CSE 291 / Math 277A - Quantum Complexity Theory November 8, 2022

Lecture 13

Lecturer: Daniel Grier *Scribe: Madhav Wagle*

1 Overview

In the last lecture, we proved that $\mathsf{NC}^0 = \mathsf{QNC}^0$; that is constant depth quantum circuits with bounded fan-in offer no advantage over constant depth classical circuits with bounded fan-in *for decision problems.* In this lecture, we will see that this equivalence does not extend to relation problems; that is, problems for which the goal is to output some bit string (instead of a single bit) from a set depending on the input. In particular, we will describe relation problem that can be computed exactly by constant-depth quantum circuits with 1- and 2-qubit geometrically local gates, but can't be computed by any constant-depth classical circuit with bounded fan-in. This will prove a separation between the classes FNC^0 and FQNC^0, the relational analogues of NC^0 and QNC^0.

The first ingredient for this result is a quantum protocol for the GHZ game (defined in the last lecture) which allows the players to win with 100% probability. Recall that in the last lecture we showed that if the players were classical, then their max probability of winning was 75%. Before we do that, we introduce a fundamental set of quantum states called graph states, which are built from Hadamard and CSIGN gates.

2 The CSIGN gate

The 2-qubit CSIGN gate (also know as Controlled-Z or CZ) gate is a 2-qubit unitary gate such that,

$$\text{CSIGN}\,|xy\rangle = (-1)^{x\cdot y}\,|xy\rangle$$

where $x, y \in \{0,1\}$. Notice that the CSIGN gate is symmetric with respect to its input qubits, which is reflected in how it's typically drawn in a circuit diagram:

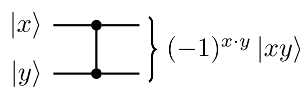

The CSIGN and CNOT gates are equivalent up to conjugation of the target by Hadamard:

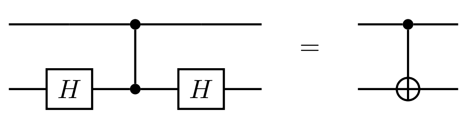

In particular, this implies that CISGN is a Clifford gate.

3 Graph States

A graph state is a special quantum state where each qubit is associated with a vertex in a graph. Given a graph $G = (V, E)$ with $|V| = n$, the n-qubit graph state corresponding to G can be constructed as follows:

1. For every vertex $v \in V$, create a new qubit in the $|+\rangle$ state.
2. For every edge $e = (u, v) \in E$, apply a CSIGN between the qubits corresponding to u and v.

This yields a state of the following form,

$$|G\rangle := \frac{1}{\sqrt{2^n}} \sum_{x\in\{0,1\}^n} (-1)^{\sum_{(u,v)\in E} x_u x_v}\,|x\rangle$$

Figure 1 shows the circuit for the triangle graph, i.e., the graph which is a cycle on 3 vertices.

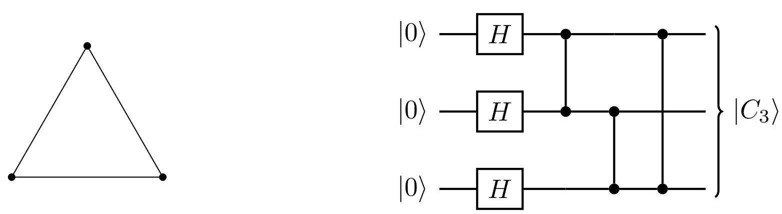

Figure 1: (*left*) triangle graph C_3 and (*right*) the circuit that generates the triangle graph state.

4 Quantum Advantage for the GHZ game

In the last lecture, we defined the GHZ game and proved that there exists no classical probabilistic protocol/algorithm that can provably win the GHZ game with probability at least 0.75. In this section, we will see a quantum strategy that always wins.

Theorem 1 (Greenberger-Horne-Zeilinger [GHZ89])**.** *There exists a quantum strategy for the GHZ game, such that the three players win with probability 1.*

Proof. The three players Alice, Bob and Charles share a three-qubit entangled state

$$|\text{GHZ}\rangle = \frac{1}{\sqrt{2}}(|0_A 0_B 0_C\rangle + |1_A 1_B 1_C\rangle).$$

Each player holds one qubit from the above shared state. All three players use the following strategy:

Quantum strategy for a player

On receiving bit s from the referee:

- $s = 0$: apply a Hadamard gate and measure (equivalently, measure in the X basis), and return the outcome to the referee.
- $s = 1$: apply a phase gate followed by a Hadamard gate and measure (equivalently, measure in the Y basis) and return the outcome to the referee.

