https://pubsonline.informs.org/journal/ijds

INFORMS JOURNAL ON DATA SCIENCE

Vol. 00, No. 0, Xxxxx 0000, pp. 000–000

Submitted to *INFORMS Journal on Data Science*

ISSN 2694-4022, EISSN 2694-4030

How INFORMS can Contribute to the Second Quantum Revolution

Sridhar Tayur
Quantum Technologies Group, Tepper School of Business, Carnegie Mellon University, Pittsburgh PA 15213, stayur@cmu.edu

Abstract. The Second Quantum Revolution—comprising Quantum Computing, Quantum Communications, and Quantum Sensing—holds the promise of further abilities to improve the human condition, security, and sustainability, and hence mass prosperity, across the world, in a variety of ways, including improving health (through better diagnosis as well as new drug development, via exciting applications of Quantum Machine Learning and Quantum Sensing) as well as by enhancing communication and cybersecurity. For certain computational workloads and sensing tasks, quantum systems have demonstrated or projected energy advantages over purely digital methods at equivalent precision, though system-level comparisons remain an active and unsettled research area. Seen through a data-science lens, quantum technologies are not only physical systems but also *data-generating systems* whose calibration, validation, and control depend on statistical inference, hypothesis testing, and learning under uncertainty. Integer Programming, Queuing, Markov Decision Processes, and Semidefinite Programs are some of the fundamental methodologies in Operations Research (OR) and Management Science (MS) that are used to tackle practical applications from Business (Supply Chain, Finance), Engineering (Communication Networks), and Medicine (Cancer Genomics, Image Recognition). At the Tepper Quantum Technologies Group, we are exploring the twin questions: (a) what can quantum do for OR/MS and (b) what can OR/MS do for quantum. This article—a companion to my 2025 INFORMS Keynote and not intended as a comprehensive survey of the field, but rather as a selective perspective organized around illustrative examples from our group—provides a brief summary of our research spanning Algorithms, Hardware, and Applications (AHA). I hope that it helps illustrate how we can contribute to the Second Quantum Revolution.

Key words: Quantum Mechanics; Semidefinite Programming; Queueing; Markov Decision Process; Integer Programming; Data Science

1. Introduction

Quantum Mechanics (QM) began in 1900, matured in 1925, and 2025 was labeled the International Year of Quantum Science and Technology by UNESCO. INFORMS also celebrated the 100th anniversary of QM with a variety of activities, including my Keynote at our annual meeting in Atlanta (October 2025). The purpose of this brief companion article is to provide illustrative examples on how INFORMS research[1] can contribute to the Second Quantum Revolution that is underway, with the goal of encouraging more researchers to join this exciting opportunity to advance the frontiers of scientific and technological capability.

Scope of this article. This paper is not a systematic survey [2] of the field. It is the author's selective perspective, organized around a small map from four OR/MS methodology families to several quantum-technology themes, and illustrated with examples drawn primarily from the Tepper Quantum Technologies Group. For completeness, we note several important quantum-operational areas not covered here: quantum error correction and fault-tolerant operations (scheduling, decoding, resource allocation for logical-qubit codes); device calibration and closed-loop quantum control; and verification, validation, and benchmarking protocols for quantum hardware. These are active research directions where OR/MS contributions are also natural.

A central theme is that Operations Research and Management Science are not merely downstream beneficiaries of quantum technologies, but must be foundational contributors to their success. As quantum systems scale, they confront challenges that are not primarily physical but *decisional*: fragility under noise, non-additivity of resources, stochastic degradation over time, and the need for optimal control under severe uncertainty. These are precisely the regimes where OR/MS has historically made its greatest contributions. The Second Quantum Revolution will therefore not be realized by physics alone, but by the co-evolution of physical insight and decision science, with INFORMS methodologies playing a central and enabling role.

Quantum randomness and data science. A feature of quantum systems that is essential for OR/MS researchers is that observed outcomes are *random variables*. Quantum measurement is probabilistic at a fundamental level: even a perfectly executed unitary evolution yields a stochastic outcome upon measurement. In real devices, additional randomness arises from noise, drift, and finite sampling. It is therefore important to distinguish three sources of stochasticity: (i) deterministic unitary evolution, which is reversible and noise-free; (ii) stochastic measurement outcomes, which are irreducible and governed by Born's rule; and (iii) additional stochasticity from device noise

---

[1] I do not cover the business of quantum computing here: see Sodhi and Tayur (2023) and Velu et al. (2025). Our community has already begun to recognize that Operations Research can play an important role in advancing Quantum Computing; see Coffrin et al. (2025).

