N-representability is QMA-complete

Yi-Kai Liu
Institute for Quantum Information
Caltech
(joint work w/ Matthias Christandl and Frank Verstraete)

Computational complexity of the N-representability problem

- Is it easy or hard?

- A theoretical study
  - Fundamental questions
    - Can we solve it exactly, or only approximately?
    - Can we solve it always, or only sometimes?
    - Theory complements practice

Computational complexity of the N-representability problem

- N-representability is QMA-complete
  - Solving it exactly in the worst case is “hard”

- Perhaps we can say more…
  - It’s also related to many other problems
    => ideas for solving it
    - p-positivity conditions are related to SDP hierarchies
      for Vertex Cover, Max Cut, etc.
  - Complexity theory is a unifying tool
    => don’t study a problem in isolation

This talk

- Complexity theory
  - NP-completeness, approximation algorithms
  - QMA-completeness

- N-representability is QMA-complete
  - Efficient algorithm for N-representability => efficient algorithm for any QMA problem
  - Technical tools:
    - Convex optimization using a membership oracle (a.k.a., the shallow-cut ellipsoid method)
    - Mapping from qubits to fermions

Complexity theory

PH, #P, PSPACE, …

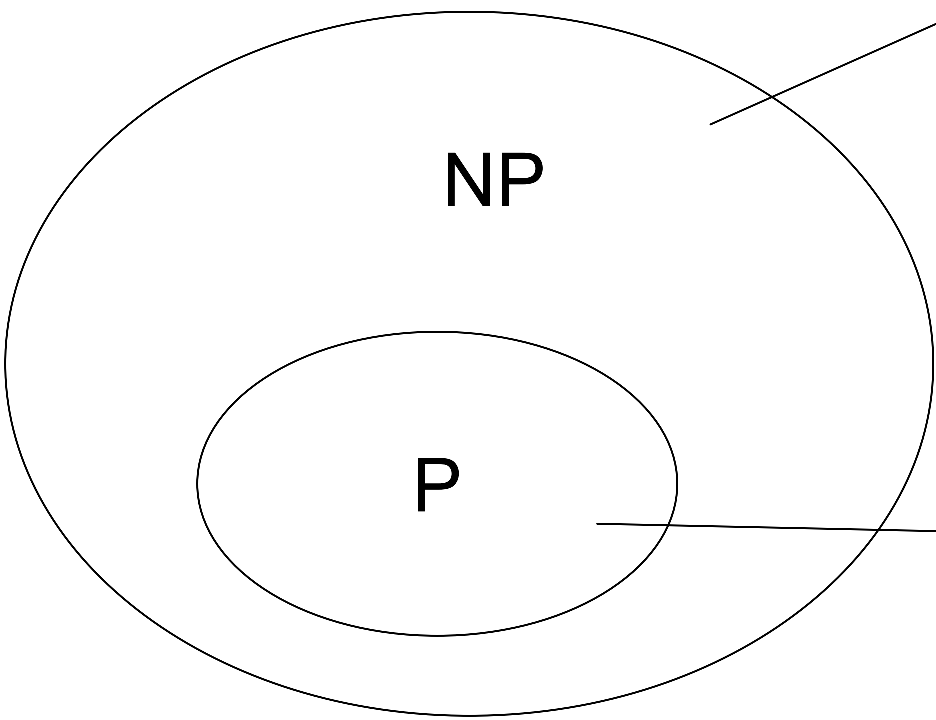

Problems that can be solved in **nondeterministic polynomial time**

Problems that can be solved in **polynomial time** = "easy"

