Interactive quantum advantage with noisy, shallow Clifford circuits

Daniel Grier Nathan Ju Luke Schaeffer

A major goal in quantum complexity theory is identifying problems which are efficiently solvable by quantum computers and not efficiently solvable by classical computers. If willing to believe certain conjectures, one can be convinced of this separation by the discovery of quantum algorithms that solve classically hard problems. For example, the belief that classical computers cannot efficiently factor integers contrasts with Shor's algorithm for factoring integers on a quantum computer [Sho97]. However, a demonstration of Shor's algorithm on instances that are not efficiently solvable by classical computers would require quantum resources far out of reach of near-term capabilities. This has spurred developments in devising sampling problems that separate efficient, near-term quantum computers and classical computers like IQP circuit sampling [BJS10], BosonSampling [AA11], and random circuit sampling [Boi+18]. However, convincing evidence that *noisy* quantum computers outperform classical computers in these tasks suffers from the necessity of assuming some non-standard complexity-theoretic conjectures that are often native to each proposal.

Surprisingly, if you restrict to the setting of constant-depth circuits, a noisy, unconditional separation *is* possible. At first, these unconditional separations were known only in the noiseless setting. That line of work was initiated by pioneering work of Bravyi, Gosset, and König [BGK18] who showed a strict separation between constant-depth quantum circuits (QNC^0) and constant-depth classical circuits with bounded fan-in gates (NC^0). The separation is based on the relation problem[1] associated with measuring the outputs of a shallow Clifford circuit, which they called the Hidden Linear Function (HLF) problem due to certain algebraic properties of the output. Furthermore, they show that NC^0 cannot even solve this problem on average, a result which was later strengthened in several ways [CSV18; Le 19; Ben+19]. Nevertheless, these works still assumed that the quantum circuit solving the task was noise free.

Follow-up work of Bravyi, Gosset, König, and Tomamichel [Bra+20] showed that it was possible to encode the qubits of the quantum circuit in such a way that it both preserved the separation and also allowed the quantum circuit to err with some constant probability. Interestingly, this was accomplished not by explicitly carrying out the quantum error correction procedure, but by simply measuring the syndrome qubits of the code and requiring the classical circuit to do the same. In fact, their procedure provided a more-or-less general recipe for taking a constant-depth quantum/classical circuit separation and turning it into a separation in which the quantum circuit was also allowed noise.

This raises an obvious question: how many circuit separations can we upgrade in this way? NC^0 circuits are fairly weak—they cannot even compute the logical AND of all input bits—and so we would like to show that even larger classes of classical devices cannot solve a problem that a noisy shallow quantum circuit can.

As a warm-up, we first consider the separation of Bene Watts, Kothari, Schaeffer, and Tal [Ben+19], which shows that constant-depth classical circuits with *unbounded* fan-in gates (AC^0)

[1] Generally speaking, a relation problem is defined by a relation $R \subseteq \Sigma^* \times \Sigma^*$. Given an input x, the task is to find some y such that $(x, y) \in R$.

cannot solve the HLF problem on average.[2] Combining this result with the general error-correction recipe for relation problems, we arrive at the following result:

Theorem 1. *There is a relation task solved by a noisy*[3] QNC^0 *circuit with probability* $1-o(1)$ *on all inputs. On the other hand, any* AC^0 *circuit can solve the problem on at most a* $\exp(-n^\alpha)$ *fraction of inputs for some constant* $\alpha > 0$.

The result of Bene Watts et al. is the strongest known low-depth separation of its kind, but stronger separations are known for tasks which admit some amount of interactivity. Consider the shallow Clifford measurement problem discussed above, where the measurements are made in two rounds. In the first round, the quantum device is given the bases in which to measure some of the qubits and returns their measurement outcomes; and in the second round, the quantum device is given bases in which to measure the remaining qubits and returns their measurement outcomes. Grier and Schaeffer [GS20] show that any classical device which can solve such problems must be relatively powerful. More specifically, if the initial Clifford state is a constant-width grid state, then the classical device can be used to solve problems in NC^1 (log-depth circuits of bounded fan-in gates); and if the starting state is a poly-width grid state, then the classical device can be used to solve problems in $\oplus\mathsf{L}$ (which can be thought of as the complete class for Clifford circuits, or more specifically, poly-depth circuits of CNOT gates).[4] Because $\mathsf{AC}^0[p] \subsetneq \mathsf{NC}^1$ unconditionally, the above interactive task can be solved by a QNC^0 circuit but not an $\mathsf{AC}^0[p]$ circuit, i.e., an AC^0 circuit with unbounded MOD_p gates[5] for some prime p.

One of the contributions of this work is massaging the noisy circuit separation recipe for relations problems into a recipe for *interactive* problems as well. Starting from an interactive protocol which exhibits a separation with a noise-free quantum circuit, there are three key steps to upgrade the separation to the noisy setting:

1. *Augment the interactive protocol with the surface code encoding of Bravyi et al. [Bra+20].* This is straightforward, but it's worth noting that it changes the problem definition—not just because there are more physical qubits due to the encoding, but because we cannot prepare the initial state exactly or decode the syndrome in constant depth. As for the relational case, the burden of these steps is offloaded into the problem definition.
2. *Show classical average-case hardness.* That is, show that even when the classical circuit simulating the interactive protocol is allowed to err on some constant fraction of its inputs, it can still be leveraged to solve a hard problem (e.g., a problem in NC^1 or $\oplus\mathsf{L}$). This step is the most involved, and new ideas will be required to upgrade existing interactive separations in this way.
3. *Connect to separations of classical complexity classes.* In some cases, this will lead to an unconditional separation between noisy shallow quantum circuits and shallow classical circuits, and in some cases this will lead to a conditional separation. We note that these separations will not be identical to those obtained in Ref [GS20] due to the fact that we use quasipolynomial-size circuits to decode the syndrome qubits of the surface code.

