Symmetries, graph properties, and quantum speedups

Shalev Ben-David   Andrew M. Childs   András Gilyén
William Kretschmer   Supartha Podder   Daochen Wang

One of the most fundamental problems in the field of quantum computing is the question of when quantum algorithms can substantially outperform classical ones. While polynomial quantum speedups are known in many settings, super-polynomial quantum speedups are known (or even merely conjectured) for only a few select problems. Crucially, exponential quantum speedups only occur for certain "structured" problems such as period-finding (used in Shor's factoring algorithm [Sho94]) and Simon's problem [Sim97], in which the input is known in advance to have a highly restricted form. In contrast, for "unstructured" problems such as black-box search or NP-complete problems, only polynomial speedups are known (and in some models, it can be formally shown that only polynomial speedups are possible). In this work, we are interested in formalizing and characterizing the structure necessary for fast quantum algorithms. In particular, we study the *types of symmetries* a function can have while still exhibiting super-polynomial quantum speedups. We show that (hyper)graph properties in the adjacency matrix model cannot have exponential quantum speedup; that such hypergraph symmeteries are essentially the only symmetries that forbid exponential speedup; but that exponential speedup is possible for graph property testing in the adjacency list model.

1 Background

Despite the widely held intuition that structure is necessary for super-polynomial quantum speedups, only a handful of works have attempted to formalize this and characterize the required structure. All of them study the problem in the query complexity (black-box) model[1] of quantum computation, which is a natural framework in which both period-finding and Simon's problem can be formally shown to give exponential quantum speedups.

Beals, Buhrman, Cleve, Mosca, and de Wolf [BBC$^+$01] showed that all *total Boolean functions* $f\colon \Sigma^n \to \{0,1\}$ have a polynomial relationship between their classical and quantum query complexities (which we denote $\mathrm{R}(f)$ and $\mathrm{Q}(f)$, respectively). This means that super-polynomial speedups are *only* possible in query complexity if we impose a promise on the input: that is, we define $f\colon P \to \{0,1\}$ with $P \subset \Sigma^n$. But not all partial functions exhibit super-polynomial quantum speedups. The question, then, is what we can say about the structure necessary for a partial Boolean function f to exhibit a super-polynomial quantum speedup. Towards this end, Aaronson and Ambainis [AA14] showed that *symmetric* functions do not allow super-polynomial quantum speedups, even with a promise. Chailloux [Cha18] improved this result by reducing the degree of the polynomial relationship between randomized and quantum algorithms for symmetric functions.

A class of problems with significant symmetry, though much less than full permutation symmetry, is the class of graph properties. For such problems, the input describes a graph, and the output depends only on the isomorphism class of that graph. In other words, the output is invariant under permuting the input in a way that is consistent with relabelling the vertices.

The setting of graph property *testing* provides a natural class of partial graph properties. Here we are promised that the input graph either has a property, or is ϵ-far from having the property, meaning

[1] In the query complexity model, the goal is to compute a Boolean function $f\colon \Sigma^n \to \{0,1\}$ using as few queries to the bits of the input $x \in \Sigma^n$ as possible, where Σ is some finite alphabet.

that we must change at least an ϵ fraction of the edges to make the property hold. Graph property testing has been extensively studied since its introduction by Goldreich, Goldwasser, and Ron [GGR98].

The behavior of classical graph property testers can differ substantially depending on the model in which the input graph is specified. For example, in the adjacency matrix model, Alon and Krivelevich [AK02] proved that bipartiteness can be tested in $\tilde{O}(1/\epsilon^2)$ queries, which is surprisingly independent of the size of the input graph. In contrast, in the adjacency list model for bounded-degree graphs, Goldreich and Ron [GR97] proved that $\Omega(\sqrt{n})$ queries are needed to test bipartiteness of n-vertex graphs.

Quantum algorithms for testing properties of bounded-degree graphs in the adjacency list model were studied by Ambainis, Childs, and Liu [ACL11]. They gave upper bounds of $\tilde{O}(n^{1/3})$ for testing bipartiteness and expansion, demonstrating polynomial quantum speedups. Furthermore, they showed that at least $\Omega(n^{1/4})$ quantum queries are required to test expansion, ruling out the possibility of an exponential quantum speedup. This work naturally raises the question (also highlighted by Montanaro and de Wolf [MdW16]) of whether there can ever be exponential quantum speedup for graph property testing, or for graph properties more generally.