Analysis: Note that all the operators in both cases of the strategy are Clifford operators. Thus, we can use the stabilizer group representation to analyze the strategy. We will need two lemmas, the proofs of which are given in Section 7.

Lemma 1. *A set of generators of the stabilizer group of the GHZ state is* $\{XXX, ZZI, IZZ\}$.

Lemma 2. *Let* $|\psi\rangle$ *be any Clifford state whose stabilizer group contains a Pauli element that is a tensor product of* Z *and* I *elements. That is,* $|\psi\rangle$ *is stabilized by* $P = \alpha P_1 \otimes \cdots \otimes P_n$ *such that* $P_i \in \{Z, I\}$ *and* $\alpha = \{\pm 1\}$. *Measure* $|\psi\rangle$ *in the computational basis, but consider only the measurements on qubits* i *such that* $P_i = Z$. *If* $\alpha = 1$, *then the parity of the measurement results is even; otherwise* ($\alpha = -1$), *the parity is odd.*

Now we will perform a case analysis on the values of the inputs (a, b, c) to show that our strategy always succeeds. We only need to consider the cases $(a, b, c) \equiv (0, 0, 0)$ and $(a, b, c) \equiv (1, 1, 0)$. The cases $(a, b, c) \equiv (0, 1, 1)$ and $(a, b, c) \equiv (1, 0, 1)$ follow from symmetry. In each case, we would keep track of the evolution of the stabilizer group generators and then apply Lemma 2.

- $(a, b, c) \equiv (0, 0, 0)$: All three players measure in the X basis:

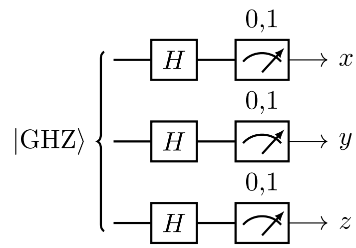

Using Lemma 1 to obtain the generators of our starting state, their evolution is given by

$$\underbrace{\begin{Bmatrix} XXX \\ Z\,Z\,I \\ Z\,I\,Z \end{Bmatrix}}_{\text{Stabilizers of } |\text{GHZ}\rangle} \xrightarrow{H^{\otimes 3}} \begin{Bmatrix} Z\,Z\,Z \\ XX\,I \\ X\,I\,X \end{Bmatrix}$$

As the stabilizer group of the state being measured contains the Pauli string ZZZ, Lemma 2 states that the output must have even parity, i.e., $x \oplus y \oplus z = 0$

- $(a, b, c) \equiv (1, 1, 0)$: Alice and Bob (who received bit a and b) measure their qubits in the Y basis. Charlie measures in the X basis. The circuit representation in this case is

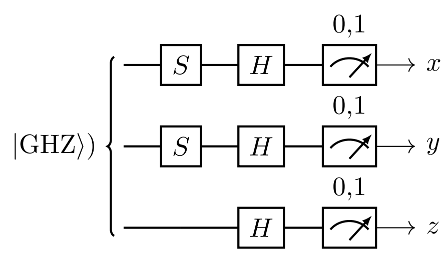

We get

$$\underbrace{\begin{Bmatrix} XXX \\ Z\,Z\,I \\ Z\,I\,Z \end{Bmatrix}}_{\text{Stabilizers of } |\text{GHZ}\rangle} \xrightarrow{S\otimes S\otimes I} \begin{Bmatrix} YYX \\ Z\,Z\,I \\ Z\,I\,Z \end{Bmatrix} \xrightarrow{H\otimes H\otimes H} \begin{Bmatrix} YY\,Z \\ XX\,I \\ X\,I\,X \end{Bmatrix}$$

Using group closure properties, $-ZZZ = (YYZ)(XXI)$ is in the stabilizer group of the state being measured. Therefore, by Lemma 2 the measurement output must have odd parity, ie. $x \oplus y \oplus z \oplus = 1$.

