Quantum Information Theory in Quantum Hamiltonian Complexity

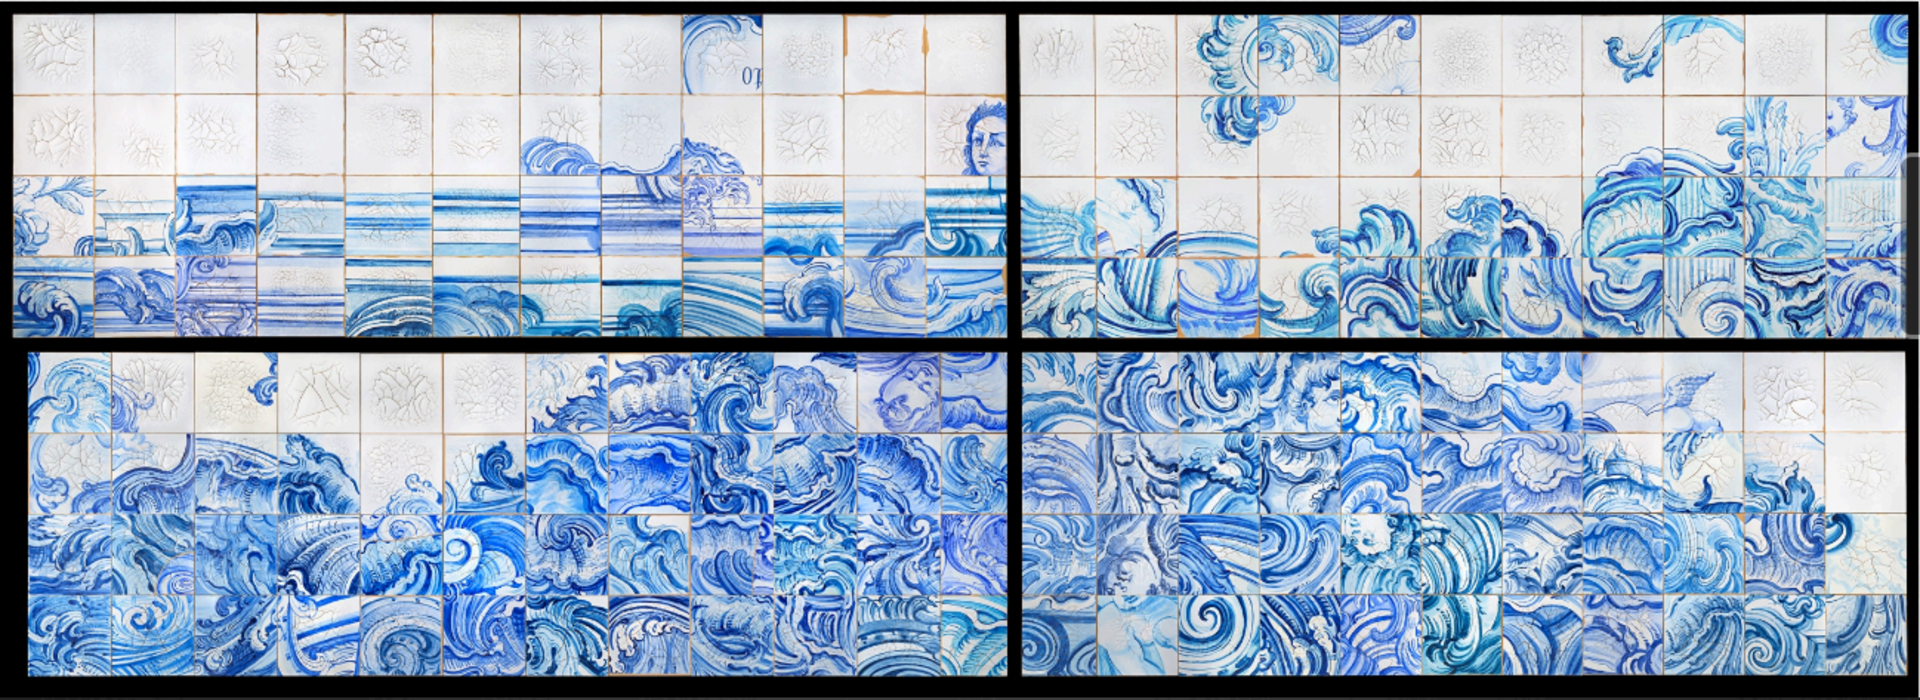

Aram Harrow (MIT)
Simons Institute
16 Jan 2014

Entanglement

Original motivation for quantum computing [Feynman '82]

Nature isn't classical, dammit, and if you want to make a simulation of Nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.

N systems in product state → O(N) degrees of freedom

N entangled systems → exp(N) degrees of freedom

Describes cost of simulating dynamics or even describing a state.

