

ARTICLE OPEN

Quantum gate learning in qubit networks: Toffoli gate without time-dependent control

Leonardo Banchi¹, Nicola Pancotti^{2,3} and Sougato Bose¹

We put forward a strategy to encode a quantum operation into the unmodulated dynamics of a quantum network without the need for external control pulses, measurements or active feedback. Our optimisation scheme, inspired by supervised machine learning, consists in engineering the pairwise couplings between the network qubits so that the target quantum operation is encoded in the natural reduced dynamics of a network section. The efficacy of the proposed scheme is demonstrated by the finding of uncontrolled four-qubit networks that implement either the Toffoli gate, the Fredkin gate or remote logic operations. The proposed Toffoli gate is stable against imperfections, has a high fidelity for fault-tolerant quantum computation and is fast, being based on the non-equilibrium dynamics.

npj Quantum Information (2016) 2, 16019; doi:10.1038/npjqi.2016.19; published online 19 July 2016

INTRODUCTION

Computational devices based on the laws of quantum mechanics hold promise to speed up many algorithms known to be hard for classical computers.¹ The implementation of a full-scale computation with existing technology requires an outstanding ability to maintain quantum coherence (i.e., isolation from the environment) without compromising the ability to control the interactions among the qubits in a scalable way. Among the most successful paradigms of quantum computation, there is the ‘circuit model’, in which the algorithm is decomposed into an universal set of single- and two-qubit gates,² and, to some extent, the so-called adiabatic quantum computation (AQC),³ in which the output of the algorithm is encoded in the ground state of an interacting many-qubit Hamiltonian. A different approach⁴ is based on the use of always-on interactions, naturally occurring between physical qubits, to accomplish the computation. Compared with the circuit model, this scheme has the advantage of requiring minimal external control and avoiding the continuous switch off and on of the interactions between all but two qubits, whereas compared with AQC it has the advantage of being faster, being based on the non-equilibrium evolution of the system. Quantum computation with always-on interactions is accomplished by combining the natural couplings with a moderate external control, e.g., with a smooth shifting of Zeeman energies,⁵ via feed-forward techniques,⁶ using measurement-based computation⁷ or quantum control.^{8,9} Most of these approaches are based on the assumption that the natural couplings are fixed by nature and not tunable, whereas local interactions can be modulated with external fields. However, the amount of external control required can be minimised if the couplings between the qubits can be statically tuned¹⁰—e.g., during the creation of the quantum device.

The recent advances in the fabrication of superconducting quantum devices has opened up to the realisation of interacting quantum networks. In a superconducting device, the qubits are built with a Josephson tunnel element, an inductance and a

capacitor,¹¹ whereas local operations and measurements are performed by coupling the qubit to a resonator.¹² The interactions can be designed using lithographic techniques by jointly coupling two qubits via a capacitor¹³ or an inductance,¹⁴ and can be modelled via an effective two-body Hamiltonian $\sum_{\alpha} J_{\alpha} \sigma_{\alpha} \otimes \sigma_{\alpha}$ ^{15,16} where σ_{α} are the Pauli matrices. Because of the flexibility in wiring the pairwise interactions among the qubits, it is possible to arrange them in a planar graph structure, namely a collection of vertices and links, in which the vertices correspond to the qubits and the links correspond to the two-body interactions between them. Moreover, thanks to the development of three-dimensional superconducting circuits,¹⁷ it may be possible in the near future to wire also non-planar configurations, namely a general qubit network.

Motivated by the above, we ask the following question: is it possible to encode a quantum algorithm into the unmodulated dynamics of a suitably large quantum network of pairwise interacting qubits? This would be extremely interesting, as it would enable quantum computation by simply ‘waiting’, without the need of continuously applying external control pulses or measurements. Even when sequential operations cannot be avoided, our scheme can enable the in-hardware implementation of recurring multi-qubit operations of a quantum algorithm (see e.g., Figure 1), such as quantum arithmetic operations,¹⁸ and possibly also the quantum Fourier transform or error-correcting codes.¹ We focus on two-body interactions, as they are the most common in physical setups, and we consider an enlarged network in which auxiliary qubits enrich the quantum dynamics. The important question analysed in this paper is as follows: given a target unitary operation U_Q on a given set of qubits Q , we consider an extended network QUA in which A is a set of auxiliary qubit (ancillae), and we ask whether it is possible to engineer the pairwise interactions in QUA , modelled by the time-independent Hamiltonian H_{QA} , such that $e^{iH_{QA}t} = U_Q \otimes V_A$ after some time t (V_A may be an extra unitary operation on the auxiliary space). More generally, the target operation can depend also on the ancillae

¹Department of Physics and Astronomy, University College London, London, UK; ²Theory Division, Max-Planck-Institut für Quantenoptik, Garching, Germany and ³Dipartimento di Fisica,

Sapienza Università di Roma, Roma, Italy.

