Noise-Aware Quantum Approximate Optimization for Carbon-Aware Scheduling of Renewable Generation, Battery Storage and Industrial Loads

Multiobjective QUBO formulation, depth selection under depolarizing noise, and validation on IEEE 14-bus derived price and carbon signals

[1]Saswat Swain, [2]Chandrasekhar Panda, [3]Pravat Satpathy
[1]Ph.D Scholar, [2]Researcher, [3]Researcher
[1]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[2,3]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[1]saswat.mck@gmail.com, [2]contact@chandrasekharpanda.com, [3]analytics.pravat@gmail.com

***Abstract*—**Carbon-aware scheduling of flexible plant assets couples a price signal and a time-varying grid carbon intensity to the dispatch of battery storage and deferrable industrial loads, which turns an otherwise small unit-scheduling task into a constrained binary program whose landscape is dominated by penalty terms. This paper examines whether the Quantum Approximate Optimization Algorithm (QAOA) can be made a credible solver for that program on noisy devices, with the negative findings reported alongside the positive ones. A multiobjective quadratic unconstrained binary optimization (QUBO) model is derived in which energy cost, a carbon price applied to the grid carbon intensity, a quadratic peak-demand surrogate and an end-of-horizon valuation of stored energy share one quadratic form; terminal battery neutrality is imposed in cycle-count space so that round-trip losses do not distort the constraint, and a state-of-charge excursion regularizer is shown to have a usable weight window within which the penalized optimum is feasible and within 0.08 % of the constrained optimum. Price and carbon signals are derived from hourly economic dispatch of the IEEE 14-bus generators, which exposes a marginal emission factor that is almost flat under coal-first merit order; the average emission factor is therefore used as the primary signal and the difference is quantified. For noise, a closed-form proposition shows that global depolarizing noise contracts the QAOA energy without moving its optimizers, which converts depth selection into a one-line rule validated numerically; for local depolarizing noise, re-optimization on a common-random-number trajectory surrogate recovers at most 0.029 in normalized gain over transferred angles once optimizer convergence is controlled for. On 12-, 15- and 18-qubit instances, QAOA at depth six concentrates between 34 % and 68 % of its probability on the feasible subspace and a feasibility-filtered decoder with 1000 shots reaches the constrained optimum to within 0.15 %; simulated annealing with the same sample budget reaches it exactly. The results support QAOA as a structured sampler for this problem class while placing its present advantage, if any, outside the regime reachable by exact simulation.

***Index Terms*—**Quantum approximate optimization algorithm, QUBO, carbon-aware scheduling, smart grid, battery energy storage, demand response, depolarizing noise, multiobjective optimization, IEEE 14-bus system, sustainability computing.

I. Introduction

Decarbonizing industrial electricity use is increasingly a scheduling problem rather than a procurement problem. Once a plant has rooftop or captive photovoltaic (PV) capacity, a share of wind, a battery energy storage system (BESS) and at least one process whose start time is negotiable, the hourly decision of when to draw from the grid determines both the bill and the emissions attributable to the site. The two signals that drive that decision, the time-of-day tariff and the grid carbon intensity, are not aligned in general: in coal-dominated systems the cheapest hours are frequently the most carbon-intensive, so a carbon-aware scheduler has to resolve a genuine trade-off rather than exploit a correlation [1], [2].

The resulting problem is a mixed-integer program with binary charge, discharge and run decisions per slot, a state-of-charge (SoC) recursion with round-trip losses, duration requirements on deferrable loads and a quadratic peak-demand component. Such penalty-dominated binary programs are a frequently cited target for quantum heuristics [3]–[5]. The Quantum Approximate Optimization Algorithm (QAOA) [6] is the natural candidate for gate-based hardware because its ansatz is built directly from the problem Hamiltonian, and its variants have been applied to unit commitment [7], [8], energy-community operation [9] and microgrid management [10]. What those studies rarely do is separate three questions that have different answers: whether the QUBO encoding is faithful to the constrained problem; whether the QAOA angles obtained on a noiseless simulator survive hardware noise; and whether the sampled bitstrings, after infeasible ones are removed, are better than what a classical sampler returns for the same budget.

This paper addresses the three questions on a carbon-aware scheduling model whose price and carbon signals are derived from the IEEE 14-bus test system, so that the benchmark link is explicit rather than nominal. The emphasis on noise is deliberate. Reported gate fidelities on

superconducting devices imply per-layer error budgets that scale with the hundreds of two-qubit gates a dense scheduling QUBO requires [11], [12], and the practical question for an energy engineer is not the asymptotic power of QAOA but how deep a circuit is still worth running and whether its parameters have to be re-tuned on the device. Both are answered here with a mixture of closed-form analysis and exact simulation.

The contributions are: (i) a multiobjective QUBO for renewable, storage and industrial-load scheduling in which cost, priced carbon, a peak surrogate and an end-of-horizon SoC valuation form one quadratic objective, with terminal neutrality encoded in cycle-count space and an SoC excursion regularizer whose weight window keeps the penalized optimum feasible; (ii) a signal set derived from hourly economic dispatch of the IEEE 14-bus generators, with marginal and average emission factors obtained analytically, the near-flat marginal factor under coal-first merit order being reported as a finding; (iii) a proposition showing that global depolarizing noise leaves QAOA optimizers invariant and makes the optimal depth follow a ratio rule, with numerical confirmation, and a controlled study of noise-aware re-optimization under local noise that separates the gain due to noise awareness from that due to optimizer convergence; (iv) a feasibility-filtered decoder whose expected best-of-N cost is computed exactly from the QAOA distribution, allowing QAOA, simulated annealing and uniform sampling to be compared without sampling error. All code, data and figure scripts are released in a reproducibility bundle.

II. Related Work

Table I summarizes the studies closest to the present one. Early quantum-annealing and QAOA treatments of unit commitment [7], [13] established the QUBO encodings of on/off and ramping decisions and reported that penalty weights dominate the landscape. Salgado *et al.* [8] split commitment and power levels between a depth-one QAOA and a classical optimizer, reaching 5.1 % of a reference solution, while Halffmann *et al.* [14] exploited QAOA sampling to extract robust rather than merely optimal schedules for plants and electric-vehicle charging. On the demand side, QUBO models of energy communities [9] and a QAOA microgrid scheduler validated on real PV, wind and storage data [10] show that the encoding scales to practical assets, but neither quantifies feasibility of the sampled solutions nor treats noise beyond a qualitative remark. Reviews of quantum computing for power systems [3]–[5] converge on the same two open items: constraint handling and device noise.

The noise literature on QAOA is more developed but largely detached from applications. Xue *et al.* [15] showed numerically that depolarizing and dephasing noise flatten the QAOA landscape without moving its extrema, and Marshall *et al.* [16] characterized local noise in QAOA circuits; Wang *et al.* [17] proved that local noise induces barren plateaus whose severity grows with depth. Parameter-setting heuristics such as the INTERP and FOURIER strategies [18], annealing-inspired initialization [19] and parameter concentration [20] reduce the classical cost of training, and noise-aware classical optimizers for QAOA have recently been proposed [21]. Carbon-aware computing and demand response [1], [2] supply the signal definitions used here, and the marginal-versus-average distinction follows [22]. The gap this paper fills is the joint treatment: a benchmark-grounded carbon-aware scheduling QUBO, an analysis of which part of the noise-induced loss is recoverable by parameter re-tuning, and a decoder-level comparison with classical sampling at equal budget.

