

Why is quantum mechanics so hard to describe?

A computational lens on the problem



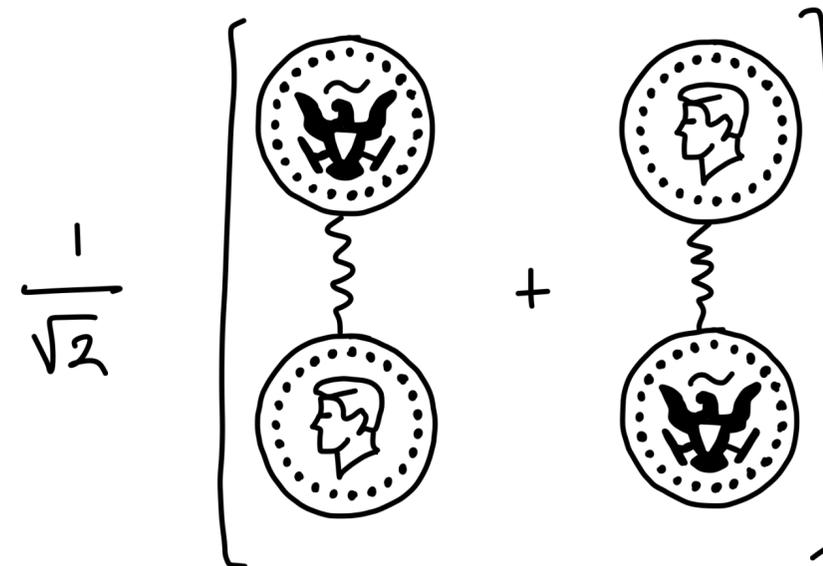
Quantum states are exponentially complex

a classical system



$|0\rangle |1\rangle$

a quantum system



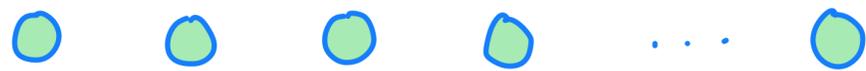
$$\frac{|0\rangle |1\rangle + |1\rangle |0\rangle}{\sqrt{2}}$$

quantum entanglement

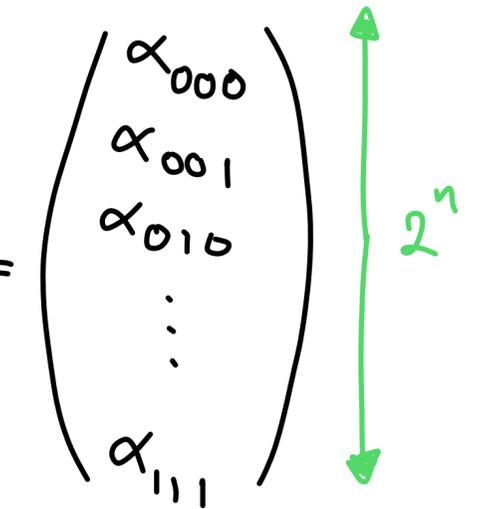
doesn't suffice to describe each individual particle

Quantum states are exponentially complex

n particles
(each with 2 states)

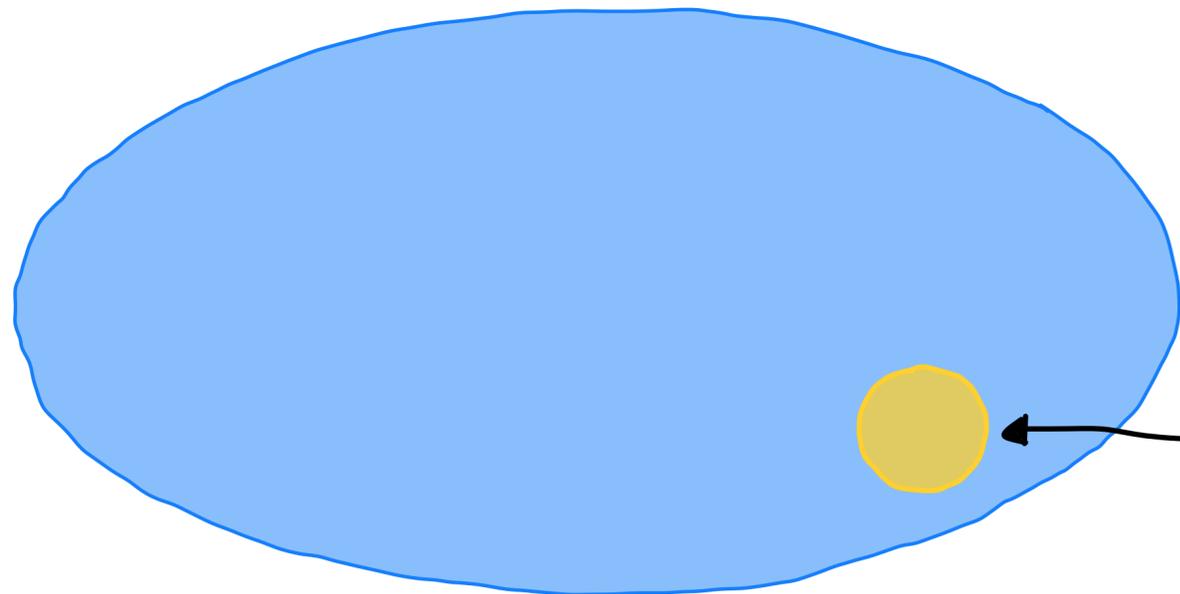


$$\text{quantum state} = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle =$$



Sum over all n -bit strings

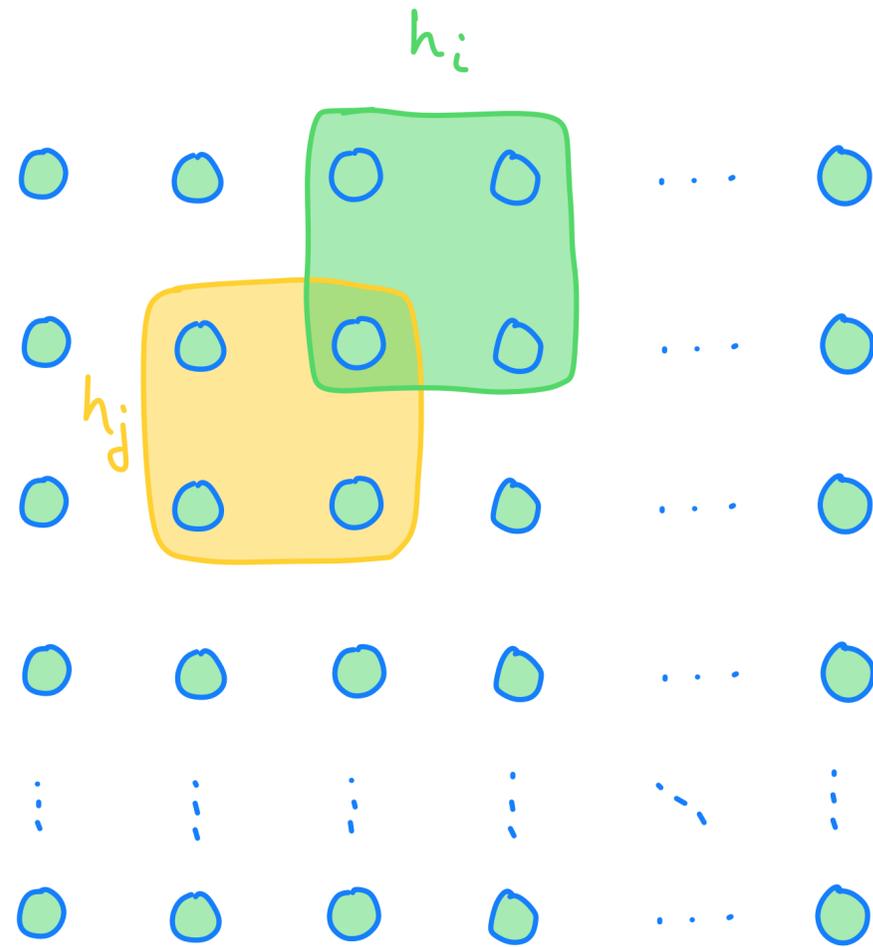
space of n particle states = \mathbb{C}^{2^n}



physically relevant corner

How is it possible to represent quantum states without exponential complexity?

The physically relevant corner



defined by local interactions

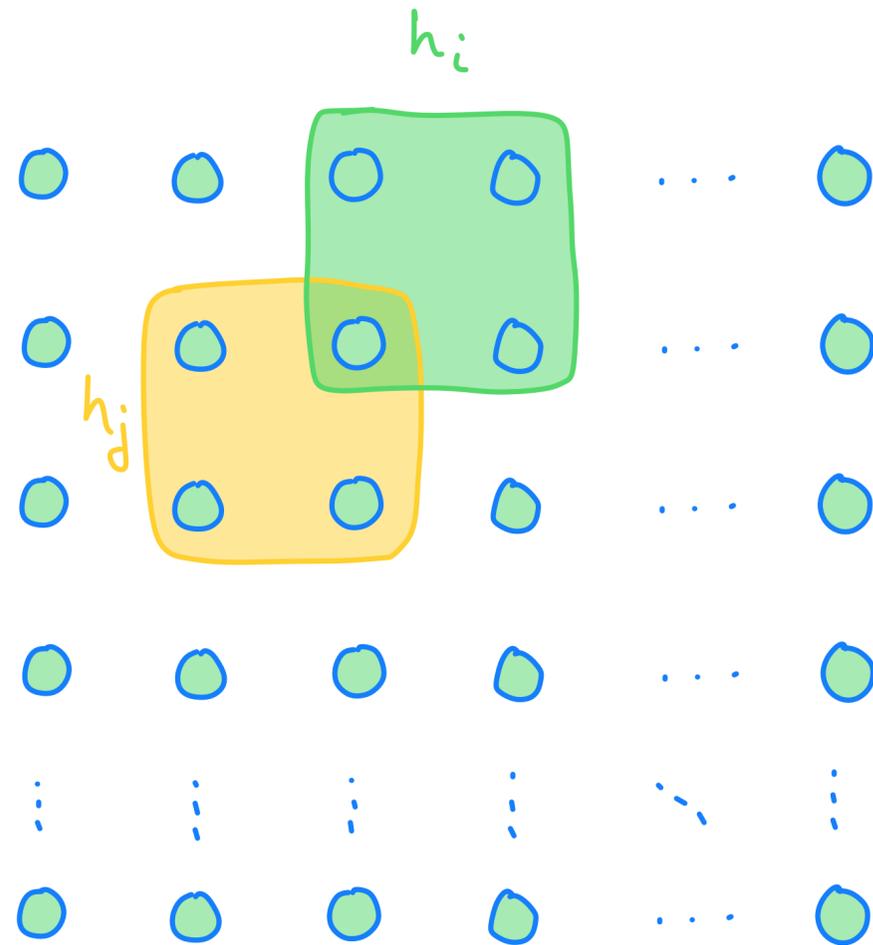
each k -local interaction is described by

Hamiltonian $h_i = \begin{pmatrix} \dots & & \\ \vdots & \ddots & \vdots \\ \dots & & \dots \end{pmatrix}$

The matrix is shown with a horizontal double-headed arrow below it labeled 2^k and a vertical double-headed arrow to its right labeled 2^k , indicating its dimensions.

The physically relevant corner

Defined by local interactions



each k -local interaction is described by

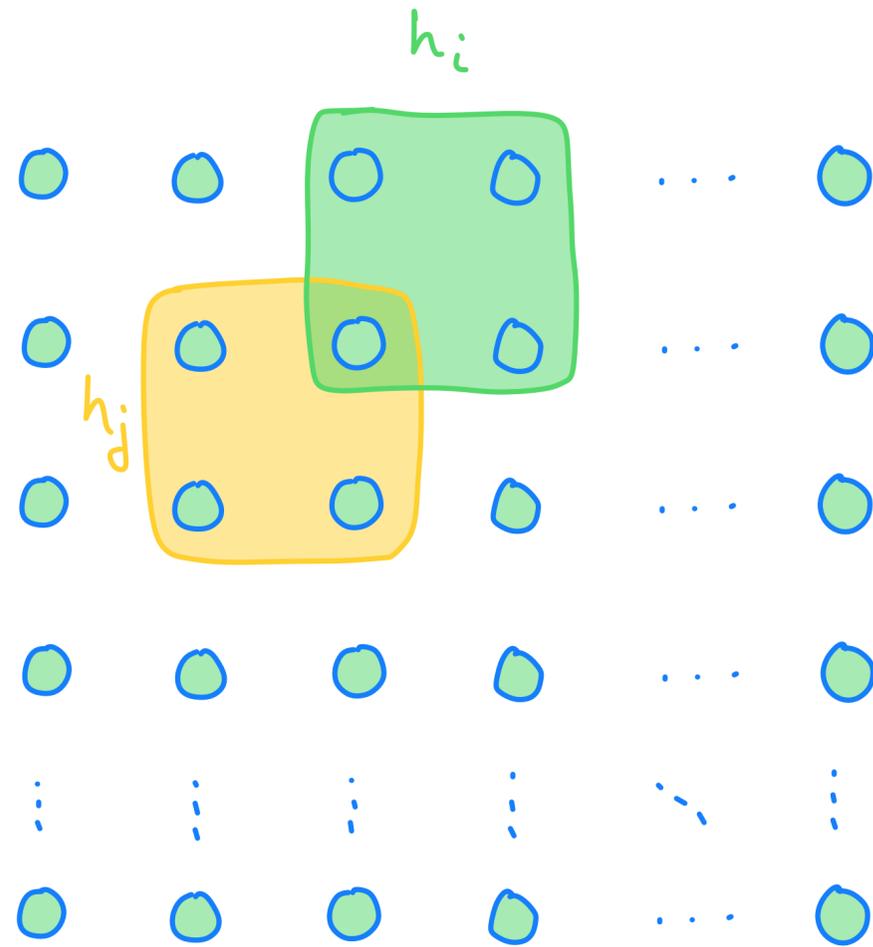
$$\text{Hamiltonian } h_i = \begin{pmatrix} \dots & & \\ \vdots & \ddots & \vdots \\ \dots & & \dots \end{pmatrix}$$

← 2^k →

↑ 2^k ↓

The physically relevant corner

Defined by local interactions

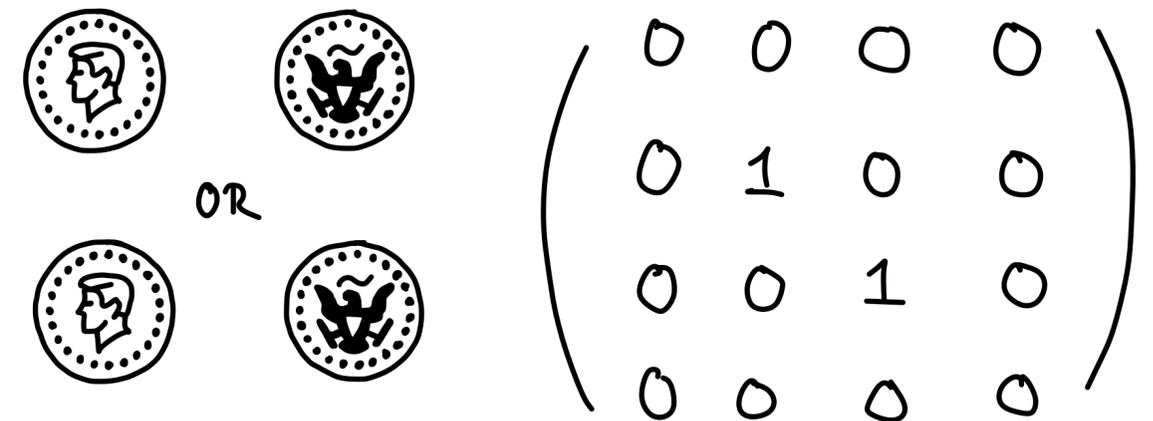


each k -local interaction is described by

Hamiltonian $h_i = \begin{pmatrix} \dots & & \\ & \ddots & \\ \dots & & \dots \end{pmatrix}$

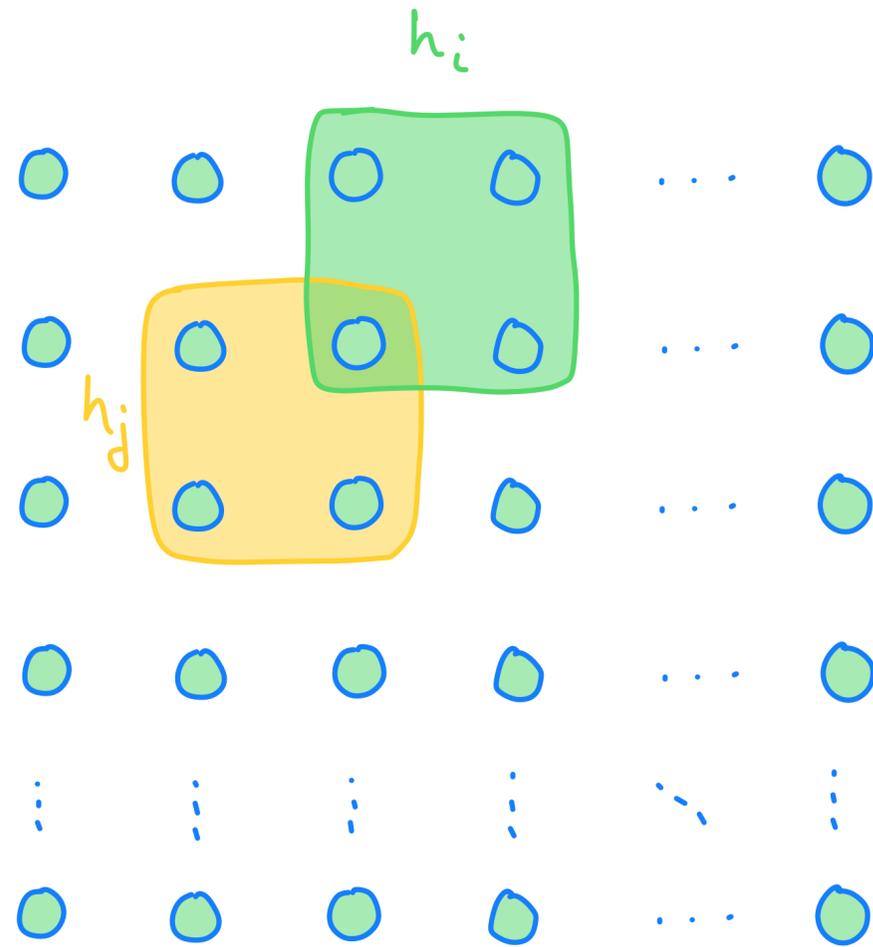
$\uparrow 2^k$
 $\downarrow 2^k$
 $\leftarrow 2^k \rightarrow$

example h_i prefers



The physically relevant corner

Defined by local interactions



each k -local interaction is described by

$$\text{Hamiltonian } h_i = \left(\begin{array}{c} \dots \\ \vdots \\ \vdots \\ \dots \end{array} \right) \begin{array}{l} \uparrow 2^k \\ \downarrow 2^k \end{array}$$

