Scalable Reconstruction of Density Matrices

T. Baumgratz,[1,2] D. Gross,[3] M. Cramer,[1,2] and M. B. Plenio[1,2]

[1]*Institut für Theoretische Physik, Albert-Einstein-Allee 11, Universität Ulm, 89069 Ulm, Germany*
[2]*Center for Integrated Quantum Science and Technology, Universität Ulm, 89069 Ulm, Germany*
[3]*Physikalisches Institut, Hermann-Herder-Straße 3, Albert-Ludwigs Universität Freiburg, 79104 Freiburg, Germany*

(Received 6 July 2012; published 11 July 2013)

Recent contributions in the field of quantum state tomography have shown that, despite the exponential growth of Hilbert space with the number of subsystems, tomography of one-dimensional quantum systems may still be performed efficiently by tailored reconstruction schemes. Here, we discuss a scalable method to reconstruct mixed states that are well approximated by matrix product operators. The reconstruction scheme only requires local information about the state, giving rise to a reconstruction technique that is scalable in the system size. It is based on a constructive proof that generic matrix product operators are fully determined by their local reductions. We discuss applications of this scheme for simulated data and experimental data obtained in an ion trap experiment.

DOI: 10.1103/PhysRevLett.111.020401 PACS numbers: 03.65.Wj, 03.67.−a

The complexity of many-body systems is one of the most intriguing, but at the same time daunting, features of quantum mechanics. The curse of dimensionality, namely the exponential growth of the descriptive complexity of even pure states, is a property of quantum mechanics which clearly distinguishes it from classical physics. Therefore, in general, the number of variables required to uniquely determine a quantum state increases in accordance with the growth of the Hilbert space exponentially.

Quantum state tomography addresses the problem of completely characterizing a state of a physical system by measuring a complete set of observables that determine the state uniquely [1]. As the complexity of quantum operations implemented in laboratories steadily increases [2–5], the demand for a reliable and scalable tomography [6,7] of prepared states is high and of considerable importance for the future of quantum technologies. The ability to store and manipulate interacting quantum many-body systems, such as linearly arranged ions in an ion trap, was enhanced rapidly during the last years. Soon, if not already, the number of particles controllable in such systems will cross the threshold for which conventional methods of full quantum state tomography fail due to both the limited time that is realistically available for the experiment and the limitations to the resources that are available for the classical postprocessing of the experimental data [2]. Further, while most experiments have so far focused on the controlled creation of pure states and scalable reconstruction methods have been tailored to the pure setting [6,7], experimental simulations of open system dynamics have begun to emerge [8], calling for efficient tomography of mixed states.

The experimental time requirement is defined by the total number of measurements which have to be done to reconstruct the state faithfully; i.e., one has to consider the system size and the number of repetitions to obtain sufficient statistics [9]. The postprocessing resources are determined by the individual tomography scheme and in particular by the representation of the state. Clearly, full quantum state tomography where the state is represented by an exponentially large number of variables will require an exponentially increasing computational power and is hence infeasible already for, e.g., trapped ion experiments available today [4]. But many naturally occurring quantum states and many states of interest for quantum information tasks are completely characterized by a number of variables scaling moderately in the number of particles: ground states of gapped local Hamiltonians [10–12], thermal states of local Hamiltonians [12,13], the W state, the GHZ state, and cluster states are all matrix product operators (matrix product states if they are pure) of low dimension, or very well approximated by them. These states are parametrized by a linear number of matrices of low bond dimension. The key insight here is not that these states are matrix product operators or states (any state is a matrix product operator, respectively state) but that the matrix dimension is low, in particular independent of the system size. This solves one issue concerning the postprocessing side of the problem mentioned above as these states may be stored efficiently on a classical computer. On the other hand, as we will see, generic matrix product operators are not only completely determined by a linear number of local observables but may also be efficiently reconstructed from such local measurements, which makes the formalism we present here a powerful tool for quantum state tomography of mixed states [14].