Thus, we conclude that whenever the parity of the input is even (i.e., $a \oplus b \oplus c = 0$), the three players return bits x, y, z such that $x \oplus y \oplus z = a \vee b \vee c$. In other words, they always answer correctly to the referee. □

5 Generalizing the non locality in the GHZ game

Recall that we want to show that there is some relation problem that a constant-depth circuit can solve that any classical circuit fails to solve with high probability. At first glance, the GHZ game might not seem particularly relevant to that question since it applies in a different communication-restricted setting. However, consider the following implication of the GHZ game: every 3-input, 3-output classical circuit solving the GHZ relation must have some input bit whose light cone contains a different output bit. Otherwise, the output bits are only functions of their particular input bits, which mimics the scenario in which the players cannot communicate!

Of course, this insight only takes us so far. Any 3-input, 3-output circuit can trivially be simulated in constant depth. To leverage this idea further, we will need to generalize the GHZ game.

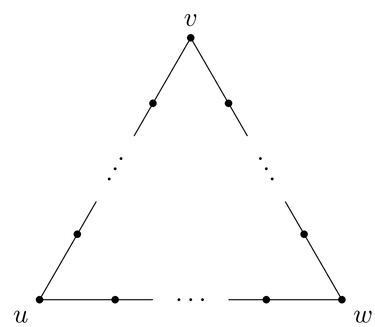

Figure 2: The big triangle graph.

5.1 Triangle graph state and the GHZ game

Our first observation is that the GHZ state is locally-equivalent to the 3-qubit triangle graph state. By *locally-equivalent*, we mean that it can be converted into the GHZ state (and vice versa) by applying a tensor product of 3 single-qubit Clifford unitaries.

To see this, recall that the generators for the stabilizer group of $|\mathrm{GHZ}\rangle$ are $\{XXX, ZZI, ZIZ\}$ since the GHZ state is just the 3-qubit cat state. The generators for the triangle graph state are given by

$$\begin{Bmatrix} ZII \\ IZI \\ IIZ \end{Bmatrix} \xrightarrow{H^{\otimes 3}} \begin{Bmatrix} XII \\ IXI \\ IIX \end{Bmatrix} \xrightarrow{\text{CSIGN gates}} \begin{Bmatrix} XZZ \\ ZXZ \\ ZZX \end{Bmatrix}$$

It remains to check that

$$\begin{Bmatrix} XXX \\ ZZI \\ ZIZ \end{Bmatrix} \xrightarrow{U\otimes V\otimes V} \begin{Bmatrix} XZZ \\ ZXZ \\ ZZX \end{Bmatrix}$$

where $U = HS^\dagger H$ and $V = SH$, which we leave as an exercise.

In other words, the triangle graph state has the same kind non-locality as the GHZ state since single qubit Cliffords don't affect the entanglement between the qubits. Because of this, we can define a new relational problem with the same gap in the winning probabilities as the GHZ game. The benefit of this new framing, however, is that it is a bit clearer how we can generalize—just start with a larger graph state.

5.2 The Big Triangle Game

In this section, we will define the $\mathsf{BIGTRIANGLE}$ problem as a candidate generalization of the GHZ game that will be hard to simulate. Unfortunately, it will turn out that this problem in itself does not separate FNC^0 and FQNC^0. Nevertheless, we will show that solving it requires a certain kind of geometric non-locality that will be critical later.

Let C_n be the cycle graph on n vertices and $|C_n\rangle$ be the corresponding n-qubit graph state. Notice that we can also view C_n as a big triangle as shown in Figure 2. We label the corners of the triangle as u, v, w.

Definition 2 ($\mathsf{BIGTRIANGLE}$ problem)**.** *Given a big triangle graph C_n with corner vertices u, v, w, and a vector $b = (b_u, b_v, b_w) \in \{0,1\}^3$ as input. Report any measurement outcome on the vertices in the graph state $|C_n\rangle$ such that the qubits on the sides of the triangle are measured in the X basis and the corner vertices are measured according to b (measure vertex a in X basis if $b_a = 0$, measure vertex a in Y basis if $b_a = 1$).*

$\mathsf{BIGTRIANGLE}$ is clearly a relation problem defined specifically to be solved by a constant-depth quantum circuit, but let's first just check this claim:

Theorem 3 ($\mathsf{BIGTRIANGLE} \in \mathsf{FQNC}^0$)**.** *There is a uniform family of bounded fan-in, constant-depth quantum circuits $\{Q_n\}_n$ that are geometrically-local on the triangle that solve the big triangle problem.*

