the following choice of Φ :

$$\Phi_1(U) = \text{Tr}[U(T)\rho(0)U^\dagger(T)\Theta], \quad (2)$$

where $\rho(0)$ is the initial density matrix of the system, Θ is an arbitrary observable operator, and T is the final dynamical time.

The optimal control of quantum unitary transformations has recently received increasing attention due to its applications to the field of quantum information processing (QIP). Over the past several years, it has become clear that the physical implementation of logical gates in QIP, which are represented by unitary propagators, may be facilitated by optimal control theory (OCT) [6–9]. However, as we will see below, the optimal control of unitary transformations also has applications to population transfer in atoms and molecules. Within OCT, the problem of maximizing the fidelity of a dynamical transformation W can be framed using [6]:

$$\Phi_2(U) = \sum_{i,j} |W_{ij} - U_{ij}(T)|^2 \tag{3}$$

$$= 2N - 2\text{ReTr}[W^\dagger U(T)] \tag{4}$$

where W is the target unitary transformation.

Solutions to these optimal control problems correspond to the condition $\delta J/\delta \varepsilon(t) = 0$. In this section, we assume $\lambda = 0$, and show that under this assumption many properties of the critical points of the functionals Φ_1 and Φ_2 can be characterized analytically. An infinitesimal functional change in the Hamiltonian $\delta H(t)$ produces an infinitesimal change in the dynamical propagator $U(t, 0)$ as follows:

$$\delta U(t, 0) = \frac{i}{\hbar} \int_0^t U(t, t') \delta H(t') U(t', 0) dt'. \tag{5}$$

It can be shown that for the objective functions Φ_1 and Φ_2 the respective changes in Φ are $\delta \Phi_1 = -\text{Tr}[\Theta(T), \rho(0)] U^\dagger(T, 0) \delta U(T, 0)$, where $\Theta(T) = U(T) \Theta U^\dagger(T)$, and $\delta \Phi_2 = \text{Tr}[(W^\dagger U - U^\dagger W) U^\dagger(T, 0) \delta U(T, 0)]$. In the special case of the electric dipole approximation, the Hamiltonian has the form $H(t) = H_0 - \mu \cdot \varepsilon(t)$, where H_0 is the internal Hamiltonian of the system and μ is the electric dipole operator. Then $\delta H(t) = \nabla_\varepsilon H(t) \cdot \delta \varepsilon(t)$, $U^\dagger(T, 0) \delta U(T, 0) = -\frac{i}{\hbar} \int_0^T U^\dagger(t, 0) \nabla_\varepsilon H(t) U(t, 0) \cdot \delta \varepsilon(t) dt$, and the gradients of the control objective functionals can be written:

$$\frac{\delta \Phi_1}{\delta \varepsilon(t)} = \text{Tr}([\Theta(T), \rho(0)] B(t)), \tag{6}$$

and

$$\frac{\delta \Phi_2}{\delta \varepsilon(t)} = \text{Tr}((W^\dagger U - U^\dagger W) B(t)), \tag{7}$$

where

$$B(t) \equiv \frac{i}{\hbar} U^\dagger(t, 0) \nabla_\varepsilon H(t) U(t, 0), \tag{8}$$

and we have adopted the shorthand notation $U \equiv U(T)$. Within the dipole approximation

$$B(t) = -\frac{i}{\hbar} U^\dagger(t, 0) \mu U(t, 0). \quad (9)$$

The variational problems of optimal control theory admit two types of minimizers[†] according to the chain rule,

$$\frac{\delta J}{\delta \varepsilon(t)} = \frac{dJ}{dU} \cdot \frac{\delta U}{\delta \varepsilon(t)}. \quad (10)$$

The first type of minimizer corresponds to those control Hamiltonians that are critical points of the control objective functional, but are not critical points of the map between control fields and associated dynamical propagators (i.e., points at which $dJ/dU=0$, while the Frechet derivative mapping from the control variation $\delta\varepsilon(t)$ to $\delta U(T)$ at $t=T$ is surjective). The second type corresponds to critical points of the latter map (i.e., points at which the mapping from $\delta\varepsilon(t)$ to $\delta U(T)$ is not locally surjective) [10]. In this section, we consider the first type of critical point, which are referred to as kinematic critical points or normal extremal controls. The second type, which are called abnormal extremal controls, are comparatively rare in quantum control problems, and we defer their study to section 7.

The class of normal extremal controls is Hamiltonian-independent and captures the most universal topological features of general quantum control landscapes. Because the map $\varepsilon(t) \rightarrow U(T)$ is locally surjective at these points, each of the matrix elements U_{pq} must be uniquely addressable by the control field $\varepsilon(t)$ for all p and q values in keeping with U being unitary, i.e., the set of N^2 functions $\delta U_{pq}/\delta \varepsilon(t)$ should be linearly independent. Hence, the critical condition for normal extremal controls is equivalent to $dJ/dU=0$ for arbitrary $\delta U/\delta \varepsilon(t)$. Moreover, it can be shown (see Appendix A.1.1) that the optimality status (i.e., minimum, maximum or saddle point) of a critical control field $\varepsilon(t)$ will be equivalent to that of the resulting propagator $U(T)$ on the unitary group $U(N)$. The local surjectivity of $\varepsilon(t) \rightarrow U(T)$ has important connections to the controllability of the quantum system, discussed in. In what follows, we provide expressions for the gradient and Hessian of the above objective functions on both the domain of control fields and unitary propagators, and use these results to characterize their critical topologies.

2.1. Observable maximization

In 1937, John von Neumann first addressed the critical topology of a problem that has direct implications for optimizing quantum observables [11, 12]. Although it is unclear

[†]In mechanics, the Lagrangian functional that determines the equations of motion is uniquely determined by symmetries of the system. However, in optimal control theory, the objective functional, which determines the control law, is chosen by the controller. This distinction lends an additional component to the study of the topology of optimal control problems, namely the topology of the map between dynamical propagators and associated values of the chosen objective function.

whether von Neumann anticipated the applications of his work to quantum control, this paper may be considered the first work in the theory of quantum optimal control landscapes. This work was recently extended by several authors [13, 14].

Within the electric dipole approximation, the gradient (4) can be explicitly written [15]:

$$\begin{aligned} \frac{\delta\Phi_1}{\delta\varepsilon(t)} &= -\frac{i}{\hbar} \text{Tr}[[\Theta(T), \rho(0)]\mu(t)] \\ &= \frac{i}{\hbar} \sum_i \rho(0)\langle i|\Theta(T)\mu(t) - \mu(t)\Theta(T)|i\rangle \\ &= \frac{i}{\hbar} \sum_{i=1}^n p_i \sum_{j=1}^N [\langle i|\Theta(T)|j\rangle\langle j|\mu(t)|i\rangle - \langle i|\mu(t)|j\rangle\langle j|\Theta(T)|i\rangle] \end{aligned} \tag{11}$$

where the initial density matrix is given as $\rho(0) = \sum_{i=1}^n p_i|i\rangle\langle i|, p_1 > \dots > p_n > 0, \sum_{i=1}^n p_i = 1$, and $\mu(t) \equiv U^\dagger(t)\mu U(t)$. The local surjectivity of $\varepsilon(t) \rightarrow U(T)$ at normal extremal controls implies that the N^2 functions of time $\langle i|\mu(t)|j\rangle$ are linearly independent. As discussed above, under this assumption the critical condition is equivalent to that on the domain of unitary propagators, $d\Phi_1/dU=0$. Because the gradient $\delta\Phi_1/\delta\varepsilon(t)$ depends on the eigenvalue spectra of $\rho(0)$ and Θ , it is convenient to simplify the analysis by investigating the critical topology on the domain $U(N)$. Expanding the argument of the objective function by $U \rightarrow U \exp(iAs)$, where s parametrizes an arbitrary curve in the Lie algebra of $U(N)$ in the direction A , the critical condition $d\Phi_1/dU=0$ can be expressed as

$$\frac{d\Phi_1}{ds} = i\text{Tr}(A[U^\dagger\Theta U, \rho(0)]) = 0. \tag{12}$$

The maximal subset of $U(N)$ which satisfies this condition is composed of the matrices of the form

$$\hat{U}_l = QP_lR^\dagger \tag{13}$$

where $P_l, l=1, \dots, N!$ is an N -fold permutation matrix whose nonzero entries are complex numbers $e^{i\phi_1}, \dots, e^{i\phi_N}$ of unit modulus, and $\rho(0) = Q^\dagger \varepsilon Q$ and $\Theta = R^\dagger \lambda R$. $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_N$ and $\lambda_1, \lambda_2, \dots, \lambda_N$ are the eigenvalues of $\rho(0)$ and Θ with associated unitary diagonalization transformations Q and R , respectively. The critical set is the union of N -tori $\bigcup_l T_l^N$ with each torus corresponding to a distinct permutation P_l .

The number of critical submanifolds corresponding to suboptimal landscape values scales factorially with system dimension for fully nondegenerate $\rho(0)$ and Θ . If $\rho(0)$ and Θ have arbitrary numbers of degeneracies D_1, \dots, D_m , and E_1, \dots, E_n , respectively, it can be shown [16] that the critical submanifold dimension on the domain $U(N)$ for a

particular critical manifold M_k is

$$d(M_k) = \sum_{s=1}^m D_s^2 + \sum_{t=1}^n E_t^2 - \sum_{l=1}^r o_l^2. \quad (14)$$

where the o_i 's are the overlap numbers (numbers of overlapping elements) between the degenerate blocks D_1, \dots, D_m and E_1, \dots, E_n , for the permutation matrix corresponding to that manifold. These degeneracies cause neighbouring submanifolds to merge, and give rise to subspaces that are invariant to eigenvalue permutations. If $\rho(0)$ or Θ is a pure state projector, the number of critical manifolds scales linearly with Hilbert space dimension N [17]. In the limiting case where both $\rho(0)$ and Θ are pure state projectors, Φ at the extrema only has the values zero and unity, corresponding respectively to no control or perfect control (i.e., the landscape is convex).

In the electric dipole approximation, the Hessian of the objective function can be written

$$\mathcal{H}^\varepsilon(t, t') = -\frac{1}{\hbar^2} \text{Tr}([\Theta(T), \mu(t)], \mu(t')) \rho(0) \quad (15)$$

The Hessian quadratic form (HQF), defined as

$$\langle \omega | \mathcal{H} | \omega \rangle = \int_0^T \int_0^T \omega(t) \frac{\delta^2 \Phi_1}{\delta \varepsilon(t) \delta \varepsilon(t')} \omega(t') dt dt' \quad (16)$$

where $\omega(t)$ is an arbitrary real function, is a polynomial representation of the Hessian that facilitates the identification of increasing, null and decreasing directions at each critical point. The explicit representation of the HQF for the general case of nondegenerate ρ and Θ is complicated and is reviewed in Appendix A.1.2. Based on this representation, it can be shown that all of the suboptimal critical submanifolds corresponding to Φ values less than the global maximum are saddle regions, and thus will not act as traps for optimal control searches.

As we will show in section 3 below, the dimension of the global optimum on $U(N)$ is useful for exploring the degeneracy of solutions to quantum observable control problems. This number ranges from N , for fully nondegenerate $\rho(0)$ and Θ , to $N^2 - (2N - 2)$, when $\rho(0)$ and Θ are both pure state projectors. In the former case, the number of decreasing directions at the global maximum is the greatest ($N^2 - N$), whereas in the latter case it is the smallest ($2N - 2$). Note that the landscape mapping analysis in Appendix A.1.1 indicates that the number of positive and negative principal axis directions of the Hessian matrix will be preserved upon passage from the domain $U(N)$ to the domain $\varepsilon(t)$, with the remainder of the directions on $\varepsilon(t)$ being flat. A recent numerical analysis [18] confirmed these predictions, by considering the problem of optimizing the expectation value of a pure state projector over a $N=4$ quantum system, initialized in a pure state, with

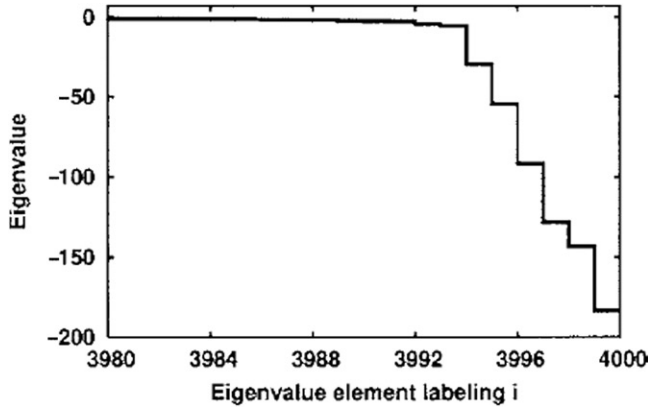


Figure 3. The dominant eigenvalues of the Hessian matrix for state-to-state population transfer in a simple four level system. Only the last 20 eigenvalues are shown, and the remaining ones are essentially zero. For this system with dimension $N=4$, it is evident that the $2N-2$ rule is satisfied with six nonzero eigenvalues being present. (From [18].)

the goal of $|0\rangle \rightarrow |3\rangle$ population transfer. In this case, the Hessian can be expanded on a basis of $2N-2$ linearly independent functions β_i , as

$$\mathcal{H}_{k'k} = \frac{\partial^2 \Phi_1}{\partial \eta_{k'} \partial \eta_k} \tag{17}$$

$$= \int_0^T \int_0^T dt dt' \frac{\delta \varepsilon(t')}{\partial \eta_{k'}} \frac{\delta^2 \Phi_1}{\delta \varepsilon(t') \delta \varepsilon(t)} \frac{\partial \varepsilon(t)}{\partial \eta_k} \tag{18}$$

$$= - \sum_{l=1}^{2N-2} \beta_l(k') \beta_l(k), \tag{19}$$

where η is a vector of appropriate control parameters. Diagonalization of the Hessian evaluated at the landscape maximum revealed $10 = 4^2 - 2 \times 4 + 2$ zero-valued eigenvalues, as expected (figure 3). Importantly, when the dipole coupling strengths were reduced to negligible values for all but two of the levels, the number of zero-value Hessian eigenvalues dropped to $2 = 2^2 - 2 \times 2 + 2$, revealing the reduction of the four-level to a quasi-two-level system. Thus, the robustness of observable maximization control solutions depends not only on the actual Hilbert space dimension, but also on the effective number of states that are accessible to the dynamics.

A related work recently studied the critical topology of observable maximization on the domain $SU(N)$ (instead of $U(N)$) [19]. It was shown that the set of maxima decomposes into two connected components, and an explicit description of both components was derived. Separately, Glaser *et al.* [20] examined the topology of observable maximization for non-Hermitian observables, such as those operators that arise in quadrature detection. These works, as well as others [8, 9], were concerned with

the problem of identifying optimal controls for quantum observables in two steps: first, solving numerically for the set of unitary matrices U that maximize the expectation value $\text{Tr}(U\rho(0)U^\dagger\Theta)$ of the observable Θ , and then, finding a control field $\varepsilon(t)$ that produces the quantum gate U at time $t = T$. We discuss analytical solutions to the latter problem in section 3.

2.2. Quantum gate control

In the electric dipole approximation, the critical point condition corresponding to the gradient (5) can be explicitly written [21, 22]

$$\frac{\delta\Phi_2(U)}{\delta\varepsilon(t)} = -\frac{i}{\hbar} \sum_i \sum_j (W^\dagger U - U^\dagger W)_{ij} \langle j | \mu(t) | i \rangle = 0. \quad (20)$$

The critical topology of the gate fidelity cost function on $U(N)$ was first studied by Frankel [23]. Assuming local surjectivity of $\varepsilon(t) \rightarrow U(T)$, a necessary and sufficient condition for the critical points is $W^\dagger U = U^\dagger W$ or

$$(W^\dagger U)^2 = I. \quad (21)$$

The solutions to this equation are the roots of I , i.e.,

$$W^\dagger U = \text{diag}(\lambda_1, \dots, \lambda_n, \lambda_i = (-1)^{n_i}), \quad n_i = 0, 1.$$

These solutions fall into equivalence classes indexed by the number of eigenvalues $\lambda_i = 1$. Thus, there are a total of $N + 1$ critical manifolds, taking on Φ values of $0, 4, \dots, 4N$. The number of suboptimal critical regions hence grows only linearly with respect to the system Hilbert space dimension N , a slower scaling than that of the landscape for observable maximization when both ρ and Θ are nondegenerate full rank matrices.

The Hessian at the critical points is

$$\mathcal{H}(t, t') = \frac{1}{\hbar^2} \text{Tr}\{W^\dagger U(\mu(t)\mu(t') + \mu(t')\mu(t))\} \quad (22)$$

which can be expanded as [22]

$$\begin{aligned} \mathcal{H}(t, t') = & \frac{2}{\hbar^2} \sum_i (-1) \langle i | \mu(t) | i \rangle \langle i | \mu(t') | i \rangle + \frac{2}{\hbar^2} \sum_i \sum_{j>i} [(-1)^{n_i} + (-1)^{n_j}] \\ & \times [\text{Re}(\langle i | \mu(t) | j \rangle) \times \text{Re}(\langle i | \mu(t') | j \rangle) + \text{Im}(\langle i | \mu(t) | j \rangle) \times \text{Im}(\langle i | \mu(t') | j \rangle)] \quad (23) \end{aligned}$$

The number of positive, negative and null directions at the critical points can be determined by simple inspection of the Hessian quadratic form. The expression for the HQF is

$$\begin{aligned} \langle \omega | \mathcal{H} | \omega \rangle &= \frac{2}{\hbar^2} \sum_i (-1)^{n_i} \left(\int_0^T \langle i | \mu(t) | i \rangle \omega(t) dt \right)^2 + \frac{2}{\hbar^2} \sum_i \sum_{j>i} [(-1)^{n_i} + (-1)^{n_j}] \\ &\times \left[\left(\int_0^T \text{Re}(\langle i | \mu(t) | j \rangle \omega(t) dt) \right)^2 + \text{Im}(\langle i | \mu(t) | j \rangle \omega(t) dt)^2 \right] \end{aligned} \quad (24)$$

At the suboptimal critical points, there are $N - m$ even and m odd integers n_i . It can be shown that the number of positive and negative Hessian eigenvalues equals the number of odd and even n_i , respectively, and that the remaining eigenvalues are zero [21]. The numbers of positive and negative directions at a critical point m are thus

$$h_+ = m^2; h_- = (N - m)^2, \quad (25)$$

whereas all the remaining principal axis directions are flat. In particular, all the local suboptima are saddle manifolds, and we see again that there are no local traps in the quantum control landscape. In contrast to the multiplicity of unitary matrices that solve the observable maximization problem, the kinematic critical regions of the landscape corresponding to global optima are isolated unitary matrices [21], although an infinite number of controls may steer the system to those matrices.

2.3. Continuous variable quantum control

The kinematic landscape critical topology for controlling continuous variable quantum dynamical transformations for systems with quadratic Hamiltonians was recently studied [24]. These systems are relevant for the implementation of continuous variable quantum information processing [25]. Continuous variable transformations can be realized by harmonic oscillators, molecular rotors, or coupled modes of the electromagnetic field. Dynamical transformations for such systems can be represented by symplectic propagators (Appendix A.8).

The critical topology of such landscapes offer insight into the differences between control landscapes for discrete and continuous quantum systems. Since the (quantum) symplectic gate U is a faithful unitary representation of a symplectic matrix S , it is reasonable to define the gate fidelity analogously to that for discrete gates as

$$\mathcal{J}[\varepsilon(t)] = \text{Tr}(S - W)^T(S - W) + (s - w)^T(s - w), \quad S_s \in \text{ISp}(2N, \mathbb{R}), \quad (26)$$

where s and w denote phase space displacements. Importantly, the symmetries of this objective functional again permit an analytical characterization of the critical topology, although this topology is more complex than that of the control landscapes for discrete

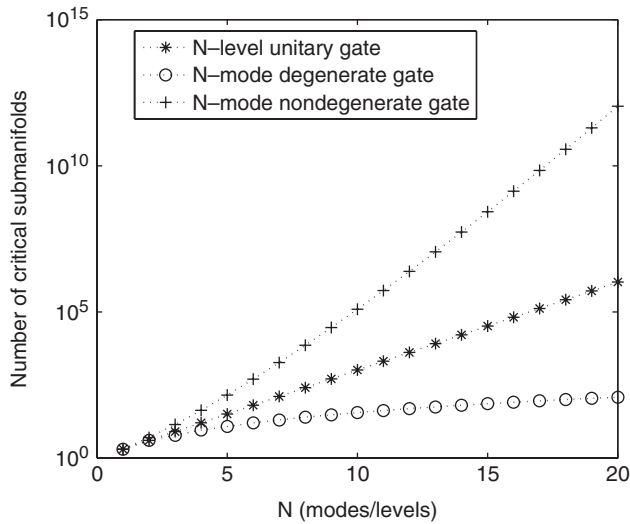


Figure 4. The scaling of the numbers of critical submanifolds for discrete quantum (unitary) and continuous quantum (symplectic) gate control landscapes with system size [24]. Degenerate/nondegenerate refer to the singular values of the symplectic matrix representing the continuous quantum propagator. Landscapes for the control of discrete quantum propagators all have identical critical topologies, whereas those for continuous quantum propagators are target-dependent. See the text for definitions of quantum gate terminology.

quantum systems. If we write the singular value decomposition of W as $W = UEV$, the critical submanifolds can be expressed as [24]:

$$S^* = R^T D R, \quad R \in \text{Stab}(E), \quad (27)$$

where R is an arbitrary orthogonal symplectic matrix in the stabilizer of E in $\text{OSp}(2N, \mathbb{R})$:

$$\text{Stab}(E) = \{R \in \text{OSp}(2N, \mathbb{R}) \mid R^T E R = E\} \quad (28)$$

$$= \text{OSp}(2n_0) \times O(n_1) \times \cdots \times O(n_s). \quad (29)$$

The characteristic matrix D consists of different operations on the separate modes represented by diagonal blocks of three different types, depending on the singular values of the target gate W . In particular, D , and hence the critical topology, differs depending on the degeneracy of the singular values. Again, all suboptimal critical points are found to be saddle manifolds, indicating that the control landscape for these infinite-dimensional quantum gates is devoid of local traps.