Correspondence: L Banchi (benkj.it@gmail.com)

Received 26 November 2015; revised 5 April 2016; accepted 24 May 2016

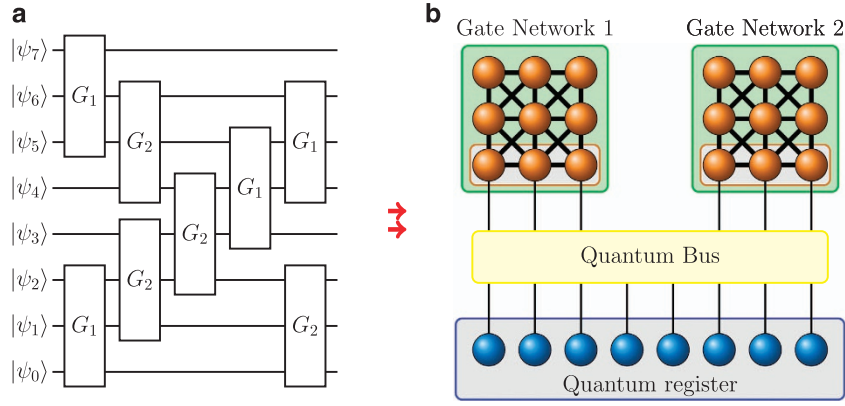


Figure 1. Schematic transposition of a quantum circuit to a trained quantum network. **(a)** An example seven-qubit circuit in which the gates G_1 and G_2 are sequentially applied. Many important circuits can be cast into the scheme **a**, such as those for quantum arithmetics.¹⁸ **(b)** Quantum network implementation of **a**: each qubit network in the green boxes implements either G_1 or G_2 on the input/output qubits (the three qubits in the bottom row). The quantum bus sequentially transfers the state of three qubits from the register, on which the gate G_1 or G_2 has to be applied, to the input/output qubits of the gate network. After the gate network has implemented its transformations, the state is transferred back to the original three qubits of the register.

initial state: if $e^{iH_Q} = \sum_n U_Q^{(n)} \otimes |A_n\rangle\langle A_n|$, where $\{|A_n\rangle\}$ forms a basis of the ancilla Hilbert space and, e.g., $U_Q = U_Q^{(1)}$, then the target operation is implemented when A is initialised in $|A_1\rangle$. Our method is particularly useful for implementing quantum gates, which requires k -body interactions ($k > 2$), such as the Toffoli or Fredkin gates^{1,19,20} where $U_Q \neq e^{iH_Q}$ for any two-local H_Q , and for remote logic, namely for applying a gate to qubits that are not directly connected but are rather interacting via intermediate systems. Our approach is completely different from the simulation of k -local Hamiltonians with pairwise interactions discussed in the AQCL literature,^{21,22} being based on the unmodulated dynamics. Moreover, being based on unmodulated (time-independent) interactions and ancillary qubits, it is significantly different from quantum optimal control.²³

Our quantum network design procedure is inspired by supervised learning in feed-forward networks,²⁴ in which the training procedure involves the optimisation of the network couplings (i.e., the weights between different nodes) such that the output corresponding to some input data has a desired functional form (e.g., for data classification). Although there are many recent developments about using a quantum device to speed up machine learning algorithms^{25–29} or storing data,³⁰ our optimisation procedure is entirely classical, but specifically developed for quantum hardware design. Our scheme is completely different from other recent proposals^{31–33} because it avoids measurements or active feedbacks and requires minimal external control.

RESULTS

Supervised quantum network design

Supervised learning is all about function approximation: given a training set $\{(I_1, O_1), (I_2, O_2), \dots\}$, namely a collection of inputs I_k and the corresponding known outputs O_k , the goal is to find a function f with two desired properties: (i) $O_k \simeq f(I_k)$ for any training pair, and (ii) f should be able to infer the unknown output of an input not contained in the training set. In classical feed-forward networks, the function f is approximated with a directed graph organised in layers, where the first layer is the input register and the last one encodes the output. The value $s_k^{(\ell)}$ of the k -th node in layer ℓ is updated via the equation $s_k^{(\ell)} = A^\ell[\sum_j \lambda_{kj}^{(\ell-1)} s_j^{(\ell-1)}]$, where A^ℓ is an appropriate (typically non-linear) activation function and

$\lambda_{kj}^{(\ell-1)}$ is the weight between node k in layer ℓ and node j in $\ell-1$. The training procedure consists in finding the optimal weights λ by minimising a suitable cost function such as $\mathcal{C} = \sum_k |O_k - f(I_k)|^2$.

A quantum network consists, on the other hand, of an undirected graph (V, E) of vertices V and links E described by a two-local Hamiltonian

$$\mathcal{H} = \sum_{(n,m) \in E} \sum_{\alpha, \beta} J_{nm}^{\alpha\beta} \frac{\sigma_n^\alpha \sigma_m^\beta}{4} + \sum_{n \in V} \sum_{\alpha} h_n^\alpha \frac{\sigma_n^\alpha}{2}, \quad (1)$$

where σ_n^α , $\alpha = x, y, z$, are the Pauli matrices acting on qubit n and, to simplify the notation, we call $\lambda = \{J_{nm}^{\alpha\beta}, h_n^\alpha\}$ the set of parameters. The vertices are composed of two disjoint sets $V = Q \cup A$, where Q consists of register qubits and A consists of auxiliary qubits. Given a separable initial state $|\psi_Q\rangle \otimes |\psi_A\rangle$, the time evolution according to Hamiltonian (1) generates a quantum channel¹ $\mathcal{E}_\lambda[|\psi_Q\rangle\langle\psi_Q|] = \text{Tr}_A[e^{-i\mathcal{H}\tilde{t}} |\psi_Q\rangle\langle\psi_Q| \otimes |\psi_A\rangle\langle\psi_A| e^{i\mathcal{H}\tilde{t}}]$ on subsystem Q —as we are interested in a fixed operational time \tilde{t} for simplicity we set $\tilde{t} = 1$, reabsorbing \tilde{t} into the definition of the definition of \mathcal{H} . Depending on the flexibility of the experimental apparatus in reliably initialising the auxiliary qubits, one can add $|\psi_A\rangle$ to the set λ . Network design consists in the following procedure: given a target unitary operation U_Q that we want to implement, the goal is to find the parameters λ , if they exist, such that $\mathcal{E}_\lambda[\rho_Q] = U_Q \rho_Q U_Q^\dagger$ for any ρ_Q . To simplify the notation, we assume that the gate output is encoded in Q but it is straightforward to generalise the formalism when the output sites differ from the input ones.