2 Our contributions

In this work, we extend the results of Aaronson-Ambainis and Chailloux to other symmetry groups. We characterize general symmetries of functions as follows.

Definition 1. *Let $f\colon P \to \{0,1\}$ be a function with $P \subseteq \Sigma^n$, where Σ is a finite alphabet and $n \in \mathbb{N}$. We say that f is* symmetric *with respect to a permutation group G acting on domain $[n]$ if for all $x \in P$ and all $\pi \in G$, the string $x \circ \pi$ defined by $(x \circ \pi)_i \coloneqq x_{\pi(i)}$ satisfies $x \circ \pi \in P$ and $f(x \circ \pi) = f(x)$.*

The case where $G = S_n$, the fully-symmetric permutation group, is the scenario handled in Chailloux's work [Cha18]: he showed that $\mathrm{R}(f) = O(\mathrm{Q}(f)^3)$ if f is symmetric under S_n. Aaronson and Ambainis [AA14] required an even stronger symmetry property. We note that when Σ is large, say $|\Sigma| = n$ or larger, the class of functions symmetric under S_n is already highly nontrivial: among others, it includes functions such as Collision, an important function whose quantum query complexity was established in [AS04]; k-Sum, whose quantum query complexity required the negative-weight adversary to establish [BŠ13]; and k-Distinctness, whose quantum query complexity is still open [BKT18].

In this work, we examine what happens when we relax the full symmetry S_n to smaller symmetry groups G. We introduce some tools for showing that particular classes of permutation groups G do not allow super-polynomial quantum speedups; that is, we provide tools for showing that every f symmetric with respect to G satisfies $\mathrm{Q}(f) = \mathrm{R}(f)^{\Omega(1)}$. Our first main result is the following theorem, in which G is the *graph symmetry*: the permutation group acting on strings of length $\binom{n}{2}$ (which represent the possible edges of a graph), which includes all permutations of the edges induced by the S_n permutations (or relabelings) of the n vertices. Boolean functions that are invariant under permuting the input string by the graph symmetry group are called graph properties in the adjacency matrix model.

Theorem 2. *Any (partial) graph property f in the adjacency matrix model has $\mathrm{R}(f) = O(\mathrm{Q}(f)^6)$.*

This theorem holds even when the input alphabet of f is non-Boolean and extends to other types of graph symmetries, including hypergraph symmetries and bipartite graph symmetries. The proof uses the argument of Chailloux [Cha18] and a reduction to the $G = S_n$ case. We introduce a notion of "well-shuffling" that characterizes permutation groups in which it is hard for a quantum algorithm to distinguish a random permutation from a random small-range function. We apply Chailloux's argument to show that $R(f)$ and $Q(f)$ are polynomially related for f symmetric under well-shuffling G and apply the reduction argument to lift the well-shuffling property of S_n, which follows from the collision lower bound, onto the graph symmetry group.

Next, we extend Theorem 2 to give a more general characterization of how symmetries relate to (super-)polynomial quantum speedups. Our main result in this regard is a dichotomy theorem for *primitive* permutation groups. Primitive permutation groups are sometimes described as the "building blocks" of permutation groups, because arbitrary permutation groups can always be decomposed into primitive groups (in a certain formal sense). We give a complete classification of which primitive groups G allow super-polynomial quantum speedups and which do not, in terms of the *minimal base size* $b(G)$. The minimal base size is a key quantity in computational group theory that roughly captures the size of a permutation group G relative to the number of points it acts on. Thus, the following theorem amounts to showing that symmetries of "large" primitive permutation groups do not allow large quantum speedups, while symmetries of "small" primitive permutation groups do.

Theorem 3. *Let G be a primitive permutation group acting on $[n]$, and let $b(G)$ denote the size of a minimal base for G. If $b(G) = n^{\Omega(1)}$, then $R(f) = Q(f)^{O(1)}$ for every f that is symmetric under G. Otherwise, if $b(G) = n^{o(1)}$, then there exists a partial function f that is symmetric under G such that $R(f) = Q(f)^{\omega(1)}$.*

By decomposing an arbitrary permutation group into primitive groups, we can extend the above theorem to show that if a permutation group G does not allow super-polynomial speedups, then it must be constructed out of a constant number of primitive groups H that satisfy $b(H) = n^{\Omega(1)}$, where H acts on n points. Put another way, if any of the primitive factors H of G satisfy $b(H) = n^{o(1)}$, or if G contains more than a constant number of primitive factors, then we exhibit a function that is symmetric under G with a super-polynomial quantum speedup.

