Optimal control-based efficient synthesis of building blocks of quantum algorithms: A perspective from network complexity towards time complexity

T. Schulte-Herbrüggen,[1,*] A. Spörl,[1] N. Khaneja,[2] and S. J. Glaser[1]
[1]*Department of Chemistry, Technical University Munich, Lichtenbergstrasse 4, D-85747 Garching, Germany*
[2]*Division of Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
(Received 30 June 2005; published 27 October 2005)

In this paper, we demonstrate how optimal control methods can be used to speed up the implementation of modules of quantum algorithms or quantum simulations in networks of coupled qubits. The gain is most prominent in realistic cases, where the qubits are not all mutually coupled. Thus the shortest times obtained depend on the coupling topology as well as on the characteristic ratio of the time scales for local controls vs nonlocal (i.e., coupling) evolutions in the specific experimental setting. Relating these minimal times to the number of qubits gives the tightest known upper bounds to the actual time complexity of the quantum modules. As will be shown, time complexity is a more realistic measure of the experimental cost than the usual gate complexity. In the limit of fast local controls (as, e.g., in NMR), time-optimized realizations are shown for the quantum Fourier transform (QFT) and the multiply controlled NOT gate ($C^{n-1}NOT$) in various coupling topologies of n qubits. The speed-ups are substantial: in a chain of six qubits the quantum Fourier transform so far obtained by optimal control is more than eight times faster than the standard decomposition into controlled phase, Hadamard and SWAP gates, while the $C^{n-1}NOT$ gate for a completely coupled network of six qubits is nearly seven times faster.

DOI: 10.1103/PhysRevA.72.042331 PACS number(s): 03.67.Lx, 03.65.Yz, 82.56.Jn, 82.56.Dj

I. INTRODUCTION

A key motivation for using experimentally controllable quantum systems to perform computational tasks or to simulate the behavior of other quantum systems [1,2] roots in reducing the complexity of the problem when going from a classical setting to a quantum setting. The most prominent pioneering example being Shor's quantum algorithm of prime factorization [3,4]. While known classical prime factorization algorithms are of nonpolynomial complexity [5], Shor's quantum algorithm brings it down into a class of polynomial complexity. Another celebrated example is Grover's quantum search algorithm [6,7], which allows for searching in an unstructured data base of n qubits with $N=2^n$ items in $O(\sqrt{N})$ quantum steps instead of $O(N)$ classical ones.

As a matter of fact, many quantum algorithms [8,9] can be subsumized as solving hidden subgroup problems in an efficient way [10]. In the Abelian case, the speed-up hinges on the quantum Fourier transform (QFT): while the network complexity of the fast Fourier transform (FFT) for n classical bits is of the order $O(n2^n)$ [11,12], the QFT for n qubits shows a complexity of order $O(n^2)$. For implementing a quantum algorithm or a quantum simulation in an experimental setup, it is customary to break it into universal elementary quantum gates [13]. Common sets comprise, e.g., (i) local operations such as the Hadamard gate, the phase gate, and (ii) the entangling operations CNOT, controlled-phase gates, $\sqrt{\text{SWAP}}, i$ SWAP, as well as (iii) the SWAP operation. The number of elementary gates required for implementing a quantum module then gives the network or gate complexity.

However, gate complexity often translates into too coarse an estimate for the actual time required to implement a quantum module (see, e.g., [14–16]), in particular, if the time scales of a specific experimental setting have to be matched. Instead, effort has been taken to give upper bounds on the actual time complexity [17], which is a demanding goal both from the Lie-algebraic [18–21] and geometric [22] points of view. With the time required for implementing a module in a specific experimental setting as the most realistic measure of cost, here we use methods of optimal control theory to find the minimum time by trying to solve the time-optimal control problem. The solution is hard to come by in general, so here we resort to numerical algorithms. The shortest times obtained depend on the coupling topology as well as on the characteristic ratio of the time scales for local controls vs nonlocal (i.e., coupling) evolutions and thus embrace the specific experimental setting. Relating these minimal times to the number of qubits gives the tightest known upper bounds to estimating the actual time complexity of the quantum modules, here for small numbers of qubits in a realistic experimental setup.

Moreover, as will be discussed, in the generic case there is no simple one-to-one relation between time complexity and network complexity, because of different time scales between local and nonlocal controls, different coupling topologies allowing for different degrees of parallelization, and different types of coupling interactions matching different sets of elementary gates.

Thus here we leave the usual approach of decomposing gates into sets of discrete universal building blocks. Instead, the scope is to exploit the differential geometry of the unitary group for optimization [23,24] when using the power of quantum control in order to obtain constructive bounds to minimal time both as close to the experimental setting and as

*Electronic address: tosh@ch.tum.de

tight as possible. In the limit of zero cost for the fast local controls (as in NMR) compared to the slow coupling interactions, we give decompositions for the QFT and the multiply-controlled NOT-gate C^{n-1}NOT that are dramatically faster than the fastest decompositions into standard gates known so far.

The paper is organized as follows. The first focus is on the somewhat surprising fact that for time-optimal decompositions of a desired unitary gate into a sequence of evolutions of experimentally available controls the global phase does play a role. This is the case when, e.g., there are different time scales for local versus nonlocal controls. However, global phases can readily be absorbed by shifting gradient flows from unitary to projective unitary groups. Then numerical time-optimal control provides the currently best upper bounds to the minimal time required for quantum modules such as the QFT or the C^{n-1}NOT gate in various coupling topologies of n qubit systems. Here we present examples with n up to seven. For $n \geqslant 3$, the resulting times are bounded from above by *KAK*-type decompositions taken to sub-Riemannian regimes [18–21].

