

REPORT DOCUMENTATION PAGE				Form Approved OMB NO. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 04-02-2008		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Aug-2004 - 31-Jan-2008	
4. TITLE AND SUBTITLE QuaCGR Fellowship: Adiabatic Quantum Algorithms				5a. CONTRACT NUMBER W911NF-04-1-0374	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHORS Yi-Kai Liu				5d. PROJECT NUMBER 411359	
				5e. TASK NUMBER 411359	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of California - San Diego Office of Contract & Grant Administration 9500 Gilman Drive, Mail Code 0934 La Jolla, CA 92093 -0934				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S) ARO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 46436-PH-QC.1	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT We studied the quantum adiabatic algorithm for combinatorial optimization. We obtained a simpler proof of the adiabatic theorem. However, we did not make much progress in developing new tools for rigorous analysis of the algorithm's performance on realistic optimization problems. This analysis seems to be substantially more difficult than for classical algorithms such as simulated annealing, because the ground state does not have a simple form. We also proved some interesting results about the complexity of the Local Consistency problem (deciding whether local					
15. SUBJECT TERMS Final Report					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON David Meyer
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER 858-534-5524

Report Title

QuaCGR Fellowship: Adiabatic Quantum Algorithms

ABSTRACT

We studied the quantum adiabatic algorithm for combinatorial optimization. We obtained a simpler proof of the adiabatic theorem. However, we did not make much progress in developing new tools for rigorous analysis of the algorithm's performance on realistic optimization problems. This analysis seems to be substantially more difficult than for classical algorithms such as simulated annealing, because the ground state does not have a simple form.

We also proved some interesting results about the complexity of the Local Consistency problem (deciding whether local density matrices that describe small pieces of a quantum system are consistent with a single overall state). In particular, we showed that this problem is QMA-complete. We also showed that N-representability, an important problem in quantum chemistry, is QMA-complete.

List of papers submitted or published that acknowledge ARO support during this reporting period. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Y.-K. Liu, M. Christandl and F. Verstraete, "N-representability is QMA-complete," Phys. Rev. Lett. 98, 110503 (2007); Arxiv preprint: quant-ph/0609125.

Number of Papers published in peer-reviewed journals: 1.00

(b) Papers published in non-peer-reviewed journals or in conference proceedings (N/A for none)

Number of Papers published in non peer-reviewed journals: 0.00

(c) Presentations

Y.-K. Liu, "Gibbs States and the Consistency of Local Density Matrices," Arxiv preprint: quant-ph/0603012. Previously presented as a poster at the SQuInT workshop, Albuquerque, NM, Feb. 17-19, 2006.

Number of Presentations: 1.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts): 0

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Y.-K. Liu, "Consistency of Local Density Matrices is QMA-complete," Proc. RANDOM 2006, pp.438-449; Arxiv preprint: quant-ph/0604166.

Y.-K. Liu, V. Lyubashevsky and D. Micciancio, "On Bounded Distance Decoding for General Lattices," Proc. RANDOM 2006, pp.450-461.

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts): 2

(d) Manuscripts

Y.-K. Liu, "The Complexity of the Consistency and N-representability Problems for Quantum States," Ph.D. thesis, Univ. of California, San Diego, 2007; preprint ArXiv:0712.3041.

Y.-K. Liu, "Adiabatic quantum algorithms," research exam (technical report), Univ. of California, San Diego, 2004.

Number of Manuscripts: 2.00

Number of Inventions:

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Yi Kai Liu	1.00
FTE Equivalent:	1.00
Total Number:	1

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:.....	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:.....	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:	0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PhDs

NAME

Yi Kai Liu

Total Number:

1

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

QuaCGR Fellowship: Adiabatic Quantum Algorithms

Final Report (8/1/2004 - 8/31/2007)

Yi-Kai Liu*

1 Problem Statement

We studied the quantum adiabatic algorithm for combinatorial optimization, first proposed by [8]. The algorithm works as follows. One constructs a time-dependent Hamiltonian $H(t)$, $0 \leq t \leq T$, which can be efficiently implemented on a quantum computer. $H(t)$ is designed so that, at time 0, the ground state of $H(0)$ is known and can be prepared efficiently, while at time T , the ground state of $H(T)$ encodes the solution of some combinatorial optimization problem. Furthermore, $H(t)$ changes “slowly” over time.

The idea is that, if we prepare the system in the ground state at time 0, and allow it to evolve until time T , then the system will follow the ground state of $H(t)$, and end up finding the solution to our optimization problem. This will work provided the change in $H(t)$ is sufficiently slow; in particular, by the adiabatic theorem, it suffices to have $T \gg \text{poly}(1/g_{\min})$, where $g_{\min}$ is the minimum gap between the two lowest eigenvalues of $H(t)$ (for all times t). Thus, the spectral gap of $H(t)$ essentially determines the running time of the algorithm. (Other factors, such as the norm of $H(t)$ and its derivatives, also play a role, but they are usually less significant.)

