Perspective

Theoretical and Experimental Perspectives of Quantum Verification

Jose Carrasco,[1] Andreas Elben,[2,3] Christian Kokail,[2,3] Barbara Kraus,[1] and Peter Zoller[2,3,*]

[1] *Institute for Theoretical Physics, University of Innsbruck, Innsbruck, Austria*
[2] *Center for Quantum Physics, University of Innsbruck, Innsbruck, Austria*
[3] *Institute for Quantum Optics and Quantum Information of the Austrian Academy of Sciences, Innsbruck, Austria*

(Received 19 January 2021; published 3 March 2021)

In this perspective we discuss verification of quantum devices in the context of specific examples formulated as proposed experiments. Our first example is verification of analog quantum simulators as Hamiltonian learning, where the input Hamiltonian as the design goal is compared with the parent Hamiltonian for the quantum states prepared on the device. The second example discusses cross-device verification on the quantum level (i.e., by comparing quantum states prepared on different quantum devices). We focus in particular on protocols using randomized measurements, and we propose establishing a central data repository, where existing experimental devices and platforms can be compared. In our final example, we address verification of the output of a quantum device from a computer science perspective, addressing the question of how a user of a quantum processor can be certain of the correctness of its output, and propose minimal demonstrations on present-day devices.

DOI: 10.1103/PRXQuantum.2.010102

I. INTRODUCTION

The dream and vision of now more than two decades to build quantum computers and quantum simulators has materialized as nascent programmable quantum devices in today’s laboratories [1–3]. While first-generation experiments focused on the basic demonstration of building blocks of quantum information processing, quantum laboratories now host programmable intermediate-scale quantum devices, which—while still imperfect and noisy—open the perspective of building quantum machines, which fulfill the promise of becoming more powerful than their classical counterparts. Significant advances in building small-scale quantum computers and quantum simulators have been reported with various physical platforms, from atomic and photonic systems to solid-state devices. A central aspect in further developments is verification of proper functioning of these quantum devices, including cross-device and cross-platform verification. Quantum verification is particularly challenging in regimes where comparison with classical simulation of quantum devices is no longer feasible.

Quantum characterization, validation, and verification is a well-developed field in quantum information theory, and we refer to reviews [4–6] and a tutorial [7] on this topic. The challenge in designing practical techniques to characterize quantum processes on intermediate-scale and large-scale quantum devices is related to the (in general) exponential scaling of the number of experiments and digital postprocessing resources with system size, as is manifest in quantum process tomography or state tomography. Exponential resources can be circumvented by extracting partial information about quantum processes providing a figure of merit, such as a process fidelity. However, such protocols also face the requirement of decoupling the state preparation and measurement errors from a process fidelity. Applications of well-established protocols in experimental settings, for example, as randomized or cycle benchmarking of quantum computers [8] or verifiable measurement-based quantum computation [9], have been reported.

In this perspective we look forward to possible near-future experiments addressing verification of quantum computers and quantum simulators, and in particular venturing into less explored territories. We illustrate aspects of verification that are physically relevant and conceptually complementary to previous work by describing three experimental scenarios as “proposed experiments.” Our discussion aims at connecting recent theoretical results with possible implementation of verification protocols in existing experimental settings. Clearly, different communities from quantum experimentalists to theorists and computer scientists look at perspectives on verification from quite different angles, and our examples are chosen to reflect this diversity.

[*] peter.zoller@uibk.ac.at

Our first example illustrates verification of analog quantum simulators [3,10] via Hamiltonian learning [11–13]. The central idea is to verify the analog quantum simulator by comparing the desired many-body Hamiltonian (i.e., the Hamiltonian to be implemented) with the actual, physically realized Hamiltonian, which can be efficiently reconstructed from measurements of quantum states prepared on the quantum device. This is applicable to and immediately relevant for present analog quantum simulation experiments for spin and Hubbard models with atoms and ions, and superconducting qubits [14–24].

In our second example we address cross-device and cross-platform verification as applicable to quantum computers and quantum simulators. Here the goal is the pairwise comparison of quantum states implemented on different quantum devices on the level of the full many-qubit wave function, or for reduced density matrices of subsystems. To this end, results of randomized measurements, performed on each device separately, can be classically correlated to estimate the fidelity of two quantum states, with efficiency scaling better with (sub)system size than what is achieved in quantum state tomography [25,26]. We envision a community effort where data from randomized measurements are uploaded to a central data repository, enabling the direct comparison of multiple quantum devices for a defined set of quantum problems, specified either as quantum circuits and algorithms or Hamiltonian evolution.

Finally, in our third example we move on to verification from a computer scientist perspective, and address the question of how a user of a quantum processor can be certain of the correctness of its output. This question becomes particularly important when the user of a quantum device does not have direct access to it (e.g., cloud computing). Is it even possible for users to rely on the result if they cannot verify it efficiently themselves? This question has been answered in the affirmative in the case where the user has access to a limited amount of quantum resources [27–34]. Such a verification of the output is feasible even via purely classical means [35]. However, not very surprisingly, the resources required to implement such a verification protocol are beyond reach with current technology. Because of rapid technological developments and the accompanying need for the ability to verify the output of a computation, we propose here a proof-of-principle experiment to implement such a verification protocol that is feasible with current technologies.

II. VERIFICATION OF ANALOG QUANTUM SIMULATORS VIA HAMILTONIAN LEARNING

The goal of quantum simulation is to solve the quantum many-body problem [10], from strongly correlated quantum materials in condensed matter physics [15] to quantum field theories in high-energy physics [36], or modeling of complex molecules and their dynamics in quantum chemistry [20,37]. Building an analog quantum simulator amounts to realizing in the laboratory synthetic, programmable quantum matter as an isolated quantum system. Here, first of all, a specified many-body Hamiltonian H must be implemented faithfully in a highly controllable quantum system with given physical resources. Furthermore, quantum states of matter must be prepared on the physical quantum device corresponding to equilibrium phases (e.g., as ground states) or must represent nonequilibrium phenomena as in quench dynamics.

Remarkable progress has been made recently in building analog quantum simulators to emulate quantum many-body systems. Examples include the realization of lattice spin models with trapped ions [22,23], Rydberg tweezer arrays [16–18], superconducting devices [24], and Hubbard models with ultracold bosonic or fermionic atoms in optical lattices [15,19,21]. While analog quantum simulation can be viewed as special-purpose quantum computing with the rather focused task of emulating a many-body system via a specified H, the unique experimental feature is the ability to scale to rather large particle numbers. This is in contrast to present-day quantum computers, which provide a high-fidelity universal gate set for a small number of qubits.

Today's ability of analog quantum simulators to prepare and store on a scalable quantum device a highly entangled many-body state, while solving a quantum problem of physical relevance, fulfills one of the original visions of Feynman's proposal of quantum simulation. However, this also raises the question of verification in regimes where comparison with classical computations with controlled error, such as tensor network techniques, is no longer possible. This includes also higher-dimensional lattice models, with fermionic particles, and quench dynamics.

