

Globally optimal control of quantum dynamics

Denys I. Bondar^{1,*}, Llorenç Balada Gaggioli^{2,3,†}, Georgios Korpas^{4,2,5,‡}, Jakub Marecek^{2,§},
Jiri Vala^{6,7,8,||} and Kurt Jacobs^{9,10,¶}

¹Department of Physics and Engineering Physics, *Tulane University*, New Orleans, Louisiana 70118, USA

²Faculty of Electrical Engineering, *Czech Technical University in Prague*, Prague 121 35, Czech Republic

³LAAS-CNRS, *Université de Toulouse*, Toulouse, France

⁴Quantum Technologies Group, *HSBC Lab*, 20 Pasir Panjang Road, 117439 Singapore, Singapore

⁵Archimedes Research Unit on AI, Data Science and Algorithms, *Athena Research Center*, 1 Artemidos Street, 15125 Marousi, Greece

⁶Department of Physics, *Maynooth University*, Maynooth, Ireland

⁷School of Theoretical Physics, *Dublin Institute for Advanced Studies*, Dublin, Ireland

⁸*Tyndall National Institute*, Cork, Ireland

⁹U.S. *DEVCOM Army Research Laboratory*, Adelphi, Maryland 20783, USA

¹⁰Department of Physics, *University of Massachusetts at Boston*, Boston, Massachusetts 02125, USA



(Received 31 March 2023; accepted 17 October 2025; published 21 November 2025)

Optimization of constrained quantum control problems powers quantum technologies. This task becomes very difficult when these control problems are nonconvex and plagued with dense local extrema. For such problems, current optimization methods must be repeated many times to find good solutions, each time requiring many simulations of the system. Here, we present *quantum control via polynomial optimization* (QCPOP), a method that eliminates this problem by directly finding globally optimal solutions. The resulting increase in speed, which can be a thousandfold or more, makes it possible to solve problems that were previously intractable. This remarkable advance is due to global optimization methods recently developed for polynomial functions. We demonstrate the power of this method by showing that it obtains an optimal solution in a single run for a problem in which local extrema are so dense that gradient methods require thousands of runs to reach a similar fidelity. Since QCPOP is able to find the global optimum for quantum control, we expect that it will not only enhance the utility of quantum control by making it much easier to find the necessary protocols, but also provide a key tool for understanding the precise limits of quantum technologies. Finally, we note that the ability to cast quantum control as polynomial optimization resolves an open question regarding the computability of exact solutions to quantum control problems.

DOI: [10.1103/g4fb-xm13](https://doi.org/10.1103/g4fb-xm13)

I. INTRODUCTION

Controlling the evolution of microscopic quantum systems is essential for the development of future quantum technologies [1–6]. The primary method of implementing such control is to change the Hamiltonian of the system with time [7–13]. This change can be implemented by changing classical control fields, although the number of these “control knobs” is usually very limited. Fortunately, the algebra of linear operators is sufficiently complex that even changing the size of a single term in the Hamiltonian, given a sufficient time, will usually

provide universal control of a system. The value(s) of the control field(s) as a function of time are referred to as the *control functions* or simply *controls*.

It has been observed that in the absence of significant constraints, finding controls to generate a particular final state or a unitary operation with high fidelity can be achieved easily using standard gradient search methods [12,14]. This implies that the search “landscape” in the space of controls is largely free of local extrema. However, constraints are usually an important part of control problems for applications in quantum technology. The most ubiquitous is a constraint on the duration of the control; given ever-present dissipation and decoherence processes, it is usually important to achieve the goal in the shortest possible time. Specific classes of control problems require other kinds of constraints, examples of which are hybrid quantum-classical optimization and variational quantum algorithms [15,16].

The methods most commonly employed to solve quantum control problems use gradient search methods. The simplest approach is to discretize the control function in some way and use a state-of-the-art gradient search method such as the Broyden-Fletcher-Goldfarb-Shanno algorithm [17]. Alternatively, a number of methods that add some additional

*Contact author: dbondar@tulane.edu

†Contact author: llorenç.balada.gaggioli@fel.cvut.cz

‡Contact author: georgios.korpas@hsbc.com

§Contact author: jakub.marecek@fel.cvut.cz

||Contact author: jiri.vala@mu.ie

¶Contact author: dr.kurt.jacobs@gmail.com

sophistication prior to applying a gradient search are commonly used: GRAPE (GRAdient Ascent Pulse Engineering) [18–21], CRAB (Chopped RANdom-Basis) [22–24], and the Krotov method [8,25,26]. For control problems with a significant density of local extrema (“traps”), these methods must be run many times, and if the density is too high, they simply fail to find good solutions. Stochastic search algorithms [27] such as simulated annealing, a plethora of reinforcement and machine learning methods [28–31], and a Bayesian approach [32] have all been employed to address this issue. While these machine learning search methods have achieved some significant improvements over gradient methods, fundamentally all of them merely search for local optima, and thus must cover a significant fraction of the search space to raise the likelihood of finding a global optimum.

In the last few years, tremendous advances have been made in solving optimization problems for polynomials (see Appendix A for a review). The methods are remarkable in their ability to bypass local optima and locate the global optimum. These methods have two stages. In the first stage, they locate a global optimum, and, in particular, are able to *guarantee* that the optimum is a global optimum. In the second stage, the methods find an argument to the function that achieves this optimum. The second stage is not yet as robust as the first stage, but the implementations we employ have been able to find arguments that realize the global optimum in all the problems we have tried.

To apply polynomial optimization methods to quantum control, we develop a procedure to approximate quantum control problems with polynomial optimization problems. Note that even with such a procedure, it is not guaranteed that polynomial optimization will provide a tremendous increase in speed and optimality over traditional search methods: It could be that to achieve a sufficiently good polynomial approximation results in a polynomial optimization problem that is so complex as to require very large numerical resources even for the powerful algorithms that are able to locate the global optima. Here, we demonstrate, at least for the common nontrivial examples we consider, that this is not the case; not only does a single run of a polynomial optimization algorithm take a time similar to that of state-of-the-art search methods, but with that single run it finds an optimum solution, whereas search methods may require thousands of runs if they are able to find good solutions at all. Not only does this tremendously increase the speed at which quantum control protocols can be obtained, but makes it possible to solve problems that were previously infeasible. As we show, they are also able to find optimal solutions where previous methods could not.

II. BACKGROUND: QUANTUM CONTROL

Consider a finite-level quantum system described by a Hilbert space $\mathcal{H}$. Both a quantum state $|\psi\rangle$ and the unitary operator, $U(T)$, that evolves the system to time T , evolve under Schrödinger equation:

$$\partial_t |\psi(t)\rangle = A(t) |\psi(t)\rangle, \quad |\psi(0)\rangle = |\psi_0\rangle, \quad (1)$$

$$\partial_t U(t) = A(t)U(t), \quad U(0) = \mathbb{1}. \quad (2)$$

Here, $A(t)$ is anti-Hermitian and given by $A(t) = -iH$, where H is the Hamiltonian of the system. (We have set $\hbar = 1$ for convenience.) To control the system, we change one (or more) terms in the Hamiltonian with time. Therefore, we set $H = H_0 + E(t)V$, where $E(t)$ is the control function. The constant part of the Hamiltonian, H_0 , is called the drift and the time-varying part, V is called the control Hamiltonian.