Motivated by the similarity with classical supervised learning, where the weights λ are tuned to maximise the ability of the network to reproduce a known output given the corresponding input, we create a training set \mathcal{T} with a random set of initial input states. For each input $|\psi\rangle \in \mathcal{T}$ the expected known output is $U_Q |\psi\rangle$, whereas the output of the network evolution is $\mathcal{E}_\lambda[|\psi\rangle\langle\psi|]$. The ‘learning’ procedure involves the minimisation of the difference between the output of the network and the expected output, and corresponds to the maximisation of the fidelity

$$\mathcal{F} = \sum_{|\psi\rangle \in \mathcal{T}} \frac{\mathcal{F}_\psi}{|\mathcal{T}|}, \quad \mathcal{F}_\psi = \langle \psi | U_Q^\dagger \mathcal{E}_\lambda[|\psi\rangle\langle\psi|] U_Q | \psi \rangle. \quad (2)$$

If the average is performed over all possible states, then equation (2) can be substituted by the average gate fidelity $\overline{\mathcal{F}} = \int \mathcal{F}_\psi d\psi$

where the formal integration can be explicitly evaluated^{10,34,35} yielding

$$\overline{\mathcal{F}} = \frac{1}{D+1} + \frac{1}{D(D+1)} \sum_{ijkl} U_{ik}^* \mathcal{E}_\lambda^{ij,kl} U_{jl}, \quad (3)$$

where $\mathcal{E}_\lambda^{ij,kl} = \langle q_i | \mathcal{E}_\lambda [|q_k\rangle \langle q_l|] |q_j\rangle$, $U_{ij} = \langle q_i | U_Q |q_j\rangle$ and $\{|q_j\rangle\}$ form the computational basis of the D -dimensional Hilbert space of qubits Q . The typical value of the fidelity for a random non-optimal evolution of the qubit network is $\overline{\mathcal{F}} = D^{-1}$, obtained using Haar integration techniques.³⁶ This value is independent of the details of the ancillae, as it depends only on the dimension of the target Hilbert space, and provides an estimate for the initial fidelity of an untrained network.

The gate-learning procedure corresponds to a global maximisation of the fidelity (3). However, because of the many parameters in the Hamiltonian (1), $\overline{\mathcal{F}}$ can have many local maxima, making the global optimisation extremely complicated. As most global optimisation algorithms introduce stochastic strategies, rather than introducing unphysical random jumps, we take advantage of the explicit stochastic nature of the problem ($\overline{\mathcal{F}}$ is a uniform average over random states) and we propose the following learning algorithm to design the interactions of the quantum network.

- 1: Choose an initial parameter set λ (e.g., at random), and choose an initial learning rate ϵ ;
- 2: Repeat
- 3: Generate a random $|\psi\rangle$;
- 4: Update L times the coupling strengths as

$$\lambda \rightarrow \lambda + \epsilon \nabla_\lambda \langle \psi | U_Q^\dagger \mathcal{E}_\lambda [| \psi \rangle \langle \psi |] U_Q | \psi \rangle; \quad (4)$$

- 5: Decrease ϵ (see Materials and methods);
- 6: Until convergence (or maximum number of operations).

Specifically, we combine the above algorithm with the maximisation of the average fidelity (see below) and we observe a drastic speed up of the optimisation process. The parameter L tunes the number of deterministic steps in the learning procedure, and can be set to the minimum value $L=1$, so that after each interaction the state is changed or to a higher value. In our simulations, we use $L=1$, for simplicity. Our algorithm is an application of the stochastic gradient descent (SGD) method³⁷ to the maximisation of the function (2). In classical feed-forward networks, SGD is the *de facto* standard algorithm for network training^{24,37} and is specifically used for large training data sets, when the evaluation of the cost function and its gradient are computationally intensive. On the other hand, the average in equation (2) can be evaluated explicitly over a uniform distribution of an infinite number of initial states, giving equation (3). Although \mathcal{F}_ψ is easier than $\overline{\mathcal{F}}$ to compute, the major advantage of SGD for quantum network design comes from its ability to escape local maxima. The crucial observation to show the latter point is that the statistical variance over random states $\text{Var} \mathcal{F} = \overline{\mathcal{F}^2} - \overline{\mathcal{F}}^2$ vanishes when $\overline{\mathcal{F}} = 1$ (see e.g., refs 34 and 35)—indeed, intuitively, as both $\overline{\mathcal{F}}$ and \mathcal{F}_ψ are bounded in $[0, 1]$, $\overline{\mathcal{F}}$ can achieve its maximum only if $\mathcal{F}_\psi = 1$ for all the states, apart from a set of measure zero. On the other hand, if $0 < \overline{\mathcal{F}} < 1$, then $\text{Var} \mathcal{F} > 0$ and the fluctuations can be so high that a local maximum of $\overline{\mathcal{F}}$ may not correspond to a maximum of \mathcal{F}_ψ for some state ψ . This is indeed shown in Figure 2 with a real example for the implementation of the Toffoli gate (see the application section below). In Figure 2, the average fidelity $\overline{\mathcal{F}}$ has three local maxima at λ_k^{loc} ($k=1, 2, 3$) and a single global maximum at λ^{gl} , namely the optimal parameters, whereas the fidelities \mathcal{F}_ψ for different random states ψ have a more complicated behaviour. In view of the argument discussed above, all the state fidelities \mathcal{F}_ψ have a global maximum at λ^{gl} , whereas, remarkably, at least one fidelity \mathcal{F}_ψ has no local maximum at λ_k^{loc} . Our stochastic learning algorithm uses a gradient descent technique for locally maximising the function




Figure 2. Average fidelity $\overline{\mathcal{F}}$ and fidelity \mathcal{F}_ψ for some random states ψ for implementing the Toffoli gate in a four-qubit network (see discussion in the text). All the parameters are set to the optimal ones, except J_{35}^{xx} in the abscissa. The region around the only global peak is filled in green.

