Efficient Classical Simulation of Slightly Entangled Quantum Computations

Guifré Vidal

Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

(Received 25 February 2003; published 1 October 2003)

We present a classical protocol to efficiently simulate any pure-state quantum computation that involves only a restricted amount of entanglement. More generally, we show how to classically simulate pure-state quantum computations on n qubits by using computational resources that grow linearly in n and exponentially in the amount of entanglement in the quantum computer. Our results imply that a necessary condition for an exponential computational speedup (with respect to classical computations) is that the amount of entanglement increases with the size n of the computation, and provide an explicit lower bound on the required growth.

DOI: 10.1103/PhysRevLett.91.147902 PACS numbers: 03.67.Lx, 03.65.Ud, 03.67.Hk, 03.67.Mn

Identifying nontrivial quantum dynamics that are efficiently simulatable on a classical computer is important in quantum many-body physics and in quantum information science. On the one hand, our understanding of quantum many-body dynamics is severely hindered by the fact that, generically, the description of the state of n interacting quantum systems requires $\mathcal{O}(\exp(n))$ parameters, rendering the simulation of such systems intractable. Consequently, those exceptional cases where an efficient simulation is possible are a precious source of insight into the physics of quantum many-body systems.

In quantum information [1], on the other hand, uncovering simulatable quantum dynamics is of interest in order to characterize the essential resources required for *genuine* quantum computation. Clearly, if a quantum device is to offer an exponential speedup with respect to classical computations, then it must involve dynamics that *cannot* be efficiently simulated classically. Thus by detecting simulatable quantum dynamics we learn about which systems may be used as a quantum computer. Examples of quantum evolutions that can be efficiently simulated are given by n qubits prepared in a computational-basis state and acted upon by gates from the Clifford group, as established by the Gottesman-Knill theorem [2], and by n bosonic (fermionic) modes evolving in a Gaussian state [3] (respectively, [4]). Also, Jozsa and Linden [5] have recently shown how to efficiently simulate any quantum computation on n qubits such that their state factors, at all times, into a direct product of states of at most a constant (i.e., independent of n) number of qubits.

In this Letter we show that any quantum computation with pure states can be efficiently simulated with a classical computer provided the amount of entanglement involved is sufficiently restricted. Here entanglement is quantified by a most natural (yet not standard) measure and may involve all n qubits of the quantum computer in a single, nonfactorizable state. More generally, we give an upper bound on the computational speedup achievable by a quantum computation, in terms of how the entanglement in the system scales with the number n of qubits. Thereby we (i) establish that entanglement is a *necessary* (but in general not *sufficient*) resource for genuine quantum computation with pure states and (ii) provide an explicit lower bound for the entanglement required in quantum computational speedups. (See also [6–8] for related discussions on the case of mixed-state quantum computations.) In addition, the simulation protocol presented here can be adapted as to apply to a class of quantum many-body systems, including spin chains, as explored in depth in a future contribution.

Measure of entanglement.—Let $|\Psi\rangle \in \mathcal{H}_2^{\otimes n}$ denote a pure state of n qubits, A a subset of the n qubits, and B the rest of them. The Schmidt decomposition (SD) [9] of $|\Psi\rangle$ with respect to the partition $A{:}B$ reads

$$|\Psi\rangle = \sum_{\alpha=1}^{\chi_A} \lambda_\alpha |\Phi_\alpha^{[A]}\rangle \otimes |\Phi_\alpha^{[B]}\rangle, \tag{1}$$

where $|\Phi_\alpha^{[A]}\rangle$ ($|\Phi_\alpha^{[B]}\rangle$) is an eigenvector of the reduced density matrix $\rho^{[A]}$ ($\rho^{[B]}$) with eigenvalue $|\lambda_\alpha|^2 > 0$, and the Schmidt coefficient λ_α follows from the relation $\langle\Phi_\alpha^{[A]}|\Psi\rangle = \lambda_\alpha|\Phi_\alpha^{[B]}\rangle$. The rank χ_A of ρ_A is a natural measure of the entanglement between the qubits in A and those in B [10]. Therefore, we could use χ,

