

Controlling Qubit Networks in Polynomial Time

Christian Arenz and Herschel Rabitz

Frick Laboratory, Princeton University, Princeton, New Jersey 08544, USA

 (Received 11 October 2017; published 30 May 2018)

Future quantum devices often rely on favorable scaling with respect to the number of system components. To achieve desirable scaling, it is therefore crucial to implement unitary transformations in a time that scales at most polynomial in the number of qubits. We develop an upper bound for the minimum time required to implement a unitary transformation on a generic qubit network in which each of the qubits is subject to local time dependent controls. Based on the developed upper bound, the set of gates is characterized that can be implemented polynomially in time. Furthermore, we show how qubit systems can be concatenated through controllable two body interactions, making it possible to implement the gate set efficiently on the combined system. Finally, a system is identified for which the gate set can be implemented with fewer controls.

DOI: 10.1103/PhysRevLett.120.220503

Achieving accurate control and scalability lie at the heart of every functioning quantum information processing device. Thus, a vital goal is to design algorithms that can be implemented efficiently. In particular, in the gate model of quantum information processing an efficient algorithm should scale polynomially in the number of gates used to carry out the computation. Through a *universal gate set* every algorithm described by a unitary transformation can be implemented up to some degree of accuracy. However, a simple counting argument shows that most of the unitary transformations cannot be implemented efficiently [1]. Quantum control theory allows for implementing the final unitary transformation directly through optimized classical control fields [2–4]. This has the advantage that, if the procedure can be done efficiently, there is no need for constructing gate sequences. Instead, optimization algorithms such as a gradient based search [5,6], learning control [7,8], or genetic algorithms [9,10], may be used to precalculate or learn the classical control fields that implement the desired unitary transformation. In fact, it has been shown that the complexity of both approaches, i.e., calculating control pulses and designing gate sequences is the same [11,12].

Similar to a universal gate set, for a fully controllable system every unitary transformation contained in the special unitary group $SU(2^n)$ is reachable through switchable controls. In order to implement a goal unitary gate U_g efficiently, it is crucial that the length of the control pulses, henceforth referred to as the minimum gate time T , scales at most polynomially in the number of qubits. Unfortunately, the determination of the minimum gate time has remained a major technical challenge to overcome for moving the field towards practical applications. In this Letter we make a step towards solving this problem by developing an upper bound for the minimum gate time under the assumption

that sufficient control resources are available. As illustrated in Fig. 1, this allows for determining the set of gates that *provably* can be implemented efficiently.

Although substantial progress has recently been made by characterizing graphs that can be controlled efficiently [13], the characterization of the set of gates that can be reached in polynomial time and the corresponding number of controls required is still unknown. Moreover, it remains challenging to identify physical models that obey the criteria developed in Ref. [13].

The main quantitative result of this Letter is the development of the upper bound

$$T(\mathbf{a}) \leq \frac{l(\mathbf{a})}{J} \left(\|\mathbf{a}\|_\infty + \frac{\pi l(\mathbf{a})(l(\mathbf{a}) - 1)(n - 2) \|\mathbf{a}\|_\infty^2}{2\sqrt{2}\epsilon} \right), \quad (1)$$

for the minimum gate time to implement a goal unitary transformation $U_g(\mathbf{a})$ up to some error ϵ for a generic n -qubit




FIG. 1. Illustration of one of the main results: (a) for a generic qubit network (3) in which each qubit is subject to two local controls (4) (gray arrows), (b) the set of gates $\mathcal{R}_{\text{poly}}$ (white area) that (provably) can be implemented in a time that scales at most polynomially in the number of qubits is characterized [see Eq. (2)]. The dark gray area represents the set of gates that can be reached with time optimal methods in polynomial time.

graph (3) in which each of the qubits is subject to two local controls (4). As illustrated in Fig. 1(a), the qubits (black circles) interact via two body interactions (solid lines) where J is the smallest coupling constant present in the graph. The generator of the goal unitary transformation is characterized by $l(\mathbf{a})$ real parameters summarized in the vector \mathbf{a} with $\|\mathbf{a}\|_\infty$ being the vector infinity norm, i.e., the largest parameter.

