Average-Case Quantum Advantage for Shallow Circuits

François Le Gall
Kyoto University

ArXiv: 1810.12792

Quantum Information and String Theory 2019
14 June 2019

Establishing Quantum "Supremacy"

it's OK if the task is not really useful

in the circuit model

Find some computational task that can be solved easily in the quantum setting, even with current (or near-future) quantum technology, but are hard to solve in the classical setting

We

Can we establish quantum supremacy without relying on any conjecture or assumption?

- ✓ quantum key distribution
- ✓ quantum games (Bell inequalities)
- ✓ quantum fingerprinting (and many tasks involving quantum communication)

Many good candidates:

Boson sampling, instantaneous quantum polynomial-time computation, random circuit sampling,…

In all these results, the proof of the classical hardness relies on conjectures (e.g., anti-concentration conjecture) or complexity-theoretic assumptions (e.g., generalization of P ≠ NP)

Establishing Quantum "Supremacy"

it's OK if the task is not really useful

in the circuit model

Find some computational task that can be solved easily in the quantum setting, even with current (or near-future) quantum technology, but are hard to solve in the classical setting

Can we establish quantum supremacy without relying on any conjecture or assumption?

Theorem ([Bravyi, Gosset, König 17])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it

Remark: similar results have been obtained independently by many other researchers
[Bravyi, Gosset, König 18], [Bene Watts, Kothari, Schaeffer, Tal 19], [Coudron, Stark, Vidick 18]
(comparison given in later slides)

Our result:

There exists a computational problem such that:

average−case classical hardness

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) no shallow classical circuit can solve it on a non-negligible fraction of inputs.

Graph States

Definition (Graph State)

The graph state corresponding to a graph G is the state obtained by the following process:

1. Prepare one qubit in state $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ for each node of G
2. Apply a controlled-Z operation on the qubits corresponding to each edge of G

$$CZ = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

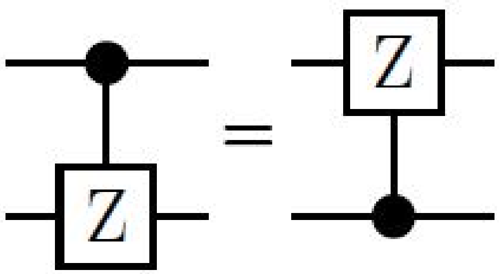

Example:

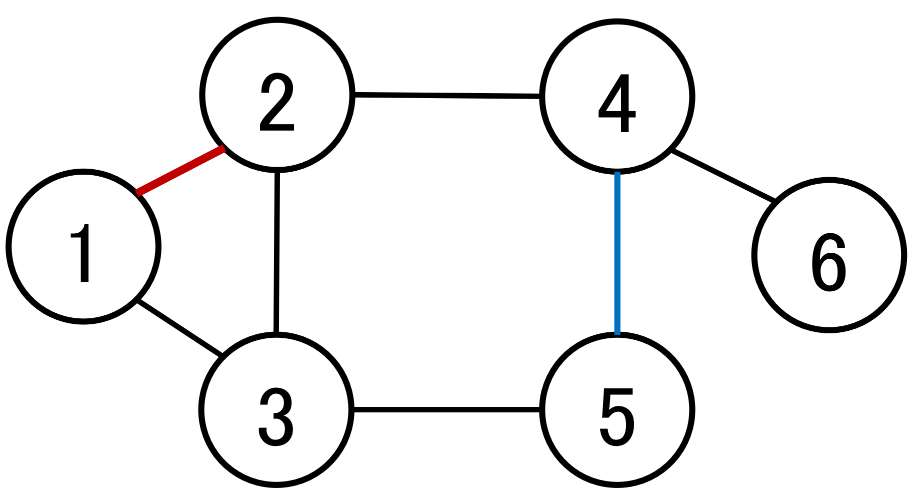

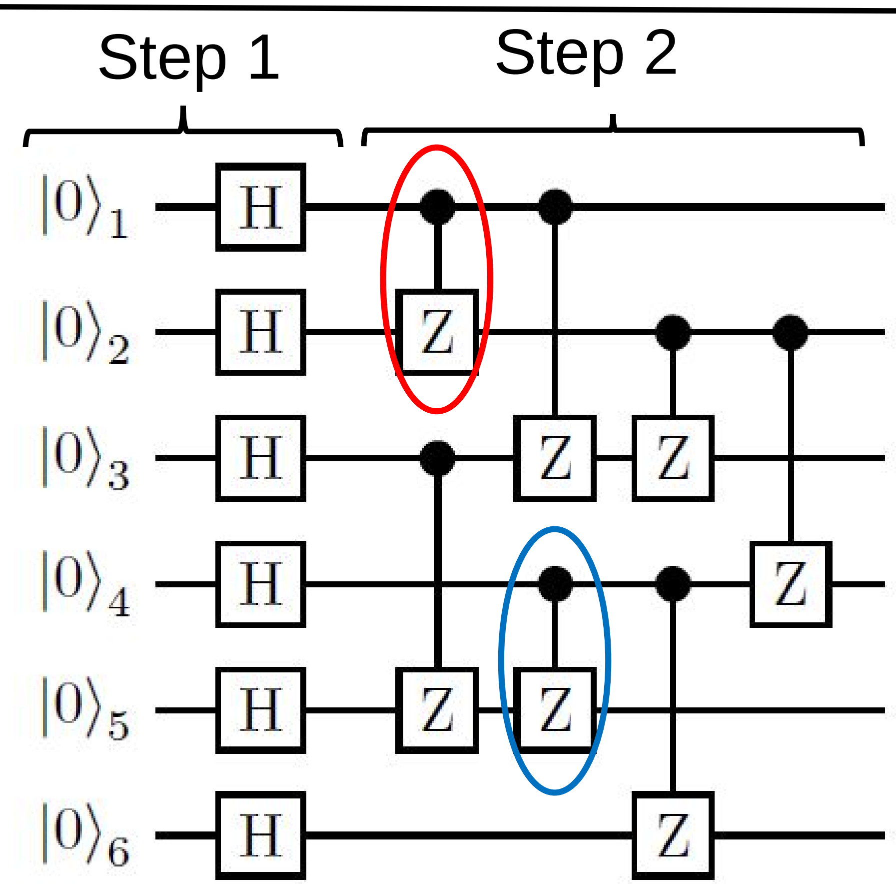

Key Prior Work [Barrett, Caves, Eastin, Elliot, Pironio 07]

multiple of 3

Consider a ring of size n (seen as a triangle)

Each "corner" gets a bit as input

Each node will output one bit

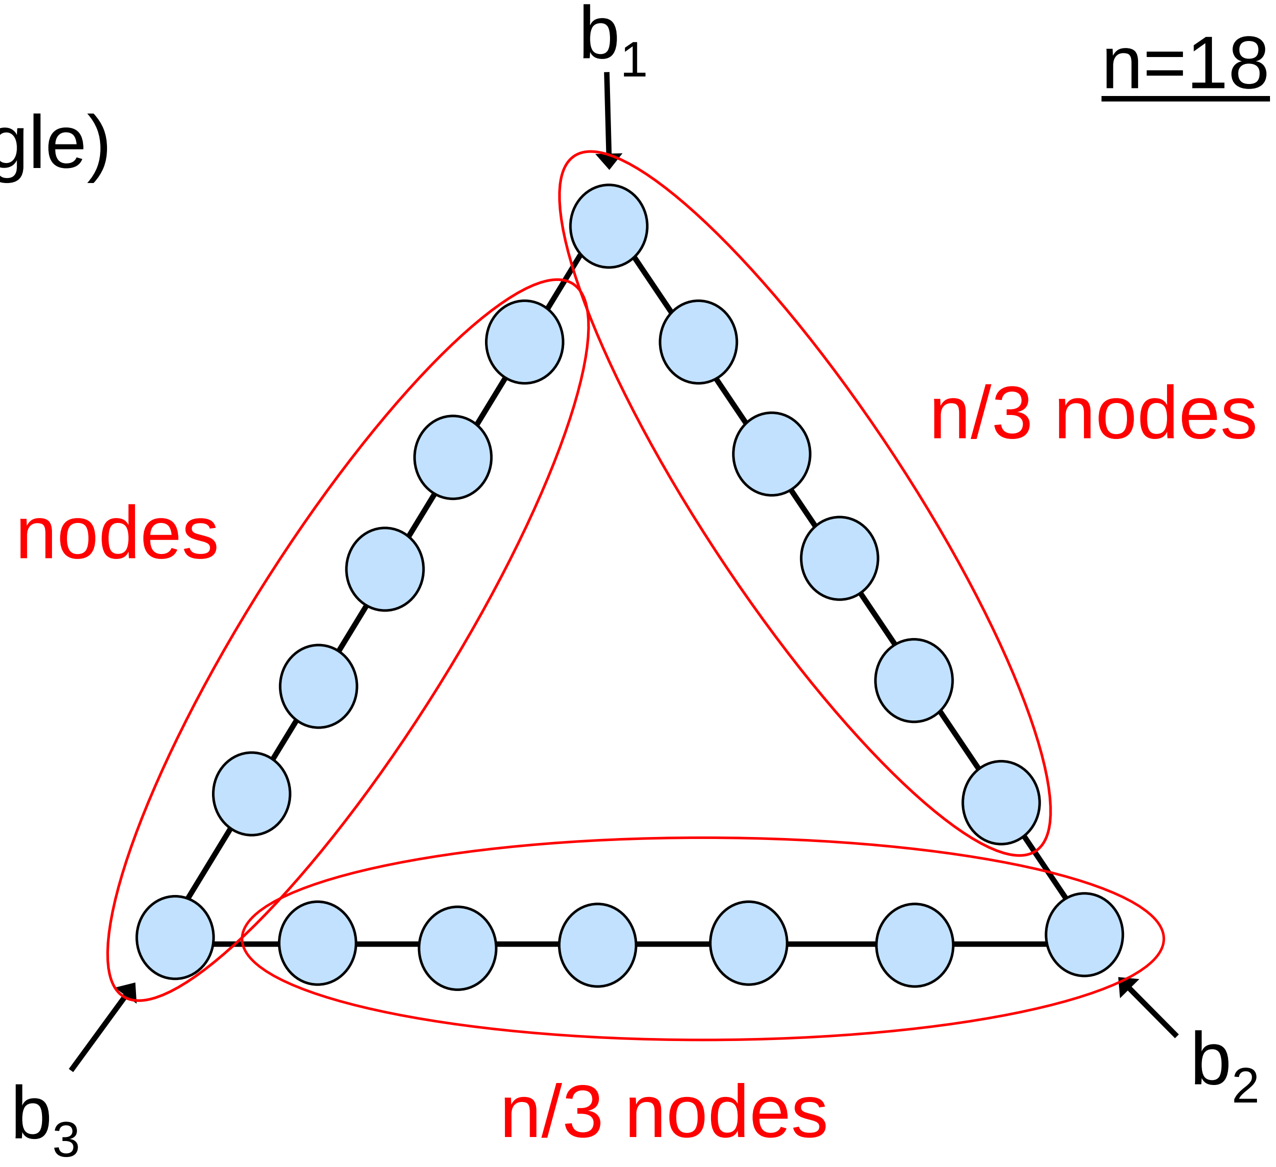

Key Prior Work [Barrett, Caves, Eastin, Elliot, Pironio 07]

multiple of 3

Consider a ring of size n (seen as a triangle)

