World Journal of Advanced Engineering Technology and Sciences
eISSN: 2582-8266
Cross Ref DOI: 10.30574/wjaets
Journal homepage: https://wjaets.com/

(Review Article)

Check for updates

Quantum-Enhanced Travel Procurement: Hybrid Quantum–Classical Optimization for Enterprise Travel Management

Prajkta Waditwar *

Senior Manager- Strategic Sourcing at Box, Inc. (Independent Researcher)

World Journal of Advanced Engineering Technology and Sciences, 2025, 17(03), 375-386

Publication history: Received 16 November 2025; revised on 23 December 2025; accepted on 25 December 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.17.3.1572

Abstract

Corporate travel procurement is a multi-objective, constraint-dense decision domain spanning strategic supplier portfolio selection, travel policy design, and operational disruption response. Classical methods—mixed-integer programming (MIP), constraint programming, metaheuristics, and machine learning—deliver strong results but face scaling and responsiveness challenges as procurement objectives broaden to include compliance, traveler experience, sustainability, and resilience. Quantum optimization methods, particularly quantum annealing and gate-based variational algorithms (e.g., QAOA), have been proposed for combinatorial problems with binary decisions and complex interaction terms. Yet empirical evidence across optimization domains often shows classical solvers match or exceed quantum approaches on meaningful instance sets, motivating a hybrid quantum–classical posture rather than replacement.

This paper formalizes quantum-enhanced travel procurement as a hypothesis-driven, hybrid decision-support approach in which quantum routines contribute to solution search, diversification, or time-to-decision for selected combinatorial cores. We present canonical procurement formulations, show explicit mappings to QUBO/Ising models, and propose hybrid architectures for supplier portfolio design, airline share allocation under commitments, policy parameter tuning, and disruption re-accommodation. To strengthen accessibility and rigor, we provide (i) a worked QUBO example with explicit coefficients and variable counts; (ii) an operational benchmark protocol with representative instance sizes, solver baselines, runtime assumptions, and statistical reporting; and (iii) an enterprise governance view including post-quantum cryptography readiness across vendor integrations.

Keywords: Quantum Computing; Corporate Travel Procurement; Combinatorial Optimization; QUBO; Hybrid Quantum–Classical Algorithms; Travel Policy; Sustainability; Post-Quantum Cryptography

1. Introduction

Enterprise travel procurement influences cost control, employee productivity, duty of care, and supplier relationships. Unlike many procurement categories, travel is characterized by (i) heterogeneous offerings (fares/rates with rules), (ii) demand volatility, (iii) policy-driven compliance and leakage dynamics, and (iv) operational disruption risk. Procurement decisions are therefore multi-layered and coupled: strategic supplier portfolios shape tactical compliance levers, which interact with operational booking and re-accommodation decisions.

Quantum computing has renewed attention on combinatorial optimization. Transportation-adjacent studies (e.g., flight gate assignment) demonstrate the feasibility of extracting real instances and expressing them as QUBOs for quantum annealers, while also highlighting practical issues such as coefficient precision and instance sizing. Meanwhile, benchmarking studies in other domains (e.g., portfolio optimization) report that classical MIP and tailored heuristics

* Corresponding author: Prajkta Waditwar.

can dominate quantum methods on realistic workloads, reinforcing the need for careful baseline comparisons and a hybrid approach.

1.1. Contributions

1.1.1. A procurement-specific framing of quantum enhancement as hybrid decision support.

This paper defines *quantum-enhanced travel procurement* as an enterprise decision-support paradigm, not as a claim of quantum supremacy or replacement of classical optimization. Specifically, we frame quantum components as search accelerators and diversification engines embedded within a classical procurement workflow that includes forecasting, constraint governance, auditing, and stakeholder negotiation. This reframing matters for travel procurement because procurement decisions must be explainable, contract-aligned, and operationally deployable (e.g., implementable in online booking tools and policy systems). By explicitly aligning quantum optimization with procurement decision cycles (strategic, tactical, operational) and procurement outcomes (cost, compliance, traveler experience, sustainability, resilience), the paper provides a domain-grounded interpretation of where quantum methods may create value and how success should be measured (e.g., scenario throughput, time-to-decision, diversity of negotiable portfolios), even in the absence of proven quantum advantage.

