

Quantum Computational Complexity of the N -Representability Problem: QMA Complete

Yi-Kai Liu,¹ Matthias Christandl,² and F. Verstraete^{3,4}

¹*Computer Science and Engineering, University of California, San Diego, California, USA*

²*Centre for Quantum Computation, Centre for Mathematical Sciences, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

³*Facultät für Physik, Universität Wien, Boltzmanngasse 5, A-1090 Wien, Austria*

⁴*Institute for Quantum Information, Caltech, Pasadena, California, USA*

(Received 17 September 2006; published 16 March 2007)

We study the computational complexity of the N -representability problem in quantum chemistry. We show that this problem is quantum Merlin-Arthur complete, which is the quantum generalization of nondeterministic polynomial time complete. Our proof uses a simple mapping from spin systems to fermionic systems, as well as a convex optimization technique that reduces the problem of finding ground states to N representability.

DOI: [10.1103/PhysRevLett.98.110503](https://doi.org/10.1103/PhysRevLett.98.110503)

PACS numbers: 03.67.Lx, 31.25.-v, 89.70.+c

The central theoretical problem in the field of many-body strongly correlated quantum systems is to find efficient ways of simulating Schrödinger's equations. The main difficulty is the fact that the dimension of the Hilbert space describing a system of N quantum particles scales exponentially in N . This makes a direct numerical simulation intractable: every time an extra particle is added to the system, the computational resources would have to be doubled.

The situation is not hopeless, however, as in principle it could be that all physical wave functions, i.e., the ones that are realized in nature, have very special properties and can be parameterized in an efficient way. The idea would then be to propose a variational class of wave functions that capture the physics of the systems of interest, and then do an optimization over this restricted class. This approach has proven to be very successful, as witnessed by mean field theory and renormalization group methods. However, it is still an open problem to find an efficient variational class to describe complex wave functions such as those arising in quantum chemistry.

One of the basic problems in quantum chemistry is to find the ground state of a Hamiltonian describing the many-body system of an atom or molecule. These Hamiltonians are very ungeneric, because they contain at most 2-body interactions. This implies that the number of free parameters in such Hamiltonians scales at most quadratically in the number of particles or modes, and hence the ground states of all such systems form a small-dimensional manifold.

For a Hamiltonian with only 2-body interactions, the energy corresponding to a wave function is completely determined by its 2-body correlation functions, and as a consequence the ground state will be the one with extremal 2-body reduced density operators. This fact was realized a long time ago, and led Coulson [1] to propose the following problem: given a set of N quantum particles, can we characterize the allowed sets of 2-body correlations or density operators between all pairs of particles?

If the particles under consideration are fermions, this has been called the N -representability problem [2]. Here, we consider the reduced density operators acting on pairs of fermions, and we want to decide whether they are consistent with some global state over N fermions. An efficient solution to the N -representability problem would be a huge breakthrough, as it would (for example) allow us to calculate the binding energies of all molecules. Therefore, a very large effort has been devoted to solving this problem [3–5].

Here we will give strong evidence that the N -representability problem is intractable, as it is QMA (quantum Merlin-Arthur) complete and hence NP hard. By “intractable,” we mean that, for large N , solving the problem in the worst case requires a number of operations that grows exponentially in N . The complexity class QMA is the natural generalization of the class NP (nondeterministic polynomial time) to the setting of quantum computing. Colloquially, a problem is in QMA if there exists an efficient quantum algorithm that, when given a possible solution to the problem, can verify whether it is correct; here the “solution” may be a quantum state on polynomially many qubits. A problem is QMA hard if it is at least as hard as any other problem in QMA; that is, given an efficient algorithm for this problem, one could solve every other problem in QMA efficiently. We say that a problem is QMA complete if it is in QMA and it is also QMA hard.