Although this landscape is devoid of traps, it can be shown [26] that the lower symmetry of the continuous variable fidelity function, compared to that of the discrete quantum fidelity function, can result in a more rapid scaling of the number of critical manifolds. Moreover, the critical topology is dependent on the target gate (figure 4).

The dynamical transformations of classical systems with quadratic Hamiltonians can also be represented by symplectic propagators. However, unlike quantum observables, classical observables may be expressed as arbitrary smooth functions on phase space (and hence the space of symplectic propagators); thus, unlike quantum control landscapes, classical observable control landscapes have no universally characterizable critical topology, either for quadratic or more general classes of Hamiltonians.

For each of the control landscapes discussed above, an important issue is the size of the attracting regions of these critical manifolds and the behaviour of the gradient flows of the objective function around them. Unlike the critical topology, these gradient flows (which represent a geometric property of the landscape) are Hamiltonian-dependent, and we will revisit them, including their connection to topology, in section 7.

Although the above analytical results were derived under the assumption that the fluence penalty coefficient $\lambda = 0$, they remain valid in many cases in the presence of a significant cost on the field fluence. In particular, in the case of observable maximization where $\rho(0)$ and Θ are pure state projectors, it was shown for the model system described in section 2.1 [18] that the Hessian retains $2N - 2$ nonzero eigenvalues in the presence of substantial fluence costs. In the next section, we discuss how the imposition of such constraints facilitates the identification of analytical solutions for certain (low-dimensional) observable and gate control problems in quantum mechanics.

3. Analytical features of quantum control landscape geometry

3.1. The role of analysis in exploring control landscape geometry

The previous sections showed that the critical point *topology* of the most common quantum optimal control problems can be established analytically, and display properties favourable for optimal search. In this section, we show that analytical statements can also be made regarding the *geometry* of quantum control landscapes. The geometry of a control landscape can be broken down into two components: (1) the relationship among controls producing the same objective function value (level sets), and (2) the search trajectory followed in locating the optimal objective function value.

The geometry of quantum control landscapes is Hamiltonian-dependent. The universal monotonicity of quantum control landscapes ensures the convergence of local algorithms, but does not provide a direct measure of the search effort involved in finding optimal solutions. A further reduction in the search effort involved in locating optimal controls could be aided by analytical landscape geometry. Although these analytical results may not fully identify the solution set to an optimal control problem, they may nonetheless restrict its structure.

In the case that no constraints are placed directly on the controls or on the time required for reaching the objective, an infinite number of solutions exist to quantum control problems. As such, analytical results pertaining to the geometry of the landscape are restricted to statements regarding generic features of the controls. When constraints or auxiliary costs are imposed in the objective function, it is possible in the case of many low-dimensional problems to explicitly solve for the optimal controls. In

these cases, multiple solutions may still exist, but they are distributed among distinct unitary propagators.

Studying analytical solutions to quantum control problems provides insight into the geometry of control landscapes in several ways. First, there can exist quantum symmetries that reduce the dimensionality of the domain of control fields $\varepsilon(t)$ over which the control search must be carried out. From either a computational or experimental perspective, this means that simpler parametrizations of the control fields can be used in optimizations, thereby reducing search effort.

Second, several important low-dimensional quantum optimal control problems are analytically soluble (integrable) once auxiliary constraints are imposed on the objective function. In particular, the problem of realizing a quantum unitary transformation in minimal time or with minimal fluence can be solved for Hilbert spaces of several different dimensions. As shown in section 2, a continuous submanifold of unitary matrices is associated with the maximization of any observable expectation value. The set of points on this submanifold can be identified (either analytically or numerically) at minimal cost. Therefore, a submanifold of the set of all solutions (one corresponding to each unitary matrix) to any observable control problem for such systems can be obtained through a combination of analytical and numerical methods with substantially reduced search effort, due to the existence of analytical solutions to gate control problems [8, 9]. This submanifold consists of all solutions that minimize the auxiliary cost (e.g., field fluence). Thus, analytical solutions to gate control problems provide a means of further delineating the level set geometry of observable control landscapes.

For these reasons, we review in this section analytical results pertaining to the solution of quantum control problems. We discuss (1) analytical results pertaining to control mechanisms in arbitrary Hilbert space dimension; (2) the integrability of low dimensional problems (which in some cases relies on (1)); and (3) the reasons that integrability breaks down in higher dimensions, whereas mechanistic simplicity is retained.

The basic theorem of optimal control theory used for solving problems of this type is the Pontryagin maximum principle, reviewed in Appendix A.2. Consider the control problem of minimizing a cost associated with steering the system

$$\dot{x} = f(x, \varepsilon), \quad x \in \mathbb{R}^n, \quad \varepsilon \in \Omega \subset R^k \quad (30)$$

from some initial state $x(0) = x_0$ to some final state x_1 . For quantum gate and state control, we are dealing with the right-invariant control systems

$$\dot{U} = f(U, \varepsilon) = -\frac{i}{\hbar}[H_d + \mu\varepsilon(t)]U \quad \text{and} \quad \dot{\psi} = f(\psi, \varepsilon) = -\frac{i}{\hbar}[H_d + \mu\varepsilon(t)]\psi, \quad (31)$$

respectively (see Appendix A.6 for a definition of right-invariance). Note that these equations can be generalized to the case of m -independent controls, which could take the form of, e.g., components of the time-dependent electromagnetic field coupled to independent Pauli spin operators in an NMR control experiment. When analytically

solving for optimal controls satisfying Pontryagin’s maximum principle (PMP), the maximization of the function Φ in equation (1) is often imposed as a fixed constraint in addition to the Schrödinger equation, and the ‘cost’ to be minimized takes the form of the last term in this equation. Such problems are framed most conveniently in Hamiltonian form. If we denote the cost as $\int_0^T f^0(x, \varepsilon)dt$, then the PMP-Hamiltonian function is defined as:

$$h(x, \lambda, \varepsilon) = \langle \lambda, f(x, \varepsilon) \rangle + \lambda_0 f^0(x, \varepsilon) \tag{32}$$

where the first term on the RHS is either a matrix or vector inner product, depending on whether the problem is defined on the space of state vectors or dynamical propagators, and the λ play the role of PMP-conjugate momenta (Appendix A.2). The PMP-Hamiltonian function takes on the following form for gate and state control, respectively:

$$h(M, \lambda_0, u) \equiv \text{Tr} \left[M \left(U_0^\dagger(t) (H_d + \mu \varepsilon(t)) U_0(t) \right) \right] + \lambda_0 f^0(t) \tag{33}$$

$$h(P, \lambda_0, u) \equiv \langle P, (H_d + \mu \varepsilon(t)) \psi(t) \rangle + \lambda_0 f^0(t) \tag{34}$$

where M is the conjugate PMP-momentum for gate control, P is the conjugate PMP-momentum for state control, and where we have considered only the case of pure state population transfer in the latter case, for simplicity. The Pontryagin maximum principle (Appendix A.2) then specifies the PMP-Hamiltonian equations of ‘motion’ for the control system; the solutions to these equations of motion correspond to the solutions to the control problem, i.e., trajectory/control couples $(x(t), \varepsilon(t))$.

The auxiliary cost $\int_0^T f^0 dt$ can take on several canonical forms. The most common are (assuming m independent controls): (1) The field fluence for fixed transfer time T , $E = \int_0^T \sum_{i=1}^m \varepsilon_i^2 dt$, and (2) The total transfer time, with fluences either unconstrained or subject to the constraint $\int_0^T \sum_{i=1}^m \varepsilon_i^2 \leq C$, with C an arbitrary positive constant, on the field amplitudes. Note that in the case of (2), the final time in $\Phi(T)$ can be taken to be a variable rather than a fixed parameter. An important distinguishing feature between solutions corresponding to these different auxiliary costs is the time-dependent structure of the corresponding optimal control fields. When controls are bounded, the optimal fields are typically resonant with the system transition frequencies, as discussed further below. By contrast, when controls are unbounded, the optimal fields are often singular (i.e., short sequences of hard pulses) [8].

Since landscape geometry is Hamiltonian-dependent, it is important to specify the class of Hamiltonians when studying analytical solutions. We restrict the analysis primarily to state and gate control problems defined on the special unitary group $SU(N)$. Consider the following right-invariant control system on $U(N)$ with m controls:

$$\dot{U}(t) = -\frac{i}{\hbar} \left[H_d + \sum_{i=1}^m \mu_i \varepsilon_i(t) \right] U(t) \tag{35}$$

The matrices H_d , μ_i , $i=1, \dots, m$ are skew-Hermitian matrices. If we write $\bar{H}_d = D_{H_d} + H_d$ and $\bar{\mu}_i = D_{\mu_k} + \mu_k$, $i=1, \dots, m$, with

$$D_{H_d} = \text{diag}\left(\frac{1}{2}\text{Tr}(\bar{H}_d), \frac{1}{2}\text{Tr}(\bar{H}_d)\right), \quad D_{\mu_i} = \text{diag}\left(\frac{1}{2}\text{Tr}(\bar{\mu}_k), \frac{1}{2}\text{Tr}(\bar{\mu}_i)\right), \quad i=1, \dots, m, \quad (36)$$

the matrices D_{H_d} , D_{μ_i} give a pure phase contribution to the solution. These terms do not contribute to the relative phases of the components of the state vector and therefore can be neglected since states that differ only by a phase are physically indistinguishable. Thus, we consider the physically equivalent problem on the domain of traceless skew-Hermitian matrices, the Lie group $SU(N)$.

Just as analytical solutions to the Schrödinger equation exist only for the simplest quantum systems, formal analytical solutions to the PMP for quantum control problems involving arbitrary Hamiltonians are very scarce. However, for certain classes of Hamiltonians, or under certain physically reasonable approximations, control problems of real practical interest may be integrable. In what follows, we assume that the rotating wave approximation (RWA), reviewed in Appendix A.3, holds for the dynamics under consideration. It is important to note that this latter condition is often not satisfied, necessitating the use of numerical methods to solve the most general class of control problems (section 4). However, in several of the most commonly encountered quantum control problems, the RWA does hold to a reasonable approximation.

The methods of geometric control theory and sub-Riemannian geometry [8, 27] provide a means of obtaining, in certain specific cases, analytical solutions for the optimal control fields reaching a given objective. Although geometric control theory was originally developed in the context of classical control, it has recently been shown that the Pontryagin maximum principle in a geometric framework can be used to obtain analytical solutions for optimal control fields for such low-dimensional quantum control problems. We examine these specific solutions after briefly reviewing analytical results pertaining to the control mechanisms for fluence-minimizing state controls.

3.2. Analytical solutions to state control problems

For discrete quantum control problems, the existence of symmetries on the Hilbert space of states often allows a significant reduction in the dimensionality of the problem and the parameterization of the controls. This feature extends beyond the limited subset of low-dimensional problems with analytical solution to discrete quantum control problems in arbitrarily high dimensions. Let $V(t) = \sum_{i=1}^m \varepsilon_i(t)\mu_i$ denote the total time dependent control Hamiltonian. A problem of particular interest for chemical applications is where control laser fields couple only neighbouring energy levels of the system, i.e., $V_{j,k} = V_{k,j} = 0$ if $j \neq k \pm 1$. This is a common scenario in strong field control experiments. The optimal controls for these problems are often in resonance with the transition frequencies of the uncontrolled system. It can be shown that for auxiliary cost 1 above (i.e., fluence minimization) with Hamiltonians of this form, the controls will always satisfy a more general

condition of ‘weak’ resonance. In either case, the search space of the problem is then reduced from the Hilbert sphere S^{2N-1} to S^N .

Definition 1: (Resonance, weak resonance of optimal controls) A control $V_{j,k}(t)$ is said to be resonant with respect to the uncontrolled system with state function ψ if it has the following form:

$$V_{j,k}(t) = A_{j,k}(t) \exp(i[(E_j - E_k)t + \pi/2 + \phi_{j,k}]) \tag{37}$$

where $A_{j,k}(\cdot) : [0, T] \rightarrow \mathbb{R}$, $A_{j,k} = -A_{k,j}$, $\phi_{j,k} \equiv \arg(\psi_j(0)) - \arg(\psi_k(0)) \in [-\pi, \pi]$. Physically, this means that the lasers oscillate with frequency $(E_j - E_k)/2\pi$; $A_{j,k}$ describes the field amplitudes. A control $V_{j,k}$ is weakly-resonant if it is resonant in each interval of time in which the states that it is coupling (i.e., ψ_j and ψ_k) are different from zero [28].

Let us denote these latter intervals $I_{j,k,l}$ where j, k indexes the matrix elements of the Hamiltonian and l indexes the time interval where $\psi_j, \psi_k \neq 0$. Then according to the terminology above, $V_{j,k}$ is weakly resonant if

$$H_{j,k}(t)|_{I_{j,k,l}} = A_{j,k,l}(t) \exp i\phi_{j,k,l}$$

$$A_{j,k,l}(\cdot) : I_{j,k,l} \rightarrow \mathbb{R}, \quad A_{j,k,l} = -A_{k,j,l}. \tag{38}$$

The phenomenon of resonance may be viewed geometrically as originating from a rotational symmetry in Hilbert space, under the transformation $Rot_\alpha : (\psi_1, \dots, \psi_n) \rightarrow (e^{i\alpha_1} \psi_1, \dots, e^{i\alpha_n} \psi_n)$. The two admissible curves $\psi(\cdot) = (\psi_1(\cdot), \dots, \psi_n(\cdot))$ and $Rot_\alpha(\psi(\cdot))$ on $[0, T]$ have the same cost. In particular, any point within the set T_{ψ^2} generated by the action of any element of $Rot_\alpha(\psi(\cdot))$ on the state vector ψ^2 can be reached at the same cost from any point within the set T_{ψ^1} , defined analogously. Let us represent the controls as

$$V_{j,k}^l(t) \equiv V_{j,k}(t)|_l \equiv (u_{j,k}^l(t) + iv_{j,k}^l(t)) \exp(i\beta_{j,k}^l(t)), \tag{39}$$

decomposing them into non-resonant and resonant time-dependent parts. If $\psi(\cdot) : [0, T] \rightarrow S^{2n-1}$ is a minimizing trajectory between sets M_{ψ^1}, M_{ψ^2} , then the transversality condition of the maximum principle (Appendix A.2) implies that $\langle P(t), TM_{\psi(t)} \rangle = 0$. We write, $\dot{\psi}_j = \sum_k (u_{j,k}^l F_{j,k}^l(\psi) + v_{j,k}^l G_{j,k}^l(\psi))$ where $F(\psi)$ and $G(\psi)$ are vector fields subsuming the action of the resonant contribution to the (real, imaginary) control Hamiltonians on the state vector ψ . It can then be shown [29] that $G_{j,k}^l(\psi)$ is always tangent to a submanifold of S^{2n-1} whose points are reached with the same cost, i.e., $G_{j,k}^l(\psi(t)) \in TM_{\psi(t)}, \forall t$.

Therefore, in the maximum principle, $\langle P(t), G_{j,k}^l(\psi(t)) \rangle = 0$, and the maximality condition of the maximum principle implies that $v_{j,k}(t) = 0$. It follows that it is possible

to join any two eigenstates ψ_j, ψ_k by a trajectory that is in resonance. For states that are not eigenstates, the weak resonance condition holds.

State control and gate control for discrete quantum systems can both be framed in terms of the identification of geodesic trajectories under suitable metrics. The existence of a set of symmetries is important for the identification of analytical solutions to these problems. In particular, for state control problems, the reduction in control dimensionality following from resonance is essential for obtaining analytical solutions in low Hilbert space dimensions.

Analytical solutions to problems of the general class described above can be obtained for population transfer in two- and three-level quantum systems, for off-diagonal control Hamiltonians whose Lie algebra spans the entire dynamical group. For instance, for population transfer in three-level systems using fluence as the cost with two controls, the PMP-Hamiltonian (13) becomes

$$h(P, \lambda_0, u) := \langle P, (H_d + \mu_1 \varepsilon_1(t) + \mu_2 \varepsilon_2(t)) \psi(t) \rangle + \frac{1}{2} \lambda_0 (\varepsilon_1^2(t) + \varepsilon_2^2(t)). \quad (40)$$

Since our primary focus here is the relationship between multiple quantum control solutions, we relegate a summary of this problem to Appendix A.4. For our present purpose, its most important features are that the assumption of resonant controls permits the (sub-Riemannian) problem to be mapped from S^5 to S^3 , and that the resulting reduced Hamiltonian system is integrable. In four dimensions, the corresponding state control problem resides on S^7 . It appears that the Hamiltonian system given by the maximum principle is not integrable in this case, but the resonance condition still holds, and simplifies numerical search in this and higher dimensions.

Although the the class of systems described above – where controls couple only two neighbouring levels – is practically important, for more general systems the optimal controls may not be resonant or weakly resonant. Control mechanisms for such systems have been studied using numerical methods described in section 4.

3.3. Analytical solutions to gate control problems

As discussed above, quantum gate control solutions can be used to significantly reduce the search effort required to obtain solutions to the corresponding observable control problems. Analytical solutions have been found for gate control problems in dimensions 2, 4 and 8 for important classes of Hamiltonians, under the assumption of unbounded controls [8, 9]. These problems can be framed as so-called adjoint control problems, a type of sub-Riemannian control problem where the optimal control minimizes length on a geometric space under constraints on the possible paths. An essential prerequisite for their analytical solution is that the Riemannian space display certain symmetries, which in this case are endowed by the geometry of the special unitary group.

Instead of seeking fluence minimizing controls (cost 1) we consider here the minimization of the transfer time as the auxiliary cost (cost 2 above), with unbounded controls. This problem is particularly important for minimizing the time required for coherence transfer in nuclear magnetic resonance (NMR) experiments with

radiofrequency pulses; the resultant time optimal pulses outperform those typically used in NMR by a significant margin [8, 9]. In the previous section we focused attention on systems whose control Hamiltonians span the entire dynamical group. Here, we examine the more common case where the control Hamiltonians span only a subgroup of the dynamical group; if controls are unbounded, the latter assumption is required for a lower bound on the evolution time to exist, since unbounded controls can attain the target in arbitrarily small time. The PMP-Hamiltonian is then

$$h(M, \lambda_0, u) = \text{Tr} \left[M \left(U_0^\dagger(t) (H_d + \mu \varepsilon(t)) U_0(t) \right) \right] + \frac{1}{2} \lambda_0. \quad (41)$$

Note (Appendix A2) that the maximum principle for time optimal control differs slightly from that for the problem of minimal field fluence with fixed final time. In fact, for time-optimal control problems where the controls do not span the dynamical group, framing the problem in terms of an adjoint control system (and an associated ‘adjoint maximum principle’) facilitates solution, as shown below.

To understand this approach, let G denote the special unitary group $SU(N)$. Under the assumption of full controllability of the system, the algebra s generated by the entire control system $\{H_d, \mu_1, \dots, \mu_m\}$ is equal to the Lie subalgebra $su(N)$, and the corresponding group S is equal to G . We call the subalgebra generated by the controls $\{\mu_1, \dots, \mu_m\}$ l , and the corresponding subgroup K . Since the controls are unbounded, any element of K can be reached in arbitrarily small time. Consider the problem of driving the evolution from U_1 to U_2 in the shortest possible time, and let the coset $kKU_1 = \{KU_1 | k \in K\}$. We then need to find the fastest way to move from KU_1 to KU_2 , since the time required to travel anywhere within a coset is negligible.

If we decompose $G = p \oplus l$ such that p is orthogonal to l , then p represents all possible directions to move in G/K . All directions in this set can be generated by using the control Hamiltonians to place the system at appropriate starting points $k \in K$, from which the drift Hamiltonian moves the system in the directions given by $k_1^\dagger H_d k_1$. However, we cannot access all these possible directions directly. All motion in G/K is generated by the drift Hamiltonian H_d . These directions are represented by

$$Ad_K(H_d) = \{Ad_{k_1}(H_d) = k_1^\dagger H_d k_1 | k_1 \in K\} \in p, \quad (42)$$

called the adjoint orbit of H_d under the action of the subgroup K . This form of direction control has been defined as an adjoint control system, reviewed in Appendix A.2.

The goal is to find the shortest path between two points in G/K under the constraint that the tangent direction must always be in the adjoint orbit (figure 5). The shortest paths between points on a manifold subject to the constraint that the tangent to the path always belongs to a subset of all permissible directions are called sub-Riemannian geodesics. The problem of finding time optimal control laws then reduces to finding sub-Riemannian geodesics in the space G/K , where the set of accessible directions is the set $Ad_K(-iH_d)$. This problem can be framed in terms of an adjoint cost function, $f(P) = \text{Tr} = (\lambda^\dagger \mathcal{H}P)$, with $P\lambda^\dagger \in p$. The adjoint-PMP Hamiltonian is

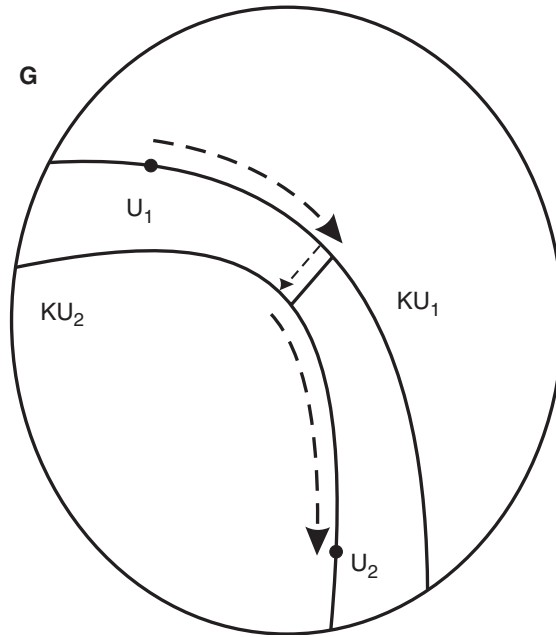


Figure 5. Time optimal path (dashed arrows) between elements U_1 and U_2 belonging to G . The long arrows depict the fast portion of the path involving movement within coset KU and correspond to the pulse; the short arrow depicts the slow portion of the path connecting different cosets and corresponds to evolution of the couplings.

