

Lower bounding the description complexity of quantum states

Chinmay Nirkhe (IBM Quantum Cambridge)

based on joint works with

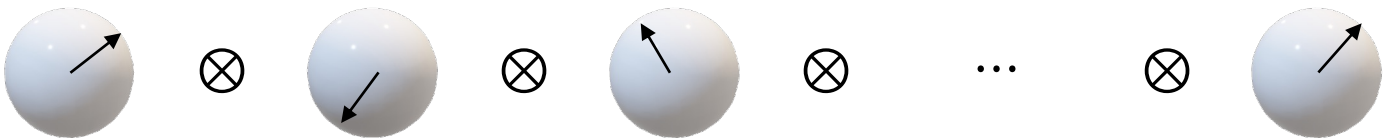
[1] A. Natarajan (MIT),

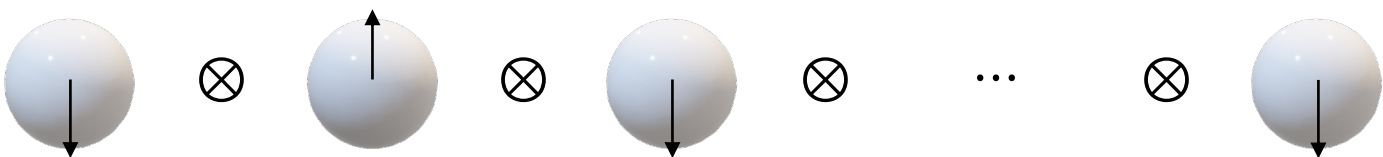
[2] A. Anshu (Harvard), N. Breuckmann (Bristol),

[3] S. Irani (Simons Inst.), A. Natarajan (MIT), S. Rao (MIT), H. Yuen (Columbia)

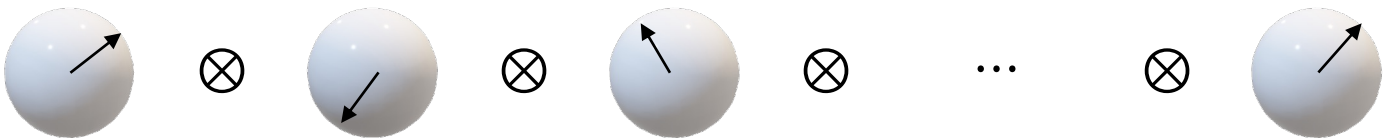


We know that most quantum states are complex...

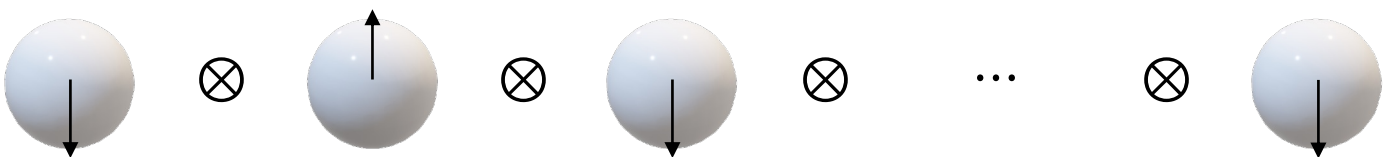
Quantum:  $\approx 2^{2^n}$ states

Classical:  $= 2^n$ states

but how many of
them are interesting
for physics?

Quantum:  $\approx 2^{2^n}$ states

The diagram shows a sequence of particles and tensor products. It starts with a sphere with an arrow pointing up and to the right, followed by a tensor product symbol (⊗), a sphere with an arrow pointing down and to the right, another tensor product symbol, a sphere with an arrow pointing up and to the left, a third tensor product symbol, an ellipsis (...), a fourth tensor product symbol, and finally a sphere with an arrow pointing up and to the right.

Classical:  $= 2^n$ states


The diagram shows a sequence of particles and tensor products. It starts with a sphere with an arrow pointing down, followed by a tensor product symbol (⊗), a sphere with an arrow pointing up, another tensor product symbol, a sphere with an arrow pointing down, a third tensor product symbol, an ellipsis (...), a fourth tensor product symbol, and finally a sphere with an arrow pointing down.

Quantifying the states that are interesting for physics

The energy operator in quantum mechanics is called the **Hamiltonian**.

Interesting physical systems are defined by Hamiltonians of a special form called **local Hamiltonians**.

$$\mathbf{H} = \sum h_i$$



$$h_i = |000\rangle\langle 000| - |111\rangle\langle 111|$$

Due to the importance of ground states in condensed matter physics,

we would really like to know if ground states of local Hamiltonians have efficient (i.e. short) and verifiable classical descriptions.

This talk:

We prove lower bounds on the classical description complexity of ground (and low-energy) states of local Hamiltonians.

Are there local Hamiltonian ground states that cannot be described by polynomial depth quantum circuits?

The converse is false [Kitaev⁰³]. For every circuit \mathcal{C} , there exists a LH with ground state $\approx \mathcal{C}|0^n\rangle \otimes |\text{junk}\rangle$.


A counting exercise ...

How many local Hamiltonians are there?

Each h_i can be described by $O(2^{2\ell} \log n)$ bits.

There are n^ℓ terms, so \mathbf{H} can be described by $O(2^{2\ell} n^\ell \log n)$ bits.

So, there are at most $2^{\text{poly}(n)}$ local Hamiltonians for $\ell = O(1)$.



