

Decoded Quantum Interferometry is Obstructed by Topological Structure

This submission is based off of [arXiv:2509.14509](#) but includes significant technical improvements that will soon appear in a v2 on the arXiv. Details on these improvements at end of abstract.

Eric R. Anschuetz, David Gamarnik, and Jonathan Z. Lu

Quantum algorithms which are known or believed to give significant speedups, such as factoring integers or solving certain systems of equations, generally have relatively rigid algebraic structure [1–4]. At the same time, a tantalizing potential application for quantum algorithms is in *optimization*. In this setting, one is interested in computing $\mathbf{z}^* = \operatorname{argmax}_{\mathbf{z} \in \{0,1\}^n} g(\mathbf{z})$ given some objective function $g : \mathbb{F}_2^n \rightarrow \mathbb{R}$. While complexity theory suggests that, at least in the

worst case, optimizing a generic objective function g to a sufficiently large constant approximation ratio is NP-hard, there is nevertheless hope of a *quantum-classical approximation gap*. This possibility is particularly attractive in the “average case,” wherein the function g is drawn from certain natural ensembles of problem instances. Namely, it is plausible that efficient quantum algorithms are able to achieve a function value g_q^* while efficient classical algorithms are only able to achieve a value $g_c^* < g_q^*$ with some non-negligible probability over the distribution of problem instances.

The classical average-case hardness of optimization has been studied through the use of statistical techniques which are themselves intimately related to learning [5]. These efforts have led to a conjecture that g_c^* is tightly characterized by the function value g_{OGP}^* at which the *overlap gap property* (OGP) manifests: a shattering of the space of near-optimal solutions into many well-separated clusters [6–11] (see Fig. 1). The onset of this topological feature is rigorously known to inhibit classical algorithms which are *Lipschitz functions* of their inputs from preparing near-optimal solutions, and is also rigorously known to be a tight characterization for such algorithms in many settings [10, 11]. The onset of the OGP also coincides with the approximation ratios achieved by the best-known efficient algorithms for random k -SAT, maximum independent set, optimization of spin glasses, and many more problem ensembles [6].

Conditioned on this well-established conjecture, then, the question of average-case quantum advantage in our setting becomes:

Question 1. For a given optimization problem, is there a quantum algorithm that overcomes the OGP barrier?




FIG. 1. The *overlap gap property* is a topological feature of the space of near-optimal solutions of combinatorial optimization problems that is widely believed to characterize their hardness. It is characterized by a clustering of near-optimal solutions (small circles) into many disconnected clusters (large circles) in the solution space under a particular metric.

Due to the ubiquity and widespread impact of combinatorial optimization problems across a plethora of fields and applications, much effort has been devoted to the development of quantum algorithms which can efficiently find solutions $\mathbf{z}^*$ whose achieved function value can exceed g_{OGP}^* [12, 13]. Unfortunately, it is now known that many of these algorithms are unable to overcome the OGP barrier [5, 14–21]; for this reason, it seems unlikely that these will exhibit any quantum advantage for typical, unstructured optimization problems.

However, a new quantum algorithm for optimization has recently been developed that is fundamentally different from the others: *Decoded Quantum Interferometry* (DQI), a quantum algorithm for MAX- k -XOR-SAT [22]. In MAX- k -XOR-SAT, one is given m linear equations in $\mathbb{F}_2$ over $n < m$ variables, and is asked to find a variable assignment $\mathbf{z} \in \mathbb{F}_2^n$ satisfying as many clauses as possible. More concretely, a problem instance is encoded as an $m \times n$, k -row sparse constraint matrix $\mathbf{B} \in \mathbb{F}_2^{m \times n}$ as well as a vector of parities $\mathbf{v} \in \mathbb{F}_2^m$, and one hopes to minimize $g(\mathbf{z}) := m - \|\mathbf{B}\mathbf{z} \oplus \mathbf{v}\|_1$ over $\mathbf{z} \in \mathbb{F}_2^n$.

