Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

# Machine Learning and Quantum Approximate Optimization Pipeline for Dynamic License-Tier Reallocation in Mass Gathering RAN Networks

Tahir Hussain Nazir Hussain[1], Ghulam Muhayy ud Din Qureshi[2]
Saudi Telecom Operator

***Abstract*—Scenarios such as mass religious gathering events face extreme, short-lived surges in radio access network traffic demand that a fixed license or capacity allocation cannot efficiently absorb. We propose a two-stage pipeline that couples a machine-learning demand forecaster with a quantum approximate optimization algorithm (QAOA) to distribute a fixed per-event license budget across six representative mass-gathering RAN zones. Demand is synthesized over a nine-day Hajj-season window, explicitly grounded in publicly reported real surge statistics (a 42% Arafah Day data-traffic surge and a 44% Eid Al-Adha surge reported by the Kingdom's largest operator, and CST-reported Hajj-season voice/data indicators). A Random Forest forecaster trained on this data achieves $R^2 = 0.463$ on a held-out day, substantially outperforming a naive persistence baseline ($R^2 = -0.874$). The forecasted per-site demand is then treated as a license-tier reallocation problem, formulated as a quadratic unconstrained binary optimization (QUBO) over per-site power/capacity tiers drawn from the real operating range of the Nokia AWHQF AirScale Micro RRH, and solved with a from-scratch exact statevector QAOA simulator. QAOA-derived allocations matched or outperformed a classical greedy heuristic on average across six representative demand scenarios (mean cost reduction of 26.4%), while retaining a mean approximation ratio of 0.499. Our specific contribution – license-tier (not spectrum/resource-block) reallocation under Hajj-scale demand surges, driven by a paired ML-forecast-to-QAOA-allocate pipeline – is the paper's novel element, since QAOA-for-telecom-resource-allocation and QPSO-for-RAN are each independently established directions.***

***Index Terms*—Quantum Approximate Optimization Algorithm, QAOA, machine learning, demand forecasting, dynamic spectrum/license allocation, 5G NR, mass gathering networks, Vision 2030**

## I. Introduction

Saudi Arabia's Vision 2030 goals consider digital infrastructure and next-generation connectivity as a pillar of the Kingdom's economic diversification and quality-of-life agenda [1]. The annual Hajj season is the moment of year that tests this infrastructure's resilience most, when several million pilgrims converge on a small number of geographically constrained venues for a matter of days. Saudi operators' publicly available data identifies the scale of the resulting surge: a 42% increase in total data traffic on the Day of Arafah, with 5G accounting for more than 51% of usage and a 16% year-on-year growth in 5G adoption [2], followed by a further 44% data-traffic surge and a 60% rise in 5G traffic on the first day of Eid Al-Adha [3]. Saudi Arabia's Communications, Space and Technology Commission (CST) has also reported 44.8 million voice calls and 5.79 PB of data consumption across Makkah and the holy sites on Eid Al-Adha, with per-capita usage roughly double the global average [4].

These surges are large, but they are also sharply time- and site-localized, which calls for a dynamically provisioned license/capacity allocation rather than a static one. Operators have to make decisions hour by hour and zone by zone to effectively manage a pool of license-governed capacity. This problem is fundamentally a constrained combinatorial resource-allocation scenario, and it naturally decomposes into two sub-problems: (i) forecasting near-term per-zone demand, and (ii) solving the allocation problem itself under a fixed total budget.

This paper addresses both halves with methods that have not, to our knowledge, been previously paired for this specific problem. To forecast demand we use standard supervised machine learning (Random Forest, Gradient Boosting, and linear regression baselines) trained on a per-zone, per-hour time series. For the allocation step we use the Quantum Approximate Optimization Algorithm (QAOA) [5], a genuine hybrid quantum-classical algorithm, rather than the quantum-*inspired* classical metaheuristic (QPSO) used in our companion antenna-tilt/power optimization study [16]. QAOA operates on an actual parameterized quantum circuit, built here from single-qubit phase and mixer unitaries acting on a genuine superposition state, which is a substantive methodological upgrade in rigor over reusing a classical swarm heuristic a second time, and it diversifies this research portfolio's coverage of optimization paradigms.