The vast majority of control problems for quantum systems are of the following four kinds. The first is called the *fixed-time* or *finite-horizon* control, and involves finding the control $E(t)$ that generates (as near as possible) a “target” unitary, U^* , at the end of a given evolution time T . Specifically,

$$\begin{aligned} &\underset{E(t)}{\text{minimize}} && F[U(T), U^*] \\ &\text{subject to} && \partial_t U(t) = A(t)U(t), \quad U(0) = \mathbb{1}, \\ &&& A(t) = -i(H_0 + E(t)V), \end{aligned} \quad (3)$$

where F is a function whose minimum is at $U = U^*$. Most forms for F that are typically used are compatible with formulating a polynomial optimization problem. Two common forms are $F = \|U(T) - U^*\|^2$ and $F = 1 - |\text{Tr}[U(T)^\dagger U^*]|/N$ [33]. Here, N is the dimension of the system and $\|M\| \equiv \sqrt{\text{Tr}[M^\dagger M]}$ is the Hilbert-Schmidt norm of an operator, M (also called the Frobenius or $L_{2,2}$ norm). The second form we will refer to as infidelity, since it is unity minus the fidelity of $U(T)$ with respect to U^* [34].

The second type of control problem is identical to Eq. (3), except that instead of using the evolution operator $U(t)$, the wave function formulation $|\psi(t)\rangle$ can be employed with the specified initial condition $|\psi(0)\rangle = |\psi_0\rangle$. In this case, rather than achieving a specific final state $|\psi^*\rangle$ at final time T , we might want to maximize some property of the final state, such as the expectation value of an operator $\hat{O}$. The objective function would then be chosen as $F(\hat{O}, |\psi^*\rangle) = -\langle \psi^* | \hat{O} | \psi^* \rangle$.

The third kind of control problem is the minimum-time control. In this case, we still want to achieve a target unitary, but we want to do so in the minimum time. The statement of this problem is

$$\begin{aligned} &\underset{E(t)}{\text{minimize}} && T \\ &\text{subject to} && T \geq 0, \quad F[U(T), U^*] \leq \varepsilon^2 \\ &&& \partial_t U(t) = A(t)U(t), \quad U(0) = \mathbb{1} \\ &&& A(t) = -i(H_0 + E(t)V). \end{aligned} \quad (4)$$

The fourth and final problem we include here is the same as Eq. (4), but with the unitaries replaced by quantum states. This minimum-time problem for state evolution is the quantum version of the brachistochrone problem [35,36].

While in the rest of the paper we focus on the single control, for completeness, let us state the generalization of problem (3) to the case of multiple controls

$$\begin{aligned} &\underset{\{E_k(t)\}_{k=1}^c}{\text{minimize}} && F[U(T), U^*] \\ &\text{subject to} && \partial_t U(t) = A(t)U(t), \quad U(0) = \mathbb{1} \\ &&& A(t) = -i\left(H_0 + \sum_{k=1}^c E_k(t)V_k\right). \end{aligned} \quad (5)$$

III. THE QCPOP METHOD

We obtain a polynomial formulation of our four control problems using three steps. First, we employ the celebrated Magnus expansion to $U(T)$ [37]:

$$U(T) = \lim_{n \rightarrow \infty} \exp[\Lambda_n], \quad \Lambda_n = \sum_{k=1}^n \Omega_k. \quad (6)$$

Using $B(s, t) \equiv [A(s), A(t)]$, $C(s, t, u) \equiv [A(s), B(t, u)]$, the first three terms of the Magnus expansion are

$$\Omega_1 = \int_0^T A(t) dt, \quad \Omega_2 = \frac{1}{2} \int_0^T \int_0^t B(t, s) dt ds, \quad (7)$$

$$\Omega_3 = \frac{1}{3!} \int_0^T \int_0^t \int_0^s [C(t, s, u) - C(u, t, s)] dt ds du. \quad (8)$$

The Magnus series converges so long as $\int_0^T \|A(t)\| dt < \pi$. An important feature of the Magnus expansion is that $\exp[\Lambda_n]$ is unitary for every n . Since the Frobenius norm is unitary invariant, the objective function in Eq. (3) can now be approximated as

$$F[U(T), U^*] = \|U(T) - U^*\|^2 \approx \|e^{\Lambda_n/2} - e^{-\Lambda_n/2} U^*\|^2. \quad (9)$$

Our second step is to use the Chebyshev polynomial expansion for the matrix exponent [38]:

$$e^{\alpha \Lambda_n} \approx J_0(\alpha) \mathbb{1} + 2 \sum_{k=1}^p J_k(\alpha) T_k, \quad (10)$$

where $J_k(x)$ is the Bessel function and T_k is the matrix-valued Chebyshev polynomial defined via the recurrence relation $T_{k+1} = 2\Lambda_n T_k + T_{k-1}$, with $T_0 = \mathbb{1}$ and $T_1 = \Lambda_n$. Note that the Chebyshev polynomials form an especially efficient basis for the interpolation due to how they distribute approximation errors: Instead of allowing large errors in some regions while achieving perfect accuracy elsewhere, the Chebyshev polynomials spread the error evenly across the entire interpolation interval. This “equioscillation” property means the maximum error is minimized, making them nearly optimal for polynomial approximation.

The final step is to use a series representation for the control function, $E(t)$, truncated to a fixed number of terms. Expansion of the control in *any* basis $\{f_k(t)\}_{k=0}^{m-1}$ will suffice, i.e.,

$$E(t) = \sum_{k=0}^{m-1} x_k f_k(t). \quad (11)$$

In our examples below, we choose a power series so that

$$E(t) = \sum_{k=0}^{m-1} x_k t^k. \quad (12)$$

Finding the optimal control function, therefore, involves finding the optimal values of the coefficients x_k . All four of our control problems will be polynomial optimization problems so long as $F[U(T), U^*]$ is a polynomial in x_k . Substituting the series representation for $E(t)$ into $A(t)$ and the latter into the integrals for the terms Ω_n in the Magnus expansion, we see that these are polynomials in x_k and so are Λ_n . Using the

Chebyshev expansion for the Magnus approximation to $U(T)$, Eq. (10), then gives us $U(T)$ as a polynomial in x_k . Finally, since $\|U(T) - U^*\|^2$ is a polynomial of the elements of $U(T)$, our four control problems are now polynomial optimization problems.

The chain of approximations (expansions) we have used to write our control problems as polynomial optimizations is characterized by the order at which we truncate each of the three expansions. Thus, an implementation of quantum control via polynomial optimization (QCPop) is characterized by the tuple (n, p, m) , where n is the number of terms used in the Magnus expansion, p is the number of terms in the Chebyshev expansion, and m is the number of terms in the control function. Using problems from Secs. IV and V A as benchmarks, we have tested the performance of QCPop for different values of n, p . We have empirically found that $n = 3$ and $p = 5$ (see Appendices B and C for details) are the optimal combinations, meaning that using larger values does not yield any improvement.