DQI acts as a *quantum reduction* from instances of the MAX-XOR-SAT optimization problem to the bounded distance decoding of a classical linear code with parity check matrix $\mathbf{B}^\top \in \mathbb{F}_2^{m \times m}$. Unlike previous quantum heuristics, the average performance of DQI can actually be *guaranteed* given a classical decoder known to efficiently decode up to

ℓ errors in the m -bit code specified by $\mathbf{B}^\top$. The expected fraction of satisfied clauses is given by a semicircle law [22]:

$$\mu_{\text{DQI}} := \frac{\langle g \rangle_{\text{DQI}}}{m} = \left(\sqrt{\frac{\ell}{2m}} + \sqrt{\frac{1}{2} \left(1 - \frac{\ell}{m} \right)} \right)^2. \quad (1)$$

Armed with this performance guarantee, one can search for structured distributions over MAX-XOR-SAT for which the corresponding decoding problem is easy in order to find problems for which μ_{DQI} is large. For instance, one can consider a special case of MAX-XOR-SAT called optimal polynomial intersection (OPI), where the decoding problem reduces to decoding a Reed–Solomon code [23] (a highly structured algebraic code constructed from polynomials). It is conjectured that DQI exhibits a quantum advantage on OPI [22]. If true, this conjecture remains consistent with the status quo intuition that quantum advantage arises from algebraic structure.

However, in average-case analysis, we instead consider the setting in which the MAX-XOR-SAT instance is chosen randomly from an unstructured ensemble. Correspondingly, the decoding problem becomes that of a random code; we restrict to the case when $\mathbf{B}$ is k -sparse since efficient good decoders for random codes are only known for LDPC codes. A crucial question then is whether one should expect a quantum advantage in this randomized case. Preliminary diagnostics suggest that a quantum advantage may be unlikely: using state-of-the-art LDPC decoders known as *belief propagation* (BP), DQI numerically achieves approximation ratios below classical heuristics [22]. One might hope, however, that (a) quantum advantage could arise on instances of k, n beyond the range explored by current numerics, and (b) even if the *expected* approximation ratio achieved may not strictly beat classical algorithms, the best approximation ratio achievable *with noticeable probability* may still have an advantage (which could then be exploited by running DQI many times and taking the best solution). Thus, these preliminary diagnostics nevertheless give compelling motivation for a more rigorous tight characterization of the performance of DQI with noticeable probability, relative to the best classical algorithms. In particular:

Question 2. Can DQI exceed the OGP threshold for unstructured MAX- k -XOR-SAT instances?

Contributions. We show that, like other quantum algorithms for optimization, DQI *cannot* surpass the OGP barrier characterizing the performance of optimal efficient classical optimizers. Specifically, we consider a definition of Lipschitz quantum algorithms that is due to Ref. [5]. It is defined with respect to the *quantum Wasserstein distance* d_W [24]; informally, the quantum Wasserstein distance is an “earth-mover’s” metric in that states which differ by a p -qubit quantum channel differ by at most $O(p)$ in quantum Wasserstein distance.

Definition 3 (Stable quantum algorithms, informal [5]). Let $\mathcal{A}$ be a quantum algorithm, i.e., a map from problem instances $\mathbf{X}$ to quantum states $\mathcal{A}(\mathbf{X})$. We say $\mathcal{A}$ is *stable* if there exist L and sufficiently small f such that, with high probability over $\mathbf{X}$, $d_W(\mathcal{A}(\mathbf{X}), \mathcal{A}(\mathbf{X}')) \leq f + L \|\mathbf{X} - \mathbf{X}'\|_1$.

We show that the OGP inhibits stable quantum algorithms for MAX-XOR-SAT. A similar statement was shown in [5] for quantum spin glass problems exhibiting a certain quantum modification of the OGP; we show that stable quantum algorithms are also obstructed by the classical OGP in classical combinatorial optimization problems.