The proper functioning of a quantum simulator can be ensured by comparing experiment with theory [38] or predictions from two different experimental quantum devices. This can be done on the level of comparing expectation values of relevant observables (e.g., on the most elementary level by comparing phase diagrams [38] or the increasingly complex hierarchies of correlation functions [39]). We return to approaches for directly comparing quantum states in Sec. III.

Verification by Hamiltonian Learning: Instead, we rephrase here "verification of an analog quantum simulator" as comparing the "input" Hamiltonian, specified as the design goal for the quantum simulator, with the actual Hamiltonian realized on the physical device. This latter, experimental Hamiltonian can be determined via "Hamiltonian tomography," or "Hamiltonian learning"; that is, inferring from measurements under certain conditions the parent Hamiltonian underlying the experimentally prepared quantum state [11,12].

Hamiltonians of many-body physics consist of a small set of terms which are (quasi)local and consist of few-body interactions (i.e., $H = \sum_i h_i$, with h_i quasilocal terms). Thus, for a given H, only a small set of physical parameters determines the accessible quantum states and their entanglement structure: for example, as a ground state, $H|\Psi_G\rangle = E_G|\Psi_G\rangle$; as a finite-temperature state in the form of a Gibbs ensemble $\sim \exp(-\beta H)$; or as generator of the quench dynamics with an initial (pure) state $|\Psi_0\rangle$ evolving in time as $|\Psi_t\rangle = \exp(-iHt)|\Psi_0\rangle$.

Remarkably, as shown recently [11–13,40], it is the local and few-body structure of physical Hamiltonians in operator space that allows efficient Hamiltonian tomography via measurements from experimentally prepared (single) quantum states on the quantum simulator. These states include the ground state, a Gibbs state, and states produced in quench dynamics. It is thus the restricted operator content of Hamiltonians that promises scalable Hamiltonian learning with system size (i.e., makes Hamiltonian tomography efficient).

Here we outline “Hamiltonian verification” for a Fermi-Hubbard model. This can be implemented with atoms in an optical lattice and can be observed with a quantum gas microscope [15,19]. To be specific, we apply the protocol in Ref. [11] for reconstruction of the parent Hamiltonian from an experimentally prepared ground state. Similar results apply to energy eigenstates, thermal states, or any stationary state. We simulate experimental runs of the protocols including the measurement budget, thus assessing accuracy and convergence [41].

The protocol in Ref. [11] describes learning of local Hamiltonians from local measurements. The starting point is the assumption of an experimentally prepared stationary state ρ, as described above. The protocol finds the parent Hamiltonian H from ρ via the steady-state condition $[H,\rho] = 0$. As ρ is stationary under H, so is the expectation value of any observable A: $\partial_t\langle A\rangle = \langle -i[A,H]\rangle = 0$. The latter equation can be used to obtain a set of linear constraints from which H can be reconstructed. Consequently, for lattice systems the algorithm can be summarized as follows [see also Fig. 1(a)]:

(1) Expand the Hamiltonian H in a local operator basis $H = \sum_{m=1}^{M} c_m S_m$. For a k-local H, $M \sim \mathcal{O}(L^k)$ basis elements are required, with L the number of lattice sites.
(2) Select a set of $N_C > M$ linearly independent operators (constraints) $\{A_n\}_{n=1}^{N_C}$.
(3) Construct a system of equations $\langle -i[A_n,H]\rangle = \sum_m c_m \langle -i[A_n,S_m]\rangle = K\mathbf{c} = 0$ by measuring the matrix elements $K_{nm} = \langle -i[A_n,S_m]\rangle$.
(4) Find the lowest right singular vector $\tilde{\mathbf{c}}$ of K that minimizes the norm $\|K\mathbf{c}\|$.

As stated in Ref. [12], the locality of H implies that such a Hamiltonian reconstruction will be unique. The reconstructed parameters $\tilde{\mathbf{c}}$ can be cross-checked with respect to the parameters of an input Hamiltonian and serve as a quantifier for the verification of the quantum simulator. The number of experimental runs required is controlled by the gap of the correlation matrix $\mathcal{M} = K^T K$, which strongly depends on the type and number of constraints [11]. In the limit of all possible constraints, the matrix $\mathcal{M}$ coincides with the correlation matrix defined by Qi and Ranard [12]. The lowest eigenvalue of this matrix corresponds to the Hamiltonian variance measured on the input state, which was used previously for experimental verification of variationally prepared many-body states [22].

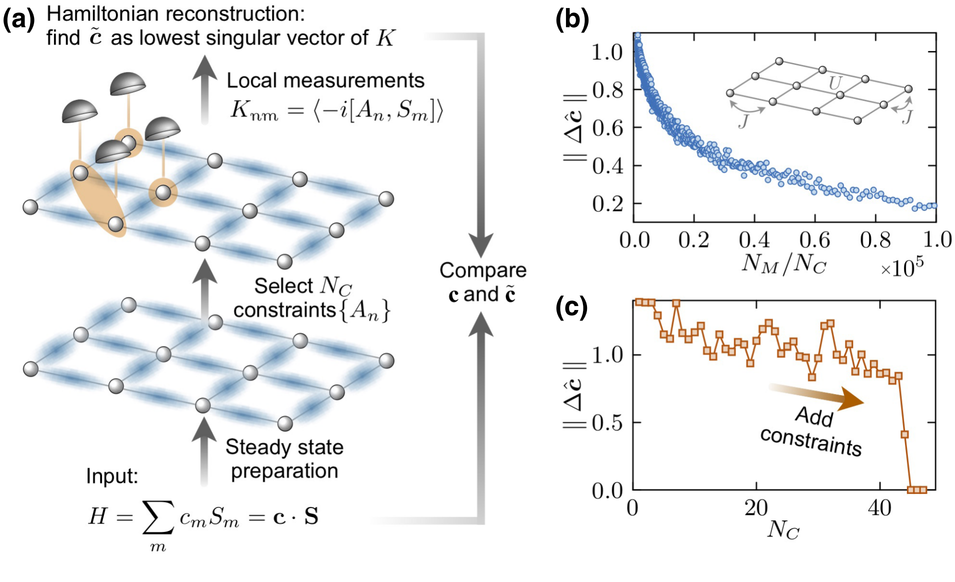

FIG. 1. (a) The protocol for verifying a local Hamiltonian from local measurements. (b) Number of measurements (number of experimental runs) required per constraint to achieve a certain parameter distance $\|\Delta\hat{\mathbf{c}}\| = \|\hat{\mathbf{c}} - \tilde{\hat{\mathbf{c}}}\|$ for a 3×4 isotropic Fermi-Hubbard model with hole doping ($N_f = 10$ fermions). Here $\hat{\mathbf{c}}$ denotes the normalized vector $\hat{\mathbf{c}} = \mathbf{c}/\|\mathbf{c}\|$, with $\mathbf{c} = \{c_m\}$ the expansion coefficients of the Hamiltonian H in the local operator basis: $H = \sum_m c_m S_m$. (c) Smallest achievable parameter distance $\|\Delta\hat{\mathbf{c}}\|$ as a function of the number of constraints N_C.

