How Much Information Is In A Quantum State?

ρ

Scott Aaronson

Andrew Drucker

Computer Scientist / Physicist Nonaggression Pact

You tolerate these complexity classes:

NP coNP BQP QMA BQP/qpoly QMA/poly

And I don't inflict these on you:

#P AM AWPP LWPP MA PostBQP PP CH
PSPACE QCMA QIP SZK NISZK EXP NEXP UP
PPAD PPP PLS TFNP ⊕P Mod_kP

So, how much information **is** in a quantum state?

So, how much information **is** in a quantum state?

An infinite amount, of course, if you want to specify the state exactly…

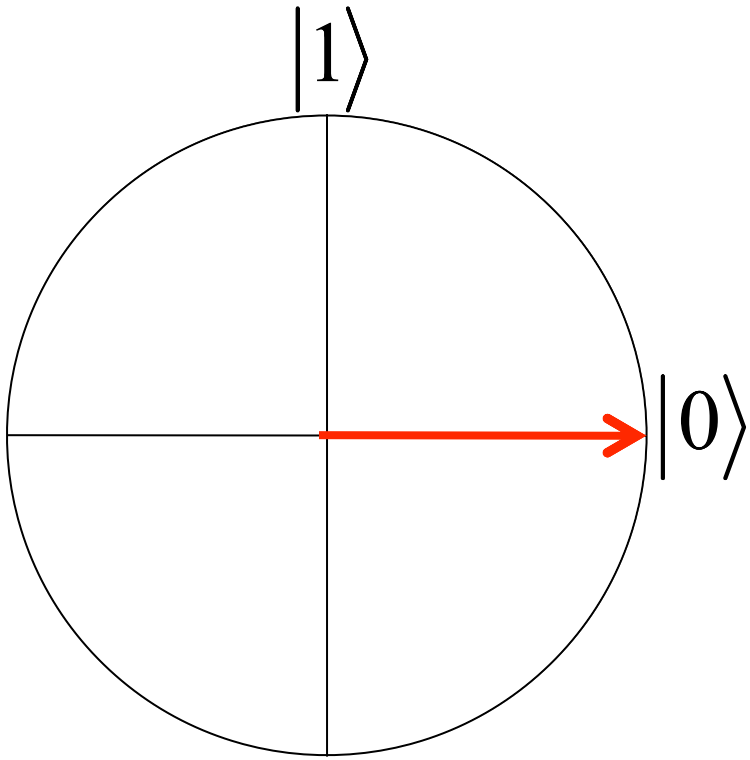

$$|C| = 2^{\aleph_0}$$

So, how much information is in a quantum state?

An infinite amount, of course, if you want to specify the state exactly…

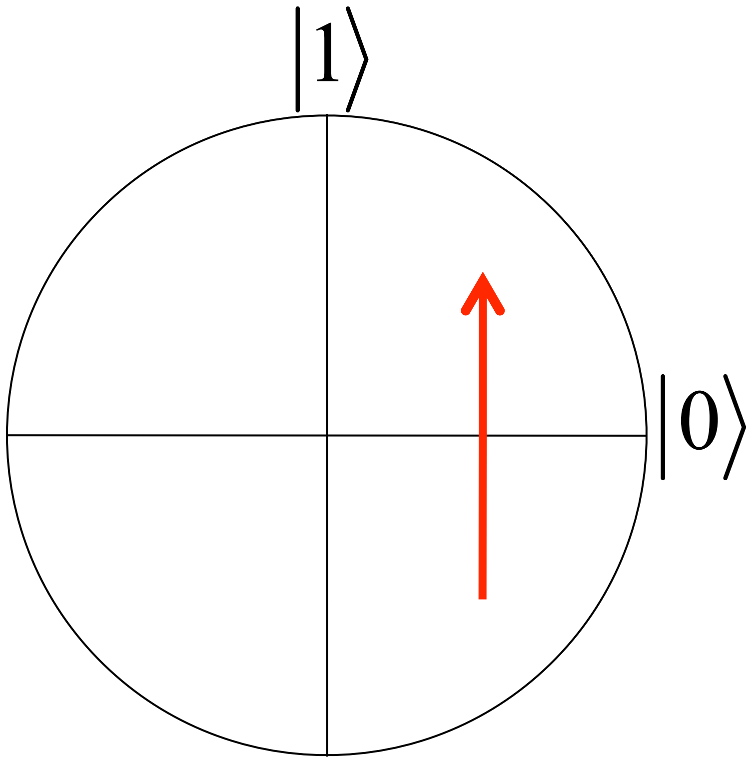

$$|C| = 2^{\aleph_0}$$

So, how much information is in a quantum state?

An infinite amount, of course, if you want to specify the state exactly…

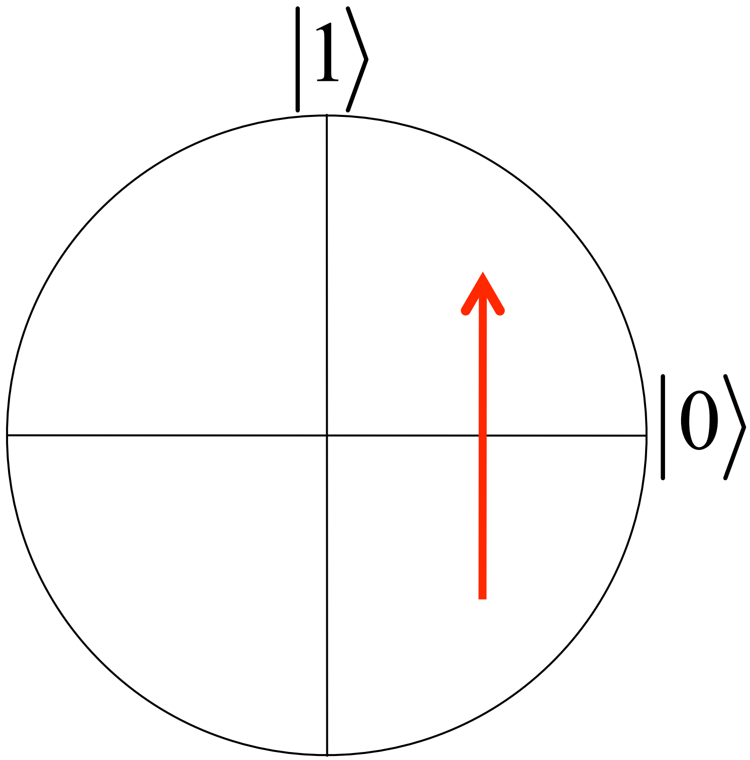

$$|C| = 2^{\aleph_0}$$

Life is too short for infinite precision

A More Serious Point

$$|\psi\rangle = \sum_{x\in\{0,1\}^n} \alpha_x |x\rangle$$

A More Serious Point

In general, a state of n possibly-entangled qubits takes **$\sim 2^n$** bits to specify, even approximately

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$$

A More Serious Point

In general, a state of n possibly-entangled qubits takes

~2^n bits to specify, even approximately

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$$

<u>To a computer scientist, this is arguably **the** central fact about quantum mechanics</u>

A More Serious Point

In general, a state of n possibly-entangled qubits takes **~2^n** bits to specify, even approximately

$$|\psi\rangle = \sum_{x\in\{0,1\}^n} \alpha_x |x\rangle$$

<u>To a computer scientist, this is arguably **the** central fact about quantum mechanics</u>

But why should we worry about it?

Answer 1: Quantum State Tomography

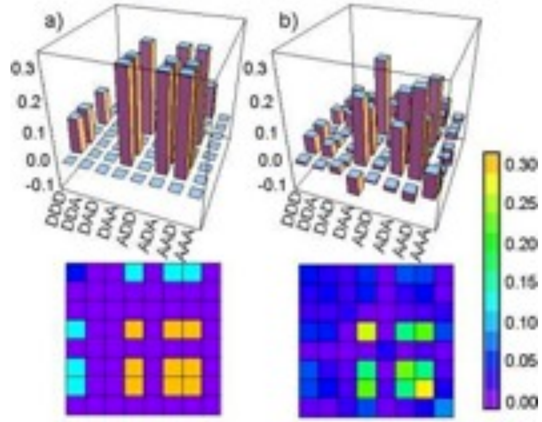

Answer 1: Quantum State Tomography

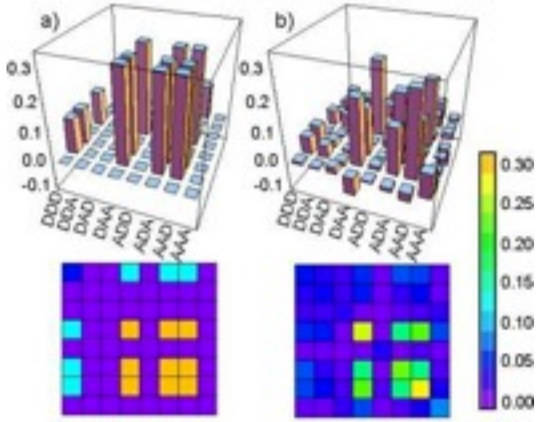

Task: Given lots of copies of an unknown quantum state ρ, produce an approximate classical description of ρ

Answer 1: Quantum State Tomography

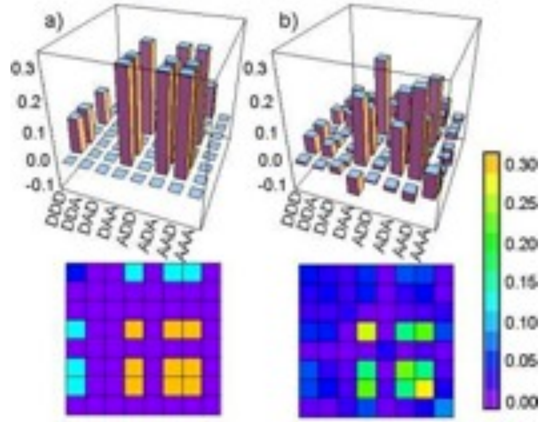

Task: Given lots of copies of an unknown quantum state ρ, produce an approximate classical description of ρ

Not something I just made up!

