Approximation Guarantees for the Local Hamiltonian Problem

…and limitations for qPCPs

Fernando G.S.L. Brandão
ETH

Aram Harrow
MIT

QIP 2013, Beijing

Constraint Satisfaction Problems

k-arity CSP:

Variables $\{x_1, \ldots, x_n\}$ in Σ^n

Alphabet Σ

Constraints $\{c_j\}_j$

$$c_j : \Sigma^k \rightarrow \{0,1\}$$

Unsat := $\min_{x \in \Sigma^n} \sum_j c_j(x_{j_1}, \ldots, x_{j_k})$

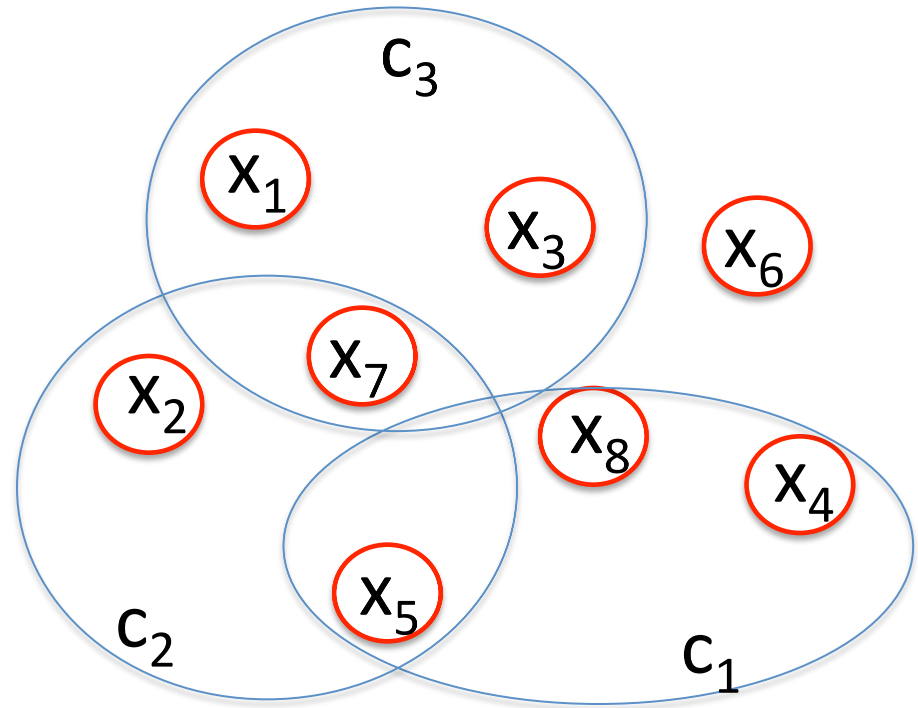

Include 3SAT, max-cut, vertex cover, …

NP-complete to compute Usant

Constraint Satisfaction Problems

…as an eigenvalue problem

Hamiltonian $H = \sum_{j=1}^{m} C_j \in \left(R^d\right)^{\otimes n},\ d = |\Sigma|$

Local Terms $C_j := \sum_{z \in \Sigma^k : c_j(z)=1} |z_1,\ldots,z_k\rangle\langle z_1,\ldots,z_k|$

Unsat = minimum eigenvalue of H

(Hamiltonian for classical spins: Ising model, Pott's model)

Quantum CSPs, aka Local Hamiltonians

k-local Hamiltonian: $H = \sum_{i=1}^{m} H_i \in \left(C^d\right)^{\otimes n}$

Local Terms: $H_i = H_{i_1,\ldots,i_k} \otimes I_{rest}, \;\; H_{i_1,\ldots,i_k} \in Herm(C^{\otimes k})$

qUnsat = $E_0(H)$: E_0 : minimum eigenvalue

Optimal assignment: Groundstate of the model

How hard are qCSP?

Quantum Hamiltonian Complexity addresses this question

The Local Hamiltonian Problem

Problem

Given a local Hamiltonian H, decide if $E_0(H)=0$ or $E_0(H)>\Delta$

$E_0(H)$: minimum eigenvalue of H

The Local Hamiltonian Problem

Problem

Given a local Hamiltonian H, decide if $E_0(H)=0$ or $E_0(H)>\Delta$

$E_0(H)$: minimum eigenvalue of H

Thm (Kitaev '99) The local Hamiltonian problem is QMA-complete for $\Delta = 1/poly(n)$

(analogue Cook-Levin thm)

QMA is the quantum analogue of NP, where the proof and the computation are quantum