Each "corner" gets a bit as input

Each node will output one bit

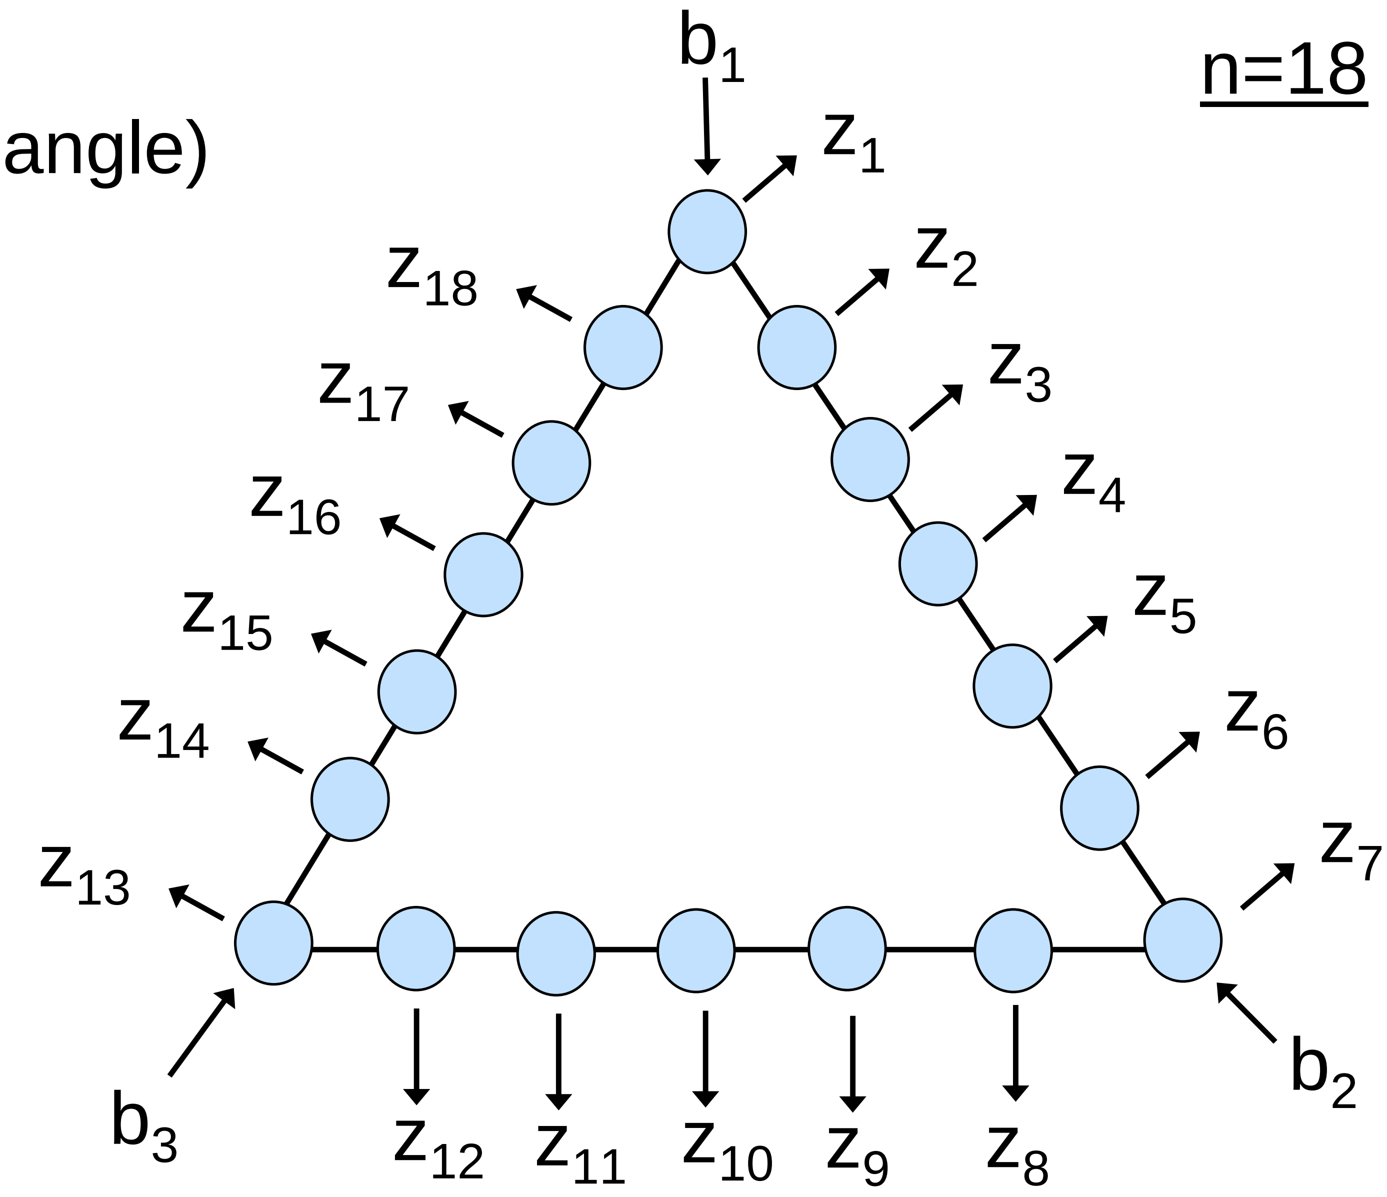

PROCESS

1. The nodes prepare the graph state corresponding to the whole triangle (each node only needs to communicate with its two nearest neighbors)
2. Each non-corner node measures its qubit in the X basis and then outputs the bit corresponding to the measurement outcome (no communication)
3. Each corner node measures its qubit in the X basis if its input bit is 0, or measures it in the Y basis if its input bit is 1, and then outputs the bit corresponding to the measurement outcome (no communication)

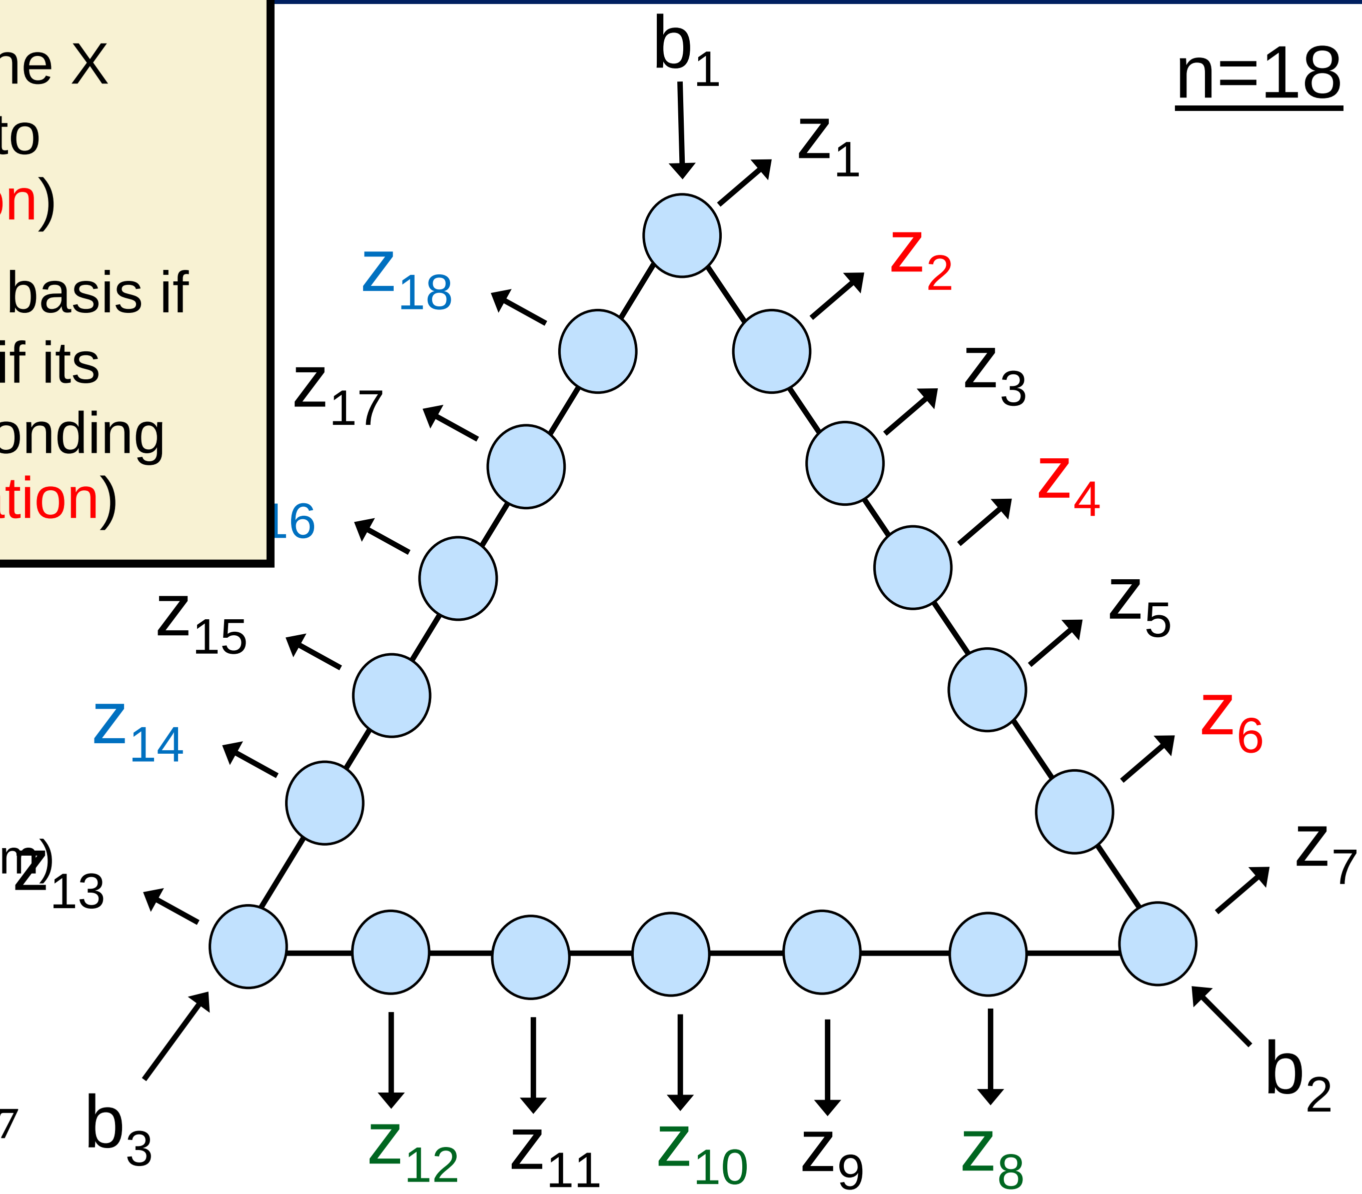

$m_R = z_2 \oplus z_4 \oplus z_6$

(parity of the outputs of the nodes of even index on the right)

$m_B = z_8 \oplus z_{10} \oplus z_{12}$

(parity of the outputs of the nodes of even index on the bottom)

$m_L = z_{14} \oplus z_{16} \oplus z_{18}$

(parity of the outputs of the nodes of even index on the left)

$m_{odd} = z_1 \oplus z_3 \oplus z_5 \oplus z_7 \oplus z_9 \oplus z_{11} \oplus z_{13} \oplus z_{15} \oplus z_{17}$

(parity of the outputs of all the nodes of odd index)

Claim 1: This quantum process samples from the uniform distribution over all binary strings $(z_1, z_2, \ldots, z_n) \in \{0,1\}^n$ satisfying the following condition:

$$\begin{cases} m_{odd} = 0 & \text{if } (b_1, b_2, b_3) = (0,0,0) \\ m_{odd} \oplus m_R = 1 & \text{if } (b_1, b_2, b_3) = (1,1,0) \\ m_{odd} \oplus m_B = 1 & \text{if } (b_1, b_2, b_3) = (0,1,1) \\ m_{odd} \oplus m_L = 1 & \text{if } (b_1, b_2, b_3) = (1,0,1) \end{cases}$$

Claim 2: Any classical protocol that samples (even approximately) from the same distribution requires long-distance communication.