This talk: can we do better when a system is only lightly entangled?

success story: quantum circuits

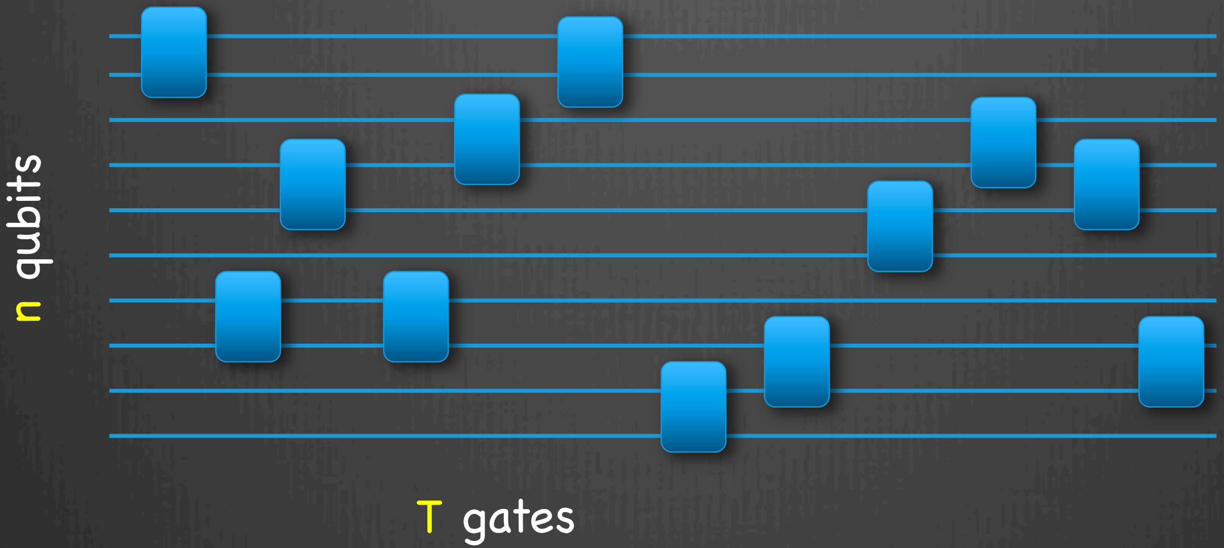

Classical simulation possible in time O(T)·exp(k), where

- k = treewidth [Markov-Shi '05]
- k = max # of gates crossing any single qubit [Yoran-Short '06, Jozsa '06]

+ Complexity interpolates between linear and exponential.

– Treating all gates as "potentially entangling" is too pessimistic.

success story: 1-D systems

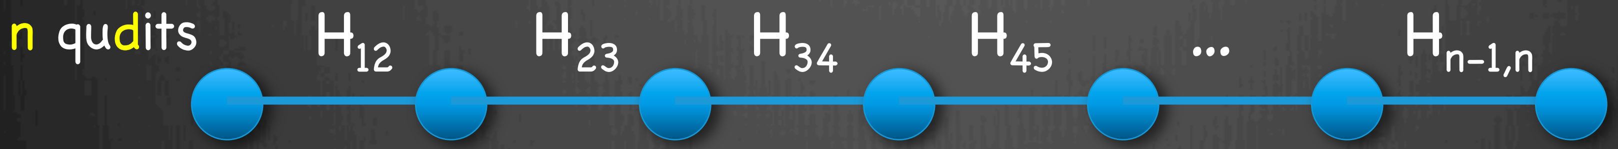

$H = H_{12} + H_{23} + \ldots + H_{n-1,n}$

Classically easy to minimize energy, calculate tr $e^{-H/T}$, etc.

Quantumly QMA-complete to estimate ground-state energy (to precision 1/poly(n) for H with gap 1/poly(n)).

[Landau-Vazirani-Vidick, '13]
n qudits with gap λ and precision ε →
runtime $\exp(\exp(d/\lambda)\log(n))\ \mathrm{poly}(1/\varepsilon)$

Extension to trees:
[Caramanolis, Hayden, Sigler]

intuition:

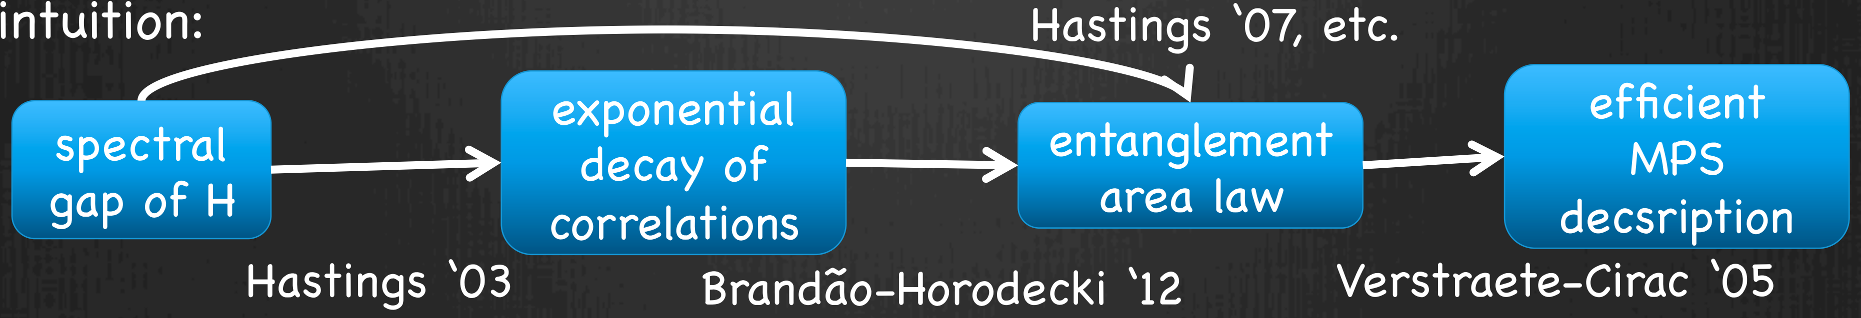
spectral gap of H → exponential decay of correlations → entanglement area law → efficient MPS decsription