A. Modification for piecewise-constant control

If we want a control function that is piecewise constant, meaning that it consists of contiguous segments with a (possibly different) constant value on each, then the polynomial approximation simplifies. In this case, we do not need the Magnus expansion. If we have m segments in which the j th segment has duration t_j , then the unitary evolution is given by $U(T) = \prod_{j=1}^m U_j$, where

$$U_j = e^{-i(H_0 + E_j V)t_j}. \quad (13)$$

Here, the m values E_j are the parameters to be optimized. This time we use the Chebyshev expansion, Eq. (10), for each of the unitaries U_j . Substituting $U(T) = \prod_{j=1}^m U_j$ into the objective function $F[U(T), U^*]$ results in a polynomial optimization problem.

B. Polynomial optimization

As mentioned above, the advantage of formulating quantum control as polynomial optimization is the remarkably powerful global optimization methods for this task that have been developed in the last several years. Here, we use two of these methods, which are already available as libraries in Julia. The first is called the “moment sum-of-squares hierarchy method” (for which the Julia library is called TSSOS [39,40]) and the second is the “homotopy continuation method” (for which the Julia library is HomotopyContinuation [41]). In Appendix A, we give a pedagogical introduction to both the moment sum-of-squares hierarchy method and the homotopy continuation method. Empirically, we have found that homotopy continuation is better when minimizing infidelity ($F = 1 - |\text{Tr}[U(T)^\dagger U^*]|/N$) and the sum-of-squares method is better when minimizing the Hilbert-Schmidt distance ($F = \|U(T) - U^*\|^2$).

IV. EVALUATING THE PERFORMANCE OF QCPOP

We first apply QCPop to a special and simple example in which the global maxima and local traps of the objective




FIG. 1. The control problem of Pechen and Tannor [42], which has a trap at the point where the control variable, λ , vanishes. Using QCPOP, the trap poses no problem for the homotopy continuation method, which returns all the extrema (shown by the red crosses). (a) The solid blue line is the value of the objective function (the expectation value of O), being a symmetric function of the control variable, λ . The green (orange) dashed line is the objective function resulting from using the Taylor (Chebyshev) approximation in QCPOP (to tenth order). (b) A diagram of the control problem. The system has three levels, each with a different value of the observable, O . The control Hamiltonian, V , couples $|1\rangle$ to $|2\rangle$ and $|2\rangle$ to $|3\rangle$.

function are known. We do this to demonstrate without any ambiguity that the optimization algorithm locates global maxima regardless of the presence of traps. The example we use was originally introduced by Pechen and Tannor in Ref. [42], and involves a three-level system that is depicted in Fig. 1(b). The drift Hamiltonian, H_0 , has eigenstates $|1\rangle$, $|2\rangle$, and $|3\rangle$, and the system starts in state $|1\rangle$. The goal is to maximize the expectation value of an observable, O , after an evolution time, $T = 1$, where the values of O for each of the states $|1\rangle$, $|2\rangle$, and $|3\rangle$ are, respectively, 0.1, 0, and 1. Since the system starts in state $|1\rangle$, achieving the goal requires turning on the control Hamiltonian to transfer the system to state $|3\rangle$. The control Hamiltonian, V , does not couple $|1\rangle$ directly to $|2\rangle$, but couples $|1\rangle$ to $|3\rangle$ and $|3\rangle$ to $|2\rangle$ as shown in Fig. 1(b). (In particular, all elements of V are zero except $v_{13} = v_{23} = v_{31} = v_{32} = 1$.) The Hamiltonian is $H = H_0 + \lambda V$, and we fix λ to be constant over the full evolution from 0 to 1. We therefore have just a single control variable (the value of λ).

The reason why the above control problem has a local maximum (trap) at $\lambda = 0$ is quite simple. Since the control must take the system from $|1\rangle$ to $|3\rangle$ via $|2\rangle$, for small values of λ the system has a higher probability of being in $|2\rangle$ than in $|3\rangle$, thus *reducing* the expectation value of O below that of $|1\rangle$. The objective function thus initially decreases as $|\lambda|$ moves away from zero before it increases to the global optimum at $|3\rangle$, hence producing the local maxima at $\lambda = 0$. In Fig. 1, we plot the objective as a function of $|\lambda|$ as well as the approximation to the objective produced by QCPOP with the Chebyshev polynomial truncated at $p = 10$. The homotopy continuation optimization method applied to QCPOP finds *all* the extrema and is therefore unaffected by the presence of local maxima [43]. Gradient search methods, if started at or near zero, find only the local maximum.

We now test QCPOP on a set of nontrivial problems of the kind that quantum control is often employed to solve, and compare it directly to two state-of-the-art methods, GRAPE




FIG. 2. Violin plots of the infidelity between the target unitary, U^* , and the realized unitary, $\hat{U}$, defined in Sec. IV, for each of 1000 target unitaries and for three quantum control methods: GRAPE, CRAB, and QCPOP. Note that f_{PSU} denotes the fidelity that is defined as $f_{\text{PSU}} = |\text{Tr}(\hat{U}^\dagger U^*)|/3$, whereas $|1 - f_{\text{PSU}}|$ is the infidelity.

[18–21] and CRAB [22–24]. The set of problems we consider is that of generating a specified “target” unitary for a three-dimensional system with drift $H_0 = \text{diag}(0, 0.515916, 1)$ and control Hamiltonian with nonzero elements $\langle 1|V|2\rangle = \langle 2|V|1\rangle = 1/\sqrt{2}$ and $\langle 2|V|3\rangle = \langle 3|V|2\rangle = 1$. We restrict the control function, $\lambda(t)$, to be piecewise constant with a specified number of segments (degrees of freedom). We first generate a set of 1000 unitary operators to use as the “target” unitaries. We generate these by evolving the system under the Hamiltonian $H = H_0 + \lambda(t)V$ for a time $T = 0.5$ with randomly selected control functions $\lambda(t)$. We select the control functions by setting $\lambda = x_0 + x_1 t + x_2 t^2$ and choosing each of the three variables x_0 , x_1 , and x_2 independently and uniformly distributed on the interval $[-1, 1]$ (see Appendix B for further discussion of these illustrations).

We compare GRAPE, CRAB, and QCPOP on control problems defined by generating each of the 1000 unitaries in time $T = 0.5$ using two-segment piecewise-constant control [44] (see also Appendix B). We use QuTIP implementations of GRAPE and CRAB, and our own Julia implementation of QCPOP with the TSSOS polynomial optimization library. All three methods require approximately 1–2 s to find a solution. After obtaining the control field, we compute the evolution operator $\hat{U}$ by solving Eq. (2). Figure 2 shows the infidelity between the target unitary U^* and the synthesized $\hat{U}$. We test CRAB and GRAPE with four initial guess types: ZERO, SINE, LIN, and RND [45]. CRAB frequently becomes trapped in local minima and shows poor performance, while GRAPE performs well. However, QCPOP achieves even better results than GRAPE.