- ✓ In any classical protocol in which no communication occurs between two nodes located at distance ≥ n/6:

  m_R is an affine function o

  m_B is an affine function o

  m_L is an affine function of b_1 and b_3

  m_{odd} is an affine function of b_1, b_2 and b_3

- ✓ Such functions cannot satisfy all the linear conditions of Claim 1

any nearest-neighbor classical circuit sampling from this distribution must have Ω(n) depth

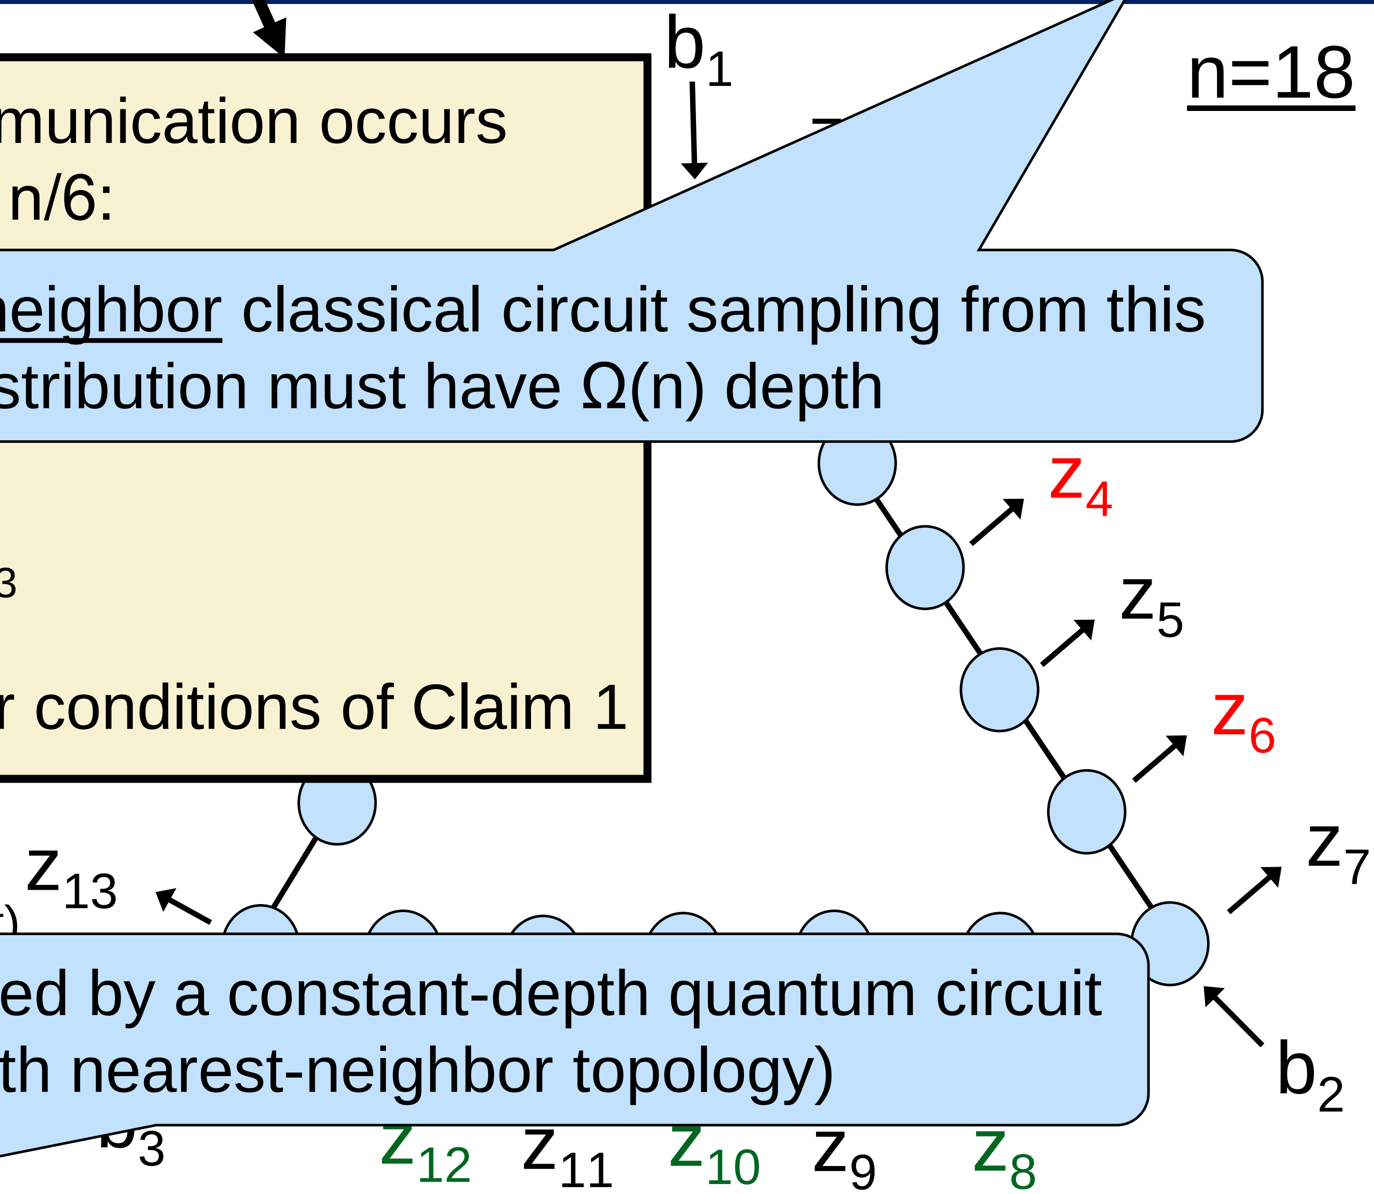

$m_R = z_2 \oplus z_4 \oplus z_6$

(parity of the outputs of the nodes of even index on the right)

$m_{odd} = z_1 \oplus z_3 \oplus z_5 \oplus z_7 \oplus z$

(parity of the outputs of all the nodes of odd ind

can be implemented by a constant-depth quantum circuit (even with nearest-neighbor topology)

Claim 1: This quantum process samples from the uniform distribution over all binary strings $(z_1, z_2, \ldots, z_n) \in \{0,1\}^n$ satisfying the following condition:

$$\begin{cases} m_{odd} = 0 & \text{if } (b_1, b_2, b_3) = (0,0,0) \\ m_{odd} \oplus m_R = 1 & \text{if } (b_1, b_2, b_3) = (1,1,0) \\ m_{odd} \oplus m_B = 1 & \text{if } (b_1, b_2, b_3) = (0,1,1) \\ m_{odd} \oplus m_L = 1 & \text{if } (b_1, b_2, b_3) = (1,0,1) \end{cases}$$

Advantage against Arbitrary Classical Circuits

Theorem ([Bravyi, Gosset, König 17])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) any classical circuit that solves it on all inputs has depth $\Omega(\log n)$.

not only restricted to nearest-neighbor topology

Consider a square grid of n nodes

Let m be the number of edges ($m = \Theta(n)$)

The input of the computational problem is a pair $(a,c) \in \{0,1\}^n \times \{0,1\}^m$

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

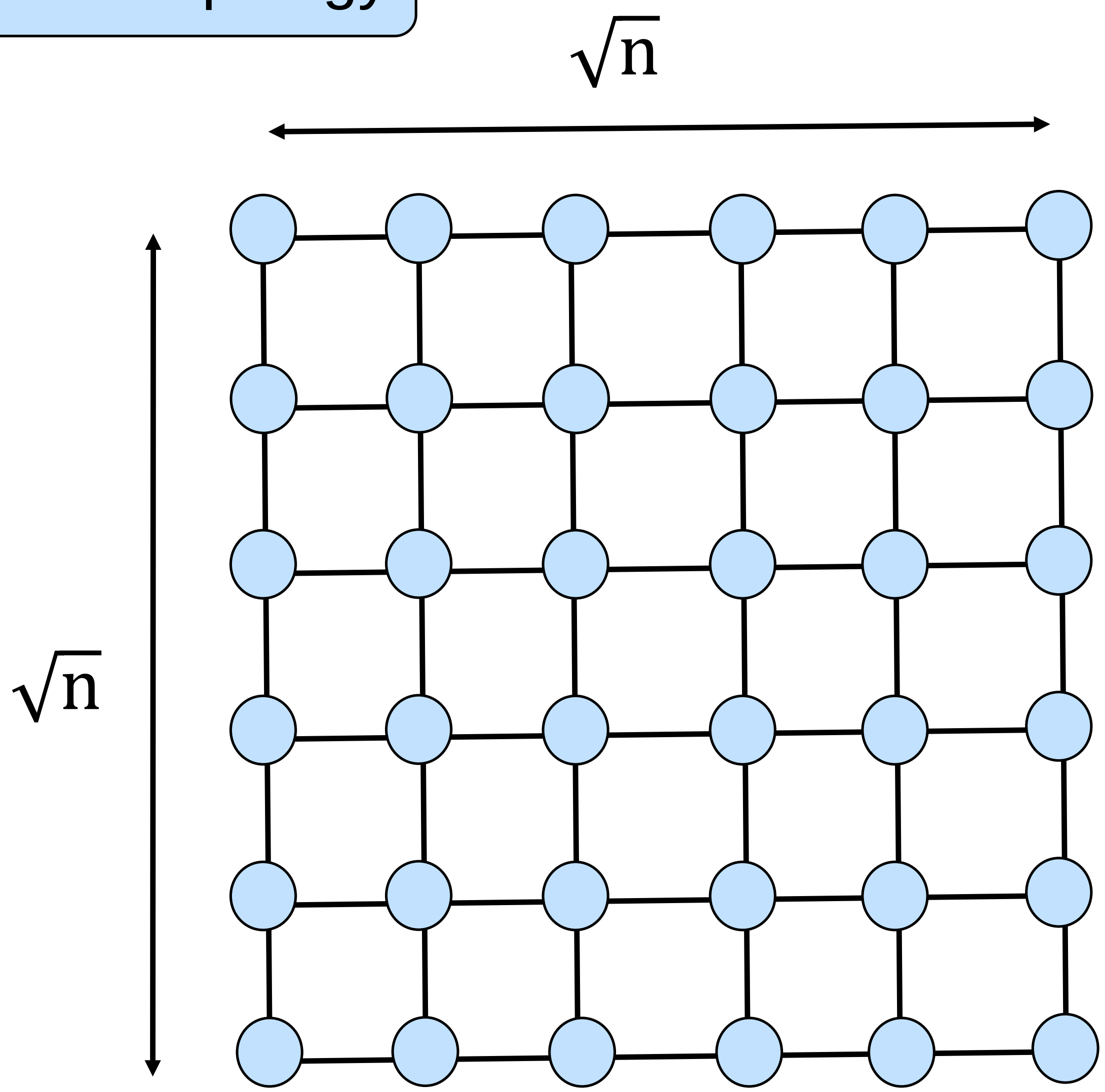

example: $n = 9$

$a = 110010100$ → measure nodes 1,2,5,7 in the Y basis, and the others in the X basis