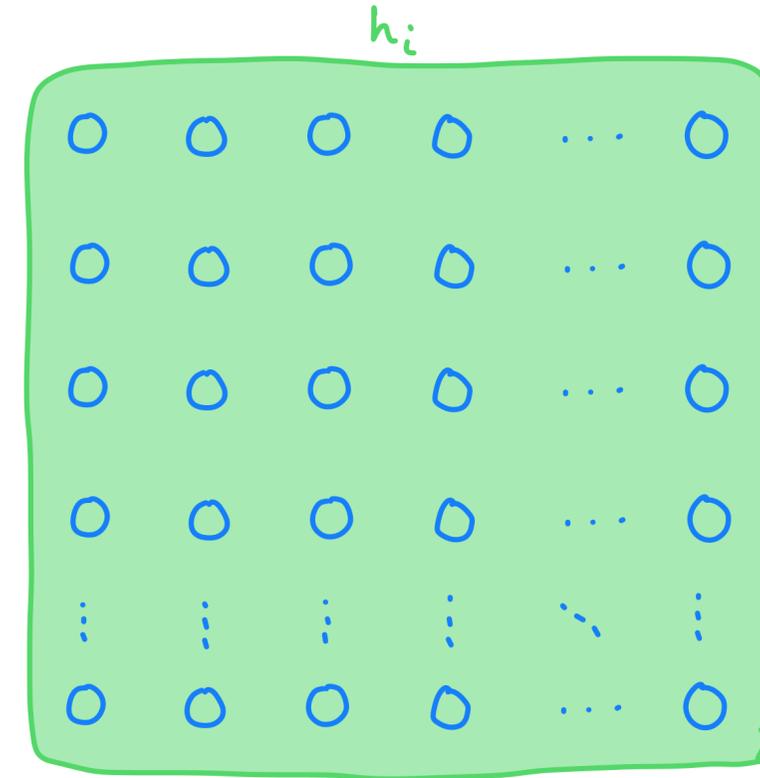
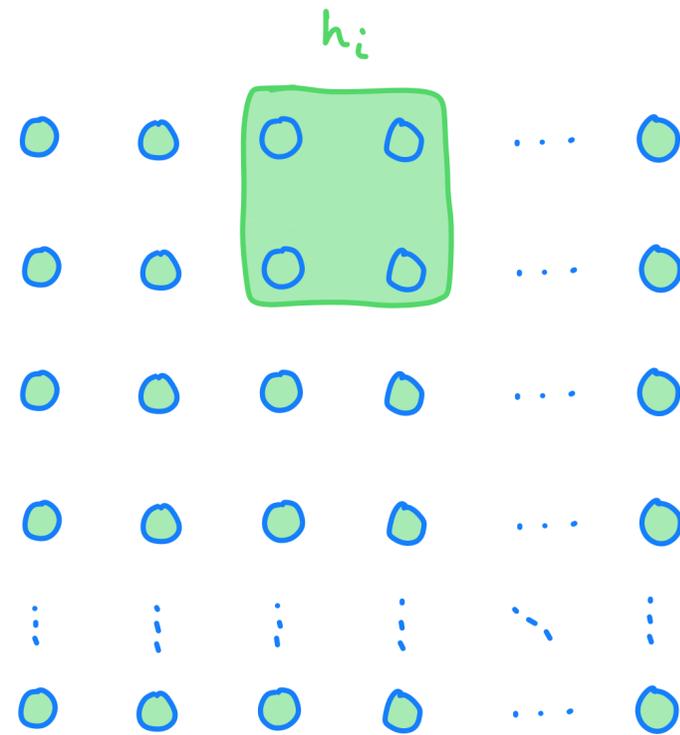
$$\leftarrow 2^k \rightarrow$$

example h_i prefers $\frac{1}{\sqrt{2}} \left[\begin{array}{c} \text{Obama} \\ \text{Obama} \end{array} \right] + \left[\begin{array}{c} \text{Clinton} \\ \text{Clinton} \end{array} \right]$

$$\begin{pmatrix} -\frac{1}{2} & 0 & 0 & -\frac{1}{2} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\frac{1}{2} & 0 & 0 & -\frac{1}{2} \end{pmatrix}$$

A global view on interactions

Local \rightarrow Global Phenomenon



Hamiltonian $h_i = \begin{pmatrix} \dots & & \\ \vdots & \times & \vdots \\ \dots & & \dots \end{pmatrix}$

Horizontal dimension: 2^k

Vertical dimension: 2^k

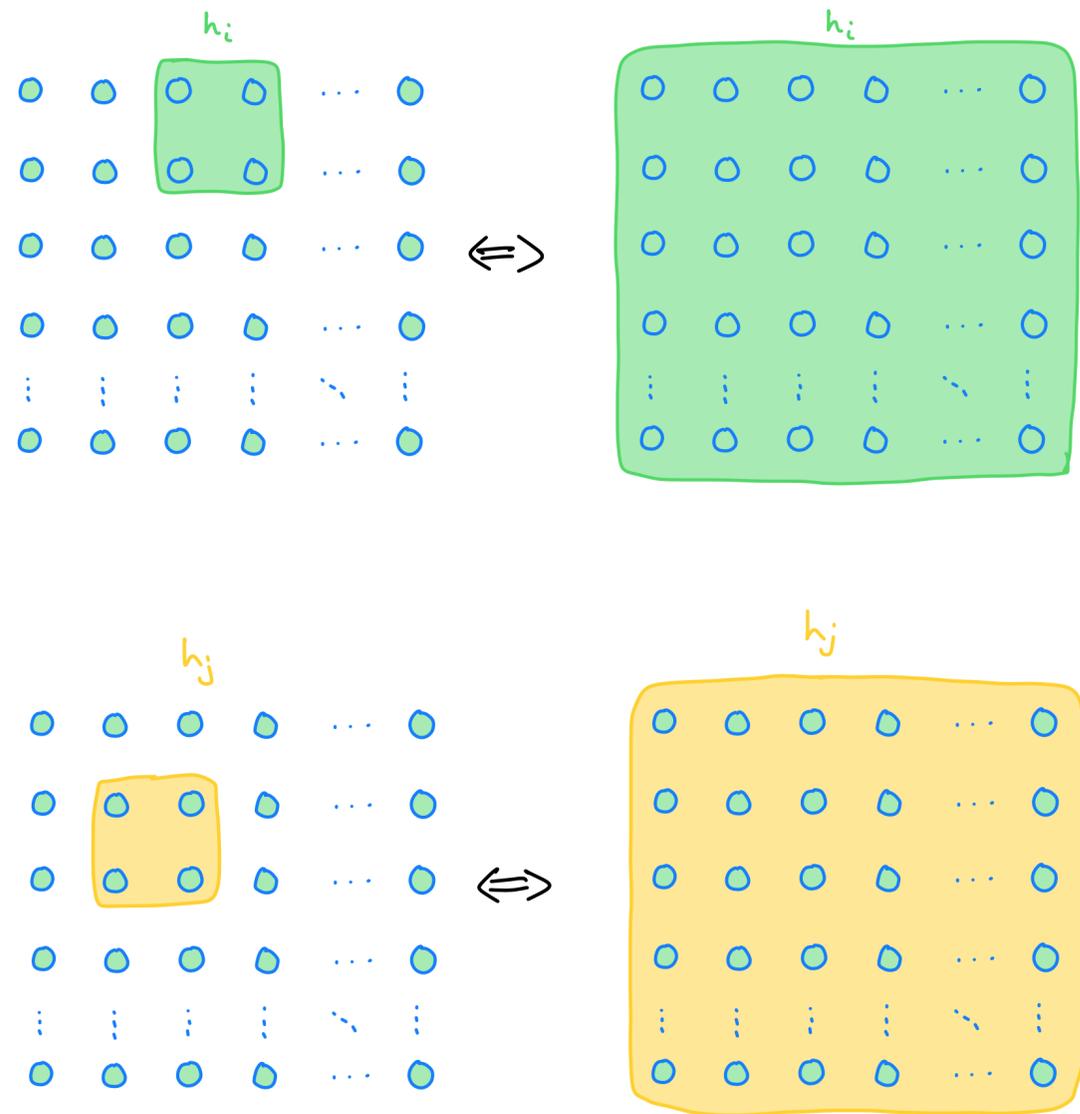
$\mathbb{1} \otimes h_i = \begin{pmatrix} h_i & & & \\ & h_i & & \\ & & \dots & \\ & & & h_i \end{pmatrix}$

Horizontal dimension: 2^n

Vertical dimension: 2^n

A global view on interactions

Local → Global Phenomenon



Each local interaction can be viewed globally
as a matrix $h_i \in \mathbb{C}^{2^n \times 2^n}$.

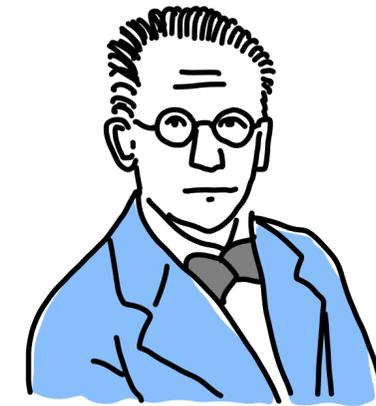
$$\text{Together } H = \sum_i h_i$$

is called the local Hamiltonian and
is the matrix governing physical interactions

A global view on interactions

Local → Global Phenomenon

- Physical systems are governed by local Hamiltonians \mathbf{H}
- Schrödinger equation:



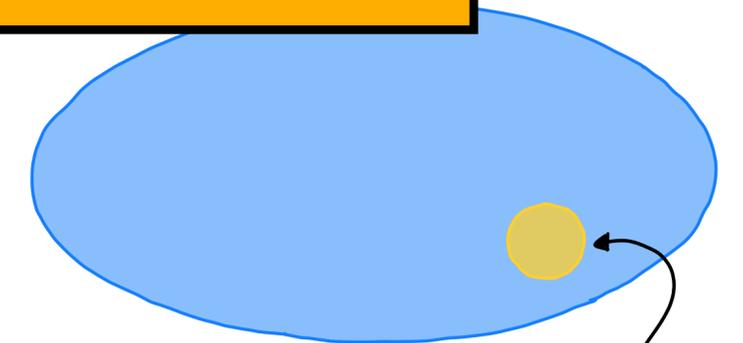
Erwin Schrödinger

$$\frac{d|\psi\rangle}{dt}$$

Do we still need $\exp(n)$ many bits to describe ground states?
Or is there a better way to describe the “physically relevant corner”?

- The relevant physical states in physics are the eigenstates of \mathbf{H} .
These are the stable solutions: $\mathbf{H}|\psi\rangle = E_\psi|\psi\rangle \implies |\psi(t)\rangle = e^{iE_\psi t}|\psi(0)\rangle$
- In particular, physics is interested in understanding the **ground states of \mathbf{H}** (and any low-energy approximations)
- **Only a tiny fraction of all states** are also ground states of local Hamiltonians. This gives a formal definition to the “physically relevant corner”.

states = \mathbb{C}^{2^n}



~~physically relevant corner~~
ground states of Hamiltonians
(and their approximations)

A global view on interactions

A brief primer on ground states

- The local Hamiltonian \mathbf{H} is a Hermitian matrix ($\mathbf{H} = \mathbf{H}^\dagger$) as it is the sum of local interactions each describable by a Hermitian matrix
- Every Hermitian matrix can be *diagonalized* into its eigenbasis

• In the bra-ket physics notation, we write this as $\mathbf{H} = \sum_{i=1}^{2^n} E_i |\psi_i\rangle\langle\psi_i|$

- The ground energy E_{\min} is $:= \min_i E_i$ and the ground states are any corresponding vectors/states
 - Equivalently, $|\psi_{\min}\rangle$ is a ground state iff $\mathbf{H} |\psi_{\min}\rangle = E_{\min} |\psi_{\min}\rangle \iff \langle\psi_{\min}| \mathbf{H} |\psi_{\min}\rangle = E_{\min}$.
 - Mathematically, E_{\min} is the minimum eigenvalue of \mathbf{H} and $|\psi_{\min}\rangle$ is the partner eigenvector.

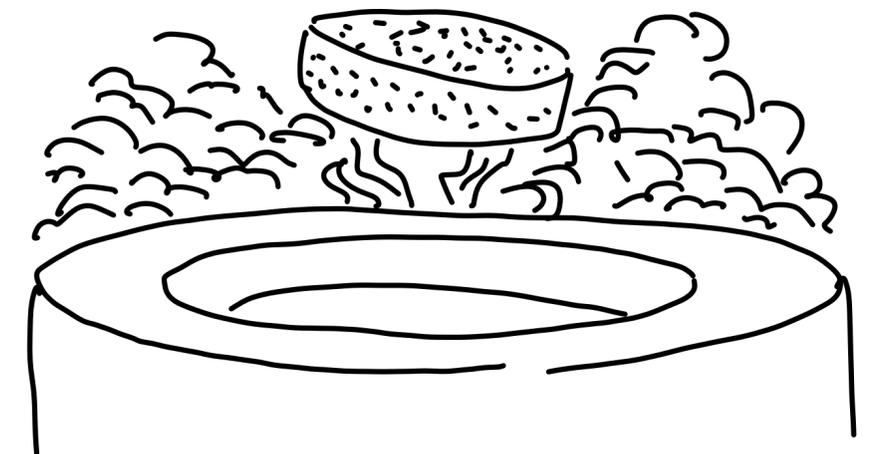
$$H = \begin{pmatrix} E_1 & & & \\ & E_2 & & \\ & & \ddots & \\ & & & E_{2^n} \end{pmatrix}$$

when written in eigenbasis

The central problem in condensed matter physics

Given the description of a local Hamiltonian system \mathbf{H} , calculate a description of the ground state as well as calculate properties of the ground state.

- Spectral gap of the Hamiltonian
- Correlations between different particles
- Entropy structure of ground states and complexity of many-body entanglement
- Topological symmetries / phases of matter
- Superfluidity/superconductivity and other phenomena



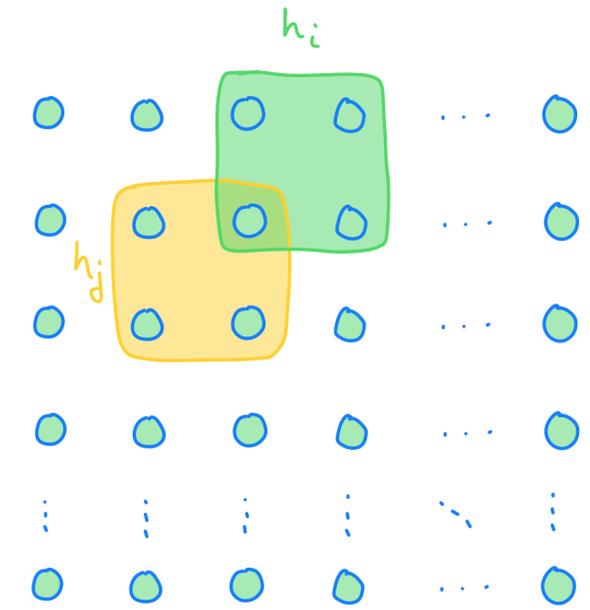
The central problem in condensed matter physics

Given the description of a local Hamiltonian system \mathbf{H} ,
calculate a description of the ground state as well as
calculate properties of the ground state.

Calculating ground states of n -qubit systems involves "stitching"
together local solutions of the individual local terms $h_i \in \mathbb{C}^{2^k \times 2^k}$ for $k = O(1)$

The complexity of ground states

- Let's start with an "easier" question.
- **Input:** The classical description of a local Hamiltonian \mathbf{H}
 - An m -term and k -local Hamiltonian system can be classically described in $O(m \cdot 2^{2k})$ bits
 - When $m = \text{poly}(n)$ and $k = O(1)$, this is $\text{poly}(n)$
- **Output:** The classical value of the ground energy: E_{\min} .
- **First question:** Why is this the right computational task to consider first? What do algorithms for this task tell us?



$$\begin{aligned} E_{\min} &= \lambda_{\min}(H) \\ &= \lambda_{\min}\left(\sum_i h_i\right) \end{aligned}$$

The complexity of ground energy

- **Input:** The classical description of a local Hamiltonian \mathbf{H} .
Output: The classical value of the ground energy: E_{\min} .
- This is an example of a classical input \rightarrow classical output problem
- There is a naive classical algorithm for it as well!
 - Since \mathbf{H} is a $2^n \times 2^n$ matrix, we can calculate E_{\min} (on a classical computer) by diagonalizing.
 - This algorithm runs in time $\exp(n)$. Is there a better algorithm (perhaps a quantum algorithm)?

The complexity of ground energy

- **Input:** The classical description of a local Hamiltonian \mathbf{H} .
Output: The classical value of the ground energy: E_{\min} .
- If there was a quantum algorithm which given \mathbf{H} can *construct* the ground state, $|\psi_{\min}\rangle$, then there is a quantum algorithm for computing E_{\min} .
- **Proof:** There is a well-known algorithm (phase estimation) for estimating (to arbitrary accuracy) $\langle \psi_{\min} | \mathbf{H} | \psi_{\min} \rangle$ (which equals E_{\min}).
- **Corollary:** A *lower bound* on the complexity of computing ground energy implies a *lower bound* on the complexity of generating ground states.