Theorem 4 (OGP inhibits stable quantum algorithms). *Let μ_{OGP} be the satisfied fraction at which the OGP occurs for MAX-XOR-SAT. Then, stable quantum algorithms can achieve no better than $g^* \leq \mu_{\text{OGP}}$ satisfied clauses with any constant probability.*

We note that Theorem 4 also applies to quantum algorithms beyond DQI which were previously shown to be stable, including standard quantum algorithms such as log-depth QAOA and phase estimation [5].

We then focus on the setting where the LDPC parity check matrix $\mathbf{B}^\top$ is drawn from the standard Gallager ensemble of LDPC codes on m bits with k -local checks and rate $r = 1 - n/m$. We then consider the following assumption on the decoder DQI utilizes, which holds for all best-known efficient decoders for LDPC codes of which we are aware.

Assumption 5 (DQI decoder conditions). The decoder satisfies one of the following two conditions:

1. The decoder is arbitrary, and is capable of correcting errors of relative weight at most:

$$\frac{\ell^*}{m} = \frac{1}{k\lambda}, \quad (2)$$

This is the best-known provably achievable rate achieved by efficient decoders on random LDPC codes [25–29]. We show explicitly that this holds for state-of-the-art exact decoders based on linear programming relaxations.

2. The decoder satisfies a certain Lipschitz condition. We prove that this class includes belief propagation/message-passing decoders in the standard density evolution model.

We show that, under the Gallager ensemble of parity check matrices and either of these two assumptions, the quantum state prepared by DQI is stable according to Definition 3 when the associated decoder corrects $\ell \leq \ell^*$ errors. In the case of the second assumption, our argument follows directly from the Lipschitzness of the utilized decoder and bypasses the choice of ℓ^* above. In the case of the first assumption, our argument for the stability of DQI follows from the observation that Gallager codes are what we call *locally restrictable*: these are ensembles of codes which are still good codes when ignoring all but an ϵ -fraction of syndromes. As stability depends only on the state prepared by the quantum algorithm, this property is independent of the decoder used for DQI, or indeed whatever methods are used to prepare the DQI state.

Theorem 6 (Stability of DQI). *Let $\mathbf{B}^\top$ be a parity check matrix drawn from an ϵ -locally restrictable code family such as the Gallager ensemble. Assume the DQI decoder satisfies Assumption 5. Then, DQI is stable for this ensemble of MAX-XOR-SAT instances.*

Next, we compute statistical properties of MAX- k -XOR-SAT at fixed k and clause density $\lambda = m/n$, including the maximum achievable satisfied fraction as well as the overlap gap threshold. These are the first times these two quantities have been computed for MAX- k -XOR-SAT when k is held fixed. Combining our results, we are able to show that DQI is topologically obstructed.

Theorem 7 (DQI is topologically obstructed for MAX- k -XOR-SAT, informal). *Assume the DQI decoder satisfies Assumption 5. DQI does not succeed in sampling a bit string $\mathbf{z} \in \mathbb{F}_2^n$ achieving a satisfied fraction:*

$$\frac{g(\mathbf{z})}{m} > \mu_{\text{DQI}} = \frac{1}{2} + \frac{1}{2\sqrt{\lambda}} \max \left((1 + o_k(1)) \sqrt{\frac{2 \ln(k)}{\ln(2)k}}, P_k^{\text{ALG}} \right) \quad (3)$$

with any constant probability over both the randomness of the sampling and the distribution of problem instances. Here, P_k^{ALG} is the k th algorithmic Parisi constant [30], which we numerically fit:

$$P_k^{\text{ALG}} \approx \sqrt{\frac{3.530 \ln(k)}{k}}. \quad (4)$$

This topological obstruction limits the best performance DQI can achieve *with substantial probability* and is thus a stronger upper bound than performance *in expectation* given by the semicircle law. We can compare the above bound with that of the semicircle law under the threshold from Eq. (2), in the large k limit:

$$\frac{\langle g \rangle_{\text{DQI}}}{m} = \left(\sqrt{\frac{\ell}{2m}} + \sqrt{\frac{1}{2} \left(1 - \frac{\ell}{m} \right)} \right)^2 \leq \frac{1}{2} + (1 + o_k(1)) \sqrt{\frac{1}{k\lambda}}. \quad (5)$$