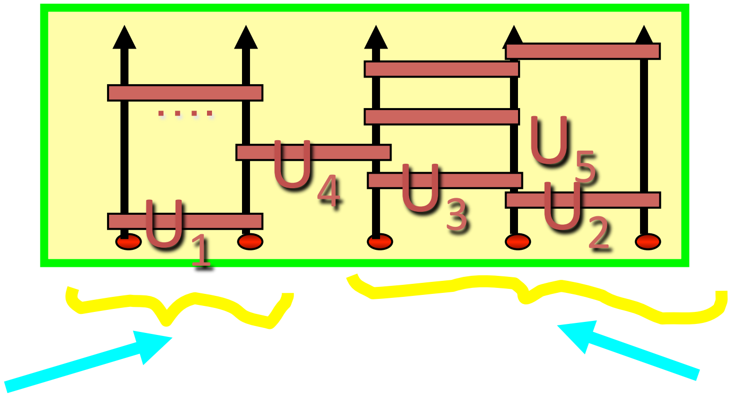

Input

Witness

The meaning of it

It's believed QMA ≠ NP

Thus there is generally no efficient classical description of groundstates of local Hamiltonians

(Even very simple models are QMA-complete
(Aharonov, Gottesman, Irani, Kempe '07) 1D,
(Oliveira, Terhal '05) 2D on qubits, …)

The meaning of it

It's believed QMA ≠ NP

Thus there is generally no efficient classical description of groundstates of local Hamiltonians

(Even very simple models are QMA-complete
(Aharonov, Gottesman, Irani, Kempe '07) 1D,
(Oliveira, Terhal '05) 2D on qubits, ...)

What's the role of the promise gap Δ on the hardness?

.... But first, what happens for CSP?

PCP Theorem

PCP Theorem (Arora et al ’98, Dinur ‘07): There is a $\varepsilon > 0$ s.t. it’s NP-complete to determine whether for a CSP with m constraints, $\text{Unsat} = 0$ or $\text{Unsat} > \varepsilon m$

PCP Theorem

PCP Theorem (Arora et al '98, Dinur '07): There is a $\varepsilon > 0$ s.t. it's NP-complete to determine whether for a CSP with m constraints, $\text{Unsat} = 0$ or $\text{Unsat} > \varepsilon m$

- NP-hard even for $\Delta = \Omega(m)$

PCP Theorem

PCP Theorem (Arora et al '98, Dinur '07): There is a $\varepsilon > 0$ s.t. it's NP-complete to determine whether for a CSP with m constraints, $\text{Unsat} = 0$ or $\text{Unsat} > \varepsilon m$

- NP-hard even for $\Delta=\Omega(m)$

- Equivalent to the existence of **P**robabilistically **C**heckable **P**roofs for NP.

PCP Theorem

PCP Theorem (Arora et al '98, Dinur '07): There is a $\varepsilon > 0$ s.t. it's NP-complete to determine whether for a CSP with m constraints, $\text{Unsat} = 0$ or $\text{Unsat} > \varepsilon m$

- NP-hard even for $\Delta=\Omega(m)$
- Equivalent to the existence of **P**robabilistically **C**heckable **P**roofs for NP.
- Central tool in the theory of hardness of approximation (optimal threshold for 3-SAT ($7/8$-factor), max-clique ($n^{1-\varepsilon}$-factor))

Quantum PCP?

The qPCP conjecture: There is $\varepsilon > 0$ s.t. the following problem is QMA-complete: Given 2-local Hamiltonian H with m local terms determine whether

(i) $E_0(H)=0$ or (ii) $E_0(H) > \varepsilon m$.

Quantum PCP?

The qPCP conjecture: There is $\varepsilon > 0$ s.t. the following problem is QMA-complete: Given 2-local Hamiltonian H with m local terms determine whether

(i) $E_0(H)=0$ or (ii) $E_0(H) > \varepsilon m$.

- (Bravyi, DiVincenzo, Loss, Terhal '08) Equivalent to conjecture for O(1)-local Hamiltonians over qdits.

Quantum PCP?

The qPCP conjecture: There is $\varepsilon > 0$ s.t. the following problem is QMA-complete: Given 2-local Hamiltonian H with m local terms determine whether

(i) $E_0(H)=0$ or (ii) $E_0(H) > \varepsilon m$.

- (Bravyi, DiVincenzo, Loss, Terhal '08) Equivalent to conjecture for $O(1)$-local Hamiltonians over qdits.

- Equivalent to estimating mean groundenergy to constant accuracy ($e_o(H) := E_0(H)/m$)

Quantum PCP?

The qPCP conjecture: There is $\varepsilon > 0$ s.t. the following problem is QMA-complete: Given 2-local Hamiltonian H with m local terms determine whether

(i) $E_0(H)=0$ or (ii) $E_0(H) > \varepsilon m$.

- (Bravyi, DiVincenzo, Loss, Terhal '08) Equivalent to conjecture for O(1)-local Hamiltonians over qdits.
- Equivalent to estimating mean groundenergy to constant accuracy ($e_o(H) := E_0(H)/m$)
- And to estimate the energy at constant temperature

Quantum PCP?

The qPCP conjecture: There is $\varepsilon > 0$ s.t. the following problem is QMA-complete: Given 2-local Hamiltonian H with m local terms determine whether

(i) $E_0(H)=0$ or (ii) $E_0(H) > \varepsilon m$.

- (Bravyi, DiVincenzo, Loss, Terhal '08) Equivalent to conjecture for O(1)-local Hamiltonians over qdits.
- Equivalent to estimating mean groundenergy to constant accuracy ($e_o(H) := E_0(H)/m$)
- And to estimate the energy at constant temperature
- At least NP-hard (by PCP Thm) and in QMA

Quantum PCP?