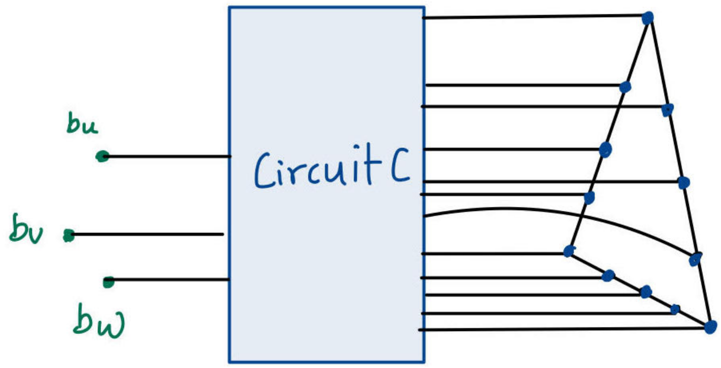

Figure 3: A circuit C that solves the big triangle problem, it takes 3 inputs on the left which specify the measurement basis for the corner vertices u, v, w. On the right side is the output of the circuit, one output for each of the n vertices of the triangle graph. (The triangle frame going through the output bits is drawn for clarity)

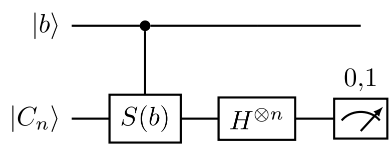

Figure 4: Constant depth circuit Q_n that solves BIGTRIANGLE.

Proof. The circuit in Figure 4 solves BIGTRIANGLE in constant depth. Notice first that $|C_n\rangle$ can be constructed in constant depth because each vertex in the triangle graph has degree 2. After the graph state is created we just need to measure in the appropriate basis. The bits $b \in \{0,1\}^n$ specify the measurement basis for each vertex of the triangle. If b_v is 1 a phase gate is applied followed by Hadamard before measuring in the Z basis. Similarly, if b_v is 0 only a Hadamard is applied before measuring in the Z basis. It follows from the definition of BIGTRIANGLE that this circuit always computes a correct answer. □

On the other hand, geometrically local classical circuits aren't so lucky...

Lemma 4 ([BCE+07, BGK18]). *If a classical probabilistic circuit C solves* BIGTRIANGLE *with high probability then the light cone of at least one of the input bits (i.e., b_u, b_v, b_w) in circuit C, must contain an output bit that is at a distance of at least $D/2$ where D is the minimum length of an edge in the big triangle.*[1]

While we will not prove this lemma, the proof strategy is identical to that of the GHZ game. It is just more involved. This has the following immediate consequence:

Theorem 5. *No bounded fan-in classical circuit which is geometrically local on the triangle can solve the* BIGTRIANGLE *game with high probability.*

Proof. The geometric locality condition combined with constant depth forces the light cone of each input bit b_u, b_v, b_w to only cover outputs for vertices that are at a constant distance from the input vertex along each edge of the big triangle. See Figure 5. This violates Lemma 4. □

Of course, this result still falls a bit short of what we'd like to prove. In particular, the geometric locality constraint on the classical circuits feels a bit contrived. We will remove it in the next section.

[1] If we want to be fully rigorous, we must impose the additional condition that the side-lengths of the triangle are even. This is mostly just an annoyance for the rest of the proof, so we will drop it. Full details in [BGK18].

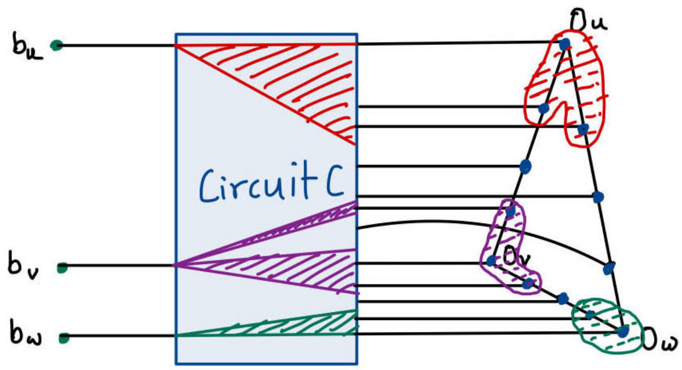

Figure 5: In this circuit, the light cone of each input bit is shown. By Lemma 4 this circuit can't compute BIGTRIANGLE with high success probability, as the light cones for all the input bits are confined in close neighborhood of their respective corner output bit.

