On the power of interleaved low-depth quantum and classical circuits

by

Abhishek Anand

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science (Quantum Information)

Waterloo, Ontario, Canada, 2022

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

- Chapter 3, 4 and 5 consist of exposition of earlier results.
- Chapter 6 is in collaboration with Prof. Shalev Ben-David and contains new interleaving model definitions, new results and the rephrasing of existing results in this new framework.
- Section 6.2.2 contains a protocol for simulating TC^0 using an untrusted low-depth quantum device which is due to Joseph Slote's observation. It combines a new result from Section 6.2.2 with an existing result of Slote.

Abstract

Low-depth quantum circuits are a well-suited model for near-term quantum devices, given short coherence times and noisy gate operations, making it pivotal to examine their computational power. It was already known as early as 2004 that simulating such low-depth quantum circuits is classically hard under complexity-theoretic assumptions. Later, it was shown that low-depth quantum circuits interleaved with low-depth classical circuits can perform approximate quantum Fourier transform, the quantum subroutine of Shor's algorithm.

Given these salient features of low-depth quantum models, Terhal and DiVincenzo, Aaronson, and Jozsa have all independently conjectured regarding the elusive power of combining low-depth quantum circuits with efficient classical computation. However, much has remained unresolved in this interleaved setting. Therefore, in this thesis, we tackle the question of characterizing the computational power of interleaved low-depth quantum and classical circuits. We first review existing separations in the low-depth setting. Then, we formally define two interleaving models based on whether the quantum device is permitted to make subset measurements (weak interleaving) or must measure all qubits together (strict interleaving).

By combining existing techniques from quantum fan-out constructions, teleportation-based quantum computation, and Clifford + T circuit synthesis, we show several results regarding the power of variants of constant-depth quantum circuits (QNC^0) strictly and weakly interleaved with constant-depth classical parity circuits ($\oplus_\mathsf{f}$). Our main new result is that QNC^0 with access to cat states strictly interleaved with $\oplus_\mathsf{f}$ can simulate constant-depth threshold circuits (TC^0), which neither of the classes can do on their own. This strictly separates this interleaved class from constant-depth classical circuits with unbounded fan-in mod p and OR gates ($\mathsf{AC}^0[p]$).

Acknowledgements

First and foremost, I am grateful to my supervisor, Prof. Shalev Ben-David, for his guidance and support over the past two years. When I joined the program in Fall of 2020, I was stuck back home in India due to travel restrictions. In fact, I had to undertake the first 15 months of the program remotely. Nevertheless, Prof. Shalev took that in stride, and never let me feel that I was missing out on anything.

Through our meetings and his course, Prof. Shalev introduced me to various topics in quantum and classical complexity, and suggested wonderful problems to work on. For our first project, I worked on the approximate degree of functions – something that was conceptually, as well as, mathematically completely new to me. He was extremely patient with me and stewarded me through the project – teaching me both technical and research skills. In particular, he taught me the value of critically and rigorously defining meaningful research questions before beginning to answer them. My second project (which forms the basis of this thesis) began as a project in his course on quantum query and communication complexity. Even though I was embarking on a completely different topic now, Prof. Ben-David encouraged me enthusiastically to chase my intellectual curiosity.

While learning about existing results about low-depth circuits and considering the power of the strictly interleaving model, I often got stuck. Prof. David Gosset, Adam Bene Watts, Daniel Grier, and Luke Schaffer graciously helped me whenever this was the case. They provided numerous insights and background information, and suggested various extensions. It was an immeasurable joy to discuss problems with them.

I also want to thank Joseph Slote for organizing the monthly meetings that bring together people from University of Waterloo, University of Maryland and Caltech interested in low-depth circuits. I had many fruitful discussions with the attendees of the meeting, especially with Natalie Parham, Suchetan Dontha and Prof. Matthew Coudron. Moreover, it was in one of these meetings that Joe noted and explained to me that my strict interleaving protocol can be made into an interactive protocol for TC^0 using his existing result. I am thankful to Joe for letting me write about that observation in this thesis.

Very close to the thesis submission deadline, I had some questions on Clifford + T gate synthesis and concerns regarding how I was using it in my thesis. I emailed Yuan Su, one of the co-authors of the paper I was considering. I am grateful for their responses - Yuan answered my doubts patiently.

I am also thankful to Prof. Richard Cleve and Prof. David Gosset, for agreeing to read my thesis. I am grateful to Rory Soiffer and the other members from the Ben-David group for fun and insightful discussions about all things quantum. Additionally, my roommates at Columbia Lake Village made my (although short) time at Waterloo comfortable and enjoyable.

I would like to thank my high school and undergraduate friends, for their encouragement and support throughout the past two years. Specifically, I thank Shraddha Anand for her feedback on drafts of the thesis that has been extremely useful in enhancing its clarity. Lastly, I would like to thank my sister, my parents, and my family for their eternal support and wisdom.

Dedication

This is dedicated to Bruno, Raji papa and daddyji.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Discovering separations between the computational power of quantum and classical devices is a fundamental problem at the intersection of quantum information science and computational complexity theory. Several tasks have been developed to provide such separations:

- Solving problems such as factoring integers [Sho99] for which efficient quantum algorithms exist but no known efficient classical algorithms do
- Sampling from distributions induced by certain circuit families [BJS11, AA11, BFNV19] and Hamiltonian evolution [GWD17, BVHS+18]
- Performing interactive cryptographic tasks [BCM+21] and tasks based on non-local games [KLVY22]

We seek such separations not only to identify useful tasks for which quantum computers provide a speed-up over their classical counterparts but also because such separations help us verify that untrusted quantum devices such as those currently being built in experimental labs are, in fact, *quantum*. However, current devices have access to a limited number of qubits that have short coherence times and noisy operations. It is unclear whether the above-listed separations can be used to verify such noisy devices that are not fault tolerant.

Due to the short coherence times, low-depth quantum circuits, in which the circuit depth does not (or only very weakly) depend on the size of the task input, may provide a fitting model of computation for such devices. Therefore, to understand the power of near-term devices, it becomes important to examine the computational prowess of low-depth quantum circuits and characterize it with regard to the computational power of corresponding classical circuits. While some of the above-listed separations can be ported to the low-depth setting, they all require reasonable complexity-theoretic assumptions. In fact, the only known unconditional separations in the single-device setting occur while comparing the power of low-depth quantum and classical circuits. This alludes to another fundamental question: how far can *unconditional* separations take us? In the first part of the thesis (Chapters 3, 4 and 5), we will bring together and review all existing unconditional low-depth separations for decision and relational problems.

Even though we can show that low-depth quantum circuits are more powerful than their corresponding low-depth classical circuits, given the low depth of the quantum device, it can be easily seen that the power of these devices has a clear limit in the decision problem setting – they cannot compute functions where the the output bit depends on more than a constant fraction of the input. In the latter half of the thesis (Chapter 6), we will look at whether the power of these devices can be enhanced beyond this limit by providing them with classical resources (in particular, by strictly interleaving quantum circuits with low-depth classical circuits). Given the omnipresence of efficient classical computation, this setting might be experimentally easier to implement compared to quantum circuits with super-constant depth. To motivate this setting further, let us first briefly review existing unconditional low-depth separations:

1.1 Existing low-depth separations

A crucial tool for us are the known strict inclusions between different families of constant-depth classical circuits. This is because a recurring technique for establishing separations is to simulate a classical circuit family using a *weaker* quantum circuit family. We will introduce the required classical and quantum circuit families as well as the strict inclusions in this section as needed. Please refer to Chapter 2 for a complete review.

Quantum circuit complexity theory originated in [Yao93], where it was shown that the circuit model is universal for quantum computation. Development in finding separations can be pedagogically divided into three phases:

- **Chapter 3:** The most fundamental family of classical circuits is NC^i: $(\log(n))^i$-depth circuits with bounded fan-in classical gates and the unbounded fan-out gate. Note that n refers to the input size. One of the first line of inquiry in quantum circuit complexity was regarding how to define the analogous quantum class. QNC^i was first defined in [MN98, MN01] as $(\log(n))^i$-depth circuits with bounded fan-in quantum gates from any universal gate set. The debate was regarding the unbounded fan-out gate – in classical circuits, this gate copies the input bit an unbounded number of times so that it can be subsequently used in multiple operations in parallel. In classical circuits, fan-out is assumed to be "free". However, in the quantum setting, due to the no cloning theorem, such an operator is not possible. [MN98, Moo99, MN01] considered the following *quantum* fan-out operation for single qubit:

$$(\alpha |0\rangle + \beta |1\rangle) |0 \dots 0\rangle \to \alpha |0 \dots 0\rangle + \beta |1 \dots 1\rangle \tag{1.1}$$

  and denoted $\mathsf{QNC}^0_\mathsf{f}$ as QNC^0 with access to unbounded quantum fan-out. In [Moo99], the quantum fan-out gate was shown to be equivalent to the unbounded fan-in quantum parity gate. Note that the quantum version of a classical gate f can be defined by the following action on input state $|x\rangle$:

$$\mathsf{Q}f |x\rangle |0\rangle = |x\rangle |f(x)\rangle \tag{1.2}$$

Moreover, it was shown that all unbounded fan-in quantum mod p gates could be performed in constant depth with access to quantum fan-out. On the other hand, in the classical setting, it is known that $\mathsf{NC}^0[p]$ (NC^0 with access to mod p gates) cannot simulate mod q for $p \neq q$. This, therefore, showed the first separation between two corresponding quantum and classical circuit classes:

$$\mathsf{QNC}^0_\mathsf{f} = \mathsf{QNC}^0[2] \not\subset \mathsf{NC}^0[2] \tag{1.3}$$

Strikingly, the fan-out operation turned out to be really powerful: [HŠ05] showed that $\mathsf{QNC}^0_\mathsf{f}$ can approximate unbounded fan-in quantum OR and threshold gates. Most surprisingly, they showed that $\mathsf{QNC}^0_\mathsf{f}$ can approximate quantum Fourier transform, the *quantum* subroutine of Shor's algorithm. Without fan-out, quantum Fourier transform requires log depth. Later, [TT16] provided exact simulations of OR and threshold gates using $\mathsf{QNC}^0_\mathsf{f}$, showing that:

$$\mathsf{QTC}^0_\mathsf{f} \subseteq \mathsf{QNC}^0_\mathsf{f} \tag{1.4}$$

where TC^0 (QTC^0) are constant-depth circuits with (quantum) NOT and unbounded fan-in (quantum) threshold gates. It is known that $\mathsf{NC}^0 \subsetneq \mathsf{AC}^0 \subsetneq \mathsf{AC}^0[p] \subsetneq \mathsf{TC}^0$ where AC^0 are constant-depth circuits with NOT and unbounded fan-in OR gates. This provided an even stronger separation:

$$\mathsf{QNC}^0_\mathsf{f} \not\subset \mathsf{AC}^0[p] \tag{1.5}$$

- **Chapter 4:** [GC99, RB01] laid the foundations for a different universal paradigm of quantum computation: measurement-based quantum computation. Given a quantum circuit to be simulated, here, an entangled resource state is prepared in constant depth. This is followed by adaptive one or two-qubit measurements - whose basis depends on the gates in the circuit and prior measurement outcomes. In [GC99, RBB03], it was shown that any Clifford operation can be simulated in one measurement step, that is, by applying only one layer of gates, followed by one layer of computational basis measurements, followed by one layer of classically-controlled (on the measurement outcomes) single-qubit gates. Given that the quantum fan-out gate is a Clifford operation, and that $\mathsf{QNC}^0_\mathsf{f}$ consists of only a constant number of layers of quantum fan-out gates, [BKP10] noted that we can replace each one of them with the above procedure. This provides another, perhaps cleaner, and more experimentally feasible way to look at the above results: adaptive constant-depth quantum circuits can simulate $\mathsf{QTC}^0_\mathsf{f}$ and perform approximate quantum Fourier transform.

- **Chapter 5:** More recently, a separation was shown for QNC^0 without requiring the quantum fan-out gate or adaptivity, albeit in the relational setting. [BGK18] construct the two-dimensional hidden linear function problem (2D HLF) and show that it can be solved by a QNC^0 circuit but no NC^0 circuit can solve it. This was done by exploiting the non-local relations that exist between the inputs and outputs of 2D HLF problem. These non-local relations were shown to be unrealizable by a constant-depth classical circuit based on ideas from non-local games. Moreover,

using modified lower-bound techniques from classical complexity theory, [WKST19] improved this by showing that no circuit from AC^0 can solve the problem:

$$\mathsf{QNC}^0 \not\subset \mathsf{AC}^0 \tag{1.6}$$

Additionally, average-case results and results taking into account certain noise models were also reached in [Gal18, BGKT20, CSV21, HG21a].

In this thesis, we will first briefly survey the tools and techniques behind these three lines of results.

1.2 Towards new interleaving models and separations

In summary, we have:

Model	Power
QNC^0	$\not\subset \mathsf{AC}^0$

adaptive-$\mathsf{QNC}^0 \supseteq \mathsf{QNC}^0_\mathsf{f}$	$\not\subset \mathsf{AC}^0[p]$ $= \mathsf{QTC}^0_\mathsf{f}$ $=$ poly-size quantum circuits with constant T-depth (including approximate QFT)

Table 1.1: Low-depth quantum circuit models and their power

The following question remains: how *realistic* are the latter two models? The quantum fan-out gate, as discussed in Section 3.5, requires log depth to be implemented using bounded fan-in gates. Therefore, we require constant-time device-dependent (native) implementations for quantum fan-out if there is any hope for $\mathsf{QNC}^0_\mathsf{f}$ to be considered a realistic model for near-term devices. Proposals for such implementations are discussed in Section 3.5.3. Remarkably though, we get the same power by using adaptive constant-depth quantum circuits. In particular, any poly-size Clifford circuit and the approximate quantum Fourier transform can be implemented by constant-depth quantum circuits with access to low-depth classical circuits. A nice feature of such hybrid models is that they cleanly separate out the *quantum* subroutine of the algorithm. Considering such parallelization results, Jozsa conjectured [Joz06] that all problems in BQP can be solved by log-depth quantum circuits interleaved with polynomial-time classical computation:

Conjecture 1.1 (from [Joz06])**.** $\mathsf{BQP} = \left(\mathsf{QNC}^1\right)^{\mathsf{BPP}}$

Similarly, hinting at the unresolved power of such interlveaing models, Aaronson conjectured [Aar05, Aar10, Aar11, Aar14] in the opposing direction:

Conjecture 1.2 (from [Aar05])**.** *There is an oracle that separates* BQP *from* $\mathsf{BPP}^{\mathsf{QNC}^1}$.

Recently, some progress has been made on this conjectures – Aaronson's conjecture has been proven, which provides evidence against Jozsa's conjecture [CM20, CCL20, AGS22, CH22, HG22]. However, the question remains unresolved in the non-relativized setting. Moreover, this was not the first time that the power of such interleaved models was considered. In 2004, [TD02] showed that under complexity-theoretical assumptions, constant-depth circuits cannot be efficiently exactly simulated by classical devices. Noting this, they wondered:

Question 1.3 (from [TD02]). *Is there a problem in* $\mathsf{BPP}^{\mathsf{QNC}^0}$ *that is not in* BPP*? That is, is* $\mathsf{BPP}^{\mathsf{QNC}^0} \not\subset \mathsf{BPP}$*?*

In the latter part of the thesis, we examine the power and experimental feasibility of interleaved models where both the quantum and classical part is low depth. Our work is directed towards resolving the constant-depth variants of the above conjectures and questions.

1.2.1 Weak interleaving

In Section 6.1, we consider weakly interleaved models and their power. We first introduce notation: $[\mathsf{Q};\mathsf{C}]^d$ refers to the model where circuits from Q can measure a subset of qubits and call circuits from C on the measurement outcomes, d times. We call this *weak* interleaving.

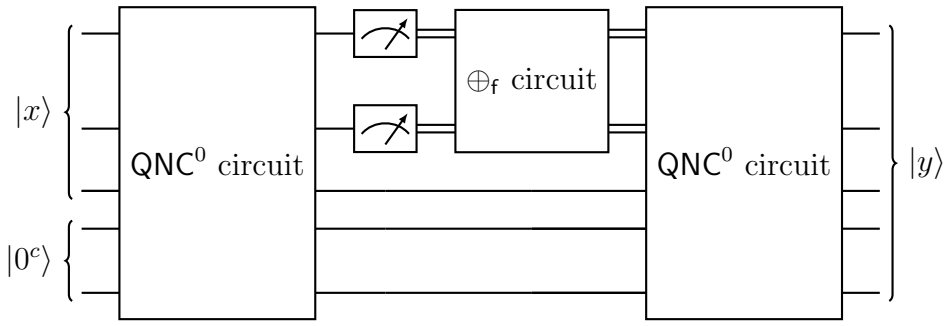

Figure 1.1: $[\mathsf{QNC}^0;\oplus_\mathsf{f}]^2$: weakly interleaving QNC^0 with $\oplus_\mathsf{f}$ circuits

We observe some basic properties of these models and then rephrase existing adaptive-QNC^0 results in this setting. Our starting point, coming from measurement-based quantum computation, will be that any poly-size Clifford circuit can be simulated in $[\mathsf{QNC}^0;\oplus_\mathsf{f}]^{O(1)}$ where $\oplus_\mathsf{f}$ refer to constant-depth circuits with classical parity and classical fan-out gates. Now, as $\mathsf{QTC}^0 \subseteq \mathsf{QNC}^0_\mathsf{f}$ and because quantum fan-out is Clifford, we can rephrase Equation 1.4 as:

$$\mathsf{QTC}^0 \subseteq [\mathsf{QNC}^0;\oplus_\mathsf{f}]^{O(1)} \tag{1.7}$$

Generalizing this to arbitrary poly-size quantum circuits, we have:

Theorem 1.4 (Informal)**.** *Any* poly*-size quantum circuit with* T*-depth* t *can be simulated in* $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(t)}$.

Note that recently several *quantum* tasks have been shown to be implementable using adaptive-QNC^0 ($[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$) circuits such as tests of quantumness [HG21b, LG21] and generation of long-range entangled states [LLKH22, TVV22]. These results eventually rely on the task having constant T-depth. We discuss such applications further in Section 6.1.2.

1.2.2 Strict interleaving

In the weakly interleaved model, note that we are required to measure a subset of qubits and maintain coherence in the remaining qubits while the classical circuit performs its operations – equivalently, we require the quantum device to have quantum memory to store the state of the unmeasured qubits. Such mid-circuit measurements enhance the experimental complexity of the setup and are out of reach of most current experiments [GDC+20].

One may therefore ask how much we need to "enhance" the powers of QNC^0 or "reduce" the powers of $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$ to simulate TC^0 and beyond. In other words:

Question 1.5. *Do there exist constant-depth models that are experimentally feasible but also provide us with separations beyond* AC^0*?*

One straight-forward way to improve experimental feasibility is to drop the quantum memory requirement – before interleaving with a classical circuit, we must measure all the qubits. That is, the interleaving is *strict*. In Section 6.2, we formally define $[\mathsf{Q}, \mathsf{C}]^d$ as this model where we strictly interleave quantum circuits from Q and classical circuits from C, d times and consider its power.

In this setting, instead of letting $\mathsf{Q} = \mathsf{QNC}^0$, we afford QNC^0 an extra power: the ability to generate cat states. That is, before looking at the input, each instance of the constant-depth quantum circuit can generate cat states whose size only depends on the input size. We require these cat states to parallelize subroutines in a manner similar to the fan-out constructions. Intuitively, if we want to simulate beyond NC^0, we need a way for each input to interact with all other inputs. In our constructions, we enable this by letting each input interact with a qubit of the cat state. Therefore, the cat state, as a whole, interacts with all inputs. We then use measurements and unbounded fan-in classical gates to reduce the the cat state down to a single bit to solve the decision problem. In the real world, we can imagine a centralized quantum device optimized for generating and storing *only* cat states – and thereby, being experimentally easier to implement than a device that can prepare any log-depth state and store it coherently. However, note that having access to cat states is strictly weaker than having access to quantum fan-out gates as discussed in Section 6.2.2.

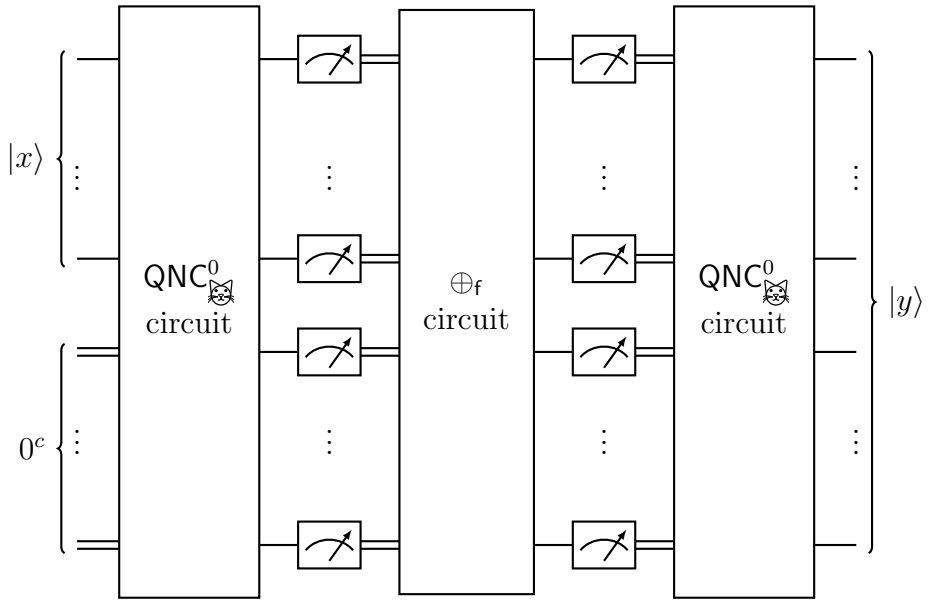

Figure 1.2: $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^2$: strictly interleaving $\mathsf{QNC}^0_{\text{cat}}$ with $\oplus_{\mathsf{f}}$ circuits

To pin down the utility of the interleaving in this hybrid circuit model, we can further ask whether there are separations that necessitate interleaving. More precisely:

Question 1.6. *Is there a decision problem that can be solved by strictly interleaving a quantum circuit class and a classical circuit class which cannot be solved by the two classes individually?*

For the strictly interleaved model we are considering, this question can also be phrased as the low-depth version of Question 1.3 from [TD02]:

Question 1.7. *Is a problem in* $(\oplus_{\mathsf{f}})^{\mathsf{QNC}^0_{\text{cat}}}$ *that is neither in* $\oplus_{\mathsf{f}}$ *nor in* $\mathsf{QNC}^0_{\text{cat}}$*?*

In fact, already in [WKST19], as a step towards separations beyond AC^0, the first example of such a problem was shown by demonstrating that if quantum circuit has access to a single cat state, it can compute the relational version of mod 3 with high probability. Therefore, showing that:

$$\mathsf{QNC}^0_{\text{cat}} \not\subset \mathsf{AC}^0[2] \tag{1.8}$$

We can very simply restate this as a decision separation by adding a layer of Hadamards, a measurement and a classical parity gate at the end of the circuit from [WKST19] to get that:

$$[\mathsf{QNC}^0_{\text{cat}}, \oplus]^2 \not\subset \mathsf{AC}^0[2] \tag{1.9}$$

It follows that this is an answer to Question 1.7 as QNC^0 circuits with arbitrary advice cannot compute mod 3. Using techniques from the fan-out constructions, we strengthen

this result and provide a harder problem with an affirmative answer to Question 1.6. We show that the unbounded fan-in OR and threshold gate can be computed *exactly* using $[\mathsf{QNC}^0_{\not\subset}, \oplus_\mathsf{f}]^2$. Therefore, using this model, we can simulate circuits from TC^0 exactly:

Theorem 1.8 (Informal)**.** $\mathsf{TC}^0 \subseteq [\mathsf{QNC}^0_{\not\subset}, \oplus_\mathsf{f}]^{O(1)}$.

One way to interpret our results is as follows – we have replaced the fan-out gate with a "one-way fan-out gate" and still retained the power to simulate TC^0. The "one-way fan-out gate" has the ability to generate cat states and compute classical parity but not the ability to "destroy" the cat state back to a single qubit.

Moreover, we combine this with a recent result by Slote [Slo22] that shows that there is an interactive proof for TC^0 with a TC^0 prover and AC^0 verifier. Their technique can be adapted to the above result to show that:

$$\mathsf{TC}^0 \subseteq \mathsf{IP}[\mathsf{QNC}^0_{\not\subset}, \mathsf{AC}^0[2]] \tag{1.10}$$

thereby providing an unconditional delegated computation protocol for TC^0 with an untrusted low-depth quantum device which can be passed without requiring the quantum device to make subset measurements.