[2] Readers seeking a comprehensive view should consult the systematic mapping study by Gomes et al. (2024) and the references therein.

Table 1 Map between INFORMS methodologies and Second Quantum Revolution topics.

INFORMS Methodology	Second Quantum Revolution Topic
Semidefinite Programs (SDPs)	Quantum Information Science / Quantum Sensing
Queueing Models	Buffering in Quantum Communications
Markov Decision Process (MDP)	Entanglement Switch
Integer Programming	Quantum (and Quantum-inspired) Computing

and finite sample sizes, which introduces estimation uncertainty on top of fundamental quantum randomness. The practical consequence is that many quantities used in sensing, networking, and near-term quantum computing—fidelities, capacities, error rates, and decoherence timescales—are not directly observed but must be *inferred* from finite samples. Hypothesis testing, detection, and estimation are therefore first-class primitives[3] in quantum information science, not mere post-processing steps. This data-scientific character of quantum systems recurs throughout the sections below and represents a natural entry point for the INFORMS community.

Structure of the paper. Table 1 provides a map between OR/MS methodologies and the topics of the Second Quantum Revolution. Section 2 discusses how Semidefinite Programs provide a natural bridge to Quantum Information Science (QIS), forming the basis for many Quantum Sensing applications. Section 3 illustrates two applications in Quantum Communications: buffering (using queueing theory) and the Entanglement Switch (modeled as a Markov Decision Process). Section 4 covers quantum and quantum-inspired computing and demonstrates how nonlinear binary constrained optimization models can be approached in entirely new ways. Section 5 concludes with suggestions for further reading.

1.1. The First Quantum Revolution

The First Quantum Revolution, spanning roughly from the 1940s to the 1980s, transformed modern civilization through technologies that exploited quantum mechanical principles without requiring direct manipulation of individual quantum states. The transistor, invented in 1947 at Bell Labs, used quantum mechanics to control electrical current flow through semiconductors, replacing bulky vacuum tubes and enabling the miniaturization of electronics. This breakthrough paved the way for the integrated circuit in 1958, which packed multiple transistors onto a single silicon chip, exponentially increasing computing power while reducing size and cost—a trend that would follow Moore's Law for decades. Meanwhile, the laser, demonstrated in 1960, harnessed quantum principles of

---

[3] For a signal-processing treatment that addresses quantum measurement, detection, and estimation directly, see Eldar and Oppenheim (2002). For a modern algorithmic primitive that illustrates the breadth of quantum algorithm design, see Motlagh and Wiebe (2024).

stimulated emission to produce coherent light with unprecedented precision, revolutionizing communications, manufacturing, medicine, and research. The charge-coupled device (CCD), invented in 1969, exploited quantum effects in semiconductors to convert light into electronic signals with remarkable sensitivity, becoming the cornerstone of digital imaging in cameras, telescopes, and scientific instruments. Finally, magnetic resonance imaging (MRI), developed for medical use in the 1970s, applied the quantum mechanics of nuclear spin to create detailed images of soft tissues inside the human body, transforming medical diagnostics without invasive procedures or harmful radiation. Together, these innovations laid the technological foundation for the information age and modern healthcare.

1.2. Quantum Information Science and the Second Quantum Revolution

Quantum Information Science (QIS) emerged in the late twentieth century when researchers realized that quantum mechanics could itself be harnessed as a computational and informational resource. Pioneering work in the 1980s and 1990s revealed that quantum phenomena like *superposition*, where particles exist in multiple states simultaneously, and *entanglement*, where particles remain correlated across vast distances, could enable fundamentally new ways of processing and transmitting information. This theoretical foundation sparked the Second Quantum Revolution, characterized by technologies that actively create, manipulate, and exploit individual quantum states rather than merely applying quantum principles in bulk materials.

