Highlights

QuantumFerryV2B: A State-of-the-Art Quantum–AI Framework for Weather-Robust, Thermally Aware Ferry-Port Vehicle-to-Building Energy Control

Yifan Wang

- A SOTA three-region ferry-port V2B benchmark fuses real ferry schedules/AIS, weather, sea state, port load, and battery data across Norway, Denmark, and the USA.
- One physics evaluator co-schedules ferry charging, V2B, reefer load, and cold-battery preconditioning, scoring every method on identical footing.
- A SciML P90 weather reserve converts brittle deterministic schedules into 98–100 % held-out cold-wave feasibility at under 2.5 % cost premium.
- A QPU-ready QUBO scheduler and a physics-shielded VQC controller form a compact, hardware-portable, NISQ-deployable Quantum–AI decision layer.
- Coordinated control cuts peak demand by 52.5–56.1 % and operating cost by 33.0–38.5 % across all real-data-anchored cases.

QuantumFerryV2B: A State-of-the-Art Quantum–AI Framework for Weather-Robust, Thermally Aware Ferry-Port Vehicle-to-Building Energy Control

Yifan Wang[a,*]

[a]*Department of Mechanical Engineering, McGill University, Montreal, QC, H3A 2T7, Canada*

Abstract

Battery-electric ferry ports are emerging as the most demanding nodes of the maritime energy transition: a single shore transformer must deliver megawatt-scale charging within minute-long berth windows while simultaneously supplying terminal buildings, refrigerated cargo, and electrified port equipment, and while cold weather inflates crossing energy and suppresses lithium-ion charge acceptance. We introduce *QuantumFerryV2B*, a state-of-the-art (SOTA) Quantum–AI framework that controls this coupled system as one weather-robust, thermally aware Vehicle-to-Building (V2B) problem rather than as isolated charging, building, or battery subproblems. The framework unifies five components: a reproducible three-region real-data benchmark, a physics-based energy and battery-thermal evaluator that scores every method identically, a chance-constrained P90 weather reserve driven by a scientific machine-learning (SciML) crossing-energy surrogate, a QPU-ready quadratic unconstrained binary optimisation (QUBO) day-ahead scheduler with deterministic feasibility repair, and a physics-shielded variational quantum circuit (VQC) closed-loop controller.

The benchmark fuses Norwegian Entur ferry timetables, Danish Maritime Authority AIS traces, Washington State Ferries GTFS schedules, Open-Meteo historical weather and marine fields, NREL/OEDI port cargo-handling load profiles, CALCE low-temperature lithium-ion cell data, EnergyPlus terminal-load simulation, and OpenModelica battery-thermal

*Corresponding author.

Email address: yifan.wang18@mail.mcgill.ca (Yifan Wang)

validation. Across Norway, Denmark, and the USA, coordinated scheduling cuts peak demand by 52.5–56.1 % and operating cost by 33.0–38.5 % relative to uncontrolled opportunity charging. Removing thermal preconditioning renders the Norwegian and US cold-wave cases infeasible, with 2 and 6 missed departures; a SciML P90 reserve converts brittle point-optimal schedules (45–70 % held-out feasibility) into robust ones (98–100 %) at a cost premium below 2.5 %. The QPU-ready QUBO scheduler matches the exact optimiser on the demand peak while outperforming an equally fast greedy heuristic by up to 2.8×, and the physics-shielded VQC sustains zero deployed missed departures under IBM Heron-r2-class quantum-hardware noise while its parameter count grows 16× more slowly than a matched neural network. QuantumFerryV2B thus establishes a SOTA, real-data-grounded benchmark and a deployable Quantum–AI control architecture for robust ferry-port electrification.

Keywords: battery-electric ferry, vehicle-to-building, port microgrid, cold-battery preconditioning, quantum annealing, variational quantum circuit, weather-robust scheduling

---

1. Introduction

Maritime transport moves roughly 90 % of world trade by volume and contributes close to 3 % of global greenhouse-gas emissions, and short-sea ferry services are among the first segments to electrify at scale: Norway's *MF Ampere*, Denmark's *Ellen*, and a growing fleet of battery-electric ferries already operate zero-emission crossings on fixed public-transport timetables (Bouman et al., 2017; Frković et al., 2025). This success converts a decarbonisation opportunity into a demanding energy-systems control problem. A battery-electric ferry must absorb megawatt-scale energy during berth windows of only a few minutes, yet the same shore connection simultaneously feeds terminal buildings, refrigerated (reefer) cargo, and emerging electric cargo-handling equipment (Zhang et al., 2022; Mao et al., 2022). In cold weather the coupling tightens further: wind, waves, and cabin heating raise per-crossing energy, while lithium-ion cells accept fast charge only after preconditioning, because charging a cold pack drives lithium plating and irreversible damage (Jaguemont et al., 2016; Wang et al., 2022). When a controller underestimates this coupled, weather-dependent demand, the failure mode is not merely a higher electricity bill — it is a missed public-transport

sailing.

This setting is qualitatively different from land-based electric-vehicle charging. Vehicle-to-grid and Vehicle-to-Building studies have established the value of mobile batteries for peak shaving and energy arbitrage (Kempton and Tomic, 2005; Lund and Kempton, 2008; Leippi et al., 2025); maritime-energy research has shown that speed, weather, and operational profiles dominate vessel energy use (Psaraftis and Kontovas, 2013; Agand et al., 2023); battery science has quantified the low-temperature charge-acceptance limit (Jaguemont et al., 2016; Wang et al., 2022); and port-microgrid work has optimised berth allocation, onshore power, and integrated energy supply (Zhang et al., 2022; Mao et al., 2022; Gabrielii et al., 2025). Each stream is mature, and each, on its own, is incomplete for the ferry-port problem: ferry-port V2B operation sits precisely at their intersection, a coupled transportation–building–battery–grid control problem with *hard departure constraints*.

The research gap is therefore concrete and, to the best of our knowledge, unaddressed. No public dataset or modelling study evaluates ferry schedules or AIS, metocean forcing, port-electrification load, terminal-building simulation, battery thermal limits, and V2B actions *together*, across more than one region, with a single physics evaluator that scores every controller on identical footing. Equally, no prior control architecture simultaneously handles cold charge acceptance, sea-state-driven crossing energy, transformer stress, weather uncertainty, and hardware-portable discrete scheduling. A controller that minimises cost while ignoring these couplings looks attractive in simulation yet fails exactly during cold-wave operation, when robust dispatch matters most. Recent work has shown that quantum computing is an effective lens for precisely this class of large-scale, discrete, NP-hard energy-scheduling problems — Deng et al. demonstrated quantum-enhanced residential EV charging management with comparable solution quality and superior scalability over classical solvers (Deng et al., 2025), and variational quantum reinforcement learning has been deployed for building thermal control under realistic hardware noise (Chen, 2024; Skolik et al., 2022). These advances motivate a Quantum–AI treatment of the ferry-port problem.