$h(P(t), \lambda(t), \mathcal{H}(t)) = \text{Tr}(\lambda^\dagger(t)\mathcal{H}(t)P(t))$. Solutions to the control problem follow from the adjoint maximum principle (Appendix A2).

For one and two spin systems, G/K is a Riemannian symmetric space. In this case, the decomposition

$$g = p \oplus l, \quad p = l^\perp \tag{43}$$

satisfies the commutation relations $[l, l] \subset l$, $[p, l] = p$, $[p, p] \subset l$. This property implies that the tangent vectors to the path through G/K must commute, since if they do not, a component of the path must lie within K , and hence the path cannot be time optimal. Let $h \subset p$ denote a subspace of maximally commuting directions or generators in G/K space. Any unitary propagator U_F can then be written $U_F = k_2 \exp(Y)k_1$, where $Y \in H$. According to the time-optimal torus theorem, the fastest way to reach U_F is on the shortest path between the identity and the propagator $\exp(Y)$ such that all tangent directions commute. If we express Y as

$$Y = \sum_{i=1}^p \alpha_i Ad_{k_i}(H_i), \quad \alpha_i > 0, \tag{44}$$

it can be shown that the shortest path corresponds to the choice of α_i with the smallest value of $\sum_{i=1}^p \alpha_i$. For two-dimensional systems, G/K is of rank 1, so $Y = \alpha Ad_k(H_d)$, $\alpha > 0$ for some $k \in K$, and the time optimal path to coset $\exp(Y)$ is to flow along $Ad_k(H_d)$ for time α . Qualitatively, the optimal controls are pulse-drift-pulse sequences, i.e., hard pulses followed by evolution under drift and then some hard pulses again. For four-dimensional systems, it is necessary to pulse the controls intermittently to generate new k s, to create a chained pulse-drift-pulse sequence.

Appendix A.4 reviews the explicit construction of the geodesic trajectories and associated minimal times for the problems in dimensions 4 and 8, based on Pontryagin's maximum principle and the time-optimal torus theorem. By first numerically determining the set of unitary propagators mapping to a particular observable expectation value, and then applying these analytical results to identify the optimal field producing that propagator in minimal time, we can obtain analytical insight into the relationship among time-minimizing controls on any given level set of a quantum observable control landscape.

In higher dimensions, G/K is no longer a Riemannian symmetric space, and it is necessary to move back and forth in noncommuting directions to obtain the optimal path through G/K . In these cases, the approach of representing the invariant control system on a Lie group as an adjoint control system is still applicable, but analytic solutions have not yet been found. Nonetheless, the application of sub-Riemannian geometry to problems of quantum gate control is currently a topic of intense interest, and promises to afford additional analytical insights into the geometry of quantum control landscape level sets.

Further analytical studies on the control of unitary transformations (in two-level quantum systems) were carried out by D'Alessandro and Dahleh [30]. These authors studied the related problem of fluence minimizing controls driving a two-level quantum system to a target unitary propagator at fixed final time T . In two dimensions, this problem also has analytical solutions, for both single and multi-input control systems. The resulting optimal controls have a more complicated temporal structure; it was shown that the optimal fields are always Jacobi elliptic functions [30]. Thus, the fluence minimizing solution does not have the simple singular behaviour of the unbounded time-minimizing solutions, consistent with the phenomenon of resonance discussed in the previous section.

4. Numerical exploration of quantum control landscape level sets

4.1. Algorithms for level set exploration

We have seen that for certain classes of low-dimensional quantum optimal control problems, analytical solutions for the control fields exist. However, as the system dimension increases, analytical solutions become increasingly difficult to obtain. Moreover, we have seen that the optimal controls for these problems display particularly simple mechanistic properties, such as resonance with the transition frequencies of the system.

For the problems studied above, auxiliary costs were imposed on the controls, such as minimal time or minimal fluence. In the absence of these constraints, quantum optimal

control problems generally possess an infinite number of solutions [1]. A quantum control level set (figure 2) consists of the collection of all fields that produce a particular value for the target observable, regardless of the intervening temporal dynamics (i.e., control mechanism). A natural question concerns the relationship between these degenerate solutions, and whether their associated control mechanisms retain the simplicity of those for low-dimensional systems. From a practical standpoint, a high degeneracy of solutions will enable control fields to be tailored for specialized applications in quantum technology, through the imposition of auxiliary costs. Level set degeneracies also improve the robustness of control solutions. There will inevitably be physical inaccuracies in the experimental implementation (owing to the presence of noise, decoherence) of a particular solution, and one would like the nonideal fields to also produce dynamics that reach the objective.

Conventional algorithms for optimal control, being designed for identification of the optima of the objective function, are not well-suited to exploring the level sets of quantum control landscapes. For this purpose, Rothman *et al.* [31–33] developed a diffeomorphic homotopy procedure for systematically exploring diverse control fields on a landscape level set, referred to as diffeomorphic modulation under observable-response-preserving homotopy (D-MORPH), which we summarize here. The procedure can selectively explore the control fields on a level set that display desired properties. For example, the algorithm allows one to numerically explore the control fields producing the various unitary propagators on a level set with minimal fluence or in minimal time, for control systems that are analytically intractable.

It is convenient to parametrize the field and its variation by the exploration variable s :

$$\varepsilon(t) \Rightarrow \varepsilon(s, t) \quad (45)$$

$$d\varepsilon(t) \Rightarrow d\varepsilon(s, t) \quad (46)$$

where $0 \leq s \leq 1$. Since the goal is to explore the set of control fields that are compatible with a given observable expectation value, the solutions $\varepsilon(s, t)$ satisfy the nonlinear equation

$$F(s) = \langle \Theta(s) \rangle_T - C_T \quad (47)$$

$$= \langle \Theta([\varepsilon(s, t), H_d(s), \mu(s)], T) \rangle - C_T = 0, \quad (48)$$

as a function of s , where C_T is the desired observable expectation value.

The maintenance of in $\langle \Theta \rangle$ over an infinitesimal step ds through the level set can be written

$$\frac{d}{ds} \langle \Theta \rangle = \int_0^T \frac{\delta \langle \Theta \rangle}{\delta \varepsilon(s, t)} \frac{\partial \varepsilon(s, t)}{\partial s} dt = 0. \quad (49)$$

The neglected higher-order terms only become relevant near an extremum, where $\delta\langle\Theta\rangle/\delta\varepsilon(t) = 0$. The relationship in equation (49) is highly underspecified for determining $\varepsilon(s, t)$ as s traverses a level set. As shown in Appendix A.5, the integral equation may be expressed as an equivalent initial value problem

$$\frac{\partial\varepsilon(s, t)}{\partial s} = S(t) \left\{ f(s, t) - \frac{\gamma(s)}{\Gamma(s)} a_0(s, t) \right\}, \quad s \geq 0 \quad (50)$$

where

$$\begin{aligned} a_0(s, t, T) &= \frac{\partial\langle\Theta\rangle}{\partial\varepsilon(s, t)} \\ &= -\frac{1}{i\hbar} \langle\psi_0|[U^\dagger(T, 0)\Theta U(T, 0), U^\dagger(t, 0)\mu U(t, 0)]\psi_0 \rangle \end{aligned} \quad (51)$$

Here $S(t)$ is an arbitrary weight function (e.g., it can bias the control field towards a short pulse that approaches zero at the endpoints of the time interval),

$$\gamma(s) = \int_0^T S(t)f(s, t)a_0(s, t)dt \quad \text{and} \quad \Gamma(s) = \int_0^T S(t)a_0(s, t)a_0(s, t)dt.$$

The ability to freely choose the function $f(s, t)$ permits exploration of the multiplicity of solutions to the original integral equation. Regardless of the choice of $f(s, t)$, $\langle\Theta(s)\rangle$ will remain invariant over the $s \geq 0$ trajectory.

4.2. Quantum control mechanisms and robustness

Perhaps the most compelling reason to explore quantum control level sets is the insight they offer into control mechanisms. The ability to transform one successful control into another, and therefore one control mechanism into another, must be considered when seeking to establish the mechanism of any particular quantum control problem. Indeed, before making any definitive statements about mechanisms, it is necessary to understand the diversity of controls on a level set. How diverse can the solutions be, given the relatively simple structure of the optimal controls for integrable problems?

In order to investigate this question, Rothman *et al.* [33] applied the D-MORPH technique to an eight-level Hamiltonian with nondegenerate energy levels and with couplings only between adjacent, next-nearest and next-next nearest states. The control objective was state-state population transfer $|1\rangle \rightarrow |8\rangle$; level sets of both high and low yield were explored (figure 6). In the case of a choice of $f(s, t)$ corresponding to fluence minimization, the control field asymptotes as $s \rightarrow \infty$ towards a field of minimal fluence, although a different asymptotic field is produced for each initial field $\varepsilon(0, t)$. By contrast, for a fluence maximizing function $f(s, t)$, the distance between $\varepsilon(0, t)$ and $\varepsilon(s, t)$ increases without bound. In the latter case, the use of multiple transition pathways over the s interval suggests that the level sets are rich with fields producing vastly different dynamics. The observation that controls of minimal fluence often involve simpler mechanisms is consistent with the mechanisms apparent in the fluence and

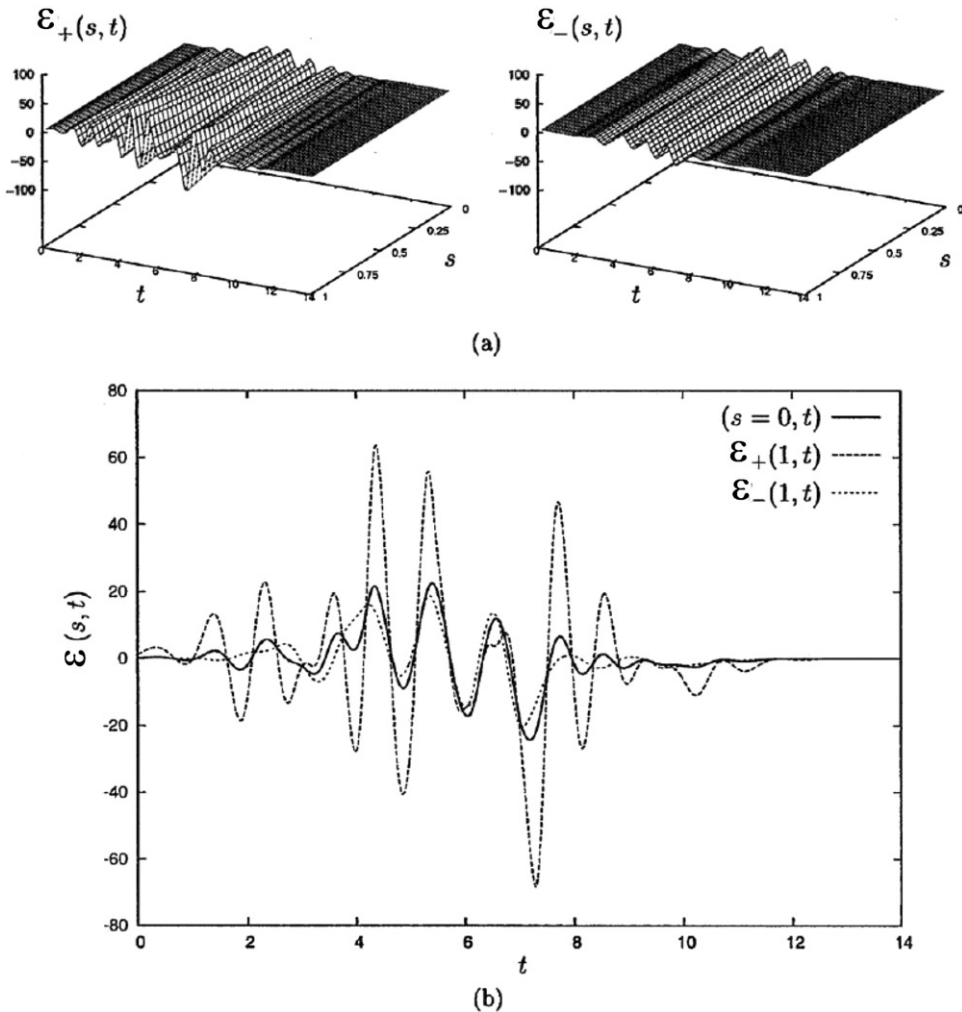


Figure 6. Starting from an initial control field $\varepsilon(s=0, t)$, the control fields $\varepsilon_{\pm}(s, t)$ are evolved on the interval $s \in [0, 1]$ subject to the functions f_{\pm} , where \pm refer to fluence maximization and minimization, respectively. (a) Under fluence maximization the control field grows in amplitude and incorporates complex structure; under fluence minimization the progression of fields decreases in amplitude. (b) Cross-sections of the fields in (a) are plotted. The field at $s=0$ is the same for both $\varepsilon_{\pm}(s, t)$, but different free functions $f(s, t)$ cause the fields to evolve in dramatically different ways with s . (From [33].)

time-minimizing analytical control solutions discussed in the previous section. In order to assess the robustness of control fields to noise along the level set, the Hessian (6) was evaluated at various points along the trajectory. In the case of fluence maximization, the trace of the Hessian was not preserved, indicating varying degrees of robustness of $\langle \Theta(T) \rangle$ to noise in the control field $\varepsilon(s, t)$.

In the weak-field regime, optimal fields for control of discrete quantum systems are typically in resonance with the transition frequencies of the system. By contrast, for continuous quantum systems, the fields are usually not simply related to natural

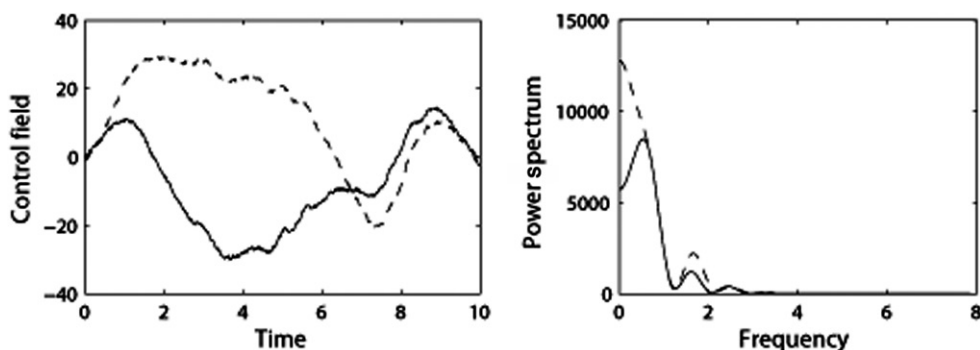


Figure 7. The optimal control fields and corresponding Fourier power spectra for continuous quantum SUM gate control in a controllable system, using two distinct control Hamiltonians. The solid and dashed lines depict the associated control fields in arbitrary units. Compare to the characteristic beat field structure in figure 6 representative of optimal controls for discrete quantum systems. (From [24].)

resonant frequencies. Wu, Chakrabarti and Rabitz compared the mechanisms of optimal control for discrete and continuous variable quantum gates [24]. Figure 7 depicts the optimal control fields for achievement of the SUM logic gate (a four-dimensional symplectic matrix) for a model two-mode continuous variable system. The resulting optimal control fields typically display complicated Fourier spectra that suggest a richer variety of possible control mechanisms. The larger diversity of mechanisms at work in continuous variable quantum control may have implications for the comparative effort of locating continuous variable versus discrete quantum control, since in the latter case the dimension of the search space cannot be reduced by requiring that the optimal controls adopt a canonical shape.

4.3. Hamiltonian-dependence of landscape level set geometry

An important quantum control goal is to discern the distinct controls that can achieve the same objective in each member of a set of similar quantum systems. A common outcome might be, for example, breaking the same type of bond in a set of molecules or creating an analogous excited state in a family of related systems. The notion of families of reactants in chemistry can be given a rigorous meaning in terms of the similarity of optimal control fields driving systems with related internal Hamiltonians to the same final state.

The diffeomorphic homotopy approach described above for level set exploration can be extended to study the relationship among controls producing the same expectation value for homologous quantum systems [31]. Diffeomorphic changes in the system Hamiltonian are introduced by scanning over a homotopy parameter s and then monitoring the control field response needed to maintain the value of a specified target observable. The time-dependent Hamiltonian is written as a function of the homotopy parameter s as

$$H(s, t) = H_d(s) - \mu(s)\varepsilon(s, t). \quad (52)$$

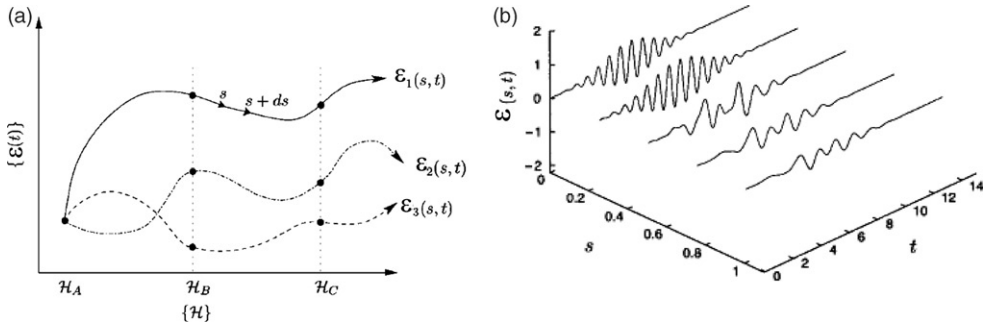


Figure 8. (a) Starting from quantum system A with Hamiltonian H_A , three paths are shown passing through systems B and C with their associated Hamiltonians H_B and H_C . Three particular control fields $\varepsilon_i(s, t)$, $i = 1, 2, 3$ characterize the pathways along which the common observable $\langle \Theta(T) \rangle$ is preserved; these distinct observable-preserving controls are specified by the auxiliary functions $f_i(s, t)$, $i = 1, 2, 3$. (b) Control fields as functions of s and t for the combined internal Hamiltonian/dipole diffeomorphism example discussed in the text (with fluence minimization). Target observable population in state $|3\rangle$ is preserved at $t = T$ for all s . (From [31].)

Figure 8 schematically displays the concept of a trajectory through Hamiltonian space.

It is possible to derive a differential equation for $\partial \varepsilon(s, t) / \partial s$ for remaining on the level set under s -dependent changes in the system Hamiltonian. The following two terms

$$a_1(s, t, T) = -\frac{1}{i\hbar} \langle \psi_0 | [U^\dagger(T, 0) \Theta U(T, 0) U^\dagger(t, 0) \frac{d\mu(s)}{ds} U(t, 0)] | \psi_0 \rangle \quad (53)$$

$$a_2(s, t, T) = -\frac{1}{i\hbar} \langle \psi_0 | [U^\dagger(T, 0) \Theta U(T, 0), U^\dagger(t, 0) \frac{dH_d(s)}{ds} U(t, 0)] | \psi_0 \rangle \quad (54)$$

analogous to a_0 above account for changes in the internal Hamiltonian and dipole operator, respectively, along the trajectory. As summarized in Appendix A.5, we obtain the explicit initial value problem:

$$\frac{\partial \varepsilon(s, t)}{\partial s} = S(t) \left[f(s, t) + \frac{(b(s, T) - \gamma(S)) a_0(s, t, T)}{\Gamma(s)} \right], \quad s \geq 0. \quad (55)$$

Rothman *et al.* carried out numerical D-MORPH simulations across a family of related three-level model systems forming a homologous set. In these studies, transfer of pure state population was considered, although the D-MORPH methodology is applicable to arbitrary observable maximization problems originating from arbitrary mixed states. The dipole moment operator and internal Hamiltonian were varied both independently and in unison.

In the case of independent dipole diffeomorphism, the goal was to move the population from state $|1\rangle$ to state $|3\rangle$. The direct transition was initially allowed, by setting $\mu_{13}(0) \neq 0$, but finally forbidden, $\mu_{13}(1) = 0$, while along the alternative

dynamical route the opposite situation exists, i.e., $\mu_{12}(0) = \mu_{23}(0) = 0$ and $\mu_{12}(1) \neq 0$ and $\mu_{23}(1) \neq 0$. Therefore, the population transfer occurs by two different mechanisms at $s = 0.0$ and $s = 1.0$. Two different trajectories were followed through Hamiltonian space, one along a straight path between the two dipole operators, and one along a curved path. Figure 8 depicts the variations in the control field required to preserve the observable across these homologous quantum systems.

It is possible to redefine the problem of level set exploration for families of related quantum systems, by identifying the set of dynamically homologous quantum systems that produce the same expectation value of a quantum observable when subjected to a fixed time-dependent control field [34]. Rather than tracking over an arbitrary path in Hamiltonian space and determining the change in the field that preserves the observable expectation value, in this case the path in Hamiltonian space is determined by the constraint that the control field does not change. An infinite number of such paths are possible, just as in the former problem. Topologically connected and disconnected families of homologous Hamiltonians have been shown to exist under various conditions. Numerical calculation of the Hessian of the associated cost functional indicates that the critical topology of this landscapes displays remarkably similar features to that for observable expectation value control; in particular, the critical points appear to be saddles rather than local traps, and the rank of the Hessian at the critical points displays the same behaviour with respect to the degeneracies in the matrices ρ , Θ .

5. Experimental exploration of quantum control landscapes

5.1. Level sets

Experimental methods are currently being developed for exploring control landscape level sets. Because the domain of control fields is infinite dimensional, the experimental investigation of quantum control landscapes requires careful choice of parametrization of the field such that the landscape can be sampled sufficiently.

