

Capacity-Constrained Multi-Robot Task Allocation via Quantum Approximate Optimisation: A Reproducible Benchmark on a Solomon VRPTW C101-Derived Instance

Sandeep Gupta^{1,2}[0000–0002–2046–115X], Sazzad Hossain³[0009–0003–1351–7287],
Akmal Akhatov⁴[0000–0003–3834–854X], Nurmamatov
Mekhriddin⁵[0000–0002–9619–5791], and Kanad Ray^{6,7,*}[0000–0001–9324–2778]

- ¹ Department of Maritime Engineering, Indian Institute of Technology Gandhinagar, Gandhinagar 382355, Gujarat, India
- ² Department of Electronics and Communication Engineering, Poornima College of Engineering (Affiliated to Rajasthan Technical University, Kota), Jaipur 302022, Rajasthan, India
sandeep.gupta@iitgn.ac.in
- ³ Department of System Management and Information Security, Samarkand State University, Samarkand 140104, Uzbekistan
sazzad69@gmail.com
- ⁴ Faculty of Artificial Intelligence and Digital Technologies, Samarkand State University, Samarkand 140104, Uzbekistan
akmalar74@gmail.com
- ⁵ Department of System Management and Information Security, Samarkand State University, Samarkand 140104, Uzbekistan
mehridinnur@gmail.com
- ⁶ Amity Cognitive Computing and Brain Informatics Centre, Amity University Rajasthan, Jaipur 303002, Rajasthan, India
kray@jpr.amity.edu
- ⁷ Faculty of Artificial Intelligence and Digital Technologies, Samarkand State University, Samarkand 140104, Uzbekistan

Abstract. Assigning tasks to robots under per-robot capacity limits is a combinatorially hard coordination problem central to logistics and warehouse automation. Quantum variational solvers have recently attracted attention for such problems, yet rigorous benchmarks with full reproducibility remain scarce. We address this gap with three contributions. First, a slack-free Quadratic Unconstrained Binary Optimisation (QUBO) formulation for capacity-constrained MRTA that encodes coverage and capacity simultaneously as equality constraints, requiring only NK binary variables with no ancilla overhead. Second, a closed-form sufficient condition for the penalty coefficient ($\lambda > c_{\text{opt}}/N$), with $\lambda = c_{\text{greedy}}$ as a parameter-free, $O(NK \log N)$ computable choice. Third, the first fully reproducible Quantum Approximate Optimisation Algorithm (QAOA)-based MRTA benchmark derived from Solomon VRPTW C101 spatial coordinates (customer locations and depot only; time windows, route sequencing, and service durations are not modelled—the problem

is capacity-constrained task assignment). Across depths $p \in \{1, 2, 3\}$ on a 16-qubit instance, QAOA achieves 100% constraint satisfaction; $p = 2$ attains approximation ratio $r = 1.0047$ (mean 1.0051 ± 0.0008 over ten seeds), outperforming both greedy ($r = 1.016$) and random ($r = 1.007$) baselines. All artefacts are deposited on Zenodo [1].

Keywords: Quantum Approximate Optimisation Algorithm · Multi-Robot Task Allocation · QUBO · Ising Hamiltonian · Capacity Constraints · Near-Term Quantum Computing · Dynamic Penalty Scaling · Reproducibility

1 Introduction

Deploying a team of robots to service a set of spatially distributed tasks requires deciding which robot handles which task, subject to physical limits such as payload capacity and travel range. When each robot may serve at most a fixed number of tasks, the resulting optimisation is equivalent to a Generalised Assignment Problem (GAP), a class that is NP-hard in the general case [4]. As instance sizes grow, classical methods ranging from integer linear programming to evolutionary algorithms either forgo optimality guarantees or scale poorly.

Quantum computing offers a qualitatively different approach. In the current Noisy Intermediate-Scale Quantum (NISQ) era [6,7], devices of 20–100 qubits are available but noise-limited, making variational hybrid algorithms the most practical route. The QAOA [8] encodes a combinatorial cost as a quantum Hamiltonian, prepares a parameterised ansatz through alternating cost and mixing unitaries, and recovers near-optimal solutions by repeated measurement. Applications of QAOA to MaxCut [8], scheduling [9], traffic routing [10], and capacitated vehicle routing [11] have been reported, yet *capacity-constrained multi-robot task assignment* with equality constraints, an asymmetric inter-robot cost structure, and a principled penalty-tuning procedure on a standard benchmark has not previously been demonstrated.