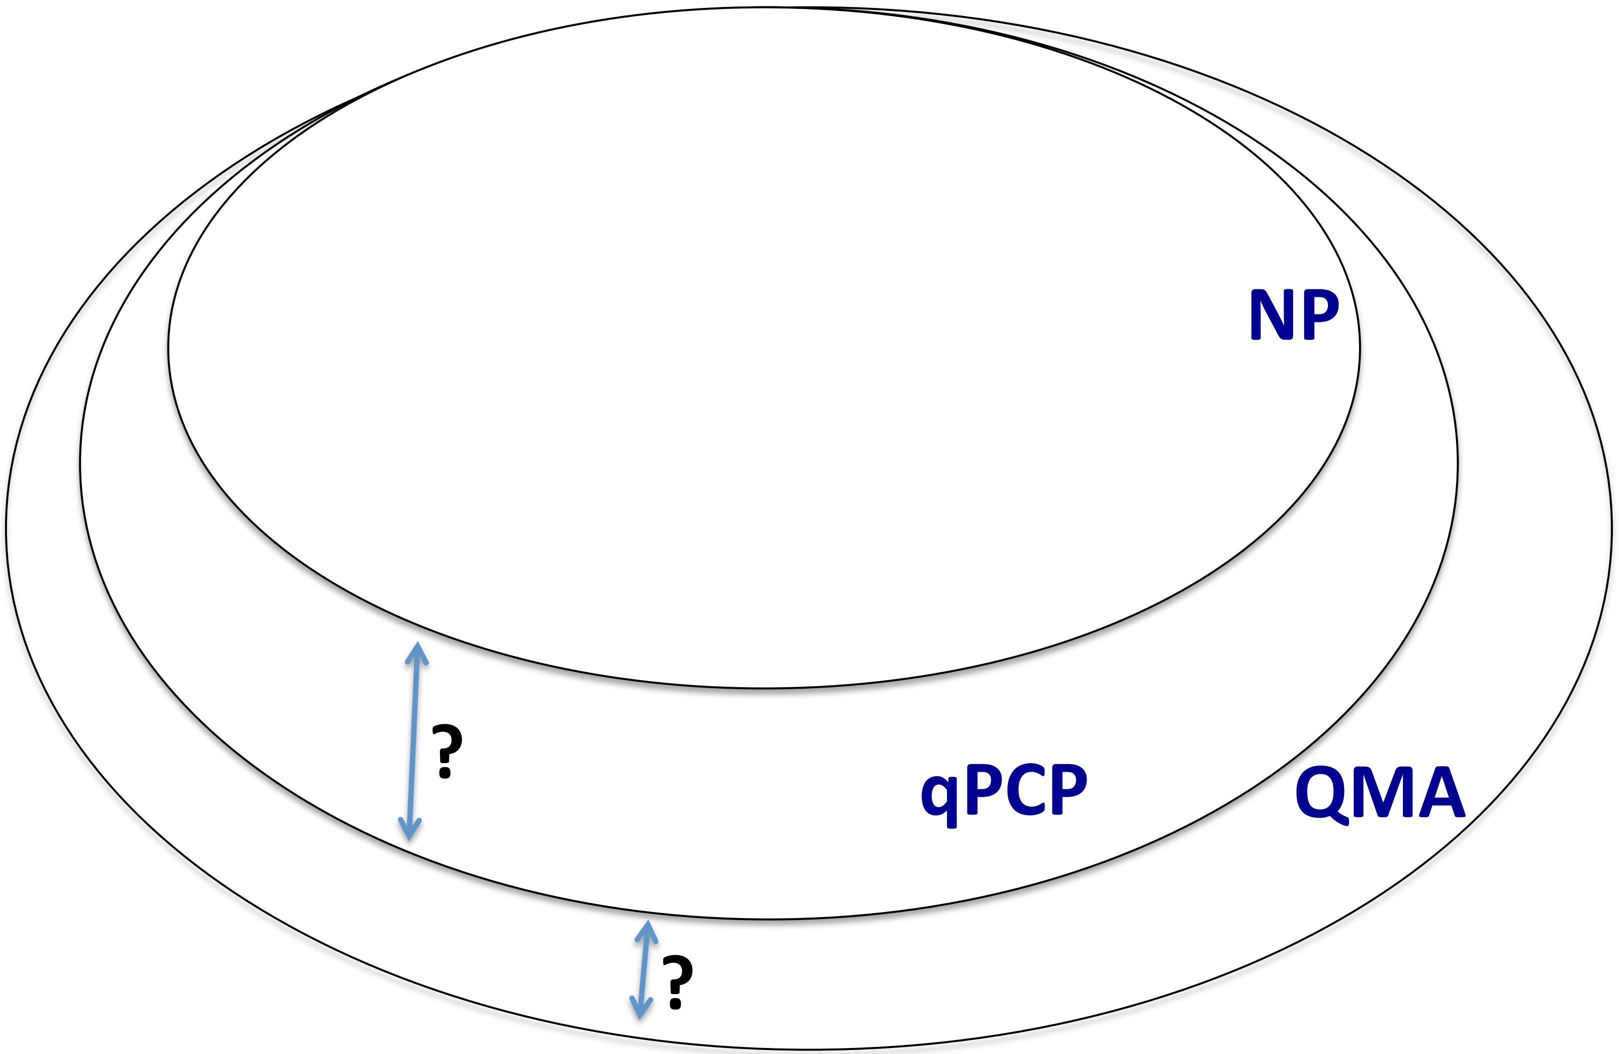

Previous Work and Obstructions

(Aharonov, Arad, Landau, Vazirani '08)
Quantum version of 1 of 3 parts of Dinur's proof of the PCP thm (gap amplification)