We now provide a true test of the ability of QCPOP to find global optima by applying it to a quantum control problem in which traps are so dense that it is virtually impossible to find solutions by using a gradient search method directly. This problem consists of transferring the state of one linear resonator to another by controlling the size of a linear coupling

between the two. This has applications to cooling [46] and information processing. The Hamiltonian is

$$H = \hbar\omega a^\dagger a + \hbar\lambda(t)(a + a^\dagger)(b + b^\dagger) + \hbar\omega b^\dagger b, \quad (14)$$

in which ω is the respective frequency of the two oscillators and a and b are their respective annihilation operators. Here, the “drift” Hamiltonian consists of the first and last term of H and the control Hamiltonian is the second term.

Given limits on the size (energy) of either or both of the drift and control Hamiltonians, for any nontrivial unitary, there is a minimum time in which it can be realized exactly. We may wish to find a control protocol that will realize the unitary approximately in a shorter time, because the resulting reduction in decoherence during the protocol can provide an overall increase in the achieved fidelity. Placing this restriction on the duration of the protocol is one situation in which a high density of traps can appear in the search space.

In Ref. [47], the authors show that for the linear state transfer problem a gradient search fails when the protocol duration T is less than $\tau/3$, where $\tau = 2\pi/\omega$ is the oscillators’ period. This indicates that a sharp transition occurs in the search landscape with traps becoming very dense below this point. Running 2000 searches geometrically equi-spaced between $T = \tau/100$ and $T = \tau$ achieved only a couple of solutions below $\tau/3$. Solutions can be found for all T in this region using a gradient search method by augmenting it with a technique called “path tracing,” which requires repeating a gradient search many times (on the order of a thousand were used in Ref. [47]) and, in general, it will only find local optima [47,48]. Results from the use of the above method show that there is a piecewise constant control of five segments that will perform an (imperfect) state transfer with an infidelity of 1% in a time $T = \tau/50$. Applying QCPOP to this problem for $T = \tau/50$ with a control function with only three segments, we find as expected that it obtains a solution in a single run. This solution also has an infidelity of 1% [49]. So long as the polynomial approximation used is sufficiently good, the polynomial optimization guarantees that this solution gives the best possible fidelity for $T = \tau/50$. QCPOP also shows us that there are no solutions for transfer in this time with a two-segment control function; three-segment control function is the simplest possible.

V. FURTHER APPLICATIONS

A. Hamiltonian identification

We have shown that the ability to write quantum control problems as polynomial optimization is very powerful. We now show that Hamiltonian identification is another problem for which polynomial optimization provides a solution. If we have a Hamiltonian $H = \sum_j c_j H_j + \lambda(t)V$ in which the parameters $\{c_j\}$ are uncertain, we can find them by: (1) choosing a control function $\lambda(t)$ and (2) letting the system evolve for a time T and measuring the result to determine the unitary that it generates, which we will call U^* . (3) The parameters $\{c_j\}$ can now be determined by minimizing $F(\{c_j\}) = \|U(T) - U^*\|^2$ over $\{c_j\}$, where $U(T)$ is the unitary generated by $H = \sum_j c_j H_j + \lambda(t)V$. As we have shown above, this minimization problem can be well approximated

by a polynomial optimization. In Appendix C, we demonstrate this method with a numerical example.

B. Computability of quantum control problems

As a final note, we observe that the ability to cast quantum control problems as polynomial optimization problems answers an open question in the theory of computing. In Ref. [50], it was shown that the problem of determining whether a given quantum control task is possible (that is, whether there exists a set of allowed controls that will solve the problem exactly) is not computable on a Turing machine. The proof was based on the fact that digitized quantum control can solve Diophantine equations, which are polynomial equations in integer-valued unknowns. The latter problem is not Turing decidable [51].

Our work here shows that if we shift the computing paradigm from that of the Turing machine, which involves discrete computations (finite precision), to that of the Blum-Shub-Smale (BSS) machine [52,53], which employs analog computations with an arbitrary precision, quantum optimal control problems become computable. Notice that the use of the BSS machine also circumvents issues [54–56] of bit complexity, which hinder claims of polynomial runtime in the dimensions of SDP relaxations on Turing machines. The run time on the BSS machine is thus determined by the rate of convergence of the Magnus expansion, which is well understood [57], and the rate of convergence of the moment/sums of-squares (SOS) hierarchy (or rather its sparsity-exploiting variants), whose understanding is only very recent [58,59]. We have shown here that such an approach is not only theoretically possible, but also a powerful computational tool.

VI. CONCLUSION

Here, we have shown that by formulating quantum optimal control as a polynomial optimization problem, control problems that previously required thousands of optimization runs to obtain even reasonable solutions can be solved in a single run. This tremendous increase in speed not only greatly reduces the numerical overhead of quantum control, but makes it possible to solve quantum control problems for large systems, for which it has hitherto been infeasible to perform thousands of optimizations. In addition, QCPOP is able to find the true optimal performance of quantum control by obtaining the global optimum as opposed to a sufficiently good local optimum. We have also shown that, more generally, polynomial optimization is a powerful tool and has applications in quantum technologies beyond quantum control. We expect that polynomial optimization will prove useful in other contexts in quantum engineering.

ACKNOWLEDGMENTS

D.I.B. was supported by the Army Research Office (ARO) (Grant No. W911NF-23-1-0288, program manager Dr. James Joseph) and by a cooperative agreement with DEVCOM ARL (Grant No. W911NF-21-2-0139). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARO or the U.S. Government. The

U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. L.B.G. has been supported by European Union's HORIZON–MSCA-2023-DN-JD programme under the Horizon Europe (HORIZON) Marie Skłodowska-Curie Actions, Grant Agreement No. 101120296 (TENORS). J.M. and G.K. have been supported by the Czech Science Foundation (Grant No. 23-07947S). G.K. was also supported by OP VVV Project No. CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics.”

This paper was prepared for information purposes, and is not a product of HSBC Bank Plc. or its affiliates. Neither HSBC Bank Plc. nor any of its affiliates make any explicit or implied representation or warranty and none of them accept any liability in connection with this paper, including, but limited to, the completeness, accuracy, reliability of information contained herein and the potential legal, compliance, tax, or accounting effects thereof. Copyright HSBC Group 2023.

DATA AVAILABILITY

The data that support the findings of this article are openly available [60].