Hastings '03

Brandão-Horodecki '12

Verstraete-Cirac '05

Hastings '07, etc.

meta-strategy

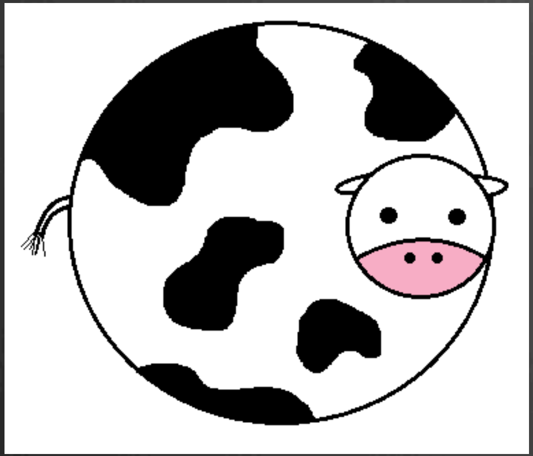

1. solve trivial special case (e.g. non-interacting theory)

2. treat corrections to theory as perturbations

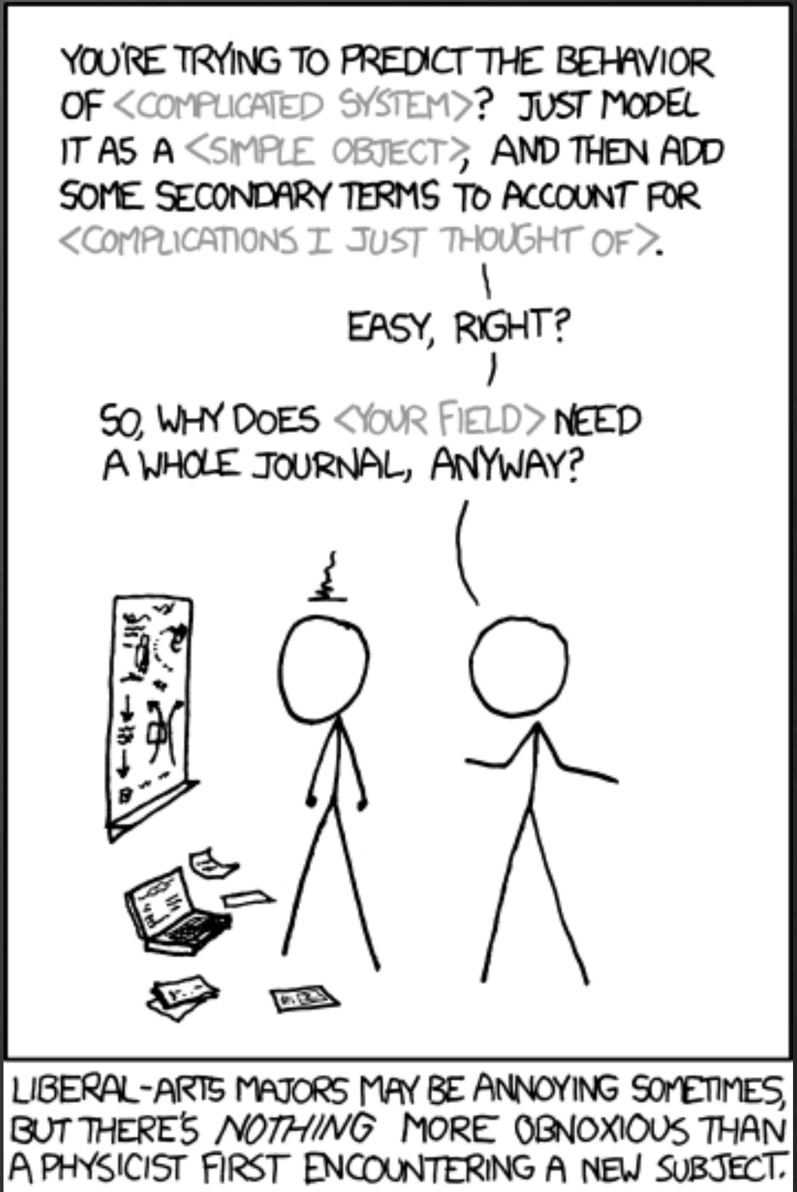

partial success: stabilizer circuits

exact version:
Clifford gates on n qubits = $\{U$ s.t. $UPU^\dagger$ is a Pauli for all Paulis $P\}$
Generated by various single-qubit gates and CNOTs.