Quantum sensing leverages the extreme sensitivity of quantum systems to external perturbations, enabling measurements of magnetic fields, gravity, time, and other quantities with unprecedented precision for applications ranging from brain imaging to underground resource detection. Quantum communication exploits entanglement and quantum key distribution to create provably unbreakable encryption, since any attempt to intercept quantum-encoded messages inevitably disturbs the quantum states and reveals the eavesdropping. Quantum computing seeks to harness superposition and entanglement to perform *certain structured problems* exponentially faster than classical computers—with proven advantages for problems such as Hamiltonian simulation and the hidden subgroup problem—promising breakthroughs in drug discovery, materials science, and cryptography. The degree to which quantum computers will benefit combinatorial optimization problems, which are central to OR/MS, remains an active and unresolved question; we return to this in Section 4.

Unlike the First Quantum Revolution's passive use of quantum effects, this second wave represents humanity's attempt to engineer and control the quantum world with exquisite precision. Because quantum devices generate stochastic measurement records whose statistical analysis is inseparable from the underlying physics, the data-scientific perspective, statistical inference, learning, and decision-making under uncertainty, is not incidental but structural.

2. SDPs as a Natural Pathway for INFORMS Researchers to Enter QIS

Among the many entry points into Quantum Information Science, semidefinite programming occupies a special position. This is not accidental. The near-term quantum landscape is dominated by noisy, intermediate-scale quantum (NISQ) devices where uncertainty, mixed states, and partial information are the norm rather than the exception. In such regimes, exact characterizations are often unattainable, while bounds, relaxations, and certificates become essential. SDPs provide a mathematically principled and computationally tractable framework for reasoning about precisely these questions, making them a natural bridge between OR/MS and quantum technologies in the NISQ era. In practice, the power of semidefinite relaxations in QIS is amplified by data-driven estimation: noise models, density matrices, and confidence bounds are inferred from finite, imperfect experimental data rather than assumed *a priori*.

Semidefinite programs thus provide an accessible entry point for Operations Research and Computer Science professionals to contribute meaningfully to Quantum Information Science without requiring deep mastery of quantum mechanics. As highlighted in the tutorial by Siddhu and Tayur (2022), SDPs extend the familiar framework of linear programming by replacing element-wise non-negative vector variables with positive semidefinite matrices, preserving many desirable properties such as duality theory and efficient solvability through interior-point methods. What makes SDPs particularly relevant to QIS is that quantum states are mathematically represented as positive semidefinite density operators, and many fundamental problems in quantum computing and communication—including state discrimination, channel capacity, and entanglement detection—naturally formulate as SDPs. This alignment creates a "trading zone" where INFORMS researchers can leverage their existing expertise in convex optimization, cutting-plane methods, and computational techniques to tackle pressing challenges in quantum technologies.

A practical caveat deserves mention. Representing an n-qubit state as a full density matrix requires matrices of dimension 2^n, so solving SDPs over the full state space becomes infeasible beyond roughly 20 qubits. The value of OR/MS algorithmic contributions in this area therefore lies not in solving exponentially large problems directly, but in developing *compact, polynomially scalable formulations* that exploit problem structure to extend tractability. The example below illustrates this point.

2.1. Example: Entanglement Detection Using the PST SDP Hierarchy

The paper by Peña et al. (2025) on entanglement detection exemplifies how INFORMS expertise in first-order methods (FOMs) and interior-point methods (IPMs) can advance quantum information science by making previously intractable problems computationally feasible. The authors introduce the PST hierarchy (named after Peña, Siddhu, and Tayur)—a new SDP hierarchy offering tighter approximations. By developing compact, polynomially scalable formulations using partition mappings and operators, they transform entanglement detection into optimization problems well-suited to both FOMs (Frank–Wolfe, projected gradient, and fast projected gradient) and a custom primal-dual IPM. Their numerical experiments demonstrate that these tailored algorithmic approaches enable solving deeper levels of the SDP hierarchy (up to $k = 15$) than previously possible with off-the-shelf solvers, successfully detecting entanglement in challenging quantum states that were previously intractable. The practical gain is a modest but meaningful extension of the frontier: for structured high-dimensional or noisy states relevant to quantum communication and near-term computation, entanglement can now be certified where it could not be before. This work not only provides practical tools for the quantum computing community to verify entanglement in high-dimensional or noisy systems—a critical requirement for quantum communication and computation—but also demonstrates how careful integration of problem structure with algorithmic design can overcome computational barriers, showcasing the unique value that INFORMS methodologies bring to the Second Quantum Revolution.

