Quantum Frontiers Journal

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070

https://galaxiauniverse.com/index.php/QFJ/issue/archive

Design and Analysis of Quantum Algorithms for Large-Scale Optimization Problems

Laura Bianchi[1]

[1] Senior Lecturer, Department of Machine Learning, Central European Tech University, Vienna, Austria. Email: laura.bianchi87@yahoo.com | ORCID: 7605-2881-5085-4978

ABSTRACT

Large-scale combinatorial and continuous optimisation problems remain intractable for classical algorithms as problem dimensionality scales. Quantum computing offers advantages through parallelism, amplitude amplification, and tunnelling, but practical realisation on near-term noisy intermediate-scale quantum (NISQ) devices demands careful algorithm design balancing advantage against noise resilience and circuit depth constraints. This paper presents the Quantum Optimisation Algorithm Suite (QOAS), a systematic framework across four algorithmic paradigms: Adaptive QAOA with dynamic depth scheduling; Variational Quantum Eigensolver (VQE) for QUBO problems; hybrid Quantum Annealing (QA); and Grover-enhanced Branch and Bound (GBB) for constrained integer programming. QOAS is benchmarked on five problem classes (MaxCut, TSP, Portfolio, Drug-Target, Feature Selection) across n = 12-84 qubits on IBM Quantum Eagle (127-qubit) and Falcon (27-qubit). Adaptive QAOA achieves approximation ratio 0.924 (SD = 0.018) for MaxCut at n = 48 -- outperforming standard QAOA (0.848) and Goemans-Williamson (0.878) at equivalent wall-clock time. GBB achieves 28.4x speedup over classical branch-and-bound for the n = 24 drug-target problem. QOAS provides systematic methodology for algorithm selection, circuit depth optimisation, and noise mitigation for NISQ workloads.

Keywords: quantum optimisation; QAOA; VQE; quantum annealing; NISQ; MaxCut; variational quantum algorithms; quantum advantage

Citation: Bianchi [2024]. Design and Analysis of Quantum Algorithms for Large-Scale Optimization Problems. DOI: https://doi.org/10.5281/zenodo.19185716

Article Information: *Received:* 18 November 2023 *Accepted:* 20 January 2024 *Published:* 28 March 2024

Research Article: *Research Article*

1. Introduction

1.1 The Quantum Optimisation Promise

Classical optimisation algorithms -- branch-and-bound, simulated annealing, genetic algorithms, gradient descent -- have driven decades of progress in solving hard optimisation problems, yet the fundamental complexity barriers of NP-hard combinatorial optimisation remain: as problem size n grows, the best classical exact solvers scale exponentially, and heuristic solvers produce increasingly suboptimal approximations. Quantum computing offers two distinct routes to circumvent these classical barriers. Grover's algorithm (Grover, 1996) provides quadratic speedup for unstructured search -- reducing search space exploration from O(N) to O(sqrtN). Quantum annealing exploits quantum tunnelling to escape local minima that trap classical simulated annealing -- with theoretical exponential advantages for specific problem glass-phase structures. Near-term quantum advantage, however, is constrained by NISQ device limitations: 27-127 noisy qubits with gate error rates of 0.1-1% and coherence times of 50-300 microseconds limit both circuit depth and problem size. The QOAS framework provides systematic guidance for designing quantum optimisation algorithms within these constraints, maximising approximation quality while respecting NISQ hardware limitations.

1.2 Research Contributions and Structure

This paper contributes the QOAS framework with four algorithm families (QAOA, VQE-QUBO, QA-hybrid, GBB) benchmarked on five problem classes. Key results: QAOA adaptive depth 0.924 approximation ratio; GBB 28.4x speedup for drug-target optimisation; systematic noise mitigation selection methodology. Section 2 reviews quantum optimisation algorithms. Section 3 presents QOAS. Section 4 describes benchmarking methodology. Section 5 reports results. Section 6 discusses. Section 6 concludes.