$\mathcal{F}_\psi(\lambda)$. Therefore, if we are around the slopes of a local maximum of $\mathcal{F}_\psi(\lambda)$ (say λ_k^{loc} from the previous example) and the state $|\psi\rangle$ is randomly changed to $|\phi\rangle$, that local maximum may disappear from $\mathcal{F}_\psi(\lambda)$, allowing the algorithm to escape from this non-optimal region when the parameters are updated via equation (4). On the other hand, when the algorithm is probing the neighbourhood of a true optimal point for which $\overline{\mathcal{F}}(\lambda) = 1$ (e.g., λ^{gl} in the previous example), then the maximum of $\mathcal{F}_\psi(\lambda)$ does not disappear when the state ψ is changed, allowing the ‘climbing’ procedure to continue.

The above stochastic algorithm may be combined with a deterministic maximisation of equation (3). In our simulations, we use stochastic learning for the initial global span of the parameter manifold, and if it reaches a suitably high fidelity (e.g., $\overline{\mathcal{F}} > 95\%$) then it is reasonable to suppose that the algorithm has found a global maximum. Starting from this point, we perform a local maximisation of equation (3), and if $\overline{\mathcal{F}} \simeq 1$ is reached the learning has been successful. Otherwise, we repeat the procedure.

It is worth emphasising that given a target gate U it is an open question to understand *a priori* whether a solution may exist for a graph with a certain set of interactions (e.g., Heisenberg, Ising and so on). Unlike in quantum control, where given a time-dependent Hamiltonian $\mathcal{H}(t) = \mathcal{H}_0 + \beta(t)\mathcal{V}$ one can check in advance whether $U = T[\exp(-i \int_0^1 \mathcal{H}(t) dt)]$ for some control profile $\beta(t)$: such a profile can exist only if U is contained in the group associated with the algebra generated by the repeated commutators of \mathcal{H}_0 and \mathcal{V} . Although no complete algebraic characterisation is known for our case (see, however, the Materials and methods for a necessary condition) and we have to study each problem numerically, in the next sections we find some structures that enable the implementation of important quantum gates. All numerical simulations have been obtained in a laptop computer using QuTiP.³⁸

Application: Toffoli gate

The Toffoli gate is a key component for many important quantum algorithms, notably the Shor algorithm,³⁹ quantum error correction,²⁰ fault-tolerant computation⁴⁰ and quantum arithmetic operations,¹⁸ and, together with the Hadamard gate, is universal for quantum computation.⁴¹ Experimental implementations of this gate have been obtained with trapped ions,⁴² superconducting circuits^{19,43} or photonic architectures.⁴⁴ Toffoli gate is a controlled-controlled-not (CCNOT) operation acting on three qubits. It can be implemented in a circuit using five two-qubit gates,¹ or it can be obtained in coupled systems via quantum control techniques.^{45,46}




Figure 3. Network implementing the Toffoli gate. The gate acts on the three external qubits (the top ones being the control qudits, and the bottom one being the target), and has an additional auxiliary qubit in the centre.

Efficient schemes require higher dimensional system (i.e., qudits).⁴⁴ On the other hand, the direct implementation using natural interactions is complicated, as the Hamiltonian $\mathcal{H}_{\text{CCNOT}}$ corresponding to the gate, i.e., $\text{CCNOT} = e^{i\mathcal{H}_{\text{CCNOT}}}$, has three-body interactions, which are unlikely to appear in nature.

By applying our quantum hardware design procedure, we show that the Toffoli gate can be implemented in a four-qubit network using only pairwise interactions and constant control fields. Our findings enable the construction of a device that implements the Toffoli gate with a fidelity $\overline{\mathcal{F}} = 99.98\%$ by simply ‘waiting’ for the natural dynamics to occur, without the need for external control pulses. We consider a four-qubit network as displayed in Figure 3, in which the control qubits are labelled by the indices 1,2, the target is qubit 3 and the ancilla is qubit 4. We start our analysis by considering a fully connected graph in which each qubit interacts with the others using XX - and ZZ -type pairwise interactions, as this kind of interaction can be obtained in superconducting circuits.¹⁵ Because of the symmetries of the Toffoli gate (see Materials and methods), we consider the two control qubits to be equally coupled to the target and the ancilla: $J_{1m}^{ab} = J_{2m}^{ab}$, for $m=3, 4$ and similarly we set $h_1^a = h_2^a$. Moreover, as the Toffoli gate is real, we only consider local fields in the X and Z directions and set $|\psi_A\rangle = \cos\eta|\uparrow\rangle + e^{i\xi}\sin\eta|\downarrow\rangle$. By combining SGD with the maximisation of equation (3), we find the following optimal parameters,

$$\begin{aligned} J_{12}^{zz} &= -8.940, & J_{13}^{zz} &= -4.957, & J_{14}^{zz} &= -5.657, \\ h_1^z &= -2.428, & h_3^z &= J_{13}^{zz}, & h_4^z &= -0.165, \\ h_3^x &= -19.08, & h_4^x &= -4.267, & J_{34}^{xx} &= 15.06, \\ \eta &= 0.8182, & \xi &= 0.0587, \end{aligned} \quad (5)$$