Although our applications refer to time-optimized quantum computing in the NMR-type limit of fast local controls, the methods introduced are very general and apply to all systems whose dynamics can be cast into the closed form of finite-dimensional Lie algebras (to sufficient approximation). So finally we give an outlook from spin to pseudospin systems with comparable time scales of local and nonlocal operations. Time-optimal control thus contributes to the growing toolbox of control methods for steering quantum systems (see e.g., Refs. [25–34]).

II. CONTROLLABLE SYSTEMS

A. Spin and pseudospin systems

Here we address fully controllable [35–40] quantum systems represented as spin or pseudospin systems, i.e., those in which—neglecting decoherence—for any initial state represented by its density operator A, the entire unitary orbit $U(A) := \{UAU^{-1} \,|\, U \text{ unitary}\}$ can be reached [41]. In systems of n qubits (e.g., spins-$\frac{1}{2}$), this is the case under the following mild conditions [23,40,42]: (1) the qubits have to be inequivalent i.e., distinguishable and selectively addressable and (2) they have to be pairwise coupled (e.g., by Ising interactions), where the coupling topology may take the form of any connected graph.

B. Time scales for local and nonlocal controls

Let the quantum system evolve in a time interval t_k under combinations of piecewise constant control Hamiltonians $\{H_j\}$ each of which may be "switched on" with its control amplitude $u_j \in \mathbb{R}$ and the drift H_d, i.e., the free-evolution Hamiltonian, according to

$$H^{(k)} = H_d + \sum_j u_j^{(k)} H_j^{(k)}. \tag{1}$$

In NMR spin dynamics [43], for instance, the local controls on qubit ℓ are represented by a linear combination of the Pauli matrices $\{\sigma_{\ell x}, \sigma_{\ell y}\}$. And the drift term is governed by the weak scalar couplings (reminiscent of Ising interactions)

$$H_d = \pi \sum_{\ell<m} J_{\ell m} \frac{1}{2} \sigma_{\ell z} \otimes \sigma_{mz}, \tag{2}$$

provided the couplings between spins are much smaller than the difference between the eigenfrequencies (shifts Ω) of the respective spins $|J_{\ell m}| \ll |\Omega_\ell - \Omega_m|$. This is the case in heteronuclear spin systems. In quantum control even the homonuclear ones can be designed such as to meet this greatly simplifying approximation [44].

For the system to be fully controllable in the sense outlined above, $\{H_d\} \cup \{H_j\}$ has to form a generating set of the Lie algebra su(N) by way of the Lie bracket. Often the time scale for local controls is considerably faster than for the costly slow coupling evolutions.

III. TIME-OPTIMAL QUANTUM CONTROL

In order to control a quantum system of n qubits (spins-$\frac{1}{2}$) such as to realize a quantum gate or module of some quantum algorithm given by the unitary propagator $U_G \in U(2^n)$ in minimal time, one has to decompose

$$U_G \sim U(T) = e^{-i(t_M - t_{M-1})H^{(M)}} \cdots e^{-i(t_k - t_{k-1})H^{(k)}} \cdots e^{-i(t_1 - t_0)H^{(1)}} \tag{3}$$

—up to a global phase factor—into a time-optimal sequence $[T := \Sigma_k (t_k - t_{k-1}) \stackrel{!}{=} \min, \text{ vide infra}]$ of evolutions under piecewise constant Hamiltonians $H^{(k)}$.

A. Relevance of global phases

However, as propagators generated by the traceless spin Hamiltonians are elements of the respective *special* unitary groups, the quantum gates U_G are realized by $U(T)$ just up to global phases ϕ_p

$$U_G = e^{-i\phi_p} U(T); \tag{4}$$

so $U_G \in U(N)$, while $U(T) \in \mathrm{SU}(N)$. For n spins-$\frac{1}{2}$ read $N := 2^n$ henceforth. Note that with the center of SU(N) being

$$\mathbb{Z}_N := \{e^{i(2\pi p/N)} \mathbb{1}_N \,|\, p = 0, 1, \ldots, N-1\}, \tag{5}$$

one has a choice of N such phases

$$\phi_p \in \{\phi_0 + 2\pi p/N \,|\, p = 0, 1, \ldots, N-1\}, \tag{6}$$

where ϕ_0 shall be the smallest angle $\phi_0 \in [0, \pi]$ so that $\det\{e^{i\phi_0} U_G\} = +1$. Although global phases are clearly immaterial to quantum evolutions $\rho_0 \mapsto U \rho_0 U^{-1}$, it is important to note they do in fact contribute to the over-all time needed to implement U_G: consider, e.g.,

$$e^{-i(\pi/2)\sigma_{\ell z} \otimes \sigma_{mz}} = e^{i(\pi/2)} e^{-i(\pi/2)(\sigma_{\ell z} \otimes \mathbf{1} + \mathbf{1} \otimes \sigma_{mz})}, \tag{7}$$

where the nonscalar part of the right-hand side can be realized solely by (fast) local controls, while the left-hand side hinges on nothing but (slow) coupling evolution. In Fig. 1 this is further illustrated for the three-qubit QFT imple-