[Gottesman-Knill '98] Clifford circuits simulable in time $\tilde{O}(nT)$.
intuition: Paulis $\cong \mathbb{F}_2^{2n}$, Cliffords $\cong Sp_{2n}(\mathbb{F}_2)$

interpolation theorem [Aaronson-Gottesman '04]
Circuits with k non-Clifford gates simulable in time $\tilde{O}(nT \exp(k))$.

+ Can simulate some highly entangled computations including most quantum error-correction schemes.

− Almost all single-qubit gates are non-Clifford gates.

partial success: high-degree graphs

Theorem [Brandão-Harrow, 1310.0017]

If H is a 2-local Hamiltonian on a D-regular graph of n qudits with $H = \mathbb{E}_{i\sim j} H_{i,j}$ and each $\|H_{i,j}\| \leq 1$, then there exists a product state

$|\psi\rangle = |\psi_1\rangle \otimes \ldots \otimes |\psi_n\rangle$ such that

$$\lambda_{\min} \leq \langle\psi|H|\psi\rangle \leq \lambda_{\min} + O(d^{2/3} / D^{1/3})$$

Corollary

The ground-state energy can be approximated to accuracy $O(d^{2/3} / D^{1/3})$ in NP.

interpretation: quantum PCP [tomorrow] impossible unless $D = O(d^2)$

intuition from physics: mean-field approximation

used in limit of high degree, e.g.

1-D

2-D

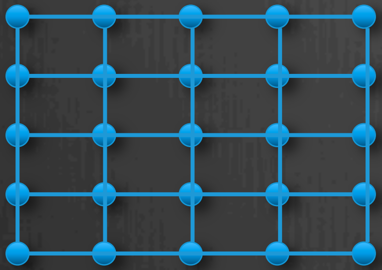

3-D

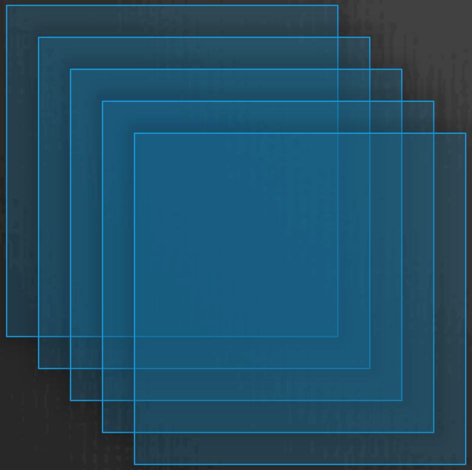

∞-D

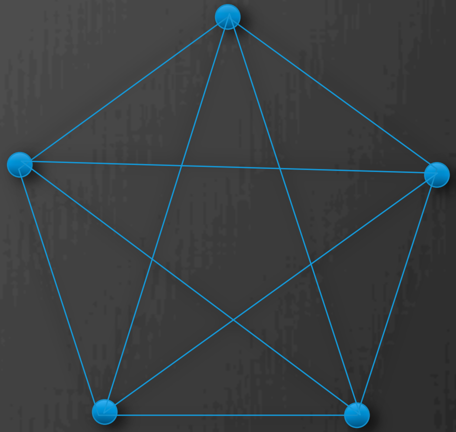

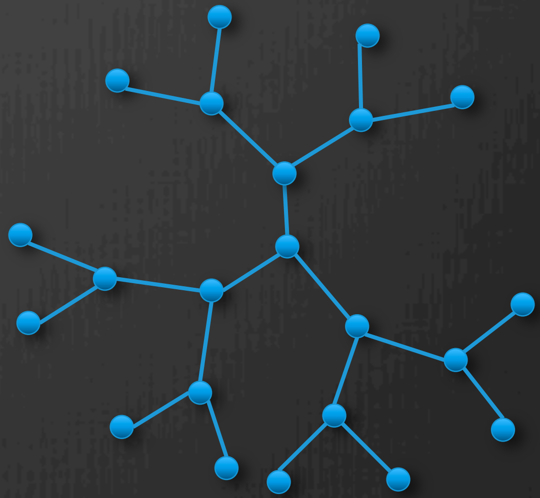

Bethe
lattice
:=
Cayley
graph

clustered approximation

Given a Hamiltonian H on a graph G with vertices partitioned into m-qudit clusters (X_1, …, $X_{n/m}$), can approximate $\lambda_{\min}$ to error

$$9\left(d^2\,\mathbb{E}_i[\Phi(X_i)]\,\frac{1}{D}\,\mathbb{E}_i\,\frac{S(X_i)_{\psi_0}}{m}\right)^{1/3}$$

with a state that has no entanglement between clusters.

