



Machine-learning engineering of quantum currents

Tobias Haug ¹, Rainer Dumke,^{1,2,3} Leong-Chuan Kwek,^{1,3,4,5} Christian Miniatura ^{3,1,6,7,8,9} and Luigi Amico^{1,3,10,11,12}¹Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543, Singapore²Division of Physics and Applied Physics, Nanyang Technological University, 21 Nanyang Link, Singapore 637371, Singapore³MajuLab, CNRS-UCA-SU-NUS-NTU International Joint Research Unit, Singapore⁴School of Electrical and Electronic Engineering, 50 Nanyang Avenue, Singapore 637553, Singapore⁵National Institute of Education, Nanyang Technological University, 1 Nanyang Walk, Singapore 637616, Singapore⁶Department of Physics, National University of Singapore, 2 Science Drive 3, Singapore 117542, Singapore⁷School of Physical and Mathematical Sciences, Nanyang Technological University, 637371 Singapore, Singapore⁸Yale-NUS College, 16 College Avenue West, Singapore 138527, Singapore⁹Université Côte d'Azur, CNRS, Institut de Physique de Nice, 1361 route des Lucioles, 06560 Valbonne, France¹⁰Dipartimento di Fisica e Astronomia, Via S. Sofia 64, 95127 Catania, Italy¹¹CNR-MATIS-IMM & INFN-Sezione di Catania, Via S. Sofia 64, 95127 Catania, Italy¹²LANEF "Chaire d'excellence," Université Grenoble-Alpes & CNRS, F-38000 Grenoble, France

(Received 9 January 2020; revised 6 July 2020; accepted 8 November 2020; published 12 January 2021)

The design, accurate preparation, and manipulation of quantum states in quantum circuits are essential operational tasks at the heart of quantum technologies. Nowadays, circuits can be designed with physical parameters that can be controlled with unprecedented accuracy and flexibility. However, the generation of well-controlled current states is still a nagging bottleneck, especially when different circuit elements are integrated together. In this work, we show how machine learning can effectively address this challenge and outperform the current existing methods. To this end, we exploit deep reinforcement learning to prepare prescribed quantum current states in circuits composed of lumped elements. To highlight our method, we show how to engineer bosonic persistent currents as they are relevant in different quantum technologies as cold atoms and superconducting circuits. We demonstrate the use of deep reinforcement learning to rediscover established protocols, as well as solving configurations that are difficult to treat with other methods. With our approach, quantum current states characterized by a single winding number or entangled currents of multiple winding numbers can be prepared in a robust manner, superseding the existing protocols.

DOI: [10.1103/PhysRevResearch.3.013034](https://doi.org/10.1103/PhysRevResearch.3.013034)

With the advent of quantum technologies, new forms of quantum circuits have emerged. The architecture and circuit performance depend on the specific physical implementation and the type of “quantum fluid” operating in the quantum network. Prominent candidates are atomtronic circuits involving neutral matter waves of cold atoms in optically generated structures with micrometric resolution [1]. Other examples range from electronic and superconducting circuits [2] based on charged matter-wave on nanolithography to photonic circuits employing photons in fiber optics [3]. These systems allow for precise control over the circuit properties such as interactions or particle statistics (fermions/bosons). In addition, it is possible with the latest achievements in the field, particularly in atomtronics, to dynamically adjust the spatial features of the circuit locally, while avoiding cross-talk effects [4–6]. Finally, quantum circuits with increasingly complex

architecture and hybrid systems, in which different technologies are interfaced, are at a mature stage of technological readiness [7–11]. Although important for the very definition of the quantum circuits, the generation and control of current states remains a difficult task to achieve. Particularly for cold atom quantum technology, matter-wave currents have been imparted so far only in continuous atomic rings [12]. In fact, currents in even a simple circuit made of a lattice ring have never been achieved. Such a problem represents a well-known bottleneck in the field, particularly urgent for the progress of atomtronics in which neutral matter currents are needed to flow in complex networks [13,14].

In the present work, we demonstrate that the problem can be overcome by applying machine learning. In this way, we can engineer fast and high fidelity current states in circuits with lumped parameters.

Machine learning with deep reinforcement learning (RL) was recently recognized as a powerful tool to engineer dynamics in quantum systems [15–21]. Here, we guide quantum systems by reconfiguring deep RL protocols that have trained artificial intelligence agents to master complex decision processes [22]. We demonstrate this approach to prepare quantum current states describing the flow of a coherent matter-wave

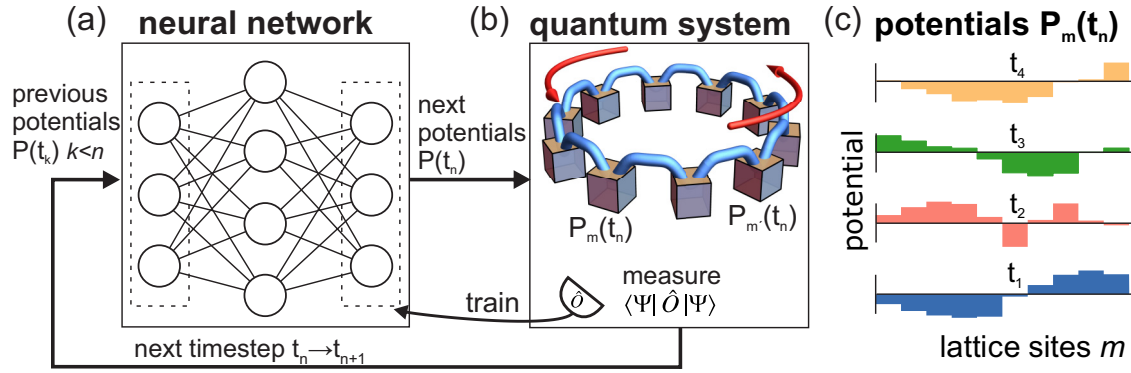


FIG. 1. (a), (b) Deep reinforcement learning to optimize generation of quantum states of currents. The quantum system is a ring-shaped circuit comprising L lumped parameters schematized as lattice sites. The wave function resides at site m with a local potential $P_m(t)$ that can be changed in time. The full control protocol (FCP) adjusts the potential at all sites individually in discrete time steps t_n to generate a current efficiently. To optimize the protocol, the neural network takes the potentials $P(t < t_n)$ at earlier time steps and returns the potential $P_m(t_n)$ to be applied to the quantum system in the next time step. A measure $\langle \Psi | \hat{O} | \Psi \rangle$ given by observable \hat{O} is measured and used to train the neural network. This process is repeated until convergence. For further details, see Appendix F. (c) Example driving potential found by the optimization algorithm.

in closed circuits: persistent currents [23]. Persistent currents are a direct manifestation of the phenomenon of quantum coherence and are therefore of central interest in fundamental aspects of many-body physics such as superfluidity, superconductivity, and mesoscopic physics [24]. At the same time, such concepts play a vital role for important emerging applications such as rotation sensors based on guided matter-waves in atomtronic circuits [25], which employ similar ideas to their superconducting counterpart, namely superconducting quantum interference devices (SQUIDs) [26]. Persistent currents have been the object of intense studies in different contexts of quantum technologies such as cold atoms [12,27], superconducting circuits [28], optical cavities [29], optomechanical cavities [30], and tailored reservoirs [31].

While charged or neutral matter-wave persistent currents have been obtained in simplified situations, protocols for more general settings are still missing. We consider quantum systems that can be mapped onto ring-shaped circuits comprising local units that we call sites. These systems encompass most of the general features and challenging aspects for the generation of quantum currents in closed circuits. We propose to create current states by locally driving the circuit parameters. If the system can be driven by changing a few control parameters, state engineering can be carried out through optimal control theory [32]. For increasingly large numbers of parameters, however, the circuit driving cannot be handled with standard means. In this work, we employ deep RL to implement current state engineering by driving each lattice site of the ring circuit independently. With our approach, we demonstrate that persistent currents with specific winding number can be imparted, on a timescale that is much shorter than other known protocols. Additionally, we can create entangled current states of up to three winding numbers, for which there is no known protocol. Our protocols can be readily trained and applied in experiments.

In Appendix A, we introduce complementary theoretical definitions for the quantum phase model and phase windings that we used in the main text. In Appendix B, we compare the fidelity with the certification measure, which can be used

to experimentally verify the quality of the prepared state. In Appendix C, we compare our results with deep reinforcement learning with the standard method GRAPE. In Appendix D, we investigate the robustness of our protocols against noise in the driving parameters. In Appendix E, we present additional numerical results on the quantum phase model and the generation of superposition states using full control and barrier protocols. In Appendix F, we introduce the deep reinforcement learning protocol in full detail. In Appendix G, we show how the deep learning protocol improves with training time. In Appendix H, we show the statistical fluctuations of different repetitions of the deep learning protocol. In Appendix I, we show experimental considerations for how to implement our protocols with cold atoms and superconducting circuits.