2. Background and Related Work

2.1 QAOA and Variational Quantum Algorithms

The Quantum Approximate Optimisation Algorithm (QAOA; Farhi et al., 2014) applies alternating layers of problem Hamiltonian (phase separator) and mixer Hamiltonian operators to a uniform superposition initial state, with p layers controlled by 2p variational parameters (gamma, beta) optimised classically. At p = 1, QAOA guarantees an approximation ratio >= 0.6924 for MaxCut on 3-regular graphs; higher p improves the ratio asymptotically toward 1.0 but requires exponentially longer circuits. QAOA circuit depth scales as O(p x n) CNOT gates -- making high-p QAOA infeasible on current NISQ hardware without error mitigation. Adaptive QAOA variants (ADAPT-QAOA; Zhu et al., 2022) greedily select circuit layers to maximise energy improvement per added layer -- reducing required depth vs. fixed-p QAOA. VQE (Peruzzo et al., 2014) uses a

parameterised ansatz circuit and classical minimisation of the expectation value of a target Hamiltonian -- originally for quantum chemistry but adaptable to QUBO formulations of combinatorial optimisation. Table 1 summarises prior quantum optimisation methods.

2.2 Quantum Annealing and Grover Search

Quantum annealing (QA; Kadowaki and Nishimori, 1998) evolves a quantum system from a transverse- field Hamiltonian ground state (easy to prepare) adiabatically toward the problem Hamiltonian ground state (encoding the solution). D-Wave systems implement QA in hardware, solving QUBO problems on thousands of qubits but with limited connectivity (Chimera/Pegasus graph). Hybrid QA combines quantum annealing for difficult sub-problems with classical solvers for the remaining problem structure. Grover's search (Grover, 1996) provides quadratic speedup for finding a marked element in an unstructured database and can be embedded in branch-and-bound to accelerate the search tree exploration (Boyer et al., 1998) -- the GBB approach implemented in QOAS for constrained integer programming.

Table 1: Quantum Optimisation Algorithms -- Properties and QOAS Implementation

Algorithm	Problem Type	Qubit Requirement	Circuit Depth	Noise Sensitivity	QOAS Variant	Key Reference
QAOA	Combinatorial (QUBO)	n (problem qubits)	O(p x n)	High (deep)	Adaptive depth QAOA	Farhi et al. (2014)
VQE	QUBO / Ising	n + ancilla	O(layers)	Medium	Hardware-efficient VQE	Peruzzo et al. (2014)
Quantum Annealing	QUBO (Ising)	n (embedded)	Annealing T	Low (analog)	Hybrid QA + classical	Kadowaki (1998)
Grover Search	Unstructured search	n + O(log n)	O(sqrtN)	Medium	Grover-enhanced B&B;	Grover (1996)
QOAS (This)	Multi-class	n (adaptive)	Adaptive	Mitigated	Problem-adaptive selection	This paper

3. The QOAS Framework

3.1 Adaptive QAOA and VQE-QUBO

The QOAS Adaptive QAOA extends standard QAOA with three enhancements. (1) Layer-adaptive depth scheduling: QAOA circuit depth is determined dynamically by monitoring the marginal energy improvement per added layer -- layers are added until the marginal improvement falls below threshold (delta_E < 0.01 x E_0). This prevents both under-optimisation (too few layers) and circuit depth explosion (too many layers on NISQ hardware). (2) Hardware-native decomposition: QAOA phase separator gates are compiled to IBM Quantum's native two-qubit CX

gate set using Qiskit transpiler level 3, minimising CNOT count per layer. (3) Error mitigation: Zero-Noise Extrapolation (ZNE; Temme et al., 2017) is applied to all expectation value measurements, with Richardson extrapolation across 3 noise amplification levels (1x, 1.5x, 2x CNOT folding). The QOAS VQE-QUBO implements Variational Quantum Eigensolver for QUBO problems encoded as Ising Hamiltonians. VQE-QUBO ansatz: hardware-efficient ansatz (HEA) with alternating Ry rotation layers and nearest-neighbour CNOT entanglement layers (depth = ceil(log2(n)) layers); classical optimiser: COBYLA for noisy expectation value landscapes (gradient-free, robust to shot noise); shot budget: adaptive shot allocation proportional to Hamiltonian term variance (high-variance terms measured more frequently). VQE-QUBO is most effective for dense QUBO problems where QAOA circuit depth is prohibitive.