Complexity theory

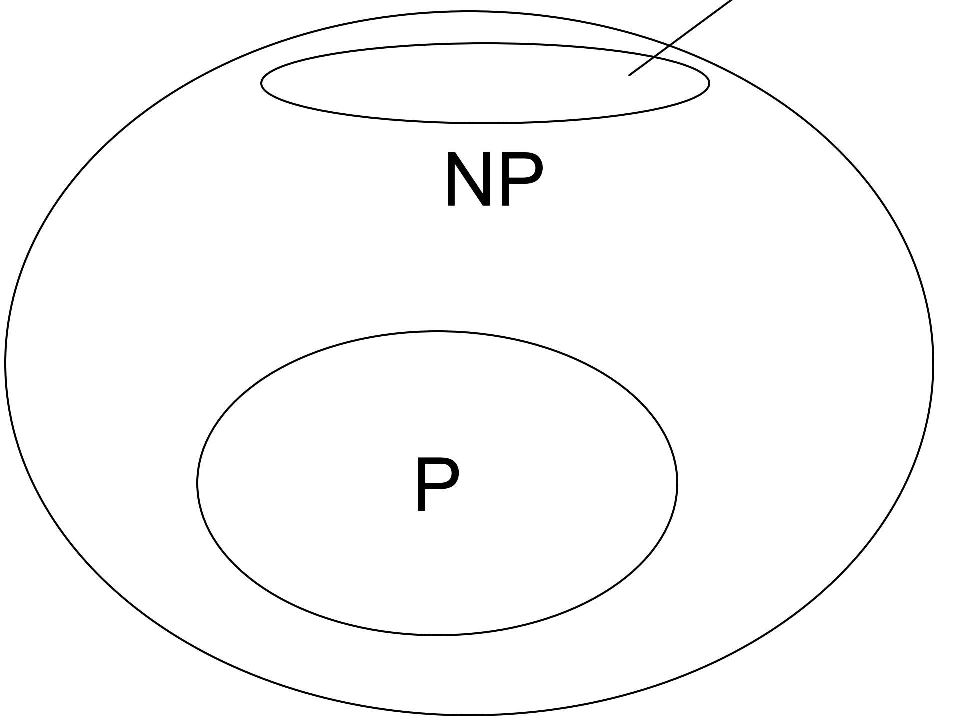

NP-complete problems: at least as hard as any other problem in NP

Cannot be solved efficiently, unless P=NP

Examples: traveling salesman problem, graph coloring, vertex cover, satisfiability of Boolean formulas, …

NP-completeness

- Note about terminology
  - **NP-hard** = as hard as any problem in NP
  - **NP-complete** = belongs to NP, and is NP-hard
- A very successful theory
  - Describes the complexity of many problems (either in P, or NP-complete)
  - Problems in P are not always easy in practice, but at least there is hope
  - NP-complete problems are not always hard in practice, but there exist hard instances

NP-completeness

- Still, some NP-complete problems seem harder than others…

Problem	**Practical experience**	**Complexity theory**
Traveling salesman	Can usually find near-optimal solutions using simple methods, e.g., local search	
Graph coloring	Hard to find high-quality solutions	

NP-completeness

- Still, some NP-complete problems seem harder than others…

Problem	**Practical experience**	**Complexity theory**
Traveling salesman	Can usually find near-optimal solutions using simple methods, e.g., local search	Good **approximation algorithms**, e.g., polynomial-time approximation schemes (PTAS)
Graph coloring	Hard to find high-quality solutions	Finding solutions within a constant factor of optimal is NP-hard! (the **PCP theorem**)

Complexity theory

- Still, there are some things we don’t understand…
- **Average-case** complexity
  - How hard is a problem for “typical” instances chosen from some distribution?
  - We can study **random** instances of a problem…
  - But most of the time, we don’t even know how to ask the right question
  - What is the distribution of instances that appear in practice?
- **Average-case** performance of algorithms
  - Theory: what is the distribution?
  - Experiment: do the results generalize?

The class NP ("nondeterministic polynomial time")

- A problem is in NP if its solution can be "verified" efficiently
  - There exists a poly-time algorithm V:

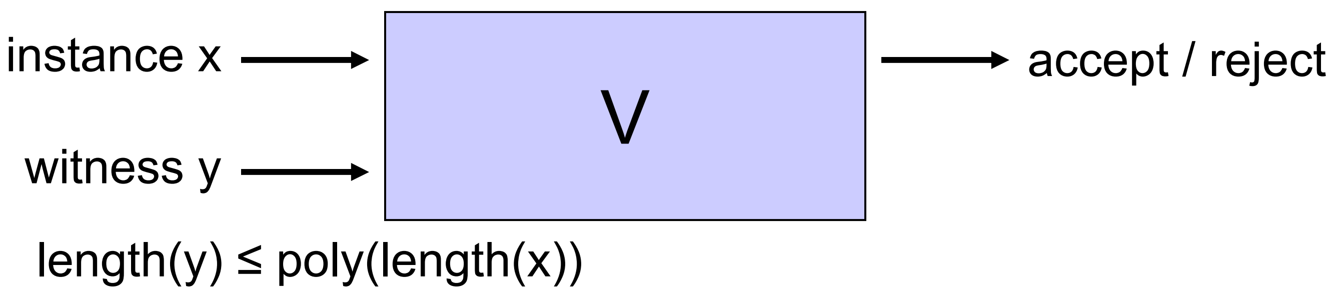

  - If the answer is "YES," there exists some y such that V accepts
  - If the answer is "NO," then for all y, V rejects