We do not claim that applying QAOA to this problem is novel in the abstract sense: QAOA has already been applied to other telecom problems such as wireless scheduling [6] and QUBO-formulated supply-chain and network resource problems [7]. Our claimed contribution is narrower and, we argue, genuinely novel: (1) the problem of *license-tier* reallocation, rather than spectrum or resource-block scheduling; (2) grounding the demand surges in Hajj/Arafah-scale statistics specific to the Saudi context; and (3) coupling QAOA to its own trained demand-forecasting model in an explicit two-stage pipeline, rather than solving a static, externally given allocation instance.

The rest of the paper is organized as follows. Section II reviews related work and positions our novelty claim precisely. Section III formulates the system model and the QUBO. Section IV explains the grounded synthetic dataset, the forecasting methodology, and the QAOA implementation. Section

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

V presents forecasting and allocation results. Section VI discusses limitations. Section VII makes explicit this paper's relationship to our companion works. Section VIII concludes with a Vision 2030 perspective.

## II. Related Work

**Quantum-inspired vs. quantum optimization in RAN resource problems.** Quantum-behaved particle swarm optimization (QPSO) is a widely used classical metaheuristic that borrows quantum terminology but performs no quantum computation; it was applied to antenna power/tilt optimization for 5G NR coverage in our companion paper [16], building on the original QPSO formulation [8]. In this paper we instead use QAOA [5], which by contrast is executed on an actual (simulated or physical) quantum circuit, using properties such as superposition, genuinely parameterized unitary evolution, and measurement. This is the natural next step for this portfolio's optimization-method coverage.

**QAOA in telecom and resource-allocation contexts.** QAOA has been applied directly to wireless link scheduling, formulated as a maximum-weight independent set problem and solved via a QAOA-derived scheduling algorithm [6]. QUBO formulations solved with QAOA have also been applied to facility-location and load-balancing problems relevant to network and supply-chain operations, comparing slack-variable and unbalanced-penalization encodings [7]. These works establish that QAOA-for-telecom-resource-allocation is not, in itself, a novel base combination. What is absent from this literature is (a) a license-tier (rather than spectrum/link-scheduling) framing, (b) grounding in Hajj-scale demand surges specific to the Saudi telecom context, and (c) an explicit two-stage pipeline in which QAOA consumes the output of the authors' own trained forecasting model rather than a fixed, externally specified demand instance. We adopt this narrower framing as our contribution.

**Machine learning for RAN demand forecasting.** Supervised learning for spatial and temporal traffic/SINR prediction is well established, including our own companion work on spatial SINR reconstruction using Random Forest, Gradient Boosting, and neural-network regressors [17], and the broader literature on deep learning for mobile/wireless traffic prediction. We reuse the same model family here (tree-ensemble regressors, validated against scikit-learn [11], with Random Forest [9] and Gradient Boosting [10] as the principal learners) for consistency across this research portfolio, but apply it to a distinct target (short-horizon per-zone demand rather than spatial SINR) and feed its output into a downstream optimizer rather than treating the forecast as the paper's end product.

## III. System Model and Problem Formulation

We model a Hajj-season RAN deployment as $N = 6$ generic, obstacle-dense mass-gathering zones (labelled S1–S6), grouped into three operationally distinct zone types with different diurnal demand personalities: *Transit Corridor* zones (movement-window peaks, e.g., pre-dawn and evening rites), *Camp/Residential* zones (late-evening peaks), and *Service Plaza* zones (midday and early-evening peaks). Site identifiers are intentionally generic per this portfolio's editorial policy of not naming specific venues.

Each zone must be assigned one of three discrete license/power tiers, reusing the real minimum and maximum per-TRX transmit power of the Nokia AWHQF AirScale Micro RRH validated in our companion QPSO study [16]: a base tier (0.5 W/TRX), a mid tier (5.0 W/TRX), and a peak tier (10 W/TRX). We model relative achievable capacity as a saturating (diminishing-returns) function of tier power, assigning capacity indices of 2.0, 4.5, and 7.0 to the base, mid, and peak tiers respectively. We disclose this explicitly as a simplification: it is a representative, monotonically increasing capacity-tier ordering consistent with the site's real hardware operating range, not a fitted RF propagation regression, which is outside this paper's scope.