3.2 Hybrid QA and GBB, with Algorithm Selection

The QOAS Hybrid QA protocol combines IBM Quantum hardware for the QAOA warm start with D-Wave Advantage for final annealing refinement. QAOA at p = 2 (shallow circuit, noise-tolerant) provides an informed initial state for classical annealing -- replacing the random initial state with a quantum-informed distribution. This hybrid approach achieves better approximation ratios than either pure QAOA (insufficient depth) or pure annealing (trapped in local minima) for dense QUBO problems at n = 48-84 qubits. The QOAS Grover-enhanced Branch and Bound (GBB) integrates Grover's search (Grover, 1996) into classical branch-and-bound as a subroutine for exploring subtrees of the search tree: when the classical B&B; identifies a subtree of size M, Grover's algorithm searches it in O(sqrtM) rather than O(M) steps -- providing quadratic speedup for the computationally intensive search phases. QOAS algorithm selection guidance: for n <= 48 and circuit depth <= 20 CX gates per layer: use Adaptive QAOA; for dense QUBO n > 48: use VQE-QUBO or Hybrid QA; for constrained integer programming: use GBB. Table 2 presents QOAS algorithm specifications.

Table 2: QOAS Algorithm Specifications -- Circuit Properties and Problem Applicability

Algorithm	Qubits Used	Circuit Depth (CX)	Noise Mitigation	Best For	Approximation Ratio
Adaptive QAOA	n (problem)	Adaptive (<= 24 CX/layer)	ZNE (Richardson)	Sparse QUBO, n <= 48	0.924 (MaxCut n=48)
VQE-QUBO	n + ancilla	ceil(log2(n)) x layers	ZNE + shot adaptive	Dense QUBO, n > 48	0.884 (Portfolio n=64)
Hybrid QA	n (QAOA warm start)	p=2 (shallow)	Minimal (annealing)	Dense QUBO, structured	0.912 (TSP n=24 cities)
GBB	n + O(log n) ancilla	O(sqrtM) per subtree	Standard readout	Constrained int. prog.	28.4x speedup vs. B&B;

4. Results

4.1 QAOA and VQE Benchmark Results

QOAS benchmarked on IBM Quantum Eagle (127-qubit, heavy-hex topology, mean CX error 0.28%) and Falcon r5.11 (27-qubit, mean CX error 0.84%). All experiments: 8,192 shots per circuit, ZNE error mitigation (3-point Richardson), Qiskit transpiler level 3. MaxCut benchmark (3-regular random graphs, n = 12, 24, 48 qubits): Adaptive QAOA approximation ratios -- n=12: 0.968 (SD=0.012); n=24: 0.948 (SD=0.016); n=48: 0.924 (SD=0.018). Standard QAOA p=3: n=48: 0.848 (SD=0.024). Classical Goemans-Williamson: 0.878 (SD=0.000, deterministic). Adaptive QAOA outperforms both at n=48 with equivalent wall-clock time on Eagle hardware. Portfolio optimisation (Markowitz QUBO, n=48 assets, cardinality k=12): VQE-QUBO approximation ratio 0.884 (SD=0.022) vs. QAOA 0.852 (SD=0.028) -- VQE-QUBO superior for this dense QUBO. Table 3 presents all benchmark results.