But: The other two parts (alphabet and degree reductions) involve massive copying of information; not clear how to do it with a highly entangled assignment

Previous Work and Obstructions

(Aharonov, Arad, Landau, Vazirani '08)
Quantum version of 1 of 3 parts of Dinur's proof of the PCP thm (gap amplification)

But: The other two parts (alphabet and degree reductions) involve massive copying of information; not clear how to do it with a highly entangled assignment

(Bravyi, Vyalyi '03; Arad '10; Hastings '12; Freedman, Hastings '13; Aharonov, Eldar '13, …)
No-go for large class of *commuting* Hamiltonians and almost commuting Hamiltonians

But: Commuting case might always be in NP

Approximation in NP

(B., Harrow '12) Let H be a 2-local Hamiltonian on qudits with interaction graph G(V, E) and |E| local terms.

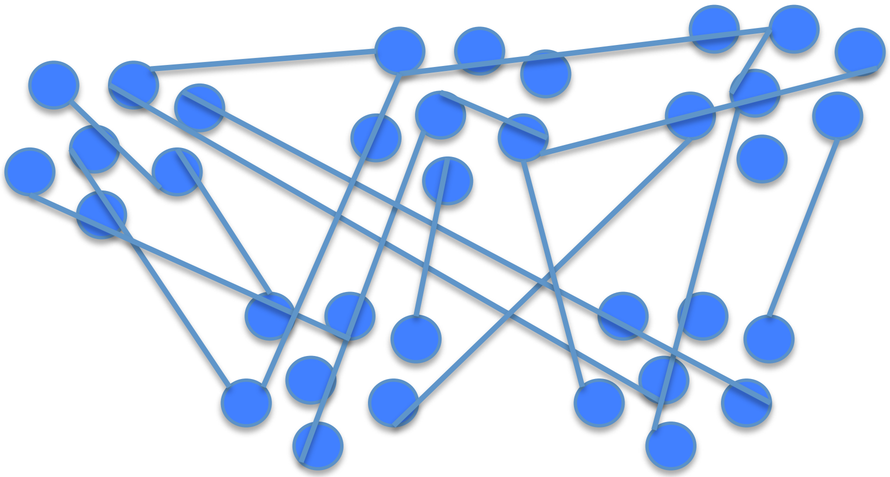

Approximation in NP

(B., Harrow '12) Let H be a 2-local Hamiltonian on qudits with interaction graph $G(V, E)$ and $|E|$ local terms.

Let $\{X_i\}$ be a partition of the sites with each X_i having m sites.

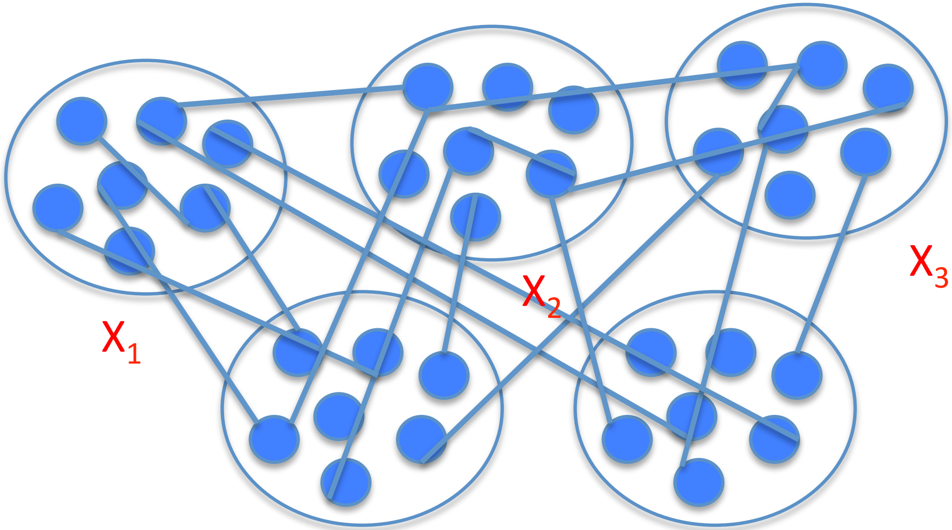

Approximation in NP

(B., Harrow '12) Let H be a 2-local Hamiltonian on qudits with interaction graph G(V, E) and |E| local terms.

