Quantum Overlapping Tomography

Jordan Cotler[1,*] and Frank Wilczek[2,3,4,5,6,†]

[1]Stanford Institute for Theoretical Physics, Stanford University, Stanford, California 94305, USA
[2]Center for Theoretical Physics, MIT, Cambridge, Massachusetts 02139, USA
[3]T. D. Lee Institute, Shanghai, China
[4]Wilczek Quantum Center, Department of Physics and Astronomy, Shanghai Jiao Tong University, Shanghai 200240, China
[5]Department of Physics, Stockholm University, Stockholm, Sweden
[6]Department of Physics and Origins Project, Arizona State University, Tempe, Arizona 25287, USA

(Received 22 August 2019; accepted 10 February 2020; published 10 March 2020)

It is now experimentally possible to entangle thousands of qubits, and efficiently measure each qubit in parallel in a distinct basis. To fully characterize an unknown entangled state of n qubits, one requires an exponential number of measurements in n, which is experimentally unfeasible even for modest system sizes. By leveraging (i) that single-qubit measurements can be made in parallel, and (ii) the theory of perfect hash families, we show that *all* k-qubit reduced density matrices of an n qubit state can be determined with at most $e^{\mathcal{O}(k)}\log^2(n)$ rounds of parallel measurements. We provide concrete measurement protocols which realize this bound. As an example, we argue that with near-term experiments, every two-point correlator in a system of 1024 qubits could be measured and completely characterized in a few days. This corresponds to determining nearly 4.5 million correlators.

DOI: 10.1103/PhysRevLett.124.100401

Introduction.—Recently there have been remarkable advances in the construction and control of intermediate-scale quantum systems containing several hundred or even thousands of entangled qubits [1–5]. The qubits come from a variety of systems, including interacting electronic spins, quantized fluxes, and spatial modes of photons. But what about measuring the state of such systems, and documenting their correlations and entanglement?

To characterize an unknown n-qubit state completely using quantum tomography requires a number of parallel measurements which grows exponentially with n [6,7]. That exponential growth renders quantum tomography for many-body systems completely impractical even for modest system sizes. Indeed, full quantum tomography has not been performed for more than 10 qubits [8]. Some limited classes of quantum states featuring *a priori* constrained patterns of entanglement allow tomography with parametrically fewer measurements (for instance, see Refs. [9,10]), but most experimental systems do not produce states of those kinds. There are ingenious protocols which can characterize expectation values of an unknown quantum state more efficiently [11], but they require entangled nondemolition measurements and are not experimentally realistic for appreciably sized systems. Thus, there is a significant gap between our ability to produce massively correlated states in controlled settings, and our ability to characterize those correlations quantitatively.

What can be done is to address the individual qubits of a system in parallel, and to measure each in a chosen basis of $\mathbb{C}^2$. Suppose we want to measure all k-qubit reduced density matrices of an n-qubit system. Access to these density matrices would enable us to completely characterize all k-qubit correlations present in the n-qubit system. There are $\binom{n}{k}$ such k-qubit reduced density matrices, and if k is small relative to n then $\binom{n}{k} \sim n^k$. Performing a k-qubit tomography requires $e^{\mathcal{O}(k)}$ measurements, and so naïvely we require $e^{\mathcal{O}(k)}\binom{n}{k} \sim e^{\mathcal{O}(k)}n^k$ measurements to obtain all k-qubit reduced density matrices. Even for $k = 2$, it would not be practical to make so many measurements once n exceeds a hundred qubits.

This count, however, ignores the power of parallelism. If we measure nonoverlapping k-qubit subsystems in parallel we can get by with fewer measurements, but that only reduces the total number of required measurements by a multiplicative factor of n/k. At first sight, it appears problematic that the set of all k-qubit subsystems is highly overlapping. In fact, it is a tremendous advantage. Measuring a particular k-qubit subsystem provides us information about all other k-qubit subsystems which overlap with it. Here we present a method to organize that information. We call it "quantum overlapping tomography" (QOT). Using QOT, we can measure *all* k-qubit reduced density matrices with at most $e^{\mathcal{O}(k)}\log^2(n)$ measurements. Our QOT protocols only require measuring each qubit in a distinct basis (i.e., a product measurement) in parallel, with judiciously chosen measurement settings. The measurements can be efficiently postprocessed to reconstruct all k-qubit reduced density matrices. QOT easily adapts to qudits (i.e., d-level systems) in place of qubits. (n, k)

families of perfect hash functions [12–14] will be a crucial tool in our measurement procedure. The theory of perfect hash families has been well studied in theoretical computer science for over thirty years, and is used in database management [12–25].

