Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Hybrid Classical-Quantum Learning for Space Based Data Centers: A CUDA-Q Study of Variational and Photonic Backends

Santanu Ganguly

ganguly.santanu@kingston.ac.uk

Kingston University

Research Article

Keywords: datacenter in space, hybrid quantum-classical systems, CUDA-Q, VQC, boson sampling, photonic computing, graph neural networks, message passing neural networks, ORCA photonics, routing under radiation/SEU stress

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9487599/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Hybrid Classical-Quantum Learning for Space Based Data Centers: A CUDA-Q Study of Variational and Photonic Backends

Santanu Ganguly[1,2][0000-0003-0141-9228]

[1] Quantum AI Research Group, School of Computer Science and Mathematics, Kingston University, Kingston upon Thames, KT1 2EE, United Kingdom
[2] OnQAI Ltd., 128 City Road, London, United Kingdom, EC1V, United Kingdom
ganguly.santanu@kingston.ac.uk
santanu.gangoly@gmail.com

Abstract. We study hybrid classical–quantum routing for a *10-node optical space-based data center* (SBDC) with 30-second routing-table refresh, asking whether quantum-enhanced subroutines improve representation learning under realistic latency constraints. Our model uses a message-passing neural backbone and compares five backends within a common CUDA-Q workflow: a classical baseline, a photonic sampling-inspired node encoder, a VQC node encoder on qpp-cpu, the same VQC on nvidia, and an ORCA-style photonic edge augmentation. We evaluate routing accuracy, route success, latency per refresh, successful-route queue and congestion reduction, robustness under radiation/SEU-style stress, and backend-aware quantum diagnostics including node-state fidelity, expectation drift, boson-output statistics, and logical resource counts. The results are mixed but informative. The photonic backend achieves the highest held-out routing accuracy in the low-latency regime (about 0.60), while the classical and ORCA-edge variants remain close (roughly 0.56–0.58) at similarly low latency. By contrast, the VQC backends incur two to three orders of magnitude larger per-refresh latency for lower routing accuracy, although they produce richer quantum-state structure. Overall, the main result is a backend-grounded demonstration that classically controlled photonic and variational subroutines measurably alter routing representations and latency–utility trade-offs in optical LEO networking.

Keywords: datacenter in space**;** hybrid quantum-classical systems; CUDA-Q; VQC; boson sampling; photonic computing; graph neural networks; message passing neural networks; ORCA photonics; routing under radiation/SEU stress.

1 Introduction

Recent industry developments suggest that *orbital space based data centers* (SBDCs) [60, 61] are moving from speculative concept toward early platform development. NVIDIA's 2026 space-computing announcement places Starcloud among the partners using NVIDIA accelerated-computing platforms for orbital data centers and next-generation space missions, while NVIDIA's earlier profile of Starcloud and Starcloud's

own public materials describe a longer-term architecture based on continuous solar energy, radiative cooling, and space-based AI compute [33, 33, 52]. Although these announcements do not yet imply a mature SBDC ecosystem, they do suggest that future orbital compute platforms may host not only AI inference and training workloads, but also backend-accelerated scientific and quantum-simulation workflows [33, 33, 52]. This trend makes networking in space newly relevant as part of a broader systems stack for *space-native compute fabrics* rather than as a communications problem in isolation.

At the same time, low Earth orbit (LEO) satellite networks are increasingly moving toward optical inter-satellite and space-to-ground links as a high-throughput communications substrate. NASA's optical communications program explicitly positions laser links as a means to increase space-link data return, and the TBIRD mission demonstrated 100–200 Gbps laser communication from LEO, making optical data planes a credible abstraction for future space networking studies [30, 48]. In such environments, routing is inherently a *graph-structured decision problem*: the preferred next hop depends on local node state, neighbouring links, and destination-conditioned topology. This makes graph-based learning, and especially *message passing neural networks* (MPNNs), a natural foundation for learned routing policies [4, 14] in their discussion of relational inductive biases and graph networks. MPNNs update node representations through neighbourhood aggregation and therefore encode the relational bias required for topology-aware control [4, 14, 31, 49].