Ground states of local Hamiltonians

- The Local Hamiltonian problem
  - Given a local Hamiltonian $H = H_1 + \ldots + H_m$ for a system of n qubits
  - Given real numbers a, b, with $b - a \geq 1/\text{poly}(n)$
  - If ground state energy of H is $\leq a$, answer "YES"
  - If ground state energy of H is $\geq b$, answer "NO"

- Local Hamiltonian is not in NP
  - But it would be if we allowed quantum witnesses…

The class QMA ("quantum Merlin-Arthur")

- A problem is in QMA if:
  - There exists a poly-time **quantum** algorithm V:

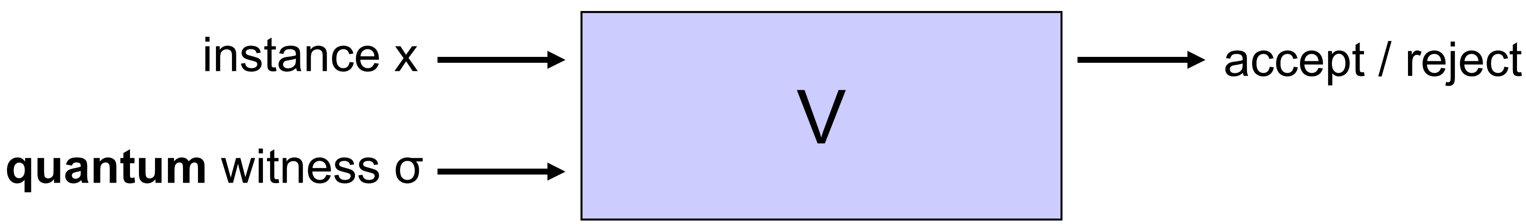

number of qubits in σ is ≤ poly(length(x))

  - If the answer is "YES," there exists some σ such that, **with prob. ≥ 2/3**, V accepts
  - If the answer is "NO," then for all σ, **with prob. ≥ 2/3**, V rejects

Local Hamiltonian is QMA-complete

- Local Hamiltonian is in QMA
- Local Hamiltonian is QMA-hard
  - Given an oracle that solves Local Hamiltonian, can solve any other QMA problem in polynomial time
  - Holds even for restricted classes of Hamiltonians:
    - 5-body terms [Kitaev ’99]
    - 2-body terms [Kempe, Kitaev & Regev ’04]
    - Nearest-neighbors on a 2D lattice [Oliveira & Terhal ’05]
    - 1D chain of qudits [Aharonov et al ’07]
    - and many more…

Our results – part I

- Consistency of local density matrices is QMA-complete [Liu '06]
  - Qubit version of N-representability

Consistency of local density matrices

- Consider a system of n qubits
- Given local density matrices $\rho_1,\ldots,\rho_m$
  - Describe subsets of qubits $C_1,\ldots,C_m$, $|C_i| \leq 2$
- Is there a global state σ (on all n qubits) that agrees with all of the ρ_i?
  - **YES:** there exists σ s.t., for all i, $\mathrm{tr}_{\{1..n\}\backslash C_i}(\sigma) = \rho_i$
  - **NO:** for all σ, there exists i s.t. $||\mathrm{tr}_{\{1..n\}\backslash C_i}(\sigma) - \rho_i||_1 \geq \beta$

$||A||_1 = \mathrm{tr}\,|A|$, and $\beta \geq 1/\mathrm{poly}(n)$

Suggested by D. Aharonov

Consistency is in QMA

- "Solution can be verified efficiently"
- Say $\rho_1,\ldots,\rho_m$ are consistent with some global state σ
  - Witness: many copies of σ
  - Verifier does measurements on the subset of qubits C_i
  - Need polynomially many measurements

- Suppose $\rho_1,\ldots,\rho_m$ are not consistent
  - Verifier should always reject
  - What if the witness is an arbitrary entangled state?
  - Verifier still won't accept; bound using Markov's inequality [Aharonov-Regev 2003]