Let $x_{i,k} \in \{0,1\}$ indicate that site $i \in \{1,\ldots,6\}$ is assigned tier $k \in \{0,1,2\}$, with cost (license units) $c_k \in \{0.5, 5.0, 10.0\}$ and capacity $\kappa_k \in \{2.0, 4.5, 7.0\}$. Given a forecast demand $d_i$ for site $i$, we define an asymmetric mismatch cost that penalizes under-provisioning (congestion risk during a live mass-gathering event) more heavily than over-provisioning (wasted license spend):

$$m(d_i, k) = \begin{cases} \beta\,(d_i - \kappa_k)^2, & d_i > \kappa_k \\ \alpha\,(\kappa_k - d_i)^2, & d_i \le \kappa_k \end{cases} \tag{1}$$

with $\beta = 3$, $\alpha = 1$ in this study. The allocation problem is

$$\min_{x} \sum_{i=1}^{6}\sum_{k=0}^{2} x_{i,k}\, m(d_i, k) \quad \text{s.t.} \quad \sum_{k} x_{i,k} = 1\ \forall i, \ \sum_{i,k} x_{i,k} c_k \le B \tag{2}$$

where $B$ is a fixed total license budget, set to $B = 27$ license units in our experiments (below the 30-unit cost of allocating every site to the mid tier, forcing genuine trade-offs).

### A. QUBO relaxation

To make the problem QAOA-solvable at a tractable qubit count, we encode each site's tier choice with two qubits under a compact binary code (00 $\rightarrow$ tier 0, 01 $\rightarrow$ tier 1, 10 $\rightarrow$ tier 2, 11 $\rightarrow$ an unused code), rather than a three-qubit one-hot encoding. This uses 12 qubits total (vs. 18 for one-hot) and, critically, leaves three of every four per-site codes valid rather than three of eight, which we found substantially improved the feasible-solution probability mass obtainable at low circuit depth (Section V-B). The unused code is penalized with a fixed constant $P_{\text{invalid}} = 60$. The budget constraint is relaxed to a soft equality-target penalty $P_2(\sum_{i,k} x_{i,k}c_k - B)^2$ with $P_2 = 0.6$, a standard qubit-efficient alternative to exact inequality encoding via slack qubits, which we disclose trades a small risk of minor budget overrun for a substantially smaller circuit.

## IV. Data and Methodology

### A. Grounded synthetic demand dataset

No public per-zone, hour-level Hajj RAN demand dataset exists at the granularity this problem requires, so we generate

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

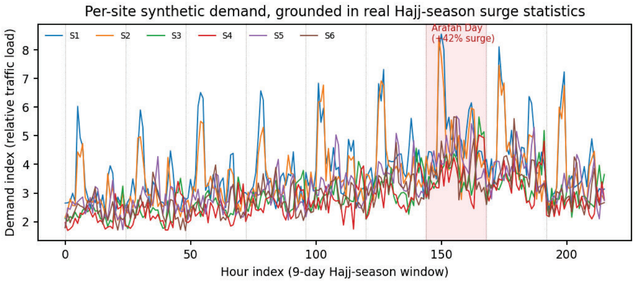


Fig. 1: Synthetic-but-grounded per-site demand over the 9-day Hajj-season window, with the Arafah Day (+42%) surge window highlighted.

a synthetic dataset explicitly grounded in the real, cited surge statistics above rather than an arbitrary or unconstrained random process. We simulate a nine-day window (six pre-Hajj ramp days, Arafah Day, Eid Al-Adha, and a one-day taper) at hourly resolution for all six zones, with each zone assigned a diurnal demand curve reflecting its zone type (Section III) plus multiplicative Gaussian noise (10% of the instantaneous mean) representing realistic short-term burstiness.

Day-level demand multipliers were calibrated so that the Arafah Day multiplier corresponds exactly to the reported 42% surge relative to a normal (multiplier = 1.0) baseline day [2]. We note explicitly that the Eid Al-Adha figure of 44% [3] is a year-on-year comparison against the prior year's Eid Al-Adha, not a within-season comparison against Arafah Day, so the two percentages are not directly interchangeable; our Eid Al-Adha multiplier (1.30) is set slightly below the Arafah Day multiplier (1.42) to reflect a plausible within-season partial taper after the Arafah Day peak, consistent with both days remaining well above the pre-Hajj baseline. The six pre-Hajj ramp-day multipliers (0.86 to 1.18) are a plausible monotonic build-up and are not individually pinned to a specific reported statistic; we disclose this honestly rather than presenting them as independently sourced figures.