“As seen in *Science* & *Nature*”

Answer 1: Quantum State Tomography

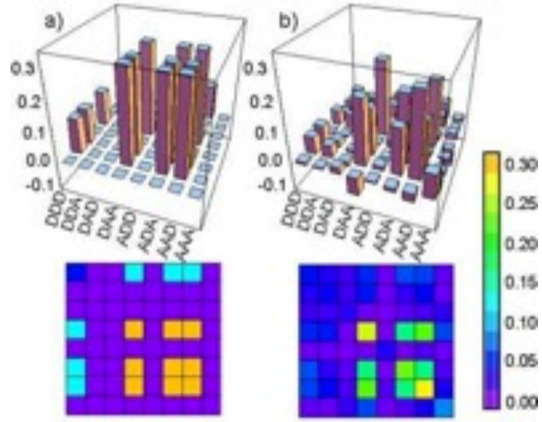

Task: Given lots of copies of an unknown quantum state ρ, produce an approximate classical description of ρ

Not something I just made up!

“As seen in *Science* & *Nature*”

Well-known problem: To do tomography on an entangled state of n spins, you need ~c^n measurements

Current record: 8 spins / ~656,000 experiments (!)

Answer 1: Quantum State Tomography

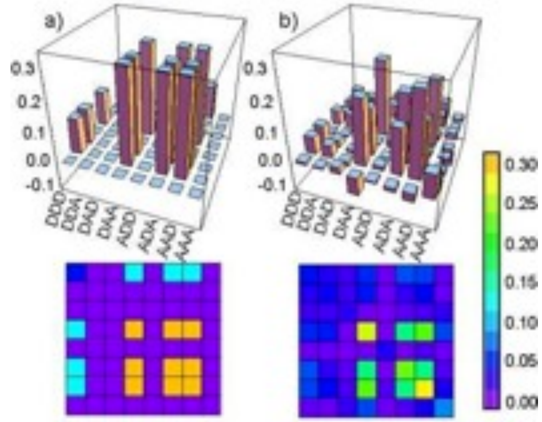

Task: Given lots of copies of an unknown quantum state ρ, produce an approximate classical description of ρ

Not something I just made up!

"As seen in *Science* & *Nature*"

Well-known problem: To do tomography on an entangled state of n spins, you need ~c^n measurements

Current record: 8 spins / ~656,000 experiments (!)

This is a conceptual problem—not just a practical one!

Answer 2: Quantum Computing Skepticism

Levin

Goldreich

‘t Hooft

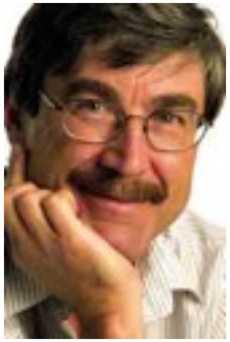

Davies

Wolfram

Some physicists and computer scientists believe quantum computers will be impossible for a fundamental reason

Answer 2: Quantum Computing Skepticism


Levin


Goldreich


‘t Hooft

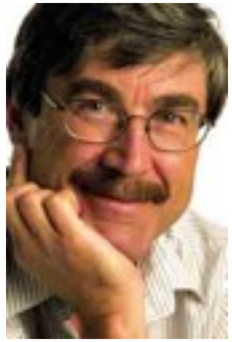
Davies


Wolfram

Some physicists and computer scientists believe quantum computers will be impossible for a fundamental reason

For many of them, the problem is that a quantum computer would “manipulate an exponential amount of information” using only polynomial resources

Answer 2: Quantum Computing Skepticism

Levin

Goldreich

't Hooft

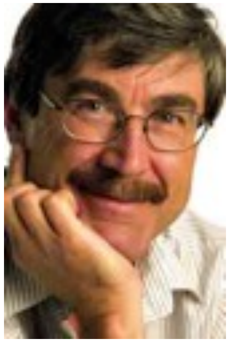

Davies

Wolfram

Some physicists and computer scientists believe quantum computers will be impossible for a fundamental reason

For many of them, the problem is that a quantum computer would "manipulate an exponential amount of information" using only polynomial resources

But is it **really** an exponential amount?

Today we'll tame the exponential beast

Today we’ll tame the exponential beast

Idea: “Shrink quantum states down to reasonable size” by viewing them operationally

Today we'll tame the exponential beast

Idea: "Shrink quantum states down to reasonable size" by viewing them operationally

Analogy: A probability distribution over n-bit strings *also* takes ~2^n bits to specify. But that fact seems to be "more about the map than the territory"

Today we'll tame the exponential beast

Idea: "Shrink quantum states down to reasonable size" by viewing them operationally

Analogy: A probability distribution over n-bit strings *also* takes ~2^n bits to specify. But that fact seems to be "more about the map than the territory"

- Describing a state by postselected measurements [A. 2004]
- "Pretty good tomography" using far fewer measurements [A. 2006]
  - Numerical simulation [A.-Dechter]
- Encoding quantum states as ground states of simple Hamiltonians [A.-Drucker 2009]

The Absent-Minded Advisor Problem

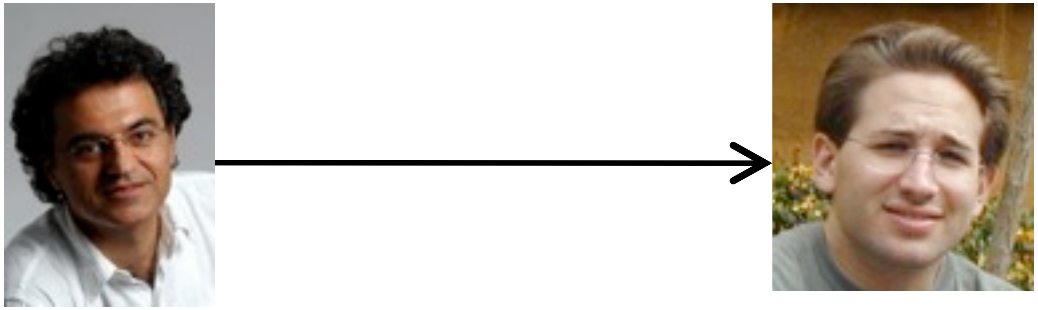

The Absent-Minded Advisor Problem

Can you give your graduate student a quantum state ρ with n qubits (or 10n, or n^3, …)—such that by measuring ρ in a suitable basis, the student can learn your answer to any **one** yes-or-no question of size n?

The Absent-Minded Advisor Problem

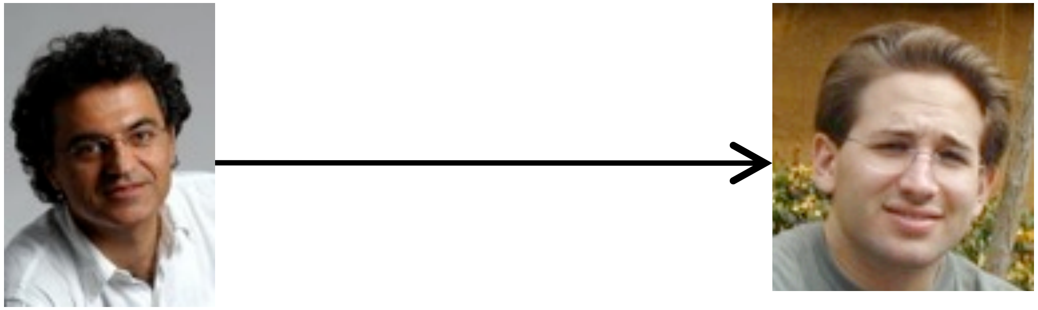

Can you give your graduate student a quantum state ρ with n qubits (or 10n, or n^3, …)—such that by measuring ρ in a suitable basis, the student can learn your answer to any **one** yes-or-no question of size n?