For example, suppose we want to minimize the function f on the Boolean hypercube $\{0,1\}^n$. The adiabatic algorithm might use a Hamiltonian of the form:

$$H(t) = (1 - t/T)H_0 + (t/T)H_1$$
$$H_0 = \sum_{i=1}^n \sigma_x^{(i)}$$
$$H_1 = \sum_{z_1, \dots, z_n \in \{0,1\}} f(z_1, \dots, z_n) |z_1\rangle\langle z_1| \otimes \cdots \otimes |z_n\rangle\langle z_n|.$$

Here, the system consists of n qubits, $\sigma_x^{(i)}$ denotes the σ_x Pauli operator acting on the i 'th qubit, and $|0\rangle$ and $|1\rangle$ denote the up and down σ_z eigenstates for a single qubit.

*Computer Science and Engineering, University of California, San Diego. E-mail: y9liu@cs.ucsd.edu, yikailiu@caltech.edu.

The adiabatic algorithm was originally proposed to solve problems such as 3-SAT, but no rigorous analysis of the running time is known in that case. Rigorous analyses are only available for very simple “toy problems” (see, e.g., [22, 7, 18]). These give some insight into the algorithm—in particular, they show how the adiabatic algorithm can succeed by tunnelling through potential barriers, and how it can fail by getting drawn into local minima. But it remains difficult to understand the performance of the adiabatic algorithm on real optimization problems, because the situation there is so much more complicated. At best, there are heuristic arguments for those cases (see, e.g., [20]).

Our goal was to develop new techniques for rigorous analysis of the adiabatic algorithm, with the hope of understanding the performance of the algorithm on some moderately complicated, quasi-realistic problems. A rigorous analysis of this type would be useful in confirming (or refuting) the conclusions obtained from heuristic arguments. We remark that this sort of analysis has been carried out, to some extent, for classical algorithms such as simulated annealing (see, e.g., [21, 9, 1, 5, 6, 4]).

2 Summary of Results

The following is a summary of our main results. For more details, see the annual progress reports, and the reprints of papers that were submitted to ARO.

2.1 The Adiabatic Algorithm

We obtained a simple proof of the adiabatic theorem, using a martingale argument [11]. Another more powerful proof was discovered independently by Ambainis and Regev [2]. These results give some intuition about why the spectral gap is significant: a small spectral gap corresponds to a rapid change in the ground state of the Hamiltonian, which in turn implies a longer running time, to ensure that the system remains in the ground state. In many cases, the system seems to undergo a quantum phase transition, where there is an “avoided crossing” between the two lowest energy eigenvalues, coinciding with a sudden change in the ground state.

Unfortunately, we made little progress in developing new tools for rigorous analysis of the problem. The main difficulty is that, generally speaking, the ground state of the Hamiltonian does not have a simple, explicit description. The only known exceptions are when: (1) the Hamiltonian has a symmetry that effectively reduces the dimension of the problem, e.g., it is symmetric under permutations of the qubits [8, 22]; (2) the Hamiltonian has a special structure that can be diagonalized exactly, e.g., it can be transformed into a system of noninteracting fermions [18]; or (3) the ground state has relatively uncomplicated correlations and entanglement, e.g., certain 1-D spin chains that can be analyzed using the martingale method [17].

For comparison, note that in classical simulated annealing, the stationary distribution of

the Metropolis random walk does have a simple description, even in a complicated system—it is the Gibbs distribution. This makes it possible to calculate the second eigenvalue, for instance by variational methods, and thus bound the spectral gap.

Perhaps heuristic arguments and numerical methods, developed in condensed matter physics, are a more viable approach to understanding the adiabatic algorithm. For instance, Schutzhold and Schaller have an interesting interpretation of the algorithm in terms of first- and second-order quantum phase transitions [19]. Also, it may be possible to simulate some instances of the adiabatic algorithm using non-exact methods such as matrix-product states [3].

2.2 The Local Consistency Problem

We also obtained several results on the complexity of the “Local Consistency” problem. (This is not related to the adiabatic algorithm, but it is nonetheless interesting.) The problem is as follows: given a collection of density matrices $\rho_1, \dots, \rho_m$, where ρ_i describes a subset of qubits C_i in an n -qubit system, with $|C_i| \leq O(1)$, decide whether there exists an n -qubit state σ that agrees with all of the ρ_i on the subsets C_i .

In a certain sense, this problem is complementary to the Local Hamiltonian problem (finding the ground state energy of a local Hamiltonian), which is QMA-complete [10]. We showed that Local Consistency is QMA-complete, via a novel reduction using convex optimization with a membership oracle [12]. We also showed that N -representability, an important problem in quantum chemistry, is QMA-complete [16]. In addition, we proved a structural property of Local Consistency (that if a solution exists, it can be chosen to be a Gibbs state) [13]; and we identified special cases of Local Consistency, for 1-D and stoquastic systems, whose complexity appears to be strictly easier than QMA [15]. These results are described in greater detail in my thesis [14].