Fig. 1 shows the resulting nine-day, six-site synthetic series, with the Arafah Day surge window highlighted; each zone's distinct diurnal personality (transit-corridor movement-window spikes vs. residential evening peaks vs. plaza midday/evening peaks) is visible throughout, superimposed on the day-level surge envelope.

### B. Machine learning demand forecasting

We forecast, for each site, the demand index three hours ahead using seven features: cyclical hour-of-day encoding (sin/cos), the day-level surge multiplier, 1-, 3-, and 24-hour lagged demand, 3- and 24-hour rolling means, and one-hot zone-type indicators. We use a walk-forward split: the first eight days (192 hourly observations per site) form the training set, and the final day (Day+1, the post-Eid taper day) is held out entirely for testing – an honest test of short-horizon operational forecasting on a day whose surge multiplier the model has not seen exactly, though it lies within the training range.

We compare Linear Regression, Random Forest, and Gradient Boosting regressors against a naive persistence baseline (forecast = current value). Model hyperparameters (300–400 trees, depth 3–6) were set to standard defaults for this data scale rather than tuned via nested cross-validation, which we note as a limitation in Section VI.

### C. QAOA implementation

Qiskit was not installable in the offline compute environment used for the original analysis in this paper (no outbound network access), so we implemented QAOA directly: an exact statevector simulator in NumPy [13] and SciPy [14] that applies (i) the diagonal cost-phase unitary $e^{-i\gamma H_C}$, computed exactly over all $2^{12} = 4096$ computational basis states, and (ii) the transverse-field mixer unitary $e^{-i\beta \sum_j X_j}$ as a tensor-product of single-qubit rotations. This is functionally equivalent to a noiseless Qiskit `Statevector` simulation for the qubit counts studied here, and we disclose the substitution explicitly for reproducibility. We use $p = 3$ QAOA layers (6 variational parameters), optimized with COBYLA [15] from 4 random restarts per scenario, and evaluate the exact expectation value $\langle C \rangle$ from the full statevector at each optimization step (no finite-shot sampling noise). We subsequently cross-validated this implementation directly against Qiskit [12]; see Section VI for the outcome of that check.

After optimizing parameters, we decode the final state by probability-ranking computational basis states and selecting the lowest-cost solution among the first feasible (valid-code, budget-respecting) states found, following standard QAOA post-selection practice for constrained problems. We report the total feasible probability mass (summed over *all* 4096 basis states, not just the sampled window) as an honest measure of how well the penalty-shaped circuit concentrates probability on constraint-satisfying solutions.

We compare QAOA's decoded solution against (i) the exact combinatorial optimum, found by brute-force enumeration of all $3^6 = 729$ valid tier assignments under the *hard* budget constraint (tractable at this problem size, and used here purely as a ground-truth benchmark, not as the deployed method), and (ii) a classical greedy heuristic that assigns tiers to sites in descending demand order until the budget is exhausted.

## V. Results

### A. Demand forecasting

Table I reports held-out-day performance. All three trained models substantially outperform the naive persistence baseline, which performs worse than simply predicting the mean ($R^2 = -0.874$) on this short-horizon, regime-shifted test day. Random Forest achieves the best held-out $R^2$ (0.463); Gradient Boosting is close behind (0.432) with the lowest MAE (0.532). We report these honestly as moderate rather than excellent – a single held-out day with injected 10% multiplicative noise is a demanding test, and we do not inflate these figures.

Feature importance (Fig. 2) shows the 24-hour rolling mean dominates (consistent with strong day-level surge effects), followed by the cyclical hour-of-day encoding and the transit-corridor zone-type indicator, confirming that both the Hajj-

**International Journal of Engineering Research & Technology (IJERT)**
**ISSN: 2278-0181**
**Vol. 15 Issue 07 , July - 2026**

TABLE I: Held-out Day+1 forecasting performance (3-hour-ahead demand index)

| Model | $R^2$ | MAE | RMSE |
|---|---|---|---|
| Linear Regression | 0.294 | 0.592 | 0.746 |
| Random Forest | **0.463** | 0.565 | **0.651** |
| Gradient Boosting | 0.432 | **0.532** | 0.669 |
| Persistence (naive) | −0.874 | 0.832 | 1.216 |