TABLE I
RELATED STUDIES ON QUANTUM OPTIMIZATION FOR ENERGY SCHEDULING AND ON NOISY QAOA

Study (year)	Problem / method	Strength	Limitation relative to this work
Koretsky et al. [7] (2021)	Unit commitment, QAOA with classical outer loop	Early QAOA formulation of commitment with ramping	Noiseless; no feasibility statistics
Nikmehr et al. [13] (2022)	Distributed unit commitment, quantum annealing	Decomposition keeps subproblems embeddable	Annealer-specific; no gate-model noise
Salgado et al. [8] (2024)	Constrained unit commitment, depth-1 QAOA + classical power levels	5.1 % of reference; polynomial hybrid scaling	Depth fixed at one; noise not modelled
Halffmann et al. [14] (2024)	Robust plant / EV scheduling, QAOA and annealing sampling	Uses sub-optimal samples constructively	Robustness, not noise; no benchmark grid
Energy-community QUBO [9] (2024)	Community operation as QUBO	Practical asset model	No quantum results under noise
Liu et al. [10] (2025)	Microgrid scheduling, QAOA on real PV / wind / storage data	Full hybrid loop on field data	No feasibility rate, noise or equal-budget baseline
Xue et al. [15] (2021)	QAOA under depolarizing / dephasing noise	Landscape flattening without extremum shift	Generic Max-Cut; no application
Dzahini et al. [21] (2025)	Noise-aware subspace optimizer for QAOA	Shot-efficient trust-region training	Measurement noise, not channel noise
This work (2026)	Carbon-aware storage / load QUBO; QAOA with global and local depolarizing noise	Benchmark signals; invariance proof; exact best-of-N decoder; equal-budget SA	Exact simulation limits n to 18; no hardware run

III. Problem Statement and System Model

A. Scheduling Problem

A plant is connected to the distribution grid at a single point of common coupling. The horizon of 24 h is divided into T slots of length $\Delta = 24/T$ h. In slot t the non-deferrable base load is d_t (kW), on-site renewable generation is $g_t = \mathrm{PV}_t + \mathrm{W}_t$, the retail tariff is τ_t (USD/MWh) and the grid carbon intensity is κ_t (kg$\mathrm{CO_2}$/kWh). A BESS of capacity E_B and rated

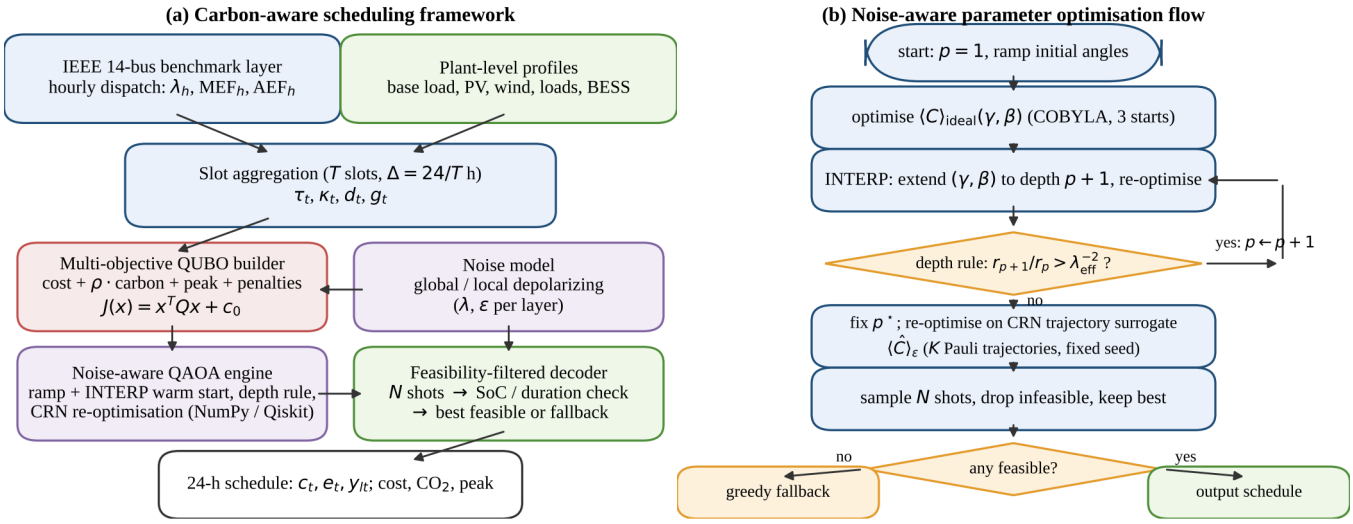


Fig. 1. Proposed framework. (a) Block diagram from the IEEE 14-bus dispatch layer to the decoded 24-h schedule. (b) Flowchart of the noise-aware QAOA parameter optimization: ramp initialization at depth one, INTERP extension, the depth rule of Proposition 1, common-random-number (CRN) re-optimization under local noise, and feasibility-filtered decoding with a greedy fallback.

power P_B is operated in one of three modes per slot through two binaries, c_t (charge) and e_t (discharge). Each deferrable industrial load l draws $P_{L,l}$ when running and must run for exactly D_l slots; its run indicator is $y_{l,t}$. Grid exchange, positive for import, is

(1)

and the SoC recursion with charging and discharging efficiencies η_c, η_d is

(2)

The constrained problem minimizes the daily operating objective

(3)

subject to $c_t e_t = 0$, $S_{min} \le S_t \le S_{max}$, $\Sigma_t\, y_{l,t} = D_l$ and the terminal condition $S_T \ge S_0 - \tfrac{1}{2} P_B \Delta$. Here ρ is a carbon price in USD/tCO_2 that converts emissions into the cost unit and sweeps the trade-off, w_p weights a quadratic surrogate of the demand charge, and π_s, set to the horizon-mean effective energy price $(\tau_t + \rho\kappa_t)/1000$ in USD/kWh, values any net depletion of stored energy so that a schedule cannot appear cheap by ending the day with an empty battery. Net metering is assumed because an asymmetric export tariff adds a non-quadratic max(·, 0) term; the quadratic peak surrogate penalizes all large imports rather than only the largest and is therefore conservative relative to a monthly demand charge.

B. Benchmark-Derived Price and Carbon Signals

The IEEE 14-bus system is a power-flow benchmark and carries no fuel or emission data, so it is used here for what it does contain: five generators with quadratic cost curves and capacity limits as distributed with MATPOWER [23]. Fuel types and emission factors are assigned by assumption (Table II) and the system is dispatched hour by hour against the case load scaled by an industrial daily shape, net of embedded utility-scale PV and wind. For demand D_h the quadratic economic dispatch

(4)

has the system marginal cost $\lambda_h = 2a_i P_i + b_i$ for every unconstrained unit $i \in U_h$. Differentiating the optimality conditions with respect to demand gives the share of an increment taken by each unconstrained unit, $(1/2a_i)/\Sigma_{j \in U}(1/2a_j)$, and hence the marginal emission factor and the average emission factor of delivered energy

(5)

where ε_i is the emission factor of unit i. The retail tariff is $\tau_h = \lambda_h + \delta_h$ with a regulated time-of-day adder δ_h of 8, 15 and 30 USD/MWh in the night, day and evening-peak blocks, mirroring the structure of industrial time-of-day tariffs. Hourly signals are averaged over each slot to obtain τ_t and κ_t. Because MEF is the correct signal for a marginal consumption decision while AEF is what carbon accounting and most public intensity feeds report [22], both are computed and the scheduling consequences of the choice are reported in Section VI-G.