in which the other XX - and ZZ -type interactions not displayed in equation (5) are found to be zero by the learning algorithm, so the optimal configuration is the one summarised in Figure 3, where the XX coupling is only between qubits 3 and 4. In more physical terms, if the maximal allowed coupling is fixed to $J/2\pi \approx 40$ MHz, then we find a gate time of 60 ns and

$$\begin{aligned} J_{12}^{zz} &= -149.2 \text{ MHz}, & J_{13}^{zz} &= -82.71 \text{ MHz}, & J_{14}^{zz} &= -94.39 \text{ MHz}, \\ h_1^z &= -40.52 \text{ MHz}, & h_3^z &= J_{13}^{zz}, & h_4^z &= -2.751 \text{ MHz}, \\ h_3^x &= -318.4 \text{ MHz}, & h_4^x &= -71.2 \text{ MHz}, & J_{34}^{xx} &= 251.3 \text{ MHz}. \end{aligned} \quad (6)$$

With the optimal parameters of equations (5) and (6), we obtain an average gate fidelity of 99.98%, above the threshold for topological fault tolerance for single- and two-qubit gates, whereas by avoiding the extra phase fixing $\xi=0$ we still obtain $\overline{\mathcal{F}} = 99.92\%$. Moreover, our gate fidelity is above the Toffoli gate

accuracy threshold ($755/756 \simeq 99.87\%$) for fault-tolerant computation in the limit in which Clifford gate errors are negligible.⁴⁷

The optimal parameters (5) and (6) are stable against an imperfect tuning of the interactions. Indeed, we considered a perturbation $\lambda_k \rightarrow \lambda_k + \epsilon r_k$, $r_k \in [0, 1]$ being a random number and ϵ being the strength of the static perturbation, and found that $\overline{\mathcal{F}} > 99.9\%$ if $\epsilon < 0.04$ ($\epsilon < 0.7$ MHz) and $\overline{\mathcal{F}} > 99\%$ if $\epsilon < 0.18$ ($\epsilon < 3$ MHz).

Application: Fredkin gate

Fredkin gate is a controlled-swap (CSWAP) operation acting on three qubits, which is universal for reversible computation.¹ We found that this gate can be obtained with perfect fidelity (up to the numerical precision) in a four-qubit network with Hamiltonian (1) where $J_{12}^{xx} = J_{13}^{xx} = 13.60$ (227.0 MHz), $J_{23}^{aa} = -4.712$ (-78.62 MHz), $J_{24}^{xx} = J_{34}^{xx} = 8.400$ (140.2 MHz), $J_{12}^{zz} = J_{13}^{zz} = 11.15$ (186.1 MHz), $h_4^z = 1.025$ (17.11 MHz), $h_1^z = \pi$ (54.42 MHz). The values in MHz correspond to a gate time of 60 ns. Moreover, $e^{-i\mathcal{H}} = \text{CSWAP}_{123} \otimes U_4$ so the gate is independent of the initial state of the ancilla. As for the Toffoli gate, this optimal configuration has been obtained by starting the training procedure with a fully connected graph with all the interactions, and thus the fact that some interactions are zero is a result of the optimisation process.

Application: remote logic

We study a qubit network that implements a maximally entangling gate between two sites that are not directly coupled. Remote logic has been studied extensively in spin chains for achieving entangling operations between the boundary sites,^{4,10,48,49} and it is a building block for a proposed architecture for solid-state quantum computation at room temperature.⁵⁰ For simplicity, we consider a $SU(2)$ invariant four-qubit network, interacting with a Heisenberg Hamiltonian $H = \sum_{i \neq j=1}^4 \sum_{\alpha=x,y,z} J_{ij} \sigma_i^\alpha \sigma_j^\alpha / 4$ where there is no direct coupling between qubits 1 and 4 ($J_{14} = 0$). Applying our learning algorithm, we found that the $\sqrt{\text{SWAP}}$ gate, which is universal for quantum computation when paired with single qubit operations,¹ can be achieved between qubits 1 and 4 with unit fidelity with different choices of $J_{12} = J_{24}$, $J_{13} = J_{34}$ and J_{23} when the initial state of ancilla is $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) / \sqrt{2}$. Given this simplification, one can then find a solution analytically: $J_{12} = a + \pi \sqrt{(2n)^2 - 1} / \sqrt{8}$, $J_{13} = a - \pi \sqrt{(2n)^2 - 1} / \sqrt{8}$, $J_{23} = a + (-1)^n \pi$, where n is an integer. We find analytically that irrespective of a the above choice gives perfect fidelity. Our strategy has not found any three-qubit configurations that implement a remote $\sqrt{\text{SWAP}}$ gate, and thus the four-qubit network is the minimal non-trivial example. Remarkably, some of our four-qubit configurations are more stable to noise than the direct implementation of the gate in a two-qubit system (namely when $J_{14} = \pi/2$ and the other couplings are zero). For instance, if $J_{ij} = J_{ij}^{\text{optimal}} + \epsilon$, J_{ij}^{optimal} being the optimal value for implementing the gate, we found that when ϵ is randomly distributed in $[0, 1/2]$, then the four-qubit system with $n=1$ still has, on average, $\overline{\mathcal{F}} \simeq 99.1\%$, whereas the direct two-qubit case has $\overline{\mathcal{F}} \simeq 98.8\%$.