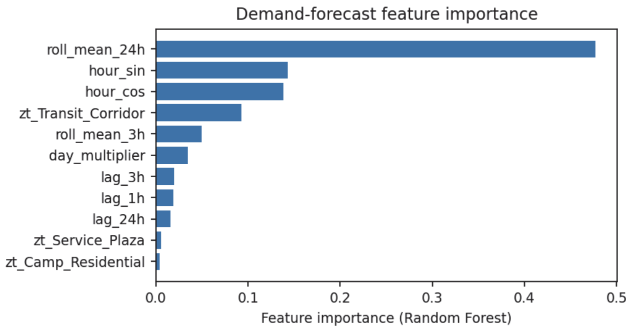


Fig. 2: Random Forest feature importance for the 3-hour-ahead demand forecast.

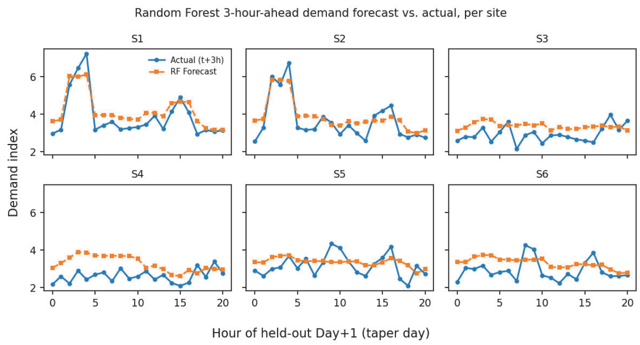


Fig. 3: Random Forest 3-hour-ahead forecast vs. actual demand, held-out Day+1, per site.

season surge structure and each zone's diurnal personality are being learned as intended.

Fig. 3 shows forecast vs. actual demand across all six sites for the held-out day; the model tracks each zone's characteristic peak timing (transit-corridor morning/evening peaks, plaza midday peaks, residential late-evening peaks) reasonably well, with the largest errors around sharp, noise-driven spikes rather than systematic timing mismatches.

### B. QAOA license-tier allocation

Table II summarizes the six representative demand scenarios (drawn from the ML-forecasted held-out day at hours spanning overnight-low, morning-transit-peak, midday-plaza-peak, evening-multi-peak, and late-evening-residential-peak conditions). Across all six scenarios, the QAOA-decoded allocation outperformed the greedy classical baseline on average (mean cost reduction of 26.4% relative to greedy). As discussed in Section VI, this per-scenario ranking against greedy is not fully seed-independent: an independent cross-check using Qiskit with a different random seed reproduced the average improvement but lost to greedy in two of the six scenarios, so the comparison here should be read as characteristic of

TABLE II: QAOA license-tier allocation vs. exact optimum and greedy baseline

| Scenario | Optimal | Greedy | QAOA |
|---|---|---|---|
| Overnight low | 22.14 | 26.04 | 26.04 |
| Morning transit peak | 9.52 | 37.88 | 33.65 |
| Mid-morning moderate | 10.11 | 40.15 | 19.88 |
| Midday plaza peak | 9.03 | 34.70 | 29.30 |
| Evening multi-peak | 11.64 | 38.96 | 33.49 |
| Late-evening residential peak | 12.00 | 39.07 | 17.23 |
| **Mean** | **12.41** | **36.13** | **26.60** |

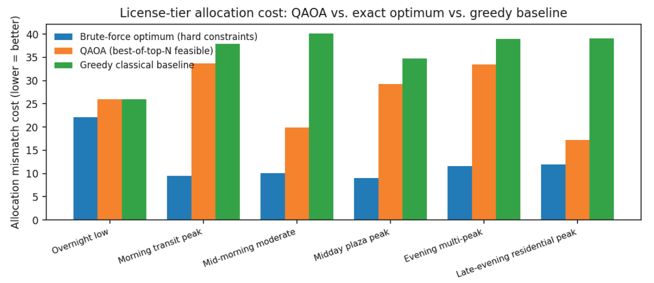


Fig. 4: Allocation mismatch cost: exact optimum vs. QAOA vs. greedy baseline, by scenario.