IV. Proposed Methodology

A. Framework

Fig. 1(a) shows the pipeline. The benchmark layer produces hourly λ_h, MEF_h and AEF_h; plant-level profiles and the slot aggregator produce the per-slot parameters; the QUBO builder assembles the quadratic form; the QAOA engine, parameterized by a noise model, returns a sampling distribution; and a feasibility-filtered decoder converts shots into a schedule or, if no shot is feasible, into a rule-based fallback. Fig. 1(b) details the parameter flow that constitutes the noise-aware part of the method: layer-wise noiseless training with INTERP warm starts, a depth rule derived in Section IV-D, and an optional re-optimization on a trajectory surrogate when the noise is local.

B. QUBO Encoding

All decision variables are binary, $x = (c, e, y) \in \{0,1\}^n$ with $n = (2 + L)T$, and G_t and S_t are affine in x, so every term of (3) is at most quadratic. Constraints are added as penalties:

(6)

Two choices distinguish (6) from a textbook encoding. First, terminal neutrality is imposed on the count of charge and discharge slots rather than on $S_T - S_0$: with $\eta_c\eta_d < 1$ a SoC-based penalty can never be zero, so it would bias the optimizer against cycling by an amount that depends on efficiency rather than on economics; the residual drift of $(1/\eta_d - \eta_c)P_B\Delta$ per cycle is instead priced by π_S in (3) and bounded by the terminal condition. Second, the SoC box constraints, which would need slack qubits, are replaced by a quadratic excursion regularizer of weight μ (Section VI-B); any remaining violation is removed by the decoder. The penalty weights A, B, C are set to a fraction f of the range $R = \max J_0 - \min J_0$ of the unpenalized objective over all bitstrings; $f = 1$ is sufficient by the usual argument that a violation can gain at most R, and the smallest f that still yields a constraint-satisfying optimum is found empirically. Expanding (6) gives $J(x) = x^{\mathrm{T}}Qx + c_0$ with an upper-triangular Q that is recovered by probing single and paired bits and verified against enumeration to 10^{-10}.

C. QAOA Ansatz

With C the diagonal operator whose eigenvalues are $J(x)$, the depth-p state and the trained quantity are

(7)

(8)

where $X = \Sigma_k X_k$ is the transverse-field mixer and $\hat{C} = (C - \bar{C})/\sigma_C$ is the cost operator shifted and scaled to zero mean and unit variance over the computational basis, which keeps the useful range of γ of order one for instances whose raw coefficients span four decades. The scaling is a relabelling of γ and does not change the ansatz. Performance is reported through the normalized gain

(9)

which is zero for uniform sampling and one at the ground state and, unlike the conventional approximation ratio, is well defined when C takes both signs. The probability P_{feas} assigned to the feasible subspace of the constrained problem and the ground-state probability P_{opt} are reported alongside.

D. Noise Models and Depth Selection

Two noise models bracket the behaviour of a gate-model device. The global depolarizing channel applies $\rho \to \lambda\rho + (1 - \lambda)I/2^n$ after every cost and mixer layer, with survival probability λ. The local depolarizing channel applies, after every layer and independently on each qubit, a uniformly random Pauli with probability ε; it is simulated exactly in expectation by averaging pure-state trajectories in which Pauli errors are sampled, a standard unravelling of a Pauli channel. For the dense circuits of Section VI-H one layer corresponds to several hundred two-qubit gates, so ε should be read as the accumulated per-qubit error of a layer, not as a gate error.

Proposition 1. Under the global depolarizing channel with survival λ per layer, the depth-p noisy energy satisfies $F_\lambda(\gamma,\beta) = \lambda^{2p}F(\gamma,\beta) + (1 - \lambda^{2p})\bar{C}$. Consequently (i) the minimizers of F_λ coincide with those of F for every $\lambda \in (0,1]$; (ii) the noisy gain is $r_\lambda(p) = \lambda^{2p}r(p)$; and (iii) increasing the depth from p to p+1 improves the noisy gain if and only if $r(p+1)/r(p) > \lambda^{-2}$.

Proof. The channel commutes with every unitary, since $U(\lambda\rho + (1-\lambda)I/2^n)U^\dagger = \lambda U\rho U^\dagger + (1-\lambda)I/2^n$. Applying it $2p$ times therefore yields $\rho_\lambda = \lambda^{2p}|\psi\rangle\langle\psi| + (1 - \lambda^{2p})I/2^n$, and $\mathrm{Tr}(C\rho_\lambda)$ gives the stated form because $\mathrm{Tr}(C)/2^n = \bar{C}$. Items (i)–(iii) follow by substituting into (9), the factor λ^{2p} being positive and independent of the angles. □

Two consequences follow. Angles trained on a noiseless simulator transfer without loss to a device whose dominant noise is close to global depolarizing, so no device-side training is needed; and depth should be increased only while the noiseless gain grows faster than λ^{-2} per layer, a check that needs only the noiseless sweep and an estimate of λ from layer benchmarking. Fig. 1(b) applies the rule with an effective λ_{eff} that, for local noise, is taken as $(1 - \varepsilon)^n$ and treated as a heuristic whose accuracy is examined in Section VI-E.

E. Noise-Aware Re-optimization under Local Noise

Local depolarizing noise does not commute with the circuit, so the noise-invariance of Proposition 1 is not guaranteed and the noisy optimum may shift. The method re-optimizes the transferred angles directly on the noisy objective, estimated as a trajectory average with a fixed random seed,

(10)

where $\psi_k(\omega_k)$ is the state of trajectory k at angles (γ,β) and ω_k its Pauli-error pattern. Holding the K patterns fixed across angle evaluations (common random numbers) makes the surrogate a smooth deterministic function of the angles, which a derivative-free optimizer can descend with a small trust region; evaluation is then repeated with fresh patterns and a larger K. The gain attributable to noise awareness is isolated by also evaluating the re-optimized angles on the noiseless objective: if the noiseless gain also rises, the improvement came from continued optimization rather than from noise adaptation. This control is applied throughout.

F. Feasibility-Filtered Decoder

Each shot is a bitstring; shots violating the SoC bounds, the duration or the terminal condition are discarded and the lowest J_0 among the survivors is returned, or the greedy schedule if none survives. Because the simulator yields the exact distribution $q(x)$, the expected value of this decoder for N shots is computed without sampling. Let $F_f(v) = \Sigma_{x\ feasible,\ J0(x)\le v}\ q(x)$; then

(11)

summed over the distinct feasible objective values, with J_{fb} the fallback cost. Because a strong fallback would flatter any sampler that returns nothing feasible, the headline decoder metric is the relative gap to the constrained optimum of $E[J_{best}]$ conditioned on at least one feasible shot, with the no-feasible probability $(1 - P_{feas})^N$ reported separately; this allows QAOA, uniform sampling and simulated annealing (for which the distribution is replaced by its empirical sample) to be compared at identical N.

G. Algorithm and Complexity

Algorithm 1 collects the steps. Classical cost is dominated by the simulator: one noiseless evaluation of depth p costs

$O(p\ n\ 2^n)$ time and $O(2^n)$ memory, and a local-noise evaluation multiplies the time by K. On hardware the circuit depth per layer after routing is the relevant cost (Section VI-H). Building Q requires $O(n^2)$ objective probes and the decoder is $O(N\ n)$. Enumeration, used only as ground truth, is $O(n\ 2^n)$.