2.3 Publications and Demographic Data

The following information is for the period from 8/1/2006 to 8/31/2007 (since last year’s progress report). Papers during this reporting period: [16] (published in journal), [14] (PhD dissertation). Demographic data: 1 full-time graduate student supported by this agreement. Inventions: none. Technology transfer: none. Other activities: Yi-Kai Liu gave a talk at RANDOM 2006; visited the Institute for Quantum Computation (IQC) at Waterloo in January 2007; gave an invited talk at QIP 2007; and attended the SQuInT 2007 workshop.

References

- [1] D. Aldous and U.V. Vazirani. “Go with the winners” algorithms. In *Foundations of Computer Science (FOCS)*, pages 492–501, 1994.

- [2] A. Ambainis and O. Regev. An elementary proof of the quantum adiabatic theorem. ArXiv:quant-ph/0411152, 2004.
- [3] M.C. Banuls, R. Orus, J.I. Latorre, A. Perez, and P. Ruiz-Femenia. Simulation of many-qubit quantum computation with matrix product states. *Phys. Rev. A*, 73:022344, 2006.
- [4] T. Carson and R. Impagliazzo. Hill-climbing finds random planted bisections. In *Symposium on Discrete Algorithms (SODA)*, pages 903–909, 2001.
- [5] T. Dimitriou and R. Impagliazzo. Towards an analysis of local optimization algorithms. In *Symposium on the Theory of Computing (STOC)*, pages 304–313, 1996.
- [6] T. Dimitriou and R. Impagliazzo. Go with the winners for graph bisection. In *Symposium on Discrete Algorithms (SODA)*, pages 510–520, 1998.
- [7] E. Farhi, J. Goldstone, and S. Gutmann. Quantum adiabatic evolution algorithms versus simulated annealing. ArXiv:quant-ph/0201031, 2002.
- [8] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution. ArXiv:quant-ph/0001106, 2000.
- [9] M. Jerrum and G.B. Sorkin. The Metropolis algorithm for graph bisection. *Discrete Applied Mathematics*, 82(1-3):155–175, 1998.
- [10] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalı. *Classical and quantum computation*. American Mathematical Society, Providence, RI, 2002.
- [11] Y.-K. Liu. Adiabatic quantum algorithms. Research exam (technical report), UCSD, 2004.
- [12] Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *Approximation, Randomization and Combinatorial Optimization (APPROX + RANDOM '06)*, pages 438–449, 2006. LNCS 4110, Springer.
- [13] Y.-K. Liu. Gibbs states and the consistency of local density matrices. ArXiv:quant-ph/0603012, 2006.
- [14] Y.-K. Liu. *The Complexity of the Consistency and N-representability Problems for Quantum States*. PhD thesis, Univ. of California, San Diego, 2007. Will be posted on Arxiv.
- [15] Y.-K. Liu. The local consistency problem for stoquastic and 1-d quantum systems. ArXiv:0712.1388 [quant-ph], 2007.
- [16] Y.-K. Liu, M. Christandl, and F. Verstraete. Quantum computational complexity of the N-representability problem: QMA complete. *Phys. Rev. Lett.*, 98(11):110503, 2007.
- [17] B. Nachtergaele. The spectral gap for some spin chains with discrete symmetry breaking. *Commun. Math. Phys.*, 175:565–606, 1996. ArXiv:cond-mat/9410110.

- [18] B.W. Reichardt. The quantum adiabatic optimization algorithm and local minima. In *Symposium on the Theory of Computing (STOC) 2004*.
- [19] R. Schutzhold and G. Schaller. Adiabatic quantum algorithms as quantum phase transitions: first versus second order. ArXiv:quant-ph/0608017, 2006.
- [20] V.N. Smelyanskiy, S. Knysh, and R.D. Morris. Quantum adiabatic optimization and combinatorial landscapes. ArXiv:quant-ph/0402199, 2004.
- [21] G.B. Sorkin. Efficient simulated annealing on fractal energy landscapes. *Algorithmica*, 6(3):367–418, 1991.
- [22] W. van Dam, M. Mosca, and U. Vazirani. How powerful is adiabatic quantum computation? In *Foundations of Computer Science (FOCS) 2001*, pages 279–287. ArXiv:quant-ph/0206003.

N -representability is QMA-complete

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

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

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

³*Institute for Quantum Information, Caltech, Pasadena, US*

⁴*Facultät für Physik, Universität Wien, Boltzmanngasse 5, A-1090 Wien, Austria*

(Dated: September 19, 2006)

We study the computational complexity of the N -representability problem in quantum chemistry. We show that this problem is QMA-complete, which is the quantum generalization of NP-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.