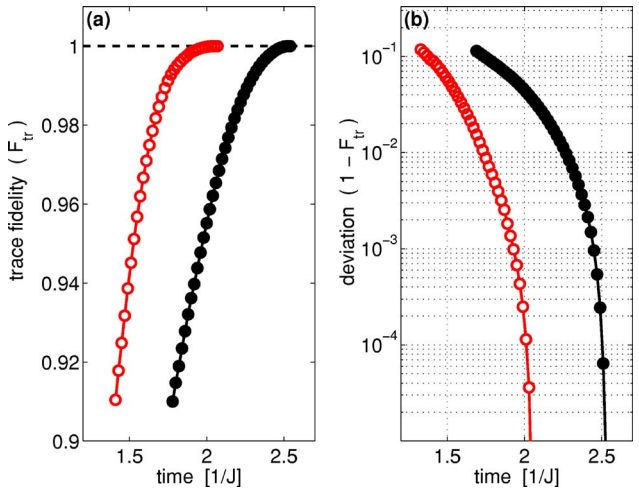

FIG. 1. (Color online) Global phase dependence of the times needed to implement the three-qubit QFT on a linear chain (L_3) of nearest-neighbor interactions with uniform weak scalar J couplings. The right curves (●) show the special unitary implementation of a QFT with the smallest global phase $\phi_0=1\pi/16$, where it takes $2.53\,J^{-1}$ to reach a fidelity $F_{\mathrm{tr}}:=(1/N)\mathrm{Re\ tr}\{e^{-i\phi_p}U_G^{\dagger}U(T)\}$ $\geqslant 0.99999$. The left curves (○) display the fastest QFT implementations obtained. They are attained with the global phase $\phi_1=5\pi/16$. Fidelities $\geqslant 0.99999$ are reached after $2.05\,J^{-1}$. Times for local controls are assumed to be negligible in this limit matching the typical NMR scenario, where the time cost is determined by coupling evolutions. (a) Gives the trace fidelities against time, while (b) shows the devitiations from full fidelity in a semilogarithmic way.

mented on a chain of three spins connected by nearest-neighbor interactions of weak scalar coupling in the NMR limit of zero time cost for the fast local controls.

B. Optimal control on projective groups

For a given unitary quantum gate U_{G} and propagators $U=U(t)$ describing the evolution of the quantum system, there are the two geometric tasks, one that explicitly carries the phase, while the other one automatically absorbs it as desired: (1) minimize the distance $\|U-e^{i\phi_p}U_{\mathrm{G}}\|_2$ by maximizing $\Phi_1:=\mathrm{Re\ tr}\{e^{-i\phi_p}U_{\mathrm{G}}^{\dagger}U\}$ for fixed ϕ_p, and (2) minimize the angle $\angle(U,U_{\mathrm{G}})\mathrm{mod}(\pi)$ by maximizing $\Phi_2:=|\mathrm{tr}\{U_{\mathrm{G}}^{\dagger}U\}|^2$.

(1) Clearly, the first task is related to the trace fidelity $F_{\mathrm{tr}}:=(1/N)\mathrm{Re\ tr}\{e^{-i\phi_p}U_G^{\dagger}U(T)\}$. In terms of control theory, it means maximizing the quality functional $\Phi_1[U(t)]=\mathrm{Re\ tr}\{e^{-i\phi_p}U_{\mathrm{G}}^{\dagger}U(T)\}$ with $0\leqslant t\leqslant T$ subject to the equation of motion $\dot{U}(t)=-iHU(t)$ (with $H=H_d+\Sigma_j u_j H_j$) and the initial condition $U(0)=\mathbb{1}$, whereas the final condition $U(T)$ is free at an appropriately fixed final time T (*vide infra*). As usual, the problem is readily solved by introducing the operator-valued Lagrange multiplier $\lambda(t)$ satisfying $\dot{\lambda}(t)=-iH\lambda(t)$ and a scalar-valued Hamiltonian function

$$h[U(t_k)]=\mathrm{Re\ tr}\left\{\lambda^{\dagger}(t_k)\left[-i\left(H_d+\sum_j u_j H_j\right)\right]U(t_k)\right\}. \quad (8)$$

Then, Pontryagin's maximum principle [45] may be exploited in a quantum setting [46,47] by requiring for the control amplitudes u_j to be optimal that

$$\frac{\partial h[U(t_k)]}{\partial u_j}\equiv -\,\mathrm{Im\ tr}\{\lambda^{\dagger}(t_k)H_jU(t_k)\}\;=\;0 \quad (9)$$

as well as for the final condition of the adjoint system

$$\lambda(T)=-\frac{\partial\Phi_1(T)}{\partial U(T)}=-e^{i\phi_p}U_{\mathrm{G}} \quad (10)$$

thus allowing one to implement a gradient-flow based recursion [47]. For the amplitude of the jth control in iteration $r+1$ at time interval t_k one finds with α as a suitably chosen step size (e.g., as input to conjugate gradients)

$$u_j^{(r+1)}(t_k)=u_j^{(r)}(t_k)+\alpha\frac{\partial h[U(t_k)]}{\partial u_j}. \quad (11)$$

The procedure is then repeated for a set of decreasing final times T up to a minimal time τ still allowing to get sufficient fidelity (compare Fig. 1).