Consistency is QMA-hard

- Given an oracle that solves Consistency, we can solve Local Hamiltonian efficiently
  - (as well as any other QMA problem)
- Idea: use convex optimization
  - Find local density matrices $\rho_1,\ldots,\rho_m$ that minimize $\mathrm{tr}(H_1\rho_1) + \ldots + \mathrm{tr}(H_m\rho_m)$, such that $\rho_1,\ldots,\rho_m$ are **consistent**
  - Use the oracle to test whether the consistency constraint is satisfied

Convex optimization using a membership oracle

- General problem:
  - K = convex set in R^n
  - Given an oracle O, where O(x) = [1 if x in K, and 0 o.w.]
  - Given a unit vector v in R^n, and a starting point p in K
  - Find x in K that minimizes $v \cdot x$

- Algorithms
  - Shallow-cut ellipsoid method [Yudin & Nemirovskii, 1978]
  - Random sampling algorithm [Bertsimas & Vempala, 2004]
  - Assume the set K is full-dimensional
  - Find exact solutions in polynomial time

Convex optimization using a membership oracle

- One novel feature: **approximate optimization**
  - Oracle makes errors of size ±1/poly(n), want to find a solution with similar accuracy
  - Need a stronger condition on the set K: assume that R/r ≤ poly(n)

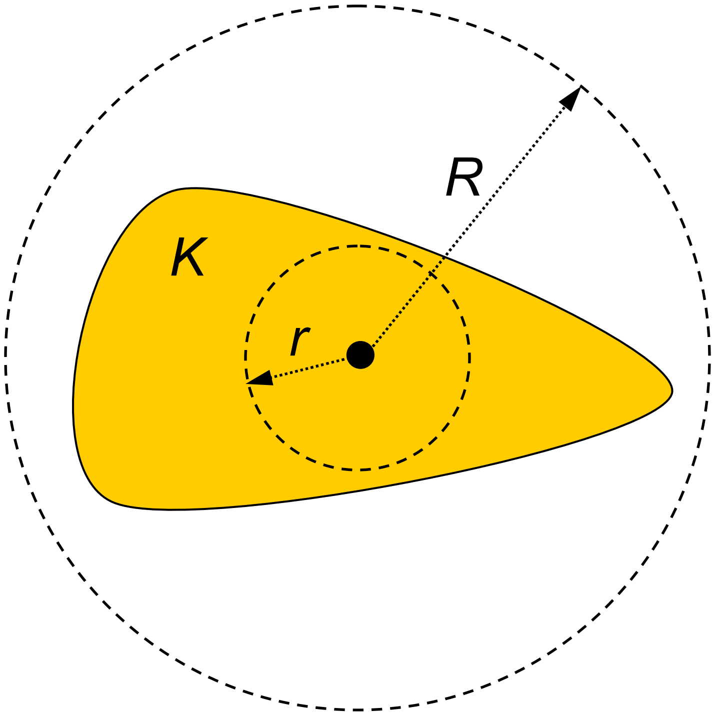

Consistency is QMA-hard

- Representing the set *K* of consistent local density matrices $(\rho_1,\ldots,\rho_m)$
  - Write down the expectation values of all Pauli operators on the subsets $C_1,\ldots,C_m$
  - Then *K* is full-dimensional and $R/r \leq \mathrm{poly}(n)$

- So an efficient algorithm for Consistency…
  - => an efficient algorithm to test membership in K
  - => an efficient algorithm to optimize over K
  - => an efficient algorithm to find ground state energies
  - => an efficient algorithm for Local Hamiltonian

Approximate optimization

- Precise statement of our result:
  - For any ε, there exists δ ≥ poly(1/n, ε, r/R), such that: if oracle has error ≤ δ, then algorithm has error ≤ ε
  - "logarithmically many bits of precision"

- The polynomial relating δ and ε is pretty awful…
  - But it's much better if one has a **separation oracle**
    - Returns a separating hyperplane when x not in K
    - Most known N-representability conditions also give this
  - **Finding separating hyperplanes** for Consistency / N-representability is QMA-hard for much larger δ