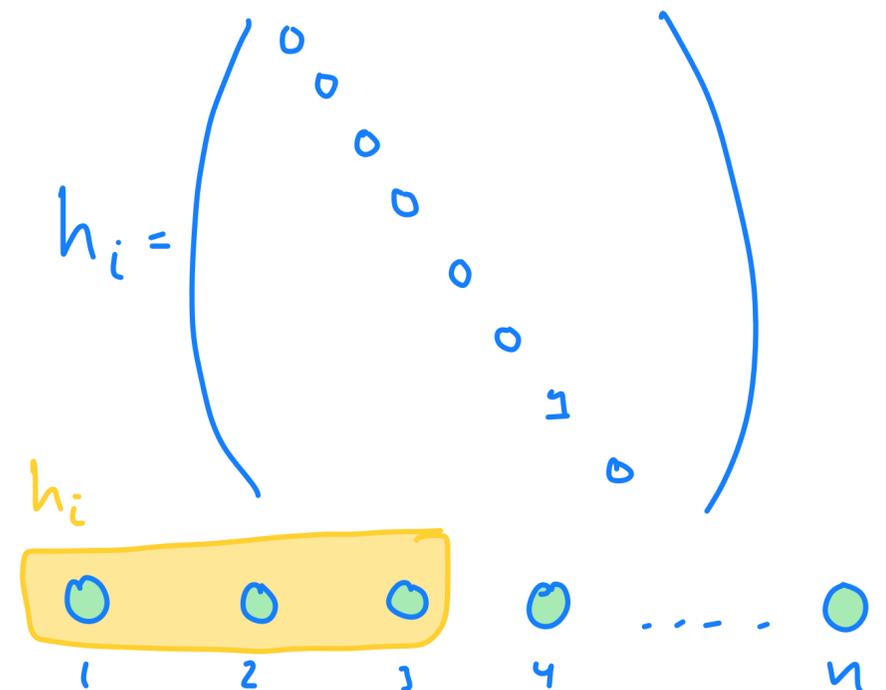
The complexity of ground energy

Lower bounds

- **Input:** The classical description of a local Hamiltonian \mathbf{H} .
Output: The classical value of the ground energy: E_{\min} .
- If we consider \mathbf{H} where every local term h_i is a diagonal matrix, then this is a generalization of the 3-SAT problem from classical computer science. Famously NP-complete!
 - Any algorithm for computing ground energy also must be able to solve 3-SAT —this is a *lower bound* on the computational complexity.
 - $3\text{-SAT} \leq_p \text{LH}$

3-SAT clause

$$\varphi_i = (x_1 \vee x_2 \vee \neg x_3)$$



The complexity of ground energy

Tight characterization

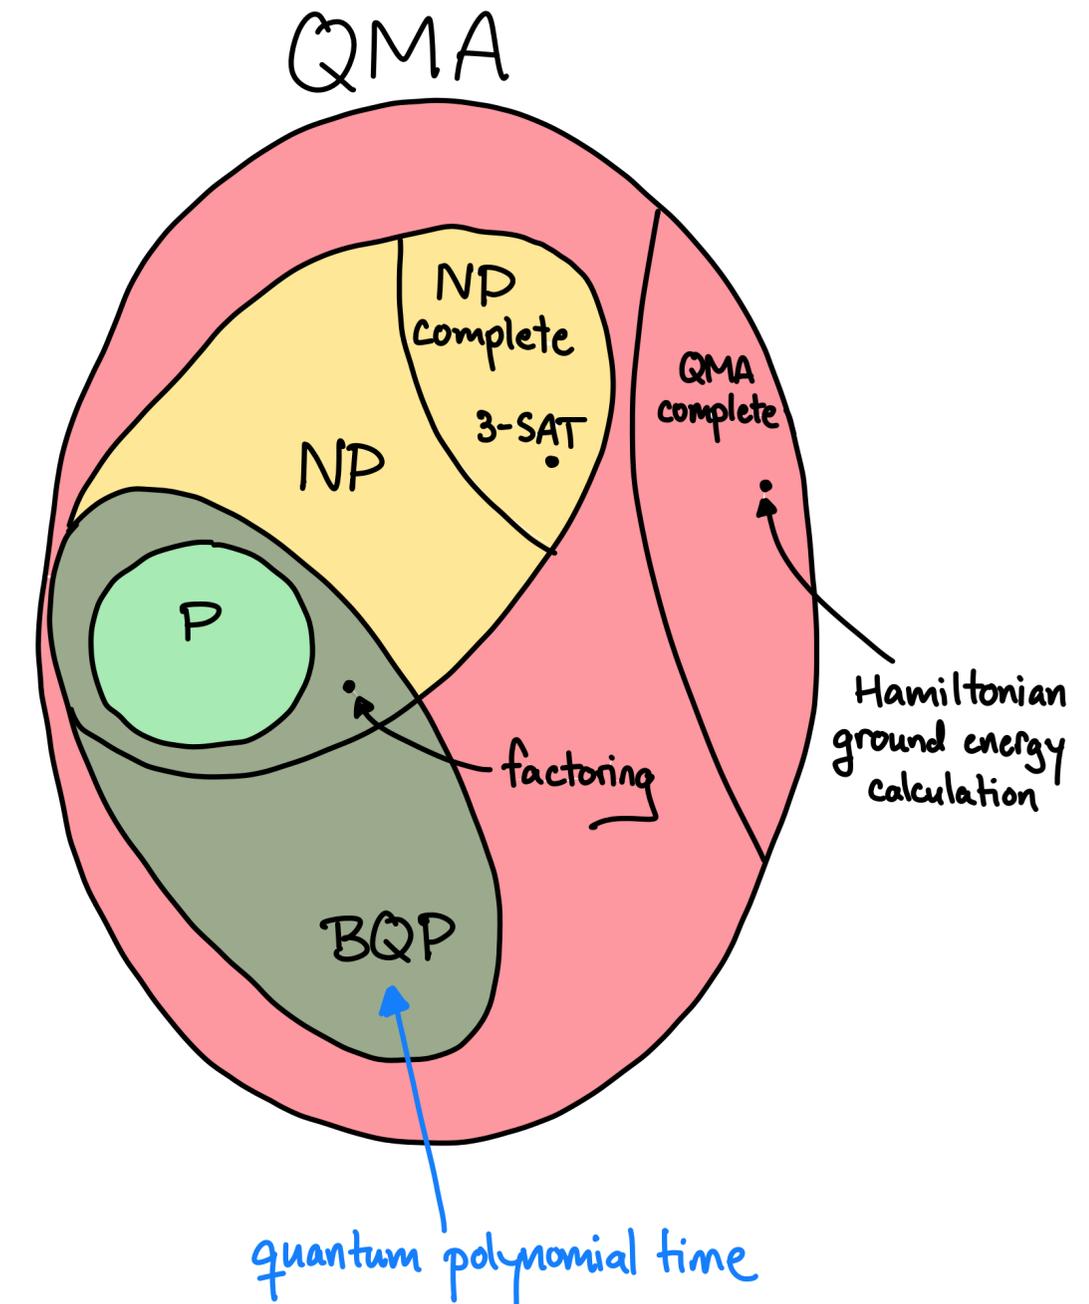


Alexei Kitaev

- **Input:** The classical description of a local Hamiltonian \mathbf{H} .
Output: The classical value of the ground energy: E_{\min} .
- In 2002, Kitaev, proved that any *non-deterministic quantum computation* (quantum analog of NP) can be reduced to calculating ground energy.
 - I.e., an algorithm (either quantum or classical) for computing ground energy would imply an algorithm for any problem that is checkable with a “quantum proof”
 - In particular, proves a distinct *lower bound* (i.e., *hardness result*) for calculating ground states or their descriptions
 - This proves that the central problem of condensed matter physics is very hard!

Exactly how hard?

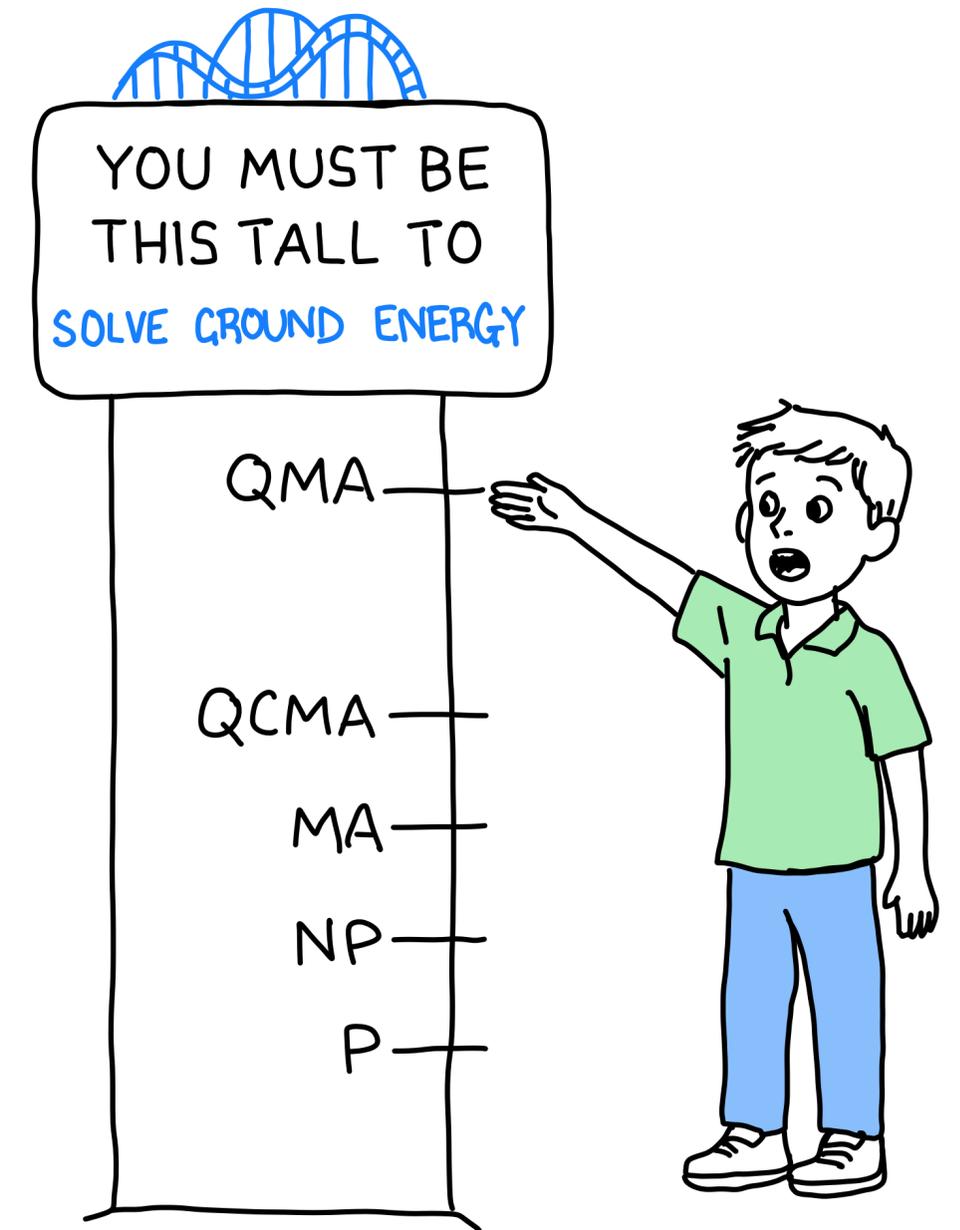
- Kitaev proved calculating ground energy is intractable for quantum computers
- However, as any computer scientist will tell you, intractability is relative
- Specifically, Kitaev proved that calculating ground energy exactly is QMA – complete
 - QMA stands for “Quantum Merlin-Arthur”
 - It is a quantum analog of non-deterministic polynomial time (NP)
- Two natural follow up questions have been driving forces of my research:
 - **Is there an easier (classical or quantum) algorithm for approximating the ground energy? Or an impossibility result?**
 - **How hard are ground states to describe and can we describe ground states efficiently (such as a circuit preparing the state)?**



Part II: Complexity theory

To classify the complexity of problems relating to computing ground energy, we need to spend a moment describing complexity classes.

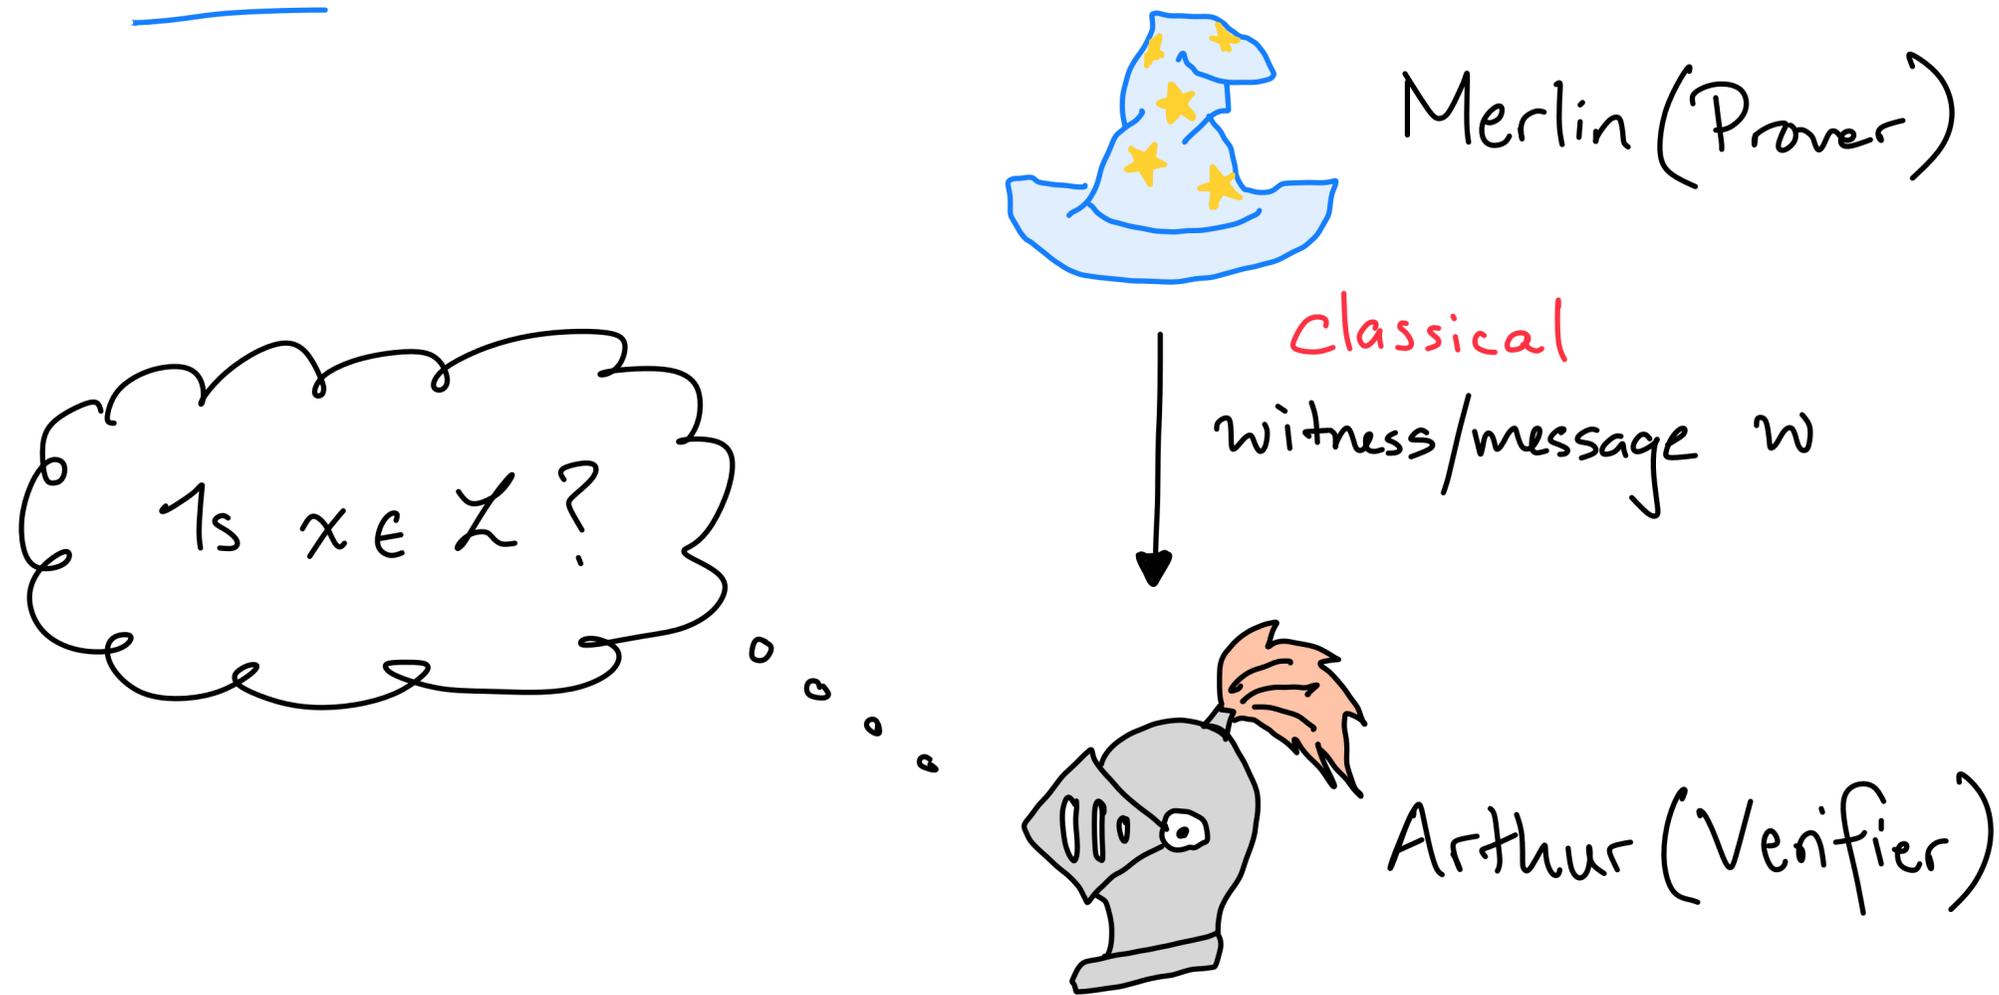
This will give us a good yard stick for measuring the hardness of problems.



NP = “non-deterministic polynomial time”

MA = “Merlin-Arthur”

NP/MA



Cook-Levin '71, Karp '72:

3-SAT, 3-Color, Traveling Salesman, Vertex Cover, Hamiltonian Path, et. Are NP-complete.

Arthur is a *classical* polytime computer and computes a fn $v(x, w)$

There are two appropriate quantum analogs to NP

QCMA

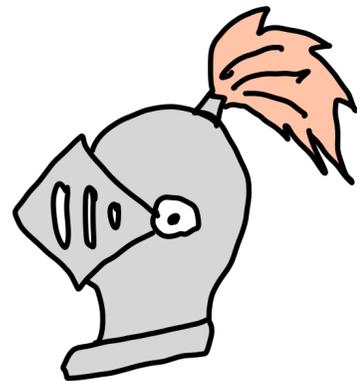
quantum classical
Merlin Arthur

Is $x \in \mathcal{L}$?



Merlin (Prover)