APPENDIX A: POLYNOMIAL OPTIMIZATION

Polynomial optimization is a fundamental problem in mathematics that frequently appears in physics. Polynomial optimization problems (POPs) involve finding the global optimum of a polynomial function subject to polynomial constraints. Mathematically, a POP can be formulated as follows:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{subject to} \quad & g_i(x) \geq 0, \quad i = 1, \dots, m, \end{aligned} \quad (\text{A1})$$

where $f(x)$ and $g_i(x)$ are polynomial functions in $x = (x_1, \dots, x_n) \in \mathbb{R}^n$. This problem is generally computationally hard, and finding the global optimum was considered impossible until relatively recently. The fundamental papers [61,62], building upon the work of Putinar [63], introduced the moment/SOS hierarchy, which provides a sequence of increasingly tighter lower bounds on the optimal value of the original polynomial optimization problem. (See Ref. [64] for a survey.) This results in a hierarchy $\{P_0, \dots, P_k\}$ of semidefinite programming (SDP) problems, with P_0 denoting the standard moment relaxation, such that for large enough k one obtains the global optimum (although in certain problems this requires $k \rightarrow \infty$).

Let us formalize the above by considering the POP (A1) that we refer to as P . In order to construct the moment/SOS hierarchies, we need to define the basic semialgebraic set $K = \{x \in \mathbb{R}^n \mid g_j(x) \geq 0\}$, such that P can be reformulated as

$$P: \quad f^* = \min \{f(x) \mid x \in K\}. \quad (\text{A2})$$

Then, let us define the family of problems

$$P_k: \quad f^* = \sup \{\lambda \mid f(x) - \lambda \geq 0, \quad \forall x \in K\}. \quad (\text{A3})$$

Lassère [61] and Parrilo [62] provided a systematic algorithmic procedure to construct certificates of positivity on K for

$f(x) - \lambda$, which are then used to obtain the solution of Prob. (A3) in practice.

Concretely, Putinar's Positivstellensatz states the following: If a real polynomial f is strictly positive on K then it can be written as

$$f(x) = \sigma_0(x) + \sigma_1(x)g_1(x) \cdots + \sigma_m(x)g_m(x), \quad (\text{A4})$$

for any real vector x , where σ_j are real SOS polynomials. Testing whether $f(x)$ can be written as in Eq. (A4), signifying the desired result $f(x) > 0$ for any $x \in K$, amounts to solving an SDP. The dual of this SDP, which features prominently in the approach of Lasserre [61], is the problem whether given a real sequence $y = \{y_\alpha\}$ with $\alpha \in \mathbb{N}^n$ a multiindex, if there exists a probability measure μ on K such that

$$y_\alpha = \int_K x_1^{\alpha_1} \cdots x_n^{\alpha_n} d\mu, \quad (\text{A5})$$

for any α , then y represents a measure supported on K . Then, one introduces a linear function L_y that maps real polynomial functions to real numbers defined as

$$f = \sum_{\alpha \in \mathbb{N}^n} f_\alpha x^\alpha \mapsto L_y(f) = \sum_{\alpha \in \mathbb{N}^n} f_\alpha y_\alpha. \quad (\text{A6})$$

The sequence y above has a representing measure on K if and only if for every real polynomial h it holds that

$$L_y(h^2) \geq 0, \quad L_y(h^2 g_j) \geq 0, \quad (\text{A7})$$

for $j = 1, \dots, m$. If Eq. (A7) holds for all polynomials h with maximum degree d , certain $m+1$ moment and localizing matrices, with entries linear in $\{y_\alpha\}$, are positive semidefinite:

$$M_d(y) \geq 0, \quad M_d(g_j y) \geq 0, \quad (\text{A8})$$

for $j = 1, \dots, m$. The latter problem, Prob. (A8), can be modeled as a so-called “generalized moment problem”:

$$\inf_{\mu_i} \left\{ \sum_{i=1}^{\ell} \int_{K_i} f_i d\mu_i \mid \sum_{i=1}^{\ell} h_{ij} d\mu_i \geq b_j \right\}, \quad (\text{A9})$$

where $K_i \subset \mathbb{R}^{n_i}$ is defined as before.

From either point of view, the moment/SOS approach amounts to solving a hierarchy of (convex) semidefinite programs of increasing size:

$$f_d^* = \sup_{\lambda, \sigma_j} \left\{ \lambda \mid \sigma_0 + \sum_{j=1}^m \sigma_j g_j, \quad \text{with} \quad \deg(\sigma_j g_j) \leq 2d \right\}. \quad (\text{A10})$$

The sequence $\{f_d^*\}$, where $d \in \mathbb{N}$, is monotone and nondecreasing. When K is compact, one obtains the global optimum of Prob. (A3) as the limit

$$f^* = \lim_{d \rightarrow \infty} f_d^*, \quad (\text{A11})$$

where finite convergence is generic.

1. Exploiting sparsity

Polynomial optimization over a basic semialgebraic set is nondeterministic polynomial-time (NP)-hard [64]. While the moment/SOS hierarchy of relaxations, which was discussed above, establishes a framework for tackling such problems,

they remain hard and despite the popularity and power of the framework, certain limitations to the scalability of the moment/SOS hierarchy exist. Specifically, for the number n of variables of the POP at hand, the size of the matrices appearing in the d th step of the moment/SOS relaxation is of the order of $\binom{n+d}{n}$.

To address this challenge, several proposals have been put forth. Putinar [63] suggested removing monomials not appearing in the moment/SOS decomposition. Lasserre [65] suggested exploiting possible sparsity patterns observed in the polynomial optimization problem. In some sense, this amounts to reindexing the SDP matrices involved in the moment/SOS relaxation as follows: consider subsets $I_1, \dots, I_p \subseteq \{1, \dots, n\}$ of the initial input SDP variables, with cardinalities $\{|I_1|, \dots, |I_p|\}$, respectively. The sparse moment/SOS hierarchies [39,40,65], consider (quasi) block-diagonal SDP matrices where the size of block j is given by $|I_j|$, $1 \leq j \leq p$. With this reformulation of the original SDP, if the cardinalities above are small compared to the initial variables, that is, if $|I_j| \ll n$ for all j , applying the moment/SOS relaxations achieves significant computational savings, which in turn enables scalability. Under mild assumption, theoretical guarantees on the global convergence of the sparse moment/SOS relaxations exist.

2. Numerical algebraic geometry

An alternative approach leverages a large body of work within numerical algebraic geometry [66]. Let us consider the polynomial optimization problem (A1) and its Lagrangian:

$$L(x, \lambda) := f(x) + \sum_{j=1}^m \lambda_j g_j(x). \quad (\text{A12})$$

The first-order optimality conditions, known as the Lagrange system

$$F(x, \lambda) := \left\{ \frac{\partial L(x, \lambda)}{\partial x_i} \mid i = 1, \dots, n \right\}, \quad (\text{A13})$$

is a system of polynomial equations, whose solutions (“roots”) are sought in numerical algebraic geometry. If one could enumerate all roots of the algebraic system, one could find the global optimum by a linear scan.

To introduce the so-called homotopy-continuation approach [67–69] within numerical algebraic geometry, one can consider the so-called straight-line homotopy. There, we see F as the “target system,” pick a suitable “start system” G that has the same or higher number of roots, known to us, and consider a homotopy:

$$H(x, \lambda, t) = tF(x, \lambda) + (1 - t)G(x, \lambda). \quad (\text{A14})$$

As one traces the path $t \in [0, 1]$ for all roots of the start system, one may be able to arrive at all the roots of the target system. Traditionally, the method is presented in terms of complex-valued polynomials, but one can clearly filter for real-valued roots. While this may seem straightforward, several challenges have made numerical algebraic geometry into an active area of research over the past three decades.