Let $\{X_i\}$ be a partition of the sites with each X_i having m sites. Then there are products states ψ_i in X_i s.t.

$$\frac{1}{|E|}\langle \psi_1,\ldots,\psi_m | H | \psi_1,\ldots,\psi_m \rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \frac{1}{\deg(G)} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

E_i : expectation over X_i

deg(G) : degree of G

$\Phi(X_i)$: expansion of X_i

$S(X_i)$: entropy of groundstate in X_i

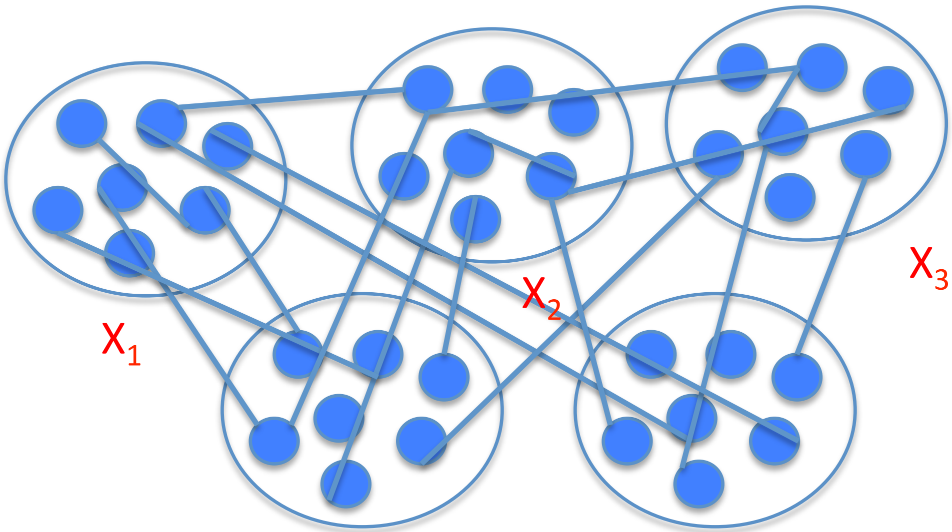

Approximation in NP

(B., Harrow '12) Let H be a 2-local Hamiltonian on qudits with interaction graph G(V, E) and |E| local terms.

Let $\{X_i\}$ be a partition of the sites with each X_i having m sites.
Then there are products states ψ_i in X_i s.t.

$$\frac{1}{|E|}\left\langle \psi_1,\ldots,\psi_m \left| H \right| \psi_1,\ldots,\psi_m \right\rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \frac{1}{\deg(G)} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

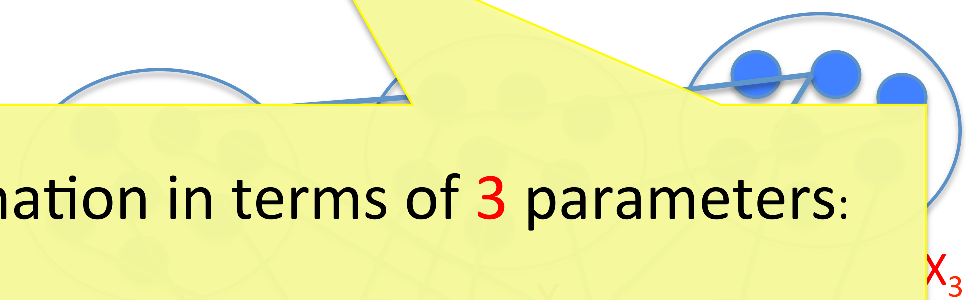

E_i : expectation over X_i
deg(G) : degree of G
$\Phi(X_i)$: expansion of X_i
$S(X_i)$: entropy of groundstate in X_i

Approximation in terms of 3 parameters:

1. Average expansion
2. Degree interaction graph
3. Average entanglement groundstate

1. Approximation in terms of average expansion

$$\frac{1}{|E|}\left\langle \psi_1,\ldots,\psi_m \middle| H \middle| \psi_1,\ldots,\psi_m \right\rangle \leq e_0(H) + \Omega\left(d^6 \boxed{E_i \Phi(X_i)} \frac{1}{\deg(G)} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

Average Expansion: $E_i \Phi(X_i) = E_i \Pr_{(u,v)\in E}\left(v \notin X_i \mid u \in X_i \right)$

1. Approximation in terms of average expansion

$$\frac{1}{|E|}\left\langle \psi_1,\ldots,\psi_m \middle| H \middle| \psi_1,\ldots,\psi_m \right\rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \frac{1}{\deg(G)} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

Average Expansion: $E_i \Phi(X_i) = E_i \Pr_{(u,v)\in E}\left(v \notin X_i \mid u \in X_i \right)$