Recently, it has been demonstrated that the reconstruction of pure quantum states for large systems can be possible with the knowledge of local information only [6]. The scheme presented in the latter reference relies on an efficient version of an iterative method first introduced in the context of matrix completion. While the original

method comes with a convergence proof, this guarantee is lost in the efficient version. Here, we take a different approach, which works provably under a mild technical assumption on the state to be reconstructed and is not restricted to pure states. We present extensive numerical results for states not meeting the assumption guaranteeing uniqueness of the reconstructed density matrix. In this Letter we consider N d-level subsystems aligned in a one-dimensional geometry, e.g., a chain of qubits ($d=2$). The aim is to reconstruct a mixed state $\hat{\varrho}$ from local information only. The local information we have in mind are estimates to all reductions of the state to a fixed number R of contiguous sites. These may be obtained by estimating the expectation values of an informationally complete set of observables on the R sites. Note that we do not require the estimates to the reductions to be states; i.e., empirical estimates to the expectation values of local observables suffice and postprocessing such as maximum likelihood estimation is not required. As R is fixed, this corresponds to an experimental effort that is linear in the system size N. We present a computationally cheap tomography scheme which scales polynomially in the system size N and succeeds provably under a certain technical invertibility condition on $\hat{\varrho}$. We demonstrate numerically and at the hand of experimental data (a W state on 8 qubits created in an ion trap experiment [2]) that it still works well when this condition is not necessarily met.

We begin our exposition of the tomography scheme by introducing some notation. We denote the to-be-reconstructed state by $\hat{\varrho}$. The input to the reconstruction scheme are estimates of expectation values which completely specify all reductions of $\hat{\varrho}$ to R contiguous sites. We denote these reductions by $\hat{\varrho}_k$, $k=1,\ldots,N-R+1$. Put mathematically, $\hat{\varrho}_k = \mathrm{tr}_{\{1,\ldots,k-1\}\cup\{k+R,\ldots,N\}}[\hat{\varrho}]$, i.e., the trace over all but the R sites $\{k,\ldots,k+R-1\}$. Now let $\{\hat{P}_i^{(\alpha)}\}_{\alpha=1,\ldots,d^2}$ be a complete operator basis for the site i. A common choice for spin-1/2 particles is given by $\hat{P}_i^{(1)} = \mathbb{1}_i/\sqrt{2}$, $\hat{P}_i^{(2)} = \hat{\sigma}_i^x/\sqrt{2}$, $\hat{P}_i^{(3)} = \hat{\sigma}_i^y/\sqrt{2}$, $\hat{P}_i^{(4)} = \hat{\sigma}_i^z/\sqrt{2}$, i.e., the orthonormal Pauli spin basis. We may then write

$$\hat{\varrho}_k = \sum_{\alpha_1,\ldots,\alpha_R} \langle \hat{P}_k^{(\alpha_1)} \cdots \hat{P}_{k+R-1}^{(\alpha_R)} \rangle_{\hat{\varrho}} \hat{P}_k^{(\alpha_1)} \cdots \hat{P}_{k+R-1}^{(\alpha_R)}; \quad (1)$$

i.e., the $\hat{\varrho}_k$ are completely specified by the local expectation values $\langle \hat{P}_k^{(\alpha_1)} \cdots \hat{P}_{k+R-1}^{(\alpha_R)} \rangle_{\hat{\varrho}} = \mathrm{tr}[\hat{\varrho} \hat{P}_k^{(\alpha_1)} \cdots \hat{P}_{k+R-1}^{(\alpha_R)}]$, estimates to which are the input to our tomography scheme.

The $\hat{\varrho}_k$ completely specify the state $\hat{\varrho}$ if a certain technical invertibility condition is met. The proof is constructive and gives an explicit method for obtaining $\hat{\varrho}$ from the $\hat{\varrho}_k$. It is partly inspired by the characterization of finitely correlated states (as opposed to C^*-finitely correlated states) on infinite spin chains as described in the early literature (see Proposition 2.1 of Ref. [15]). In addition to the fact that we are working in a finite and nontranslation invariant setting, the main novel technical point here is that we only use local information, provided by the $\hat{\varrho}_k$. To state the invertibility condition, we first need to establish some notation. We collect the N sites of the one-dimensional system in the set $\mathcal{N} = \{1,\ldots,N\}$. For $\mathcal{I} \subset \mathcal{N}$, we define the complex vector spaces $V_{\mathcal{I}}$ spanned by

$$\left\{ \prod_{i\in\mathcal{I}} \hat{P}_i^{(\alpha_i)} \right\}_{\alpha_i=1,\ldots,d^2}. \quad (2)$$