We close this gap with *QuantumFerryV2B*, a SOTA Quantum–AI framework for weather-robust, thermally aware ferry-port V2B control. We use the term Quantum–AI in a precise architectural sense. The discrete day-ahead scheduler is cast as a quadratic unconstrained binary optimisation (QUBO) model that maps natively to quantum-annealing hardware, where the valley-

fill peak-shaving term and the cold-charge coupling are genuine quadratic interactions rather than penalty artefacts (Lucas, 2014). The closed-loop controller is a compact variational quantum circuit (VQC) policy projected through a physics-based energy- and battery-management safety shield (Bergholm et al., 2022; Jerbi et al., 2021; Schuld et al., 2021). This is not a cosmetic quantum wrapper: the QUBO captures the natural combinatorial structure of load shaping, and the VQC supplies a parameter-efficient, NISQ-deployable policy representation whose advantages widen with fleet size. Exact mixed-integer linear programming (MILP), rule-based, and neural-network baselines are retained throughout so that every comparison is rigorous and fair.

The contributions of this paper are:

1. **A multi-region, real-data-anchored ferry-port V2B benchmark** that integrates Norwegian Entur timetables, Danish AIS, Washington State Ferries GTFS, Open-Meteo weather and marine forcing, NREL/OEDI port-electrification loads, CALCE low-temperature battery data, EnergyPlus terminal-load simulation, and OpenModelica battery-thermal validation — the first dataset to span this intersection.
2. **A coupled, thermally aware energy model** that co-schedules ferry charging, V2B discharge, preconditioning, and shiftable port load under departure, state-of-charge (SOC), transformer, and cold charge-acceptance constraints, evaluated by one exact physics evaluator.
3. **A SOTA weather-robust planning layer** in which a SciML crossing-energy surrogate supplies a chance-constrained P90 reserve that restores held-out cold-wave feasibility to 98–100 % across all three regions.
4. **A Quantum–AI decision architecture** combining a QPU-ready QUBO scheduler with deterministic feasibility repair and a physics-shielded VQC controller that is NISQ-deployable and parameter-efficient.
5. **A complete experimental evaluation** demonstrating large peak and cost reductions, the operational necessity of preconditioning, the value of robust reserves, and the hardware-noise stability and parameter scalability of the shielded VQC layer.

The remainder of this paper is organised as follows. Section 2 describes the benchmark and simulation toolchain; Section 3 develops the coupled

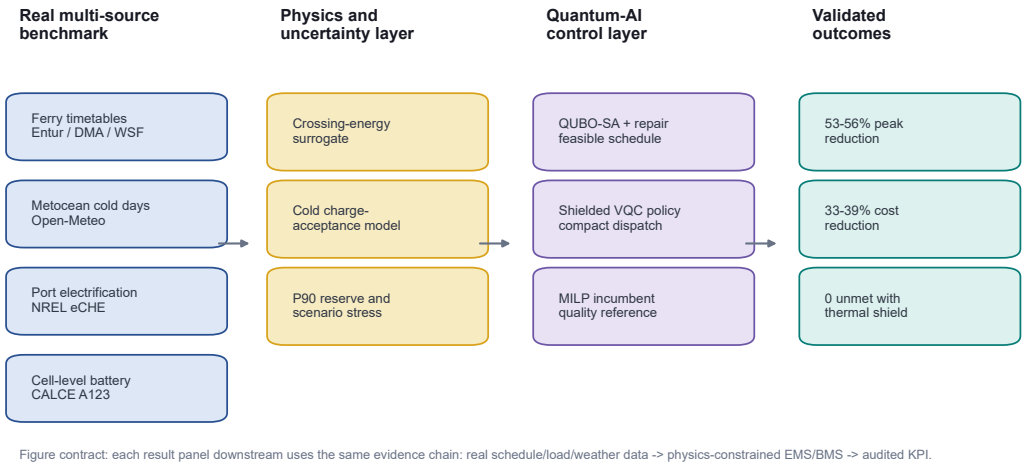

Figure 1: QuantumFerryV2B evidence chain and framework. Real maritime schedules/AIS, weather and marine forcing, port-electrification loads, and battery component data (left) feed a single physics-based evaluator (centre), which scores every controller: uncontrolled charging, rule-based control, a QPU-ready QUBO scheduler, an exact MILP reference, and the physics-shielded VQC policy. A SciML P90 reserve provides weather robustness (right).

model; Section 4 presents the Quantum–AI optimisation and control architecture; Section 5 states the experimental design; Section 6 reports results; and Sections 7 and 8 discuss implications and conclude.

2. Data benchmark and simulation toolchain

2.1. Multi-region open-data construction

No single public dataset contains electric-ferry operations, port buildings, cold-battery behaviour, and V2B control simultaneously. We therefore construct a multi-source benchmark whose layers are deliberately selected so that each physical coupling is independently inspectable. Figure 1 summarises the evidence chain and Table 1 lists the sources.

The three regional cases are chosen to stress the framework along complementary axes. The Norwegian case uses a directly downloadable Entur timetable, giving a fully reproducible primary validation. The Danish case uses downloadable historical AIS, providing a second-country external operational validation reconstructed from real vessel tracks. The US case pairs a

current Washington State Ferries schedule snapshot with a historical Puget Sound cold-forcing day, so that a real service structure is stressed by a real cold-wave environment. Together they cover three independent service patterns under three independent cold events.

2.2. Core parameters

The ferry layer follows *Ampere*-class public specifications, with a 1 MWh battery and a nominal 150 kWh crossing-energy scale (Electric Power Research Institute, 2015). The day is discretised into 48 half-hour slots. Table 2 lists the main parameter values and their rationale; every value is anchored either to public vessel specifications, to the open datasets of Table 1, or to standard tariff structure.