(2) The second task amounts to maximizing $\Phi_2[U(t)]=|\mathrm{tr}\{U_{\mathrm{G}}^{\dagger}U(T)\}|^2$, which relates to the square of the trace fidelity and is easy to handle by gradient-flow methods. This problem, however, can readily be reduced to task (1): observe that to $U\in\mathrm{SU}(N)$,

$$\hat{U}:=U^*\otimes U \quad (12)$$

is a representation of the corresponding element of the projective special unitary group (definition, e.g., in Ref. [48])

$$\mathrm{PSU}(N)\overset{\mathrm{iso}}{=}\frac{\mathrm{SU}(N)}{\mathbb{Z}_N}\overset{\mathrm{iso}}{=}\frac{\mathrm{U}(N)}{\mathrm{U}(1)} \quad (13)$$

embedded in $\mathrm{SU}(N^2)$. Hence this representation is highly reducible yet very convenient, because

$$\Phi_1[\hat{U}(t)]=\mathrm{Re\ tr}\{\hat{U}_{\mathrm{G}}^{\dagger}\hat{U}(T)\}=|\mathrm{tr}\{U_{\mathrm{G}}^{\dagger}U(T)\}|^2=\Phi_2[U(t)]. \quad (14)$$

Hence one may adopt the previous results to obtain the gradient flow on $\mathrm{PSU}(N)$ just by using

$$\frac{\partial h[\hat{U}(t_k)]}{\partial u_j}\equiv -\,2\,\mathrm{Im}(\mathrm{tr}\{\lambda^t(t_k)H_j^tU^*(t_k)\}\,\mathrm{tr}\{\lambda(t_k)^{\dagger}U(t_k)\}) \quad (15)$$

in Eq. (11) (cp. also Ref. [47]). Thus an explicit tensor product never enters the algorithm. And the final condition of the adjoint system does not require any prior knowledge or screening of the global phase ultimately giving the fastest implementation as has been the case in previous settings, e.g., Ref. [49], because embedding $\mathrm{PSU}(N)$ in $\mathrm{SU}(N^2)$ enforces a global phase of zero for $\hat{U}$. Absorbing the phases cuts the number of computations for n-qubit systems by a factor $N=2^n$.

Having reduced task (2) to task (1) also saves all the convergence and step-size considerations [24] from $\mathrm{SU}(N^2)$ to apply to $\mathrm{PSU}(N)$. With these stipulations, the Hamiltonians H_k according to Eq. (1), and the numbering as in Eq. (3), the iterations r of Eq. (11) can be used in the following algorithmic scheme [47].

(1) Set initial controls $u_j^{(0)}(t_k)$ for all times t_k with $k=1,2,\ldots,M$ at random or by guess.
(2) Starting from $U_0=\mathbb{1}$, calculate for all $t_1,t_2,\ldots,t_k$ the forward propagation

$$U^{(r)}(t_k)=e^{-i(t_k-t_{k-1})H_k^{(r)}}e^{-i(t_{k-1}-t_{k-2})H_{k-1}^{(r)}}\cdots e^{-i(t_1-t_0)H_1^{(r)}}U_0. \tag{16}$$

(3) Likewise, starting with $T=t_M$ and $\lambda(T)$ from Eq. (10), compute for all $t_M,t_{M-1},\ldots,t_k$ the back propagation

$$\lambda^{(r)}(t_k)=e^{i(t_k-t_{k-1})H_k^{(r)}}e^{i(t_{k+1}-t_k)H_{k+1}^{(r)}}\cdots e^{i(t_M-t_{M-1})H_M^{(r)}}\lambda(T). \tag{17}$$

(4) Calculate $\partial h[U(t_k)]/\partial u_j$ according to Eq. (9).
(5) With $u_j^{(r+1)}(t_k)$ from Eq. (11) update all the piecewise Hamiltonians to $H_k^{(r+1)}$ and return to step 2.

IV. APPLICATIONS

For simplicity, the coupling strengths in all the subsequent examples are assumed uniform, thus enabling to give the times in units of J^{-1}. However, all our algorithms can equally well cope with nonuniform coupling constants directly matching the experimental settings.

Computational settings. The algorithmic scheme outlined above was implemented as a C++ program following previous work [47]. Typically, 10 random initial control sequences of a fixed final time T (step 1) are followed through 10 000 iterations of steps 1 to 5. This turned out to be sufficient and superior to using more random initial conditions at fewer number of iterations. Out of the 10 runs up to 10 000 cycles, the successful candidates typically show already a trace fidelity between 95 and 99 % (if the final time T is large enough). Out of them the best one is selected for further iterations up to fidelities $F_{\mathrm{tr}}>0.999\,99$. This procedure is typically repeated some five times for each fixed final time T. In order to give the shortest final time still allowing to obtain $F_{\mathrm{tr}}>0.999\,99$ with a resolution of $\pm 0.01\ J^{-1}$, some 20 to 50 candidates of T have to be calculated.

A. Towards a time-optimal quantum Fourier transform

The quantum Fourier transform (QFT) is central to all quantum algorithms of Abelian hidden subgroup type [50,51]. The time required to implement this module in n-qubit systems obviously depends on the topology of the coupling interactions. Figure 2 shows some of the topologies for the couplings of four qubits and the respective times (best numerical results from optimal control) for implementing the four-qubit QFT. Clearly, the complete coupling topology corresponds to the maximal graph K_n and thus allows for fast implementations, while the chain of nearest-neighbor interactions L_n is the minimal connected graph entailing slow implementations. Note, however, that the minimal times also depend on the ordering in the graph, because permutations (carried out by transpositions) may call for timewise costly SWAPs.