We will begin by reviewing quantum tomography, and then provide a probabilistic argument for the scaling of our measurement procedure. We then explain the measurement procedure in explicit mathematical detail for $k = 2$, and more briefly for $k > 2$. Then we describe its possible realization to measure all 2-qubit correlators in a system of ultracold atoms, and conclude with a summary and forward-looking discussion.

Review of quantum tomography.—Here we review the basic essentials of quantum tomography (see, for instance, Ref. [26]). We focus on the standard experimental protocol which only requires product measurements, i.e., measuring each qubit independently, since more sophisticated schemes involving entangled measurements are not presently experimentally feasible. It will be useful to be very concrete about the measurement procedure. To begin, we will explicitly explain how to do quantum tomography for a 2-qubit density matrix.

Suppose we have a 2-qubit density matrix ρ, and that we want to perform a quantum tomography of it. To do so, we must be able to produce many copies of ρ, via some state preparation procedure, quantum source, etc. Let σ_i^α denote a Pauli operator on the ith site, where here $i = 1, 2$ in the 2-qubit case. We have $\alpha = 0, 1, 2, 3$ where $\sigma^0 = \mathbf{1}$, $\sigma^1 = \sigma^x$, $\sigma^2 = \sigma^y$, and $\sigma^3 = \sigma^z$, as is standard. We can write ρ as

$$\rho = \frac{1}{4}\sum_{\alpha,\beta=0}^{3} \mathrm{tr}\{(\sigma_1^\alpha \otimes \sigma_2^\beta)\rho\}\sigma_1^\alpha \otimes \sigma_2^\beta, \tag{1}$$

and so to perform a quantum tomography we need to measure the expectation values $\mathrm{tr}\{(\sigma_1^\alpha \otimes \sigma_2^\beta)\rho\}$, of which there are $4 \times 4 = 16$. We do not need to measure the $\alpha = \beta = 0$ expectation value, since it is guaranteed to be $\mathrm{tr}\{(\mathbf{1} \otimes \mathbf{1})\rho\} = 1$ since ρ has unit trace. Thus, we only need to measure 15 expectation values.

First we consider the 1-site expectation values, for which either $\alpha = 0$, or $\beta = 0$. For example, suppose we want to measure $\mathrm{tr}\{(\sigma_1^0 \otimes \sigma_2^2)\rho\}$ where we recall that $\sigma_1^0 = \mathbf{1}$. Then we only need to measure the second qubit in the y basis, and find that

$$\mathrm{tr}\{(\sigma_1^0 \otimes \sigma_2^2)\rho\} \approx \frac{1}{M}[\mathcal{N}_2(\uparrow_y) - \mathcal{N}_2(\downarrow_y)], \tag{2}$$

where $\mathcal{N}_2(\uparrow_y)$ is the number of times we measure the second qubit to be up in the y basis, and $\mathcal{N}_2(\downarrow_y)$ is defined similarly. The other 1-site expectation values can be obtained in similar fashion.

Now we turn to 2-site expectation values for which neither α nor β equal zero. As an example, considering the expectation value $\mathrm{tr}\{(\sigma_1^1 \otimes \sigma_2^2)\rho\}$, we need to measure the first qubit in the x basis, and concurrently the second qubit in the y basis. Let $\mathcal{N}_{12}(\uparrow_x, \uparrow_y)$ be the number of times we measure both the first qubit to be up in the x basis, *and* the second qubit to be up in the y basis. The quantities $\mathcal{N}_{12}(\uparrow_x, \downarrow_y)$, $\mathcal{N}_{12}(\downarrow_x, \uparrow_y)$, and $\mathcal{N}_{12}(\downarrow_x, \downarrow_y)$ are defined similarly. If we make a total number of measurements M, then we can approximate

$$\begin{aligned}\mathrm{tr}\{(\sigma_1^1 \otimes \sigma_2^2)\rho\} \approx \frac{1}{M}[&\mathcal{N}_{12}(\uparrow_x, \uparrow_y) - \mathcal{N}_{12}(\uparrow_x, \downarrow_y) \\ &- \mathcal{N}_{12}(\downarrow_x, \uparrow_y) + \mathcal{N}_{12}(\downarrow_x, \downarrow_y)]\end{aligned} \tag{3}$$

which becomes exact in the limit of a large number of measurements M. All other 2-site expectation values, for which neither α nor β equal zero, can be obtained in an analogous manner.

