Computational complexity of interacting electrons and fundamental limitations of density functional theory

Norbert Schuch[1]★ and Frank Verstraete[2]★

One of the central problems in quantum mechanics is to determine the ground-state properties of a system of electrons interacting through the Coulomb potential. Since its introduction[1,2], density functional theory has become the most widely used and successful method for simulating systems of interacting electrons. Here, we show that the field of computational complexity imposes fundamental limitations on density functional theory. In particular, if the associated 'universal functional' could be found efficiently, this would imply that any problem in the computational complexity class Quantum Merlin Arthur could be solved efficiently. Quantum Merlin Arthur is the quantum version of the class NP and thus any problem in NP could be solved in polynomial time. This is considered highly unlikely. Our result follows from the fact that finding the ground-state energy of the Hubbard model in an external magnetic field is a hard problem even for a quantum computer, but, given the universal functional, it can be computed efficiently using density functional theory. This work illustrates how the field of quantum computing could be useful even if quantum computers were never built.

The difficulty of finding the ground-state properties of a large system of interacting electrons originates both from the exponential dimension of the underlying Hilbert space and from the fermionic nature of the wavefunction. It is a problem encountered virtually everywhere in quantum chemistry as well as in condensed-matter physics: for instance, the spatial configuration of a molecule is the one for which the energy of the interacting electrons moving in the nuclear potential, together with the electrostatic energy of the nuclei, becomes minimal. Similarly, a rich variety of phenomena in solid-state physics, in particular conductance and magnetic phenomena, can be understood by considering electrons moving in the periodic lattice potential, including such exciting phenomena as high-temperature superconductivity and the fractional quantum Hall effect.

A system of N electrons is described by the Hamiltonian

$$H = -\frac{1}{2}\underbrace{\sum_{i=1}^{N}\Delta_i}_{=:T} + \underbrace{\sum_{1\leq i<j\leq N}\frac{\gamma}{|r_i-r_j|}}_{=:I} + \sum_i V(x_i) \qquad (1)$$

($\gamma > 0$, and $x_i = (r_i, s_i)$ with r_i position and s_i spin), where the potential V contains both an electrostatic field $\phi(r)$ and a magnetic field $\mathbf{B}(r)$, which couples to the spin (the coupling to the orbit can be ignored for our purposes; see Supplementary Information), and the problem is to find the ground state within the set of fermionic

(that is, antisymmetric) quantum states. Following the early work of ref. 3, it was shown[1,2] that this problem could be rephrased as a single-particle minimization problem, for the reason that the only problem-dependent part is the external potential V, whose expectation value only depends on the local density, whereas the kinetic and interaction terms T and I are fixed and universal for all systems. Thus, the ground-state energy is given by

$$E_0 = \min_\rho\{\mathrm{tr}(V\rho) + F[\rho]\} \qquad (2)$$

where ρ is a single-electron density, and the functional F contains the problem-independent minimization over T and I,

$$F[\rho] = \min_{\Omega\to\rho}\mathrm{tr}[(T+I)\Omega] \qquad (3)$$

Here, the minimization runs over all N-electron density operators Ω that give rise to the reduced density ρ. The central requirement for a good density functional theory (DFT) algorithm is to find a suitable approximation to the universal functional, and indeed better and better techniques have been developed, making DFT the most widely used and most successful algorithm for treating interacting electrons.

However, as we show in this letter, there exist fundamental limits that constrain the ability to find a generally applicable and efficiently computable approximation to the universal functional, and thus put bounds on the applicability of DFT. To this end, we consider the two-dimensional (2D) Hubbard model with local magnetic fields, which arises from the problem of interacting electrons for a specifically chosen lattice potential, and can thus be simulated using DFT. We first determine the computational complexity of solving the Hubbard model and show that it is among the hardest problems in the complexity class QMA, Quantum Merlin Arthur. QMA contains problems that are believed to be hard to solve even by quantum computers, but once a solution is found it can be checked efficiently by a quantum computer. Thus, QMA encompasses the complexity class NP. We compare this with the difficulty of solving the Hubbard model using DFT with a suitable approximation of the functional at hand, and find that in that case the Hubbard model can be solved by a classical computer in a time polynomial in the number of electrons. This means that the existence of an efficient approximation to the functional would imply QMA = P, that is, computing the functional to polynomial accuracy in the number of electrons is a QMA-hard problem, which poses fundamental limitations on the ability to approximate the functional in DFT. Of