1.1.2. Procurement-to-QUBO mappings with constraint penalty design patterns.

The paper contributes a structured translation layer from real procurement constructs—preferred supplier selection, market coverage rules, share commitments, sustainability targets, and policy exceptions—into QUBO/Ising-compatible formulations. Rather than presenting QUBO as an abstract mathematical model, we provide reusable modeling patterns that procurement teams and researchers can apply across problem families, including:

- Cardinality constraints ("select exactly K preferred hotels") using squared penalties,
- Market exclusivity constraints ("choose one preferred airline per market"),
- Coverage constraints ("ensure at least one preferred hotel in each zone"),
- Caps and diversity constraints ("no more than M properties from a chain"), and
- Interaction modeling for redundancy penalties and complementarity bonuses. We also discuss practical penalty coefficient considerations (e.g., feasibility vs. numerical stiffness), which is critical because procurement models often combine coefficients with very different magnitudes (rates, demand weights, rebates, emissions multipliers). These patterns establish a repeatable methodology for building procurement QUBOs suitable for quantum annealing or QAOA while remaining comparable to classical MIP baselines.

1.1.3. A worked numerical QUBO example for preferred-hotel portfolio selection.

To improve accessibility and technical transparency, we include a fully specified, small-scale QUBO (10 binary variables) for a preferred-hotel selection problem. The worked example explicitly defines:

- Decision variables,
- Objective coefficients (cost, sustainability bonuses),
- Constraints (exactly k selections and neighborhood coverage),
- Interaction penalties (redundancy overlap), and
- Example penalty weights.

The purpose is twofold: (i) to show concretely how procurement objectives become QUBO terms and (ii) to make hardware constraints tangible by illustrating how variable counts and constraint interactions scale. This example serves as a template that researchers can extend to city-level procurement subproblems, and it enables reviewers to validate that the modeling claims are operational rather than purely conceptual.

1.1.4. A concrete evaluation methodology: benchmark instances, sizes, runtimes, and solver comparisons.

The paper operationalizes evaluation by specifying a benchmark design that procurement and quantum researchers can reproduce. This includes:

- Instance families (Hotels and Airlines) with representative ranges of cities/markets, candidates per market, constraint densities, and scenario volumes,

- Typical variable counts and scaling behavior (e.g., 60-candidate city subproblems vs. multi-city coupled formulations),
- Baseline solvers that must be included for credibility (MIP, CP-SAT, metaheuristics such as LNS/tabu/SA, and quantum/hybrid solvers),
- Runtime budgets aligned to procurement reality (seconds for disruption response, minutes for tactical tuning, hours for strategic portfolios), and
- Reporting standards (objective decomposition, feasibility/repair rates, runtime distributions such as median/p90/p99, robustness across scenarios, and diversity of near-optimal candidates).

This contribution addresses a common weakness in quantum application papers—high-level evaluation plans without executable detail—by providing a reviewer-acceptable experimental protocol that supports fair comparison to strong classical approaches.

1.1.5. Governance guidance incorporating post-quantum cryptography migration planning for procurement ecosystems.

Unlike many quantum optimization papers that focus narrowly on algorithms, this paper adds an enterprise governance dimension that is highly relevant to travel procurement because procurement systems manage PII, payment data, traveler itineraries, and supplier contracts and integrate across a broad vendor ecosystem (TMCs, airlines, hotels, payment providers, identity systems). We incorporate post-quantum cryptography (PQC) readiness as a parallel requirement for organizations exploring quantum technologies. Concretely, the paper outlines governance actions including cryptographic inventory, cryptographic agility requirements, vendor security roadmaps, and contract language that anticipates algorithm migration. This contribution ensures the quantum-enhanced procurement discussion is enterprise-complete: it recognizes that adopting quantum capabilities intersects with security obligations and long-lived data risks, and it provides a practical starting point for procurement/IT alignment.

2. Background and Related Work

2.1. Quantum Optimization Paradigms

Quantum annealing (QA) targets minimizing an Ising/QUBO objective. It is attractive for binary decision problems but constrained by embedding overhead, coefficient precision, connectivity, and hardware noise.

Gate-based variational algorithms, such as QAOA, use a quantum-classical loop that tunes circuit parameters to minimize an objective. QAOA performance is sensitive to noise, optimizer choice, and ansatz depth, motivating research on optimizer strategies and metaheuristics.

2.2. Transportation and Logistics Evidence

Work on flight gate assignment shows a realistic quadratic assignment structure and documents extraction of instance sizes small enough for hardware, plus mitigation strategies for coefficient precision. This is relevant because travel procurement includes similar portfolio/assignment substructures.