Related work

- Earlier work (using exact convex optimization)
  - Separability of bipartite quantum states is NP-hard [Gurvits ’02]
- Strong NP-hardness (using approximate optimization)
  - Separability is strongly NP-hard [Gharibian ’08]
  - Testing membership in the set of local operations w/ shared entanglement is strongly NP-hard [Gutoski ’09]
- Reverse reduction (Consistency ≤ Local Ham.)
  - Consistency is the dual problem to Local Ham. [Liu ’07]

Our results – part II

- N-representability is QMA-complete
  [Liu, Christandl and Verstraete ’06]
  - Fermionic analogue of the previous result

Ground states of molecules

- Fermionic Local Hamiltonian problem:
  - N electrons and d modes, $d \leq \text{poly}(N)$
  - Hamiltonian consists of identical pairwise interactions between electrons:

    $H = \sum_{1 \leq i < j \leq N} H^{(2)}_{ij}$

    $H^{(2)}$ has dimension (d-choose-2) x (d-choose-2)

  - Estimate the ground state energy with precision $\pm 1/\text{poly}(N)$

N-representability

- *N*-representability problem:
  - *N* electrons and *d* modes, $d \leq \text{poly}(N)$
  - Given a 2-electron density matrix ρ
  - If there exists an *N*-electron state σ s.t. $\text{tr}_{3,\ldots,N}(\sigma) = \rho$, answer “YES”
  - If for all *N*-electron states σ, $||\text{tr}_{3,\ldots,N}(\sigma) - \rho||_1 \geq \beta$, answer “NO”

  - Here, $||A||_1 = \text{tr}\ |A|$, and $\beta \geq 1/\text{poly}(N)$

These are QMA-complete

- These problems are in QMA
  - Witness is a fermionic state
  - Encode into qubits using the Jordan-Wigner transform

- These problems are QMA-hard
  - Qubit Local Ham. ≤ Fermionic Local Ham.
  - Use fermions to simulate qubits

  - Fermionic Local Ham. ≤ N-representability
  - Convex optimization w/ membership oracle

Qubit Local Ham. ≤ Fermionic Local Ham.

- Given a 2-local Hamiltonian on N qubits
- Construct a 2-local fermionic Hamiltonian
  - N fermions, $d = 2N$ modes
  - Qubit $\boldsymbol{i}$ is in state **|0)** => mode $\boldsymbol{a_i}$ is occupied
    Qubit $\boldsymbol{i}$ is in state **|1)** => mode $\boldsymbol{b_i}$ is occupied

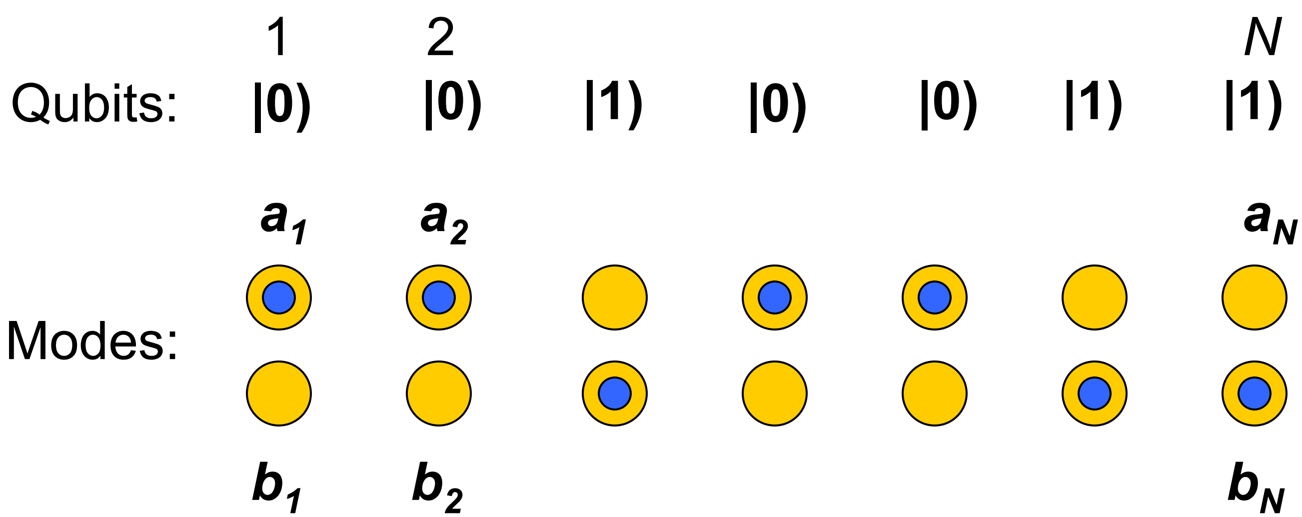

Qubit Local Ham. ≤ Fermionic Local Ham.