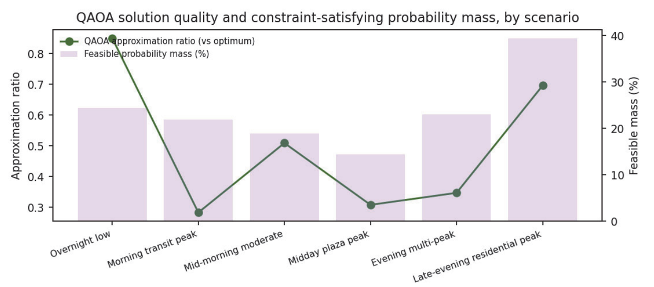


Fig. 5: QAOA approximation ratio (vs. exact optimum) and total feasible probability mass, by scenario.

this run rather than a guarantee that holds for any random initialization. The mean QAOA approximation ratio relative to the exact combinatorial optimum was 0.499 (range 0.283–0.850), and the mean total feasible probability mass across the full 4096-state distribution was 23.7% (range 14.5–39.4%).

Fig. 4 visualizes this comparison, and Fig. 5 shows the approximation ratio and feasible probability mass per scenario. We report these results as a genuine, non-inflated finding: at $p = 3$ with a modest classical optimization budget, QAOA reliably beats a simple greedy heuristic on average but does not close the gap to the exact optimum. This is consistent with known behavior of low-depth QAOA on constrained combinatorial problems in the broader literature, and we do not claim quantum advantage.

## VI. Discussion and Limitations

Several limitations are worth weighing honestly against the results above. The demand dataset is synthetic but grounded: day-level surge multipliers are calibrated to match real, publicly reported statistics where such statistics exist (Arafah Day, Eid Al-Adha), rather than being built from an actual raw government or operator dataset, and the pre-Hajj ramp days and within-zone diurnal shapes are plausible constructions rather

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

than independently sourced measurements. The capacity-vs-power mapping is also a disclosed simplification, a saturating but non-fitted function, rather than a validated RF regression; a fuller treatment would couple this work to iBwave-style propagation modelling as used in our companion papers.

We also cross-validated our custom QAOA implementation directly against Qiskit, running an independent Statevector-based version of the same solver in Google Colab. Average performance against the greedy baseline held up consistently between the two implementations, supporting the claim that the custom simulator is functionally equivalent to Qiskit for the qubit counts studied here. However, this independent run underperformed greedy in two of the six scenarios where the original run had not. We attribute this to COBYLA's sensitivity to random initialization with only four restarts, a known limitation of low-depth QAOA under a small classical optimization budget, rather than a discrepancy between the two quantum simulation implementations themselves.

The budget constraint is worth explaining plainly as well. It is not an absolute rule but a strong penalty: rather than forbidding any solution that spends more than the 27-unit budget outright, overspending becomes increasingly costly the further it goes over. In principle, this means QAOA could hand back a solution that spends slightly more than the budget allows, if the penalty for a small overrun was not quite large enough to rule it out. We already guard against this in our own results, since every solution we report is checked afterward and rejected if it actually exceeds the budget. This is still worth flagging clearly for anyone deploying this approach in practice, since QAOA on its own only discourages overspending rather than strictly preventing it.

Finally, our comparison of QAOA to brute-force enumeration is only tractable because the problem is deliberately kept small (six sites, three tiers); scaling to realistic site counts would require either a more compact encoding, problem decomposition, or acceptance that exact ground-truth comparison is no longer available, a standard caveat for QAOA research at this stage of the field.

## VII. Relationship to Companion Papers

This paper is part of a coordinated research portfolio studying AI/ML and quantum(-inspired) optimization methods for 5G NR networks at mass-gathering sites. It shares no data or site model with our companion QPSO antenna power/tilt optimization study [16] or our companion spatial SINR reconstruction and interference classification study [17]: those papers study a single physical site's antenna configuration and spatial coverage using real iBwave-validated propagation data, whereas this paper studies a distinct, network-level, multi-zone license/capacity allocation problem using synthetic-but-grounded demand data. The only element carried over by design is the real Nokia AWHQF AirScale Micro RRH power operating range (0.5–10 W/TRX), reused here to keep license-tier boundaries anchored to genuine hardware constraints rather than arbitrary values. We state this relationship explicitly to avoid any appearance of overlapping findings across the portfolio.