PACS numbers: 03.67.a, 31.25.-v

The central theoretical problem in the field of many-body strongly correlated quantum systems is to find efficient ways of simulating Schrödinger's equations: it is very easy to write down those equations, but notoriously difficult to solve them or even to find approximate solutions. 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 wavefunctions, 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 wavefunctions 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. So far, an efficient variational class to describe complex wavefunctions such as those arising in quantum chemistry has not been found.

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. The essential element that makes typical Hamiltonians very ungeneric is the fact that at most 2-body interactions occur. 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.

In the case of a Hamiltonian with only 2-body interactions, the energy corresponding to a wavefunction 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 and there exists a large literature on this subject (see e.g. [3, 4, 5]).

Here we will show that the N -representability problem is intractable, as it is QMA-complete and hence NP-hard. The complexity class QMA (Quantum Merlin-Arthur) 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 [23]. 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 [24]. 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, 8], and even in the case of geometrically local interactions [9].

Because the Hamiltonians under consideration are local, the corresponding ground states have extremal local properties. The dual problem of determining the ground state energy of a local Hamiltonian is to decide whether

a given set of local density operators can be realized as the reduced density operators of the same global state. Checking the consistency of such a set of local density operators $\rho^{(ij)}$, i.e., checking whether there exists a global state σ compatible with those local density operators, is a QMA-complete problem itself [10]. In the present paper, we will use very similar techniques to prove that N -representability, which is the fermionic version of that problem, is also QMA-complete.

To sketch the main ideas of the proof, we will first consider the classical marginal problem: suppose we have N random variables that can take the values ± 1 , according to some joint probability distribution, and we define the 2-variable marginal distributions

$$p(a_i, a_j) = \sum_{a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_{j-1}, a_{j+1}, \dots, a_N} p(a_1, a_2, \dots, a_N).$$

Given a set of marginal probability distributions, we would like to check whether they are consistent, i.e., whether there exists a global probability distribution whose marginals are equal to the ones we were given.

Suppose that there exists an efficient algorithm (whose running time is polynomial in N) to solve this problem. Then it would be possible to identify ground states of all Ising spin glasses. The strategy would be as follows: since the set of consistent marginals is convex and the energy is a linear function of the marginals, the problem amounts to minimizing a linear function subject to a set of convex constraints that can be checked in polynomial time. This can be done in polynomial time using the ellipsoid method [11]. It therefore follows that an efficient algorithm for checking consistency allows to find ground states of Ising spin glasses, a problem that is known to be NP-hard [12]. Hence the problem of determining whether a set of binary marginals is consistent is itself NP-hard.

Let us now formulate the N -representability problem in the context of quantum chemistry. Electrons and nuclei tend to arrange themselves such as to minimize their energy, and the binding energy of a molecule can be determined by calculating the minimal energy of the corresponding Hamiltonian. In practice, the nuclei are fairly well localized and can be treated as classical degrees of freedom, and the wavefunction 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 wavefunction must be antisymmetric, and this is most easily taken into account by working in the formalism of second quantization:

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

Here $a_j^\dagger$ is the creation operator for the j 'th mode, and $|\Omega\rangle$ represents the vacuum state without fermions. The

creation and annihilation operators obey the following anticommutation relations:

$$\{a_i, a_j\} = 0 = \{a_i^\dagger, a_j^\dagger\} \quad \{a_i, a_j^\dagger\} = 2\delta_{ij}.$$

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 $|\psi\rangle$ with exactly N fermions. In the language of second quantization, 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 (with d modes) 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 $\frac{d(d-1)}{2} \times \frac{d(d-1)}{2}$, where each entry is specified with $\text{poly}(N)$ bits of precision. In addition, we are given a real number $\beta \geq 1/\text{poly}(N)$, specified with $\text{poly}(N)$ bits of precision. The problem is to distinguish between the following two cases:

- There exists an N -fermion state σ such that $\text{tr}_{3, \dots, N}(\sigma) = \rho$. In this case, answer “YES.”
- 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 do not insist on solving the problem exactly; we allow an error tolerance of $\beta \geq 1/\text{poly}(N)$. We use the ℓ_1 matrix norm or trace distance, $\|A\|_1 = \text{tr}|A|$, to measure the distance between σ and ρ .

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 oracle 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, defined on a system of N qubits, to a 2-local Hamiltonian on fermions, with $d = 2N$ modes; this is the opposite of what has been done in [13]. The idea is to 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:

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

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:

$$\sigma_i^x \leftrightarrow a_i^\dagger b_i + b_i^\dagger a_i, \quad \sigma_i^y \leftrightarrow i \left(b_i^\dagger a_i - a_i^\dagger b_i \right), \quad \sigma_i^z \leftrightarrow \mathbb{1} - 2b_i^\dagger b_i. \quad (2)$$

It can easily be checked that these fermionic operators obey the Pauli commutation relations within the subspace spanned by states of the form (1).