Scope clarification. Our benchmark derives robot start positions, task locations, and the depot from the Solomon VRPTW C101 dataset [25] as a widely recognised source of reproducible spatial data. However, the problem we solve is *not* the Vehicle Routing Problem with Time Windows: time window constraints, customer service durations, and vehicle route sequencing are entirely absent. The resulting formulation is a *capacity-constrained task assignment* problem (ST-SRIA with capacity, in the taxonomy of Gerkey and Mataric [2]).

Figure 1 provides an end-to-end visual overview.

Contributions.

- (i) **Compact equality-constraint QUBO** (Sec. 3.2): NK binary variables, no slack qubits, exact coverage and capacity enforcement.




Fig. 1. Graphical abstract. (Left) Eight customer tasks (blue circles), depot (black star), Robot-0 at SW (32,42) and Robot-1 at NE (48,58)—coordinates taken from Solomon C101; no time windows are modelled. (Centre) QAOA circuit with 16 qubits, $p \in \{1, 2, 3\}$ cost/mixing layers. (Right) Near-optimal assignment: Robot-0 serves 3 tasks (red), Robot-1 serves 5 tasks (green); $r = 1.0047$, 8/8 tasks, 0 violations.

- (ii) **Sufficient penalty condition with proof** (Sec. 3.3): revised proof considers all infeasible assignments; $\lambda = c_{\text{greedy}}$ is proposed as a parameter-free choice.
- (iii) **Asymmetric robot-cost instance** (Sec. 3.4): non-degenerate geometry ensures all baselines produce meaningfully different results.
- (iv) **Fully reproducible quantum benchmark** (Sec. 6): complete open-access package on Zenodo [1], including multi-seed results and polynomial-time classical baselines for context.

2 Background and Related Work

2.1 Multi-Robot Task Allocation

The formal taxonomy of MRTA proposed by Gerkey and Mataric [2] classifies problems along three axes: single-task vs. multi-task robots; single-robot vs. multi-robot collaboration; and instantaneous vs. time-extended assignment. Our setting is ST-SR-IA with a per-robot capacity budget, yielding a GAP [4]. Korsah et al. [3] extend this taxonomy to cross-schedule interdependencies.

Complexity note. In our specific formulation, all tasks carry unit demand and robot capacities are fixed non-negative integers. Under these conditions the problem reduces to capacitated bipartite b -matching, which admits polynomial-time solutions via min-cost flow [5] in $O(N^2 K \log N)$. The NP-hardness of the general

GAP arises when task demands are heterogeneous and robot capacities interact non-trivially; our experiments use the unit-demand instance as a clean, reproducible testbed for evaluating the QAOA formulation at small scale. Physical multi-robot systems motivating this study include Bluetooth-enabled formation control platforms reported in [12].

2.2 QUBO and Ising Models

Lucas [13] showed that a wide class of NP problems can be expressed as Ising minimisations:

$$H_{\text{Ising}} = - \sum_{i < j} J_{ij} s_i s_j - \sum_i h_i s_i, \quad s_i \in \{-1, +1\} \quad (1)$$

via the substitution $x_i = (1 - s_i)/2$. The QUBO form $H_{\text{QUBO}} = \mathbf{x}^\top \mathbf{Q} \mathbf{x}$ maps directly to a quantum cost Hamiltonian:

$$\hat{H}_C = \sum_i h_i \hat{Z}_i + \sum_{i < j} J_{ij} \hat{Z}_i \hat{Z}_j \quad (2)$$

Glover et al. [14] provide a practitioner’s guide to QUBO construction, highlighting sensitivity to the penalty coefficient λ —a challenge our dynamic scheme addresses.

2.3 The Quantum Approximate Optimisation Algorithm

QAOA at depth p [8] prepares:

$$|\gamma, \beta\rangle = \prod_{l=1}^p e^{-i\beta_l \hat{B}} e^{-i\gamma_l \hat{H}_C} |+\rangle^{\otimes n} \quad (3)$$

where $\hat{B} = \sum_i \hat{X}_i$ is the mixing Hamiltonian and $|+\rangle^{\otimes n}$ is the uniform superposition. The expected cost is minimised classically over the $2p$ angles. Wurtz and Love [18] prove monotonic improvement with p under exact parameter optimisation; Bengtsson et al. [19] confirm this requires sufficient classical budget—consistent with our non-monotonic $p = 3$ result. The barren plateau phenomenon [16] does not affect our experiments at $p \leq 3$, $n = 16$.

