Research Summary

Subhasree Patro

One of the major challenges in the field of complexity theory (both classical and quantum) is the inability to prove unconditional time lower bounds. One way around this is the study of fine-grained complexity, where we use special reductions to prove time lower bounds for many problems in P based on the conjectured hardness of some key problems like the Satisfiability (SAT) problem, 3SUM or APSP. The situation in the quantum regime is no better; almost all known lower bounds are defined in terms of query complexity, which is not very useful for problems whose best-known algorithms take superlinear time. Therefore in order to understand the quantum time complexity of certain problems, employing fine-grained reductions in the quantum setting seems a natural way forward, which most of my PhD work has been about.

1 Contributions in Quantum Fine-Grained Complexity

We studied some existing classical reductions from problems like CNF-SAT, 3SUM and APSP and observed that translating these classical fine-grained reductions directly into the quantum regime is not always trivial. Fortunately, we were able to develop frameworks and proof strategies, using which we are able to circumvent these obstacles and were able to comment on the time complexity of a lot of string comparison, computational geometry, graph theoretic and other related problems.

1.1 A Framework of Quantum Strong Exponential-Time Hypotheses

The strong exponential-time hypothesis (SETH) is a commonly used conjecture in the field of complexity theory. It essentially states that determining whether a CNF formula is satisfiable cannot be done faster than exhaustive search over all possible assignments. This hypothesis and its variants gave rise to a fruitful field of research obtaining (mostly tight) lower bounds for many problems in P whose unconditional lower bounds are very likely beyond current techniques.

Our contribution In this joint work with Harry Buhrman and Florian Speelman, we introduce an extensive framework of Quantum Strong Exponential-Time Hypotheses, as quantum analogues to what SETH is for classical computation. Using the QSETH framework, we are able to translate quantum query lower bounds on black-box problems to conditional quantum time lower bounds for many problems in P. As an example, we provide a conditional quantum time lower bound of $\Omega(n^{1.5})$ for the Longest Common Subsequence and Edit Distance problems. We also show that the n^2 SETH-based lower bound for a recent scheme for Proofs of Useful Work carries over to the quantum setting using our framework.

The conference version of this paper is published in STACS 2021 and also was presented in the non-proceedings track of TQC 2020.

1.2 Fine-Grained Complexity via Quantum Walks

A fundamental conjecture in the classical setting states that the 3SUM problem cannot be solved by (classical) algorithms in time $O(n^{2-\epsilon})$ for an $\epsilon > 0$. Consequently, lower bounds for many problems were concluded based on this conjecture.

Our contribution In this joint work with Harry Buhrman, Bruno Loff, and, Florian Speelman, we further extend the theory of fine-grained complexity in the quantum regime. We formulate an analogous conjecture, the Quantum-3SUM-Conjecture, which states that there exist no sublinear $O(n^{1-\alpha})$ time quantum algorithms for the 3SUM problem. Based on the Quantum-3SUM-Conjecture, we show new lower-bounds on the time complexity of quantum algorithms for several computational problems. These results are proven by adapting to the quantum setting known classical fine-grained reductions from the 3SUM problem. This adaptation is not trivial, since the original classical reductions require sorting the input, and sorting provably cannot be done in sublinear quantum time. We overcome this bottleneck by combining a quantum walk with a classical dynamic datastructure having a certain "history-independence" property. This (general) proof strategy allows us to prove tight lower bounds on several computational geometry problems, on Convolution-3SUM and on the 0-Edge-Weight-Triangle problem, conditional on the Quantum-3SUM-Conjecture.

The conference version of this work appeared in ITCS 2022 and was presented at QIP 2022 and also in the non-proceedings track of TQC 2021.

1.3 Memory Compression with Quantum Random-Access Gates

The quantum fine-grained results from our earlier work lead us to some interesting observations: There are quite a few quantum walk based algorithms that are space *efficient* but are complicated and extremely non-trivial to construct. There are however space *inefficient* analogous algorithms that have a much simpler structure and are far more easy to construct. Unfortunately, making a quantum algorithm space inefficient is not at all appealing especially when maintaining even a few hundred qubits coherently seem like a daunting task. In this project, we present techniques using which we can compress a space inefficient quantum algorithm into a space efficient one with only slight increase in time complexity and slight worsening of error.

Our contribution In this joint work with Harry Buhrman, Bruno Loff and Florian Speelman, we prove the following result. If we have a quantum algorithm that runs in time T and uses M qubits, and which is such that the state of the memory, at any time step, is supported on vectors of Hamming weight at most m, then it may be simulated by another algorithm which uses only $O(m \log M)$ memory. Using some known quantum and classical techniques, we obtain an algorithms running in time $O(T \log T \log M)$. We show how this theorem can be used, in a black-box way, to dramatically simplify several papers, including our 3SUM paper (mentioned above). Broadly speaking, when there exists a need for a space-efficient history-independent quantum data-structure, it is much simpler to construct a space-inefficient, yet sparse, quantum data structure.