Towards a scalable architecture for quantum computation

Current architectures for quantum computation, e.g., with superconducting qubits⁵¹ or ion traps,⁵² are based on an arrays of interacting qubits that are continuously controlled via external pulses to implement the desired operation. This approach may suffer from scalability issues because, even assuming the ability to maintain quantum coherence for a long time, extremely large (classical) control units will be necessary to generate the

sophisticated pulse sequences required to implement a full-scale quantum algorithm. On the other hand, the approach that we have in mind shares more similarities with integrated circuits in present-day electronics, where a set of special-purpose logic units (modules) are wired together to achieve computation (or other tasks). In our vision, different modules can be fabricated with qubit networks designed to produce a specific logic task, namely a quantum gate, automatically without the need of external control. As in Figure 1, the different logic and memory units can be reciprocally connected using a quantum bus, whose purpose is to transfer the qubit states between the quantum registers/memory and the input/output qubits of the modules. In Figure 1, for simplicity the input/output ports of the modules are designed in the same physical qubits, although this can be easily extended to more general cases. The quantum bus can be realised with different technologies, e.g., with microwave resonators,⁵³ or it can also be implemented via quantum-state transfer in a qubit network.⁵⁴ The modules shown in Figure 1 can be designed to produce either simple basic operations, such as the CNOT or the Toffoli gate, or, in principle, they can directly implement larger components of a quantum algorithm such as the Quantum Fourier Transform or error-correcting codes.¹ In this respect, to treat systems with many parameters, one can easily combine our optimisation strategy, based on fidelity statistics, with metaheuristic strategies,⁵⁵ which simultaneously deals with many candidate solutions and are known to be fast in global optimisation with high-dimensional parameter spaces. Moreover, highly optimised deep learning algorithms are already used to train neural networks with 60 millions of parameters.⁵⁶ However, given the difficulty in numerically simulating large quantum systems, this approach may be reasonable for networks up to, say, 20–30 qubits.

DISCUSSION

Inspired by classical supervised learning, we have proposed an optimisation scheme to encode a quantum operation into the unmodulated dynamics of a qubit register, which is part of a bigger network of pairwise interacting qubits. Our strategy is based on the static engineering of the pairwise couplings, and it enables the creation of a quantum device, which implements the desired operation by simply waiting for the natural dynamics to occur, without the need of external control pulses. Our findings show that machine learning-inspired techniques can be combined with quantum mechanics not only for data classification speed up^{25,26} or quantum black-box certification,^{57,58} but also for quantum hardware design.

This paper opens up the topic of encoding quantum gates and operations into the unmodulated dynamics of qubit networks. Although we have focused on small systems, larger networks can be considered using more efficient training schemes. These would enable the simulation of larger components of a quantum algorithm, as different multi-qubit gates can be combined into a unique quantum operation, which can be simulated in a large quantum network. Moreover, when combined with a quantum bus as in Figure 1, our strategy can provide an alternative approach to universal quantum computation, which avoids the decomposition of the algorithm into one- and two-qubit gates. Note that most quantum algorithms take classical inputs, and thus the extra control required for initialisation demands the further ability to fully polarise globally the spins. The latter step is, however, typically much easier than the implementation of entangling gates, which has been considered in this paper. Moreover, in view of the recent experimental measures of the average gate fidelity,⁵⁹ it is tempting to predict an all-quantum version of our learning procedure, where \mathcal{F} is not classically simulated, but rather directly measured. This would require a further highly controlled system to infer the optimal parameters of an uncontrolled quantum network, which can be used to

industrialise the production of unmodulated quantum devices implementing the desired algorithm.

Our results demonstrate the efficacy of the proposed scheme in designing four-qubit networks that implement the Toffoli and Fredkin gates or remote logic operations. The proposed Toffoli gate is fast, has high fidelity for fault-tolerant computation and only uses static XX - and ZZ -type interactions, which can be achieved in superconducting systems.¹⁵ The key advantage of our method is in exploiting all the permanent interactions in the qubit network without trying to suppress some of them sequentially to implement pairwise gates. Moreover, being based on non-equilibrium dynamics, our gate is fast: if $J/2\pi \approx 40$ MHz, then the total operation time is ~ 60 ns, which matches the current gate times for single- and two-qubit operations.⁵¹

MATERIALS AND METHODS

Learning rate

The choice of the learning rate ϵ is crucial. If the initial learning rate is too small, it might not escape from the different 'local maximum' points, whereas if it is too large it will continue to randomly jump without even seeing the local maxima. To maximise the speed and precision of SDG, the learning rate ϵ has to decrease as a function of the steps, a common choice being $\epsilon \propto m^{-1/2}$ where m is the step counter.³⁷ However, when the gradient in equation (4) cannot be performed analytically, one can use more sophisticated techniques⁶⁰ in which both the learning rate and the finite difference approximation of the gradient change as a function of m .

Symmetries

In the design of the quantum network and its couplings, the number of parameters can be drastically reduced if the target unitary operation U_Q has some symmetries, namely if there exists some unitary matrix S such that $[U_Q, S_Q] = 0$. This condition requires the quantum channel $\mathcal{E}_\lambda[\rho] = \text{Tr}_A[e^{-iH}\rho \otimes \rho_A e^{iH}]$ to satisfy $\mathcal{E}_\lambda(S_Q \rho S_Q^\dagger) = S_Q \mathcal{E}_\lambda(\rho) S_Q^\dagger$ for each state ρ , e.g., $[H, S_Q \otimes \mathbb{1}_A] = 0$. Conversely, if the interaction type is fixed by nature (for instance, only Ising or Heisenberg interactions are allowed), then one has to check whether the Lie algebra spanned by the operators in \mathcal{H} contains the generators of U_Q .

Bottom-up construction: Lie algebraic characterisation

All the numerical results presented in the main text are obtained using a top-down approach: after selecting the interaction types (e.g., XX , ZZ , Heisenberg and so on), the algorithm starts with a zero-bias fully connected configuration in which all the qubit pairs of the network interact with all possible interactions, each weighted with a different parameter, and different local fields. As a result of the training procedure, we found numerically that most of these parameters are indeed zero. However, for larger networks it is better to use a bottom-up approach, in which one starts with a minimal set of parameters and then adds other parameters until a solution is found.

