Efficient Solvability of Hamiltonians and Limits on the Power of Some Quantum Computational Models

Rolando Somma,[1] Howard Barnum,[1] Gerardo Ortiz,[1] and Emanuel Knill[2]

[1]*Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

[2]*Mathematical and Computational Sciences Division, National Institute of Standards and Technology, Boulder, Colorado 80305, USA*

(Received 23 January 2006; revised manuscript received 17 May 2006; published 9 November 2006)

One way to specify a model of quantum computing is to give a set of control Hamiltonians acting on a quantum state space whose initial state and final measurement are specified in terms of the Hamiltonians. We formalize such models and show that they can be simulated classically in a time polynomial in the dimension of the Lie algebra generated by the Hamiltonians and logarithmic in the dimension of the state space. This leads to a definition of Lie-algebraic "generalized mean-field Hamiltonians." We show that they are efficiently (*exactly*) solvable. Our results generalize the known weakness of fermionic linear optics computation and give conditions on control needed to exploit the full power of quantum computing.

DOI: 10.1103/PhysRevLett.97.190501 PACS numbers: 03.67.Lx, 03.67.Mn, 03.65.Ud, 05.30.−d

To solve a given problem, such as determining the ground state (GS) energy of a quantum system with arbitrary precision, various strategies can be considered. Each strategy is described by a sequence of elementary instructions defining an algorithm, of classical or quantum nature, and uses a particular model of computation (an abstract counterpart to a particular type of computer). Although it is widely believed that quantum computers are more powerful than classical computers, there are two situations where quantum computations have been shown to be efficiently simulatable by classical means: one is when only limited entanglement is used [1,2], and the other is when the set of gates or the quantum control used is far from universal. The latter situation includes fermionic linear optical [3,4] and Clifford group computing [5]. In general, it is desirable to determine conditions on the computational state space and the implementable operations that are needed to make available the full power of quantum computing. Such conditions can guide the choice of devices for realizing quantum computers and help to recognize quantum algorithms that are efficiently classically simulatable.

Here we consider models of quantum computing based on continuous control. The standard model of quantum computing (SQC) is equivalent to one such model. By focusing on the Lie algebra generated by the controllable Hamiltonians, we obtain the model of Lie-algebraic quantum computing (LQC). This model allows us to distinguish between the algebraic properties of the available control, which are independent of the physical realization, and the state space that is being acted on. The physical motivation for LQC is that Lie algebras are a natural way to represent restrictions on the dynamics of quantum devices that are used to implement computations. Thus they provide a theoretical tool for evaluating proposals for quantum computers based on particular physical systems and can be used to provide necessary conditions for achieving the full power of quantum computing.

Because LQC initial states are GSs of controllable Hamiltonians [6,7], LQC lets us derive a sufficient criterion for the efficient (or *exact*) solvability of Hamiltonians (determination of an arbitrarily chosen eigenvalue and eigenvector). We prove that quantum models with generalized mean-field Hamiltonians (GMFHs) [8] are efficiently solvable and do not provide a stronger-than-classical model of computing. In particular, a quantum computing device whose gates are generated by GMFHs can be efficiently classically simulated.

Ignoring precision, SQC can be specified as the model of computing whose control Hamiltonians include the Pauli matrices and products of Pauli matrices on two qubits. A computation begins with N qubits in the state $|0\cdots0\rangle$, which is the GS of $-\sum_{\mathbf{i}}\sigma_{\mathbf{i}}^{z}$, where $\sigma_{\mathbf{i}}^{z}$ acts on the $\mathbf{i}$th qubit. It continues by applying a sequence of unitary gates e^{-iH}, where H is a control Hamiltonian. The last step is a measurement, say, of σ_{1}^{z}, where it is sufficient to learn its expectation with constant precision independent of N or the length of the computation. As Feynman observed [9] for quantum physics simulations, the obvious strategy for classically simulating SQC is inefficient. Exponential resources are required to represent states of N qubits, and explicitly evolving them is costly. This inefficiency is witnessed by the exponential size of the algebra $\mathfrak{su}(2^N)$ generated by the control Hamiltonians.