Since the semicircle bound is on average rather than with non-negligible probability, it is smaller than μ_{DQI} . The asymptotic gap (in k), however, suggests that DQI may not perform optimally over unstructured ensembles even among stable algorithms. Indeed, we show that a depth-1 QAOA [12] achieves the same k scaling as μ_{DQI} for MAX- k -XOR-SAT, beyond what DQI can achieve in expectation:

$$\frac{\langle g \rangle_{\text{QAOA}}}{m} = \frac{1}{2} + (1 + o_k(1)) \sqrt{\frac{\ln(k)}{4ek\lambda}}. \quad (6)$$

Finally, we also consider what is optimally achieved by classical algorithms. The classical algorithm known as approximate message passing (AMP)—which is widely believed to saturate the OGP bound—achieves a value:

$$\mu_{\text{AMP}} = \frac{1}{2} + \frac{1}{2\sqrt{\lambda}} P_k^{\text{ALG}}. \quad (7)$$

for random MAX- k -XOR-SAT instances [11, 30, 31]. Our Theorem 7 then suggests that $\mu_{\text{DQI}} \leq \mu_{\text{AMP}}$, and in particular DQI is unable to outperform AMP on typical MAX- k -XOR-SAT instances under Assumption 5.

Outlook. On the OGP side, key open questions include rigorous proofs for stronger versions of the OGP known as branching OGPs, as well as whether OGPs hold for more general irregular LDPC ensembles. On the stability side, a compelling open direction is whether coherent *quantum* decoding [32] can far exceed the stability thresholds here, and whether DQI can be modified to utilize quantum decoding and thereby surmount the OGP barrier.

Updates from arXiv version. Our original arXiv preprint used an assumption that the efficient decoding threshold $1/k\lambda$, which was based on various estimates on state-of-the-art classical decoders across the literature. In this submission, we show this more rigorously: for state-of-the-art exact decoders based on linear programming, we show that $1/k\lambda$ upper bounds the threshold. For state-of-the-art approximate decoders such as belief propagation, we prove an upper bound of $\ln(k)/k\lambda$ under the standard density evolution analysis model, but use a different technique based on Lipschitzness to complete the proof of stability. A v2 with these results will appear on arXiv soon.