2.4 Quantum Approaches to Routing and Logistics

Neukart et al. [10] applied quantum annealing to traffic optimisation. Feld et al. [11] used QAOA on the Capacitated VRP with $O(N^2 K)$ route-encoding variables, reporting approximation ratios of 1.1–1.3 on six-customer instances. Quetschlich et al. [21] addressed a VRPTW variant on four customers (12 qubits).

Distinction from prior work. Relative to [11]: (i) we solve task assignment, not route sequencing, reducing variable count from $O(N^2K)$ to $O(NK)$; (ii) equality constraints give exact satisfaction without soft penalty loss; (iii) we provide full reproducibility on a standard spatial benchmark. Relative to [21]: larger scale (8 vs. 4 customers), asymmetric costs, and a systematic penalty theorem.

Broader human-centred AI context. Beyond computational performance, intelligent and autonomous systems increasingly need to address transparency, privacy, accessibility, and human-centred operation. Mahmud and collaborators have explored explainable and privacy-preserving personalised AI [22], privacy-preserving multi-agent healthcare systems [23], and inclusive explainable digital-human modelling [24]. Although developed in different domains, these studies highlight design principles that may also inform future multi-robot systems operating in assistive, healthcare, and human-robot collaborative settings.

3 Problem Formulation

3.1 Problem Definition

Let $\mathcal{J} = \{0, \dots, N-1\}$ be a task set and $\mathcal{R} = \{0, \dots, K-1\}$ a set of robots with integer capacities $\text{cap}_k \geq 1$, satisfying the balance condition $\sum_k \text{cap}_k = N$. The travel-plus-return cost for robot k serving task j is:

$$c_{kj} = \|\mathbf{r}_k - \mathbf{t}_j\|_2 + \|\mathbf{t}_j - \mathbf{d}\|_2 \quad (4)$$

The optimisation is:

$$\begin{aligned} \min_{x_{kj}} \quad & \sum_{k,j} c_{kj} x_{kj} & (5) \\ \text{s.t.} \quad & \sum_k x_{kj} = 1, \forall j & (\text{Coverage}) \\ & \sum_j x_{kj} = \text{cap}_k, \forall k & (\text{Capacity}) \\ & x_{kj} \in \{0, 1\}, \forall k, j & (\text{Binary}) \end{aligned}$$

3.2 QUBO Encoding

Converting the equality constraints to quadratic penalties:

$$H_{\text{QUBO}}(\mathbf{x}) = \sum_{k,j} c_{kj} x_{kj} + \lambda \sum_j \left(\sum_k x_{kj} - 1 \right)^2 + \lambda \sum_k \left(\sum_j x_{kj} - \text{cap}_k \right)^2 \quad (6)$$




Fig. 2. QUBO formulation pipeline. *Left:* Bipartite assignment graph (Robot-0 cap= 3, Robot-1 cap= 5, 8 tasks). *Centre:* 16×16 QUBO matrix Q ; orange diagonal: objective costs; blue off-diagonal: penalty couplings. *Right:* 16-qubit Ising Hamiltonian arranged as a 4×4 grid.

This requires exactly NK binary variables—no slack variables because the balance condition ensures consistency of both equality families. The QUBO matrix entries are:

$$Q_{ii} = c_{kj} + \lambda(1 - 2 - 2 \text{cap}_k) \quad (7)$$

$$Q_{ij} = 2\lambda, \quad \text{if } i, j \text{ share a task or a robot} \quad (8)$$

giving Ising coefficients $h_i = -\frac{1}{2} \sum_j Q_{ij}$ and $J_{ij} = \frac{1}{4} Q_{ij}$ ($i \neq j$).

Figure 2 illustrates the complete pipeline from the bipartite assignment graph through the 16×16 QUBO matrix to the 16-qubit Ising Hamiltonian.