NO [Ambainis, Nayak, Ta-Shma, Vazirani 1999]

Indeed, quantum communication is no better than classical for this problem as $n \rightarrow \infty$.

(Earlier, Holevo showed you need n qubits to send n bits)

On the Bright Side…

On the Bright Side…

Suppose Alice wants to describe an n-qubit state ρ to Bob, well enough that for any 2-outcome measurement E, Bob can estimate Tr(Eρ)

Then she'll need to send **$\sim c^n$** bits, in the worst case.

But… suppose Bob only needs to be able to estimate Tr (Eρ) for every measurement E in a finite set S.

On the Bright Side…

Suppose Alice wants to describe an n-qubit state ρ to Bob, well enough that for any 2-outcome measurement E, Bob can estimate Tr(Eρ)

Then she'll need to send **$\sim c^n$** bits, in the worst case.

But… suppose Bob only needs to be able to estimate Tr(Eρ) for every measurement E in a finite set S.

Theorem (A. 2004): In that case, it suffices for Alice to send **$\sim n \log n \cdot \log|S|$** bits


$|\psi\rangle$
ALL MEASUREMENTS

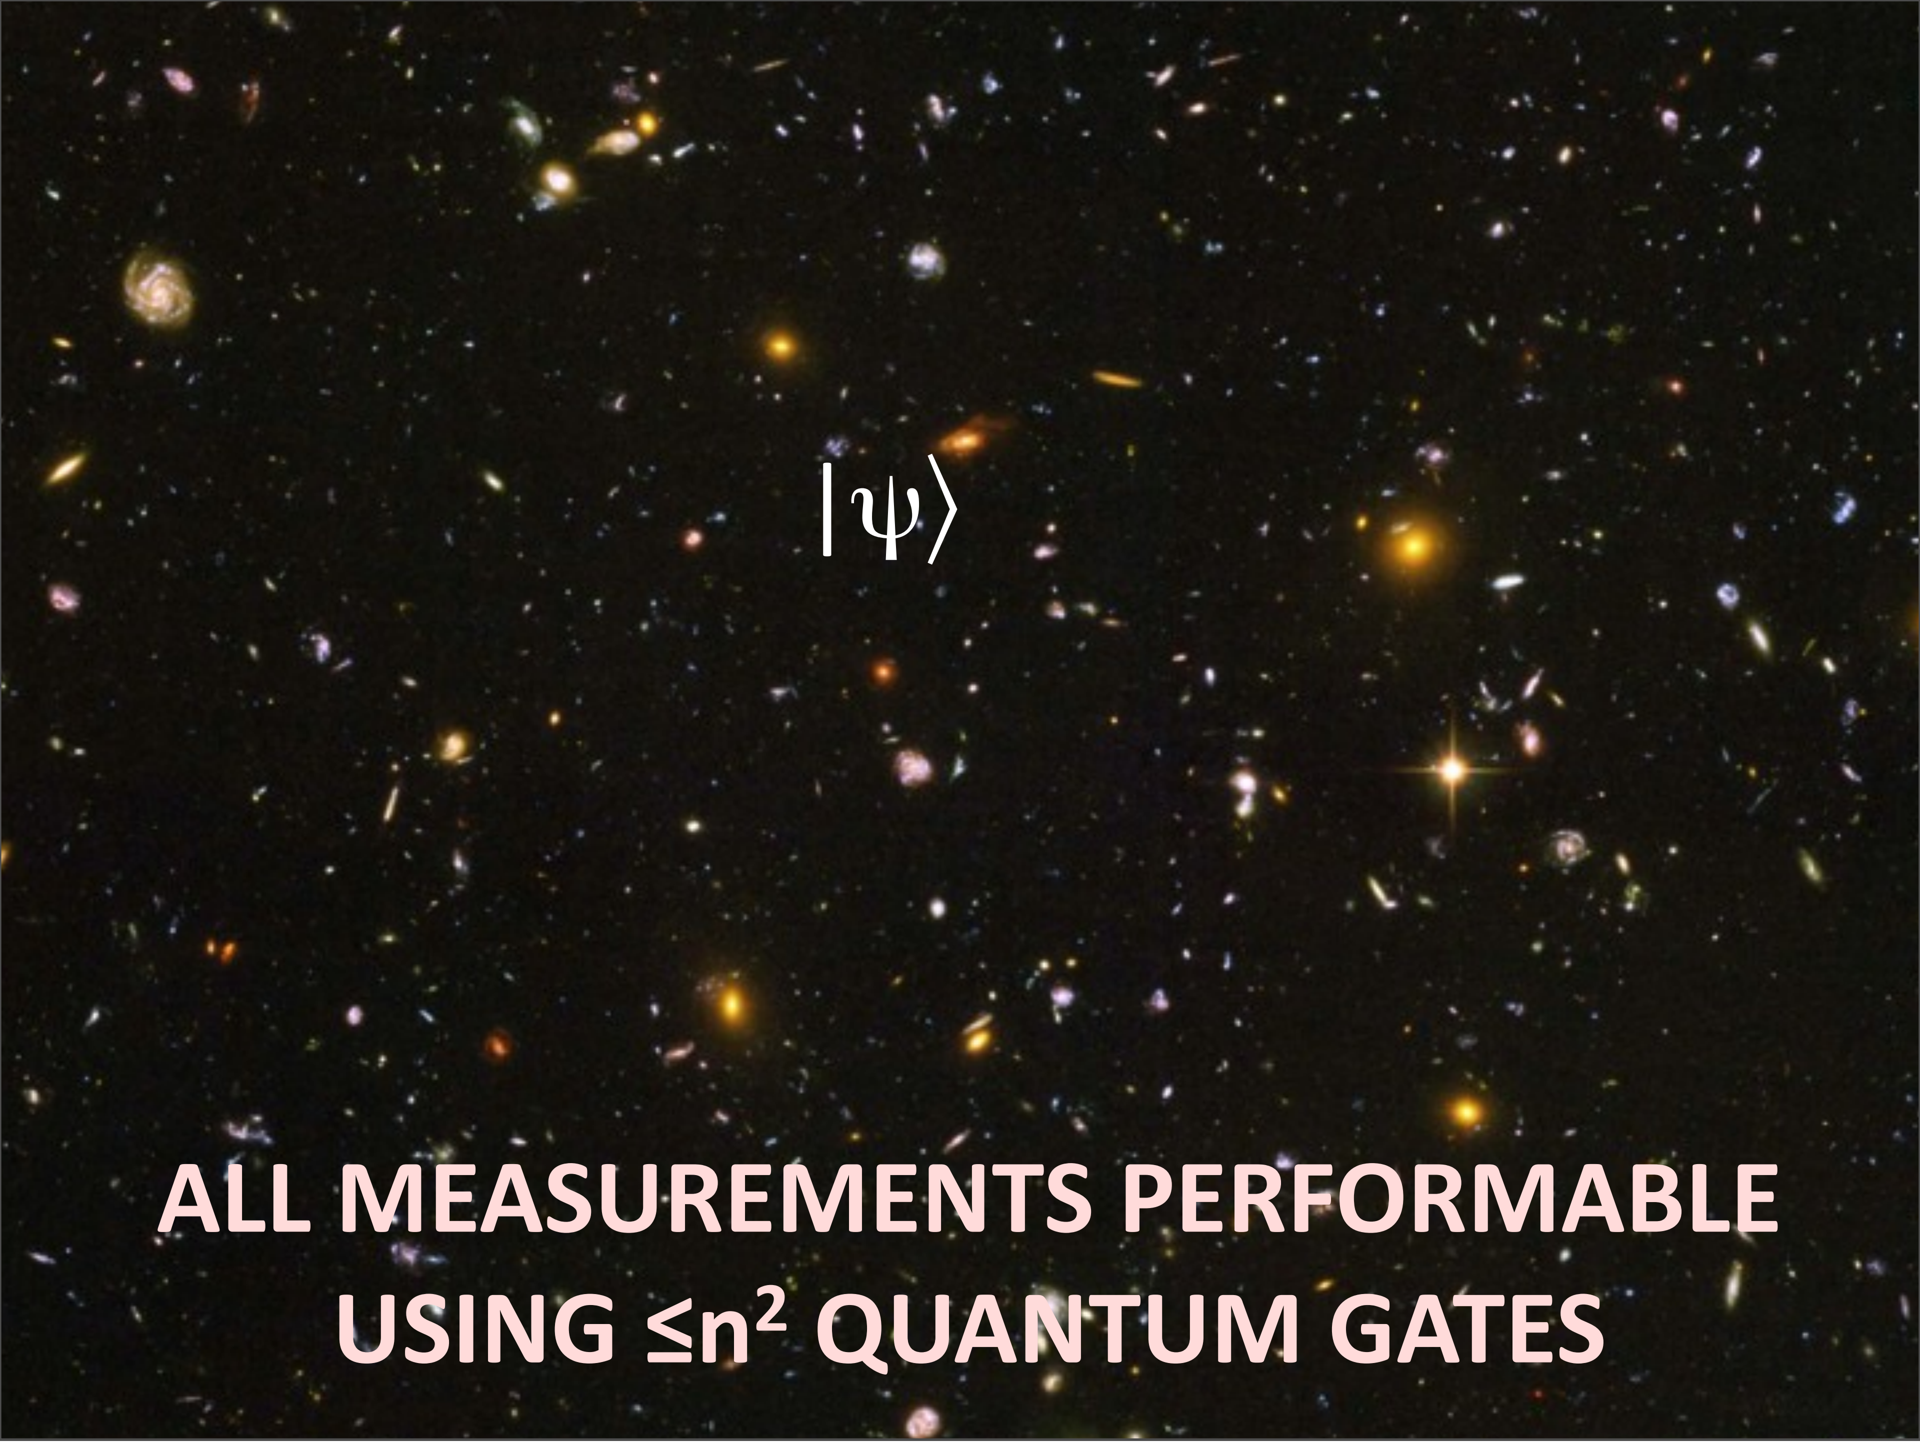
$|\psi\rangle$
ALL MEASUREMENTS PERFORMABLE
USING $\leq n^2$ QUANTUM GATES