$$\chi \equiv \max_A \chi_A, \tag{2}$$

i.e., the maximal value of χ_A over all possible bipartite splittings $A{:}B$ of the n qubits, to quantify the entanglement of state $|\Psi\rangle$. Here we mainly use the related entanglement measure E_χ,

$$E_\chi \equiv \log_2(\chi), \tag{3}$$

which fulfills the following appealing properties: (i) $0 \leq E_\chi \leq n\log_2(d)/2$ for $|\Psi\rangle \in \mathcal{H}_d^{\otimes n}$, with $E_\chi = 0$ if and only if $|\Psi\rangle$ is a product (i.e., completely unentangled) state [13]; (ii) E_χ is an entanglement monotone [11] that decreases both under deterministic and stochastic local manipulations of the system; (iii) E_χ is additive under tensor products, $E_\chi(\Psi \otimes \Psi') = E_\chi(\Psi) + E_\chi(\Psi')$; and (iv) in the bipartite setting E_χ upper bounds the more standard measure *entropy of entanglement* [12].

State decomposition.—Let us now consider the expansion of $|\Psi\rangle \in \mathcal{H}_2^{\otimes n}$ in the computational basis,

$$|\Psi\rangle = \sum_{i_1=0}^{1} \cdots \sum_{i_n=0}^{1} c_{i_1 \cdots i_n} |i_1\rangle \otimes \cdots \otimes |i_n\rangle. \quad (4)$$

The key ingredient of our simulation protocol is a particular decomposition of coefficients $c_{i_1 i_2 \cdots i_n}$, namely,

$$c_{i_1 i_2 \cdots i_n} = \sum_{\alpha_1, \cdots, \alpha_{n-1}} \Gamma^{[1] i_1}_{\alpha_1} \lambda^{[1]}_{\alpha_1} \Gamma^{[2] i_2}_{\alpha_1 \alpha_2} \lambda^{[2]}_{\alpha_2} \Gamma^{[3] i_3}_{\alpha_2 \alpha_3} \cdots \Gamma^{[n] i_n}_{\alpha_{n-1}}. \quad (5)$$

This decomposition employs n tensors $\{\Gamma^{[1]}, \cdots, \Gamma^{[n]}\}$ and $n-1$ vectors $\{\lambda^{[l]}, \cdots, \lambda^{[n-1]}\}$, whose indices i_l and α_l take values in $\{0, 1\}$ and $\{1, \cdots, \chi\}$, respectively. Therefore, (5) reexpresses the 2^n coefficients $c_{i_1 i_2 \cdots i_n}$ of (4) in terms of about $(2\chi^2 + \chi)n$ parameters, a number that is linear in n and exponential in E_χ,

$$\begin{matrix} n \text{ qubit} \\ \text{state} \end{matrix} \leftrightarrow \begin{matrix} n \exp(E\chi) \\ \text{parameters.} \end{matrix} \quad (6)$$

For a generic state $|\Psi\rangle$, E_χ is of the order n and the decomposition in terms of Γ's and λ's is uninteresting, for it employs $\mathcal{O}(n \exp(n))$ parameters. However, notice that if E_χ scales as $\mathcal{O}(\log(n))$, then only poly(n) parameters are required, leading to an efficient description of $|\Psi\rangle$.

Decomposition (5) depends on the particular way qubits have been ordered from 1 to n, and essentially consists of a concatenation of $n-1$ SDs. More explicitly, we first compute the SD of $|\Psi\rangle$ according to the bipartite splitting of the systems into qubit 1 and the $n-1$ remaining qubits (from now on we omit the tensor product symbol),

$$|\Psi\rangle = \sum_{\alpha_1} \lambda^{[1]}_{\alpha_1} |\Phi^{[1]}_{\alpha_1}\rangle |\Phi^{[2 \cdots n]}_{\alpha_1}\rangle \quad (7)$$

$$= \sum_{i_1, \alpha_1} \Gamma^{[1] i_1}_{\alpha_1} \lambda^{[1]}_{\alpha_1} |i_1\rangle |\Phi^{[2 \cdots n]}_{\alpha_1}\rangle, \quad (8)$$