In Figs. 1(b) and 1(c) we illustrate Hamiltonian learning for a Fermi-Hubbard model

$$H = -J \sum_{\langle ij \rangle \sigma} \left(c^{\dagger}_{i\sigma} c_{j\sigma} + \text{H.c.} \right) + U \sum_{i} n_{i\uparrow} n_{i\downarrow} \qquad (1)$$

on a two-dimensional square lattice [41]. Here $c^{\dagger}_{i\sigma}$ ($c_{i\sigma}$) denotes creation (annihilation) operators of spin-1/2 fermions at lattice sites i, and $n_{i\sigma} = c^{\dagger}_{i\sigma} c_{i\sigma}$. Consequently, in this example the local basis $\{S_m\}_{m=1}^{M}$ consists of hopping operators for all bonds (i,j) $(c^{\dagger}_{i\sigma} c_{j\sigma} + \text{H.c.})$ for each spin component σ and of operators counting double occupancies on the individual sites i, $n_{i\uparrow} n_{i\downarrow}$. In the case of the 3×4 lattice studied in Fig. 1, the operator basis therefore includes $M = 46$ elements. As an input state for the protocol we take the ground state in the strongly repulsive regime ($J = 1$, $U = 8$) and introduce a small hole doping of $n = 0.83$. As a set of constraints we adopt the operators $A_{ijk} = i(c^{\dagger}_{i\sigma} c_{j\sigma} - \text{H.c.}) n_{k\sigma'}$, in which i, j, and k are nearest-neighbor sites [42]. The particular combinations of sites $\{i, j, k\}$ is chosen in such a way that the rows of the matrix K are linearly independent. Obtaining the matrix elements $K_{nm} = \langle -i[A_n, S_m] \rangle$ requires the measurement of locally resolved atomic currents $\mathcal{J}^{\sigma}_{i \leftrightarrow j} = i(c^{\dagger}_{i\sigma} c_{j\sigma} - \text{H.c.})$, where j can be located within two lattice constants around i. In experiments with atoms in optical lattices, these currents can be accessed by inducing superexchange oscillations accompanied by spin-resolved measurements in a quantum gas microscope [43,44].

Figure 1(b) shows the relation between the distance of the exact versus the reconstructed Hamiltonian parameters $\| \Delta \hat{\mathbf{c}} \|$ and the number of measurements per constraint on a 3×4 Hubbard lattice. Figure 1(c) displays the improvement in quality of the Hamiltonian reconstruction as additional constraints A_{ijk} are added to the system of equations $K\mathbf{c} = 0$. As can be seen, the Hamiltonian can be recovered exactly as the number of constraints N_C approaches the number of elements M in the operator basis $\{S_m\}_{m=1}^{M}$. The total measurement budget can be optimized via arrangement of the operators $[A_n, S_m]$ into commuting groups such that they can be evaluated from the same measurement outcomes [41].

In the Hamiltonian learning protocol outlined above, the number of measurements required to obtain a fixed parameter distance $\| \Delta \hat{\mathbf{c}} \|$ scales polynomially with the system size [11]. As pointed out in Ref. [41], the method described here is robust against uniform measurement errors $K \to (1 - \epsilon) K$ (where ϵ is the error rate) since the Hamiltonian learning procedure is not influenced when the matrix K is multiplied by a global sclaing factor. Recent work demonstrates that the method can be extended to recover Linbladians from steady states, potentially allowing efficient recovery of dissipative processes [45]. Future investigations will have to include the relation of the type and number of constraints to the gap of the correlation matrix that determines the total number of experimental runs required, as well the role of measurement errors and decoherence (see, e.g., Ref. [46]).

An entirely different verification protocol, which can also be applied to quantum simulation, is cross-device verification, which is described in the following section. There, verification is achieved by cross-checking of the results from two quantum simulators simulating the same physics by measurement of overlaps of quantum states on the level of reduced density operators for various subsystem sizes.

III. CROSS-DEVICE VERIFICATION OF QUANTUM COMPUTATIONS AND QUANTUM SIMULATIONS

In the previous section, we presented the verification of an analog quantum simulator by comparing the Hamiltonian actually realized in the device with the input or target Hamiltonian. A different approach to verification, aiming to gain confidence with regard to the output of a quantum simulation or quantum computation is to run the simulation or computation on various different quantum devices and compare the outcomes with each other and—if available—with an idealized theoretical simulation [for an illustration see Fig. 2(a)]. Such a cross-comparison can be implemented on different levels of sophistication. While quantum simulations have been compared on the level of low-order observables [38], for instance, order parameters characterizing phase diagrams, recent protocols aim to compare full quantum states [25,47–49].

A convenient measure of comparison of two, possibly mixed, quantum states ρ_1 and ρ_2 is the quantum fidelity $F(\rho_1, \rho_2)$. While various definitions of fidelities exist, they have in common that $F(\rho_1, \rho_2) = 1$ for identical states $\rho_1 = \rho_2$ and that $F(\rho_1, \rho_2) = |\langle \psi_1 | \psi_2 \rangle|^2$ for pure states $\rho_1 = |\psi_1\rangle\langle\psi_1|$ and $\rho_2 = |\psi_2\rangle\langle\psi_2|$. A standard choice is provided by the Uhlmann fidelity $F_U(\rho_1, \rho_2) = \text{Tr}\sqrt{\sqrt{\rho_1}\rho_2\sqrt{\rho_1}}$ [50], which is in general difficult to access experimentally and also via numerical calculations [51]. Simpler alternatives include Hilbert-Schmidt fidelities [51], such as

$$F_{\max}(\rho_1, \rho_2) = \frac{\text{Tr}(\rho_1 \rho_2)}{\max\{\text{Tr}\,\rho_1^2, \text{Tr}\,\rho_2^2\}}, \qquad (2)$$

where $F_{\max}$ defines a metric on the space of density matrices [51] and coincides with the Uhlmann fidelity if one state is pure [51].

To measure quantum fidelities, various approaches exist. A pure quantum protocol would establish a quantum link, teleport quantum states, and compare states locally, for instance, via a SWAP test [53–55]. While such overlap measurements have been demonstrated locally in seminal experiments [56–58], a quantum link teleporting large