[1]Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, D-85748 Garching, Germany, [2]Fakultät für Physik, Universität Wien, Boltzmanngasse 5, A-1090 Wien, Austria. *e-mail: norbert.schuch@gmail.com; fverstraete@gmail.com.
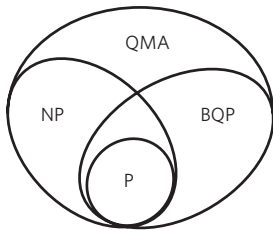
Figure 1 | The relevant complexity classes and their relationships. P and BQP are the classes of problems efficiently solvable by classical and quantum computers, respectively; NP (QMA) contains decision problems that are likely to be hard to solve by classical (quantum) computers, but where, for positive instances, classical (quantum) proofs exist that can be checked efficiently by a classical (quantum) computer. All inclusions are believed to be strict. We show that solving the Hubbard model is among the hardest problems in QMA, whereas the existence of an efficient description of the universal functional in DFT would put it in P, leading to the collapse of all aforementioned complexity classes. This puts tight bounds on the existence of such functionals.

course, this does not mean that DFT is not applicable in practice: much lower (for example constant) accuracies will typically suffice, and DFT is indeed a highly successful method.

The 2D Hubbard model[4,5] describes a system of fermions hopping on a lattice. Although it typically appears as a phenomenological model for strongly bound electrons in solid-state physics[6], it can be derived rigorously from (1) for an appropriate potential, as we show in the Supplementary Information. The Hubbard model with local magnetic fields is given by the Hamiltonian

$$H_{\text{Hubb}} = -t \sum_{\langle i,j \rangle, s} a_{i,s}^{\dagger} a_{j,s} + U \sum_{i} n_{i,\uparrow} n_{i,\downarrow} - \sum_{i} \boldsymbol{\sigma}_{i} \cdot \mathbf{B}_{i} \qquad (4)$$

where $a_{i,s}^{\dagger}$ creates an electron of spin $s \in \{\uparrow, \downarrow\}$ on lattice site i, $\langle i,j \rangle$ denotes nearest neighbours on the 2D square lattice, $n = a^{\dagger}a$, $\boldsymbol{\sigma}_{i} = (\sigma^{x,i}, \sigma^{y,i}, \sigma^{z,i})$ and $\sigma^{\alpha,i} = \sum_{s,s'} \sigma_{ss'}^{\alpha} a_{i,s}^{\dagger} a_{i,s'}$ with σ^{α} the Pauli matrices. The first term describes an electron tunnelling at rate t from one site to the adjacent one without changing its spin and the second the on-site Coulomb repulsion U between two electrons of different spin sitting on the same site, and the rightmost term contains the contribution from the magnetic field, which imposes a local field $\mathbf{B}_{i}$ at each site i—this is the only term that we can tune locally.

The 2D Hubbard model is of high interest on its own, as it is the minimal model that is believed to describe the physics arising in high-temperature superconductivity, quantum magnetism and heavy fermions. Indeed, it is one of the most intensively studied models in solid-state physics, making the investigation of its computational complexity interesting on its own. In the following, we show that computing its ground-state energy up to polynomial accuracy is complete for the complexity class QMA, the quantum analogue of NP. (See Fig. 1 for the relation of the relevant complexity classes.) A decision problem is in QMA if—although possibly hard to solve even by a quantum computer—every positive instance has a quantum proof that can be checked efficiently by a quantum computer. In particular, finding the ground-state energy of a local spin system with an accuracy polynomial in the lattice size is in QMA: the ground state serves as a proof, as expectation values of local Hamiltonians can be estimated efficiently. Conversely, it has been shown that any circuit verifying a QMA proof can be encoded as a ground-state problem[7,8]; that is, ground-state problems are QMA-complete. (A problem is called complete for a class if it is among the hardest problems in this class, that is, if any problem in the
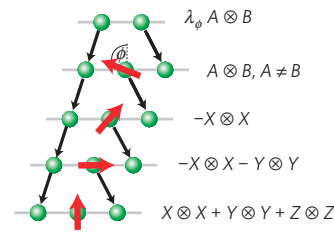
class can be reduced to it.) Using the same argument as before, finding the ground-state energy of the Hubbard model is inside QMA, because it can be mapped to a spin system through the Jordan–Wigner transform: this enables us to specify its ground state using spins, in such a way that it is possible to measure the ground-state energy efficiently[9].