I. LOCAL DRIVE OF BOSE-HUBBARD RING CIRCUITS

As sketched in Fig. 1(b), our model system is a ring circuit comprised of L sites, a natural architecture to consider to generate persistent currents, described with the Bose-Hubbard Hamiltonian. N_p interacting bosonic particles are filling the ring lattice and can hop between nearest-neighbor sites j and $j + 1$ with an amplitude J , and they interact on-site with each other with strength U . The ring lattice can be locally driven by externally varying in time each on-site potential $P_j(t)$,

$$\mathcal{H}_{\text{BH}} = \sum_{j=1}^L \left[-J (\hat{a}_j^\dagger \hat{a}_{j+1} + \hat{a}_{j+1}^\dagger \hat{a}_j) + P_j(t) \hat{n}_j + \frac{U}{2} \hat{n}_j (\hat{n}_j - 1) \right]. \quad (1)$$

Here \hat{a}_j , \hat{a}_j^\dagger , and $\hat{n}_j = \hat{a}_j^\dagger \hat{a}_j$ are the usual bosonic creation, annihilation, and number operators on site j , satisfying the commutation relation $[\hat{a}_i, \hat{a}_j^\dagger] = \delta_{ij}$ and periodic boundary conditions $\hat{a}_{L+1}^\dagger = \hat{a}_1^\dagger$.

In the limit of a large average number of particles per site $N_s = N_p/L \gg 1$, the Bose-Hubbard Hamiltonian effectively

reduces to the so-called quantum phase model (QPM)

$$\mathcal{H}_{\text{QP}} = \sum_{j=1}^L \left[-2J_E \cos(\hat{\phi}_j - \hat{\phi}_{j+1}) + P_j(t) \hat{Q}_j + \frac{U}{2} \hat{Q}_j^2 \right], \quad (2)$$

where $J_E = JN_s$, $\hat{Q}_j = \hat{n}_j - N_s$ is the on-site particle number fluctuations, and ϕ_j are the phase operators [33,34]. The operators satisfy the commutation relations $[\hat{\phi}_i, \hat{Q}_j] = i\hbar\delta_{ij}$. Hamiltonians (1) and (2) describe a wide class of different physical quantum systems ranging from one-dimensional arrays of Josephson junctions and qubits [35] to atomtronic circuits.

II. QUANTUM CURRENT STATES

In a coherent quantum circuit, the current states in the ring are quantized as the phase along a closed path can only change by integer multiples of 2π . We describe these winding numbers k by defining the single-particle winding state $|k\rangle \doteq \hat{b}_k^\dagger |\text{vac}\rangle$, where $|\text{vac}\rangle$ denotes the vacuum state and $\hat{b}_k^\dagger = \frac{1}{\sqrt{L}} \sum_n e^{i2\pi kn/L} \hat{a}_n^\dagger$ is the quasimomentum creation operator (details are in Appendix A). In a ring system, the quasimomentum corresponds to the winding number and therefore the quantized current in the ring. We choose $\Omega = \{k_1, k_2, \dots, k_{N_C}\}$ as a set of N_C winding numbers that we want to prepare in an entangled superposition state. The generation of such quantum current states is one of the defining goals of quantum technology, and these states are notoriously difficult to generate. We consider states in the form $|\Psi_{\text{EC}}\rangle = \frac{1}{\sqrt{N_C}} \sum_{k \in \Omega} e^{i\phi_k} |k\rangle^{\otimes N_p}$ consisting of N_C winding numbers with arbitrary phase ϕ_k . Important examples of entangled current (EC) states that we will specifically consider in the present work are the NOON-state ($N_C = 2$) and the W-state ($N_C = 3$). We characterize the ability of our protocols to generate these states with the fidelity $F = |\langle \Psi(T) | \Psi_{\text{EC}} \rangle|^2$. We also refer to a certification measure that is related to observables in cold-atom settings (time-of-flight measurements) [36,37] and reflects the behavior of the fidelity: $W_\Psi \propto \prod_{k \in \Omega} \langle \Psi | \hat{n}_k \hat{n}_k | \Psi \rangle$ (see Appendix B).

We introduce the full control protocol (FCP) to generate currents efficiently. In such a protocol, the potential at each lattice site is driven freely within a range $|P_j| < P_{\text{max}}$. The total driving time T is discretized into N_T time steps of equal length $\Delta t = T/N_T$. Within each time step, the system evolves under constant parameters and we assume that the potential parameters change instantaneously between two time steps. We use deep RL with proximal policy optimization using the actor-critic method, and implementation in TENSORFLOW to optimize the protocols [22,38,39] (a sketch of algorithm is in Fig. 1; details are in Appendix F). Here, we note that our scheme relies on a model-free optimization algorithm: The learning algorithm does not make any assumptions about the specific system, nor does it know about quantum mechanics. As a possible application, the algorithm could be supplied with experimental data (in our case the potential parameters and measurement observables) to optimize the experiment directly.

We compare two fundamental approaches to generate quantum current states. First, we investigate the established method of stirring with a barrier, but here applied to a discrete ring lattice. Then, we investigate the FCP protocol.

III. STIRRING A LOCALIZED BARRIER POTENTIAL

Transforming the nonrotating ground state to a specific rotating state requires perturbing the state in a manner that explicitly breaks time-reversal symmetry. This has been implemented in cold-atom settings [12,40] for continuous systems, but it has not been investigated for discrete settings such as a ring lattice. Here, we move a potential barrier initially localized at one site of amplitude P_B to the next site at a fixed frequency ν . We target generating current states with one winding number, starting from the ground state of the system. We find that a high fidelity is reached by driving the ring for a time $T \approx 20/J$ for different particle numbers N_p [rightmost curves in Fig. 2(a)]. Matching calculations for continuous systems, we find that the first rotational state is created best by stirring with about a speed of half the desired atom velocity $v \approx 0.5/J$ [see Fig. 2(d) [41]]. In a rotationally invariant system, the energies for consecutive winding numbers k_1, k_2 are degenerate when the barrier is driven at the mean speed $(v_2 - v_1)/2$. A localized barrier splits the degeneracy and introduces nonadiabatic transitions (Landau-Zener) between the two states. A similar mechanism works also for a ring condensate interrupted by three barriers [42].

The actual implementation of the barrier protocol implies a tradeoff between achieving high fidelity F and short protocol time T : With increasing P_B the maximal achievable fidelity decreases, however it is reached in a shorter time T [see Fig. 2(c)]. To shorten the protocol run time without sacrificing fidelity, more complex protocols are required.

As a benchmark, we apply the deep learning algorithm on a setting with simple instructions (see Fig. 3). The ring is initialized with a single barrier with strength $P = J$. The neural network can at each time step choose to either move the barrier forward by one site or keep the barrier at the current position. In this way, we see that the deep RL is able to find a protocol with nearly constant velocity [dashed line Fig. 3(b)] that is very similar to the analytic solution [solid line in Fig. 3(b)]. This demonstrates the ability of the learning algorithm to reconstruct the known solution by starting from a general control scheme.

IV. LOCAL CONTROL OF THE CIRCUIT

To achieve a higher control over the dynamics of the quantum system, we now apply the FCP in which each lattice site is driven individually by varying the potential in a continuous fashion between $-J < P_j < J$. In Fig. 2(a), we compare stirring (right-hand side) and FCP (left-hand side). The FCP reaches a better fidelity than the barrier driving protocol, requiring only half the time or less. An example protocol that optimizes the state generation is shown in Fig. 1(c). It is very remarkable that the FCP works also in the limit of many particles by employing the QPM Eq. (2) in the regime of intermediate interaction [see Figs. 2(a) and 2(b)]. In experimental settings, the driving potential is subject to fluctuations due to