In a seminal work, Kitaev [6] proved that the local Hamiltonian problem—determining the ground state energy of a spin Hamiltonian that is a sum of 5-body terms (on n qubits), with accuracy $\pm\epsilon$ where ϵ is inverse polynomial in n —is QMA complete. In fact, it was later shown that this problem remains QMA complete when restricted to 2-body interactions [7], and even in the case of geometrically local interactions [8].

Another problem is to decide whether a given set of local density operators is consistent, i.e., whether they can be realized as the reduced density operators of the same global state. This is in some sense the dual of the local Hamiltonian problem, and it was recently shown to be

QMA complete [9]. In the present Letter, we will use very similar techniques to prove that N representability, which is the fermionic version of the consistency problem, is also QMA complete.

Let us now formulate the problem. In a molecule, the nuclei are fairly well localized and can be treated as classical degrees of freedom, and the wave function of the N electrons can be approximated as a linear combination of tensor products of the d single-particle modes in the system (those form the basis set). As electrons are fermionic, the complete wave function must be antisymmetric, and this is most easily taken into account by working in the formalism of second quantization:

$$|\psi\rangle = \sum_{\substack{i_1, \dots, i_d \in \{0,1\} \\ i_1 + \dots + i_d = N}} c_{i_1, \dots, i_d} (a_1^\dagger)^{i_1} \dots (a_d^\dagger)^{i_d} |\Omega\rangle.$$

Here a_j^\dagger is the creation operator for the j th mode, and $|\Omega\rangle$ represents the vacuum state without fermions. The operators obey the following anticommutation relations: $\{a_i, a_j\} = 0 = \{a_i^\dagger, a_j^\dagger\}$ and $\{a_i, a_j^\dagger\} = \delta_{ij}$. Note that we restrict ourselves to the subspace of states with exactly N fermions. d denotes the number of modes, which is typically much larger than N . The number of degrees of freedom is $\binom{d}{N}$, which grows exponentially in N when $d \geq cN$ for some constant $c > 1$.

In the case of quantum chemistry, the Hamiltonian typically contains only one- and two-body interactions between all modes, and so it can be written as a linear combination of terms of the form $a_i^\dagger a_j$ and $a_i^\dagger a_j^\dagger a_l a_k$.

The 2-fermion reduced density matrix (2-RDM) is calculated by tracing out all but two of the fermions:

$$\rho^{(2)} = \text{Tr}_{3, \dots, N} \rho^{(N)},$$

where $\rho^{(N)}$ is a mixture of states with exactly N fermions. The matrix elements of the 2-RDM are given by

$$\rho_{ijkl}^{(2)} = \frac{1}{N(N-1)} \langle a_k^\dagger a_l^\dagger a_j a_i \rangle.$$

The N -representability problem can now be stated as follows: Consider a system of N fermions and d modes, $d \leq \text{poly}(N)$. (For purposes of complexity, we consider N to be the “size” of the problem.) We are given a 2-fermion density matrix ρ , of size $\frac{d(d-1)}{2} \times \frac{d(d-1)}{2}$, and a real number $\beta \geq 1/\text{poly}(N)$. All numbers are specified with $\text{poly}(N)$ bits of precision. The problem is to distinguish between the following two cases:

(i) There exists an N -fermion state σ such that $\text{Tr}_{3, \dots, N}(\sigma) = \rho$. In this case, answer “YES.”

(ii) For all N -fermion states σ , $\|\text{Tr}_{3, \dots, N}(\sigma) - \rho\|_1 \geq \beta$. In this case, answer “NO.”

If neither of these cases applies, then one may answer either “YES” or “NO.” Note that we allow an error tolerance of $\beta \geq 1/\text{poly}(N)$, measured using the trace distance $\|A\|_1 = \text{Tr}|A|$.

We will show that N representability is QMA complete. The proof consists of two parts. First, we show that any 2-local Hamiltonian of spins can be simulated using a 2-local Hamiltonian of fermions with $d = 2N$. Using techniques of convex programming, we show that an efficient algorithm for N representability would allow us to estimate the ground state energies of 2-local Hamiltonians; thus, N representability is QMA hard.