One way to obtain an upper bound on T is to find a specific way to implement a generic unitary transformation and upper bound the corresponding time. The procedure used here can be summarized by the following steps, with details found below and in the Supplemental Material [14]: (1) Because of the assumption that each qubit is subject to two unconstrained orthogonal controls, a decoupling sequence allows us to select arbitrary two body interactions instantaneously [see Eq. (5)]. (2) A sequence formed by such two qubit unitaries [see Eq. (6)] allows for creating unitary operations which are generated by k -body interaction terms in a time that scales linearly in k [see Eq. (7)]. (3) Finally, unitary transformations that are generated by linear combinations of l , k -body interaction terms can be created (up to an error ϵ) using a Trotter sequence [see Eq. (8)].

As illustrated in Fig. 1(b), the bound (1) allows us to conclude that the gate set,

$$\mathcal{R}_{\text{poly}} = \{U(\mathbf{a}) \in \text{SU}(2^n) | l(\mathbf{a}), \|\mathbf{a}\|_\infty \leq \mathcal{O}(\text{poly}(n))\}, \quad (2)$$

can be implemented on a qubit graph in which each qubit is subject to two local controls in a time that scales at most polynomially in the number of qubits n , and, moreover, enables for characterizing the Hamiltonians that can be simulated efficiently. We, furthermore, show that for a specific system the gate set $\mathcal{R}_{\text{poly}}$ can be implemented with less controls. Moreover, a strategy is presented for efficiently scaling the system by controlling two body interactions (schematically represented in Fig. 2). We remark that our findings are a proof of feasibility rather than a strategy to implement gates in a time optimal manner, which remains a practical challenge.

A quantum control problem can be expressed as follows. The system of interest is described by a time dependent Hamiltonian of the form $H(t) = H_0 + H_c(t)$, where H_0 is referred to as the drift Hamiltonian and the controls enter in $H_c(t)$ via time dependent functions. The aim of quantum control is then to steer the system towards a desired target by shaping the control functions. Here we are interested in implementing a generic target unitary transformation U_g on an n qubit system. The first question to consider is whether every U_g can be reached, i.e., whether the system is fully controllable. When control enters in a bilinear way in $H_c(t)$ [17], known as the Lie rank criterion [2], the system is fully controllable if and only if the controls and drift generate the full algebra (see, e.g., Refs. [18–24] and references therein

for examples). More formally, if the system is fully controllable there exist controls that allow implementing every $U_g = \exp(\Theta)$ with $\Theta \in \mathfrak{su}(2^n)$ up to arbitrarily high precision in finite time. Throughout this work the special unitary algebra $\mathfrak{su}(2^n)$ is expressed in terms of the Pauli operator basis $\{B_i\}_{i=1}^{2^n-1}$, in which each B_i corresponds to a string of Pauli operators. Every $\Theta \in \mathfrak{su}(2^n)$ can be written as $\Theta(\mathbf{a}) = \sum_{i=1}^{l(\mathbf{a})} a_i B_i$, where the real coefficients are summarized in the vector \mathbf{a} and we denote by $l(\mathbf{a}) \leq 2^n - 1$ the number of its nonzero elements. Except for low dimensional systems [25–30], the minimum gate time $T(\mathbf{a})$ needed to implement $U_g(\mathbf{a})$ up to some accuracy is not known.

Consider a connected graph $G(V, E)$ where the vertices V and edges E represent qubits and two body interactions, respectively. The most general form of such an n -qubit graph is described by the drift Hamiltonian

$$H_0 = \sum_{\substack{i \in V \\ \alpha \in \{x,y,z\}}} \omega_\alpha^{(i)} \sigma_\alpha^{(i)} + \sum_{\substack{(i,j) \in E \\ \alpha, \beta \in \{x,y,z\}}} g_{\alpha\beta}^{(i,j)} \sigma_\alpha^{(i)} \sigma_\beta^{(j)}, \quad (3)$$