Roslund and Rabitz [35] explored the level set surfaces for second harmonic generation and related nonresonant two-photon absorption. The second harmonic spectral field is given by

$$E_2(\Omega_2) \sim \int_{-\infty}^{\infty} E_1(\Omega') E_1(\Omega_2 - \Omega') d\Omega' \tag{56}$$

where $E_2(\Omega_2)$ and $E_1(\Omega)$ are the complex spectral envelopes of the second harmonic and control pulses, respectively, and the frequencies of these envelopes are relative to their spectral centre, i.e., $\Omega = \omega - \omega_0$ and $\Omega_2 = \omega - 2\omega_0$. The time integrated signal, given by

$$S \propto \int_{-\infty}^{\infty} |E_1(t)|^4 dt = \int_{-\infty}^{\infty} |E_2(\Omega_2)|^2 d\Omega_2, \tag{57}$$

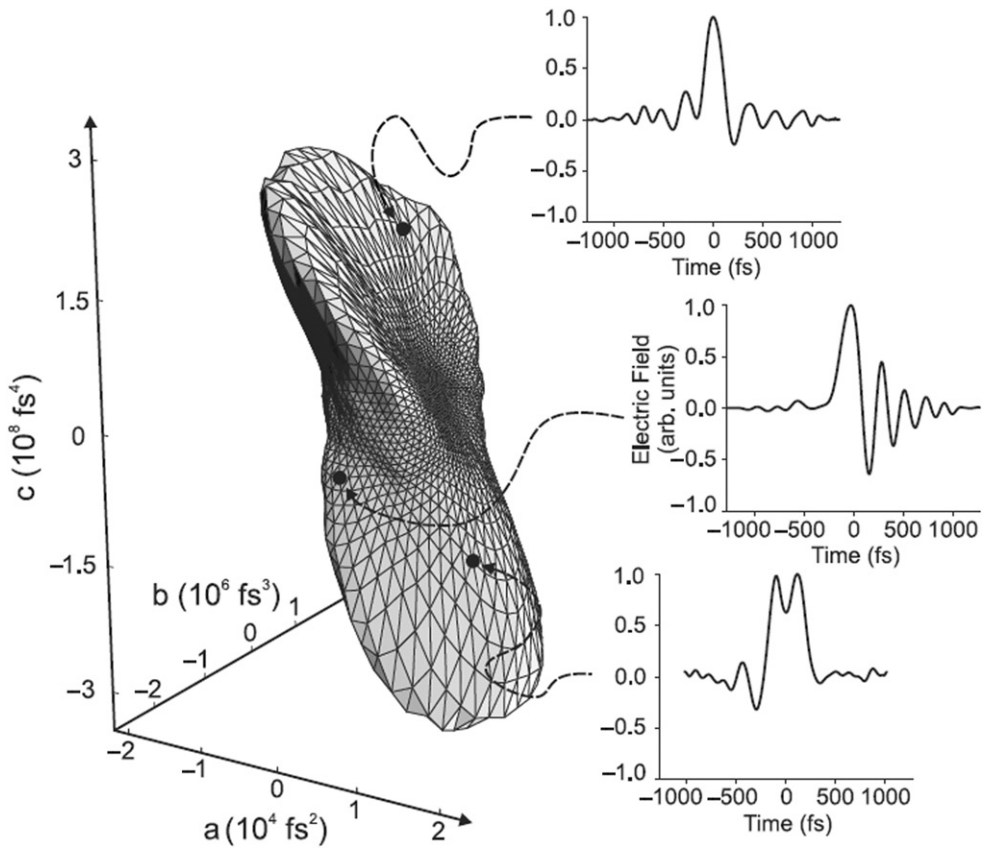


Figure 9. Experimental second harmonic generation (SHG) level set surface for a yield of 50%. Included are three experimentally retrieved control fields located on the surface. (From [35].)

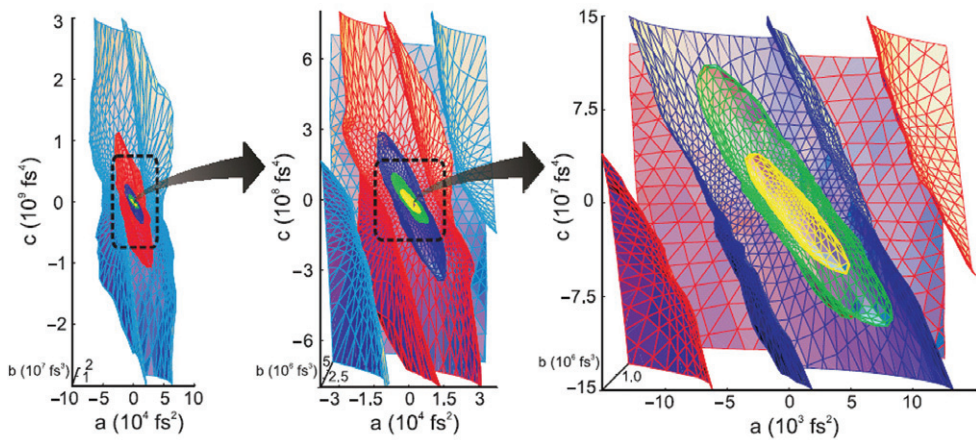


Figure 10. Experimental second harmonic generation (SHG) level set surfaces with increasing magnification moving to the right for $\alpha = 0.10$ (light blue), 0.25 (red), 0.50 (dark blue), 0.75 (green) and 0.90 (gold). Each surface is sliced along its a - c plane at $b = 0$ so that the interior is visible. (From [35].)

is measured. The spectral phase $\phi(\Omega)$ serving as the control is a truncated Taylor expansion around the centre frequency ω_0 ,

$$\phi(\Omega) = \frac{a}{2}\Omega^2 + \frac{b}{6}\Omega^3 + \frac{c}{24}\Omega^4 \quad (58)$$

where the zeroth and first-order terms are discarded because they simply correspond to an arbitrary constant phase and shift in the time origin of the pulse, respectively. The level sets were expressed on the domain of phase parameters a, b, c ; multiple level sets were identified, with one surface at 50% yield shown in figure 9. The level sets of different yield were nested in the phase parameter space (figure 10). Each of the continuously varying control fields over a given level set preserves the observable value by its own distinct manipulation of constructive and destructive quantum interferences. Thus, the richness of quantum control fields meeting a particular observable value is accompanied by an equally diverse family of control mechanisms.

For these systems, the level sets were shown to be closed surfaces in the control parameter space. In order to explore the origin of phenomenon, the Hessian of the cost functional was computed numerically at several points progressively farther away from the global optimum. Based on the positive-definiteness of the Hessian, it was shown that the level sets are predicted to be ellipsoids (see below), consistent with the experimental observation. In general, however, level sets can be unbounded in extent.

5.2. Landscape topology

In section 2, we demonstrated that formally, no traps exist in quantum control landscapes in the absence of direct costs or constraints placed on the control field.† Roslund and Rabitz [36] made the first explicit experimental demonstration of the trap-free, monotonic behaviour of unconstrained quantum control landscapes in the case of two systems, unfiltered and filtered second harmonic generation (SHG). These landscapes were randomly sampled and interpolated up to landscape level of data noise. In order to explore the topology of the landscapes, 1500 trajectories originating at random points were propagated along the gradient flow of the objective functional, according to the equation

$$\mathbf{r}(s) = \mathbf{r}(0) + \int_0^s \nabla S_t[\mathbf{r}(s')] ds' \quad (59)$$

where the gradient in the integrand was determined from the laboratory SHG landscape data. 3% (48) of these trajectories did not converge to the global optimum, and ended up distributed among two other local maxima, but these latter maxima were demonstrated to be artifacts due to noise in the control apparatus.

†Again, this statement holds rigorously in the absence of abnormal extremal controls. See section 7 for a detailed discussion.

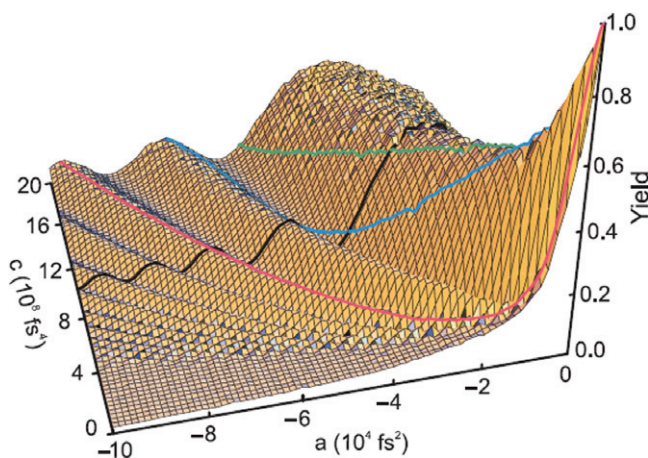


Figure 11. [Colour online] Experimental, unsmoothed quantum control landscape for filtered second harmonic generation (SHG). Four trajectories are shown; the landscape possesses a single global optimum that may be reached monotonically by the curvilinear channels that slice through the landscape. (From [36].)

Figure 11 depicts several possible search trajectories along the filtered SHG landscape. (Filtering the SHG signal to be evaluated at $\Omega_2 = 2\omega_0$ removes the dependence of the signal on the cubic Taylor coefficient in the polynomial basis.) Several of these paths correspond to simple parameterizations of the control field. As can be seen, restricted parameterization of the control field will generally produce artificial structure by forcing projections of the original full infinite dimensional control space.

Although linear trajectories in the polynomial phase representation considered above are incapable of following the gradient flow trajectory, they do display the favourable property of preserving the intrinsic topology of the landscape. This can be verified by perturbation analysis near the global maximum of the filtered SHG landscape. Expanding the exponential phase to second order around the optimal solution $\phi(\omega) = 0$, it may be shown [36] that the perturbative signal $\delta S_f(\Omega) = S_f^* - S_f(\Omega)$, where S_f^* is the transform limited signal, indicates that the level sets are ellipsoidal. In addition, assuming a spectral amplitude of the form $A(\Omega) = \exp(-\Omega^2/2\Delta^2)$, the normalized signal variation under a Taylor expansion of the phase functional preserves this landscape topology. In general, however, the appropriate choice of control parameters or variables that preserve landscape topology may not be apparent *a priori*. Methodologies exist [37] for transforming to an optimal local basis set of laboratory control parameters. A physically convenient parameterization is first chosen prior to the outset of the experiment, followed by a Hessian analysis to determine a locally separable representation. This methodology has been successfully illustrated using the example of molecular fragmentation of CH_2BrI .

A possible cause for the appearance of local traps in quantum control landscapes is the existence of costs or constraints on the controls, which in some cases may be impossible to avoid. Numerical results [18] indicate that landscape topology should be reasonably preserved even in the presence of small penalties on the field fluence. These

results suggest that in weak-field experiments, where the control field constraints may not be particularly limiting, the topology of experimental quantum control landscapes should remain monotonic.

In the strong field regime, however, local traps may appear more readily. For example, Wells *et al.* [38] applied adaptive search algorithms to the fragmentation of a complex molecule, octahedral sulfur (S_8), investigating several different control field parametrizations. The structure of the optimal pulses obtained using these various parameterizations were considerably different, although they produced comparable signal enhancements. The fluence of the control fields were constrained considerably, corresponding to large values of λ in equation (1). Although the sampling was not exhaustive enough to rigorously establish level set structure or landscape topology, the results suggested that local maxima existed in the S_8 fragmentation landscape for the parameterizations employed. However, even in the strong field regime, the fundamental monotonicity of the landscape need not be compromised. In related work [39], the control landscape for the strong-field ionization of potassium atoms was sampled by phase modulated pulses. The intensity of the Autler–Townes components in the photoelectron spectra were controlled by a sinusoidal phase modulation function. The use of a two-dimensional parameter space enabled constrained, but effective sampling of the control landscape. The maxima and minima of the landscape were identified, with clear level sets, and no evidence for local traps was found.

In most cases of practical interest in the weak-field regime, unconstrained quantum control landscapes possess no suboptimal traps. Thus, local experimental search algorithms should be effective in locating optimal controls; moreover, since gradient-based algorithms can take advantage of landscape structure, they may perform better than ‘blind’ algorithms (e.g., genetic algorithms) under suitable conditions. In the aforementioned work of Roslund and Rabitz, only the total SHG yield was measured; the gradient of the objective was determined based on radial basis function interpolation. More recent work [40] has addressed the question of how to experimentally measure and follow the gradient flow of the observable maximization objective function, given that the basis must change at each step along the curvilinear path. The gradient of the filtered SHG objective was measured using a moment-based method with only 30 observable measurements on a 128-dimensional parameter space. Despite the statistical uncertainty in the measurements and the presence of noise, following this flow resulted in convergence to >90% achievement in half the number of steps that were required for a GA. Note that accurate estimation of the gradients of observable expectation values with respect to control field parameters is properly a subject of quantum statistical inference [41]. Quantitative assessment of the statistical uncertainties associated with these estimates, as a function of the number of measurements made, is essential for determining the speedup that can be achieved by using gradient versus adaptive search algorithms.

Following the gradient flow of the objective function exploits the favourable topology of quantum control landscapes, but does not explicitly exploit their geometry. In particular, the best path to the global maximum in figure 11 is not necessarily the gradient flow path. A further possibility is to use local algorithms that make use of local gradient information, but do not follow the gradient flow directly, instead tracking

alternative observable expectation value paths. In this regard, an important question is whether certain observable paths are expected to display more rapid convergence to the global optimum than others. Section 8 discusses how geometric features of quantum control landscapes indicate that certain observable paths may in fact be globally more efficient than the gradient flow. Because they do not follow the path of steepest ascent, these algorithms may require more accurate estimation of the gradient. Such methodologies could also be applied to explicit experimental tracking of predetermined level set trajectories [42].

In summary, improved methods for the measurement of the gradient in the presence of noise, and the accessibility of control field parameterizations capable of tracking the gradient are the central challenges for implementing experimental control algorithms that exploit the favourable topological and geometric structure of quantum control landscapes. Note that the analytical results of section 3 pertaining to dominant quantum control mechanisms can also be used to help choose appropriate control field parameterizations that simplify experimental landscape search.

6. Quantum system controllability and landscape structure

As discussed in section 2, the full controllability of a quantum system is a necessary condition for the nonexistence of local traps in the control landscape. If the target transformation is reachable at the final time T , but intermediate dynamical propagators are not, the path to the target may be plagued with local traps. It is therefore important to review the conditions for controllability of quantum systems at a fixed dynamical time T . Moreover, since the choice of T is to some extent arbitrary, it is important to identify the possible choices of T that lead to full controllability.

The conditions for the controllability of finite-dimensional quantum systems were established [43] based upon earlier work on the controllability of systems on compact Lie groups [44]. In [43], an easily implementable algorithm for ascertaining the controllability of a finite-level quantum system was provided. Consider the right-invariant system described by Definition 7 (Appendix A.6); let S be the subgroup of the dynamical group G generated by the internal and control Hamiltonians (with corresponding Lie algebra s), and let L be the subgroup generated by the control Hamiltonians alone (with corresponding Lie algebra l). The following theorem establishes sufficient conditions for the full controllability of such systems.

Theorem 1 (Controllability of right-invariant systems on Lie groups): *The reachable set from the identity matrix in G is contained in S . If S is compact then the reachable set from the identity matrix equals S . In particular, if the dimension of the Lie algebra l equals the dimension of the ambient Lie group G , and the Lie group is compact, then the control system is controllable. Furthermore, in this case it is possible to reach any matrix with an admissible control which is bounded in amplitude.*

Controllability can be checked using the rank condition, which states that if the dimension is N^2 for the Lie algebra spanned by H_d , μ_i and their commutators such as $[H_d, \mu_i]$, $[H_d, [H_d, \mu_i]]$, $[\mu_i, [H_d, \mu_i]]$, etc., then the system is controllable. This is

equivalent to requiring that l be the Lie algebra of all $N \times N$ skew-Hermitian matrices, which in turn is equivalent to requiring that the dimension of l as a vector space over the real numbers is precisely N^2 . This criterion for controllability of quantum dynamical propagators extends to that of states; all coherent superpositions of states can be achieved if S equals $U(N)$. The fact that the controllability of a linear control system can be checked via a simple rank criterion which, in addition, does not vary from point to point, is an important property which is generally not valid for a nonlinear control system. Usually this condition guarantees only accessibility [43].

6.1. Exact-time controllability of discrete quantum systems

The above theorem establishes the necessary conditions for the existence of a time T at which the system is controllable, but does not constructively define T . Landscape search would be simplest if the quantum control system satisfied the conditions for strong controllability.

Definition 2: A control system $F=(A, B_i, u_i(t))$ is strongly controllable over a subgroup M if for any $T > 0$ any point of M is reachable from any other point by F in T or fewer units of time. A control system is said to be strongly controllable if the property of strong controllability holds for the entire dynamical group G . A control system is exact time controllable at time T if any point of M is reachable from any other point by F in exactly T units of time.

Just as there exist analytical solutions to the Pontryagin maximum principle for finite-dimensional quantum systems (i.e., right-invariant systems on a compact Lie group), there are also powerful general theorems for establishing exact-time and strong controllability of these systems. Also by analogy, there are differences in the conditions establishing exact-time and strong controllability of dynamical transformations versus that of quantum states.

For finite-dimensional quantum systems, strong controllability can be guaranteed if two controls are used, and these controls span the whole Lie algebra of the dynamical group. In the more common case of one control, strong controllability cannot be guaranteed, but exact-time controllability across a wide range of times T can be straightforwardly established.

D'Alessandro and Dahleh [45] have studied the exact-time controllability of two-qubit gates. They showed that if a two qubit gate is controllable at time T_1 (called the critical time), then it is also controllable at any time $T_2 > T_1$. We summarize their proof here because of its possible extensions to higher-dimensional quantum systems. Let $R(T)$ be the reachable set from I , i.e., the set of possible values for $X(T)$ obtained by varying the controls u_1, \dots, u_m within the set of continuous functions defined on $[0, T]$. We also define the sets

$$\mathbb{R}(\leq T) = \bigcup_{0 \leq t \leq T} R(t) \quad (60)$$

$$\mathbb{R} = \bigcup_{0 \leq t \leq \infty} R(t). \quad (61)$$

Because of right-invariance, $R(I, T)S = R(S, T)$ for every $S \in SU(2)$ and every T . Therefore, it suffices to consider the (exact-time) controllability properties of the set reachable from the identity. For quantum systems of arbitrary dimension, the following theorem establishes the existence of a critical time beyond which the reachable set $\mathbb{R}(T)$ is equal to the entire unitary group $U(N)$ of dynamical propagators.

Theorem 2 [44]: *Let S denote the subalgebra generated by the controls A, B_1, \dots, B_m . If $(A, B_i, u_i(t))$ is a right-invariant control system on a Lie group G and S is compact, (i) $\mathbb{R} = S$; (ii) there exists a $T > 0$ such that $\mathbb{R}(\leq T) = \mathbb{R}$.*

Powerful exact-time controllability results may be proven in dimension 2 because a Lie algebra isomorphism γ exists between $su(2)$ and $so(3)$. It follows from Lie's third theorem [45], that ρ induces a homomorphism γ' mapping $SU(2)$ onto $SO(3)$. It can be shown that if a system is controllable on $SO(3)$, any element of the subgroup $SO(2)$ can be reached in arbitrarily small time (i.e., the system is small-time controllable on $SO(2)$). It is then straightforward to demonstrate that on either $SO(3)$ or $SU(2)$, if $T_1 \leq T_2$, then $R(T_1) \subseteq R(T_2)$ and therefore $\mathbb{R}(\leq T) = R(T)$ for each $T \geq 0$ [45]. In other words, if $T_1 < T_2$, all the points reachable at T_1 are reachable at T_2 . Combining this result with Theorem 2, we obtain the the following exact-time controllability result in dimension 2:

Theorem 3: *There exists a time T_c , such that $R(T) = SU(2)$ for every $T > T_c$. The critical time T_c is the least time such that for every $T > T_c$, it is possible to drive the system from the identity to an arbitrary matrix in $SU(2)$.*

An important question is whether this property can be extended to finite quantum systems of arbitrary dimension N , i.e., whether the above equivalence between the sets R and \mathbb{R} holds in general. If this is the case, quantum control landscapes for larger systems will display a homogeneous structure for all T above a critical time, such that simulations and experiments need not sample extensively over the time T in order to obtain a landscape with simple topology.

6.2. Controllability of classical and continuous variable quantum systems

For the noncompact Lie groups describing the evolution of classical or continuous variable quantum systems, strong controllability is not established by the above Lie algebra rank condition. This implies that the critical topology of control landscapes for such systems may change considerably as the final dynamical time T is varied, underscoring the comparative simplicity of discrete quantum control landscapes.

For continuous variable (infinite-dimensional) quantum systems, the rank condition is also sufficient for establishing controllability on noncompact symplectic groups (i.e., those with quadratic Hamiltonians) in the common case where H_0 is compact, but there is no guarantee of exact-time controllability, i.e., some particular gates may only be reachable after an extremely long time. However, exact-time controllability can be achieved at arbitrary positive times (i.e., $R(T) = \mathbb{R}(\leq T) = \mathbb{R} = U(N)$ for every $T > 0$) if we can employ two control Hamiltonians that span the whole Lie algebra of the group of dynamical propagators. Thus, the topology of control landscapes for the subset of infinite dimensional quantum gates described in section 2 will be largely insensitive to the final time T if two independent controls are used.