where in the last line we have expanded each Schmidt vector $|\Phi^{[1]}_{\alpha_1}\rangle = \sum_{i_1} \Gamma^{[1] i_1}_{\alpha_1} |i_1\rangle$ in terms of the basis vectors $\{|0\rangle, |1\rangle\}$ for qubit 1. We then proceed according to the following three steps: (i) we expand each Schmidt vector $|\Phi^{[2 \cdots n]}_{\alpha}\rangle$ in a local basis for qubit 2,

$$|\Phi^{[2 \cdots n]}_{\alpha_1}\rangle = \sum_{i_2} |i_2\rangle |\tau^{[3 \cdots n]}_{\alpha_1 i_2}\rangle; \quad (9)$$

(ii) then we write each (possibly unnormalized) vector $|\tau^{[3 \cdots n]}_{\alpha_1 i_2}\rangle$ in terms of the *at most* χ Schmidt vectors $\{|\Phi^{[3 \cdots n]}_{\alpha_2}\rangle\}_{\alpha_2=1}^{\chi}$ (i.e., the eigenvectors of $\rho^{[3 \cdots n]}$) and the corresponding Schmidt coefficients $\lambda^{[2]}_{\alpha_2}$,

$$|\tau^{[3 \cdots n]}_{\alpha_1 i_2}\rangle = \sum_{\alpha_2} \Gamma^{[2] i_2}_{\alpha_1 \alpha_2} \lambda^{[2]}_{\alpha_2} |\Phi^{[3 \cdots n]}_{\alpha_2}\rangle; \quad (10)$$

(iii) finally we substitute Eq. (10) in Eq. (9) and the latter in Eq. (8) to obtain

$$|\Psi\rangle = \sum_{i_1, \alpha_1, i_2, \alpha_2} \Gamma^{[1] i_1}_{\alpha_1} \lambda^{[1]}_{\alpha_1} \Gamma^{[2] i_2}_{\alpha_1 \alpha_2} \lambda^{[2]}_{\alpha_2} |i_1 i_2\rangle |\Phi^{[3 \cdots n]}_{\alpha_2}\rangle. \quad (11)$$

Iterating steps (i)–(iii) for the Schmidt vectors $|\Phi^{[3 \cdots n]}_{\alpha_2}\rangle, |\Phi^{[4 \cdots n]}_{\alpha_3}\rangle, \cdots, |\Phi^{[n]}_{\alpha_{n-1}}\rangle$, one can express state $|\Psi\rangle$ as in (5).

A useful feature of the description of $|\Psi\rangle$ in terms of the Γ's and λ's of Eq. (5) is that it readily gives the SD of $|\Psi\rangle$ according to $[1 \cdots l]{:}[(l+1) \cdots n]$, i.e., the bipartite splitting $A{:}B$ such that A contains the first l qubits and B the rest of them,

$$|\Psi\rangle = \sum_{\alpha_l} \lambda^{[l]}_{\alpha_l} |\Phi^{[1 \cdots l]}_{\alpha_l}\rangle |\Phi^{[(l+1) \cdots n]}_{\alpha_l}\rangle. \quad (12)$$

Indeed, it can be checked by induction over l that

$$|\Phi^{[1 \cdots l]}_{\alpha_l}\rangle = \sum_{\alpha_1, \cdots, \alpha_{l-1}} \Gamma^{[1] i_1}_{\alpha_1} \lambda^{[1]}_{\alpha_1} \cdots \Gamma^{[l] i_l}_{\alpha_{l-1} \alpha_l} |i_1 \cdots i_l\rangle, \quad (13)$$

whereas by construction we already had that

$$|\Phi^{[(l+1) \cdots n]}_{\alpha_l}\rangle = \sum_{\alpha_{l+1}, \cdots, \alpha_n} \Gamma^{[l+1] i_{l+1}}_{\alpha_l \alpha_{l+1}} \cdots \lambda^{[n-1]}_{\alpha_{n-1}} \Gamma^{[n] i_n}_{\alpha_{n-1}} |i_{l+1} \cdots i_n\rangle. \quad (14)$$