where $\omega_\alpha^{(i)}$, $g_{\alpha\beta}^{(i,j)}$ are energy splittings and coupling constants, respectively. Here the notation refers to $\sigma_\alpha^{(j)} \equiv 1 \otimes \sigma_\alpha \otimes \mathbb{1}$ where σ_α with $\alpha \in \{x, y, z\}$ are Pauli spin operators. That is, $\sigma_\alpha^{(i)}$ acts only nontrivially on the i th qubit. We assume that each qubit is subject to two local controls $\{\sigma_x^{(i)}, \sigma_y^{(i)}\}$ such that

$$H_c(t) = \sum_{i \in V} [f_i(t) \sigma_x^{(i)} + h_i(t) \sigma_y^{(i)}], \quad (4)$$

where $f_i(t)$, $h_i(t)$ are the corresponding control fields that are assumed to be unconstrained. This is a typical assumption in the context of quantum control theory and dynamical decoupling and its crucial for the development of the upper bound below (see the note [31]). Before relaxing the assumption of two orthogonal controls on each qubit, we first describe in more detail how the upper bound on $T(\mathbf{a})$ can be derived for this control system. Further details of the derivation can be found in Ref. [14].

The analysis starts with the form of $H_c(t)$, allowing for having two orthogonal controls on each qubit, such that every single qubit gate can be implemented instantaneously [26,27]; moreover, the system is fully controllable [23]. Using a decoupling sequence [32,33] formed by the controls permits instantaneously selecting arbitrary two body interaction terms [34]. Thus, we can implement every unitary transformation

$$U_{\alpha\beta}^{(i,j)}(k) = e^{\pm ik \sigma_\alpha^{(i)} \sigma_\beta^{(j)}}, \quad k \in \mathbb{R}_+, \quad \alpha, \beta \in \{x, y, z\}, \quad (5)$$

in a time $t = k/g_{\alpha\beta}^{(i,j)}$ [14], noting that each Pauli operator can be rotated instantaneously to a generic Pauli operator using local operations. The following analysis makes use of the fact that every basis operator B_i can be created by a nested commutator of the form $[\dots, [S_1, [S_2, S_3]]]$, where $S_k \in \mathcal{S} = \{i\sigma_\alpha^{(i)}\sigma_\beta^{(j)}\}$, which are referred to as a generating set and we refer to the length of the nested commutator as the *depth* D with $[S_1, S_2]$ being a commutator of $D = 1$. Using a sequence of the form

$$U_{x,z}^{(2,3)\dagger}(\pi/4)U_{z,y}^{(1,2)}(k)U_{x,z}^{(2,3)}(\pi/4) = \exp(ik\sigma_z^{(1)}\sigma_z^{(2)}\sigma_z^{(3)}), \quad (6)$$

and introducing the smallest coupling constant $J = \min_{i,j,\alpha,\beta}\{g_{\alpha\beta}^{(i,j)}\}$ present in H_0 , an upper bound for the time Δt to create a unitary operation generated by a commutator of depth 1, in Eq. (6) a 3-body interaction term, can be found, i.e., $\Delta t \leq (\pi/2J)$ [14].

There are other sequences that allow for increasing or decreasing the length of a Pauli string [35]. However, due to the form of the construction (6), a unitary operation generated by a nested commutator of depth D will then take at most time $D\Delta t$. Thus, the time $\tau(a_i B_i)$ to implement a unitary operation $U_g = \exp(a_i B_i)$ is upper bounded by

$$\tau(a_i B_i) \leq \frac{1}{J} \left(D(B_i) \frac{\pi}{2} + |a_i| \right), \quad (7)$$

which is compared with known results in the online material [14]. Through a Trotter-Suzuki sequence [36] we can further upper bound the time it takes to generate a unitary operator generated by linear combinations of the basis operators up to an error ϵ . We find

$$T(\mathbf{a}) \leq \frac{1}{J} \left(\|\mathbf{a}\|_1 + \frac{\pi K(\mathbf{a}) \sum_{i=1}^{l(\mathbf{a})} D(B_i)}{4\sqrt{2}\epsilon} \right), \quad (8)$$