$b = 110110101110$ → keep only the edges number 1,2,4,5,7,9,10,11

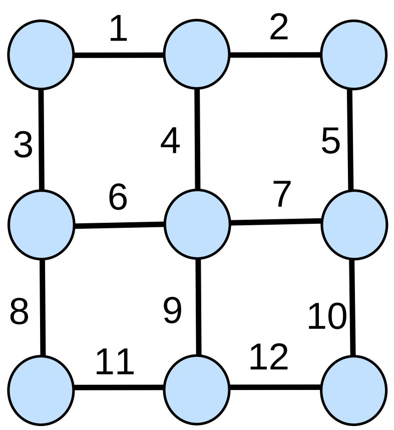

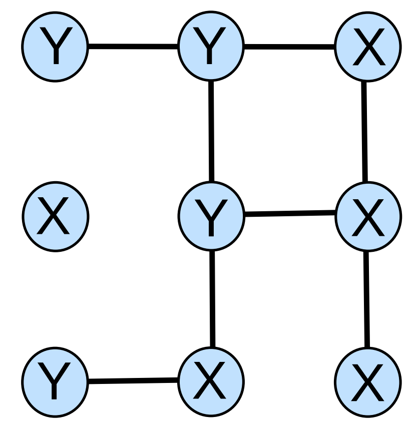

$m = 12$

Consider a square grid of n nodes

Let m be the number of edges ($m = \Theta(n)$)

The input of the computational problem is a pair $(a,c) \in \{0,1\}^n \times \{0,1\}^m$

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

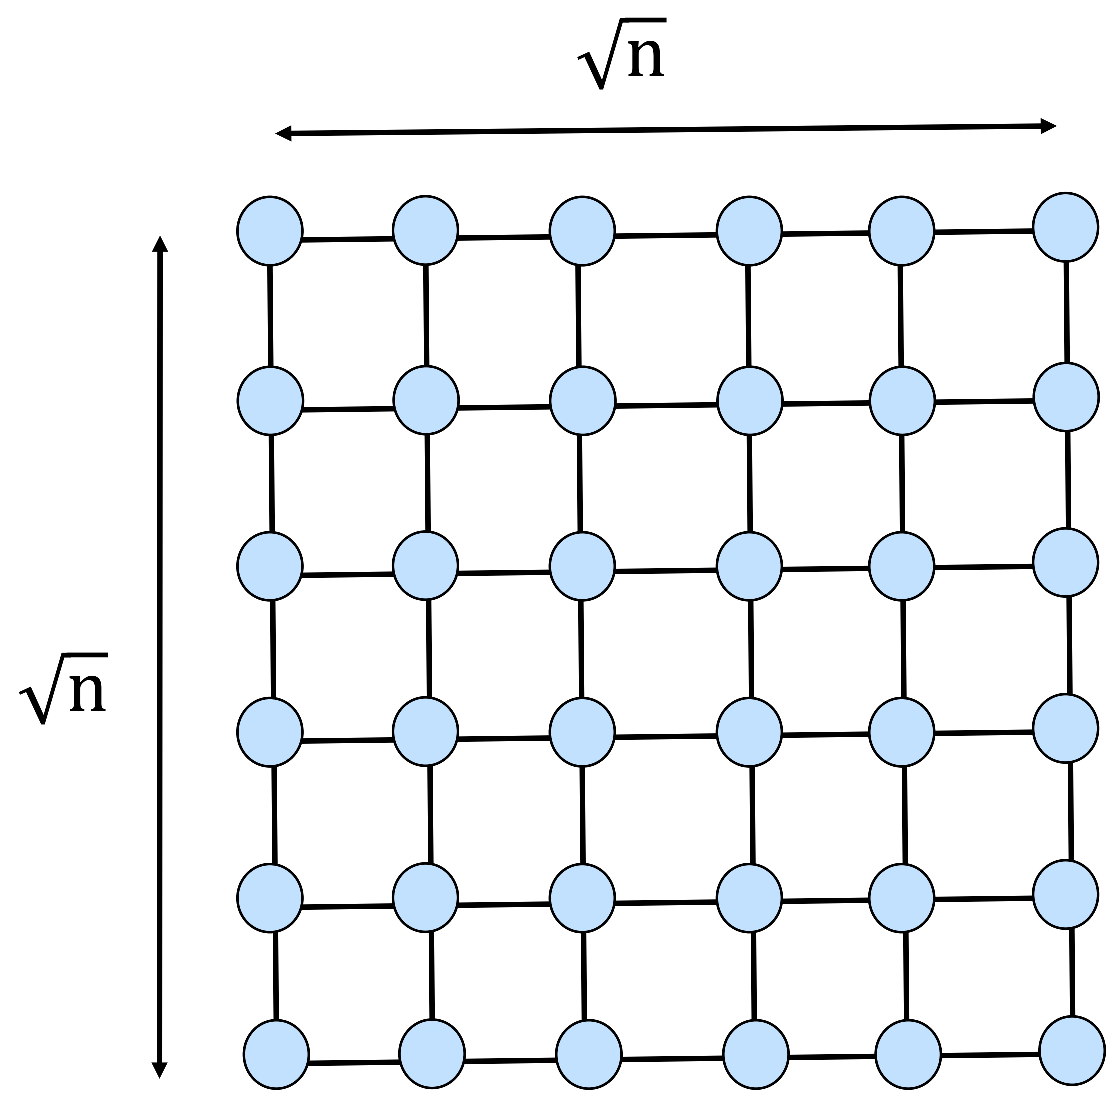

Advantage against Arbitrary Classical Circuits

straightforward: graph states on constant-degree graphs can be constructed by a constant-depth quantum circuit

Theorem ([Bravyi, Gosset, König 17])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) any classical circuit that solves it on all inputs has depth $\Omega(\log n)$.

Consider a square grid of n nodes

Let m be the number of edges ($m = \Theta(n)$)

The input of the computational problem is a pair $(a,c) \in \{0,1\}^n \times \{0,1\}^m$

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

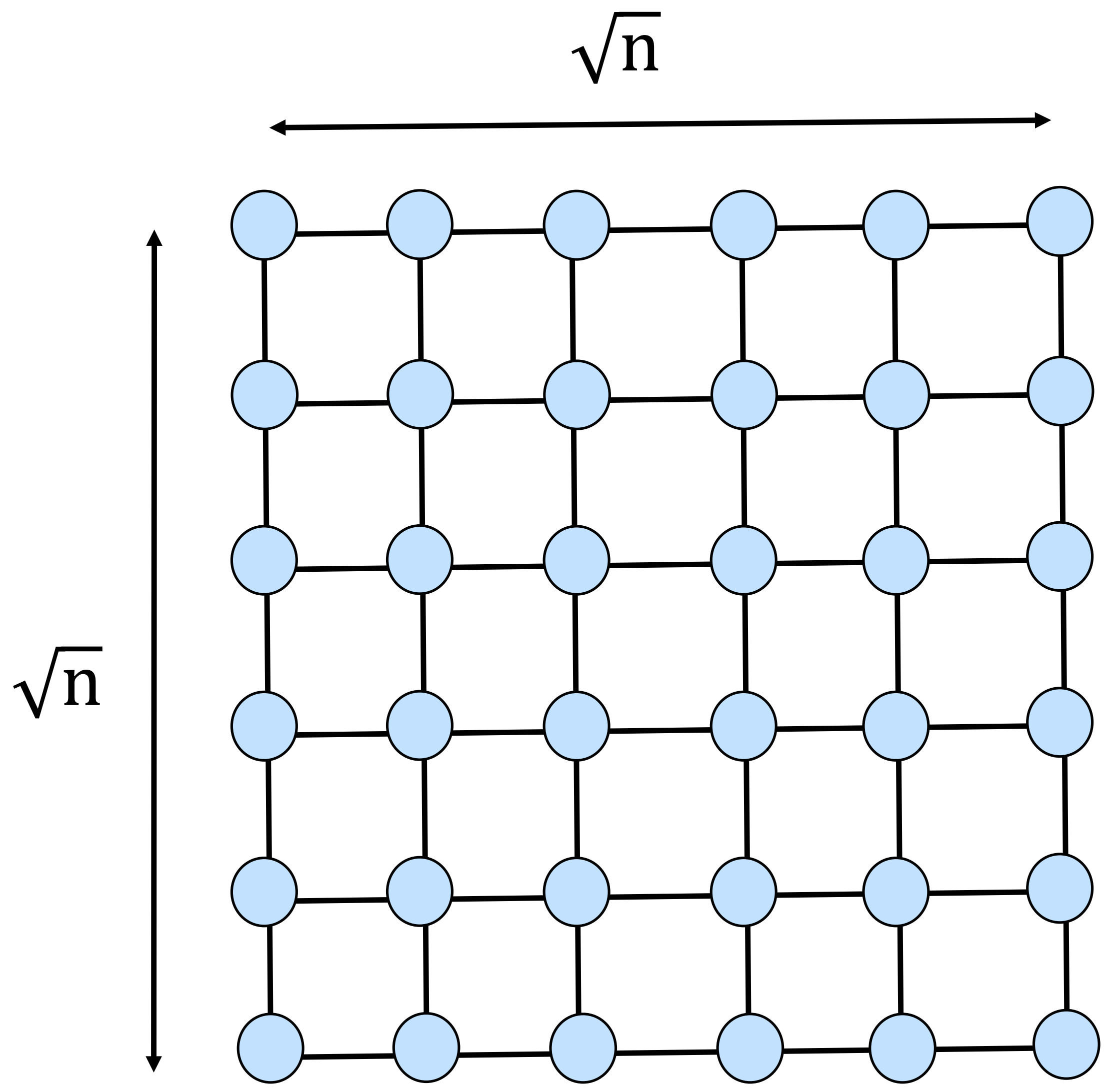

Proof of the C

Claim (trivial): In a classical circuit of small depth any input bit can contribute only to a small amount of output bits

- ✓ Consider any classical circuit of small depth that solves our problem
- ✓ The circuit has n + m input wires and n output wires
- ✓ Associate to each of the first n input wires the corresponding node of the grid
  Associate to each of the n output wires the corresponding node of the grid

Consider a square grid of n nodes

Let m be the number of edges ($m = \Theta(n)$)

The input of the computational problem is a pair $(a,c) \in \{0,1\}^n \times \{0,1\}^m$