In complete coupling topologies K_n, the time for the decomposition into standard gates (controlled phase gate, Hadamard, and SWAP) can easily be expressed in units of J^{-1} as a function of the number qubits (compare Refs. [52,53])

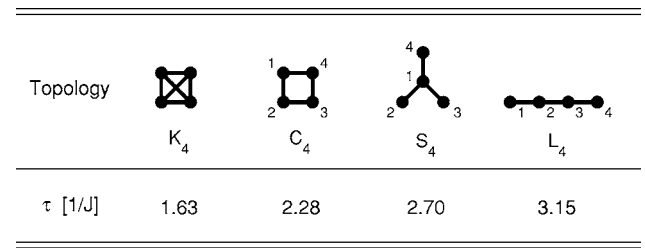

FIG. 2. Four ordered connected graphs with four vertices representing the topology of pairwise couplings (edges) between four qubits (vertices). Times rounded to $0.01J^{-1}$ are given for the shortest QFT-realizations obtained thus far by numerical time-optimal control with trace fidelities $(1/N)|\mathrm{tr}\{U_G^{\dagger}U(T)\}|\geqslant 0.99999$.

$$\tau(n)=\frac{1}{4}(n+3), \tag{18}$$

where the final SWAP introduces a constant time offset.

For complete topologies there is not much time to gain by using numerical time-optimal control. However, as is shown in Fig. 3 and Table I, for linear spin chains (L_n) with nearest-neighbor Ising interactions, time-optimal control provides decompositions of the QFT that are significantly faster than the corresponding decompositions into standard gates would impose: in six qubits, for instance, the speed-up is more than eightfold and in seven qubits approximately ninefold.

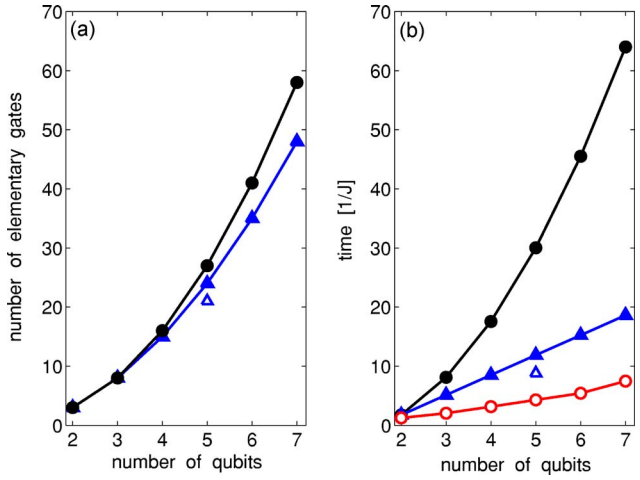

FIG. 3. (Color online) (a) Gate complexity of the QFT in linear coupling topologies L_n. Standard-gate decomposition (●) [58] and optimized scalable gate decomposition (▲) [54]. (b) Time complexity of the QFT in linear coupling topologies. Upper traces give analytical times associated with the decompositions of part (a): standard-gate decompositions (●) [58] and optimized scalable gate decompositions (▲) [54]; (△) gives a special (i.e., nonscalable) five-qubit decomposition into standard gates obtained by simulated annealing [54]. Lowest trace: speed-up by time-optimal control with shortest numerical realizations obtained (○) rounded to $0.01J^{-1}$. Further details in Table I.

TABLE I. Speed-up of the quantum Fourier transform on linear spin chains L_n.

qubits	Stand. QFT[a] $\tau\,[1/J]$[d]	Blais[b] $\tau\,[1/J]$[d]	Best results[c] $\tau\,[1/J]$[d]	Speed-ups Stand.	Blais
2	1.75	1.75	1.25	1.40	1.40
3	8.13	5.13	2.05	3.94	2.50
4	17.56	8.50	3.15	5.58	2.70
5	30.03	11.88(8.81)	4.30	6.98	2.76(2.05)
6	45.52	15.25	5.43	8.38	2.81
7	64.01	18.63	7.47	8.57	2.49

[a]Analytical times for decomposition into standard gates [58].
[b]Reference [54] in brackets: the nonscalable special five-qubit QFT.
[c]Upper bounds to minimal time for achieving a trace fidelity of $\geqslant 0.99999$ [>0.9975 for 7 qubits] by numerical optimal control.
[d]Times τ are rounded to $0.01\,J^{-1}$.

For a fair comparison, however, note that Blais [54] (in contrast to Fowler *et al.* [55]) also suggests to permute the output qubits which saves one SWAP of $1.5J^{-1}$. Clearly, searching through $n!$ qubit permutations or even $(2^n)!$ levels [56] can provide even faster realizations, as will be pursued elsewhere.

B. Towards a time-optimal C^{n-1}NOT

Likewise, one may strive to implement the C^{n-1}NOT gate in a time-optimal way. In a complete coupling topology of n qubits, the algorithmic complexity was described by Barenco *et al.* [57] as increasing exponentially up to six qubits, whereas the increase from seven qubits onwards is quadratic. Again, time-optimal control provides a dramatic speed-up in this case as well, see Fig. 4 and Table II as well as the controls in Fig. 5.