-
- [1] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM review* **41**, 303 (1999).
 - [2] C. Gyurik, A. Schmidhuber, R. King, V. Dunjko, and R. Hayakawa, Quantum computing and persistence in topological data analysis, arXiv preprint arXiv:2410.21258 (2024).
 - [3] D. W. Berry, Y. Su, C. Gyurik, R. King, J. Basso, A. D. T. Barba, A. Rajput, N. Wiebe, V. Dunjko, and R. Babbush, Analyzing prospects for quantum advantage in topological data analysis, *PRX Quantum* **5**, 010319 (2024).
 - [4] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Physical review letters* **103**, 150502 (2009).
 - [5] E. R. Anschuetz, Quantum glassiness from efficient learning, *Commun. Math. Phys.* **407**, 30 (2026).
 - [6] D. Gamarnik, The overlap gap property: A topological barrier to optimizing over random structures, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2108492118 (2021).
 - [7] W.-K. Chen, D. Gamarnik, D. Panchenko, and M. Rahman, Suboptimality of local algorithms for a class of max-cut problems, *Ann. Probab.* **47**, 1587 (2019).
 - [8] D. Gamarnik, E. C. Kızıldağ, W. Perkins, and C. Xu, Algorithms and barriers in the symmetric binary perceptron model (2022), arXiv:2203.15667 [cs.CC].
 - [9] D. Gamarnik and E. C. Kızıldağ, Algorithmic obstructions in the random number partitioning problem, *Ann. Appl. Probab.* **33**, 5497 (2023).
 - [10] B. Huang and M. Sellke, Tight Lipschitz hardness for optimizing mean field spin glasses, *Commun. Pure Appl. Math.* **78**, 60 (2025).
 - [11] H. E. Cheairi and D. Gamarnik, Algorithmic universality, low-degree polynomials, and Max-Cut in sparse random graphs (2024), arXiv:2412.18014 [cs.DS].
 - [12] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm (2014), arXiv:1411.4028 [quant-ph].
 - [13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, Variational quantum algorithms, *Nature Reviews Physics* **3**, 625 (2021).
 - [14] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case (2020), arXiv:2004.09002 [quant-ph].
 - [15] E. R. Anschuetz, Critical points in quantum generative models, in *International Conference on Learning Representations*, edited by K. Hofmann, A. Rush, Y. Liu, C. Finn, Y. Choi, and M. Deisenroth (OpenReview, 2022).
 - [16] E. R. Anschuetz and B. T. Kiani, Quantum variational algorithms are swamped with traps, *Nat. Commun.* **13**, 7760 (2022).
 - [17] J. Basso, D. Gamarnik, S. Mei, and L. Zhou, Performance and limitations of the QAOA at constant levels on large sparse hypergraphs and spin glass models, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (2022) pp. 335–343.
 - [18] A. Anshu and T. Metger, Concentration bounds for quantum states and limitations on the QAOA from polynomial approximations, *Quantum* **7**, 999 (2023).
 - [19] A. Chen, N. Huang, and K. Marwaha, Local algorithms and the failure of log-depth quantum advantage on sparse random CSPs (2023), arXiv:2310.01563 [quant-ph].
 - [20] M. Goh, The overlap gap property limits limit swapping in the qaoa (2025), arXiv:2404.06087 [quant-ph].
 - [21] E. R. Anschuetz, A unified theory of quantum neural network loss landscapes, in *International Conference on Learning Representations*, edited by Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (OpenReview, 2025).
 - [22] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar, and R. Babbush, Optimization by Decoded Quantum Interferometry (2025), arXiv:2408.08292 [quant-ph].
 - [23] I. S. Reed and G. Solomon, Polynomial codes over certain finite fields, *Journal of the society for industrial and applied mathematics* **8**, 300 (1960).
 - [24] G. De Palma, M. Marvian, D. Trevisan, and S. Lloyd, The quantum Wasserstein distance of order 1, *IEEE Trans. Inf. Theory* **67**, 6627 (2021).
 - [25] V. V. Zyablov and M. S. Pinsker, Estimation of the error-correction complexity for Gallager low-density codes, *Probl. Peredachi Inf.* **11**, 23 (1975).
 - [26] T. Richardson and R. Urbanke, *Modern coding theory* (Cambridge university press, 2008).
 - [27] J. Feldman, M. Wainwright, and D. R. Karger, Using linear programming to decode linear codes, in *37th annual Conference on Information Sciences and Systems (CISS'03)* (2003).
 - [28] R. Koetter, W. Li, P. Vontobel, and J. Walker, Characterizations of pseudo-codewords of ldpc codes, arxiv report, arXiv preprint cs.IT/0508049 (2005).
 - [29] B. Ghazi and E. Lee, Lp/sdp hierarchy lower bounds for decoding random ldpc codes, *IEEE Transactions on Information Theory* **64**, 4423 (2017).
 - [30] A. El Alaoui, A. Montanari, and M. Sellke, Optimization of mean-field spin glasses, *Ann. Probab.* **49**, 2922 (2021).
 - [31] K. Marwaha and S. Hadfield, Bounds on approximating Max k XOR with quantum and classical local algorithms, *Quantum* **6**, 757 (2022).
 - [32] A. Chailloux and J.-P. Tillich, The quantum decoding problem, arXiv preprint arXiv:2310.20651 (2023).