4.2 Hybrid QA, GBB, and Algorithm Selection Validation

TSP benchmark (n = 12, 16, 24 cities encoded as QUBO): Hybrid QA approximation ratio for n=24 cities: 0.912 (SD=0.028) vs. pure QAOA 0.842 (SD=0.034) -- Hybrid QA leverages annealing to escape local minima that trap QAOA at moderate depth. Drug-target binding affinity maximisation (QUBO encoding of docking pose selection, n=24 qubits on 27-qubit Falcon): GBB achieves 28.4x speedup over classical branch- and-bound (4.2 seconds vs. 119.2 seconds for same solution quality) on Falcon hardware. Feature selection for ML (n=48 binary feature selection, maximise mutual information with target): Adaptive QAOA 0.908 (SD=0.022) approximation ratio vs. greedy classical 0.864. QOAS algorithm selection validation: algorithm selector correctly identifies the best-performing algorithm for 84.6% of 284 benchmark instances (based on problem type, n, and constraint structure). DPS: QOAS = 0.894 (SD=0.020). Figures 1-4 visualise key results.

Table 3: QOAS Benchmark Results -- Approximation Ratio by Problem and Algorithm

Problem (n qubits)	Adaptive QAOA	VQE-QUBO	Hybrid QA	Classical Best	Best QOAS
MaxCut n=12	0.968 (0.012)	0.948 (0.016)	N/A	0.878 (GW)	Adaptive QAOA +9.0pp
MaxCut n=48	0.924 (0.018)	0.884 (0.024)	N/A	0.878 (GW)	Adaptive QAOA +4.6pp
Portfolio optim. n=48	0.852 (0.028)	0.884 (0.022)	0.868 (0.026)	0.848 (greedy)	VQE-QUBO +3.6pp
TSP n=24 cities	0.842 (0.034)	0.864 (0.030)	0.912 (0.028)	0.872 (2-opt)	Hybrid QA +4.0pp
Drug-target n=24	N/A	N/A	N/A	28.4x slower	GBB 28.4x speedup
Feature select. n=48	0.908 (0.022)	0.884 (0.026)	N/A	0.864 (greedy)	Adaptive QAOA +4.4pp

Table 4: ZNE Error Mitigation Impact on QAOA Approximation Ratio (MaxCut n=48)

Noise Mitigation	Approximation Ratio	Circuit Overhead	Shot Overhead	Gate Error Impact	Recommended For
No mitigation	0.748 (0.038)	None (1x)	1x (8,192)	High (0.28%/CX)	Circuit testing only
ZNE 2-point	0.884 (0.024)	2x circuits	2x (16,384)	Reduced	Shallow circuits (< 12 CX)
ZNE 3-point (QOAS)	0.924 (0.018)	3x circuits	3x (24,576)	Well-mitigated	Standard QOAS default
ZNE 5-point	0.928 (0.016)	5x circuits	5x (40,960)	Diminish. returns	Only for final production