Inspired by this setting, we study routing on a 10-node clustered optical satellite ring and augment a classical graph-learning backbone with *quantum-enhanced feature mechanisms*. The objective is not to replace the routing controller with a stand-alone quantum system, but to use photonic computation and variational quantum circuits [12] as *structured representation layers* inside a topology-aware classical model. In particular, we use a node-local photonic sampling feature map whose outputs are converted into quantum-enhanced node features before message passing. This design is supported by prior work on bosonic photonic sampling and graph learning [15] where *Gaussian boson sampling* was introduced as a physically meaningful photonic sampling model, and later, it was shown [50] that Gaussian-boson-sampling statistics can induce graph kernels and graph similarity measures. In this sense, the photonic layer is used here as a stochastic feature generator whose distributional outputs may encode graph structure in a way that is difficult to reproduce with a purely deterministic classical map [15, 50].

1.1 Contributions.

Our **main contributions** are fivefold: **(i)** we formulate optical LEO routing as a hybrid graph-learning problem with 30-second snapshot refresh, a message-passing backbone, and explicit route-validity masking; **(ii)** we integrate two quantum-enhanced mechanisms within one routing stack, namely a boson-sampling-inspired photonic feature generator and a node-local VQC; **(iii)** we provide a backend-aware CUDA-Q implementation comparing qpp-cpu, nvidia, mqpu, and orca-photonics in a single workflow; **(iv)** we *benchmark* the resulting models using route-level, systems-level, and quantum-level diagnostics, including routing accuracy, route success, latency, successful-route queue/congestion reduction, radiation/SEU robustness, node-state fidelity, expectation

maps, boson-output statistics, and logical resource counts; and **(v)** we conclude that the strongest design principle supported by the experiments is *classical control for routing with quantum subroutines for representation and stochastic structure*, rather than fully quantum-native routing control [34, 25, 55].

The remainder of this paper is **organized as follows**: Section 2 describes the background and motivation; Section 3 addresses the methods; Section 4 the results; Section 5 the conclusion and outlook, followed by references, and Section 6, which is the Appendix.

2 Background and Motivation

Graph learning and message passing. MPNNs provide a general framework for learning on graphs by iteratively exchanging and aggregating messages between neighboring nodes before producing a task-specific prediction [14]. More broadly, graph networks are understood as architectures with a strong relational inductive bias, meaning that they are designed to exploit the structured relationships among entities and edges rather than ignoring that structure [4]. For routing problems, this is important because the next-hop action at node v depends not only on v's own features, but also on neighboring satellites, link conditions, and destination-conditioned graph context. Accordingly, MPNNs are a natural choice for destination-aware routing in satellite rings [21].

2.1 Photonic sampling as a feature map

One of the main motivations is to test whether photonic sampling at the network edge can improve routing in optical space networks, where local context matters and node state may be noisy, uncertain, or slowly varying. In a clustered optical ring, many nodes occupy similar but not identical structural positions, so routing depends on subtle local distinctions rather than coarse topology alone. Message passing provides the appropriate relational inductive bias because it allows each node to aggregate information from its neighbours instead of acting in isolation [4, 14]. The photonic layer then serves as a structured stochastic feature map, transforming local graph structure into a richer representation than a purely hand-engineered classical encoding.

Bosonic photonic samplers, and Gaussian boson sampling in particular, have been studied not only as candidates for near-term quantum advantage, but also as devices whose output statistics encode graph structure [15, 50] showed that samples from a Gaussian boson sampler can be used to construct graph feature vectors and graph kernels, demonstrating a direct connection between photonic output distributions and graph similarity [45, 17]. This motivates our use of a photonic sampling feature map at each node: instead of requiring a fully quantum graph learner, we use the photonic sampler to produce stochastic node-local features that are then fused with classical state and propagated by message passing.

2.2 Variational quantum algorithms (VQA)

On the quantum side, variational quantum algorithms has been a leading near-term approach for hybrid quantum-classical learning [7], where a parameterized quantum circuit produces observables that feed a classical outer loop [8, 17]. This is the precise role of the VQC in our architecture: it is not a standalone routing controller, but a trainable embedding subroutine whose outputs are fused into the classical routing model. On the photonic side, boson sampling has been formalized as a model in which identical photons interfere inside a linear-optical network and are then measured at the output, with output probabilities determined by the interferometer [3]. This makes boson sampling a principled source of structured stochastic features, rather than arbitrary randomness.