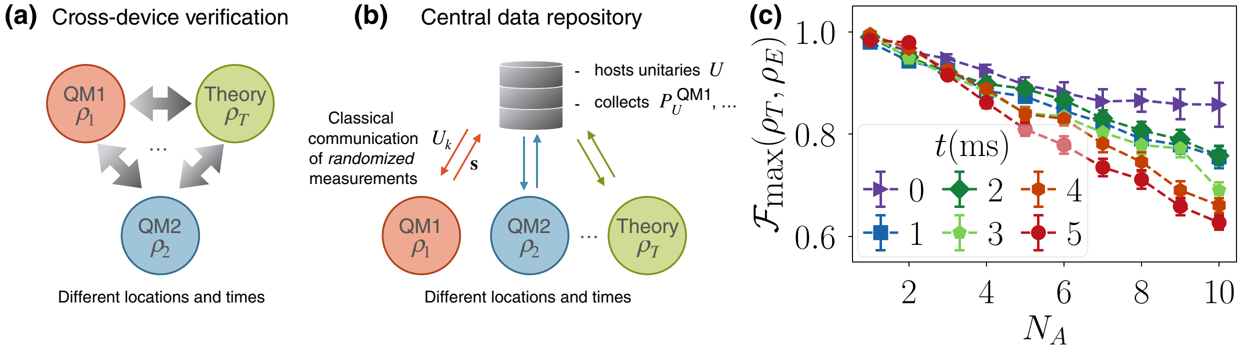

FIG. 2. (a) Cross-device verification of unknown quantum states ρ_1, ρ_2, etc. prepared on different quantum machines QM1, QM2, etc. based on various platforms via (b) a central, standardized data repository and classical communication of randomized measurements (local random unitaries and measurement outcomes). In regimes where a classical simulation is possible, the implemented states can additionally be compared with a theoretical target state ρ_T. (c) Experiment-theory fidelities between a quantum state prepared in a trapped ion quantum simulator and its classical simulation as a function of the subsystem size N_A (the total system consists of ten qubits) for various evolution times (different colors) after a quantum quench in a long-range Ising model [52]. Reprinted from Ref. [25].

quantum states of many particles with high accuracy between two quantum devices is not expected to be achievable in the near future.

Today, protocols relying on classical communication between many-body quantum devices are thus required. Here, ultimate brute force tests include quantum state and quantum progress tomography, which aim for a full classical reconstruction, allowing a classical comparison, of quantum states or processes. Even with incorporation of recent advances, such as compressed sensing for density matrices with low rank [59], such approach requires, however, at least 3^N measurements to accurately determine an arbitrary N-qubit state. Efficient methods, such as tensor network [60,61] or neural network tomography [62], have been developed, but rely on a special structure of the states of interest.

A direct approach toward fidelity estimation is provided by randomized measurements [25,49,52,63–67]. Here, a randomized measurement on an N-qubit quantum state ρ is performed by the application of a unitary U, chosen at random from a tomographically complete set and a subsequent measurement in the computational basis $\{|\mathbf{s}\rangle\}$. Statistical correlations of such randomized measurements, performed sequentially on a single quantum device, allow tomographic reconstruction of the quantum state [49,65,68] but also give direct access to nonlocal and nonlinear (polynomial) functionals of density matrices such as Rényi entropies [49,63,64]. In particular, recent work [49] combined randomized measurements with the notion of shadow tomography [69], which aims to predict directly expectation values of arbitrary observables, instead of reconstructing the full density matrix. Using insights from the stabilizer formalism [70], Huang *et al.* [49] devised an efficient implementation of shadow tomography via randomized measurements that enables one to estimate expectation values of arbitrary (multicopy) observables with high precision and rigorous performance guarantees [49]. This allows one, in particular, to estimate the fidelity between the quantum state ρ and a known theoretical target. It complements methods such as direct fidelity estimation [47,48] and randomized benchmarking [8,71–75], which use the absolute knowledge of the theoretical target to be efficient for certain target states and processes.

Cross-device verification with randomized measurements: In a very general setting, one faces the situation where two unknown quantum states have been prepared on two separate quantum devices, potentially at very different points in space and time [Fig. 2(a)]. In Ref. [25] (see also Ref. [26]), it was proposed to measure the cross-device fidelity $F_{\max}(\rho_1, \rho_2)$ of two unknown quantum states, described by (reduced) density matrices ρ_1 and ρ_2 and prepared on two separate devices. To this end, randomized measurements are implemented with the same random unitaries U on both devices. Facilitating the direct experimental realization, these unitaries U can be local, $U = \bigotimes_{k=1}^{N} U_k$, with U_k acting on qubit k and sampled from a unitary 2 design [72,76] defined on the local Hilbert space $\mathbb{C}^2$. From statistical cross-correlation and autocorrelation of the outcome probabilities $P_U^{(i)}(\mathbf{s}) = \mathrm{Tr}(U \rho_i U^{\dagger} |\mathbf{s}\rangle \langle \mathbf{s}|)$ of the randomized measurements, the overlap $\mathrm{Tr}\, \rho_1 \rho_2$ and purities $\mathrm{Tr}\, \rho_1^2$ ($\mathrm{Tr}\, \rho_2^2$), and thus $F_{\max}(\rho_1, \rho_2)$, are estimated via

$$\mathrm{Tr}\,(\rho_i \rho_j) = 2^N \sum_{\mathbf{s},\mathbf{s}'} (-2)^{-\mathcal{D}[\mathbf{s},\mathbf{s}']}\, \overline{P_U^{(i)}(\mathbf{s}) P_U^{(j)}(\mathbf{s}')} \qquad (3)$$

for $i,j = 1,2$. Here $\overline{\ldots}$ denotes the ensemble average over local random unitaries, and the Hamming distance

$\mathcal{D}[\mathbf{s}, \mathbf{s}']$ between two strings $\mathbf{s}$ and $\mathbf{s}'$ is defined as $\mathcal{D}[\mathbf{s}, \mathbf{s}'] \equiv \left|\left\{k \in \{1, \ldots, N\} \mid s_k \neq s'_k\right\}\right|$.

In the regime where a classical simulation of the output is possible, this protocol can also be used for an experiment-theory comparison (c.f. direct fidelity estimation [47,48] and classical shadow tomography [49]). In Fig. 2(c), experiment-theory fidelities between highly entangled quantum states prepared via quench dynamics in a trapped ion quantum simulator [52] and the theoretical simulation are shown [25]. Such experiment-theory comparisons for simple (product) states can also be used to identify and mitigate errors resulting from imperfect measurements [25,77,78].

On the basis of numerical simulations, it was found in Ref. [25] that the number of experimental runs necessary to estimate the fidelity $F_{\max}$ up to a fixed statistical error scales exponentially with the subsystem size, approximately 2^{bN}. In comparison with quantum state tomography, exponents $b \lesssim 1$ are, however, favorable, enabling fidelity estimation for (sub)systems consisting of a few tens of qubits with state-of-the art quantum devices. For two very large quantum devices, consisting of several tens of qubits to a few hundred qubits, the present protocol thus allows one to estimate fidelities of possibly disconnected subsystems only up to a given size, which is determined by the available measurement budget. These data represent very fine-grained local information on fidelities of subsystems. It remains an open question whether this information can be combined with additional knowledge of a few global properties to obtain (at least bounds on) the total system fidelity.

While we outlined above protocols to cross-check two individual devices, we envision a community effort where specific quantum problems, either as quantum circuits and algorithms or for quantum simulation, are defined and data from theoretical simulations as well as measurement data from quantum devices are uploaded to a central data repository [for an illustration see Fig. 2(b)]. In regimes where a classical simulation is possible, an ultimate reference could here be represented by a theory target state. For larger quantum devices, reference operations and circuits could be executed and density matrices of (sub)systems could be compared with each other. This would allow a standardized, pairwise cross-check of multiple quantum devices representing various platforms.