Algorithm 1 Noise-aware QAOA scheduling

Input: hourly λ, MEF, AEF (benchmark dispatch); plant profiles; T, L; ρ, w_p; noise model (λ or ε); N, K.

1: aggregate signals to slots; set π_S; build J_0, calibrate R, set A = B = C = fR, μ in its window; form C = diag J(x).

2: p ← 1; (γ,β) ← ramp; minimize F(γ,β) by COBYLA from three starts.

3: repeat: (γ,β) ← INTERP(γ,β); minimize F at depth p+1; compute r(p+1)/r(p); if ratio > λ_eff^{−2} then p ← p+1 else stop.

4: if noise is local: minimize $\hat{F}_\varepsilon$ of (10) from the transferred angles with fixed seed; verify with fresh seed and larger K.

5: sample N shots; discard infeasible; return the best feasible schedule or the greedy fallback; report E[J_best] by (11).

V. Simulation Setup

All results were produced in Python 3.11 with NumPy 2.5, SciPy 1.18 and Qiskit 2.6.1 on a single CPU container (no GPU); the exact QAOA simulator is a purpose-written NumPy statevector code whose agreement with the Qiskit Statevector backend on the same circuit is reported in Section VI-H. Qiskit is also used to transpile the depth-one circuits onto a 19-qubit heavy-hex coupling map. A MATLAB Optimization Toolbox script that solves the exported Q with the toolbox QUBO solver and with a linearized intlinprog model is included in the bundle for independent cross-checking; it was not executed in the reported build and no result in this paper depends on it. The global random seed is 2026 for every experiment.

Three instance sizes are used, all with one deferrable load ($L = 1$, $P_L = 300$ kW, $D = 2$ slots): $T = 4$, 5 and 6 slots giving $n = 12$, 15 and 18 qubits. The BESS has $E_B = 1000$ kWh, initial SoC 50 %, bounds 5–90 %, $\eta_c = \eta_d = 0.95$ and a rated power $P_B = 0.4E_B/\Delta$ so that one slot moves 40 % of the capacity irrespective of T; a second consecutive charge or discharge therefore violates the SoC bounds and the feasible fraction of the search space is small (90, 310 and 945 feasible bitstrings for $n = 12, 15, 18$). Table II lists the remaining parameters. The noiseless depth sweep runs to $p = 6$; global noise is evaluated in closed form for $\lambda \in \{0.995, 0.99, 0.98, 0.95, 0.90\}$; local noise is simulated for $\varepsilon \in \{0.002, 0.005, 0.01, 0.02\}$ with $K = 48$ trajectories during re-optimization and $K = 400$ for evaluation ($K = 32$ and 160 for $n = 18$, where only $\varepsilon \in \{0.005, 0.01\}$ and $p \le 3$ were affordable).

TABLE II
SIMULATION PARAMETERS

Group	Parameter	Value
IEEE 14-bus	Generators (bus: Pmax MW; a, b)	1: 332.4; 0.0430, 20 • 2: 140; 0.25, 20 • 3, 6, 8: 100; 0.01, 40
	Assumed fuel / εi (kgCO2/kWh)	G1 coal 0.95; G2 CCGT 0.40; G3 OCGT 0.58; G6 oil 0.75; G8 biomass 0.05
	System load; renewables	259 MW × 1.25 peak × daily shape; PV 150 MW, wind 60 MW
	Tariff adder δh (USD/MWh)	8 (00–05), 15 (06–16, 21–23), 30 (17–20)
Plant	Base load; PV; wind (kW)	420–600 industrial shape; 450 kWp; 160 kW
	BESS	1000 kWh; P_B = 0.4E_B/Δ; SoC 5–90 %; η = 0.95
	Deferrable load	300 kW, 2 slots; default: first two slots
Objective	ρ; w_p; π_S	80 USD/t (sweep 0–600); 4 USD/MW²; 0.099 USD/kWh (n = 18)
QUBO	A, B, C; μ	0.5R (R = 779 USD, n = 18); 12 USD
QAOA	Optimizer; init; depth	COBYLA, 3 starts, ramp + INTERP, refined to $\Delta r < 10^{-4}$; p = 1–6
Baselines	SA; uniform; greedy; relaxation	single-flip Metropolis, 50–1000 sweeps, N ∈ {100, 1000}; same N; rule-based; continuous battery QP

VI. Results and Discussion

A. Benchmark Signals

Fig. 2(a) shows the hourly dispatch. The coal unit G1 carries the base and, together with the CCGT unit, is on the margin for 23 of 24 hours; the three 40 USD/MWh units enter only at the evening peak. The marginal emission factor is therefore 0.869 $kgCO_2$/kWh for most of the day and drops to 0.494 $kgCO_2$/kWh only in the hour where the oil, OCGT and biomass units share the margin. This is not an artefact: it is what the case14 cost ordering implies, and it is the typical situation in coal-first systems where the cheapest energy is also the dirtiest at the margin. The average emission factor of delivered energy (Fig. 2(b)) varies between 0.41 and 0.76 $kgCO_2$/kWh because the embedded PV displaces thermal output at midday. The tariff peaks at 70 USD/MWh in the evening and is lowest, 37 USD/MWh, at night. With the AEF signal the cheapest slot (night) is not the cleanest slot (midday), so the trade-off is real; with the MEF signal it degenerates (Section VI-G). Fig. 2(c) shows the plant profiles and their 4-h slot means.

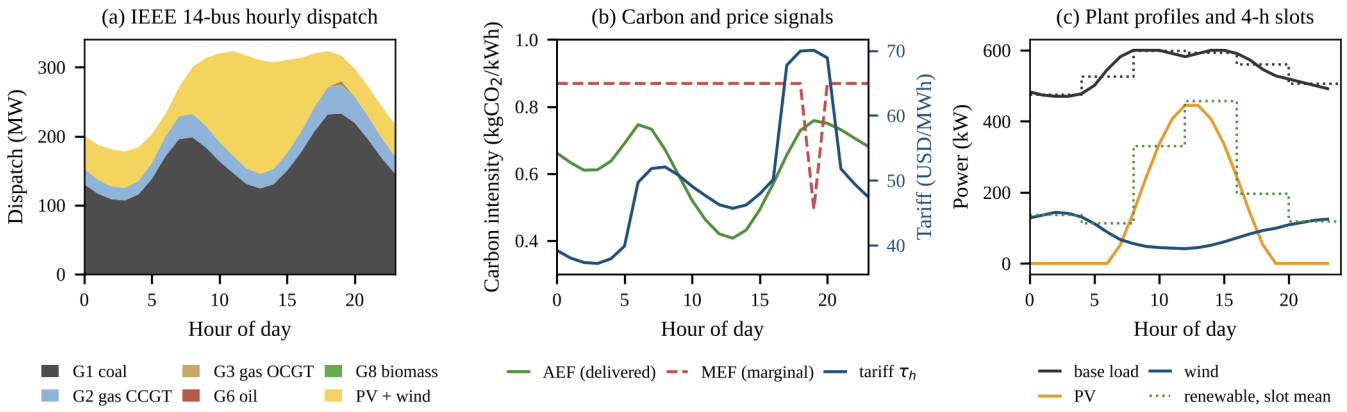


Fig. 2. IEEE 14-bus derived signals. (a) Hourly economic dispatch of the five case14 generators with embedded PV and wind. (b) Marginal (MEF) and average (AEF) emission factors with the retail tariff. (c) Plant base load, PV and wind and the 4-h slot means used for n = 18.