3. Quantum Communication: Beyond Classical Information Transfer

Quantum communication replaces the classical transmission of bits with the transmission of quantum bits (qubits), whose ability to exist in superposition and become entangled over long distances enables correlations and security guarantees impossible in any classical network. Because qubits cannot be copied or amplified without disturbance, building a quantum internet requires fundamentally new infrastructure, notably quantum repeaters that create and swap entanglement across many short links to span continental scales (Wehner et al., 2018). Such a network would not replace today's internet but augment it, allowing tasks such as provably secure key distribution, private access to remote quantum computers, ultra-precise clock synchronization, distributed sensing, and coordinated computation between distant quantum processors. Even modest end nodes handling only single qubits can already unlock advantages, since the power of the quantum internet derives from entanglement itself rather than large local quantum computers, making the stepwise development of multinode quantum networks a realistic path toward a future global quantum internet.

Quantum key distribution (QKD) provides provably unbreakable encryption by leveraging the fact that measuring quantum states inevitably disturbs them, making eavesdropping detectable. The capacity of quantum channels depends critically on whether classical or quantum information is being sent, and whether product or entangled encoding strategies are employed. Understanding these capacity limits through quantum Shannon theory is essential for designing practical quantum networks.

A defining challenge of quantum communication systems is the breakdown of assumptions that underpin classical information theory. Noise is no longer independent and identically distributed; memory effects cannot be ignored; and operational decisions—such as buffering, scheduling, or switching—directly influence physical decoherence. This makes queueing theory and Markov decision processes not merely convenient abstractions, but structurally necessary tools for understanding and optimizing quantum communication networks. As quantum networks evolve from point-to-point links to operational systems, their performance becomes inseparable from the data they generate—arrival processes, waiting times, fidelity decay—making statistical modeling and online inference essential complements to information-theoretic analysis.

We briefly discuss two applications[4] next. Taken together, queue-channel models and quantum switch control illustrate a broader lesson: in quantum networks, information-theoretic performance is inseparable from temporal dynamics. Waiting times, storage decisions, and control policies shape the effective noise experienced by quantum states, creating feedback loops between physical degradation and operational choices. These interactions are largely invisible in static or asymptotic models, but become central once quantum communication is treated as a networked, resource-constrained system.

3.1. Quantum Queue-Channels

The paper by Siddhu et al. (2026) addresses a critical yet often overlooked aspect of quantum communication networks: the decoherence that qubits experience while waiting in buffers before transmission. Traditional quantum Shannon theory assumes independent and identically distributed (i.i.d.) noise across channel uses, but in realistic quantum networks, qubits waiting in transmission buffers continue interacting with their environment, experiencing waiting-time-dependent decoherence that introduces correlations in the effective noise process. The authors characterize the classical capacity of unital qubit queue-channels—where the noise depends on queueing dynamics governed

---

[4] For an advanced application, Quantum Money, see Singh et al. (2025).

by arrival and service rates—and demonstrate that, surprisingly, simple product (non-entangled) encoding and decoding strategies achieve maximum capacity for this broad class of channels. We note that the capacity result applies to the setting where the receiver has knowledge of the queue state (i.e., waiting times are known at the decoder as side information); the case where no such side information is available is a harder open problem and a natural direction for future work.

Their key insight is that every unital qubit queue-channel has an induced binary symmetric queue-channel whose Shannon capacity equals the quantum channel's classical capacity, providing explicit capacity-achieving measurement strategies. For the important special case of symmetric generalized amplitude damping (GAD) queue-channels, which model realistic decoherence in optical fibers and quantum memories at finite temperature, they derive closed-form capacity expressions. A crucial design insight emerges: simply maximizing qubit preparation rates can be counterproductive, as high rates cause longer buffer waiting times and increased decoherence, with optimal operating points depending on the physical characteristics of the buffer environment. This work demonstrates that explicitly modeling buffering effects through queue-channel analysis is essential for practical quantum network design, revealing fundamental tradeoffs between throughput and decoherence that are invisible in idealized i.i.d. models. From a data-science perspective, the waiting-time statistics that drive decoherence are themselves random variables to be estimated from observed network traffic, connecting queueing analysis directly to statistical inference.