For given $\hat{O} \in V_{\mathcal{N}}$ and $\mathcal{I},\mathcal{J} \subset \mathcal{N}$ we define the linear map $E_{\mathcal{I}}^{\mathcal{J}}: V_{\mathcal{J}} \to V_{\mathcal{N}\setminus\mathcal{I}}$ as

$$\hat{X} \mapsto E_{\mathcal{I}}^{\mathcal{J}}(\hat{X}) = \mathrm{tr}_{\mathcal{N}\setminus\mathcal{I}}[\hat{X}\hat{O}]. \quad (3)$$

We note that the map $E_{\mathcal{I}}^{\mathcal{J}}$ depends only on the reduction $\hat{O}_{\mathcal{I}\cup\mathcal{J}} = \mathrm{tr}_{\mathcal{N}\setminus\mathcal{I}\cup\mathcal{J}}[\hat{O}]$ of $\hat{O}$ to sites $\mathcal{I}\cup\mathcal{J}$ as

$$E_{\mathcal{I}}^{\mathcal{J}}(\hat{X}) = \mathrm{tr}_{\mathcal{J}}[\hat{X}\hat{O}_{\mathcal{I}\cup\mathcal{J}}]; \quad (4)$$

this is illustrated in Fig. 1. Note that from now on we will only consider cases where $\mathcal{I}\cup\mathcal{J}$ is connected.

Definition 1 (Invertibility): Let l, $r \in \mathbb{N}$, $2 \le l+r \le N-2$. If $\hat{O}$ is such that for all $k \in \mathbb{N}$, $l \le k \le N-r-1$, the equality

$$\mathrm{rank}\left[E_{\{k-l+1,\ldots,k\}}^{\{k+1,\ldots,k+r\}}\right] = \mathrm{rank}\left[E_{\{1,\ldots,k\}}^{\{k+1,\ldots,N\}}\right] \quad (5)$$

holds, we call $\hat{O}$ (l,r) invertible.

We may now state the main theorem, a proof of which may be found in Sec. A of the Supplemental Material [16].

Theorem 1: Let l, $r \in \mathbb{N}$ such that $2 \le l+r \le N-2$. Let $\hat{O} \in V_{\mathcal{N}}$ be (l,r) invertible. Then, for all $\hat{X}_i \in V_{\{i\}}$, the equality

$$\mathrm{tr}_{\mathcal{N}}[\hat{X}_1 \cdots \hat{X}_N \hat{O}] = \mathrm{tr}_{\mathcal{N}}[\hat{X}_1 \cdots \hat{X}_l \hat{Y}_l \hat{O}] \quad (6)$$

holds. Here, the $\hat{Y}_l \in V_{\{l+1,\ldots,l+r\}}$ are recursively defined as follows. We set $\hat{Y}_{N-r} = \hat{X}_{N-r+1} \cdots \hat{X}_N$ and

$$\hat{Y}_{k-1} = \bar{E}_{\{k-l,\ldots,k-1\}}^{\{k,\ldots,k+r-1\}}\left(E_{\{k-l,\ldots,k-1\}}^{\{k,\ldots,k+r\}}(\hat{X}_k \hat{Y}_k)\right) \quad (7)$$

for $k = l+1,\ldots,N-r$. Here, the bar indicates the Moore-Penrose pseudoinverse.

Note that for Eq. (6) the reduction of $\hat{O}$ to sites $\{1,\ldots,l+r\}$ is needed; for the inverse we require the reduction of $\hat{O}$ to sites $\{k-l,\ldots,k+r-1\}$, and for $E_{\{k-l,\ldots,k-1\}}^{\{k,\ldots,k+r\}}(\hat{X}_k\hat{Y}_k)$ we require the reduction of $\hat{O}$ to sites