Remarkably, the proof of Theorem 3 also includes a strong characterization of what the primitive permutation groups G that satisfy $b(G) = n^{\Omega(1)}$ look like. Indeed, we show that as a consequence of the classification of finite simple groups, all such groups essentially look like hypergraph symmetries or minor extensions thereof. Thus, in some sense, permutation groups built out of hypergraph symmetries are the *only* permutation groups that are inconsistent with super-polynomial speedups. We consider this one of the most surprising consequences of our work; a priori, one could expect there to be many different kinds of symmetries that disallow super-polynomial speedups.

Finally, we return to graph properties. While the aforementioned results show that no exponential speedup is possible for graph properties in the adjacency *matrix* model, they do not resolve the original open question of Ambainis, Childs, and Liu [ACL11], which specifically addressed graph property *testing* in the adjacency *list* model. Graph symmetries in this model manifest themselves in a different way that is not captured by Definition 1, so Theorem 2 does not apply.[2] In this model, we show the following:

Theorem 4. *There exists a graph property testing problem in the adjacency list model for which there is an exponential quantum speedup.*

The proof uses the fact that we can classically test whether a graph is a welded-trees graph [CCD+03] given advice $b_v \in \{0, 1\}$ indicating whether each vertex v is a leaf of not. For each v, we hide b_v as the parity of two bits stored at the two roots of another copy of the welded-trees graph. Then b_v can be efficiently read by a quantum computer that can traverse the welded trees but cannot be efficiently read by a classical computer.

Of course, Theorem 4 is in stark contrast to Theorem 2, the situation in the adjacency matrix case. Together, Theorem 2 and Theorem 4 fully settle the open question about the possibility of exponential quantum speedup for graph properties, showing that its answer is highly model-dependent: exponential speedup is impossible in the adjacency matrix model, but is possible in the adjacency list model, even in the restrictive setting of graph property testing.

[2] Given a graph $x\colon [n] \times [d] \to [n] \cup \{*\}$ in the adjacency list model with n vertices and bounded degree d, its isomorphism class essentially consists of graphs of the form $\pi^{-1} \circ x \circ (\pi \times \mathrm{id}_{[d]})$ where $\pi \in S_n$ is a relabeling of vertices and π^{-1} is defined as the inverse of π but also maps $*$ to $*$. A function f is a graph property in this model if and only if its domain P is a union of such isomorphism classes and f is constant on each isomorphism class.

References

[AA14] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. *Theory of Computing*, 10:133–166, 2014. arXiv:0911.0996

[ACL11] Andris Ambainis, Andrew M. Childs, and Yi-Kai Liu. Quantum property testing for bounded-degree graphs. In *Proceedings of the 15th International Workshop on Randomization and Computation (RANDOM)*, volume 6845 of *Lecture Notes in Computer Science*, pages 365–376. Springer, 2011. arXiv:1012.3174

[AK02] Noga Alon and Michael Krivelevich. Testing k-colorability. *SIAM Journal on Discrete Mathematics*, 15(2):211–227, 2002.

[AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, July 2004.

[BBC$^+$01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS'98. arXiv:quant-ph/9802049

[BKT18] Mark Bun, Robin Kothari, and Justin Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. In *Proceedings of the 50th ACM Symposium on the Theory of Computing (STOC)*, pages 297–310, 2018. arXiv:1710.09079

[BŠ13] Aleksandrs Belovs and Robert Špalek. Adversary lower bound for the k-sum problem. In *Proceedings of the 4th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 323–328, 2013. arXiv:1206.6528

[CCD$^+$03] Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by quantum walk. In *Proceedings of the 35th ACM Symposium on the Theory of Computing (STOC)*, pages 59–68, 2003. arXiv:quant-ph/0209131

[Cha18] André Chailloux. A note on the quantum query complexity of permutation symmetric functions. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 19:1–19:7, 2018. arXiv:1810.01790

[GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.

[GR97] Oded Goldreich and Dana Ron. Property testing in bounded degree graphs. In *Proceedings of the 29th ACM Symposium on the Theory of Computing (STOC)*, page 406–415, 1997.

[MdW16] Ashley Montanaro and Ronald de Wolf. *A survey of quantum property testing*. Number 7 in Graduate Surveys. Theory of Computing Library, 2016. arXiv:1310.2035

[Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 124–134, 1994. arXiv:quant-ph/9508027

[Sim97] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.