3.2. Quantum Switches

Kumar et al. (2023) investigate optimal entanglement distillation policies for quantum switches in entanglement distribution networks. The authors model a quantum switch with two clients that generates elementary entanglement (Bell pairs) probabilistically, which are stored in noisy quantum memories where they decohere over time. Using a Markov Decision Process framework, they determine optimal policies for when the switch should wait, perform entanglement distillation (combining two noisy Bell pairs into one higher-quality pair), or perform entanglement swapping (connecting Bell pairs from both clients to create end-to-end entanglement). We note that the optimality conclusions are derived under a specific memory and decoherence model (Markovian depolarizing noise); conclusions may differ under alternative physical models, and the robustness of these policies to model misspecification is an important open question.

The key finding is that for intermediate fidelity thresholds, allowing entanglement distillation significantly improves throughput compared to policies that only use swapping, as distillation helps preserve link quality in the face of decoherence. However, at very low or very high fidelity thresholds, distillation provides minimal benefit. Through simulations, they demonstrate improvements in average throughput, average fidelity, and timing jitter when using the optimal distillation-enabled policy, establishing the value of distillation in mitigating decoherence effects in quantum switches. From a data-science perspective, optimal switch policies are not merely computed offline but must be *learned and adapted* online, as empirical transition probabilities and decoherence rates evolve over time and must be continuously re-estimated from observed operations.

4. Quantum Computing

Quantum computing differs from sensing and communication in one important respect: many of its near-term target problems overlap with Operations Research. Optimization, search, and decision-making under constraints are familiar territory. What changes in the quantum setting is not the objective or the constraints, but the computational substrate: energy landscapes replace feasible regions, noise replaces numerical precision, and driving computational performance is the practical focus. This shift plays directly to OR/MS strengths in modeling, algorithm design, and empirical evaluation.

We emphasize an important caveat: no quantum algorithm with a proven superpolynomial speedup over the best known classical algorithms is currently known for generic NP-hard optimization problems. The “killer applications” of quantum computing with the strongest theoretical backing—proven exponential speedups—involve Hamiltonian simulation and the hidden subgroup problem. For combinatorial optimization, which is central to OR/MS, the picture is more nuanced: quantum approaches show empirical promise on certain structured instances, but unconditional quantum advantage for optimization has not been established. The value of the quantum and quantum-inspired approaches described below therefore lies in their empirical performance on structured instances relevant to practice, rather than in theoretical guarantees.

4.1. Tutorial: A Rapid Introduction to Quantum Computing

The tutorial by Mazumder and Tayur (2025) provides a hands-on introduction to solving Quadratic Unconstrained Binary Optimization (QUBO) models using current quantum computing platforms. This bridges the gap between theoretical quantum computing concepts and practical implementation by covering both IBM’s gate-circuit architecture and D-Wave’s quantum annealing systems. The tutorial is structured to reach undergraduate and graduate students as well as industry professionals, offering five companion notebooks in an associated GitHub repository with detailed implementation guides.

The tutorial demonstrates QUBO formulations for three canonical NP-complete problems and two practical applications. The canonical problems include number partitioning, max-cut, and minimum vertex cover. The practical applications include order partitioning for A/B testing in hedge fund portfolio management, and cancer genomics pathway identification using mutation data from The Cancer Genome Atlas.

After providing a comprehensive foundation in quantum computing (qubits, superposition, quantum gates, Hamiltonians), the tutorial introduces three solution approaches. Simulated annealing serves as a classical baseline, using temperature schedules and probabilistic acceptance to escape local minima. Quantum annealing leverages the adiabatic theorem and quantum tunneling to navigate energy landscapes, gradually evolving from a simple driver Hamiltonian to a problem-encoding cost Hamiltonian. The Quantum Approximate Optimization Algorithm (QAOA) represents a gate-based approach that discretizes quantum annealing through trotterization, alternating between cost and mixer operators across multiple layers while optimizing parameters classically.