- Translate qubit operators into fermionic operators
  - σ_i^x becomes $a_i^+ b_i + b_i^+ a_i$
  - σ_i^y becomes $i\,(b_i^+ a_i - a_i^+ b_i)$
  - σ_i^z becomes $1 - 2\, b_i^+ b_i$

Operators on different qubits commute, since they are quadratic

2-local qubit operators become 2-local fermionic operators

- Want exactly one particle in each pair of modes a_i , b_i
  - Add terms of the form $\beta\,[(2\, a_i^+ a_i - 1)\,(2\, b_i^+ b_i - 1) + 1]$
  - These commute with the other terms in the Hamiltonian
  - Set β = a constant times the norm of the Hamiltonian

Fermionic Local Ham. ≤ N-representability

- Given a fermionic Hamiltonian $\boldsymbol{H} = \sum_{i \neq j} \boldsymbol{H}^{(2)}_{ij}$
- Want to estimate the ground state energy of ***H***, using an oracle for *N*-representability
  - Interesting behavior of ***H*** occurs in the subspace of *N*-particle states

- Solve this convex program:
  - Find a 2-electron density matrix $\boldsymbol{\rho}$ that minimizes $\mathbf{Tr}(\boldsymbol{H}^{(2)}\boldsymbol{\rho})$, such that $\boldsymbol{\rho}$ is *N*-representable

Fermionic Local Ham. ≤ N-representability

- Use second-quantized notation
  - $\boldsymbol{H}$ contains terms of the form $a_i^+ a_j$ and $a_i^+ a_j^+ a_l a_k$
  - 1- and 2-electron reduced density matrices:
  - $\rho^{(1)}_{ij} = N^{-1} \langle a_j^+ a_i \rangle$
  - $\rho^{(2)}_{ijkl} = (N(N-1))^{-1} \langle a_k^+ a_l^+ a_j a_i \rangle$

- Only consider states with exactly *N* particles
  - So we can write $a_i^+ a_j = (N-1)^{-1}\, a_i^+ (\Sigma_k\, a_k^+ a_k)\, a_j$

Fermionic Local Ham. ≤ N-representability

- Technical issues
  - Geometry of the convex set of N-representable states
  - Numerical precision
- Let $\boldsymbol{S}$ be a complete set of 2-electron observables
  - For all pairs of modes $I = \{i_1, i_2\}$, $i_1 < i_2$, define $a_I = a_{i2} a_{i1}$
  - List all pairs of modes in some order
  - Define the observables:
    $X_{IJ} = a_I^+ a_J + a_J^+ a_I$, for all $I < J$
    $Y_{IJ} = -i a_I^+ a_J + i a_J^+ a_I$, for all $I < J$
    $Z_I = a_I^+ a_I$, for all I except the last one
  - These operators are Hermitian, with eigenvalues in [-1,1]

Fermionic Local Ham. ≤ N-representability

- How to represent the 2-electron state ρ?
  - Take expectation values for the set of observables S
  - This is a vector α of dimension $\ell = |S| \leq \text{poly}(d)$

- K_N = {vectors α that are N-representable}
- Claim: K_N is contained in a ball of radius $\sqrt{\ell}$, and contains a ball of radius $1/\text{poly}(\ell)$

Fermionic Local Ham. ≤ N-representability

- K_N contains a ball of radius 1/poly(ℓ)