The protocols outlined rely on classical communication of randomized measurement results and are restricted, due to an exponential scaling of the number of experimental runs required, to (sub)systems of a few tens of qubits. To overcome this challenge, we expect that in the future quantum state transfer protocols will become available to develop efficient fully quantum protocols in addition to hybrid quantum-classical protocols for cross-checking quantum devices.

IV. VERIFICATION OF THE OUTPUT OF AN UNTRUSTED QUANTUM DEVICE

In the validation procedures considered above, the person testing the quantum processor (the user) either has direct access to the device or trusts the person operating it. Computer scientists are often concerned about a very different notion of verification: the verification of the output of a computation performed by an untrusted device. Such a verifiability demand will become particularly relevant once quantum devices that reliably process hundreds of qubits become usable as cloud computers.

To demonstrate the need for these verification protocols, let us consider the various kinds of problems such cloud computers could be used for. If the user uses a quantum computer to solve a problem within nondeterministic polynomial time, such as factoring a large number into its prime factors, the solution to the problem is simple: knowing the factors, the output can be efficiently verified with a classical computer. However, it is believed that quantum computers are capable of efficiently solving problems that can no longer be efficiently verified classically, such as simulating quantum many-body systems. How can one then rely on the output, given that the quantum computer (or the person operating it) might be malicious and want to convince the user that the answer to, for example, a decision problem is “yes” when it is actually “no”? Hence, harnessing the full power of a quantum device that is not directly accessible to the user brings with it the necessity to derive protocols for verifying its output. The aim here is to derive quantum verification protocols that allow a computationally limited (e.g., a classical) user to verify the output of a (powerful) quantum computer. Complicating matters is the need to ensure that a honest prover (the quantum computer) can convince the user of the correct outcome efficiently [79]. To express things more simply, we refer now to the user (called the “verifier”) as Alice (A) and to the prover as Bob (B).

Verification protocols [80] where A has access to limited quantum resources [27–34] or is able to interact with two noncommunicating provers have been derived [81]. In a recent breakthrough Mahadev [35] showed that even a purely classical user can verify the output of a quantum processor. In contrast to the verification protocols mentioned before, this protocol relies on a computational assumption: the existence of trapdoor functions that are postquantum secure [82]. These functions are hard to invert even for a quantum computer. However, the possession of additional information (trapdoor) enables one to compute the preimages of the function efficiently. Use of the notion of postquantum secure trapdoor functions in combination with powerful previously derived findings led to the surprising result that a classical user can indeed verify the output of a quantum computer, as will briefly explain below. The notion and techniques developed in Refs. [35,83,84] were recently utilized to propose protocols with, for example, zero-knowledge polynomial-time verifiers [85] and noninteractive classical verification [86].

At first glance it simply seems impossible to efficiently verify the output of a much more powerful device (even if it were classical) if one is just given that output and is prevented from testing the device. The key idea here is to use interactive proofs. The exchange of messages allows A to test B and to eventually be convinced that B's claim is indeed correct or to mistrust him and reject the answer. The graph nonisomorphism problem is a simple example of a task where the output (of a powerful classical device) can be verified with an interactive proof [87].

To explain the general idea of how to verify the output of a quantum device, we assume that B possesses a quantum computer, whereas A has only classical computational power. A asks B to solve a decision problem (within bounded-error quantum polynomial time, i.e., a problem that can be solved efficiently by a quantum computer) and wants to verify the answer. Of particular importance here is that one can show that the outcome of such a decision problem can be encoded in the ground-state energy of a suitable, efficiently computable, local Hamiltonian H [88]. This implies that if B claims that the answer to the decision problem is "yes" [89], he can convince A of this fact by preparing a state with energy (with respect to H) below a certain value, which would be impossible if the correct answer were "no." An instance of such a state is the so-called clock state, $|\eta\rangle$ [90,91], which can be prepared efficiently by a quantum computer. Hence, the output of the quantum computer can be verified by determining the energy of the state prepared by B. This can be achieved by performing measurements only in the X basis as well as the Z basis [92–94]. It remains to ensure that A can delegate these measurements to B without revealing the measurement basis. The important contribution of Mahadev [35] is the derivation of such a measurement protocol (see Fig. 3). The properties of postquantum secure trapdoor functions are exploited precisely at this point to ensure that B cannot learn whether a qubit is measured in the Z basis or the X basis, which prevents him from cheating.

For reasonable choices of the security parameters, the realization of such a verification protocol is, even without considering faulty devices, not feasible with current technology (on B's side). Already the number of auxiliary qubits required in the measurement protocol would be too demanding [83,93]. Nevertheless, because of rapid technological development and the accompanying need for these kinds of verification protocols, we present here a proposal for a proof-of-principle experiment. The minimal example explained here can already be performed with a total of seven qubits. First, the original decision problem is mapped to a local Hamiltonian problem. B prepares the corresponding state, $|\eta\rangle$ (consisting of four qubits in this example), whose energy needs to be determined. Because of the linearity of the measurement protocol, it is sufficient to demonstrate how A can delegate the measurement in the X basis or Z basis on a single-qubit state $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ (belonging to $|\eta\rangle$) without revealing

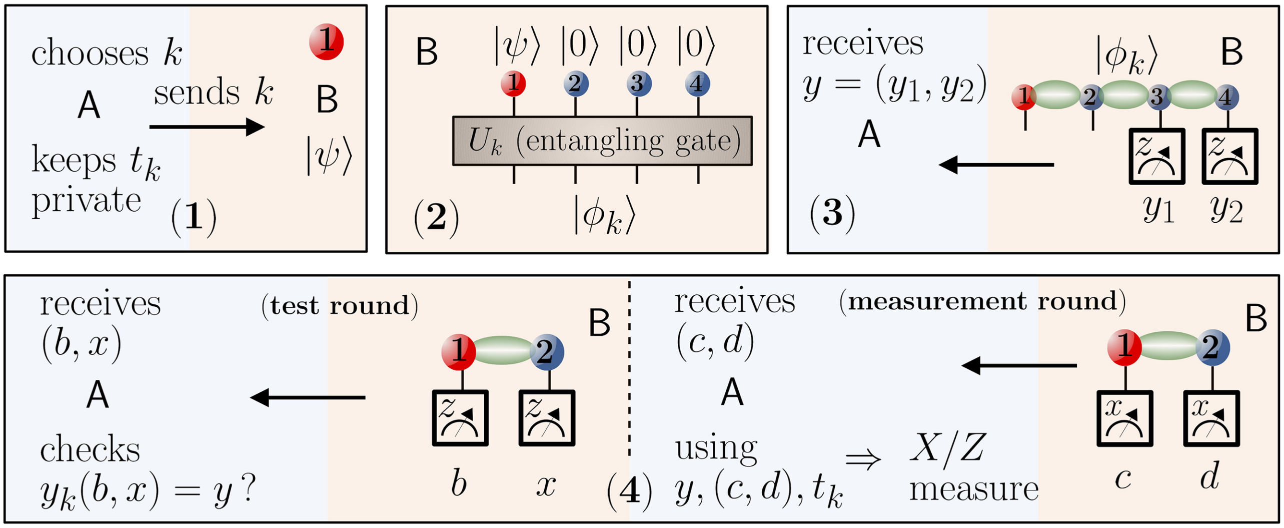

FIG. 3. Sketch of steps 1–4 in the measurement protocol [35] (see the main text). Alice (A) represents a classical verifier, Bob (B) represents an efficient quantum prover, and we use arrows to emphasize the flow of classical information. A performs a delegated measurement on a single-qubit state $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ in either the Z basis or the X basis without revealing the measurement basis to B. An appropriate combination of the statistics of the results of such measurements allows A to calculate the energy of the state prepared by B and thereby verify the outcome of the original decision problem.