1.3 Summary of results

The following table summarizes the two existing discussed in this thesis in the new framework and how our new results fits in:

Model	Power
QNC^0	$\not\subset \mathsf{AC}^0$
$[\mathsf{QNC}^0_{\not\subset}, \oplus_\mathsf{f}]^{O(1)}$	$\not\subset \mathsf{AC}^0[p]$ $\supseteq \mathsf{TC}^0$
$[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)} = \mathsf{QNC}^0_\mathsf{f}$	$\not\subset \mathsf{AC}^0[p]$ $= \mathsf{QTC}^0_\mathsf{f}$ $=$ poly-size quantum circuits with constant T-depth (including approximate QFT)

Table 1.2: Interleaved low-depth quantum-classical circuit models and their power. $\oplus_f$ is the family of constant-depth classical circuits with parity and fan-out gates. $[\mathsf{QNC}^0_{\not\subset}, \oplus_\mathsf{f}]^{O(1)}$ refers to $\mathsf{QNC}^0_{\not\subset}$ *strictly* interleaved with $\oplus_\mathsf{f}$ circuits, a constant number of times. $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$ refers to QNC^0 circuits *weakly* interleaved with $\oplus_\mathsf{f}$ circuits, a constant number of times.

The thesis ends with a discussion regarding open questions and future topics.

Chapter 2

Mathematical preliminaries

Classical and quantum circuits

Basic notation: $[n]$ for $n \in \mathbb{N}$ refers to the set $\{1, 2, \ldots, n-1, n\}$. For $x \in \{0,1\}$, $\bar{x} = 1 - x$. For $x \in \mathbb{R}$, $\log(x)$ refers to logarithm base 2. $\log^*(x)$ is the maximum number of iterated logarithms that can be applied to x such that the output exists and is real.

Classical circuits: A classical circuit on n bits applies a series of gates on input state $x \in \{0,1\}^n$. The NOT has the following action on single bit x_i for $i \in [n]$: $\mathsf{NOT}(x_i) = \bar{x}_i$. We will be using the following mult-bit classical gates:

Definition 2.1 (Multi-bit classical gates)**.** *The following table defines the operation of classical gate g on subset $s \subseteq [n]$ of x, $x_{[s]}$ for $|s| = a$:*

Classical gate g	$g(x_{[s]}) = 1$ **iff**
OR_a	$\lvert x_{[s]}\rvert \geq 1$
$\mathsf{MOD}_{p,a}$ (mod p gate)	$\lvert x_{[s]}\rvert \mod p \neq 0$
$\mathsf{EX}_{t,a}$ (exact gate)	$\lvert x_{[s]}\rvert = t$
$\mathsf{T}_{t,a}$ (threshold gate)	$\lvert x_{[s]}\rvert \geq t$

The output of the circuit is the final bit string (or its subset) after all gates have been applied. A simple universal gate set is $\{\mathsf{NOT} \circ \mathsf{OR}_2\}$.

Fourier analysis of functions: Every function in the $\{\pm\}$ basis has an expansion in terms of only the parity of subsets of the input. This is called the *Fourier expansion* of the function. In particular, we have that for $f : \{\pm 1\}^n \to \{\pm 1\}$ and for $x \in \{\pm 1\}^n$:

$$f(x) = \sum_{s \subseteq [n]} c_s \prod_{i \in s} x_i \tag{2.1}$$

where $\{c_s\}$ are the Fourier coefficients of f. Note that $\prod_{i \in s} x_i$ precisely computes parity over subset $s \subseteq [n]$ in the $\{\pm\}$ basis. Refer to O'Donnell's excellent book on the analysis of Boolean functions [O'D14] to learn more.

Quantum circuits: Analogous to classical circuits, in a quantum circuit on $n \in \mathbb{N}$ qubits, we apply quantum gates on a quantum state, which is a norm-1 vector of the 2^n-dimensional Hilbert space $\mathcal{H}$. Quantum gates on a qubits are $2^a \times 2^a$ unitary operations (operators U such that $U^\dagger U = U^\dagger U = \mathsf{I}$) and hence, are always reversible. We will be using the following Pauli quantum gates:

$$\mathsf{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \ \mathsf{Z} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \ \mathsf{Y} = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \tag{2.2}$$

and their corresponding rotation operators:

$$\mathsf{R}_\mathsf{X}(\theta) = \exp\left(-i\mathsf{X}\frac{\theta}{2}\right) = \begin{bmatrix} \cos\left(\frac{\theta}{2}\right) & -i\sin\left(\frac{\theta}{2}\right) \\ -i\sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \end{bmatrix} \tag{2.3}$$

$$\mathsf{R}_\mathsf{Z}(\theta) = \exp\left(-i\mathsf{Z}\frac{\theta}{2}\right) = \begin{bmatrix} \exp\left(-i\frac{\theta}{2}\right) & 0 \\ 0 & \exp\left(i\frac{\theta}{2}\right) \end{bmatrix} \tag{2.4}$$

Moreover, we will be using Clifford gates:

$$\mathsf{H} = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \ \mathsf{S} = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} = \sqrt{\mathsf{Z}}, \ \mathsf{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{2.5}$$

and the T gate :

$$\mathsf{T} = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{bmatrix} \tag{2.6}$$

as Clifford + T forms a simple universal gate set. The controlled-U gate for some gate U acts on two qubits and applies U to the second register (target) if and only if the first register (control) is in state $|1\rangle$. We will also be using the following fact from linear algebra: if two unitary matrices commute with each other, then they are simultaneously diagonalizable. That is, they share eigenvectors.

To reach a classical output from a quantum circuit, we measure qubits in the computational basis once all gates have been applied. Note that the classical gates in Definition 2.1 are irreversible. However, we can define their reversible versions (denoted by Q preceding the classical gate name). That is, on input $x \in \{0,1\}^n$ and auxiliary qubit state $|s\rangle$ with $s \in \{0,1\}$, for classical gate g acting on subset $x_{[s]}$ of x ($s \subseteq [n]$):

$$\mathsf{Q}g\,|x\rangle\,|s\rangle = |x\rangle\,\big|s \oplus g(x_{[s]})\big\rangle \tag{2.7}$$

Pauli and Clifford groups: The Pauli group on n qubits, $\mathbb{P}_n$ is the group of $2^n \times 2^n$ unitary matrices generated by the tensor product of Pauli gates and the identity. The Clifford group on n qubits is defined as the set of $2^n \times 2^n$ unitary operators that commute with the Pauli group (or in other words, normalize the Pauli group). That is, a unitary U belongs to the Clifford group on n qubits if and only if for all $P \in \mathbb{P}_n$, we have that:

$$UPU^\dagger = P' \tag{2.8}$$

and $P' \in \mathbb{P}_n$. The Clifford gates are the generators of the Clifford group. Moreover, the Gottesman–Knill theorem [Got98] says that the action of an element from the Clifford group on computational basis states can be efficiently simulated classically.

Circuit complexity classes

A classical (or quantum) circuit C consists of input bits (or qubits), some gates and output bits (qubits). Consider a graph $G(C)$ where the vertices are the input and output nodes. There is an edge between an input and output node if they lie on the same wire on the circuit. We also add edges between two vertices if they interact with each other i.e. when a gate acts on both vertices. Circuit depth is defined as follows:

Definition 2.2 (Circuit depth)**.** *Circuit depth of a circuit C is the maximum path length from an input node to an output node in $G(C)$.*

Equivalently, we can also think of circuit depth as the number of layers of gates in a circuit such that all gates in each layer act on disjoint sets of bits (qubits). In this sense, circuit depth can be used as a proxy for computation time of running the circuit if we have the ability to perform gates on disjoint sets of qubits simultaneously. For intuition, consider the following ladder circuit:

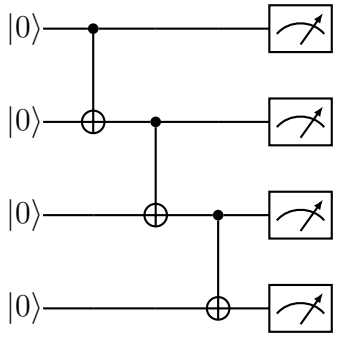

Figure 2.1: Ladder circuit demonstrating linear depth

Even though for each input qubit, we apply single gate and then measure the qubit, the depth of this circuit is 3 and in general, linear in the input size. The third CNOT can only be applied after the second is applied and so on.

Let us examine the different constant-depth classical circuit classes that we shall be comparing the quantum classes to:

- $\mathsf{NC}^i : O\left((\log n)^i\right)$-depth poly-size circuits with bounded fan-in OR and NOT gates
- $\mathsf{AC}^i : O\left((\log n)^i\right)$-depth poly-size circuits with unbounded fan-in OR and NOT gates
- $\mathsf{AC}^0[p]$: AC^0 with unbounded fan-in mod p gate for prime p
- $\mathsf{ACC}^0 = \bigcup_i \mathsf{AC}^0[i]$
- $\mathsf{TC}^i : O\left((\log n)^i\right)$-depth poly-size circuits with unbounded fan-in threshold and NOT gates

Note that we will be talking about these classes either as a set of circuits (or Turing machines) or interchangeably, as the problems computable by the respective circuits (or Turing machines). For decision version of this classes, the following inclusions and subset relations are known:

$$\mathsf{NC}^0 \subsetneq \mathsf{AC}^0 \subsetneq \mathsf{AC}^0[p] \subsetneq \mathsf{ACC}^0 \subseteq \mathsf{TC}^0 \subseteq \mathsf{NC}^1 \tag{2.9}$$

[Vol99] contains proofs for most of these relations. Additionally, refer to the wonderful Complexity Zoo [Com] to learn more about these classes and their properties.

Computational tasks

We will be considering different computational tasks to be performed by the interleaved low-depth quantum-classical circuits:

- **Decision problem:** The task is defined by a function $f : \{0,1\}^* \to \{0,1\}$. For given input $x \in \{0,1\}^n$, we must return $f(x)$.
- **Relational problem:** The task is defined by a function $f : (\{0,1\}^* \times \{0,1\}^*) \to \{0,1\}$. For given input $x \in \{0,1\}^n$, we must return a string $y \in \{0,1\}^m$ such that $f(x,y) = 1$.
- **Circuit simulation:** Given a classical (quantum) circuit C (Q) on n bits (qubits) and an input state $x \in \{0,1\}^n$ ($|\psi_x\rangle$, a 1-norm vector in $\mathcal{H}^{2^n}$), we must return $y \in \{0,1\}^n$ ($|\psi_y\rangle$, a 1-norm vector in $\mathcal{H}^{2^n}$) such that $y = Cx$ ($|\psi_y\rangle = Q\,|\psi_x\rangle$).
- **Quantum circuit sampling:** Given a quantum circuit Q on n qubits and an input state $|\psi_x\rangle$, a 1-norm vector in $\mathcal{H}^{2^n}$, we must return $y \in \{0,1\}^n$ at random according to the distribution over bit strings defined by $Q\,|\psi_x\rangle$.

Chapter 3

Constant-depth quantum circuits with unbounded quantum fan-out

In classical circuits, the unbounded fan-out gate generates a polynomial number of copies of the input bit, and it is assumed to be "free" even in NC^0. Fan-out is crucial in constructing low-depth classical circuits for various tasks as it parallelizes subroutines that act on the same bit. However, a quantum analogue of this operation cannot exist as no unitary operation can copy the state of a qubit due to the no-cloning theorem. One way to increase the number of operations that can act on one qubit simultaneously is to first entangle the qubit with a set of auxiliary qubits and then perform operations on each of them separately at the same time. [MN98, MN01] first formally defined this operation as the unbounded quantum fan-out gate:

Definition 3.1 (Quantum fan-out gate)**.** *The quantum fan-out gate of size $a \in O(\mathsf{poly}(n))$ (where n is the input size), F_a, is defined by the following action on data qubit in state $\alpha|0\rangle + \beta|1\rangle$ and auxiliary qubit basis state $|x\rangle$, with $x \in \{0,1\}^a$:*

$$\mathsf{F}_a\left((\alpha|0\rangle + \beta|1\rangle)|x\rangle\right) = \alpha|0\rangle|x\rangle + \beta|1\rangle|\bar{x}\rangle \tag{3.1}$$

[MN98, Moo99, MN01] constructed the first parallelization techniques enabled by the unbounded quantum fan-out gate and defined circuit classes in which quantum fan-out was a primitive. [HŠ05, TT16] further expanded on these techniques to show that constant-depth quantum circuits with quantum fan-out gates, i.e. the class $\mathsf{QNC}^0_\mathsf{f}$, is much more potent than its classical counterpart: $\mathsf{QNC}^0_\mathsf{f}$ can simulate all decision problems in TC^0, strictly separating it from $\mathsf{AC}^0[p]$. In this chapter, we will build toward this result.

3.1 Parallelization lemma

This section discusses the critical construction using quantum fan-out gates that lets us parallelize certain operations to constant depth. The subsequent observations and lemmas,

along with the gadgets used to derive them were first shown in [MN98, MN01]. Consider circuits with the following structure:

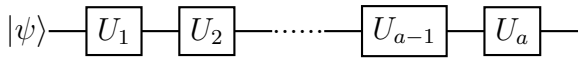

where $\{U_i\}_{i\in[a]}$ are mutually commuting unitary operators, each implementable in QNC^0. Using quantum fan-out, we will be able to generate the output state of the circuit $U_aU_{a-1}\dots U_2U_1\,|\psi\rangle$ in constant depth. The intuition behind the lemma is that as these operators commute, the circuit has no temporal structure. Therefore, if it were not for them acting on the same set of qubits, we could apply them simultaneously in a single layer.

Lemma 3.2. *Let $\{U_i\}_{i\in[a]}$ be mutually commuting unitary operations for $a \in \mathsf{poly}(n)$ with each operator implementable in QNC^0. Then, we can simulate the circuit that applies them to input state $|\psi\rangle$ using a $\mathsf{QNC}^0_\mathsf{f}$ circuit.*

Proof. As $\{U_i\}_{i\in[a]}$ mutually commute, they are simultaneously diagonalizable. Let $U_i = ED_iE^\dagger$ for all $i \in [a]$ where $\{D_i\}_{i\in[a]}$ are the simultaneously diagonalized matrices while E consists of the simultaneous eigenvectors of $\{U_i\}_{i\in[a]}$. Now, as $EE^\dagger = I$, using this eigendecomposition, we have that:

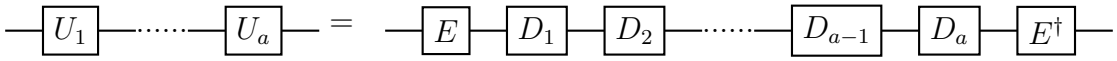

Now for some basis state $|x\rangle$, let the state after gate E be $\sum_y \alpha_y\,|y\rangle$. Then, the final state is given by $E^\dagger\left(\sum_y \prod_i D_{i,y}\alpha_y\,|y\rangle\right)$ where $D_{i,y}$ refers to the element of the D_i diagonal operator associated with the $|y\rangle$ basis state. Using quantum fan-out of size a on each qubit of the initial state $|\psi\rangle$, consider the following construction:

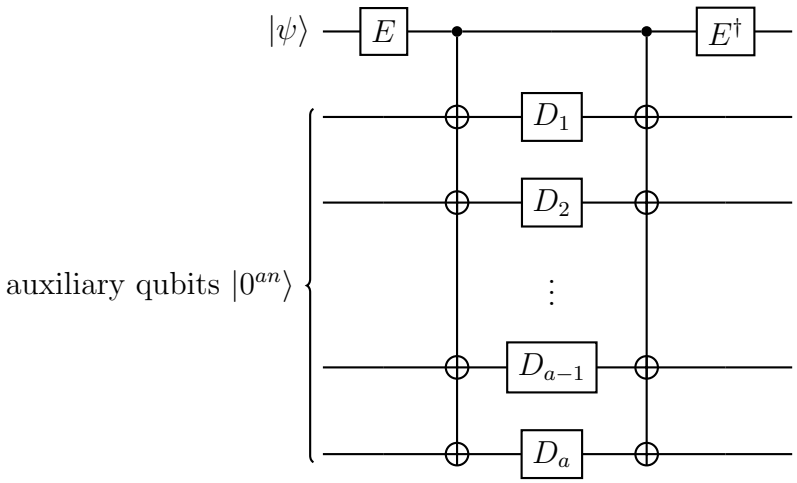

Figure 3.1: Parallelizing commuting gates using quantum fan-out

Consider the action of this circuit on a basis state $|x\rangle$. After the quantum fan-out, we are in state $\sum_y \alpha_i |y^{a+1}\rangle$. Note that every wire of the auxiliary qubits represents n qubits (a "copy" of $\sum_y \alpha_i |y\rangle$). After application of the diagonal D gates, we are in state $\sum_y \left(\prod_i D_{i,y}\alpha_y |y^{a+1}\rangle\right)$. Now using the fact that the quantum fan-out gate is its own inverse, after the second quantum fan-out, we are in state $\left(\sum_y \prod_i D_{i,y}\alpha_y |y\rangle\right)|0^{an}\rangle$. Note that for the second quantum fan-out to work cleanly, it was crucial that the D operators were diagonal and only impacted the coefficients of each $|y\rangle$. Applying $E^\dagger$ gives us the desired output state. □

In the above construction, the quantum fan-out gate can be replaced by any gate that constructs a cat state based on the data qubit when the auxiliary qubits are in the all-zero state. Therefore, another way to state the capability of the quantum fan-out gate in this construction is that it can generate and degenerate cat states. To examine how important it is to have *both* these capabilities, in Chapter 6, we will consider models where cat states are provided as auxiliary qubits but one cannot reduce them back into a single qubit in constant depth.

3.2 Simulating $\mathsf{QAC}^0[2]$

QNC^0 circuits can use bounded fan-in Toffoli gates to simulate the action of reversible versions of NOT and bounded fan-in OR gates. Now with access to the quantum fan-out gates, we can also simulate classical fan-out. Therefore, we have that $\mathsf{NC}^0 \subseteq \mathsf{QNC}^0_\mathsf{f}$.

In order to simulate AC^0 circuits, we will need to simulate unbounded fan-in OR. In [HŠ05], Høyer and Špalek provided a $\mathsf{QNC}^0_\mathsf{f}$ implementation that computes unbounded fan-in OR approximately and accompanied this with a $\log^*$-depth exact implementation. More recently, Takahashi and Tani [TT16] made the $\mathsf{QNC}^0_\mathsf{f}$ construction exact by exploiting the Fourier expansion of the OR function and using the fact that $\mathsf{QNC}^0_\mathsf{f}$ can compute unbounded fan-in parity.

We begin by showing that unbounded fan-in quantum parity and unbounded quantum fan-out are equivalent, which was first observed in [Moo99].

3.2.1 Implementing unbounded fan-in quantum parity

Consider the circuit computing $\mathsf{QMOD}_{2,\mathsf{a}}$ and F_a in linear depth using CNOT gates:

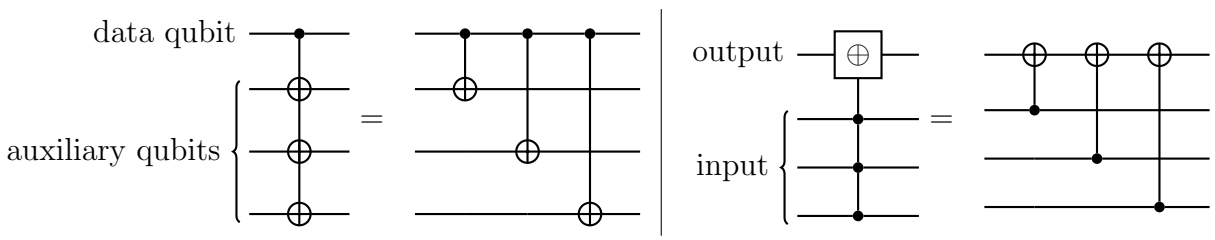

Figure 3.2: Linear-depth implementation of F_3 and $\mathsf{QMOD}_{2,3}$

Now, we use the fact that sandwiching a CNOT gate with a layer of Hadamard gates on either side switches the target and control qubit to see the relation between F_3 and $\mathsf{QMOD}_{2,3}$:

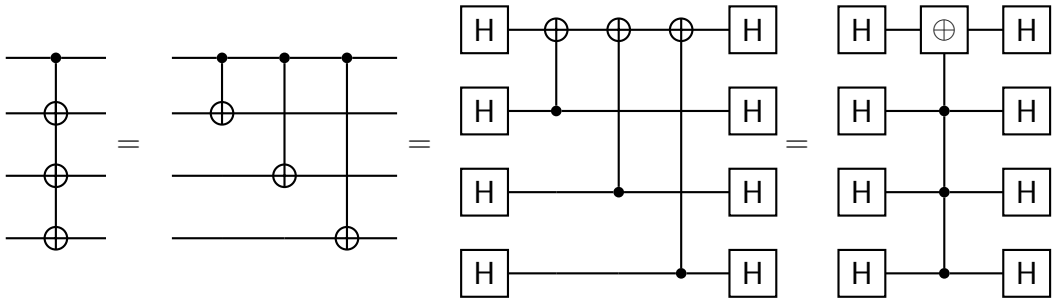

Figure 3.3: Equivalence of F_3 and $\mathsf{QMOD}_{2,3}$

This gives us the following:

Observation 3.3. *In any model with access to Hadamard gates, unbounded quantum fan-out gate and unbounded fan-in* QMOD_2 *gates are depth-equivalent:*

$$\mathsf{H}^{\otimes(a+1)}\mathsf{F}_a\mathsf{H}^{\otimes(a+1)} = \mathsf{QMOD}_{2,\,\mathsf{a}} \tag{3.2}$$

We have another characterization of $\mathsf{QNC}^0_\mathsf{f}$ as $\mathsf{QNC}^0_\mathsf{f} = \mathsf{QNC}^0[2]$.

3.2.2 Implementing unbounded fan-in QOR

Implementing unbounded fan-in QOR is the trickiest problem in this series of results. It was already shown in [HŠ05] that if we could do this, we could also compute unbounded fan-in quantum threshold gates. Only later in 2012 did Takahashi and Tani discover the approach described in this section [TT16]. To compute QOR on $a \in \mathsf{poly}(n)$ qubits, a protocol from [HŠ05] is first used to reduce the problem to computing QOR on $\log(a)$ qubits. This is then combined with a circuit from [TT16] that computes QOR using a constant-depth but exponential-size circuit (as it uses exponentially-large quantum fan-out gates). The circuit construction exploited the Fourier expansion of the OR function. [HCLB11] had first introduced this approach of using a function's expansion to construct a exponential-size quantum circuit.

Step 1: Reducing to log size

Lemma 3.4. *Computing* OR_a *can be reduced to computing* $\mathsf{OR}_{\lceil\log(a)\rceil}$ *for* $a \in \mathsf{poly}(n)$ *in* $\mathsf{QNC}^0_\mathsf{f}$.

Proof. Let $|x\rangle$ with $x \in \{0,1\}^a$ be a basis state on which we want to compute OR_a. Consider the following $m+1$ test quantum states for $m = \lceil \log(a) \rceil$:

$$\forall 0 \leq i \leq m,\ |t_i\rangle \propto |0\rangle + \exp\left(i\pi \frac{|x|}{2^i}\right)|1\rangle \tag{3.3}$$

Note that if $|x| = 0$, $|t_i\rangle \propto |0\rangle + |1\rangle$ for all i and hence, $(\mathsf{H})^{\otimes(m+1)}|t\rangle = |0^{m+1}\rangle$ for $|t\rangle = \prod_i |t_i\rangle$. Therefore, for this case: $\mathsf{OR}_a|x\rangle = \mathsf{OR}_m(\mathsf{H})^{\otimes(m+1)}|t\rangle$. Otherwise, we have that $0 < |x| \leq a$. Consider the unique prime factorization of $|x|$ as $2^b \times$ (some odd natural number) for $b \in \mathbb{N}$. Note that as $|x| \leq a$, it must be that $0 \leq b \leq \log(a)$. Therefore:

$$|t_b\rangle \propto |0\rangle + \exp\left(i\pi \frac{2^b \times (\text{some odd natural number})}{2^b}\right)|1\rangle \tag{3.4}$$

$$= |0\rangle - |1\rangle \tag{3.5}$$

Therefore, $\mathsf{H}|t_b\rangle = |1\rangle$ and hence, each individual basis state in $(\mathsf{H})^{\otimes m}|t\rangle$ has at least one $|1\rangle$ and we have that

$$\mathsf{OR}_a|x\rangle = \mathsf{OR}_{m+1}(\mathsf{H})^{\otimes(m+1)}|t\rangle \tag{3.6}$$

in all cases. Now all that is left to show is that $|t\rangle$ can be prepared in constant depth. Consider the following $\mathsf{poly}(a)$-depth circuit that computes $|t\rangle$ for $a = 4$:

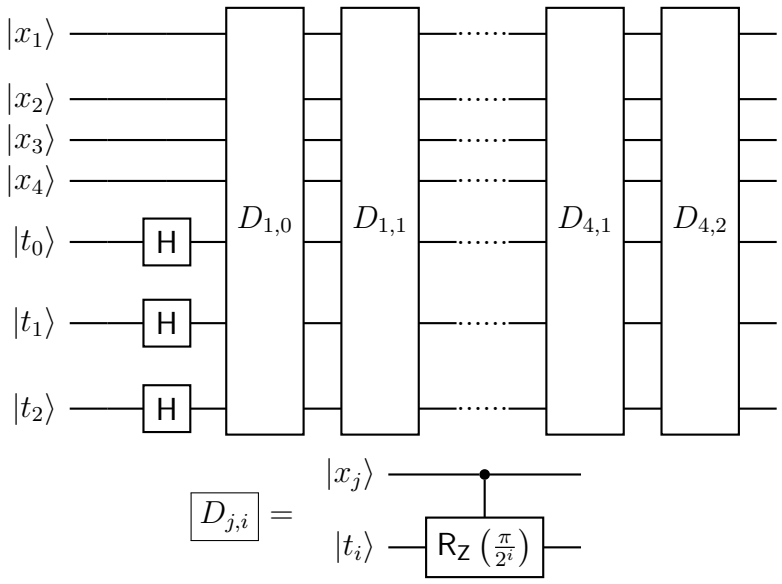

Figure 3.4: poly-depth implementation of OR's reduction to log size

As the Z rotation happens only with $|x_j\rangle = |1\rangle$, the phase in each $|t_i\rangle$ state will be proportional to $|x|$. Notice that there are $O(a\log(a))$ D gates, and their only action is applying a Z rotation to one of the $|t\rangle$ qubits – these gates commute with each other. Therefore, we can apply Lemma 3.2 with $\{D_{j,i}\}$ as the mutually commuting operations to get a $\mathsf{QNC}^0_\mathsf{f}$ circuit that outputs $|t\rangle$. □

Now by repeating this procedure until the number of qubits under consideration is two and then using the bounded fan-in quantum gates to compute the reversible OR_2, we can recover Høyer and Špalek's $\log^*$-depth implementation.