Well known fact: 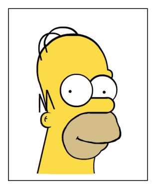 's divide and conquer

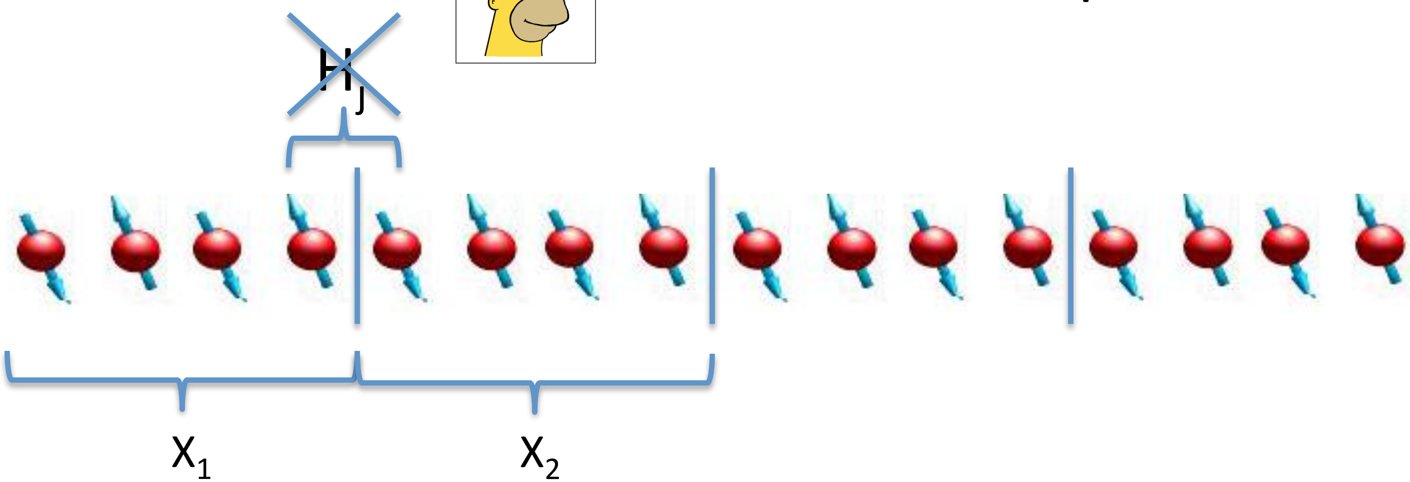

Potential hard instances must be based on highly expanding graphs

2. Approximation in terms of degree

$$\frac{1}{|E|}\langle \psi_1,\ldots,\psi_m | H | \psi_1,\ldots,\psi_m \rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \boxed{\frac{1}{\deg(G)}} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

More surprising, no classical analogue:

(PCP + parallel repetition) For all α, β, γ > 0 it's NP-complete to determine whether a CSP *C* is s.t.

Unsat = 0 or Unsat > α $\Sigma^{\beta}/\deg(G)^{\gamma}$

Parallel repetition: *C* -> *C'*

i. deg(G') = $\deg(G)^k$
ii. Σ' = Σ^k
ii. Unsat(G') > Unsat(G)

(Raz '00) even showed Unsat(G') approaches 1 exponentially fast

2. Approximation in terms of degree

$$\frac{1}{|E|}\langle\psi_1,\ldots,\psi_m|H|\psi_1,\ldots,\psi_m\rangle \le e_0(H)+\Omega\left(d^6 E_i \Phi(X_i)\boxed{\frac{1}{\deg(G)}}E_i\frac{S(X_i)}{m}\right)^{1/8}$$

More surprising, no classical analogue:

(PCP + parallel repetition) For all α, β, γ > 0 it's NP-complete to determine whether a CSP *C* is s.t.

Unsat = 0 or Unsat > $\alpha\, \Sigma^{\beta}/\deg(G)^{\gamma}$

Contrast: It's in NP determine whether a Hamiltonian H is s.t

$E_0(H)=0$ or $E_0(H) > \alpha d^{3/4}/\deg(G)^{1/8}$

Quantum generalizations of PCP *and* parallel repetition *cannot* both be true (assuming QMA not in NP)

2. Approximation in terms of degree

$$\frac{1}{|E|}\langle \psi_1,\ldots,\psi_m | H | \psi_1,\ldots,\psi_m \rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \boxed{\frac{1}{\deg(G)}} E_i \frac{S(X_i)}{m} \right)^{1/8}$$

More surprising, no classical analogue:

(PCP + parallel repetition) For all α, β, γ > 0 it's NP-complete to determine whether a CSP *C* is s.t.