Finally, insight into the meaning of decomposition (5) may be gained by defining χ^2 unnormalized states $|\varphi^{[l]}_{\alpha\alpha'}\rangle = \sum_i \Gamma^{[l] i}_{\alpha\alpha'} |i\rangle$ for qubit l and expanding $|\Psi\rangle$ as a linear combination of χ^{n-1} product states,

$$|\Psi\rangle = \sum_{\alpha_1, \cdots, \alpha_{n-1}} |\varphi^{[1]}_{\alpha_1}\rangle \lambda^{[1]}_{\alpha_1} |\varphi^{[2]}_{\alpha_1 \alpha_2}\rangle \cdots \lambda^{[n-1]}_{\alpha_{n-1}} |\varphi^{[n]}_{\alpha_{n-1}}\rangle, \quad (15)$$

where the sum over α's is what accounts for the correlations between qubits [14].

Update of the decomposition.—The following lemmas explain how to update the description of state $|\Psi\rangle$ when a one-qubit gate or a two-qubit gate (acting on consecutive qubits) is applied to the system. Remarkably, the computational cost of the updating is independent of the total number n of qubits and grows in χ only as a polynomial of low degree.

Lemma 1. *Updating the Γ's and λ's of state $|\Psi\rangle$ after a unitary operation Uacts on qubit l involves transforming only $\Gamma^{[l]}$. The incurred computational cost is of $\mathcal{O}(\chi^2)$ basic operations.*

Proof. *In the SD according to the splitting $[1 \cdots (l-1)]{:}[l \cdots n]$, a unitary operation U on qubit l does not modify the Schmidt vectors for part $[1 \cdots (l-1)]$ and therefore $\Gamma^{[j]}$ and $\lambda^{[j]}$ $(1 \le j \le l-1)$ remain the same. Similarly, by considering the SD for the splitting $[1 \cdots l]{:}[(l+1) \cdots n]$, we conclude that also $\Gamma^{[j]}$ and $\lambda^{[j-1]}$ $(l+1 \le j \le n)$ remain unaffected. Instead, $\Gamma^{[l]}$ changes according to*

$$\Gamma'^{[l] i}_{\alpha\beta} = \sum_{j=0,1} U^i_j \Gamma^{[l] j}_{\alpha\beta} \forall\, \alpha, \beta = 1, \cdots, \chi. \quad (16)$$

That is, for each value of α and β, a small matrix

multiplication is required, leading to a total of $\mathcal{O}(\chi^2)$ *basic operations.*

Lemma 2. *Updating the* Γ*'s and* λ*'s of state* $|\Psi\rangle$ *after a unitary operation* V *acts on qubits* l *and* $l+1$ *involves transforming only* $\Gamma^{[l]}$, $\lambda^{[l]}$, *and* $\Gamma^{[l+1]}$. *This can be achieved with* $\mathcal{O}(\chi^3)$ *basic operations.*

Proof. *In order to ease the notation we regard* $|\Psi\rangle$ *as belonging to only four subsystems,*

$$\mathcal{H} = \mathcal{J} \otimes \mathcal{H}_C \otimes \mathcal{H}_D \otimes \mathcal{K}. \tag{17}$$

Here, $\mathcal{J}$ *is spanned by the* χ *eigenvectors of the reduced density matrix*

$$\rho^{[1\cdots(l-1)]} = \sum_\alpha |\alpha\rangle\langle\alpha|, \qquad |\alpha\rangle \equiv \lambda_\alpha^{[l-1]} |\Phi_\alpha^{[1\cdots(l-1)]}\rangle; \tag{18}$$

and, similarly, $\mathcal{K}$ *is spanned by the* χ *eigenvectors of the reduced density matrix*

$$\rho^{[(l+2)\cdots n]} = \sum_\gamma |\gamma\rangle\langle\gamma|, \qquad |\gamma\rangle \equiv \lambda_\gamma^{[l+1]} |\Phi_\gamma^{[(l+2)\cdots n]}\rangle; \tag{19}$$