The first challenge is related to the scalability and exploitation of a sparsity in the system. To some extent, this

can be addressed by the choice of the start system. The theorem of Bernstein [70–72], also known as the Bernstein-Khovanskii-Kushnirenko (BKK) theorem, provides an upper bound on the number of isolated solutions a sparse polynomial system, which is exact, generically. This underlies the polyhedral homotopy algorithm of Huber and Sturmfels [73]. HomotopyContinuation utilizes the BKK theorem to bound the number of roots of the target system and pick the matching start system.

The second challenge is related to numerical stability of tracing the paths (“path-tracking”) from the start system to the target system. As a path approaches a singular endpoint, one needs to modify the continuation process to improve numerical stability. Morgan *et al.* [74–76] coined the term “end-game” for approaches based on the uniformization theorem for germs of one-dimensional analytic sets, which have been subsequently extended by numerous other researchers [77, e.g.]. We employ the end-game based on power series [76] and Cauchy [75] endgame, as implemented in HomotopyContinuation [41].

The third challenge is related to the *post hoc* certification of the correctness and completeness of the set of roots. Leading approaches in this direction [78,79] are based on the point-estimation theory of Smale ([80], also known as the α -theory) and interval arithmetics [81], respectively. HomotopyContinuation implements a technique of [79] based on interval arithmetics.

APPENDIX B: ILLUSTRATIONS

In this section, we provide numerical illustrations for the polynomial formulation (3) of the terminal quantum control problem of synthesizing a desired unitary gate. The code used in the analysis is available as a Julia Jupyter notebook [82]. In particular, we utilize the drift and control Hamiltonians of the form

$$H_0 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0.515916 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad V = \begin{pmatrix} 0 & 0.707107 & 0 \\ 0.707107 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad (\text{B1})$$

which are scaled versions of the Hamiltonians for IBM Q 2-qubit devices [83]. The analysis is started by randomly drawing components of the vector $\mathbf{x}^*$ ($m = 3$) from a uniform distribution on the interval $[-1, +1]$. In total, 1000 samples are generated. Then, for each sample, we perform the following steps: Setting $\mathbf{x} = \mathbf{x}^*$ in Eq. (12), we obtain the control field used to get the target unitary U^* by solving the Schrödinger equation (1). The terminal time is chosen at $T = 0.5$. Hence, we know that $\mathbf{x}^*$ is an exact solution to the terminal quantum control problem (3). Figure 3(a) shows that the inequality $\int_0^T \|A(t)\| dt < \pi$ is satisfied [Eq. (8)], therefore justifying the usage of the truncated Magnus series to approximate the evolution operator. The polynomial formulation (9) with $n = 3$ and $p = 5$ is then solved using the TSSOS Julia package [39,40]. This library extracts the minimizer $\hat{\mathbf{x}}$, i.e., the value of $\mathbf{x}$ yielding the global minimum




FIG. 3. Histogram plots for Magnus series convergence tests [Eq. (8)] (a) when exact solutions $\mathbf{x}^*$ for quantum control problem (3) are used and (b) when global minimizers $\hat{\mathbf{x}}$ for Eq. (9), obtained by the TSSOS Julia package, are employed.

estimate for Eq. (9). It is noteworthy that the obtained global minimum estimate via TSSOS is larger than the value of the objective function (9) at $\mathbf{x}^*$ (see Fig. 4) as expected from




FIG. 4. Comparing the global minima for (9), obtained by the TSSOS Julia package, with the value of the polynomial objective function (9) when an exact solution $\mathbf{x}^*$ for quantum control problem (3) is used.




FIG. 5. A relation between the quantum control problem (3) and its approximate polynomial formulation (9). $\hat{U}$ is the evolution operator obtained via polynomial optimization (9).

the monotone convergence property of the hierarchy of SDP relaxations (see Sec. A). Hence, the polynomial optimization yields $\hat{\mathbf{x}} \neq \mathbf{x}^*$. Notice also the well-known fact that quantum control problems typically have nonunique solutions [12,14]. Figure 3(b) shows that the Magnus expansion for $\hat{\mathbf{x}}$ converges. The control field [Eq. (12)] obtained from $\hat{\mathbf{x}}$ generates the evolution operator $\hat{U}$ that is supposed to be very close to the target U^* . Figure 5 indeed confirms that.

We conclude that the approximate polynomial formulation (9) provides a viable alternative for the original quantum control problem (3).

APPENDIX C: HAMILTONIAN IDENTIFICATION

Let us now consider an easy starting point in this direction. For convenience, assume that we know the drift Hamiltonian H_0 (B1), but we do not know the values of nonzero coupling constants in the interaction Hamiltonian V (B1). Hence, V is




FIG. 6. An illustration of the effectiveness of polynomial optimization formulation (C2) for the problem of Hamiltonian identification. Distribution of the error between the exact $\mathbf{z}$ and recovered $\hat{\mathbf{z}}$ nonzero element of the interaction Hamiltonian (C1).

assumed to be of the form

$$V = \begin{pmatrix} 0 & z_1 & 0 \\ z_1 & 0 & z_2 \\ 0 & z_2 & 0 \end{pmatrix}, \quad (\text{C1})$$

and our goal is to find the unknown coupling constants $\mathbf{z} = (z_1, z_2)$. We note that this is a physically important formulation of the problem of Hamiltonian parameter identification. The diagonal elements correspond to eigenenergies of the nondriven quantum system, which are easy to measure, and thence can be assumed to be known accurately. The positions of zero elements in the interaction Hamiltonian V (B1) follow from the symmetry of physical systems, e.g., via the Wigner–Eckart theorem.

To determine the unknown coupling constants $\mathbf{z}$, we drive the quantum system with a known control signal $\mathbf{x}$, repeatedly, and perform the tomography at the end of the evolution to determine the synthesized unitary U^* . Experimentally, we show that it is often sufficient to know only a single pair $\mathbf{x} \rightarrow U^*$, in order to very accurately recover the unknown couplings $\mathbf{z}$ by solving the following unconstrained polynomial optimization problem:

$$\underset{\mathbf{z}}{\text{minimize}} \|e^{\Lambda_n/2} - e^{-\Lambda_n/2} U^*\|^2. \quad (\text{C2})$$

The formulation (C2) is a minor modification of the polynomial optimization of quantum control (9). The difference lies in the fact that in the latter, the control signal $\mathbf{x}$ is to be found, using a known Hamiltonian (i.e., known $\mathbf{z}$). In the former (C2), we wish to find parameters of the interaction Hamiltonians, using the known control signal.