Recall that the qubit Hamiltonian can be written as a linear combination of Pauli operators. So the fermionic Hamiltonian can now be created by rewriting these Pauli operators with respect to the above mapping (note that only bilinear terms occur, and hence operators on different sites commute). Notice that if the qubit Hamiltonian only contains 2-body terms, then the fermionic Hamiltonian 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 the following projectors as extra terms in the fermionic Hamiltonian:

$$P_i = (2a_i^\dagger a_i - \mathbb{1})(2b_i^\dagger b_i - \mathbb{1}).$$

All the P_i are biquadratic and commute with all the operators introduced in (2), and hence the complete Hamiltonian will be block diagonal. By making the weights of these projectors 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 [25].

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 a 2-local Hamiltonian on qubits, a problem which is known to be QMA-hard [8]. We start by transforming the Hamiltonian H_{qubit} on N qubits with 2-particle interactions into a Hamiltonian H_{fermi} on $d = 2N$ fermionic modes with 2-particle interactions (as described above). The problem of finding the ground state energy of H_{fermi} can be

expressed as a convex program, which can be solved in polynomial time using an oracle for N -representability (this is what we will show below). Our approach is similar to [10], though in this case we are dealing with fermions rather than qubits.

The basic idea is to construct a convex program that finds a 2-particle density matrix ρ which is N -representable, and which minimizes the expectation value of H_{fermi} . This program has polynomially many variables, and it is easy to see that the set of N -representable states is convex, and $\langle H_{\text{fermi}} \rangle$ is a linear function of ρ .

There are two minor complications. First, H_{fermi} describes a system with $2N$ modes and an arbitrary number of particles, whereas N -representability pertains to a system with exactly N particles. But this is not a problem, because the interesting behavior in H_{fermi} occurs in a subspace where all the states have exactly N particles. So, in our convex program, we only need to consider states ρ that have exactly N particles.

This lets us simplify H_{fermi} as follows. Since we are only interested in how it acts on N -particle states, we can write $a_i^\dagger a_j = \frac{1}{N-1} a_i^\dagger (\sum_k a_k^\dagger a_k) a_j$. So we can assume that all the terms in H_{fermi} are of the form $a_i^\dagger a_j^\dagger a_l a_k$.

Second, 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 [26], 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. Note that the N -representability oracle allows us to 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); the coefficients γ_S can be computed using the Gram-Schmidt procedure. 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 (see [11], and references cited therein). We mention a few technical details. The algorithm requires some additional information about K , namely 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 . (Also, the running time of the algorithm grows polynomially in $\log(R/r)$.) In our case, we can set $R = \sqrt{\ell}$, and $r = 1/\text{poly}(\ell)$ [27].

There is also the issue of numerical precision. Our oracle for N -representability has limited accuracy—it may give incorrect answers for points near the boundary of K . On the other hand, we only need an approximate

solution to the convex program, in order to solve the Local Hamiltonian problem. It turns out that $1/\text{poly}(N)$ precision is sufficient. The ellipsoid algorithm still works in this setting; see [11] for details.

Note that, in place of the ellipsoid algorithm, we could have used a different algorithm based on random walks in convex bodies [14]; this was the approach used in [10]. However, it is not clear if we could use one of the interior-point methods for convex optimization; these methods are substantially faster, but they usually require an explicit description of the constraints, not just a membership oracle.

This completes the proof that N -representability is QMA-hard. As a corollary, we have also proven that estimating the ground state energy for local fermionic Hamiltonians is QMA-hard. Note that the use of the ellipsoid algorithm to reduce a convex optimization problem to a convex membership problem is not new; see [11] for references to related work.

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 $\tilde{O} = A_i^\dagger A_j^\dagger A_l A_k + A_k^\dagger A_l^\dagger A_j A_i$, which is a tensor product of many single-qubit observables and one four-qubit observable.

We claim that the expectation value $\langle \tilde{O} \rangle$ can be estimated efficiently. Without loss of generality, assume that the eigenvalues of $\tilde{O}$ lie in the interval $[0, 1]$. Then there is an efficiently-implementable measurement which outputs “1” with probability $\langle \tilde{O} \rangle$ [28]. By repeating this measurement on multiple copies of the same state, we can estimate $\langle \tilde{O} \rangle$; for our purposes, it is enough to have polynomially many copies of the state.

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. Then V checks whether they match the expectation values of ρ . One problem arises: 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 [15]. This suffices to show that N -representability is in QMA, and hence finishes the proof.