with $\|\cdot\|_1$ being the vector-1 norm, $K(\mathbf{a}) = (1/\sqrt{2^n}) \sum_{j>k} |a_j a_k| \| [B_j, B_k] \|$, and $\|\cdot\|$ is the Hilbert-Schmidt norm. The scaling in ϵ , explicitly given in Ref. [14], can be traced back to the use of the Suzuki-Trotter series, and the scaling can be improved using more sophisticated sequences [37]. An algorithm finding the “shortest” path, possibly weighted by the coupling constants, to create a B_i would produce the tightest bound. However, it takes a nested commutator of depth $(n-2)$ to create a basis operator that contains n Pauli operators $\sigma_\alpha^{(1)}\sigma_\beta^{(2)}\dots\sigma_\delta^{(n)}$. From this operator it takes another $(n-2)$ commutators to create any B_i . For an illustration we refer to the Lie tree diagram in Ref. [14]. Thus, the depth is upper bounded by $D(B_i) \leq 2(n-2)$, yielding the bound (1). Provided that $\|\mathbf{a}\|_\infty$ scales at most polynomially in the number of qubits,

we then have as a sufficient criterion for efficiently implementing a goal unitary U_g the following result. For the control system (3) and (4), a unitary gate $U_g(\mathbf{a})$ that is parametrized through $l(\mathbf{a})$ parameters can be implemented in a time that is at most polynomial in the number of qubits n if $l(\mathbf{a}) \leq \mathcal{O}(\text{poly}(n))$. Thus, for the control system in Eqs. (3) and (4) the set of gates $\mathcal{R}_{\text{poly}}$ given by Eq. (2) can be reached in a time that scales at most polynomially in the number of qubits. In particular for $\|\mathbf{a}\|_\infty = \mathcal{O}(1)$ and $l(\mathbf{a}) = \mathcal{O}(n)$ every U_g can be implemented in a time at most of the order $\mathcal{O}(n^4)$. However, in general for $l(\mathbf{a}) = 2^{2^n} - 1$ the upper bound scales exponentially $T(\mathbf{a}) \leq \mathcal{O}(n2^{6n})$. The bound (8) can be directly applied to efficiently simulating the dynamics with Hamiltonians [38–40]. For the control system expressed in Eqs. (3) and (4) every Hamiltonian $H = -i\Theta(\mathbf{a})$ consisting of $l(\mathbf{a}) \leq \mathcal{O}(\text{poly}(n))$ k -body interaction terms can be simulated efficiently. Since the strategy to obtain (8) is not necessarily time optimal, the actual set of gates that can be reached in polynomial time may be larger. It would be interesting to see how much the set can be increased using time optimal control methods [41]. However, the set $\mathcal{R}_{\text{poly}}$ can certainly be increased by considering the full expression in Eq. (8). Moreover, one can easily determine the maximum time needed to implement $U_g(\mathbf{a}) = \exp(\Theta(\mathbf{a}))$ by expanding Θ in the Pauli operator basis and calculating (8).

For $l(\mathbf{a}) = 1$ the target unitary operation is given by $U_g = \exp(a_i B_i)$ and it follows from Eq. (7) that the time to implement such an operation is upper bounded by $\tau(a_i B_i) \leq (1/J)[\pi(n-2) + |a_i|]$. For instance, every two qubit gate corresponding to a basis operator with two Pauli operators can be implemented in a time that scales at most linearly in the number of qubits. Moreover, gates corresponding to basis operators with n Pauli operators, i.e., n body interaction terms of the form $\sigma_\alpha^{(1)}\sigma_\beta^{(2)}\dots\sigma_\delta^{(n)}$, can be implemented in linear time as well. The bound can be tightened by introducing the geodesic path distance $d(i, j)$ between two qubits i and j as the smallest number of edges in a path connecting the two considered qubits. For example, it follows that the time to create a CNOT gate between qubit i and j is upper bounded $T_{\text{CNOT}} \leq \pi(\{[d(i, j) - 1]/J\} + 1/4J)$. Since every two qubit gate can be implemented with at most three CNOT gates [42], up to local unitary rotations, we have $T_{2\text{qubit}} \leq 3\pi(\{[d(i, j) - 1]/J\} + 1/4J)$. This bound is tighter than the bound that would be obtained by simply implementing a CNOT gate on two nearest neighbor qubits followed by SWAP operations [26,27]. The upper bound for $T_{2\text{qubit}}$ describes how much time is maximally needed in order to implement a generic two qubit gate on a qubit graph (3), provided each qubit can be instantaneously controlled locally. Therefore, the bound for $T_{2\text{qubit}}$ characterizes the time scale for entangling two qubits in a generic qubit network.