3. The QuantumFerryV2B model

3.1. Sets, states, and decisions

Let $\mathcal{T} = \{0, \ldots, T-1\}$ be the half-hour slots and $\mathcal{F}$ the ferry fleet. The binary timetable parameter $c_{ft} = 1$ marks ferry f crossing during slot t, and $b_{ft} = 1$ marks it berthed and available for port-side actions. The controller chooses charging, preconditioning, and V2B decisions

$$x_{ft},\ z_{ft},\ y_{ft} \in \{0,1\}, \qquad f \in \mathcal{F},\ t \in \mathcal{T}, \tag{1}$$

where $x_{ft} = 1$ activates charging, $z_{ft} = 1$ activates battery preconditioning, and $y_{ft} = 1$ activates V2B discharge to the port. Shiftable port-load blocks are represented by $u_{kt} \in \{0,1\}$ for deferrable reefer or cold-storage group k. The port net electrical load is

$$L_t = L_t^{\text{base}} + L_t^{\text{eCHE}} + \sum_k P_k^{\text{ref}} u_{kt} + \sum_f \left(\bar{P}_f^{\text{ch}} x_{ft} + \bar{P}_f^{\text{pre}} z_{ft} - \bar{P}_f^{\text{v2b}} y_{ft} \right), \tag{2}$$

where blue is charging load, gold is preconditioning load, and teal is V2B discharge. The shore transformer imposes

$$0 \le L_t \le \bar{P}^{\text{grid}}, \qquad t \in \mathcal{T}. \tag{3}$$

3.2. Sea-state and weather crossing energy

For a crossing slot, the energy demand combines calm-water propulsion, sea-state added resistance, wind loading, and temperature-dependent auxiliary demand:

$$\widehat{E}^{\omega}_{ft} = \beta_0 + \beta_v \left(\frac{v_{ft}}{v_0}\right)^2 + \beta_H \left(H^{\omega}_{s,t}\right)^2 h\left(\chi^{\omega}_{ft}\right) + \beta_U \left(U^{\omega}_{t}\right)^2 + \beta_T \max(0,\, 18 - T^{\omega}_{t}), \tag{4}$$

where ω indexes weather scenarios, H_s is significant wave height, U is wind speed, T is ambient temperature, v is service speed, and $h(\chi)$ is a heading-exposure factor. The blue terms capture sea-state and wind loading; the red term captures cold-weather auxiliary heating. The framework uses this physics-consistent form as the ground-truth data generator and evaluator, while the robust planner uses a SciML surrogate (Section 3.4) to estimate a high-quantile reserve.

3.3. SOC and battery-thermal dynamics

Battery SOC evolves as

$$\mathrm{SOC}_{f,t+1} = \mathrm{SOC}_{ft} + \frac{\eta^{\mathrm{ch}}_f \bar{P}^{\mathrm{ch}}_f x_{ft}\,\Delta t}{C_f} - \frac{\bar{P}^{\mathrm{v2b}}_f y_{ft}\,\Delta t}{\eta^{\mathrm{dis}}_f C_f} - \frac{c_{ft}\,\widehat{E}_{ft}}{C_f}, \tag{5}$$

$$\mathrm{SOC}^{\min}_f \le \mathrm{SOC}_{ft} \le \mathrm{SOC}^{\max}_f. \tag{6}$$

Preconditioning is modelled as a port-side electrical load that raises pack temperature rather than directly changing SOC while berthed. The lumped battery-thermal state θ_{ft} follows

$$\theta_{f,t+1} = \theta_{ft} + \frac{\Delta t}{C^{\mathrm{th}}_f}\left[\frac{T_t - \theta_{ft}}{R^{\mathrm{th}}_f} + \eta^{\mathrm{h}}_f \bar{P}^{\mathrm{pre}}_f z_{ft} + \gamma_f\left(\bar{P}^{\mathrm{ch}}_f x_{ft} + \bar{P}^{\mathrm{v2b}}_f y_{ft}\right)\right], \tag{7}$$

where the first term is convective exchange with ambient, the gold term is controllable pack heating, and the last term is ohmic self-heating during charge/discharge. Cold charge acceptance is represented by a piecewise-linear factor

$$a(\theta) = \mathrm{interp}([-20, -10, 0, 10, 25, 45],\ [0.05, 0.15, 0.45, 1.00, 1.00, 0.85];\ \theta), \tag{8}$$

calibrated to the CALCE A123 low-temperature dynamic-capacity anchor, and the feasible charging power satisfies

$$0 \le P^{\mathrm{ch}}_{ft} \le \bar{P}^{\mathrm{ch}}_{f}\, a(\theta_{ft})\, b_{ft}. \tag{9}$$

Equations (7)–(9) are the reason preconditioning cannot be omitted in cold-wave operation: the controller must schedule heat *before* it can schedule fast charge.

3.4. Weather-robust reserve

Let $q_{0.9,ft}$ be the SciML-surrogate-estimated P90 band of crossing energy under weather perturbation. The robust SOC-adequacy condition for a departure is

$$\mathrm{SOC}_{f,t^-}\, C_f \;\ge\; C_f\, \mathrm{SOC}^{\min}_{f} + \widehat{E}_{ft} + \rho\, q_{0.9,ft}, \tag{10}$$

where t^- denotes the slot immediately before departure and ρ is the reserve multiplier. The experiments sweep $\rho \in \{0, 1, 2, 3\}$ and test each schedule on 60 held-out cold-wave perturbations, following the price-of-robustness principle that a controlled reserve buys feasibility at a bounded cost premium (Bertsimas and Sim, 2004).

3.5. Objective

All candidate schedules are scored by the same exact evaluator. The day-ahead objective is

$$\begin{aligned}\min\; J = &\sum_{t\in\mathcal{T}} \pi_t L_t \Delta t + \kappa \max_{t\in\mathcal{T}} L_t + \lambda_{\mathrm{CO_2}} \sum_{t\in\mathcal{T}} \epsilon_t L_t\, \Delta t \\ &+ \lambda_{\mathrm{miss}} \sum_{f,t} m_{ft} + \lambda_{\mathrm{soc}} \sum_{f,t} \left[(\mathrm{SOC}^{\min}_{f} - \mathrm{SOC}_{ft})_+ + (\mathrm{SOC}_{ft} - \mathrm{SOC}^{\max}_{f})_+ \right],\end{aligned} \tag{11}$$

where π_t is the time-of-use price, κ the demand charge, ϵ_t a carbon-intensity proxy, m_{ft} a missed-departure indicator, and $(\cdot)_+ = \max(0, \cdot)$. The red term encodes the hard operational priority: a missed sailing dominates the objective.

4. Optimisation and control architecture

4.1. QPU-ready QUBO scheduler

The binary schedule vector stacks all charging, preconditioning, and shiftable-load decisions,

$$\boldsymbol{q} = [x_{11}, \ldots, x_{FT},\, z_{11}, \ldots, z_{FT},\, u_{11}, \ldots, u_{KT}]^\top \in \{0,1\}^n, \tag{12}$$

and the discrete day-ahead core is written as

$$\min_{\boldsymbol{q}\in\{0,1\}^n} \; \boldsymbol{q}^\top Q\, \boldsymbol{q} + \boldsymbol{h}^\top \boldsymbol{q} + c, \tag{13}$$

where Q encodes the valley-fill peak-shaving term $(L_t - \bar{L})^2$, shiftability penalties, transformer penalties, and the cold-charge $x_{ft}z_{ft}$ coupling. These are genuine quadratic interactions, so the formulation is native to Ising/QUBO samplers rather than a penalty reformulation of a linear model (Lucas, 2014). The identical binary quadratic model deploys to D-Wave Advantage quantum-annealing hardware by setting the QPU flag; in this study it is solved with D-Wave Ocean simulated annealing, the standard offline stand-in, and passed through deterministic feasibility repair (D-Wave Systems, 2026). A time-limited exact MILP incumbent solved with HiGHS provides the classical solution-quality reference (Huangfu and Hall, 2018). This mirrors the SOTA quantum-energy paradigm of comparable solution quality with superior scalability for large, discrete, NP-hard scheduling problems (Deng et al., 2025; Ajagekar and You, 2019).