In general, an LQC computation involves a Lie algebra $\hat{\mathfrak{h}}$ of control Hamiltonians acting on a Hilbert space $\mathcal{H}$, where the initial state is a GS of a control Hamiltonian and the final measurement is an expectation of either a control Hamiltonian or a unitary operator in the group generated by $\hat{\mathfrak{h}}$. We show (Theorem 3 below) that an LQC computation can be simulated efficiently in the dimension of $\hat{\mathfrak{h}}$ and the logarithm of the dimension of $\mathcal{H}$, a significant improvement over the obvious strategy whose complexity must be at least linear in the dimension of $\mathcal{H}$ due to the need for explicitly representing states. Furthermore, the

precision with which the final measurement can be simulated is also exponential, whereas a physical realization can only obtain polynomial precision.

Our results imply that the power of SQC depends on the exponential dimension of $\mathfrak{su}(2^N)$. From this point of view, the size of the Hilbert space is incidental. Nevertheless, one can consider the features of the computationally accessible states in $\mathcal{H}$ that result in efficient classical simulations of LQC computations. The initial state of an LQC computation is a "generalized coherent state" (GCS) [10]. Such states may be identified with generalized unentangled states [6,7] of the system. In fact, generalized entanglement plays a decisive role in our analysis, and our result can be viewed as confirmation that generalized entanglement is required for exploiting the advantages of quantum computation.

Important special cases of LQC are fermionic linear optical and matchgate computing, which were previously known to be efficiently classically simulatable [3,11,12]. Their Lie algebras are $\mathfrak{so}(2N)$ and $\mathfrak{so}(2N+1)$, represented on fermionic modes. The classical simulation algorithms for fermionic linear optics are based on combinatorics [3] or Wick's theorem [11]. Our simulations of LQC computations are much more general and applicable to all compact Lie algebras, including the exceptional ones, and any finite-dimensional representation or physical realization [13,14]. In particular, the classical Lie algebras $\mathfrak{su}(D)$, $\mathfrak{so}(D)$, $\mathfrak{sp}(D)$, and their sums can be simulated. For example, consider ensemble quantum computers [15] consisting of a large number K of D-dimensional subsystems such as the nuclear spins of identical molecules. The control Hamiltonians and weakly measurable operators are in $\mathfrak{su}(D)$, expressible as sums of identical operators acting on each subsystem. In view of the fact that state preparation can often be used to turn a computationally weak into a powerful system [16], one might hope that by starting with a special not necessarily symmetric state of the ensemble, one could greatly enhance the power of the system. Our results imply that if this state is a GS of a typical control Hamiltonian (one without GS degeneracies in irreducible representations), the power can increase at most polynomially in K. Note that the preparation of other states requires interactions involving operators not in $\mathfrak{su}(D)$.

The efficient classical simulatability of LQC computations is related to the Gottesman-Knill theorem [5,17], according to which Clifford group computations are efficiently classically simulatable, and similar results on bosonic linear optical quantum computing with homodyne measurements [4,18]. Like LQC, both models of computing are defined by the set of allowed operations and compatible measurements acting on a special initial state. In both cases, the proofs of simulatability involved showing that the set of accessible states is small in either number or dimension. Our results show that such results may depend less on the particular state space used than on the algebraic properties of the allowed operations.

An algorithm is efficient if the resources required to solve problem instances of size, or specification complexity (SC), N are polynomial in N [poly(N)]. The relevant resources are time and space. We want to determine the output of a given LQC computation using an efficient classical algorithm. The problem specification includes the information needed to "run" the LQC computation. The SC includes the size in bits of the information and the expected resources required. The latter is included to enable the comparison of models of computing. The specification of an LQC computation contains a description of a Lie algebra $\hat{\mathfrak{h}}$, the state space it acts on, the initial state, the sequence of evolutions to be performed, and the final measurement and its precision. In the following we use basic results of Lie theory; see [13,19,20]. We assume $\hat{\mathfrak{h}}$ is an M-dimensional Lie algebra of skew-Hermitian operators acting on a finite-dimensional Hilbert space $\mathcal{H}$, with Lie bracket $[\hat{X}, \hat{Y}] := \hat{X}\hat{Y} - \hat{Y}\hat{X}$. Without loss of generality, the action is irreducible. The control Hamiltonians are operators in $\sqrt{-1}\hat{\mathfrak{h}} \equiv i\hat{\mathfrak{h}}$. The initial state is the unique GS of a given $\hat{L} \in i\hat{\mathfrak{h}}$, written as $|\mathrm{lw}\rangle$ (lw stands for "lowest weight"). Gates are of the form $e^{\hat{X}}$ for $\hat{X} \in \hat{\mathfrak{h}}$. Final measurements yield $\langle \hat{W} \rangle$ for $\hat{W} \in i\hat{\mathfrak{h}}$ or $|\langle \hat{W} \rangle|$ for $\hat{W} \in e^{\hat{\mathfrak{h}}}$ with specified precision.