Second, we show that N representability is in QMA; specifically, we construct a quantum verifier that can check whether a 2-particle state is N representable, given a suitable witness.

Let us first show how to map a 2-local Hamiltonian, H_{qubit} , defined on a system of N qubits, to a 2-local Hamiltonian on fermions, H_{Fermi} , with $d = 2N$ modes, such that the ground state energy remains the same; this is the opposite of what has been done in [10]. We represent each qubit i as a single fermion that can be in two different modes a_i, b_i ; so each N -qubit basis state corresponds to the following N -fermion state:

$$|z_1\rangle \otimes \dots \otimes |z_N\rangle \mapsto (a_1^\dagger)^{1-z_1} (b_1^\dagger)^{z_1} \dots (a_N^\dagger)^{1-z_N} (b_N^\dagger)^{z_N} |\Omega\rangle.$$

Also, all the relevant single-qubit operators should correspond to bilinear functions of the creation and annihilation operators (this construction guarantees that operators on different qubits commute). This can be achieved by the following mapping:

$$\begin{aligned} \sigma_i^x &\rightarrow a_i^\dagger b_i + b_i^\dagger a_i; & \sigma_i^y &\rightarrow i(b_i^\dagger a_i - a_i^\dagger b_i); \\ \sigma_i^z &\rightarrow \mathbb{1} - 2b_i^\dagger b_i. \end{aligned}$$

H_{Fermi} is now obtained by rewriting these Pauli operators with respect to the above mapping and will only contain terms with at most 2 annihilation and 2 creation operators. The only thing that is left to do is to guarantee that there is exactly one fermion on every site i . This can be achieved by adding to H_{Fermi} terms of the following form: $P_i = \mathbb{1} + (2a_i^\dagger a_i - 1)(2b_i^\dagger b_i - \mathbb{1})$. All the P_i are biquadratic and commute with all the operators introduced earlier, and hence the complete Hamiltonian will be block diagonal. By making the weights of those terms large enough (a constant times the norm of the Hamiltonian, which is at most polynomial in N), we can always guarantee that the ground state of the full Hamiltonian will have exactly one fermion per site.

Let us now assume that we have an efficient algorithm for N representability. We claim that this allows us to efficiently determine the ground state energy of any 2-local Hamiltonian on qubits, a problem which is known to be QMA hard [7]. We start by transforming the Hamiltonian H_{qubit} into H_{Fermi} , as described above. We then construct a convex program that finds a 2-fermion density matrix ρ that is N representable, and that minimizes the expectation value of H_{Fermi} . Note that this program has polynomially many variables, the set of N -representable states is convex, and $\langle H_{\text{Fermi}} \rangle$ is a linear function of ρ . Assuming that we

have an efficient algorithm for N representability, we can solve this convex program in polynomial time; this tells us the ground state energy of H_{Fermi} , and hence of H_{qubit} .

Note that the interesting behavior in H_{Fermi} occurs in the subspace of states with exactly N particles. Restricting ourselves to this subspace, we have the identity $a_i^\dagger a_j = \frac{1}{N-1} a_i^\dagger (\sum_k a_k^\dagger a_k) a_j$, and we can write all the terms in H_{Fermi} in the form $a_i^\dagger a_j^\dagger a_l a_k$.

Convex optimization algorithms usually require that the set K of feasible solutions be full dimensional; i.e., K cannot lie in a lower-dimensional subspace. So we have to represent the state ρ in such a way that there are no redundant variables. Let \mathcal{S} be a complete set of 2-particle observables [11], and let $\ell = |\mathcal{S}|$; note that $\ell \leq \text{poly}(d) \leq \text{poly}(N)$. We represent ρ in terms of its expectation values $\alpha_S = \text{Tr}(S\rho)$ for all $S \in \mathcal{S}$; let $\vec{\alpha} \in \mathbb{R}^\ell$ denote the vector of these expectation values. We define K to be the set of all $\vec{\alpha}$ such that the corresponding state ρ is N representable. The N -representability algorithm lets us test whether a given point $\vec{\alpha}$ is in K . We write our Hamiltonian in the form $H = \sum_{S \in \mathcal{S}} \gamma_S S$ (plus a constant term). It is easy to see that $\text{Tr}(H\rho) = \sum_{S \in \mathcal{S}} \gamma_S \alpha_S$. Our convex program is as follows: find some $\vec{\alpha} \in K$ that minimizes $f(\vec{\alpha}) = \sum_{S \in \mathcal{S}} \gamma_S \alpha_S$.