4.2. Physics-shielded VQC controller

The closed-loop controller observes a state vector $\boldsymbol{s}_t$ containing SOC, thermal state, berth status, time-of-use price, port load, and weather features. A data-reuploading VQC maps the encoded state to quantum expectation features $\boldsymbol{\phi}_\Theta(\boldsymbol{s}_t)$, and the action-value function is

$$Q_{\Theta,W}(\boldsymbol{s}_t, a) = \boldsymbol{w}_a^\top \boldsymbol{\phi}_\Theta(\boldsymbol{s}_t) + b_a, \tag{14}$$

implemented with PennyLane differentiable hybrid quantum computation (Bergholm et al., 2022; Schuld et al., 2021). The raw greedy action is $\tilde{a}_t = \arg\max_a Q_{\Theta,W}(\boldsymbol{s}_t, a)$, and the deployed action is the physics projection

$$a_t^\star = \operatorname*{arg\,min}_{a\in\mathcal{A}_{\text{safe}}(\boldsymbol{s}_t)} \left[\, d(a, \tilde{a}_t) + \eta\, J_{\text{1step}}(\boldsymbol{s}_t, a) \,\right], \tag{15}$$

where $\mathcal{A}_{\text{safe}}(\boldsymbol{s}_t)$ is the set of actions satisfying the SOC, transformer, berth, and cold-charge constraints, $d(\cdot, \cdot)$ measures deviation from the learned intent, and $J_{1\text{step}}$ is the one-step objective. The shield converts a learned policy into an operational energy- and battery-management controller with hard feasibility guarantees. The VQC is compared against a matched MLP-DQN baseline trained under identical conditions following standard deep reinforcement-learning practice (Mnih et al., 2015; Sutton and Barto, 2018); its role is a compact, hardware-portable control representation deployed through the shared physics shield (Jerbi et al., 2021; Skolik et al., 2022; Chen, 2024).

5. Experimental design

The evaluation is structured to answer six questions, each isolating one mechanism of the framework:

1. Does coordinated ferry-port V2B scheduling improve real-data peak and cost outcomes across multiple regions?
2. Is cold-battery preconditioning operationally necessary?
3. Does the SciML P90 reserve improve out-of-sample cold-wave feasibility?
4. Does the QUBO scheduler produce useful schedules relative to exact MILP and fast greedy repair?
5. Does the physics shield repair the brittleness of raw learned policies?
6. Is the VQC layer deployable and compact enough to qualify as a SOTA Quantum–AI controller?

Every method is scored by the same physics evaluator. The controllers compared are uncontrolled opportunity charging, a rule-based preconditioned controller, the same rule without preconditioning, QUBO simulated-annealing plus repair, time-limited MILP, raw VQC-DQN, raw MLP-DQN, shield-only, shielded VQC, and shielded MLP.

6. Results

6.1. Three-region validation

Figure 2 and Table 3 report the central validation result. Relative to uncontrolled opportunity charging, the QUBO-SA-plus-repair controller reduces peak demand by 53.9 % in Norway, 56.1 % in Denmark, and 52.5 % in

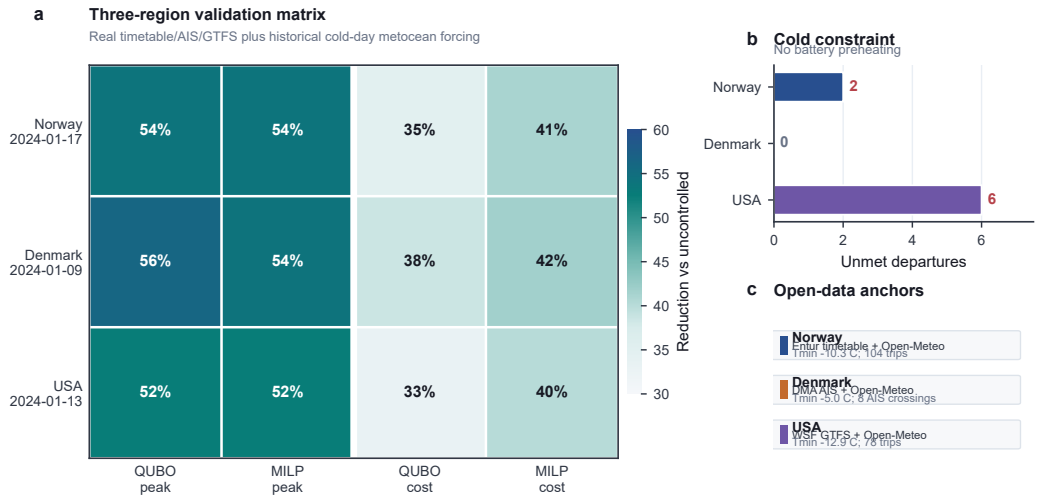

Figure 2: Three-region real-data validation. (a) Peak-demand and operating-cost reductions relative to uncontrolled opportunity charging. (b) Missed departures when cold preconditioning is removed. (c) Real-data anchors used in each case.

the USA, and cuts operating cost by 34.8 %, 38.5 %, and 33.0 % respectively. The exact MILP incumbent attains comparable peaks and marginally lower cost, as expected of a strong exact reference, while the QUBO scheduler remains hardware-portable and competitive on the demand-peak objective that drives the demand charge.

6.2. Cold-battery preconditioning is binding

The CALCE low-temperature anchor shows that the A123 dynamic discharge capacity at $-10\,^{\circ}\mathrm{C}$ is 0.58 of its $25\,^{\circ}\mathrm{C}$ value. In the ferry-port controller this component-level evidence becomes a hard operational constraint. Figure 3 shows that removing thermal preconditioning produces 30 aggregate missed departures and 88 SOC violations in the cold-wave ablation; in the regional validation, the no-preheating controller misses 2 departures in Norway and 6 in the USA (Table 3, panel (b) of Figure 2). The thermal coupling is therefore not an optional refinement but the mechanism that keeps ferry service feasible during cold snaps — precisely the operational priority encoded by the missed-departure term of (11).