$$h_i = |000\rangle\langle 000| - |111\rangle\langle 111|$$

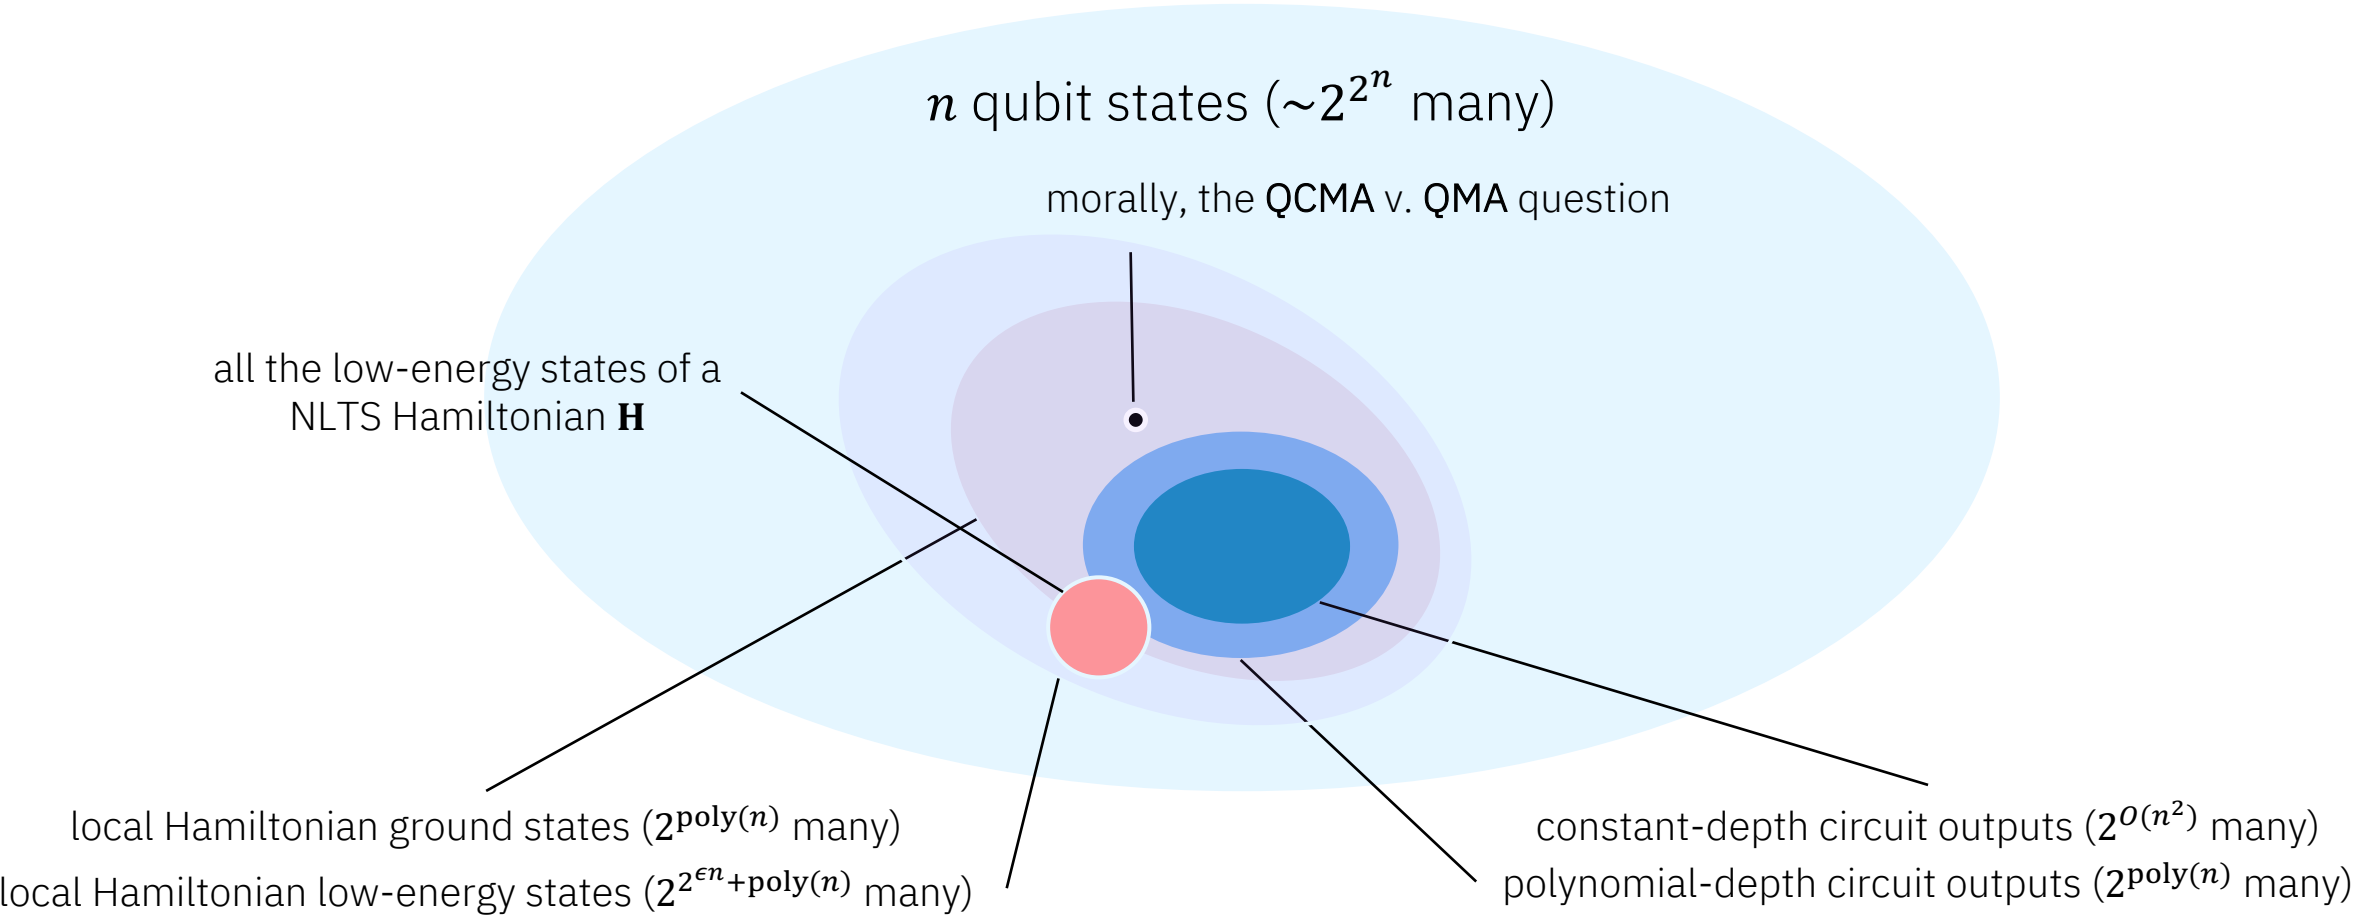
How many depth t quantum circuits are there?

Each row of the quantum circuit can be described by a matching and $n/2$ 4×4 matrices.

So, there are $2^{O(t n^2)}$ quantum circuits of depth t .

Therefore, counting arguments alone don't show separation and, secondly, both the number of local Hamiltonians and "accessible" quantum states are \ll number of quantum states (2^{2^n}).

The space of quantum states



Outline of today's talk

Oracle separations between **QMA** and **QCMA** [1]

*i.e., proving polynomial lower bounds for ground state descriptions
(for a generalization of local Hamiltonians)*

The **NLTS** problem [2]

*i.e., proving logarithmic-depth circuit lower bounds for
all low-energy states of some local Hamiltonians*

Quantum search-to-decision reductions [3]

*i.e., why we cannot study the complexity of quantum states via
the study of quantum decision problems*

[1] A classical oracle separation between QMA and QCMA (2022).