Unsat = 0 or Unsat > $\alpha\ \Sigma^{\beta}/\deg(G)^{\gamma}$

Bound: $\Phi_G < 1/2 - \Omega(1/\deg)$ implies

Highly expanding graphs (Φ_G -> 1/2) are not hard instances

2. Approximation in terms of degree

…shows mean field becomes exact in high dim

1-D

2-D

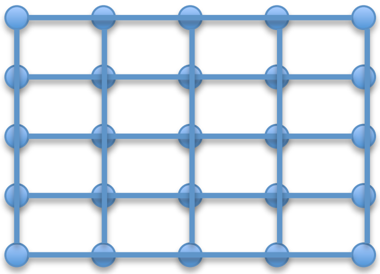

3-D

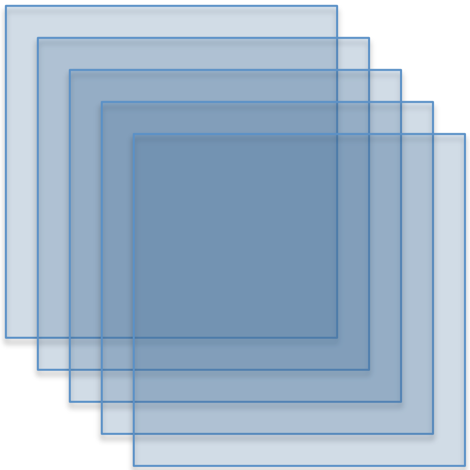

∞-D

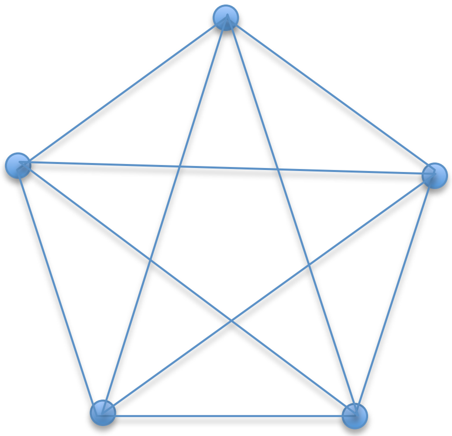

Rigorous justification to folklore in condensed matter physics

3. Approximation in terms of average entanglement

$$\frac{1}{|E|}\left\langle \psi_1,\ldots,\psi_m \middle| H \middle| \psi_1,\ldots,\psi_m \right\rangle \leq e_0(H) + \Omega\left(d^6 E_i \Phi(X_i) \frac{1}{\deg(G)} \boxed{E_i \frac{S(X_i)}{m}} \right)^{1/8}$$

The problem is in NP if entanglement of groundstate satisfy a subvolume law:

$$E_i \frac{S(X_i)}{m} = o(1)$$

Connection of amount of entanglement in groundstate and computational complexity of the model

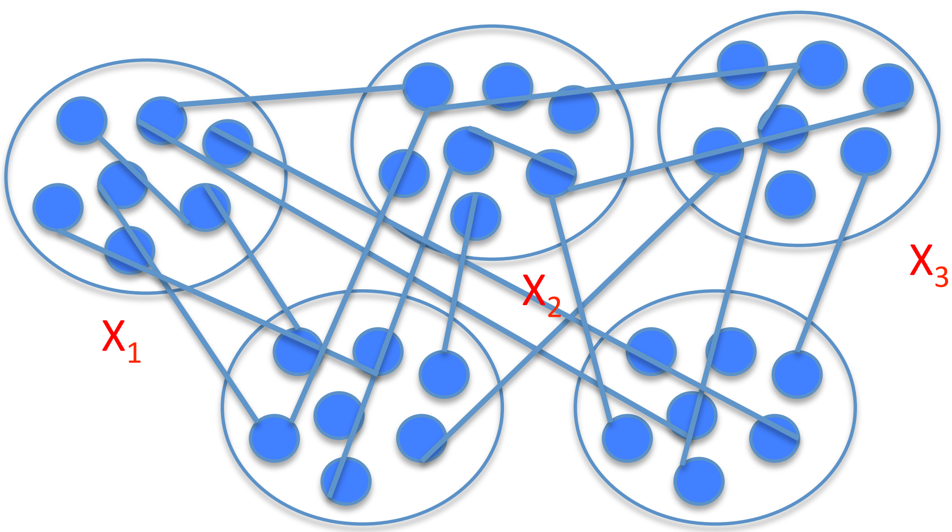

Intuition: Monogamy of Entanglement

Quantum correlations are non-shareable

e.g. $\left(|\uparrow\rangle|\downarrow\rangle - |\downarrow\rangle|\uparrow\rangle\right)/\sqrt{2}$