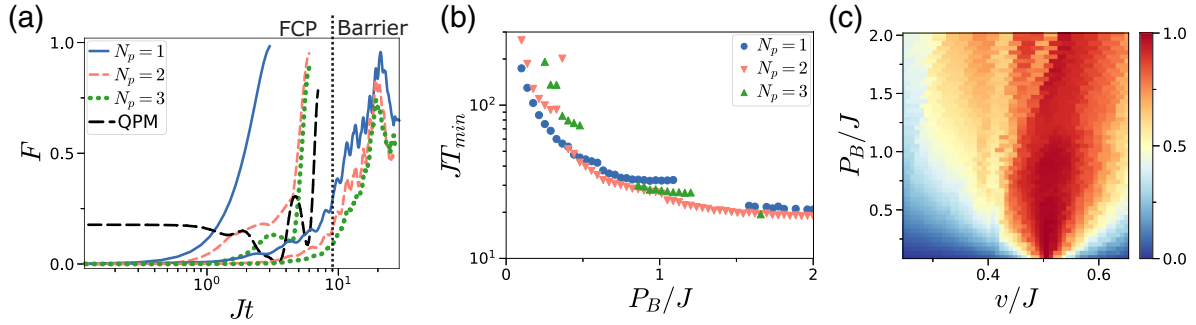


FIG. 2. (a) Generating currents by driving the potential of a ring circuit. We plot the time evolution of the fidelity $F = |\langle \Psi(T) | \Psi_{\text{target}} \rangle|^2$ obtained with two different protocols, comparing different numbers of particles [Bose-Hubbard model Eq. (1) with $L = 12$ sites, $U = J$, target state $|\Psi_{\text{target}}\rangle = |k = 1\rangle^{\otimes N_p}$]. The first protocol (barrier) stirs the wave function by moving a single-site barrier ($P_B = 2J$ for $N_p = 1, 2$, $P_B = 1.6J$ for $N_p = 3$) at a constant speed $v = 0.53J$ (see the curves on the right of the vertical dotted line), with the Bose-Hubbard model ($L = 12$ ring sites, interaction $U = J$). The second protocol [full control protocol (FCP)] is fully controlling the potential at every lattice site in time (see the curves on the left of the vertical dotted line) with $N_T = 6$ control steps: solid blue curves, $N_p = 1$; dashed red line, $N_p = 2$; dotted green curves, $N_p = 3$. Limit of large particle number with quantum phase model (QPM) Eq. (2) with $L = 7$, $U = 3J_E$, and target state $|\Psi_{\text{target}}\rangle = |\Psi_{\text{QP}}(\Phi_1)\rangle$ (see Appendix A): long-dashed black curve. (b) Minimal time T_{min} required to create rotational states above a threshold fidelity ($F_{\text{min}} = 0.95$ for $N_p = 1$; for more elusive higher particle number states, $F_{\text{min}} = 0.85$) for different values of barrier amplitude P_B . (c) Maximal fidelity achieved when rotating the barrier with amplitude P_B and speed v for $N_p = 1$ particles. We find that the best rotation speed of the barrier is at $v \approx 0.5J$.

environmental noise and experimental imperfections. We find that the FCP scheme is robust to noise in the driving potentials to up to 20% (see Appendix D).

To go beyond quantum current states composed of a single winding number, we employ FCP to engineer entangled superposition of winding numbers. In this way, we demonstrate the preparation of entangled superposition states of different currents for up to three winding numbers for which we are not aware of any protocol for their generation [see Fig. 4(a)]. The fidelity improves over protocol time T and eventually reaches a plateau. A minimal number of protocol time steps are required to reach sufficient fidelity. For interacting systems, we achieve best results for $N_T \geq 4$ [Fig. 4(b)]. For the QPM, higher fidelity can be achieved with increasing interaction U , however the initial fidelity is higher, and the momentum distribution of the ground state becomes broader [Fig. 4(c)]. Further data and a comparison with the alternative optimization method GRAPE can be found in Appendixes C and D

[43,44]. Finally, we investigate protocols with discretized control amplitudes in Fig. 4(d). The driving amplitude at every




FIG. 3. Current generation by moving a single δ -like potential (with constant amplitude $P = J$). We compare two protocols: Moving potential at constant speed $v = 0.52J$ (solid line) or optimize with deep RL (dashed line). Over $n_t = 100$ time steps, the RL agent can either move the barrier by one site, or keep it at the current position, trained over 50 000 epochs. (a) Fidelity of the best found protocol to create the phase winding state $\Omega = 1$ for $N_p = 1$. (b) Barrier position in time.

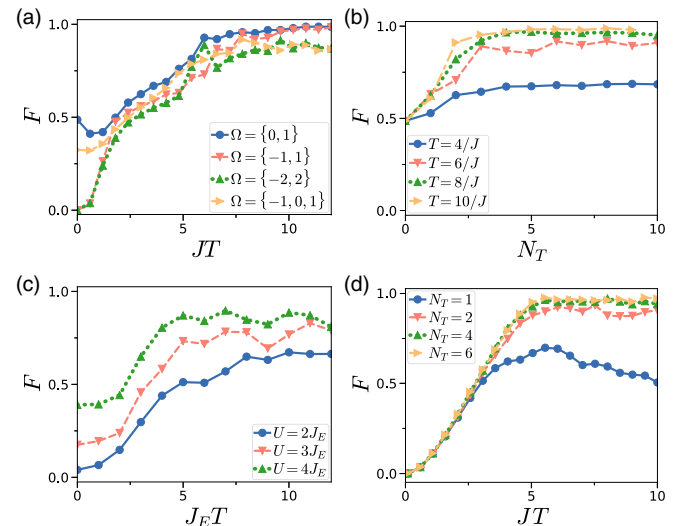


FIG. 4. (a)–(c) Full control with continuous control: Potential at every site can take arbitrary values between $-J < P_j < J$. Generation of entangled superpositions of current states of type $|\text{EC}\rangle = \frac{1}{\sqrt{N_C}} \sum_{k \in \Omega} |k\rangle^{\otimes N_p}$ of a set of winding numbers $\Omega = \{k_1, k_2, \dots, k_{N_C}\}$ using deep RL. (a) Fidelity as a function of protocol time T for $N_p = 2$ particles, $N_T = 6$ time steps, and $U = J$. (b) Fidelity for varying time steps N_T ($N_p = 2$, $U = J$, $\Omega = \{0, 1\}$). (c) Fidelity for the limit of large particle number (QPM) with FCP protocol for $N_T = 6$, $L = 7$, $|P_m(t)| \leq J_E$ and limiting number fluctuations at $\Delta \hat{Q}_m \leq 2$ and target state $|\Psi_{\text{target}}\rangle = |\Psi_{\text{QP}}(\Phi_1)\rangle$. The ground state of the QPM has a broad winding number distribution and thus there is a finite initial fidelity at $T = 0$. (d) Full control protocol with discretized amplitudes: Local potential can take either the value $P_j(t) = 0$ or $P_j(t) = J$. Fidelity to create phase winding state $\Omega = 1$ for $N_p = 1$, for varying protocol time T and number time steps N_T .

site can take only the values $P_j = 0$ or $P_j = J$, and they have a similar fidelity to the continuous case. They give rise to simplistic driving protocols (see Appendix E).

V. DISCUSSION

In this work, we demonstrated how to use machine learning for the efficient generation of currents in closed quantum circuits. The essential features of this problem are captured by a ring consisting of lumped elements, which we schematize as lattice sites and can be modeled using the Bose-Hubbard and quantum phase model, paradigmatic frameworks for the physics of cold atoms, superconducting circuits, and photonic waveguides.

We introduced the full control protocol (FCP), where all sites of the lattice are driven individually. This protocol paves the way for enhanced control over the many-body system. Inspired by recent advances in deep RL, we use neural networks to learn these complex driving protocols. In contrast to other optimization methods such as GRAPE, deep RL is agnostic of the underlying physical system as a kind of hybrid quantum-classical optimizer [45–47]. Thus, it can directly improve experiments, optimizing the fidelity with an experimentally observable measure W_ψ as input to the learning algorithm, while achieving comparable results to those with GRAPE (see Appendix B). We envision that first the time-consuming training of the protocol is performed on the experiment, after which the protocol can be readily applied whenever a specific current state is needed.

We benchmarked our algorithm with the stirring protocol, the standard protocol used for cold-atom technology [12]. Remarkably, deep learning is able to rediscover the established protocol. With the FCP, we show how to generate current states consisting of a single winding number more than twice as fast and with higher fidelity compared to the standard stirring scheme used so far by the cold-atom community; see Fig. 2(a). Furthermore, we show how to produce entangled current states such as NOON states as well as W-type states involving three winding numbers, for which no protocol has been known so far (see Fig. 4). This scheme is robust to noise in the driving parameter to up to 20% (see Appendix D). We find that the complexity of the driving protocol (protocol time and number of time steps) depends on the number of particles: Noninteracting systems can generate currents much faster and with simpler protocols compared to the interacting many-body system. Both small and large numbers of particle regimes (through Bose-Hubbard and quantum phase dynamics, respectively) were explored.

Our deep-learning method can be also applied to problems with discretized driving amplitudes, where gradient-based methods like GRAPE are difficult to apply [16,19] [see Fig. 4(d), Appendix E]. These protocols take quite a simplistic form, which we conjecture could be related to general classes of driving protocols for current generation. Based on the physical mechanism behind the optimal stirring protocols for a single barrier and three barriers (see the Landau-Zener argument above), it is tempting to conclude that the FCP generates the desired current state by looking at the self-avoiding crossing in the energy landscape and by optimizing the transition amplitudes between states with different winding numbers.

Our findings are of direct relevance in different contexts of quantum technology. Especially for cold-atom quantum technology, our approach clearly shows a path to solve important challenges in realizing atomtronic circuits. In particular, we note that our approach does not rely on artificial gauge fields, which most other methods require to generate currents. This feature leads to a simplification of the experimental apparatus. Similarly, the FCP could be exploited in superconducting circuits to achieve controlled electronic currents. Our approach can be extended to other quantum many-body systems, for current states in more complex circuit geometries and hybrid quantum networks, as well as for quantum-enhanced sensing [48].