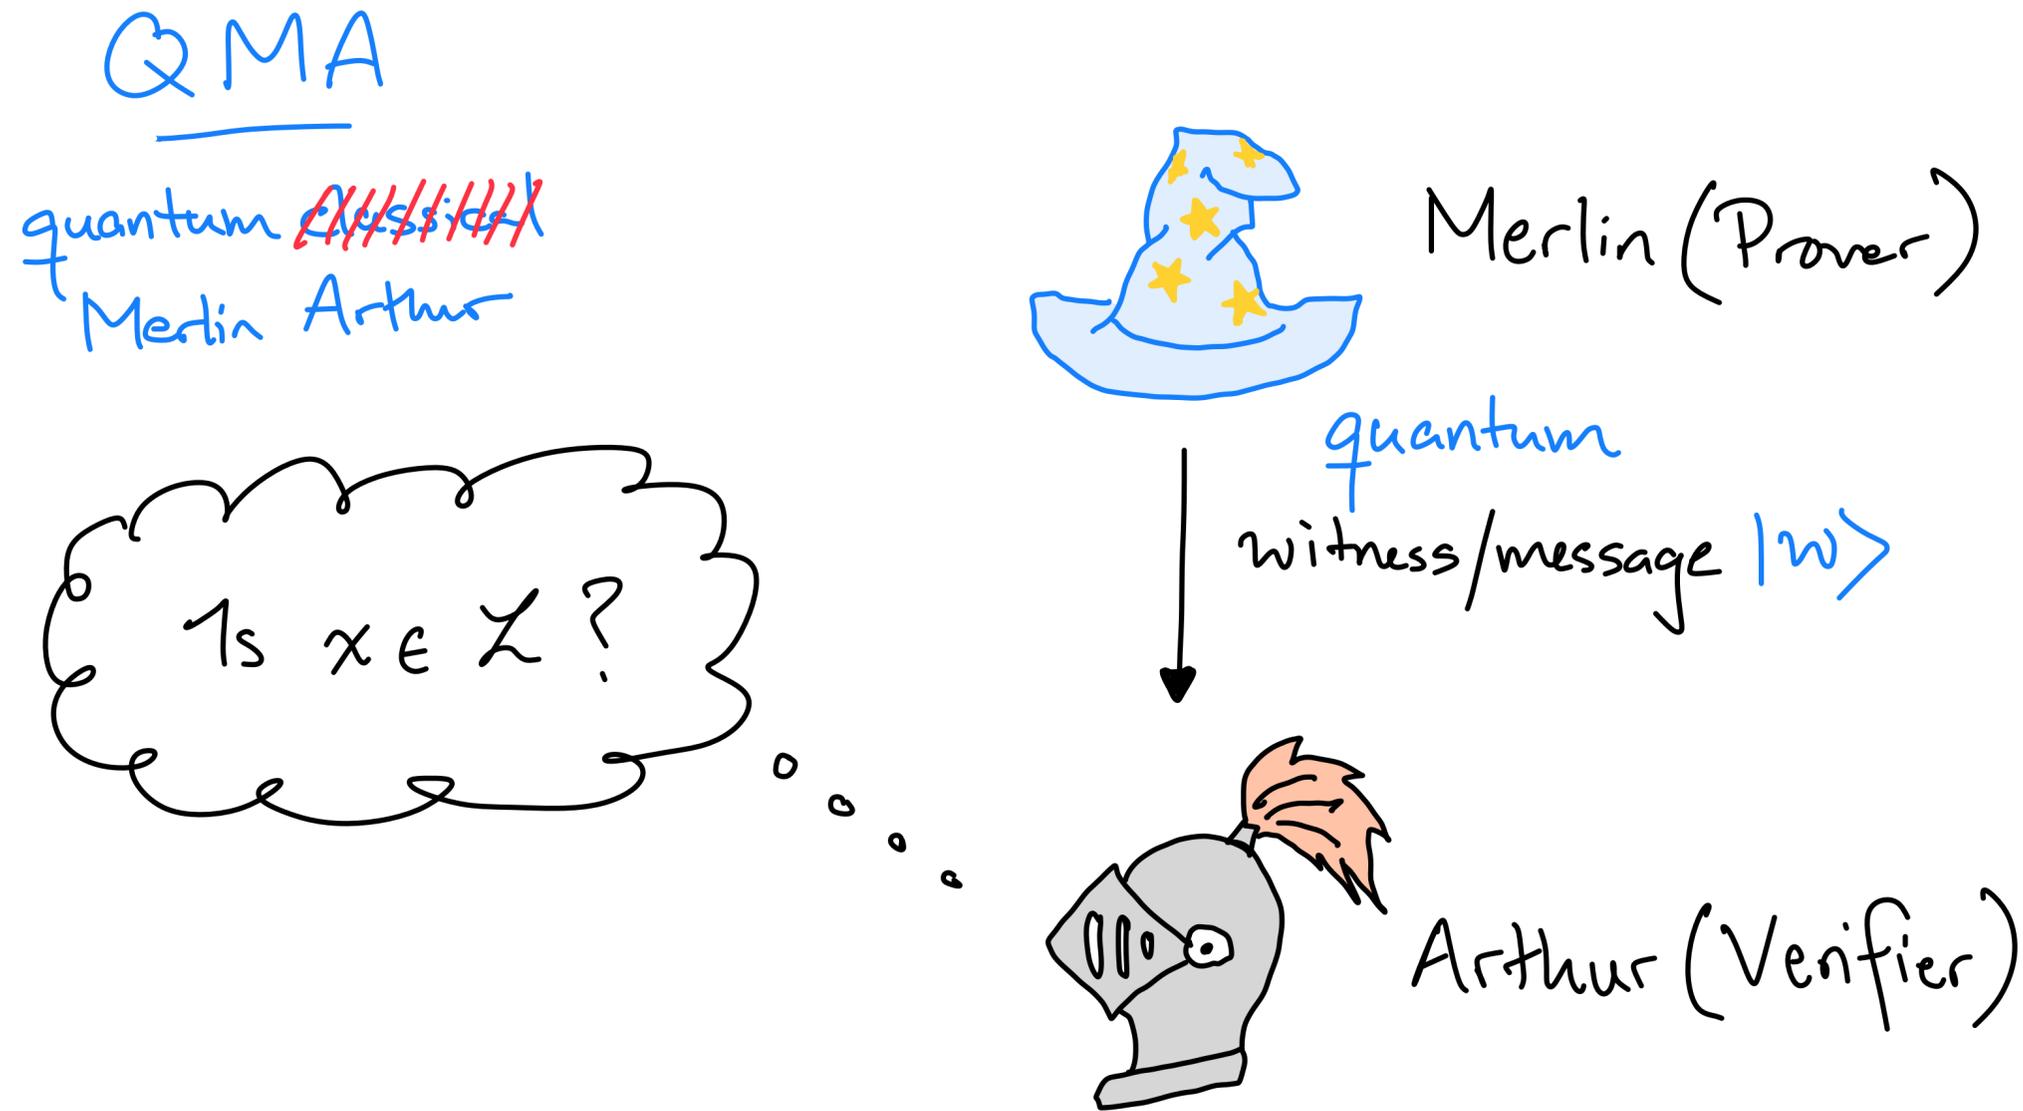
classical
witness/message w



Arthur (Verifier)

Arthur is a quantum polytime
computer and computes a fn
 $v(x, w)$

There are two appropriate quantum analogs to NP



Kitaev 2002:

Computing ground energy is QMA-complete.

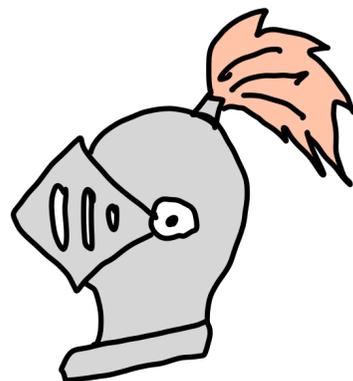
Arthur is a quantum polytime computer and computes a fn $V(x, w)$

There are two appropriate quantum analogs to NP

NP / MA



Classical
 w

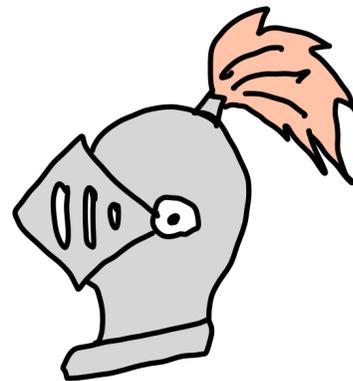


classical

QCMA



Classical
 w

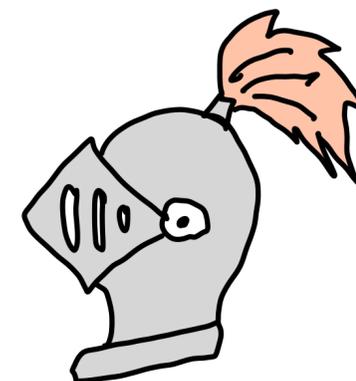


quantum

QMA



quantum
 $|w\rangle$



quantum

Why should you care?

Quantum states are hard to describe



Say Arthur wants to understand the ground state $|\psi\rangle$ of a local Hamiltonian system \mathbf{H} .

Computing the energy of $|\psi\rangle$ is **QMA**-complete.

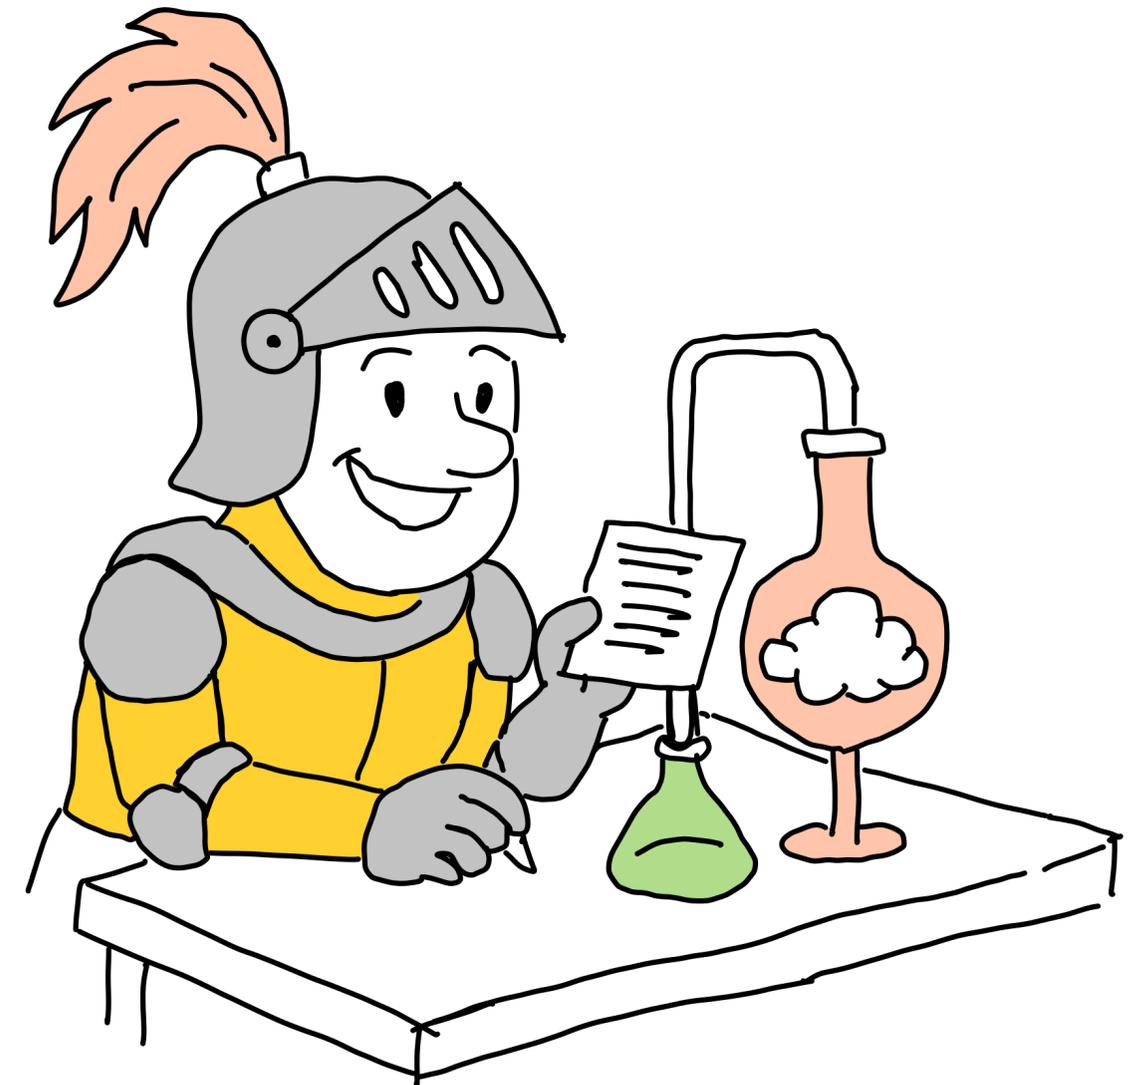
But assuming **BQP** \neq **QCMA**,

Arthur is not able to write down classical descriptions of $|\psi\rangle$ even if he has a copy of $|\psi\rangle$ in his lab.

Why care?

Assuming $\text{QCMA}=\text{QMA}$,

Merlin can write a classical message about $|\psi\rangle$ which Arthur can use to verify that \mathbf{H} has a low-energy ground state.



Why care?

However, if **QCMA** \neq **QMA**,

Merlin struggles to write a description of $|\psi\rangle$ that Arthur will be able to understand.

Meaning that $|\psi\rangle$ can be verified but not described.

The only hope for Arthur is if Merlin sends him a copy of $|\psi\rangle$ which we will be able to verify but not write down a classical description of.

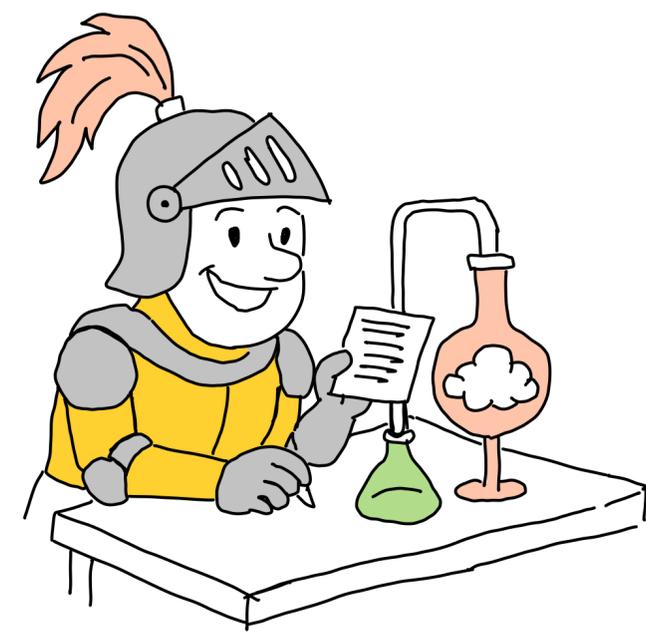


Resolving QMA v QCMA

Is deciding which of the two worlds we live in



OR



I.e, it also explains whether ground states can or cannot be efficiently described.

Evidence that QCMA \neq QMA

Bostanci, Haferkamp, *Nirkhe*, Zhandry (STOC 2026).

- Proving QCMA \neq QMA requires proving P \neq PSPACE (open since 1960's)
- Instead, we prove that the classes are not equal in a query complexity model
- Many interpretations of this result
 - QCMA \neq QMA for problems about *classical* black-box functions/oracles
 - QCMA \neq QMA when Arthur's verification algorithm runs in logarithmic-time
 - QCMA cannot even approximate ground energy (QMA-complete) unless the algorithm strongly depends on the structure $\mathbf{H} = \sum_i h_i$.



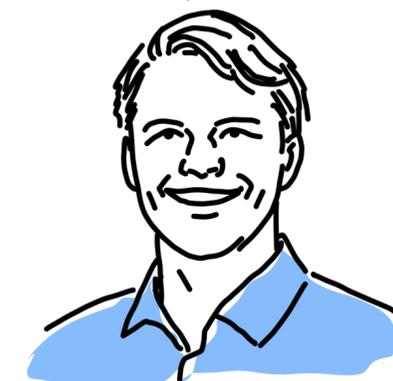
John Bostanci



Jonas Haferkamp



Chinmay Nirkhe



Mark Zhandry

The power of quantum witnesses

- What differentiates a quantum state $|\psi\rangle$ from a classical string w ?
- **No-cloning theorem.** There is no general transform cloning quantum states:

$$|\psi\rangle \not\rightarrow |\psi\rangle \otimes |\psi\rangle$$

- Whereas given a classical string w , we can make as many copies as we want!
- This is a way in which classical witnesses are more powerful than quantum ones!
- What we prove is that, paradoxically, this increase in power is too good to be true and therefore there exist quantum problems without classical witnesses!

The power of quantum witnesses

- What makes quantum and classical witness different?
- Quantum witnesses are use-once objects. Classical witnesses can be reused.

Given witness $|\psi\rangle$ and q. circuit V we can

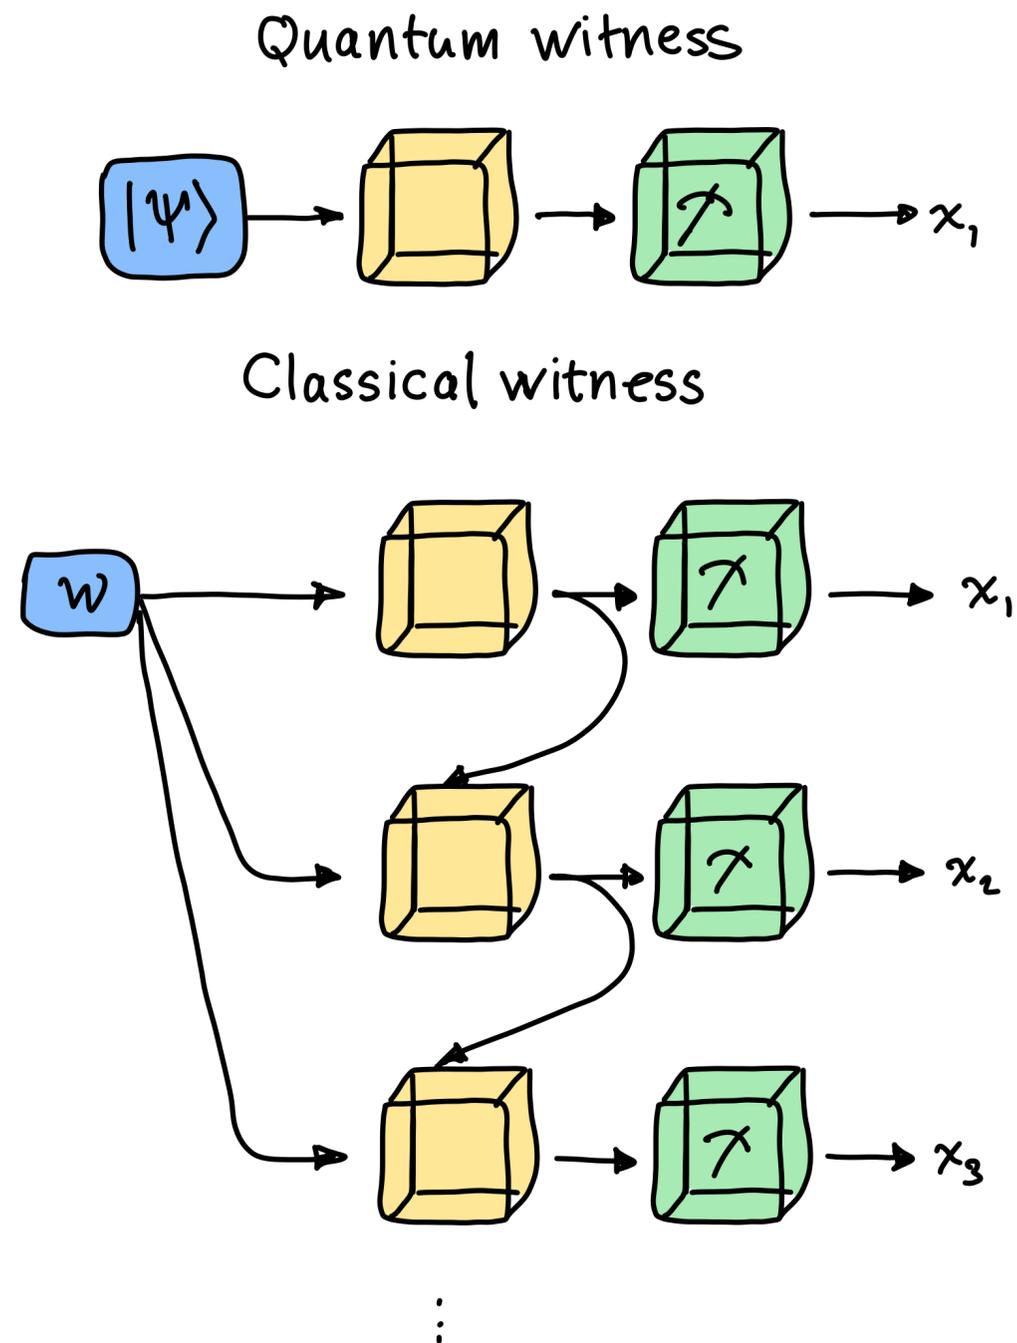
produce 1 sample from $p(x) = |\langle x|V|\psi\rangle|^2$

Because we can clone w
and can't clone $|\psi\rangle$.