[2] NLTS Hamiltonians from quantum codes (2022).

[3] Quantum search-to-decision and the state synthesis problem (2022).


A classical oracle
separations for
QCMA and QMA

The connection to quantum complexity theory

The energy operator in quantum mechanics is called the **Hamiltonian**.

Interesting physical systems are defined by Hamiltonians of a special form called **local Hamiltonians**.

$$\mathbf{H} = \sum h_i$$



$$h_i = |000\rangle\langle 000| - |111\rangle\langle 111|$$

The problem of calculating the ground energy

$$E = \min_{|\psi\rangle} \langle \psi | \mathbf{H} | \psi \rangle$$

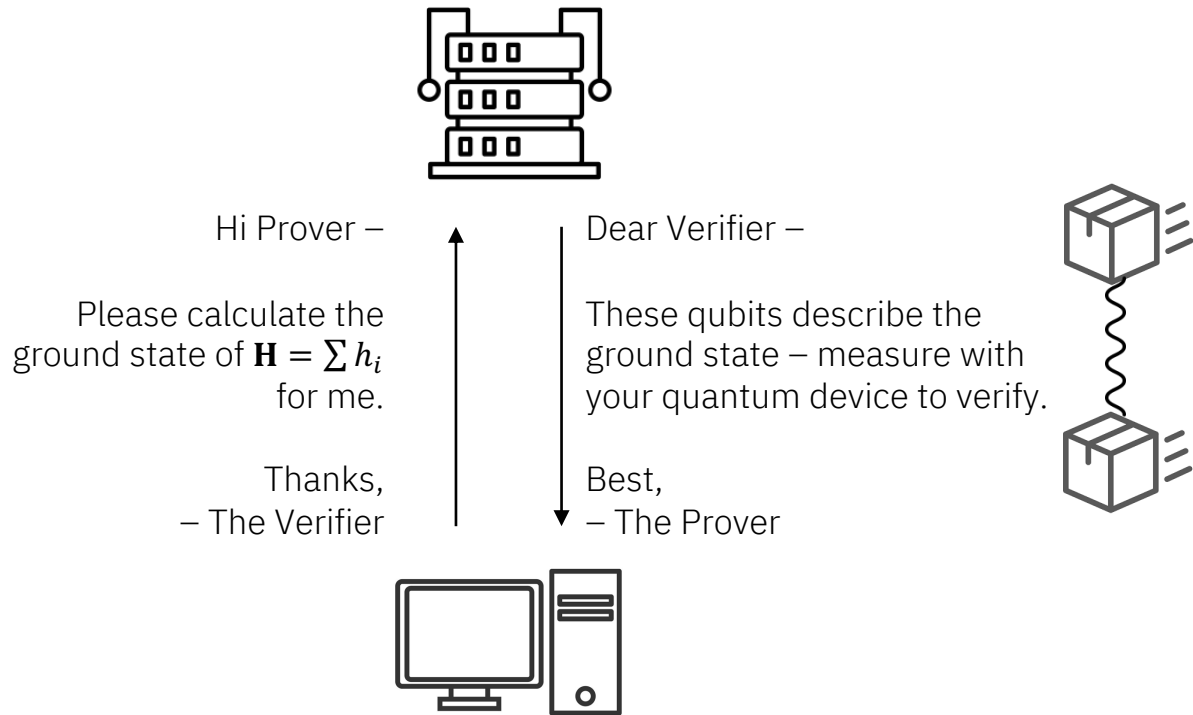
famously connects physics to computer science as the problem is *complete* for the class **QMA**.

QMA = Quantum Merlin-Arthur

If the ground state can always be *verifiably* classically described, then the problem is complete for the class **QCMA**.

QCMA = Quantum-Classical Merlin-Arthur

QMA



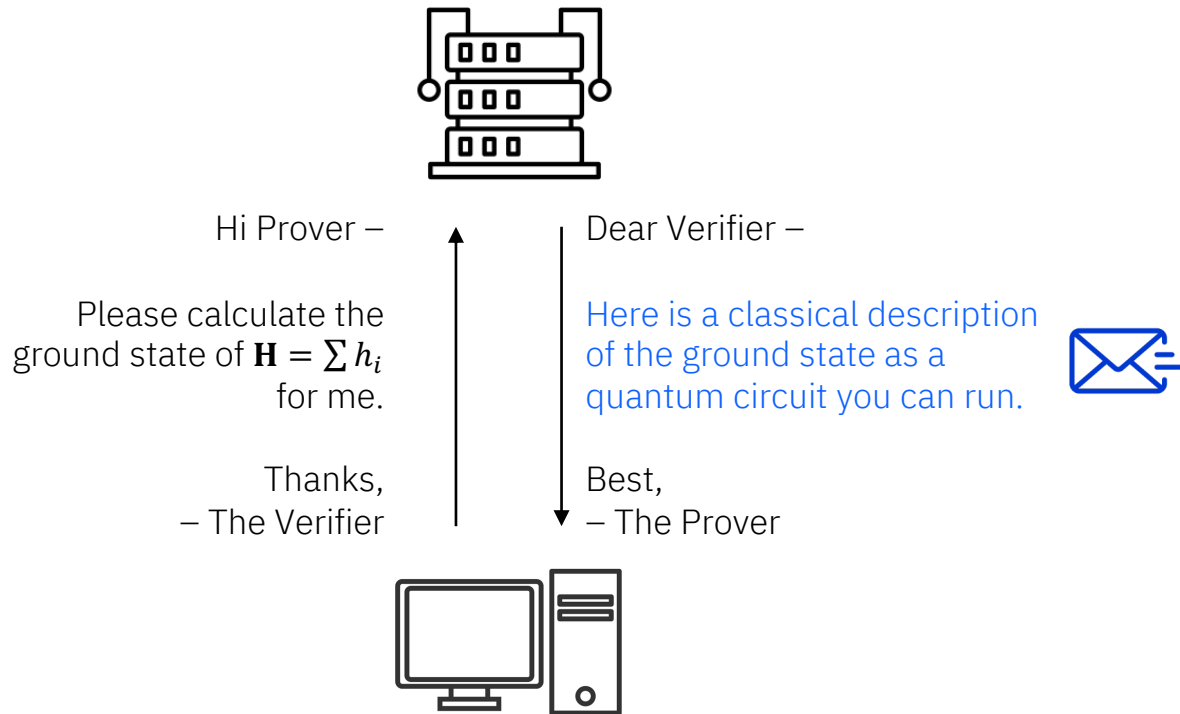
QMA: The set of problems that can be verified by a prover sending a quantum state to the verifier which the verifier uses in a quantum circuit.