2.3. Benchmarking Evidence and Implications

Large-scale benchmarks in other optimization domains caution against assuming near-term quantum advantage. A recent extensive benchmark in quantum portfolio optimization concluded that classical MIP solves all instances quickly and that problem-tailored heuristics outperform quantum approaches under fixed runtimes. Consequently, quantum-enhanced procurement should be framed as hypothesis-driven and evaluated against strong classical baselines with transparent reporting.

3. Travel Procurement as an Optimization Problem

3.1. Decision Layers

- **Strategic (months/quarters):** preferred supplier selection by market, contract rate/discount design, volume commitments, policy architecture.
- **Tactical (weeks):** policy parameter tuning, supplier performance monitoring, market overrides, demand shaping.

- **Operational (minutes/hours):** offer selection, re-shopping, disruption re-accommodation and duty-of-care constraints.

3.2. Objective Functions

A procurement objective is typically multi-objective, e.g.:

$$\min \underbrace{C(\mathbf{x})}_{\text{expected cost}} + \lambda_1 \underbrace{T(\mathbf{x})}_{\text{time/friction}} + \lambda_2 \underbrace{V(\mathbf{x})}_{\text{noncompliance/leakage}} + \lambda_3 \underbrace{E(\mathbf{x})}_{\text{emissions}} - \lambda_4 \underbrace{R(\mathbf{x})}_{\text{rebates/benefits}}$$

where $\mathbf{x}$encodes supplier selections, policy rules, or allocations.

4. QUBO Formulation of Travel Procurement

4.1. Binary Decision Variables

Binary encoding covers many procurement choices:

- hotel $i$is preferred in city c: $x_{c,i} \in \{0,1\}$
- airline $a$is preferred in market m: $y_{m,a} \in \{0,1\}$
- policy rule enabled (binary switch): $z_k \in \{0,1\}$

4.2. Constraint Encoding via Penalties

QUBO requires unconstrained form; constraints are incorporated using penalty terms. Selecting exactly Kitems:

$$P(\sum_{i=1}^{N} x_i - K)^2$$

where $P > 0$is tuned to discourage infeasible solutions.

4.3. Quadratic Interactions in Procurement

Procurement often has genuine pairwise interactions: overlap penalties (redundant hotels), complementarity bonuses (coverage), and policy conflicts (combinations of rules that induce leakage). These map naturally to quadratic terms $Q_{ij}x_i x_j$.

4.4. Worked Numerical Example: Preferred-Hotel Portfolio (10 Variables)

This section provides a concrete QUBO instance with explicit coefficients to reduce abstraction.

4.4.1. Scenario

A company selects **exactly two** preferred hotels in a city. There are **five candidate hotels** (H1–H5). Procurement wants:

- minimize expected nightly cost (lower is better),
- ensure **coverage** of two neighborhoods (A and B) using “at least one hotel in each neighborhood,”
- encourage sustainability certification,
- penalize selecting two hotels that are geographically redundant (overlap penalty).

4.4.2. Variables (5 selection + 5 neighborhood slack binaries = 10 total)

Selection variables:

$$x_i \in \{0,1\}, i \in \{1, \dots, 5\}$$

Neighborhood constraint uses binary slack variables $s_A, s_B \in \{0,1\}$to represent “at least one” constraints in squared form (standard trick for QUBO). To keep the worked example at 10 variables, we encode coverage with one slack per neighborhood plus three auxiliary "balancing" binaries u_1, u_2, u_3 for coefficient scaling in the cardinality penalty (illustrative). Total variables:

$$\mathbf{v} = [x_1, x_2, x_3, x_4, x_5, s_A, s_B, u_1, u_2, u_3]$$

4.4.3. Data (illustrative coefficients)

- Expected nightly costs (scaled):
  - H1: 160, H2: 140, H3: 150, H4: 130, H5: 170
- Sustainability bonus (negative cost):
  - Certified: H2, H4 ⇒ bonus = 15 each
- Neighborhood membership:
  - A: {H1, H2, H4}, B: {H3, H5}
- Redundancy overlap penalty:
  - (H1,H2) overlap high ⇒ +20 if both selected
  - (H2,H4) overlap medium ⇒ +10 if both selected

4.4.4. QUBO Objective

We minimize:

$$\min \underbrace{\sum_{i=1}^{5} c_i x_i}_{\text{cost}} - \underbrace{\sum_{i=1}^{5} b_i x_i}_{\text{sustainability}} + \underbrace{P_K\Big(\sum_{i=1}^{5} x_i - 2\Big)^2}_{\text{pick exactly 2}} + \underbrace{P_A\Big(1 - \sum_{i\in A} x_i - s_A\Big)^2}_{\text{cover A}} + \underbrace{P_B\Big(1 - \sum_{i\in B} x_i - s_B\Big)^2}_{\text{cover B}} + \underbrace{20x_1x_2 + 10x_2x_4}_{\text{overlap}}$$

Let $P_K = 60$, $P_A = P_B = 80$. Sustainability bonuses: $b_2 = b_4 = 15$, others 0.