To specify $\hat{\mathfrak{h}}$, $\mathcal{H}$, and $|\mathrm{lw}\rangle$ we use a Cartan-Weyl (CW) basis [13]. Such a basis consists of l *lowering* (*raising*) operators $\hat{e}^-_{\alpha_j}$ ($\hat{e}^+_{\alpha_j} \equiv \hat{e}^{-\dagger}_{\alpha_j}$) of $\mathbb{C}\hat{\mathfrak{h}}$ (the complex linear combinations of operators in $\hat{\mathfrak{h}}$) and r generators $\hat{h}_k$ of a Cartan subalgebra (CSA, maximal set of commuting Hermitian operators) of $i\hat{\mathfrak{h}}$ with *roots* $\alpha_j = (\alpha_j^1, \ldots, \alpha_j^r) \in \mathbb{R}^r$. The dimension of $\hat{\mathfrak{h}}$ satisfies $M = 2l + r$. We can assume that $\hat{L}$ is in the CSA and has the property that $[\hat{L}, e^+_{\alpha_j}] = c_j e^+_{\alpha_j}$ with $c_j > 0$. The *weights* of $|\mathrm{lw}\rangle$ are defined by $\hat{h}_k|\mathrm{lw}\rangle = w(\hat{h}_k)|\mathrm{lw}\rangle$ and are integral. The lowering operators annihilate $|\mathrm{lw}\rangle$, and the raising operators map $|\mathrm{lw}\rangle$ to eigenstates of $\hat{L}$ with higher eigenvalues. According to the basic theory of Lie algebras, $\hat{\mathfrak{h}}$, $\mathcal{H}$, and $|\mathrm{lw}\rangle$ are completely determined by the abstract CW basis, the structure constants (which express the commutators of CW basis elements in terms of the CW basis), and the weights. Furthermore, by appropriate choice of basis, the structure constants and weights have bit complexity polynomial in M and $\log \max_k(|w(\hat{h}_k)|)$. The latter is bounded above by a constant plus $\log(d)$, with d the dimension of $\mathcal{H}$. The SC of $\mathfrak{h}$ and $|\mathrm{lw}\rangle$ is the total bit complexity of the structure constants and weights. One can also consider the situation where $\mathfrak{h}$ is given in terms of an arbitrary basis and its structure constants. Since there are algorithms for finding a suitable CW basis, we do not explicitly consider this case.

The basic gates of an LQC computation are unitaries $e^{t\hat{X}}$, with $\hat{X} = i\hat{h}_k$ or $\hat{X} = \sqrt{\mp 1}(\hat{e}^+_{\alpha_j} \pm \hat{e}^-_{\alpha_j})$. The SC of a gate is the number of bits b necessary to represent t and the choice of gate. The time resource required by a gate is $|t|$, which we add to the SC. More generally, gates can be any $e^{\hat{H}}$ with

$\hat{H} \in \hat{\mathfrak{h}}$. The SC of such a gate is that of $\hat{H}$. As for any operator in the Lie algebra, this is the number of bits used to express its coefficients in the CW basis. The time resource required by the gate is the sum of the absolute values of the coefficients of $\hat{H}$ in the CW basis, and this is also added to the SC of the generalized gate. An LQC computation includes a sequence of gates; its SC is the sum of the SCs of the gates.