ACKNOWLEDGMENTS

We thank D. Basko, V. Bastidas, M. Holzmann, D. Hutchinson, A. Minguzzi, and B. Munro for discussions. The Grenoble LANEF framework (ANR-10-LABX-51-01) is acknowledged for its support with mutualized infrastructure. The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore [49].

APPENDIX A: THEORETICAL MODELS

1. Winding number states

To define our current states, we transform the ring Hamiltonian Eq. (1) in the main text of L sites with $U = 0$ and $P_j = 0$ by Fourier transforming the operators into the quasimomentum basis

$$\mathcal{H}_{\text{FT}} = \sum_{k=0}^{L-1} -2J \cos\left(\frac{2\pi k}{L}\right) \hat{n}_k, \quad (\text{A1})$$

where $\hat{n}_k = \hat{b}_k^\dagger \hat{b}_k$, with $\hat{b}_k^\dagger = \frac{1}{\sqrt{L}} \sum_n e^{i2\pi kn/L} \hat{a}_n^\dagger$. As the wave function around the ring is continuous, the wave function must be the same after going once around the ring. Thus, we demand $\exp(i\frac{2\pi k}{L}n) = \exp(i\frac{2\pi k}{L}(n+L))$, which is only fulfilled if k is an integer number, which describes how often the phase of the wave function winds by 2π around the ring. The state with winding number k for a single particle is defined as $\hat{b}_k^\dagger|\text{vac}\rangle = |k\rangle$, where $|\text{vac}\rangle$ denotes the vacuum state. Many-body states are generated as tensor products of particles, e.g., a state with N_p particles with winding number k is given by $|\Psi_k\rangle = |k\rangle^{\otimes N_p}$. For this state, the expectation value of the number of particles with winding number k is given by $\langle \hat{n}_k \rangle = \langle \Psi_k | \hat{n}_k | \Psi_k \rangle = N_p$.

2. Quantum phase model fidelity

To define the current state for the quantum phase model (QPM), we use the QPM with an applied artificial magnetic field Φ_M ,

$$\mathcal{H}_{\text{QP}}(\Phi_M) = \sum_{j=1}^L \left[-2JN_s \cos(\hat{\phi}_j - \hat{\phi}_{j+1} - \Phi_M) + P_j(t) \hat{Q}_j + \frac{U}{2} \hat{Q}_j^2 \right]. \quad (\text{A2})$$

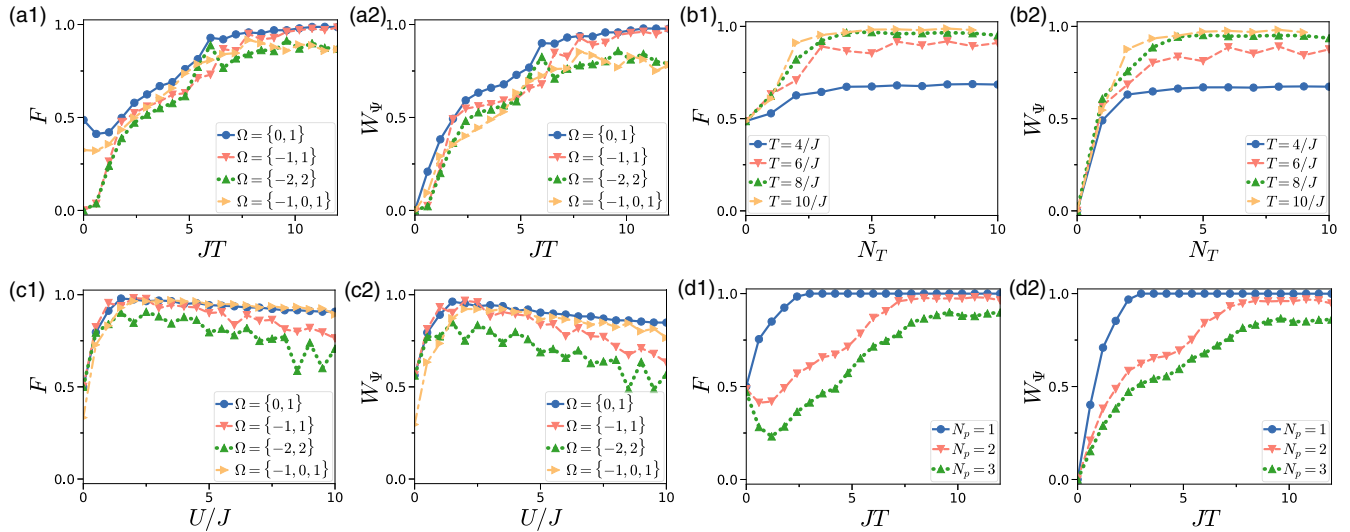


FIG. 5. Comparison of fidelity with entangled state $F = |\langle \Psi | \text{ES} \rangle|^2$ (for figure index 1) and certification measure W_Ψ (for index 2) for $L = 12$ sites for various parameters. Same parameters as in Fig. 4. We optimize for equal weight entangled states of N_C winding number k of type $|\text{ES}\rangle = \frac{1}{\sqrt{N_C}} \sum_{k \in \Omega} |k\rangle^{\otimes N_p}$ of a set of winding number $\Omega = \{k_1, k_2, \dots, k_{N_C}\}$. (a) Varying time T to generate different entangled states for $N_p = 2$ particles, $N_T = 6$ time steps, and $U = J$. (b) For varying time steps N_T to reach state $\Omega = \{0, 1\}$ for $N_p = 2$ particles. (c) Interaction U dependence for different types of states for $N_T = 6$ time steps and protocol time $T = 9/J$. (d) Total protocol time T for $N_T = 4$ time steps to generate entangled superposition state of winding number $\Omega = \{0, 1\}$.

To define the target state that carries a current, we refer to the ground state with a winding number distribution that is centered around a specific winding number, depending on Φ_M . For $\Phi_k = \frac{2\pi k}{L}$, the ground state $|\Psi_{\text{QP}}(\Phi_k)\rangle$ winding number distribution is centered around the winding number k . We define the fidelity as $F = |\langle \Psi | \Psi_{\text{QP}}(\Phi_k) \rangle|^2$.

APPENDIX B: FIDELITY AND CERTIFICATION MEASURE

In the main text, we discuss entangled current states using the fidelity F . However, fidelity is not an observable that can be easily measured in experiments. We proposed a certification measure W_Ψ that is an experimental observable [50],

$$W_\Psi = \frac{N_C^{N_C}}{N_p^{2N_C}} \prod_{k \in \Omega} \langle \Psi | \hat{n}_k \hat{n}_k | \Psi \rangle. \quad (\text{B1})$$

This measure behaves similarly to the standard fidelity for the entangled quantum current states $F = |\langle \Psi | \Psi_{\text{EC}} \rangle|^2$. For the target state $|\Psi_{\text{EC}}\rangle$ it reaches a maximal value $W_\Psi = 1$ and gives similar results for the state design. In contrast to the fidelity F , however, we note that W_Ψ is related to particle densities and therefore it is an observable. In a cold-atom setting, for example, W_Ψ can be accessed by measuring the number of particles in a specific momentum mode, which can be achieved by time-of-flight measurements [36,37] (see below). In Fig. 5, we show results on the fidelity for the same parameters as in the main text in Fig. 4. We find that the fidelity behaves similar to the certification measure. The certification measure is always zero for the initial state (e.g., seen for data points with $T = 0$), however the fidelity can be nonzero. This is because the initial state has in some cases a finite overlap with the target state. In contrast, the certification

measure is constructed such that this initial overlap does not affect it.

To characterize the entangled current states in an experimental setting, we defined the certification measure Eq. (B1), which is a product of expectation values of observables. Experimentally, one is required to measure the square of the particle-number operator $\langle \hat{n}_k^2 \rangle$ of the winding number mode k . For cold-atom condensates, this measure can be determined from time-of-flight measurements, where the prepared state is expanded in free space [36]. After free expansion, the density of the atoms develops a hole in the center, which is proportional to the winding number k [12]. $\langle \hat{n}_k^2 \rangle$ can be calculated by repeating the time-of-flight measurement several times and calculating the fluctuations in the winding numbers.

For superconducting circuits, the expectation value of the square of the particle number of a specific winding number k can be derived from the expectation value of fourth-order correlators between different qubits. We find

$$\hat{n}_k^2 = \frac{1}{L^2} \sum_{n,m,r,s} e^{i2\pi k(n+r-m-s)/L} \hat{a}_n^\dagger \hat{a}_m^\dagger \hat{a}_r \hat{a}_s,$$

where the correlators can be derived by Fourier transforming the annihilation and creation operators of the operator.