Complete problem: Deciding if there exists a quantum state $|\psi\rangle$ such that

YES: $\langle\psi|\mathbf{H}|\psi\rangle \leq a$

NO: or else, $\langle\psi|\mathbf{H}|\psi\rangle \geq a + 1/n^3$ for all states

QCMA



QCMA: The set of problems that can be verified by a prover sending a **classical string** to the verifier which the verifier uses in a quantum circuit.

Complete problem: Deciding if there exists a quantum state $|\psi\rangle$ such that

YES: $\langle\psi|\mathbf{H}|\psi\rangle \leq a$

NO: or else, $\langle\psi|\mathbf{H}|\psi\rangle \geq a + 1/n^3$ for all states

where in both cases, $|\psi\rangle$ is the output of a quantum circuit of polynomial depth.


The question, of whether $QCMA$ and QMA are equal, is a major open question in quantum complexity theory.

QCMA \neq QMA

implies

that not all local Hamiltonians
ground states have short verifiable
classical descriptions.

Contrapositive: If short descriptions *always* exist, then prover can *always* send the short description instead of the quantum state.



vs.



DALL-E 2 renderings.

Theorem [Natarajan-Nirkhe²²]: There is a black-box distribution D for which we can prove that


$$\text{QCMA}^D \neq \text{QMA}^D.$$

This is the strongest evidence yet that ground states cannot be classically described.

Proving $\text{QCMA} \neq \text{QMA}$ outright would have incredible implications for complexity theory. Consequently, it requires truly novel techniques.




Instead, the best we can do is slightly generalize the notion of local Hamiltonians until we can prove $\text{QCMA} \neq \text{QMA}$.




Generalizing sparse Hamiltonians to graphs

$$\mathbf{H} = \begin{pmatrix} 0 & \alpha & 0 & 0 & 0 & \theta & 0 & 0 \\ \alpha & 0 & 0 & 0 & 0 & \epsilon & 0 & 0 \\ 0 & 0 & 0 & \beta & 0 & 0 & \delta & 0 \\ 0 & 0 & \beta & \lambda & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \gamma & 0 & 0 & \iota \\ \theta & \epsilon & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta & 0 & 0 & 0 & 0 & \zeta \\ 0 & 0 & 0 & 0 & \iota & 0 & \zeta & 0 \end{pmatrix}$$

n -qubit sparse Hamiltonian



2^n vertex weighted graph



2^n vertex unweighted graph



$$\mathbf{H} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$

Adjacency matrix of graph

Low energies of sparse graph Hamiltonians

If \mathbf{H} is the Hamiltonian corresponding to a d -regular sparse graph,

then $-\mathbf{H}$ is a Hamiltonian with ground energy equal to $-d$ with a ground state of $\sum_{x \in \{0,1\}^n} |x\rangle$, the uniform super position.

What is the second smallest eigenvalue?

If the graph has ≥ 2 connected components, it is also $-d$.

If the graph is α -expanding, it is $-d + \alpha d$.


Theorem 1: Deciding if a graph (given as black-box sparse adjacency list) either has multiple connected components or is 0.01 -expanding is in QMA.

This also holds for certain distributions over graphs.

Theorem 2: For the same set of distributions, we can show that no efficient QCMA algorithm exists.

Together they prove $\text{QCMA}^D \neq \text{QMA}^D$.

The QMA algorithm



$$|\xi_1\rangle = \frac{|3\rangle + |4\rangle + |5\rangle + |7\rangle + |8\rangle}{\sqrt{5}}$$

$$|\xi_2\rangle = \frac{|1\rangle + |2\rangle + |6\rangle}{\sqrt{3}}$$

Easy to check that

$$|\xi_{\text{solution}}\rangle = \frac{\sqrt{5}}{\sqrt{8}} |\xi_2\rangle - \frac{\sqrt{3}}{\sqrt{8}} |\xi_1\rangle$$

is a eigenvector of eigenvalue $-d$ as well and is orthogonal to $\sum_{x \in \{0,1\}^n} |x\rangle$.


Quantum solution is to provide $|\xi_{\text{solution}}\rangle$ which proves that there are 2 eigenvectors of eigenvalue $-d$.

If the graph is α -expanding, this test **fails** with probability $\alpha/4$ as there is only 1 eigenvector of eigenvalue $-d$ and the next eigenvalue is $-d + \alpha d$.

Sketch of QCMA impossibility result

In the QMA algorithm, we saw that the solution state only depends on the *vertices* in a connected component.

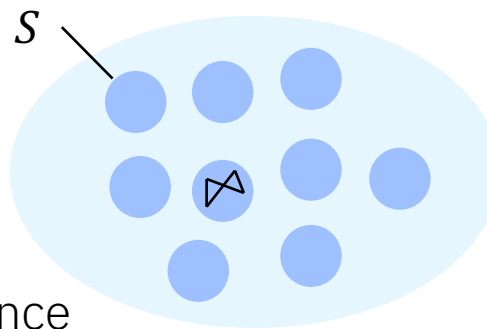
Let's first try to prove an easier subproblem: an impossibility result for all classical proofs that only depend on the *vertices* in one connected component.




Notation: $N = 2^n$ (number of vertices).
 $z = N^{9/10}$ (size of each component).

Setup: YES instance graphs are promised to have N/z expanding connected components each of size z . NO instance graphs have 1 expanding component.

Assume that for every set $S \subset [N]$ of size z , there exists exactly one YES instance that has S has a connected component and that the proof for said YES instance depends only on S .



YES instance



NO instance

Many YES instances correspond to the same proof π

Assume there exists a **bijection** between [YES instance graphs] and [sets of size z].

Then, there are at least $\binom{N}{z} \geq 2^{N^{0.9}}$ YES instances.

And, the number of proofs of length q is at most 2^q .

Pidgeon-hole principle: If there was a QCMA algorithm for all YES instances with proof of length q , then there exists a popular proof $\pi \in \{0,1\}^q$ such that at least $2^{N^{0.9}-q} \geq 2^{N^{0.8}}$ graphs correspond to π .

By the **bijection**, $2^{N^{0.8}}$ [sets of size z] correspond to π .


Ramsey theory: “If the set is large enough, there exists a subset which is structured”

Within the unstructured set of YES instances that correspond to π , there exists a set of YES instances that are structured: they form a sunflower:

Sunflower: A collection Σ of sets of size z such that

- (1) (core) $\exists F \subset [N]$ such that $\forall S \in \Sigma, F \subset S$.
- (2) $\forall x \in (\bigcup_{S \in \Sigma} S) \setminus F$, then

$$\Pr_{S \in \Sigma} [x \in S] \leq \left(\frac{z}{N}\right)^{0.99}.$$



Fixing the proof of a **QCMA** algorithm yields a **BQP** algorithm

Consider the **BQP** algorithm \mathcal{A} generated by fixing the popular proof π into the **QCMA** algorithm.