The final step of an LQC computation is the measurement. If the measurement is of $\hat{A} \in \mathbb{C}\hat{\mathfrak{h}}$, the SC is that of $\hat{A}$. If it is of $\hat{V} = e^{\hat{A}} \in e^{\mathbb{C}\hat{\mathfrak{h}}}$, the SC is the same as that of $\hat{V}$ used as a gate. We do not include the desired precision ϵ in specification of the computation, but make it a separate part of the problem. A physical realization of an LQC computation normally requires a resource overhead of the order of $1/\sqrt{\epsilon}$ because the desired precision is obtained by independent repetitions of measurements of an observable. Our simulations achieve b bits of precision with a poly(b) overhead, with $b = 1/\epsilon$. Note that SQC augmented with the hypothetical ability to determine expectation values with b bits of precision using poly(b) resources can efficiently solve problems in $\#P$, the class associated with counting the number of solutions to NP-complete problems [21]. This implies that our simulations of LQC computations are more powerful than their direct quantum physical realizations.

The first problem encountered when trying to simulate an LQC computation is that d may be exponentially large. The key idea behind our results is that when $M = \dim(\mathfrak{h})$ is *small*, small dimensional faithful representations of $\hat{\mathfrak{h}}$ can be used for the purposes of simulation. We use the *adjoint* representation by $(M \times M)$ matrices. To distinguish the different versions of $\hat{\mathfrak{h}}$ we use $\mathfrak{h}$ for the abstract Lie algebra represented by $\hat{\mathfrak{h}}$ in $\mathcal{H}$ and $\bar{\mathfrak{h}}$ for its adjoint representation. Objects related to the Lie algebra are similarly distinguished, so that for $A \in \mathfrak{h}$, $\hat{A}$ ($\bar{A}$) is a member of $\hat{\mathfrak{h}}$ ($\bar{\mathfrak{h}}$).

The main problem is then to relate the effects of a simulation in the adjoint representation to the desired effects in $\mathcal{H}$; representation theory enables us to do this efficiently by calculating in the weight space. Our first two results show how to do this for computing the expectation of an operator in $\mathbb{C}\hat{\mathfrak{h}}$ with respect to a specified mixture of GCSs. For details of proofs not given here see [22,23].

Theorem 1 Let $\rho = \sum_s p_s e^{\hat{A}_s} |\mathrm{lw}\rangle\langle \mathrm{lw}| e^{-\hat{A}_s}$ with $\hat{A}_s \in \hat{\mathfrak{h}}$ and $\hat{W} \in \mathbb{C}\hat{\mathfrak{h}}$. Then $\langle \hat{W} \rangle = \mathrm{Tr}(\hat{W}\rho)$ can be classically computed to precision ϵ in time polynomial in $\log(1/\epsilon)$ and the sum of the SCs of $\mathfrak{h}$, $|\mathrm{lw}\rangle$, W, A_s and p_s.

Proof (outline).—Define $\hat{W}_s = e^{-\hat{A}_s} \hat{W} e^{\hat{A}_s}$. Using the adjoint representation, expand it in the CW basis. Since $\hat{e}^-_{\alpha_j} |\mathrm{lw}\rangle = 0$, we have $\langle \hat{W} \rangle = \sum_s p_s \sum_{k=1}^r u_k^s w(\hat{h}_k)$, with $u_k^s \in \mathbb{C}$ given by the matrix projection of $\bar{W}_s$ onto $\bar{h}_k$. □

The next result determines the expectation of $e^{\hat{H}}$ and may be viewed as an algebraic analogue of the quantum field-theoretic problem of computing "vacuum-to-vacuum transition probabilities," where $|\mathrm{lw}\rangle$ is the vacuum.

Theorem 2 Let ρ be as in Theorem 1 and $\hat{W} = e^{\hat{H}}$, with $\hat{H} \in \mathbb{C}\hat{\mathfrak{h}}$. Then $\langle \hat{W} \rangle = \mathrm{Tr}(\hat{W}\rho)$ can be classically computed to precision ϵ in time polynomial in $\log(1/\epsilon)$ and the sum of the SCs of $\mathfrak{h}$, $|\mathrm{lw}\rangle$, W, A_s, and p_s.