APPENDIX C: COMPARISON TO GRAPE

A standard approach to optimize quantum dynamics is GRAPE [43,44]. It relies on calculating the gradients of the control unitaries, then updating the driving parameters with gradient descent. As such, it requires access to the full model parameters as well as the underlying unitaries, in contrast to deep RL, which suffices with observables only and is model-free. We compare GRAPE in Fig. 6 with the data generated with deep RL in Fig. 5. GRAPE achieves the same or a bit higher fidelity in most cases. However, this is not surprising

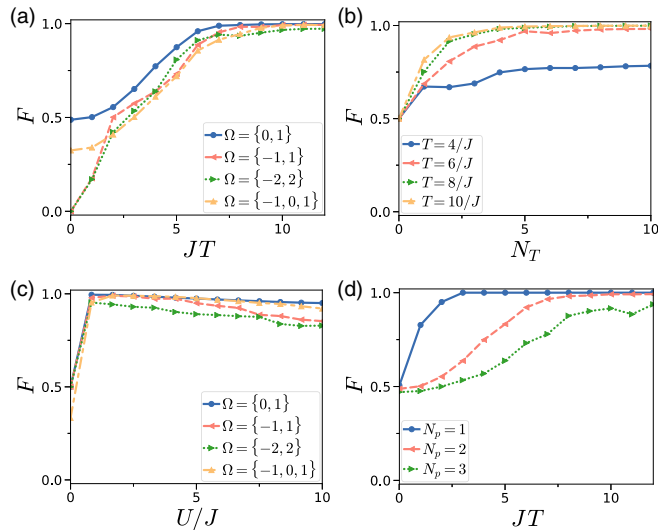


FIG. 6. Generation of entangled superpositions of current states of type $|\text{EC}\rangle = \frac{1}{\sqrt{N_C}} \sum_{k \in \Omega} |k\rangle^{\otimes N_p}$ of a set of winding numbers $\Omega = \{k_1, k_2, \dots, k_{N_C}\}$ using GRAPE. (a) Fidelity as a function of protocol time T for $N_p = 2$ particles, $N_T = 6$ time steps, and $U = J$. (b) Fidelity for varying time steps N_T ($N_p = 2$, $U = J$, $\Omega = \{0, 1\}$). (c) Fidelity as a function of interaction U for different types of states ($N_p = 2$, $N_T = 6$, $T = 9/J$). (d) Fidelity as a function of protocol time T for three different particle numbers N_p ($N_T = 4$, $U = J$, $\Omega = \{0, 1\}$). All data for $L = 12$ sites. The GRAPE algorithm is run multiple times to avoid solutions that became stuck in local extrema.

since it has access to all the information about the quantum system itself. GRAPE relies on gradients, and as such cannot optimize problems with discretized driving parameters as considered in Fig. 4(d).

APPENDIX D: ROBUSTNESS TO NOISE IN DRIVING

Experimental realizations are afflicted by experimental uncertainties and noise. In particular, for the full control protocol, where the potential at each site is controlled in time, the actual potential can only be set with finite accuracy. Experimental imperfections may also lead to noise, such that the real value used in the experiment is perturbed from the desired value. We simulate a random fluctuation of the potential in Fig. 7. For each time step and site, a random potential $P_{i,\text{actual}}(t_n) = P_{i,\text{desired}}(t_n) + \Delta P$ is added, which is sampled from a uniform distribution between $\Delta P \sim [-\delta P, \delta P]$. We observe that for small $\delta P < 0.25J$, noise has barely any influence on the resulting dynamics. We observe that increasing time T negatively affects robustness [see Figs. 7(a), 7(b), and 7(d)], and the standard deviation of the fidelity increases. Increasing the number of time steps N_T improves the robustness [see Fig. 7(c)]. We verify that these results persist for both one and two particles [for two particles, see Figs. 7(c) and 7(d)], as well as for entangled [for entangled current, see Fig. 7(d)] and nonentangled winding states.

APPENDIX E: SUPPORTING NUMERICAL RESULTS

Here, we present further data to support our findings. To solve the quantum phase model numerically, we restrict the

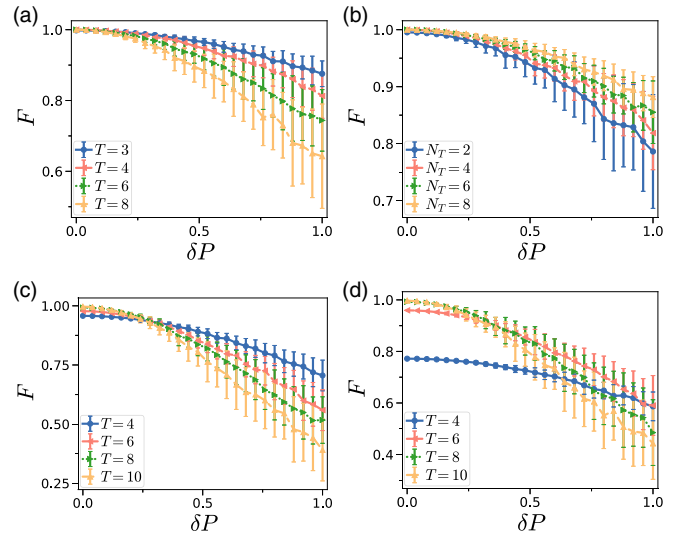


FIG. 7. Full control driving; local potential can take continuous values between $P = -J$ and $P = J$, changing in discrete time steps. The protocol is optimized with GRAPE. Then, the local potential of the optimal protocol is perturbed with random noise, sampled from uniform distribution between $[-\delta P/2, \delta P/2]$, sampled for every site and time step. Each point is sampled from 100 noise realizations; dots show the mean value and error bars show the standard deviation of fidelity F . (a) Fidelity to create the phase winding state $\Omega = 1$ for $N_p = 1$, for varying random perturbation δP and protocol time T . Robustness against noise decreases with increasing time T . (b) Varying protocol time steps N_T increases with the number of time steps for $\Omega = 1$ and $N_p = 1$. (c) Varying T for $N_p = 2$, $U = J$, $N_T = 6$ and target state $\Omega = 1$. (d) Same parameters for entangled superposition state of $\Omega = \{0, 1\}$.

fluctuations around the mean particle number to $\Delta \hat{Q}_m$. Here, we increase the number of particles allowed to fluctuate. In Fig. 8, we show the fidelity $\Delta \hat{Q}_m = 4$ and $L = 5$.

Next, we show the full control protocol with constrained control amplitudes. In Fig. 9, the local potential can either assume $P_i(t) = 0$ or $P_i(t) = J$. An example potential is shown




FIG. 8. Fidelity of generating the current state for the quantum phase model. The initial state is the ground state without flux, while the target state is the ground state of the model with one flux quantum. The plot shows the fidelity of reaching the target state $|\Psi_{\text{QP}}(\Phi_1)\rangle$ for different protocol times T . We restrict the local Hilbert space $\Delta \hat{Q}_m = 4$. We choose $N_T = 8$, $L = 5$, and potential $|P_{\text{max}} < J_E|$.




FIG. 9. (a) Full control protocol with discretized amplitudes: The local potential can take either the value $P_j(t) = 0$ or $P_j(t) = J$. Fidelity to create the phase winding state $\Omega = 1$ for $N_p = 1$, for varying protocol time T and number of time steps N_T . (b) Resulting protocol for discretized driving amplitudes.

in Fig. 9(b). The resulting protocol takes a very simplistic form.

In the main text, we demonstrate the generation of a state with winding number $k = 1$ using the barrier protocol. The same protocol can also generate superposition states of $k = 0$ and 1. To generate the entangled state, the barrier is rotated at the same speed as for the $k = 1$ case with shorter time T . In Fig. 10, we study the dynamics of creating entangled superposition states and compare the FCP against driving a barrier.

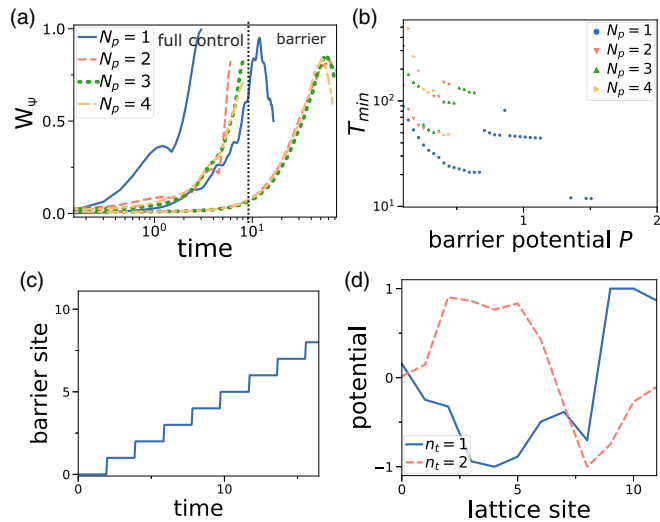


FIG. 10. Generation of (entangled) superposition states of zero and one rotational quantum $\Omega = \{0, 1\}$ in a ring lattice with $L = 12$ sites. Evolution of certification measure [Eq. (B1)] during driving. We compare two different protocols: A barrier localized at a single site moving at constant speed [right curves in (a)] or fully controlling the potential (FCP) of every lattice individually [left curves of (a)]. (b) Minimal time T_{\min} required to create rotational states above a threshold fidelity ($W_{\min} = 0.95$ for $N_p = 1$, otherwise $W_{\min} = 0.8$) for different values of barrier amplitude. We find the best rotation speed of the barrier is at $v \approx 0.5J$. (c) Best protocol for rotating barrier. Curve shows the barrier position over time. (d) Best protocol for full control over lattice potentials for a protocol of two time steps n_t . Barrier and full control protocols shown calculated for $N_p = 1$ particles.