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

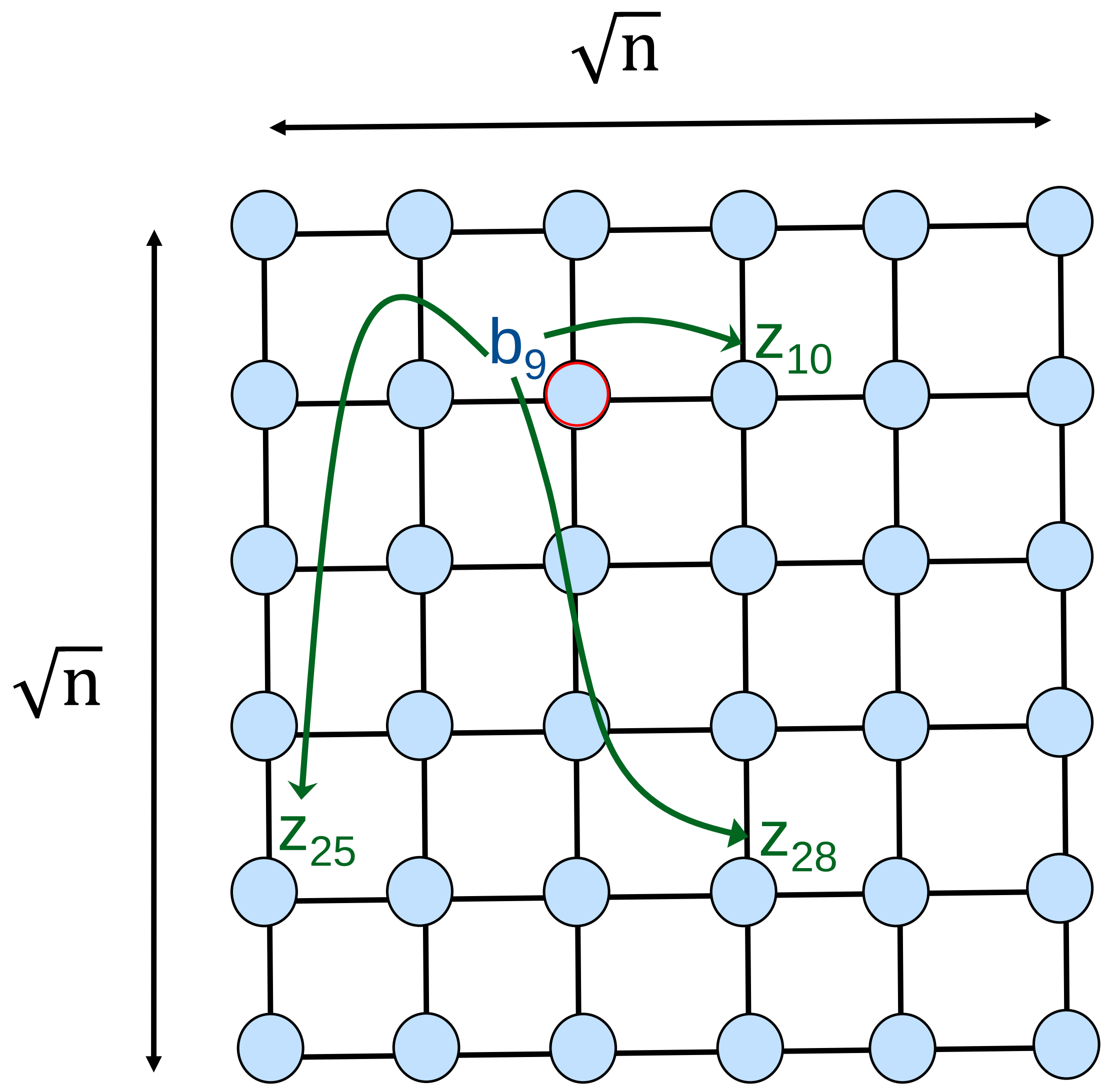

Proof of the C

Claim (trivial): In a classical circuit of small depth any input bit can contribute only to a small amount of output bits

- ✓ Consider any classical circuit of small depth that solves our problem
- ✓ The circuit has n + m input wires and n output wires
- ✓ Associate to each of the first n input wires the corresponding node of the grid
  Associate to each of the n output wires the corresponding node of the grid

There may be long-distance communication, but not for too many pairs

⇨ There exists a long cycle on which no long-distance communication occurs

- ✓ Consider the string a that specifies this long cycle
- ✓ The circuit cannot work for all strings c, from the argument from the first part of the talk

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

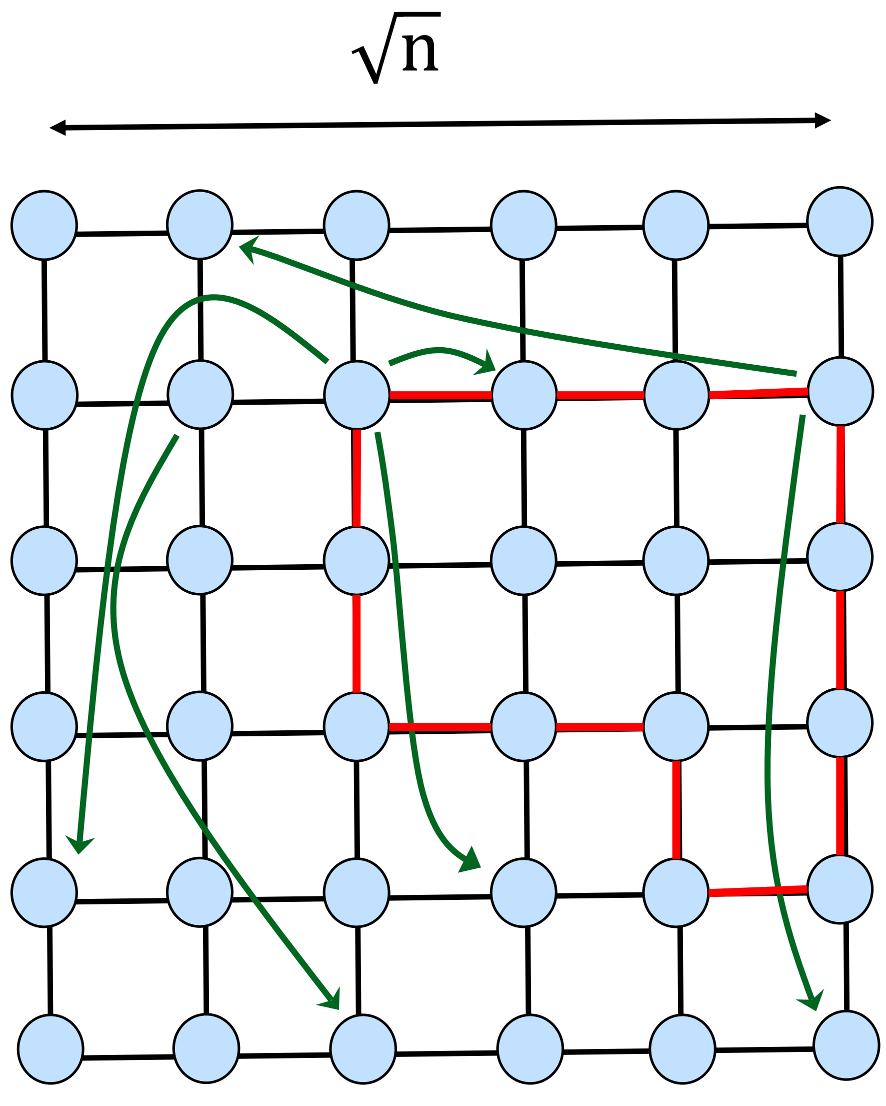

Advantage against Arbitrary Classical Circuits

straightforward: graph states on constant-degree graphs can be constructed by a constant-depth quantum circuit

Theorem ([Bravyi, Gosset, König 17])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) any classical circuit that solves it on all inputs has depth $\Omega(\log n)$.

for any small-depth classical circuit there exists an input (a,c) such that the circuit does not work

worst−case classical hardness

Consider a square grid of n nodes

Let m be the number of edges ($m = \Theta(n)$)

The input of the computational problem is a pair (a,c) $\in \{0,1\}^n \times \{0,1\}^m$

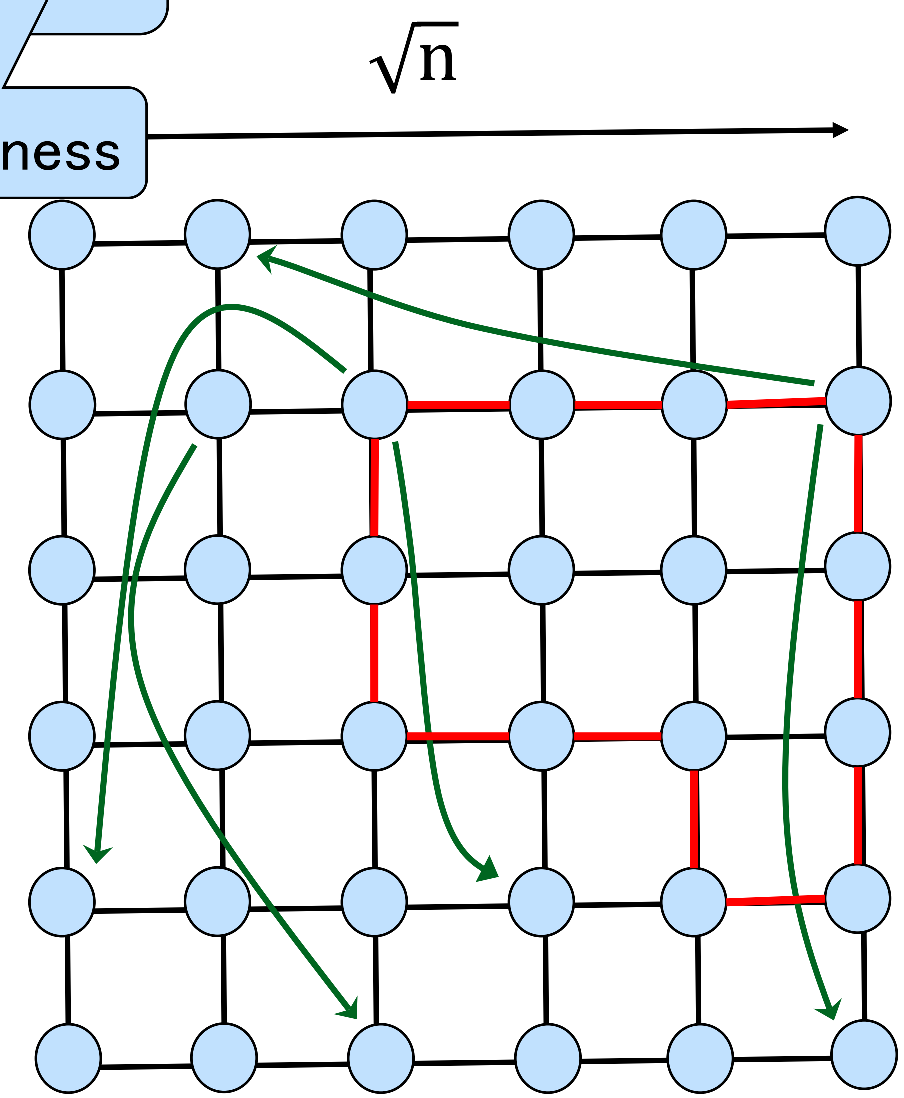

The computational problem asks to sample from the distribution corresponding to measuring the graph state specified by the string a in the basis specified by the string c