Proof (outline).—Define $\hat{\Pi} = |\mathrm{lw}\rangle\langle \mathrm{lw}|$, $\hat{W}_s = e^{-\hat{A}_s} \hat{W} e^{\hat{A}_s}$, and $\hat{O}_{s,s'} = \hat{\Pi} \hat{W}_s \hat{\Pi} \hat{W}_{s'}^\dagger \hat{\Pi}$. Then $|\langle \hat{W} \rangle|^2 = \sum_{s,s'} p_s p_{s'} \mathrm{Tr} \hat{O}_{s,s'}$. The operator $\hat{O}_{s,s'}$ is proportional to $\hat{\Pi}$ and its trace is the constant of proportionality. We can express $\hat{\Pi}$ as a limit of operators in $e^{\mathbb{C}\hat{\mathfrak{h}}}$. Redefine $\hat{L} = \sum_{k=1}^r w(\hat{h}_k) \hat{h}_k$ and define ω by $\hat{L}|\mathrm{lw}\rangle = \omega |\mathrm{lw}\rangle$. Then $\langle \psi | \hat{L} | \psi \rangle > \omega$ for $|\psi\rangle \neq |\mathrm{lw}\rangle$, so $\hat{\Pi} = \lim_{t \to \infty} e^{t\omega} e^{-t\hat{L}}$ [6]. Because the eigenvalues of $\hat{L}$ are integral, convergence is exponentially fast in t. Let $\hat{E}(t) = \sum_{s,s'} p_s p_{s'} e^{3\omega t} \times e^{-t\hat{L}} \hat{W}_s e^{-t\hat{L}} \hat{W}_{s'}^\dagger e^{-t\hat{L}}$, which converges to $\sum_{s,s'} p_s p_{s'} \hat{O}_{s,s'}$ as $t \to \infty$. For a given t, we can compute $\bar{E}(t)$ by computing exponentials and multiplying matrices in the adjoint representation. Observe that the maximum eigenvalue $\kappa(t)$ of $\hat{E}(t)$ converges exponentially fast to $|\langle \hat{W} \rangle|^2$. To compute $\kappa(t)$ we first determine $\bar{Q}(t) \in i\bar{\mathfrak{h}}$ such that $\bar{E}(t) = e^{\bar{Q}(t)}$ and $\hat{E}(t) = e^{\hat{Q}(t)}$. The maximum eigenvalue $q(t)$ of $\hat{Q}(t)$ can be obtained by unitary transformation to the CSA via an efficient diagonalization procedure for $\bar{Q}(t)$ (e.g., [24]) and analyzing the structure of roots and their relationship to the representation $\hat{\mathfrak{h}}$. We then obtain $\kappa(t) = e^{q(t)}$. The necessary classical computations can be realized efficiently in the SCs and the number of digits of precision of $\kappa(t)$. □

Theorem 3 The output of the final measurement of an LQC computation can be computed classically with precision ϵ in time polynomial in its SC and $\log(1/\epsilon)$.

Proof.— Let the sequence of gates of the LQC computation be $e^{\hat{A}_1}, \ldots, e^{\hat{A}_t}$, with $\hat{A}_m \in \hat{\mathfrak{h}}$. The final state before the measurement is $|\phi\rangle = \hat{U}|\mathrm{lw}\rangle$ with $\hat{U} = \prod_{m=1}^t e^{\hat{A}_m}$. Let $\hat{W}$ be the operator whose expectation is measured. Then $\langle \hat{W} \rangle = \langle \mathrm{lw} | \hat{U}^\dagger \hat{W} \hat{U} | \mathrm{lw} \rangle$, and the desired output is either $\langle \hat{W} \rangle$ if $\hat{W} \in \mathbb{C}\hat{\mathfrak{h}}$, or $|\langle \hat{W} \rangle|$ if $\hat{W} \in e^{\mathbb{C}\hat{\mathfrak{h}}}$. The algorithms in the proofs of Theorems 1 and 2 can be used. It suffices to compute $\bar{U}$ in the adjoint representation using matrix exponentiation and multiplication and use $\bar{U}$ in place of $e^{\bar{A}_s}$. □