Suppose we require M measurements of each expectation value to obtain ample statistics. If we want to do a tomography of ρ, which requires measuring 15 expectation values with M measurements each, then naïvely we require $15M$ measurements to determine ρ. However, note that when we measure 2-site expectation values for which α and β are both nonzero, we can use this data to extract 1-site expectation values. For instance, upon collecting data to construct $\mathrm{tr}\{(\sigma_1^1 \otimes \sigma_2^2)\rho\}$, we can use that same data to construct both $\mathrm{tr}\{(\sigma_1^1 \otimes \mathbf{1})\rho\}$ and $\mathrm{tr}\{(\mathbf{1} \otimes \sigma_2^2)\rho\}$. Thus, instead of measuring all 15 expectation values to determine ρ, we effectively only need to measure 9 expectation values (i.e., the 2-site expectation values where neither α nor β is zero), since we can reuse their measurements to reconstruct the other 6 expectation values. In summary, we only require $9M$ measurements to fully determine ρ.

Now, suppose we have a k-qubit density matrix ρ'. Writing ρ' as

$$\rho' = \frac{1}{2^k}\sum_{i_1,\ldots,i_k=0}^{3} \mathrm{tr}\{(\sigma_1^{i_1} \otimes \cdots \otimes \sigma_k^{i_k})\rho'\}\sigma_1^{i_1} \otimes \cdots \otimes \sigma_k^{i_k}, \tag{4}$$

we evidently need to determine $4^k - 1$ expectation values, where we have subtracted 1 since we already know $\mathrm{tr}\{(\mathbf{1} \otimes \cdots \otimes \mathbf{1})\rho'\} = 1$. Since we obtain each expectation value by multiplying the outputs of k 2-outcome measurements, we need the probability that each measurement is faulty to be sufficiently small. In particular, if the probability of a faulty measurement is Δ, then we want $\Delta \sim 1/k$ so that $k\Delta \sim \mathcal{O}(1)$.

Using a similar procedure as in the 2-qubit case, we only need to perform $M3^k$ total measurements, comprised of all combinations of x-basis, y-basis, and z-basis measurement settings for the k sites, each repeated M times to gain ample statistics. If we want our approximations to all terms

$\mathrm{tr}\{(\sigma_1^{i_1} \otimes \cdots \otimes \sigma_k^{i_k})\rho'\}$ to be within ε of their true values with constant probability close to 1, then by the Chernoff-Hoeffding inequality and a union bound, we require M to be at most $\sim 4\log(2)k/\varepsilon^2$. We will review the Chernoff-Hoeffding inequality in Appendix A of Supplemental Material [27]. If we want our reconstructed density matrix to be ε-close to the true density matrix in the 1-norm, then would can require M to be at most $\sim 4\log(2)k8^k/\varepsilon^2$ due to the norm inequalities $\|A\|_1 \leq d^{1/2}\|A\|_2 \leq d^{3/2}\max_{i,j}|A_{i,j}|$ for $d \times d$ matrices. (If we want to guarantee that our approximation to the true density matrix is itself positive definite, Hermitian, and has unit trace, then we can use a protocol such as Ref. [30], although this will not change our bounds very much.) Using either definition of closeness, we require a total of $e^{\mathcal{O}(k)}$ measurements to perform a quantum tomography on k qubits.

Probabilistic argument.—In the last section, we saw that to perform a quantum tomography on k qubits, we needed to perform measurements for all combinations of the measurement settings (either the x basis, y basis, or z basis for each qubit), i.e., varying the measurement basis of each qubit independently. Since there are three bases for each qubit and k total qubits, we required 3^k measurements, times a multiplicative factor of M to build up enough statistics.