- \mathcal{A} must answer YES (with high probability) on every graph corresponding to a subset S in the sunflower Σ .
- \mathcal{A} must answer NO (with high probability) on every NO instance graph.

We will show that any algorithm \mathcal{A} achieving these two tasks requires exponential time.

Ramsey theory: “If the set is large enough, there exists a subset which is structured”


Within the unstructured set of YES instances that correspond to π , there exists a set of YES instances that are structured: they form a sunflower:

Sunflower: A collection Σ of sets of size z such that

(1) (core) $\exists F \subset [N]$ such that $\forall S \in \Sigma, F \subset S$.

(2) $\forall x \in (\bigcup_{S \in \Sigma} S) \setminus F$, then

$$\Pr_{S \in \Sigma} [x \in S] \leq \left(\frac{z}{N}\right)^{0.99}.$$



The adversary method

\mathcal{A} must answer YES (with high probability) on every graph corresponding to a subset S in the sunflower Σ .

Sunflower: A collection Σ of sets of size z such that

(1) (core) $\exists F \subset [N]$ such that $\forall S \in \Sigma, F \subset S$.


(2) $\forall x \in (\bigcup_{S \in \Sigma} S) \setminus F$, then

$$\Pr_{S \in \Sigma}[x \in S] \leq \left(\frac{z}{N}\right)^{0.99}.$$

Ideal sunflower corresponding to F :

$$\Sigma_{\text{ideal}} = \{S \subset [N] \mid |S| = z, F \subset S\}.$$


Lemma: Any quantum algorithm cannot distinguish a graph sampled from Σ and a graph sampled from Σ_{ideal} without making an exponential number of queries to the graph.



Then by assumption, \mathcal{A} must answer YES (with high probability) on every graph corresponding to a subset S in the sunflower Σ_{ideal} .

Random walk sampling reduction

\mathcal{A} must answer YES (with high probability) on every graph corresponding to a subset S in the sunflower Σ_{ideal} .



We now generate an algorithm \mathcal{A}' which has F hard-coded. \mathcal{A}' will not be time-efficient (just query-efficient).

Starting with a random point v_1 , \mathcal{A}' applies a random walk sampling to sample points

$$V = \{v_1, v_2, \dots, v_{|F|}\}$$

from the same connected component.


\mathcal{A}' picks a permutation $\sigma: [N] \rightarrow [N]$ mapping $V \rightarrow F$.

Then, \mathcal{A}' simulates \mathcal{A} except when a query to vertex w is made, \mathcal{A}' queries $\sigma^{-1}(w)$.

By the correctness of \mathcal{A} and the random walk being a good sampler (our YES instances are expanding within components), \mathcal{A}' must also succeed with high probability.

This gives a **BQP** algorithm \mathcal{A}' for deciding the original problem (modulo a few technicalities).

Appeal to known quantum query lower bounds



[Ambainis-Childs-Liu¹¹] proved a exponential q. query lower bound for the problem (without proof).

Overall, we have shown that any *classical proof* depending only on the vertices in one connected component cannot help with this problem.

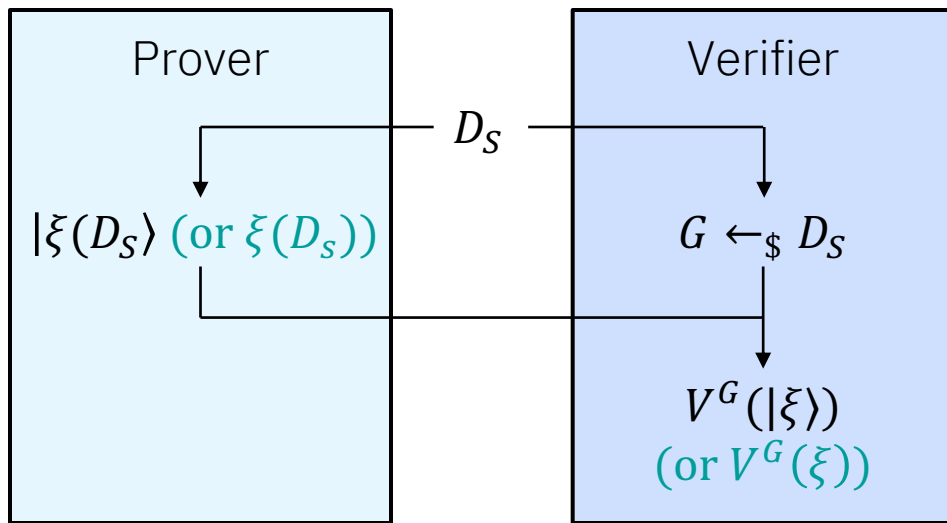
In other words, a “*near-sighted*” classical prover – **one who can only see the overall structure of the graph (connected components) but not internal structures (such as edges, triangles, etc.)** – cannot help.

But recall from our **QMA** algorithm, a “near-sighted” quantum prover can help!

Can we “blind” the classical prover so that she is near-sighted? Blinding the quantum prover has no effect.

Blinding the classical prover

Let D_S be the uniform distribution over all graphs that have a connected component S for S of size z .



Meaning, the prover only knows the graph up to the distribution. This is equivalent to the prover only answering based on the connected component S – this is the same as blinding the prover until it is “near-sighted”.

Blinding the classical prover yields our impossibility result while blinding the quantum prover has no effect.

This proves that there is a *distribution over graphs* which separates **QMA** and **QCMA**.

Removing the “blinding” step seems difficult as there are known obstacles that need avoiding.

The combination of multiple query lower bounds

At a high level, the proof relies on showing that two distributions D_1 and D_4 over boolean fns. cannot be distinguished by quantum query algorithms without *exponentially many* q. queries.

This is proven in 3 steps.

The adversary method D_1 to D_2

This is useful for proving a lightly structured distribution is indistinguishable from a fully structured distribution.

The most famous use of the adversary method is to prove the [BBBV⁹⁷] lower bound for unconstrained search. I.e., distinguishing functions with Hamming weight 1 from the 0 function.

We use the adversary method to connect the sunflower derived from Ramsey theory to the ideal sunflower.

Statistical distance D_2 to D_3

This is useful for proving information theoretic indistinguishability between two distributions.

The techniques used here are Pinsker's inequality and KL divergence bounds.

We use statistical distance to argue that sampling the graph followed by sampling a set of points from the graph is indistinguishable from sampling the points first and then a corresponding graph.

The polynomial method D_3 to D_4

This is useful for proving that two structured distributions whose supports are far from each other are indistinguishable.