The characterization of the set of gates that can be reached in polynomial time (2) relied on the assumption that each qubit is subject to two orthogonal controls. A natural question is whether the number of controls can be reduced while still being able to implement $\mathcal{R}_{\text{poly}}$ in a time that scales at most polynomially in the number of qubits. Before presenting an n -qubit graph for which this is the case with only $n + 1$ controls, we address the question regarding how qubit systems can be concatenated in order to implement $\mathcal{R}_{\text{poly}}$ on the total system.

Concatenating systems.—Suppose we have two n -qubit graphs $G_1(V_1, E_1)$ and $G_2(V_2, E_2)$ for which the time to implement a generic two qubit unitary on each of the graphs is upper bounded by T_c . Now, as represented in Fig. 2, connect the two graphs with a single controllable two-body interaction, say $\sigma_z^{(i)} \sigma_z^{(j)}$ with $i \in V_1$ and $j \in V_2$. Importantly, $\{i\sigma_\alpha^{(i)} \sigma_\beta^{(i')}, i\sigma_\gamma^{(j)} \sigma_\delta^{(j')}, i\sigma_z^{(i)} \sigma_z^{(j)}\}$ with $i, i' \in V_1$ and $j, j' \in V_2$ forms a generating set \mathcal{S} .

Thus, every basis operator $B_i \in \mathfrak{su}(2^{2n})$ for the total system can be created through a nested commutator formed by the elements of \mathcal{S} . The time to create a unitary $U_g = \exp(a_i S_i)$ with $S_i \in \mathcal{S}$ is upper bounded by T_c . Therefore, it takes $\Delta t \leq 2T_c$ to produce a nested commutator of depth 1 and the depth for a generic basis operator $D(B_i)$ is upper bounded by $2(2n - 1)$. Consequently, the time to implement a unitary transformation $U_g = \exp(a_i B_i)$ on the total system is upper bounded by $\tau(a_i B_i) \leq T_c[4(2n - 1) + 1]$. As in the previous paragraph, a Trotter sequence yields a generic $U_g(\mathbf{a}) \in \text{SU}(2^{2n})$ up to an error ϵ so that for the combined system an upper bound on $T(\mathbf{a})$ is obtained, where the explicit form is given in Ref. [14]. For $\|\mathbf{a}\|_\infty^2 = \mathcal{O}(1)$ and $l(\mathbf{a}) = \mathcal{O}(n)$, we conclude that $T(\mathbf{a}) \leq \mathcal{O}(T_c(2n)^5)$. It immediately follows that upon combining L qubit graphs, each consisting of n qubits, through $L - 1$ controllable two body interactions, then the time $T(\mathbf{a})$ to implement $U_g(\mathbf{a}) \in \text{SU}(2^{Ln})$ scales at most as $\mathcal{O}(T_c(Ln)^5)$. Thus, as a sufficient criterion for a qubit system being scalable we have the following result.




FIG. 2. Illustration of how L qubit graphs, each consisting of n qubits, can efficiently be concatenated through controllable two body interactions (dotted lines). Assuming that on each qubit graph any two qubit gate can be implemented in a time smaller than T_c , for $\|\mathbf{a}\|_\infty^2 = \mathcal{O}(1)$ and $l(\mathbf{a}) = \mathcal{O}(n)$, we then have for the total system $T(\mathbf{a}) \leq \mathcal{O}(T_c(Ln)^5)$.