Now we turn to constructing *all* $\binom{n}{k}$ of the k-qubit reduced density matrices of the n-qubit system. To formalize the problem, suppose we have a family of N functions $f_1, \ldots, f_N$, each taking $[n] \to [k]$, where $[n] \coloneqq \{1, \ldots, n\}$ and $[k]$ is defined similarly. These functions form an (n, k) family of perfect hash functions if for any subset S of $[n]$ where $|S| = k$ (i.e., S contains k elements), there is some f_i in the family which is injective on S [12–14]. For us, this means that for any given subsystem of k qubits, there is at least one function f_i in the family which assigns each qubit in that subsystem to a distinct number 1 through k.

Given such a family of functions $f_1, \ldots, f_N$, the approach of QOT is to use each f_i to partition the qubits into k disjoint subsets [i.e., $f_i^{-1}(j)$ for $j = 1, \ldots, k$], and for each of the N partitions perform all parallel measurements over single-qubit measurement bases (x, y, or z) such that all qubits in the same subset of the partition have identical measurement settings. (Note that there are 3^k such parallel measurements for each partition.) These measurements are all repeated M times each. This entails making a total of $NM3^k$ total measurements, and allows us to determine *all* k-qubit reduced density matrices. This procedure will be explained in more detail in the following sections. Then a crucial question is, what is the smallest N for which we can construct an (n, k) family of perfect hash functions?

To construct a bound on N, we consider a probabilistic argument in which each f_i is chosen randomly, i.e., f_i assigns each qubit to a number 1 through k uniformly at random. Suppose we want the probability that $f_1, \ldots, f_N$ does *not* form an (n, k) family of perfect hash functions to be less than a small parameter δ. Then we need N to be at most (see Appendix B of Supplemental Material [27])

$$N < e^{\mathcal{O}(k)}\left(\frac{1}{k}\log(1/\delta) + \log(n)\right). \tag{5}$$

This implies that we require $Me^{\mathcal{O}(k)}\log(n)$ measurements to determine all k-qubit reduced density matrices of an n-qubit system using QOT. Using the Chernoff-Hoeffding inequality and a union bound (see Appendix A of Supplemental Material [27]), if we want to determine all terms $\mathrm{tr}\{(\sigma_1^{i_1} \otimes \cdots \otimes \sigma_k^{i_k})\rho'\}$ within ε of their true values with constant probability close to 1, then we require $M \sim k\log(n)/\varepsilon^2$. Therefore, the total number of measurements is $e^{\mathcal{O}(k)}\log^2(n)$. As mentioned above, if we wanted our reconstructed density matrix to be ε-close to the true density matrix in the 1-norm we would gain a multiplicative factor of 8^k which would still give us a total of $e^{\mathcal{O}(k)}\log^2(n)$ measurements.

There is a substantial literature which constructs explicit and efficiently computable (n, k) families of perfect hash functions which satisfy the bound in Eq. (5), such as Refs. [12–25]. In the next section, we explain the simplest example, namely an explicit $(n, 2)$ family of perfect hash functions of size $\lceil \log_2(n) \rceil$, which is well known.

QOT for $k = 2$.—In this section, we provide a QOT procedure for measuring all 2-qubit reduced density matrices with only $(3M + 6M\lceil \log_2(n) \rceil)$ measurements, for $M \sim 256\log(n)/\varepsilon^2$ to achieve ε-closeness to the true density matrix in the 1-norm. We consider a simple but very useful example of an $(n, 2)$ family of perfect hash functions, comprised of $q = \lceil \log_2(n) \rceil$ functions $f_1, \ldots, f_q$ each taking $[n] \to \{0, 1\}$. (In our previous notation, we would have said that the functions take $[n] \to [2] = \{1, 2\}$, but here we instead use $\{0, 1\}$ as the codomain for convenience.) The function f_i is defined by

$$f_i(j) = i\text{th digit in the binary expansion of} (j - 1). \tag{6}$$

Here we are implicitly representing $(j - 1)$ by a q-bit string, and by the ith digit we mean the ith most significant digit. For instance, if we consider a (16,2) family so that $q = 4$, then $f_1(5) = 0$, $f_2(5) = 1$, $f_3(5) = 0$ and $f_4(5) = 0$. This follows from the fact that $4 = 5 - 1$ can be expressed as the q-bit string 0100. The functions f_1, f_2, f_3, f_4 are depicted in Fig. 1.