How does the theorem work?

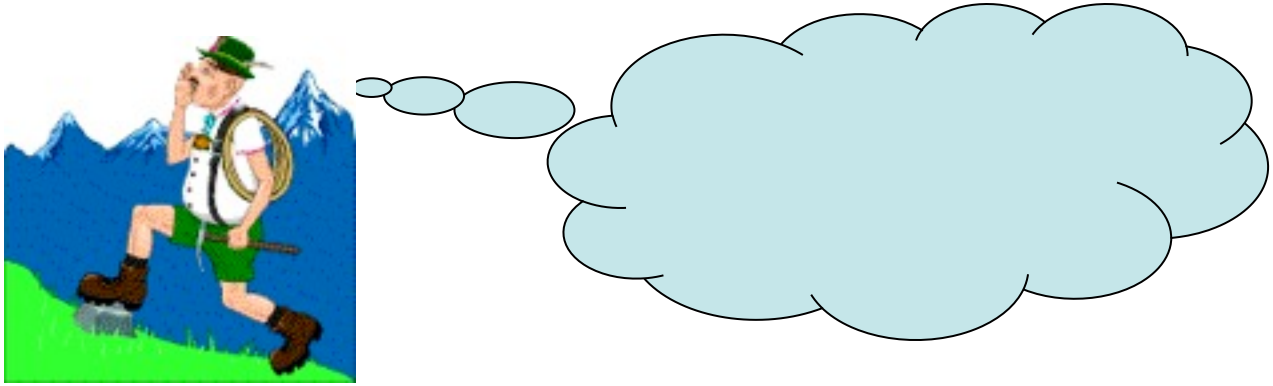

How does the theorem work?

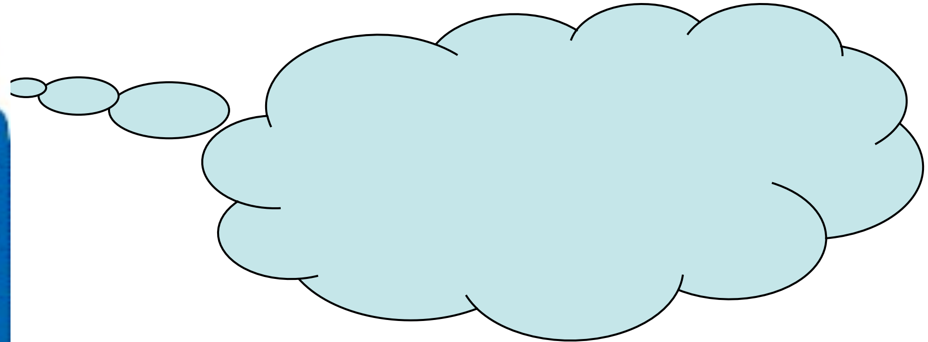

Alice is trying to describe the quantum state ρ to Bob

How does the theorem work?

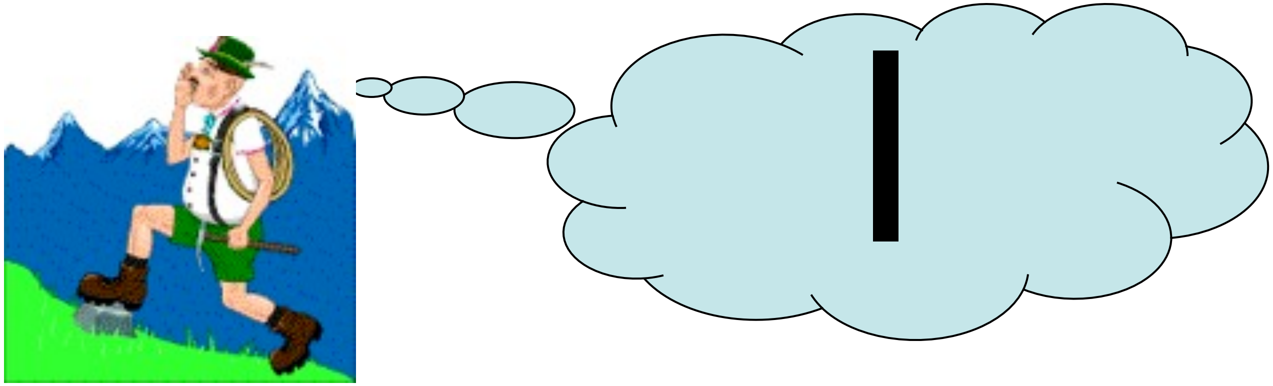

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state ρ_0=I

How does the theorem work?

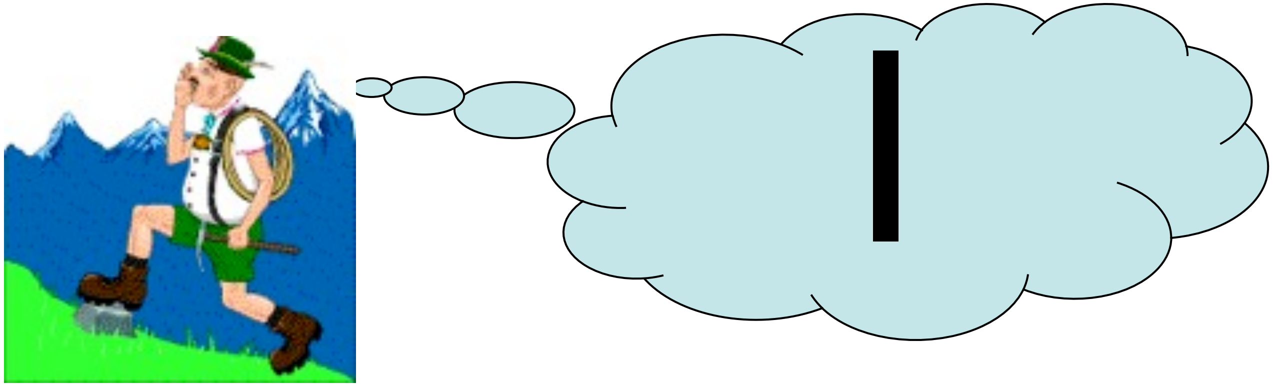

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state $\rho_0=I$

Then Alice helps Bob **improve**, by repeatedly telling him a measurement $E_t \in S$ on which his current guess ρ_{t-1} badly fails

How does the theorem work?

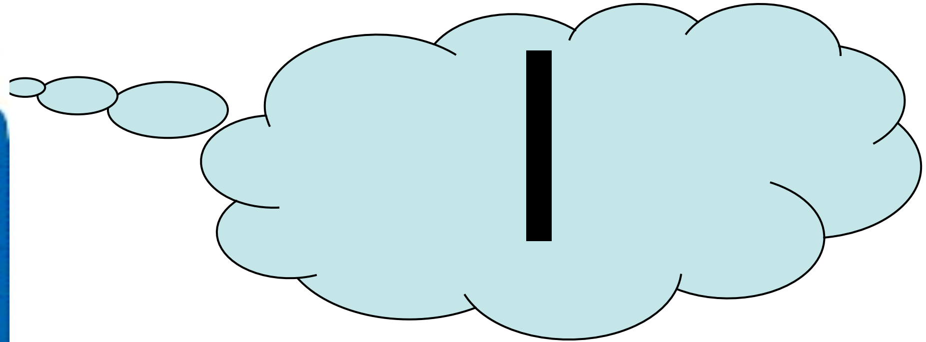

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state ρ_0=I

Then Alice helps Bob **improve**, by repeatedly telling him a measurement $E_t \in S$ on which his current guess ρ_{t-1} badly fails

Bob lets ρ_t be the state obtained by starting from ρ_{t-1}, then performing E_t and **postselecting** on getting the right outcome

How does the theorem work?