The polynomial method requires that the distributions have simple generating functions and is hard to apply, but very powerful when applicable.

We use this to argue that a distribution almost entirely supported on NO instances is indistinguishable from a distribution supported on YES instances.

Theorem [Natarajan-Nirkhe²²]: There is a black-box distribution D over boolean functions for which we can prove that

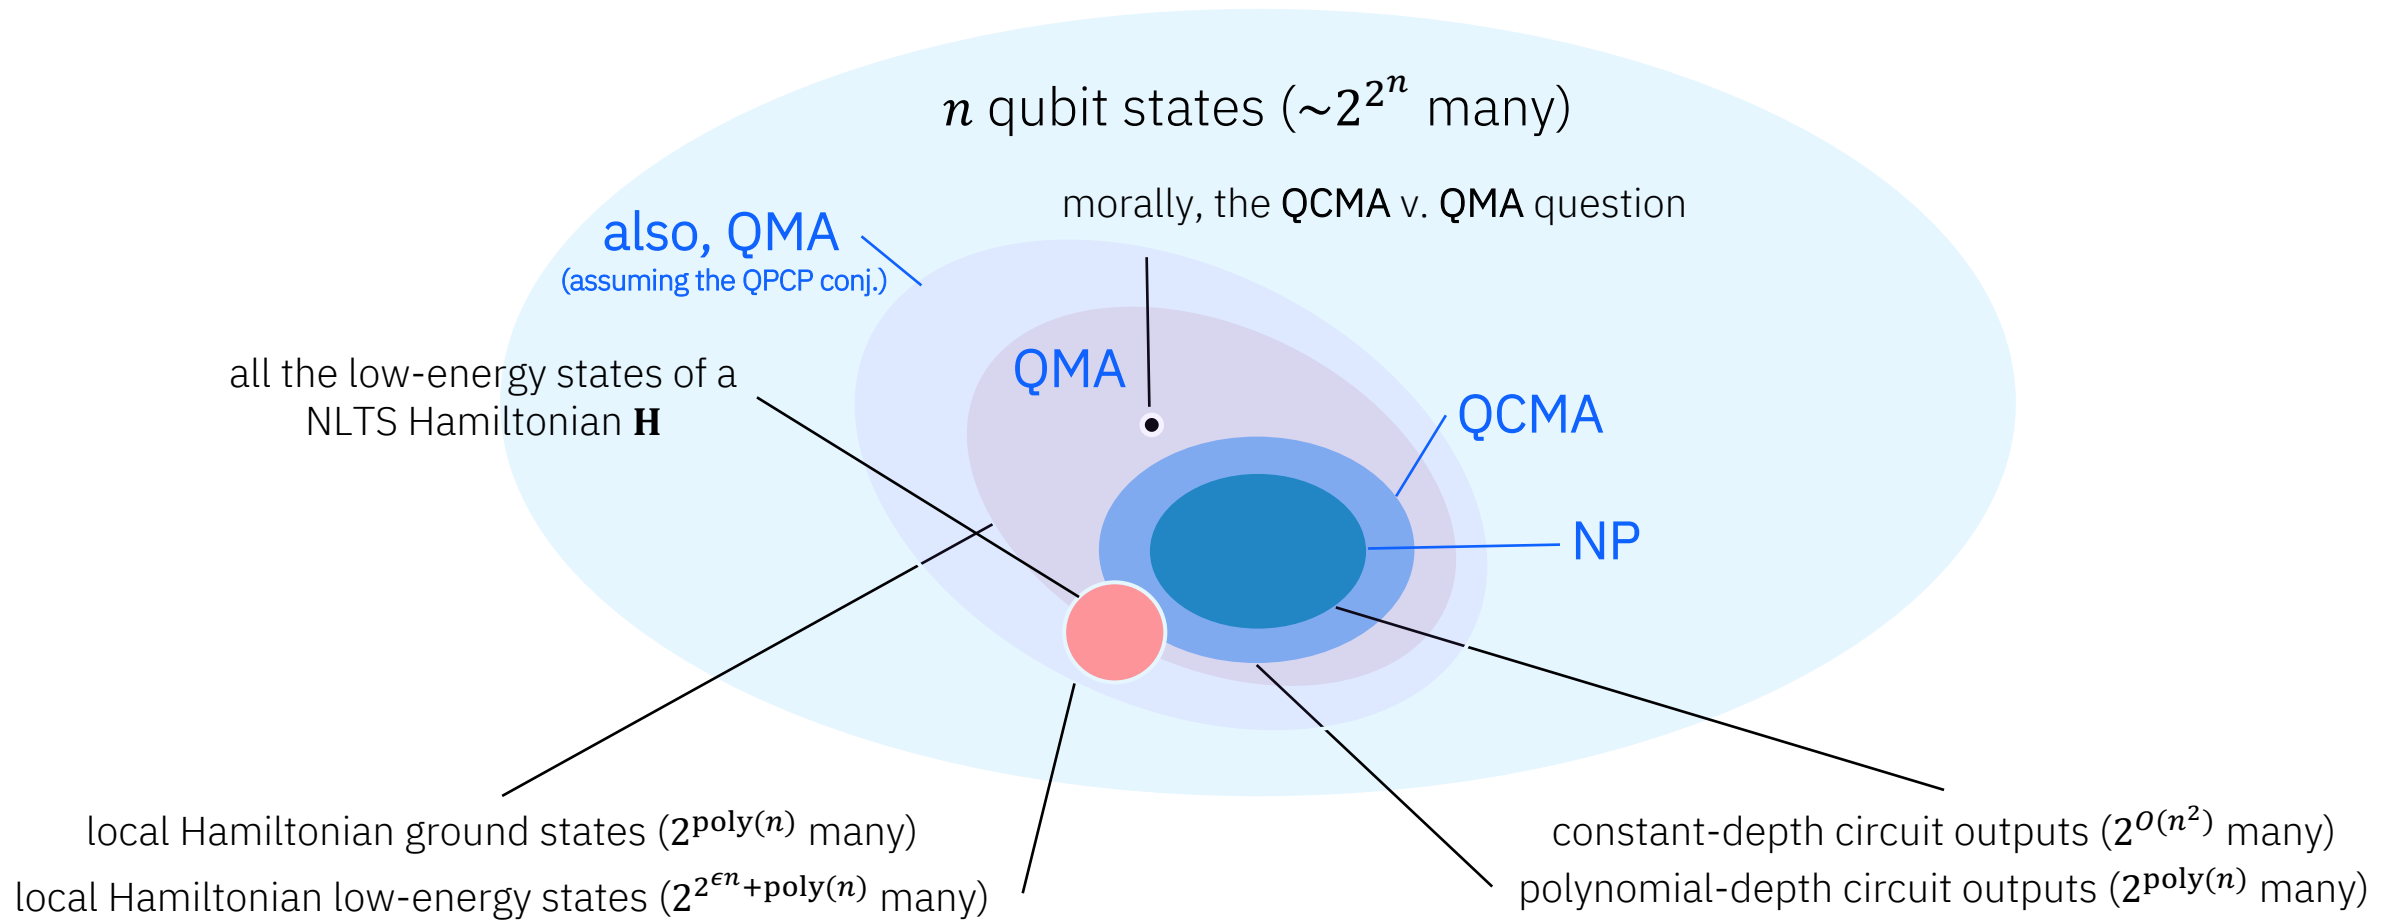
$$\text{QCMA}^D \neq \text{QMA}^D.$$

This is the strongest evidence yet that ground states cannot be classically described.

