Universiteit Leiden
The Netherlands

Separating quantum and classical computing: rigorous proof and practical application

Marshall, S.C.

Citation

Marshall, S. C. (2025, May 27). *Separating quantum and classical computing: rigorous proof and practical application*. Retrieved from https://hdl.handle.net/1887/4247215

Version:	Publisher's Version
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/4247215

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 7

Conclusion

In this chapter we present the conclusions of the thesis. First, we will recall the key research questions and detail the answers. Second, this chapter will provide details on the limitations of the methods discussed. Finally, we will reoutline promising directions for future work.

7.1 Research Overview

This thesis has centred on how we can differentiate classical and quantum computing, either practically or theoretically. The first of the research questions given in the introduction tackled the theoretical aspect of this:

Research question 1
What is the strongest theoretical basis for the claim "In polynomial time quantum computers can perform computations that classical computers cannot"?

This was the question addressed by the first two chapters, in chapter 2 we addressed this question directly by casting it into a comparison between FBQP and FBPP. Chapter 2 strengthened the separation between these classes to a new state of the art, i.e. that they are separate unless the polynomial hierarchy collapses to the 2nd level or two conjectures about permanents turn out to be incorrect. The chapter discussed the significance of this improvement, as the separation between FBQP and FBPP is in some sense the most rigorous separation of quantum and classical computing currently known.

Chapter 3 answered a similar research question as chapter 2, but now in the setting of advice.

Research question 2
If polynomial-time quantum computers can give better advice than polynomial-time classical computers, can we find such an advice-generating algorithm?

By focusing on the question of bounded advice, we found that if a quantum-classical separation could be reduced to advice then quantum and classical computing could not be separated in exponential time. This concept was put to the test in chapter 6, which developed a machine learning algorithm capable of exploiting any advantage posed by quantum advice. This algorithm proved successful both practically and for trap-door functions.

The remainder of the chapters focused on a more practical question than the first two. Instead of asking about the improvements offered by quantum computing in the limit of growing size, we asked how these advantages could be gained earlier. In particular, if we can lower the burden of running a particular algorithm on a smaller quantum computer.

Research question 3
Do there exist methods to reduce the number of qubits required to run a given machine learning algorithm?

In chapter 4 we demonstrated one potential algorithm that was able to replicate the performance of a higher-qubit-count machine learning algorithm with fewer qubits, suggesting the question may resolve to yes. Chapter 5 demonstrated the limitations of this scheme and all other cut-local schemes, this chapter proved that as long as $\mathsf{BQP} \neq \mathsf{BPP}$ then the form of circuit cutting used in the preceding chapter could never successfully cut **all** circuits.

7.1.1 Future work

The chapter on the separation of FBQP and FBPP highlighted further improvements, particularly lowering the collapse of the polynomial hierarchy even further. In the similar case of **exactly sampling** from a quantum computer Fuji et al. [21] were able to collapse the hierarchy to $\mathsf{AM} \cap \mathsf{coAM}$ if the task was possible classically. While their techniques did not extend to the practically relevant case captured by SampBQP, they still provide information on what might be possible with further work. Indeed, our collapse is a much better starting point for further collapses.

For advice and advice following/giving models, it would be interesting to see where the advice classes could be deployed. While the surrogate model proposed by chapter 6 is in some sense optimal up to a polynomial factor this polynomial factor could be large. It would therefore be interesting to see other advice-following/generating models be developed and to compare how these perform in real-world tasks. Further work stemming from chapter 3 could focus further on developing the connections between bounded advice classes, as opposed to linking back to other classes as we did.

While chapter 5 proved that cut-local schemes would never be able to remove even a single qubit efficiently, there is still work to be done to generalise this result. While a fully general theorem is essentially the field of circuit complexity it may be possible to produce better bounds on the ability of more complicated circuit-cutting algorithms. This type of research will naturally also suggest how a more efficient machine learning algorithm could be developed to combine multiple small quantum circuits into outputs that would traditionally need larger circuits.