Using $L - 1$ controllable two body interactions every $U_g(\mathbf{a}) \in \text{SU}(2^{Ln})$ can be implemented on a Ln -qubit network in a time which scales at most polynomially in L if $\|\mathbf{a}\|_\infty, l(\mathbf{a}) \leq \mathcal{O}(\text{poly}(Ln))$.

Concatenating blocks of qubits through controllable two body interactions allows for scaling the total system so that the gate set $\mathcal{R}_{\text{poly}}$ can be implemented efficiently on the combined system. This situation emphasizes the importance of being able to control two body interactions. However, an allied question is whether a qubit graph exists for which a few local controls are sufficient to implement $\mathcal{R}_{\text{poly}}$ efficiently. To address this goal requires identifying a system and a number of controls that allow for implementing each two qubit unitary (5) in a time that scales at most polynomially in the number of qubits. Based on a decoupling scheme, for an n -qubit system in the previous paragraph, this goal is always possible using $2n$ controls. Now we show that for a star shaped graph the number of controls can be reduced to $n + 1$.

Reducing the number of controls.—Consider a star shaped graph described by the drift Hamiltonian $H_0 = J \sum_{i=2}^{N+1} (\sigma_x^{(1)} \sigma_x^{(i)} + \sigma_y^{(1)} \sigma_y^{(i)}) + J \sum_{i=2}^{N+1} \sigma_y^{(i)}$, where for the sake of simplicity the couplings and the energy splittings are assumed to be all given by J . Control is exerted through $\{\sigma_x^{(1)}, \sigma_y^{(1)}, \sigma_z^{(i)}\}$, $i = 2, \dots, N + 1$. For an illustration of such a graph we refer to the online material [14]. Through decoupling using a string of $\sigma_z^{(i)}$ we can instantaneously implement unitaries corresponding two body interaction terms $H_k = (\sigma_x^{(1)} \sigma_x^{(k)} + \sigma_y^{(1)} \sigma_y^{(k)} + \sigma_y^{(k)})$. Further decoupling with $\sigma_x^{(1)}, \sigma_z^{(k)}$ and instantaneous local rotations of qubit 1 and qubit k yield unitaries corresponding to $\sigma_\alpha^{(1)} \sigma_x^{(k)}, \sigma_\beta^{(1)} \sigma_y^{(k)}$ and $\sigma_y^{(k)}$. Recall that $\Delta t = (\pi/2J)$ units of time are needed to create a unitary operation generated by $[\sigma_\alpha^{(1)} \sigma_x^{(k)}, \sigma_y^{(k)}]$. Further note that a unitary operation $U_g = \exp(a_i S_i)$ with $S_i \in \{i\sigma_\alpha^{(1)} \sigma_\beta^{(k)}\}$ takes at most $T_c = (1/J)[(\pi/2) + |a_i|]$ time, where $\{i\sigma_\alpha^{(1)} \sigma_\beta^{(k)}\}$ forms a generating set. In order to obtain a unitary operation corresponding to a commutator $[S_1, S_2]$ requires $\Delta t \leq (3\pi/J)$ units time. Thus, the time $\tau(a_i B_i)$ to create $U_g = \exp(a_i B_i)$ with $B_i \in \mathfrak{su}(2^{N+1})$ is upper bounded by $\tau(a_i B_i) \leq D(B_i)(3\pi/J) + (1/J)[(\pi/2) + |a_i|]$. By upper bounding the depth we then find $\tau(a_i B_i) \leq (1/J)[(\pi/2)(12(n - 2) + 1) + |a_i|]$, where $n = N + 1$ is the total number of qubits. Therefore, every $U_g = \exp(a_i B_i)$ can be implemented in a time that scales at most linearly in the number of qubits. Again using a Trotter sequence finally permits concluding that for the star shaped graph every $U \in \mathcal{R}_{\text{poly}}$ can be implemented efficiently up to some error ϵ using only $n + 1$ controls. The star shaped graph model is of particular importance since it is used to describe the interaction of an electron spin in a nitrogen vacancy center with the surrounding nuclear spins [43–46], and, in general, quantum dots in a spin bath [47–49].