3.3 Penalty Sufficient Condition

Theorem 1 (Sufficient Penalty Condition). *If $\lambda > c_{\text{opt}}/N$, then every optimal solution of H_{QUBO} is feasible.*

Proof. Partition binary strings into $\mathcal{F}$ (feasible) and $\mathcal{I}$ (infeasible). For any $\mathbf{x} \in \mathcal{I}$, at least one constraint among the N coverage equations or the K capacity equations is violated. A single violated coverage constraint $(\sum_k x_{kj} - 1)^2 \geq 1$ adds penalty at least λ to $H_{QUBO}(\mathbf{x})$. Hence $H_{QUBO}(\mathbf{x}) \geq \lambda$ for all $\mathbf{x} \in \mathcal{I}$. For the assignment objective we have $c_{\text{feas}} \leq c_{\text{opt}} \cdot N/N = c_{\text{opt}}$ for any feasible solution. Setting $\lambda > c_{\text{opt}}/N$ ensures $\lambda N > c_{\text{opt}} \geq c_{\text{feas}}$, so $H_{QUBO}(\mathbf{x}) \geq \lambda > c_{\text{opt}}/N$ for all $\mathbf{x} \in \mathcal{I}$, while the feasible optimum satisfies $H_{QUBO}(\mathbf{x}^*) = c_{\text{opt}}$. Therefore every QUBO optimum is feasible. $\square$

Remark. Theorem 1 provides a *sufficient* condition; a tight lower bound on λ requires characterising the minimum penalty across all infeasible configurations and is left for future work.

We propose $\lambda^* = c_{\text{greedy}} \geq c_{\text{opt}}$, satisfying the condition without knowledge of the true optimum, computable in $O(NK \log N)$ before the quantum run.

3.4 Instance Construction

We extract the first $N = 8$ customer spatial coordinates from Solomon VRPTW C101 [25], using only (x, y) positions. **Time windows, service durations, vehicle return requirements, and any routing constraints from C101 are not incorporated.** The depot is placed at $\mathbf{d} = (40, 50)$; robots are positioned asymmetrically:

$$\mathbf{r}_0 = (32, 42) \text{ [SW]}, \quad \mathbf{r}_1 = (48, 58) \text{ [NE]} \quad (9)$$

Capacities $\text{cap} = [3, 5]$ prevent the trivially-degenerate all-to-Robot-1 assignment. Figure 3 shows the instance geometry and optimal assignment.

4 QAOA Implementation

4.1 Quantum Circuit Architecture

The QAOA circuit (Fig. 4) consists of four stages: (1) Hadamard layer preparing $|+\rangle^{\otimes 16}$; (2) cost layer applying $e^{-i\gamma_l \hat{H}_C}$ per depth l ; (3) mixing layer applying $e^{-i\beta_l \hat{B}}$; (4) measurement of all 16 qubits, repeated $S = 4096$ times.

Total two-qubit gate depth scales as $O(p \cdot |\mathcal{E}|)$, where $|\mathcal{E}| \approx 100$ off-diagonal entries in the 16×16 QUBO matrix give ~ 300 two-qubit gates at $p = 3$.

Simulation vs. real hardware. All experiments use the Qiskit Aer statevector simulator on a T4 GPU, which is noise-free. Reported runtimes (0.79–23.67 s) reflect classical simulation cost, which grows *exponentially* with qubit count; on physical NISQ hardware, circuit execution would grow only as $O(p \cdot |\mathcal{E}|)$ gate operations but approximation ratios would degrade due to decoherence and gate error. Experiments on real hardware are identified as future work (Sec. 7).

4.2 Parameter Optimisation

The variational parameters $\boldsymbol{\theta} = (\boldsymbol{\gamma}, \boldsymbol{\beta}) \in \mathbb{R}^{2p}$ are optimised using COBYLA [26], a gradient-free derivative-free method suited to noisy quantum objectives. Zhou et al. [27] confirm COBYLA finds near-optimal QAOA parameters for small instances in polynomial time. Initial values are $\gamma_l = \beta_l = 0.5$ for all l ; the budget is 300 function evaluations at $S = 4096$ shots. Multi-seed experiments (Sec. 5.4) use 10 independent random initialisations per depth.




Fig. 3. Physical instance. Depot (black star) at (40, 50); Robot-0 (red triangle, SW) at (32, 42); Robot-1 (green triangle, NE) at (48, 58); 8 customer tasks (blue circles C1–C8). Customer locations are taken from Solomon C101 coordinates only; no time windows are modelled. Optimal total cost $c_{\text{opt}} = 275.61$.

4.3 Algorithm

5 Experimental Results

5.1 Setup

Figure 5 shows the complete seven-stage pipeline. Table 1 summarises numerical parameters.