APPENDIX F: DEEP REINFORCEMENT LEARNING

Here, we describe our machine-learning algorithm in more detail. A detailed figure describing the neuronal network structure is shown in Fig. 11, and the pseudocode of the algorithm is shown in Table I. We learn the driving protocol via a deep Q -learning network [38], utilizing the actor-critic method acting on a continuous action space. Our method is using Proximal Policy optimization [22], and the implementation is based on TENSORFLOW [39]. The quantum system is controlled by an agent, which depending on the state s_t of the system acts with an action a_t using the probabilistic policy $\pi(a_t|s_t)$. The idea of Q -learning is to find the Q -function $Q_\pi(s_t, a_t)$ that estimates the future reward that is paid out at the end of the full protocol with this policy. The goal is to learn a policy that can realize long-term rewards over smaller short-term gains. The optimal Q -function is determined by the Bellman equation

$$\begin{aligned} Q(s_t, a_t, \pi) &= \mathbb{E}[r_t + \gamma Q(s_{t+1}, a_{t+1}, \pi)] \\ &= \mathbb{E}[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots], \end{aligned}$$

where $\mathbb{E}[\cdot]$ indicates sampling over many instances. $\gamma \leq 1$ is a discount factor that weighs future rewards against immediate rewards. The input to the neural network is the Hamiltonian parameters at previous time steps, and it outputs the parameters for the policy $\pi(a_t|s_t, \mu, \sigma)$, where the actions are sampled from a normal distribution with mean value μ and width σ . μ is determined by the neural network, and σ is optimized as a global variable and decreases during the optimization procedure. We constrain the possible output values for the potential by mapping values outside of the constraint to the maximally allowed value. Proximal policy optimization is based on the actor-critic method. The idea is to have two neural networks: a policy network and a value network. The policy network (actor) decides on the next action by determining the parameters of the policy. The value-based network (critic) evaluates the action taken on how well it solves the task and estimates the future expected reward. It is used as an input to train the policy network. The two networks are trained at the same time using Adam [51]. Better performance can be achieved if the Q -function is split into two parts [52]: $Q(s_t, a_t) = A(s_t, a_t) + V(s_t)$, where $A(s_t, a_t)$ is the advantage function and $V(s_t)$ the value function. $V(s_t)$ gives the expected future reward averaged over the possible actions according to the policy. This is the output of the critic network. $A(s_t, a_t)$ gives the improvement in reward for action a_t compared to the mean of all choices. We estimate the Q -function from the value function with $Q(s_t, a_t) = r_t + \gamma V(s_{t+1})$, where r_t is the reward given out under action a_t , and $V(s_{t+1})$ is the value function for the next time step. We then minimize the square of the difference of the value function of the network and the predicted reward in the next time step $L_V(\theta) = \mathbb{E}_t[[V_\theta(s_t) - y_t]^2]$, where θ are the current network parameters, and $y_t = r_t + V(s_{t+1})$ is the calculated reward of the next time step. The advantage function $A(s_t, a_t) = Q(s_t, a_t) - V(s_t)$ tells us how good a certain action a_t is compared to other possible actions. Using the above estimation of the Q -function, the advantage function can be approximated. The advantage function is the input to train the policy network (the actor). Following the idea of

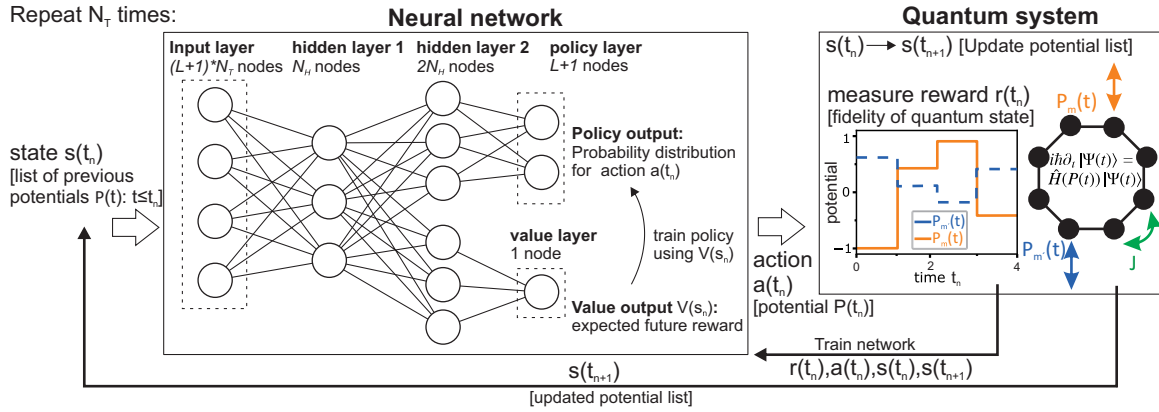


FIG. 11. Neural network to optimize protocols to generate quantum states. Deep learning relies on representing a highly complex function (e.g., the quality of the driving protocol) with a neural network, and optimizes it using observable data (e.g., measurement outcomes). The quantum system is a lattice ring with L sites where particles can hop between neighboring sites with strength j . Each site m has a local potential $P_m(t_n)$ that can be modulated in discrete time steps t_n . The neural network controls the evolution of the quantum system by adjusting $P_m(t_n)$ and optimizes the parameters over many runs. The neural network performs stepwise evolution of the quantum system in N_T discrete time steps t_n over total runtime T . It uses the chosen potentials of previous time steps as an input [state $s(t_n)$], and it returns the potentials to be chosen at the next step [action $a(t_n)$] by sampling them from a Gaussian distribution. The training is performed by using a measure for the quantum state [reward $r(t_n)$].

proximal policy optimization [22], the goal is to maximize

$$L_p(\theta) = \mathbb{E}_t \left[\frac{\pi_\theta(s_t, a_t)}{\pi_{\theta_{\text{old}}}(s_t, a_t)} A(s_t, a_t) \right], \quad (F1)$$

where θ are the network parameters and θ_{old} are the network parameters of a previous instance. Maximizing $L_p(\theta)$ for the network parameters θ over many sampled instances guides the distribution $\pi_\theta(s_t, a_t)$ such that it returns actions a_t with maximal advantage. However, the ratio

$$b_t(\theta) = \frac{\pi_\theta(s_t, a_t)}{\pi_{\theta_{\text{old}}}(s_t, a_t)}$$

can acquire excessively large values, causing too large changes in the policy in every training step and making convergence difficult. It was proposed to use a clipped

ratio [22]

$$L_p(\theta) = \mathbb{E}_t [\min\{b_t(\theta)A(s_t, a_t), \text{clip}(r_t(\theta), \times 1 - \epsilon, 1 + \epsilon)A(s_t, a_t)\}],$$

such that the update at each step stays in reasonable bounds. We use $\epsilon = 0.1$. We optimize the neural network over many epochs N_E . For our results, we show the best protocol that was achieved during the optimization process. We update the network by randomly sampling N_{train} past iterations from a memory (replay buffer B) that stores the last N_{memory} epochs. To reduce premature convergence, we add the entropy of the normal distributions of the policy to the loss function $L_S(\theta) = \mathbb{E}_t [\frac{1}{2} \sigma \ln(2\pi e)]$. This contribution slows down optimization to avoid convergence to a local minimum. The final loss function to optimize is $L(\theta) = L_p + c_v L_V + c_s L_S$, where c_s and c_v are hyperparameters. We find $c_s = 0.02$ and