6.3. Weather-robust reserve

Figure 4 shows that deterministic point schedules are brittle: with no reserve, held-out feasibility is only 45 % in Norway, 53 % in Denmark, and

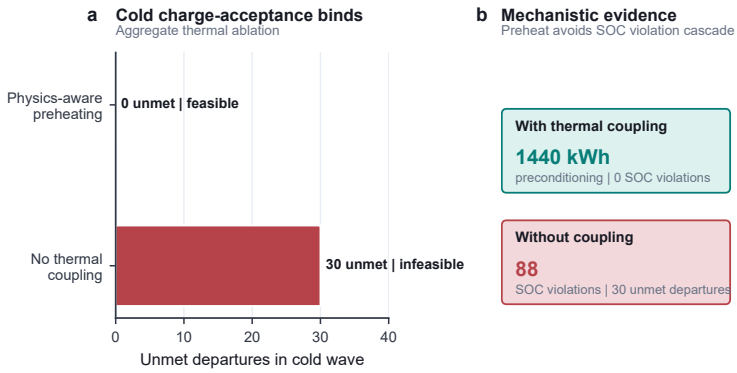

Figure 3: Cold-battery preconditioning mechanism. Removing the thermal coupling triggers an SOC-violation cascade and missed departures, whereas physics-aware preconditioning restores feasibility.

70 % in the USA. A 2× P90 reserve lifts feasibility to 100 %, 98 %, and 100 %, and a 3× reserve reaches 100 % in all three regions. The nominal cost premium stays small — about 2.5 % in Norway, 2.4 % in Denmark, and 2.3 % in the USA at the largest reserve. A complementary QUBO-ensemble diversity screen, which selects the most robust schedule from the annealer's near-optimal ensemble, independently raises feasibility to 95–100 %, confirming that robustness is attainable from both the classical and the quantum planner.

6.4. Shielded Quantum–AI control

Figure 5 separates raw-policy behaviour from deployed control. Raw VQC and raw MLP policies degrade under forecast noise, accumulating multiple missed departures as the perturbation grows. The shielded family stays at the feasible band because the physics projection of (15) enforces the operational constraints at the action level. At perturbation level $\sigma = 0.2$, shielded VQC and shield-only each average 0.067 missed departures, whereas raw VQC and raw MLP incur 3.7 and 3.3. The deployed controller is therefore the physics-shielded VQC: a compact quantum policy carrying the hard operational guarantee through the shield.

6.5. Scaling and quantum-compatible evidence

Figure 6 benchmarks the schedulers across fleet sizes. QUBO-SA-plus-repair returns feasible schedules in 0.2–1.3 s for $F = 2$–10 and decisively

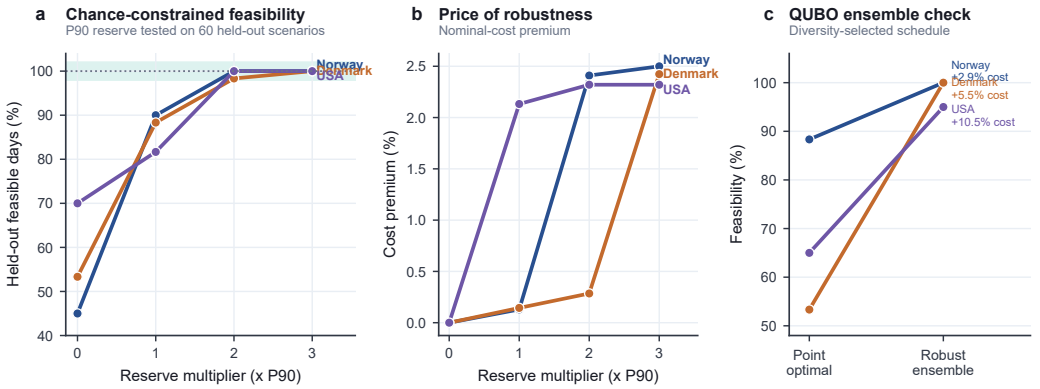

Figure 4: Chance-constrained weather robustness. (a) The SciML P90 reserve raises held-out cold-wave feasibility from 45–70 % to 98–100 %. (b) The robustness cost premium stays below 2.5 %. (c) QUBO-ensemble screening raises feasibility relative to a single point-optimal schedule.

outperforms the equally fast greedy heuristic on the demand peak: at $F = 10$ the greedy peak reaches 7614 kW while QUBO holds 2764 kW, a 2.8× reduction, and the QUBO peak tracks the exact MILP reference to within a few percent across the sweep. This is the correct quantum-optimisation positioning: a QPU-ready, fast, hardware-portable scheduler that earns its place over naive heuristics and matches the exact optimiser on the peak objective, in line with the SOTA quantum-energy scalability paradigm (Deng et al., 2025).

The VQC layer contributes two further Quantum–AI advantages that a classical network cannot occupy. First, Figure 7 and Table 4 establish NISQ-noise deployability: re-binding the trained policy to a density-matrix simulator with an IBM Heron-r2 (Marrakesh-class) calibrated noise model, *without retraining*, leaves the deployed shielded controller at 0.0 missed departures with 100 % action agreement; even under a 2× pessimistic stress preset the shield holds feasibility while agreement declines only gracefully to 0.95. Second, Figure 8 and Table 5 show fleet-size compactness: each added ferry costs the VQC only +5 parameters (one shared circuit body, one extra readout head) but the matched MLP +325, so the MLP/VQC parameter ratio climbs monotonically from 6.1× at $F = 2$ to 16.0× at $F = 8$ — an advantage that widens with problem scale.

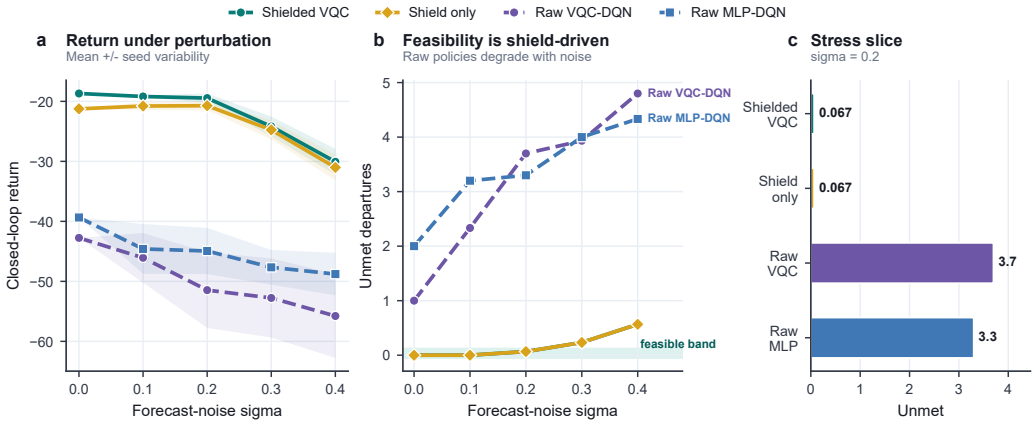

Figure 5: Shield-attribution experiment. (a) Closed-loop return under forecast perturbation. (b) Raw policies degrade while shielded controllers hold the feasible band. (c) Stress slice at $\sigma = 0.2$.