4.4.5. Expanded Interpretation (what the QUBO "means")

- The **cardinality penalty** forces exactly two hotels.
- The **coverage penalties** force at least one from neighborhood A and at least one from neighborhood B (otherwise the squared term adds ≥80).
- Overlap terms discourage redundant pairs.
- Sustainability lowers effective cost for certified hotels.

4.4.6. Solution Insight (manual reasoning)

Neighborhood B has only H3 or H5. So one selected must be H3 or H5. Neighborhood A requires H1, H2, or H4. So the decision reduces to selecting one from A and one from B. Evaluate plausible pairs (ignoring penalties since feasible pairs satisfy constraints):

- H4 + H3: (130 −15) + 150 = **265**
- H2 + H3: (140 −15) + 150 = **275**
- H1 + H3: 160 + 150 = **310**
- H4 + H5: (130 −15) + 170 = **285**
- H2 + H5: (140 −15) + 170 = **295**
- H1 + H5: 160 + 170 = **330**

Overlap penalties do not apply (no H1+H2, no H2+H4) in these pairs except if A-pair overlap were used; here each pair is A+B so overlap is irrelevant. The best feasible pair is **H4+H3** with objective ≈ **265**, illustrating how procurement preferences become a small QUBO.

4.4.7. Why this example matters for hardware intuition

Even this tiny instance already uses **10 binary variables**. Scaling is rapid:

- If a city has 60 candidate hotels and you select K=6, that is 60 decision variables plus multiple constraint/auxiliary variables.
- Across 30 cities, a naïve monolithic formulation quickly reaches **thousands of binaries**, motivating decomposition (solve per city) and hybrid architectures.

---

5. QUBO vs. MIP: Practical Trade-Offs for Travel Procurement

5.1. When MIP is Strong

MIP excels when constraints are naturally linear and instances fit memory/time budgets. Modern solvers provide optimality bounds and strong feasibility guarantees—highly valuable for procurement defensibility.

5.2. When QUBO Can Be Attractive

QUBO is convenient when the core decision is binary with many interaction terms and when one seeks rapid approximate solutions or diversified candidate portfolios. QUBO also aligns with quantum annealing hardware and QAOA formulations.

5.3. Procurement-Relevant Implication

For near-term enterprise use, the most credible role of QUBO/quantum is:

- candidate generation (diverse near-optimal portfolios),
- acceleration of repeated scenario solving under tight time windows,
- solving decomposed subproblems embedded within a classical governance loop.
- Claims of superiority should be framed as empirical hypotheses and tested against MIP/CP and tailored heuristics.

---

6. Hybrid Quantum–Classical Architectures

6.1. Pattern 1: Quantum-Assisted Candidate Generation

Classical preprocessing filters options; quantum (annealing or QAOA) generates multiple low-energy solutions; classical post-processing repairs constraints and ranks candidates by full KPI scorecard (cost, compliance, traveler friction, emissions).

6.2. Pattern 2: Decomposition + Quantum Subproblems

Decompose by city, region, top OD markets, or traveler personas. Solve sub-QUBOs (city-level portfolios, market-level airline mixes) while classical coordination enforces global commitments (e.g., total share targets).

6.3. Pattern 3: Operational Disruption Re-Accommodation

In irregular operations, allocate travelers to rebooking/hotel options under duty-of-care, fairness, and budget constraints. Classical filtering reduces candidate sets; quantum explores combinatorial allocations; classical integration executes bookings.

---

7. Priority Use Cases

This section identifies procurement use cases in which quantum-enhanced optimization is most plausibly valuable, based on three criteria: (i) intrinsic combinatorial complexity with binary decision structure, (ii) dense interaction effects that are difficult to linearize cleanly, and (iii) decision contexts where near-optimal solutions or rapid scenario exploration are more valuable than proven global optimality.

7.1. Preferred Hotel Portfolio Design

(Coverage, sustainability, redundancy control)

Preferred hotel portfolio design is a core strategic procurement activity in which organizations select a limited number of hotels per city or region to be designated as "preferred" within booking and policy systems. The objective is to minimize expected lodging cost while ensuring adequate geographic coverage, traveler satisfaction, sustainability alignment, and contractual diversity.