An application of Theorem 3 is to representations of $\mathfrak{so}(2N+1)$ of dimension $M = N(2N+1)$. One is generated by $\{c_\mathbf{i}^\dagger, c_\mathbf{i}, c_\mathbf{i}^\dagger c_\mathbf{j}, c_\mathbf{i}^\dagger c_\mathbf{j}^\dagger, c_\mathbf{i} c_\mathbf{j}\}$, where $c_\mathbf{i}$ and $c_\mathbf{i}^\dagger$ are the annihilation and creation operators for spinless fermions on N modes. The dimension of $\mathcal{H}$ is $d = 2^N$, which is exponential in M. Other interesting representations for which classical simulations are not obvious are obtained via Jordan-Wigner mappings [14]. For example, consider a computation that applies gates that are exponentials of $\hat{H}_I = \sum_{\mathbf{i}=1}^N (g_x \sigma_\mathbf{i}^x \sigma_{\mathbf{i}+1}^x + g_y \sigma_\mathbf{i}^y \sigma_{\mathbf{i}+1}^y + \sigma_\mathbf{i}^z) + b\sigma_1^x \in \mathfrak{so}(2N+1)$ to qubits initially in the state $|0 \cdots 0\rangle$, and then measures the expectation of $\hat{H}_I$ or its exponential. This expectation can be computed classically in time polynomial in N and the number of digits of precision.

The algorithms above can be used to analyze certain physical models. We use the term GMFH [8] for Hamiltonians $\hat{H}_{\mathrm{MF}} \in i\hat{\mathfrak{h}}$, where $\hat{\mathfrak{h}}$ is any one of a family of M-dimensional Lie algebras acting on d-dimensional Hilbert spaces with M at most polylog(d). A GMFH must be specified in terms of a basis of $\mathfrak{h}$ that can be efficiently transformed to a CW basis. $\hat{H}_I$ above is an example. We say a Hamiltonian $\hat{H}$ acting on $\mathcal{H}$ can be efficiently (*exactly*) solved when an arbitrarily chosen eigenvalue of $\hat{H}$, and an appropriate description of the corresponding eigenstate, can be obtained and represented to precision ϵ by means of a classical algorithm efficient in $\log(d)$ and $1/\epsilon$. This definition, motivated by complexity theory, yields a sufficient criterion for exact solvability.

Theorem 4 GMFHs can be efficiently (exactly) solved.

Proof.—Let $\hat{H}_{\mathrm{MF}}$ be given in a CW basis. We show that to solve $\hat{H}_{\mathrm{MF}}$ it suffices to diagonalize it according to $\hat{H}_D = \hat{U}\hat{H}_{\mathrm{MF}}\hat{U}^\dagger = \sum_{k=1}^{r} \varepsilon_k \hat{h}_k$, with $\varepsilon_k \in \mathbb{R}$ and $\hat{U} \in e^{\hat{\mathfrak{h}}}$ unitary. The eigenvalues of $\hat{H}_{\mathrm{MF}}$ are those of $\hat{H}_D$, and its eigenspaces are those of $\hat{H}_D$ transformed by $\hat{U}^\dagger$, described by a sequence of LQC gates. The eigenspaces of $\hat{H}_D$ are those of $\hat{h}_k$ (weight states), obtained from $|\mathrm{lw}\rangle$ by applying raising operators: $\prod_{j=1}^{l} (\hat{e}^+_{\alpha_j})^{n_j} |\mathrm{lw}\rangle$. The eigenvalues are $\lambda(\hat{H}_D) = \sum_k \varepsilon_k \lambda(\hat{h}_k)$, with $\lambda(\hat{h}_k) = w(\hat{h}_k) - \sum_j n_j \alpha_j^k$ defining the corresponding weights, and $n_j \geq 0$ integers. To efficiently diagonalize $\hat{H}_{\mathrm{MF}}$ and obtain $\hat{U}$, we apply a generalization of the Jacobi method [24,25] to $\hat{H}_{\mathrm{MF}}$. It yields an exponentially converging diagonalization and an expression for $\hat{U}$ as a product of unitaries $\hat{U}_j(x) = e^{(x\hat{e}^+_{\alpha_j} - x^* \hat{e}^-_{\alpha_j})} \in e^{\hat{\mathfrak{h}}}$. □

Theorem 4 encompasses the solution of $\hat{H}_I$ via a nonlinear Bogoliubov transformation [26] and of other GMFHs solved by standard linear ones, but it is much more general since it is applicable to any representation of any compact Lie algebra, for most of which no Bogoliubov transformation is known.

Corollary 1 Let $|\phi\rangle$ be the GS of a GMFH. Then $|\phi\rangle$ can be prepared efficiently on a quantum computer if $|\mathrm{lw}\rangle$ can.