Performance comparisons on small number partitioning instances show that simulated annealing and quantum annealing outperform QAOA in both approximation ratio and runtime, though all current approaches face limitations from noise, decoherence, and hardware constraints in the NISQ era. The authors identify critical future research directions including circuit decomposition methods, error mitigation strategies, hybrid quantum-classical algorithms, hardware-aware compilation, and the development of quantum-inspired classical methods.

4.2. Quantum-Inspired Computing

Some early large-scale commercial impacts may come not from fully universal quantum computers, but from quantum-inspired algorithms and specialized analog hardware, which have demonstrated competitive empirical performance on certain structured combinatorial problems while avoiding the engineering challenges of maintaining coherent quantum states at scale. For OR/MS, this represents a particularly fertile near-term opportunity. We stress that quantum-inspired algorithms are, at the mathematical level, classical algorithms; the gains arise not from quantum mechanics *per se*, but from the structural properties of problem formulations that quantum physics naturally suggested—such as Ising encodings, annealing schedules, and energy-landscape traversal—and, in the case of photonic hardware, from specialized physical dynamics. For problems without this structure, there is no general reason to expect quantum-inspired methods to outperform the best classical algorithms.

4.2.1. Application and Algorithm: FRNDP Karahalios et al. (2024) introduce the First Responder Network Design Problem (FRNDP), which addresses a critical challenge in disaster management: determining which road lanes should be reserved exclusively for first responders (FRs) following a disaster, while minimizing the total evacuation time for civilians. The problem arises from a proposal by Turkey's Ministry of Transportation and Infrastructure to designate specific lanes on certain road segments for FR use after disasters such as earthquakes. The authors formulate FRNDP as a bilevel mixed-integer nonlinear program where the outer problem selects links for lane reservation to ensure FRs can reach all demand points from entry nodes, while the inner problem models evacuees' selfish routing behavior under user equilibrium traffic conditions.

To solve this computationally challenging problem, the authors develop GAGA, a novel quantum-inspired bilevel optimization algorithm. GAGA employs the Graver Augmented Multiseed Algorithm (GAMA) at both the outer level (to search over feasible FR paths) and the inner level (to approximately solve the traffic assignment problem), followed by an exact gradient descent method to refine the final solution. GAMA uses Graver basis elements—candidate improving directions from algebraic geometry—as a test set for iterative augmentation. The algorithm generates partial Graver bases by converting problem constraints into QUBO problems, solvable by simulated or quantum annealing; see Tayur and Tenneti (2024) and Guddanti et al. (2024).

Computational experiments on both synthetic graph instances (10–30 nodes) and realistic case studies based on Istanbul's Avcılar district demonstrate GAGA's effectiveness. Comparisons with state-of-the-art branch-and-bound show that GAGA generally produces superior solution quality with competitive or better runtime performance: for Istanbul instances with 179 nodes and 234 links, GAGA achieved objective values 5–15% better than branch-and-bound within the same time limit. The results reveal that approximately 60% of potential FR road segments were selected for lane reservation, and that the lane reservation strategy remained remarkably consistent regardless of evacuation levels.

4.2.2. Hardware: Photonic Ising Machines Photonic Ising Machines[5] (PIMs) raise familiar but newly urgent questions from an Operations Research perspective: When does specialized hardware outperform general-purpose solvers? How do solution quality, time-to-solution, and energy consumption trade off as problem size grows? And how should such devices be benchmarked fairly against classical and quantum alternatives?

[5] PIMs are specialized analog processors that use interacting laser or optical oscillator networks to physically emulate Ising spin systems, enabling extremely fast approximate solutions to hard combinatorial optimization problems by letting the system relax to a low-energy configuration. Ising formulations of many NP problems are catalogued in Lucas (2014).

Prabhakar et al. (2023) present two types of PIMs as alternatives to quantum and simulated annealing for solving NP-hard optimization problems. The authors demonstrate how QUBO problems can be mapped to Ising Hamiltonians and solved using photonic systems. An important clarification is needed regarding comparisons with solvers: the comparisons are in terms of *time to a good solution* (approximation quality within a fixed wall-clock budget), not in terms of proven optimality or optimality certificates. Gurobi is used as a general-purpose benchmark; it is likely that specialized classical MaxCut solvers (such as QuBowl or BiqMac) would outperform both Gurobi and, potentially, the PIM on many instances—this comparison was not performed in that work and remains an important open empirical question.