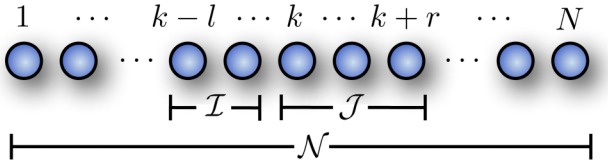

FIG. 1 (color online). Definition of the sets $\mathcal{N} = \{1,\ldots,N\}$ and $\mathcal{I}$, $\mathcal{J} \subset \mathcal{N}$. The linear map $E_{\mathcal{I}}^{\mathcal{J}}(\hat{X}) = \mathrm{tr}_{\mathcal{J}}[\hat{X}\hat{O}_{\mathcal{I}\cup\mathcal{J}}]$ maps operators $\hat{X}$ (e.g., observables) living on set $\mathcal{J}$ into operators on set $\mathcal{I}$ by means of the reductions of $\hat{O}$ (e.g., the state) to $\mathcal{I}\cup\mathcal{J}$.

$\{k-l,\ldots,k+r\}$. Hence, expectation values of the form $\mathrm{tr}_{\mathcal{N}}[\hat{X}_1\cdots\hat{X}_N\hat{O}]$ are completely specified by reductions of $\hat{O}$ to the sites $\{k-l,\ldots,k+r\}$, $k=l+1,\ldots,N-r$, i.e., by all reductions to $R=r+l+1$ contiguous sites. By choosing the $\hat{X}_i$ to be the basis operators $\hat{P}_i^{\alpha_i}$, this implies that $(l,\ r)$-invertible operators $\hat{O}$ may be fully reconstructed from their reductions to R consecutive sites, which is the same as knowing the expectation values

$$\mathrm{tr}[\hat{P}_k^{\alpha_k}\cdots\hat{P}_{k+R-1}^{\alpha_{k+R-1}}\hat{O}],\quad \alpha_i=1,\ldots,d^2, \tag{8}$$

for all $k=1,\ldots,N-R+1$.

One can prove that a vast majority of matrix product operators fulfil the invertibility condition; i.e., a vast majority of matrix product operators may be reconstructed from local reductions alone (see Sec. B of the Supplemental Material [16] for a technical proof). As noted above, practically relevant states are (well approximated by) matrix product operators of low dimension; i.e., we expect the scheme to work for a large class of mixed states. Now, of course, experimentally, the exact expectation values even for states satisfying the invertibility condition are only known to within a certain statistical error (e.g., the estimated standard deviation of the mean after a finite number of measurements). This error propagates into the singular values of the map $E^{\{k,\ldots,k+r-1\}}_{\{k-l,\ldots,k-1\}}$. As this map needs to be inverted, even small errors on singular values close to zero will lead to a large error in the reconstruction. This issue may be avoided by using stochastic robust approximation techniques [17–19] (see Sec. C of the Supplemental Material [16] for technical details). Before we apply the reconstruction scheme to experimental data, we present numerical results for states that do not necessarily fulfil the invertibility condition and for which the local expectation values are subject to inevitable statistical noise.

We restrict our attention to qubits $d=2$, and illustrate the behavior of the tomography scheme for thermal states of the Ising Hamiltonian at its quantum critical point

$$\hat{H}=-\sum_{i=1}^{N-1}\hat{\sigma}_i^x\hat{\sigma}_{i+1}^x-\sum_{i=1}^{N}\hat{\sigma}_i^z. \tag{9}$$

We obtain the thermal states by an imaginary time evolution [20,21] using the time evolving block decimation (TEBD) algorithm. We simulate the measurements in the following way. We first compute the exact local expectation values $p^k_{\alpha_1,\ldots,\alpha_R}=\langle\hat{\sigma}_k^{(\alpha_1)}\cdots\hat{\sigma}_{k+R-1}^{(\alpha_R)}\rangle_{\hat{\varrho}}$, $\alpha_i=0,x,y,z$, for all k. Statistical noise is then simulated by adding random numbers (drawn from a Gaussian distribution with zero mean and standard deviation σ) to them. The resulting $\bar{p}^k_{\alpha_1,\ldots,\alpha_R}$ then serve as the input to our reconstruction scheme. We compare the reconstructed state $\hat{\varrho}_{\mathrm{rec}}$ to the exact state $\hat{\varrho}$ by computing the Hilbert-Schmidt norm difference $D(\hat{\varrho},\hat{\varrho}_{\mathrm{rec}})=\|\hat{\varrho}_{\mathrm{rec}}-\hat{\varrho}\|^2/\|\hat{\varrho}\|^2$. To obtain meaningful results, we have rescaled the norm such that the deviations are measured in units of $\|\hat{\varrho}\|^2$, the natural length scale of the state to be learned. In Fig. 2, we show the norm difference for the exact and the reconstructed states as a function of the system size N and the error σ. It indicates that, for given N, the error $D(\hat{\varrho},\hat{\varrho}_{\mathrm{rec}})$ scales roughly as σ; similarly, for given σ, it scales roughly as N. In Sec. D of the Supplemental Material [16] we provide further numerical experiments analyzing the performance of the algorithm for thermal states of random next-neighbor Hamiltonians and mixed states obtained by tracing out parts of a matrix product state in a larger Hilbert space. Again, these numerical results suggest that the scaling of our scheme is polynomial in both N and σ.