Fortunately, it was shown in Ref [GS20] that Step 2 holds[6] for the NC^1-hardness result. We immediately obtain the following separation:

[2] The authors of that paper refer to their task as the "Relaxed Parity Halving Problem," but it is still essentially the problem of measuring the outputs of a constant-depth Clifford circuit.

[3] We employ the same local stochastic noise model used in [Bra+20]. See the main text for details.

[4] It is known that $\mathsf{NC}^1 \subseteq \oplus\mathsf{L} \subseteq \mathsf{NC}^2$.

[5] The MOD_p gates outputs 1 iff the sum of the inputs bits is 0 mod p.

[6] Although it is not strictly required, we prove a slightly stronger average-case hardness result in the paper.

Theorem 2. *There is a two-round interactive task solved by a noisy* QNC^0 *circuit with probability* $1 - o(1)$ *on all inputs. Any* $\mathsf{AC}^0[p]$ *circuit (for primes* $p \geq 2$*) fails the task with some constant probability.*

Unfortunately, Step 2 is left as an open question in Ref [GS20] for the $\oplus\mathsf{L}$-hardness result. The second major contribution of this paper is to show that we can, in fact, obtain average-case hardness for this setting:

Theorem 3. *There is a two-round interactive task solved by a* QNC^0 *circuit with certainty. There exists a constant* $\delta > 0$ *such that any sufficiently powerful classical device which solves the task with probability at least* $1 - \delta$ *can also solve problems in* $\oplus\mathsf{L}$*. That is,*

$$\oplus\mathsf{L} \subseteq (\mathsf{BPAC}^0)^{\mathcal{R}}$$

where $\mathcal{R}$ *is an oracle for the classical solution.*

The proof of this theorem borrows an idea from cryptography called *randomized encodings*. In particular, we will employ the construction of Applebaum, Ishai, and Kushilevitz [AIK06] which randomizes instances of the following problem—given a layered DAG, determine the parity of the number of paths from vertex s to vertex t. In fact, we will use that this problem reduces to the $\oplus\mathsf{L}$-hardness result in [GS20]. Importantly, we show that when we compose the randomized encoding with the rest of the reduction, the distribution over inputs in the promise will be fairly uniform. This leads to a general way to boost the randomization in worst-to-average-case reductions using the framework in [GS20].

Using the recipe for interactive circuit separations, we obtain the following consequence:

Theorem 4. *There is a two-round interactive task solved by a noisy* QNC^0 *circuit with probability* $1 - o(1)$ *on all inputs. Assuming* $\oplus\mathsf{L} \not\subseteq (\mathsf{qBPAC}^0)^{\mathsf{L}}$*, any log-space machine fails the task with some constant probability.*

Let us briefly unpack the $\oplus\mathsf{L} \not\subseteq (\mathsf{qBPAC}^0)^{\mathsf{L}}$ assumption. First, consider the plausible assumption that $\oplus\mathsf{L} \not\subseteq \mathsf{L}$. An L machine is deterministic, while a $\oplus\mathsf{L}$ machine is non-deterministic and accepts if the parity of accepting paths is zero. On the other hand, it is well-known that the parity function is not in qBPAC^0 (i.e., random AC^0 circuits of quasipolynomial size). Therefore, one might also expect that $(\mathsf{qBPAC}^0)^{\mathsf{L}}$ is insufficiently powerful to compute $\oplus\mathsf{L}$ functions.

We do not attempt to give an exhaustive list of separations obtainable from Theorem 3. Much like the results of [GS20], there is an inherent tradeoff to the separation. We can weaken the assumption at the expense of weakening the separation.

Finally, we explore the regime between the NC^1-hardness result (i.e., an interactive task on constant-width grids) and the $\oplus\mathsf{L}$-hardness result (i.e., an interactive task on poly-width grids). To this end, we consider the interactive task on general width-w grids, and connect them to the problem of solving width-w permutation branching programs. We prove the analogue of Theorem 3 in this setting, which once again leads to conditional separations between noisy shallow quantum circuits and complexity classes solved by width-w permutation branching programs.

References

[Sho97] Peter W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM J. Comput.* 26.5 (Oct. 1997), pp. 1484–1509.

[AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. "Cryptography in NC^0". In: *SIAM J. Comput.* 36.4 (2006), pp. 845–888.

[BJS10] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. "Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2010, pp. 459–472.

[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*. STOC '11. Association for Computing Machinery, 2011, pp. 333–342. ISBN: 9781450306911.

[Boi+18] S. Boixo et al. "Characterizing quantum supremacy in near-term devices". In: *Nature Physics* 14.6 (2018), pp. 595–600.

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

[CSV18] Matthew Coudron, Jalex Stark, and Thomas Vidick. "Trading locality for time: certifiable randomness from low-depth circuits". In: *arXiv preprint arXiv:1810.04233* (2018).

[Ben+19] Adam Bene Watts et al. "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*. STOC '19. Association for Computing Machinery, 2019.

[Le 19] François Le Gall. "Average-case quantum advantage with shallow circuits". In: *34th Computational Complexity Conference (CCC 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2019.

[Bra+20] Sergey Bravyi et al. "Quantum advantage with noisy shallow circuits". In: *Nature Physics* (2020), pp. 595–600.

[GS20] Daniel Grier and Luke Schaeffer. "Interactive shallow Clifford circuits: quantum advantage against NC^1 and beyond". In: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*. STOC '20. Association for Computing Machinery, 2020.