To construct a minimal set of parameters, one can use a Lie algebraic characterisation inspired by quantum control. We write the Hamiltonian as $\mathcal{H} = \sum_j \lambda_j \mathcal{O}_j$, where λ_j is the independent parameter and \mathcal{O}_j is the operator. If the parameters are time dependent, then there exist suitable pulses $\lambda_j(t)$ such that the dynamics implements the target gate G only if $\log(G)$ is contained in the algebra generated by the repeated commutators $[\mathcal{O}_j, [\mathcal{O}_k, \dots]]$. As our scheme is based on the particular choice where $\lambda_j(t)$ is constant, the above characterisation still provides a necessary condition. As an example, we consider the Toffoli gate and the solution equation (5) where $\mathcal{O}_1 = \sigma_1^z \sigma_2^z$, $\mathcal{O}_2 = \sigma_1^z \sigma_3^z + \sigma_2^z \sigma_3^z + 2\sigma_3^z$, $\mathcal{O}_3 = \sigma_1^z \sigma_4^z + \sigma_2^z \sigma_4^z$, $\mathcal{O}_4 = \sigma_3^z \sigma_4^z$, $\mathcal{O}_5 = \sigma_1^z + \sigma_2^z$, $\mathcal{O}_7 = \sigma_4^z$, $\mathcal{O}_8 = \sigma_3^z$, $\mathcal{O}_9 = \sigma_4^x$. It is simple to check that $\log(G)$ (up to an irrelevant constant factor) is contained in the algebra generated by the operators \mathcal{O}_j , whereas this is not the case if the operator \mathcal{O}_8 is removed from the Hamiltonian. Therefore, no solution is possible if $\lambda_8 \equiv 0$.

Inspired by the above example, the bottom-up approach consists in the following steps: (i) based on the symmetries of the target gate and on the physically allowed interactions, one defines an initial set of operators; (ii) other operators are added to the set until the dynamical algebra contains $\log(G)$; (iii) one starts the numerical parameter training to check for convergence (different runs may be required). Until the solution is found, one can either add new operators or changes the previous ones.

ACKNOWLEDGEMENTS

L.B. and S.B. acknowledge the financial support by the ERC under Starting Grant 308253 PACOMANEDIA. We thank P. Wittek, A. Monras and J.I. Cirac for their valuable comments and suggestions.

COMPETING INTERESTS

The authors declare no conflict of interest.