$$\Phi(X_i) = \Pr_{(u,v)\in E}(v \notin X_i | u \in X_i)$$

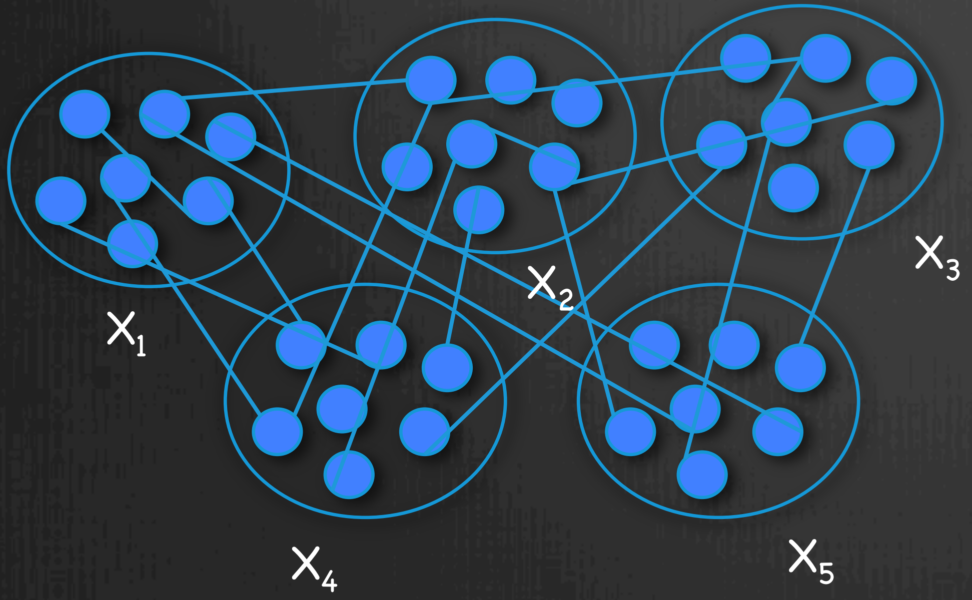

good approximation if

1. expansion is o(1)
2. degree is high
3. entanglement satisfies subvolume law

proof sketch

mostly following [Raghavendra-Tan, SODA '12]

Chain rule Lemma:

$I(X:Y_1\ldots Y_k) = I(X:Y_1) + I(X:Y_2|Y_1) + \ldots + I(X:Y_k|Y_1\ldots Y_{k-1})$

→ $I(X:Y_t|Y_1\ldots Y_{t-1}) \leq \log(d)/k$ for some $t \leq k$.

Decouple most pairs by conditioning:

Choose $i, j_1, \ldots, j_k$ at random from $\{1, \ldots, n\}$

Then there exists $t<k$ such that

$$\mathop{\mathbb{E}}_{i,j,j_1,\ldots,j_t} I(X_i : X_j | X_{j_1} \ldots X_{j_t}) \leq \frac{\log(d)}{k}$$

Discarding systems $j_1,\ldots,j_t$ causes error $\leq k/n$ and leaves a distribution q for which

$$\mathop{\mathbb{E}}_{i,j} I(X_i : X_j)_q \leq \frac{\log(d)}{k} \quad \Longrightarrow \quad \mathop{\mathbb{E}}_{i\sim j} I(X_i : X_j)_q \leq \frac{n}{D}\frac{\log(d)}{k}$$

Does this work quantumly?

What changes?

😊 Chain rule, Pinsker, etc, still work.

😟 Can't condition on quantum information.

😥 $I(A{:}B|C)_\rho \approx 0$ doesn't imply ρ is approximately separable [Ibinson, Linden, Winter '08]

Key technique: informationally complete measurement maps quantum states into probability distributions with poly(d) distortion.

$$\underbrace{d^{-3} \| \rho - \sigma \|_1}_{\text{quantum trace distance}} \le \underbrace{\| M(\rho) - M(\sigma) \|_1}_{\text{classical variational distance}} \le \underbrace{\| \rho - \sigma \|_1}_{\text{quantum trace distance}}$$

Proof of qPCP no-go

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

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

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

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

NP vs QMA

Can you give me some description I can use to get a 0.1% accurate estimate using fewer than 10^{50} steps?

NO.

IGS!
IS T
STA
J, U
KS.

I CAN, HOWEVER, GIVE YOU MANY PROTONS, WHOSE MASS YOU CAN MEASURE.

better approximation?

– There is no guaranteed way to improve the approximation with a larger witness.

Approximation quality depends on:

- degree (fixed)
- average expansion (can change, but might always be high)
- average entropy (can change, but might always be high)

SDP hierarchy:

variables = {density matrices for all sets of ≤k qubits}

constraints = overlap compatibility + global PSD constraint (tomorrow)

Can prove this finds a good product state when k ≫ poly(threshold rank).

Clearly converges to the true ground state energy as k→n.