Step 2: Computing OR in exponential size and constant depth

As discussed in Chapter 2, all Boolean functions on input y of size n in the $\{\pm 1\}$ basis can be uniquely represented as multi-linear polynomials:

$$f(y) = \sum_{s \subseteq [n]} c_s \prod_{i \in s} y_i \tag{3.7}$$

where $\{c_s\}$ are the real Fourier coefficients of f. Note that $\prod_{i \in s} x_i$ precisely computes parity over subset $s \subseteq [n]$ but in the $\{\pm 1\}$ basis. We can convert this statement to the $\{0/1\}$ basis (input x) using the transformations $\prod_{i \in s} y_i = 1 - 2\mathsf{MOD}_{2,|s|}(x_s)$ and $f(y) = 1 - 2f(x)$ for corresponding x and y inputs. We get that:

$$f(x) = \sum_{s \subseteq [n]} c'_s \mathsf{MOD}_{2,|s|}(x_s) + c \tag{3.8}$$

for some reals $\{c'_s\}$ and some constant c. Moreover, for a function which evaluates to 0 on the all-zero input, $c = 0$. Using induction, we can quickly show that for any subset $s \subseteq [n]$, $c'_s = \frac{1}{2^{n-1}}$ for OR_n [TT16]. The base case (for $n = 1$) is trivial. Let this be true for inputs of size $n = k$. Then, for the $k+1$ case, we see that OR_{k+1} has the expected behaviour:

$$\mathsf{OR}_{k+1}(x_1, ..., x_{k+1}) = \frac{1}{2^k} \sum_{s \subseteq [k+1]} \mathsf{MOD}_{2,|s|}(x_s) \tag{3.9}$$

$$= \begin{cases} \frac{1}{2^k} \sum_{s \subseteq [k]} 2\mathsf{MOD}_{2,|s|}(x_s) & \text{if } x_{k+1} = 0 \\ \frac{1}{2^k} \sum_{s \subseteq [k]} \left(\mathsf{MOD}_{2,|s|}(x_s) + \left(1 - \mathsf{MOD}_{2,|s|}(x_s)\right)\right) & \text{if } x_{k+1} = 1 \end{cases} \tag{3.10}$$

$$= \begin{cases} \mathsf{OR}_k(x_1, ..., x_k) \text{ if } x_{k+1} = 0 \\ 1 \text{ if } x_{k+1} = 1 \end{cases} \tag{3.11}$$

The second step follows by considering subsets of $[k+1]$ containing x_{k+1} and those that do not separately. The last step follows from the induction hypothesis. Now, for a given basis state $|x\rangle$, we can compute OR_n by simply computing each of these parity terms separately (in parallel) and then combining them:

Theorem 3.5. *There is a constant-depth circuit for OR_n using bounded fan-in quantum gates and exponential-size quantum fan-out gates.*

Proof. Using the exponential-size quantum fan-out gate, we first fan-out the classical input 2^n times. Therefore, our state is:

$$\mathsf{F}_{2^n} |x\rangle \left|0^{2^n}\right\rangle = |x\rangle^{2^n+1} \tag{3.12}$$

Using the 2^n copies, we compute $\mathsf{MOD}_2(x_s)$ for every $s \subseteq [n]$ using the construction from Observation 3.3 in parallel. Now our state is:

$$|x\rangle \prod_{s\subseteq[n]} \left|\mathsf{MOD}_{2,|s|}(x_s)\right\rangle \tag{3.13}$$

In another auxiliary qubit, we construction $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and use the quantum fan-out operation to create a 2^n-size cat state. Then, we perform the following circuit in one layer:

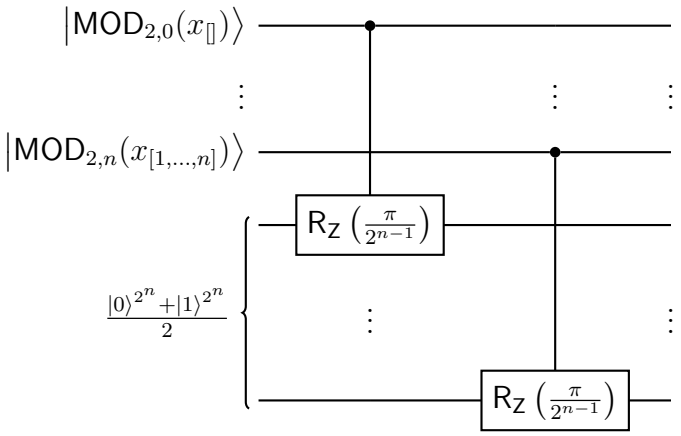

Figure 3.5: Computing OR_n in phase of cat state

The cat state is now:

$$\frac{|0\rangle^{2^n} + \exp\left(i\frac{\pi}{2^{n-1}}\sum_{s\in[n]}\mathsf{MOD}_{2,|s|}(x_s)\right)|1\rangle^{2^n}}{\sqrt{2}} \tag{3.14}$$

Next, we use quantum fan-out to invert this back to a single qubit. Then, by applying a Hadamard gate and using OR_n's Fourier expansion, we get:

$$\mathsf{H}\left(\frac{|0\rangle + \exp\left(i\frac{\pi}{2^{n-1}}\sum_{s\in[n]}\mathsf{MOD}_2(x_s)\right)|1\rangle}{\sqrt{2}}\right) = \mathsf{H}\left(\frac{|0\rangle + \exp\left(i\pi\mathsf{OR}_n(x)\right)|1\rangle}{\sqrt{2}}\right) = |\mathsf{OR}_n(x)\rangle \tag{3.15}$$

Note that we only used a constant number of layers of bounded fan-in gates and exponential-size quantum fan-out gates. □

Moreover, all auxiliary qubits that are still entangled with the input $|x\rangle$ (for example, the $\prod_{s\subseteq[n]}\left|\mathsf{MOD}_{2,|s|}(x_s)\right\rangle$ states) at the end can be uncomputed in constant depth after generating the $|\mathsf{OR}_n(x)\rangle$. Therefore, we can apply this procedure for inputs in superposition and implement QOR_n. We have now shown that $\mathsf{QAC}^0[2] \subseteq \mathsf{QNC}^0_\mathsf{f}$ and consequentially, $\mathsf{AC}[2] \subseteq \mathsf{QNC}^0_\mathsf{f}$.

3.3 Simulating $\mathsf{QAC}^0[p]$

The $\mathsf{QNC}^0_\mathsf{f}$ implementation of $\mathsf{QMOD}_{p,a}$ was first shown in [Moo99] and is a direct consequence of the parallelization trick (Lemma 3.2). This was the first surprising result at the time, alluding to the power of unbounded quantum fan-out. This is because we cannot compute $\mathsf{MOD}_{q,a}$ even approximately with $\mathsf{AC}^0[2]$ circuits for any prime q. However, in the quantum case, $\mathsf{QAC}^0[2]$ can compute $\mathsf{QMOD}_{p,a}$ exactly for any prime p. Moreover, there is nothing special about $\mathsf{QMOD}_{2,a}$; [GHMP01] later showed the opposite containment: $\mathsf{QMOD}_{2,a}$ can be exactly computed by a $\mathsf{QAC}^0[p]$ circuit for any prime p.

Corollary 3.6. *$\mathsf{QMOD}_{p,a}$ for $a \in \mathsf{poly}(n)$ can be computed by a circuit in $\mathsf{QNC}^0_\mathsf{f}$.*

Proof. We treat p as a constant. Let the input be some basis state $|x\rangle$ for $x \in \{0,1\}^a$. Consider the unitary operation defined by the following action on $|t\rangle$ with $0 \leq t < p$:

$$\mathsf{PLUS}_p |t\rangle = |(t+1) \mod p\rangle \tag{3.16}$$

$|t\rangle$ can be represented using $m = \lceil \log(p) \rceil$ qubits via its binary representation. Consider the following circuit acting on $|x\rangle |t = 0\rangle$:

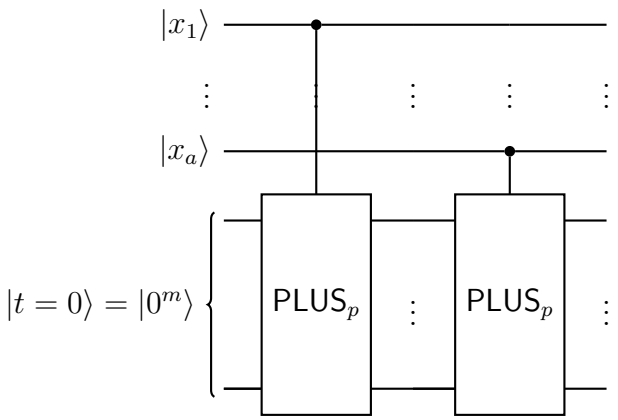

Figure 3.6: Computing $\mathsf{MOD}_{q,a}$ in auxiliary qubits

As we increment t controlled on x_i for $i \in [a]$, the auxiliary qubits are in final state $|t_f\rangle = |0^m\rangle$ if and only if $|x| \mod p = 0$. Moreover, the controlled-PLUS_q operations commute with each other, and each of them can be implemented in constant depth as they act on a constant number of qubits. Therefore, we can apply Lemma 3.2 with the controlled-PLUS_q operations as our mutually commuting operations to compute the final state of the circuit in $\mathsf{QNC}^0_\mathsf{f}$. Now we simply use bounded fan-in QOR on the auxiliary qubits to check whether $|t_f\rangle = |0^m\rangle$ and hence, $\mathsf{QOR}_m |t_f\rangle |0\rangle = |t_f\rangle |\mathsf{MOD}_{q,a}(x)\rangle$. Using the uncomputing trick, this procedure works for a superposition of basis states as input. □

This shows us that $\mathsf{QACC}^0 \subseteq \mathsf{QNC}^0_\mathsf{f}$ and separates $\mathsf{QNC}^0_\mathsf{f}$ from all $\mathsf{AC}^0[p]$ for prime p.

3.4 Simulating QTC^0

We have all the tools required to simulate unbounded fan-in quantum threshold gates in $\mathsf{QNC}^0_\mathsf{f}$ using the approach from [HŠ05].

Theorem 3.7. *The unbounded fan-in quantum threshold gate $\mathsf{QTH}_{t,a}$ for $a \in \mathsf{poly}(n)$ can be computed by a circuit in $\mathsf{QNC}^0_\mathsf{f}$.*

Proof. Consider $\mathsf{TH}_{t,a}$ for $0 \leq t \leq a$. We will be building the threshold gate by combining circuits for the exact gates. Consider the following state:

$$\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle \tag{3.17}$$

If $|x| \geq t$ then, $\mathsf{EX}_{|x|,a}|x\rangle = |1\rangle$ (and the rest are in state $|0\rangle$) and if not then $\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle = |0^{a-t+1}\rangle$. Therefore, by computing the parity of state $\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle$, we can compute the threshold gate:

$$\mathsf{MOD}_{2,a-t+1}\left(\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle\right) = |\mathsf{TH}_{t,a}(x)\rangle \tag{3.18}$$

For now, let us assume that there is a $\mathsf{QNC}^0_\mathsf{f}$ procedure A_k that generates the state $|\mathsf{EX}_{k,a}(x)\rangle$. We first fan-out $a-t+1$ copies of $|x\rangle$ to generate $(\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle)$ in parallel and then apply the $\mathsf{MOD}_{2,a-t+1}$ gate. Here is a circuit summarizing the procedure:

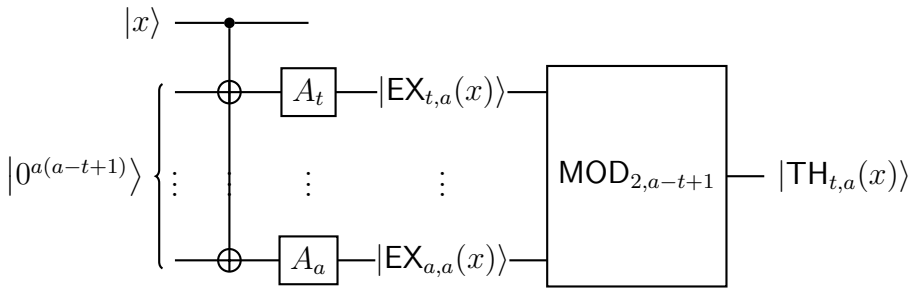

Figure 3.7: Computing $\mathsf{TH}_{t,a}$ using exact gates

As $\mathsf{MOD}_{2,a-t+1} \in \mathsf{QNC}^0_\mathsf{f}$ and $(\prod_{k=t}^{a} |\mathsf{EX}_{k,a}(x)\rangle)$ can be generated in $\mathsf{QNC}^0_\mathsf{f}$ via the A_k operations, we can compute $\mathsf{TH}_{t,a}(x)$ in $\mathsf{QNC}^0_\mathsf{f}$.

Therefore, all there is left is to construct the required A_k. We will slightly modify the $|t\rangle$ construction from Lemma 3.4. For $0 \leq j \leq m = \lceil \log(a) \rceil$, we define:

$$|t_j\rangle \propto |0\rangle + \exp\left(i\pi \frac{|x|-k}{2^j}\right)|1\rangle \tag{3.19}$$

Note that if $|x|-k=0$, all $|t_j\rangle \propto |0\rangle+|1\rangle$ and hence, $(\mathsf{H})^{\otimes(m+1)}|t\rangle = |0^{m+1}\rangle$ for $|t\rangle = \prod_j |t_j\rangle$. Otherwise, we have that $|x|-k \leq a$. Consider the unique prime factorization of $|x|-k$ as

$2^b \times$ (some odd integer) for some natural number b. Note that as $|x| - k \leq a$, it must be that $0 \leq b \leq \log(a)$. Therefore:

$$|t_b\rangle \propto |0\rangle + \exp\left(i\pi \frac{2^b \times (\text{some odd integer})}{2^b}\right)|1\rangle \tag{3.20}$$

$$= |0\rangle - |1\rangle \tag{3.21}$$

Therefore, $H|t_b\rangle = |1\rangle$ and hence, each individual basis state in $(\mathsf{H})^{\otimes(m+1)}|t\rangle$ has at least one $|1\rangle$ and we have that

$$|\mathsf{EX}_{k,a}(x)\rangle = \mathsf{NOT}\,\mathsf{OR}_{m+1}\,(\mathsf{H})^{\otimes(m+1)}|t\rangle \tag{3.22}$$

in all cases. NOT and OR_{m+1} are in $\mathsf{QNC}^0_\mathsf{f}$. Finally, to generate $|t_j\rangle$, we append the procedure in Lemma 3.4 by a $\mathsf{R}_z\left(\pi\frac{-k}{2^j}\right)$ (uncontrolled) rotation before applying Lemma 3.2.

As before, all entangled auxiliary qubits can be uncomputed at the end, while MOD and OR can be upgraded to their quantum versions to make this procedure work for quantum inputs.

□

We have now shown our final inclusion result that $\mathsf{TC}^0 \subseteq \mathsf{QNC}^0_\mathsf{f}$. Moreover, we have that $\mathsf{QTC}^0 \subseteq \mathsf{QTC}^0_\mathsf{f} = \mathsf{QAC}^0_\mathsf{f} = \mathsf{QNC}^0_\mathsf{f}$. Therefore, the hierarchy of constant-depth quantum circuits with access to unbounded quantum fan-out collapses.

3.5 Implementing quantum fan-out

This section discusses the limitations of implementing unbounded quantum fan-out using bounded fan-in gates and Hamiltonian-based schemes proposed to counter these limitations.

3.5.1 Log-depth implementation of quantum fan-out

Notice that in Figure 3.2, if the auxiliary qubits were in state $|0^{a-1}\rangle$, after the first CNOT, the top two qubits are indistinguishable: they are in the state $\alpha|00\rangle + \beta|11\rangle$. Therefore, the following CNOTs can be controlled by either of the two qubits; hence, we can now apply the next two CNOT gates in a single layer. By repeating this procedure, in one layer, we can double the size of the cat state. The cat state of size a, $\alpha|0^a\rangle + \beta|1^a\rangle$ can be generated in $\log(a-1)$ depth with access with $a-1$ additional auxiliary qubits. Once we form this cat state, we can apply the F_a operation to the original auxiliary qubit state by applying a CNOT gates, each controlled by one of the qubits in the cat state with a qubit in the original auxiliary state as the target.

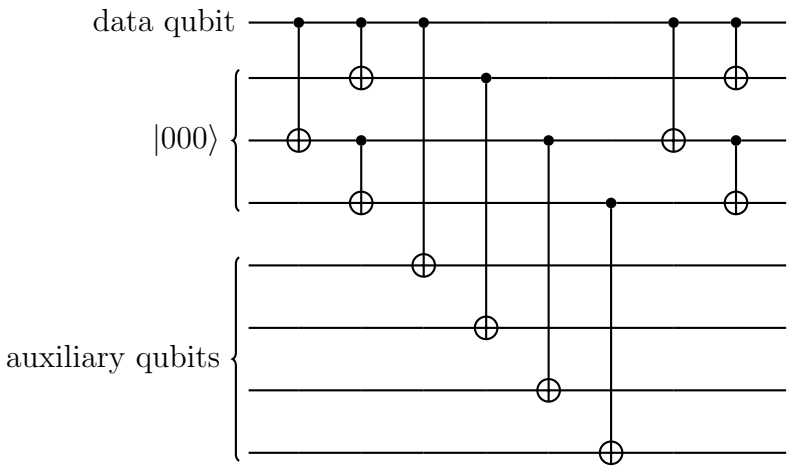

Figure 3.8: Log-depth implementation of fan-out gate F_4

For a basis auxiliary state $|x\rangle$ with $x \in \{0,1\}^a$, we are now in state: $\alpha |0\rangle |0^{a-1}\rangle |x\rangle + \beta |1\rangle |1^{a-1}\rangle |\bar{x}\rangle$. We then apply the inverses of the gates used to create the cat state to reset it back to state $|0^{a-1}\rangle$:

Observation 3.8. *Quantum fan-out operation to $a \in O(\mathsf{poly}(n))$ auxiliary qubits can be implemented in $O(\log(a))$-depth using bounded fan-in quantum gates.*

3.5.2 Log-depth lower bound of quantum fan-out

We defer the corresponding lower bound to Chapter 6 where we prove a stronger result using light cone arguments: any circuit that only uses bounded fan-in quantum gates (even with access to arbitrary auxiliary states) requires $O(\log(n))$-depth to compute unbounded fan-in parity (Theorem 6.19). Due to the relationship between parity and unbounded quantum fan-out, this implies that the depth of the construction in the previous section is optimal. Note that this gives us a strict separation: $\mathsf{QNC}^0 \subsetneq \mathsf{QNC}^0_{\mathsf{f}}$.

3.5.3 Proposals for Hamiltonian-based implementations

Given this limitation of models based on bounded fan-in quantum gates, ever since the power of the unbounded quantum fan-out gate was discovered, an important task has been to identify whether there are other *realistic* models of quantum computation where quantum fan-out can be implemented in constant-depth natively. Here we discuss some implementations for models based on Hamiltonian evolution. The rough intuition behind why this might be possible is that if we have two gates $A = \exp(ia)$ and $B = \exp(ib)$ that commute but act on the same sets of qubits, the corresponding Hamiltonians that implement these gates must commute as well. Moreover, in devices with programmable control over the global Hamiltonian of the system, A and B can be implemented simultaneously by evolving the system under the joint Hamiltonian $\exp(i(a+b))$ in one step. This can

be extended to an unbounded number of commuting gates, thus implementing quantum fan-out.

More specifically, [Fen03, FZ04] provide a constant-time Hamiltonian that implements quantum fan-out. However, one drawback of this approach is that it requires an all-to-all interaction graph with equal interaction strength among all pairs of qubits. This requirement may not be satisfied by scalable architectures. Nevertheless, refer to [GKH+21] for a review of proposals and experimental demonstrations of this approach on trapped ion, nuclear magnetic resonance, neutral atom, superconducting circuit and nitrogen-vacancy center platforms.

If we restrict the interaction graph of the Hamiltonian to only nearest-neighbours on a d-dimensional lattice, [Ros12] provides a no-go theorem as they prove a $n^{\frac{1}{d}}$ lower bound for implementing fan-out. However, a more realistic setting would be where all qubits interact, but the interaction strength decays with distance. [FW22] provides a characterization of interaction strength variations that can be withstood by the [Fen03, FZ04] approach with some examples showing that interaction strength can decay with distance and still permit constant-depth quantum fan-out. More recently, [GDC+20] provides a different protocol on a d-dimensional lattice that performs quantum fan-out in constant depth if the interaction strength scale as $\frac{1}{r^a}$ for $a < d$ where r is the distance between the two qubits.

Chapter 4

Adaptive constant-depth quantum circuits

Compared to the implementations for bounded fan-in quantum gates, Hamiltonian-based proposals for quantum fan-out discussed at the end of the previous chapter require programmable control over the global Hamiltonian and hence, increased experimental complexity. Therefore, scalable constant-depth quantum fan-out implementations have remained elusive. In 2009, Browne, Kashefi, and Perdrix [BKP10] provided a pathway towards remedying these challenges by constructing a way to implement quantum fan-out in constant-depth using only *adaptive* bounded fan-in quantum gates. That is, QNC^0 with intermediate measurements where measurement outcomes can classically control the application of subsequent gates.

[BKP10] does so by exploiting a critical connection between measurement-based quantum computation and Clifford operations. In particular, it has been well known since the advent of the measurement-based paradigm of quantum computation that all Clifford operations can be parallelized to constant depth, provided the quantum circuit is interspersed with measurements and efficient classical computation and the quantum gates are *adaptive* i.e. they can depend on the output of preceding classical gates [GC99]. Now, as quantum fan-out is Clifford, this gives us with a more winsome model, adaptive constant-depth quantum circuits (adaptive-QNC^0), that has the power of $\mathsf{QNC}^0_\mathsf{f}$ at minimum.

In this chapter, first, we briefly review a particular form of measurement-based quantum computation, teleportation-based quantum computation and then, we show how to perform quantum fan-out (or any other Clifford operation) using an adaptive constant-depth quantum circuit.

4.1 Teleportation-based quantum computation

Our starting point is the quantum state teleportation protocol which was first discussed in [BBC+93]. We are given qubit A in some state $|\psi_0\rangle_A = \alpha|0\rangle + \beta|1\rangle$ and qubits B and C that share a cat state $|\psi_0\rangle_{BC} = \frac{|00\rangle+|11\rangle}{\sqrt{2}}$. Now, by an explicit calculation, it can be shown

that by measuring qubits A and B in the Bell basis and then, by performing certain Pauli gates on qubit C conditioned on the measurement outcomes, we can "teleport" $|\psi_0\rangle_A$, that is, qubit C now holds state $\alpha|0\rangle + \beta|1\rangle$.

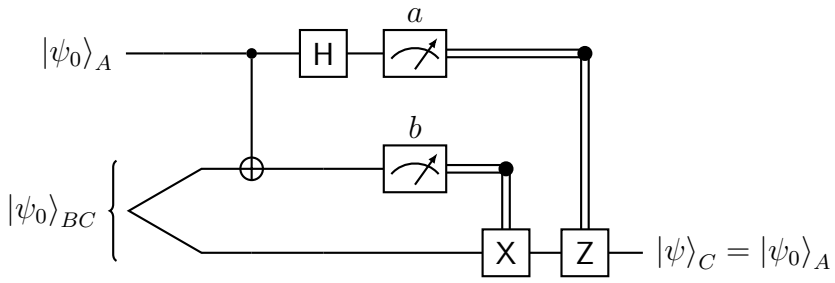

Figure 4.1: Quantum state teleportation

Here, the $<$ symbol represents the creation of a cat state as below:

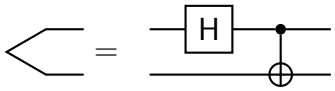

and $a, b \in \{0, 1\}$ are the measurement outcomes. This is called teleportation, as we can imagine the following scenario: after the qubits B and C were entangled to create the cat state, we take them far apart to two different locations. Then, we perform a joint measurement of one of them with the data qubit and send the measurement outcomes (using a *classical* communication channel) to the other qubit of the cat state that then performs the required corrections. Through this process, we were able to move the data qubit's quantum state between the two locations. We can use this procedure as a primitive anywhere in a quantum circuit to effectively "move" a qubit state.