Suppose that we have n qubits, and that we want to perform quantum tomography on every 2-qubit reduced density matrix using QOT. We consider an $(n, 2)$ family of perfect hash functions given by Eq. (6) with $q = \lceil \log_2(n) \rceil$. The procedure is as follows:

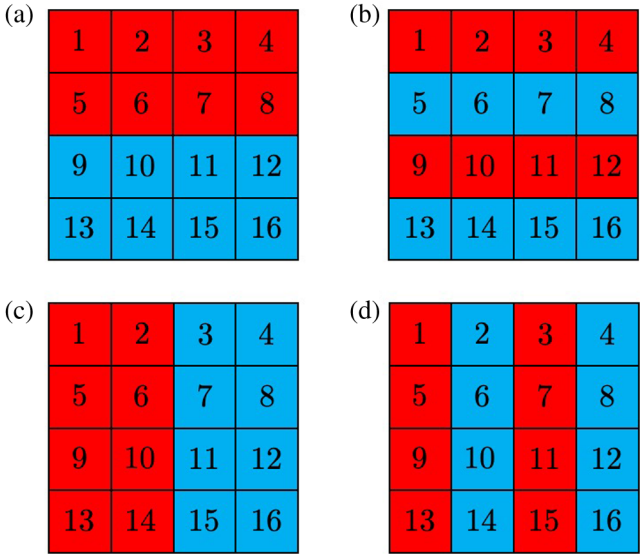

FIG. 1. A visual depiction of the (16,2) family of perfect hash functions given by f_1, f_2, f_3, f_4 from Eq. (6). The four functions in the family are displayed in order in (a)–(d), where red corresponds to 0 and blue corresponds to 1. Note that for any pair (i, j) for $1 \leq i, j \leq 16$ and $i \neq j$, there is at least one function for which i and j are assigned distinct colors.

Step 1: Measure all qubits in the x basis, y basis, and z basis, each M times. This corresponds to $3M$ measurements of all qubits in parallel.

Step 2: This step will be divided into q substeps, $2.1, \ldots, 2.q$. For each $j = 1, \ldots, q$, step $2.j$ is as follows. Consider the function f_j. If a qubit is assigned to 0 by f_j, then we call the qubit "red." Similarly, if a qubit is assigned to 1 by f_j, then we call the qubit "blue." Then we perform the following nine measurements, M times each:

(1) measure each red qubit in the **x** *basis*, and each blue qubit in the **y** *basis*;

(2) measure each red qubit in the **y** *basis*, and each blue qubit in the **x** *basis*;

(3) measure each red qubit in the **x** *basis*, and each blue qubit in the **z** *basis*;

(4) measure each red qubit in the **z** *basis*, and each blue qubit in the **x** *basis*;

(5) measure each red qubit in the **y** *basis*, and each blue qubit in the **z** *basis*;

(6) measure each red qubit in the **z** *basis*, and each blue qubit in the **y** *basis*.

Because of parallelization, each step $2.j$ corresponds to $6M$ measurements, and thus $6Mq = 6M\lceil \log_2(n) \rceil$ measurements total for all of step 2.

Step 3: Steps 1 and 2 collect all of the data we need, and only require a total of $3M + 6M\lceil \log_2(n) \rceil$ measurements. Suppose we want to reconstruct the reduced density matrix ρ_{rs} of the rth qubit and the sth qubit, for $1 \leq r, s \leq n$ and of course $r \neq s$. Note that the q-bit binary representation of $(r-1)$ and $(s-1)$ must differ on at least one bit, since r and s are distinct numbers. Suppose that $(r-1)$ and $(s-1)$ differ on their tth bits. Then

(7) to obtain $\mathrm{tr}\{(\mathbf{1}_r \otimes \sigma_s^x)\rho_{rs}\}$, $\mathrm{tr}\{(\mathbf{1}_r \otimes \sigma_s^y)\rho_{rs}\}$, $\mathrm{tr}\{(\mathbf{1}_r \otimes \sigma_s^z)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^x \otimes \mathbf{1}_s)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^y \otimes \mathbf{1}_s)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^z \otimes \mathbf{1}_s)\rho_{rs}\}$, we use the data collected from steps 1 and 2;

(8) to obtain $\mathrm{tr}\{(\sigma_r^x \otimes \sigma_s^x)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^y \otimes \sigma_s^y)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^z \otimes \sigma_s^z)\rho_{rs}\}$, we use the data collected from step 1;

(9) to obtain $\mathrm{tr}\{(\sigma_r^x \otimes \sigma_s^y)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^y \otimes \sigma_s^x)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^x \otimes \sigma_s^z)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^z \otimes \sigma_s^x)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^y \otimes \sigma_s^z)\rho_{rs}\}$, $\mathrm{tr}\{(\sigma_r^z \otimes \sigma_s^y)\rho_{rs}\}$, we use the data collected from step $2.t$.