TABLE I. Pseudocode for our proximal policy optimization algorithm to generate quantum states.

```

Randomly initialize critic  $V(s|\theta)$  and actor  $\mu(s|\theta)$  with weights  $\theta$ 
Initialize replay buffer  $B$ 
for epoch=1,  $N_E$  do
  Input initial state  $s_1$ 
  for  $t = 1, N_T$  do
    Sample action  $a_t = \mu(s_t|\theta)$  from probability distribution
    Execute action  $a_t$ , receive reward  $r_t$  and next state  $s_{t+1}$ 
    Sample random batch of  $N_{\text{train}}$  transitions  $(s_t, a_t, r_t, s_{t+1})$  from  $B$ 
    Set  $y_t = r_t + \gamma V(s_{t+1})$ 
    Update critic by minimizing loss  $L = \frac{1}{M} \sum_i [y_i - V(s_i|\theta)]^2$ 
    Calculate advantage function  $A(s_t, a_t) = Q(s_t, a_t) - V(s_t) = r_t + \gamma V(s_{t+1}) - V(s_t)$ 
    Calculate probability ratio  $r_p = P(a_t|\mu(s_t|\theta)) / P(a_t|\mu(s_t|\theta^{\text{old}}))$  of current policy  $\mu(s_t|\theta)$  and previous policy  $\mu(s_t|\theta^{\text{old}})$ 
    function  $\text{clip}(r_p, c)$  clips  $r_p$  between  $1 - c < r_p < 1 + c$ 
    Update actor policy with clipped loss  $L = \min[r_p A(s_t, a_t), A(s_t, a_t) \text{clip}(r_p, c)]$ 
  end for
  Store  $(s_t, a_t, r_t, s_{t+1}) \forall t$  in replay buffer  $B$ 
end for

```




FIG. 12. Optimization of the FCP protocol by the neuronal network over the number of epochs (number of protocol runs). We show exemplary data that generated the protocol shown in Fig. 2(a). The dots indicate fidelity achieved during a particular run, while the red line is the moving average over the results. (a) $N_p = 1$, (b) $N_p = 2$, (c) $N_p = 3$, and (d) quantum phase model.

$c_V = 0.5$ to be good choices. A sketch of our neural network is shown in Fig. 11. The protocol solves the Schrödinger equation for a total time T with N_T discrete time steps of width Δt , with respective times t_n . The network determines the Hamiltonian parameters at different sites m of total L sites. For one epoch, the system runs the network N_T times. Input are the potentials used at previous time steps t_n , and the parameters are returned to be used for the $n + 1$ time step. The input vector has length $(L + 1)N_T$; it lists the parameters and the corresponding times t_k used up to current time step $t_k : t_1, \dots, t_n$. The network propagates through two hidden layers of fully connected neurons of size N_H with ReLu activation functions. The output layer has size $L + 1$ and uses a linear activation function. For the value function (critic), the output of the last hidden layer is collected to a single node, which represents the value function $V(s_t)$. For the policy (actor), L outputs determine the mean values of the normal distribution that generates the potential at the next time step t_{n+1} of the protocol. The neuronal network is trained with the loss function after calculating the full time evolution to time T and measuring all the rewards.

For the actual implementation, we choose the following parameters: learning rate with Adam $\alpha = 0.0002$, $N_H = 200$ neurons in the hidden layer, training over $N_E = 120\,000$ epochs, training with a randomly sampled batch size 500, and a replay buffer B of $N_{\text{train}} = 500N_T$ previous results.

APPENDIX G: OPTIMIZATION RUNTIME

The machine-learning algorithm starts with a randomly initialized neuronal network that generates the driving sequence. By running repeatedly, the network generates better driving sequences. Finally, we plot the best driving sequence found during the training epochs. Here, we show results on the training procedure. The fidelity achieved increases during the training over many epochs. In Fig. 12, we show the fidelities during the training procedure for the FCP protocols that were used for Fig. 2(a) of the main text.

APPENDIX H: STATISTICS

Our goal is to optimize a high-parameter space driving protocol. In general, the optimization landscape is complex, with many local minima. We run the machine-learning algorithm several times, and we look at the convergence of the certification measure. As the algorithm is nondeterministic and not guaranteed to converge to the global minimum,

each run can yield different end results. In Fig. 13, we show the minimal and maximal certification measure achieved for 20 runs to create entangled states. For reaching $\Omega = \{0, 1\}$ we see only a small variation between minimal and maximal achieved certification measure; see Fig. 13(a). Thus, in one run of the algorithm we can be sure that a very good solution is found. However, we see a significant spread in certification measure results for higher particle number and more complex entangled states, e.g., $\Omega = \{-1, 1\}$ and $N_p = 3$ particles; see Fig. 13(b). Thus, for this parameter set to find the best result, several runs have to be performed. This implies that the complexity and difficulty of the optimization problem to generate entangled states are highly dependent on the parameters of the problem. We took care to check that the variance of the solutions is within reasonable bounds for our results.

APPENDIX I: EXPERIMENTAL CONSIDERATIONS

For a cold-atom implementation, the driving of the ring lattice can create excitations. Within the Bose-Hubbard approximation, only the first Bloch band is considered. It is assumed that higher Bloch bands are far-detuned in energy and thus do not contribute. In most experiments, the energy gap between the Bloch bands within the harmonic approximation of the lattice sites is given by $E_{\text{lattice}} = 2\sqrt{V_0}E_R$, where V_0 is the potential energy of a sinusoidal confinement and E_R is the recoil energy [53]. For typical $V_0 = 10E_R$, we find $E_{\text{lattice}} = 6.3E_R$. The nearest-neighbor coupling J can be

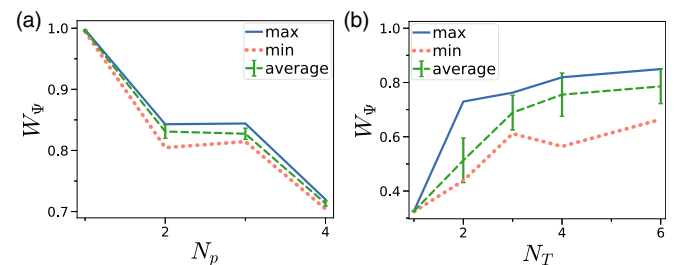


FIG. 13. Statistics (minimum, maximum, and average certification measure W_Ψ) over 20 repeated runs of the algorithm for different parameter sets. (a) Different particle numbers N_p for $N_T = 4$, $T = 8/J$, $U = J$, and $\Omega = \{0, 1\}$. (b) Different protocol steps N_T for a total protocol length of $T = 10/J$, $\Omega = \{-1, 1\}$, $L = 12$ sites for $N_p = 3$ particles, and $U = J$. Driving with local potential $|P| < P_{\text{max}} = J$.

approximated as $J/E_R = \frac{4}{\sqrt{\pi}} \left(\frac{V_0}{E_R}\right)^{3/4} \exp(-2\sqrt{\frac{V_0}{E_R}}) \approx 0.02$ [36]. Thus, the energy separation between the first and second Bloch band is $\Delta E = E_{\text{lattice}}/J \approx 315$. From first-order perturbation theory, we know that the overlap with higher-order states scales as $P/\Delta E$, where P_{max} is the strength of the perturbation. The perturbation of the potential that we apply is on the order of $P_{\text{max}} = J$, which is much smaller than the energy gap. Thus, we can safely ignore excitations to higher Bloch bands. For the quantum phase model, the nearest-neighbor coupling strength scales as $J_E = JN_p/L$. To justify the one Bloch band approximation, J_E has much less than the energy gap between the Bloch bands. Thus, $J_E = \frac{JN_p}{L} \ll E_{\text{lattice}}$.

For cold atoms, the control parameter to change the local potential P_j is V_0 . In our protocol, we change $P_j(t)$ and thus V_0 in time on the order of the nearest-neighbor coupling J . However, J is actually a function of V_0 and thus may change

due to the driving. However, as $V_0 \gg J$, changing V_0 on the order of J has a negligible effect on J .

We assume a stepwise control of the potential, with sharp changes in the potential. The steps that change the potential have an amplitude on the order of J , at timescales of J . We detail how this can be realized in experiment. In the case of cold atoms, the potential is generated by laser pulses. Light-shaping techniques can modify the potential with a frequency of about 20 kHz [54]. The relevant timescale of the experiment is on the order of J , which is far smaller. For superconducting circuits, the potential is controlled by microwave pulses that modify the circuit potential. The circuit potential can be modulated on the order of 35 MHz, while the nearest-neighbor couplings are far slower with $J \approx 4$ MHz [28]. Thus, we conclude that the driving parameters can be feasibly modulated on the timescales we consider.