6 Separation between FNC^0 and FQNC^0

The problem with the BIGTRIANGLE problem is that are only a few inputs, corresponding to the measurement choices for the vertices at the corners of the triangle. A classical circuit just needs to have answers for those few inputs. To trip up a classical simulator, we need to make sure there are a lot of possible big triangle instances. If the classical circuit is constant depth, it won't be able to coordinate its answers to all of these big triangle problem instances simultaneously. To this end, we define a generalization of the big triangle problem on a 2D grid.

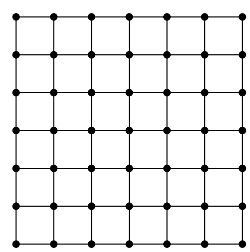

Figure 6: The 7×7 2D grid graph.

Definition 6 (GRID problem)**.** *Let $G = (V, E)$ be the $N \times N$ grid graph. The input is a vector $A \in \{0,1\}^{|E|}$ that specifies a subgraph of the grid G and a vector $b \in \{0,1\}^{|V|}$ corresponding to the vertices of that matrix. The goal is to output any measurement result on the graph state with edges specified by A and measurements bases specified by b (in X basis if $b_v = 0$, in Y basis if $b_v = 1$).*

Note that GRID is a relational problem. For every vertex v in the grid there is a corresponding input qubit b_v giving the measurement basis, and an output qubit M_v reporting the measurement outcome on the qubit corresponding to v.

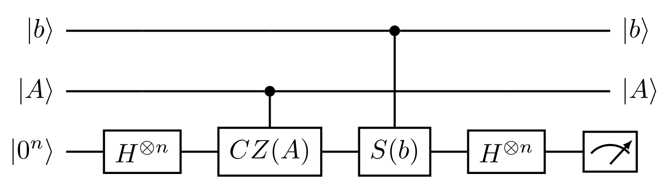

Figure 7: Constant depth quantum circuit that solves GRID [BGK18].

Theorem 7 (GRID $\in$ FQNC^0)**.** *There is a uniform family of bounded fan-in, constant-depth quantum circuits $\{Q_N\}_N$ that are geometrically-local on the $N \times N$ 2D grid that solve the* GRID *problem.*

Proof. The quantum circuit in Figure 7 solves the GRID problem in constant depth. The first two layers (Hadamard and CSIGN gates) of the circuit construct the graph state specified by the vector A. Since G has degree at most 4, so does the graph specified by A, so this state can be constructed in constant depth. After the graph state is prepared, we simply measure according to the bases specified by b. □

Theorem 8 (GRID $\notin$ FNC^0)**.** *Let $\{C_N\}_{N=1}^{\infty}$ be any family of classical bounded fan-in circuits solving the* GRID *problem with high probability. Then, the depth of C_N is at least $\Omega(\log(N))$.*

Proof. For simplicity, let's assume that the fan-in of the circuit is 2, though the result generalizes to higher constant fan-in circuits. We will show that there exists adversarial instances of GRID that this circuit family cannot compute. In particular, we will show that for every sufficiently large N there exists some instance of BIGTRIANGLE on which the circuit C_N is effectively geometrically local. This will violate our result (Theorem 5) that no geometrically local constant-depth circuit can solve the BIGTRIANGLE problem. Specifically, we will show (as illustrated in Figure 8) that

1. We can find three vertices u, v, w in the grid with small light cones. Furthermore, the light cone from each input bit will not touch any of the output bits which are close to the other vertices.
2. There exists a subgraph of the grid that is a triangle with u, v, w as the corner vertices.

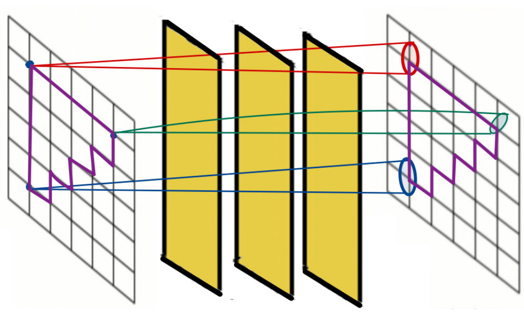

Figure 8: The leftmost grid represents the circuit input (purple: big triangle instance); the center panels in yellow represent the layers of the constant-depth bounded fan-in circuit C_N; and the rightmost panel is the circuit output.