Then we can reconstruct ρ_{rs} using

$$\rho_{rs} = \frac{1}{4} \sum_{\alpha,\beta=0}^{3} \mathrm{tr}\{(\sigma_r^\alpha \otimes \sigma_s^\beta)\rho_{rs}\}\sigma_r^\alpha \otimes \sigma_s^\beta. \tag{7}$$

Once we have all of the 2-qubit reduced density matrices at hand, we can analyze their bipartite entanglement. For instance, there are explicit formulas for computing the entanglement of formation [31] and related quantities [32,33]. One can then study, for example, how entanglement varies as the qubits comprising the 2-qubit subsystem are chosen to be further apart in space.

QOT for arbitrary k.—To perform QOT to determine all k-qubit reduced density matrices of an n-qubit system, one proceeds in the same way as in the previous section, but instead utilizing an (n, k) family of perfect hash functions. In the language of the previous section, each function f_i in the family assigns each qubit to one of k "colors," i.e., red, blue, green, etc. The procedure generalizes in the obvious way. Then the total number of required measurements scales as $Me^{\mathcal{O}(k)} \log(n) \sim e^{\mathcal{O}(k)} \log^2(n)$, which has an n dependence significantly better than even shadow tomography applied to measuring subsystems [11]. Such a shadow tomography would require $\mathcal{O}(n\mathrm{polylog}(n))$ measurements.

For $k > 2$, constructing (n, k) families of perfect hash functions which contain as few functions as possible can be a difficult task. Luckily, there is an extensive literature on constructing such families, and we refer the reader to Refs. [12–25]. Also, there is a web page providing a list of the smallest known (n, k) families for various values of n and k [34].

Experimental prospects.—Here we estimate the practical potential of QOT based on near-term technology. Consider an ultracold atom system with spin-1/2 degrees of freedom, which we can prepare in the ground state of a local Hamiltonian and then probe with a quantum gas microscope (for a review, see Ref. [35]). In arrays of neutral atoms configured using optical tweezers [5,36] or in arrays of optically trapped ions [37], each measurement round takes at most a few hundred milliseconds. A single cycle of the experimental protocol can be significantly faster, even for systems sizes up to hundreds or even thousands of

atoms. We will consider on the order of a thousand atoms, which should be possible in the near future. We also suppose that there is apparatus to individually rotate each of the atoms. Concretely, let us choose $n = 1024$ qubits, subsystems of size $k = 2$, and $M = 15500$, so that around 95% of the time *all* measured two-point expectation values are within 0.05 of their true values. (See Appendix A of Supplemental Material [27] for more details on the estimation of M.) Then the $k = 2$ QOT protocol in the section "*QOT for* $k = 2$" requires 1 000 000 measurement rounds. Assuming 250 ms per measurement round, $k = 2$ QOT could be performed in a block of three days. By contrast, the naïve strategy of measuring nonoverlapping 2-qubit subsystems in parallel with similar error probabilities also requires around $M = 15\,500$ and thus $9M\binom{n}{2}/(n/2) \approx 10^9$ total measurement rounds, which would take nearly 60 weeks nonstop. Thus, QOT would enable the measurement of all 4715008 two-point correlators of a many-body quantum state. (It would be especially interesting to use 2D or 3D arrays, since in the higher-dimensional setting it is easier to trap a large number of atoms, and also more interesting to characterize low-energy eigenstates of experimentally realizable Hamiltonians.)

Summary and discussion.—QOT provides efficient protocols to measure many-body correlations in systems with large numbers of degrees of freedom. We anticipate that QOT will be a useful tool for experimental characterization of many-body quantum states.