5.2 Penalty Ablation

Table 2 validates Theorem 1. With $\lambda \leq 50$, the all-zero QUBO value ($8 \times 50 = 400$) still exceeds many feasible costs (≈ 275), producing vacuous results. With $\lambda = 150$, the constraint is satisfied, confirming the sufficient condition is met slightly below c_{greedy} . Our dynamic choice $\lambda^* = 280.13 = c_{\text{greedy}}$ provides a conservative, parameter-free guarantee.




Fig. 4. QAOA circuit (8 qubits shown for clarity; full experiments use 16). Orange ZZ boxes: cost-layer entangling gates. Blue R_x boxes: mixing-layer single-qubit rotations. Right: COBYLA classical optimisation loop.

5.3 Main Algorithm Comparison

Table 3 summarises results (seed 42 run). All three exact classical methods confirm $c_{\text{opt}} = 275.61$. QAOA $p = 2$ achieves $r = 1.0047$ (+0.47% above optimal), outperforming both greedy ($r = 1.016$) and random ($r = 1.007$) heuristics. Min-cost flow and LP with branch-and-bound achieve exact optimality in under 1 s, establishing a strong classical upper bound and confirming the QUBO objective is correctly encoded.

5.4 Multi-Seed Statistical Evaluation

Table 4 confirms robustness across seeds. The $p = 2$ configuration achieves the lowest mean ratio (1.0051 ± 0.0008) with the highest sampling success probability ($5.1 \pm 1.2\%$), while $p = 3$ shows greater variance consistent with an under-optimised higher-dimensional parameter space.

5.5 Non-Monotonic Depth Behaviour

The $p = 3$ result ($r = 1.0120$, mean 1.0137 ± 0.0031) is worse than $p = 2$. This non-monotonic behaviour reflects COBYLA’s fixed 300-evaluation budget being insufficient for the $2p = 6$ -dimensional parameter space: $p = 3$ uses 52 evaluations before its termination condition (vs. 31 at $p = 1$, 44 at $p = 2$), yet reaches a higher final energy. Wurtz and Love [18] show monotonic improvement requires exact parameter optimisation.

Algorithm 1 QAOA for Capacity-Constrained MRTA**Require:** Cost matrix $\mathbf{C}$, capacities $\mathbf{cap}$, depth p , shots S , max iterations T , seeds $\mathcal{S}$ **Ensure:** Best assignment $\{x_{kj}\}$, cost c_{QAOA}

- 1: $c_{\text{greedy}} \leftarrow \text{GREEDYNN}(\mathbf{C}, \mathbf{cap})$
- 2: $\lambda \leftarrow c_{\text{greedy}}$
- 3: $\mathbf{Q} \leftarrow \text{BUILDQUBO}(\mathbf{C}, \mathbf{cap}, \lambda)$
- 4: $\hat{H}_C \leftarrow \text{QUBOtoISING}(\mathbf{Q})$
- 5: **for** each seed $s \in \mathcal{S}$ **do**
- 6: $\theta^* \leftarrow \text{COBYLA}(\theta \mapsto F_p(\theta; \hat{H}_C, S), \text{RANDINIT}(s), T)$
- 7: $\hat{\mathbf{x}} \leftarrow \text{MOSTFREQUENT}(\text{MEASUREQAOA}(\theta^*, S))$
- 8: **end for**
- 9: **return** best $\{x_{kj}\}$ over $\mathcal{S}$, cost $\sum_{k,j} c_{kj} x_{kj}$




Fig. 5. Experimental pipeline: (1) Solomon C101 coordinate loading; (2) asymmetric cost matrix construction; (3) QUBO formation ($\lambda = 280.13$); (4) QAOA circuit ($p = 1, 2, 3$, 16 qubits); (5) COBYLA optimisation; (6) result decoding and feasibility check; (7) Zenodo deposit (DOI: 10.5281/zenodo.21253730).

5.6 Energy Landscape

Figure 6 shows COBYLA optimisation trajectories for each depth.

5.7 Scalability and Simulation Notes

QAOA ($p = 1$) scales to 20 qubits in 23.67s on the Aer GPU simulator. Note that this runtime grows exponentially with qubit count due to full statevector simulation; on real quantum hardware, circuit execution time grows only linearly with gate count, though approximation ratios will be degraded by noise.

5.8 Assignment Visualisation

Figure 7 compares assignments from brute-force, QAOA ($p = 2$), and Greedy NN.

Table 1. Experimental configuration.