Note that if in Figure 4.1 we applied some single-qubit gate U before the CNOT between qubit A and B, we would simply teleport $\mathsf{U}|\psi_0\rangle_A$. Moreover, $\mathsf{U}, \mathsf{CNOT}$ and H can be "absorbed" into the measurement if we allow arbitrary two-qubit measurements. The measurement basis would be determined by U – it would be the Bell basis rotated by U on the first qubit. Therefore, we can apply arbitrary single-qubit gates using existing entangled pairs (a cat state), two-qubit measurements and adaptive Pauli gates.

Gottesman and Chuang [GC99] noticed that the CNOT gate could also be implemented in this manner. This is crucial as CNOT and single-qubit gates form a universal gate set for quantum computation. Their construction is a simple adaptation of the above protocol: consider qubits A, A', B, B', C and C' where we want to apply CNOT between A and A'. By the construction in Figure 4.1, the following circuit performs the required CNOT operation:

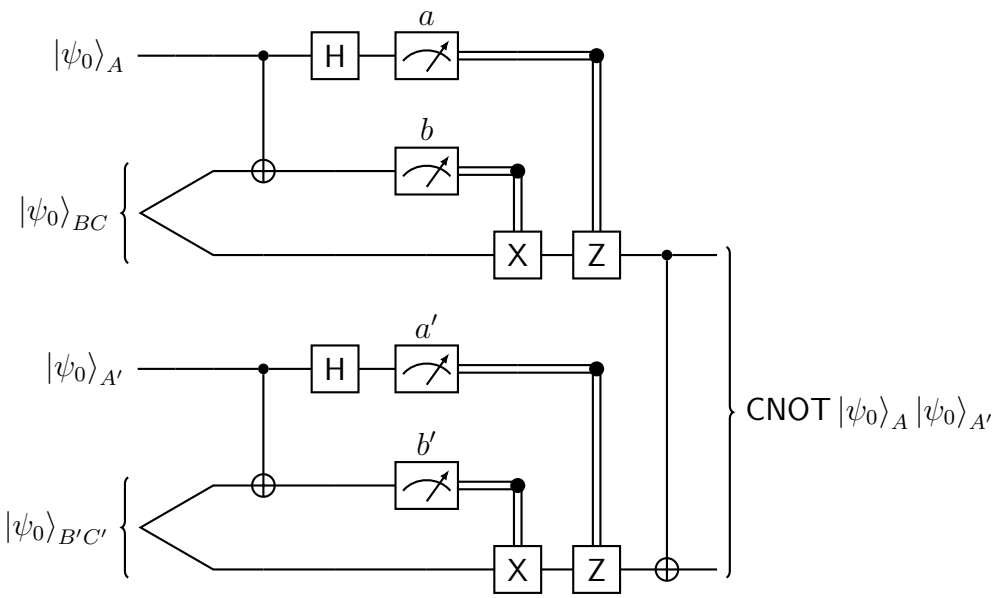

Figure 4.2: Applying CNOT after teleporting input states

This is because after the teleportation gadget, qubits C and C' hold the input states to the CNOT. Now, we can use the fact that CNOT commutes with the Pauli group. Explicitly, we have the following relations for elements P_1, P_2 from the Pauli group:

P_1	$P_2 = \mathsf{CNOT}\, P_1\, \mathsf{CNOT}$
$\mathsf{X} \otimes \mathsf{I}$	$\mathsf{X} \otimes \mathsf{X}$
$\mathsf{Z} \otimes \mathsf{I}$	$\mathsf{Z} \otimes \mathsf{I}$
$\mathsf{I} \otimes \mathsf{X}$	$\mathsf{I} \otimes \mathsf{X}$
$\mathsf{I} \otimes \mathsf{Z}$	$\mathsf{Z} \otimes \mathsf{Z}$

Table 4.1: CNOT-Pauli commutation relations

which can be verified by an explicit calculation. Now in Figure 4.2, for every controlled Pauli P_1 that appears before the CNOT, we can use the $P_1\mathsf{CNOT} = \mathsf{CNOT}P_2$ relation to replace P_1 with P_2 and move the CNOT to the very beginning – before the application of the Pauli corrections. In fact, as the CNOT gate act on qubits C and C' on which no other operations are applied apart from the Pauli corrections, we can apply the CNOT in the state preparation step – even before we are given qubits A and A'.

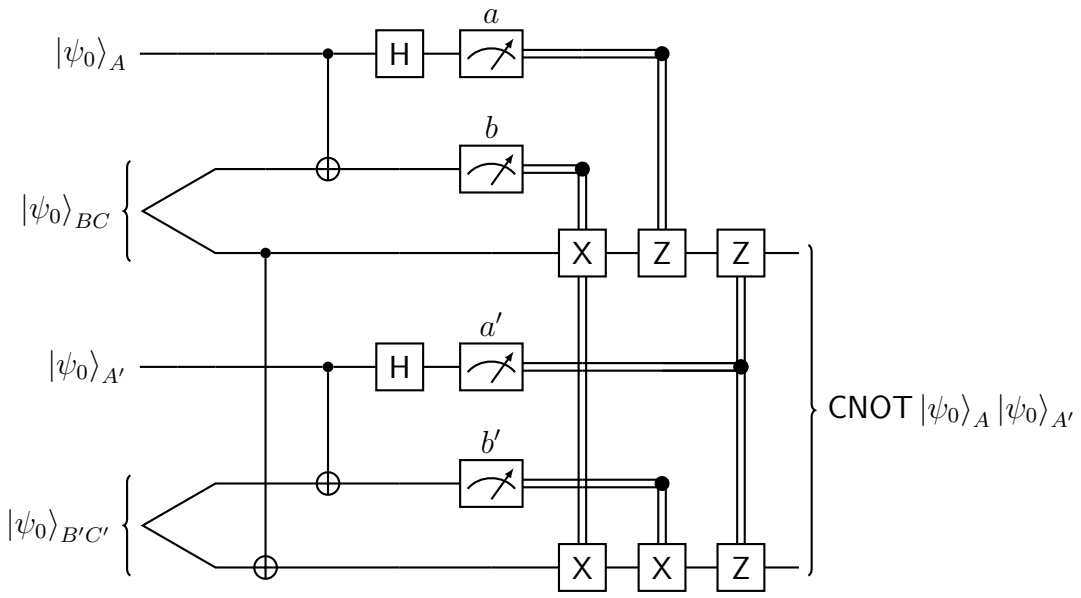

Figure 4.3: Implementing CNOT using measurements and classical feedback

This shows us that teleportation can be viewed as a computational primitive for universal quantum computation:

Observation 4.1. *We can perform universal quantum computation using constant-depth quantum advice, two-qubit measurements and classically-controlled Pauli gates.*

where the quantum advice we require are the two-size cat states and states of the form:

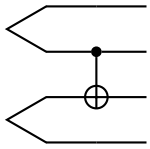

which can both be prepared in constant depth. Note that for qubit C and C', the Pauli correction to be performed now depend on the outcome of both sets of measurements. In fact, they can be expressed as $\mathsf{X}^b\mathsf{Z}^{a+a'}$ for C and $\mathsf{X}^{b+b'}\mathsf{Z}^{a'}$ for C'. Moreover, further work shows that we can also eliminate the quantum advice and Pauli gates requirements in Observation 4.1 [Nie03, Leu01, Leu04]. They do so by providing schemes for constructing two-size cat states and Pauli gates only using 2-qubit measurements. Other measurement-based quantum computation paradigms exist, for example, the one-way model [RBB03] and the hypergraph state based model [GGM19]. In these models, an entangled graph or hypergraph state is prepared initially (akin to the cat states in the teleportation-based model), and only single-qubit measurements suffice for universality. Refer to [Joz06, GGM19] for a discussion on the relationship between the three models.

4.2 Simulating Clifford operations using teleportation

This section will review how all *poly*-size Clifford circuits can be performed by adaptive-QNC^0.

We consider the setting in which we want to perform layers of gates $L_d \dots L_1$ where each layer L_i consists of single-qubit and CNOT gates acting on disjoint sets of qubits. We can first perform L_1 via teleportation as described in Section 4.1 and then use the "new" wires (for example, qubits C and C' in Figure 4.3) as inputs to L_2's teleportation-based implementation. However, we cannot make the required measurements for L_1 and L_2 simultaneously as the Pauli corrections caused due to implementing L_1 must be performed before performing L_2. There is no decrease in depth through a teleportation-based implementation compared to the circuit model. In principle, we perform the Pauli corrections for every layer, each incrementing the depth of the adaptive quantum circuit.

Note that these corrections are the only way the input wires of a layer (qubit C and C') directly interact with the input wires of an earlier layer (qubit A and A'). Therefore, if we do *not* perform the Pauli corrections, the quantum circuit will be constant depth. Explicitly, let us consider the depth d circuit on a single qubit $\mathsf{U}_n \dots \mathsf{U}_1 |\psi_0\rangle$ and its construction via teleportation but ignoring the Pauli corrections:

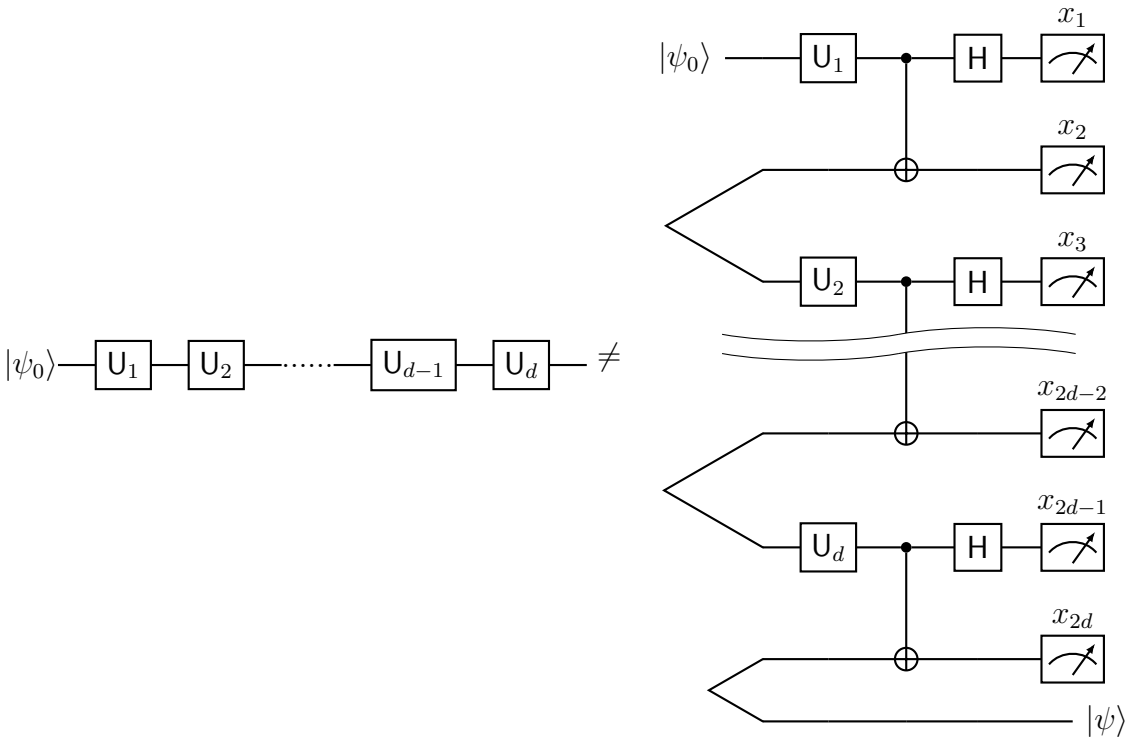

Figure 4.4: Teleportation-based construction of a d-depth quantum circuit without Pauli corrections

The output wire of the final layer will no longer hold the state $U_1 U_2 \dots U_n |\psi_0\rangle$. Note that we effectively performed the U_2 gate on $\mathsf{X}^{x_2}\mathsf{Z}^{x_1}\mathsf{U}_1 |\psi_0\rangle$ and so on. Therefore, instead,

the final state is:

$$|\psi\rangle = \mathsf{X}^{x_{2d}}\mathsf{Z}^{x_{2d-1}}\mathsf{U}_d\mathsf{X}^{x_{2d-2}}\mathsf{Z}^{x_{2d-3}}\mathsf{U}_{d-1}\dots U_3\mathsf{X}^{x_4}\mathsf{Z}^{x_3}\mathsf{U}_2\mathsf{X}^{x_2}\mathsf{Z}^{x_1}\mathsf{U}_1|\psi_0\rangle \tag{4.1}$$

Generalising this to the case of a general circuit on n qubits, we can use the CNOT construction in Figure 4.3 to produce a similar constant-depth circuit that outputs the state:

$$|\psi\rangle = \left(\prod_i^n \mathsf{X}_i^{x_{2d,i}}\mathsf{Z}_i^{x_{2d-1,i}}\right) L_d \left(\prod_i^n \mathsf{X}_i^{x_{2d-2,i}}\mathsf{Z}_i^{x_{2d-3,i}}\right) L_{d-1}\dots \tag{4.2}$$

$$\dots L_3 \left(\prod_i^n \mathsf{X}_i^{x_{4,i}}\mathsf{Z}_i^{x_{3,i}}\right) L_2 \left(\prod_i^n \mathsf{X}_i^{x_{2,i}}\mathsf{Z}_i^{x_{1,i}}\right) L_1|\psi_0\rangle \tag{4.3}$$

where X_i (Z_i) represents X (Z) acting on qubit i and $x_{j,i}$ for $j \in [2d]$ is either a measurement outcome or sum of two measurement outcomes (coming from the CNOT case in Figure 4.3). In general, we don't know how to decipher the actual output state $L_d \dots L_1|\psi_0\rangle$ from $|\psi\rangle$. However, if we restrict to the setting where each layer only uses Clifford gates, we can exploit the following fact discussed in Chapter 2:

Fact 4.2. *The Clifford group and Pauli group commute with each other.*

In particular, using the commutation relations (Table 4.1) and the relation $\mathsf{XZ} = -\mathsf{ZX}$, we can commute $\left(\prod_i^n \mathsf{X}_i^{x_{2,i}}\mathsf{Z}_i^{x_{1,i}}\right)$ across L_2:

$$L_2\left(\prod_i^n \mathsf{X}_i^{x_{2,i}}\mathsf{Z}_i^{x_{1,i}}\right) L_1|\psi_0\rangle = \left(\prod_i^n \mathsf{X}_i^{y_{2,i}}\mathsf{Z}_i^{y_{1,i}}\right) L_2L_1|\psi_0\rangle \tag{4.4}$$

for some $\{y_1\}, \{y_2\}$ where each element is sum of elements from $\{x_1\}$ and $\{x_2\}$. Now as $\mathsf{X}^2 = \mathsf{Z}^2 = \mathsf{I}$, we can combine these coefficients with the coefficients of $\left(\prod_i^n \mathsf{X}_i^{x_{4,i}}\mathsf{Z}_i^{x_{3,i}}\right)$ and get that:

$$\left(\prod_i^n \mathsf{X}_i^{x_{4,i}}\mathsf{Z}_i^{x_{3,i}}\right) L_2\left(\prod_i^n \mathsf{X}_i^{x_{2,i}}\mathsf{Z}_i^{x_{1,i}}\right) L_1|\psi_0\rangle = \left(\prod_i^n \mathsf{X}_i^{x_{4,i}\oplus y_{2,i}}\mathsf{Z}_i^{x_{3,i}\oplus y_{1,i}}\right) L_2L_1|\psi_0\rangle \tag{4.5}$$

By using the above two transformations repeatedly, we can write state $|\psi\rangle$ as the required output state followed by some Pauli gates:

$$|\psi\rangle = \left(\prod_i^n \mathsf{X}_i^{p_i}\mathsf{Z}_i^{q_i}\right) L_dL_{d-1}\dots L_2L_1|\psi_0\rangle \tag{4.6}$$

where each element in $\{p\}$ and $\{q\}$ is the parity of a subset of measurement outcomes of the constant-depth circuit. Moreover, as there are only $2n$ terms that we need to compute, we can first use $2n$-size classical fan-out to make $2n$ copies of each of the measurement outcomes and then use a single layer of classical parity gates to compute $\{p\}$ and $\{q\}$. Then, in two layers, we can apply Pauli gates controlled by $\{p\}$ and $\{q\}$ to "correct" $|\psi\rangle$. The following figure demonstrates this procedure for the one-qubit case:

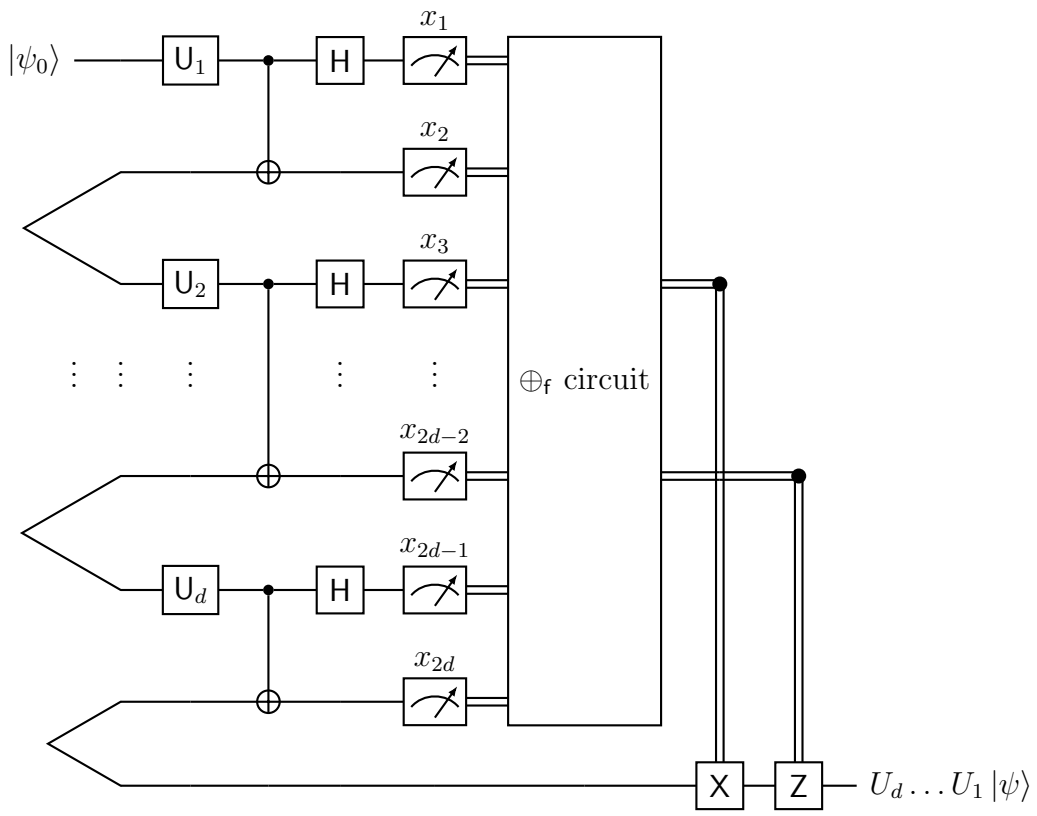

Figure 4.5: Adaptive constant-depth circuit for d-depth Clifford circuit

Theorem 4.3. poly*-size Clifford circuits can be simulated by adaptive constant-depth quantum circuits – in particular, by* poly*-size constant-depth quantum circuits interleaved with a constant-depth* $\oplus_{\mathsf{f}}$ *circuit.*

We will make a more rigorous statement in Chapter 6 after formally defining an interleaving model inspired by the above construction.

Sampling from Clifford circuits without adaptivity: [Joz06] noted that if we merely want to sample from the Clifford circuit, we do not even need to apply the Pauli corrections. We first measure the output wires of the final layer ($|\psi\rangle$). The Z corrections do not impact the probability distribution, while X corrections will just flip some of the measurement outcomes. Therefore, to faithfully sample from the Clifford circuit, we just compute the $\{p\}$ terms and flip the measurement outcome corresponding to qubit i if $p_i = 1$. Therefore, we can sample from poly-size Clifford circuits using a constant-depth quantum circuit followed by classical post-processing. Note that this sampling protocol differs from the one in Gottesman-Knill theorem [Got98] in one fundamental way – the initial state does not have to be a computational basis state. This sampling problem with non-computational basis states is known to be classically hard under complexity-theoretic assumptions [BJS11, JN13, Koh15, MSM17, BFK17, YJS19, PBG20]. Therefore, the above approach provides a classically hard sampling task that can be performed by a constant-depth quantum circuit. However, this reduction in depth comes at the expense of increasing the number of qubits.

Chapter 5

Constant-depth quantum circuits

In a breakthrough result, [BGK18] discovered a relational problem (the two-dimensional hidden linear function problem) that can be solved exactly by a QNC^0 circuit but not by any NC^0 circuit. This showed the first unconditional separation between the vanilla versions of two corresponding quantum and classical circuit classes:

$$\mathsf{QNC}^0 \not\subset \mathsf{NC}^0 \tag{5.1}$$

This result has been extended in two directions:

- It was made robust to two different noise models:
  - [BGKT20] showed that the above separation is sustained even if the quantum circuit suffers from local stochastic noise: at every layer of the circuit, for some error rate p, there is a p^k probability of a random k-qubit Pauli error. This was done by equipping the original protocol with single-shot fault-tolerant state preparation and error correction codes that only used adaptive constant-depth Clifford circuits.
  - [HG21a] considered a somewhat non-standard noise model, showing that if a constant fraction of the qubits have been manipulated by noise, the relational problem on the remaining qubits is still hard for NC^0. This was achieved by considering a modified version of the two-dimensional hidden linear function problem, the graph state sampling problem on expander graphs.
- [BGK18, Gal18, CSV21] supplemented this result with an average-case statement: no NC^0 circuit can solve the problem on a small fraction of the inputs, or equivalently, on a random input. Later, [WKST19] strengthened the result beyond NC^0, showing that no AC^0 circuit can solve the two-dimensional hidden linear function problem on a random even-parity input:

$$\mathsf{QNC}^0 \not\subset \mathsf{AC}^0 \tag{5.2}$$

The [WKST19] result subsumes the existing results from [BGK18, Gal18, CSV21]. Therefore, this chapter will briefly review the technical steps towards the AC^0 lower-bound following the approach in [WKST19]. In particular, for simplicity, instead of proving the above result for the two-dimensional hidden linear problem, we will consider variants of *the parity halving problem*. However, as the two-dimensional hidden linear problem can be reduced to these problems, all results also apply to the former.

We will first show how the parity halving problem separates AC^0 from QNC^0 circuits with access to cat states. Then, we will observe how to modify the problem to achieve a separation for QNC^0 without quantum advice.

5.1 Separations with quantum advice

QNC^0 with cat state advice provides us with access to cat states of size $c \in \mathsf{poly}(n)$. However, critically, the cat state (or its size) cannot depend on the input x (it can only depend on the input size). We can think of this as a pre-processing step. This increases the QNC^0 circuit's power and enables parallelization of certain operations (by letting all input qubbits interact amongst each other in constant depth). We will revisit cat state advice and its power in Section 6.2.1.

5.1.1 Parity halving problem

Definition 5.1 (Parity halving problem)**.** *For a given $x \in \{0,1\}^n$ with even parity, to solve the parity halving problem with m outputs ($\mathsf{PHP}_{n,m}$) we must return $y \in \{0,1\}^m$ such that:*

$$|y| = \frac{|x|}{2} \ (mod\ 2) \tag{5.3}$$

Quantum upper bound: This problem for $n = m$ has been studied as a non-local game in [Mer90, BBT04] in the setting with n players that cannot communicate with each other. Player i gets x_i and responds with y_i. The perfect quantum strategy [Mer90, BBT04] for this game can be implemented as a $\mathsf{QNC}^0_{\text{🙀}}$ circuit as follows.