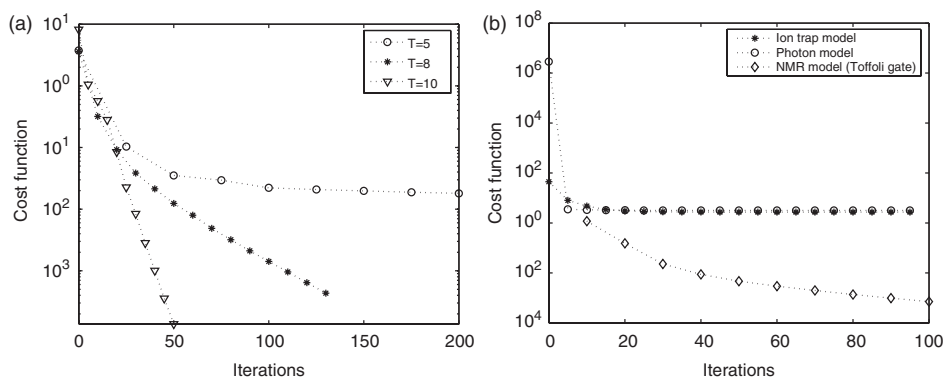


Figure 12. Effects of exact-time controllability on the optimal control fidelity of discrete vs. continuous quantum systems. (a) The convergence of control field searches for the two-qubit continuous quantum SUM gate, with a weakly controllable continuous variable system, using conjugate gradient algorithms and different final times. (b) The convergence of optimal searches for a three-qubit SUM gate with a weakly controllable system (photon model), three-qubit SUM gate with an uncontrollable system (ion-trap model) and three-qubit discrete quantum Controlled-CNOT gate with a standard NMR spin coupling model. See text for definitions of quantum gate terminology. (From [24].)

The strong controllability of general infinite-dimensional quantum systems (i.e., those with nonquadratic Hamiltonians) was studied by Wu and Tarn [46]. Such systems were shown to be associated with dynamical symmetries represented by noncompact Lie groups with infinite-dimensional unitary representations. A criterion for approximate strong controllability, called smooth controllability, was given, showing that such systems, which possess an uncountable number of levels, can be well manipulated using a finite number of control fields.

The effects of exact-time controllability of discrete versus continuous variable quantum systems on the simplicity of control field search were studied by Wu, Chakrabarti and Rabitz [24]. Figure 12 compares the convergence of optimal searches for achieving the CV SUM logic gate (symplectic propagator) versus the computationally equivalent discrete CNOT gate (unitary propagator), using identical gradient-based algorithms. The CV gates were implemented using several models with varying degrees of controllability, including a weakly controllable system (employing ion trap interactions) and an uncontrollable system (employing photon-atomic spin interactions). As can be seen, the maximal achievable fidelity is highly sensitive to the choice of final time for uncontrollable or weakly controllable systems. Moreover, even when the gate was reachable, the search effort required for convergence was found to increase with decreasing controllability. By contrast, a randomly chosen final time was sufficient for achieving near perfect fidelity in the discrete quantum system. The weaker controllability of continuous systems can also result in control fields with more complicated Fourier power spectra, as shown in figure 13.

7. Computational complexity of quantum control landscapes

The scaling of the expense of quantum simulation with system dimension is of fundamental importance in quantum chemistry. Similarly, the scaling of the expense for

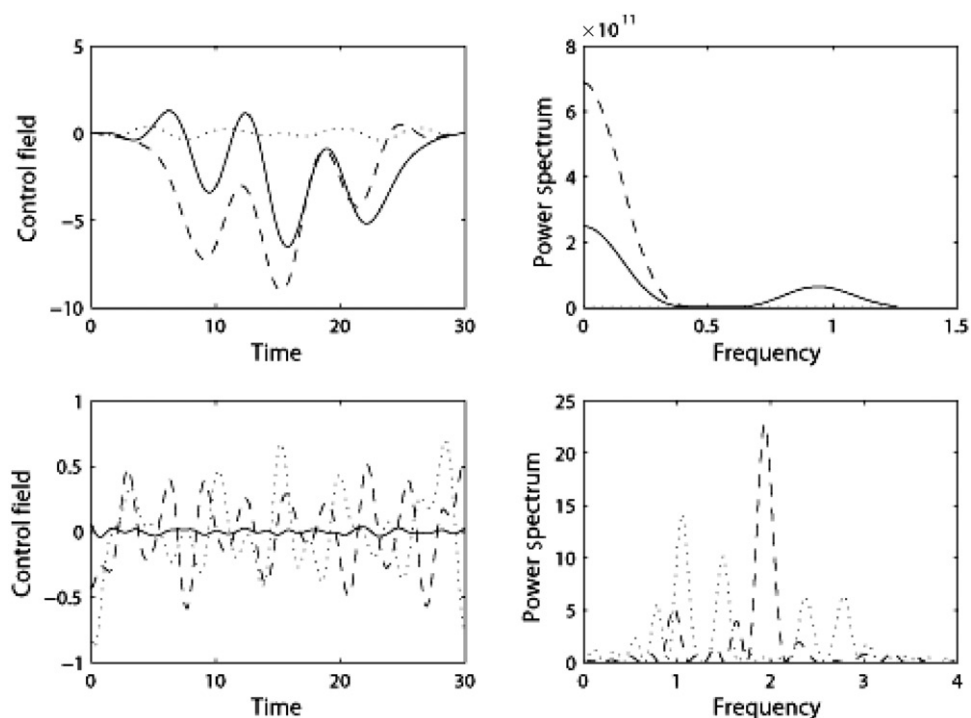


Figure 13. Comparison of the optimal control fields for discrete and continuous variable quantum gates depicted in figure 12(b). Three independent controls were employed, represented by solid, dashed, and dotted lines, respectively. (a) Optimal fields for three-qubit continuous quantum SUM gate control in a weakly controllable system; (b) Optimal fields for three-qubit discrete quantum controlled-CNOT gate control in a standard NMR spin coupling model. Control field amplitudes are arbitrary. (From [24].)

quantum control search – either OCT or OCE – lies at the heart of the applicability of quantum control to the large molecules of practical interest in many applications. Whether the search is carried out numerically or experimentally, this scaling is referred to as the problem’s computational complexity. Compared to other central optimization problems in quantum technology and quantum information, such as state or process reconstruction [47], the complexity of quantum control problems is more difficult to assess since the optimization is carried out over an infinite-dimensional parameter space.

We have seen that analytical solutions to quantum optimal control problems appear to exist only for low-dimensional systems, and in some cases only for restricted Hamiltonians. Nonetheless, the existence of analytical solutions is rare in classical OCT, and begs the question of how the difficulty of finding solutions to nonintegrable quantum control problems increases with system dimension. Of course, this complexity is algorithm-dependent. It is also a function of the system Hamiltonian, but may display homogeneous features across families of Hamiltonians. Recall that in section 3, we showed that for certain classes Hamiltonians, analytical results pertaining to control mechanisms permits reduction in the dimension of the search space.

The search effort required in OCT calculations for observable maximization is remarkably insensitive to the dimension of the quantum system. A multitude of such calculations have been reported in the literature [48–50], on systems ranging in dimension from 2 up to more than 10^2 . Even for the largest systems studied, the number of iterations required for convergence seldom exceeds 10^3 , with most calculations requiring between 10^2 and 10^3 iterations [51, 52]. Much larger systems have been subjected to OCE studies; in the common case of selective bond dissociation, a continuum of states is accessed. Nonetheless, the search effort reported in most OCE studies is of the same order of magnitude as in OCT.

Since the majority of quantum OCT algorithms are based on the gradient of the objective function, and the gradient is relatively straightforward to implement in OCE, it is natural to ask about the complexity of optimal control when the search is carried out using these algorithms.

7.1. Gradient flows and search complexity

In section 2, we showed how the critical topology of the most common objective functions in quantum optimal control can be determined analytically and display features favourable optimal search. In section 3, we established several important analytical results pertaining to the geometry of control landscapes, in particular their level sets, finding that in certain cases, control mechanisms can be exploited to reduce the dimensionality of the control search space. Here we examine another feature of the geometry of control landscapes that permits analytical investigation, namely the kinematic gradient flows of the objective function Φ . On the domain $\mathcal{U}(N)$, these gradient flows themselves represent integrable dynamical systems. As a result of this feature, it is possible to identify the system-independent contribution to the scaling with system dimension of the search effort for locating quantum optimal controls.

The gradient flow is the trajectory followed by the search algorithm when the algorithmic step is defined according to the differential equation

$$\frac{dU}{ds} = -\nabla_U J(U) \quad (62)$$

The unitary gradient flow equations for observable maximization and gate optimization are then, respectively,

$$\left(\frac{dU}{ds}\right)_1 = -U[\rho(0), U^\dagger \Theta U], \quad (63)$$

$$\left(\frac{dU}{ds}\right)_2 = W - UW^\dagger U. \quad (64)$$

In the case that $\rho(0)$ has only one nonzero eigenvalue, corresponding to an initial pure state, it was shown that under the change of variables $\rho(T, s) = |\psi(s)\rangle\langle\psi(s)|$,

$|\psi(s)\rangle = (c_1(s), \dots, c_N(s))$, $x(s) \equiv (|c_1(s)|^2, \dots, |c_N(s)|^2)$, the gradient flow of Φ_1 can be explicitly integrated to give [42]:

$$x(s) = \frac{1}{\sum_{i=1}^N |c_i(0)|^2 e^{2s\lambda_i}} e^{2s\Theta} \cdot (|c_1(0)|^2, \dots, |c_N(0)|^2) \tag{65}$$

$$= \frac{1}{\sum_{i=1}^N |c_i(0)|^2 e^{2s\lambda_i}} (e^{2s\lambda_1} |c_1(0)|^2, \dots, e^{2s\lambda_N} |c_N(0)|^2) \tag{66}$$

where $\lambda_1, \dots, \lambda_N$ denote the eigenvalues of Θ . The explicit solution for the gradient trajectory of objective functional Φ_2 was shown to be

$$W^\dagger U(s) = (\sinh(s) + \cosh(s)W^\dagger U_0)(\cosh(s) + \sinh(s)W^\dagger U_0)^{-1} \tag{67}$$

where the initial condition is $U_0 = U(0)$ [42].

Chakrabarti *et al.* [53] have calculated upper bounds of the convergence times of these unitary gradient flows into a ball of radius ϵ around the solution. For the class of observable maximization problems above, this bound was found to be

$$t_{c,1}(H) = \max \leq \frac{1}{2\mu} \left[\ln \left(\frac{2Nk}{\epsilon^2} \right) + 2 \ln \frac{(N-k-2)\lambda_{k+1}}{k(\lambda_{(1)} - \lambda_{k+1})} \right]. \tag{68}$$

Where N is the Hilbert space dimension, k is the degeneracy of the largest eigenvalue of the observable operator θ and μ is the absolute value of the difference of the two largest eigenvalues of Θ . The upper bound on the

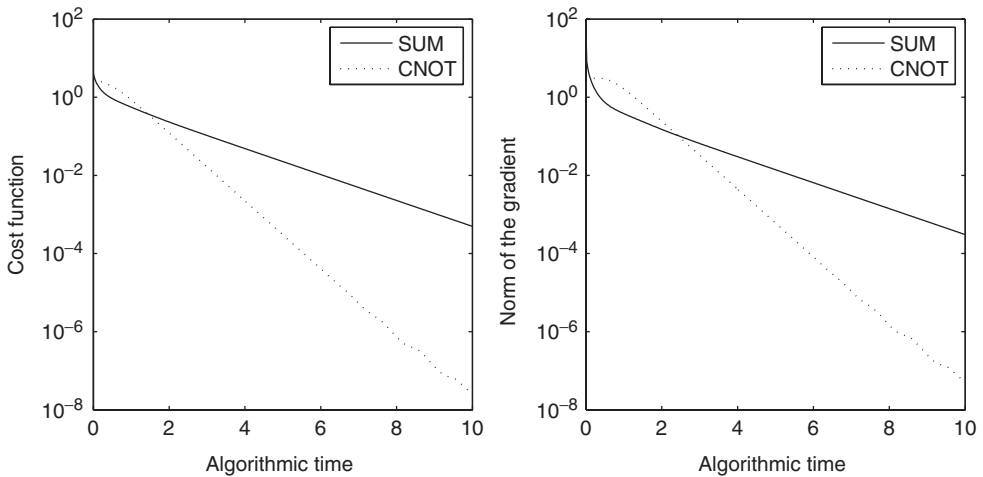


Figure 14. The convergence of the kinematic gradient flows for optimal search of the SUM gate on $Sp(4, \mathbb{R})$ and CNOT gate on $\mathcal{U}(4)$. (From [24].)

convergence time for the unitary gradient flow of the gate fidelity function was shown to be

$$t_{c,2}(H) = t_{c,2}(\epsilon) \leq \frac{1}{2} \ln \left(\frac{4N}{a^2 \epsilon} \right) \tag{69}$$

where $a = \sin \theta_0 / (1 - \cos \theta_0)$, for small ϵ . Since both of these critical times scale logarithmically with the Hilbert space dimension, the problem of optimizing the objective functions on the domain of unitary propagators (via gradient algorithms) belongs to a logarithmic analogue complexity class, referred to as CLOG within the analogue complexity literature [54].

By contrast, the kinematic gradient flows for optimization of the corresponding objective functions for classical or continuous variable systems are not integrable. The kinematic flows for optimization of the continuous variable SUM gate and the discrete variable CNOT gate are shown in figure 14 [24], revealing that the kinematic contribution to control optimization is less favourable for continuous variable systems.

The integrated flow derived above for the problem of observable maximization applies only to the case where $\rho(0)$ is a pure state. For the more general problem of a mixed initial state, analytic solutions are more difficult to obtain [52]. However, insight into the scaling of the search effort for these problems may be obtained from numerical simulations, as well as consideration of the dimension of the manifold on which the gradient flow evolves.

As such, numerical optimizations [55] of the observable expectation value function on the domain of unitary propagators were carried out for $\rho(0)$ and Θ operators of various possible ranks and degeneracies, for Hilbert space dimensions ranging from 2 to 40. From a kinematic perspective, the optimization problem is symmetric with respect to these two operators. It was found that the scaling was completely determined by the number of nondegenerate eigenvalues of $\rho(0)$; in particular, if $\rho(0)$ is a pure state, the spectrum of Θ does not alter the complexity class (or vice versa). In the latter case, the scaling of effort was observed to be roughly logarithmic in the Hilbert space dimension, consistent with the analytical result above. In the limiting case where $\rho(0)$ is a full rank matrix with nondegenerate eigenvalues, the effort was found to scale linearly with system dimension.

The origin of the observed numerical scaling has been probed [22] by examining the dimension of the subspace on which the gradient evolves. It is found that this dimension is identical on the domain of unitary propagators and control fields, although in the latter case the basis functions change continuously along the optimization trajectory. The gradient can be expanded in terms of at most $N(N - 1)$ linearly independent functions of the time-dependent dipole operator $\mu(t)$. Let ρ consist of r subsets of degenerate eigenvalues p_1, \dots, p_r with multiplicities n_1, \dots, n_r , and write $\rho(0) = \sum_{i=1}^r p_i |i\rangle\langle i|$. The expression (6) for the gradient from section 2 can be expanded to give

$$\begin{aligned} \frac{\delta \Phi_1}{\delta \epsilon(t)} &= \frac{i}{\hbar} \sum_{k=1}^r p_k \sum_{i=s_k+1}^{s_{k+1}} \left[\sum_{j=1}^{s_k} + \sum_{j=s_{k+1}+1}^N \right] \{ \langle i | \Theta(T) | j \rangle \langle j | \mu(t) | i \rangle - \langle i | \mu(t) | j \rangle \langle j | \Theta(T) | i \rangle \} \\ &+ \frac{i}{\hbar} \sum_{k=1}^r p_k \sum_{i=s_k+1}^{s_{k+1}} \sum_{j=s_k+1}^{s_{k+1}} \{ \langle i | \Theta(T) | j \rangle \langle j | \mu(t) | i \rangle - \langle i | \mu(t) | j \rangle \langle j | \Theta(T) | i \rangle \} \end{aligned} \tag{70}$$

where n_i are the degeneracies of the eigenvalues p_i of $\rho(0)$, $s_1=0$, $s_k = \sum_{i=1}^{k-1} n_i$, $k=2, \dots, r+1$, $s_{r+1}=n$. The terms in the second summation, of which there are $\sum_i n_i^2$, add to zero, from which it can be shown that the dimension of the subspace of skew-Hermitian matrices upon which the gradient flow evolves is [22]

$$D = N^2 - (N - n)^2 - \sum_{i=1}^r n_i^2 = n(2N - n) - \sum_{i=1}^r n_i^2. \quad (71)$$

Therefore, when $\rho(0)$ is full rank and nondegenerate, the gradient can be expressed in terms of a linear combination of $N(N-1)$ basis functions, irrespective of the spectrum of Θ . Increasing degeneracy in the spectrum of $\rho(0)$ reduces the dimension of the subspace on which the gradient evolves, such that when $\rho(0)$ has n degenerate eigenvalues, this dimension is equal to $2n(N-n)$. Although this result does not establish the kinematic contribution to the scaling of search effort with Hilbert space dimension, it reveals that the dimension of the subspace on which the gradient evolves for generic observable maximization problems scales less favourably with N when ρ and Θ have more nondegenerate eigenvalues.

The existence of a low-dimensional basis set of functions upon which the gradient can be expanded is especially useful given the difficulty of implementing effective high-dimensional control field parameterizations in the experimental setting (section 5). Although this basis set varies from point to point along the landscape, recent work [56] suggests that in many cases, it can be remarkably homogeneous, providing a rational means of estimating the minimal control parameterization dimensionality needed to effectively climb the landscape.

7.2. Relation between dynamic and kinematic gradient flows

The integrated U -gradient flows of the observable maximization and gate fidelity cost functions identify Hamiltonian-independent contributions to the scaling of quantum control search effort when using gradient algorithms. In this section, we derive the Hamiltonian-dependent relationship between these U -gradient flows and the ε -gradient flows that are followed by OCT and OCE algorithms.

The ε -gradient flows are the solutions to the differential equations

$$\frac{d\varepsilon(s, t)}{ds} = \nabla\Phi(\varepsilon(t)) = \alpha \frac{\delta\Phi(s, T)}{\delta\varepsilon(s, t)} \quad (72)$$

where s is a continuous variable parametrizing the algorithmic time evolution of the search trajectory, and α is an arbitrary scalar that we will set to 1. The gradient on $\varepsilon(t)$ is related to the gradient on $\mathcal{U}(N)$ through

$$\frac{\delta\Phi}{\delta\varepsilon(t)} = \sum_{i,j} \frac{\delta U_{ij}}{\delta\varepsilon(t)} \frac{d\Phi}{dU_{ij}}. \quad (73)$$

Now suppose that we have the gradient flow of $\varepsilon(s, t)$ that follows (22) and let $U(s)$, the system propagator at time T driven by $\varepsilon(s, t)$, be the projected trajectory on the

unitary group $\mathcal{U}(N)$. The (algorithmic) time derivative of $U(s)$ is then

$$\frac{dU_{ij}(s)}{ds} = \int_0^T \frac{\delta U_{ij}(s)}{\delta \varepsilon(s, t)} \frac{\partial \varepsilon(s, t)}{\partial s} dt \quad (74)$$

which, combined with (22) and (23), gives

$$\frac{dU_{ij}(s)}{ds} = \int_0^T \frac{\delta U_{ij}(s)}{\delta \varepsilon(s, t)} \sum_{p, q} \frac{\delta U_{pq}(s)}{\delta \varepsilon(s, t)} \frac{d\Phi}{dU_{pq}} dt. \quad (75)$$

It is convenient to write this equation in vector form, replacing the $N \times N$ matrix $U(s)$ with the N^2 dimensional vector $\mathbf{u}(s)$:

$$\frac{d\mathbf{u}(s)}{ds} \left[\int_0^T \frac{\delta \mathbf{u}(s)}{\delta \varepsilon(s, t)} \frac{\delta \mathbf{u}^T(s)}{\delta \varepsilon(s, t)} dt \right] \nabla \Phi[\mathbf{u}(s)] \equiv G[\varepsilon(s, t)] \nabla \Phi[\mathbf{u}(s)] \quad (76)$$

where the superscript T denotes the transpose. This relation implies that the variation of the propagator in $\mathcal{U}(N)$ caused by the natural gradient flow in the space of control field is Hamiltonian-dependent, where the influence of the Hamiltonian is contained in the N^2 -dimensional symmetric matrix $G[\varepsilon(s, t)]$.

Thus, although the convergence times for the U -gradient flows above scale favourably with system size, the ε -gradient flows do not generally follow the same paths, and Hamiltonian-dependent effects may dominate the scaling when following the gradient on the domain of control fields. OCT calculations suggest a difference in the scaling of observable maximization and gate control search effort when using local gradient-based algorithms [6], although the U -gradient flow scalings are similar for these two problems. Systematic dynamical OCT studies have been carried out for unitary gate optimization on systems of dimension ranging from 2 to 32, using iterative algorithms. The computational effort was found to scale exponentially in the Hilbert space dimension.

A natural question is whether the Hamiltonian-dependent unfavourable scaling of local OCT or OCE algorithms can be mitigated by employing global algorithms whose optimization trajectories are less sensitive to the system Hamiltonian. In the next section, we discuss such global search algorithms and the properties of quantum control landscapes that render these algorithms effective.

8. Global search algorithms for quantum control

8.1. Scalar and matrix tracking algorithms

As shown above, the convergence time of gradient OCT or OCE algorithms is dependent on the system Hamiltonian. By contrast, it is possible to employ algorithms

that follow a predetermined track of observable expectation values, independent of the system Hamiltonian. An efficient algorithm for following a predetermined track

$$P(s) = \langle \psi(s, T) | \Theta | \psi(s, T) \rangle, \quad 0 \leq s \leq 1 \quad (77)$$

for the expectation values of a target observable operator Θ at final time T may be derived from the diffeomorphic homotopy formalism described in section 4, originally developed for level set exploration. By making the substitution

$$\int_0^T a_0(s, t, T) \frac{\partial \varepsilon(s, t)}{\partial s} dt = b(s, T) + \frac{d\langle \Theta(s, T) \rangle}{ds},$$

in equation (49), we obtain the following expression for the algorithmic step for the control field (Appendix A.5):

$$\frac{\partial \varepsilon(s, t)}{\partial s} = S(t) \left[f(s, t) + \frac{(b(s, T) + \frac{dP}{ds} - \gamma(s)) a_0(s, t, T)}{\Gamma(s)} \right]. \quad (78)$$

If the Hamiltonian is kept fixed during tracking, $b(s, T) = 0$. As mentioned in section 4, such algorithms may also be implemented experimentally, but require a precise measurement of the gradient. Analogous equations can be derived for following a predetermined track of gate fidelity function values.