To numerically illustrate the effectiveness of formulation (C2), we perform an analysis similar to that of Sec. B. We start by randomly drawing components of the vector $\mathbf{x}$ ($m = 3$) from a uniform distribution on the interval $[-1, +1]$. In total, 1000 samples are generated. Then, for each sample, we perform the following steps: We obtain the terminal unitary U^* at $T = 0.5$ by solving the Schrödinger equation (1) using the exact Hamiltonian (B1). The polynomial minimization (C2) with $n = 3$ and $p = 5$ is solved via the TSSOS Julia package [39,40], yielding the estimate $\hat{\mathbf{z}}$ of the global minimizer for $\mathbf{z}$. The obtained estimate is then refined by performing a local minimization. The code used for this analysis can be found in the accompanying Julia Jupyter notebook [84]. The results are shown in Fig. 6, where the distribution of the errors between the exact couplings $\mathbf{z}$ and recovered couplings $\hat{\mathbf{z}}$ is shown for all samples. As can be readily seen, the Hamiltonian is recovered with a very high accuracy even from a single pair $\mathbf{x} \rightarrow U^*$.

- [1] J. P. Palao and R. Kosloff, Quantum computing by an optimal control algorithm for unitary transformations, *Phys. Rev. Lett.* **89**, 188301 (2002).
- [2] M. A. Nielsen, M. R. Dowling, M. Gu, and A. C. Doherty, Quantum computation as geometry, *Science* **311**, 1133 (2006).
- [3] T. F. Roque, A. A. Clerk, and H. Ribeiro, Engineering fast high-fidelity quantum operations with constrained interactions, *npj Quantum Inf.* **7**, 28 (2021).
- [4] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly nonlinear qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [5] J. Atalaya, S. Zhang, M. Y. Niu, A. Babakhani, H. C. H. Chan, J. M. Epstein, and K. B. Whaley, Continuous quantum error correction for evolution under time-dependent Hamiltonians, *Phys. Rev. A* **103**, 042406 (2021).
- [6] E. v. d. Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [7] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe, *EPJ Quantum Technol.* **9**, 19 (2022).
- [8] O. V. Morzhin and A. N. Pechen, Krotov method for optimal control of closed quantum systems, *Russ. Math. Surv.* **74**, 851 (2019).
- [9] J. Zhang, Y.-x. Liu, R.-B. Wu, K. Jacobs, and F. Nori, Quantum feedback: Theory, experiments, and applications, *Phys. Rep.* **679**, 1 (2017).
- [10] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen *et al.*, Training Schrödinger’s cat: Quantum optimal control, *Eur. Phys. J. D* **69**, 279 (2015).
- [11] I. Petersen and D. Dong, Quantum control theory and applications: A survey, *IET Control Theory Appl.* **4**, 2651 (2010).
- [12] C. Brif, R. Chakrabarti, and H. Rabitz, Control of quantum phenomena: Past, present and future, *New J. Phys.* **12**, 075008 (2010).
- [13] D. d’Alessandro, *Introduction to Quantum Control and Dynamics* (CRC Press, New York, 2007).
- [14] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Quantum optimally controlled transition landscapes, *Science* **303**, 1998 (2004).
- [15] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, From pulses to circuits and back again: A quantum optimal control perspective on variational quantum algorithms, *PRX Quantum* **2**, 010101 (2021).
- [16] X. Ge, R.-B. Wu, and H. Rabitz, The optimization landscape of hybrid quantum–classical algorithms: From quantum control to NISQ applications, *Annu. Rev. Control* **54**, 314 (2022).
- [17] J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer, New York, NY, 2006).
- [18] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [19] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [20] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov, Second order gradient ascent pulse engineering, *J. Magn. Reson.* **212**, 412 (2011).

- [21] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, *Nat. Commun.* **8**, 94 (2017).
- [22] P. Doria, T. Calarco, and S. Montangero, Optimal control technique for many-body quantum dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [23] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [24] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape, *Phys. Rev. A* **92**, 062343 (2015).
- [25] V. F. Krotov and I. N. Feldman, An iterative method for solving problems of optimal control, *Eng. Cybern.* **21**, 123 (1983).
- [26] M. H. Goerz and K. Jacobs, Efficient optimization of state preparation in quantum networks using quantum trajectories, *Quantum Sci. Technol.* **3**, 045005 (2018).
- [27] R. S. Judson and H. Rabitz, Teaching lasers to control molecules, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [28] M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement learning in different phases of quantum control, *Phys. Rev. X* **8**, 031086 (2018).
- [29] P. Palittapongarnpim, P. Wittek, E. Zahedinejad, S. Vedaie, and B. C. Sanders, Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics, *Neurocomput.* **268**, 116 (2017).
- [30] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 33 (2019).
- [31] M. Dalgaard, F. Motzoi, J. J. Sørensen, and J. Sherson, Global optimization of quantum dynamics with AlphaZero deep exploration, *npj Quantum Inf.* **6**, 6 (2020).
- [32] R. Mukherjee, H. Xie, and F. Mintert, Bayesian optimal control of Greenberger-Horne-Zeilinger states in Rydberg lattices, *Phys. Rev. Lett.* **125**, 203603 (2020).
- [33] D. Basilewitsch, C. P. Koch, and D. M. Reich, Quantum optimal control for mixed state squeezing in cavity optomechanics, *Adv. Quantum Technol.* **2**, 1800110 (2019).
- [34] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary Edition (Cambridge University Press, New York, 2011).
- [35] X. Wang, M. Allegra, K. Jacobs, S. Lloyd, C. Lupo, and M. Mohseni, Quantum brachistochrone curves as geodesics: Obtaining accurate minimum-time protocols for the control of quantum systems, *Phys. Rev. Lett.* **114**, 170501 (2015).
- [36] A. R. Kuzmak and V. M. Tkachuk, The quantum brachistochrone problem for an arbitrary spin in a magnetic field, *Phys. Lett. A* **379**, 1233 (2015).
- [37] W. Magnus, On the exponential solution of differential equations for a linear operator, *Commun. Pure Appl. Math.* **7**, 649 (1954).
- [38] H. Tal-Ezer and R. Kosloff, An accurate and efficient scheme for propagating the time dependent Schrödinger equation, *J. Chem. Phys.* **81**, 3967 (1984).
- [39] J. Wang, V. Magron, and J.-B. Lasserre, TSSOS: A moment-SOS hierarchy that exploits term sparsity, *SIAM J. Optim.* **31**, 30 (2021).
- [40] J. Wang, V. Magron, and J.-B. Lasserre, Chordal-TSSOS: A moment-SOS hierarchy that exploits term sparsity with chordal extension, *SIAM J. Optim.* **31**, 114 (2021).
- [41] P. Breiding and S. Timme, HomotopyContinuation.Jl: A package for homotopy continuation in Julia, in *International Congress on Mathematical Software* (Springer, South Bend, 2018), pp. 458–465.
- [42] A. N. Pechen and D. J. Tannor, Are there traps in quantum control landscapes? *Phys. Rev. Lett.* **106**, 120402 (2011).
- [43] Details are given in the Jupyter notebook, https://github.com/dibondar/QControlPolyOpt/blob/master/Three_level_system_with_trap--piecewise_constant.ipynb.
- [44] Optimization details are provided in the Jupyter notebook at https://github.com/dibondar/QControlPolyOpt/blob/master/QuantumControlPolyOpt_PiecewiseConst.ipynb.
- [45] See https://qutip-qtrl.readthedocs.io/en/latest/apidoc/qutip_qtrl.pulseoptim.html.
- [46] X. Wang, S. Vinjanampathy, F. W. Strauch, and K. Jacobs, Ultraefficient cooling of resonators: Beating sideband cooling with quantum control, *Phys. Rev. Lett.* **107**, 177204 (2011).
- [47] K. Jacobs, R. Wu, X. Wang, S. Ashhab, Q.-M. Chen, and H. Rabitz, Fast quantum communication in linear networks, *Europhys. Lett.* **114**, 40007 (2016).
- [48] K. W. Moore Tibbetts, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T.-S. Ho, R.-B. Wu, and H. Rabitz, Exploring the tradeoff between fidelity and time optimal control of quantum unitary transformations, *Phys. Rev. A* **86**, 062309 (2012).
- [49] Details are given in the Jupyter notebook, https://github.com/dibondar/QControlPolyOpt/blob/master/State_transfer_in_linear%20networks.ipynb.
- [50] D. I. Bondar and A. N. Pechen, Uncomputability and complexity of quantum control, *Sci. Rep.* **10**, 1195 (2020).
- [51] Y. V. Matiyasevich, What can and cannot be done with diophantine problems, *Proc. Steklov Inst. Math.* **275**, 118 (2011).
- [52] L. Blum, M. Shub, and S. Smale, On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines, *Bull. Amer. Math. Soc.* **21**, 1 (1989).
- [53] L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and Real Computation* (Springer Science & Business Media, New York, 1998).
- [54] S. Gribling, S. Polak, and L. Slot, A note on the computational complexity of the moment-SOS hierarchy for polynomial optimization, in *Proceedings of the 2023 International Symposium on Symbolic and Algebraic Computation* (ACM Digital Library, Tromsø, 2023), pp. 280–288.
- [55] P. Raghavendra and B. Weitz, On the bit complexity of sum-of-squares proofs, in *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, New York, 2017).
- [56] R. O'Donnell, SOS is not obviously automatizable, even approximately, in *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Wadern, 2017).
- [57] S. Blanes, F. Casas, J.-A. Oteo, and J. Ros, The Magnus expansion and some of its applications, *Phys. Rep.* **470**, 151 (2009).
- [58] C. Schlosser and M. Tacchi, Specialized effective positivstellensätze for improved convergence rates of the moment-SOS hierarchy, *IEEE Control Syst. Lett.* **8**, 2319 (2024).