Proof.—According to Theorem 4, the GS of $\hat{H}_{\mathrm{MF}}$ (to precision ϵ) can be obtained by applying polylog(d) + poly($1/\epsilon$) gates of the form $\hat{U}_j(x)$ to the state $|\mathrm{lw}\rangle$. □

Our results cast light on why quantum computation may be more powerful than classical. In general, the poly(N) generators of the gates of a quantum computation can generate an exponential-dimensional Lie algebra acting on an exponentially large space, as in SQC, where single- and two-qubit interactions generate the $(2^{2N} - 1)$-dimensional algebra $\mathfrak{su}(2^N)$. By contrast, when the gate generators induce a poly(N)-dimensional Lie algebra, a computation with such gates and compatible state preparations and measurements can be efficiently classically simulated by working in a low-dimensional faithful representation of the algebra. Is LQC with von Neumann measurements and feedforward control efficiently classically simulatable? Can a general Lie-theoretic approach unify our results with ones on Clifford-like groups (using Lie algebras over finite fields) and bosonic computation (using nonsemisimple Lie algebras)? What other algebraically constrained models of quantum computation are efficiently classically simulatable? Such structures may underlie the efficient solvability of classes of many-body Hamiltonians beyond GMFHs, such as those solvable via a Bethe ansatz.

We thank L. Viola and L. Gurvits for discussions and pointing out [25], and the U.S. DOE and NSA for support.

[1] R. Jozsa and N. Linden, Proc. R. Soc. A **459**, 2011 (2003).
[2] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003).
[3] L. G. Valiant, in *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computation (STOC'01)* (ACM Press, El Paso, TX, 2001), pp. 114–123.
[4] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, Phys. Rev. Lett. **88**, 097904 (2002).
[5] D. Gottesman, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 1997.
[6] H. Barnum, E. Knill, G. Ortiz, and L. Viola, Phys. Rev. A **68**, 032308 (2003).
[7] H. Barnum, E. Knill, G. Ortiz, R. Somma, and L. Viola, Phys. Rev. Lett. **92**, 107902 (2004).
[8] R. Somma, G. Ortiz, H. Barnum, E. Knill, and L. Viola, Phys. Rev. A **70**, 042311 (2004).
[9] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[10] A. Perelomov, Commun. Math. Phys. **26**, 222 (1972).
[11] B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **65**, 032325 (2002).
[12] E. Knill, Los Alamos National Laboratory Technical Report No. LAUR-01-4472, 2001.
[13] J. F. Cornwell, *Group Theory in Physics* (Academic Press, London, 1989).
[14] C. Batista and G. Ortiz, Adv. Phys. **53**, 1 (2004).
[15] D. G. Cory, A. F. Fahmy, and T. F. Havel, Proc. Natl. Acad. Sci. U.S.A. **94**, 1634 (1997).
[16] E. Knill, R. Laflamme, and W. Zurek, Proc. R. Soc. A **454**, 365 (1998).
[17] S. Aaronson and D. Gottesman, Phys. Rev. A **70**, 052328 (2004).
[18] S. D. Bartlett and B. C. Sanders, Phys. Rev. Lett. **89**, 207903 (2002).
[19] J. Fuchs, *Affine Lie Algebras and Quantum Groups* (Cambridge University Press, Cambridge, 1992).
[20] A. Knapp, *Lie Groups: Beyond an Introduction* (Birkhäuser, Boston, 1996).
[21] S. Fenner, F. Green, S. Homer, and R. Pruim, Proc. R. Soc. A **455**, 3953 (1999).
[22] R. Somma, H. Barnum, G. Ortiz, and E. Knill, quant-ph/0601030.
[23] R. D. Somma, Ph.D. thesis, Instituto Balseiro, Argentina, 2005, quant-ph/0512209.
[24] N. J. Wildberger, Proc. Am. Math. Soc. **119**, 649 (1993).
[25] M. Kleinsteuber, U. Helmke, and K. Hüper, SIAM J. Matrix Anal. Appl. **26**, 42 (2004).
[26] H. Fukutome, M. Yamamura, and S. Nishiyama, Prog. Theor. Phys. **57**, 1554 (1977).