Let us finally apply the reconstruction scheme to experimental data obtained in an ion trap experiment in a full quantum state tomography setting. The considered state is a W state implemented on $N=8$ qubits with local phases [2], i.e.,

$$|W(\phi)\rangle=[|0\ldots001\rangle+e^{i\phi_1}|0\ldots010\rangle+\cdots+e^{i\phi_{N-1}}|1\ldots000\rangle]/\sqrt{N}. \tag{10}$$

The available experimental data are the set of relative frequencies corresponding to 100 measurements in each of the 3^N different basis rotations (measurements along the X, Y, and Z directions). From these, we obtain maximum likelihood estimates to the reduced density matrices on all blocks of R sites [22]. As described in the Supplemental Material [16], we apply a stochastic robust approximation technique to avoid difficulties in ill-conditioned inversion problems making use of the Fisher information matrix of the local estimates [22]. Let us stress that the input to the reconstruction scheme are merely the relative frequencies corresponding to the measurements on all subsystems of R contiguous sites and the total number of measurements. Absolute values of the reconstructed density matrices for

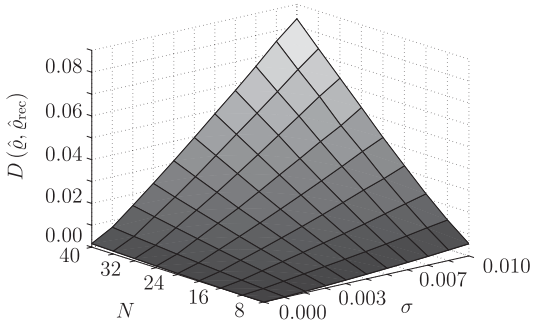

FIG. 2. Quality of our reconstruction scheme for thermal states of the Ising Hamiltonian in Eq. (9) for $\beta=5$ and $R=5$, i.e., the state is reconstructed from local expectation values on five consecutive sites. For each pair $(N,\ \sigma)$, the plot shows the mean of the norm difference obtained from 100 realizations and renormalized by the purity of the exact state, i.e., $D(\hat{\varrho},\hat{\varrho}_{\mathrm{rec}})=\|\hat{\varrho}_{\mathrm{rec}}-\hat{\varrho}\|^2/\|\hat{\varrho}\|^2$. This corresponds to 100 experiments, each of which carries an uncertainty of σ about the local expectation values.