Theorem 5.2. *A circuit $\mathsf{QNC}^0_{\text{🙀}}$ solves $\mathsf{PHP}_{n,m}$ exactly for $m \geq n$ and $m \in \mathsf{poly}(n)$.*

Proof. Let us restrict to the setting $m = n$. The quantum strategy involves all players sharing a n-size cat state. Player i conditionally applies a phase to wire i of the cat state, based on input x_i. Then, all players apply a H gate to and measure their cat state wire. The measurement outcome is returned as the response. Here is the circuit version of the strategy:

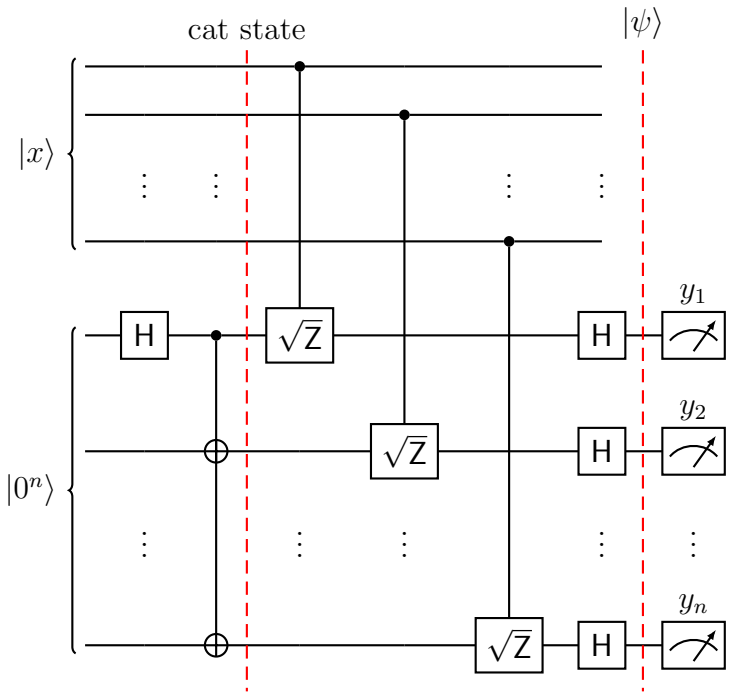

Figure 5.1: $\mathsf{QNC}^0_{\text{cat}}$ circuit for parity halving problem for $n = m$

Note the action of a layer of Hadamards on a cat state:

$$\mathsf{H}^{\otimes n}\left(\frac{|0^n\rangle + |1^n\rangle}{\sqrt{2}}\right) = |\text{even}_n\rangle \tag{5.4}$$

$$\mathsf{H}^{\otimes n}\left(\frac{|0^n\rangle - |1^n\rangle}{\sqrt{2}}\right) = |\text{odd}_n\rangle \tag{5.5}$$

where $|\text{even}_n\rangle$ ($|\text{odd}_n\rangle$) is a uniform superposition over all even (odd) bit strings of size n. Using this, we have that state $|\psi\rangle$ is given by:

$$|\psi\rangle = (\mathsf{H})^{\otimes n}\left(|0^n\rangle + i^{|x|}\,|1^n\rangle\right) = \begin{cases} |\text{even}_n\rangle & \text{if } |x| = 0 \bmod 4 \\ |\text{odd}_n\rangle & \text{if } |x| = 2 \bmod 4 \end{cases} \tag{5.6}$$

Therefore, the parity of y is 0 if $\frac{|x|}{2} = 0 \pmod 2$ and 1 if $\frac{|x|}{2} = 1 \pmod 2$, as desired. For $m > n$, we can use the same strategy but with a cat state of size m (with no controlled gates on the last $m-n$ wires) and verify that $|\psi\rangle$ will still satisfy Equation 5.6. Therefore, as long as $m \in \mathsf{poly}(n)$, there is a $\mathsf{QNC}^0_{\text{cat}}$ circuit that solves $\mathsf{PHP}_{n,m}$ exactly. □

Note that we are proving separations in the relational setting. This is because computing $|y|$ directly as a bit is not in $\mathsf{QNC}^0_{\text{cat}}$ as it requires one output qubit to depend on all input bits. In fact, $\mathsf{QNC}^0_{\text{cat}}$ cannot even compute parity (Theorem 6.19). Therefore, in general, to prove separations for $\mathsf{QNC}^0_{\text{cat}}$ in the decision setting, we must provide the

quantum circuit access to quantum fan-out (Chapter 3) or classical fan-out (Chapters 4 and 6).

1-locality NC^0 lower bound: [Mer90, BBT04] also provided a classical lower bound but for a very weak family of circuits – where each output can depend on only one input. For an output bit, we define its light cone as the input bits that can affect it. That is, the input bits to which there is a path from the output bit via gates. A circuit is l-local if the size of all output bit light cones is bounded by l. Therefore, the lower bound induced from the non-local game apply to 1-local classical circuits. Additionally, note that for a given input, we only care about deterministic strategies as a probabilistic strategy is a distribution over deterministic ones (we can choose the one that performs the best on a given input).

Theorem 5.3. *No 1-local NC^0 circuit can solve $\mathsf{PHP}_{n,n}$ with probability greater than $\frac{1}{2} + 2^{-\lceil \frac{n}{2} \rceil}$.*

Proof. We outline the proof strategy from [BBT04]. For player i, there are only four classical strategies as the size of the question and answer set is 2. Therefore, it can be defined by function $f(i,j) = \pm 1$ for $j \in \{0,1\}$. Player i responds with 0 on input j, if and only if $f(i,j) = 1$. Moreover, the parity of the response string is $\prod_{i=1}^{n} f(i,x_i)$ in the $\{\pm\}$ basis. As per [BBT04], we consider the expression:

$$c = \prod_j \left(f(j,0) + if(j,1)\right) \tag{5.7}$$

Expanding this out, we have a $i^{|x|}\prod_j f(j,x_j)$ term for every $x \in \{0,1\}^n$. Moreover, $i^{|x|} = 1$ if and only if $x = 0 \bmod 4$. Therefore, $\sum_x i^{|x|}\prod_j f(j,x_j)$ is equal to (number of input strings for which we won) - (number of input strings for which we lost). Each $(f(j,0) + if(j,1))$ is one of the fourth roots of unity with magnitude $\sqrt{2}$. When n is even, their product will be either $2^{\frac{n}{2}} \times \pm 1$ or $2^{\lfloor \frac{n}{2} \rfloor} \times \pm i$. While if n is odd, it is $2^{\frac{n}{2}}\left(\frac{\pm 1 \pm i}{\sqrt{2}}\right)$. In all cases, the real part of c is less than $2^{\lfloor \frac{n}{2} \rfloor}$. Moreover, the success probability can be easily computed using c to be bounded by $\frac{1}{2} + 2^{-\lceil \frac{n}{2} \rceil}$. □

Note that if some x_k was fixed,

$$c = i^{x_k} f(k,x_k) \prod_{j \in [n]-k} \left(f(j,0) + if(j,1)\right) \tag{5.8}$$

In both the even n and odd n cases, $i^{x_k} f(k,x_k)$ will not impact the real value of the product of the remaining terms, giving us:

Corollary 5.4. *No 1-local NC^0 circuit can solve $\mathsf{PHP}_{n,n}$ in which d of the input bits are fixed with probability greater than $\frac{1}{2} + 2^{-\lceil \frac{n-d}{2} \rceil}$.*

NC^0 lower bound: [WKST19] now used the locality of NC^0 circuits to show that if these circuits could solve $\mathsf{PHP}_{n,m}$ on a random input with probability p, then there is a strategy for the restricted version of parity halving non-local game with success probability p.

Theorem 5.5. *No l-local NC^0 circuit can solve the $\mathsf{PHP}_{n,m}$ outputs with probability greater than $\frac{1}{2} + 2^{-\Omega\left(\min\left(n, \frac{n^2}{l^2 m}\right)\right)}$.*

Proof sketch. Let there be a l-local NC^0 circuit that solves the $\mathsf{PHP}_{n,m}$ outputs with probability p on random input. The key idea is that due to the limited locality, there is always a set S of $\Omega\left(\min\left(n, \frac{n^2}{l^2 m}\right)\right)$ input bits that impact disjoint set of output bits. To see this, notice that each input bit is, on average, in the light cone of $\frac{lm}{n}$ output bits. Therefore, if we consider the graph where an input bit is connected to another input bit if they share an output bit in their light cones, the average degree of this graph will be $\frac{l^2 m}{n}$ as each output has at most l input bits in its light cone. Now an application of Turán's theorem gives us the required S.

To reduce this l-local circuit to a 1-local circuit, we fix the input bits that do not belong to S such that the circuit still succeeds with probability p on the restricted problem. We assign a player to every input bit in S. Each player can follow the NC^0 circuit to get the value of the output bits in their input bit's light cone without communicating with any other player. Any output bits not in the light cones of the players can be added to the output of any player's 1-local circuit as everyone knows the values of the fixed bits. Then, each player returns the parity of the output bits in their light cone. The parity of the player responses is the same as the parity of the original circuit's output. Therefore, this is a strategy for the non-local game that succeeds with probability p. Using Corollary 5.4 gives us the desired result. □

In the setting of NC^0 circuits, $l = O(1)$, as they have bounded fan-in. We have that, as long as $m < n^2$, the probability of success for NC^0 is exponentially close to $\frac{1}{2}$. Therefore, we have a window of separation: for $n \leq m < n^2$, $\mathsf{PHP}_{n,m}$ on a random input can be solved by a $\mathsf{QNC}^0_{\text{cat}}$ circuit exactly, but all NC^0 circuits can only solve it with probability at most $\frac{1}{2} + \exp(-n^\alpha)$ for some α.

AC^0 lower bound: To extend the above average-case NC^0 separation to AC^0, [WKST19] use a tool from classical circuit complexity, Håstad's switching lemma [Hås14, Ros17] which provides the following reduction. Informally, if an AC^0 circuit can solve $\mathsf{PHP}_{n,m}$ on a random input with probability p, then an NC^0 circuit can solve $\mathsf{PHP}_{n,m^{1+\epsilon}}$ for some small ϵ on a random input with probability p. The details are quite technical, so we omit them in this thesis. Using this, the above separation is ported to AC^0.

5.1.2 Relational MOD_3

As a step beyond AC^0 separations, [WKST19] also showed that through a similar construction to Figure 5.1, $\mathsf{QNC}^0_{\text{cat}}$ circuit can compute the relational version of unbounded fan-in MOD_3 with probability $\frac{3}{4}$:

Theorem 5.6. *For input $x \in \{0,1\}^n$, there exists a $\mathsf{QNC}^0_{\text{cat}}$ circuit such that its output $y \in \{0,1\}^n$ satisfies:*

$$\mathsf{MOD}_{2,n}(y) = \mathsf{MOD}_{3,n}(x) \tag{5.9}$$

with probability at least $\frac{3}{4}$.

Proof. Consider the following $\mathsf{QNC}^0_{\text{cat}}$ circuit:

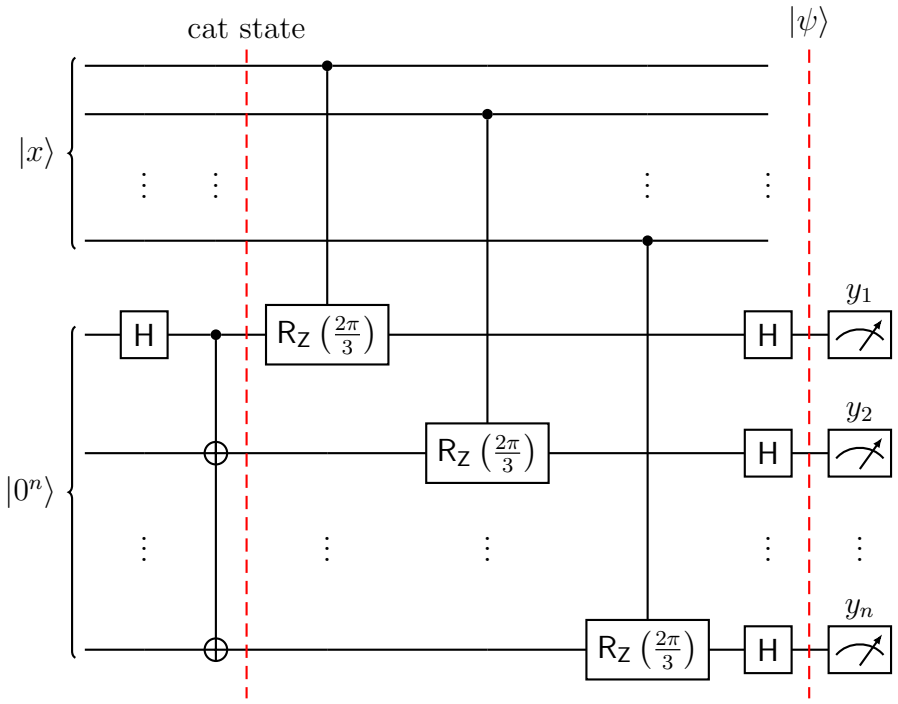

Figure 5.2: $\mathsf{QNC}^0_{\text{cat}}$ circuit for the relational version of MOD_3

We apply a controlled $\mathsf{R_Z}\left(\frac{2\pi}{3}\right)$ rotations with x_i as the control and wire i of the cat state as target:

$$|\psi\rangle = \mathsf{H}^{\otimes n}\left(\frac{1}{\sqrt{2}}\left(|0\rangle^n + \exp\left(\frac{2\pi i|x|}{3}\right)|1\rangle^n\right)\right) \tag{5.10}$$

Using Equations 5.4 and 5.5:

$$\mathsf{P}(\mathsf{MOD}_{2,n}(y)=0) = \left|\left(\frac{\langle 0|^n + \langle 1|^n}{\sqrt{2}}\right)\frac{1}{\sqrt{2}}\left(|0\rangle^n + \exp\left(\frac{2\pi i|x|}{3}\right)|1\rangle^n\right)\right|^2 = \cos^2\left(\frac{\pi|x|}{3}\right) \tag{5.11}$$

Therefore, this gives us:

- If $|x| = 0 \pmod 3$, $\mathsf{MOD}_{2,n}(y) = \mathsf{MOD}_{3,n}(x)$ with probability 1
- If $|x| = 1$ or $2 \pmod 3$, $\mathsf{MOD}_{2,n}(y) = \mathsf{MOD}_{3,n}(x)$ with probability $\frac{3}{4}$

□

This is interesting because if a (poly-size) $\mathsf{AC}^0[2]$ circuit could do the above task, then it could compute the decision version of MOD_3 with probability $\frac{3}{4}$ by adding a parity gate at the end. However, $\mathsf{AC}^0[2]$ circuits that compute MOD_3 with high probability must have exponential size [Smo87]. Therefore, giving us:

$$\mathsf{AC}^0[2] \not\subset \mathsf{QNC}^0_{\text{🙀}} \tag{5.12}$$

In fact, this setup is the inspiration behind the strictly interleaved models we define in Section 6.2.1.

5.2 Separations without quantum advice

Unfortunately, the cat states required in the above constructions cannot be constructed in QNC^0. If we could that, we would be able to implement quantum fan-out in constant-depth, which is not possible (Section 3.5.2). To port the above separations to QNC^0 (without advice), [WKST19] noted that even though we cannot create cat states, we can create $\frac{|z\rangle+|\bar{z}\rangle}{\sqrt{2}}$ for random $z \in \{0,1\}^n$ in constant depth:

- Initialize $2n-1$ qubits to $|0\rangle$ and apply H to the first n qubits:

$$\sum_{z\in\{0,1\}^n} |z\rangle |0\rangle^{n-1} \tag{5.13}$$

- Compute quantum parity between qubits i and $i+1$ with target qubit $n+i$ for $i \in [n-1]$:

$$\sum_{z\in\{0,1\}^n} \left(|z\rangle \prod_i^{n-1} |z_i \oplus z_{i+1}\rangle \right) \tag{5.14}$$

- Measure the last $n-1$ qubits and let d be the classical outcome.

Note that d uniquely identifies the $|z\rangle$ it was associated with up to bit flips on all qubits. To see this notice that given z_1, d_1 fixes z_2 and given z_2, d_2 fixes z_3 and so on. Therefore, after this procedure we are in state $\frac{|z\rangle+|\bar{z}\rangle}{\sqrt{2}}$ for some $z \in \{0,1\}^n$ that can be inferred from d. This procedure can be adapted to construct such states for other underlying topologies: in step two, we compute the quantum parity for two qubits (i,j) if they share an edge. For the rest of this section, we will generate $\frac{|z\rangle+|\bar{z}\rangle}{\sqrt{2}}$ based on a $\sqrt{n} \times \sqrt{n}$ two-dimensional grid. Note that now d is of size $2n - 2\sqrt{n}$.

Consider what happens if we use this state instead of the cat state in Figure 5.1. Before the layer of Hadamards, we are in state:

$$\frac{i^{\sum_i x_i z_i} |z\rangle + i^{\sum_i x_i \bar{z}_i} |\bar{z}\rangle}{\sqrt{2}} \tag{5.15}$$

Noting that $x_i\bar{z}_i = x_i - x_i z_i$, taking out the global phase, and after applying the Hadamards:

$$|\psi\rangle = \mathsf{H}^{\otimes n}\left(\frac{|z\rangle + (-1)^{\frac{|x|}{2}-\sum_i x_i z_i}|\bar{z}\rangle}{\sqrt{2}}\right) \tag{5.16}$$

$$= 2^{-\frac{n}{2}}\sum_{y\in\{0,1\}^n}\left(\frac{(-1)^{\sum_i y_i z_i} + (-1)^{\sum_i y_i \bar{z}_i + \frac{|x|}{2} - \sum_i x_i z_i}}{\sqrt{2}}\right)|y\rangle \tag{5.17}$$

Note that $|y\rangle$ survives only if the two coefficients are equal (mod 2). Therefore, the string y always satisfies:

$$\sum_i y_i z_i = \sum_i y_i (1 - z_i) + \frac{|x|}{2} - \sum_i x_i z_i \pmod 2 \tag{5.18}$$

$$\implies |y| = \frac{|x|}{2} - \sum_i x_i z_i \pmod 2 \tag{5.19}$$

Using this, we can define the relaxed parity halving problem on the two-dimensional grid that a QNC^0 circuit can solve:

Definition 5.7 (Relaxed parity halving problem on the 2D grid)**.** *For a given $x \in \{0,1\}^n$ with even parity, to solve the relaxed parity halving problem on the two-dimensional grid we must return $y \in \{0,1\}^n$ and $d \in \{0,1\}^{2n-2\sqrt{n}}$ such that there is a $z \in \{0,1\}^n$ that satisfies:*

$$|y| = \frac{|x|}{2} - \sum_i x_i z_i \ (mod\ 2) \tag{5.20}$$

and that for an edge (i,j) in the two-dimensional grid, $d_{i,j} = z_i \oplus z_j$.

Now all that remains to show is that this problem cannot be solved by an AC^0 circuit on a random input. We do so by showing a reduction from $\mathsf{PHP}_{n,O(n^{\frac{3}{2}})}$, which we know is hard for AC^0 circuits (Section 5.1.1), to the relaxed 2D grid version.

Theorem 5.8. $\mathsf{PHP}_{n,O(n^{\frac{3}{2}})}$ *is NC^0-reducible to the relaxed parity halving problem on the 2D grid.*

Proof. We want to solve $\mathsf{PHP}_{n,O(n^{\frac{3}{2}})}$ on some input $x \in \{0,1\}^n$. Let us provide x as input to our oracle that solves the relaxed parity halving problem on the 2D grid instead. By definition, we receive $y \in \{0,1\}^n$ and $d \in \{0,1\}^{2n-2\sqrt{n}}$ such that there is a $z \in \{0,1\}^n$ that satisfies:

$$|y| = \frac{|x|}{2} - \sum_i x_i z_i \pmod 2 \tag{5.21}$$

and that for an edge (i,j) in the two-dimensional grid, $d_{i,j} = z_i \oplus z_j$. The idea is to use y and d to construct a string y' whose parity is $\sum_i x_i z_i \pmod 2$ as then we can send y along

with y' and we know that $|y + y'| = \frac{|x|}{2} \pmod 2$. To do this, we first notice that $\bar{z}$ must satisfy Equation 5.21 as well. So without loss of generality, we can assume that $z_1 = 0$. Given this, we can compute z_i by the parity of the elements of d that trace a path from z_1 to z_i in the 2D grid. Moreover, note that in the 2D grid, the path length between any two bits is at most $O(\sqrt{n})$. Therefore, we can write:

$$\sum_i x_i z_i = \sum_i \sum_{j \in D_i} x_i d_j \pmod 2 \tag{5.22}$$

where each D_i has size at most $O(\sqrt{n})$. As we have access to x and d, with unbounded classical fan-out, we can compute each $x_i d$ term in parallel as they depend on only two bits. If some d_j appears with multiple x_i terms, we can use unbounded classical fan-out first to copy it n times. We return these $O(n^{\frac{3}{2}})$ bits as y' along with y, thereby solving $\mathsf{PHP}_{n,O(n^{\frac{3}{2}})}$. □

Chapter 6

Interleaved constant-depth quantum-classical circuits

Inspired by the power of adaptive-QNC^0 circuits in Chapter 4, we begin this chapter by defining a general model of computation where quantum and classical circuits interleave with each other and prove some basic properties of such models. Then, we restrict ourselves to the setting of Chapter 4, where we interleave QNC^0 and $\oplus_\mathsf{f}$ circuits, and rephrase Theorem 4.3 under this new framework. We generalize the result to arbitrary poly-size quantum circuits and discuss some existing applications.

Next, we focus on the experimentally more feasible model of interleaving, where all the qubits of the quantum circuit must be measured before interleaving with classical computation. To enable parallelization, inspired by the fan-out constructions in Section 5.1, we allow cat state advice to QNC^0. We then "de-fanout" certain constructions of Chapter 3 to show that strictly interleaved $\mathsf{QNC}^0_{\bowtie}$ and $\oplus_\mathsf{f}$ circuits can simulate TC^0 circuits. Using a protocol by Slote [Slo22], we show that we can simulate TC^0 even if the quantum device is untrusted if we strictly interleave it with $\mathsf{AC}^0[2]$ circuits instead.

Note that we are considering the power of quantum-classical circuits to reduce the *quantum* circuit depth needed for a particular task. Similarly, works such as [BSS16, DGC18, PHOW20] have earlier considered the power of quantum-classical models with the aim of reducing the *size* of the quantum circuit.

6.1 Weak interleaving

When considering poly-size quantum circuits, it does not make sense to provide them access to efficient classical computation as the quantum device can simulate it. However, in the near term, where we may be able to implement only constant-depth quantum circuits, there may be operations that these quantum devices cannot perform, but an efficient classical device can. Therefore, given the ubiquity of efficient classical computation, in the near-term, we care about the power of QNC^0 interacting with a BPP machine.

In Chapter 4, we have already seen an advantage such interaction provides: with classical fan-out and parity gates, constant-depth quantum circuits can perform the powerful quantum fan-out gate, which they cannot on their own. Based on this construction, let us formally define the weakly interleaved circuit model as follows:

Definition 6.1 (Weakly interleaved circuits with d layers)**.** *For two circuit classes,* A *and* B*,* $[\mathsf{A};\mathsf{B}]^d$ *consists of circuits with the following action on some input* x_1 *on* n *bits or quantum state* $|x_1\rangle$ *on* n *qubits (depending on whether* A *is classical or quantum):*

- *Let* $\phi_0 = \psi_0$ *be initialized empty.*
- *For* $i \in [d-1]$*:*
  - *Apply* $A_i \in \mathsf{A}$ *on* $x_i 0^{m_i} \psi_{i-1}$ *(*$|x_i 0^{m_i}\rangle |\psi_{i-1}\rangle$*) for* $m_i \in \mathsf{poly}(|x_i|)$ *where* A *is a classical (quantum) circuit class. Measure a subset of bits (qubits). Let* y_i *denote the classical outcome and let* ψ_i *(*$|\psi_i\rangle$*) denote the state of the unmeasured bits (qubits).*
  - *Apply* $B_i \in \mathsf{B}$ *on* $y_i 0^{k_i} \phi_{i-1}$ *(*$\left|y_i 0^{k_i}\right\rangle |\phi_{i-1}\rangle$*) for* $k_i \in \mathsf{poly}(|y_i|)$ *where* B *is a classical (quantum) circuit class. Measure a subset of bits (qubits). Let* x_{i+1} *denote the classical outcome and let* ϕ_i *(*$|\phi_i\rangle$*) denote the state of the unmeasured bits (qubits).*
- *Apply* $A_d \in \mathsf{A}$ *on* $x_d 0^{m_d} \psi_{d-1}$ *(*$|x_d 0^{m_d}\rangle |\psi_{d-1}\rangle$*) for* $m_d \in \mathsf{poly}(|x_d|)$ *where* A *is a classical (quantum) circuit class. Let* y_d *(*$|y_d\rangle$*) be the output of the interleaved circuit.*

In essence, we can think of first applying a circuit from A, measuring some of the output and then applying a circuit from B on the measurement outcomes of that circuit and so on. Note that at every step, both models are permitted to use a polynomial number (in the size of the input) of auxiliary qubits. Moreover, crucially, both types of circuits have *memory* — in round k, when we apply B_k on the measurement outcome of the circuit from A_{k-1}, B_k also has as input access to the post-measurement output state of the previous round.