Given witness w and q. circuit V we can

produce as many samples as we want from $p(x) = |\langle x|V|w\rangle|^2$

The power of quantum witnesses

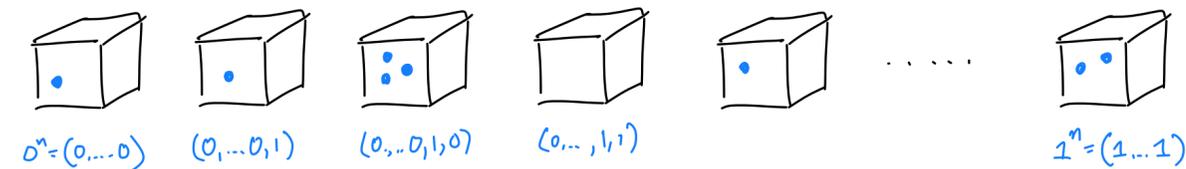


Proof sketch of QMA v QCMA query separation

Identify a distribution \mathcal{D} that is provably hard to produce $\exp(n)$ many samples from

Argue that if $\text{QCMA} = \text{QMA}$ (in query model) that we could construct such a sampler, a contradiction

Distribution is derived from the position vs momentum uncertainty principle for bosons

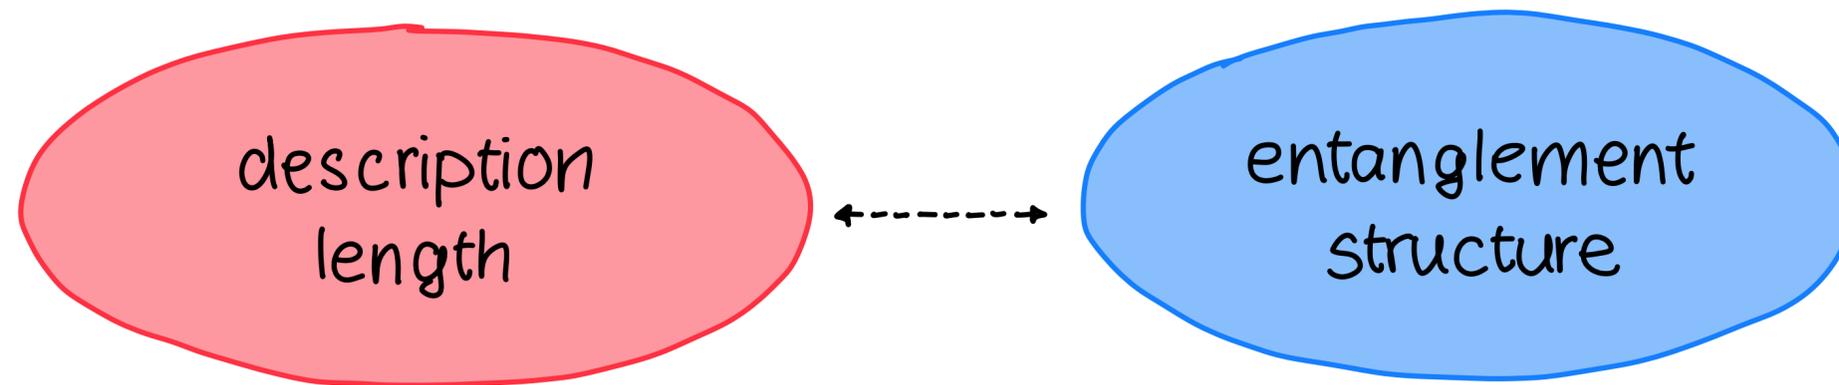


Part III: Robust entanglement

Ground energy approximations

- What have we seen so far?
 - Computing ground energy is QMA-complete due to Kitaev
 - And from $\text{QMA} \neq \text{QCMA}$ query proof, good evidence that ground states cannot be efficiently described classically
- **How hard is it to approximately compute the ground energy?**
 - How hard is it to describe approximations of the ground state?
 - Kitaev's lower bounds only extend to show hardness for describe approximations that are $O(1/n^3)$ -close

The relationship between description complexity and entanglement



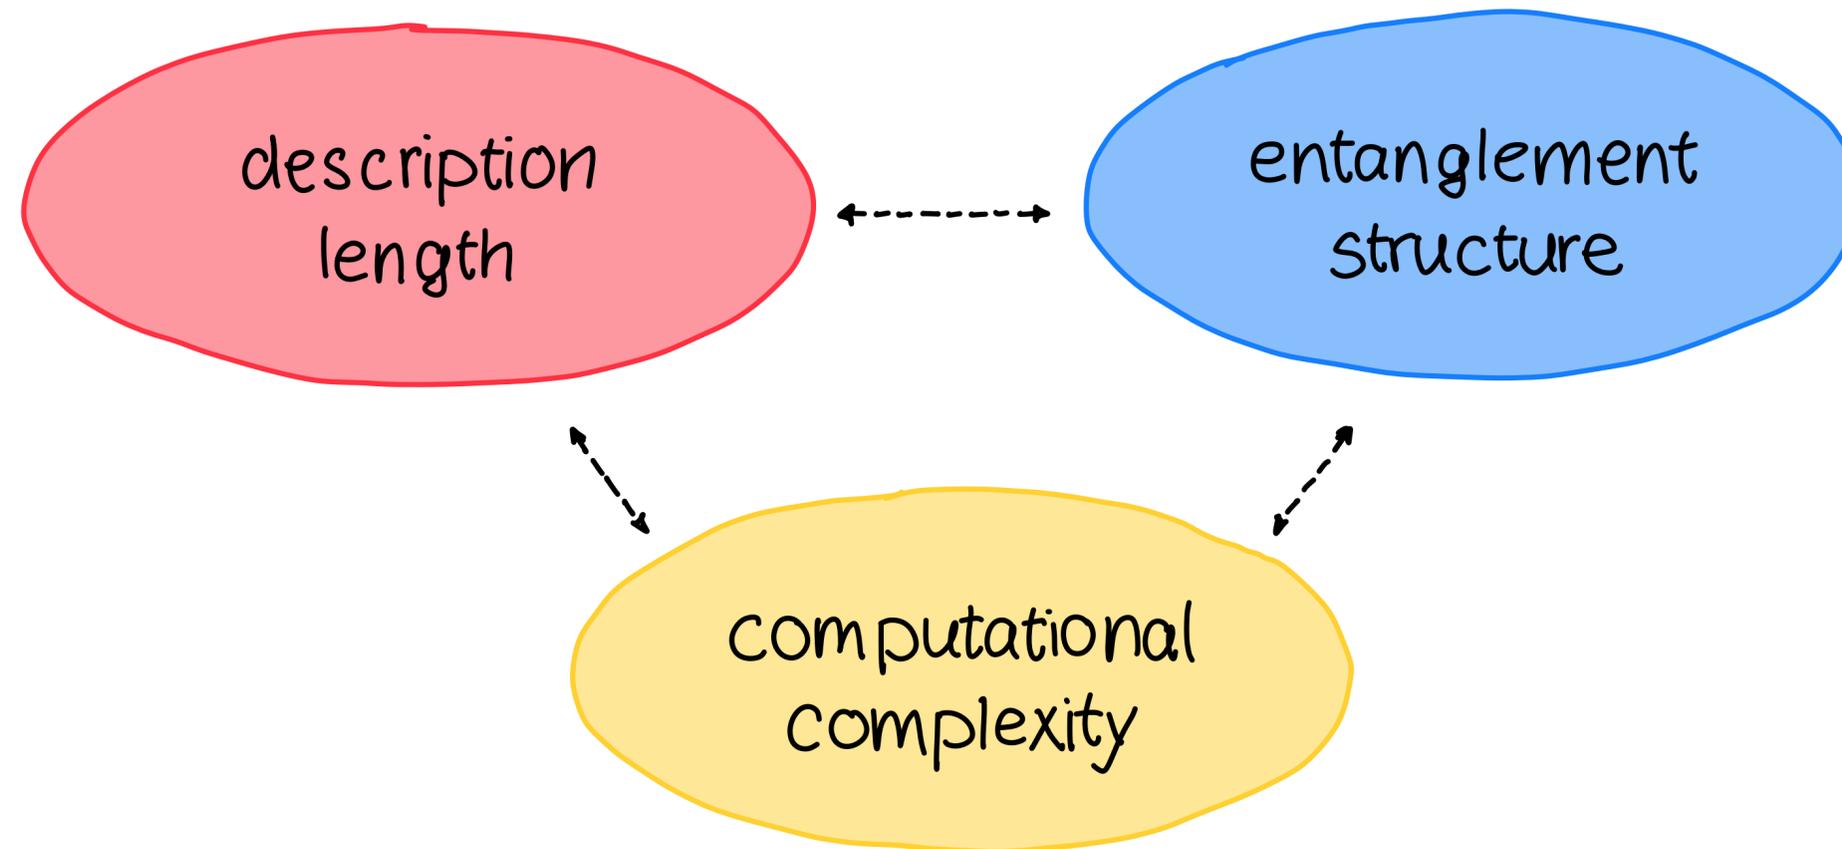
quantum entanglement

doesn't suffice to describe each individual particle

the longer the description \implies the more "complex" the entanglement

description length is a "measure" of entanglement complexity

The relationship between description complexity and entanglement

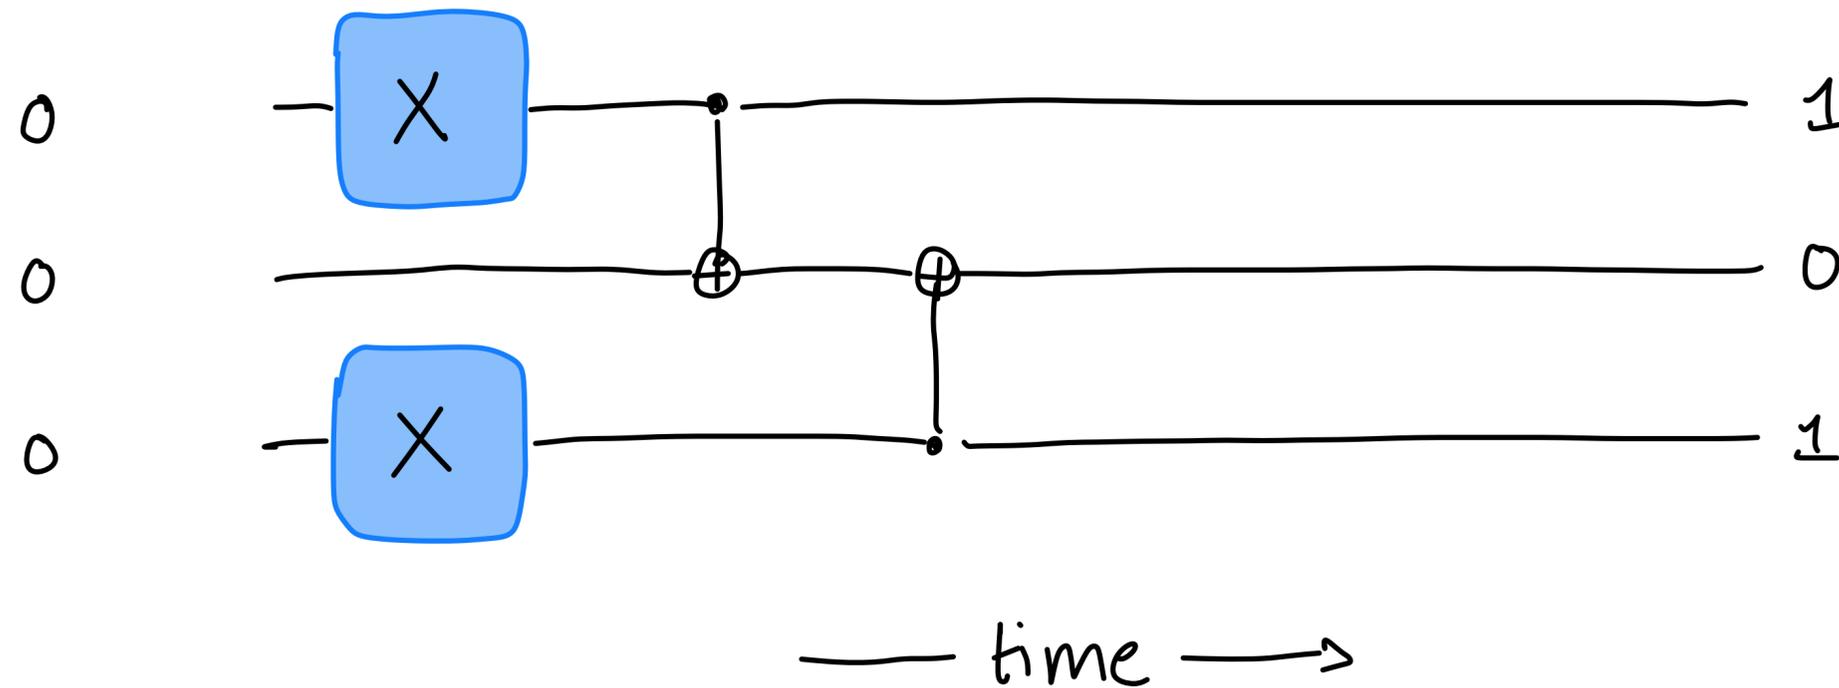


How do description length & computational complexity intersect?

The Circuit Model

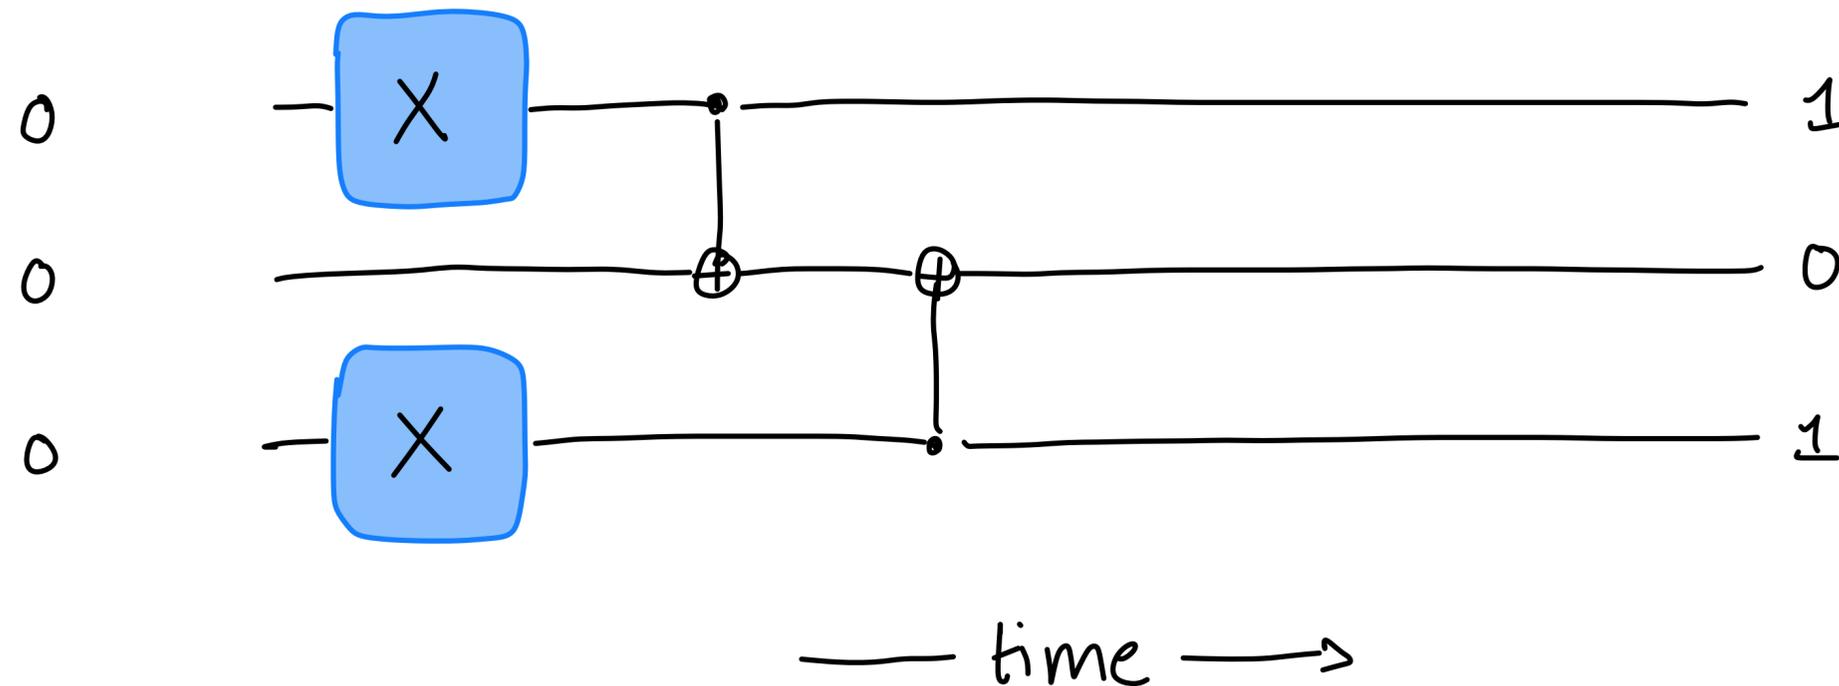
The circuit model

Computation can be described in the *circuit* model



The circuit model

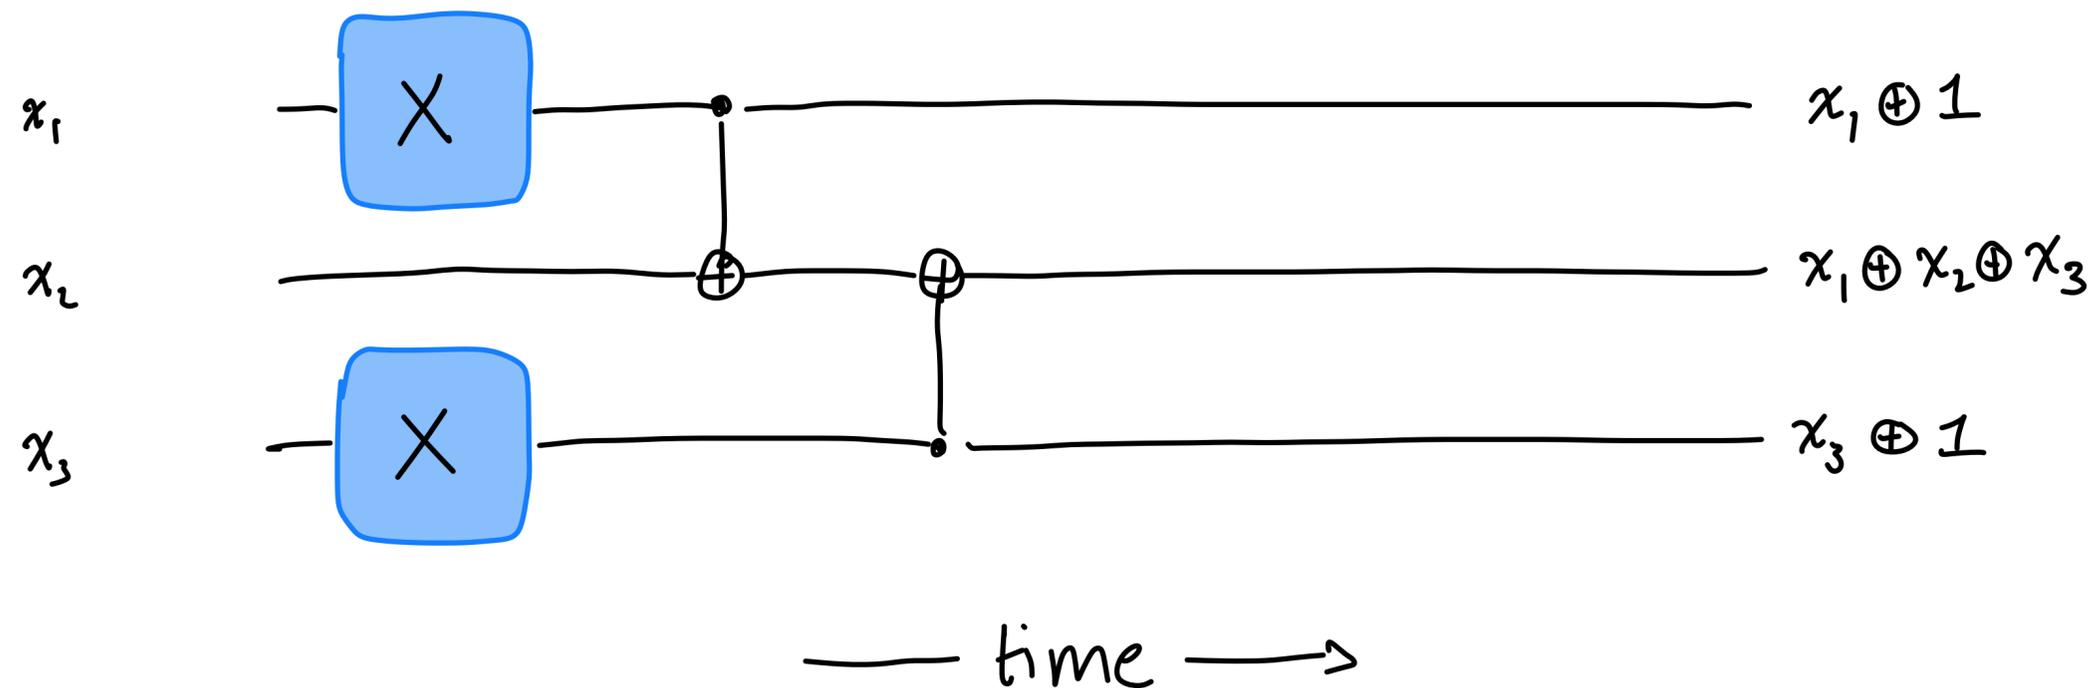
Computation can be described in the *circuit* model