At each step in this approach, the observable expectation value is specified, but the unitary propagator is not; many unitary propagators will map to the same expectation value. The tracking errors that occur in this approach will be system-specific, depending on the system dimension and Hamiltonian. The method was applied to a five-level quantum system initially in state $|1\rangle$, with the goal of transferring population to state $|5\rangle$. Figure 15 depicts the changes in the control field along the track for $P(s) = |\langle 5 | \psi(s=0, T) \rangle|^2 + \sin(2\pi s)$. In this case, the dipole operator was also morphed along the track, in order to demonstrate the feasibility of simultaneous Hamiltonian variation.

It is natural to consider the prospects of tracking paths in the space of dynamical propagators, rather than observable expectation values. Doing so may permit a more system-independent definition of quantum control complexity. Although this may be difficult to achieve in OCE using current technology, it can be easily implemented in simulations. Indeed, diffeomorphic homotopy provides a natural means of tracking paths $U(s, T)$, $0 \leq s \leq 1$, in the group of quantum dynamical propagators. Consider the problem of tracking the unitary gradient flow for observable expectation value maximization, integrated in section 7.

In order for the projected flow from $\varepsilon(t)$ onto $U(T)$ to match the integrated gradient flow on $U(T)$, the quantity $\partial \varepsilon(s, t) / \partial s$ that corresponds to movement in each step must satisfy a generalized differential equation:

$$\frac{dU(s)}{ds} = \int_0^T \frac{\delta U(s)}{\delta \varepsilon(s, t)} \frac{\partial \varepsilon(s, t)}{\partial s} dt = \nabla \Phi[U(s)]. \quad (79)$$

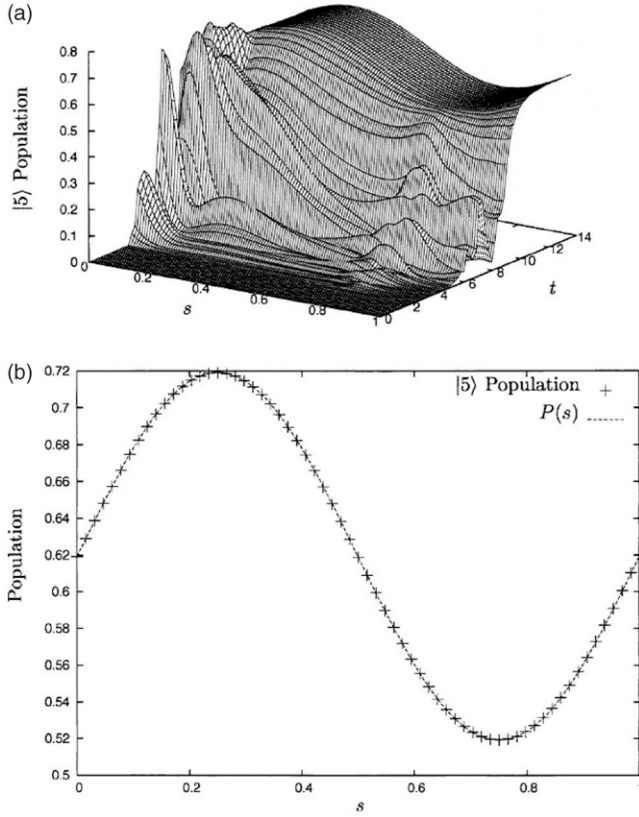


Figure 15. (a) Population surface of state $|5\rangle$ as a function of s and t for the tracking example described in the text ($P(s) = |\langle 5|\psi(s=0, T)\rangle|^2 + \sin(2\pi s)$). The elements μ_{34} and μ_{14} of the dipole operator were simultaneously modulated according to $\mu_{34}(s) = s$, $\mu_{14}(s) = 0.5(1 - s)$, with other elements kept fixed. At $s=0$, the only available direct transition pathway to $|5\rangle$ is $|1\rangle \rightarrow |4\rangle \rightarrow |5\rangle$, whereas at $s=1$ the only possible pathway is ladder climbing through all states. The change of the temporal dynamics with s reflects the use of different dynamical pathways throughout the s interval. (b) Cross-section of (a) at time $t = T$. The calculated observable population follows the imposed track very well. (From [32].)

In the dipole approximation, this relation becomes the following matrix integral equation:

$$\int_0^T \mu(s, t) \frac{\partial \varepsilon(s, t)}{\partial s} dt = U^\dagger(s) \nabla \Phi[U(s)], \quad (80)$$

where $\mu(s, t) \equiv U^\dagger(s, t) \mu U(s, t)$. When Φ is the observable expectation value objective function, we have

$$\int_0^T \mu(s, t) \frac{\partial \varepsilon(s, t)}{\partial s} dt = [\rho, U^\dagger(s) \Theta U(s)]. \quad (81)$$

On the basis of eigenstates, the matrix integral equation is written

$$\int_0^T \mu_{ij}(s, t) \frac{\partial \varepsilon(s, t)}{\partial s} dt = i\hbar \langle i | U^\dagger(s, T) \nabla \Phi[U(s, T)] | j \rangle. \quad (82)$$

To solve this equation, we first note that the flexibility in the choice of the representation of the variation in $\varepsilon(s, t)$ allows us to expand it on the basis of functions $\mu_{ij}(s, t)$, as

$$\frac{\partial \varepsilon(s, t)}{\partial s} = \sum_{i,j} x_{ij} \mu_{ij}(s, t). \quad (83)$$

Inserting this expansion into the above equation produces

$$\sum_{p,q} x_{pq}(s) \int_0^T \mu_{ij}(s, t) \mu_{pq}(s, t) dt = i\hbar \langle i | U^\dagger(s, T) \nabla \Phi[U(s, T)] | j \rangle. \quad (84)$$

If we denote the correlation matrix $G(s)$ as

$$G_{ij,pq}(s) = \int_0^T \mu_{ij}(s, t) \mu_{pq}(s, t) dt \quad (85)$$

$$= \int_0^T \langle i | \mu(s, t) | j \rangle \langle p | \mu(s, t) | q \rangle dt, \quad (86)$$

(as in equation (86) above, but now specifically in the case of the dipole approximation) and define

$$\Delta_{ij}(s) \equiv i\hbar \langle i | U^\dagger(s, T) \nabla \Phi[U(s, T)] | j \rangle, \quad (87)$$

it can be shown [42, 57] that the matrix integral equation (82) can be converted into the following N^2 -dimensional algebraic (nonsingular) differential equation (through a procedure analogous to that described in Appendix A.5 for scalar tracking):

$$\frac{\partial \varepsilon}{\partial s} = f_s + (\nu(\Delta) - \alpha)^T G^{-1} \nu(\mu(t)) \quad (88)$$

where $f_s = f_s(t)$ is a ‘free’ function resulting from the solution of the homogeneous differential equation (analogous to f_s in Appendix A.5), the operator ν vectorizes its matrix argument (as in equation (86)) and $\alpha \equiv \int_0^T \nu(\mu(t)) f_s dt$. This equation is similar to equation 30 in Appendix A.5, with the notable distinctions that the scalar functions $a_0(s, T)$, $a_1(s, T)$, $a_2(s, T)$ and Γ are now replaced by N^2 -dimensional vector and N^4 -dimensional matrix counterparts, respectively.

The computational overhead required for implementing such unitary matrix tracking, compared to scalar tracking, scales (roughly) as N^4 , the expense of inverting the matrix G . However, if this track can faithfully be followed, the scaling of the path length and convergence time to the optimal control solution could mirror that of the gradient flow of the objective function on the domain of unitary propagators. In the case of observable maximization, the analysis in the previous section suggests that for mixed initial states, the optimization trajectory followed by the unitary gradient flow will scale unfavourably for highly nondegenerate $\rho(0)$.

Importantly, it is possible to choose a global path that is even more favourable than the integrated unitary gradient flow of the gate objective function. Indeed, unitary matrix tracking algorithms have recently been developed [57] that follow the shortest path between the initial guess and the target matrix, namely the geodesic between these points in the unitary group. Across a wide variety of target gates and system dimensions, these algorithms were capable of tracking geodesic paths in $U(N)$ with almost negligible error. As the geodesic is the globally optimal path in $U(N)$, this indicates that it may be possible to define a system-independent search complexity for quantum optimal control problems in terms of the scaling of the length of the near-geodesic path with Hilbert space dimension.

The ability to track globally optimal paths in $U(N)$, originating in the favourable properties of the input-state map for discrete quantum systems (below), may also be useful for design of more efficient experimental quantum control algorithms. In particular, it is possible to specify the observable expectation value path that corresponds to the geodesic path in $U(N)$ for a given $\rho(0)$ and Θ . Tracking this path experimentally may cause the control system to follow a path in $U(N)$ that is globally more optimal, and more system-independent, than the projected path of the ε -gradient, permitting a more system-invariant definition of landscape complexity for OCE.

In this approach [42], the unitary propagator W that maximizes the observable expectation value is first determined numerically. This can be achieved at minimal computational cost if $\rho(0)$ and Θ are known. The observable track that corresponds to the geodesic $U(s) = \exp(i \log(W^\dagger U_0) s)$ that connects U_0 and W , i.e., $\langle \Theta(s) \rangle = \text{Tr}(U(s)^\dagger \rho(0) U(s) \Theta)$, is compatible with an infinite number of paths $V(s)$ in $U(N)$, and the set of these paths can be written

$$M_T = \left\{ V(s) \mid \text{Tr}(V(s) \rho(0) V(s)^\dagger \Theta) = \text{Tr}(U(s) \rho(0) U(s)^\dagger \Theta) = \langle \Theta(s) \rangle \right\}. \quad (89)$$

In order to implement a global OCE search along the desired observable track, computational overhead must be paid in order to reconstruct the initial density matrix of the system. Research in quantum statistical inference [41] has demonstrated that this reconstruction can generally be achieved with high statistical certainty at comparatively low cost. As in the case of gradient measurements, n repeated observations are made on identically prepared quantum systems. Each quantum measurement is described by a positive operator-valued measure (POVM) [58]. Denoting by \mathcal{F}_i the POVM corresponding to the i -th observation, the likelihood functional

$$L(\rho(t)) = \prod_{i=1}^n \text{Tr}(\rho(t) F_i) \quad (90)$$

describes the probability of obtaining the set of observed outcomes for a given density matrix $\rho(t)$. The method of maximal likelihood estimation (MLE) maximizes this function over the set of density matrices [58]. An effective parameterization of $\rho(t)$ is

$\rho(t) = T^\dagger T$, which guarantees positivity and Hermiticity, and the condition of unit trace is imposed via a Lagrange multiplier λ , to give

$$L(T) = \sum_{i=1}^n \ln \text{Tr}(T^\dagger T F_i) - \lambda \text{Tr}(T^\dagger T). \quad (91)$$

Standard numerical techniques, such as Newton–Raphson or downhill simplex algorithms, may be used to search for the maximum over the N^2 parameters of the matrix T . Note that other methods for state reconstruction, such as the maximum entropy method or Bayesian quantum state identification [41], can alternatively be used.

It is possible (and in many cases imperative for good results) to constrain the optimization trajectory to follow still more precise paths in $\mathcal{U}(N)$, by specifying observable tracks $\langle \Theta_i(s) \rangle$ corresponding to the desired unitary track $U(s)$ for multiple observable operators Θ_i . In the limit that the set $\{\Theta_i\}$ constitutes a complete tomographic set of observables ($N^2 - 1$ orthogonal observables Θ_i), and the initial density matrix is nondegenerate, the observable tracking algorithm becomes effectively identical to the unitary matrix tracking algorithm described above. For $1 \leq i \leq N^2 - 1$ observables, the geodesic track $U(s)$ is approximated to progressively greater accuracy, for only incrementally greater cost. Moreover, it can be shown [42] that the k gradients $\delta \langle \Theta_i(T) \rangle / \delta \varepsilon(t)$ are closely related, such that the overhead required to statistically sample $\delta \langle \Theta_1(T) \rangle / \delta \varepsilon(t)$ assists in the determination of $\delta \langle \Theta_2(T) \rangle / \delta \varepsilon(t), \dots, \delta \langle \Theta_k(T) \rangle / \delta \varepsilon(t)$.

The accuracy with which globally optimal paths in $U(N)$ can be tracked depends on the properties of the input-state map $M: \varepsilon(t) \rightarrow U(T)$. In particular, if the matrix G is singular or close-to-singular, the tracking errors will be greater and the performance of these globally optimal algorithms will be compromised. It can be shown that the requirement that G is nonsingular amounts to a more demanding form of controllability than full controllability, i.e., from any given unitary propagator $U(T)$, it must be possible to track to any neighbouring $U'(T)$ to first order in $\varepsilon(t)$. An important advantage of (orthogonal) observable tracking compared to unitary matrix tracking is that these algorithms place less stringent demands on controllability while reaping many of the benefits of following a globally direct path to the solution [42].

In the next section, we review properties of discrete quantum control systems that render them less likely to encounter tracking problems, compared to classical control systems.

8.2. Extremals of the input-state map

In section 2, we saw that the critical points of quantum optimal control variational problems that are not critical points of the map between control fields and unitary propagators can be identified analytically. These critical points are called normal (or regular) extremals in the terminology of control theory. In this section, we examine the properties of the so-called input-state map $\varepsilon(t) \rightarrow U(T)$ between control fields and associated dynamical propagators, comparing to the classical case. The critical points of

this map are called abnormal (or singular) extremals of the control problem (see Appendix A.2 for a formal definition).

Given the simple critical topology of normal extremals for quantum optimal control problems, and the integrability of the gradient flows of the objective functions on the domain of dynamical propagators, the properties of the input-state map play a dominant role in determining the effort required for locating optimal controls. As such, the behaviour (and design) of advanced search algorithms for quantum optimal control are largely governed by the properties of this map.

It can be shown that abnormal extremals correspond to control fields $\varepsilon(t)$ which produce a singular matrix G [57]. Because singular extremals correspond to places where algorithms that track paths in the space of dynamical propagators break down, their abundance plays a central role in determining the maximum possible efficiency of optimal control search.

For general problems of sub-Riemannian geometry [59], abnormal extremals exist in abundance. In particular, a dynamical system can be strongly controllable and still possess abnormal extremals. For discrete quantum systems, however, it is possible to prove certain analytical results pertaining to the existence of abnormal extremals that suggest that these extremals are particularly rare.

D'Alessandro and Dahleh [30] gave a complete classification of normal and abnormal extremals for the (single-input) optimal gate control problem on $SU(2)$, with field fluence as the cost. It was shown that the only abnormal extremal in this case is the control $\varepsilon(t) = -\text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$. For two control fields, if we write $A := a_1 \mu_1 + a_2 \mu_2 + a_3 [\mu_1, \mu_2]$, the only abnormal extremal is $\varepsilon_1 = -a_1$, $\varepsilon_2 = -a_2$. For three controls, there are no abnormal extremals.

We sketch the proof of this result for a single control field. Recall that in this case, the PMP Hamiltonian function for the maximum principle can be written

$$h(M, \lambda_0, \varepsilon) \equiv \text{Tr}[M(U_0^*(t)(H_d + \mu \varepsilon(t))U_0(t))] + \frac{1}{2} \lambda_0 \varepsilon^2(t) \tag{92}$$

as above, where M again plays the role of the conjugate momentum corresponding to the dynamical propagator state variable. We first demonstrate that all optimal controls except $\varepsilon(t) = -\text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$ are normal. It is useful to define $[[H_d, \mu], \mu] = c_1 H_d + c_2 \mu$ $[[H_d, \mu], H_d] = d_1 H_d + d_2 \mu$ where $c_2 = -d_1$, $d_1/c_1 = \text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$.

Now, if $\lambda_0 = 0$ (i.e., if the control is an abnormal extremal), the maximizer of the PMP-Hamiltonian must satisfy $\text{Tr}(M, U_0(t) \mu U_0(t)) = 0$. Differentiating twice, we obtain

$$c_1 \varepsilon(t) \text{Tr}(M U_0^*(t) H_d U_0(t)) + x_2 \varepsilon(t) \text{Tr}(M U_0^*(t) \mu U_0(t)) + d_1 \text{Tr}(M U_0^*(t) U_0(t)) + d_2 \text{Tr}(M U_0^*(t) \mu U_0(t)) = 0 \tag{93}$$

which implies

$$(c_1 \varepsilon(t) + d_1) \text{Tr}(M U_0^*(t) H_d U_0(t)) = 0. \tag{94}$$

It can be shown that $\varepsilon(t) \neq \text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$ implies $c_1 + \varepsilon(t) d_1 \neq 0$. Thus,

$$\text{Tr}(MU_0^*(t)H_d U_0(t)) \equiv 0.$$

It follows that M must also equal zero, which contradicts the supposition that the minimizer is abnormal. Thus, all optimal controls except $\varepsilon(t) \neq \text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$ are normal. Conversely, it can be verified that $\varepsilon = -\text{Tr}(H_d \mu) / \text{Tr}(\mu \mu)$ satisfies the maximum principle with $\lambda_0 = 0$ and M chosen such that $\text{Tr}(M \mu) = \text{Tr}(M, [H_d, \mu]) = 0$. This implies that it is an abnormal extremal, and since all extremals that are not of this form are normal, this is the only abnormal extremal. Note that the only abnormal extremal in this case is a constant control. For higher-dimensional gate control problems using one control field, it can be shown that there also exist constant controls that are abnormal extremals, although it is not yet clear whether additional abnormal extremals exist.

For the important problem of laser control of population transfer examined in section 3.2, where the control Hamiltonian couples only neighbouring energy levels (and the internal Hamiltonian can be eliminated), the existence of abnormal extremals has been studied for discrete quantum systems of arbitrary dimension; in particular, for any time t on the domain of a given solution to any such state control problem, there exists an interval $[t_1, t_2]$ around t where the solution is not strictly abnormal [29]. This property might be termed ‘weak normality’, by analogy to the phenomenon of weak resonance discussed above.

The need to partition the time domain of the solutions in this way arises for the same reason as in section 3.2, i.e., because the dynamics has singularities whenever a state vector coefficient is zero. For example, if $\psi_1 = \psi_2 = 0$, the control $V_{1,2}$ has no effect on dynamics, i.e., $V_{1,2} \psi_2 = V_{2,1} \psi_1 = 0$. In order to prove the property of weak normality, one can define a subspace of S^{n-1} on which the corresponding control problem does not encounter singularities. An auxiliary control problem can then be framed on this subspace, which can be shown straightforwardly to possess no abnormal extremals. Solutions to that control problem can be lifted to S^{n-1} to produce normal solutions to the original control problem.

We can accomplish these steps by first defining a partition $I \cup J$ of $\{1, \dots, n\}$ satisfying the following condition:

$$j \in I \leftrightarrow \psi_j(t) = 0 \quad \forall t \in [t_1, t_2] \quad (95)$$

$$j \in J \leftrightarrow \psi_j(t) \neq 0 \quad \forall t \in [t_1, t_2]. \quad (96)$$

In addition, it is necessary to subdivide J into subspaces connected by controls, because only in these subspaces are the vector fields $F_{j,k}(\psi) = V_{j,k} \psi_k$ corresponding to the controls identically nonzero. Two indices j, k of J are connected ($j \sim k$) if there exists a sequence j_1, \dots, j_s of indices of J such that $j = j_1, k = j_s$, and $\forall r < s, V_{j_r, j_{r+1}}$ is a control, i.e., the two energy levels E_j and E_k are connected if there exists path through state space where successive states are coupled by control matrix elements. Denote by K_1, \dots, K_r the equivalence classes defined by \sim . Let m_1, \dots, m_r denote their respective cardinalities, and define $M_0 = 0, M_l = \sum_{k \leq l} m_k$. For convenience, we reorder the indices in the

partition such that $\forall l \leq r, K_l = \{M_{l-1} + 1, \dots, M_l\}$ and $I = \{M_r + 1, \dots, n\}$. This simply shifts the indices responsible for singularities to the upper end of the spectrum.

The essential property of discrete quantum state control systems that permits proof of weak normality is that it is possible to frame this associated auxiliary control problem on an analytical submanifold of the domain of the original problem, S^{n-1} . Because $\sum_{j \in K_l} |\hat{\psi}_j(t)|^2$ is constant on $[t_1, t_2]$, we can define the analytic submanifold on which $\psi(t)$ evolves for $t \in [t_1, t_2]$ in terms of the $\psi_j(t_1)$ s as:

$$X = S^{m_1-1}(C_1) \times \dots \times S^{m_r-1}(C_r) \times \prod_{j \in I} \{\psi_j(t_1)\} \quad (97)$$

where $C_l = \sqrt{\sum_{j \in K_l} |\psi_j(t_1)|^2}$. In other words, because the subspaces labelled by l are connected by matrix elements of the control Hamiltonian, the corresponding components of the state vector undergo unitary (state vector norm-preserving) evolution on those subspaces. It follows that any extremal $\hat{\psi}$ of the original control problem is also an extremal of the auxiliary control problem on X , since $\hat{\psi}$ remains the same if the controls $\hat{V}_{j,k}$, where at least one of the indices j, k is in I , are set to zero.

It is then straightforward to show [29] that the cardinality of each subspace Δ_l generated by the control elements $V_{j,k}$ with $j, k \in K_l$ is equal to the cardinality of the Hilbert sphere $S^{m_l-1}(C_l)$, and hence that the cardinality of $\Delta = \bigoplus_l \Delta_l$ is equal to the cardinality of J . As such, no singularities can exist on this domain and the auxiliary control problem has no abnormal extremals. We refer the reader to [29] for details and for a discussion of how the normality of solutions to the auxiliary control problem is preserved upon lifting to S^{n-1} . Note that the conditions for normality are equivalent whether the problem is formulated on the space of quantum states or unitary propagators. Extending these results to the whole domain of the solution is as of yet an unsettled question, but this work represents the first step in that direction.