whereas $\mathcal{H}_C$ *and* $\mathcal{H}_D$ *correspond, respectively, to qubits* l *and* $l+1$. *In this notation we have*

$$|\Psi\rangle = \sum_{\alpha,\beta,\gamma=1}^{\chi} \sum_{i,j=0}^{1} \Gamma_{\alpha\beta}^{[C]i} \lambda_\beta \Gamma_{\beta\gamma}^{[D]j} |\alpha i j \gamma\rangle, \tag{20}$$

and, reasoning as in the proof of Lemma 1, when applying unitary V *to qubits* C *and* D *we need update only* $\Gamma^{[C]}$, λ, $\Gamma^{[D]}$. *We can expand* $|\Psi'\rangle \equiv V|\Psi\rangle$ *as*

$$|\Psi'\rangle = \sum_{\alpha,\gamma=1}^{\chi} \sum_{i,j=0}^{1} \Theta_{\alpha\gamma}^{ij} |\alpha i j \gamma\rangle, \tag{21}$$

where

$$\Theta_{\alpha\gamma}^{ij} = \sum_\beta \sum_{kl} V_{kl}^{ij} \Gamma_{\alpha\beta}^{[C]k} \lambda_\beta \Gamma_{\beta\gamma}^{[D]l}. \tag{22}$$

By diagonalizing $\rho'^{[D\mathcal{K}]}$,

$$\rho'^{[D\mathcal{K}]} = \mathrm{tr}_{JC} |\Psi'\rangle\langle\Psi'| = \sum_{j,j',\gamma,\gamma'} \left(\sum_{\alpha,i} \langle\alpha|\alpha\rangle \Theta_{\alpha\gamma}^{ij} (\Theta_{\alpha\gamma'}^{ij'})^* \right) |j\gamma\rangle\langle j'\gamma'|, \tag{23}$$

we obtain its eigenvectors $\{|\Phi_\beta'^{[D\mathcal{K}]}\rangle\}$, *which we can expand in terms of* $\{|j\gamma\rangle\}$ *to obtain* $\Gamma'^{[D]}$,

$$|\Phi_\beta'^{[D\mathcal{K}]}\rangle = \sum_{j,\gamma} \Gamma_{\beta\gamma}'^{[D]j} |j\gamma\rangle. \tag{24}$$

The eigenvectors of $\rho'^{[JC]}$ *and* λ' *follow then from*

$$\lambda'_\beta |\Phi_\beta'^{[JC]}\rangle = \langle\Phi_\beta'^{[D\mathcal{K}]}|\Psi'\rangle \tag{25}$$

$$= \sum_{i,j,\alpha,\gamma} (\Gamma_{\beta\gamma}'^{[D]j})^* \Theta_{\alpha\gamma}^{ij} \langle\gamma|\gamma\rangle |\alpha i\rangle, \tag{26}$$

and by expanding each $|\Phi_\beta'^{[JC]}\rangle$,

$$|\Phi_\beta'^{[JC]}\rangle = \sum_{i\alpha} \Gamma_{\alpha\beta}'^{[C]i} |\alpha i\rangle, \tag{27}$$

we also obtain $\Gamma'^{[C]}$. *In the above manipulations, the largest tensors contain* $\mathcal{O}(\chi^2)$ *components. The most expensive manipulations are the following: in Eq.* (22), *for each value of* α *and* γ *the sum over* β *requires summing up* $\mathcal{O}(\chi)$ *terms, leading to a total of* $\mathcal{O}(\chi^3)$ *basic operations; similarly* $\mathcal{O}(\chi^3)$ *operations are required in Eq.* (23) *and in Eq.* (26); *finally, the same cost applies to the diagonalization of* $\rho'^{[J\mathcal{C}]}$, *a matrix of* χ^2 *elements, so that the overall update can be done with* $\mathcal{O}(\chi^3)$ *basic operations.*