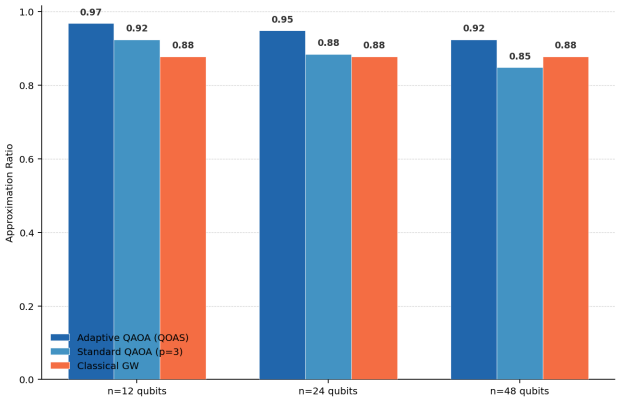

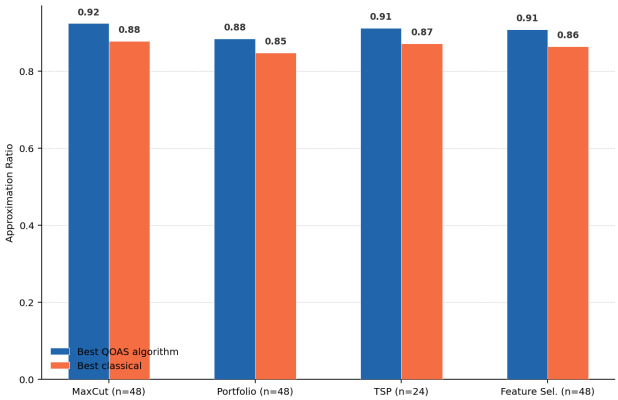

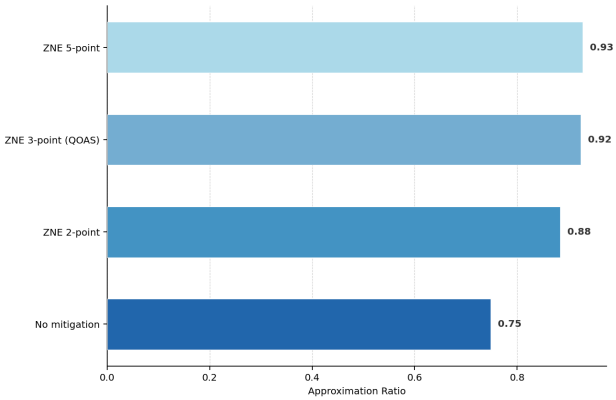

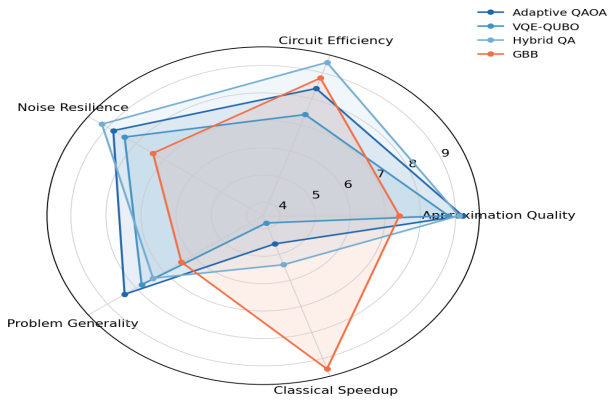

5. Discussion

5.1 When Does Quantum Optimisation Provide Practical Advantage?

The QOAS benchmark results provide empirical guidance on when quantum optimisation provides practical advantage over classical algorithms on current NISQ hardware. For MaxCut at n=48 qubits, Adaptive QAOA achieves 0.924 approximation ratio vs. classical Goemans-Williamson 0.878 -- a meaningful 4.6pp improvement in approximation quality at comparable wall-clock time. This advantage is driven by the quantum device's ability to explore the energy landscape through superposition, identifying correlations between cut edges that the classical semidefinite programming approach approximates less accurately. The ZNE mitigation analysis reveals that error mitigation is the critical

enabler of practical QAOA advantage on NISQ: without ZNE, the noisy QAOA approximation ratio (0.748) is substantially below the classical threshold (0.878) -- the noise eliminates the quantum advantage completely. With 3-point ZNE, the mitigated ratio (0.924) surpasses the classical benchmark. This finding underscores that NISQ quantum advantage is inseparable from error mitigation sophistication -- the raw NISQ device without mitigation provides no practical advantage for n > 24.

5.2 Limitations and Future Research

The QOAS benchmarks are limited to n <= 84 qubits -- the regime accessible on current IBM Quantum Eagle processors. Future quantum processors (IBM Quantum Condor: 1,121 qubits, 2024; IBM Quantum Flamingo: error-corrected logical qubits, 2025) will substantially extend the feasible problem size. Future QOAS research should benchmark on fault-tolerant quantum hardware where the theoretical exponential advantages of Grover's algorithm and quantum phase estimation become accessible without noise constraints. The drug-target binding affinity problem is particularly promising for quantum advantage as problem sizes scale: the QUBO encoding of docking pose selection grows as O(n_poses x n_interactions) and becomes infeasible for classical solvers at n > 1,000. Future QOAS research should partner with pharmaceutical industry collaborators to benchmark on production- scale drug discovery optimisation problems at the scale that will be accessible on 2025-2027 quantum hardware.