To show that there are u, v, w that satisfy these conditions, we will use the probabilistic method. That is, we will show that if we choose u, v, w randomly according to some process, there will be some non-zero probability for which the conditions hold. This, in turn, implies that there is at least one instance which satisfies the conditions. It remains to show how we go about sampling u, v, w. To start, we label 3 squares of size $N/3 \times N/3$ in the grid as $\mathcal{U}, \mathcal{V}, \mathcal{W}$ as shown in Figure 9.

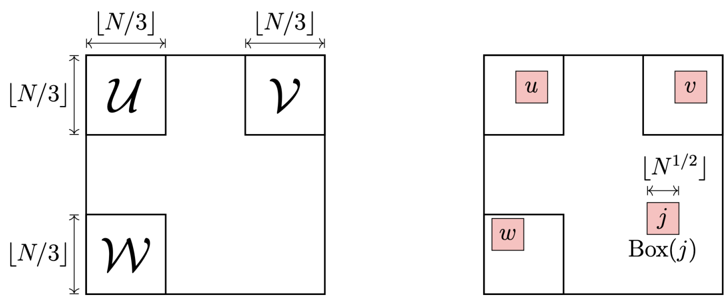

Figure 9: (*left*) the three squares $\mathcal{U}, \mathcal{V}, \mathcal{W}$ in the grid (*right*) the vertices u, v, w and the boxes associated with each vertex. The boxes are supposed to represent output bits of vertices that are near u, v, w [BGK18]

Notice that each region contains $\Omega(N^2)$ vertices. Remarkably, because the circuit is low-depth and has bounded fan-in, one can show each region also contains $\Omega(N^2)$ vertices that correspond to input bits have small light cones.

Lemma 9. *Within each region* $\mathcal{U}, \mathcal{V}, \mathcal{W}$, *there are* $\Omega(N^2)$ *inputs with light cones of size at most* $N^{1/4}$.

Proof. Our goal will be to bound the number of inputs which have light cones larger than $N^{1/4}$. To do this, consider a bipartite graph where one side represents the input bits in a region and the other side represents the output bits for that region. We have an edge between two vertices if the output bit is in the light cone of the input bit.

First notice that each output bit is connected most 2^d input bits, where 2 is the fan-in of the circuit and d is the depth. Since there are $N^2/9$ output bits in each region, this implies an upper bound on the number of edges in the graph:

$$2^d N^2/9 \leq N^{2+\epsilon}$$

where we've used the fact that $d = o(\log N)$ implies $d \leq \epsilon \log N$ for any constant ϵ and sufficiently large N. On the other hand, each input bit which has a light cone larger than $N^{1/4}$ contributes at least $N^{1/4}$ edges to the graph. Combining the above observations, we get

$$(\text{\# of vertices with large light cones}) \cdot N^{1/4} \leq (\text{\# of edges in the graph}) \leq N^{2+\epsilon}$$

which implies that

$$(\text{\# of vertices with large light cones}) \leq N^{2-1/4+\epsilon} = o(N^2).$$

In other words, there are $\Omega(N^2)$ vertices in each region, but only a tiny fraction of them can have a large light cone. □

We now know that each region mostly contains vertices with relatively small light cones. For our probabilistic method argument, let's choose a u, v, w randomly amongst these vertices. We will also draw boxes of size $\sqrt{N} \times \sqrt{N}$ around each of u, v, w (see Figure 9). We are going to show that the light cone of each vertex doesn't contain *any* of the vertices in the other boxes.

To do this, let's first focus on a single randomly chosen vertex u and the box define by the randomly chosen vertex v.