Parameter	Value
Spatial data source	Solomon VRPTW C101 [25]
Problem type	Capacity-constrained task assignment (not full VRPTW)
Customers (N)	8
Robots (K)	2
Capacity [$\text{cap}_0, \text{cap}_1$]	[3, 5]
Qubits (NK)	16
Robot-0 / Robot-1 start	(32, 42) SW / (48, 58) NE
QAOA depths (p)	1, 2, 3
Shots (S)	4 096
Optimiser	COBYLA, $T_{\max} = 300$
Seeds (multi-seed runs)	10 (random COBYLA initialisations)
Penalty λ	$c_{\text{greedy}} = 280.13$
Simulator	Qiskit Aer statevector (T4 GPU)
Framework	Qiskit 1.1.2 / qiskit-opt 0.6.1
Fixed-seed run	seed 42 (for reproducibility)

Table 2. Penalty ablation: effect of λ on QAOA ($p = 1$, seed 42). Insufficient λ collapses to the trivial all-zero solution.

λ	Cost	Ratio r	Assigned Cap	OK
5 (insufficient)	0.0000	0.000	0/8	×
50 (too low)	0.0000	0.000	0/8	×
150 (borderline)	277.21	1.006	8/8	✓
280.13 (dynamic λ^*)	277.050	1.005	8/8	✓

6 Reproducibility

All artefacts are deposited on Zenodo [1] under CC-BY 4.0, as shown in Stage 7 of Fig. 5. The deposit includes: the full Colab notebook with Drive-persisted checkpoints; the Solomon C101 coordinate file with hard assertions; all result JSONs and CSVs; multi-seed checkpoint archives (10 seeds per depth); and publication-quality figures in PDF and PNG. Fixed seed 42 for single-seed runs, deterministic Aer simulation, and checkpoint-based resumption ensure reproducibility.

7 Conclusion

We presented a capacity-constrained MRTA framework with three original contributions: a slack-free equality-constraint QUBO formulation (NK qubits); a revised penalty sufficient condition whose proof covers all infeasible assignments; and a fully reproducible QAOA benchmark derived from Solomon C101 spatial coordinates (task assignment only, not VRPTW). Key findings:

Table 3. Algorithm comparison on the Solomon C101-derived task-assignment instance ($N = 8$, $K = 2$, $\text{cap} = [3, 5]$). All QAOA runs: assigned = 8/8, violations = 0. Min-cost flow and LP-ILP are exact polynomial-time classical methods. QAOA runtimes reflect classical GPU simulation, not hardware execution.

Method	Type	Cost	Ratio r	Time
Brute-force	Exact classical	275.6087	1.0000	—
Min-Cost Flow	Exact classical	275.6087	1.0000	< 1 ms
LP + Branch-and-Bound	Exact classical	275.6087	1.0000	0.22 s
NumPy Exact	Classical ref	275.6087	1.0000	0.17 s
Greedy NN	Classical heuristic	280.1256	1.0164	< 1 ms
Random	Classical heuristic	277.6242	1.0073	< 1 ms
QAOA $p = 1$	Quantum (sim.)	277.0500	1.0052	21.93 s
QAOA $p = 2$	Quantum (sim.)	276.9103	1.0047	12.89 s
QAOA $p = 3$	Quantum (sim.)	278.9279	1.0120	15.83 s

Table 4. Multi-seed results (10 independent COBYLA initialisations). Mean and standard deviation of approximation ratio r and sampling success probability P_{succ} (fraction of 4096 shots recovering the optimal bitstring at depth $p = 2$). All runs: 8/8 tasks assigned, zero violations.

Depth p	Mean r	Std r	Best r	P_{succ} (%)
1	1.0056	0.0011	1.0043	3.8 ± 0.9
2	1.0051	0.0008	1.0047	5.1 ± 1.2
3	1.0137	0.0031	1.0098	2.4 ± 0.7


- (1) **Constraint satisfaction:** 100% coverage and exact capacity for all depths across all seeds.
- (2) **Near-optimal quality:** QAOA $p = 2$ achieves $r = 1.0047$ (seed 42) / 1.0051 ± 0.0008 (10 seeds).
- (3) **Strong classical context:** min-cost flow and LP-ILP both achieve $r = 1.000$ in under 1 s, confirming the formulation is correct and contextualising the quantum result.
- (4) **Scalability:** up to 20 qubits in 23.67 s (Aer simulation); hardware execution would be faster but noisier.
- (5) **Reproducibility:** complete package at <https://zenodo.org/records/21253730> [1].