Simulation protocol.—We consider a pure-state quantum computation using n qubits, initially in state $|0\rangle^{\otimes n}$, and consisting of poly(n) one- and two-qubit gates and a final measurement in the computational basis. The simulation protocol works as follows. We use tensors Γ's and λ's to store the initial state $|0\rangle^{\otimes n}$ and update its description as the gates are applied [15]. Recall that decomposition (5) assumes a specific ordering of the qubits. In order to update $|\Psi\rangle$ according to a two-qubit gate between nonconsecutive qubits C and D, we first simulate $\mathcal{O}(n)$ swap gates between adjacent qubits to bring C and D together. Notice that a swap gate does not modify χ. The update of the decomposition after an arbitrary two-qubit gate has been applied can thus be always achieved with at most $\mathcal{O}[n\mathrm{poly}(\chi)]$ basic operations. From the Γ's and λ's it is possible to compute, at a cost growing as poly(χ), the outcome probabilities of a measurement on qubit l and the new Γ's and λ's for the postselected states of the remaining $n-1$ qubits. Thus simulating the outcome probabilities of the final measurement on the computational basis requires $n\mathrm{poly}(\chi)$ basic operations. We can therefore state the main results of the Letter.

Theorem. *A pure-state quantum computation on* n *qubits and consisting of* poly(n) *elementary gates can be classically simulated with a cost in computational time and memory space given by* poly(n)$\exp(E\chi)$, *where* $E\chi$ *denotes the maximal value of* $E\chi(\Psi)$ *achieved during the computation,*

$$\begin{matrix} n \text{ qubit} \\ \text{computation} \end{matrix} \leftrightarrow \begin{matrix} \mathrm{poly}(n)\exp(E\chi) \\ \text{time and memory.} \end{matrix} \tag{28}$$

Corollary. *If* $E\chi$ *scales at most as* log(n), *then a classical simulation can be accomplished with* poly(n) *computational resources.*

Discussion.—The above results establish a clear connection between the cost of classically simulating a pure-state quantum computation and the amount of entanglement involved in the computation. The amount of entanglement determines the computational cost incurred when using the specific protocol described in this Letter. In particular, the corollary states that any *slightly* entangled quantum computation, such that $E\chi$ is upper bounded by $k\log(n)$ for some $k > 0$, can be *efficiently*

simulated in a classical computer. On the other hand, also some highly entangled quantum computations can be simulated efficiently [2–4] through other simulation schemes. In other words, a small amount of entanglement is a sufficient condition for an inexpensive classical simulation, but not a necessary one.

The above considerations translate into a lower bound on the amount of entanglement that is produced during a pure-state quantum computation leading to a computational speedup. Indeed, it follows from the theorem that if a strictly exponential speedup is to occur, then $E\chi$ must grow linearly in n. Similarly, a computational speedup of, say, $\exp(\sqrt{n})$, requires that $E\chi$ grow at least as $\sqrt{n}$, and so on.

The tools presented in this Letter can also be applied to simulate continuous-time dynamics in certain quantum many-body systems, as further discussed in a future contribution. Roughly, time is divided into small steps so as to approximate the continuous-time evolution by a sequence of small gates, and then the present simulation protocol is used to simulate the sequence of gates. As above, the computational cost of a given simulation depends on the amount of entanglement in the system. Efficient simulations correspond to a very restricted subset of states [16] and one could have expected not to find physical systems of interest where the simulation protocol is efficient. However, it turns out that the amount of entanglement in many one-dimensional many-body systems, such as quantum spin chains at zero temperature [17], is typically sufficiently small so that an efficient classical simulation is possible.

Thus, the uncovered connection between entanglement and the cost of classical simulations does not only contribute to our understanding of the essential ingredients of quantum computation, but it also stimulates the development of new techniques for the study of realistic quantum many-body phenomena.

The author thanks Dave Bacon, Ignacio Cirac, Ann Harvey, Richard Jozsa, and Debbie Leung for valuable advice. Support from the U.S. National Science Foundation under Grant No. EIA-0086038 is acknowledged.

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[2] D. Gottesman, in *Proceedings of the 23rd International Colloquium on Group Theoretical Methods in Physics* (International Press, Cambridge, 1999), p. 23.