We can imagine this setting as there being two players: Alice with the computational power of A and Bob with the computational power of B. The initial input is given to Alice, who applies A_1, measures and shares x_1 with Bob through a *classical* communication channel. Critically, when Bob replies with y_2, Alice applies A_2 on the product state of x_2 and the output state of the previous circuit A_1. Note that if A was quantum, this would imply that Alice would have to maintain coherence in the post-measurement state after the first measurement while Bob is performing its circuit. Therefore, Alice must have access to quantum memory. Here is a diagram depicting the model:

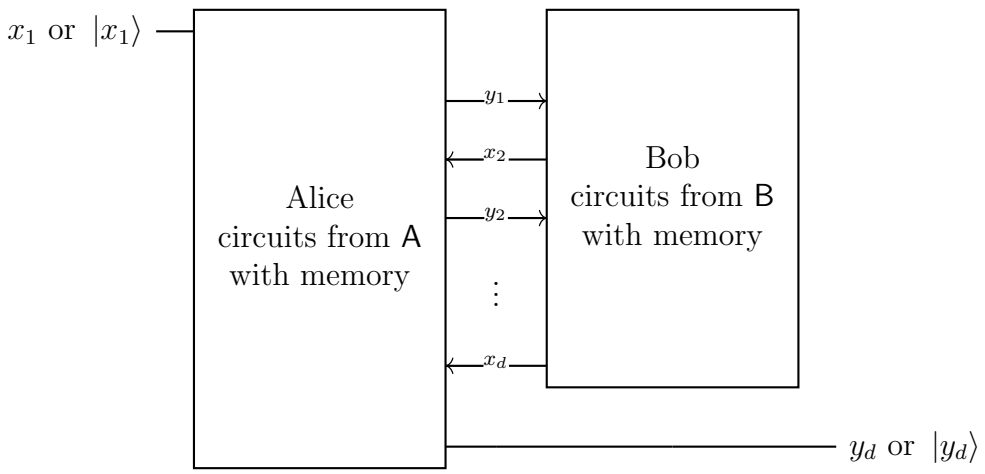

Figure 6.1: $[\mathsf{A};\mathsf{B}]^d$: Circuit classes A and B weakly interleaved with d layers

Before showing how this model can simulate different classes of quantum circuits, we establish some basic observations.

Observation 6.2. *With circuit classes* Q *(quantum) and* C *(classical), and where the '+' symbol means "followed by", we have:*

(i) $(QC_1 \in [\mathsf{Q};\mathsf{C}]^n + QC_2 \in [\mathsf{Q};\mathsf{C}]^m) \in [\mathsf{Q};\mathsf{C}]^{n+m+1}$

(ii) $(QC_1 \in [\mathsf{Q};\mathsf{C}]^n + (Q_1 \in \mathsf{Q} \text{ or } C_1 \in \mathsf{C})) \in [\mathsf{Q};\mathsf{C}]^{n+1}$

(iii) $((Q_1 \in \mathsf{Q} \text{ or } C_1 \in \mathsf{C}) + QC_1 \in [\mathsf{Q};\mathsf{C}]^n) \in [\mathsf{Q};\mathsf{C}]^{n+1}$

Proof. (i) Let the variables for QC_1 be $\{x, y\}$ and the variables for QC_2 be $\{x', y'\}$. We can simply let Alice save $|y_n\rangle$ and set $y_n = x_{n+1}$ to empty. Bob does nothing. Now, Alice performs the operation specified by the first circuit of QC_2 on $|y_n\rangle$ instead of x_{n+1} and then we continue as usual according to QC_2. We added one extra interaction; hence, the total number of layers now is $n + m + 1$.

(ii)-(iii) Similarly, if we follow QC_1 with Q_1, we can let Alice save $|y_n\rangle$ and set $y_n = x_{n+1}$ to empty. Bob does nothing. Alice applies Q_1 to $|y_n\rangle$. If we follow QC_1 with C_1, Bob performs circuit C_1 using y_n as the input and sends the output to Alice as x_{n+1} who does nothing and reports the output $y_{n+1} = x_{n+1}$. Hence, the number of layers has increased by 1. Similar arguments work for the third observation. □

6.1.1 Simulating poly-size quantum circuits using $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^d$

Now let us restrict our attention to the setting associated with the teleportation-based schemes for implementing Clifford circuits from Chapter 4. In particular, we let $\mathsf{A} = \mathsf{QNC}^0$

and $\mathsf{B} = \oplus_{\mathsf{f}}$. Note that in this case, as B is classical, Bob does not need to have memory. In each round, it can simply send all of its output bits to Alice, and if Bob needs them for a subsequent $\oplus_f$ circuit, Alice can send the relevant bits as part of the interaction message. Therefore, the $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^d$ model reduces to QNC^0 having access to $d-1$ calls to a $\oplus_{\mathsf{f}}$ oracle. In particular, if the number of layers d is a constant, we have that:

$$[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^{O(1)} = \left(\mathsf{QNC}^0\right)^{\oplus_{\mathsf{f}}} \tag{6.1}$$

Here is a circuit that belongs to $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$:

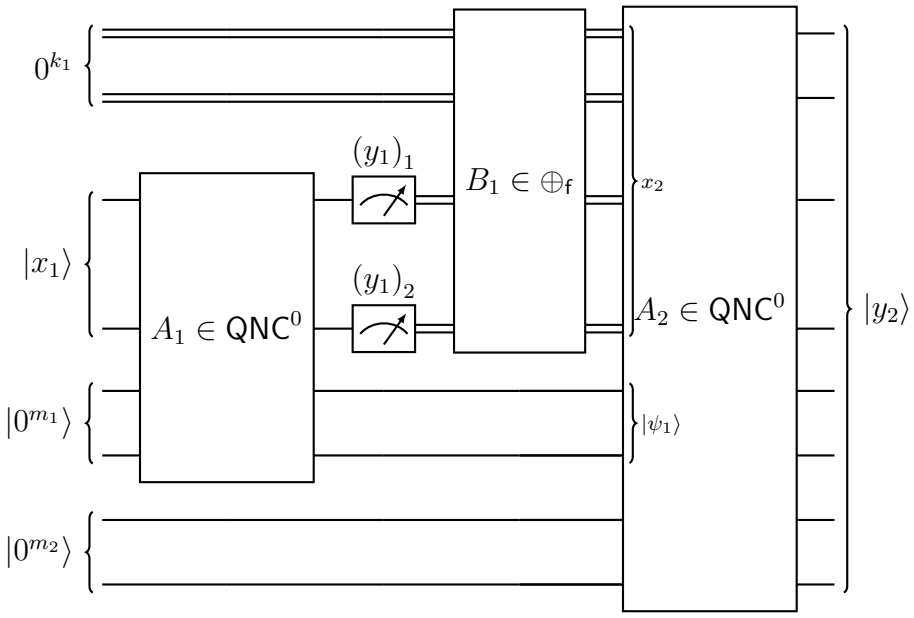

Figure 6.2: Structure of a circuit from $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$

Variants of this model have been considered in earlier work. In particular:

- Our model is most similar to the $(n, O(1), O(\log(n)), O(1), O(\mathsf{poly}(n)))$-jozsa-quantum circuit family from [CM20]. These are poly-size constant-depth quantum circuits interleaved with log-depth classical circuits, a constant number of times. Now as $\oplus_{\mathsf{f}} \subset \mathsf{NC}^1$, $[\mathsf{QNC}^0, \oplus_{\mathsf{f}}]^{O(1)}$ is a subset of this circuit family.
- Models where quantum circuits are interleaved with poly-depth classical computation have also been considered with different restrictions on the quantum depth. In 2004, [TD02] defined adaptive quantum computation QC_{ad} based on [GC99] as this model with poly-size quantum circuits. Later, [Joz06] considered this model with to poly-size log-depth quantum circuits. More recently, the $O(1)$-QC circuit family from [CCL20] refers to the same model restricted to constant-depth quantum circuits. Our model, $[\mathsf{QNC}^0, \oplus_{\mathsf{f}}]^{O(1)}$ is a subset of all these families.

Now, we can finally rephrase Theorem 4.3 (restated below) under the framework of interleaved models.

Theorem 6.3 (Theorem 4.3 restated)**.** poly*-size Clifford circuits can be simulated by adaptive constant-depth quantum circuits – in particular, by* poly*-size constant-depth quantum circuits interleaved with a constant-depth* $\oplus_{\mathsf{f}}$ *circuit.*

Using this, we directly get that:

Corollary 6.4. *Any* poly*-size Clifford circuit can be simulated by a* $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$ *circuit.*

By comparing the teleportation-based construction in Figure 4.5 with the structure of a $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$ circuit in Figure 6.2, we note that all Clifford gates are applied in parallel using circuit $A_1 \in \mathsf{QNC}^0$. Then, we apply Pauli corrections through circuit $A_2 \in \mathsf{QNC}^0$ based on measurement outcomes $\{y_1\}$. Circuit $B_1 \in \oplus_{\mathsf{f}}$ determines which Pauli gates to apply.

Moreover, as Clifford + T is a universal gate set, this gives us a more general statement for any quantum poly-size quantum circuit:

Corollary 6.5. *Any* poly*-size circuit with* T*-depth* t *can be simulated by* $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^{O(t)}$

Proof. Let Q be a poly-size circuit with T-depth t. From the definition of T-depth, we can write Q as:

$$Q = C_{t+1}T_tC_tT_{t-1}C_{t-1}\dots T_2C_2T_1C_1 \tag{6.2}$$

where $\{C_i\}$ are poly-size Clifford circuits and $\{T_i\}$ is a single layer consisting of T gates only. Crucially, for all $i \in [t]$, T_i is a circuit in QNC^0 as it only consists of a single layer of gates. For every C_i with $i \in [t+1]$, we can simulate its action by a $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$ circuit by Corollary 6.4. Therefore, Q can be written as a $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$ operation followed by QNC^0 circuit followed by a $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^2$ and so on, $O(t)$ times. Therefore, by using Observation 6.2(*ii*)-(*iii*) repeatedly, we have that Q can be simulated by a circuit from $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^{O(t)}$. □

6.1.2 Simulating TC^0 and approximate quantum Fourier transform using $[\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^{O(1)}$

We use the results from the last section along with the fact that quantum fan-out is a Clifford operation:

Corollary 6.6. $\mathsf{QNC}^0_{\mathsf{f}} \subseteq [\mathsf{QNC}^0; \oplus_{\mathsf{f}}]^{O(1)}$

Proof. Every quantum fan-out in $\mathsf{QNC}^0_{\mathsf{f}}$ can be replaced by a log-depth circuit consisting of just Clifford gates. Therefore, the "quantum fan-out free" circuit is still at most log-depth with constant T-depth. We apply Corollary 6.5 to this circuit with $t = O(1)$. □

Now, by combining this with $\mathsf{QTC}^0_\mathsf{f} \subseteq \mathsf{QNC}^0_\mathsf{f}$ (Theorem 3.7), we get:

Corollary 6.7. $\mathsf{TC}^0 \subseteq \mathsf{QTC}^0_\mathsf{f} \subseteq [\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$

This is an example where $[\mathsf{A}; \mathsf{B}]^{O(1)}$ can simulate circuits that A and B cannot on their own. We will see the proof of why $\mathsf{TC}^0 \not\subset \mathsf{QNC}^0$ in the next section (Theorem 6.19) and shows us a log-depth separation between interleaved and non-interleaved QNC^0. Moreover, the flagship application of such interleaved circuits is definitely the approximate quantum Fourier transform (QFT):

Corollary 6.8. *Approximate quantum Fourier transform can be simulated by a circuit in* $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$.

Approximate QFT is the *quantum* subroutine of Shor's algorithm. That is, given this result, we can implement Shor's algorithm using a classical computer with oracle access to a $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$ machine in polynomial time! Corollary 6.8 follows as [HŠ05] showed that $\mathsf{QNC}^0_\mathsf{f}$ circuits can simulate approximate quantum Fourier transform with arbitrary moduli. Here, we briefly discuss their construction for the case where the modulus is 2^n. For a given quantum basis state $|x\rangle$, we want to compute the state $|\Psi_x\rangle = \frac{1}{\sqrt{2^n}} \sum_y^{2^n-1} \omega^{xy} |y\rangle$ where ω is the 2^n-root of unity. [CW00] provide a log-depth quantum circuit for approximate quantum Fourier transform, which consists of two key steps. Using quantum fan-out, both of these steps can be compressed to constant depth:

- **Step 1:** Construct the state $|x\rangle |\Psi_x\rangle^{O(\log(n))}$. It can be shown that $|x\rangle |\Psi_x\rangle$ can be constructed using n^2 commuting operators acting on $|x\rangle |0\rangle$. Using the parallelization trick of fan-out constructions (Lemma 3.2), we can construct each $|\Psi_x\rangle$ in constant-depth with unbounded quantum fan-out. Moreover, we can first fan-out $|x\rangle$ $O(\log(n))$ times to compute $O(\log(n))$ copies of $|\Psi_x\rangle$ and then uncompute all the extra $|x\rangle$ copies.
- **Step 2:** Use $|\Psi_x\rangle^{O(\log(n))} |0\rangle$ to construct $|\Psi_x\rangle^{O(\log(n))} |x\rangle$ with high probability and then use the constructed $|x\rangle$ to "uncompute" the original $|x\rangle$. $|x\rangle$ can be estimated by measuring $|\Psi_x\rangle$ in different bases. However, to approximate $|x\rangle$ with error at most ϵ, we need to do this procedure on $\log\left(\frac{n}{\epsilon}\right)$ copies of $|\Psi_x\rangle$ and then take a majority vote. With unbounded quantum fan-out, we can compute the majority function in constant depth (Theorem 3.7) and by setting $\epsilon = \frac{1}{\mathsf{poly}(n)}$, $O(\log(n))$ copies suffice.

Other applications of weakly interleaved QNC^0 circuits:

- **Cryptographic applications:** [BCM+21] introduced cryptographic tests of quantumness: interactive protocols that can be passed by an efficient quantum device but by no efficient classical device under cryptographic assumptions. Moreover, in contrast to circuit sampling tasks, it can be efficiently determined whether a device passed the protocol based on the device's responses. Such tasks have been used as a primitive for many cryptographic applications. [HG21b, LG21] independently show

that $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$ circuits can pass these tests of quantumness. This is shown by noting that the task requires the quantum device to perform a NC^0 computation (evaluating the randomized encodings of trapdoor claw-free functions) which can be computed in $\mathsf{QNC}^0_\mathsf{f}$. This, therefore, has applications in verifying the quantumness of current experimental setups.

As another application of $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$, [BY22] shows that another cryptographic primitive, generating randomized encodings for quantum circuits, can be performed by $\mathsf{QNC}^0_\mathsf{f}$ circuits. For quantum circuit C, a randomized function f is a randomized encoding of it, if on some input x, Cx can be inferred from $f(x)$ but $f(x)$ provides no more information about C or x.

- **Generating long-range entanglement:** Given the inherent locality of QNC^0 circuits, long-range entangled states such as cat states, those induced by conformal field theory and those exhibiting various forms of topological order cannot be generated by QNC^0 circuits. However, if these states can be prepared using constant T-depth, they admit interleaved QNC^0 realizations using Corollary 6.5. [LLKH22, TVV22] review and provide additional examples of using interleaved QNC^0 to generate such states by exploiting intermediate measurements. Moreover, building on these results [BKKK22] shows that all operations required for topological quantum computation based on non-abelian anyons can be performed using interleaved QNC^0 circuits.

6.2 Strict interleaving

Even though $[\mathsf{QNC}^0; \oplus_\mathsf{f}]^{O(1)}$ is constant-depth *quantumly*, its implementation poses serious experimental challenges. In particular, the model requires the ability to measure only a certain subset of qubits and keep the unmeasured qubits coherently stored while the $\oplus_f$ circuit runs. Such subset measurements and robust quantum memories may not be available in near-term experiments [GDC+20] – intuitively, one reason for this is because measurements on a subset may disturb the remaining qubits. However, note that some recent experiments have demonstrated this ability at a small scale using trapped ions [WKE+19, NNZ+21, PDF+21, ZKML+21], superconducting circuits [CTI+21, RRG+22], neutral atoms [DLH+22, SBA+22], and silicon quantum dots [OCTW+22].

To mitigate these concerns, we consider the computational power of a restriction of the above model that is experimentally easier to implement. In this section, we consider *strictly* interleaving circuits, that is, circuits in which one must measure all qubits before interleaving with the classical circuit. Formally:

Definition 6.9 (Strictly interleaved circuits with d layers)**.** *For two circuit classes,* A *and* B, $[\mathsf{A}, \mathsf{B}]^d$ *consists of circuits with the following action on some input* x_1 *on* n *bits or quantum state* $|x_1\rangle$ *on* n *qubits (depending on whether* A *is classical or quantum):*

- *For* $i \in [d-1]$*:*

– *Apply* $A_i \in \mathsf{A}$ *on* $x_i 0^{m_i}$ $(|x_i 0^{m_i}\rangle)$ *for* $m_i \in \mathsf{poly}(|x_i|)$ *where* A *is a classical (quantum) circuit class. Measure all bits (qubits) and let* y_i *denote the classical outcome.*

– *Apply* $B_i \in \mathsf{B}$ *on* $y_i 0^{k_i}$ $(|y_i 0^{k_i}\rangle)$ *for* $k_i \in \mathsf{poly}(|y_i|)$ *where* B *is a classical (quantum) circuit class. Measure all bits (qubits) and let* x_{i+1} *denote the classical outcome.*

- *Apply* $A_d \in \mathsf{A}$ *on* $x_d 0^{m_d}$ $(|x_d 0^{m_d}\rangle)$ *for* $m_d \in \mathsf{poly}(|x_d|)$ *where* A *is a classical (quantum) class. Let* y_d $(|y_d\rangle)$ *be the output of the interleaved circuit.*

Firstly, we can notice that for two circuit classes A and B:

$$[\mathsf{A}, \mathsf{B}] \subseteq [\mathsf{A}; \mathsf{B}] \tag{6.3}$$

This model is the same two-player protocol we discussed in Figure 6.1 with the restriction that players do not have memory between the different rounds. The observations we noted earlier all pass down to the *strict* interleaving model if the input and output of the interleaved circuits are classical:

Observation 6.10. *With complexity classes* A *and* B*, if all inputs and outputs are classical, and where the '+' symbol means "followed by", we have:*

(i) $(AB_1 \in [\mathsf{A}, \mathsf{B}]^n + AB_2 \in [\mathsf{A}, \mathsf{B}]^m) \in [\mathsf{A}, \mathsf{B}]^{n+m+1}$

(ii) $(AB_1 \in [\mathsf{A}, \mathsf{B}]^n + (A_1 \in \mathsf{A}) \text{ or } (B_1 \in \mathsf{B}))) \in [\mathsf{A}, \mathsf{B}]^{n+1}$

(iii) $((A_1 \in \mathsf{A}) \text{ or } (B_1 \in \mathsf{B})) + AB_1 \in [\mathsf{A}, \mathsf{B}]^n) \in [\mathsf{A}, \mathsf{B}]^{n+1}$

(iv) $(AB_1 \in [\mathsf{A}, \mathsf{B}]^n) \in [\mathsf{B}, \mathsf{A}]^{n+1}$

Proof. (i) Let the variables for AB_1 be $\{x, y\}$ and the variables for AB_2 be $\{x', y'\}$. We let Bob accept y_n, and send it back to Alice with $x_{n+1} = y_n$ (identity circuit). Now, Alice and Bob perform the operation specified by the circuits of AB_2. We increased the number of interactions by 1; hence, the number of layers now is $n + m + 1$.

(ii)-(iii) Similarly, if we follow AB_1 with A_1, then Bob accepts y_n and sends it to Alice who performs A_1 on it. If we follow AB_1 with B_1, Bob performs B_1 with y_n as the input and sends the output to Alice as x_{n+1} who reports the output $y_{n+1} = x_{n+1}$. Hence, the number of layers has increased by 1. Similar arguments work for the third observation.

(iv) We want Bob to take the input and give the output. Therefore, identity circuits at the beginning and the end suffice. □

The circuit for $[\mathsf{A}, \mathsf{B}]^d$ with a classical input can be expressed as:

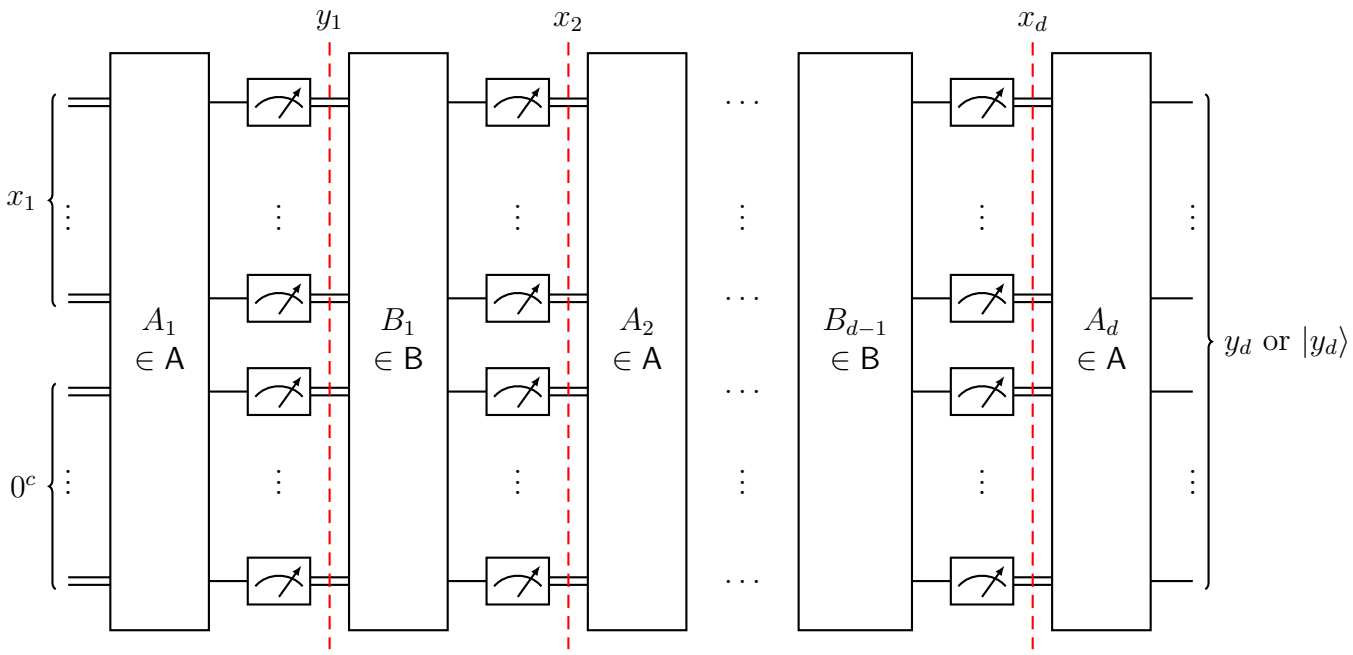

Figure 6.3: Structure of a circuit from $[\mathsf{A}, \mathsf{B}]^d$

where we have provided the auxiliary bits of size $c \in \mathsf{poly}(|x_1|)$ for all circuits at the start.

6.2.1 Computing MOD_3 in $[\mathsf{QNC}^0_{\text{cat}}, \oplus]^2$

Section 5.1 provides a protocol that can be rephrased as solving a decision problem using a strictly interleaved circuit. We restrict to the relevant case, $\mathsf{A} = \mathsf{QNC}^0_{\text{cat}}$, $\mathsf{B} = \oplus_{\mathsf{f}}$ and $d = O(1)$. Related models have been considered in the literature:

- $(n, O(1), O(\log(n)), O(1), O(\mathsf{poly}(n)))$-hybrid-quantum circuits from [CM20] are poly-size constant-depth quantum circuits strictly interleaved with log-depth classical circuits, a constant number of times. $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$ is slightly stronger than this model as we allow cat state quantum advice.
- [CCL20] considers $O(1)$-CQ circuits which are poly-size classical circuits with oracle access to QNC^0 circuits. The relation between $O(1)$-CQ and our model is unclear as the permitted resources seem incomparable.

Furthermore, it is interesting to note that in the relativized setting, the strictly interleaving $O(1)$-CQ and the weakly interleaving $O(1)$-QC have been proven to be incomparable [AGS22]. However, in our setting, as we restrict to low-depth classical computation, we have an inclusion relation between the weak and strict interleaving models (Equation 6.3). Now consider the following theorem regarding relational $\mathsf{MOD}_{3,n}$ from Section 5.1:

Theorem 6.11 (Theorem 5.6 restated)**.** *For input $x \in \{0,1\}^n$, there exists a $\mathsf{QNC}^0_{\text{cat}}$ circuit such that its output $y \in \{0,1\}^n$ satisfies:*

$$\mathsf{MOD}_{2,n}(y) = \mathsf{MOD}_{3,n}(x) \tag{6.4}$$

with probability at least $\frac{3}{4}$.

Note that now we can convert this directly to a decision problem by applying a parity gate to y:

Corollary 6.12. *A circuit $[\mathsf{QNC}^0_{\text{cat}}, \oplus]^2$ solves $\mathsf{MOD}_{3,n}$ with probability at least $\frac{3}{4}$.*

Moreover, the construction in Section 5.1 from [WKST19] provides us with an intuitive understanding of the role of the cat state and the classical parity gate. The output must depend on all inputs if we want to solve decision problems beyond NC^0. No one qubit can interact with all inputs through a QNC^0 circuit but a large cat state can interact with all inputs in parallel. However, the solution to the decision problem is now stored in the state of n maximally entangled qubits. We use measurements followed by a parity gate to reduce this state to a single bit.