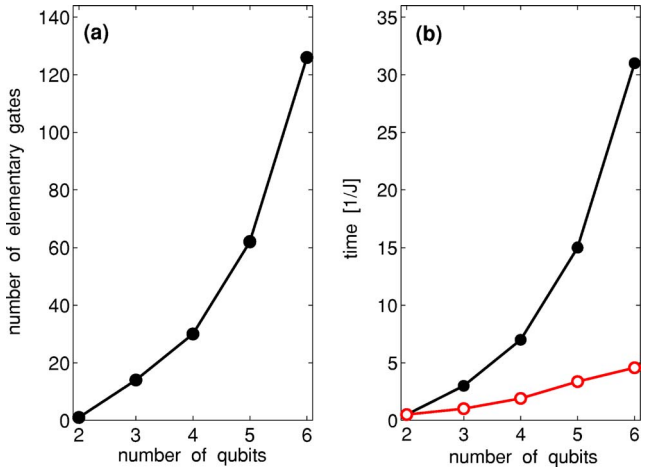

FIG. 4. (Color online) (a) Network complexity of the C^{n-1}NOT gate on complete coupling topologies K_n [57]. (b) Time complexity of the C^{n-1}NOT gate on complete coupling topologies. Upper trace: analytical times for decomposition into standard gates (●) [57]. Lower trace: speed-up by time-optimal control with shortest times (○) currently needed for realizing C^{n-1}NOT by numerical control rounded to $0.01\,J^{-1}$.

C. Beyond spins: Controlling coupled charge qubits in Josephson devices

Obviously the optimal control methods presented thus far can be generalized such as to hold for systems with finite degrees of freedom allowing for a pseudospin formulation in terms of closed Lie algebras. For instance, the standard CNOT-gate can be realized in two coupled charge qubits of a solid-state Josephson device some five times faster than in the pioneering setting of Ref. [59]. Yet one easily obtains a trace fidelity beyond $1-10^{-9}$ as shown in Ref. [60]. Similarly we found realizations of the TOFFOLI gate in three linearly coupled charge qubits that are some 13 times faster than by decomposition into nine CNOT gates of the NEC group [59] and still 2.8 times faster than nine optimized CNOTs [60].

V. DISCUSSION

The goal is to extend the optimal control methods to larger modules of quantum algorithms or simulations in order to implement them both in a time-optimal and experimentally robust way. Thus the growing set of numerical examples will hopefully provide inspiration to understand time-optimal steerings of quantum systems algebraically, which, however, seems very demanding in the cases presented here. In other instances such as for the propagator

TABLE II. Speed-up for the C^{n-1}NOT gate in complete coupling topologies of n qubits K_n.

qubits	Stand. decomposition[a] $\tau\,[1/J]$[c]	Best results[b] $\tau\,[1/J]$[c]	speed-up
2	0.5	0.50	1.00
3	3.0	1.01	2.97
4	7.0	1.90	3.68
5	15.0	3.37	4.45
6	31.0	4.59	6.75

[a]Barenco *et al.* [57].
[b]Upper bounds to minimal time for trace fidelities $\geqslant 0.99999$ [>0.999 for 6 qubits] by numerical optimal control.
[c]Times τ are rounded to $0.01\,J^{-1}$.

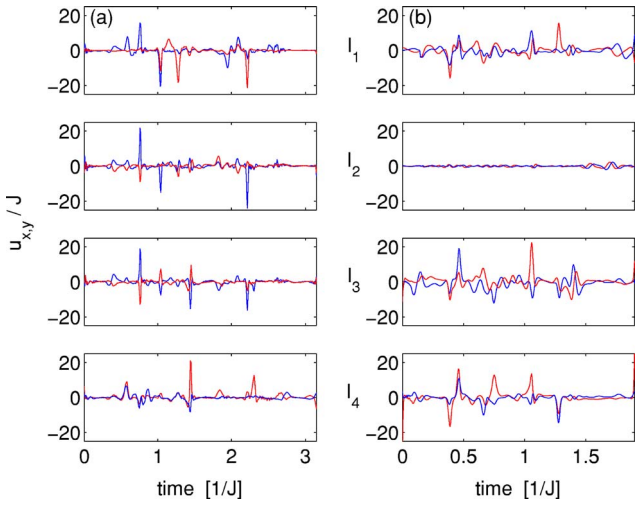

FIG. 5. (Color online) Time course of controls for the shortest realizations obtained for the following quantum modules: (a) the QFT on a linear coupled four-qubit system (L_4); (b) the C^3NOT gate on a fully coupled four-qubit system (K_4). Traces in blue: amplitudes for $\sigma_{\ell x}$ controls (x pulses); red: amplitudes for $\sigma_{\ell y}$ controls (y pulses) on the spins $\ell=1, 2, 3, 4$.

$$U(t) = e^{-i\pi Jt[(1/2)\sigma_{1z}\otimes\sigma_{2z}\otimes\sigma_{3z}]}, \tag{19}$$

the theory is fully understood, and the predictions based on sub-Riemannian geodesics [20] perfectly match (i) the numerically obtained time complexity as well as (ii) the actual time course for the controls [47] for all $\pi Jt \in [0, \pi/2]$.

Along these lines, the above controls may finally trigger a theoretical understanding. The ultimate challenge then is to extract a principle for a scalable time-optimal control scheme from the set of numerical examples.