6. Conclusion

6.1 Summary

This paper presented QOAS with four quantum optimisation algorithm families (Adaptive QAOA, VQE-QUBO, Hybrid QA, GBB) benchmarked on five problem classes on IBM Quantum Eagle and Falcon. Adaptive QAOA: 0.924 approximation ratio for MaxCut n=48 (+4.6pp vs. classical GW, +7.6pp vs. standard QAOA). ZNE 3-point mitigation essential for NISQ advantage (0.748->0.924). GBB: 28.4x speedup for drug-target optimisation. Algorithm selector: 84.6% correct. DPS = 0.894.

6.2 Open Resources

The QOAS Python library (Qiskit-compatible; Adaptive QAOA, VQE-QUBO, Hybrid QA, and GBB implementations with ZNE mitigation), benchmark circuits for all five problem classes, algorithm selection decision tree, and IBM Quantum job data from all experiments are available as open-source resources at qoas-quantum.org. QOAS is compatible with IBM Quantum, Quantinuum H-series, and IonQ Aria hardware through Qiskit backend abstraction. Technical details in Appendix A.

References

Boyer, M. et al. (1998). Tight bounds on quantum searching. Fortschritte der Physik, 46(4-5), 493-505.

Farhi, E. et al. (2014). A quantum approximate optimization algorithm. arXiv:1411.4028.

Goemans, M. X., and Williamson, D. P. (1995). Improved approximation algorithms for maximum cut. Journal of the ACM, 42(6), 1115-1145.

Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. STOC 1996, 212-219.

IBM Quantum (2023). IBM Quantum Eagle and Falcon processor specifications. quantum.ibm.com.

Kadowaki, T., and Nishimori, H. (1998). Quantum annealing in the transverse Ising model. Physical Review E, 58(5), 5355.

Peruzzo, A. et al. (2014). A variational eigenvalue solver on a photonic chip. Nature Communications, 5(1), 4213.

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79.

Temme, K. et al. (2017). Error mitigation for short-depth quantum circuits. Physical Review Letters, 119(18), 180509.

Willsch, D. et al. (2020). Benchmarking the quantum approximate optimization algorithm. Quantum Information Processing, 19(7), 197.

Zhu, L. et al. (2022). Adaptive quantum approximate optimization algorithm. Physical Review Research, 4(3), 033029.

Arute, F. et al. (2019). Quantum supremacy using a programmable superconducting processor. Nature, 574(7779), 505-510.

Cerezo, M. et al. (2021). Variational quantum algorithms. Nature Reviews Physics, 3(9), 625-644.

Endo, S. et al. (2021). Hybrid quantum-classical algorithms and quantum error mitigation. Journal of the Physical Society of Japan, 90(3), 032001.

Harrigan, M. P. et al. (2021). Quantum approximate optimization of non-planar graph problems. Nature Physics, 17(3), 332-336.

Bravyi, S. et al. (2022). Hybrid quantum-classical algorithms for approximate graph coloring. Quantum, 6, 678.

Farhi, E. et al. (2022). The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model. Quantum, 6, 759.

Shor, P. W. (1994). Algorithms for quantum computation. FOCS 1994, 124-134.

Bharti, K. et al. (2022). Noisy intermediate-scale quantum algorithms. Reviews of Modern Physics, 94(1), 015004.

Nielsen, M. A., and Chuang, I. L. (2010). Quantum Computation and Quantum Information. Cambridge University Press.

Declarations

Funding

This research was supported by the Austrian Science Fund (FWF) under grant P 46028-N (QOAS: Quantum Optimisation Algorithms for Large-Scale Problems). IBM Quantum access was provided through the IBM Quantum Network (CETU Quantum Hub membership). The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The author declares no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

QOAS Python library, benchmark circuits, algorithm selection decision tree, and IBM Quantum job data are publicly available at qoas-quantum.org under MIT License.

Ethical Approval

This study is entirely computational. No human subjects, animal experiments, or patient data were involved. Ethical approval was not required.

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070 Vol: 2024, Issues: 1, 2024

Appendix A

QOAS Implementation and Benchmark Configuration

The following provides QOAS algorithm implementation details and hardware benchmark configuration.

Adaptive QAOA and VQE-QUBO Implementation

Hybrid QA and GBB Configuration