6.2.2 Simulating TC^0 circuits using $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$

The above construction only uses a single cat state and a single classical parity gate. If we permit a constant number of cat states and $\oplus_{\mathsf{f}}$ circuits, we can use techniques from the fan-out constructions to compute the classical threshold gate:

Theorem 6.13. $\mathsf{TC}^0 \subseteq [\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$

In the rest of this section, we first provide the proof for Theorem 6.13 and then consider some consequences of the result. We will be "de-fanouting" the circuit constructions from Chapter 3. To begin doing so, first, we need to show an analogous "parallelization lemma" (corresponding to Lemma 3.2) and a construction for unbounded fan-in EX for the $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$ setting:

Lemma 6.14 (Parallelization lemma for strictly interleaved circuits)**.** *Consider a circuit on n data qubits and one auxiliary qubit $|0\rangle$. For $i \in [n]$, let $\mathsf{CR}_{\mathsf{X}}(\theta_i)$ apply to the auxiliary qubit and be controlled by the data qubit i. Then, there is a circuit in $[\mathsf{QNC}^0_{\text{cat}}, \oplus]^{O(1)}$ that can simulate the action of applying $\{\mathsf{CR}_{\mathsf{X}}(\theta_i)\}_{i\in[n]}$ and measuring the auxiliary qubit.*

Proof. Note that naively implementing $\{\mathsf{CR}_{\mathsf{X}}(\theta_i)\}_{i\in[n]}$ will have depth n for some input state $|x\rangle$:

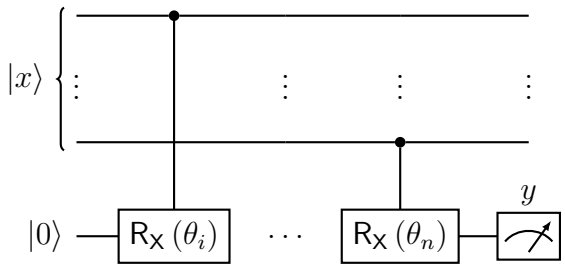

Figure 6.4: n-depth $\{\mathsf{CR_X}(\theta_i)\}_{i\in[n]}$ implementation

where $y \in \{0,1\}$ is the measurement outcome. As the operators commute with each other and H diagonalizes X gates into Z gates, we can apply Lemma 3.2 to generate the following equivalent circuit:

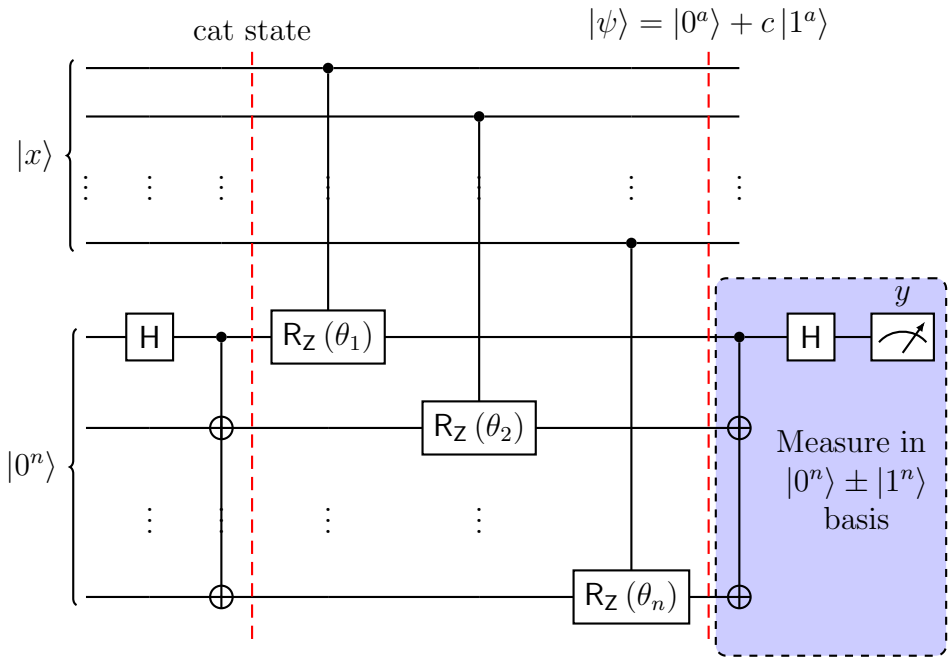

Figure 6.5: Constant-depth $\{\mathsf{CR_X}(\theta_i)\}_{i\in[n]}$ implementation using quantum fan-out

Note that the auxiliary qubits' state after the first fan-out is a cat state of size n, which we have access to. Now, if we consider the state just before the second fan-out, it is of the form $\frac{|0^n\rangle + c|1^n\rangle}{\sqrt{2}}$ for some constant c as all gates are diagonal. In particular, we have that:

$$\mathsf{P}(y=0) = \frac{1}{2}\left|\langle 0|\,\mathsf{H}\,\mathsf{F}_n\left(|0^n\rangle + c\,|1^n\rangle\right)\right|^2 \tag{6.5}$$

$$= \frac{1}{2}\left|\langle 0|\,\mathsf{H}\;\left(|0\rangle + c\,|1\rangle\right)\right|^2 \tag{6.6}$$

$$= \frac{1}{4}\left|1+c\right|^2 \tag{6.7}$$

Therefore, y's statistics reflect where we were in the Bloch sphere spanned by basis states $|0\rangle^n \pm |1\rangle^n$. Therefore, we can use:

$$\mathsf{H}^{\otimes n}\left(\frac{|0^n\rangle + |1^n\rangle}{\sqrt{2}}\right) = |\text{even}_n\rangle \tag{6.8}$$

$$\mathsf{H}^{\otimes n}\left(\frac{|0^n\rangle - |1^n\rangle}{\sqrt{2}}\right) = |\text{odd}_n\rangle \tag{6.9}$$

and parity gates to simulate y's statistics. In particular, we employ the following gadget:

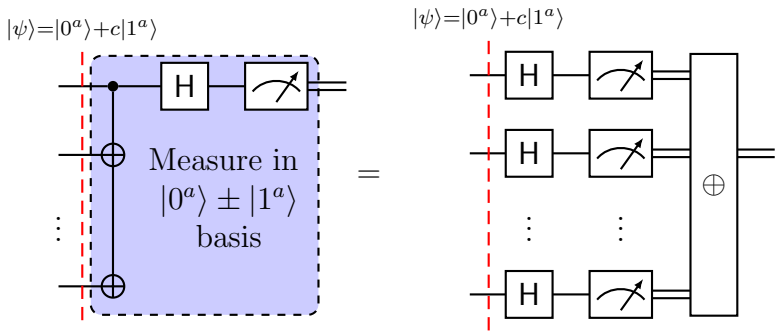

Figure 6.6: Measuring in the cat state basis using a classical parity gate

to construct the following interleaved circuit:

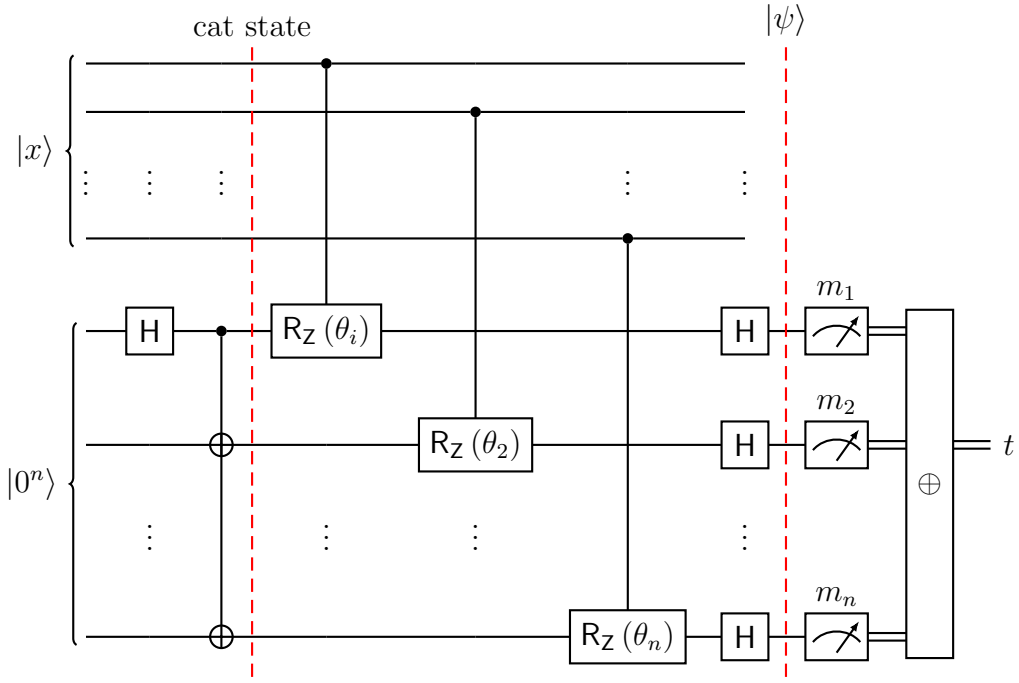

Figure 6.7: $\{\mathsf{CR}_{\mathsf{X}\theta_i}\}_{i\in[n]}$ using a cat state and a parity gate

where $m \in \{0,1\}^n$ are the measurement outcomes. We can verify the construction by an explicit calculation. The state $|\psi\rangle$ is given by:

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(\mathsf{H}^{\otimes n}\left(|0^n\rangle + c\,|1^n\rangle\right)\right) \tag{6.10}$$

$$= \frac{1}{2}\left((1+c)\,|\text{even}_n\rangle + (1-c)\,|\text{odd}_n\rangle\right) \tag{6.11}$$

Therefore, we have that:

$$\mathsf{P}(t=0) = \mathsf{P}(\mathsf{MOD}_{2,n}(m)=0) = \frac{1}{4}|1+c|^2 = \mathsf{P}(y=0) \tag{6.12}$$

□

Theorem 6.15. *Unbounded fan-in* EX *can be simulated in* $[\mathsf{QNC}^0_{\text{🐱}}, \oplus_{\mathsf{f}}]^{O(1)}$

Proof. The proof follows the same strategy as the quantum fan-out construction for OR. We will replace each use of quantum fan-out operation with either a cat state or a classical parity circuit. We first reduce $\mathsf{EX}_{c,n}$ on $x \in \{0,1\}^n$ to NOR_m on $y \in \{0,1\}^m$ where $m = \lceil \log(n) \rceil$ [HŠ05] and then use a $[\mathsf{QNC}^0_{\text{🐱}}, \oplus_{\mathsf{f}}]^{O(1)}$ circuit with exponentially-large operations to compute OR_m [TT16].

- **Step 1: Reducing to $\log(n)$ inputs:** Akin to Lemma 3.4, we have $m+1$ test quantum states:

$$\forall 0 \le j \le m,\ |t_j\rangle = \mathsf{CR}_{\mathsf{X}}\left(\frac{\pi\left(|x|-c\right)}{2^j}\right)|0\rangle \tag{6.13}$$

  Let $|t\rangle = \prod_j^m |t_j\rangle$ and let t be the outcome of measuring $|t\rangle$.

  - If $|x| = c$, then $|t\rangle = |0^m\rangle \implies \mathsf{NOR}(t) = 1$
  - If $|x| \neq c$, consider the unique factorization of $|x| - c$:

$$|x| - c = 2^b\,(\text{some odd integer})$$

    As $|x| - c < n$, we have that $0 \le b < \log(n)$. This gives us:

$$|t_b\rangle = \mathsf{CR}_{\mathsf{X}}\left(\pi\,(\text{some odd integer})\right)|0\rangle = |1\rangle$$
$$\implies \mathsf{NOR}(t) = 0$$

  Therefore, we have that $\mathsf{NOR}(t) = \mathsf{EX}_{c,n}(x)$ for all cases:

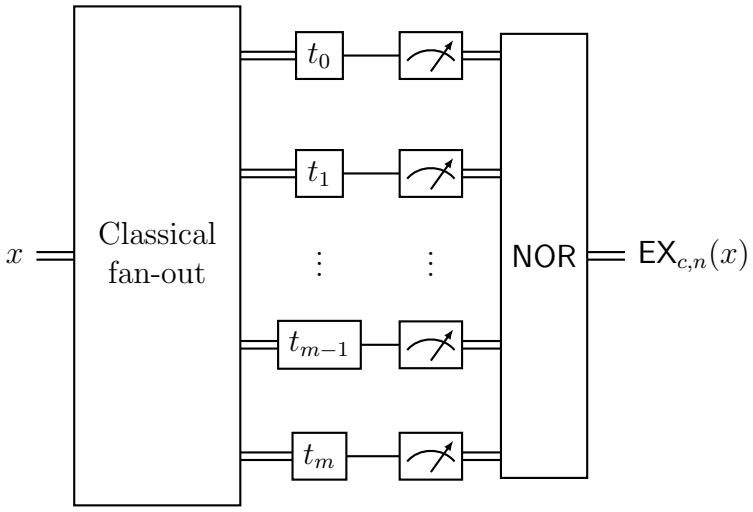

Figure 6.8: Computing EX using test state $|t\rangle$

We want to simulate the construction and measurement outcomes of $|t\rangle$ in $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$. We first use classical fan-out from $\oplus_{\mathsf{f}}$ to copy our input x, $m+1$ times and then simulate the measurement of each $|t_i\rangle$ by observing that the circuit for constructing $|t_i\rangle$ is:

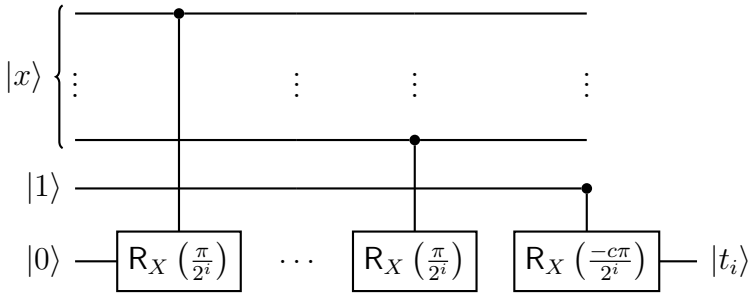

Figure 6.9: Constructing $|t_i\rangle$ using n $\mathsf{R_X}$ operations

which enables us to apply Lemma 6.14.

- **Step 2: OR_n using exponential-size operations:** We permit exponential-sized operations for this step as the input size is $O(\log(n))$. Using exponential-size fan-out, we first copy the classical input 2^n times. Then using these copies, we compute $\mathsf{MOD}_2(x_s)$ for each $s \subseteq [n]$ using classical parity gates. Now we use Lemma 6.14 and the Fourier expansion of OR. Explicitly, we now generate a 2^n-size cat state and perform the following operation in $[\mathsf{QNC}^0_{\text{cat}}, \oplus_{\mathsf{f}}]^{O(1)}$:

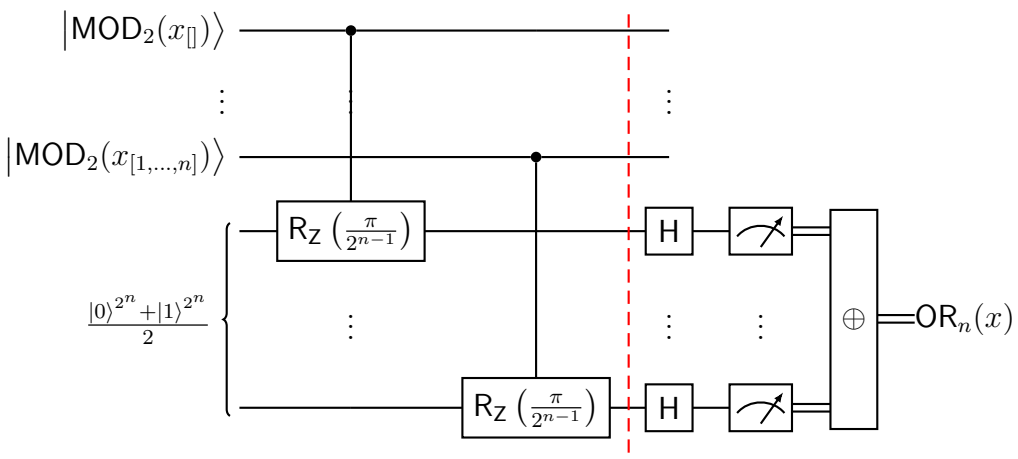

Figure 6.10: Computing OR_n in phase of exponential-size cat state

Using the Fourier expansion of OR, the cat state before the layer of Hadamards is:

$$\frac{|0\rangle^{2^n} + \exp\left(i\frac{\pi}{2^{n-1}}\sum_{s\in[n]}\mathsf{MOD}_2(x_s)\right)|1\rangle^{2^n}}{\sqrt{2}} = \frac{|0\rangle^{2^n} + \exp\left(i\pi\mathsf{OR}_n(x)\right)|1\rangle^{2^n}}{\sqrt{2}} \tag{6.14}$$

Using Equations 6.8 and 6.9, then we have that the measured string has even parity if and only if $\mathsf{OR}_n(x) = 0$.

□

To show that $\mathsf{TC}^0 \subseteq [\mathsf{QNC}^0_{\text{🐱}}, \oplus_\mathsf{f}]^{O(1)}$, we first need to show a construction for $\mathsf{TH}_{t,n}$ for $0 \le t \le a$. Using this, given a TC^0 circuit, we implement each layer using a $[\mathsf{QNC}^0_{\text{🐱}}, \oplus_\mathsf{f}]^{O(1)}$. Then using Observation 6.10(i), we can combine all circuits into one $[\mathsf{QNC}^0_{\text{🐱}}, \oplus_\mathsf{f}]^{O(1)}$ circuit as there are only a constant number of layers.

Theorem 6.16. *The threshold gate can be simulated in* $[\mathsf{QNC}^0_{\text{🐱}}, \oplus_\mathsf{f}]^{O(1)}$

Proof. Let $x \in \{0,1\}^n$ be the input to $\mathsf{TH}_{t,n}$ for $0 \le t \le n$. We will use the same protocol as Section 3.4:

- Using classical fan-out in $\oplus_\mathsf{f}$, make $n - t + 1$ copies of x. Let us call the copies $x^{(t)}, x^{(t+1)}, \ldots, x^{(n)}$.
- For i from t to n, in parallel:
  - Compute $\mathsf{EX}_{i,n}$ on $x^{(i)}$
- Compute the parity of the output of the exact gates

If $|x| \geq t$, then exactly one of $\{\mathsf{EX}_{i,n}(x)\}$ is 1, that is, $\mathsf{EX}_{|x|,n} = 1$. If $|x| < t$, all of $\{\mathsf{EX}_{i,n}(x)\}$ will be 0. Therefore, computing parity gives us the threshold gate:

$$\mathsf{TH}_{t,n} = \oplus_{t \leq i \leq n} \left(\mathsf{EX}_{i,n}(x)\right) \tag{6.15}$$

Here is the summary of the procedure:

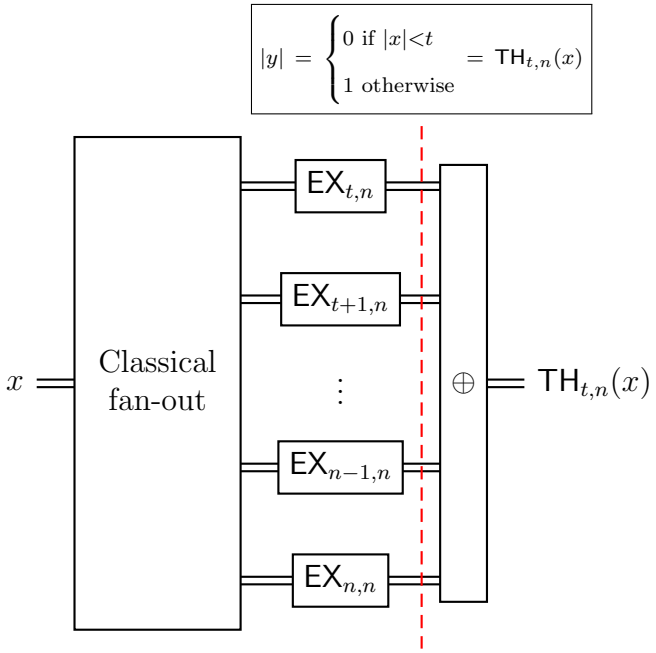

Figure 6.11: Computing threshold gate using exact gates

Theorem 6.15 shows us how to compute exact gates in $[\mathsf{QNC}^0_{\not\subset}, \oplus_{\mathsf{f}}]^{O(1)}$, completing the proof. □

Applications and interpretations:

- **Non-interleaving interpretation:** One way to interpret the result is that we have replaced quantum fan-out with a "one-way fan-out" and still retained the power to simulate TC^0 circuits:

  Definition 6.17 (One-way quantum fan-out)**.** *On input $x \in \{0,1\}^n$, one-way quantum fan-out allows us to perform the following irreversible operations:*

  - $|0^c\rangle \to \frac{|0^c\rangle + |1^c\rangle}{\sqrt{2}}$ *for $c \in \mathsf{poly}(|x|)$ – corresponding to cat state advice in the interleaved model*
  - $x \to \mathsf{MOD}_{2,n}(x)$ *and* $x \to x^c$ *where $c \in \mathsf{poly}(|x|)$ – corresponding to $\oplus_{\mathsf{f}}$ in the interleaved model*

  Corollary 6.18. $\mathsf{TC}^0 \subseteq \left(\mathsf{QNC}^0\right)^{\mathsf{one-way\ F}}$

However, crucially, we are not able to compute QTC^0 unlike $\mathsf{QNC}^0_\mathsf{f}$.

- **Interleaving is necessary:** Theorem 6.13 provides another example of a problem that cannot be solved by a quantum and classical class individually but can be solved once we strictly interleave them together. As $\mathsf{AC}^0[2] \subsetneq \mathsf{TC}^0$, threshold gates cannot be computed by $\oplus_\mathsf{f}$. Moreover, using light cone arguments:

  Theorem 6.19. *No* $\mathsf{QNC}^0_{\text{🐱}}$ *circuit can compute* $\oplus$.

  Proof. Let there be a d-depth quantum circuit Q with cat state advice that solves $\oplus$ on input $x \in \{0,1\}^n$ such that output qubit y_j for some index j of $Q\left|x\right\rangle\left|0^a\right\rangle$ for $a \in \mathsf{poly}(n)$ holds $\left|\oplus(x)\right\rangle$.

  Let the maximum fan-in of any gate be c. We define the light cone of an input qubit as the set of all output qubits that it can affect, and the light cone of an output qubit as the set of all input qubits that can affect it. Note that if there is no path (through gates) from an input qubit to an output qubit, the state of the output qubit cannot be affected by the state of the input qubit. Therefore, the light cone size for input and output qubits is at most c^d. Set all input qubits in the light cone of y_j to 0. If $d < \log_c(n)$, then there must be at least one input qubit x_i that is not in the light cone of y_j. Set all remaining input bits to 0 as well. The circuit must report the correct answer in this reduced case; that is, it must be that the state of qubit y_j is $\left|x_i\right\rangle$. However, this is impossible as x_i cannot affect the state of qubit y_j. Therefore, it must be that $d > \log_c(n)$. □

  Therefore, $\mathsf{QNC}^0_{\text{🐱}}$ cannot compute threshold gates either – the interleaving was crucial.

- **Simulating TC^0 with an untrusted quantum device:** Theorem 6.13 can be combined with an interactive proof for TC^0 using an untrusted TC^0 prover and a $\mathsf{AC}^0[2]$ verifier [Slo22]:

  Theorem 6.20. $\mathsf{TC}^0 \subseteq \mathsf{IP}[\mathsf{TC}^0, \mathsf{AC}^0$

  Proof. This proof is by Slote [Slo22]. We want to compute the $\mathsf{TH}_{t,n}$ gate for some input $x \in \{0,1\}^n$. The TC^0 prover sends a $n \times n$ tableau matrix M to the verifier with $1 \leq i, j \leq n$: $M_{i,j} = \mathsf{EX}_{j,i}(x_{[i]})$. Therefore, each row i has at most one 1, and it appears in the column j such that $|x_{[i]}| = j$. As a TC^0 device can compute exact gates, it can compute this matrix by first copying the input n^2 times using classical fan-out. Moreover, note that $\mathsf{TH}_{t,n}(x)$ is the OR of the last $n-t+1$ elements of the final row. Therefore, using this tableau, the verifier can compute threshold gates.

  We must show that a AC^0 circuit can verify that the tableau corresponds to input x. It can do so by performing the following checks in parallel:

  - If $x_1 = 0$, then OR_n of the first row must be 0. Otherwise, OR_{n-1} of the last $n-1$ elements of the first row must be 0 and $M_{1,1} = 1$.

– For $2 \leq i \leq n$:
  * If $x_i = 0$, then piece-wise XOR of row i and row $i-1$ must be 0.
  * If $x_i = 1$ and OR_n of row $i-1$ is 0: $M_{i,1} = 1$ and OR_{n-1} of the last $n-1$ elements of the first row must be 0.
  * If $x_i = 1$ and OR_n of row $i-1$ is 1: XOR of $M_{i-1,j}$ and $M_{i,j+1}$ for all $1 \leq j \leq n-1$ must be 0 and $M_{i,1} = 0$.