VI. CONCLUSION

Here we have left the usual approach of decomposing quantum modules into sets of discrete building blocks, such as elementary universal quantum gates thus expressing the cost as algorithmic network complexity. Instead we proposed to refer to time complexity as the experimentally relevant cost: it allows for exploiting the continuous differential geometry of the unitary Lie groups as well as the power of quantum control for getting constructive upper bounds to minimal times both perfectly matching the experimental setting while being as tight as possible. In particular when local and nonlocal operations are of different time scale, different global phases entail different minimal times. In the limit of zero cost for the fast local controls we gave decompositions for the QFT and the multiply controlled NOT-gate C^{n-1}NOT that are dramatically faster than the best decompositions into standard gates known so far would impose. However, there is no guarantee the ultimate time optimum is attained, also because going systematically through all permutations of the qubits may give further improvement.

Note that the time gain in decomposing unitary modules holds, no matter, whether pure states or ensembles are used. The approach also clearly shows that in the generic case there is no simple one-to-one relation between time complexity and network complexity. This is for very practical reasons: typically (1) not all the elementary gates are of the same time cost, but each experimental implementation comes with its characteristic ratio of times required for local vs nonlocal (coupling) operations. (2) Not all the elementary gates have to be performed sequentially, but can be rearranged so that some of the commuting operations (e.g., controlled phase gates between several qubits) or operations in disjoint subspaces can be parallelized. (3) The coupling topology between the qubits does not have to form a complete graph (K_n) but may be just a connected subgraph, and each graph comes with a specific potential of parallelizing timewise costly interactions; this is demonstrated for the QFT on complete coupling topologies K_n versus the linear coupling topology L_n, where the parallel performance of controlled phase gates [54] reduces quadratic time complexity to linear complexity, which, however, can be further speeded up by time-optimal control. (4) The experimental setting with its specific type and individual strengths of coupling interaction (e.g., Ising or Heisenberg-*XY* or *XYZ* type) related to the choice of universal gates for the network decomposition may introduce some arbitrariness. It is for these very reasons that time complexity is the more realistic measure of the experimentally relevant cost than network complexity is.

VII. OUTLOOK

Although extrapolation may be premature, it is fair to anticipate that in systems of some 20 qubits network decompositions will often become impractical. Thus time-optimal decompositions into controls actually available in the experimental setting promise to widen the range of experimentally accessible tasks significantly and will prove useful in many experimental implementations. Moreover, analyzing the topology dependence of minimal times while allowing for nonuniform coupling strengths will contribute valuable guidelines for designing optimized architectures of quantum computational hardware.

ACKNOWLEDGMENTS

This work was supported by *Deutsche Forschungsgemeinschaft*, DFG, through the grant Gl 203/4-2 in the incentive SPP 1078 (*Schwerpunkt-Programm "Quanteninformationsverarbeitung," QIV*). Encouraging discussion with Ivan Deutsch and Andrew Landahl at the QCMC 2004 where part of this work was presented is gratefully acknowledged. We wish to thank Timo Reiss for help with early stages of the C++ program in the collaboration of Ref. [47].