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 be able to convince himself that 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. that he gets the state $\rho \otimes \sigma$: then Arthur can calculate the expectation value of the observable $\text{tr}(\rho\sigma)$, and this can only be close to 1 when ρ is pure and $\sigma \simeq \rho$. Indeed, if $\text{tr}(\sigma^2) \leq 1 - \varepsilon$, then for all states τ , $\text{tr}(\sigma\tau) \leq \sqrt{\text{tr}(\sigma^2)\text{tr}(\tau^2)} \leq (1 - \varepsilon)^{1/2} \leq 1 - \varepsilon/2$. The problem however is that Merlin can cheat, and hand out a correlated state to Arthur, such that the above test becomes inconclusive. This is precisely the feature that distinguishes the complexity class QMA(k) from QMA [16]: in QMA(k), the verifier is promised to get a tensor product of different states, and it has been conjectured that QMA(k) is strictly larger than QMA. The above discussion shows that the pure-state N -representability problem is contained in QMA(k), and it is hence plausible that it is harder than the N -representability problem. It would be very interesting to investigate whether the problem is QMA(k)-complete.

It is remarkable that checking consistency of 2-body reduced density operators of fermionic states is so hard, while checking consistency of 1-body reduced density operators is simple [2]. This can be easily understood from the previous discussion: the extreme points of the convex set of 1-body density operators $\langle a_i^\dagger a_j \rangle$ correspond to ground states of 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, and hence the consistency problem can easily be solved. 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, where consistency only depends on the eigenvalues, have been investigated recently [17]; most remarkably, a character-

sation has been obtained for the polytope of 1-particle marginals that are N -representable under the condition that the N -particle wavefunction must be pure [18].

These results have to be contrasted with our problem of finding the N -representable 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 of deciding N -representability of 2-fermion density operators where only the diagonal elements $D_{ij} = \langle a_i^\dagger a_j^\dagger a_j a_i \rangle$ are specified. If we consider the case $d = 2N$ and the mapping discussed above, one easily finds that the extreme points of this set would be obtained by ground states of local spin Hamiltonians which only contain commuting σ^z operators. These correspond to spin-glasses, and so the problem of deciding N -representability of $\{D_{ij}\}$ is NP-hard [12]. 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's finally discuss the relevance of the above results in the context of quantum chemistry. We have shown that it is a hopeless task to determine the ground states of all local fermionic Hamiltonians, and in particular, that an approach by means of the N -representability problem is intractable, even on a quantum computer. But it is possible that other physical systems, e.g. systems with different particle statistics, additional symmetry or a limited number of modes, might allow for an efficient characterisation of the two-particle reduced density matrices, and hence an efficient calculation of the ground states of local Hamiltonians. This seems to be the case in e.g. one-dimensional translational invariant spin systems, where the density matrix renormalization group [21] allows for a systematic approximation of the allowed convex set of reduced density operators from within [22]. A very different numerical approach has been developed in the context of quantum chemistry and it known as the the contracted Schrödinger equation [4]. In that method, one approximates the convex set from the outside, and the N -representability problem is a crucial ingredient of the algorithm. The foregoing discussion shows that this approach has to break down in the most general case, and it would be very interesting to investigate the conditions under which those 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 computational complexity of this classic problem. .

Acknowledgements: Y.K.L. and M.C. thank the Institute for Quantum Information for its hospitality. Y.K.L. is supported by an ARO/DTO Quantum Computing Grad-

uate Research Fellowship. M.C. acknowledges an EPSRC Postdoctoral and a Nevile Research Fellowship which he holds at Magdalene College Cambridge, and is supported by the EU under the FP6-FET Integrated Project SCALA, 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 National Science Foundation under Grant No. PHY-0456720.

-
- [1] C. A. Coulson, Rev. Mod. Phys. **32**, 175 (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 (2000).
 - [4] J. Cioslowski, *Many-Electron Densities and Reduced Density Matrices*, Kluwer Academic (2000).
 - [5] D. A. Mazziotti, Acc. Chem. Res. **39**, 207 (2006).
 - [6] A. Yu Kitaev, A. H. Shen and M. N. Vyalai, *Classical and Quantum Computation*, AMS (2002).
 - [7] J. Kempe and O. Regev, Quant. Info. and Comput. **3**, 258 (2003).
 - [8] J. Kempe, A. Kitaev and O. Regev, SIAM J. Comput. **35**, 1070 (2006); quant-ph/0406180.
 - [9] R. Oliveira and B. Terhal, quant-ph/0504050.
 - [10] Y.-K. Liu, Proc. RANDOM 2006, 438; quant-ph/0604166.
 - [11] M. Grötschel, L. Lovász and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Springer-Verlag (1988).
 - [12] F. Barahona, J. Phys. A: Math. Gen. **15**, 3241 (1982).
 - [13] F. Verstraete and J.I. Cirac, J. Stat. Mech. P09012 (2005).
 - [14] D. Bertsimas and S. Vempala, J. Assoc. Comput. Mach. **51**, 540-556 (2004).
 - [15] D. Aharonov and O. Regev, Proc. FOCS 2003, 210; quant-ph/0307220.
 - [16] H. Kobayashi, K. Matsumoto and Tomoyuki Yamakami, quant-ph/0110006.
 - [17] A. Higuchi, A. Sudbery, J. Szulc, Phys. Rev. Lett. **90**, 107902 (2003); S. Bravyi, Quantum Inf. and Comp. **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. **315**, 80 (2005).
 - [18] A. Klyachko, J. of Physics: Conference Series **36** (1) pp.72-86 (2006).
 - [19] M. L. Yoseloff and H. W. Kuhn, Journ. of Math. Phys. **10**, 703 (1969).
 - [20] M. Deza and M. Laurent, Journ. of Comp. and Appl. Math. **55**, 217 (1994).
 - [21] U. Schollwoeck, Rev. Mod. Phys. **77**, 259 (2005).
 - [22] F. Verstraete and J.I. Cirac, Phys. Rev. B **73**, 094423 (2006).
 - [23] Technically, the problem must be formulated as a yes-or-no question, e.g., "for a given Hamiltonian H , is the ground state energy less than γ ?" If the answer is yes, there exists a "witness" (e.g., the ground state) which can prove it. Moreover, there is an efficient algorithm that can check the validity of a witness.