Conclusions.—We have characterized the set of gates that (provably) can be implemented in polynomial time on a generic qubit network where each qubit is controlled locally using time dependent fields. The characterization relied on the assumption that the control fields are unconstrained in strength. Further investigations regarding the importance of this assumption, as well as an assessment of the tightness of the derived bound can be found in the Supplemental Material [14]. The control of two body interactions allows for concatenating blocks of qubits so that the total system can be controlled efficiently, thereby paving the way towards scalable quantum devices. Moreover, we have identified a model, for which the efficiently implementable gate set can be realized with fewer local controls. An interesting goal would be the determination of the minimum number of controls required to implement $\mathcal{R}_{\text{poly}}$ and the corresponding graph topologies.

C. A. acknowledges the NSF (Grant No. CHE-1464569) and fruitful discussion with Daniel Burgarth and Benjamin Russell. H. R. acknowledges the ARO (Grant No. W911NF-16-1-0014).

-
- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, MA, 2010).
- [2] D. D'Alessandro, *Introduction to Quantum Control and Dynamics* (Chapman & Hall, New York, 2008).
- [3] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloss, I. Kuprov, B. Luy, S. Schirmer, T. H. Herbrüggen, D. Sugny, and F. K. Wilhelm, *Eur. Phys. J. D* **69**, 279 (2015).
- [4] C. Brif, R. Chakrabarti, and H. Rabitz, *New J. Phys.* **12**, 075008 (2010).
- [5] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *J. Magn. Reson.* **172**, 296 (2005).
- [6] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, and J. Biamonte, *Nat. Commun.* **5**, 3371 (2014).
- [7] H. Rabitz, R. de Vivie-Riedly, M. Motzkus, and K. Kompa, *Science* **288**, 824 (2000).
- [8] B. Dive, A. Pitchford, F. Mintert, and D. Burgarth, *arXiv:1701.01723*.
- [9] R. S. Judson and H. Rabitz, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [10] V. S. Manu and A. Kumar, *Phys. Rev. A* **86** 022324 (2012).
- [11] M. A. Nielsen, M. R. Dowling, M. Gu, and A. C. Doherty, *Science* **311**, 1133 (2006).
- [12] M. A. Nielsen, M. R. Dowling, M. Gu, and A. C. Doherty, *Phys. Rev. A* **73**, 062323 (2006).
- [13] C. Gokler, S. Lloyd, P. Shor, and K. Thompson, *Phys. Rev. Lett.* **118**, 260501 (2017).
- [14] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.120.220503>, which includes Refs. [15,16].
- [15] J. Johansson, P. Nation, and F. Nori, *Comput. Phys. Commun.* **183**, 1760 (2012).
- [16] R. Johansson, P. Nation, A. Pitchford, C. Granade, A. L. Grimsmo, Markusbaden, A. Vardhan, P. Migda, Kafischer, D. Vasilyev *et al.*, qutip/qutip: Qutip-4.0.0 (2016), URL: <https://doi.org/10.5281/zenodo.220867>.
- [17] D. L. Elliot, *Bilinear Control Systems*, Applied Mathematical Science (Springer, New York, 2000), Vol. 169.
- [18] C. Arenz, G. Gualdi, and D. Burgarth, *New J. Phys.* **16**, 065023 (2014).
- [19] X. Wang, A. Bayat, S. G. Schirmer, and S. Bose, *Phys. Rev. A* **81**, 032312 (2010).
- [20] D. Burgarth, K. Maruyama, M. Murphy, S. Montangero, T. Calarco, F. Nori, and M. B. Plenio, *Phys. Rev. A* **81**, 040303 (2010).
- [21] J. S. Hodges, J. C. Yang, C. Ramanathan, and D. G. Cory, *Phys. Rev. A* **81**, 010303 (2010).
- [22] Y. Zhang, C. A. Ryan, R. Laflamme, and J. Baugh, *Phys. Rev. Lett.* **107**, 170503 (2011).
- [23] R. Zeier and T. S. Herbrüggen, *J. Math. Phys. (N.Y.)* **52**, 113510 (2011).
- [24] Z. Zimborás, R. Zeier, T. Schulte-Herbrüggen, and D. Burgarth, *Phys. Rev. A* **92**, 042309 (2015).
- [25] G. C. Hegerfeldt, *Phys. Rev. Lett.* **111**, 260501 (2013).
- [26] N. Khaneja, R. Brockett, and S. J. Glaser, *Phys. Rev. A* **63**, 032308 (2001).
- [27] N. Khaneja, S. J. Glaser, and R. Brockett, *Phys. Rev. A* **65**, 032301 (2002).
- [28] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, *Phys. Rev. Lett.* **96**, 060503 (2006).
- [29] N. Khaneja, B. Heitmann, A. Spörl, H. Yuan, T. Schulte-Herbrüggen, and S. J. Glaser, *Phys. Rev. A* **75**, 012322 (2007).
- [30] S. Ashhab, P. C. de Groot, and F. Nori, *Phys. Rev. A* **85**, 052327 (2012).
- [31] We remark here that this assumption, and thus the ability to instantaneously implement generic single qubit gates, is a reasonable approximation in the case where the strength of the control fields can be made much larger than the energy splittings or coupling constants of the systems being considered, which, for instance, is typically achieved in spin systems. In the Supplemental Material [14] we show for a 4-spin Heisenberg chain that smooth pulses with a finite amplitude can be obtained through numerical gate optimization with a gate time below the derived bound.
- [32] L. Viola, E. Knill, and S. Lloyd, *Phys. Rev. Lett.* **82**, 2417 (1999).
- [33] J. L. Dodd, M. A. Nielsen, M. J. Bremner, and R. T. Thew, *Phys. Rev. A* **65**, 040301 (2002).
- [34] L. Viola, S. Lloyd, and E. Knill, *Phys. Rev. Lett.* **83**, 4888 (1999).
- [35] F. Motzoi, M. P. Kaicher, and F. K. Wilhelm, *Phys. Rev. Lett.* **119**, 160503 (2017).
- [36] M. Suzuki, *J. Math. Phys. (N.Y.)* **26**, 601 (1985).
- [37] H. De Raedt, in *Annual Reviews of Computational Physics IV*, edited by D. Stauffer (World Scientific, Singapore, 1996), p. 107.
- [38] S. Lloyd, *Science* **273**, 1073 (1996).
- [39] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).