Computation can be viewed as an *input/output* function

The circuit model

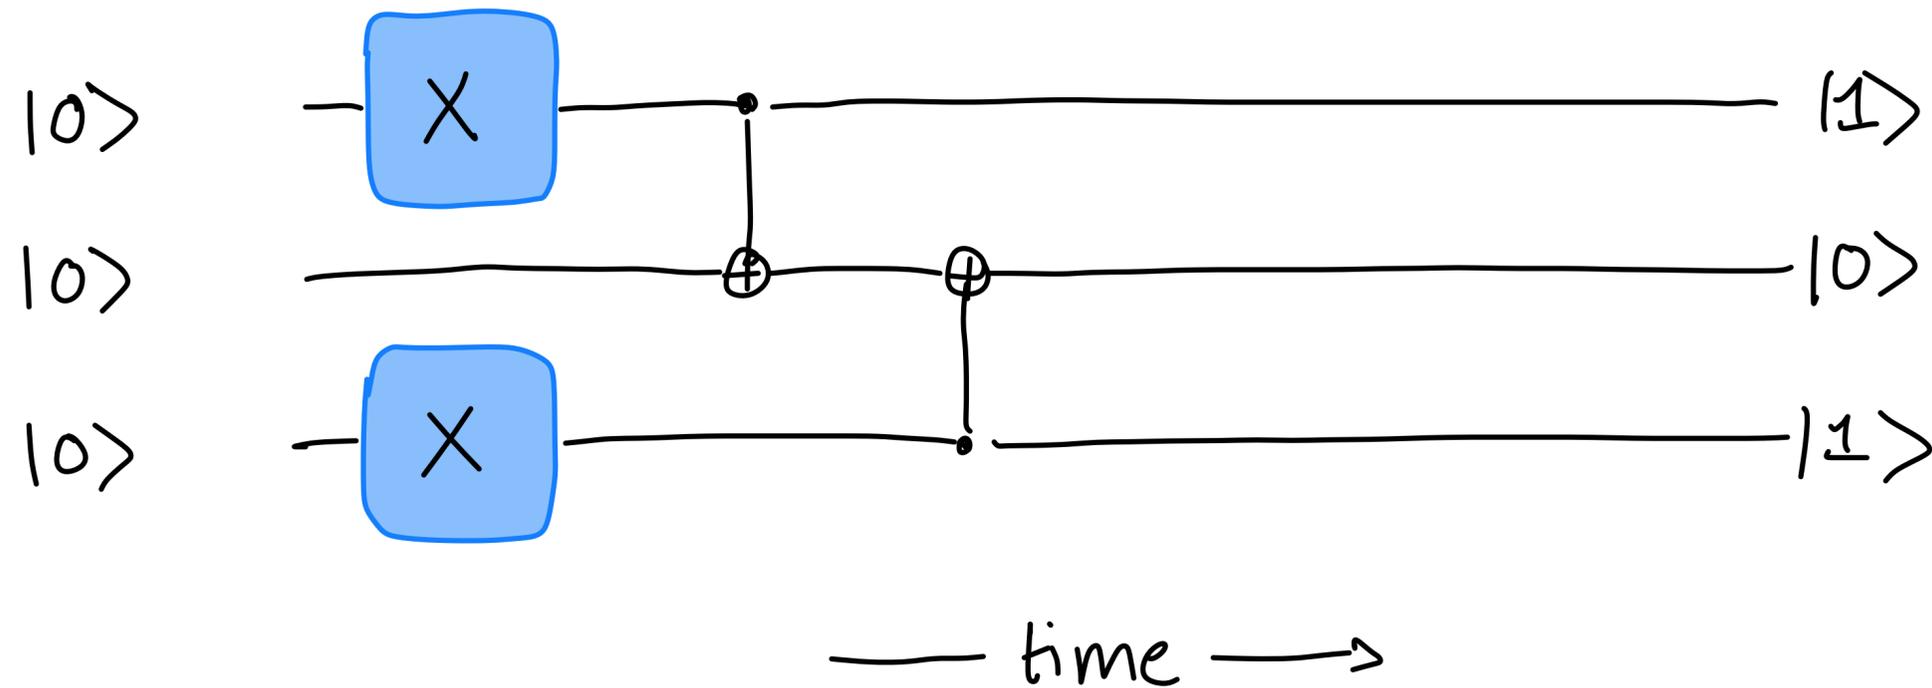
Computation can be described in the *circuit* model



Computation can be viewed as an *input/output* function

The circuit model

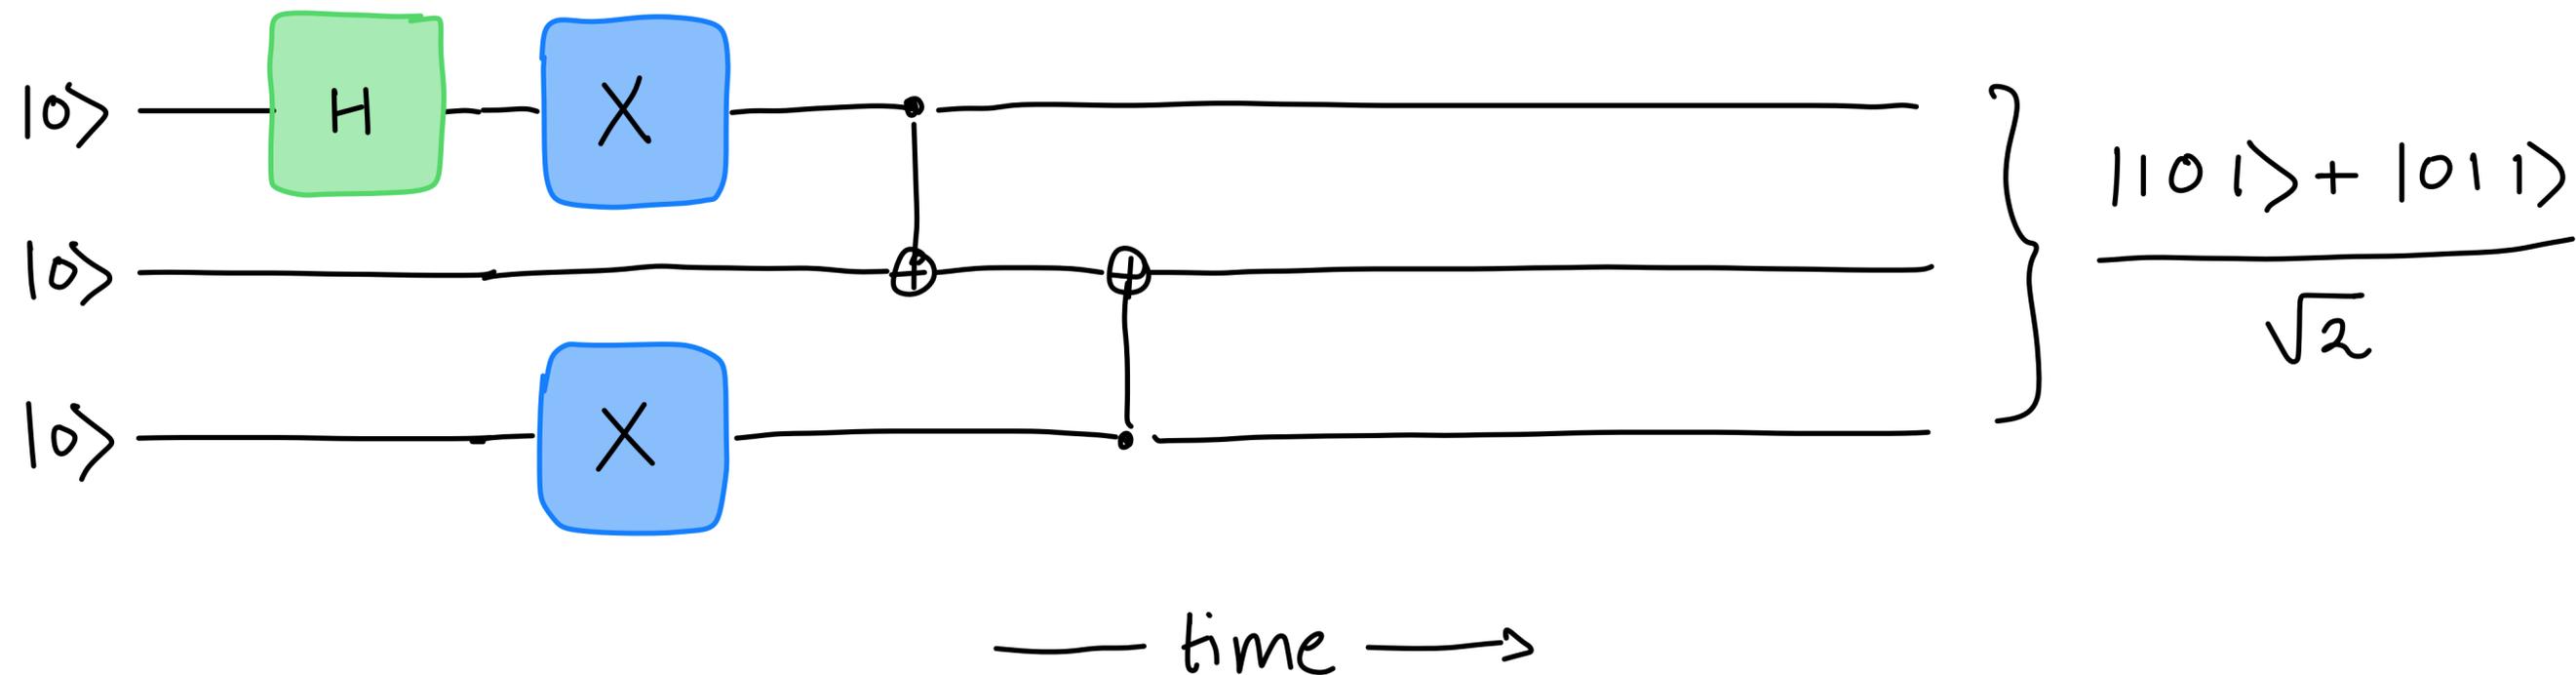
Likewise, quantum computation can be described in the *circuit* model



The circuit model

$$\boxed{\text{H}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

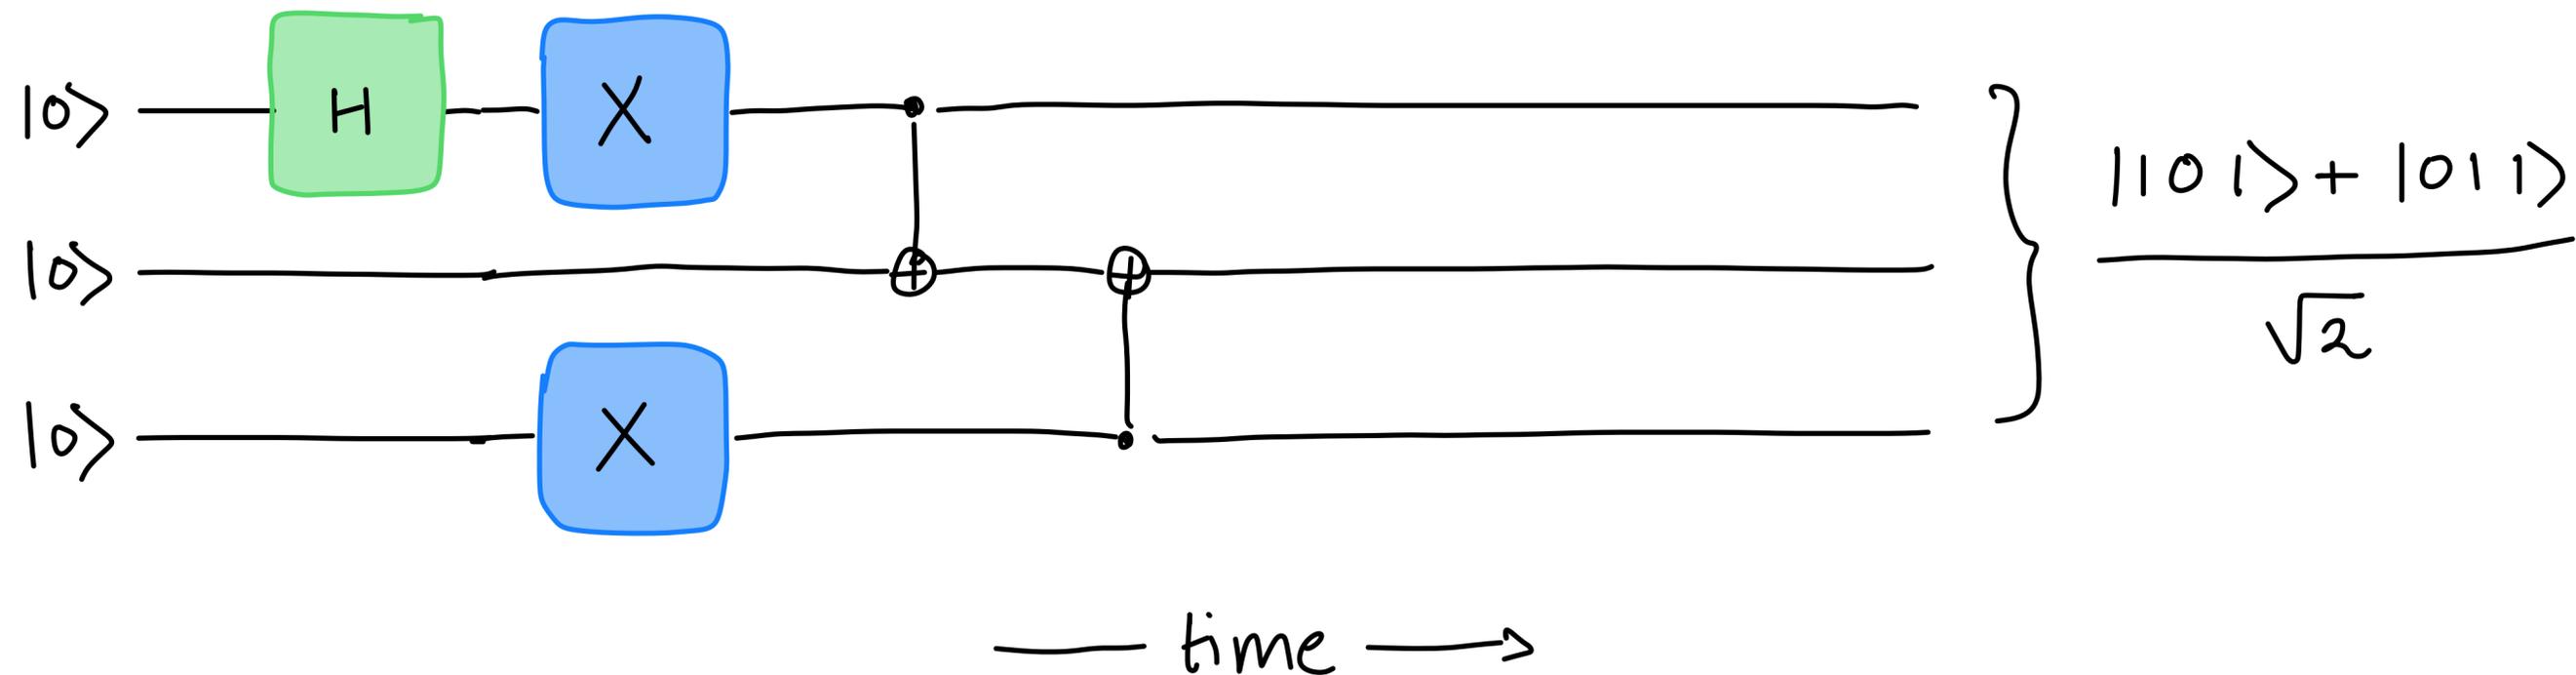
Likewise, quantum computation can be described in the *circuit* model