- [24] This includes search and optimization problems, e.g., “find the ground state energy of H .” These problems may be QMA-hard, even though they are not phrased as yes-or-no questions.
- [25] Another option would be to represent each qubit i using either zero or two fermions, e.g., for a single qubit, $|0\rangle$ and $|1\rangle$ correspond to $|\Omega\rangle$ and $a_i^\dagger b_i^\dagger |\Omega\rangle$. Define the mapping $\sigma_i^x \leftrightarrow (a_i^\dagger b_i + b_i^\dagger a_i) + (b_i a_i + a_i^\dagger b_i^\dagger)$, $\sigma_i^y \leftrightarrow i(b_i^\dagger a_i - a_i^\dagger b_i) + i(a_i^\dagger b_i^\dagger - b_i a_i)$, $\sigma_i^z \leftrightarrow \mathbb{1} - 2b_i^\dagger b_i$. Then the Hamiltonian would act identically on the subspace with one fermion per site, and on the subspace with zero or two fermions per site, and even a small constant in front of the projectors P_i would be enough to guarantee that we end up in the right subspace.
- [26] 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 $I < J$; $Y_{IJ} = -ia_I^\dagger a_J + ia_J^\dagger a_I$, 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.
- [27] It is easy to see that K is contained in a ball of radius $R = \sqrt{\ell}$, since for all $\vec{\alpha} \in K$, we have $-1 \leq \alpha_S \leq 1$.

The second claim, that K contains a ball of radius $r = 1/\text{poly}(\ell)$, is less trivial. The proof is as follows.

We will consider N -representability for different values of N ; let K_N denote the set of all vectors $\vec{\alpha}$ that are N -representable. Obviously, K_2 contains a ball of radius $1/\text{poly}(\ell)$ (this is the trivial case). We will prove that K_N contains a ball of radius $1/\text{poly}(\ell)$, for all $2 \leq N \leq d-2$.

We define “particle-hole” observables, by replacing a_i with $a_i^\dagger$, and vice versa: $X'_{IJ} = a_I^\dagger a_J^\dagger + a_J a_I$; $Y'_{IJ} = -ia_I^\dagger a_J^\dagger + ia_J a_I$; and $Z'_I = a_I^\dagger a_I$. Let $\vec{\alpha}'$ denote a vector containing expectation values for these observables; let K'_N be the set of all $\vec{\alpha}'$ that are N -representable.

First, we claim that $K_2 = K'_{d-2}$. Notice that there is a natural correspondence between the 2-particle Slater basis states and the $(d-2)$ -particle Slater basis states: the 2-particle state with modes i and j occupied corresponds to the $(d-2)$ -particle state with modes i and j empty.

So take any point $\alpha \in K_2$, which represents the expectation values of the 2-particle observables for some 2-particle state σ . Use σ to construct a $(d-2)$ -particle state τ , by replacing each 2-particle Slater basis state with the corresponding $(d-2)$ -particle Slater basis state. Then the expectation values of the 2-particle observables for σ are exactly the expectation values of the 2-hole observables for τ . So α is in K'_{d-2} . This shows that $K_2 \subseteq K'_{d-2}$. A similar argument shows that $K'_{d-2} \subseteq K_2$. So $K_2 = K'_{d-2}$.

Next, we claim that there is an invertible linear transformation A that maps K'_{d-2} to K_{d-2} .

