Graduate Research Plan Statement

Ronak Ramachandran | UT Austin | NSF GRFP 2025

Objective: Apply the tools of fine-grained complexity in quantum complexity theory to discover new quantum algorithms and conditional lower bounds on running times of quantum and classical algorithms.

Introduction and Background: Reductions are a fundamental tool of complexity theory allowing us to relate the running times of algorithms solving different problems. However, typical reductions within P, the class of efficiently solvable problems, and NP, the class of efficiently checkable problems, are coarse, in that they often ignore polynomial factors. These factors can make a big difference in practice. As the input length n of a problem grows, the difference between an algorithm running in time O(n) and $O(n^2)$ can be the difference between being possible to run on today's hardware and being practically infeasible. Fine-grained reductions let us relate running times more carefully, allowing us to see precisely how speedups in algorithms solving one problem can cascade into similar speedups for many others.

In 2005, Ryan Williams kickstarted the field of fine-grained complexity by establishing the first fine-grained reduction from an NP-complete problem, namely the satisfiability of certain Boolean formulas (CNF-SAT), to a problem in P, namely the Orthogonal Vectors problem (OV) [W05]. This meant that sub-quadratic algorithms for OV implied faster-than-brute-force algorithms for CNF-SAT, violating a widely believed conjecture known as the Strong Exponential Time Hypothesis (SETH). Since then, a mountain of follow-up work has related a small collection of similar conjectures to many other important problems in P. Computer scientists seeking speedups for these problems are now aware that doing so would necessarily be a more difficult task than disproving these widely believed conjectures.

The first applications of fine-grained complexity in quantum computing only came about in 2018, when Ambainis et. al. [ABIKPV19] gave a fine-grained reduction from Set-Cover to a minimum-finding problem, where Grover search could then yield a quadratic speedup. Shortly afterwards, Amir Abboud leveraged that reduction to develop faster quantum algorithms for problems like RNA-folding and to prove tighter conditional classical lower bounds for various algebraic and graph theoretic problems [A19]. Along a separate line of attack, two groups, Aaronson et. al. [ACLWZ20] and Buhrman et. al. [BPS21], defined and applied quantum versions of SETH, appropriately named QSETH, to prove new conditional lower bounds on the quantum complexities of various problems from Closest Pair to Longest Common Subsequence. The field is still new, however, and many open questions and research directions in quantum fine-grained complexity remain unexplored.

Proposal: I am currently exploring two related research directions in quantum fine-grained complexity.

Direction 1 (Quantum Speedups and/or Lower Bounds for String Comparison Problems): Longest Common Subsequence (LCS), and more generally, Edit Distance, are extremely well-studied string comparison problems with $O(n^2)$ time dynamic programming algorithms [WF74]. Assuming SETH, this running time is tight classically, but even assuming certain versions of QSETH, the best known quantum lower bounds are only $O(n^{1.5})$ [BPS21]. This leaves open the possibility of quantum speedups for both problems. I am currently searching for subquadratic quantum algorithms for both problems while simultaneously attempting to prove quadratic lower bounds conditioned on versions of QSETH following the framework of Buhrman, Patro, and Speelman [BPS21]. Until tight upper and lower bounds are known, concurrent work toward both contradictory goals can lead to novel insights in both directions.

Direction 2 (Quantum and Classical Lower Bounds on the Hidden Subgroup Problem): Kitaev's generalization of Shor's algorithm [K95] famously solves the Hidden Subgroup Problem (HSP) for Abelian groups in polynomial time on a quantum computer, but the best-known algorithms for general non-Abelian groups have exponential quantum circuit complexity [EHK04]. Even loose quantum lower bounds on this problem have been elusive outside of the oracle and sample complexity settings [YL22], but the generality of HSP makes reductions to it very natural. Fine-grained reductions to non-Abelian instances of HSP from problems conjectured to be quantumly hard could provide new evidence against better quantum algorithms. On the flip side, fine-grained reductions to Abelian HSP from problems conjectured to be classically hard could provide stronger evidence for quantum advantage.

Future Directions (Novel Hardness Conjectures): In addition to the hardness conjectures defined by Aaronson et. al. [ACLWZ20] and Buhrman et. al. [BPS21], which are based on the quantum hardness of circuit satisfaction problems, we might conjecture certain quantum problems are classically hard. For instance, it is still unknown whether BQP-complete problems, the hardest problems solvable efficiently by a quantum computer, require exponential time for a classical computer. If we conjecture that a BQP-complete problem takes exponential time classically, could we prove stronger conditional classical lower bounds? To do so we would need novel fine-grained reductions from such BQP-complete problems into better understood classical problems. We could also consider weaker conjectures to get more believable lower bounds. For instance, we might conjecture the hardness of minimizing the energy of a physical system with local interactions, known as the k-Local Hamiltonian Problem (k-LHP). This problem arises in nearly every area of physics, and it captures the difficulty of QMA, the quantum analog of NP. What problems admit fine-grained reductions from k-LHP? These questions, though less explored, motivate developing strong quantum fine-grained complexity techniques for my main two aims above.

Intellectual Merit

Query complexity—my most recent research area—is unable to prove superlinear lower bounds on the running times of algorithms, and **fine-grained complexity is currently one of the best methods** we have to break this barrier, albeit, assuming the truth of a few conjectures. In general, fine-grained complexity advances knowledge by consolidating the hardness of many desired algorithmic speedups into a few hard-to-prove conjectures. This allows complexity theorists to spend more time on a focused set of problems and sets up the possibility that **a few algorithmic breakthroughs might immediately cascade into speedups for many others.** Expanding this area of research into the domain of quantum complexity theory lets us use the best tools at our disposal to partially address the central question of the field: **what can quantum computers help us solve more quickly, and what problems resist speedups?**

Broader Impacts

Longest Common Subsequence and Edit Distance have **a wide range of applications in basic biological research**. Faster algorithms for the two problems can help researchers more quickly find commonalities and differences in DNA sequences, which is crucial for understanding **evolutionary relationships, genetic mutations, and protein structure.** A heuristic algorithm to approximately solve these problems in nearly linear time, the Basic Local Alignment Search Tool (BLAST), is used heavily by bioscience researchers, with more than 114,000 citations to date. Knowing whether or not quantum computers could provide speedups to these problems would motivate researchers to invest more time and effort into the best technologies to accelerate future biological breakthroughs.

Important problems like Factoring, Discrete Log, and Graph Isomorphism are all special cases of the Hidden Subgroup Problem (HSP), and these problems have **crucial applications in cryptography [RSA78][DH76].** These special cases of HSP are joined by their unique status as problems in NP that are not known to be in P or NP-complete. Proving new conditional lower bounds on such problems could give us **a clearer picture of the relative power of quantum and classical computers.**

References:

[W05] Williams. (2005). [ABIKPV19] Ambainis et. al. (2019). [A19] Abboud. (2019). [ACLWZ20] Aaronson et. al. (2020). [BPS21] Buhrman, Patro, and Speelman. (2021). [WF74] Wagner and Fischer (1974). [K95] Kitaev. (1995). [EHK04] Ettinger, Høyer, and Knill. (2004). [YL22] Ye and Li. (2022). [RSA78] Rivest, Shamir, and Adleman. (1978). [DH76] Diffie and Hellman. (1976). [KKVB02] Kashefi et. al. (2002).