From an optimization perspective, this problem exhibits several features well suited to QUBO formulations:

- **Cardinality constraints**, such as selecting exactly K hotels per city.
- **Coverage constraints**, ensuring representation across neighborhoods, office locations, or transit hubs.
- **Redundancy penalties**, discouraging the selection of multiple hotels serving essentially the same micro-market.
- **Sustainability incentives**, such as minimum shares of certified or low-emissions properties.

These elements introduce quadratic interactions among selection variables, particularly when redundancy and complementarity effects are modeled explicitly. Classical mixed-integer formulations can handle such problems but may require linearization or large numbers of auxiliary variables as interaction density increases. A hybrid quantum–classical approach can be used to generate diverse near-optimal portfolios at the city level, which are then evaluated and repaired classically to ensure global compliance with brand caps, negotiated terms, and auditability requirements.

7.2. Airline Share Allocation Under Commitments and Service Constraints

Airline procurement frequently involves allocating traveler share across preferred carriers while meeting contractual share-of-wallet commitments and maintaining acceptable service levels. Decisions are made across many origin–destination markets and are subject to constraints such as minimum global or regional share, maximum reliance on a single carrier, connection limits, and travel-time thresholds.

This use case resembles portfolio optimization but includes additional network and service-quality interactions that complicate classical linear modeling. For example, allocating additional share to a carrier in one market may improve commitment compliance but degrade traveler experience in another due to schedule or connection penalties.

Quantum-enhanced optimization may contribute by:

- Exploring **combinatorial allocations** of market-level preferences that satisfy global commitments.
- Capturing **interaction effects** between markets and carriers in a quadratic objective.
- Producing **multiple feasible share-allocation scenarios** that procurement teams can use during negotiations.

Given existing benchmark results showing strong classical performance in portfolio optimization, this use case is best framed as an exploratory hybrid application where quantum routines assist in scenario generation or diversification rather than guaranteed cost improvement.

7.3. Policy Parameter Tuning

(Discrete approximations to continuous thresholds)

Corporate travel policies include numerous parameters—such as advance purchase windows, price tolerance thresholds for non-preferred bookings, cabin eligibility rules, and approval triggers—that strongly influence cost, compliance, and traveler behavior. These parameters are often tuned manually or via trial-and-error simulation.

From a modeling standpoint, many policy parameters are **continuous** (e.g., "allow non-preferred booking if price is within δ%"), which are not directly compatible with QUBO. However, practical policy tuning often considers a **finite set of candidate values** (e.g., $\delta \in \{5\%, 10\%, 15\%, 20\%\}$), which can be encoded as one-hot binary variables.

In this discretized form, policy tuning becomes a combinatorial selection problem with interaction effects (e.g., the combined impact of price thresholds and advance purchase rules on compliance). Quantum-enhanced methods can be used to explore combinations of policy settings that balance cost savings against leakage and traveler friction, while classical simulation evaluates downstream behavioral effects. This hybrid approach allows procurement teams to test policy bundles rather than isolated parameter changes.

7.4. Negotiation Scenario Generation

(Producing multiple credible alternatives)

Negotiation with airlines and hotels is a central procurement activity, and its effectiveness depends heavily on the availability of **credible alternative portfolios**. Rather than seeking a single "optimal" solution, procurement teams benefit from understanding a frontier of near-optimal options that differ in supplier mix, coverage, and risk exposure.

This use case is particularly well aligned with quantum-enhanced optimization because:

- The goal is **solution diversity**, not a single global optimum.
- Quantum annealing and variational methods naturally generate multiple low-energy solutions.
- Interaction-rich objectives (e.g., overlapping supplier coverage, rebate thresholds) can be encoded compactly in QUBO form.

Hybrid systems can use quantum routines to generate candidate portfolios that are then evaluated classically for contractual feasibility and negotiation leverage. The resulting solution set supports data-driven negotiation strategies and sensitivity analysis ("what if we shift 3% share from carrier A to B?").

7.5. Disruption Mass Re-Accommodation

(Assignment under severe time constraints)

During large-scale disruptions—such as weather events, strikes, or system outages—organizations must rapidly re-accommodate travelers while satisfying duty-of-care obligations, minimizing incremental cost, and maintaining fairness across travelers. This is an operational assignment problem characterized by severe time pressure, rapidly changing availability, and hard constraints.

The problem structure includes:

- Binary assignment decisions (traveler-to-option).
- Capacity constraints (limited seats or hotel rooms).
- Equity and priority rules (e.g., medical needs, executive travelers).
- Time-dependent feasibility constraints.