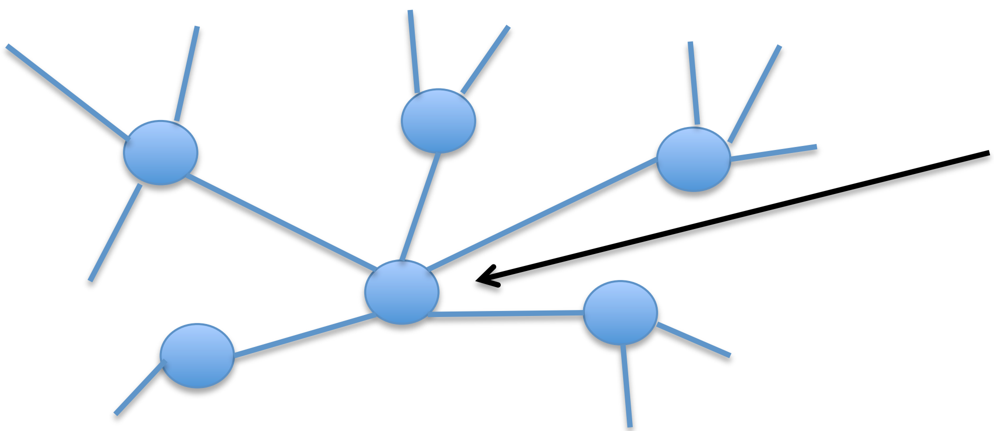

Cannot be highly entangled with too many neighbors

Tool: Information Theory

1. Mutual Information

$$I(X:Y)_p = D(p_{XY} \| p_X \otimes p_Y)$$

2. Pinsker's inequality

$$I(X:Y)_p = \frac{1}{2\ln 2}\left\| p_{XY} - p_X \otimes p_Y \right\|_1^2$$

3. Conditional Mutual Information

$$I(X:Y|Z) = I(X:YZ) - I(X:Z)$$

4. Chain Rule

$$I(X:Y_1 \ldots Y_k) = I(X:Y_1) + \ldots + I(X:Y_k | Y_1 \ldots Y_{k-1})$$

$$\Rightarrow I(X:Y_t | Y_1 \ldots Y_{t-1}) \le \log(\Sigma)/k$$ for some t<k

Conditioning Decouples

Idea that almost works (c.f. Raghavendra-Tan '11)

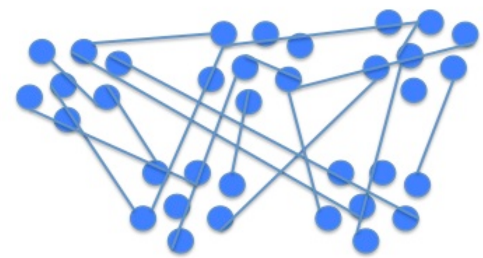

1. Choose $i, j_1, \ldots, j_k$ at random from $\{1, \ldots, n\}$
Then there exists $t<k$ such that

$$\mathop{E}_{i,j_1,\ldots,j_t} I(Z_i : Z_{j_t} \mid Z_{j_1} \ldots Z_{j_{t-1}}) \leq \frac{\log d}{k}$$

2. Conditioning on subsystems $j_1, \ldots, j_t$ causes error $<k/n$ and leaves a distribution q for which

$$\mathop{E}_{i,j} I(Z_i, Z_j)_q \leq \frac{\log d}{k} \quad \text{which implies} \quad \mathop{E}_{i \sim_G j} I(Z_i, Z_j)_q \leq \frac{n}{\deg(G)} \frac{\log d}{k}$$

By Pinsker's: $$\mathop{E}_{i \sim_G j} \left\| q_{Z_i Z_j} - q_{Z_i} \otimes q_{Z_j} \right\|_1 \leq \sqrt{\frac{1}{2\ln 2} \frac{n}{\deg(G)} \frac{\log d}{k}}$$

Does it Work Quantumly?

Good news:

- $I(A:B)$, $I(A:B:C)$ still defined
- Pinsker, chain rule, etc still hold
- $I(A:B|C)=0$ implies ρ_{AB} separable

Bad news:

- Can't condition on quantum info
- $I(A:B|C)\approx 0$ doesn't imply ρ_{AB} is close to separable in trace norm (Ibinson, Linden, Winter '08)

Good news we can use

Informatinally-complete measurement M satisfies

$$d^{-3} \| \rho - \sigma \|_1 \leq \| M(\rho) - M(\sigma) \|_1 \leq \| \rho - \sigma \|_1$$

Proof Overview

1. Measure εn qudits with M and condition on outcomes. Incur error ε.

2. Most pairs of other qudits would have mutual information $\leq \log(d) / \varepsilon \deg(G)$ if measured.

3. Thus their state is within distance $d^3(\log(d) / \varepsilon \deg(G))^{1/2}$ of product.

4. Witness is a global product state. Total error is $\varepsilon + d^6(\log(d) / \varepsilon \deg(G))^{1/2}$.
   Choose ε to balance these terms.

5. General case follows by coarse graining sites (and a few other tricks)

New Classical Algorithms for Quantum Hamiltonians

Following same approach we also obtain polynomial time algorithms for approximating the groundstate energy of

1. Planar Hamiltonians, improving on (Bansal, Bravyi, Terhal '07)
2. Dense Hamiltonians, improving on (Gharibian, Kempe '10)
3. Hamiltonians on graphs with low threshold rank, building on (Barak, Raghavendra, Steurer '10)

In all cases we prove that a product state does a good job (after coarse graining some of the sites) and use efficient algorithms for CSPs (Baker '94, Arora, Karger, Karpinski '95)

Similar techniques give new de Finetti thm for general quantum states

Conclusions

- Can approximate mean energy in terms of **degree** and **amount of entanglement:** Monogamy of entanglement in groundstates

- **Mean field exact** in the limit of **large dimensions**

- **No-go** against **qPCP + "quantum parallel repetition"**

- Tools from **information theory** are **useful**

Open Questions

- Go **beyond mean field**
- Is there a meaningful notion of **parallel repetition** for **qCSP**?
- Does every groundstate have **subvolume entanglement after constant-depth-circuit** renormalization?
- Find **more** classes of **Hamiltonians with efficient algorithms**
- (**dis**)**prove qPCP** conjecture!

Thank you!