Several adaptations of QOT may be interesting to consider. Systems with symmetry obey constraints and selection rules which might be exploited to streamline the protocol. One might also try to focus on local correlations, in systems where long-range correlations are not significant. This poses interesting mathematical problems. For example, given an n-qubit state on a lattice, how do we efficiently measure all k-qubit reduced density matrices, for k-qubit subsystems where every pair of qubits is at most a distance d apart? Taking geometric constraints into account would require a generalization of (n, k) families of perfect hash functions, entailing the additional data of (i) a weighted graph G representing the geometry, and (ii) a distance d which serves as the maximum diameter of the k-qubit subsystems.

Since QOT allows us to efficiently measure *all* k-point functions of a system, it would be natural to use QOT to diagnose long-range order and critical behavior. A modification of the QOT protocols may be useful to focus on special types of nonlocal order parameters (for instance, stringlike products) which appear in the classification of topological order (see, e.g., Refs. [38–40]).

QOT can be applied to measuring expectation values of k-local Hamiltonians, such as those which appear in quantum and classical versions of k-SAT [41,42] and in recent work on quantum machine learning [43–49]. Also, QOT can supply needed input for the quantum marginal problem (see Refs. [50,51] for recent overviews, and Ref. [52] for applications in tomography), in which one tries to determine a quantum state as well as possible given its reduced density matrices up to a given size.

A special thank you to Jian-Wei Pan and Yu-Ao Chen for suggesting the problem and discussing experimental implementations. J. C. is grateful to Reuben Saunders for earlier discussions on multiplexed RNA perturb-sequencing. We are happy to thank Noga Alon, Ryan Alweiss, and Xiaoyu He for suggesting valuable references on perfect hash functions, Soonwon Choi for guidance on ultracold atom references, and Patrick Hayden, Steve Flammia, and Daniel Ranard for discussions and feedback on the manuscript. J. C. is supported by the Fannie and John Hertz Foundation and the Stanford Graduate Fellowship program. F. W.'s work is supported by the U.S. Department of Energy under grant Contract No. DE-SC0012567, by the European Research Council under Grant No. 742104, and by the Swedish Research Council under Contract No. 335-2014-7424.

*jcotler@stanford.edu
†wilczek@mit.edu

[1] T. D. Ladd , F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, Quantum computers, Nature (London) **464**, 45 (2010).
[2] C. Monroe and J. Kim, Scaling the ion trap quantum processor, Science **339**, 1164 (2013).
[3] M. H. Devoret and R. J. Schoelkopf, Superconducting circuits for quantum information: An outlook, Science **339**, 1169 (2013).
[4] D. D. Awschalom, L. C. Bassett, A. S. Dzurak, E. L. Hu, and J. R. Petta *et al.*, Quantum spintronics: Engineering and manipulating atom-like spins in semiconductors, Science **339**, 1174 (2013).
[5] H. Bernien *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, Nature (London) **551**, 579 (2017).
[6] R. O'Donnell and J. Wright, Efficient quantum tomography, in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2016).
[7] J. Haah *et al.*, Sample-optimal tomography of quantum states, IEEE Trans. Inf. Theory **63**, 5628 (2017).
[8] C. Song *et al.*, 10-Qubit Entanglement and Parallel Logic Operations with a Superconducting Circuit, Phys. Rev. Lett. **119**, 180511 (2017).
[9] 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).
[10] B. P. Lanyon *et al.*, Efficient tomography of a quantum many-body system, Nat. Phys. **13**, 1158 (2017).
[11] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2018).
[12] K. Mehlhorn, *Data Structures and Algorithms. 1. Sorting and Searching* (Springer-Verlag, Berlin, 1984).