Quantum-enhanced optimization may assist by rapidly exploring feasible assignment combinations within a restricted candidate set generated by classical filtering. Even modest improvements in solution quality or speed can have high operational value in disruption contexts. This use case aligns with prior quantum work in assignment problems, such as gate allocation, while remaining grounded in realistic enterprise constraints.

8. Data, Modeling, and Preprocessing

8.1. Data Requirements

Booking history, supplier content (rates/fare rules), quality proxies, demand forecasts, emissions factors, and policy/compliance telemetry.

8.2. Preprocessing for Instance Reduction

Dominance pruning, market clustering, and coefficient scaling are essential for both quantum hardware limits and classical solver efficiency.

8.3. Clarifying Policy Parameter Types

- **Binary policy parameters**: on/off switches (e.g., "require preferred hotel").
- **Continuous thresholds**: e.g., "allow out-of-policy if price is within δ%."

To fit QUBO, continuous thresholds must be discretized (e.g., choose among 10–20 candidate δ values) and encoded with one-hot binaries; otherwise, a hybrid approach keeps continuous tuning in classical optimization while using quantum for discrete structure.

9. Operationalized Evaluation Methodology

9.1. Benchmark Instance Definitions (Pilot-Scale)

We recommend publishing a benchmark suite with at least two families:

9.1.1. Family H (Hotels):

- Cities: 10, 30, 100
- Candidates per city: 25, 60, 120
- K preferred per city: 3–8
- Constraints: coverage zones (2–6), sustainability share, brand caps, minimum quality, rate-fence rules
- Interaction graph density: low/medium/high redundancy

9.1.2. Family A (Airline share allocation):

- Markets: 50, 200, 800 OD markets
- Airlines per market: 3–8
- Constraints: global share commitments (region + total), service-level constraints (connection count, time), policy tiers
- Scenario sets: 20–200 demand/price scenarios

9.2. Typical Problem Sizes and Scaling

- City-level hotel subproblem (60 candidates): **60 binaries** + coverage/auxiliary variables (≈10–40) ⇒ **70–100 variables** per city.
- 30-city program: decomposed approach solves 30 subproblems of ∼100 variables each (parallelizable).
- Monolithic coupling (global brand caps, global sustainability): introduces cross-city couplings and can reach **2,000–5,000 binaries**, motivating Lagrangian relaxation or Benders-like decomposition (classical) plus quantum for subproblem solving.

9.3. Solver Baselines (Required for Credibility)

Report results against:

- **MIP** (e.g., Gurobi/CPLEX or open-source equivalents)
- **CP-SAT** (e.g., OR-Tools) for constraint-heavy variants
- **Metaheuristics**: simulated annealing, tabu search, large neighborhood search
- **Quantum annealing / hybrid annealing** (where available)
- **QAOA** (simulator + hardware runs when feasible) with optimizer variants

9.4. Runtime Budgets (Reflecting Procurement Reality)

- Strategic portfolio: minutes to hours (batch).
- Tactical policy tuning: minutes (iterative weekly).
- Operational disruption: seconds to <5 minutes, depending on severity.

9.5. Reporting Protocol

For each instance family and size:

- Solution quality (objective + KPI breakdown),

- Constraint violation rate (should be zero after repair),
- Runtime distribution (median, p90, p99),
- Robustness across scenarios,
- Number of distinct near-optimal candidates produced (for negotiation support).

10. Enterprise Governance and Security Considerations

10.1. Post-Quantum Cryptography (PQC) Readiness

Procurement systems integrate with TMCs, airlines, hotels, payment rails, identity providers, and travel risk vendors—exposing long-lived sensitive data. NIST finalized initial PQC standards in 2024, and U.S. government guidance emphasizes early migration planning (inventory, roadmap, vendor engagement) to mitigate "harvest now, decrypt later" risks.

10.2. Vendor Ecosystem Implications

A quantum-enhanced roadmap should include: cryptographic agility requirements, PQC transition plans for vendors, and contract language supporting algorithm migration.

11. Limitations and Open Challenges

Despite the potential of quantum-enhanced optimization for travel procurement, significant limitations and open challenges remain. This section explicitly articulates these challenges to frame the contribution of this work as exploratory and hypothesis-driven rather than prescriptive.

11.1. Hardware Capacity and Embedding Constraints

Current quantum hardware—both quantum annealers and gate-based devices—imposes strict limitations on problem size, connectivity, and coefficient precision. In practical travel procurement settings, even moderately scoped problems can involve hundreds to thousands of binary decision variables once supplier choices, coverage constraints, policy rules, and auxiliary variables are included.