Getting Average-Case Hardness: Our Key Construction

any graph

Given a graph

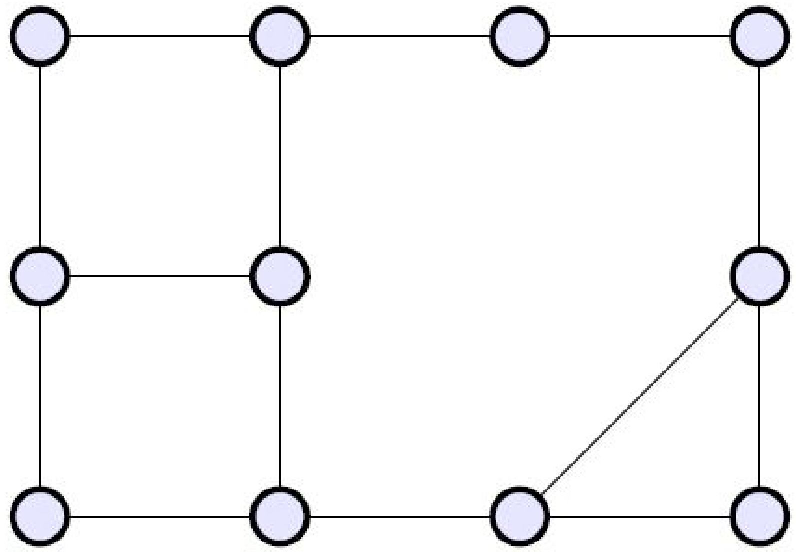

we define its “extended graph” as

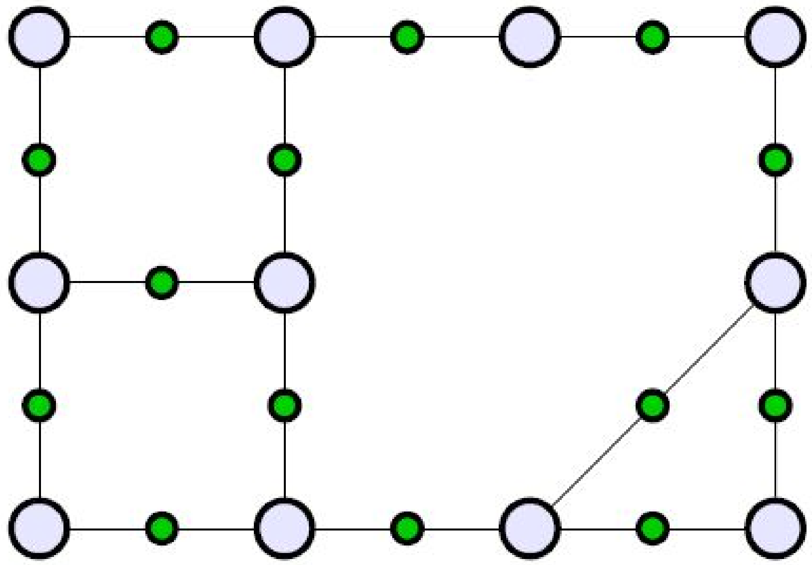

Similar construction used in, e.g., [Fujii and Morimae 2017]

1. The nodes prepare the graph state corresponding to the whole graph
2. Each non-corner node (this includes the nodes outside the cycle) measures its qubit in the X basis and then outputs the bit corresponding to the measurement outcome
3. Each corner node measures its qubit in the X basis if its input bit is 0, or measures it in the Y basis if its input bit is 1, and then outputs the bit corresponding to the measurement outcome

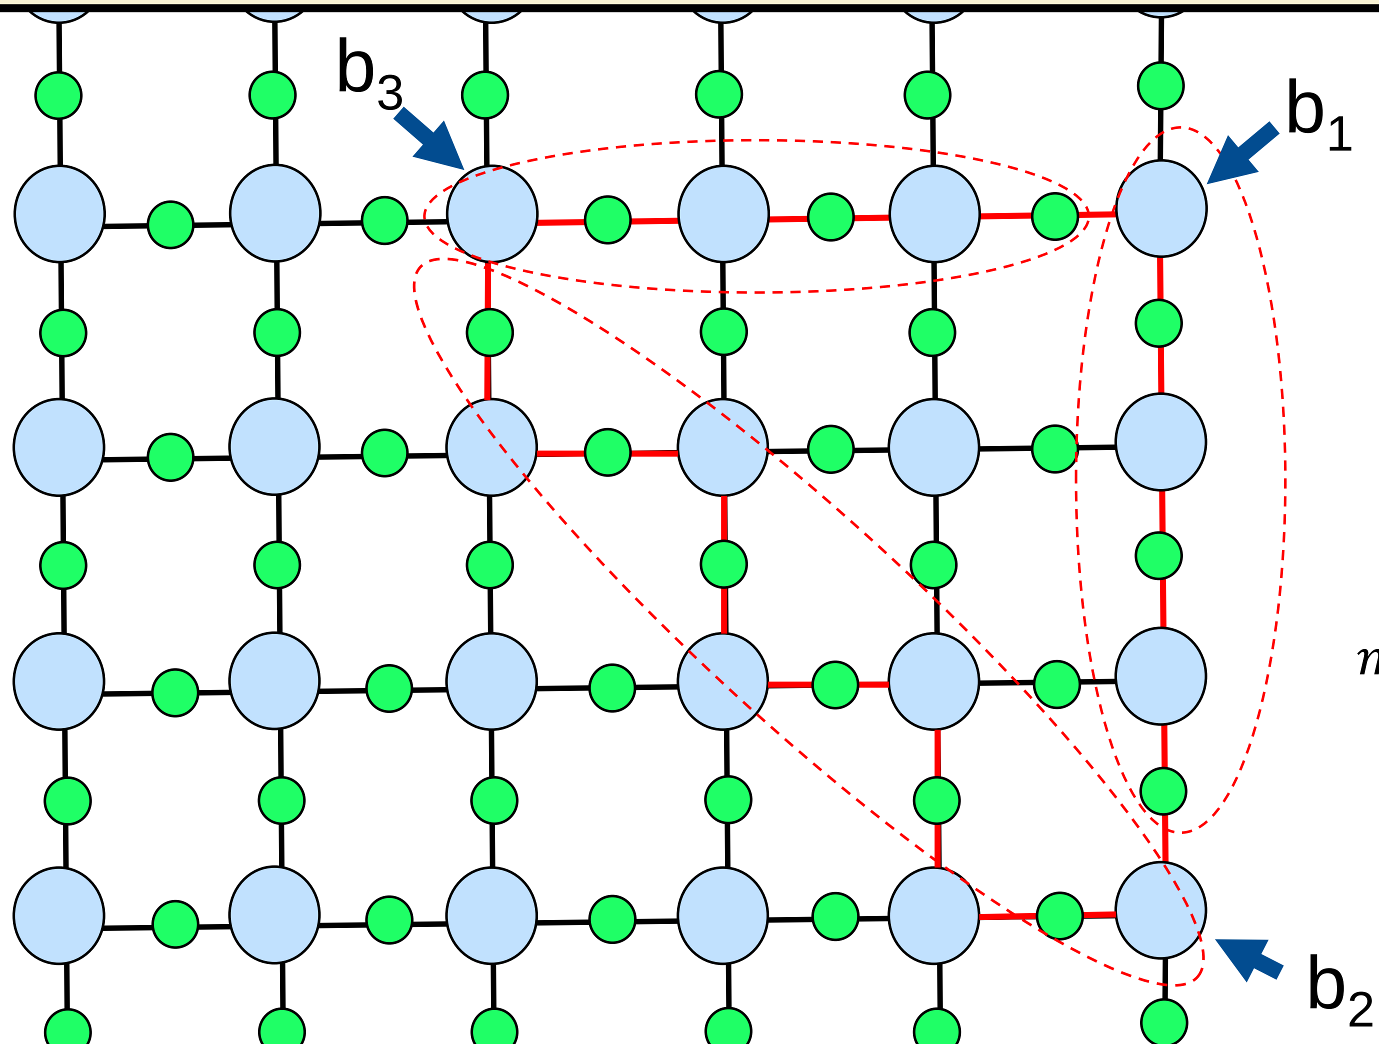

- ✓ Consider any cycle and see it as a triangle by dividing it into three parts (of roughly the same size)
- ✓ Each corner gets a bit as input
- ✓ Each node of the graph will output a bit

N: total number of vertices of the whole graph

m_{all}: parity of the outputs of all blue nodes

m_R: parity of the outputs of all green nodes in the right side of the triangle

m_T: parity of the outputs of all green nodes in the top side of the triangle

m_L: parity of the outputs of all green nodes in the left side of the triangle

Claim: This quantum process samples from the uniform distribution over all binary strings $(z_1, z_2, \dots, z_N) \in \{0,1\}^N$ satisfying the following condition:

$$\begin{cases} m_{all} = 0 & \text{if } (b_1, b_2, b_3) = (0,0,0) \\ m_{all} \oplus m_R = 1 & \text{if } (b_1, b_2, b_3) = (1,1,0) \\ m_{all} \oplus m_L = 1 & \text{if } (b_1, b_2, b_3) = (0,1,1) \\ m_{all} \oplus m_T = 1 & \text{if } (b_1, b_2, b_3) = (1,0,1) \end{cases}$$

Claim 2: Any classical protocol that samples (even approximately) from the same distribution requires long-distance communication.

- ✓ In any classical protocol in which no long-distance communication occurs between nodes on the three sides:

m_R is an affine function of b_1 and b_2
m_T is an affine function of b_1 and b_3
m_L is an affine function of b_2 and b_3
m_{all} is an affine function of b_1, b_2, b_3