- [40] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, *Phys. Rev. Lett.* **106**, 170501 (2011).
- [41] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, *Phys. Rev. A* **72**, 042331 (2005).
- [42] G. Vidal and C. M. Dawson, *Phys. Rev. A* **69**, 010301 (2004).
- [43] R. Hanson, F. M. Mendoza, R. J. Epstein, and D. D. Awschalom, *Phys. Rev. Lett.* **97**, 087601 (2006).
- [44] A. Albrecht, A. Reztker, F. Jelezko, and M. B. Plenio, *New J. Phys.* **15**, 083014 (2013).
- [45] F. Jelezko, T. Gaebel, I. Popa, A. Gruber, and J. Wrachtrup, *Phys. Rev. Lett.* **92**, 076401 (2004).
- [46] M. V. Gurudev Dutt, L. Childress, L. Jiang, E. Togan, J. Maze, F. Jelezko, A. S. Zibrov, P. R. Hemmer, and M. D. Lukin, *Science* **316**, 1312 (2007).
- [47] L. Ratschbacher, C. Sias, L. Carcagni, J. M. Silver, C. Zipkes, and M. Köhl, *Phys. Rev. Lett.* **110**, 160402 (2013).
- [48] E. M. Kessler, S. Yelin, M. D. Lukin, J. I. Cirac, and G. Giedke, *Phys. Rev. Lett.* **104**, 143601 (2010).
- [49] M. Warner, S. Din, I. S. Tupitsyn, G. W. Morley, A. M. Stoneham, J. A. Gardener, Z. Wu, A. J. Fisher, S. Heutz, C. W. M. Kay, and G. Aeppli, *Nature (London)* **503**, 504 (2013).