B. QUBO Landscape and Penalty Window

Tables III and IV report the two calibration sweeps. With μ = 0 the QUBO argmin violates the SoC bounds for $n = 15$ and 18 (two consecutive charges around midday are cheaper than one), so a decoder is indispensable; for μ between 10 and 15 USD the argmin is feasible for all three sizes with an objective within 0.08 % of the constrained optimum, and for μ ≥ 30 USD the regularizer begins to suppress storage use and the gap grows to 0.97 %. The window exists because the regularizer must exceed the economic gain of a double excursion but remain below the per-slot arbitrage value of a single one. For the hard penalties, $f \geq 0.35$ yields a constraint-satisfying optimum in every instance and $f = 0.5$ is used. The expectation that oversized penalties would flatten the landscape QAOA has to navigate [7] was not borne out within the range tested: at $n = 12$ and depth three the normalized gain is 0.760 at $f = 0.5$ and 0.786 at $f = 1.5$, the feasible mass rises from 0.271 to 0.322, and the decoder gap at 1000 shots is below 0.01 % in both cases. Below sufficiency the picture changes: at $f = 0.1$ only 0.005 of the mass is feasible and the decoder gap is 0.59 %. The operating point was therefore chosen for sufficiency with margin, not for QAOA performance, and the sweep is reported because it contradicts a common assumption.

TABLE III
CALIBRATION OF THE SOC REGULARIZER μ (GAP OF J_0 AT THE QUBO ARGMIN TO THE CONSTRAINED OPTIMUM, %; "INF": INFEASIBLE ARGMIN)

μ (USD)	n = 12	n = 15	n = 18
0	+0.00	inf -0.06	inf -0.44
5	+0.02	+0.00	inf -0.44
10	+0.02	+0.00	+0.08
12	+0.02	+0.00	+0.08
15	+0.02	+0.00	+0.08
20	+0.02	+0.18	+0.08
30	+0.02	+0.97	+0.08
40	+0.02	+0.97	+0.08

Negative gaps occur only for infeasible argmins, which exploit the violated SoC bound.

TABLE IV
CALIBRATION OF THE PENALTY FACTOR F (SAME GAP MEASURE) AND THE DEPTH-3 QAOA GAIN AT N = 12

f	n = 12	n = 15	n = 18	r (p = 3, n = 12)
0.1	inf -14.97	inf -12.64	inf -10.91	0.774
0.2	inf -14.97	+0.00	+0.08	0.766
0.35	+0.02	+0.00	+0.08	0.755
0.5	+0.02	+0.00	+0.08	0.760
1	+0.02	+0.00	+0.08	0.784
1.5	+0.02	+0.00	+0.08	0.786

C. Noiseless Depth Sweep

Fig. 3(a) gives the refined noiseless results. The normalized gain rises monotonically with depth and is remarkably size-independent, from 0.55–0.60 at $p = 1$ to 0.91–0.93 at $p = 6$. The feasible-subspace probability is the more telling quantity: uniform sampling places 0.4 % of its mass on feasible schedules at $n = 18$, whereas QAOA raises this to 2.5 % at $p = 1$ and 34.1 % at $p = 6$. The ground-state probability stays below 0.96 %, as expected for a continuous-valued objective whose excited states differ from the optimum by fractions of a dollar; the decoder, not P_{opt}, is therefore the right lens. With the default COBYLA budget the gain was under-estimated by up to 0.043, and all values are after refinement to an improvement below 10^{-4} per pass.

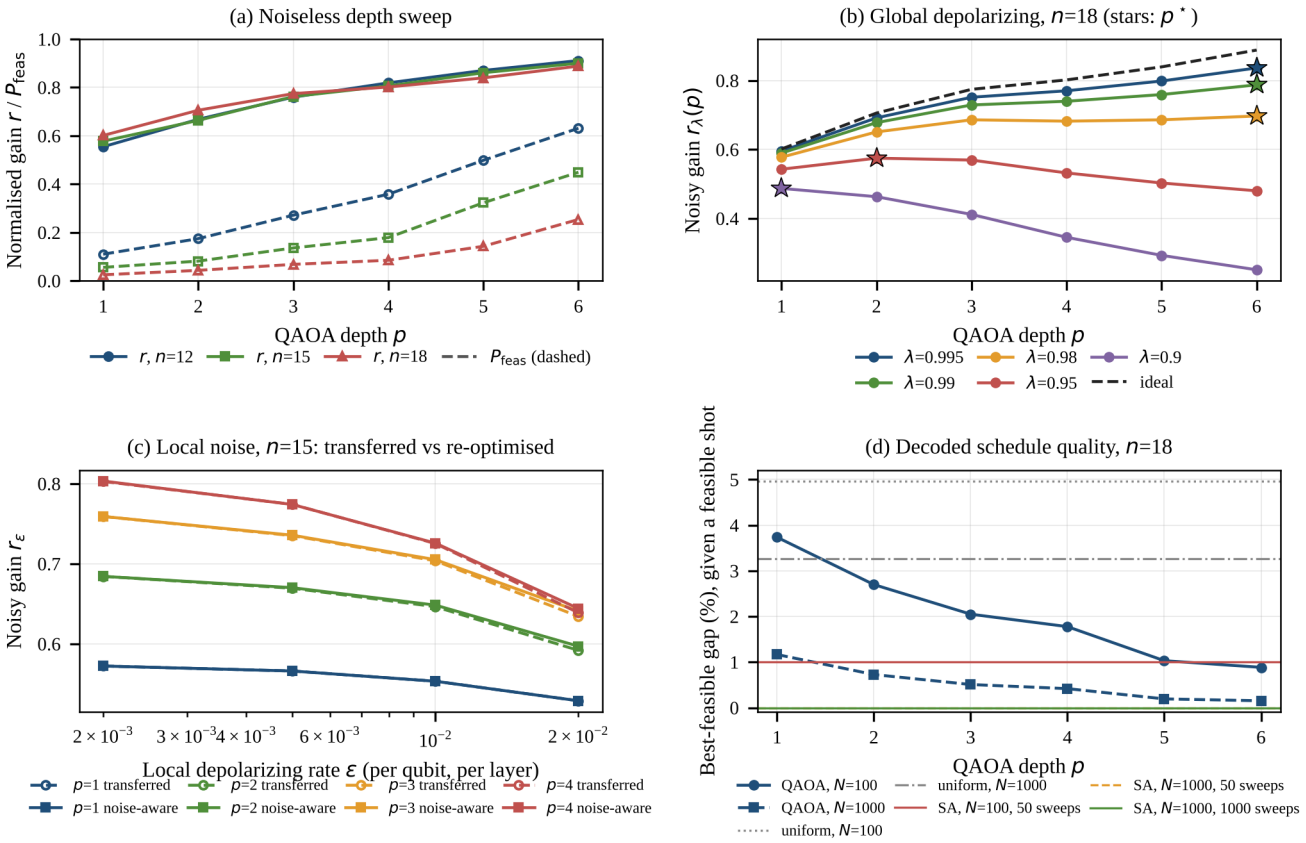


Fig. 3. QAOA performance. (a) Noiseless normalized gain r (solid) and feasible-subspace probability (dashed) versus depth for n = 12, 15, 18. (b) Gain under global depolarizing noise for n = 18 from Proposition 1; stars mark the optimal depth p*(λ). (c) Gain under local depolarizing noise for n = 15 with transferred noiseless angles (open markers, dashed) and after noise-aware re-optimization (filled). (d) Expected best-feasible gap of the decoder at n = 18 for N = 100 and 1000 shots against uniform sampling and simulated annealing at the same N (the 50-sweep and 1000-sweep SA lines at N = 1000 coincide at zero).