[13] M. L. Fredman, J. Komlós, and E. Szemerédi, Storing a sparse table with $O(1)$ access time, J. Assoc. Comput. Mach. **31**, 538 (1984).
[14] M. L. Fredman and J. Komlós, On the size of separating systems and families of perfect hash functions, SIAM J. Algebraic Discrete Methods **5**, 61 (1984).
[15] N. Alon, Explicit construction of exponential sized families of k-independent sets, Discrete Math. **58**, 191 (1986).
[16] J. Korner and K. Marton, New bounds for perfect hashing via information theory, Eur. J. Combinatorics **9**, 523 (1988).
[17] J. P. Schmidt and A. Siegel, The spatial complexity of oblivious k-probe hash functions, SIAM J. Comput. **19**, 775 (1990).
[18] N. Alon, R. Yuster, and U. Zwick, Color-coding, J. ACM **42**, 844 (1995).
[19] M. Naor, L. J. Schulman, and A. Srinivasan, Splitters and near-optimal derandomization, in *Proceedings of IEEE 36th Annual Foundations of Computer Science* (IEEE, Los Alamitos, 1995).
[20] M. Atici, S. S. Magliveras, D. R. Stinson, and W.-D. Wei, Some recursive constructions for perfect hash families, Journal of combinatorial designs **4**, 353 (1996).
[21] Z. J. Czech, G. Havas, and B. S. Majewski, Perfect hashing, Theor. Comput. Sci. **182**, 1 (1997).
[22] S. R. Blackburn and P. R. Wild, Optimal linear perfect hash families, J. Comb. Theory A **83**, 233 (1998).
[23] S. R. Blackburn, Perfect hash families: Probabilistic methods and explicit constructions, J. Comb. Theory A **92**, 54 (2000).
[24] D. R. Stinson, R. Wei, and L. Zhu, New constructions for perfect hash families and related structures using combinatorial designs and codes, Journal of combinatorial designs **8**, 189 (2000).
[25] N. Alon and S. Gutner, Balanced families of perfect hash functions and their applications, ACM Trans. Algorithms (TALG) **6**, 54 (2010).
[26] M. A. Nielsen and I. Chuang, Quantum Computation and Quantum Information (2002), Vol. **558**.
[27] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.124.100401 for details on our application of the Chernoff-Hoeffding inequality, as well as bounds involving random perfect hash families, which includes Refs. [28,29].
[28] A. Nilli, Perfect hashing and probability, Comb. Probab. Comput. **3**, 407 (1994).
[29] V. Guruswami and A. Riazanov, Beating Fredman-Komló s for perfect *k*-hashing, arXiv:1805.04151.
[30] M. Guta *et al.*, Fast state tomography with optimal error bounds, arXiv:1809.11162.
[31] W. K. Wootters, Entanglement of Formation of an Arbitrary State of Two Qubits, Phys. Rev. Lett. **80**, 2245 (1998).
[32] G. Vidal, Optimal local preparation of an arbitrary mixed state of two qubits: Closed expression for the single-copy case, Phys. Rev. A **62**, 062315 (2000).
[33] K. Osuga and D. N. Page, Harmony for 2-qubit entanglement, arXiv:1906.09273.
[34] R. Dougherty, http:/www.public.asu.edu/redoughe/phf pages/phf tables.html.
[35] S. Kuhr, Quantum-gas microscopes: A new tool for cold-atom quantum simulators, Natl. Sci. Rev. **3**, 170 (2016).
[36] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays, Science **354**, 1021 (2016).
[37] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe *et al.*, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, Nature (London) **551**, 601 (2017).
[38] M. A. Levin and X.-G. Wen, String-net condensation: A physical mechanism for topological phases, Phys. Rev. B **71**, 045110 (2005).
[39] M. A. Levin and X.-G. Wen, Detecting Topological Order in a Ground State Wave Function, Phys. Rev. Lett. **96**, 110405 (2006).
[40] X.-G. Wen, Topological order: From long-range entangled quantum matter to a unified origin of light and electrons, ISRN Condens. Matter Phys. **2013** 198710 (2013).
[41] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, SIAM J. Comput. **35**, 1070 (2006).
[42] A. Yu Kitaev *et al.*, *Classical and Quantum Computation* (American Mathematical Society, Providence, 2002), Vol. 47.
[43] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science **355**, 602 (2017).
[44] E. P. L. Van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, Nat. Phys. **13**, 435 (2017).
[45] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, Nat. Phys. **13**, 431 (2017).
[46] L. Wang, Discovering phase transitions with unsupervised learning, Phys. Rev. B **94**, 195105 (2016).
[47] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Quantum Entanglement in Deep Learning Architectures, Phys. Rev. Lett. **122**, 065301 (2019).
[48] Y. Zhang and E.-A. Kim, Quantum Loop Topography for Machine Learning, Phys. Rev. Lett. **118**, 216401 (2017).
[49] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. **14**, 447 (2018).
[50] M. Walter, Multipartite quantum states and their marginals, arXiv:1410.6820.
[51] C. Schilling, Quantum marginal problem and its physical relevance, arXiv:1507.00299.
[52] T. Xin *et al.*, Quantum State Tomography Via Reduced Density Matrices, Phys. Rev. Lett. **118**, 020401 (2017).