The circuit model

$$\boxed{\text{H}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

Likewise, quantum computation can be described in the *circuit* model

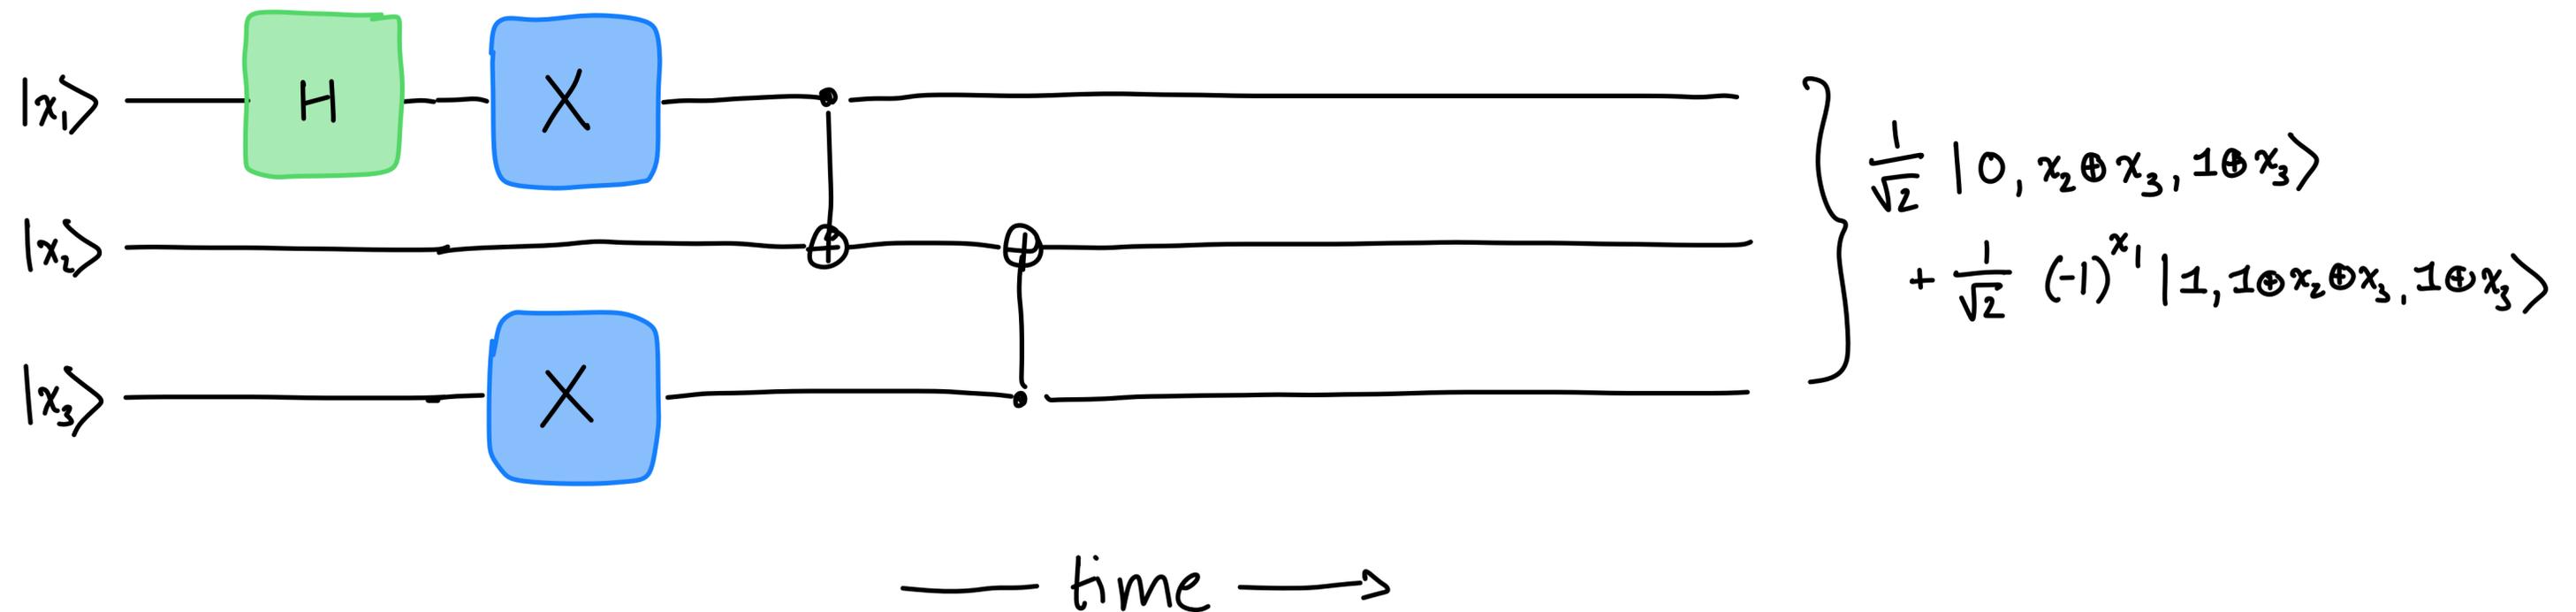


quantum computation can also be viewed as an *input/output* function

The circuit model

$$\boxed{\text{---} \text{H} \text{---}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

Likewise, quantum computation can be described in the *circuit* model

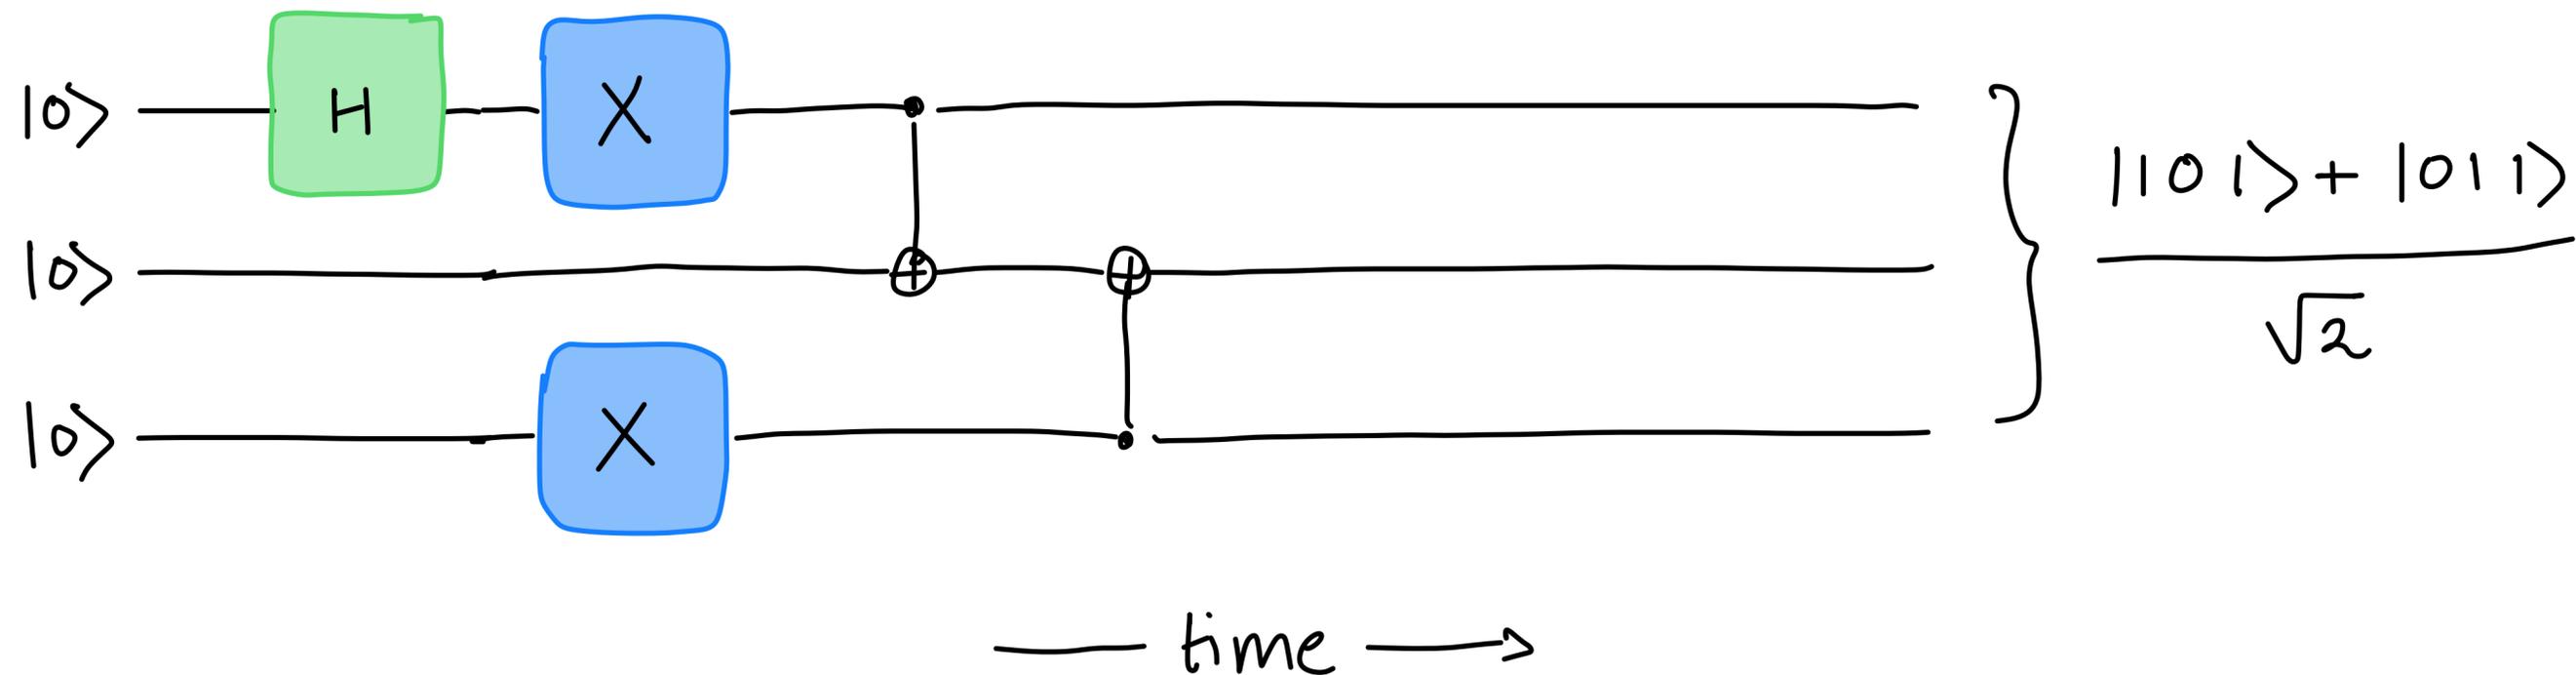


quantum computation can also be viewed as an *input/output* function

The circuit model

$$\boxed{\text{H}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

Likewise, quantum computation can be described in the *circuit* model

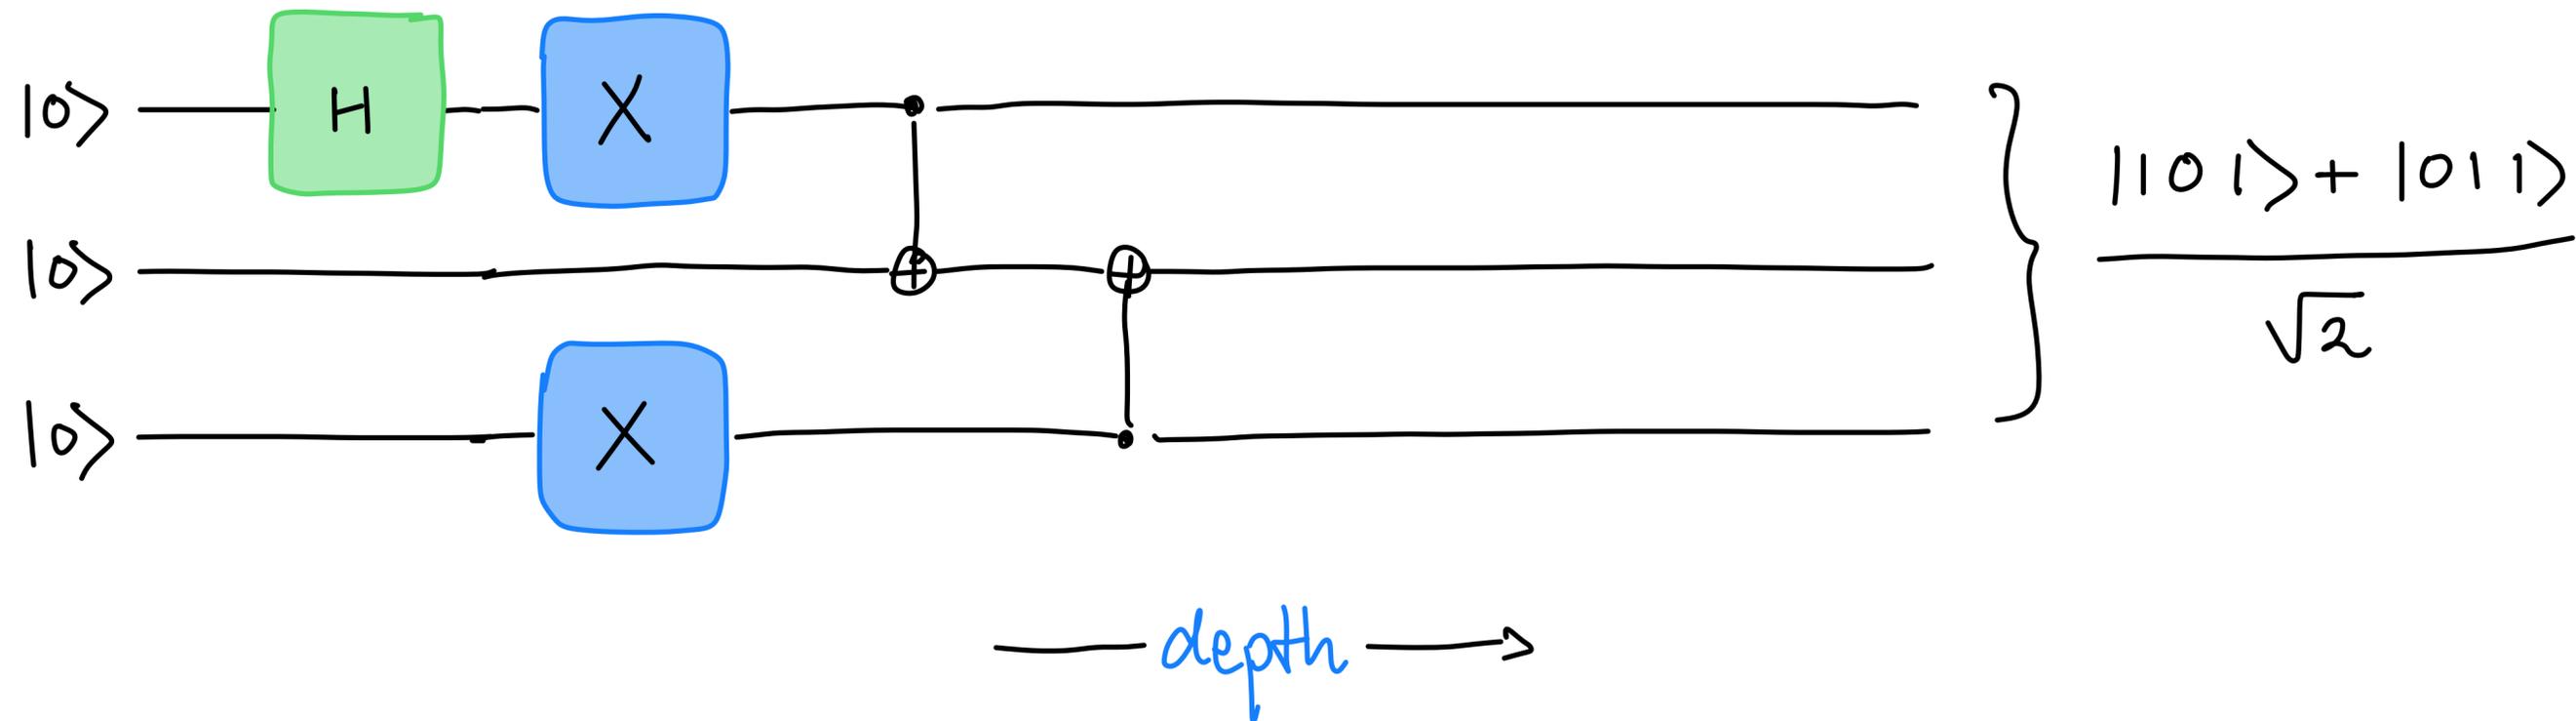


quantum computation also gives us a measure of complexity for states

The circuit model

$$\boxed{\text{---} \text{H} \text{---}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

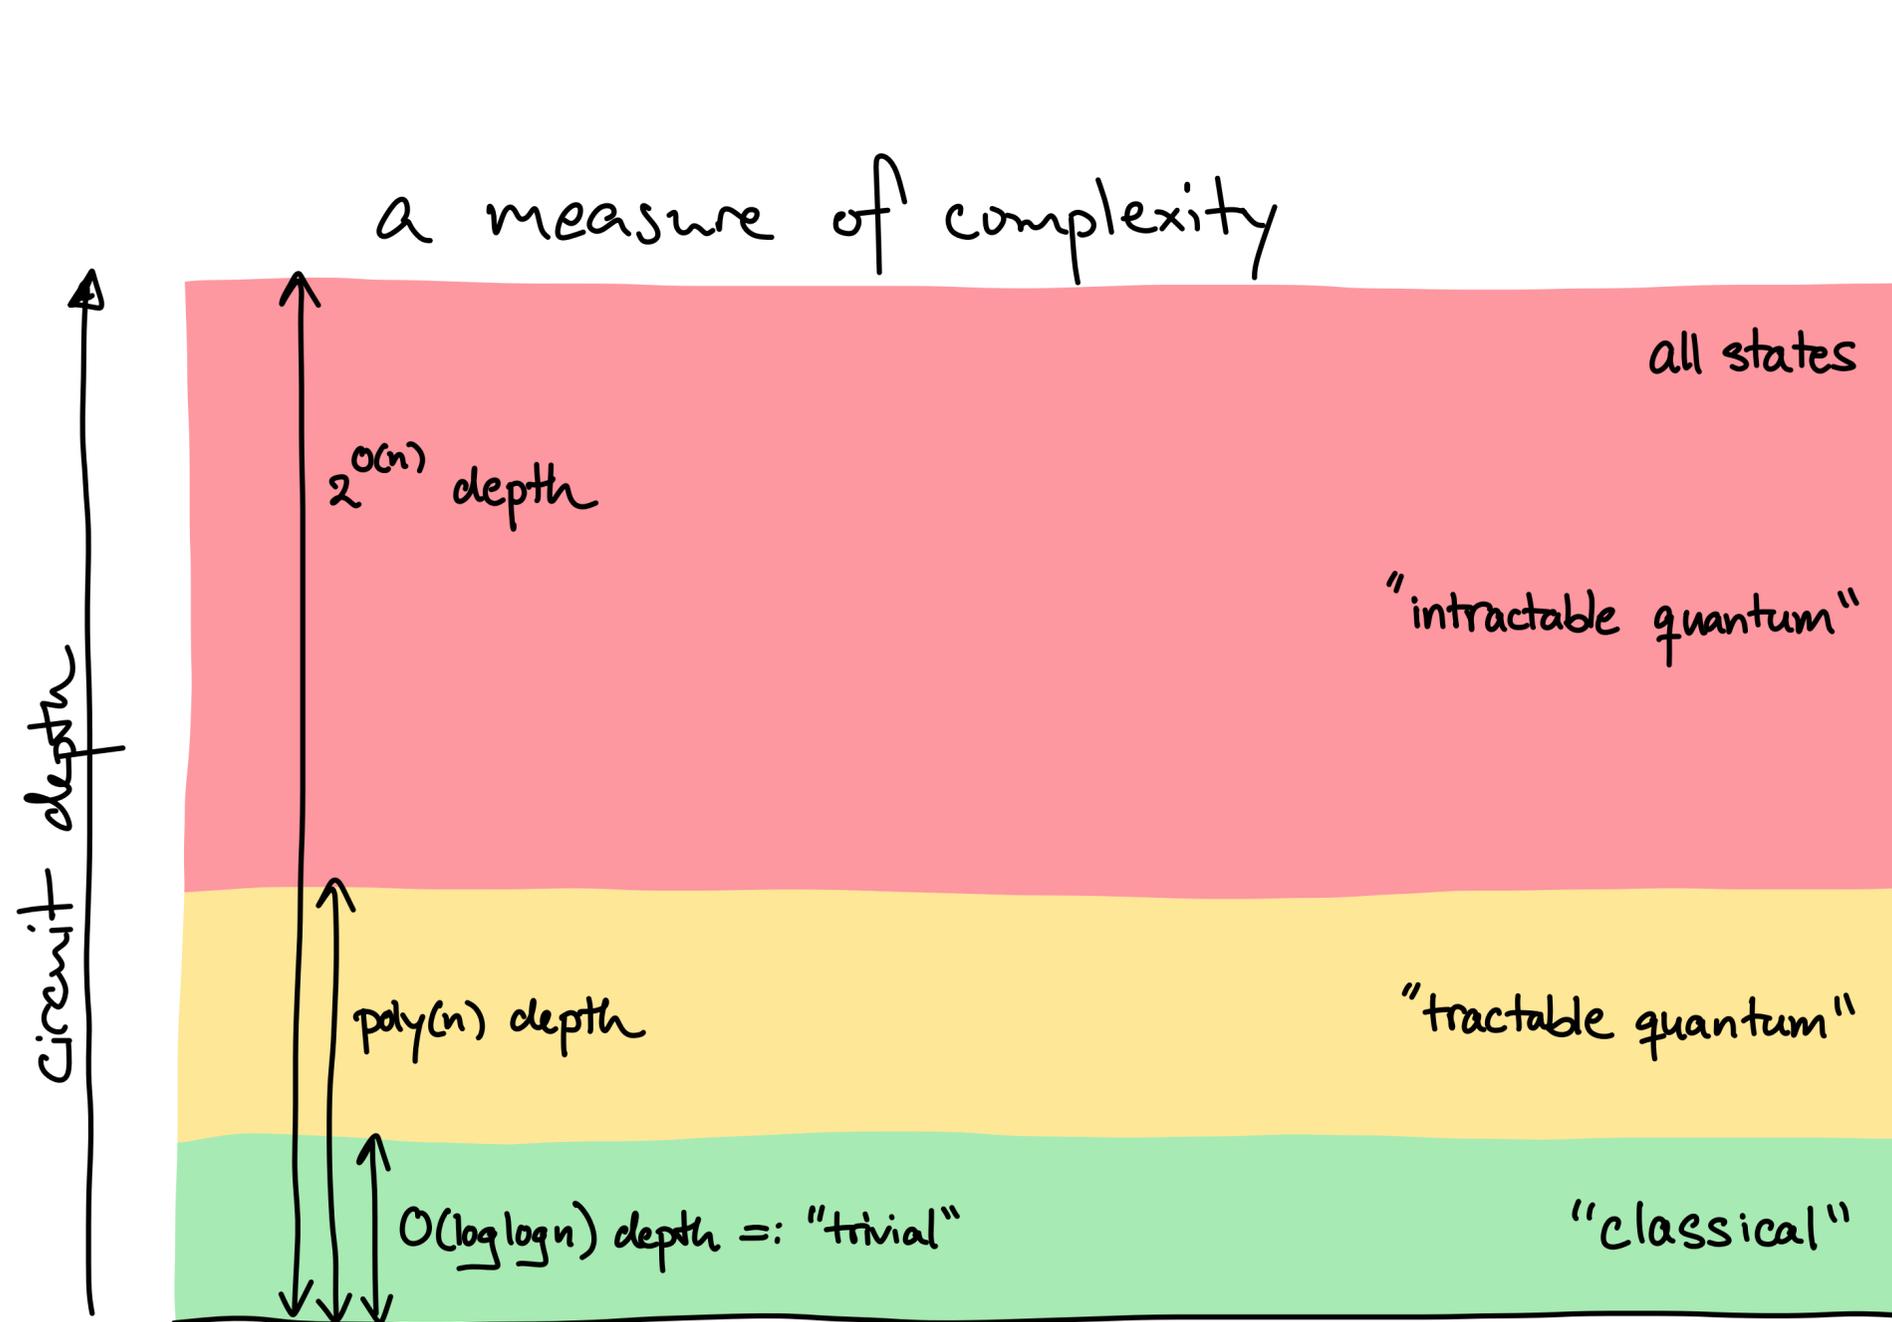
Likewise, quantum computation can be described in the *circuit* model