In the following, we show that the Hubbard model with magnetic fields is also a hard problem for QMA, and thus QMA-complete. To this end, we start from a class of Hamiltonians for which finding the ground-state energy is known to be QMA-complete— that is, as hard as finding the ground-state energy of any local Hamiltonian—and show that this problem can be reduced to finding the ground-state energy of the Hubbard model with local magnetic fields. This is accomplished by a sequence of reductions, each of which reduces the previous Hamiltonian problem to a more restricted class of Hamiltonians. Each step makes use of perturbation-theory constructions (so-called gadgets), such that the original Hamiltonian arises as the effective low-energy theory of the new Hamiltonian[13,14].
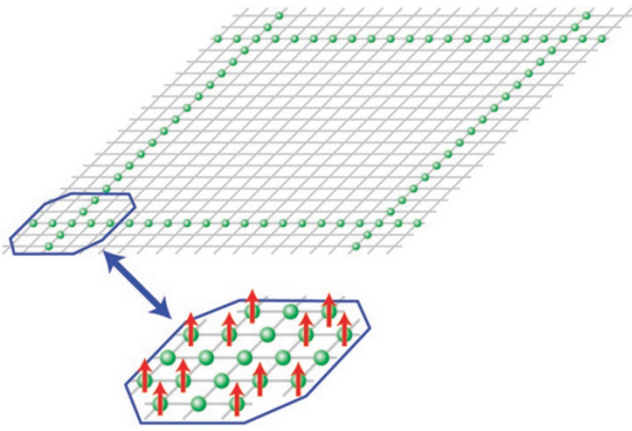
We start off with the Hamiltonian

$$H_{\text{Pauli}} = \sum_{\langle i,j \rangle} \lambda_{ij} A_{(ij)} \otimes B_{(ij)} \qquad (5)$$

defined on a 2D lattice with N spins, with $A_{(ij)}$ and $B_{(ij)}$ Pauli matrices and $|\lambda| \leq 1$, for which it has been proven that finding the ground-state energy up to a polynomial accuracy $1/q(N)$ is QMA-hard[10]. Following ref. 10, we call interactions of the form $\lambda_{ij} A_{(ij)} \otimes B_{(ij)}$ Pauli interactions. We first show how the Pauli Hamiltonian (5) can be reduced to the 2D Heisenberg lattice with local fields (see Supplementary Information for details). To this end, we use a chain of gadgets, all of which replace a two-qubit coupling by a chain of three qubits with a more restricted coupling. The idea is that by imposing a strong local field on the central ('mediator') qubit the system will essentially be in the ground state of the central qubit—but there will be second-order processes in which an excitation hops from the left qubit to the central one and then to the right, or vice versa, yielding an effective coupling between the outer qubits. (The excitation can also hop back and give an extra local term, which can however easily be compensated by adjusting the local magnetic field.) Note that similar gadgets have already been used, for example, in refs 10, 15.

The full sequence of reductions to the Heisenberg lattice is illustrated in Fig. 2. In a first step, we reduce arbitrary Pauli couplings $\lambda A \otimes B$ to Pauli couplings with constant λ and $A \neq B$. We illustrate this with a $\lambda Y_{l} \otimes Z_{r}$ coupling (we use X for σ^{x} and so on in the following), which is obtained from three qubits with couplings $Y_{l} \otimes X_{m} \otimes \mathbb{1}_{r} + \mathbb{1}_{l} \otimes Y_{m} \otimes Z_{r}$ by putting a strong field in the XY plane on the central qubit: a short calculation shows that this indeed gives a $Y_{l} \otimes Z_{r}$ coupling, where the strength is given by the angle in the XY plane. The intuition behind this is that, for example, a Y_{l} on the left

Figure 2 | Gadgets to reduce Pauli couplings to Heisenberg couplings. Each gadget works by inserting an extra spin in the middle, which is subject to a strong local field, yielding the desired interaction in second-order perturbation theory. A and B are Pauli matrices.