Lemma 10. *The probability that the light cone of* u *intersects the box of* v *is at most 1/10.*

Proof. Notice that every vertex in the grid is part of at most N boxes of size $N^{1/2} \times N^{1/2}$. Since u was chosen randomly amongst vertices which have light cones of size at most $N^{1/4}$, we get that the light cone of u intersects at most $N \cdot N^{1/4} = N^{5/4}$ boxes. On the other hand, since there are $\Omega(N^2)$ possible choices for v, there are also $\Omega(N^2)$ possible boxes that v could have defined. We get that

$$\Pr[\text{light cone of } u \text{ intersects box defined by } v] = O\left(\frac{N^{5/4}}{N^2}\right).$$

For large enough N, this probability is less that $1/10$. □

Of course, there's nothing special about u and v in the argument above. We can apply the same argument to any pair of vertices. Since there are 6 pairs of vertices in the set $\{u, v, w\}$ we get

$$\Pr[\text{light cone of any input bit intersects the box of any another vertex}] < 6/10$$

by the union bound. In other words, by taking the complement, we see that there exists at least one choice of u, v, w for which the light cone of any input bit b_u, b_v, b_w does not intersect the box of any other vertex.

We are finally ready to choose our BIGTRIANGLE instance. Our goal is to find a triangle in the grid with corners u, v, w such that the light cones of each corner bit only intersects with its respective box. We need three observations which are depicted in Figure 10:

(a) There are $N^{1/2}$ disjoint paths between the boxes for u, v, w.

(b) Since u, v, w have small light cones ($\leq N^{1/4}$), one such path does not intersect any of the light cones.

(c) We can find a path from the edges of the box to the chosen vertex within the box.

We have found a BIGTRIANGLE instance for which the light cones of the inputs only contain outputs at distance $\sqrt{N}$. Therefore, invoking Lemma 4, there must be some constant probability for which C_N fails to solve the BIGTRIANGLE instance. This completes the proof. □

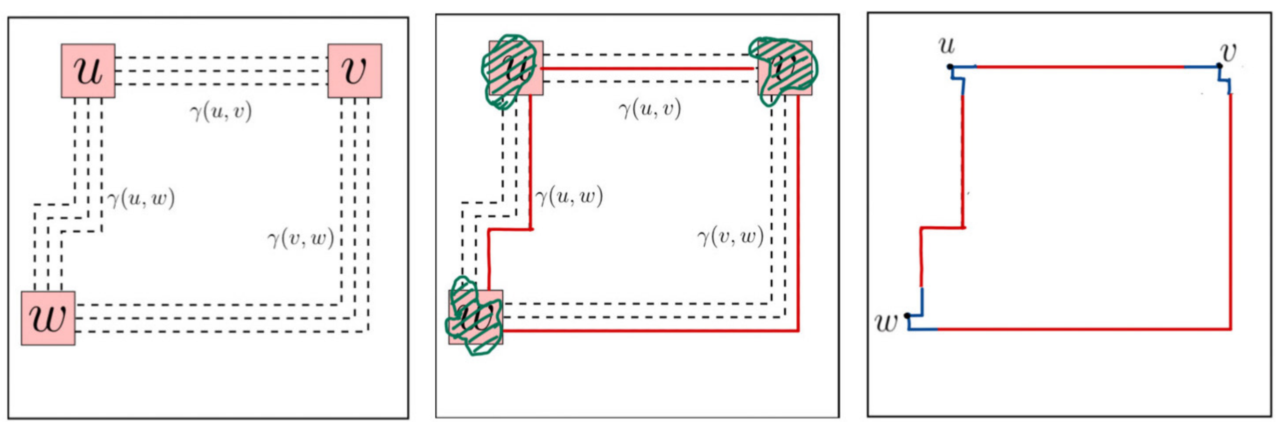

Figure 10: (a) There are $\sqrt{N}$ disjoint paths between each pair of boxes. (b) As the light cones of b_u, b_v, b_w (shown in green) are of size at most $N^{1/4}$, there exists at least one path (drawn in red) between each pair of boxes such that the output bits on the red paths are not contained in the light cones of input bits b_u, b_v, b_w. (c) Finally, we can extend the red paths inside the boxes(drawn in blue) to complete the triangle with u, v, w as vertices.

7 Appendix

Lemma 1. *A set of generators of the stabilizer group of the GHZ state is* $\{XXX, ZZI, IZZ\}$.

Proof. The GHZ state is generated by the following circuit,

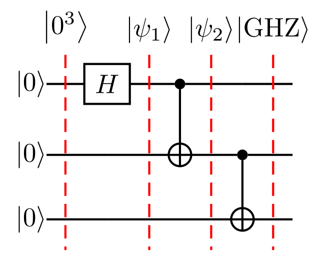

Figure 11: Generating the GHZ state.