We can solve this convex program in polynomial time, using the shallow-cut ellipsoid algorithm [12]. The algorithm requires a guarantee that K is contained in a ball of radius R centered at 0, and K contains a ball of radius r centered at some point p . [The running time of the algorithm grows polynomially in $\log(R/r)$.] In our case, we can set $R = \sqrt{\ell}$, and $r = 1/\text{poly}(\ell)$ [13] if we choose to center the ball around the maximally mixed state. There is also the issue of numerical precision. We only assumed the existence of an algorithm for N representability having precision $1/\text{poly}(N)$; the algorithm may give incorrect answers near the boundary of K . However, the ellipsoid algorithm still works in this setting [12], and this is sufficient to estimate the ground state energy with precision $1/\text{poly}(N)$, and thus solve the local Hamiltonian problem. This completes the proof that N representability is QMA hard. In addition, we have shown that estimating the ground state energy for local fermionic Hamiltonians is QMA hard. \square

Next, we show that N representability is in QMA. That is, we construct a poly-time quantum verifier V that takes two inputs: a description of the problem (that is, ρ and β); and a “witness” τ , which is a quantum state on polynomially many qubits. The verifier V should have the following property: if ρ is N representable, there exists a witness τ that causes V to output “true” with probability $\geq p_1$; if ρ is not N representable (within error tolerance β), then for all possible states τ , V outputs true with probability $\leq p_0$; and $p_1 - p_0 \geq 1/\text{poly}(N)$.

The idea is that, when ρ is N representable, the correct witness τ consists of (multiple copies of) an N -fermion state σ that satisfies $\text{Tr}_{3,\dots,N}(\sigma) = \rho$. Then the verifier can use quantum state tomography to compare σ and ρ .

We represent the N -fermion state σ using d qubits, via the following mapping: $(a_1^\dagger)^{i_1} \dots (a_d^\dagger)^{i_d} |\Omega\rangle \leftrightarrow |i_1\rangle \otimes \dots \otimes |i_d\rangle$. Call the resulting qubit state $\tilde{\sigma}$. We use the Jordan-Wigner transform to map the fermionic annihilation operators to qubit operators: $a_i \leftrightarrow A_i \equiv -(\otimes_{k<i} \sigma_k^z) \otimes |0\rangle\langle 1|_i$. Thus, an observable $O = a_i^\dagger a_j^\dagger a_l a_k + a_k^\dagger a_l^\dagger a_j a_i$ is transformed into a qubit operator, and its expectation value can be estimated efficiently using polynomially many copies of the state $\tilde{\sigma}$. See [13] for details.

We now describe the verifier V . The witness τ consists of several (i.e., polynomially many) blocks, where each block has d qubits, supposedly representing one copy of the state $\tilde{\sigma}$. On each block, V measures the observable $\sum_k |1\rangle\langle 1|_k$, and if the outcome does not equal N , V outputs “false.” This projects each block onto the space of N -particle states. Next, V performs measurements on each block, to estimate the expectation values of $\tilde{\sigma}$, for a suitable set of observables. If these match the expectation values of ρ , then V outputs “true,” and otherwise, V outputs “false.” One problem arises: when ρ is not N representable, the prover could try to cheat by entangling the different blocks of qubits. One can show that this does not fool the verifier, using a Markov argument, as was done in [14]. This suffices to show that N representability is in QMA. \square