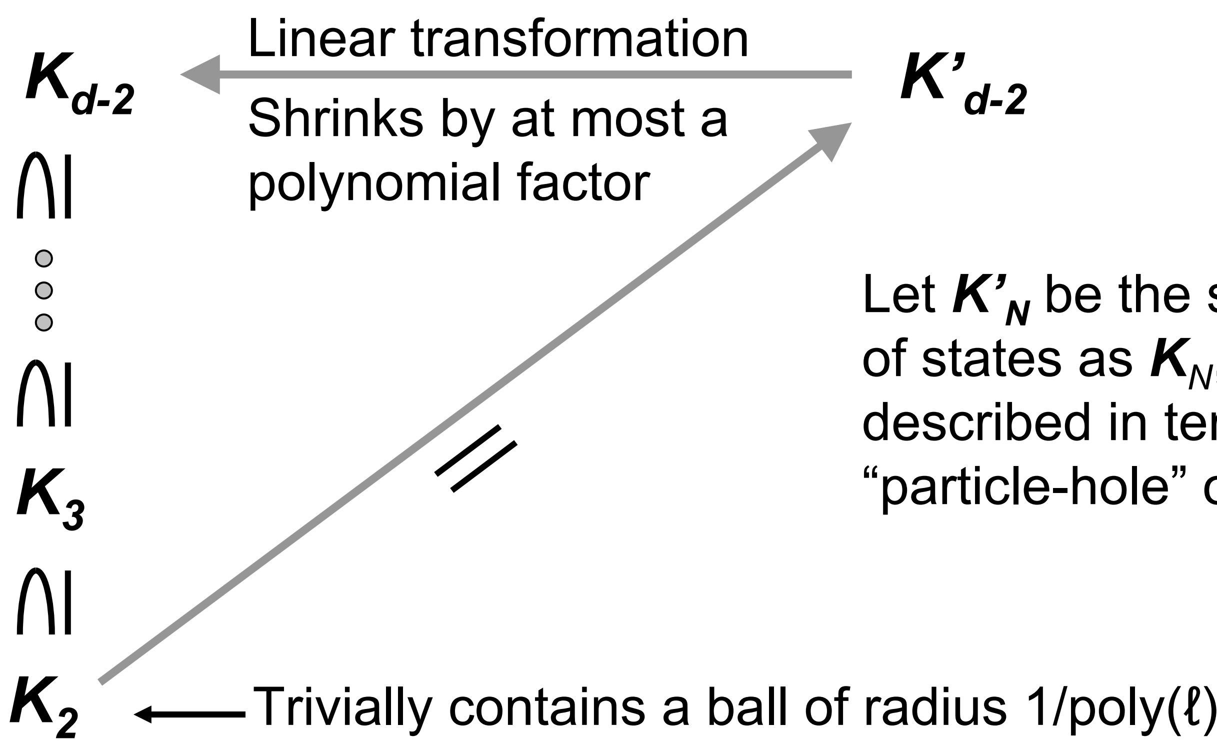

Let K'_N be the same set of states as K_N, but described in terms of "particle-hole" observables

Outlook…

- *N*-representability is hard in the worst case, asymptotically as the size of the problem grows

- However…
  - Molecular systems have special structure?
  - Can be solved by methods that exploit physical intuition?

- What about bosonic *N*-representability?
  - Also QMA-complete [Wei, Mosca & Nayak ’09]
- Reduced density matrices in translation-invariant systems?
  - Seems easier, see work by Verstraete et al…

p-positivity conditions and approximation algorithms

- *p*-positivity conditions [Erdahl, Jin, Mazziotti, … '00]
  - Hierarchy of constraints on the *p*-electron RDM
  - Approximating the *N*-representable set from the outside

- Approximation algorithms for combinatorial optimization, based on LP / SDP relaxations
  - K = convex hull of integer feasible solutions
  - Hierarchy of LP's / SDP's that approximate K from the outside [Lovasz-Schrijver, Sherali-Adams, Lasserre '90's]
    - Similarity with RDM methods? (Pironio et al)
  - Surprising negative results
    - Worst case: even going to the *n*'th level of the hierarchy is no good [Schoenebeck et al '07, Georgiou et al '07, Charikar et al '09, …]

References

- **Consistency and N-representability are QMA-complete**
  - Y.-K. Liu, *RANDOM* '06, Arxiv:quant-ph/0604166
  - Y.-K. Liu, M. Christandl and F. Verstraete, *PRL* 98, 110503 (2007), Arxiv:quant-ph/0609125
  - Y.-K. Liu, PhD Thesis, Arxiv:0712.3041
- **Complexity theory for the working physicist**
  - M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, 1979.
  - P. Crescenzi and V. Kann, *A compendium of NP optimization problems*, http://www.csc.kth.se/~viggo/problemlist/
- **Some of my more recent work**
  - **Quantum algorithms using the curvelet transform**, *STOC* 2009, Arxiv:0810.4968
  - **Low-rank quantum state tomography using compressed sensing** (w/ S. Becker, J. Eisert, S. Flammia, D. Gross, Y. Plan), in preparation