[1] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[2] R. P. Feynman, *Feynman Lectures on Computation* (Perseus Books, Reading, MA, 1996).
[3] P. W. Shor, in *Proceedings of the Symposium on the Foundations of Computer Science, 1994, Los Alamitos, California* (IEEE Computer Society Press, New York, 1994), pp. 124–134.
[4] P. W. Shor, SIAM J. Comput. **26**, 1484 (1997).
[5] C. H. Papadimitriou, *Computational Complexity* (Addison Wesley, Reading, MA, 1995).
[6] L. K. Grover, in *Proceedings of the 28th Annual Symposium on the Theory of Computing, 1996, Philadelphia, Pennsylvania* (ACM Press, New York, 1996), pp. 212–218.
[7] L. K. Grover, Phys. Rev. Lett. **79**, 325 (1997).
[8] R. Jozsa, Proc. R. Soc. London, Ser. A **454**, 323 (1998).
[9] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Proc. R. Soc. London, Ser. A **454**, 339 (1998).
[10] M. Ettinger, P. Høyer, and E. Knill, Inf. Process. Lett. **91**, 43 (2004).
[11] J. W. Cooley and J. W. Tukey, Math. Comput. **19**, 297 (1965).
[12] T. Beth, *Verfahren der schnellen Fourier-Transformation* (Teubner, Stuttgart, 1984).
[13] D. Deutsch, Proc. R. Soc. London, Ser. A **400**, 97 (1985).
[14] G. Vidal, K. Hammerer, and J. I. Cirac, Phys. Rev. Lett. **88**, 237902 (2002).
[15] A. M. Childs, H. L. Haselgrove, and M. A. Nielsen, Phys. Rev. A **68**, 052311 (2003).
[16] R. Zeier, M. Grassl, and T. Beth, Phys. Rev. A **70**, 032319 (2004).
[17] P. Wocjan, D. Janzing, and T. Beth, Quantum Inf. Comput. **2**, 117 (2002).
[18] N. Khaneja and S. J. Glaser, Chem. Phys. **267**, 11 (2001).
[19] N. Khaneja, R. Brockett, and S. J. Glaser, Phys. Rev. A **63**, 032308 (2001).
[20] N. Khaneja, S. J. Glaser, and R. Brockett, Phys. Rev. A **65**, 032301 (2002).
[21] H. N. Sá-Earp and J. K. Pachos, quant-ph/0505128 (unpublished).
[22] M. A. Nielsen, quant-ph/0502070 (unpublished).
[23] S. J. Glaser, T. Schulte-Herbrüggen, M. Sieveking, O. Schedletzky, N. C. Nielsen, O. W. Sørensen, and C. Griesinger, Science **280**, 421 (1998).
[24] U. Helmke, K. Hüper, J. B. Moore, and T. Schulte-Herbrüggen, J. Global Optim. **23**, 283 (2002).
[25] D. J. Tannor and S. A. Rice, J. Chem. Phys. **83**, 5013 (1985).
[26] S. Shi and H. Rabitz, J. Chem. Phys. **92**, 364 (1989).
[27] H. Rabitz, R. de Vivie-Riedle, M. Motzkus, and K. Kompa, Science **288**, 824 (2000).
[28] S. Lloyd, Phys. Rev. A **62**, 022108 (2000).
[29] See entire special issue "Laser Control of Quantum Dynamics," Chem. Phys. **267** 1 (2001).
[30] J. P. Palao and R. Kosloff, Phys. Rev. Lett. **89**, 188301 (2002).
[31] J. J. García-Ripoll, P. Zoller, and J. I. Cirac, Phys. Rev. Lett. **91**, 157901 (2003).
[32] M. Tesch and R. de Vivie-Riedle, J. Chem. Phys. **121**, 12158 (2004).
[33] Y. Ohtsuki, G. Turinici, and H. Rabitz, J. Chem. Phys. **120**, 5509 (2004).
[34] S. E. Sklarz and D. J. Tannor quant-ph/0404081 (unpublished).
[35] V. Jurdjevic and H. Sussmann, J. Diff. Equations **12**, 313 (1972).
[36] H. Sussmann and V. Jurdjevic, J. Diff. Equations **12**, 95 (1972).
[37] R. W. Brockett, SIAM J. Control **10**, 265 (1972).
[38] R. W. Brockett, SIAM J. Appl. Math. **25**, 213 (1973).
[39] W. M. Boothby and E. N. Wilson, SIAM J. Control Optim. **17**, 212 (1979).
[40] T. Schulte-Herbrüggen, Ph.D. thesis, Diss-ETH 12752, Zürich, 1998.
[41] F. Albertini and D. D'Alessandro, IEEE Trans. Autom. Control **48**, 1399 (2003).
[42] T. Schulte-Herbrüggen, K. Hüper, U. Helmke, and S. J. Glaser, *Applications of Geometric Algebra in Computer Science and Engineering* (Birkhäuser, Boston, 2002), Chap. Geometry of Quantum Computing by Hamiltonian Dynamics of Spin Ensembles, pp. 271–283.
[43] R. R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Clarendon Press, Oxford, 1987).
[44] R. Marx, A. F. Fahmy, J. M. Myers, W. Bermel, and S. J. Glaser, Phys. Rev. A **62**, 012310 (2000).
[45] L. S. Pontryagin, V. G. Bol'tanskii, R. S. Gamkrelidze, and E. F. Mischenko, *The Mathematical Theory of Optimal Processes* (Pergamon Press, New York, 1964).
[46] A. G. Butkovskiy and Y. I. Samoilenko, *Control of Quantum-Mechanical Processes and Systems* (Kluwer, Dordrecht, 1990).
[47] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, J. Magn. Reson. **172**, 296 (2005).
[48] B. Huppert, *Endliche Gruppen* (Springer, Berlin, 1979), Vol. I especially pp. 177 ff. and 233 ff.
[49] J. H. Wesenberg, Phys. Rev. A **69**, 042323 (2004), see in particular Eq. (14) of the extended preprint version quant-ph/0312076 (unpublished).
[50] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).
[51] G. Alber, T. Beth, M. Horodecki, P. Horodecki, R. Horodecki, M. Rötteler, H. Weinfurter, R. Werner, and A. Zeilinger, *Quantum Information: an Introduction to Basic Concepts and Experiments, Vol. 173 of Springer Tracts in Modern Physics* (Springer, Heidelberg, 2000).
[52] M. D. Price, T. F. Havel, and D. J. Cory, New J. Phys. **2**, 10.1 (2000).
[53] K. Dorai and D. Suter, quant-ph/0211030 (unpublished).
[54] A. Blais, Phys. Rev. A **64**, 022312 (2001).
[55] A. G. Fowler, S. J. Devitt, and L. C. L. Hollenberg, Quantum Inf. Comput. **4**, 237 (2004).
[56] M. Nakahara, J. Vartiainen, Y. Kondo, S. Tanimura, and K. Hata, quant-ph/0411153 (unpublished).
[57] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Phys. Rev. A **52**, 3457 (1995).
[58] A. Saito, K. Kioi, Y. Akagi, N. Hashizume, and K. Ohta, quant-ph/0001113 (unpublished).
[59] T. Yamamoto, Y. A. Pashkin, O. Astafiev, Y. Nakamura, and J. S. Tsai, Nature (London) **425**, 941 (2003).
[60] A. K. Spörl, T. Schulte-Herbrüggen, S. J. Glaser, V. Bergholm, M. J. Storcz, J. Ferber, and F. K. Wilhelm, quant-ph/0504202 (unpublished).