- ✓ Such functions cannot satisfy all the linear conditions of the claim

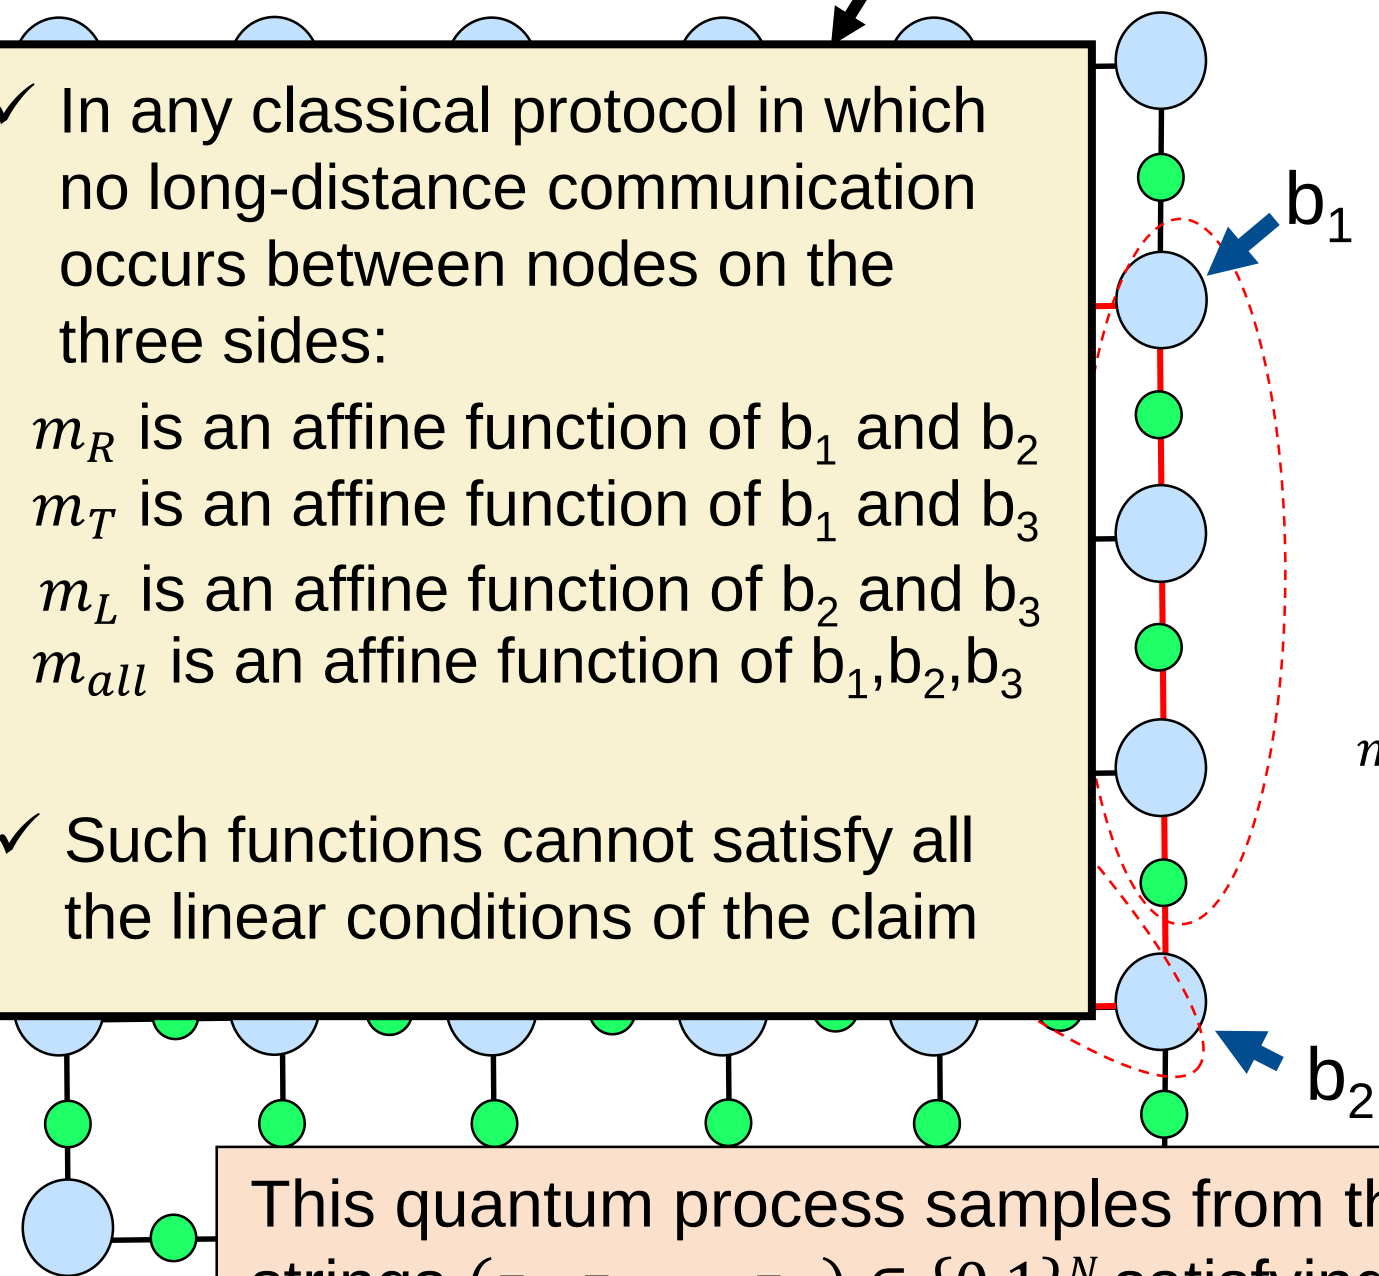

- ✓ Consider any cycle and see it as a triangle by dividing it into three parts (of roughly the same size)
- ✓ Each corner gets a bit as input
- ✓ Each node of the graph will output a bit

m_{all}: parity of the outputs of all blue nodes

m_R: parity of the outputs of all green nodes in the right side of the triangle

m_T: parity of the outputs of all green nodes in the top side of the triangle

m_L: parity of the outputs of all green nodes in the left side of the triangle

This quantum process samples from the uniform distribution over all binary strings $(z_1, z_2, \dots, z_N) \in \{0,1\}^N$ satisfying the following condition:

Claim:

$$\begin{cases} m_{all} = 0 & \text{if } (b_1, b_2, b_3) = (0,0,0) \\ m_{all} \oplus m_R = 1 & \text{if } (b_1, b_2, b_3) = (1,1,0) \\ m_{all} \oplus m_L = 1 & \text{if } (b_1, b_2, b_3) = (0,1,1) \\ m_{all} \oplus m_T = 1 & \text{if } (b_1, b_2, b_3) = (1,0,1) \end{cases}$$

Proof of the Classical Lower Bound

[Bravyi, Gosset, König 17]

✓ Consider any classical circuit of small depth that solves our problem

✓ The circuit has n + m input wires and n output wires

✓ Associate to each of the first n input wires the corresponding node of the grid
Associate to each of the n output wires the corresponding node of the grid

There may be long-distance communication, but not for too many pairs

⇨ There exists a long cycle on which no long-distance communication occurs

Needed to remove everything except this red cycle

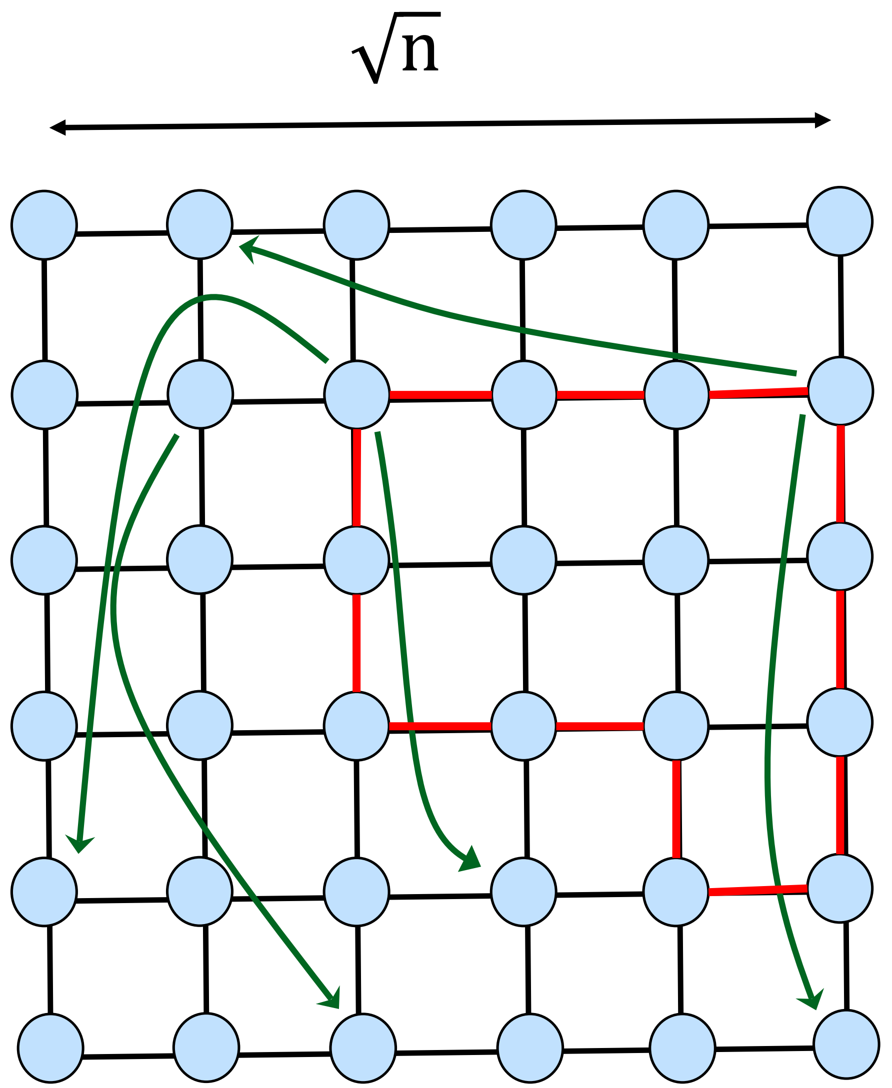

Proof of the Classical Lower Bound

[Bravyi, Gosset, König 17]

- ✓ Consider any classical circuit of small depth that solves our problem
- ✓ The circuit has n + m input wires and n output wires
- ✓ Associate to each of the first n input wires the corresponding node of the grid
  Associate to each of the n output wires the corresponding node of the grid

There may be long-distance communication, but not for too many pairs

⇨ There exists a long cycle on which no long-distance communication occurs

Needed to remove everything except this red cycle

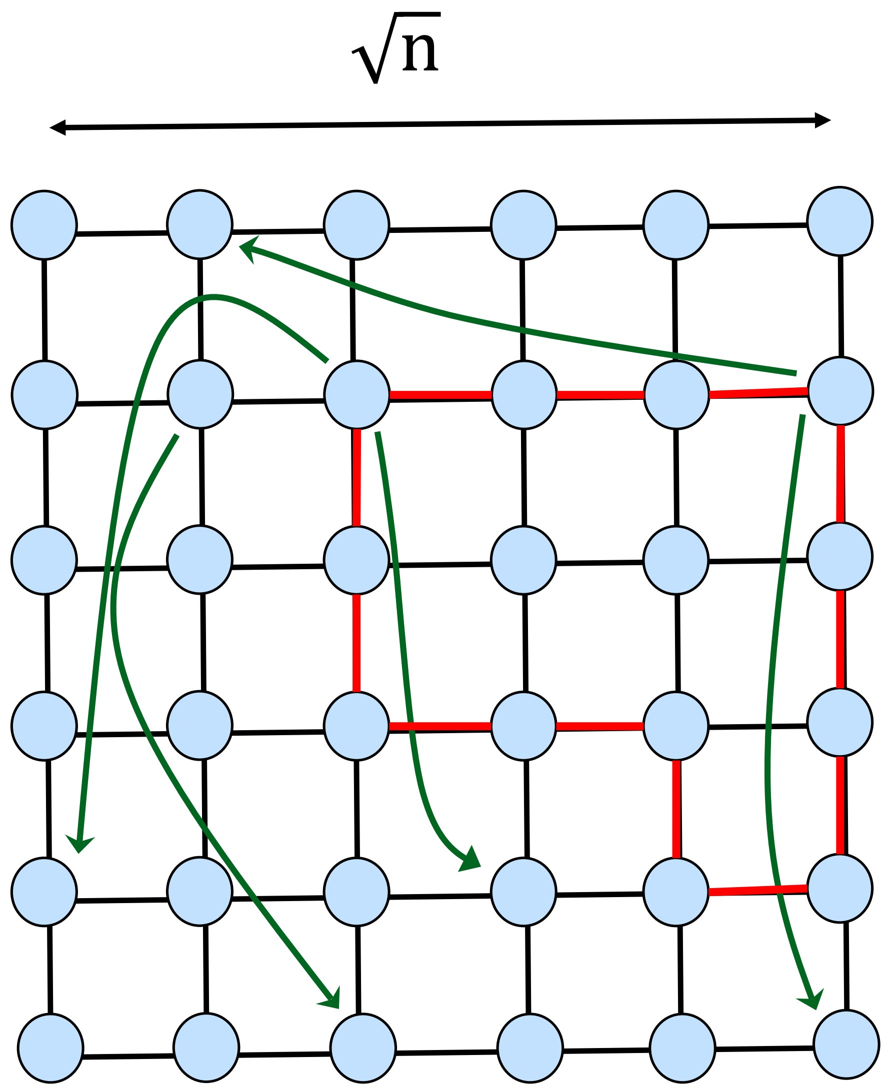

Proof of the Classical Lower Bound

with our trick

- ✓ Consider any classical circuit of small depth that solves our problem
- The input of the new computational problem is simply a string $c \in \{0,1\}^m$
- ✓ The circuit has $n + m$ input wires and n output wires
- ✓ ... node of the grid
- ... de of the grid

The new computational problem asks to sample from the distribution corresponding to measuring the extended graph state of the square grid in the basis specified by the string c

⇨ There exists a long cycle on which no long-distance communication occurs

No need to remove anything, since our new impossibility argument works even with the vertices outside the cycle