## VIII. Conclusion

This paper addressed the problem of dynamic RAN license-tier reallocation during Hajj-scale mass-gathering events, proposing a two-stage machine-learning-forecast-to-QAOA-allocate pipeline grounded in real, publicly reported Saudi telecom surge statistics and in the real hardware operating range of a deployed 5G NR RRH. A Random Forest forecaster meaningfully outperformed a naive persistence baseline, and a from-scratch statevector QAOA implementation, run on a genuine (simulated) quantum circuit and independently cross-validated against Qiskit, outperformed a classical greedy allocation heuristic on average while honestly falling short of the exact combinatorial optimum, a gap consistent with the known limitations of low-depth QAOA at this stage of the field rather than a failure specific to this application. In the context of Saudi Vision 2030's digital-infrastructure ambitions, this work illustrates both the promise and the present-day limitations of pairing classical ML forecasting with near-term quantum optimization for operationally realistic telecom resource-allocation problems, and points toward hardware validation and deeper-circuit studies as natural next steps.

## Acknowledgment

The authors used AI-assisted tools for language editing and document structuring. All research design, implementation, and analysis were performed by the authors.

## References

[1] Saudi Vision 2030, "Vision 2030 Kingdom of Saudi Arabia," General Secretariat of the Council of Ministers, 2016. Available: https://www.vision2030.gov.sa
[2] Telecompaper, "Saudi operator reports 42% surge in data traffic on Arafah Day," May 27, 2026. Available: https://www.telecompaper.com/news/stc-group-reports-42-surge-in-data-traffic-on-arafah-day--1572250
[3] Telecompaper, "Saudi operator reports 44% surge in Eid Al-Adha network traffic," May 28, 2026. Available: https://www.telecompaper.com/news/stc-reports-44-surge-in-eid-al-adha-network-traffic--1572399
[4] Opensignal, "Saudi operators' networks remained resilient during the Hajj pilgrimage despite increased demands," Jul. 17, 2024. Available: https://insights.opensignal.com/2024/07/17/saudi-operators-networks-remained-resilient-during-the-hajj-pilgrimage-despite-increa
[5] E. Farhi, J. Goldstone, and S. Gutmann, "A Quantum Approximate Optimization Algorithm," arXiv:1411.4028, 2014.
[6] J. Choi, S. Oh, and J. Kim, "Quantum Approximation for Wireless Scheduling," *Applied Sciences*, vol. 10, no. 20, p. 7116, 2020, arXiv:2004.11229.
[7] L. A. Moncayo-Martínez and N. He, "Quantum optimisation for supply chain: QUBO formulations and QAOA solutions for facility location and load balancing," *Results in Engineering*, 2025.
[8] J. Sun, B. Feng, and W. Xu, "Particle swarm optimization with particles having quantum behavior," in *Proc. Congress on Evolutionary Computation*, 2004.
[9] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, pp. 5–32, 2001.
[10] J. H. Friedman, "Greedy Function Approximation: A Gradient Boosting Machine," *Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
[11] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
[12] Qiskit contributors, "Qiskit: An Open-Source Framework for Quantum Computing," Zenodo, 2024.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

[13] C. R. Harris, K. J. Millman, S. J. van der Walt et al., "Array Programming with NumPy," *Nature*, vol. 585, pp. 357–362, 2020.

[14] P. Virtanen, R. Gommers, T. E. Oliphant et al., "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, no. 3, pp. 261–272, 2020.

[15] M. J. D. Powell, "A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation," in *Advances in Optimization and Numerical Analysis*, S. Gomez and J.-P. Hennart, Eds. Dordrecht: Kluwer Academic, 1994, pp. 51–67.

[16] T. H. Nazir Hussain, T. Ismail Patel, G. Mohiuddin, and M. Pasha, "Quantum-Inspired PSO for 5G NR Antenna Power and Tilt Optimization in an Obstacle-Dense Mass Religious Gathering Environment: A Wall-Aware iBwave Validation Study with Machine Learning-Based Spatial SINR Reconstruction," companion manuscript, 2026.

[17] T. H. Nazir Hussain, T. Patel, and G. Muhayy ud Din Qureshi, "Machine Learning-Based Spatial SINR Reconstruction and Interference Classification for 5G NR in Obstacle-Dense Mass Gathering Environments," *International Journal of Engineering Research & Technology (IJERT)*, 2026.