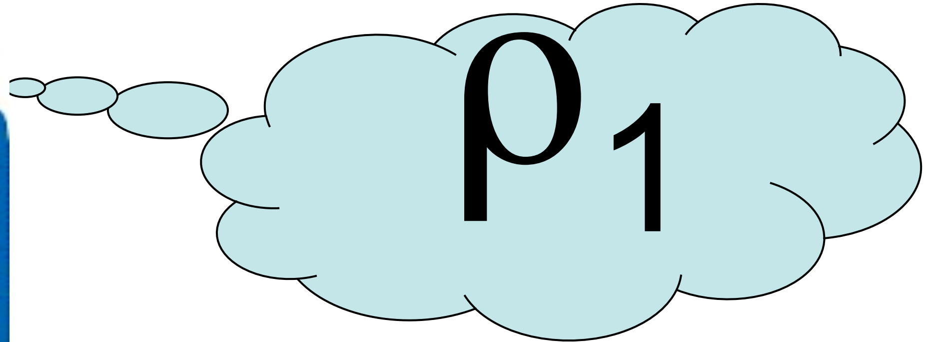

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state $\rho_0=I$

Then Alice helps Bob **improve**, by repeatedly telling him a measurement $E_t \in S$ on which his current guess ρ_{t-1} badly fails

Bob lets ρ_t be the state obtained by starting from ρ_{t-1}, then performing E_t and **postselecting** on getting the right outcome

How does the theorem work?

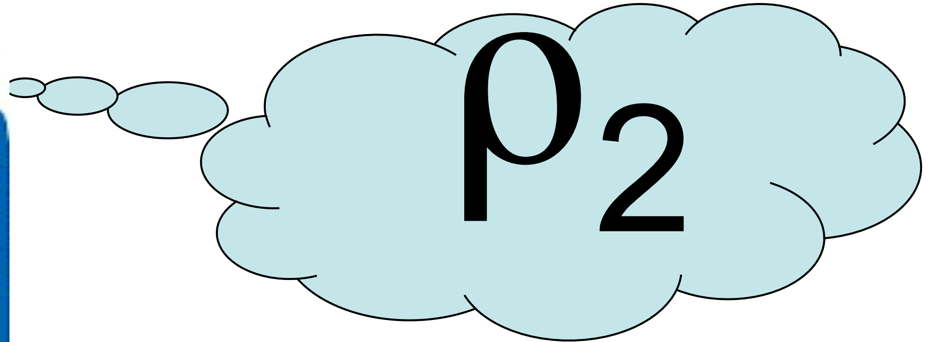

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state ρ_0=I

Then Alice helps Bob **improve**, by repeatedly telling him a measurement $E_t \in S$ on which his current guess ρ_{t-1} badly fails

Bob lets ρ_t be the state obtained by starting from ρ_{t-1}, then performing E_t and **postselecting** on getting the right outcome

How does the theorem work?

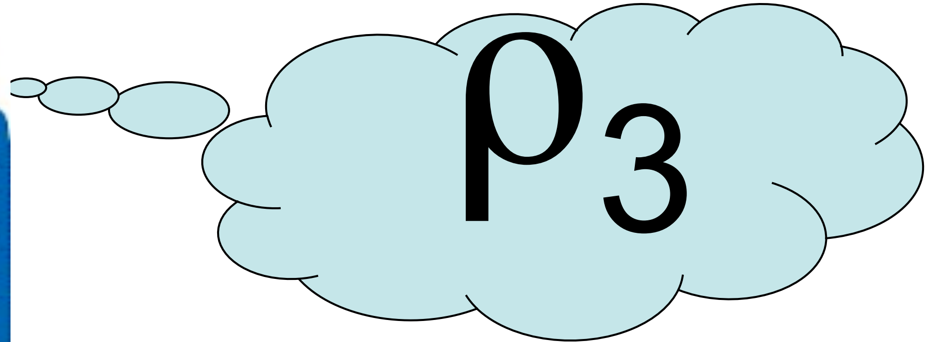

Alice is trying to describe the quantum state ρ to Bob

In the beginning, Bob knows nothing about ρ, so he guesses it's the maximally mixed state ρ_0=I

Then Alice helps Bob **improve**, by repeatedly telling him a measurement $E_t \in S$ on which his current guess ρ_{t-1} badly fails

Bob lets ρ_t be the state obtained by starting from ρ_{t-1}, then performing E_t and **postselecting** on getting the right outcome

Quantum Occam's Razor Theorem [A. 2006]

Quantum Occam's Razor Theorem [A. 2006]

Let ρ be an unknown quantum state of n spins

Suppose you just want to be able to estimate Tr(Eρ) for **most** measurements E drawn from some probability measure D

Quantum Occam's Razor Theorem [A. 2006]

Let ρ be an unknown quantum state of n spins

Suppose you just want to be able to estimate Tr(Eρ) for **most** measurements E drawn from some probability measure D

Then it suffices to do the following, for some m=O(n):

1. Choose $E_1,\ldots,E_m$ independently from D

2. Go into your lab and estimate $Tr(E_i\rho)$ for each $1 \le i \le m$

3. Find **any** "hypothesis state" σ such that $Tr(E_i\sigma) \approx Tr(E_i\rho)$ for all $1 \le i \le m$

Quantum Occam's Razor Theorem [A. 2006]

Let ρ be an unknown quantum state of n spins

Suppose you just want to be able to estimate Tr(Eρ) for **most** measurements E drawn from some probability measure D

Then it suffices to do the following, for some m=O(n):

1. Choose $E_1,\ldots,E_m$ independently from D

2. Go into your lab and estimate $Tr(E_i\rho)$ for each $1 \le i \le m$

3. Find **any** "hypothesis state" σ such that $Tr(E_i\sigma) \approx Tr(E_i\rho)$ for all $1 \le i \le m$

Quantum Occam's Razor Theorem [A. 2006]

Let ρ be an unknown quantum state of n spins

Suppose yo … (Eρ) for **most**
measureme … measure D

> **"Quantum states are PAC-learnable"**

Then it suffices to do the following, for some m=O(n):

1. Choose $E_1,\ldots,E_m$ independently from D

2. Go into your lab and estimate $Tr(E_i\rho)$ for each $1 \le i \le m$

3. Find **any** "hypothesis state" σ such that $Tr(E_i\sigma) \approx Tr(E_i\rho)$ for all $1 \le i \le m$

Numerical Simulation

[A.-Dechter]

Numerical Simulation

[A.-Dechter]

We implemented the "pretty-good tomography" algorithm in MATLAB, using a fast convex programming method developed specifically for this application [Hazan 2008]

We then tested it (on simulated data) using MIT's computing cluster

We studied how the number of sample measurements m needed for accurate predictions scales with the number of qubits n, for $n \leq 10$

Numerical Simulation

[A.-Dechter]

We implemented the "pretty-good tomography" algorithm in MATLAB, using a fast convex programming method developed specifically for this application [Hazan 2008]

We then tested it (on simulated data) using MIT's computing cluster

We studied how the number of sample measurements m needed for accurate predictions scales with the number of qubits n, for $n \leq 10$