REFERENCES

- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2000).
- Barenco, A. et al. Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457 (1995).
- Aharonov, D. et al. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Rev.* **50**, 755–787 (2008).
- Benjamin, S. C. & Bose, S. Quantum computing with an always-on heisenberg interaction. *Phys. Rev. Lett.* **90**, 247901 (2003).
- Benjamin, S. C. & Bose, S. Quantum computing in arrays coupled by “always-on” interactions. *Phys. Rev. A* **70**, 032314 (2004).
- Satoh, T. et al. Scalable quantum computation architecture using always-on ising interactions via quantum feedforward. *Phys. Rev. A* **91**, 052329 (2015).
- Li, Y., Browne, D. E., Kwek, L. C., Raussendorf, R. & Wei, T.-C. Thermal states as universal resources for quantum computation with always-on interactions. *Phys. Rev. Lett.* **107**, 060501 (2011).
- Burgarth, D. et al. Scalable quantum computation via local control of only two qubits. *Phys. Rev. A* **81**, 040303 (2010).
- Müller, M. et al. Optimizing entangling quantum gates for physical systems. *Phys. Rev. A* **84**, 042315 (2011).
- Banchi, L., Bayat, A., Verrucchi, P. & Bose, S. Nonperturbative entangling gates between distant qubits using uniform cold atom chains. *Phys. Rev. Lett.* **106**, 140501 (2011).
- Devoret, M. & Schoelkopf, R. Superconducting circuits for quantum information: an outlook. *Science* **339**, 1169–1174 (2013).
- Wallraff, A. et al. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431**, 162–167 (2004).
- Barends, R. et al. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.* **111**, 080502 (2013).
- Chen, Y. et al. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.* **113**, 220502 (2014).
- Geller, M. R. et al. Tunable Coupler for Superconducting Xmon Qubits: Perturbative Nonlinear Model. Preprint at <http://arxiv.org/abs/1405.1915> (2014).
- Neeley, M. et al. Generation of three-qubit entangled states using superconducting phase qubits. *Nature* **467**, 570–573 (2010).
- Paik, H. et al. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Phys. Rev. Lett.* **107**, 240501 (2011).
- Vedral, V., Barenco, A. & Ekert, A. Quantum networks for elementary arithmetic operations. *Phys. Rev. A* **54**, 147 (1996).
- Fedorov, A., Steffen, L., Baur, M., Da Silva, M. & Wallraff, A. Implementation of a toffoli gate with superconducting circuits. *Nature* **481**, 170–172 (2011).
- Cory, D. G. et al. Experimental quantum error correction. *Phys. Rev. Lett.* **81**, 2152 (1998).
- Bravyi, S., DiVincenzo, D. P., Loss, D. & Terhal, B. M. Quantum simulation of many-body hamiltonians using perturbation theory with bounded-strength interactions. *Phys. Rev. Lett.* **101**, 070503 (2008).
- Biamonte, J. D. & Love, P. J. Realizable hamiltonians for universal adiabatic quantum computers. *Phys. Rev. A* **78**, 012352 (2008).
- Brif, C., Chakrabarti, R. & Rabitz, H. Control of quantum phenomena: past, present and future. *New J. Phys.* **12**, 075008 (2010).
- Bishop, C. M. *Pattern Recognition and Machine Learning* (Springer, 2006).
- Wittek, P. *Quantum Machine Learning: What Quantum Computing Means to Data Mining* (Elsevier, 2014).
- Rebentrost, P., Mohseni, M. & Lloyd, S. Quantum support vector machine for big data classification. *Phys. Rev. Lett.* **113**, 130503 (2014).
- Paparo, G. D., Dunjko, V., Makmal, A., Martin-Delgado, M. A. & Briegel, H. J. Quantum speedup for active learning agents. *Phys. Rev. X* **4**, 031002 (2014).
- Wiebe, N., Kapoor, A. & Svore, K. Quantum nearest-neighbor algorithms for machine learning. *Quantum Inf. Comput.* **15**, 0318–0358 (2015).
- Lloyd, S., Mohseni, M. & Rebentrost, P. Quantum principal component analysis. *Nat. Phys.* **10**, 631–633 (2014).
- Rotondo, P., Lagomarsino, M. C. & Viola, G. Dicke simulators with emergent collective quantum computational abilities. *Phys. Rev. Lett.* **114**, 143601 (2015).
- Nagaj, D. Universal two-body-hamiltonian quantum computing. *Phys. Rev. A* **85**, 032330 (2012).
- Bang, J., Lim, J., Kim, M. & Lee, J. Quantum Learning Machine. Preprint at <https://arxiv.org/abs/0803.2976> (2008).
- Gammelmark, S. & Mølmer, K. Quantum learning by measurement and feedback. *New J. Phys.* **11**, 033017 (2009).
- Magesan, E., Blume-Kohout, R. & Emerson, J. Gate fidelity fluctuations and quantum process invariants. *Phys. Rev. A* **84**, 012309 (2011).
- Pedersen, L. H., Møller, N. M. & Mølmer, K. The distribution of quantum fidelities. *Phys. Lett. A* **372**, 7028–7032 (2008).
- Collins, B. & Śniady, P. Integration with respect to the haar measure on unitary, orthogonal and symplectic group. *Commun. Math. Phys.* **264**, 773–795 (2006).
- Bottou, L. in *Online Learning and Neural Networks* (ed. Saad, D.) (Cambridge University Press, Cambridge, UK, 1998).
- Johansson, J., Nation, P. & Nori, F. Qutip 2: A python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.* **184**, 1234–1240 (2013).
- Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484–1509 (1997).
- Dennis, E. Toward fault-tolerant quantum computation without concatenation. *Phys. Rev. A* **63**, 052314 (2001).
- Shi, Y. Both toffoli and controlled-not need little help to do universal quantum computing. *Quantum Inf. Comput.* **3**, 84–92 (2003).
- Monz, T. et al. Realization of the quantum toffoli gate with trapped ions. *Phys. Rev. Lett.* **102**, 040501 (2009).
- Reed, M. et al. Realization of three-qubit quantum error correction with superconducting circuits. *Nature* **482**, 382–385 (2012).
- Lanyon, B. P. et al. Simplifying quantum logic using higher-dimensional hilbert spaces. *Nat. Phys.* **5**, 134–140 (2009).
- Stojanović, V. M., Fedorov, A., Wallraff, A. & Bruder, C. Quantum-control approach to realizing a toffoli gate in circuit qed. *Phys. Rev. B* **85**, 054504 (2012).
- Zahedinejad, E., Ghosh, J. & Sanders, B. C. High-fidelity single-shot toffoli gate via quantum control. *Phys. Rev. Lett.* **114**, 200502 (2015).
- Gaitan, F. *Quantum Error Correction and Fault Tolerant Quantum Computing* (CRC Press, 2008).
- Yao, N. Y. et al. Quantum logic between remote quantum registers. *Phys. Rev. A* **87**, 022306 (2013).
- Banchi, L., Compagno, E. & Bose, S. Perfect wave-packet splitting and reconstruction in a one-dimensional lattice. *Phys. Rev. A* **91**, 052323 (2015).
- Yao, N. Y. et al. Scalable architecture for a room temperature solid-state quantum information processor. *Nat. Commun.* **3**, 800 (2012).
- Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500–503 (2014).
- Monz, T. et al. Realization of a scalable shor algorithm. *Science* **351**, 1068–1070 (2016).
- Mariantoni, M. et al. Implementing the quantum von Neumann architecture with superconducting circuits. *Science* **334**, 61–65 (2011).
- Nikolopoulos, G. M. & Jex, I. *Quantum State Transfer and Network Engineering* (Springer, 2014).
- Storn, R. & Price, K. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J. Global Opt.* **11**, 341–359 (1997).
- Krizhevsky, A., Sutskever, I. & Hinton, G. E. in *Advances in Neural Information Processing Systems* Vol.25 (eds Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q.) 1097–1105 (Curran Associates, Inc., 2012).
- Bisio, A., Chiribella, G., D’Ariano, G. M., Facchini, S. & Perinotti, P. Optimal quantum learning of a unitary transformation. *Phys. Rev. A* **81**, 032324 (2010).
- Wiebe, N., Granade, C., Ferrie, C. & Cory, D. Hamiltonian learning and certification using quantum resources. *Phys. Rev. Lett.* **112**, 190501 (2014).
- Lu, D. et al. Experimental estimation of average fidelity of a clifford gate on a 7-qubit quantum processor. *Phys. Rev. Lett.* **114**, 140505 (2015).
- Spall, J. C. *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control* vol. 65 (John Wiley & Sons, 2005).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2016