The first implementation is a temporal multiplexed Ising machine using the bistable response of an electro-optic modulator (Mach–Zehnder Modulator) to represent spin states. This system tackles the Max-Cut problem with up to 1000 spins, performing comparably to Gurobi and outperforming it for graphs with 100 nodes and density less than 50%. The second implementation is a spatial photonic Ising machine (SPIM) that solves the Mattis spin glass model by convolving a coherent laser wavefront with a spatial light modulator (SLM) pattern. This system tackles the number partitioning problem by encoding problem instances as phase masks on the SLM using adiabatic tuning. The SPIM successfully partitions arrays of $2^{14} = 16{,}384$ integers, significantly exceeding what Gurobi (limited to 1024 spins) and D-Wave (limited to 121 spins due to embedding constraints) could handle in these experiments. The quality of the solution generally improves with the size of the problem, suggesting favorable scaling properties for this photonic approach. Fair benchmarking across heterogeneous quantum and quantum-inspired hardware ultimately relies on data-scientific rigor—careful experimental design, uncertainty quantification, and statistically meaningful comparisons of solution quality, runtime, and energy use.

4.3. Current Activities in QC Hardware

Quantum computing's promise hinges on overcoming the fragility of qubits through robust error correction, where physical qubits are encoded into logical qubits that can detect and correct errors in real time. IBM's publicly stated roadmap (as of the writing of this article) integrates error correction into its development path: incremental hardware and software advances—including higher-connectivity processors and improved error mitigation—are planned, targeting demonstration of fault-tolerant modules and the delivery of its first full fault-tolerant system (Starling) capable of running millions of gates on hundreds of logical qubits (IBM, 2024). Alongside organic technological progress, the industry has seen strategic mergers and acquisitions aimed at accelerating development: Google purchased Atlantic Quantum; D-Wave Quantum agreed to acquire Quantum Circuits Inc. (QCI) to blend QCI's dual-rail error-detecting technology with D-Wave's control stack and speed its gate-model roadmap; and IonQ acquired SkyWater Technology to bring chip fabrication in-house and support its long-term hardware scaling (D-Wave, 2024; IonQ, 2024). All of these developments remain subject to change, and readers are encouraged to consult current company disclosures for the most up-to-date information.

5. Suggestions for Further Reading

Looking ahead, the opportunities for INFORMS researchers extend well beyond the examples discussed here. As quantum technologies mature, questions of verification, certification, resource allocation, and control under uncertainty will only grow in importance. Hybrid classical-quantum systems will demand new modeling paradigms, while fair benchmarking across heterogeneous hardware platforms will require careful experimental design and statistical rigor. By engaging early and deeply, the OR/MS community can help ensure that the Second Quantum Revolution is not only scientifically impressive, but operationally robust and societally impactful.

Beyond the tutorials mentioned above on quantum information science (Siddhu and Tayur, 2022) and quantum computing (Mazumder and Tayur, 2025), a popular reference text for a comprehensive introduction to Quantum Information Science (including Quantum Computing) is Nielsen and Chuang (2011). See McGeoch (2014) for an accessible introduction to Adiabatic Quantum Computing (AQC) and Rieffel and Polak (2014) for a gentle introduction to gate-based models. A systematic mapping study of quantum and quantum-inspired algorithms for Operations Research is Gomes et al. (2024); an illustration of Quantum Machine Learning (for image recognition) is Guddanti et al. (2024); and research on quantum approximate multi-objective optimization is Kotil et al. (2025). I hope that this invited article encourages INFORMS researchers to contribute further to the Second Quantum Revolution.

Acknowledgements

I thank EIC Yu Ding for inviting me to write this article, and Tinglong Dai, Gabriel Falcão, Jay Swaminathan, Alan Scheller-Wolf, Claudio Gomes, and João Paulo Fernandes for their feedback on an earlier draft. I also thank the two reviewers whose comments substantially improved this manuscript.