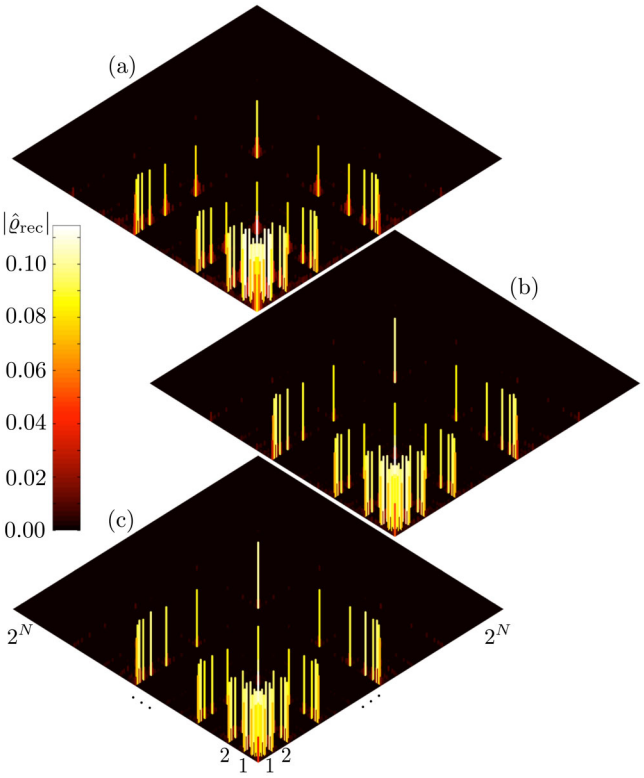

FIG. 3 (color online). Absolute value $|\hat{\varrho}_{\mathrm{rec}}|$ of the corresponding reconstructed density matrix of the experimentally realized W state. (a) Reconstructed operator using the scheme described in this Letter where the reductions to $R = 3$ sites are known. (b) Estimate with $R = 5$ sites. (c) Maximum likelihood estimate of full quantum state tomography (see also Ref. [2]). The numbers $1, 2, \ldots, 2^N$ denote the entries of the density matrix $\hat{\varrho}_{\mathrm{rec}}$.

$R = 3$ and $R = 5$ along with the maximum likelihood estimate obtained in the full tomography procedure [2] are presented in Fig. 3. Comparing the maximum likelihood estimate with our results we find for the renormalized Hilbert-Schmidt norm difference $D(\hat{\varrho}_{\mathrm{ML}}, \hat{\varrho}_{\mathrm{rec}}) = 0.087$ for $R = 3$ and $D(\hat{\varrho}_{\mathrm{ML}}, \hat{\varrho}_{\mathrm{rec}}) = 0.012$ for $R = 5$. For the full quantum state tomography experiment, maximizing the fidelity of the maximum likelihood estimate with respect to the local phases of a pure W state yields $f = \langle W(\phi_{\mathrm{opt}})|\hat{\varrho}|W(\phi_{\mathrm{opt}})\rangle = 0.722$ [2]. With the matrix product operator scheme we achieve a fidelity of $f = 0.688$ for $R = 3$ and $f = 0.718$ for $R = 5$ with respect to the optimal W state $|W(\phi_{\mathrm{opt}})\rangle$ revealing that the main contribution in our estimates stems from the same $|W(\phi_{\mathrm{opt}})\rangle$ as in [2]. We are only using local information and hence a local addressing of the ions in the trap is sufficient, resulting in the linear scaling of the scheme with the number of constituents. Further, the full maximum likelihood algorithm uses a huge amount of resources since it requires the storage and manipulation of 6^N measurement operators resulting in a time consuming postprocessing. In contrast, our reconstruction takes about one second on a laptop given the local maximum likelihood estimates and the corresponding Fisher information matrices [22].

In this Letter we have presented a scheme to reconstruct mixed states from local measurements efficiently. We have shown that, in principle, all states may be reconstructed from reductions to contiguous sets of sites alone and that the reconstruction is efficient with respect to the measurement time and the postprocessing resources for practically relevant states. It should be noted, however, that our rigorous performance guarantees apply only when the model assumption of an essentially one-dimensional structure is justified. As is the case for most statistical estimators, the scheme is not suitable for model selection; i.e., it cannot certify unconditionally from data alone that the model is valid. To investigate the latter issue, the impact of statistical noise and the performance of the reconstruction scheme for states that do not necessarily fulfil the condition which guarantees perfect reconstruction have been investigated for simulated states and experimental data in detail. For all simulations the Hilbert-Schmidt norm difference (normalized by the purity of the exact state) between the exact state and the reconstructed state was obtained and the numerical results suggest that the quality of the reconstruction scales algebraically in N and σ. The methods presented here hence pave the way for the reconstruction of mixed states of a large number of qubits.

We gratefully acknowledge H. Häffner for providing experimental data, R. Rosenbach for results of the TEBD algorithm, and K. Audenaert for fruitful discussions. Computations of the TEBD algorithm were performed on the bwGRiD [23]. This work was supported by the Alexander von Humboldt Foundation, the EU Integrated Projects QESSENCE and SIQS, the BMBF Verbundprojekt QuOReP, the Excellence Initiative of the German Federal and State Governments (Grant No. ZUK 43), and the Swiss National Science Foundation.