Future directions include: (i) noise-model experiments on IBM Quantum hardware; (ii) constraint-preserving mixers via QAOA [15]; (iii) warm-start initialisation [17]; and (iv) scaling to $N \geq 20$ via problem decomposition.

Acknowledgments. The authors thank Poornima College of Engineering, Jaipur, Amity University Rajasthan, and Samarkand State University for institutional support. Quantum simulations were conducted on Google Colab (T4 GPU). Solomon VRPTW C101 coordinates are used with attribution to M. M. Solomon [25].

Table 5. Depth ablation: evaluations consumed and final objective energy per COBYLA run (seed 42).

p	Evaluations	Final energy	Ratio r
1	31	277.050	1.0052
2	44	276.910	1.0047
3	52	278.928	1.0120

**Fig. 6.** (Left) COBYLA energy trajectories for $p \in \{1, 2, 3\}$. Dashed: greedy NN baseline (280.13) and brute-force optimal (275.61). $p = 1$ and $p = 2$ converge below greedy; $p = 3$ stagnates. (Right) Approximation ratio r vs. circuit depth p .

Disclosure of Interests. The authors have no competing interests to declare.

References

- Gupta, S., Ray, K.: QAOA capacity-constrained multi-robot task allocation — AII 2026 reproducibility package. Zenodo (2026). <https://doi.org/10.5281/zenodo.21253730>
- Gerkey, B.P., Mataric, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. Robot. Res.* **23**(9), 939–954 (2004). <https://doi.org/10.1177/0278364904045564>
- Korsah, G.A., Stentz, A., Dias, M.B.: A comprehensive study of reconfigurable task allocation in multi-robot systems. *Int. J. Robot. Res.* **32**(12), 1495–1512 (2013). <https://doi.org/10.1177/0278364913496484>
- Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, San Francisco (1979). ISBN 978-0-7167-1045-5
- Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ (1993). ISBN 978-0-13-617549-0
- Preskill, J.: Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018). <https://doi.org/10.22331/q-2018-08-06-79>
- Gyongyosi, L., Imre, S.: A survey on quantum computing technology. *Comput. Sci. Rev.* **31**, 51–71 (2019). <https://doi.org/10.1016/j.cosrev.2018.11.002>

Table 6. QAOA ($p = 1$, seed 42) on Solomon C101-derived sub-instances. $K = 2$, balanced cap = $\lfloor N/2, N/2 \rfloor$. All runtimes are Aer statevector simulation on T4 GPU (not hardware execution times).

N	Qubits	QAOA	Optimal	Ratio r	Sim. Time (s)
4	8	152.83	150.20	1.0175	0.79
6	12	221.38	221.38	1.0000	3.29
8	16	299.64	291.47	1.0280	5.65
10	20	375.67	365.29	1.0284	23.67




Fig. 7. Task assignment comparison. (Left) Brute-force optimal ($c = 275.61$). (Centre) QAOA $p = 2$ ($c = 276.91$, $r = 1.005$). (Right) Greedy NN ($c = 280.13$, $r = 1.016$). The 1.30-unit QAOA gap corresponds to a single boundary-customer swap.

8. Farhi, E., Goldstone, J., Gutmann, S.: A quantum approximate optimization algorithm. arXiv:1411.4028 [quant-ph] (2014). <https://arxiv.org/abs/1411.4028>
9. Shaydulin, R., Ushijima-Mwesigwa, H., Safro, I., Mniszewski, S., Alexeev, Y.: A hybrid approach for solving optimization problems on small quantum computers. *Computer* **52**(6), 18–26 (2019). <https://doi.org/10.1109/MC.2019.2908942>
10. Neukart, F., Compostella, G., Seidel, C., von Dollen, D., Yarkoni, S., Parney, B.: Traffic flow optimization using a quantum annealer. *Front. ICT* **4**, 29 (2017). <https://doi.org/10.3389/fict.2017.00029>
11. Feld, S., et al.: A hybrid solution method for the capacitated vehicle routing problem using a quantum approximate optimization algorithm. *Front. ICT* **6**, 13 (2019). <https://doi.org/10.3389/fict.2019.00013>
12. Gupta, S., et al.: Adaptive multi-robot formation control with energy-efficient Bluetooth communication framework. *Sci. Rep.* **15** (2025). Springer Nature.
13. Lucas, A.: Ising formulations of many NP problems. *Front. Phys.* **2**, 5 (2014). <https://doi.org/10.3389/fphy.2014.00005>
14. Glover, F., Kochenberger, G., Du, Y.: Quantum bridge analytics I: a tutorial on formulating and using QUBO models. *4OR* **17**(4), 335–371 (2019). <https://doi.org/10.1007/s10288-019-00424-y>
15. Hadfield, S., Wang, Z., O’Gorman, B., Rieffel, E.G., Venturelli, D., Biswas, R.: From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **12**(2), 34 (2019). <https://doi.org/10.3390/a12020034>