[3] S. D. Barlett *et al.*, Phys. Rev. Lett. **89**, 207903 (2002); S. D. Barlett and B. C. Sanders, *ibid.* **89**, 097904 (2002).

[4] L. Valiant, in *Proceedings of the 33rd ACM Symposium on the Theory of Computation* (ACM Press, New York, 2001), pp. 114–123; B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **65**, 032325 (2002); E. Knill, quant-ph/0108033.

[5] R. Jozsa and N. Linden, quant-ph/0201143.

[6] E. Knill and R. Laflamme, Phys. Rev. Lett. **81**, 5672 (1998).

[7] N. Linden and S. Popescu, Phys. Rev. Lett. **87**, 047901 (2001).

[8] G. Vidal, "Efficient Classical Simulation of Barely Correlated Mixed-State Quantum Dynamics" (to be published).

[9] E. Schmidt, Math. Ann. **63**, 433 (1906); A. Ekert and P. L. Knight, Am. J. Phys. **63**, 415 (1995); A. Peres, *Quantum Theory: Concepts and Methods* (Kluwer Academic Publishers, Dordrecht, 1995).

[10] The use of χ_A [equivalently, of $\log_2(\chi_A)$] as a measure of entanglement can be justified by considering how a series of nonlocal resources can be interconverted under (nonasymptotic) local operations and classical communication (LOCC). For instance, about $\log_2(\chi_A)$ EPR pairs shared between A and B [or $\log_2(\chi_A)$ Control-NOT gates involving A and B] are necessary and sufficient to prepare $|\Psi\rangle$ with the additional help of LOCC. Also, $\log_2(\chi_A)$ is the maximal number of EPR pairs that, with finite probability, can be extracted from $|\Psi\rangle$ by LOCC. Both χ_A and its logarithm can be shown not to increase (not even probabilistically) under LOCC [11], as required of any entanglement measure, and they are related to the more popular measure *entropy of entanglement* $E(\Psi) \equiv -\mathrm{tr}(\rho_A \log_2 \rho_A)$ [12] through $E(\Psi) \leq \log_2(\chi_A)$. Notice that using $E(\Psi)$ to quantify the entanglement of $|\Psi\rangle$ in the present context may not be an appropriate choice, since $E(\Psi)$ refers to asymptotic properties of $|\Psi\rangle$, i.e., to properties of $|\Psi\rangle^{\otimes N}$ in the limit $N \to \infty$, whereas here we are concerned with the case $N = 1$.

[11] G. Vidal, J. Mod. Opt. **47**, 355 (2000).

[12] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Phys. Rev. A **53**, 2046 (1996).

[13] We also note that $E\chi(\Psi)$ is not a continuous function of $|\Psi\rangle$ with respect any reasonable distance. This implies that a good approximation $|\tilde{\Psi}\rangle$ to $|\Psi\rangle$ may exist with a significantly lower value of $E\chi$, a fact that can be exploited to improve the efficiency of the simulation protocol described in this Letter.

[14] It is also insightful to derive the explicit decomposition for simple n-qubit states with small χ, such as a product state $|0\rangle^{\otimes n}$, with $\chi = 1$; a cat state, $|0\rangle^{\otimes n} + |1\rangle^{\otimes n}$ with $\chi = 2$; and a W state $|00\cdots01\rangle + |00\cdots10\rangle + \cdots + |10\cdots00\rangle$, also with $\chi = 2$.

[15] A digital computer allows only for an approximate description of gates and states, since real coefficients are truncated. See, e.g., Ref. [5] for a discussion on how to obtain efficient approximations by using rational numbers.

[16] A pure state of n qubits picked up at random from $\mathcal{H}_2^{\otimes n}$ (according, say, to the unitarily invariant probability density) will have $E_\chi = n/2$ with certainty, as implied by a simple parameter counting argument. From this perspective, states with $E_\chi = \mathcal{O}(\log(n))$ are extremely rare.

[17] G. Vidal, J. I. Latorre, E. Rico, and A. Kitaev, Phys. Rev. Lett. **90**, 227902 (2003).