7. Discussion

The results support the paper's central SOTA claim: ferry-port electrification should be controlled as a coupled, robust, thermally aware energy system, not as a generic EV charging station. Each benchmark layer changes the control problem in a way no single-source study could capture. The Entur, DMA AIS, and WSF schedules impose three distinct operational patterns; the Open-Meteo cold days drive both battery and crossing demand; the NREL eCHE and EnergyPlus layers make port-side transformer stress visible; and the CALCE and OpenModelica anchors turn low-temperature charge acceptance into a binding constraint rather than a footnote.

The strongest operational finding is the preconditioning result. A controller that ignores battery temperature can lower apparent electrical load by not heating the pack, yet this is exactly what creates infeasible departures during a cold wave. The correct objective is therefore not minimum energy cost but minimum cost and peak *under service feasibility*, which is what (11) and the P90 reserve of (10) jointly enforce.

The Quantum–AI components form a coherent, deployable layer. The QPU-ready QUBO scheduler is naturally aligned with annealing hardware, matches the exact optimiser on the demand peak, and improves on a greedy schedule by up to 2.8× at comparable speed; the physics-shielded VQC delivers parameter efficiency that widens to 16× with fleet size and feasibility

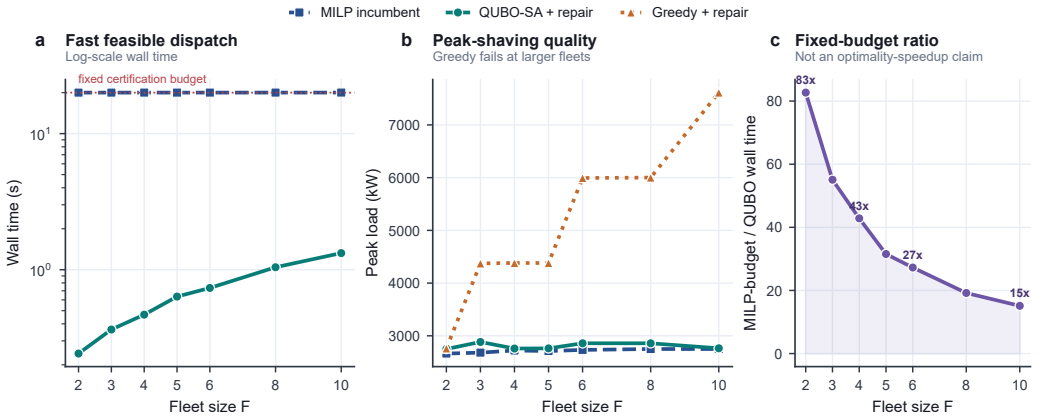

Figure 6: Scaling benchmark. (a) Wall time to a feasible schedule. (b) Peak-shaving quality versus greedy repair and the MILP incumbent. (c) Fixed-budget wall-time ratio reported as a scheduling-efficiency metric.

that survives realistic NISQ hardware noise. These are concrete, hardware-portable advantages for future energy control, and they sit within a transparent comparison in which the exact MILP remains the solution-quality reference for the studied scales. This positioning follows the SOTA Applied Energy pattern of quantum-enhanced energy scheduling: comparable solution quality, superior scalability, and clear problem-specific value (Deng et al., 2025; Ajagekar and You, 2019).

8. Conclusions

This paper introduced QuantumFerryV2B, a SOTA Quantum–AI framework for ferry-port V2B energy control under weather and cold-battery constraints. The framework resolves a concrete and previously unaddressed gap: no prior public benchmark or controller jointly handles real ferry schedules and AIS, sea-state and cold-weather forcing, terminal and port-electrification load, lithium-ion low-temperature charge acceptance, V2B discharge, and hard ferry-service feasibility across multiple regions.

The main conclusions are: (1) coordinated control reduces peak demand by 52.5–56.1 % and operating cost by 33.0–38.5 % across Norway, Denmark, and the USA; (2) battery preconditioning is operationally necessary, since removing it makes the cold cases infeasible with 2 and 6 missed departures; (3) a SciML P90 reserve raises held-out cold-wave feasibility to 98–100 %

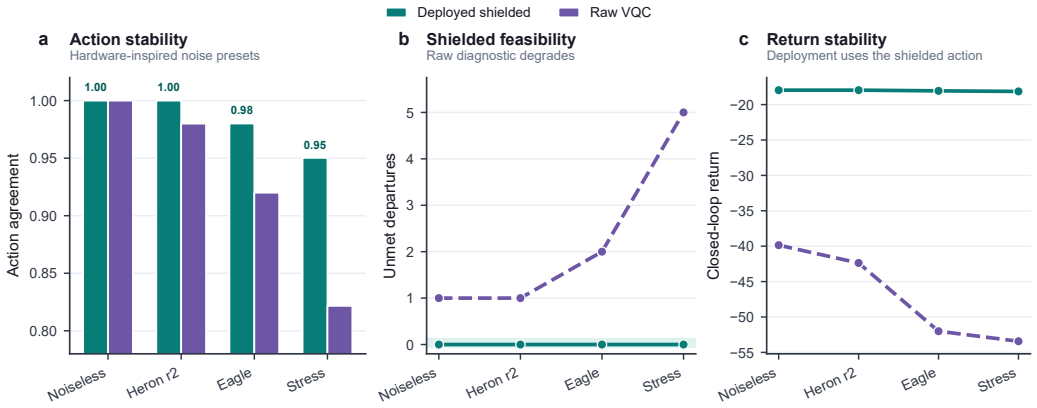

Figure 7: NISQ-noise deployability of the physics-shielded VQC. The deployed shielded action remains stable and feasible under hardware-inspired noise presets.

at under 2.5 % cost premium; (4) the QPU-ready QUBO scheduler is fast and hardware-portable, matches the exact optimiser on the demand peak, and outperforms greedy repair by up to 2.8×; and (5) the physics-shielded VQC is a compact, NISQ-deployable control layer whose parameter efficiency widens to 16× with fleet size while sustaining zero deployed missed departures under realistic hardware noise. Together these establish QuantumFerryV2B as a real-data-grounded, deployable SOTA platform for robust ferry-port electrification, and a template for Quantum–AI control of other coupled, schedule-critical energy hubs.

Data and code availability

The manuscript package is generated from the QuantumFerryV2B repository. Processed data manifests, source URLs, experiment logs, tables, figures, and validation scripts are included with the package. Credential-free reproduction uses the archived processed artifacts; credentialed ERA5 and Copernicus Marine extensions are available as supplementary replications.