Note that the proofs of these analytical results pertaining to the sparseness of abnormal extremals for discrete quantum control problems make use of the compactness of the discrete unitary group of dynamical propagators. Since this property is not shared by the noncompact classical or continuous variable quantum dynamical propagators, the more general circumstance of abundant abnormal extremals most likely holds for those cases.

Although critical points of the input-state map are unlikely to be encountered directly during the search for optimal controls, unitary propagators within a certain distance of these critical points will be associated with nearly singular matrices G , thus possibly compromising global tracking efficiency. Current numerical work focuses on identifying the radius in $(\mathcal{U}N)$ within which such ill-conditioned matrices occur, for various families of homologous Hamiltonians.

9. Open quantum systems

Quantum systems of practical interest in chemistry or physics are always exposed to some kind of environment, which can render the dynamics nonunitary and irreversible. Intuitively, environmentally induced irreversible quantum dynamics would seem to

downgrade the quality of the control outcome. It is therefore imperative to determine whether the favourable features of quantum control landscapes, derived in the context of ideal closed systems, are preserved in the presence of environmental decoherence. In this section, we examine the effects of strong decoherence on the critical topology of open quantum system control problems.

The composite of the system and environment obeys the Schrödinger equation:

$$i\hbar \frac{d\rho_{\text{total}}}{dt} = [H_{\text{total}}, \rho_{\text{total}}]. \quad (98)$$

The composite Hilbert space is, where \mathcal{H}_S and \mathcal{H}_E are the Hilbert spaces of the system and environment, respectively, and the initial state of the total system is $\rho_{\text{tot}}(0) = \rho_S \otimes \rho_E$. The problem of maximizing the expectation value of an observable of the system can be expressed in terms of the Kraus dynamical propagators K_{mn} of the quantum system in the presence of the environment (see Appendix A.7):

$$J(K) = \text{Tr} \left(\sum_{m,n=1}^{\lambda} K_{mn} \rho K_{mn}^\dagger \Theta \right) \quad (99)$$

$$= \text{Tr} \{ K(\rho(0) \otimes I_\lambda) K^\dagger(\Theta \otimes I_\lambda) \}. \quad (100)$$

Using the terminology $K = F_E(U) = U(I_N \otimes \rho_E^{1/2})$, where U is a unitary propagator in $U(\lambda N)$, we can lift the landscape topology problem onto the composite Hilbert space as

$$J_{\rho E}(U) = \text{Tr} \{ F_{\rho E}(U)(\rho(0) \otimes I_\lambda) F_{\rho E}^\dagger(U) \Theta' \} \quad (101)$$

$$= \text{Tr}(UPU^\dagger \Theta') \quad (102)$$

where $P = \rho_S \otimes \rho_E$ and $\Theta' = \Theta \otimes I_\lambda$. The composite system is assumed to be controllable over $U(\lambda N)$. Under this assumption, Wu *et al.* [60] showed that no suboptimal traps exist in the control landscape for maximization of observable expectation values. Moreover, the enhanced controllability attainable with open dynamics actually broadens the range of attainable expectation values.

Because $K[\rho, N]$ is homeomorphic to the homogeneous space of $U(\lambda N)$, a landscape mapping satisfying the conditions of Theorem 4 can be built from $U(\lambda N)$ to $K[\rho, N]$. This theorem allows one to extract the critical topology of the Kraus landscape in terms of the associated unitary landscape. In particular, since the observable maximization landscape for unitary evolution was demonstrated to have no suboptimal traps (section 2), we can immediately conclude that the solution sets to open quantum system observable maximization problems also have no (normal) traps, assuming the environment is controllable. Although the latter condition may appear difficult to achieve in practice, in most cases the interaction of a system

with its environment is dominated by local interactions, which may be straightforward to control.

We note that the global observable and matrix tracking algorithms described in section 8 can be applied to open quantum systems as well. In this case, a complete tomographic set of observables $\{\Theta_i\}$ is composed of N^4 rather than N^2 operators [47], thus increasing the expense of tracking, but not rendering it prohibitive.

10. Conclusion and future challenges

As we have seen, there are stark differences between the optimal control landscapes for quantum and classical systems. In particular, the geometric properties of the compact Lie group of finite-dimensional quantum propagators endows the corresponding control landscapes with remarkable properties that are considerably simpler than those of classical systems.† Although quantization of a finite-dimensional classical system generally produces an infinite-dimensional quantum system, the Hilbert spaces of most quantum systems of practical interest that possess an infinite number of levels can be effectively truncated to finite dimensions. Thus, counterintuitively, locating optimal quantum controls becomes in many ways easier than locating corresponding classical controls. Given the apparent favourable scaling of landscape search complexity with Hilbert space dimension, even if distant energy levels play a role in the dynamics, the effort involved in locating controls may still be minimal.

This is most important for the practical feasibility of quantum control simulations and experiments on large molecules. The scaling of the expense of quantum dynamical simulation suggests that the computational problems inherent in quantum chemistry – high-precision electronic structure calculations become prohibitively expensive for most systems of practical interest – should be exacerbated for optimal control of such systems. However, the simple features of quantum control landscapes described above indicate that optimal control search need not add additional complexity to these problems.

The simple topology and geometry of quantum control landscapes, moreover, can be exploited to develop both numerical and experimental search algorithms that may outperform local or adaptive algorithms. The further development of global experimental algorithms is of particular interest, as these would take advantage of landscape structure without suffering from the exponentially unfavourable scaling of the cost of quantum simulation with Hilbert space dimension.

Although the landscapes for control of finite-dimensional quantum systems are thus simpler than those for classical systems, the need for statistical inference of quantum observable expectation values, states, or gradients thereof adds additional overhead to the cost of identifying optimal controls in an experimental context. This overhead is exacerbated when applying global algorithms that attempt to take advantage of information regarding the quantum state or dynamical propagator at each step

†Formally, the reduction in the dimensionality (hence complexity) of the search space for discrete quantum controls originates in quantum symmetries, as shown in section 3. Future work may aim to frame this statement within the context of Noether's theorem for optimal control, which assigns so-called conserved currents to solutions to the maximum principle, based on such symmetries.

along the landscape search trajectory. An especially noteworthy challenge, therefore, is the characterization of how the emerging methodologies of quantum statistical inference may be employed to further reduce the search complexity of quantum control problems.

The other feature of quantum dynamics that might be considered prohibitive to their effective control, namely quantum decoherence, was shown to not have a significant effect on some of the most important properties of optimal control landscapes, in particular their critical topology. Future work should more thoroughly explore how the geometry of the control landscapes and the effectiveness of global search algorithms are affected by noise, the nonunitary evolution of incoherent quantum dynamics, and measurement.

Appendix: Mathematical appendices

A.1. Critical topology

A.1.1. Landscape mapping.

Theorem 4: *Suppose the function $x=f(y)$ is locally surjective near some point $y_0 \in Y$, i.e., the Jacobian has full rank:*

$$\text{rank} \frac{df}{dy} \Big|_{y=y_0} = \dim X \Big|_{x=f(y_0)}$$

in some local coordinate system. Then y_0 is critical for $L \circ f$ in Y if and only if $x_0=f(y_0)$ is critical for L in X , and they have identical numbers of positive and negative Hessian eigenvalues at y_0 and x_0 , respectively. Moreover, if the inverse image $f^{-1}(x_0)$ of every critical point x_0 is connected, then the connected components of their critical manifolds are one-to-one between the two landscapes [60].

A.1.2. Hessian quadratic form: observable maximization

On the domain of unitary propagators, the Hessian quadratic form (HQF) for observable maximization can be written

$$\mathcal{H}_A(U) = \text{Tr}[-4A^2\rho(0)U\Theta U^\dagger + 4A\rho(0)AU\Theta U^\dagger]$$

at a unitary matrix U , expanded along an arbitrary direction A in the Lie algebra of $U(N)$. Consider a particular solution U_l that generates orderings $\lambda_j \rightarrow \pi(\lambda_j)$ of the eigenvalues of Θ , where the array π_l specifies an N-index permutation mapping. If we define the matrix elements of A as $A_{ij}=\alpha_{ij} + i\beta_{ij}$, we obtain after some straightforward calculations [17]

$$\mathcal{H}_A(\hat{U}_l) = - \sum_{j < k} (\alpha_{jk}^2 + \beta_{jk}^2) [(\lambda_{\pi(j)} - \lambda_{\pi(k)})(\epsilon_j - \epsilon_k)]$$

from which the counting results presented in section 2 can be derived. The number of positive principal axis directions equals the number of (j, k) pairs for which $(\lambda_j - \lambda_k)(\epsilon_j - \epsilon_k) \geq 0$, and the number of negative principal axis directions equals the number of (j, k) pairs for which $(\lambda_j - \lambda_k)(\epsilon_j - \epsilon_k) \leq 0$.

A.2. Maximum principle and adjoint control systems

Theorem 5 (Pontryagin maximum principle): *Consider the problem of steering the control system*

$$\dot{x} = f(x, u), \quad x \in M, \quad u \in \Omega \subset \mathbb{R}^k,$$

where M is the state space of the system, from some initial state $x(0) = x_0$ to some final state x_1 while minimizing a cost of the form $\int_0^T f^0(x, u) dt$. The maximum principle states that if the couple $\bar{u}(t), \bar{x}(t)$ is optimal, there exists an absolutely continuous vector $\lambda(t) \in \mathbb{R}^n$ and a constant $\lambda_0 \leq 0$, such that the PMP-Hamiltonian function $h(x(t), \lambda(t), u(t)) = \langle \lambda(t), f(x(t), u(t)) \rangle + \lambda_0 f^0(x(t), u(t))$ satisfies

$$h(\bar{x}(t), \lambda(t), \bar{u}(t)) = \max_u h(\bar{x}(t), \lambda(t), u)$$

and

$$\lambda_j(t) = -\frac{\partial h}{\partial x_j}, \quad j \in 1, \dots, n.$$

Moreover, denoting the tangent space to the manifold M at state $x(t')$ by $T_{x(t')}M$, we have $\langle \lambda(0), T_{x(0)}M \rangle = \langle \lambda(T), T_{x(T)}M \rangle$ (transversality condition) [27]. If the final time T is fixed, $h(\bar{x}(t), \lambda(t), \bar{u}(t))$ is constant, whereas if T is allowed to vary, $h(\bar{x}(t), \lambda(t), \bar{u}(t)) = 0$.

If the control objective is to minimize the final time T instead of a cost of the form above, the optimal trajectory on $[0, T]$ is associated with the Hamiltonian $-\lambda_0 + \langle \lambda(t), f(x(t), u(t)) \rangle$. In this case, $\max_u h(\bar{x}(t), \lambda(t), u) = 0$ in $[0, T]$, and we have the additional condition that if $\lambda_0 = 0$, then $\lambda(t) \neq 0$ for any t [27].

Definition 3 (Normal, abnormal extremals): A trajectory $\bar{x}(t)$ satisfying the above condition is called an extremal. If $\lambda_0 = 0$, it is called an abnormal extremal; if $\lambda_0 < 0$, it is called a normal extremal. If an extremal is abnormal but not normal, it is called a strictly abnormal extremal.

Definition 4 (Adjoint control system): Consider the following control system F on a Lie group G :

$$\dot{U} = -\frac{i}{\hbar} \left[H_d + \sum_{j=1}^m u_j \mu_j \right] U.$$

Let K denote the subgroup spanned by the control Hamiltonians μ_j , and denote the adjoint orbit of $-iH_d$ (the internal, or drift Hamiltonian) under the action of the subgroup K by $Ad_K(-iH_d)$, i.e., $Ad_K(-iH_d) = \{k_1^\dagger(-iH_d)k_1 | k_1 \in K\}$. Then the adjoint control system of F is defined as the system $\dot{P} = \mathcal{H}P$, $\mathcal{H} \in Ad_K(-iH_d)$, $P \in G$, which evolves on the coset space G/K [8].

Definition 5 (Infimizing time): For the control system F above, let $R(I, t)$ denote the reachable set from the identity in time t . Then $t^*(U_F) = \inf\{t \geq 0 | U_F \in R(I, t)\}$ is called the infimizing time for producing the propagator U_F .

Theorem 6 (Equivalence theorem): *The infimizing time $t^*(U_F)$ for steering the system*

$$\dot{U} = \left[H_d + \sum_{j=1}^m u_j H_j \right] U.$$

From $U(0) = I$ to U_F is the same as the minimum coset time $L^(KU_F)$ for steering the adjoint system*

$$\dot{P} = \mathcal{H}P, \quad \mathcal{H} \in Ad_K(H_d)$$

from $P(0) = I$ to KU_F [8].

Theorem 7 (Adjoint maximum principle): *For the above adjoint control system, denote the time-optimal control law by $\bar{\mathcal{H}}(t)$ and the corresponding optimal trajectory by $\bar{P}(t)$. Define an adjoint auxiliary cost function, $f(P) = \text{Tr}(\lambda^\dagger \mathcal{H}P)$, $P\lambda^\dagger \in p$. The corresponding adjoint PMP-Hamiltonian is $h(P(t), \lambda(t), \mathcal{H}(t)) = \text{Tr}(\lambda^\dagger(t)\mathcal{H}(t)P(t)) \equiv \text{Tr}(N(t)\mathcal{H}(t))$. The optimal adjoint control-trajectory pairs are then the solutions to the Hamiltonian equations*

$$\frac{d\lambda(t)}{dt} = -\frac{\partial h}{\partial P} = \hat{\mathcal{H}}(t)\lambda(t).$$

The adjoint maximum principle [9] demands that there exists a $N(t) \in p$ (directions in G/K space) such that

$$\begin{aligned} \bar{\mathcal{H}}(t) &= \text{argmax}_{\mathcal{H}} \text{Tr}(\mathcal{H}N(t)), \quad \mathcal{H} \in Ad_K(-iH_d) \\ \frac{d\bar{P}(t)}{dt} &= \bar{\mathcal{H}}(t)\bar{P}(t) \\ \frac{dN(t)}{dt} &= [\bar{\mathcal{H}}(t), N(t)]. \end{aligned}$$

A.3. Rotating wave approximation

The rotating wave approximation (RWA) consists of a unitary change of coordinates (and controls) by which the internal (drift) Hamiltonian can be eliminated in problems

involving atom-electromagnetic wave interactions, by virtue of the electromagnetic radiation being nearly resonant, or where the interaction Hamiltonian couples only neighbouring states [28]. We consider the latter case. Let $\psi(t) = U(t)\psi'(t)$. Then the state vector in the rotated coordinate system satisfies the Schrödinger equation:

$$i \frac{d\psi'(t)}{dt} = H'(t)\psi'(t)$$

where the Hamiltonian in the rotated coordinate system is

$$H' = U^{-1}HU - iU^{-1} \frac{dU}{dt}.$$

In order to eliminate the internal Hamiltonian, we choose $U(t) = \exp(-iDt)$; since $H = D + V(t)$, $H' = iU^{-1}DU^{-1} + U^{-1}(D + V(t))U = \exp(iDt)V(t)\exp(-iDt)$. Redefining $\psi' \rightarrow \psi$ and $H := -iH'$, we have

$$\frac{d\psi(t)}{dt} = H(t)\psi(t),$$

where H is skew-Hermitian. The elements of this Hamiltonian are either zero or are controls; as such, the drift is eliminated. Assuming that the control Hamiltonian V is off-diagonal (i.e., $V_{ij} = 0$ only if $i = j \pm 1$), the relation between the original and ‘new’ controls $H_{j,k}(t)$ is:

$$V_{j,k}(t) = H_{j,k}(t) \exp(i[(E_k - E_j)t + \pi/2]).$$

In the more general case where the control Hamiltonian is not off-diagonal but the control fields are assumed to be roughly in resonance with the system transition frequencies, the transformed Hamiltonian takes on a similarly simple form under the approximation that rapidly oscillating terms average to zero.

A.4. Analytical solutions to state and gate control problems

A.4.1. Low-dimensional gate control problems

Consider the right-invariant control system described in Definition 4. Following section 3.3, let G denote the special unitary group $SU(N)$. Call the subalgebra generated by the controls $\{\mu_1, \dots, \mu_m\}$ and the corresponding subgroup K . If we decompose $G = p \oplus l$ such that p is orthogonal to l , then p represents all possible directions to move in G/K space. Denote by $h \subset p$ a subspace of maximally commuting directions or generators in G/K .

Specifically, in the case of two-qubit systems, $G/K = SU(4)/SU(2) \otimes SU(2)$, $g = su(4)$, and $K = SU(2) \otimes SU(2)$. In this case, it can be shown that the Lie algebras l, p , and h are

$$\begin{aligned} l &= \text{span } i\{I_x, I_y, I_z, S_x, S_y, S_z\} \\ p &= \text{span } i\{I_x S_x, I_x S_y, I_x S_z, I_y S_x, I_y S_y, I_y S_z, I_z S_x, I_z S_y, I_z S_z\} \\ h &= \text{span } i\{I_x S_x, I_y S_y, I_z S_z\}. \end{aligned}$$

Decomposing the target unitary propagator as $U_F = k_2 \exp(Y) k_1$, (section 3.3), we have

$$U_F = k_1 \exp[-i(\alpha_1 I_x S_x + \alpha_2 I_y S_y + \alpha_3 I_z S_z)] k_2, \quad k_1, k_2 \in SU(2) \otimes SU(2),$$

where the sub-Riemannian problem consists of generating $\exp(Y)$ in the fastest possible way. If we define $k_y^- = \exp(-i\pi/2I_y) \exp(-i\pi/2S_y)$ and $k_y^+ = \exp(i\pi/2I_y) \exp(-i\pi/2S_y)$, we can verify that

$$k_y^\pm \exp(-iI_z S_z) (k_y^\pm)^{-1} = \exp(\pm iI_x S_x)$$

and similarly for k_x , showing we can generate any element of the Cartan subalgebra \mathfrak{h} .

In the case of three spins coupled by local interactions such that $J_{12} = J_{23}$, $J_{13} = 0$, G/K is a nonsymmetric space, but is still a finite-dimensional Riemannian manifold. Khaneja and coworkers [9] considered the generation of unitary propagators of the form

$$U = \exp(-i\theta I_{1\alpha} I_{2\beta} I_{3\gamma}), \quad \alpha, \beta, \gamma \in \{x, y, z\},$$

which are hard to produce as they involve trilinear terms in the effective Hamiltonian (trilinear propagators). Applying the decomposition $U_F = k_2 \exp(Y) k_1$, it is sufficient to produce

$$\exp(Y) = \exp(-i\theta I_{1z} I_{2z} I_{3z}), \quad \theta \in [0, 4\pi]$$

because all other propagators belonging to the set $\exp(-i\theta I_{1\alpha} I_{2\beta} I_{3\gamma} | \alpha, \beta, \gamma \in \{x, y, z\})$ of trilinear propagators can be produced from U_F in arbitrarily small time by selective hard pulses.

The corresponding adjoint control problem (Definition 4) has $\mathcal{H} \in \text{Ad}_k(-i2\pi J(I_{1z} I_{2z} + I_{2z} I_{3z}))$. For adjoint control problems, an equivalent version of the Pontryagin maximum principle exists (Appendix A.2). It can be verified [9] that the following analytical time-dependent control satisfies this principle:

$$\bar{H}(t) = -i2\pi J \left[(I_{1z} I_{2x} + I_{2x} I_{3z}) \cos\left(\frac{\beta t}{T}\right) - (I_{1z} I_{2y} + I_{2y} I_{3z}) \sin\left(\frac{\beta t}{T}\right) \right]$$

and steers the adjoint system from $P(0) = I$ to $P(T) \in KU_F$ in minimal time. The minimum time $t^*(U_F)$ required to produce a propagator of the form $U_F = \exp(-i\theta I_{1z} I_{2z} I_{3z})$, $\theta \in [0, 4\pi]$, is then given by

$$t^*(U_F) = \frac{\sqrt{\kappa(4 - \kappa)}}{2J}$$

where $\kappa = \theta/2\pi$.

A.4.2. Low-dimensional state control problems

Analytical solutions to problems of state-to-state population transfer can be obtained for two- and three-level quantum systems, for off-diagonal control Hamiltonians (i.e., $V_{jk} = 0$ if $j \neq k \pm 1$), under the rotating wave approximation (Appendix A.3). We summarize the results for the optimal control of population transfer in three-level systems using fluence as the cost, with two controls that span the control Lie algebra [28], since this provides an example for how objective function symmetry can endow integrability to quantum control systems and more generally simplify the search for optimal controls. Let us consider the problem of transferring the population from pure state $|1\rangle$ to pure state $|3\rangle$, with $\mu_{1,12} = \mu_{2,21} = 1$, $\mu_{2,12} = \mu_{2,21} = 1$. The Schrödinger equation for the Hamiltonian in this case can be written

$$\dot{c}_1 = -i\varepsilon_1(t)c_2, \quad \dot{c}_2 = -i(\varepsilon_1(t)c_1 + \varepsilon_2(t)c_3), \quad \dot{c}_3 = -iu_2(t)c_2,$$

where c_i denote the coefficients of the wavefunction eigenstates x_i . If we set $c_1 = x_1 + ix_2$, $c_2 = x_4 - ix_3$, $c_3 = x_5 + ix_6$, we can write this concisely as $\dot{x} = \varepsilon_1 F_1 + \varepsilon_2 F_2$, where $x = (x_1, \dots, x_6)$ and $F_1 = (-x_3, -x_4, x_1, x_2, 0, 0)$ and $F_2 = (0, 0, x_5, x_6, -x_3, -x_4)$ denote the action of the control Hamiltonians μ_1 and μ_2 on the state x . This is a problem on the five-dimensional Hilbert sphere, S^5 . The initial condition for this problem is a point on the circle $S_{\text{in}}^1 \equiv \{x \in S^5 \mid x_1^2 + x_2^2 = 1\}$, whereas the target is a point on the circle $S_{\text{fin}}^1 \equiv \{x \in S^5 \mid x_5^2 + x_6^2 = 1\}$. However, the dimensionality of this problem can be reduced if we assume that the controls are resonant (section 3.2). In this case, for each $x_0 \in S_{\text{in}}^1$ the orbit $O(x_0)$ (the reachable set of states) is a two-dimensional submanifold of S^5 , and hence the system is not fully controllable (i.e., not all superpositions of states can be reached from arbitrary initial conditions). Nonetheless, arbitrary eigenstate–eigenstate transitions can be controlled. Let us define $x_0(\alpha)$ as the initial condition $x_1(0) = \cos(\alpha)$, $x_2(0) = \sin(\alpha)$, $\alpha \in [0, 2\pi]$. Then this submanifold is the two-dimensional sphere defined by the equation $x_1'^2 + x_3'^2 + x_5'^2 = 1$, where $x' = R \otimes I_3 x$. In other words, due to the isometry, all the points in S_{in} can be considered equivalently. Therefore, we can study the optimal control problem on the orbit $O(x_0)$. Let us consider the case where x_0 is defined by $x_1 = 1$; in this case, $O(x_0)$ is the sphere defined by $x_1^2 + x_3^2 + x_5^2 = 1$.