the measurement basis to B. The measurement protocol has the following high-level structure (see Fig. 3):

(1) A computes a family of postquantum secure trapdoor functions $\{y_k\}$, labeled by an index k, together with the associated information t_k (trapdoor). The functions y_k are of one of two types, either one-to-one or two-to-one functions [96]. If A wants to measure $|\psi\rangle$ in the Z basis (X basis), she chooses a label k such that y_k is a one-to-one (two-to-one) function. A keeps t_k private (this is precisely the leverage A has over B) and sends k to B. Knowing k, B can efficiently evaluate the function y_k on any input. However, it is computationally hard for him to determine which type of function y_k is. Furthermore, A can compute the preimages of y_k efficiently using t_k, while B cannot.

(2) B is asked to prepare the state $|\phi_k\rangle \propto \sum_{b,x} \alpha_b |b\rangle |x\rangle |y_k(b,x)\rangle$. This can be done efficiently by a quantum computer.

(3) B is asked to measure the last register (qubits 3 and 4 in our example) of $|\phi_k\rangle$ in the Z basis and to send the measurement outcome y to A. The state of the first and second registers (qubits 1 and 2 in our example) is then, depending on the type of y_k either (i) the product state $|b\rangle |x\rangle$ (with probability $|\alpha_b|^2$), where $y_k(b,x) = y$, or (ii) the entangled state $\alpha_0 |0\rangle |x_0\rangle + \alpha_1 |1\rangle |x_1\rangle$, where $y_k(0,x_0) = y_k(1,x_1) = y$.

(4) A randomly chooses to run either a "test" or a "measurement" round. In a "test" ("measurement") round, B is asked to measure the qubits in the first and second registers in the Z basis (X basis) and to send the outcome to A. The "test" rounds allow A to gain confidence that B has indeed prepared $|\phi_k\rangle$ by checking that $y_k(b,x) = y$. In a "measurement" round the first qubit is effectively measured in either the Z basis or the X basis, depending on the type of y_k. Using the trapdoor information t_k, A can classically postprocess the outputs to obtain the corresponding measurement outcome.

As mentioned above, a minimal, nontrivial example that can be realized with an ion-trap quantum computer [97] requires only seven qubits in total and some tens of single-qubit and two-qubit gates [98]. In this case the clock state $|\eta\rangle$ is a four-qubit state and, for this minimal example, one can choose the second and third registers to have one and two qubits, respectively (as shown in Fig. 3). Here $y_k : \{0,1\}^2 \rightarrow \{0,1\}^2$ and k labels either one of the 24 one-to-one functions or one of the 24 two-to-one functions.

We finally mention that protocols that allow the verification of the output of imperfect quantum computers were recently proposed for the case where the verifier has limited access to quantum resources [99]. Similar ideas can also be used in the purely classical verification protocol [35], ensuring that the measurements can still be performed without jeopardizing its security [97].

V. CONCLUSION AND OUTLOOK

In an era when we are building noisy intermediate-scale quantum devices, with the effort to scale them to larger system sizes and optimize their performance, verification of quantum devices is becoming a main aim in theoretical and experimental quantum information science. In this perspective on theoretical and experimental aspects of quantum verification, we discuss three examples formulated as proposed experiments. The three examples are verification of quantum simulation via Hamiltonian learning (Sec. II), cross-checking of quantum states prepared on different quantum devices (Sec. III), and addressing the question of how a user of a quantum processor can be certain of the correctness of its output (Sec. IV). While our choice of examples highlighting quantum verification is subjective and guided by personal interests, the common theme is that these "proposed experiments" can be performed with the quantum devices existing in today's quantum laboratories or near-future devices. In addition, our examples illustrate the diversity of questions in quantum verification and tools and techniques to address them, with emphasis on what we identify as problems of high relevance.

Of course, by the very nature of a perspective as forward looking, we identify interesting topics and outline possible avenues, while highlighting open issues for future theoretical and experimental work. These open problems range from technical to conceptual issues, and we summarize some of these questions within the various sections. The overarching challenge is, of course, to develop efficient and quantitative verification protocols and techniques that will eventually scale to large system sizes we envision as useful for quantum devices. In Sec. II on verification of analog quantum simulation via Hamiltonian learning, the local Hamiltonian ansatz scales, by construction, with the system size and leads, in principle, to a quantified error assessment. While one may raise issues of imperfect state preparation and measurement errors in experiments, and the measurement budget available in a specific experiment, we emphasize that these protocols also involve heavy classical postprocessing of data, which may provide limits from a practical and conceptual perspective. While this might not pose serious limitations for near-future devices, we might ask here, but also in a broader context, if some of this postprocessing can be replaced by more efficient quantum postprocessing on the device. The cross-device check of quantum states in Sec. III provides another example of this type. There, the protocol underlying the comparison of quantum states with a central data repository involves classical communication. The protocol described is much more efficient than tomography, and scales with a "friendly exponential" in system size, allowing today experimental implementation for tens of qubits. A future development of quantum state transfer as quantum communication between the devices promises to overcome these limitations. Finally, our discussion in Sec. IV on verification of the output of an untrusted quantum device presents an absolute minimal example that can be run on present quantum computers, leaving as challenges the verification of outputs of imperfect quantum devices and more advanced experimental demonstrations.

Verification of quantum processors is particularly challenging and relevant in the regimes of quantum advantage, where quantum devices outperform their classical counterparts [100,101]. As solving a "useful" computational task (such as factoring a large number) would neither be feasible with noisy intermediate-scale quantum computers nor necessary to demonstrate quantum superiority, one focuses on sampling problems [102–104] in this context. However, these approaches entail difficulties in demonstrating quantum superiority. On the one hand, the fact that the sampling was performed faithfully needs to be verified. On the other hand, one needs to show that the task is computationally hard for any classical device (taking into account that the quantum computer is imperfect). In this context, both strong complexity-theoretical evidence of classical intractability and new proposals for experimental realizations for various setups are desirable.

ACKNOWLEDGMENTS

Work at Innsbruck is supported by the European Union Horizon 2020 program under Grants Agreement No. 817482 (PASQuanS) and No. 731473 (QuantERA via QTFLAG), the US Air Force Office of Scientific Research via IOE Grant No. FA9550-19-1-7044 LASCEM, the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (651440, PZ), the Austrian Science Fund (FWF) by stand-alone project P32273-N27, SFB BeyondC, and the Institut für Quanteninformation.