SDP relaxation ≤ true ground state energy ≤ variational bounds

improves with k

need better ansatz, eg MPS

quantifying entanglement

bipartite pure states – the nice case

$$|\psi\rangle = \sum_{i=1}^{d}\sum_{j=1}^{d} c_{i,j}\,|i\rangle \otimes |j\rangle$$

$$= \sum_{i=1}^{d} \sqrt{\lambda_i}\,|a_i\rangle \otimes |b_i\rangle$$

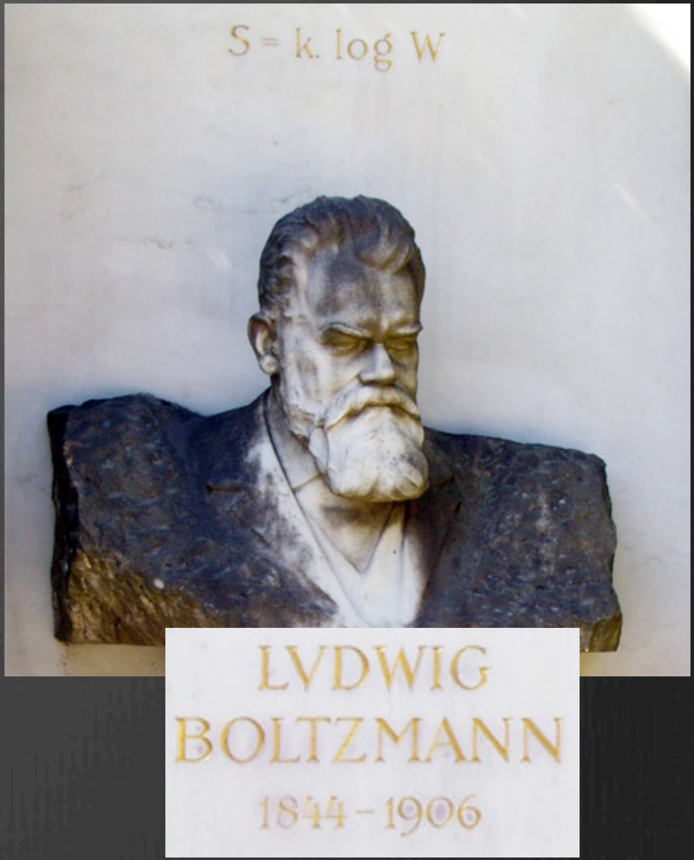

- $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_d \geq 0$ determine equivalence under local unitaries
- LOCC can modify λ according to majorization partial order
- entanglement can be quantified by [Rènyi] entropies of λ

- asymptotic entanglement determined by $H(\lambda) = S(\psi^A) = S(\psi^B)$
  "entropy of entanglement" → entanglement as resource
  [Bennett, Bernstein, Popescu, Schumacher '95]

mixed / multipartite

mixed-state and/or multipartite entanglement measures form a zoo

- relating to pure bipartite entanglement (formation/distillation)
- distance to separable states (relative entropy of entanglement, squashed ent.)
- easy to compute but not operational (log negativity, concurrence)
- operational but hard to compute (distillable key, geometric measure, tensor rank)
- not really measuring entanglement (ent. of purification, ent. of assistance)
- regularized versions of most of the above

Generally "entropic" i.e. match on pure states.
Hopefully convex, continuous, monotonic, etc.

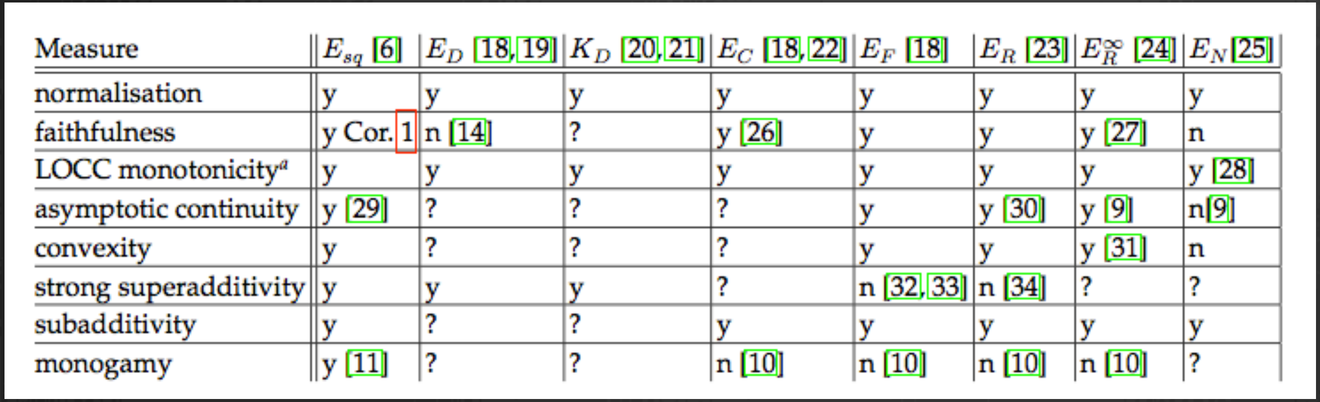

Measure	E_{sq} [6]	E_D [18, 19]	K_D [20, 21]	E_C [18, 22]	E_F [18]	E_R [23]	E_R^∞ [24]	E_N [25]
normalisation	y	y	y	y	y	y	y	y
faithfulness	y Cor. 1	n [14]	?	y [26]	y	y	y [27]	n
LOCC monotonicity[a]	y	y	y	y	y	y	y	y [28]
asymptotic continuity	y [29]	?	?	?	y	y [30]	y [9]	n[9]
convexity	y	?	?	?	y	y	y [31]	n
strong superadditivity	y	y	y	?	n [32, 33]	n [34]	?	?
subadditivity	y	?	?	y	y	y	y	y
monogamy	y [11]	?	?	n [10]	n [10]	n [10]	n [10]	?