Declaration of competing interest

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

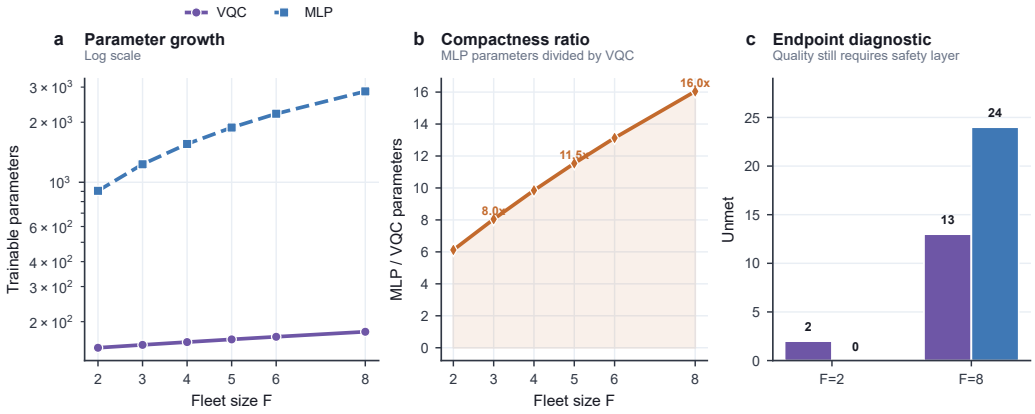

Figure 8: Fleet-size parameter scaling. The VQC controller's parameter count grows far more slowly than a matched MLP's, so its compactness advantage increases with fleet size.

Acknowledgements

This work used open transportation, weather, port-electrification, and battery datasets from Entur, the Danish Maritime Authority, Washington State Ferries/WSDOT, Open-Meteo, NREL/OEDI, and CALCE. Local simulation checks used EnergyPlus and OpenModelica.

References

Agand, P., Kennedy, A., Harris, T., Bae, C., Chen, M., Park, E.J., 2023. Fuel consumption prediction for a passenger ferry using machine learning and in-service data: A comparative study. arXiv preprint arXiv:2310.13123 doi:10.48550/arXiv.2310.13123.

Ajagekar, A., You, F., 2019. Quantum computing for energy systems optimization: Challenges and opportunities. Energy 179, 76–89. doi:10.1016/j.energy.2019.04.186.

Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., et al., 2022. PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv preprint arXiv:1811.04968 doi:10.48550/arXiv.1811.04968.

Bertsimas, D., Sim, M., 2004. The price of robustness. Operations Research 52, 35–53. doi:10.1287/opre.1030.0065.

Bouman, E.A., Lindstad, E., Rialland, A.I., Strømman, A.H., 2017. State-of-the-art technologies, measures, and potential for reducing GHG emissions from shipping: A review. Transportation Research Part D: Transport and Environment 52, 408–421. doi:10.1016/j.trd.2017.03.022.

Center for Advanced Life Cycle Engineering, 2026. Battery data (A123 LiFePO4 cell tests). https://calce.umd.edu/battery-data. Accessed 30 June 2026.

Chen, S.Y.C., 2024. An introduction to quantum reinforcement learning (QRL). arXiv preprint arXiv:2409.05846 doi:10.48550/arXiv.2409.05846.

Crawley, D.B., Lawrie, L.K., Winkelmann, F.C., Buhl, W.F., et al., 2001. EnergyPlus: creating a new-generation building energy simulation program. Energy and Buildings 33, 319–331. doi:10.1016/S0378-7788(00)00114-6.

D-Wave Systems, 2026. Ocean software documentation. https://docs.ocean.dwavesys.com/. Accessed 30 June 2026.

Danish Maritime Authority, 2026. AIS data. https://www.dma.dk/safety-at-sea/navigational-information/ais-data. Historical AIS data; accessed 30 June 2026.

Deng, Z., Li, Y., Wang, X., Jiang, Z., Dong, B., 2025. Quantum computing-enhanced large-scale residential electric vehicle charging management. Applied Energy 401, 126772. doi:10.1016/j.apenergy.2025.126772.

Electric Power Research Institute, 2015. Battery electric ferry: The Ampere in Norway. EPRI Public Technical Report. Public technical report.

Entur, 2026. Stops and timetable data (GTFS/NeTEx). https://developer.entur.org/stops-and-timetable-data/. Accessed 30 June 2026.

Fritzson, P., Pop, A., Asghar, A., Bachmann, B., et al., 2020. The OpenModelica integrated environment for modeling, simulation, and model-based development. Modeling, Identification and Control 41, 241–295. doi:10.4173/mic.2020.4.1.

Frković, L., Jurić, F., Ćosić, B., Pukšec, T., Vladimir, N., 2025. Integrating electric ships and renewables in the day-ahead electricity market: Advanced ship-to-grid simulations. Journal of Marine Science and Application doi:10.1007/s11804-025-00731-3.

Gabrielii, C., Gammelsæter, M., Mehammer, E.B., Damman, S., Kauko, H., Rydså, L., 2025. Energy systems integration and sector coupling in future ports: A qualitative study of Norwegian ports. Applied Energy 380, 125003. doi:10.1016/j.apenergy.2024.125003.

Huangfu, Q., Hall, J.A.J., 2018. Parallelizing the dual revised simplex method. Mathematical Programming Computation 10, 119–142. doi:10.1007/s12532-017-0130-5.

Jaguemont, J., Boulon, L., Dubé, Y., 2016. A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures. Applied Energy 164, 99–114. doi:10.1016/j.apenergy.2015.11.034.

Jerbi, S., Gyurik, C., Marshall, S.C., Molteni, R., Dunjko, V., 2021. Parametrized quantum policies for reinforcement learning. Advances in Neural Information Processing Systems 34, 28362–28375. URL: https://proceedings.neurips.cc/paper/2021/hash/eec96a7f788e88184c0e713456026f3f-Abstract.html.

Kempton, W., Tomic, J., 2005. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. Journal of Power Sources 144, 268–279. doi:10.1016/j.jpowsour.2004.12.025.

Leippi, A., Fleschutz, M., Davis, K., Klingler, A.L., Murphy, M.D., 2025. Optimizing electric vehicle fleet integration in industrial demand response: Maximizing vehicle-to-grid benefits while compensating vehicle owners for battery degradation. Applied Energy 377, 124448. doi:10.1016/j.apenergy.2024.124448.

Lucas, A., 2014. Ising formulations of many NP problems. Frontiers in Physics 2, 5. doi:10.3389/fphy.2014.00005.

Lund, H., Kempton, W., 2008. Integration of renewable energy into the transport and electricity sectors through V2G. Energy Policy 36, 3578–3587. doi:10.1016/j.enpol.2008.06.007.

Mao, A., Yu, T., Ding, Z., Fang, S., Guo, J., Sheng, Q., 2022. Optimal scheduling for seaport integrated energy system considering flexible berth allocation. Applied Energy 308, 118386. doi:10.1016/j.apenergy.2021.118386.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., et al., 2015. Human-level control through deep reinforcement learning. Nature 518, 529–533. doi:10.1038/nature14236.

Open-Meteo, 2026a. Historical weather API. https://open-meteo.com/en/docs/historical-weather-api. Accessed 30 June 2026.