Quantum annealers require problem graphs to be embedded onto hardware with limited qubit connectivity, often introducing substantial overhead in the form of chained qubits. This embedding overhead can multiply the effective qubit requirement, rendering otherwise moderate QUBO instances infeasible without aggressive problem reduction. Gate-based approaches face analogous constraints in circuit depth, qubit counts, and noise accumulation, particularly for higher-depth QAOA circuits needed to capture complex interaction landscapes.

As a result, direct monolithic formulations of enterprise-scale procurement problems are currently impractical. Effective application of quantum methods requires decomposition strategies (e.g., per-city, per-market, or per-persona subproblems), hybrid coordination with classical solvers, and careful selection of subproblem scope. Designing such decompositions while preserving global procurement constraints remains an open research challenge.

11.2. Sensitivity to Penalty Coefficient Tuning

QUBO formulations encode constraints via penalty terms, but the selection of penalty magnitudes is nontrivial and highly problem-dependent. If penalty coefficients are set too low, the optimization may return low-energy solutions that violate essential procurement constraints (e.g., insufficient coverage or unmet share commitments). Conversely, excessively large penalties can lead to numerical stiffness, in which the objective landscape becomes dominated by constraint terms, reducing effective resolution of cost and preference trade-offs.

This sensitivity is particularly acute in travel procurement, where coefficients naturally span multiple orders of magnitude (e.g., nightly hotel rates, annual demand weights, emissions multipliers, and rebate values). Quantum hardware further constrains feasible coefficient ranges due to precision limits. While classical solvers also require careful scaling, they typically provide richer diagnostic feedback (e.g., infeasibility certificates or dual information) than current quantum workflows.

Developing systematic, automatable penalty calibration methods—potentially leveraging classical pre-solving, adaptive penalties, or learning-based approaches—remains an important open problem for reliable quantum-enhanced procurement optimization.

11.3. Benchmark Realism and Evaluation Integrity

A persistent challenge in applied quantum optimization research is the gap between benchmark problems and real enterprise decision contexts. Small or artificially structured instances ("toy problems") may fit on current quantum hardware but fail to capture the constraint density, heterogeneity, and uncertainty inherent in real travel procurement.

Meaningful evaluation requires:

- Realistic instance sizes and constraint combinations,
- Comparison against state-of-the-art classical baselines (mip, cp-sat, and advanced heuristics),
- Transparent runtime budgets aligned with procurement decision cycles, and
- Reporting of solution robustness and feasibility, not just objective value.

Without such rigor, performance claims risk overstating practical relevance. For travel procurement in particular—where classical solvers are already highly effective—quantum-enhanced methods must demonstrate incremental value under realistic conditions, such as faster scenario exploration or improved solution diversity. Establishing shared, domain-relevant benchmark suites remains an open and necessary research direction.

11.4. End-to-End Integration and Organizational Overhead

Even if a quantum or hybrid solver performs well on an isolated optimization task, procurement value depends on successful integration into end-to-end enterprise workflows. Travel procurement systems interface with demand forecasting, booking tools, policy engines, supplier content feeds, analytics platforms, and governance processes. Each interface introduces latency, data quality issues, and operational constraints that are not captured in standalone optimization benchmarks.

Additionally, procurement decisions must be explainable to stakeholders, auditable for compliance, and implementable within existing contractual and technological frameworks. Solutions that are difficult to interpret or that require significant manual intervention may fail to deliver value regardless of computational novelty.

From an organizational perspective, adopting quantum-enhanced approaches also entails skills development, vendor coordination, and alignment with IT security and compliance requirements. These integration and change-management costs may outweigh algorithmic gains if not carefully managed. Consequently, solver novelty alone is insufficient; measurable procurement impact depends on holistic system design and organizational readiness.

12. Conclusion

This paper has examined the applicability of quantum optimization techniques to enterprise travel procurement through a pragmatic, enterprise-centered lens. Rather than positioning quantum computing as a disruptive replacement for established mixed-integer programming (MIP), constraint programming (CP), and heuristic methods, we frame quantum-enhanced travel procurement as a hybrid, hypothesis-driven decision-support paradigm. In this paradigm, quantum routines complement classical optimization by accelerating combinatorial search, improving solution diversity, or enabling faster exploration of complex scenario spaces under realistic time constraints.