Brandão-Christandl-Yard '10

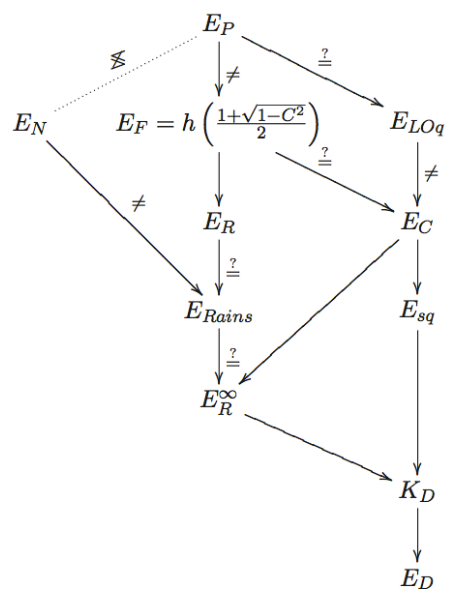

$$\begin{tikzcd}
& E_P \arrow[d, "\neq"] \arrow[rrd, "\overset{?}{=}"] \arrow[ld, dotted, no head, "\not\lesssim"'] & & \\
E_N \arrow[rdd, "\neq"] & E_F = h\left(\frac{1+\sqrt{1-C^2}}{2}\right) \arrow[d] \arrow[rrd, "\overset{?}{=}"] & & E_{LOq} \arrow[d, "\neq"] \\
& E_R \arrow[d, "\overset{?}{=}"] & & E_C \arrow[d] \arrow[lldd] \\
& E_{Rains} \arrow[d, "\overset{?}{=}"] & & E_{sq} \arrow[dd] \\
& E_R^\infty \arrow[rrd] & & \\
& & & K_D \arrow[d] \\
& & & E_D
\end{tikzcd}$$

Christandl '06

conditional mutual information and Markov states

$$
\begin{aligned}
I(A:B|C) &= H(A|C) + H(B|C) - H(AB|C) \\
&= H(AC) + H(BC) - H(ABC) - H(C) \\
&= \Sigma_c\, p(C=c)\, I(A:B)_{p(\cdot,\,\cdot|C=c)} \\
&\geq 0
\end{aligned}
$$

only true classically!

still true quantumly

Classical

TFAE:

- $I(A:B|C)=0$
- $p(a,b,c) = p_1(c)\, p_2(a|c)\, p_3(b|c)$
- $p = \exp(H_{AC} + H_{BC})$ for some H_{AC}, H_{BC} [Hammersley-Clifford]
- A & B can be reconstructed from C

Quantum

$I(A:B|C)=0$

$\Updownarrow$ [Hayden, Jozsa, Petz, Winter '04]

$$C \cong \bigoplus_i C_{A,i} \otimes C_{B,i}$$

$$\rho^{ABC} = \sum_i p_i \alpha^{AC_{A,i}} \otimes \beta^{BC_{B,i}}$$

$\Downarrow$

ρ^{AB} is separable

conditional mutual information

$I(A:B|C)=0 \Leftrightarrow \rho$ is a Markov state
$I(A:B|C)=\varepsilon \Leftrightarrow \rho$ is an approximate Markov state?

Classical

$I(A:B|C)_p = \min_{q \text{ Markov}} D(p \,||\, q)$

I(A:B|C) small → can approximately reconstruct A,B from C.

Quantum

$I(A:B|C)_\rho \leq \min_{\sigma \text{ Markov}} D(\rho \,||\, \sigma)$

I(A:B|C) can be ≪ RHS
[Ibinson, Linden, Winter '06]

ρ^{AB} can be far from separable in trace distance but not 1-LOCC distance. [Brandão,Christandl,Yard '10]

approximate reconstruction? [Winter]

application to Hamiltonians?
[Poulin, Hastings '10] [Brown, Poulin '12]