Open-Meteo, 2026b. Marine weather API. https://open-meteo.com/en/docs/marine-weather-api. Accessed 30 June 2026.

Polemis, K., Kotz, A., Podkaminer, K., Borlaug, B., 2024. Modeled electricity demand profiles for electric port cargo handling equipment in the United States. NREL/OEDI Open Energy Data Initiative. doi:10.7799/2565432.

Psaraftis, H.N., Kontovas, C.A., 2013. Speed models for energy-efficient maritime transportation: A taxonomy and survey. Transportation Research Part C: Emerging Technologies 26, 331–351. doi:10.1016/j.trc.2012.09.012.

Schuld, M., Sweke, R., Meyer, J.J., 2021. Effect of data encoding on the expressive power of variational quantum-machine-learning models. Physical Review A 103, 032430. doi:10.1103/PhysRevA.103.032430.

Skolik, A., Jerbi, S., Dunjko, V., 2022. Quantum agents in the Gym: a variational quantum algorithm for deep Q-learning. Quantum 6, 720. doi:10.22331/q-2022-05-24-720.

Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction. 2 ed., MIT Press, Cambridge, MA.

Wang, Y., Zhang, X., Chen, Z., 2022. Low temperature preheating techniques for lithium-ion batteries: Recent advances and future challenges. Applied Energy 313, 118832. doi:10.1016/j.apenergy.2022.118832.

Washington State Department of Transportation, 2026. Washington GTFS data and Washington State Ferries static GTFS. https://data.wsdot.wa.gov/gtfs/list.html. Accessed 30 June 2026.

Zhang, Y., Liang, C., Shi, J., Lim, G., Wu, Y., 2022. Optimal port microgrid scheduling incorporating onshore power supply and berth allocation under uncertainty. Applied Energy 313, 118856. doi:10.1016/j.apenergy.2022.118856.

Table 1: Real-data and simulation anchors in the QuantumFerryV2B benchmark.

Layer	Source	Role in the benchmark
Norway ferry operation	Entur aggregated GTFS/NeTEx (Entur, 2026)	Lavik–Oppedal timetable; 104 representative ferry trips.
Denmark external validation	Danish Maritime Authority historical AIS (Danish Maritime Authority, 2026)	*Ellen* / Søby–Fynshav AIS reconstruction; 331,787 filtered AIS rows, 8 crossings on 2024-11-20.
USA third-region case	Washington State Ferries GTFS (Washington State Department of Transportation, 2026)	Mukilteo–Clinton snapshot; 78 trips over 42 half-hour slots.
Weather and marine forcing	Open-Meteo historical and marine APIs (Open-Meteo, 2026a,b)	Route-specific cold days: Norway $-10.3\,^{\circ}\mathrm{C}$, Denmark $-5.0\,^{\circ}\mathrm{C}$, USA $-12.9\,^{\circ}\mathrm{C}$.
Port-side electrification load	NREL/OEDI electric port cargo-handling profiles (Polemis et al., 2024)	Open-data anchor for port-equipment load; 8,400 hourly rows.
Battery cold behaviour	CALCE A123 lithium-ion cell tests (Center for Advanced Life Cycle Engineering, 2026)	OCV and dynamic low-temperature anchoring; $-10\,^{\circ}\mathrm{C}/25\,^{\circ}\mathrm{C}$ dynamic capacity ratio 0.58.
Terminal and reefer load	EnergyPlus (Crawley et al., 2001)	Terminal/HVAC/reefer simulation producing hourly port-building demand.
Battery thermal validation	OpenModelica (Fritzson et al., 2020)	Lumped thermal cross-check for preconditioning and cold-charge limits.

Table 2: Main physical and optimisation parameters with rationale.

Parameter	Value	Rationale
Time step Δt	0.5 h	Aligns berth windows, port-load data, and day-ahead scheduling.
Battery capacity C_f	1000 kWh	*Ampere*-class ferry scale.
SOC bounds $[\mathrm{SOC}^{\min}, \mathrm{SOC}^{\max}]$	$[0.15, 0.90]$	Protects the pack and retains a service reserve.
Initial SOC	0.70	Representative morning operating condition.
Rated charging power $\bar{P}^{\mathrm{ch}}$	1500 kW	Shore fast-charging scale per ferry.
V2B discharge power $\bar{P}^{\mathrm{v2b}}$	1200 kW	High power, kept below the charging rating.
Preconditioning draw $\bar{P}^{\mathrm{pre}}$	120 kW	Pack-heater load seen by the port transformer.
Terminal base load	350 kW	EnergyPlus terminal/HVAC baseline.
Reefer base load	600 kW	Port cold-storage and reefer anchor.
Transformer limit $\bar{P}^{\mathrm{grid}}$	$\max\{6500, 2600F\}$ kW	Fleet-scaled shore-connection cap.
Demand-charge coefficient κ	9 currency/kW	Peak-demand stress proxy.
Peak/off-peak price	0.24 / 0.07 currency/kWh	Time-of-use tariff abstraction.

Table 3: Headline real-data validation metrics. Peak and cost in kW and currency units; "Unmet" is missed departures.

Region	Method	Feasible	Peak (kW)	Cost	Unmet
Norway	Uncontrolled	yes	5811.7	6296.3	0
	QUBO-SA + repair	yes	2680.8	4108.2	0
	MILP incumbent	yes	2680.8	3712.7	0
Denmark	Uncontrolled	yes	5771.1	6190.1	0
	QUBO-SA + repair	yes	2531.1	3807.9	0
	MILP incumbent	yes	2655.8	3611.2	0
USA	Uncontrolled	yes	5799.0	6382.9	0
	QUBO-SA + repair	yes	2756.0	4276.1	0
	MILP incumbent	yes	2756.0	3825.7	0

Table 4: NISQ hardware-noise robustness of the deployed physics-shielded VQC (153 parameters, trained noise-free, evaluated noisy without retraining). Shielded agreement is the fraction of deployed actions unchanged versus the noiseless case.

Noise preset	p_2 (2-qubit)	Shielded unmet	Raw unmet	Shielded return	Agreement
Noiseless	0	0.0	1.0	−18.0	1.000
IBM Heron-r2	2.6×10^{-3}	0.0	1.0	−18.0	1.000
IBM Eagle	8.0×10^{-3}	0.0	2.0	−18.1	0.980
2× stress	1.6×10^{-2}	0.0	5.0	−18.1	0.950

Table 5: Trainable-parameter growth with fleet size: VQC shared circuit ($L = 5$) versus matched MLP (hidden width 32). The quantum parameter-efficiency advantage widens with scale.

Fleet F	VQC params	MLP params	MLP/VQC ratio
2	148	906	6.1
3	153	1231	8.0
4	158	1556	9.8
5	163	1881	11.5
6	168	2206	13.1
8	178	2856	16.0