By tracing the evolution of the stabilizer group generators we get,

$$\underbrace{\begin{Bmatrix} ZII \\ IZI \\ IIZ \end{Bmatrix}}_{\text{Stab}(|0^3\rangle)} \xrightarrow{H\otimes I\otimes I} \underbrace{\begin{Bmatrix} XII \\ IZI \\ IIZ \end{Bmatrix}}_{\text{Stab}(|\psi_1\rangle)} \xrightarrow{\text{CNOT}\otimes I} \underbrace{\begin{Bmatrix} XXI \\ ZZI \\ IIZ \end{Bmatrix}}_{\text{Stab}(|\psi_2\rangle)} \xrightarrow{I\otimes\text{CNOT}} \underbrace{\begin{Bmatrix} XXX \\ ZZI \\ ZIZ \end{Bmatrix}}_{\text{Stab}(|\text{GHZ}\rangle)}$$

where $\text{Stab}(|\psi\rangle)$ is the stabilizer group of state $|\psi\rangle$, and the listed elements are a set of generators. □

Lemma 2. *Let $|\psi\rangle$ be any Clifford state whose stabilizer group contains a Pauli element that is a tensor product of Z and I elements. That is, $|\psi\rangle$ is stabilized by $P = \alpha P_1 \otimes \cdots \otimes P_n$ such that $P_i \in \{Z, I\}$ and $\alpha = \{\pm 1\}$. Measure $|\psi\rangle$ in the computational basis, but consider only the measurements on qubits i such that $P_i = Z$. If $\alpha = 1$, then the parity of the measurement results is even; otherwise ($\alpha = -1$), the parity is odd.*

Proof. Let a be a bit string $a \in \{0,1\}^n$ and P be a Pauli matrix $P \in \{X, Y, Z, I\}$. We will use the notation P^a to denote the n-qubit Pauli operator

$$P^a := P^{a_1} \otimes \cdots \otimes P^{a_n}$$

where $P^0 = I$. Using this notation and the assumption of the lemma, $|\psi\rangle$ has a stabilizer of the form $(-1)^b Z^z$ where $z \in \{0,1\}^n$ is a bit string and $b \in \{0,1\}$ is the sign (i.e., $(-1)^b Z^z |\psi\rangle = |\psi\rangle$). Notice also that for any $x \in \{0,1\}^n$, we have $|x\rangle = X^x |0\rangle$. This is because applying the Pauli string X^x corresponds to applying the Pauli X operator to qubit i if $x_i = 1$ and the identity operator if $x_i = 0$. Therefore, we have

$$\langle x|\psi\rangle = (-1)^b \langle 0|X^x Z^z|\psi\rangle$$

By (anti-)commutativity property of Pauli strings, we have $X^x Z^z = (-1)^{x\cdot z} \cdot Z^z X^x$, which then gives

$$\langle x|\psi\rangle = (-1)^{x\cdot z \oplus b} \langle 0|Z^z X^x|\psi\rangle = (-1)^{x\cdot z\oplus b}\langle 0|X^x|\psi\rangle = (-1)^{x\cdot z\oplus b}\langle x|\psi\rangle$$

where we've used that Z^a is a stabilizer of the all zeros state: $\langle 0| Z^a = \langle 0|$. Notice that if we want $\langle x|\psi\rangle$ to be nonzero (i.e., there is some chance to output measurement result $|x\rangle$), we need that $x \cdot z \oplus b = 0$; otherwise, we get that $\alpha = -\alpha$ for some non-zero complex number α. In other words, if $b = 0$, the parity of the measurement results on the non-identity elements of the stabilizer must be even; and similarly, if $b = 1$, the parity must be odd. □

References

[AB06] Simon Anders and Hans J Briegel. Fast simulation of stabilizer circuits using a graph-state representation. *Physical Review A*, 73(2):022334, 2006.

[BCE+07] Jonathan Barrett, Carlton M. Caves, Bryan Eastin, Matthew B. Elliott, and Stefano Pironio. Modeling pauli measurements on graph states with nearest-neighbor classical communication. *Phys. Rev. A*, 75:012103, Jan 2007.

[BGK18] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018.

[GHZ89] Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. *Bell's Theorem, Quantum Theory and Conceptions of the Universe*. Springer Netherlands, Dordrecht, 1989.