Full version of this paper is available on arXiv:2203.05599 and got published in the proceedings of TQC 2022.

1.4 Quantum complexity of APSP-hard problems

The All Pairs Shortest Path (APSP) problem is one of the well studied problems in graph algorithms. It is defined as follows: Given a graph G of nodes and weighted directed edges, output the shortest paths between every pair of nodes in G. The best known classical algorithm to solve APSP on a graph of n nodes runs in $O(n^3)$ time. No significant improvement to this run-time is known. The APSP problem has been useful in obtaining $\Omega(n^3)$ classical lower bounds for a variety of computational problems, for example, Matrix Multiplication over Semiring,[1] Detecting Negative Triangle, Zero Edge Weight Triangle, Matching Triangles, Triangle Collection, etc. Subcubic algorithms for any of these problems imply a subcubic algorithm for APSP. Moreover, most of these problems are cubic-equivalent, i.e., these problems either all have truly subcubic algorithms, or none of them do.

Our contribution. In this joint work with Andris Ambainis, Harry Buhrman, Keon Leinsje and Florian Speelman we attempted to understand the quantum complexities of the APSP problem and other related problems. There exists a $O(n^{2.5})$ quantum time algorithm for solving APSP. We observe that the cubic equivalence that holds in the classical setting no longer holds in the quantum case with their time complexity ranging from $O(n^{1.5})$ to $O(n^{<3})$ for these problems. We studied the classical reductions from APSP to some problems and observe that almost all those reductions can be trivially adapted to the quantum setting, thus proving tight lower bounds for nearly all of them. The matching upper bounds for most of them can be derived using Grover-like speedups, except for Δ-Matching Triangles and Triangle Collection for which we present quantum algorithms that require careful use of data structures, Ambainis' variable time search and certain other ingredients as subroutines.

Additionally, just like in the classical case, hardness results for Δ-Matching Triangles and Triangle Collection can be derived conditioned on hardness of any of the three key problems, i.e., CNF-SAT, 3SUM and APSP. Not only does this allows us to check the consistency of our approaches, but clearly for problems such as Δ-Matching Triangles and Triangle Collection a weaker quantum hardness assumption can be made to show these time lower bounds.

The online version of this work can be found at arXiv:2207.11068. We were also invited by the Hon-hai research institute in Taiwan to virtually present this result.

2 Other Research Projects

2.1 Improved Quantum Query Upper Bounds Based on Classical Decision Trees

One of the other approaches towards understanding hardness of a problem is to set the goals a little bit lower and try to understand this problem on a simpler model of computation. Perhaps the simplest model

[1] This should not be mistaken for the usual Matrix Multiplication problem which has a $O(n^\omega)$ time algorithm with ω being the matrix multiplication constant currently at 2.3728596.

of computation is the Decision Tree. Decision tree complexities in the classical, randomized, and quantum settings are well defined and very well studied.

Our contribution In this joint work with Arjan Cornelissen and Nikhil S. Mande, we first define the randomized decision tree size complexity of a relation $f \subseteq \{0,1\}^n \times R$, let's denote by $RDTsize(f)$, and explore connections with its bounded-error quantum query complexity $Q(f)$. We prove that $Q(f) \leq O(\sqrt{RDTsize(f)})$ for any relation f. We do so by giving an explicit span program and dual adversary solution witnessing the same. Lin and Lin [ToC'16] and Beigi and Taghavi [Quantum'20] showed results of a similar flavor, and gave upper bounds in terms of a quantity which we call the "guessing complexity" of a decision tree. We identify that the guessing complexity of a decision tree equals its rank, a notion introduced by Ehrenfeucht and Haussler [Inf. Comp.'89] in the context of learning theory. This answers a question posed by Lin and Lin, who asked whether the guessing complexity of a decision tree is related to any complexity-theoretic measure. We also show a polynomial separation between rank and randomized rank for the complete binary AND-OR tree. Beigi and Taghavi constructed span programs and dual adversary solutions for Boolean functions given classical decision trees computing them and an assignment of non-negative weights to its edges. We explore the effect of changing these weights on the resulting span program complexity and objective value of the dual adversary bound, and capture the best possible weighting scheme by an optimization program. We exhibit a solution to this program and argue its optimality from first principles. We also exhibit decision trees for which our bounds are asymptotically stronger than those of Lin and Lin, and Beigi and Taghavi. This answers a question of Beigi and Taghavi, who asked whether different weighting schemes could yield better upper bounds.

The full version of this paper is available on arXiv:2203.02968 and was published in the proceedings of FSTTCS 2022 and was also presented in the non-proceedings track of TQC 2022.