[1] *Quantum State Estimation*, Lecture Notes in Physics Vol. 649. edited by M. G. A. Paris and J. Řeháček (Springer, Berlin Heidelberg, 2004).
[2] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, Nature (London) **438**, 643 (2005).
[3] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, R. Reichle, and D. J. Wineland, Nature (London) **438**, 639 (2005).
[4] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, Phys. Rev. Lett. **106**, 130506 (2011).
[5] X.-C. Yao, T.-X. Wang, P. Xu, H. Lu, G.-S. Pan, X.-H. Bao, C.-Z. Peng, C.-Y. Lu, Y.-A. Chen, and J.-W. Pan, Nat. Photonics **6**, 225 (2012).

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

[7] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Phys. Rev. Lett. **105**, 150401 (2010); G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, Phys. Rev. Lett. **105**, 250403 (2010); D. Gross, IEEE Trans. Inf. Theory **57**, 1548 (2011); M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Phys. Rev. Lett. **107**, 210404 (2011); S. T. Flammia and Y.-K. Liu, Phys. Rev. Lett. **106**, 230501 (2011); S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, New J. Phys. **14**, 095022 (2012); O. Landon-Cardinal and D. Poulin, New J. Phys. **14**, 085004 (2012).

[8] P. Schindler, M. Müller, D. Nigg, J. T. Barreiro, E. A. Martinez, M. Hennrich, T. Monz, S. Diehl, P. Zoller, and R. Blatt, Nat. Phys. **9**, 361 (2013).

[9] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Phys. Rev. A **64**, 052312 (2001).

[10] M. B. Hastings, J. Stat. Mech. (2007) P08024.

[11] M. B. Plenio, J. Eisert, J. Dreißig, and M. Cramer, Phys. Rev. Lett. **94**, 060503 (2005).

[12] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).

[13] M. B. Hastings, Phys. Rev. B **73**, 085115 (2006).

[14] There is of course no contradiction between the exponential speed up of quantum computation and the fact that some states in large ambient spaces can be efficiently handled. Indeed, we believe it is a key contribution of quantum information to develop complementary tools with the eventual aim to decide which states allow for an efficient simulation (like the ones treated here) and which ones do not (e.g., those enabling universal quantum computation).

[15] M. Fannes, B. Nachtergaele, and R. F. Werner, Commun. Math. Phys. **144**, 443 (1992).

[16] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.111.020401 for proofs, additional material, and further numerical experiments.

[17] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, England, 2004).

[18] J. S. Lundeen, A. Feito, H. Coldenstrodt-Ronge, K. L. Pregnell, Ch. Silberhorn, T. C. Ralph, J. Eisert, M. B. Plenio, and I. A. Walmsley, Nat. Phys. **5**, 27 (2008).

[19] L. Zhang, H. Coldenstrodt-Ronge, A. Datta, G. Puentes, J. S. Lundeen, X.-M. Jin, B. J. Smith, M. B. Plenio, and I. A. Walmsley, Nat. Photonics **6**, 364 (2012).

[20] M. Zwolak and G. Vidal, Phys. Rev. Lett. **93**, 207205 (2004).

[21] To generate the thermal states we use a fourth order Trotter expansion and adjust the step size in the imaginary time evolution to $\delta T = 10^{-3}$ for the Ising Hamiltonian and the random next-neighbor Hamiltonians, respectively, i.e., $\hat{\varrho} = (e^{-\delta T \hat{H}})^{\beta/\delta T}/Z$. Further, we keep the 100 largest singular values in each decomposition, i.e., approximate the state as a matrix product operator with bond dimension $D = 100$.

[22] Z. Hradil, J. Řeháček, J. Fiurášek, and M. Ježek, Lect. Notes Phys. **649**, 59 (2004).

[23] The bwGRID (http://www.bw-grid.de) is a distributed aggregation of computer nodes, coupled within the framework of the German D-Grid Initiative, funded by the Ministry for Education and Research and the Ministry for Science, Research and Arts Baden-Württemberg.