Figure 3 | The sparse Heisenberg lattice as obtained from H_{2D}, equation (5), using a sequence of gadgets. It can be reduced to a 2D Heisenberg lattice using the erasure gadget, where strong local fields are used to decouple unwanted qubits to leading order, as shown in the inset.

qubit can excite the central qubit, and the excitation then hops to the right qubit as a Z_r. In order for this hopping to be possible, the central field must not be along the X or Y axis, and correspondingly the hopping amplitude is controlled by the overlap of the field with X and Y eigenvectors.

The second gadget reduces Pauli couplings $A \otimes B$ ($A \neq B$) to Ising couplings. This is achieved by essentially the same gadget as before: for $X_l \otimes Y_r$, take $X_l \otimes X_m \otimes \mathbb{1}_r + \mathbb{1}_l \otimes Y_m \otimes Y_r$ and place a strong field in $X + Y$ direction on the m qubit. In a next gadget, Ising couplings are reduced to two $X \otimes X + Y \otimes Y$ couplings: placing a strong Y field on the central qubits only allows for the hopping of excitations through the $X \otimes X$ part of the coupling. Similarly, the above coupling is reduced to the Heisenberg interaction $X \otimes X + Y \otimes Y + Z \otimes Z$: a strong Z field prohibits hopping through the $Z \otimes Z$ term, whereas $X \otimes X + Y \otimes Y$ describes hopping in the $\{|01\rangle, |10\rangle\}$ subspace, which to second order yields the very same hopping term between the two outer qubits.

Putting these gadgets together, we have managed to reduce the QMA-complete Hamiltonian (5) to the Heisenberg Hamiltonian in a magnetic field on a sparse lattice. This can, in turn, be reduced to the full 2D Heisenberg lattice by using an 'erasure gadget', illustrated in Fig. 3: putting strong fields on the qubits to be erased decouples them up to polynomial precision. We have thus shown that the 2D Heisenberg Hamiltonian in a magnetic field on a 2D square lattice,

$$H_{\text{Heis}} = J \sum_{\langle i,j \rangle} \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j - \sum_i \mathbf{B}_i \cdot \boldsymbol{\sigma}_i \qquad (6)$$

is QMA-complete, both for $J > 0$ (which we use further on) and $J < 0$. Note that the presence of a magnetic field is crucial for the construction, as it is the only set of parameters available to encode a computational problem.

The final step is to reduce the Heisenberg lattice to the Hubbard model (4). The procedure can be found, for example, in ref. 6, and has been included in the Supplementary Information: in (4), we choose an on-site repulsion U very large as compared to t, and operate the system in the so-called half-occupancy regime where there are as many electrons as sites. (Note that this implies that a polynomial scaling in the lattice size is equal to a polynomial scaling in the number of electrons.) The tunnelling is suppressed as t/U, so that in the ground state each site will be occupied by exactly one electron, providing the desired spin degree of freedom. The coupling between the spins is achieved by a second-order process where one electron tunnels to an adjacent site, interacts

with the other electron, and tunnels back. However, this can only take place if the spins form a singlet, giving rise to the effective Hamiltonian (6) up to a constant. As the process is of second order, we have that $J = t^2/U > 0$, and the error from higher-order processes is $O(N^3 t^3 / U^2)$; thus, U/t has to grow polynomially with the system size N.

All these gadgets can be combined straightforwardly: first, the gadgets in one layer do not interact, as they never share a coupling term. It can thus be checked straightforwardly that, for second-order perturbation theory, there will be no cross-talk between the gadgets. Second, all gadget layers can be applied one after another, as long as the total strength of the previous gadgets is sufficiently smaller than the strong local fields of the new gadgets. As the number of layers is constant, this can be achieved by choosing polyscale field strengths, which allows for a polynomial scaling of the interaction strength as well as an arbitrary polynomial precision in energy. (See Supplementary Information for details.)