A central contribution of this work is the grounding of abstract quantum optimization concepts in concrete procurement constructs. By translating preferred supplier selection, coverage requirements, sustainability incentives, and redundancy penalties into explicit QUBO formulations—and by providing a worked numerical example—this paper reduces the modeling gap between theoretical quantum algorithms and real procurement decision problems. This explicit grounding clarifies both the opportunities and the limitations of quantum methods when confronted with the scale, heterogeneity, and governance requirements of enterprise travel programs.

Equally important, the paper emphasizes evaluation rigor as a prerequisite for credible progress. By operationalizing benchmark instance definitions, specifying representative problem sizes, identifying mandatory classical baselines, and aligning runtime budgets with procurement decision cycles, we propose an evaluation framework that enables fair and transparent comparison between quantum-enhanced and classical approaches. This focus directly addresses a common weakness in applied quantum research—namely, the reliance on simplified benchmarks and weak baselines—and supports a more honest assessment of where quantum methods may deliver incremental value.

Beyond algorithmic considerations, the paper situates quantum-enhanced optimization within the broader enterprise ecosystem. Travel procurement systems are deeply integrated with booking platforms, supplier interfaces, analytics pipelines, and governance processes, and they manage sensitive traveler and contractual data. By incorporating post-quantum cryptography readiness into the discussion, this work recognizes that experimentation with quantum technologies intersects with long-term security and compliance obligations. This governance perspective ensures that the proposed research agenda is not only technically sound but also operationally responsible.

Taken together, the contributions of this paper define a realistic adoption pathway for quantum-enhanced travel procurement. In the near term, value is most likely to arise from hybrid applications that focus on decomposed subproblems, scenario generation, and decision-support augmentation rather than guaranteed optimality or large-scale quantum advantage. Over the longer term, advances in hardware capacity, algorithm design, and benchmarking standards may expand the scope of feasible applications. Until then, progress will depend on disciplined experimentation, transparent reporting, and close alignment with procurement realities.

In conclusion, quantum-enhanced travel procurement should be viewed not as an endpoint but as an evolving research and innovation trajectory. By articulating clear use cases, explicit models, rigorous evaluation practices, and enterprise governance considerations, this paper provides a publication-ready foundation for future research and for carefully scoped enterprise pilots that can test, refine, and validate the role of quantum optimization in corporate travel management.

References

[1] T. Stollenwerk, E. Lobe, and M. Jung, “Flight gate assignment with a quantum annealer,” 2018. https://arxiv.org/pdf/1811.09465

[2] Waditwar, P. (2024e). The intersection of strategic sourcing and artificial intelligence: a paradigm shift for modern organizations. Open Journal of Business and Management, 12(06), 4073–4085. https://doi.org/10.4236/ojbm.2024.126204

[3] Waditwar, P. (2025h). Reimagining procurement payments: From transactional bottlenecks to strategic value creation. World Journal of Advanced Research and Reviews, 28(1), 588–598. https://doi.org/10.30574/wjarr.2025.28.1.3480

[4] Lykov, D., Wurtz, J., Poole, C. et al. Sampling frequency thresholds for the quantum advantage of the quantum approximate optimization algorithm. npj Quantum Inf 9, 73 (2023). https://doi.org/10.1038/s41534-023-00718-4

[5] AI Agents Unleashed: Transforming Procurement and Supply Chains with Intelligence and Initiative: 9798308460763: Computer Science Books @ Amazon.com. (n.d.-b). https://www.amazon.com/Agents-Unleashed-Transforming-Procurement-Intelligence/dp/B0DV5FBJL6

[6] A. Pellow-Jarman et al., “QAOA performance in noisy devices: The effect of classical optimizers and ansatz depth,2023” https://arxiv.org/html/2407.12242v4

[7] Zarate-Carbajal, J., Ruiz-Cruz, R., & Sánchez-Torres, J. D. (2024). Strategic Decisions in Corporate Travel: Optimization through Decision Trees. Mathematics, 12(23), 3741. https://doi.org/10.3390/math12233741

[8] IATA, “New Distribution Capability (NDC),” International Air Transport Association, 2025. [Online]. Available: https://www.iata.org/en/programs/airline-distribution/retailing/ndc/ (accessed 2025-12-22).

[9] NIST, “NIST releases first 3 finalized post-quantum encryption standards,” National Institute of Standards and Technology, Aug. 13, 2024. [Online]. Available: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards (accessed 2025-12-22).

[10] CISA, NSA, and NIST, “Quantum-readiness: Migration to post-quantum cryptography,” Fact Sheet, Aug. 2023. [Online]. Available: https://www.cisa.gov/resources-tools/resources/quantum-readiness-migration-post-quantum-cryptography (accessed 2025-12-22).