D. Global Depolarizing Noise and the Depth Rule

Fig. 3(b) applies Proposition 1 to the $n = 18$ sweep. The factor λ^{2p} turns the monotone noiseless curve into one with an interior maximum whenever the per-layer gain ratio falls below λ^{-2}: the optimal depth is 6 for $\lambda = 0.995$, 6 for $\lambda = 0.99$, 5 for $\lambda = 0.98$, 2 for $\lambda = 0.95$ and 1 for $\lambda = 0.90$, with maximal noisy gains of 0.86, 0.81, 0.72, 0.58 and 0.49. A layer that survives with probability 0.95 already makes depth beyond 2 counter-productive. The invariance claim was tested directly at $n = 12$, $p = 3$: optimizing the noisy objective at $\lambda = 0.95$ and at $\lambda = 0.90$ from the same four random starts returned the same angle set, whose noiseless gain is 0.729 and 0.729 respectively, identical to four decimals; the survival probability therefore moved neither the optimizer nor the basin reached. That this value lies below the best multi-start noiseless gain of 0.771 at the same depth is a property of the landscape, which has several basins of comparable quality, not of the noise, and it is reported because a single-start comparison would have hidden it.

E. Local Depolarizing Noise and Re-optimization

Fig. 3(c) and Table V report the local-noise study. Transferred angles lose gain in proportion to the error rate, and the loss per unit ε grows with depth, but depth still pays at every rate tested: at $\varepsilon = 0.02$ the $n = 15$ gain rises from 0.529 at $p = 1$ to 0.639 at $p = 4$, whereas the same angles give 0.826 without noise. The depth rule with the proxy $\lambda_{eff} = (1 - \varepsilon)^n$ is therefore far too pessimistic: it predicts $p\bigstar = 2$ at $\varepsilon = 0.005$ and $p\bigstar = 1$ at $\varepsilon \geq 0.01$, whereas the measured optimum is at the largest depth tested ($p = 4$) for every ε. The proxy treats every single-qubit error as a total loss of the state, as a global channel would, whereas a Pauli error on one qubit leaves most of the amplitude structure intact; the rule is thus exact for global noise and a conservative lower bound on the worthwhile depth for local noise. Re-optimization on the CRN surrogate never improves the noisy gain by more than a few hundredths, and at $n = 18$ the change is within the trajectory-evaluation noise (±0.002 at $K = 160$); once optimizer convergence is controlled the noise-specific component is at most 0.029 ($n = 12$, $\varepsilon = 0.02$, $p = 4$), below 0.009 in 90 % of the configurations, and the angle shift has a median of 0.054 rad over all angles. Before the convergence pass the same experiment had shown apparent gains of up to 0.03 at $\varepsilon = 0.002$ that vanished once the noiseless angles were refined, an artefact an uncontrolled comparison would have published as a method. In decoder terms the distinction matters little: at $N = 1000$ the best-feasible gap under local noise stays below 0.61 % for $n \leq 15$ at every ε tested, because the decoder needs feasible mass rather than a sharp peak and local noise erodes that mass slowly (P_{feas} falls from 0.230 to 0.119 between $\varepsilon = 0.002$ and 0.02 at $n = 15$, $p = 4$). The reduced grid at $n = 18$ shows the same ordering (Table V).

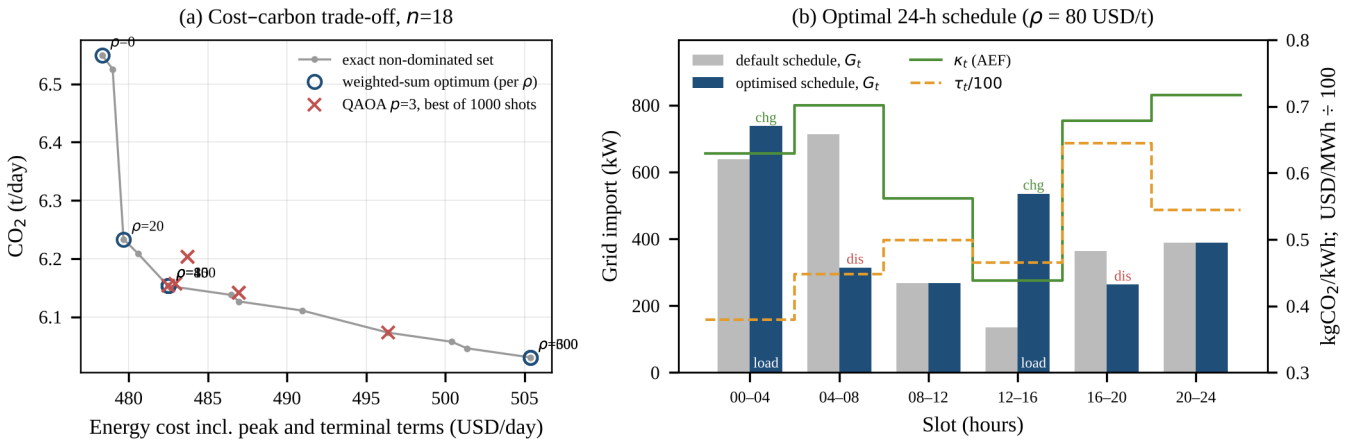


Fig. 4. Multiobjective results at n = 18. (a) Exact non-dominated set of feasible schedules, weighted-sum optima for seven carbon prices and the best feasible QAOA (p = 3) sample of 1000 shots at each price. (b) Grid import of the default and the optimal schedule (ρ = 80 USD/t, AEF signal) with battery and load actions annotated; the AEF and scaled tariff are overlaid.

TABLE V
LOCAL DEPOLARIZING NOISE: GAIN WITH TRANSFERRED (TR.) AND NOISE-AWARE RE-OPTIMIZED (RE.) ANGLES, NOISE-SPECIFIC PART OF THE IMPROVEMENT (ΔNS) AND DECODER GAP AT N = 1000

n	ε	p	r tr.	r re.	Δns	P_feas re.	gap (%)
12	0.005		0.777	0.778	0.002	0.408	0.00
12	0.02		0.614	0.629	0.029	0.232	0.00
15	0.002		0.572	0.572	0.000	0.054	0.52
15	0.002		0.803	0.803	0.000	0.235	0.12
15	0.005		0.566	0.566	0.000	0.053	0.53
15	0.005		0.774	0.774	0.002	0.203	0.14
15	0.01		0.553	0.553	0.000	0.050	0.56
15	0.01		0.725	0.726	0.001	0.176	0.17
15	0.02		0.529	0.529	0.000	0.045	0.61
15	0.02		0.639	0.644	0.014	0.119	0.26
18	0.005		0.595	0.595	0.000	0.024	1.19
18	0.005		0.698	0.698	0.001	0.046	0.79
18	0.005		0.763	0.762	0.001	0.072	0.58
18	0.01		0.585	0.585	0.000	0.023	1.22
18	0.01		0.679	0.679	0.000	0.043	0.83
18	0.01		0.735	0.736	0.004	0.063	0.64