[1] I. H. Deutsch, Harnessing the power of the second quantum revolution, PRX Quantum **1**, 020101 (2020).

[2] J. Preskill, Quantum computing in the nisq era and beyond, Quantum **2**, 79 (2018).

[3] E. National Academies of Sciences and Medicine, *Manipulating Quantum Systems: An Assessment of Atomic, Molecular, and Optical Physics in the United States* (The National Academies Press, Washington, DC, 2020).

[4] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, Nat. Rev. Phys. **2**, 382 (2020).

[5] A. Gheorghi, T. Kapourniotis, and E. Kashefi, Verification of quantum computation: An overview of existing approaches, Theory Comput. Syst. **63**, 715 (2019).

[6] I. Šupić and J. Bowles, Self-testing of quantum systems: A review, Quantum **4**, 337 (2020).

[7] M. Kliesch and I. Roth, Theory of quantum system certification – a tutorial, arXiv:2010.05925 [quant-ph] (2020).

[8] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, Nat. Commun. **10**, 5347 (2019).

[9] S. Barz, J. Fitzsimons, E. Kashefi, and P. Walther, Experimental verification of quantum computation, Nat. Phys. **9**, 727 (2013).

[10] J. I. Cirac and P. Zoller, Goals and opportunities in quantum simulation, Nat. Phys. **8**, 264 (2012).

[11] E. Bairey, I. Arad, and N. H. Lindner, Learning a Local Hamiltonian from Local Measurements, Phys. Rev. Lett. **122**, 020504 (2019).

[12] X.-L. Qi and D. Ranard, Determining a local hamiltonian from a single eigenstate, Quantum **3**, 159 (2019).

[13] Z. Li, L. Zou, and T. H. Hsieh, Hamiltonian Tomography via Quantum Quench, Phys. Rev. Lett. **124**, 160502 (2020).

[14] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, and K.-M. C. Fu *et al.*, Quantum simulators: Architectures and opportunities, arXiv:1912.06938.

[15] L. Tarruell and L. Sanchez-Palencia, Quantum simulation of the hubbard model with ultracold fermions in optical lattices, C. R. Phys. **19**, 365 (2018).

[16] A. Browaeys and T. Lahaye, Many-body physics with individually controlled rydberg atoms, Nat. Phys. **16**, 132 (2020).

[17] M. Morgado and S. Whitlock, Quantum simulation and computing with rydberg qubits, arXiv:2011.03031.

[18] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, and W. W. Ho *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, arXiv:2012.12281.

[19] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, A cold-atom fermi–hubbard antiferromagnet, Nature **545**, 462 (2017).

[20] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, and J. I. Cirac, Analogue quantum chemistry simulation, Nature **574**, 215 (2019).

[21] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu, and Y. Takahashi, Tools for quantum simulation with ultracold atoms in optical lattices, Nat. Rev. Phys. **2**, 411 (2020).

[22] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, Nature **569**, 355 (2019).

[23] C. Monroe, W. Campbell, L.-M. Duan, Z.-X. Gong, A. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, and G. Pagano *et al.*, Programmable quantum simulations of spin systems with trapped ions, arXiv:1912.07845.

[24] H.-L. Huang, D. Wu, D. Fan, and X. Zhu, Superconducting quantum computing: A review, Sci. China Inf. Sci. **63**, 1 (2020).

[25] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller, Cross-Platform Verification of Intermediate Scale Quantum Devices, Phys. Rev. Lett. **124**, 010504 (2020).

[26] S. Flammia, Quantum computer crosscheck, Physics **13**, 3 (2020).

[27] T. Morimae and K. Fujii, Blind quantum computation protocol in which alice only makes measurements, Phys. Rev. A **87**, 050301(R) (2013).

[28] T. Morimae, Verification for measurement-only blind quantum computing, Phys. Rev. A **89**, 060302(R) (2014).

[29] M. Hayashi and T. Morimae, Verifiable Measurement-Only Blind Quantum Computing with Stabilizer Testing, Phys. Rev. Lett. **115**, 220502 (2015).

[30] A. Gheorghiu, E. Kashefi, and P. Wallden, Robustness and device independence of verifiable blind quantum computing, New J. Phys. **17**, 083040 (2015).

[31] M. Hajdušek, C. A. Pérez-Delgado, and J. F. Fitzsimons, arXiv:1502.02563.

[32] J. F. Fitzsimons and E. Kashefi, Unconditionally verifiable blind quantum computation, Phys. Rev. A **96**, 012303 (2017).

[33] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev, arXiv:1704.04487.

[34] Y. Takeuchi, T. Morimae, and H. Masahito, Quantum computational universality of hypergraph states with pauli-X and Z basis measurements, Sci. Rep.-UK **9**, 13585 (2019).

[35] U. Mahadev, arXiv:1804.01082.

[36] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, and S. Montangero *et al.*, Simulating lattice gauge theories within quantum technologies, Eur. Phys. J. D **74**, 1 (2020).

[37] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, Rev. of Mod. Phys. **92**, 015003 (2020).

[38] P. Hauke, F. M. Cucchietti, L. Tagliacozzo, I. Deutsch, and M. Lewenstein, Can one trust quantum simulators?, Rep. Prog. Phys. **75**, 082401 (2012).

[39] T. Schweigler, V. Kasper, S. Erne, I. Mazets, B. Rauer, F. Cataldini, T. Langen, T. Gasenzer, J. Berges, and J. Schmiedmayer, Experimental characterization of a quantum many-body system via higher-order correlations, Nature **545**, 323 (2017).

[40] T. J. Evans, R. Harper, and S. T. Flammia, Scalable Bayesian Hamiltonian learning, arXiv:1912.07636 (2019).

[41] C. Kokail, L. Pastori, and *et al.*, (to be published).

[42] Since the ground-state wave function of the Fermi-Hubbard model and the Hamiltonian H are real valued, the constraints must be chosen as imaginary operators to obtain nonzero matrix elements $K_{nm} = \langle -i[A_n, S_m]\rangle$.

[43] C. Schweizer, M. Lohse, R. Citro, and I. Bloch, Spin Pumping and Measurement of Spin Currents in Optical Superlattices, Phys. Rev. Lett. **117**, 170405 (2016).

[44] S. Keßler and F. Marquardt, Single-site-resolved measurement of the current statistics in optical lattices, Phys. Rev. A **89**, 061601 (2014).

[45] E. Bairey, C. Guo, D. Poletti, N. H. Lindner, and I. Arad, Learning the dynamics of open quantum systems from their steady states, New J. Phys. **22**, 032001 (2020).

[46] P. M. Poggi, N. K. Lysne, K. W. Kuper, I. H. Deutsch, and P. S. Jessen, Quantifying the sensitivity to errors in analog quantum simulation, PRX Quantum **1**, 020308 (2020).