-
- [1] L. Amico, G. Birkel, M. Boshier, and L.-C. Kwek, Focus on atomtronics-enabled quantum technologies, *New J. Phys.* **19**, 020201 (2017).
- [2] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus *et al.*, Digital quantum simulation of fermionic models with a superconducting circuit, *Nat. Commun.* **6**, 7654 (2015).
- [3] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O'Brien, Silica-on-silicon waveguide quantum circuits, *Science* **320**, 646 (2008).
- [4] C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsforth, and A. Kuhn, Control and manipulation of cold atoms in optical tweezers, *New J. Phys.* **14**, 073051 (2012).
- [5] K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier, Experimental demonstration of painting arbitrary and dynamic potentials for Bose-Einstein condensates, *New J. Phys.* **11**, 043030 (2009).
- [6] H. Rubinsztein-Dunlop, A. Forbes, M. V. Berry, M. R. Dennis, D. L. Andrews, M. Mansuripur, C. Denz, C. Alpmann, P. Banzer, T. Bauer *et al.*, Roadmap on structured light, *J. Opt.* **19**, 013001 (2016).
- [7] G. Kurizki, P. Bertet, Y. Kubo, K. Mølmer, D. Petrosyan, P. Rabl, and J. Schmiedmayer, Quantum technologies with hybrid systems, *Proc. Natl. Acad. Sci. USA* **112**, 3866 (2015).
- [8] A. Aspuru-Guzik and P. Walther, Photonic quantum simulators, *Nat. Phys.* **8**, 285 (2012).
- [9] A. A. Houck, H. E. Türeci, and J. Koch, On-chip quantum simulation with superconducting circuits, *Nat. Phys.* **8**, 292 (2012).
- [10] I. Bloch, J. Dalibard, and S. Nascimbene, Quantum simulations with ultracold quantum gases, *Nat. Phys.* **8**, 267 (2012).
- [11] G. Gauthier, S. S. Szigeti, M. T. Reeves, M. Baker, T. A. Bell, H. Rubinsztein-Dunlop, M. J. Davis, and T. W. Neely, Quantitative Acoustic Models for Superfluid Circuits, *Phys. Rev. Lett.* **123**, 260402 (2019).
- [12] K. C. Wright, R. B. Blakestad, C. J. Lobb, W. D. Phillips, and G. K. Campbell, Driving Phase Slips in a Superfluid Atom Circuit with a Rotating Weak Link, *Phys. Rev. Lett.* **110**, 025302 (2013).
- [13] T. Haug, R. Dumke, L.-C. Kwek, and L. Amico, Andreev-reflection and Aharonov-Bohm dynamics in atomtronic circuits, *Quantum Sci. Technol.* **4**, 045001 (2019).
- [14] T. Haug, H. Heimonen, R. Dumke, L.-C. Kwek, and L. Amico, Aharonov-Bohm effect in mesoscopic Bose-Einstein condensates, *Phys. Rev. A* **100**, 041601(R) (2019).
- [15] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [16] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement Learning in Different Phases of Quantum Control, *Phys. Rev. X* **8**, 031086 (2018).
- [17] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 33 (2019).
- [18] A. G. R. Day, M. Bukov, P. Weinberg, P. Mehta, and D. Sels, Glassy Phase of Optimal Quantum Control, *Phys. Rev. Lett.* **122**, 020601 (2019).
- [19] X.-M. Zhang, Z. Wei, R. Asad, X.-C. Yang, and X. Wang, When does reinforcement learning stand out in quantum control? A comparative study on state preparation, *npj Quantum Inf.* **5**, 85 (2019).
- [20] H. Xu, J. Li, L. Liu, Y. Wang, H. Yuan, and X. Wang, Generalizable control for quantum parameter estimation through reinforcement learning, *npj Quantum Inf.* **5**, 82 (2019).
- [21] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- [23] U. Eckern and A. Schmid, Persistent currents of single mesoscopic rings, *Europhys. Lett.* **18**, 457 (1992).
- [24] E. Akkermans and G. Montambaux, *Mesoscopic Physics of Electrons and Photons* (Cambridge University Press, Cambridge, 2007).

- [25] C. Ryu, P. W. Blackburn, A. A. Blinova, and M. G. Boshier, Experimental Realization of Josephson Junctions for an Atom SQUID, *Phys. Rev. Lett.* **111**, 205301 (2013).
- [26] M. Tinkham, *Introduction to Superconductivity* (Courier, Mineola, New York, 2004).
- [27] S. Pandey, H. Mas, G. Drougakis, P. Thekkeppatt, V. Bolpasi, G. Vasilakis, K. Poullos, and W. von Klitzing, Hypersonic Bose-Einstein condensates in accelerator rings, *Nature (London)* **570**, 205 (2019).
- [28] P. Roushan, C. Neill, A. Megrant, Y. Chen, R. Babbush, R. Barends, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth *et al.*, Chiral ground-state currents of interacting photons in a synthetic magnetic field, *Nat. Phys.* **13**, 146 (2017).
- [29] A. Metelmann and A. A. Clerk, Nonreciprocal Photon Transmission and Amplification Via Reservoir Engineering, *Phys. Rev. X* **5**, 021025 (2015).
- [30] K. Fang, J. Luo, A. Metelmann, M. H. Matheny, F. Marquardt, A. A. Clerk, and O. Painter, Generalized non-reciprocity in an optomechanical circuit via synthetic magnetism and reservoir engineering, *Nat. Phys.* **13**, 465 (2017).
- [31] M. Keck, D. Rossini, and R. Fazio, Persistent currents by reservoir engineering, *Phys. Rev. A* **98**, 053812 (2018).
- [32] H. Jirari, F. W. J. Hekking, and O. Buisson, Optimal control of superconducting n-level quantum systems, *Europhys. Lett.* **87**, 28004 (2009).
- [33] J. J. García-Ripoll, J. I. Cirac, P. Zoller, C. Kollath, U. Schollwöck, and J. von Delft, Variational ansatz for the superfluid Mott-insulator transition in optical lattices, *Opt. Express* **12**, 42 (2004).
- [34] R. Fazio and H. Van Der Zant, Quantum phase transitions and vortex dynamics in superconducting networks, *Phys. Rep.* **355**, 235 (2001).
- [35] Y. Ye, Z.-Y. Ge, Y. Wu, S. Wang, M. Gong, Y.-R. Zhang, Q. Zhu, R. Yang, S. Li, F. Liang *et al.*, Propagation and Localization of Collective Excitations on a 24-Qubit Superconducting Processor, *Phys. Rev. Lett.* **123**, 050502 (2019).
- [36] I. Bloch, J. Dalibard, and W. Zwerger, Many-body physics with ultracold gases, *Rev. Mod. Phys.* **80**, 885 (2008).
- [37] T. Haug, J. Tan, M. Theng, R. Dumke, L.-C. Kwek, and L. Amico, Readout of the atomtronic quantum interference device, *Phys. Rev. A* **97**, 013633 (2018).
- [38] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, Human-level control through deep reinforcement learning, *Nature (London)* **518**, 529 (2015).
- [39] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from [tensorflow.org](https://www.tensorflow.org).
- [40] A. Ramanathan, K. C. Wright, S. R. Muniz, M. Zelan, W. T. Hill, C. J. Lobb, K. Helmerson, W. D. Phillips, and G. K. Campbell, Superflow in a Toroidal Bose-Einstein Condensate: An Atom Circuit with a Tunable Weak Link, *Phys. Rev. Lett.* **106**, 130401 (2011).
- [41] C. Schenke, A. Minguzzi, and F. W. J. Hekking, Nonadiabatic creation of macroscopic superpositions with strongly correlated one-dimensional bosons in a ring trap, *Phys. Rev. A* **84**, 053636 (2011).
- [42] D. Aghamalyan, N. T. Nguyen, F. Auzstol, K. S. Gan, M. M. Valado, P. C. Condylis, L.-C. Kwek, R. Dumke, and L. Amico, An atomtronic flux qubit: A ring lattice of Bose-Einstein condensates interrupted by three weak links, *New J. Phys.* **18**, 075013 (2016).
- [43] J. R. Johansson, P. D. Nation, and F. Nori, Qutip 2: A python framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **184**, 1234 (2013).
- [44] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [45] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn *et al.*, Fast machine-learning online optimization of ultra-cold-atom experiments, *Sci. Rep.* **6**, 25890 (2016).
- [46] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong *et al.*, Unsupervised machine learning on a hybrid quantum computer, [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [47] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos *et al.*, Self-verifying variational quantum simulation of lattice models, *Nature (London)* **569**, 355 (2019).
- [48] S. Ragole and J. M. Taylor, Interacting Atomic Interferometry for Rotation Sensing Approaching the Heisenberg Limit, *Phys. Rev. Lett.* **117**, 203002 (2016).
- [49] <https://www.nscclsg>
- [50] M. Gluza, M. Kliesch, J. Eisert, and L. Aolita, Fidelity Witnesses for Fermionic Quantum Simulations, *Phys. Rev. Lett.* **120**, 190501 (2018).
- [51] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [52] Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot, and N. De Freitas, Dueling network architectures for deep reinforcement learning, [arXiv:1511.06581](https://arxiv.org/abs/1511.06581).
- [53] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold Bosonic Atoms in Optical Lattices, *Phys. Rev. Lett.* **81**, 3108 (1998).
- [54] G. Gauthier, I. Lenton, N. M. Parry, M. Baker, M. J. Davis, H. Rubinsztein-Dunlop, and T. W. Neely, Direct imaging of a digital-micromirror device for configurable microscopic optical potentials, *Optica* **3**, 1136 (2016).