Let us now turn our attention back to DFT and the problem of interacting electrons. As we show in the Supplementary Information, the Hubbard model (4) with arbitrary local fields arises from (1) for an appropriately chosen V. For this particular potential, we can explicitly write down the wavefunction $w_i(r)$ of each mode a_i to sufficient precision. Thus, the ground-state wavefunction of the Hubbard model is supported by the $w_i(r)$, and consequently the single-electron density for the ground state of the Hubbard model must be of the form $\rho(r) = \sum \lambda_{i,s,s'} |w_i(r)|^2 |s\rangle \langle s'|$. As the functional (3) is convex and the physical $\lambda_{i,s,s'}$ form a convex set in $\mathbb{R}^{4N}$, the minimization in equation (2) (up to polynomial accuracy) can be carried out efficiently[16]; that is, finding the ground-state energy of the Hubbard model is in P. This implies that $P^{\text{UF}} = \text{QMA}$ with UF an oracle for the universal functional, that is, computing the functional is QMA-hard under Turing reductions.

Let us note that there are alternative ways to define $F[\rho]$, for example as the minimum over all pure N-electron states[11,12], in which case F is not convex. Yet, efficient computability (or even certifiability) of F would still imply that we could give a certificate for the ground-state energy, that is, QMA would collapse to NP. This is considered very unlikely, thus implying that any reasonably defined F cannot be computed in NP.

Finally, DFT can also be based on a functional defined on two-electron densities, which can be computed efficiently[11,12]. In this case, the QMA-hardness of the problem arises from the fact that characterizing the set of allowed two-electron reduced states, the N-representability problem, is QMA-hard[9]; in fact, this provides an alternative proof of its hardness.

References
1. Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136,** B864–B871 (1964).
2. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140,** A1133–A1138 (1965).
3. Slater, J. C. A simplification of the Hartree–Fock method. *Phys. Rev.* **81,** 385–390 (1951).
4. Hubbard, J. Electron correlations in narrow energy bands. *Proc. R. Soc. Lond. A* **276,** 238–257 (1963).
5. Hubbard, J. Electron correlations in narrow energy bands. II. The degenerate band case. *Proc. R. Soc. Lond. A* **277,** 237–259 (1964).
6. Auerbach, A. *Interacting Electrons and Quantum Magnetism* (Springer, 1994).
7. Kitaev, A. Y., Shen, A. H. & Vyalyi, M. N. *Classical and Quantum Computation* (American Mathematical Society, 2002).
8. Aharonov, D. & Naveh, T. Quantum NP—a survey. Preprint at <http://arxiv.org/abs/quant-ph/0210077> (2002).
9. Liu, Y.-K., Christandl, M. & Verstraete, F. Quantum computational complexity of the n-representability problem: QMA complete. *Phys. Rev. Lett.* **98,** 110503 (2007).

10. Oliveira, R. & Terhal, B. M. The complexity of quantum spin systems on a two-dimensional square lattice. *Quant. Inf. Comput.* **8,** 900–924 (2008).
11. Parr, R. G. & Yang, W. *Density-Functional Theory of Atoms and Molecules* (Oxford Univ. Press, 1989).
12. Dreizler, R. M. & Gross, E. K. U. *Density Functional Theory* (Springer, 1990).
13. Kempe, J., Kitaev, A. & Regev, O. The complexity of the local Hamiltonian problem. *SIAM J. Comp.* **35,** 1070–1097 (2006).
14. Bravyi, S., DiVincenzo, D. P., Loss, D. & Terhal, B. M. Simulation of many-body Hamiltonians using perturbation theory with bounded-strength interactions. *Phys. Rev. Lett.* **101,** 070503 (2008).
15. Biamonte, J. D. & Love, P. J. Realizable Hamiltonians for universal adiabatic quantum computers. *Phys. Rev. A* **78,** 012352 (2008).
16. Grötschel, M., Lovász, L. & Schrijver, A. *Geometric Algorithms and Combinatorial Optimization* (Springer, 1988).

Author contributions

All authors have contributed equally to this paper.

Additional information

Supplementary information accompanies this paper on www.nature.com/naturephysics. Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions. Correspondence and requests for materials should be addressed to N.S. or F.V.