The NLTS problem
and the QPCP
conjecture

The space of quantum states

and the complexity class that they serve as witness for



The hardness of low-energy states

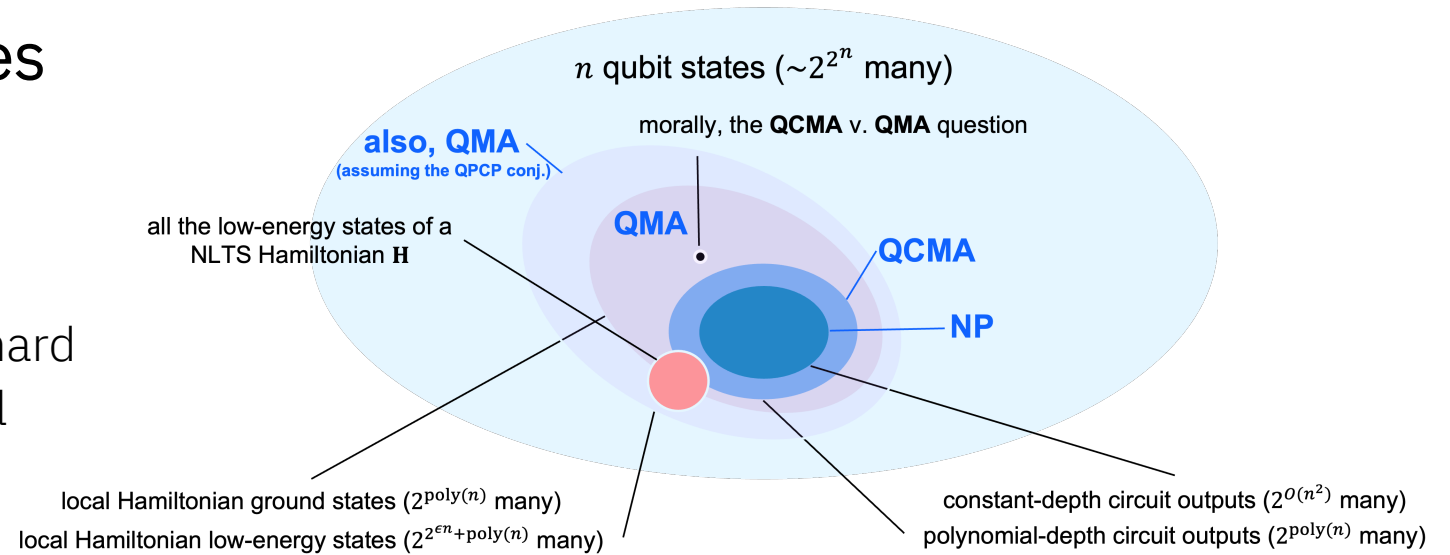
QPCP Conjecture: For some $\epsilon > 0$, it is QMA-hard to decide whether the ground energy of a local Hamiltonian $\mathbf{H} = \sum_{\{i=1\}}^m h_i$ is

(yes) small: $\langle \psi | H | \psi \rangle \leq a$

(no) or large: $\langle \psi | H | \psi \rangle \geq a + \epsilon m$

Simple consequence: Any state $|\psi'\rangle$ of energy $\leq a + \frac{\epsilon}{2}m$ is also a verifiable if the QPCP conj. is true

Therefore, the set of witnesses not only includes ground states but also low-energy states



NLTS Thm [Anshu-Breuckmann-Nirkhe²²]: There exists a local Hamiltonian s.t. no low-energy states ($\leq \epsilon n$) is the output of a constant-depth circuit.

QPCP Conj. + $NP \neq QMA \Rightarrow NLTS$

i.e., a necessary, but not sufficient, consequence

NLTS is an unconditional statement, unlike the previous QCMA v QMA separation.


Sketch of the **NLTS** proof

We first need a technique for proving logarithm-depth circuit depth lower bounds.

Local indistinguishability: Two states $|\psi\rangle$ and $|\psi'\rangle$ are d -locally indistinguishable if every reduced density matrix on d qubits is the same.


Fact: d -locally indistinguishable states have $\Omega(\log d)$ circuit-depth lower bounds.

Corollary: If measuring a quantum state yields a probability distribution p such that there exist two subsets $S_1, S_2 \subset \{0,1\}^n$ with $p(S_1), p(S_2) \geq \Omega(1)$ and the Hamming distance between S_1 and S_2 is $\Omega(n)$, then the quantum state has a $\Omega(\log n)$ circuit depth lower bound.



The challenge is to find a local Hamiltonian for which every low-energy state when measured yields such a “well-spread” distribution.

The low-energy subset of a linear-distance error-correcting code is supported like a “well-spread” dist.




Sketch of the **NLTS** proof

By considering quantum LDPC error-correcting codes of constant-rate and linear-distance (plus an additionally robustness property), we can find the “well-spread” property and construct local Hamiltonians with the **NLTS** property.

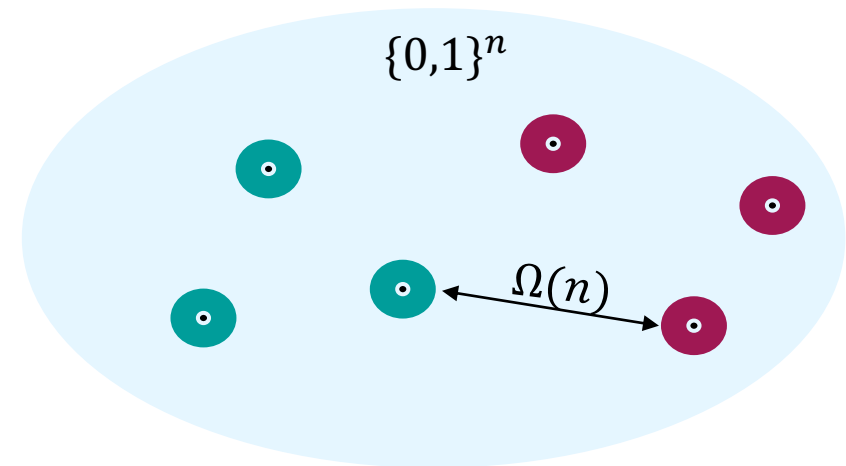
The key is using the *uncertainty lemma* to argue that for any state, measuring in either the standard or Hadamard basis must yield a well-spread distribution.