First we map K'_{d-2} to K_{d-2} . Observe that we can write each 2-particle operator as a linear combination of 2-hole operators. (This holds provided we restrict

the operators to act only on the subspace of $(d-2)$ -particle states.) For instance, when $I \cap J = \emptyset$, $a_I^\dagger a_J = a_J a_I^\dagger$. When $|I \cap J| = 1$, we write equations such as $a_i^\dagger a_l^\dagger a_l a_j = a_i^\dagger a_j a_l^\dagger a_l = -a_j a_i^\dagger (1 - a_l a_l^\dagger) = -a_j a_i^\dagger + a_j a_i^\dagger a_l a_l^\dagger = -a_j a_i^\dagger + a_l a_j a_i^\dagger a_l^\dagger$; then use the fact that $a_j a_i^\dagger = (\sum_{k \neq i,j} a_k a_k^\dagger) a_j a_i^\dagger = \sum_{k \neq i,j} a_k a_j a_i^\dagger a_k^\dagger$. When $I = J$, we have $a_I^\dagger a_I = \sum_{L \cap I = \emptyset} a_L a_L^\dagger$. (For each of these equations, it is easy to check that the left and right sides act identically on all $(d-2)$ -particle Slater basis states.)

Thus, for any $(d-2)$ -particle state σ , the expectation values of the 2-particle observables are linear functions of the expectation values of the 2-hole observables. Thus we have a linear transformation A that maps K'_{d-2} to K_{d-2} .

Similarly, we can map K_{d-2} to K'_{d-2} . We write each 2-hole operator as a linear combination of 2-particle operators (again, restricting the operators to act only on the subspace of $(d-2)$ -particle states). For instance, when $I \cap J = \emptyset$, $a_I a_J^\dagger = a_J^\dagger a_I$. When $|I \cap J| = 1$, we write equations such as $a_l a_i a_j^\dagger a_l^\dagger = a_i a_j^\dagger a_l a_l^\dagger = -a_j^\dagger a_i (1 - a_l^\dagger a_l) = -a_j^\dagger a_i + a_j^\dagger a_i a_l^\dagger a_l = -a_j^\dagger a_i + a_j^\dagger a_l^\dagger a_l a_i$; then use the fact that $a_j^\dagger a_i = (\frac{1}{d-3} \sum_{k \neq i,j} a_k^\dagger a_k) a_j^\dagger a_i = \frac{1}{d-3} \sum_{k \neq i,j} a_j^\dagger a_k^\dagger a_k a_i$. When $I = J$, we write $a_I a_I^\dagger = a_{i_1} a_{i_1}^\dagger a_{i_2} a_{i_2}^\dagger = (1 - a_{i_1}^\dagger a_{i_1})(1 - a_{i_2}^\dagger a_{i_2}) = 1 - a_{i_1}^\dagger a_{i_1} - a_{i_2}^\dagger a_{i_2} + a_I^\dagger a_I$; then use the fact that $1 = \binom{d-2}{2}^{-1} \sum_L a_L^\dagger a_L$, and $a_i^\dagger a_i = \frac{1}{d-3} \sum_{l \neq i} a_l^\dagger a_l a_i^\dagger a_i = \frac{1}{d-3} \sum_{l \neq i} a_i^\dagger a_l^\dagger a_l a_i$.

Thus, for any $(d-2)$ -particle state σ , the expectation values of the 2-hole observables are linear functions of the expectation values of the 2-particle observables. Thus we have a linear transformation B that maps K_{d-2} to K'_{d-2} , and $B = A^{-1}$.

We claim that K_{d-2} contains a ball of radius $1/\text{poly}(\ell)$. We know that K'_{d-2} contains a ball of radius $1/\text{poly}(\ell)$ (since $K_2 = K'_{d-2}$), and we will show that the map A does not shrink this too much. Write the singular value decomposition $A = UDV$, where U and V are unitary, and D is diagonal, with diagonal entries $D_{ii} > 0$. Let $B = A^{-1}$. Looking at the matrix elements of B , we can see that $\text{tr}(B^\dagger B) = \sum_{ij} |B_{ij}|^2 \leq \text{poly}(\ell)$. At the same time, $\text{tr}(B^\dagger B) = \text{tr}(UD^{-1}V V^{-1}D^{-1}U^{-1}) = \text{tr}(D^{-2}) \geq D_{ii}^{-2}$, for all i . So we have $D_{ii} \geq 1/\text{poly}(\ell)$, for all i . Thus the map A shrinks the radius of the ball by at most a $\text{poly}(\ell)$ factor. So K_{d-2} contains a ball of radius $1/\text{poly}(\ell)$.

Now we are almost done. Notice that $K_{N-1} \supseteq K_N$, for all $3 \leq N \leq d-2$ (since, if $\vec{\alpha}$ is consistent with some N -fermion state σ , then it is consistent with the $(N-1)$ -fermion state σ' that results from tracing out the last particle). So we conclude that K_N contains a ball of radius $1/\text{poly}(\ell)$, for all $2 \leq N \leq d-2$. This completes the proof.

- [28] For example, let λ_i and $|\theta_i\rangle$ be the eigenvalues and eigenvectors of $\tilde{O}$, and note that they are easy to compute. Add an ancilla qubit, perform the unitary operation $U : |\theta_i\rangle|0\rangle \mapsto |\theta_i\rangle(\sqrt{1-\lambda_i}|0\rangle + \sqrt{\lambda_i}|1\rangle)$, and measure the ancilla in the 0,1 basis.