Result of experiment: My theorem appears to be true

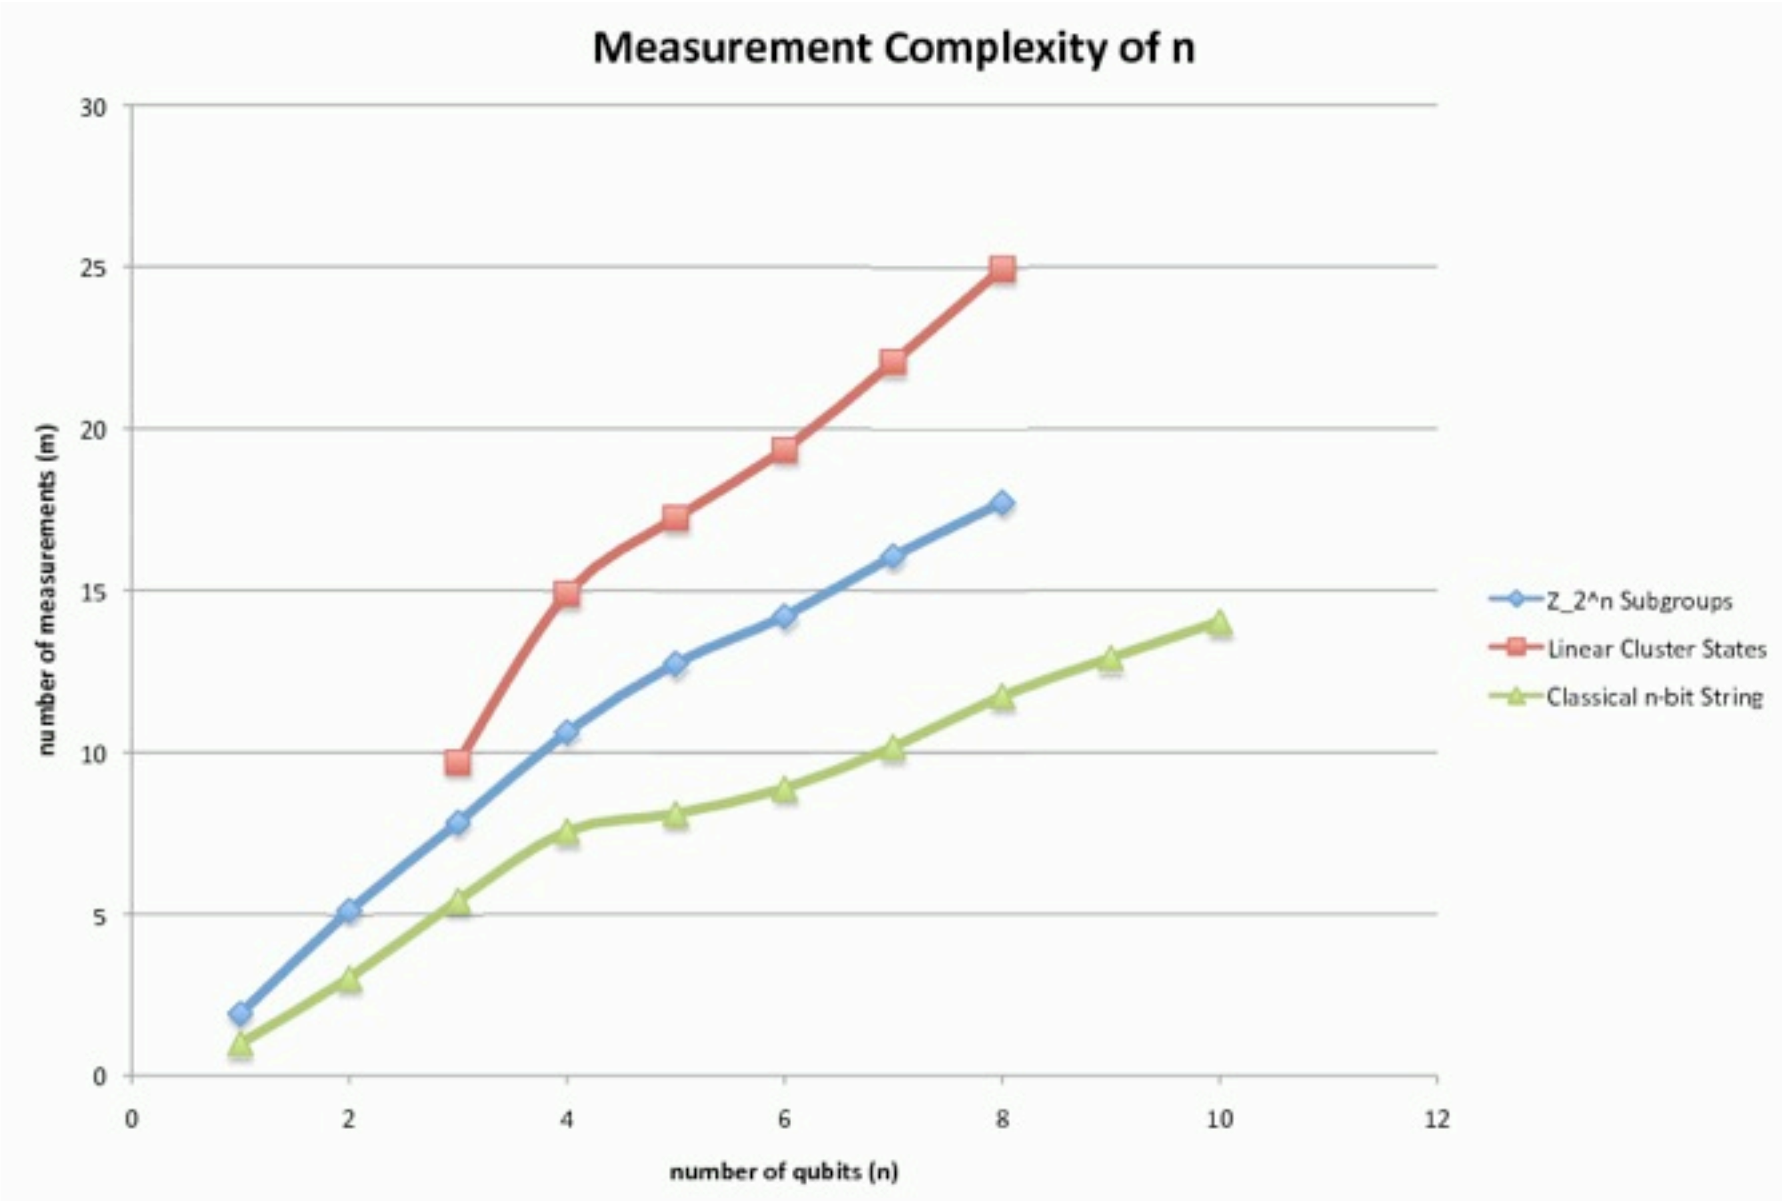
Measurement Complexity of n
30
25
20
15
10
5
0
number of measurements (m)
0
2
4
6
8
10
12
number of qubits (n)
Z_2^n Subgroups
Linear Cluster States
Classical n-bit String

Recap: Given an unknown n-qubit entangled quantum state ρ, and a set S of two-outcome measurements…

Learning theorem: "Any hypothesis state σ consistent with a small number of sample points behaves like ρ on **most** measurements in S"

Postselection theorem: "A particular state ρ_T (produced by postselection) behaves like ρ on **all** measurements in S"

Recap: Given an unknown n-qubit entangled quantum state ρ, and a set S of two-outcome measurements…

Learning theorem: "Any hypothesis state σ consistent with a small number of sample points behaves like ρ on **most** measurements in S"

Postselection theorem: "A particular state ρ_T (produced by postselection) behaves like ρ on **all** measurements in S"

Dream theorem: "Any state σ that passes a small number of tests behaves like ρ on **all** measurements in S"

Recap: Given an unknown n-qubit entangled quantum state ρ, and a set S of two-outcome measurements…

Learning theorem: "Any hypothesis state σ consistent with a small number of sample points behaves like ρ on **most** measurements in S"

Postselection theorem: "A particular state ρ_T (produced by postselection) behaves like ρ on **all** measurements in S"

Dream theorem: "Any state σ that passes a small number of tests behaves like ρ on **all** measurements in S"

[A.-Drucker 2009]: The dream theorem holds

New Result

Any quantum state can be "simulated," on all efficient measurements, by the ground state of a local Hamiltonian

New Result

Any quantum state can be "simulated," on all efficient measurements, by the ground state of a local Hamiltonian

IN OTHER WORDS…

Given any n-qubit state ρ, there exists a local Hamiltonian H (indeed, a sum of 2D nearest-neighbor interactions) such that:

For any ground state $|\psi\rangle$ of H, and measuring circuit E with ≤m gates, there's an efficient measuring circuit E' such that

$$\left|\langle\psi|E'|\psi\rangle - \mathrm{Tr}(E\rho)\right| \leq \varepsilon .$$

Furthermore, H is on poly(n,m,1/ε) qubits.

What Does It Mean?

Without loss of generality, every quantum advice state is the ground state of a local Hamiltonian

BQP/qpoly $\subseteq$ **QMA/poly**. Indeed, trusted quantum advice is *equivalent* in power to trusted classical advice combined with untrusted quantum advice.
("Quantum states never need to be trusted")

"Quantum Karp-Lipton Theorem": NP-complete problems are not efficiently solvable using quantum advice, unless some *uniform* complexity classes collapse

Intuition: We're given a black box (think: quantum state)

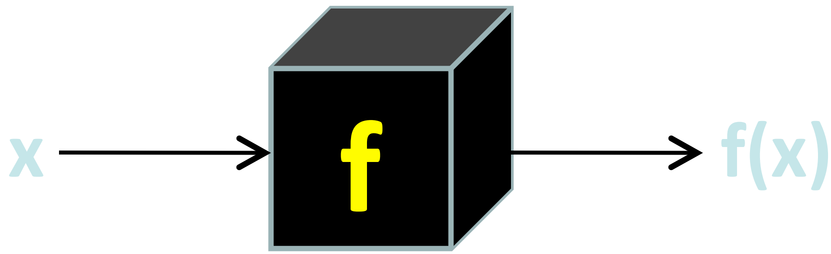

Intuition: We're given a black box (think: quantum state)

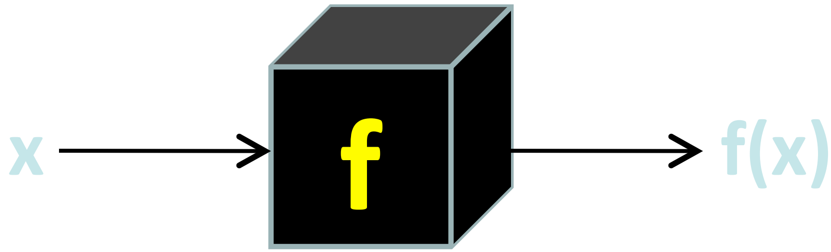

that computes some Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$ belonging to a "small" set S (meaning, of size $2^{poly(n)}$). Someone wants to prove to us that f equals (say) the all-0 function, by having us check a polynomial number of outputs $f(x_1),\ldots,f(x_m)$.