References

Coffrin C, Lobe E, Nannicini G, Parekh O (2025) Introduction to the special issue on quantum computing and operations research. *INFORMS Journal on Computing* 37(1):2.

Eldar YC, Oppenheim AV (2002) Quantum signal processing. *IEEE Signal Processing Magazine* 19(6):12–32.

Gomes C, Fernandes JP, Falcão G, Kar S, Tayur S (2024) A systematic mapping study on quantum and quantum-inspired algorithms in operations research. *ACM Computing Surveys* 57(3), Article 66, pp. 1–35.

Guddanti SS, Padhye A, Prabhakar A, Tayur S (2024) Pneumonia detection by binary classification: classical, quantum, and hybrid approaches for support vector machine (SVM). *Frontiers in Computer Science* 6:1286657.

IBM (2024) IBM Quantum Development Roadmap. *IBM Research*, `https://www.ibm.com/quantum/roadmap`.

IonQ (2024) IonQ to Acquire SkyWater Technology. Press Release, IonQ Inc., `https://ionq.com/news`.

D-Wave Quantum (2024) D-Wave to Acquire Quantum Circuits Inc. Press Release, D-Wave Quantum Inc., `https://www.dwavequantum.com/news`.

Karahalios A, Tayur S, Tenneti A, Pashapour A, Salman FS, Yildiz B (2024) A quantum-inspired bilevel optimization algorithm for the first responder network design problem. *INFORMS Journal on Computing* 36(3):711–731.

Kotil A, Pelofske E, Riedmüller S, Egger DJ, Eidenbenz S, Koch T, Woerner S (2025) Quantum approximate multi-objective optimization. *Nature Computational Science* 5(1):48–60.

Kumar V, Chandra NK, Seshadreesan KP, Scheller-Wolf A, Tayur S (2023) Optimal entanglement distillation policies for quantum switches. *Proceedings of IEEE International Conference on Quantum Computing and Engineering (QCE)*, pp. 295–305.

Lucas A (2014) Ising formulations of many NP problems. *Frontiers in Physics* 2:5.

Mazumder A, Tayur S (2025) Five starter problems: Solving QUBOs on quantum computers. *TutORials in Operations Research*, INFORMS, pp. 145–183.

McGeoch CC (2014) *Adiabatic Quantum Computation and Quantum Annealing*. Morgan & Claypool Publishers.

Motlagh D, Wiebe N (2024) Generalized quantum signal processing. *PRX Quantum* 5:020368.

Nielsen MA, Chuang IL (2011) *Quantum Computation and Quantum Information*, 10th anniversary ed. Cambridge University Press.

Peña J, Siddhu V, Tayur S (2025) Tailored first-order and interior-point methods and a new semidefinite programming hierarchy for entanglement detection. arXiv:2508.05854.

Prabhakar A, Shah P, Gautham U, Natarajan V, Ramesh V, Chandrachoodan N, Tayur S (2023) Optimization with photonic wave-based annealers. *Philosophical Transactions of the Royal Society A* 381(2241):20210401.

Rieffel E, Polak W (2014) *Quantum Computing: A Gentle Introduction*. MIT Press.

Siddhu V, Tayur S (2022) Five starter pieces: Quantum information science via semidefinite programs. *TutORials in Operations Research*, INFORMS, pp. 59–92.

Siddhu V, Chatterjee A, Jagannathan K, Mandayam P, Tayur S (2026) Unital qubit queue-channels: Classical capacity and product decoding. *IEEE Transactions on Quantum Engineering* 7:1–14.

Singh AK, Sharma N, Singh VP, Prabhakar A, Tayur S (2025) Quantum money using differential phase encoding. *Optica Quantum 2.0 Conference and Exhibition* QW3A.16.

Sodhi MS, Tayur S (2023) Make your business quantum-ready today. *Management and Business Review* 3(2):12–19.

Tayur S, Tenneti A (2024) Quantum annealing research at CMU: Algorithms, hardware, applications. *Frontiers of Computer Science* 18(6):186901.

Velu C, Norman K, Zhu Y, Putra FHR, Noble C (2025) *The Business of Quantum Technologies*. Cambridge University Press.

Wehner S, Elkouss D, Hanson R (2018) Quantum internet: A vision for the road ahead. *Science* 362(6412):eaam9288.