□

Combining this with Theorem 6.13, we have that:

Theorem 6.21. $\mathsf{TC}^0 \subseteq \mathsf{IP}[\mathsf{QNC}^0_{\textcircled{\scriptsize ⋈}}, \mathsf{AC}^0[2]]$

The only difference now is that the tableau is constructed through $O(1)$ rounds of interaction between the $\mathsf{QNC}^0_{\textcircled{\scriptsize ⋈}}$ and $\oplus_f$ device as they can compute the exact gate. Then, a $\mathsf{AC}^0[2]$ circuit verifies the tableau in the last round and computes the threshold gate. This gives a protocol for *unconditional* delegated computation of TC^0 using an untrusted low-depth quantum device with cat state advice.

Chapter 7

Future directions

In this chapter, we briefly discuss future directions and open questions relating to the study of interleaved low-depth quantum-classical circuits.

- **Further applications and separations in the current setting:** We would like to know whether the interleaved separations of Chapter 6 are optimal.
  - Since quantum circuits with constant T-gate depth can be compressed to constant *quantum* depth in the weak interleaved model, the following critical question arises:

    Question 7.1. *Which quantum tasks have constant* T*-depth implementations?*

    As we saw, surprisingly, this is true for approximate quantum Fourier transform. In general, it will be interesting to study and examine problems that cannot be compressed to constant T-depth as compared to factoring, in some sense, these problems are more *powerful* applications of quantum computation. Furthermore, we can consider the power of the weak interleaved models for relational and interactive tasks [GS20, GJS21].
  - For the strict interleaving setting, we are restricted to computing classical functions given the loss of coherence at every layer. We showed that the strongest classical inclusion (TC^0) of $\mathsf{QNC}^0_\mathsf{f}$ ports to this model (if we allow cat state quantum advice). To show that $\mathsf{QNC}^0_\mathsf{f}$ can simulate NC^1 and beyond, we must consider parallelization lemmas that act on non-commuting gates [GS20, BT88]:

    Question 7.2. *Can we reduce the depth of a series of non-commuting gates using unbounded quantum fan-out gates?*

- **Locality:** One very desirable feature of the separations present in [BGK18, WKST19] is that the quantum circuits needed to achieve them are geometrically-local and hence, easier to implement in practice. Therefore, a key question is to check whether our constructions in Chapter 6 are sustained if we restrict ourselves to geometrically-local circuits. On the other hand, [CC22, DTS+22] show that the probability of a particular output string from geometrically-local constant-depth quantum circuits on a grid can

be estimated to inverse-polynomial error in quasi-polynomial time. Therefore, another direction would be to investigate whether interleaving the geometrically-local circuits with low-depth classical computation thwarts these probability estimation methods.

- **Uniformity:** Another helpful future task is to nail down the exact complexity of generating a description of the interleaved circuits we construct. We conjecture that all constructions used are DLOGTIME-uniform. In particular, it seems that deriving the description of the $\oplus_f$ circuit when simulating Clifford operations by weak interleaving can be done in log-time using a divide-and-conquer strategy.

- **Different quantum advice and classical gates:** In all constructions, we used cat states to parallelize commuting gates. A natural question emerges: can some other state advice give qualitatively different parallelization results? To study this, first, we must determine which states are not constant-depth equivalent to cat states or the computational basis states (in the flavour of [DVC00]). Similarly, we used $\oplus_f$ in all constructions as the classical circuit model. A natural next question would be to consider whether more powerful classical circuits can strengthen the separations.

- **Quantum advantage demonstration via interleaving low-depth circuits:** One could imagine using the separations discussed to perform demonstrations of quantum advantage. Given the constant-depth of $[\mathsf{QNC}^0, \oplus_f]^{O(1)}$ and $[\mathsf{QNC}^0; \oplus_f]^{O(1)}$, practically, these schemes might be implementable in the near-term. Moreover, their separations are unconditional, and the output is efficiently verifiable as opposed to sampling-based schemes. In fact, using the two-dimensional hidden linear problem separation [BGK18], [CSV21] devised a protocol on a $2^{10} \times 2^{10}$ grid that was robust to some additive noise. Moreover, there has been an experimental demonstration using these separations on trapped ion qubits [DZA+22] with the aim of beating 1-depth classical circuits. However, as noted in [CSV21], given that these tasks are easy to compute classically (in P), the key issue is verifying that the device is of sub-logarithmic depth. This can only be done by heuristic or physical means. [CSV21] suggested doing so by considering the operation frequency of state-of-the-art classical devices. Such arguments would be even harder to make in the interleaved models, where a part of the computation is allowed to be log-depth. Nevertheless, whether the new separations can be used to reduce the number of qubits required could be studied.

  In the weakly interleaved case, both interactive cryptographic tasks [LG21, HG21b] as well as sampling tasks [BJS11, JN13, TYT14, Koh15, MSM17, BFK17, YJS19, PBG20] admit realizations that only need $[\mathsf{QNC}^0; \oplus_f]^{O(1)}$ devices. Therefore, these strategies seem to be a more robust way of demonstrating quantum advantage using interleaved models as these tasks cannot be performed efficiently by classical machines under complexity-theoretic conjectures.

  Looking beyond the near-term, for certain platforms (based on their noise characteristics), it might be easier to implement approximate quantum Fourier transform using the $[\mathsf{QNC}^0; \oplus_f]^{O(1)}$ implementation rather than the QNC^1 one.

- **Classical lower bounds on low-depth quantum circuits:** Another fruitful direction is to consider the limits of the power of interleaved low-depth circuits. Not much is known in this regard, even in the non-interleaved case. [GHMP01] show lower bounds for restricted variants of $\mathsf{QAC}^0_\mathsf{f}$: with some planarity and locality conditions imposed, $\mathsf{QAC}^0_\mathsf{f} \subseteq \mathsf{P/poly}$. On the other hand [FFG+03, Ber11, PFGT20, Ros20] studied whether QAC^0 circuits compute can quantum fan-out. It is shown that this is not possible for QAC^0 with a sublinear number of auxiliary qubits [FFG+03, Ber11], while with no restriction of the number of auxiliary qubits, this lower bound is only known for 2-depth QAC^0 circuits [PFGT20, Ros20] and "mostly classical" QAC^0 circuits [Ros20]. In general, any new techniques for proving lower bounds for low-depth circuits will be instrumental to the field.

References

[AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011. [p. 1]

[Aar05] Scott Aaronson. Ten semi-grand challenges for quantum computing theory. https://www.scottaaronson.com/writings/qchallenge.html, 2005. [p. 4]

[Aar10] Scott Aaronson. BQP and the polynomial hierarchy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 141–150, 2010. [p. 4]

[Aar11] Scott Aaronson. Projects aplenty. https://www.scottaaronson.com/blog/?p=663, 2011. [p. 4]

[Aar14] Scott Aaronson. Oracular separations between poly- and log-depth quantum circuits. Theoretical Computer Science Stack Exchange. https://cstheory.stackexchange.com/q/25145, 2014. [p. 4]

[AG04] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.

[AGS22] Atul Singh Arora, Alexandru Gheorghiu, and Uttam Singh. Oracle separations of hybrid quantum-classical circuits. *arXiv preprint arXiv:2201.01904*, 2022. [pp. 5, 49]

[AMMR13] Matthew Amy, Dmitri Maslov, Michele Mosca, and Martin Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 32(6):818–830, 2013.

[BBC+93] Charles H Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical review letters*, 70(13):1895, 1993. [p. 25]

[BBT04] Gilles Brassard, Anne Broadbent, and Alain Tapp. Recasting Mermin's multi-player game into the framework of pseudo-telepathy. *arXiv preprint quant-ph/0408052*, 2004. [pp. 33, 35]

[BCM+21] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. *Journal of the ACM (JACM)*, 68(5):1–47, 2021. [pp. 1, 46]

[Ber11] Debajyoti Bera. A lower bound method for quantum circuits. *Information processing letters*, 111(15):723–726, 2011. [p. 61]

[BFK17] Adam Bouland, Joseph F Fitzsimons, and Dax Enshan Koh. Complexity classification of conjugated Clifford circuits. *arXiv preprint arXiv:1709.01805*, 2017. [pp. 31, 60]

[BFNV19] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159–163, 2019. [p. 1]

[BGK18] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018. [pp. 3, 32, 33, 59, 60]

[BGKT20] Sergey Bravyi, David Gosset, Robert Koenig, and Marco Tomamichel. Quantum advantage with noisy shallow circuits. *Nature Physics*, 16(10):1040–1045, 2020. [pp. 4, 32]

[BJS11] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011. [pp. 1, 31, 60]

[BKKK22] Sergey Bravyi, Isaac Kim, Alexander Kliesch, and Robert Koenig. Adaptive constant-depth circuits for manipulating non-abelian anyons. *arXiv preprint arXiv:2205.01933*, 2022. [p. 47]

[BKP10] Dan Browne, Elham Kashefi, and Simon Perdrix. Computational depth complexity of measurement-based quantum computation. In *Conference on Quantum Computation, Communication, and Cryptography*, pages 35–46. Springer, 2010. [pp. 3, 25]

[BSS16] Sergey Bravyi, Graeme Smith, and John A Smolin. Trading classical and quantum computational resources. *Physical Review X*, 6(2):021043, 2016. [p. 41]

[BT88] David A Mix Barrington and Denis Therien. Finite monoids and the fine structure of NC. *Journal of the ACM (JACM)*, 35(4):941–952, 1988. [p. 59]

[BVHS+18] Juan Bermejo-Vega, Dominik Hangleiter, Martin Schwarz, Robert Raussendorf, and Jens Eisert. Architectures for quantum simulation showing a quantum speedup. *Physical Review X*, 8(2):021010, 2018. [p. 1]

[BY22] Zvika Brakerski and Henry Yuen. Quantum garbled circuits. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 804–817, 2022. [p. 47]

[CC22] Nolan J Coble and Matthew Coudron. Quasi-polynomial time approximation of output probabilities of geometrically-local, shallow quantum circuits. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 598–609. IEEE, 2022. [p. 59]

[CCL20] Nai-Hui Chia, Kai-Min Chung, and Ching-Yi Lai. On the need for large quantum depth. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 902–915, 2020. [pp. 5, 44, 49]

[CH22] Nai-Hui Chia and Shih-Han Hung. Classical verification of quantum depth. *arXiv preprint arXiv:2205.04656*, 2022. [p. 5]

[CM20] Matthew Coudron and Sanketh Menda. Computations with greater quantum depth are strictly more powerful (relative to an oracle). In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 889–901, 2020. [pp. 5, 44, 49]

[Com] Complexity Zoo. https://complexityzoo.net/. [p. 12]

[CSV21] Matthew Coudron, Jalex Stark, and Thomas Vidick. Trading locality for time: certifiable randomness from low-depth circuits. *Communications in mathematical physics*, 382(1):49–86, 2021. [pp. 4, 32, 33, 60]

[CTI+21] Antonio D Córcoles, Maika Takita, Ken Inoue, Scott Lekuch, Zlatko K Minev, Jerry M Chow, and Jay M Gambetta. Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits. *Physical Review Letters*, 127(10):100501, 2021. [p. 47]

[CW00] Richard Cleve and John Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 526–536. IEEE, 2000. [p. 46]

[Dam90] Carsten Damm. Problems complete for $\oplus\mathsf{L}$. In *International Meeting of Young Computer Scientists*, pages 130–137. Springer, 1990.

[DGC18] Vedran Dunjko, Yimin Ge, and J Ignacio Cirac. Computational speedups using small quantum devices. *Physical review letters*, 121(25):250501, 2018. [p. 41]

[DLH+22] Emma Deist, Yue-Hui Lu, Jacquelyn Ho, Mary Kate Pasha, Johannes Zeiher, Zhenjie Yan, and Dan M Stamper-Kurn. Fast non-destructive cavity readout of single atoms within a coherent atom array. *arXiv preprint arXiv:2205.14138*, 2022. [p. 47]

[DTS+22] Suchetan Dontha, Shi Jie Samuel Tan, Stephen Smith, Sangheon Choi, and Matthew Coudron. Approximating output probabilities of shallow quantum circuits which are geometrically-local in any fixed dimension. *arXiv preprint arXiv:2202.08349*, 2022. [p. 59]

[DVC00] Wolfgang Dür, Guifre Vidal, and J Ignacio Cirac. Three qubits can be entangled in two inequivalent ways. *Physical Review A*, 62(6):062314, 2000. [p. 60]

[DZA+22] Austin K Daniel, Yingyue Zhu, C Huerta Alderete, Vikas Buchemmavari, Alaina M Green, Nhung H Nguyen, Tyler G Thurtell, Andrew Zhao, Norbert M Linke, and Akimasa Miyake. Quantum computational advantage attested by nonlocal games with the cyclic cluster state. *Physical Review Research*, 4(3):033068, 2022. [p. 60]

[Fen03] Stephen A Fenner. Implementing the fanout gate by a Hamiltonian. *arXiv preprint quant-ph/0309163*, 2003. [p. 24]

[FFG+03] Maosen Fang, Stephen Fenner, Frederic Green, Steven Homer, and Yong Zhang. Quantum lower bounds for fanout. *arXiv preprint quant-ph/0312208*, 2003. [p. 61]

[FW22] Stephen Fenner and Rabins Wosti. Implementing the fanout operation with simple pairwise interactions. *arXiv preprint arXiv:2203.01141*, 2022. [p. 24]

[FZ04] Stephen A Fenner and Yong Zhang. Implementing fanout, parity, and mod gates via spin exchange interactions. *arXiv preprint quant-ph/0407125*, 2004. [p. 24]

[Gal18] François Le Gall. Average-case quantum advantage with shallow circuits. *arXiv preprint arXiv:1810.12792*, 2018. [pp. 4, 32, 33]

[GC99] Daniel Gottesman and Isaac L Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(6760):390–393, 1999. [pp. 3, 25, 26, 44]

[GDC+20] Andrew Y Guo, Abhinav Deshpande, Su-Kuan Chu, Zachary Eldredge, Przemyslaw Bienias, Dhruv Devulapalli, Yuan Su, Andrew M Childs, and Alexey V Gorshkov. Implementing a fast unbounded quantum fanout gate using power-law interactions. *arXiv preprint arXiv:2007.00662*, 2020. [pp. 6, 24, 47]

[GGM19] Mariami Gachechiladze, Otfried Gühne, and Akimasa Miyake. Changing the circuit-depth complexity of measurement-based quantum computation with hypergraph states. *Physical Review A*, 99(5):052304, 2019. [p. 28]

[GHMP01] Frederic Green, Steven Homer, Cristopher Moore, and Christopher Pollett. Counting, fanout, and the complexity of quantum ACC. *arXiv preprint quant-ph/0106017*, 2001. [pp. 20, 61]

[GJS21] Daniel Grier, Nathan Ju, and Luke Schaeffer. Interactive quantum advantage with noisy, shallow Clifford circuits. *arXiv preprint arXiv:2102.06833*, 2021. [p. 59]

[GKH+21] Pranav Gokhale, Samantha Koretsky, Shilin Huang, Swarnadeep Majumder, Andrew Drucker, Kenneth R Brown, and Frederic T Chong. Quantum fan-out: circuit optimizations and technology modeling. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 276–290. IEEE, 2021. [p. 24]

[Got98] Daniel Gottesman. The Heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*, 1998. [pp. 10, 31]

[GS20] Daniel Grier and Luke Schaeffer. Interactive shallow Clifford circuits: quantum advantage against NC^1 and beyond. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 875–888, 2020. [p. 59]

[GWD17] Xun Gao, Sheng-Tao Wang, and L-M Duan. Quantum supremacy for simulating a translation-invariant Ising spin model. *Physical review letters*, 118(4):040502, 2017. [p. 1]

[Hås14] Johan Håstad. On the correlation of parity and small-depth circuits. *SIAM Journal on Computing*, 43(5):1699–1708, 2014. [p. 36]

[HCLB11] Matty J Hoban, Earl T Campbell, Klearchos Loukopoulos, and Dan E Browne. Non-adaptive measurement-based quantum computation and multi-party Bell inequalities. *New Journal of Physics*, 13(2):023014, 2011. [p. 16]

[HG21a] Atsuya Hasegawa and François Le Gall. Quantum advantage with shallow circuits under arbitrary corruption. *arXiv preprint arXiv:2105.00603*, 2021. [pp. 4, 32]

[HG21b] Shuichi Hirahara and François Le Gall. Test of quantumness with small-depth quantum circuits. *arXiv preprint arXiv:2105.05500*, 2021. [pp. 6, 46, 60]

[HG22] Atsuya Hasegawa and François Le Gall. An optimal oracle separation of classical and quantum hybrid schemes. *arXiv preprint arXiv:2205.04633*, 2022. [p. 5]

[HHB+20] Jonas Haferkamp, Dominik Hangleiter, Adam Bouland, Bill Fefferman, Jens Eisert, and Juani Bermejo-Vega. Closing gaps of a quantum advantage with short-time Hamiltonian dynamics. *Physical Review Letters*, 125(25):250501, 2020.

[HŠ05] Peter Høyer and Robert Špalek. Quantum fan-out is powerful. *Theory of computing*, 1(1):81–103, 2005. [pp. 3, 13, 15, 16, 21, 46, 53]

[JN13] Richard Jozsa and Maarten Van den Nest. Classical simulation complexity of extended Clifford circuits. *arXiv preprint arXiv:1305.6190*, 2013. [pp. 31, 60]

[Joz06] Richard Jozsa. An introduction to measurement based quantum computation. *NATO Science Series, III: Computer and Systems Sciences. Quantum Information Processing-From Theory to Experiment*, 199:137–158, 2006. [pp. 4, 28, 31, 44]

[KLVY22] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. *arXiv preprint arXiv:2203.15877*, 2022. [p. 1]

[Koh15] Dax Enshan Koh. Further extensions of Clifford circuits and their classical simulation complexities. *arXiv preprint arXiv:1512.07892*, 2015. [pp. 31, 60]

[Leu01] Debbie W Leung. Two-qubit projective measurements are universal for quantum computation. *arXiv preprint quant-ph/0111122*, 2001. [p. 28]

[Leu04] Debbie W Leung. Quantum computation by measurements. *International Journal of Quantum Information*, 2(01):33–43, 2004. [p. 28]

[LG21] Zhenning Liu and Alexandru Gheorghiu. Depth-efficient proofs of quantumness. *arXiv preprint arXiv:2107.02163*, 2021. [pp. 6, 46, 60]

[LLKH22] Tsung-Cheng Lu, Leonardo A Lessa, Isaac H Kim, and Timothy H Hsieh. Measurement as a shortcut to long-range entangled quantum matter. *arXiv preprint arXiv:2206.13527*, 2022. [pp. 6, 47]

[Mer90] N David Mermin. Extreme quantum entanglement in a superposition of macroscopically distinct states. *Physical Review Letters*, 65(15):1838, 1990. [pp. 33, 35]

[MN98] Cristopher Moore and Martin Nilsson. Some notes on parallel quantum computation. *arXiv preprint quant-ph/9804034*, 1998. [pp. 2, 13, 14]

[MN01] Cristopher Moore and Martin Nilsson. Parallel quantum computation and quantum codes. *SIAM journal on computing*, 31(3):799–815, 2001. [pp. 2, 13, 14]

[Moo99] Cristopher Moore. Quantum circuits: Fanout, parity, and counting. *arXiv preprint quant-ph/9903046*, 1999. [pp. 2, 13, 15, 20]

[MSM17] Jacob Miller, Stephen Sanders, and Akimasa Miyake. Quantum supremacy in constant-time measurement-based computation: A unified architecture for sampling and verification. *Physical Review A*, 96(6):062320, 2017. [pp. 31, 60]

[Nie03] Michael A Nielsen. Quantum computation by measurement and quantum memory. *Physics Letters A*, 308(2-3):96–100, 2003. [p. 28]

[NNZ+21] Crystal Noel, Pradeep Niroula, Daiwei Zhu, Andrew Risinger, Laird Egan, Debopriyo Biswas, Marko Cetina, Alexey V Gorshkov, Michael J Gullans, David A Huse, et al. Observation of measurement-induced quantum phases in a trapped-ion quantum computer. *arXiv preprint arXiv:2106.05881*, 2021. [p. 47]

[OCTW+22] GA Oakes, VN Ciriano-Tejel, D Wise, MA Fogarty, T Lundberg, C Lainé, S Schaal, F Martins, DJ Ibberson, L Hutin, et al. Fast high-fidelity single-shot readout of spins in silicon using a single-electron box. *arXiv preprint arXiv:2203.06608*, 2022. [p. 47]

[O'D14] Ryan O'Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014. [p. 9]

[PBG20] Hakop Pashayan, Stephen D Bartlett, and David Gross. From estimation of quantum probabilities to simulation of quantum circuits. *Quantum*, 4:223, 2020. [pp. 31, 60]

[PDF+21] Juan M Pino, Jennifer M Dreiling, Caroline Figgatt, John P Gaebler, Steven A Moses, MS Allman, CH Baldwin, M Foss-Feig, D Hayes, K Mayer, et al. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, 2021. [p. 47]

[PFGT20] Daniel Padé, Stephen Fenner, Daniel Grier, and Thomas Thierauf. Depth-2 QAC circuits cannot simulate quantum parity. *arXiv preprint arXiv:2005.12169*, 2020. [p. 61]

[PHOW20] Tianyi Peng, Aram W Harrow, Maris Ozols, and Xiaodi Wu. Simulating large quantum circuits on a small quantum computer. *Physical Review Letters*, 125(15):150504, 2020. [p. 41]

[RB01] Robert Raussendorf and Hans J Briegel. A one-way quantum computer. *Physical review letters*, 86(22):5188, 2001. [p. 3]

[RBB03] Robert Raussendorf, Daniel E Browne, and Hans J Briegel. Measurement-based quantum computation on cluster states. *Physical review A*, 68(2):022312, 2003. [pp. 3, 28]

[Ros12] David Rosenbaum. Optimal quantum circuits for nearest-neighbor architectures. *arXiv preprint arXiv:1205.0036*, 2012. [p. 24]

[Ros17] Benjamin Rossman. An entropy proof of the switching lemma and tight bounds on the decision-tree size of AC^0, 2017. [p. 36]

[Ros20] Gregory Rosenthal. Bounds on the QAC^0 complexity of approximating parity. *arXiv preprint arXiv:2008.07470*, 2020. [p. 61]

[RRG+22] Kenneth Rudinger, Guilhem J Ribeill, Luke CG Govia, Matthew Ware, Erik Nielsen, Kevin Young, Thomas A Ohki, Robin Blume-Kohout, and Timothy Proctor. Characterizing midcircuit measurements on a superconducting qubit using gate set tomography. *Physical Review Applied*, 17(1):014014, 2022. [p. 47]

[SBA+22] Kevin Singh, Conor E Bradley, Shraddha Anand, Vikram Ramesh, Ryan White, and Hannes Bernien. Mid-circuit correction of correlated phase errors using an array of spectator qubits. *arXiv preprint arXiv:2208.11716*, 2022. [p. 47]

[Sel13] Peter Selinger. Quantum circuits of T-depth one. *Physical Review A*, 87(4):042302, 2013.

[Sho99] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999. [p. 1]

[Slo22] Joseph Slote. Personal communication., 2022. [pp. 8, 41, 57]

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82, 1987. [p. 38]

[TD02] Barbara M Terhal and David P DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *arXiv preprint quant-ph/0205133*, 2002. [pp. 5, 7, 44]

[TT16] Yasuhiro Takahashi and Seiichiro Tani. Collapse of the hierarchy of constant-depth exact quantum circuits. *computational complexity*, 25(4):849–881, 2016. [pp. 3, 13, 15, 16, 18, 53]

[TVV22] Nathanan Tantivasadakarn, Ashvin Vishwanath, and Ruben Verresen. A hierarchy of topological order from finite-depth unitaries, measurement and feedforward. *arXiv preprint arXiv:2209.06202*, 2022. [pp. 6, 47]

[TYT14] Yasuhiro Takahashi, Takeshi Yamazaki, and Kazuyuki Tanaka. Hardness of classically simulating quantum circuits with unbounded Toffoli and fan-out gates. *Quantum Information & Computation*, 14(13-14):1149–1164, 2014. [p. 60]

[Vol99] Heribert Vollmer. *Introduction to circuit complexity: a uniform approach*. Springer Science & Business Media, 1999. [p. 12]

[WKE+19] Yong Wan, Daniel Kienzler, Stephen D Erickson, Karl H Mayer, Ting Rei Tan, Jenny J Wu, Hilma M Vasconcelos, Scott Glancy, Emanuel Knill, David J Wineland, et al. Quantum gate teleportation between separated qubits in a trapped-ion processor. *Science*, 364(6443):875–878, 2019. [p. 47]

[WKST19] Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 515–526, 2019. [pp. 4, 7, 32, 33, 35, 36, 38, 50, 59]

[Yao93] A Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 352–361. IEEE, 1993. [p. 2]

[YJS19] Mithuna Yoganathan, Richard Jozsa, and Sergii Strelchuk. Quantum advantage of unitary Clifford circuits with magic state inputs. *Proceedings of the Royal Society A*, 475(2225):20180427, 2019. [pp. 31, 60]

[ZKML+21] Daiwei Zhu, Gregory D Kahanamoku-Meyer, Laura Lewis, Crystal Noel, Or Katz, Bahaa Harraz, Qingfeng Wang, Andrew Risinger, Lei Feng, Debopriyo Biswas, et al. Interactive protocols for classically-verifiable quantum advantage. *arXiv preprint arXiv:2112.05156*, 2021. [p. 47]