Intuition: We're given a black box (think: quantum state)

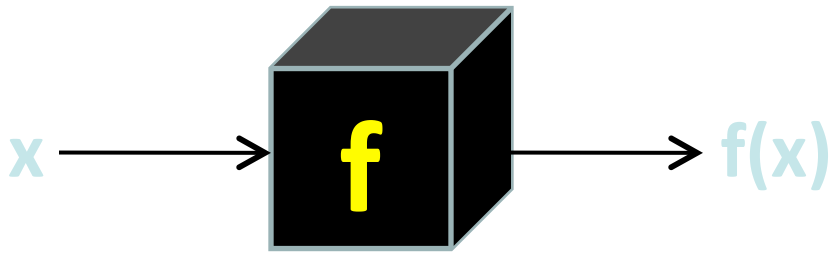

that computes some Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$ belonging to a "small" set S (meaning, of size $2^{poly(n)}$). Someone wants to prove to us that f equals (say) the all-0 function, by having us check a polynomial number of outputs $f(x_1),\ldots,f(x_m)$.

This is trivially impossible!

	f_0	f_1	f_2	f_3	f_4	f_5
x_1	0	1	0	0	0	0
x_2	0	0	1	0	0	0
x_3	0	0	0	1	0	0
x_4	0	0	0	0	1	0
x_5	0	0	0	0	0	1

Intuition: We're given a black box (think: quantum state)

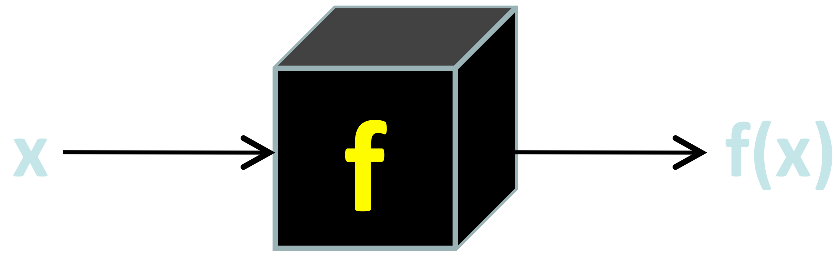

that computes some Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$ belonging to a "small" set S (meaning, of size $2^{poly(n)}$). Someone wants to prove to us that f equals (say) the all-0 function, by having us check a polynomial number of outputs $f(x_1),\ldots,f(x_m)$.

This is trivially impossible!

But … what if we get **3** black boxes, and are allowed to simulate $f=f_0$ by taking the point-wise MAJORITY of their outputs?

	f_0	f_1	f_2	f_3	f_4	f_5
x_1	0	1	0	0	0	0
x_2	0	0	1	0	0	0
x_3	0	0	0	1	0	0
x_4	0	0	0	0	1	0
x_5	0	0	0	0	0	1

Majority-Certificates Lemma

Majority-Certificates Lemma

Definitions: A *certificate* is a partial Boolean function $C:\{0,1\}^n \rightarrow \{0,1,*\}$. A Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$ is *consistent* with C, if $f(x)=C(x)$ whenever $C(x) \in \{0,1\}$. The *size* of C is the number of inputs x such that $C(x) \in \{0,1\}$.

Majority-Certificates Lemma

Definitions: A *certificate* is a partial Boolean function $C:\{0,1\}^n \rightarrow \{0,1,*\}$. A Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$ is *consistent* with C, if $f(x)=C(x)$ whenever $C(x) \in \{0,1\}$. The *size* of C is the number of inputs x such that $C(x) \in \{0,1\}$.

Lemma: Let S be a set of Boolean functions $f:\{0,1\}^n \rightarrow \{0,1\}$, and let $f^* \in S$. Then there exist $m=O(n)$ certificates $C_1,\ldots,C_m$, each of size $k=O(\log|S|)$, such that

(i) Some $f_i \in S$ is consistent with each C_i, and

(ii) If $f_i \in S$ is consistent with C_i for all i, then $MAJ(f_1(x),\ldots,f_m(x))=f^*(x)$ for all $x \in \{0,1\}^n$.

Proof Idea

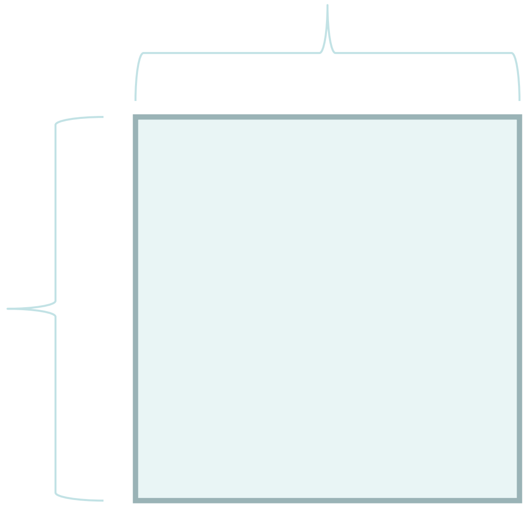

Proof Idea

By symmetry, we can assume f^* is the all-0 function. Consider a two-player, zero-sum matrix game:

Bob picks an input $x \in \{0,1\}^n$

Alice picks a certificate C of size k consistent with some $f \in S$

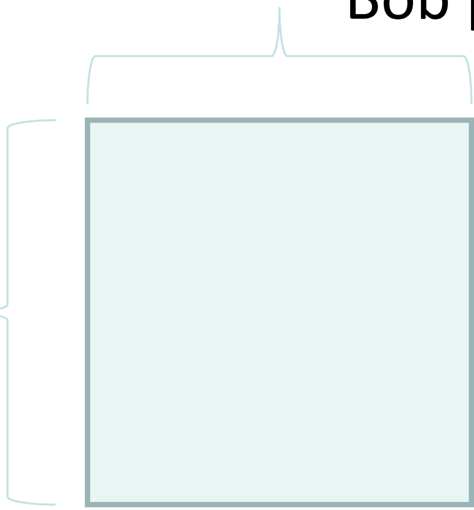

Alice wins this game if $f(x)=0$ for all $f \in S$ consistent with C.

Proof Idea

By symmetry, we can assume f^* is the all-0 function. Consider a two-player, zero-sum matrix game:

Bob picks an input $x \in \{0,1\}^n$

Alice picks a certificate C of size k consistent with some $f \in S$

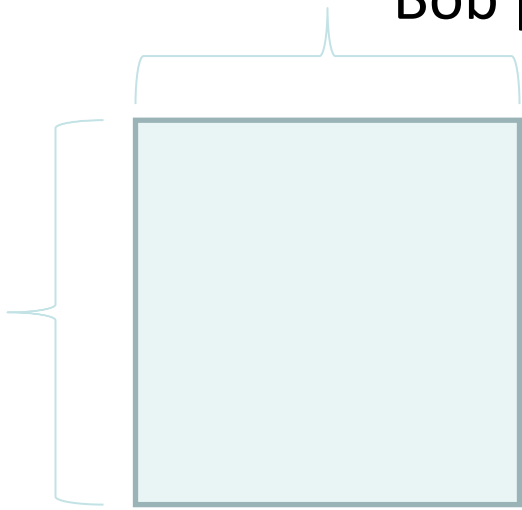

Alice wins this game if f(x)=0 for all $f \in S$ consistent with C.

Crucial Claim: Alice has a mixed strategy that lets her win >90% of the time.

Proof Idea

By symmetry, we can assume f^* is the all-0 function. Consider a two-player, zero-sum matrix game:

Bob picks an input $x \in \{0,1\}^n$

The lemma follows from this claim! Just choose certificates $C_1,\ldots,C_m$ independently from Alice's winning distribution. Then by a Chernoff bound, almost certainly $MAJ(f_1(x),\ldots,f_m(x))=0$ for all $f_1,\ldots,f_m$ consistent with $C_1,\ldots,C_m$ respectively and all inputs $x \in \{0,1\}^n$. So clearly there *exist* $C_1,\ldots,C_m$ with this property.

Alic
C

Alice wins this game if $f(x)=0$ for all $f \in S$ consistent with C.

Crucial Claim: Alice has a mixed strategy that lets her win >90% of the time.