F. Decoded Schedule Quality against Classical Sampling

Fig. 3(d) and Table VI settle the equal-budget comparison at n = 18. With N = 100 shots the probability of drawing no feasible schedule is 70 % for uniform sampling and 8 % for QAOA at p = 1, falling below 1 % from p = 2; given a feasible shot, the QAOA decoder's expected gap falls from 3.74 % at p = 1 to 0.88 % at p = 6, against 4.97 % for uniform sampling; with N = 1000 the figures are 1.17 % and 0.15 % against 3.26 %. Simulated annealing with 1000 reads finds the optimum exactly already at 50 sweeps (0.04 s on the same CPU) and at 1000 sweeps (0.68 s); with 100 reads and 50 sweeps its gap is 1.00 %. The greedy rule is 0.08 % above the optimum and the default schedule 3.63 %. The continuous-battery relaxation is 0.16 % below the binary optimum, which is what the two-bit battery encoding costs: small, but about twice the loss of the greedy rule. At 18 qubits QAOA is thus a structured sampler that beats uniform sampling clearly and approaches the optimum with a few thousand shots, while a few milliseconds of single-flip annealing on the same QUBO is better still. The economic stakes of the schedule itself are modest at this plant scale: the optimum lowers the daily bill from 481.4 to 473.5 USD and emissions from 6544 to 6102 $kgCO_2$ relative to the default schedule, a 6.8 % emission reduction obtained mostly by moving the deferrable load into the midday PV window and cycling the battery once from midday into the evening peak (Fig. 4(b)).

TABLE VI
DECODER COMPARISON AT N = 18: EXPECTED GAP OF THE BEST FEASIBLE SCHEDULE TO THE CONSTRAINED OPTIMUM

Method	N = 100	N = 1000	P_feas	Wall time
Uniform sampling	4.97 %	3.26 %	0.004	–
QAOA p = 1	3.74 %	1.17 %	0.025	126 s train
QAOA p = 3	2.05 %	0.51 %	0.086	1080 s train
QAOA p = 6	0.88 %	0.15 %	0.341	4730 s train
SA, 50 sweeps	1.00 %	0.00 %	0.474	0.04 s
SA, 1000 sweeps	0.61 %	0.00 %	0.477	0.68 s
Greedy rule	0.08 %		1	< 1 ms
Exhaustive (2^{18})	0 %			0.50 s

Wall times are for the Python implementation on the simulation CPU; QAOA training time is the cumulative statevector cost up to the stated depth and is not a hardware estimate.

G. Cost–Carbon Trade-off and the Choice of Carbon Signal

Fig. 4(a) traces the trade-off at n = 18 by sweeping the carbon price. The exact non-dominated set of the feasible schedules contains 12 points; the weighted-sum optima for ρ ∈ {0, 20, 40, 80, 150, 300, 600} USD/t recover 4 distinct supported points, which is the known limitation of scalarized multiobjective QUBOs: non-convex parts of the front are unreachable whatever the solver. QAOA at p = 3 with 1000 shots did not return the weighted-sum optimum at any of the seven prices; its best feasible sample lay between 0.16 % and 1.30 % above it in objective, consistent with the depth-3 decoder gap of Table VI, and the sampled points sit visibly inside the front rather than on it. Raising ρ from 0 to 600 USD/t reduces emissions from 6550 to 6030 $kgCO_2$ at a cost increase from 478.3 to 505.4 USD per day; the implied marginal abatement cost along the front rises from 4 to 186 USD/t.

The signal choice is consequential. Under the AEF signal the optimal schedule runs the deferrable load in the 00–04 h and 12–16 h slots and attributes 6102 $kgCO_2$ to the plant; under the MEF signal it runs the load in the 00–04 h and 04–08 h slots and attributes 8612 $kgCO_2$. Evaluated on the respective other signal, the AEF-optimal schedule emits 8612 $kgCO_2$ of marginal emissions and the MEF-optimal schedule 6524 $kgCO_2$ of average emissions. Because the MEF is nearly flat, the MEF-optimal schedule coincides with the cost-only ($\rho = 0$) schedule; a carbon-aware scheduler driven by a marginal signal in a coal-first system therefore has almost nothing to optimize, and the midday shift that the AEF schedule performs is, from a marginal perspective, a transfer of the same coal-fired MWh to another hour. Which accounting is right depends on whether the plant's reduction is meant to be attributional or consequential [22]; the model makes the difference explicit instead of choosing for the user.

H. Simulator Validation and Hardware Relevance

For each size the QAOA circuit built in Qiskit from the Ising form of Q was evaluated with the Statevector class at random angles and compared with the NumPy simulator: the energies agree to better than 10^{-10} USD and the state fidelities to $1 - 10^{-14}$. Transpiled onto the 19-qubit heavy-hex coupling map with the {rz, sx, x, ecr} basis at optimization level 3, one QAOA layer requires 171, 278 and 410 two-qubit gates and has depth 295, 466 and 626 for $n = 12$, 15 and 18, because the duration and neutrality penalties couple every pair of load bits and every pair of battery bits. At a two-qubit error of 5×10^{-3} the per-layer survival $(1 - 0.005)^{\text{gates}}$ would be about 0.42 at $n = 12$ and 0.13 at $n = 18$, far below the regime in which Fig. 3(b) still shows a gain at $p = 1$; the depth rule would then recommend $p = 1$ and the sampler would be close to uniform. This is the central negative result of the study: for this problem family the noise analysis is not a refinement but the binding constraint, and penalty-induced all-to-all couplings are the reason.

VII. Comparative Analysis and Key Outcomes

Table VII positions the work against the studies of Table I along the axes a practitioner would apply. The present method is the only one in the set that reports feasibility rates, an exact equal-budget comparison with a classical sampler and a noise analysis tied to the circuit cost of the actual instances. Its scalability is bounded by exact simulation and its reliability on hardware is predicted to be poor at today's two-qubit error rates; both statements are reported rather than softened.

TABLE VII
QUALITATIVE COMPARISON WITH EXISTING WORK

Criterion	Unit-commitment QAOA [7], [8]	Microgrid QAOA [10]	Noisy-QAOA studies [15]–[17]	This work
Accuracy vs exact optimum	5.1 % at p = 1 [8]	not reported	n/a	0.15 % (decoder, N = 1000, p = 6)
Constraint handling	penalties	penalties	n/a	penalties + cycle-count neutrality + regularizer window + feasibility filter
Noise	none	none	generic Max-Cut	global (closed form) and local (trajectories) on the application
Equal-budget classical baseline	no	no	no	SA and uniform at same N, exact expectation
Benchmark grounding	test units	field data	none	IEEE 14-bus dispatch, MEF and AEF
Scalability (this build)	small	small	small	$n \le 18$ exact; hardware cost quantified
Sustainability metric	cost	cost, efficiency	none	priced carbon, Pareto front, signal sensitivity

In engineering terms: the QUBO formulation is faithful, since for μ in its window and $f \ge 0.35$ its argmin is the constrained optimum to within 0.08 %. Global noise needs no re-training and fixes the depth by a rule; local-noise re-training is worth at most a few hundredths of normalized gain. The practical value lies less in a quantum advantage, which is not demonstrated, than in a benchmarked, reproducible yardstick for future hardware runs and in the signal-sensitivity finding, which applies to any scheduler handed a grid carbon-intensity feed.