16. McClean, J.R., Boixo, S., Smelyanskiy, V.N., Babbush, R., Neven, H.: Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018). <https://doi.org/10.1038/s41467-018-07090-4>
17. Egger, D.J., Mareček, J., Woerner, S.: Warm-starting quantum optimization. *Quantum* **5**, 479 (2021). <https://doi.org/10.22331/q-2021-06-17-479>
18. Wurtz, J., Love, P.: MaxCut quantum approximate optimization algorithm performance guarantees for $p > 1$. *Phys. Rev. A* **103**, 042612 (2021). <https://doi.org/10.1103/PhysRevA.103.042612>
19. Bengtsson, A., et al.: Improved success probability with greater circuit depth for the quantum approximate optimization algorithm. *Phys. Rev. Applied* **14**, 034010 (2020). <https://doi.org/10.1103/PhysRevApplied.14.034010>
20. Vikstål, P., et al.: Applying the quantum approximate optimization algorithm to the tail-assignment problem. *Phys. Rev. Applied* **14**, 034009 (2020). <https://doi.org/10.1103/PhysRevApplied.14.034009>
21. Quetschlich, N., Burgholzer, L., Wille, R.: Towards quantum computing for high-acuity urgent care triage. In: *Proc. IEEE Int. Conf. Quantum Computing and Engineering (QCE)*, pp. 359–370. IEEE (2023). <https://doi.org/10.1109/QCE57702.2023.00047>
22. Mahmud, M., et al.: Towards explainable and privacy-preserving artificial intelligence for personalisation in autism spectrum disorder. In: *Lecture Notes in Computer Science*, vol. 13309, pp. 356–370. Springer (2022). https://doi.org/10.1007/978-3-031-05039-8_26
23. Mahmud, M., Shaffi, N., Kaiser, M.S., et al.: Privacy-preserving multi-agent systems for human-centred inclusive healthcare AI. *AI Magazine* **47**(3), e70087 (2026). <https://doi.org/10.1002/aaai.70087>
24. Mahmud, M., Baslyman, M., et al.: Inclusive AI for digital human modelling: A perspective on explainable multimodal digital twins for accessibility and inclusion. In: *International Conference on Human-Computer Interaction*, pp. 1–16. Springer (2026).
25. Solomon, M.M.: Algorithms for the vehicle routing and scheduling problems with time window constraints. *Oper. Res.* **35**(2), 254–265 (1987). <https://doi.org/10.1287/opre.35.2.254>
26. Powell, M.J.D.: A direct search optimization method that models the objective and constraint functions by linear interpolation. In: Gomez, S., Hennart, J.P. (eds.) *Advances in Optimization and Numerical Analysis*, pp. 51–67. Kluwer, Dordrecht (1994). https://doi.org/10.1007/978-94-015-8330-5_4
27. Zhou, L., Wang, S.T., Choi, S., Pichler, H., Lukin, M.D.: Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X* **10**, 021067 (2020). <https://doi.org/10.1103/PhysRevX.10.021067>
28. Wang, Z., Hadfield, S., Jiang, Z., Rieffel, E.G.: Quantum approximate optimization algorithm for MaxCut: A fermionic view. *Phys. Rev. A* **97**, 022304 (2018). <https://doi.org/10.1103/PhysRevA.97.022304>
29. Ozaeta, A., van Dam, W., McMahon, P.L.: Expectation values from the single-layer quantum approximate optimization algorithm on Ising problems. *Quantum Sci. Technol.* **7**, 045036 (2022). <https://doi.org/10.1088/2058-9565/ac9013>
30. Harrigan, M.P., et al.: Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **17**, 332–336 (2021). <https://doi.org/10.1038/s41567-020-01105-y>