We can then execute a change of variables to $y_1 = x_1, y_2 = x_3, y_3 = -x_5$, such that the control system can be rewritten as

$$\begin{pmatrix} \dot{y}_1 \\ \dot{y}_2 \\ \dot{y}_3 \end{pmatrix} = u_1 F_1 + u_2 F_2; \quad F_1 = \begin{pmatrix} -y_2 \\ y_1 \\ 0 \end{pmatrix}, \quad F_2 = \begin{pmatrix} 0 \\ -y_3 \\ y_2 \end{pmatrix}.$$

It is useful to frame the reduced problem in spherical coordinates, where

$$y_1 = \cos(\theta) \cos(\phi)$$

$$y_2 = \sin(\theta)$$

$$y_3 = \cos(\theta) \sin(\phi).$$

The control system can then be written

$$\begin{pmatrix} \dot{\theta} \\ \dot{\phi} \end{pmatrix} = v_1 G_1 + v_2 G_2,$$

where $G_1 = \partial_\theta$, $G_2 = \tan(\theta)\partial_\phi$.

In spherical coordinates, the Hamiltonian associated with the maximum principle is

$$h(\theta, \phi, P_\theta, P_\phi, v_1, v_2) = \langle P, v_1 G_1 + v_2 G_2 \rangle + p_0(v_1^2 + v_2^2) = v_1 P_\theta + v_2 P_\phi \tan(\theta) + p_0(v_1^2 + v_2^2).$$

The maximum principle demands $\partial h/\partial v_1 = 0$, $\partial h/\partial v_2 = 0$; as such, $v_1 = P_\theta$, $v_2 = P_\phi \tan(\theta)$. The Hamiltonian corresponding to these controls is $\hat{h} = 1/2((P_\theta^2 + (\tan(\theta)P_\phi)^2)$. The Hamiltonian equations of motion following from the maximum principle are then:

$$\begin{aligned} \dot{\theta} &= \frac{\partial \hat{h}}{\partial P_\theta} = P_\theta, & \dot{\phi} &= \frac{\partial \hat{h}}{\partial P_\phi} = P_\phi \tan^2(\theta) \\ \dot{P}_\theta &= \frac{\partial \hat{h}}{\partial \theta} = -P_\phi^2 \tan(\theta)(1 + \tan^2(\theta)), & \dot{P}_\phi &= -\frac{\partial \hat{h}}{\partial \phi} = 0. \end{aligned}$$

This Hamiltonian system is Liouville integrable, since there are two independent and commuting constants of the motion \hat{h} and $P_\phi = a$. The solution for minimal fluence, with fixed transfer time T can then be shown to be [28]:

$$\min\left(\int_0^T (\varepsilon_1^2 + \varepsilon_2^2) dt\right) = \frac{3}{4}\pi^2 \frac{1}{T}.$$

A.5. Diffeomorphic homotopy on control landscapes

We sketch the derivation of the general diffeomorphic homotopy procedure for Hamiltonian morphing and observable tracking [31–33]. The condition for remaining on a designated level set of an observable control landscape is

$$\frac{dF(s)}{ds} = \frac{d\langle \Theta(s) \rangle_T}{ds} = 0$$

In the following, we use the subscript + to refer to a point infinitesimally close to the system at point s . The interaction dynamical propagator in the interaction picture can then be written $U_I(t, 0) = U^\dagger(t, 0)U_+(t, 0)$. To first order, the interaction picture

propagator at the final time T is:

$$U_I(T, 0) = I + \frac{ds}{i\hbar} \int_0^T dt \left(U^\dagger(t, 0) \frac{\partial H(s, t)}{\partial s} U(t, 0) \right).$$

If we Taylor expand the Hamiltonian to first order at algorithmic time s , i.e.,

$$H_+(s, t) = H(s + ds, t) = H(s, t) + ds \frac{\partial H(s, t)}{\partial s},$$

the expectation value of the observable of interest at the point $s + ds$ in Hamiltonian space can be expressed as:

$$\langle \Theta(s + ds) \rangle_T = \langle \Theta(s) \rangle_T + \frac{ds}{i\hbar} \left\langle \psi_0 \left| \int_0^T dt \left[\Theta(T), U^\dagger(t, 0) \frac{\partial H(s, t)}{\partial s} U(t, 0) \right] \right| \psi_0 \right\rangle. \quad (\text{A1})$$

It can then be shown that the condition for remaining on the level set can be written

$$\frac{dF(s)}{ds} = \int_0^T dt \left(a_0(s, t, T) \frac{\partial \varepsilon(s, t)}{\partial s} + a_1(s, t, T) \varepsilon(s, t) + a_2(s, t, T) \right) = 0. \quad (\text{A2})$$

where

$$\begin{aligned} a_0(s, t, T) &= -\frac{1}{i\hbar} \left\langle \psi_0 \left| [U^\dagger(T, 0) \Theta U(T, 0), U^\dagger(t, 0) \mu(s) U(t, 0)] \right| \psi_0 \right\rangle \\ a_1(s, t, T) &= -\frac{1}{i\hbar} \left\langle \psi_0 \left| \left[U^\dagger(T, 0) \Theta U(T, 0), U^\dagger(t, 0) \frac{d\mu(s)}{ds} U(t, 0) \right] \right| \psi_0 \right\rangle \\ a_2(s, t, T) &= -\frac{1}{i\hbar} \left\langle \psi_0 \left| \left[U^\dagger(T, 0) \Theta U(T, 0), U^\dagger(t, 0) \frac{dH_d(s)}{ds} U(t, 0) \right] \right| \psi_0 \right\rangle. \end{aligned}$$

A natural way to solve equation (A2) is to transform it into an initial value problem for the laser field $\varepsilon(s, t)$. Equation (A2) can be reexpressed as the differential equation

$$a_0(s, t, T) \frac{\partial \varepsilon(s, t)}{\partial s} + a_1(s, t, T) \varepsilon(s, t) + a_2(s, t, T) = f(s, t)$$

where $f(s, t)$ is an arbitrary function satisfying the constraint

$$\int_0^T f(s, t) dt = 0 \quad \forall s.$$

In general the coefficient $a_0(s, t, T)$ may vanish at some values of s and t . Since there are *a priori* no restrictions on $a_0(s, t, T)$, we cannot preclude the possibility of singular behaviour with the class of D-MORPH controls admitted by equation (A2). Singular behaviour is unattractive because it implies the possible existence of similar undesirable behaviour in the control field ε . Regardless of the behaviour of a_0 , a nonsingular class of D-MORPH solutions may be generated. Defining the integral

$$b(s, T) = - \int_0^T (a_1(s, t, T)\varepsilon(s, t) + a_2(s, t, T))dt,$$

we have

$$\int_0^T a_0(s, t, T) \frac{\partial \varepsilon(s, t)}{\partial s} dt = b(s, T),$$

which is a Fredholm integral equation of the first kind. More generally, in the case that following a track of objective function values $\langle \Theta(s) \rangle$ is desired, instead of remaining on a level set, we have

$$\int_0^T a_0(s, t, T) \frac{\partial \varepsilon(s, t)}{\partial s} dt = b(s, T) + \frac{d\langle \Theta(s, t) \rangle}{ds}.$$

Then, it can be shown [31] that this equation can be transformed into the equivalent (nonsingular) differential equation

$$\frac{\partial \varepsilon(s, t)}{\partial s} = S(t) \left[f(s, t) + \frac{b(s, T) + \left(\frac{d\langle \Theta \rangle}{ds} - \gamma(s) \right) a_0(s, t, T)}{\Gamma(s)} \right] \quad (\text{A3})$$

where $S(t)$ is an arbitrary shape function, $\Gamma(s) = \int_0^T S(t) [a_0(s, t, T)]^2 dt$ and $\gamma(s) = \int_0^T S(t) a_0(s, t, T) f(s, t) dt$ is the projection of the arbitrary function $f(s, t)$ onto $a_0(s, t, T)$.

The freedom to choose $f(s, t)$ corresponds to the multiplicity of control field solutions on the level set, and arises naturally as a consequence of the underspecified nature of the integral form of the original D-MORPH equation. In particular, in the case of fluence minimization, $f(s, t) = f_m(s, t)$ satisfies the condition

$$\int_0^T S(t) a_0(s, t, T) f_m(s, t) dt = 0.$$

It can be shown [31] that with the choice

$$f(s, t) = - \frac{1}{\Delta s} \frac{\varepsilon(s, t)}{S(t)},$$

the algorithm seeks to minimize the total field fluence at each step. Analogous free functions that correspond to other auxiliary costs (such as minimal time) in the Pontryagin maximum principle can be constructed.

A.6. Controllability on compact Lie groups

Definition 6 (Reachable sets and controllability): Consider a control system F defined on a manifold M . For each $T > 0$, and each x_0 in M , the set of points reachable from x_0 at time T , denoted by $A(x_0, T)$, is equal to the set of terminal points $x(T)$ of integral curves of F that originate at x_0 . The union of $A(x_0, T)$, for $T \geq 0$, is called the reachable set [27] from x_0 . The set of points reachable in T or fewer units of time, defined as the union of $A(x_0, t)$, $t \leq T$, is denoted $A(x_0, \leq T)$. A control system F is controllable if any point of M is reachable from any other point of M , at any time $T > 0$.

Definition 7 (Invariant control system on a Lie group): A control system on a Lie group G , where G is the Lie group associated with a Lie algebra \mathfrak{h} , is defined by the equations

$$\dot{U} = AU(t) + \sum_{i=1}^m u_i(t)B_iU(t)$$

where A and the B_i , $i = 1, \dots, m$ belong to \mathfrak{h} , $U(t)$ belongs to G , and the $u_i(t)$ are scalar functions of time which play the role of the external controls. The control system is said to be right-invariant if the following condition holds: If $U(t)$ is a solution corresponding to the initial condition equal to the identity matrix, the solution corresponding to the initial condition F is given by $U(t)F$.

A.7. Kraus superoperator formalism

Consider a composite of system and environment whose Hamiltonian H_{total} consists of the Hamiltonians of the system, environment, and their interaction. The total system evolution operator is $U_{\text{total}}(t)$ on the total Hilbert space $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_E$, where \mathcal{H}_S and \mathcal{H}_E are the Hilbert spaces of the system (of dimension N) and environment, respectively. The initial state of the total system is $\rho_{\text{tot}}(0) = \rho_S \otimes \rho_E$. We can obtain an expression for the system dynamics $\rho_S(t)$ by tracing $\rho_{\text{total}}(t) = U_{\text{total}}(t)\rho_{\text{total}}(0)U_{\text{total}}^\dagger(t)$ over the environment [60]:

$$\rho_S(t) = \text{Tr}_E\{U_{\text{total}}(t)(\rho_S \otimes \rho_E)U_{\text{total}}^\dagger(t)\}.$$

Define the λN -dimensional matrix $K(t) = \text{Tr}_E\{U_{\text{total}}(t)(I_{N\rho_E^{1/2}})\}$. Divide K into $\lambda^2 N \times N$ matrices $K_{\alpha\beta}(t) = |\beta\rangle\langle\alpha|$, ($\alpha, \beta = 1, \dots, \lambda$), where α, β constitute an arbitrary basis

for \mathcal{H}_E . These matrices form the Kraus representation of the dynamical map as

$$\rho_S(t) = \sum_{\alpha, \beta=1}^{\lambda} K_{\alpha\beta}(t)\rho(0)K_{\alpha\beta}^\dagger(t).$$

A.8. Symplectic propagators

Consider a continuous variable quantum system with a quadratic Hamiltonian $H(t)$. $H(t)$ induces a Hamiltonian vector field, which generates a one-parameter family of transformations $U(t)$ on the Hilbert space \mathcal{H} that obeys the Schrödinger equation

$$\frac{\partial U(t)}{\partial t} = \frac{i}{\hbar} H(t)U(t), \quad (\text{A4})$$

where the parameter is the time. The evolution propagator transforms the quadrature vector of position and momentum operators $\hat{z} = (\hat{q}_1, \dots, \hat{q}_N; \hat{p}_1, \dots, \hat{p}_N)^T$ linearly through

$$U : \hat{z}_\alpha \rightarrow U^\dagger(t)\hat{z}_\alpha U(t) = \sum_{\beta} S_{\alpha\beta}(t)\hat{z}_\beta,$$

where the $2N \times 2N$ matrix $S(t)$ is an element of the symplectic group $\text{Sp}(2N, \mathbb{R})$ that satisfies $S^T J S = J$, with

$$J = \begin{pmatrix} & I_N \\ -I_N & \end{pmatrix}.$$

Thus, the matrix S captures the Heisenberg equations of motion for the operators \hat{z}_i , and the unitary propagator U forms the metaplectic unitary representation of S in $\text{Sp}(2N, \mathbb{R})$ [61]. Like the infinite-dimensional unitary group (but unlike the finite-dimensional unitary group), $\text{Sp}(2N, \mathbb{R})$ is noncompact.

Whereas a logical operation on N discrete quantum bits (qubits) is represented by a 2^N -dimensional unitary matrix, the corresponding operation on N continuous quantum bits (qunits) can be represented by a $2N$ -dimensional symplectic matrix.

Acknowledgements

The authors acknowledge support from DARPA.

References

- [1] M. Demiralp and H. A. Rabitz, *Phys. Rev. A* **47**, 809 (1993).
- [2] A. Assion, T. Baumert, M. Bergt, T. Brixner, and B. Kiefer, *Science* **282**, 5390 (1998).
- [3] T. Baumert, T. Brixner, V. Seyfried, M. Strehle, and G. Gerber, *Appl. Phys. B* **65**, 779782 (1997).
- [4] J. L. Herek and J. Photochem., *Photobiol.* **180**, 225 (2006).
- [5] R. Bartels, S. Backus, E. Zeek, L. Misoguti, G. Vdovin, I. P. Christov, M. M. Murnane, and H. C. Kapteyn, *Nature* **406**, 164 (2000).
- [6] J. P. Palao and R. Kosloff, *Phys. Rev. Lett* **89**, 188301 (2002).

- [7] M. Grace, C. Brif, H. A. Rabitz, I. Walmsley, R. L. Kosut, and D. A. Lidar, in press (2007), eprint arXiv:quant-ph/0611189.
- [8] N. Khaneja, R. W. Brockett, and S. J. Glaser, *Phys. Rev. A* **63**, 032308 (2001).
- [9] N. Khaneja, S. J. Glaser, and R. W. Brockett, *Phys. Rev. A* **65**, 032301 (2002).
- [10] R. Wu and H. A. Rabitz, in preparation (2007).
- [11] J. von Neumann, *Tomsk. Univ. Rev.* **1**, 286 (1937).
- [12] J. von Neumann, Unpublished works, Institute for Advanced Study Archives (1937).
- [13] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, *Science* **303**, 1998 (2004).
- [14] M. D. Girardeau, S. G. Schirmer, J. V. Leahy, and R. M. Koch, *Phys. Rev. A* **58**, 2684 (1998).
- [15] T. S. Ho and H. A. Rabitz, *J. Photochem. Photobiol. A* **180**, 226 (2006).
- [16] R. Wu, M. M. Hsieh, and H. A. Rabitz, submitted (2007).
- [17] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, *J. Chem. Phys.* **51**, 204107 (2006).
- [18] Z. Shen, M. M. Hsieh, and H. A. Rabitz, *J. Chem. Phys.* **124**, 204106 (2006).
- [19] G. Dirr, U. Helmke, K. Huper, M. Kleinsteuber, and Y. Liu, *J. Global Opt.* **35**, 443 (2006).
- [20] S. J. Glaser, T. Schulte-Herbruggen, M. Sieveking, O. Scheletzky, N. C. Nielsen, O. W. Sorensen, and C. Griesinger, *Science* **280**, 421 (1998).
- [21] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, *Phys. Rev. A* **72**, 052337 (2005).
- [22] T. S. Ho and H. A. Rabitz, to be submitted (2007).
- [23] T. Frankel, in *Differential and Combinatorial Topology: A Symposium in Honor of Marston Morse*, edited by S. S. Cairns (Princeton University Press, Princeton, NJ, 1962), pp. 37–53.
- [24] R. Wu, R. Chakrabarti, and H. A. Rabitz, to be submitted (2007), eprint arXiv:0708.2118 [quant-ph].
- [25] S. Lloyd and S. L. Braunstein, *Phys. Rev. Lett.* **82**, 1784 (1999).
- [26] R. Wu, R. Chakrabarti, and H. A. Rabitz, to be submitted (2007), eprint arXiv:0708.3822 [math-ph].
- [27] V. Jurdjevic, *Geometric Control Theory* (Cambridge University Press, Cambridge, 1997).
- [28] U. Boscain, G. Charlot, J. P. Gauthier, S. Guerin, and H. R. Jauslin, *J. Math. Phys.* **43**, 2017 (2002).
- [29] U. Boscain and G. Charlot, to be submitted (2007).
- [30] D. D'Alessandro and M. Dahleh, *IEEE Trans. Autom. Control* **46**, 866 (2001).
- [31] A. Rothman, T. S. Ho, and H. A. Rabitz, *Phys. Rev. A* **72**, 023416 (2005).
- [32] A. Rothman, T. S. Ho, and H. A. Rabitz, *J. Chem. Phys.* **123**, 134104 (2006).
- [33] A. Rothman, T. S. Ho, and H. A. Rabitz, *Phys. Rev. A* **73**, 053401 (2006).
- [34] V. Beltrani, J. Dominy, and H. A. Rabitz, *Phys. Rev. A* **74**, 043414 (2007).
- [35] J. Roslund, M. Roth, and H. A. Rabitz, *Phys. Rev. A* **74**, 043414 (2006).
- [36] J. Roslund and H. A. Rabitz, to be submitted (2007).
- [37] D. Cardoza, C. Trallero-Herrero, F. Langhojer, H. A. Rabitz, and T. Weinacht, *J. Chem. Phys.* **122**, 124306 (2005).
- [38] E. Wells, K. J. Betsch, C. W. S. Conover, M. J. DeWitt, D. Pinkham, and R. R. Jones, *Phys. Rev. A* **72**, 063406 (2005).
- [39] M. Wollenhaupt, A. Prakelt, C. Sarpe-Tudoran, D. Liese, and R. Baumert, *J. Mod. Opt.* **52**, 2187 (2005).
- [40] J. Roslund and H. A. Rabitz, to be submitted (2007).
- [41] J. Malley and J. Hornstein, *Stat. Sci.* **8**, 433 (1993).
- [42] R. Chakrabarti, R. Wu, and H. A. Rabitz, in preparation (2007), eprint arXiv:0708.3384 [quant-ph].
- [43] V. Ramakrishna, M. V. Salapaka, M. Dahleh, and H. A. Rabitz, *Phys. Rev. A* **51**, 960 (1995).
- [44] V. Jurdjevic and J. Sussmann, *J. Diff. Equat.* **12**, 313 (1972).
- [45] D. D'Alessandro, *Syst. Control Lett.* **41**, 213 (2000).
- [46] R. Wu, T. J. Tarn, and C. W. Li, *Phys. Rev. A* **73**, 012719 (2006).
- [47] M. Jezek, J. Fiurasek, and Z. Hradil, *Phys. Rev. A* **68**, 012305 (2003).
- [48] G. Maday and G. Turinici, *J. Chem. Phys.* **118**, 8191 (2003).
- [49] G. G. Balint-Kurti, F. R. Manby, Q. Ren, M. Artamonov, T. S. Ho, and H. A. Rabitz, *J. Chem. Phys.* **122**, 084110 (2005).
- [50] B. Amstrup, G. J. Toth, G. Szabo, H. A. Rabitz, and A. Loerincz, *J. Phys. Chem.* **99**, 5206 (1995).
- [51] R. Levis, G. Menkir, and H. A. Rabitz, *Science* **292**, 709 (2001).
- [52] R. Chakrabarti, R. Wu, and H. A. Rabitz, in preparation (2007).
- [53] R. Chakrabarti, R. Wu, and H. A. Rabitz, to be submitted (2007), eprint arXiv:0708.3513 [quant-ph].
- [54] H. T. Siegelmann, A. B. Hur, and S. Fishman, *Phys. Rev. Lett.* **83**, 1463 (1999).
- [55] G. T. Riviello, K. A. Moore, and H. A. Rabitz, in preparation (2007).
- [56] K. A. Moore, M. M. Hsieh, and H. A. Rabitz, to be submitted (2007).
- [57] J. Dominy and H. A. Rabitz, in preparation (2007).
- [58] J. Rehacek, Z. Hradil, E. Knill, and A. I. Lvovsky, *Phys. Rev. A* **75**, 042108 (2007).
- [59] A. A. Agrachev and A. V. Sarychev, *J. Math. Syst. Est. Control* **5**, 1 (1995).
- [60] R. Wu, A. Pechen, H. A. Rabitz, M. M. Hsieh, and B. Tsou, submitted (2007), eprint arxiv:0708.2119 [quant-ph].
- [61] Arvind, B. Dutta, N. Mukunda, and R. Simon, *Phys. Rev. A* **52**, 1609 (1995).