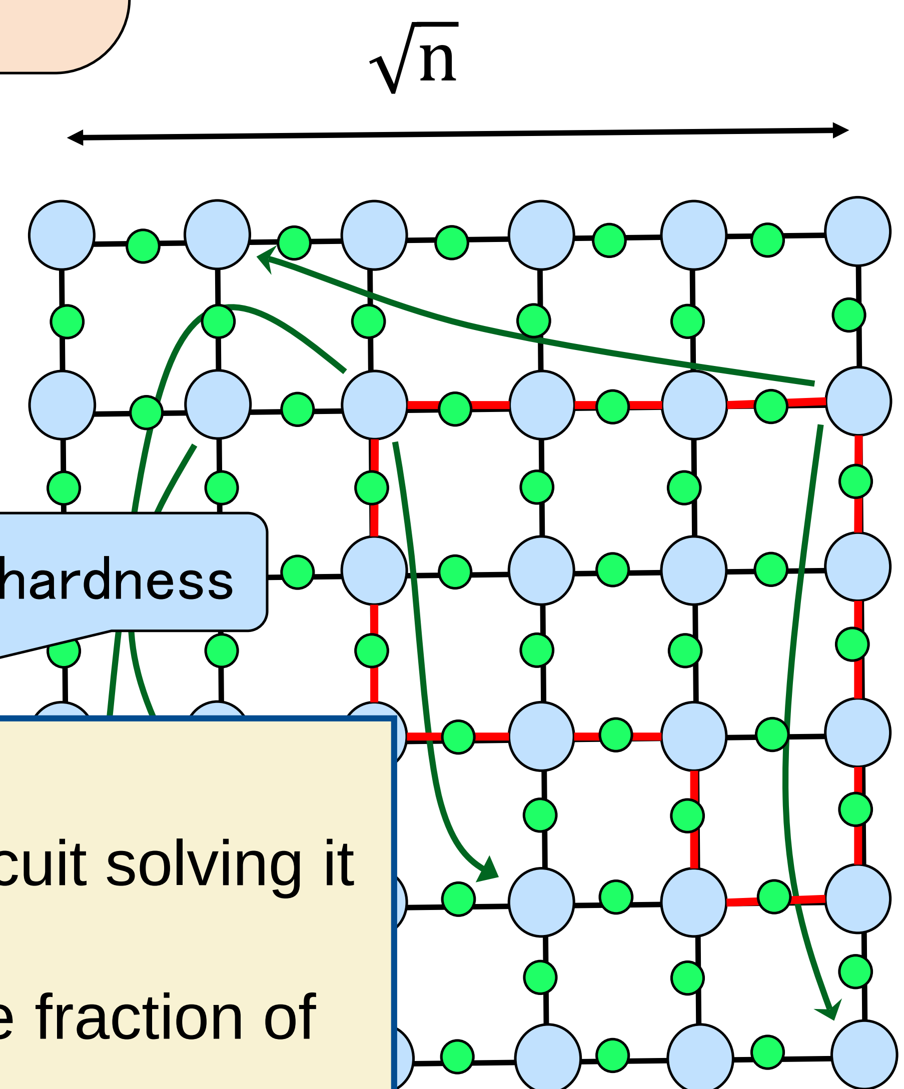

average-case classical hardness

Our result

For this computational problem:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all strings c; but

(ii) any classical circuit that solves it on a non-negligible fraction of the strings c has depth $\Omega(\log n)$.

Omitted Details

- ✓ To obtain such a strong average-case hardness result we need to use amplification (repeat the same process on several copies of the construction)

- ✓ For technical reasons we need to work a graph slightly more complicated

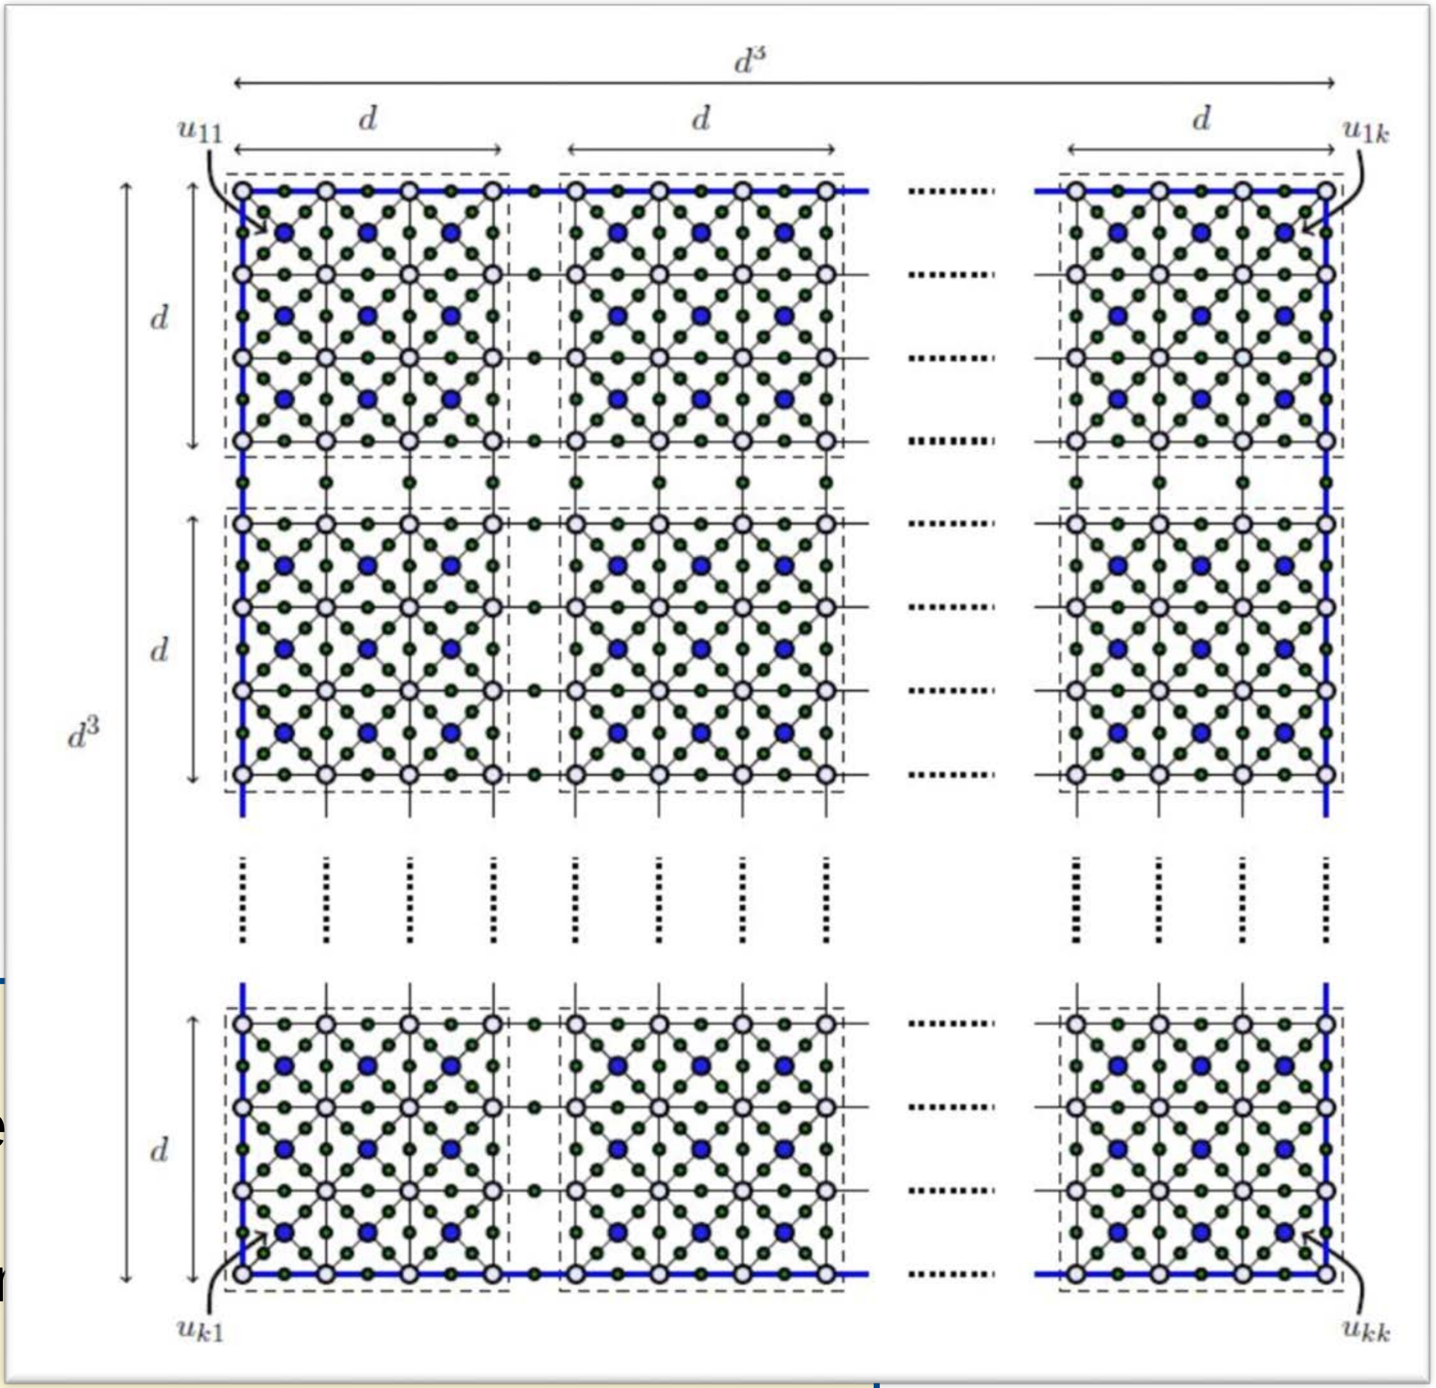

Our result

For this computational problem:

(i) there is a shallow (i.e., constant-de
on all strings c; but

(ii) any classical circuit that solves it o
the strings c has depth Ω(log n).

Relation with Concurrent Works

worst−case classical hardness

Theorem ([Bravyi, Gosset, König 17 (ArXiv version)])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) no shallow classical circuit can solve it on all inputs.

Theorem ([Bravyi, Gosset, König 18 (Supplementary materials)])

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) no shallow classical circuit can solve it on a constant fraction of inputs.

average−case classical hardness

Our result:

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) no shallow classical circuit can solve it on a non-negligible fraction of inputs.

but different construction

[Coudron, Stark, Vidick 18]: same statement as ours + application to randomness expansion

different construction

[Bene Watts, Kothari, Schaeffer, Tal (unpublished, QIP’19, STOC’19)]:
same statement as ours + holds even against classical circuits with unbounded fanin

Conclusion and Open Problems

does not rely on any conjecture or assumption

Our result: average-case quantum advantage usin

solves it only when there is no noise

There exists a computational problem such that:

(i) there is a shallow (i.e., constant-depth) quantum circuit solving it on all inputs; but

(ii) no shallow classical circuit can solve it on a non-negligible fraction of inputs.

but a logarithmic-depth classical circuit can solve it

Open problem #1: quantum supremacy with noisy quantum computation

[Bravyi, Gosset, König, Tomamichel 19] showed a noisy version of this theorem using error-correction techniques (for local noise)

What about more general versions of noise?

Open problem #2: show advantage against stronger classes of classical computation

Can we break this logarithmic barrier for a separation that does not rely on any conjecture or assumption?