What can be said about the complexity of the pure-state N -representability problem, where one has to decide whether the reduced density operators arise from a pure-state of N fermions? In that case, the verifier must check whether the state he gets is pure. This can be done when he gets two states ρ and σ that are promised to be uncorrelated; i.e., he gets the state $\rho \otimes \sigma$; then the verifier can use the “swap test” to estimate the quantity $\text{Tr}(\rho\sigma)$, and this can only be close to 1 when ρ is pure and $\sigma \approx \rho$. Indeed, if $\text{Tr}(\sigma^2) \leq 1 - \varepsilon$, then for all states τ , $\text{Tr}(\sigma\tau) \leq \sqrt{\text{Tr}(\sigma^2)\text{Tr}(\tau^2)} \leq (1 - \varepsilon)^{1/2} \leq 1 - \varepsilon/2$.

However, this test becomes inconclusive if the verifier gets a state that is not a product state $\rho \otimes \sigma$. So the above argument does not suffice to show that pure-state N representability is in QMA. However, it does show that the problem is in QMA(k), a variant of QMA where the verifier is promised to get a tensor product of k independent states [15]. It has been conjectured that QMA(k) is strictly larger than QMA, and indeed it is plausible that pure-state N representability is harder than ordinary N representability. It would be interesting to investigate whether this problem is QMA(k) complete.

It is remarkable that checking consistency of 2-body reduced density operators is so hard, while checking consistency of 1-body reduced density operators is simple [2]. This can be understood from the previous discussion: intuitively, 1-body density operators $\langle a_i^\dagger a_j \rangle$ correspond to Hamiltonians only containing bilinear terms in a_i^\dagger and a_j ; such Hamiltonians can easily be diagonalized as they represent systems of free fermions. As shown in [2], consistency can be decided in that case based solely on the

eigenvalues of the reduced density operators. A number of related problems have been investigated recently [16]; in particular, see [17].

These results have to be contrasted with our problem of deciding N representability for 2-body density operators, where the eigenvalues alone are not enough to decide consistency but also the eigenvectors are relevant. Actually, let us consider the simpler problem where only the diagonal elements of the 2-body density operators $D_{ij} = \langle a_i^\dagger a_j^\dagger a_j a_i \rangle$ are specified. Using the mapping from spins to fermions discussed above, one easily finds that these D_{ij} correspond to local spin Hamiltonians which only contain commuting σ^z operators. These are spin glasses, and so the problem of deciding N representability of $\{D_{ij}\}$ is NP hard [18]. It was indeed pointed out a long time ago that N representability restricted to the diagonal elements is equivalent to a combinatorial problem [19] that was later shown to be equivalent to the NP-hard problem of deciding membership in the boolean quadric polytope [20].

Let us finally discuss the relevance of the above results in the context of quantum chemistry. We have shown that determining ground state energies by means of the N -representability problem is intractable in the worst case, even on a quantum computer. This leaves open the possibility of finding efficient algorithms that give accurate results for particular physical systems (though they must break down in the general case). Typically, such algorithms rely on simpler approximations of the convex set of N -representable two-particle density matrices. The hope is that some physical systems may exhibit additional features that make the problem easier. This seems to be the case for, e.g., one-dimensional translational invariant spin systems, where the density matrix renormalization group allows for a systematic approximation of the convex set of allowed reduced density operators from within [21]. This might also be the case for some molecular systems, where applying positivity conditions on the 3-particle reduced density operators gives promising results (at least for molecules up to a certain size) [22]; this gives an approximation of the convex set from the outside. The preceding discussion shows that these methods cannot work in the most general case; it would be very interesting to investigate the conditions under which these kinds of approximations are justified.

In conclusion, we investigated the problem of N representability, and characterized its computational complexity by showing that it is QMA complete. Obviously, the theory of quantum computing was a prerequisite to pinpoint the complexity of this classic problem.