Proof of Claim

Proof of Claim

Use the Minimax Theorem! Given a distribution D over x, it's enough to create a *fixed* certificate C such that

$$\Pr_{x \in D}\left[\exists f \text{ consistent with } C \text{ s.t. } f(x) = 1\right] < \frac{1}{10}.$$

Proof of Claim

Use the Minimax Theorem! Given a distribution D over x, it's enough to create a *fixed* certificate C such that

$$\Pr_{x\in D}\left[\exists f \text{ consistent with } C \text{ s.t. } f(x)=1\right]<\frac{1}{10}.$$

Stage I: Choose $x_1,\ldots,x_t$ independently from D, for some t=O(log|S|). Then with high probability, requiring $f(x_1)=\ldots=f(x_t)=0$ kills off every f∈S such that

$$\Pr_{x\in D}\left[f(x)=1\right]\geq\frac{1}{10}.$$

Proof of Claim

Use the Minimax Theorem! Given a distribution D over x, it's enough to create a *fixed* certificate C such that

$$\Pr_{x \in D}\left[\exists f \text{ consistent with } C \text{ s.t. } f(x)=1\right] < \frac{1}{10}.$$

Stage I: Choose $x_1,\ldots,x_t$ independently from D, for some t=O(log|S|). Then with high probability, requiring $f(x_1)=\ldots=f(x_t)=0$ kills off every f∈S such that

$$\Pr_{x \in D}\left[f(x)=1\right] \geq \frac{1}{10}.$$

Stage II: Repeatedly add a constraint $f(x_i)=b_i$ that kills at least half the remaining functions. After ≤ $\log_2|S|$ iterations, we'll have winnowed S down to just a single function f∈S.

Proof of Claim

Use the Minimax Theorem! Given a distribution D over x, it's enough to create a *fixed* certificate C such that

$$\Pr_{x\in D}\left[\exists f \text{ consistent with } C \text{ s.t. } f(x)=1\right]<\frac{1}{10}.$$

Stage I: Choose $x_1,\ldots,x_t$ independently from D, for some $t=O(\log|S|)$. Then with high probability, requiring $f(x_1)=\ldots=f(x_t)=0$ kills off every $f\in S$ such that

$$\Pr_{x\in D}\left[f(x)=1\right]\geq\frac{1}{10}.$$

Stage II: Repeatedly add a constraint $f(x_i)=b_i$ that kills at least half the remaining functions. After $\leq \log_2|S|$ iterations, we'll have winnowed S down to just a single function $f\in S$. ☐

“Lifting” the Lemma to Quantumland

Boolean Majority-	BQP/qpoly=YQP/
Set S of Boolean	Set S of p(n)-qubit
“True” function	“True” advice state
Other functions f_1..	Other states ρ_1..
Certificate C_i to	Measurement E_i to

"Lifting" the Lemma to Quantumland

Boolean Majority-	BQP/qpoly=YQP/
Set S of Boolean	Set S of p(n)-qubit
"True" function	"True" advice state
Other functions f_1..	Other states ρ_1..
Certificate C_i to	Measurement E_i to

New Difficulty	Solution
The class of p(n)-qubit quantum states is infinitely	Result of A.'06 on learnability of quantum
Instead of Boolean functions f:{0,1}^{n}→{0,1}, now we have	Learning theory has tools to deal with this: fat-
How do we verify a quantum	**QMA**=**QMA+** (Aharonov &
What if a certificate asks us to	"Safe Winnowing Lemma"

Majority-Certificates Lemma, Real Case

Lemma: Let S be a set of functions $f:\{0,1\}^n \rightarrow [0,1]$, let $f^* \in S$, and let $\varepsilon > 0$. Then we can find $m = O(n/\varepsilon^2)$ functions $f_1,\ldots,f_m \in S$, sets $X_1,\ldots,X_m \subseteq \{0,1\}^n$ each of size

$$k = O\left(\frac{n}{\varepsilon^3}\,\mathrm{fat}_{\varepsilon/48}(S)\right),$$

and

$$\alpha = \Omega\left(\frac{\varepsilon^2}{n\,\mathrm{fat}_{\varepsilon/48}(S)}\right)$$

for which the following holds. All functions $g_1,\ldots,g_m \in S$ that satisfy

$$\max_{x \in X_i} \left| g_i(x) - f_i(x) \right| \leq \alpha$$

for all $i \in [m]$ also satisfy

satisfy

$$\max_{x \in \{0,1\}^n} \left| \frac{1}{m}\left[g_1(x) + \cdots + g_m(x) \right] - f^*(x) \right| \leq \varepsilon .$$

Friday, September 25, 2009

Theorem: BQP/qpoly ⊆ QMA/poly.

Proof Sketch: Let L∈BQP/qpoly. Let M be a quantum algorithm that decides L using advice state $|\psi_n\rangle$. Define

$$f_\rho(x) := \Pr\left[M(x,\rho)\text{accepts}\right]$$

Let S = $\{f_\rho : \rho\}$. Then S has fat-shattering dimension at most poly(n), by A.'06. So we can apply the real analogue of the Majority-Certificates Lemma to S. This yields certificates C_1, …,C_m (for some m=poly(n)), such that any states $\rho_1,\ldots,\rho_m$ consistent with $C_1,\ldots,C_m$ respectively satisfy

$$\left|\frac{1}{m}\left(f_{\rho_1}(x)+\cdots+f_{\rho_m}(x)\right)-f_{|\psi_n\rangle\langle\psi_n|}(x)\right|\leq\varepsilon$$

for all $x\in\{0,1\}^n$ (regardless of entanglement). To check the C_i's, we use the "QMA+ super-verifier" of Aharonov & Regev.

Quantum Karp-Lipton Theorem

Quantum Karp-Lipton Theorem

Karp-Lipton 1982: If **NP** ⊂ **P/poly**, then $\mathbf{coNP}^{\mathbf{NP}} = \mathbf{NP}^{\mathbf{NP}}$.

Quantum Karp-Lipton Theorem

Karp-Lipton 1982: If **NP** $\subset$ **P/poly**, then **coNP**$^{\mathbf{NP}}$ = **NP**$^{\mathbf{NP}}$.

Our quantum analogue:

If **NP** $\subset$ **BQP/qpoly**, then **coNP**$^{\mathbf{NP}}$ $\subseteq$ **QMA**$^{\mathbf{PromiseQMA}}$.

Quantum Karp-Lipton Theorem

Karp-Lipton 1982: If **NP** $\subset$ **P/poly**, then **coNP**^{NP} = **NP**^{NP}.

Our quantum analogue:

If **NP** $\subset$ **BQP/qpoly**, then **coNP**^{NP} $\subseteq$ **QMA**^{PromiseQMA}.

Proof Idea: In **QMA**^{PromiseQMA}, first guess a local Hamiltonian H whose ground state $|\psi\rangle$ lets us solve **NP**-complete problems in polynomial time, together with $|\psi\rangle$ itself. Then pass H to the **PromiseQMA** oracle, which reconstructs $|\psi\rangle$, guesses the first quantified string of the **coNP**^{NP} statement, and uses $|\psi\rangle$ to find the second quantified string.

Quantum Karp-Lipton Theorem

Karp-Lipton 1982: If **NP** $\subset$ **P/poly**, then $\mathbf{coNP}^{\mathbf{NP}} = \mathbf{NP}^{\mathbf{NP}}$.

Our quantum analogue:

If **NP** $\subset$ **BQP/qpoly**, then $\mathbf{coNP}^{\mathbf{NP}} \subseteq \mathbf{QMA}^{\mathbf{PromiseQMA}}$.

Proof Idea: In $\mathbf{QMA}^{\mathbf{PromiseQMA}}$, first guess a local Hamiltonian H whose ground state $|\psi\rangle$ lets us solve **NP**-complete problems in polynomial time, together with $|\psi\rangle$ itself. Then pass H to the **PromiseQMA** oracle, which reconstructs $|\psi\rangle$, guesses the first quantified string of the $\mathbf{coNP}^{\mathbf{NP}}$ statement, and uses $|\psi\rangle$ to find the second quantified string.

To check that $|\psi\rangle$ actually works, use the self-reducibility of **NP**-complete problems (like in the original K-L Theorem)

Summary

In many natural scenarios, the "exponentiality" of quantum states is an illusion

That is, there's a short (though possibly cryptic) classical string that specifies how a quantum state ρ behaves, on any measurement you could actually perform

Applications: Pretty-good quantum state tomography, characterization of quantum computers with "magic initial states"…

Open Problems

Find classes of quantum states that can be learned in a **computationally** efficient way

[A.-Gottesman, in preparation]: Stabilizer states

Oracle separation between **BQP/poly** and **BQP/qpoly**

[A.-Kuperberg 2007]: Quantum oracle separation

Other applications of "isolatability" of Boolean functions?

"Experimental demonstration"?