This is because the quantum code simultaneously corrects large X- and Z-errors.



The challenge is to find a local Hamiltonian for which every low-energy state when measured yields such a “well-spread” distribution.

The low-energy subset of a linear-distance error-correcting code is supported like a “well-spread” dist.

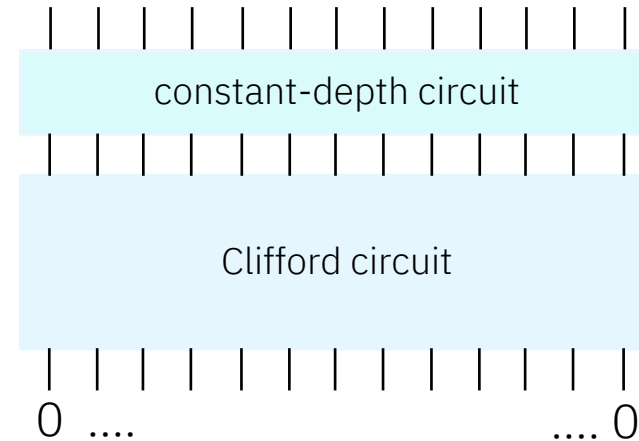


Proving stronger lower bounds than **NLTS**

Constant-depth quantum circuits are just one of *many* classical witnesses that can be provided for an **NP** proof.

QPCP Conj. + $\text{NP} \neq \text{QMA} \Rightarrow$
lower-bounds for all families of **NP** witnesses

Open question: Can we prove lower-bounds for some other families of **NP** witnesses? Is there is a family of local Hamiltonians for which all known **NP** witnesses are insufficient?



Any state of this form is also a **NP** witness.

NLTS+ conjecture: There exists local Hamiltonian such that all such states have energy $\geq \epsilon n$.

Our proof of **NLTS** does not satisfy this!

Quantum search-
to-decision
reductions

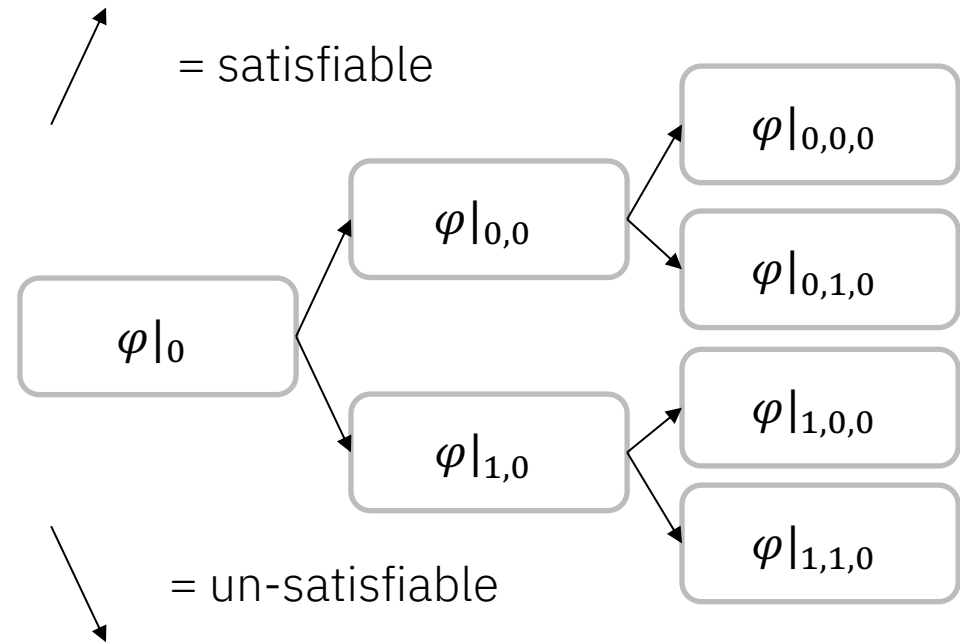
NP has search-to-decision reductions

Say there is a black box which takes as input 3SAT formulas φ and outputs (with $\text{pr} = 1$) if they are satisfiable or not...

Crucially, it does not tell you the solution (satisfying assignment) $x \in \{0,1\}^n$ s.t. $\varphi(x) = 1$.

This is a black box for NP and repeated uses can be used to build a solution $x \in \{0,1\}^n$.

Let $\varphi|_{a_1, a_2, \dots, a_k}$ be the restriction of φ on the first k variables.



After n queries, we learn a complete satisfying assignment.

$$\text{NP-Search} \subseteq \text{P}^{\text{Decision-NP}}$$

Does **QMA** have search-to-decision reductions?

In classical CS theory, defining decision problems as the *de facto* model of computation is justified by search-to-decision reductions.

What about in quantum CS? Is the same definition justified? Or do we need to rectify our *de facto* notion of computation?

[Irani-Natarajan-Nirkhe-Rao-Yuen²¹]:

Theorem 1: QMA-search is reducible to 1-query PP-decision

Theorem 2: Oracle proof that QMA-search not reducible to QMA-decision

Theorem 1: Let (\mathbf{H}, a, b) be a local Hamiltonian problem. If the problem is a YES instance, there exists a BQP algorithm making 1-query to a PP-oracle such that the output state of the algorithm has energy $\leq a + (b - a)/2$.

Theorem 2: We show that a unitary oracle corresponding to a QCMA v QMA separation also proves that QMA-search is not reducible to QMA-decision

It's open whether our distribution oracle also proves such an impossibility result.

Understanding quantum states

The difference in quantum search-to-decision reductions and classical search-to-decision reductions suggests that quantum states cannot be entirely studied through the lens of decision problems.

We need to better understand the complexity of quantum states and ideally how to prove polynomial-depth lower bounds for quantum states without oracles.

Other open questions:

How can we remove the distributions from the oracle separations between **QCMA** and **QMA**?

Can we use the techniques to prove oracle separations for other quantum complexity classes such as **QMA(2)** and **QMA**?

What is the power of **BQP^{QMA}**? It lies somewhere between **QMA** and **QCMA**.

Thank you for listening.

Chinmay Nirkhe (IBM Quantum Cambridge)

IBM