VIII. Limitations

Everything reported is simulation. No hardware execution was performed, and the noise models omit coherent errors, readout error and crosstalk; the per-layer ε is an abstraction of several hundred gates. Exact statevector simulation limits the instances to 18 qubits and one deferrable load, which is also why classical solvers dominate; the annealing comparison is at equal sample count, not equal energy or hardware wall time. The IEEE 14-bus link is to generator cost and capacity data only; no power flow or line limits are modelled, and the fuel and emission assignments are assumptions. Net metering, the quadratic peak surrogate and deterministic forecasts are simplifications; the weighted-sum QUBO cannot reach non-supported Pareto points; the MATLAB script is untested; and the re-optimization result depends on COBYLA and on $K = 48$ trajectories, so a gradient-based optimizer might extract slightly more, subject to the same control.

IX. Conclusion and Future Scope

A carbon-aware scheduling model for renewable generation, battery storage and industrial loads was cast as a

multiobjective QUBO grounded in IEEE 14-bus dispatch signals and solved with QAOA under explicit noise models. The formulation's cycle-count neutrality and SoC regularizer window keep the penalized optimum feasible and within 0.08 % of the constrained optimum; Proposition 1 shows that global depolarizing noise leaves the QAOA optimizers untouched and reduces depth selection to a gain-ratio rule, confirmed numerically; local-noise re-optimization yields a noise-specific gain of at most 0.029 once optimizer convergence is controlled; and the feasibility-filtered decoder reaches the constrained optimum to within 0.15 % with 1000 shots at depth six, concentrating up to 68 % of the QAOA mass on feasible schedules. Simulated annealing at equal budget finds the optimum exactly, and the transpiled circuit cost indicates that current gate errors would confine QAOA to depth one on these instances. Three directions follow. Constraint-preserving mixers (XY or Grover-type) would remove the penalty couplings that make the circuits dense; a hardware run on a heavy-hex device at $n = 12$ would test Proposition 1 against real noise and calibrate λ; and a rolling-horizon deployment with stochastic forecasts, an asymmetric export tariff and a true demand charge would turn the benchmark into an operational tool. Independently of the solver, the signal-sensitivity result argues for publishing marginal and average intensity side by side wherever carbon-aware demand response is proposed.

ACKNOWLEDGMENT

The authors thank the Research and Development Division of Tekstein Solutions Private Limited for computational resources. Author designations and the final author order are to be confirmed before submission.

REFERENCES

[1] A. Radovanović et al., “Carbon-aware computing for datacenters,” IEEE Trans. Power Syst., vol. 38, no. 2, pp. 1270–1280, Mar. 2023.

[2] J. Lindberg, Y. Abdennadher, J. Chen, B. Lesieutre, and L. Roald, “A guide to reducing carbon emissions through data center geographical load shifting,” in Proc. ACM e-Energy, 2021, pp. 430–436.

[3] Y. Zhou, Z. Tang, N. Nikmehr, P. Babahajiani, F. Feng, T.-C. Wei, H. Zheng, and P. Zhang, “Quantum computing in power systems,” iEnergy, vol. 1, no. 2, pp. 170–187, 2022.

[4] A. Ajagekar and F. You, “Quantum computing for energy systems optimization: Challenges and opportunities,” Energy, vol. 179, pp. 76–89, 2019.

[5] A. Giani and Z. Eldredge, “Quantum computing opportunities in renewable energy,” SN Comput. Sci., vol. 2, art. 393, 2021.

[6] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” arXiv:1411.4028, 2014.

[7] S. Koretsky, P. Gokhale, J. M. Baker, J. Viszlai, H. Zheng, N. Gurung, R. Burg, E. A. Paaso, A. Khodaei, R. Eskandarpour, and F. T. Chong, “Adapting quantum approximation optimization algorithm (QAOA) for unit commitment,” in Proc. IEEE Int. Conf. Quantum Comput. Eng. (QCE), 2021, pp. 181–187.

[8] B. Salgado, A. Sequeira, and L. P. Santos, “A hybrid classical-quantum approach to highly constrained unit commitment problems,” arXiv:2412.11312, 2024.

[9] Universidad Europea (authors to be verified), “Modeling energy community operations as a QUBO problem for quantum optimization,” 2024 (venue and authors to be verified before submission).

[10] M. Liu, M. Liao, R. Zhang, X. Yuan, Z. Zhu, and Z. Wu, “Quantum computing as a catalyst for microgrid management: Enhancing decentralized energy systems through innovative computational techniques,” Sustainability, vol. 17, no. 8, art. 3662, 2025.

[11] J. Preskill, “Quantum computing in the NISQ era and beyond,” Quantum, vol. 2, p. 79, 2018.

[12] P. D. Nation, H. Paik, A. Edwards, S. Sheldon, and B. Johnson, “Benchmarking the performance of quantum computing software,” arXiv:2409.08844, 2024.

[13] N. Nikmehr, P. Zhang, and M. A. Bragin, “Quantum distributed unit commitment: An application in microgrids,” IEEE Trans. Power Syst., vol. 37, no. 5, pp. 3592–3603, Sep. 2022.

[14] P. Halffmann, S. Lenk, and M. Trebing, “Harnessing inferior solutions for superior outcomes: Obtaining robust solutions from quantum algorithms,” in Proc. GECCO Companion, 2024, doi: 10.1145/3638530.3664160.

[15] C. Xue, Z.-Y. Chen, Y.-C. Wu, and G.-P. Guo, “Effects of quantum noise on quantum approximate optimization algorithm,” Chin. Phys. Lett., vol. 38, no. 3, art. 030302, 2021.

[16] J. Marshall, F. Wudarski, S. Hadfield, and T. Hogg, “Characterizing local noise in QAOA circuits,” IOP SciNotes, vol. 1, no. 2, art. 025208, 2020.

[17] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, “Noise-induced barren plateaus in variational quantum algorithms,” Nat. Commun., vol. 12, art. 6961, 2021.

[18] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, “Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices,” Phys. Rev. X, vol. 10, art. 021067, 2020.

[19] S. H. Sack and M. Serbyn, “Quantum annealing initialization of the quantum approximate optimization algorithm,” Quantum, vol. 5, p. 491, 2021.

[20] F. G. S. L. Brandão, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, “For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances,” arXiv:1812.04170, 2018.

[21] K. J. Dzahini, J. M. Larson, M. Menickelly, and S. M. Wild, “A noise-aware scalable subspace classical optimizer for the quantum approximate optimization algorithm,” arXiv:2507.10992, 2025.

[22] D. S. Callaway, M. Fowlie, and G. McCormick, “Location, location, location: The variable value of renewable energy and demand-side efficiency resources,” J. Assoc. Environ. Resour. Econ., vol. 5, no. 1, pp. 39–75, 2018.

[23] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, “MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education,” IEEE Trans. Power Syst., vol. 26, no. 1, pp. 12–19, Feb. 2011.

[24] A. Lucas, “Ising formulations of many NP problems,” Front. Phys., vol. 2, art. 5, 2014.

[25] F. Glover, G. Kochenberger, and Y. Du, “Quantum bridge analytics I: A tutorial on formulating and using QUBO models,” 4OR, vol. 17, pp. 335–371, 2019.

[26] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, “A review on quantum approximate optimization algorithm and its variants,” Phys. Rep., vol. 1068, pp. 1–66, 2024.

[27] A. Javadi-Abhari et al., “Quantum computing with Qiskit,” arXiv:2405.08810, 2024.

[28] S. Hadfield, Z. Wang, B. O'Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz,” Algorithms, vol. 12, no. 2, art. 34, 2019.