quantum computation also gives us a measure of complexity for states

$$\text{depth}(|\psi\rangle) = \min \text{ depth over all circuits generating } |\psi\rangle \text{ from } |0^n\rangle$$

The quantum (computational) yardstick

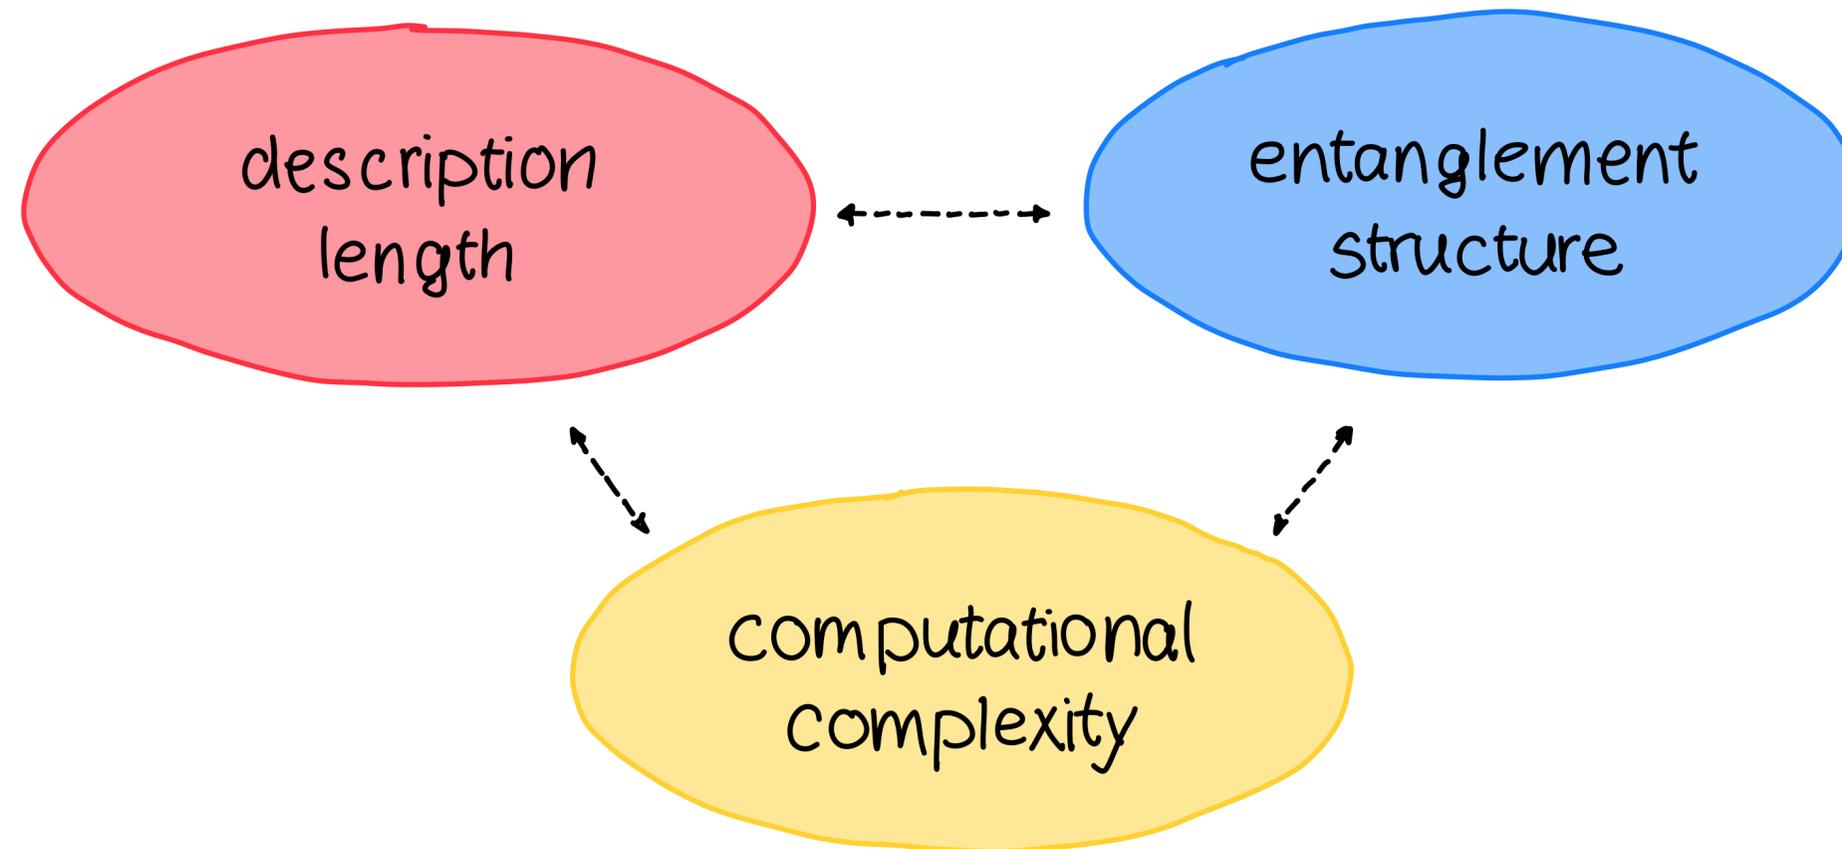


$\text{depth}(|\psi\rangle)$ = minimum depth of circuit with output $|\psi\rangle$

efficient **quantum** algorithm for calculating the energy of $|\psi\rangle$ for any local Hamiltonian H

efficient classical algorithm for calculating the energy of $|\psi\rangle$ for any local Hamiltonian H

The relationship between description complexity and entanglement

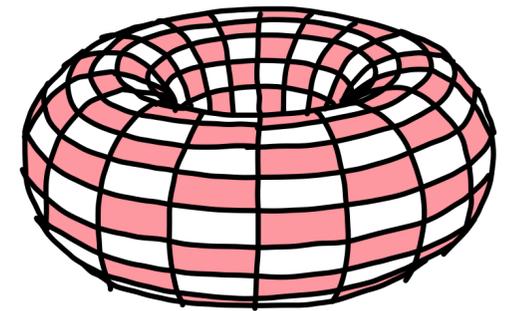


Do all local Hamiltonian systems have low depth circuit approximations for the ground state?

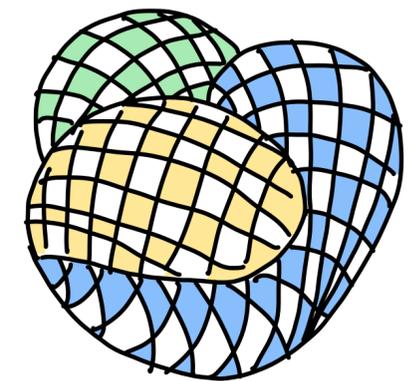
Inapproximable local Hamiltonian systems

NLTS theorem: Anshu, Breuckmann, *Nirkhe* (STOC 2023).

- Most physical systems (ex. Toric code, 2D Ising models, AKLT models, etc.) have approximations that can be described by low-depth circuits
- **Theorem:** But we can *engineer* systems for which all the ϵ – approximations of the ground state require circuit depth $\Omega(\log n)$ for n -qubit Hamiltonians
- Robust entanglement can (potentially) exist at “warm” temperatures!
- This is an *engineering* feat. — construction built from mathematical tools like error-correction and expander graphs
- We believe stronger lower bounds are possible — some evidence of this in restricted models of computation



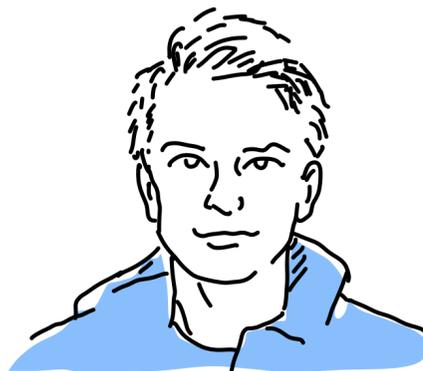
Hamiltonian on a physical system on a torus



The Hamiltonian we engineer lives on a high-dimensional manifold with an expander-like topology. This is a cartoon of it folding in on itself.



Anurag Anshu



Niko Breuckmann

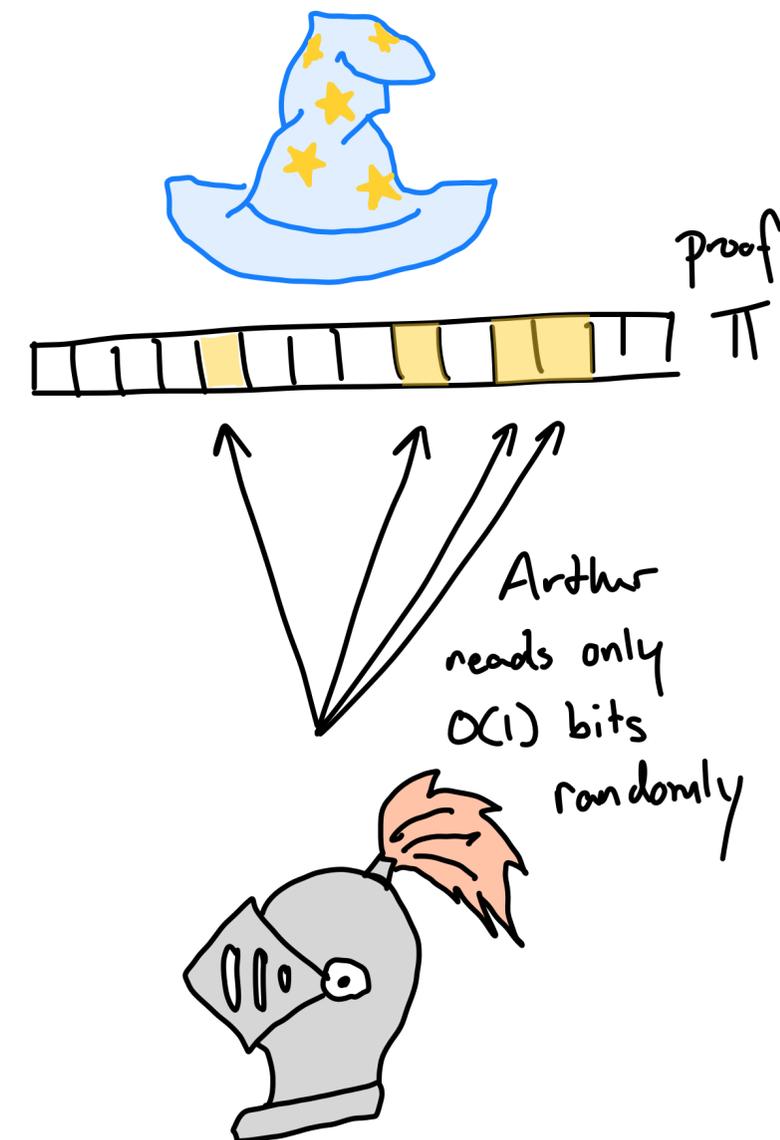


Chinmay Nirkhe

Part IV: The open questions

The Quantum PCP conjecture

- In 1992, Arora-Lund-Motwani-Sudan-Szegedy, had a breakthrough in classical computer science.
 - They were studying the complexity of approximating NP problems like 3-SAT due to breakthroughs in a field called *interactive proofs*
 - They proved that approximating the number of 3-SAT equations satisfiable to $< 12.5\%$ is as hard as solving the problem exactly
 - In other words, the approximation problem is NP-complete
- This was a major breakthrough — proving why no good approximation algorithms had been found for these fundamental problems
- This result is called the **probabilistically checkable proofs theorem** due to an alternate interpretation as Arthur verifying a classical proof sent by Merlin approximately

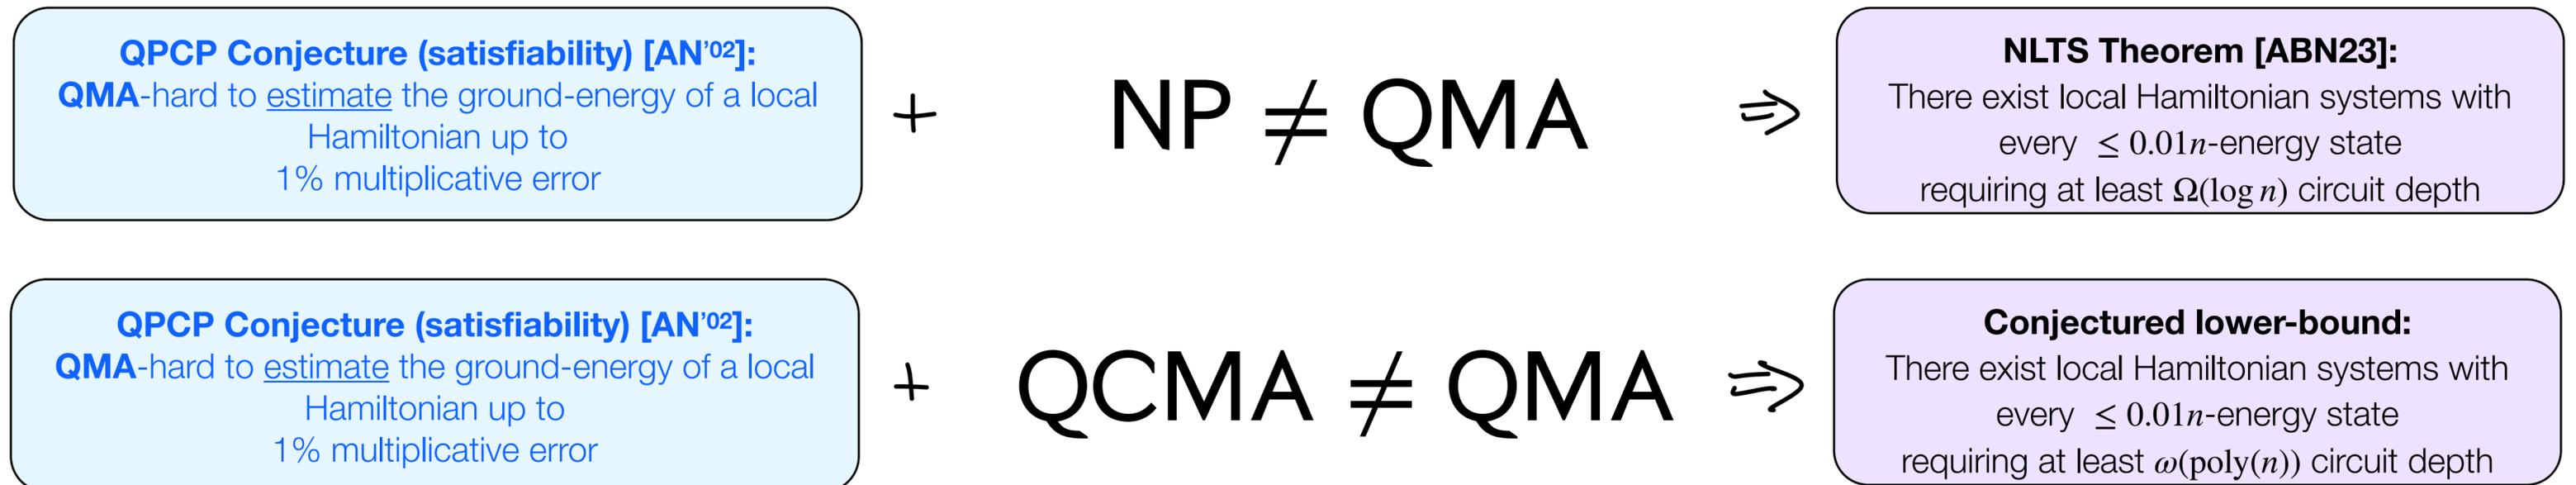


The Quantum PCP conjecture



Dorit Aharonov

- The question of whether a quantum analog of the PCP theorem is the most outstanding open question in quantum complexity theory
- First posed by Aharonov in 2002, many alternate formulations, consequences, variants have been studied (including many by myself)



Complexity theory → Cryptography & Quantum Engineering

- Complexity theory gives us evidence of complex entanglement
- Can we engineer that complex entanglement to *engineer* interesting states?
- **Quantum cryptography**
 - Use entanglement to create new cryptographic primitives
 - Key desired properties: unforgeability, unclonability, use-one, etc.
- **Quantum engineering**
 - Design robust entanglement for better quantum computations, reduce communication overhead, fault tolerance
 - Understand the nature of entanglement by explore what can be engineered



Last thing ... a shameless plug for UW

- If you are interested in researching more about quantum in general ... be it from a computational, physics, chemistry angle, the University of Washington is a great place to consider for graduate school!
- Multiple programs and a large investment by the region in quantum computation/technologies
- cs.washington.edu
theory.cs.washington.edu | quantum.cs.washington.edu



W