[47] S. T. Flammia and Y.-K. Liu, Direct Fidelity Estimation from few Pauli Measurements, Phys. Rev. Lett. **106**, 230501 (2011).

[48] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Practical Characterization of Quantum Devices Without Tomography, Phys. Rev. Lett. **107**, 210404 (2011).

[49] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).

[50] R. Jozsa, Fidelity for mixed quantum states, J. Mod. Opt. **41**, 2315 (1994).

[51] Y.-C. Liang, Y.-H. Yeh, P. E. M. F. Mendonça, R. Y. Teh, M. D. Reid, and P. D. Drummond, Quantum fidelity measures for mixed states, Rep. Prog. Phys. **82**, 076001 (2019).

[52] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing rényi entanglement entropy via randomized measurements, Science **364**, 260 (2019).

[53] A. K. Ekert, C. M. Alves, D. K. L. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek, Direct Estimations of Linear and Nonlinear Functionals of a Quantum State, Phys. Rev. Lett. **88**, 217901 (2002).

[54] D. A. Abanin and E. Demler, Measuring Entanglement Entropy of a Generic Many-Body System with a Quantum Switch, Phys. Rev. Lett. **109**, 020504 (2012).

[55] A. J. Daley, H. Pichler, J. Schachenmayer, and P. Zoller, Measuring Entanglement Growth in Quench Dynamics of Bosons in an Optical Lattice, Phys. Rev. Lett. **109**, 020505 (2012).

[56] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, Measuring entanglement entropy in a quantum many-body system, Nature **528**, 77 (2015).

[57] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, Quantum thermalization through entanglement in an isolated many-body system, Science **353**, 794 (2016).

[58] N. M. Linke, S. Johri, C. Figgatt, K. A. Landsman, A. Y. Matsuura, and C. Monroe, Measuring the rényi entropy of a two-site fermi-hubbard model on a trapped ion quantum computer, Phys. Rev. A **98**, 052334 (2018).

[59] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum State Tomography via Compressed Sensing, Phys. Rev. Lett. **105**, 150401 (2010).

[60] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, Nat. Commun. **1**, 149 (2010).

[61] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, and A. J. Daley *et al.*, Efficient tomography of a quantum many-body system, Nat. Phys. **13**, 1158 (2017).

[62] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. **14**, 447 (2018).

[63] S. J. van Enk and C. W. J. Beenakker, Measuring Trρn on Single Copies of ρ Using Random Measurements, Phys. Rev. Lett. **108**, 110503 (2012).

[64] A. Elben, B. Vermersch, M. Dalmonte, J. Cirac, and P. Zoller, Rényi Entropies from Random Quenches in Atomic Hubbard and Spin Models, Phys. Rev. Lett. **120**, 050406 (2018).

[65] A. Elben, B. Vermersch, C. F. Roos, and P. Zoller, Statistical correlations between locally randomized measurements: A toolbox for probing entanglement in many-body quantum states, Phys. Rev. A **99**, 052323 (2019).

[66] A. Ketterer, S. Imai, N. Wyderka, and O. Gühne, arXiv:2012.12176.

[67] L. Knips, A moment for random measurements, Quantum Views **4**, 47 (2020).

[68] M. Ohliger, V. Nesme, and J. Eisert, Efficient and feasible state tomography of quantum many-body systems, New J. Phys. **15**, 015024 (2013).

[69] S. Aaronson, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2018), p. 325.

[70] D. Gottesman, Ph.D. thesis, Caltech 1997.

[71] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, J. Opt. B **7**, 347 (2005).

[72] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, Phys. Rev. A **80**, 012304 (2009).

[73] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).

[74] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Symmetrized characterization of noisy quantum processes, Science **317**, 1893 (2007).

[75] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, and T. A. Ohki *et al.*, Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking, Phys. Rev. Lett. **109**, 080505 (2012).

[76] D. Gross, K. Audenaert, and J. Eisert, Evenly distributed unitaries: On the structure of unitary designs, J. Math. Phys. **48**, 052104 (2007).

[77] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, arXiv:2011.09636.

[78] D. E. Koh and S. Grewal, arXiv:2011.11580.

[79] D. Aharonov and U. Vazirani, arXiv:1206.3686.

[80] For an excellent, recent review of these protocols see Ref. [5].

[81] B. W. Reichardt, F. Unger, and U. Vazirani, arXiv:1209.0448.

[82] These functions are constructed in Refs. [35,83] on the basis of the hardness of the learning with errors problem [105].

[83] U. Mahadev, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, NY, USA, 2018), p. 332, arXiv:1708.02130.

[84] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, NY, USA, 2018), p. 320.

[85] T. Vidick and T. Zhang, Classical zero-knowledge arguments for quantum computations, Quantum **4**, 266 (2020).

[86] G. Alagic, A. M. Childs, A. B. Grilo, and S.-H. Hung, arXiv:1911.08101.

[87] T. Vidick, From operator algebras to complexity theory and black, Notices American Mathematical **66**, 1619 (2019).

[88] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local hamiltonian problem, SIAM J. Comput. **35**(5), 1070 (2005).

[89] The case in which B claims that the answer to the decision problem is "no" can be treated similarly.

[90] R. P. Feynman, Quantum mechanical computers, Found. Phys. **16**, 507 (1986).

[91] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Quantum Computation* (American Mathematical Society, Providence, 2002), Vol. 47.

[92] J. Biamonte and P. Love, Realizable hamiltonians for universal adiabatic quantum computers, Phys. Rev. A **78**, 012352(7) (2008).

[93] T. Morimae, D. Nagaj, and N. Schuch, Quantum proofs can be verified using only single-qubit measurements, Phys. Rev. A **93**, 022326(6) (2016).

[94] J. F. Fitzsimons, M. Hajdušek, and T. Morimae, Post hoc Verification of Quantum Computation, Phys. Rev. Lett. **120**, 040501(5) (2018).

[95] R. Lindner and C. Peikert, Better key sizes (and attacks) for lwe-based encryption, Cryptology ePrint Archive, Report 2010/613 (2010). https://eprint.iacr.org/2010/613.

[96] The two-to-one functions are such that $y_k(0,\cdot)$ and $y_k(1,\cdot)$ are injective.

[97] J. Carrasco *et al.*, to be published.

[98] This would correspond to a decision problem encoded in the output of a quantum circuit composed of three single-qubit and two-qubit gates acting on two qubits initialized in the state $|0\rangle^{\otimes 2}$.

[99] A. Gheorghiu, M. Hoban, and E. Kashefi, A simple protocol for fault tolerant verification of quantum computations, Quantum Sci. Technol. **4**, 015009 (2019).

[100] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[101] H.-S. Zhong *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[102] S. Aaronson and A. Arkhipov, arXiv:1011.3245.

[103] S. Boixo *et al.*, Characterizing quantum supremacy in near-term devices, Nat. Phys. **14**, 595 (2018).

[104] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, Nat. Phys. **15**, 159 (2019).

[105] O. Regev, in *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05 (Association for Computing Machinery, New York, NY, USA, 2005), p. 84.