Y.K.L. and M.C. thank the Institute for Quantum Information for its hospitality. Y.K.L. is supported by ARO/DTO QuaCGR. M.C. acknowledges support from EPSRC and Magdalene College, Cambridge, and is supported by the EU under the FP6-FET Integrated Project

SCALA, No. CT-015714. F.V. is supported by the Gordon and Betty Moore Foundation through Caltech's Center for the Physics of Information, and by the NSF under Grant No. PHY-0456720.

-
- [1] C.A. Coulson, Rev. Mod. Phys. **32**, 170 (1960); R.H. Trethold, Phys. Rev. **105**, 1421 (1957).
 - [2] A.J. Coleman, Rev. Mod. Phys. **35**, 668 (1963).
 - [3] A.J. Coleman and V.I. Yukalov, *Reduced Density Matrices: Coulson's Challenge* (Springer, New York, 2000).
 - [4] J. Cioslowski, *Many-Electron Densities and Reduced Density Matrices* (Kluwer Academic, New York, 2000).
 - [5] D.A. Mazziotti, Acc. Chem. Res. **39**, 207 (2006).
 - [6] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalıy, *Classical and Quantum Computation* (AMS, Providence, 2002).
 - [7] J. Kempe, A. Kitaev, and O. Regev, SIAM J. Comput. **35**, 1070 (2006); quant-ph/0406180.
 - [8] R. Oliveira and B. Terhal, quant-ph/0504050.
 - [9] Y.-K. Liu, in *Proceedings of RANDOM 2006* (Springer-Verlag, Berlin, 2006), p. 438; quant-ph/0604166.
 - [10] F. Verstraete and J.I. Cirac, J. Stat. Mech. (2005) P09012.
 - [11] One possible set of observables is the following: First, define $a_I = a_{i_2} a_{i_1}$, for all pairs $I = \{i_1, i_2\}$, $i_1 < i_2$. Also fix an ordering on the pairs I . We now define the following observables: $X_{IJ} = a_I^\dagger a_J + a_J^\dagger a_I$, for all $I < J$; $Y_{IJ} = -i a_I^\dagger a_J + i a_J^\dagger a_I$, for all $I < J$; and $Z_I = a_I^\dagger a_I$, for all I except the last one. These operators are Hermitian, with eigenvalues in the interval $[-1, 1]$. Taking real linear combinations, these operators form a basis for the space of 2-fermion density matrices.
 - [12] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization* (Springer-Verlag, Berlin, 1988).
 - [13] Y.-K. Liu, M. Christandl, and F. Verstraete, quant-ph/0609125.
 - [14] D. Aharonov and O. Regev, in *Proceedings of Foundations of Computer Science 2003* (IEEE Press, Los Alamitos, California, 2003), p. 210; quant-ph/0307220.
 - [15] H. Kobayashi, K. Matsumoto, and Tomoyuki Yamakami, quant-ph/0110006.
 - [16] A. Higuchi, A. Sudbery, and J. Szulc, Phys. Rev. Lett. **90**, 107902 (2003); S. Bravyi, Quantum Inf. Comput. **4**, 12 (2004); M. Christandl and G. Mitchison, Commun. Math. Phys. **261**, 789 (2006); A. Klyachko, quant-ph/0409113; S. Daftuar and P. Hayden, Ann. Phys. (N.Y.) **315**, 80 (2005).
 - [17] A. Klyachko, J. Phys. Conf. Ser. **36**, 72 (2006).
 - [18] F. Barahona, J. Phys. A **15**, 3241 (1982).
 - [19] M. L. Yoseff and H. W. Kuhn, J. Math. Phys. (N.Y.) **10**, 703 (1969).
 - [20] M. Deza and M. Laurent, J. Comput. Appl. Math. **55**, 217 (1994).
 - [21] F. Verstraete and J.I. Cirac, Phys. Rev. B **73**, 094423 (2006).
 - [22] D.A. Mazziotti, Phys. Rev. A **74**, 032501 (2006).