- [59] M. Korda, V. Magron, and R. Rios-Zertuche, Convergence rates for sums-of-squares hierarchies with correlative sparsity, *Math. Program.* **209**, 435 (2025).
- [60] D. I. Bondar, QControlPolyOpt, <https://github.com/dibondar/QControlPolyOpt>.
- [61] J. B. Lasserre, Global optimization with polynomials and the problem of moments, *SIAM J. Optim.* **11**, 796 (2001).
- [62] P. A. Parrilo, Semidefinite programming relaxations for semialgebraic problems, *Math. Program.* **96**, 293 (2003).
- [63] M. Putinar, Positive polynomials on compact semi-algebraic sets, *Indiana Univ. Math. J.* **42**, 969 (1993).
- [64] M. Laurent, Sums of squares, moment matrices and optimization over polynomials, in *Emerging Applications of Algebraic Geometry*, edited by M. Putinar and S. Sullivant (Springer, New York, NY, 2009), pp. 157–270.
- [65] J. B. Lasserre, Convergent SDP-relaxations in polynomial optimization with sparsity, *SIAM J. Optim.* **17**, 822 (2006).
- [66] J. Lindberg, L. Monin, and K. Rose, A polyhedral homotopy algorithm for computing critical points of polynomial programs, [arXiv:2302.04117](https://arxiv.org/abs/2302.04117).
- [67] A. J. Sommese and C. W. Wampler, *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science* (World Scientific, Singapore, 2005).
- [68] E. L. Allgower and K. Georg, Numerical path following, *Handbook Numer. Anal.* **5**, 3 (1997).
- [69] P. Bürgisser and F. Cucker, *Condition: The Geometry of Numerical Algorithms* (Springer Science & Business Media, New York, 2013), Vol. 349.
- [70] L. Accardi, The noncommutative Markovian property, *Funct. Anal. Appl.* **9**, 1 (1975).
- [71] A. G. Khovanskii, Newton polyhedra and the genus of complete intersections, *Funkts. Anal. Pril.* **12**, 38 (1978).
- [72] A. G. Kushnirenko, Newton polytopes and the Bezout theorem, *Funkts. Anal. Pril.* **10**, 233 (1976).
- [73] B. Huber and B. Sturmfels, A polyhedral method for solving sparse polynomial systems, *Math. Comput.* **64**, 1541 (1995).
- [74] A. P. Morgan, A. J. Sommese, and C. W. Wampler, Computing singular solutions to nonlinear analytic systems, *Numer. Math.* **58**, 669 (1990).
- [75] A. P. Morgan, A. J. Sommese, and C. W. Wampler, Computing singular solutions to polynomial systems, *Adv. Appl. Math.* **13**, 305 (1992).
- [76] A. P. Morgan, A. J. Sommese, and C. W. Wampler, A power series method for computing singular solutions to nonlinear analytic systems, *Numer. Math.* **63**, 391 (1992).
- [77] B. Huber and J. Verschelde, Polyhedral end games for polynomial continuation, *Numer. Alg.* **18**, 91 (1998).
- [78] J. D. Hauenstein and F. Sottile, Algorithm 921: alphaCertified: Certifying solutions to polynomial systems, *ACM Trans. Math. Softw.* **38**, 1 (2012).
- [79] P. Breiding, K. Rose, and S. Timme, Certifying zeros of polynomial systems using interval arithmetic, *ACM Trans. Math. Software* **49**, 11 (2020).
- [80] S. Smale, Newton’s method estimates from data at one point, in *The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics: Proceedings of a Symposium Held in Honor of Gail S. Young at the University of Wyoming, August 8–10, 1985. Sponsored by the Sloan Foundation, the National Science Foundation, and Air Force Office of Scientific Research* (Springer, New York, 1986), pp. 185–196.
- [81] R. E. Moore, A test for existence of solutions to nonlinear systems, *SIAM J. Numer. Anal.* **14**, 611 (1977).
- [82] <https://github.com/dibondar/QControlPolyOpt/blob/master/QuantumControlPolyOpt.ipynb>.
- [83] https://github.com/q-optimize/c3/blob/master/examples/two_qubits.ipynb.
- [84] <https://github.com/dibondar/QControlPolyOpt/blob/master/HamiltonianParameterIdentification.ipynb>.