approximate quantum Markov state

three possible definitions

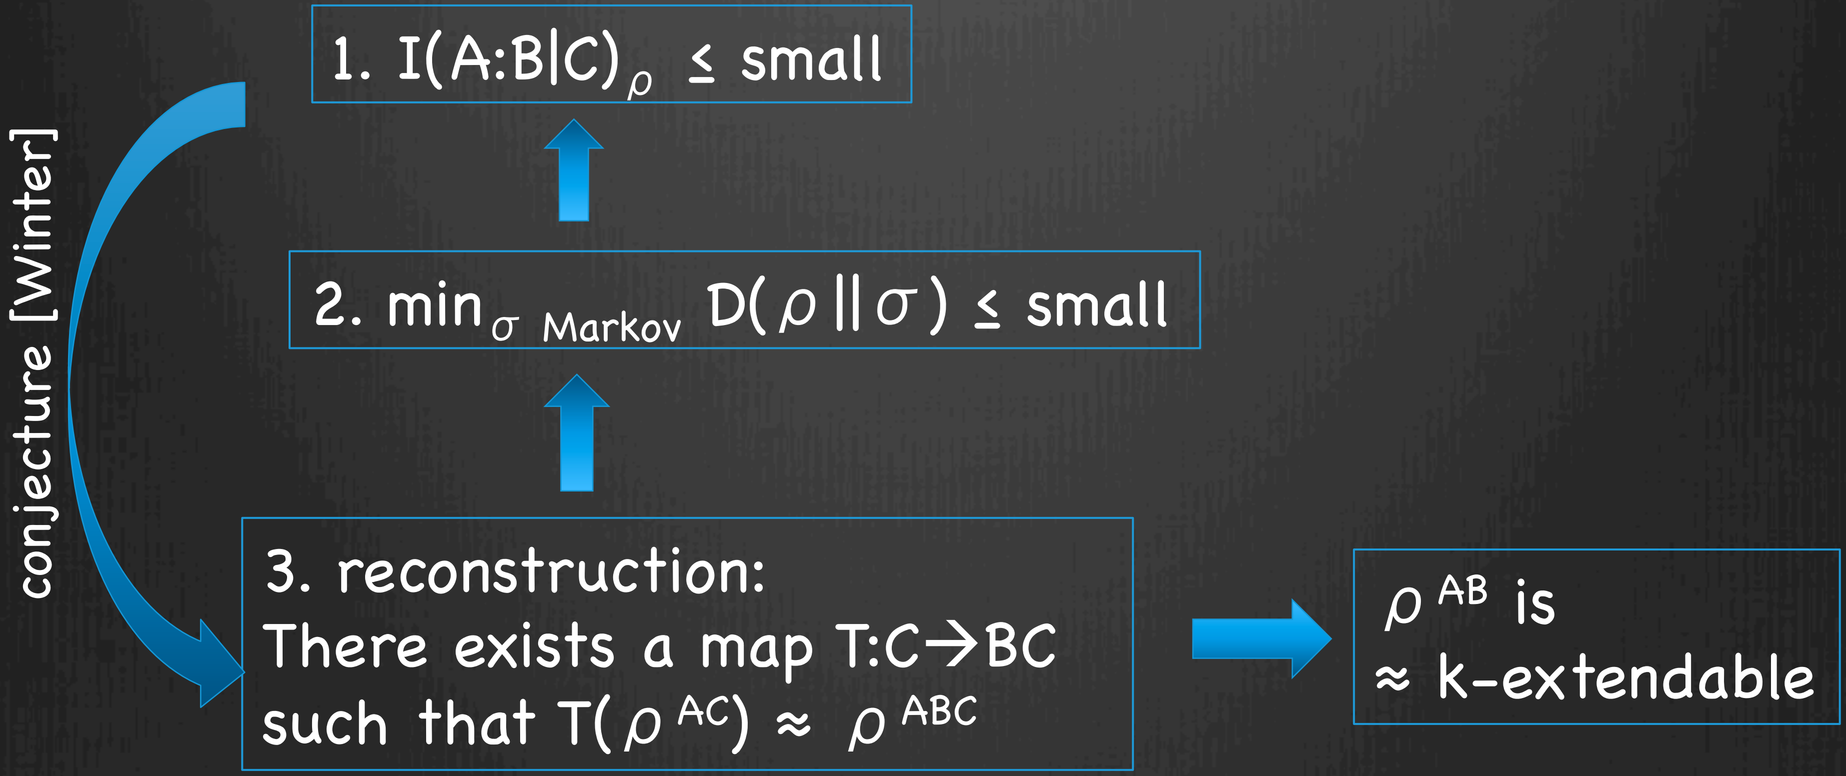

1. $I(A:B|C)_\rho \leq$ small

2. $\min_{\sigma \text{ Markov}} D(\rho || \sigma) \leq$ small

3. reconstruction:
There exists a map $T: C \rightarrow BC$
such that $T(\rho^{AC}) \approx \rho^{ABC}$

ρ^{AB} is
$\approx$ k-extendable

dynamics

Time evolution of quantum systems

$$\frac{d\rho}{dt} = -i(H\rho - \rho H) + \text{noise terms that are linear in } \rho$$

Can we simulate lightly entangled dynamics?
i.e. given the promise that entanglement is always "≤ k" is there a simulation that runs with overhead exp(k)?

noise per gate

0

ideal
QC

10^{-2}-ish

FTQC
possible

?

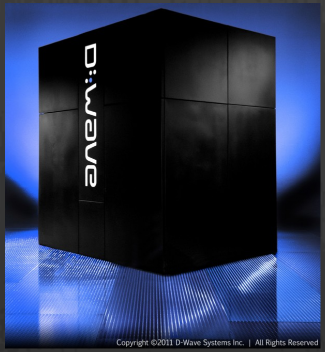

?

≈0.3

classical
simulation
possible

1

open question

If exponential quantum speedup/hardness is due to entanglement, then can we make this quantitative?

Answer may include:

- saving the theory of entanglement measures from itself
- new classical ways to describe quantum states (e.g. MPS)
- conditional mutual information
- the right definition of "approximate quantum Markov states"