3 Research Contributions during Masters at IIIT-Hyderabad

Prior to my PhD in my masters, I had co-authored the following results in the area of Quantum Information Theory as a first author and also had a small contribution to a paper that recently appeared with the title 'Teleportation of Quantum Coherence' (arXiv:2302.11499).

3.1 Impossibility of Cloning of Quantum Coherence

Our contribution. It is well known that it is impossible to clone an arbitrary quantum state. However, this inability does not lead directly to no cloning of quantum coherence. Here, in this joint work with Dhrumil Patel, Chiranjeevi Vanarasa, Indranil Chakrabarty, and, Arun Kumar Pati, we show that it is impossible to clone the coherence of an arbitrary quantum state. In particular, with an ancillary system as machine state, we show that it is impossible to clone the coherence of states whose coherence is greater than the coherence of the known states on which the transformations are defined. Also, we characterize the class of states for which coherence cloning will be possible for a given choice of machine. Furthermore, we find the maximum range of states whose coherence can be cloned perfectly. The impossibility proof also holds when we do not include machine states. Lastly, we generalize the impossibility of cloning of coherence in terms of dimension of the quantum state and coherence measure taken into consideration.

The journal version of this work appeared in Phys. Rev. A 103 in 2021.

3.2 Non-negativity of conditional von Neumann entropy and global unitary operations

Our contribution. Conditional von Neumann entropy is an intriguing concept in quantum information theory. In this joint work with Indranil Chakrabarty and Nirman Ganguly, we examine the effect of global unitary operations on the conditional entropy of the system. We start with a set containing states with a non-negative conditional entropy and find that some states preserve the non-negativity under unitary operations on the composite system. We call this class of states the absolute conditional von Neumann entropy non-negative (ACVENN) class. We characterize such states for $2\otimes2$–dimensional systems. From a different perspective the characterization accentuates the detection of states whose conditional entropy becomes negative after the global unitary action. Interestingly, we show that this ACVENN class of states forms a set which is convex and compact. This feature enables the existence of Hermitian witness operators. With these we can distinguish the unknown states which will have a negative conditional entropy after the global unitary operation. We also show that this has immediate application to superdense coding and state merging, as the negativity of the conditional entropy plays a key role in both these information processing tasks. Some illustrations followed by analysis are also provided to probe the connection of such states with absolutely separable states and absolutely local states.

The journal version of this work appeared in Phys. Rev. A 96 in 2017.

4 Future Directions

The following are a few concrete examples of problems that I currently working on and plan to explore in the immediate future.

- The $O(n^2)$ time quantum upper bound for both Edit Distance and LCS is not known to be optimal as the lower bound is still $\Omega(n^{1.5})$. We are trying to understand the exact complexities of string comparison problems including Edit Distance and LCS in the quantum setting; these problems are very relevant to the industry, especially in the field of computational biology and linguistics.

- Exploring quantum complexity of other string comparison problems such as Fréchet Distance, Dynamic Time Warping, etc, are also interesting questions.

- As another application of QSETH, we are trying to understand the quantum time complexity for post-quantum cryptography problems, more specifically for problems in lattice-based cryptography. For now, we are trying to prove quantum fine-grained lower bounds for variants of the closest vector problem (CVP) and shortest vector problem (SVP).

- There are SETH-based classical fine-grained results known for the strong simulation of quantum circuits. It will be interesting to explore similar results in the quantum setting via our QSETH framework.

- For a few computational geometry problems their exact complexities still remain unclear. It is possible that by using variants of Quantum-3SUM-Conjecture we can further improve these lower bounds. As part of my postdoc collaboration with the computer science department of Utrecht University (which has a strong computational geometry group), I plan to understand the exact complexities of such computational geometry problems better. Quantum upper bound for counting version of the 3SUM problem was recently explored in this paper; it will be interesting to study this problem from a lower bound perspective.

- Our 3SUM result and the memory compression results opens up the following question in the context of quantum memory and data structures: is it possible to have quantum data structures that are deterministic, history-independent, memory efficient and allow for insertions, deletions and lookups in poly-logarithmic time? Or, is there an impossibility result suggesting otherwise?

Additionally, I would also like to broadly work on the following topics in the coming future which (I think) naturally connects to the world of quantum fine-grained complexity.

- Expand the quantum fine-grained reductions further to include other key problems and other computational problems, as done in Schoneveld's bachelors' thesis [Sch22] that I co-supervised.

- The results mentioned here discuss the worst-case complexity of computational problems. However, in practice, it will be beneficial to also look at the average-case complexity. I plan to pursue that in the future, especially in the context of the average case fine-grained hardness of cryptography problems.

- Parameterised complexity is a field where we try to study which is the 'hard' parameter responsible for the complexity. This area is also under-explored in the quantum setting.