2.3 NVIDIA CUDA-Q.

In the context of this work, NVIDIA's CUDA-Q [9] is directly relevant because it provides a unified hybrid programming model spanning CPUs, GPUs, and QPUs, with explicit support for GPU-accelerated state-vector simulation and multi-QPU workflows. NVIDIA's current documentation states that the nvidia backend provides GPU-accelerated state-vector simulation through `cuStateVec` and the `mqpu` mode exposes asynchronous multi-QPU or multi-GPU execution patterns. Taken together, these capabilities make CUDA-Q-based simulation a plausible future workload for orbital datacenters, provided that such platforms expose compatible NVIDIA accelerated-computing environments. Even though current public evidence supports this only as a forward-looking systems possibility, the trajectory is clear enough to motivate this study: if AI-capable orbital compute platforms become part of the LEO ecosystem, then quantum simulated applications and routing architectures that combine graph learning, photonic sampling, and GPU/QPU-aware simulation will be directly relevant to how those systems are designed, tested, and deployed [35].

3 Methods

All models were run on Ubuntu 24.04 LTS utilizing CUDA-Q version 0.14, CUDA 12.6, Python 3.11, PyTorch 2.10 for cuda on RTX 3050, network 3.6.1. All trainings were run for 30 epochs.

3.1 Architectural Considerations

Routing in this setting is naturally a *graph-structured prediction problem*, since the preferred next hop at node v depends not only on the local state of v, but also on neighboring satellites, adjacent links, and the destination-conditioned structure of the ring. Our choice of *message passing neural network* (MPNN*)* provide a general framework for learning on graphs through iterative neighbourhood aggregation, while the broader graph-network perspective emphasizes the value of relational inductive biases for structured decision-making [4, 14]. In our context, this means that the routing policy can

propagate information across the ring rather than treating each node as an isolated classifier.

The architectural choice becomes more relevant in light of recent developments in *AI compute in space*. StarCloud-1 reportedly launched in November 2025 carrying an NVIDIA H100 into orbit, further strengthening the idea that data center style networking in space may become an active engineering direction, especially for AI workloads [33, 34, 52). We therefore frame the present work not only as a routing study, but also as a small step toward the broader problem of *networked AI and quantum-enhanced computation in space*.

To enrich the node representation beyond hand-engineered classical features, we augment each node with a *photonic sampling feature map*. Concretely, a local neighbourhood representation is encoded into photonic parameters, samples are drawn from the corresponding photonic distribution, and the resulting shot patterns are converted into a node-local feature vector before message passing. We deliberately keep this photonic layer general rather than tying the method to a single Gaussian-specific construction, but our design is motivated by prior work on Gaussian boson sampling (GBS), which showed that bosonic photonic output statistics can encode graph structure and can be turned into graph-feature or graph-similarity representations [15, 50]. This suggests a natural hybrid architecture in which the photonic layer supplies stochastic local features and the classical MPNN turns them into topology-aware routing actions.

3.2 The Model

The datacenter network in LEO: Topology and Hybrid Routing Model. We model the network as a 10-node clustered optical satellite ring in low Earth orbit (LEO), with nodes representing satellites and edges representing optical inter-satellite links as visualized in **Fig. 1**, and **Fig. 2** shows a high-level diagram of the model pipeline; throughout the experiments, the neighbor spacing is set to approximately 1000 km as a simulation parameter to represent a compact high-rate orbital cluster rather than a full-orbit constellation. For each routing query, a source node and destination node are specified (**Fig. 1**), and the learning problem is to predict the next-hop action at each node conditioned on the destination. We parameterize the routing policy with a *message passing neural network* (MPNN) [21] as the routing backbone, since message passing provides a natural graph-based mechanism for aggregating neighborhood information and supplies the relational inductive bias needed for structured decision-making on networks: the preferred action at node v depends on neighboring satellites, link states, and destination-conditioned graph context rather than on node-local features alone [4, 14]. The use of optical links is motivated by the role of free-space optical communications in next-generation LEO satellite networks [20], where narrow beams, large available bandwidth, and high-rate operation make optical inter-satellite links an appropriate physical-layer abstraction, making an aggressive optical data-plane abstraction physically plausible even though the topology here remains a simulation construct [10, 48]. **Fig. 16** of Section 6.5 of the Appendix shows the corresponding topology graphs of `classical_identity` and `photonic_sampler`.

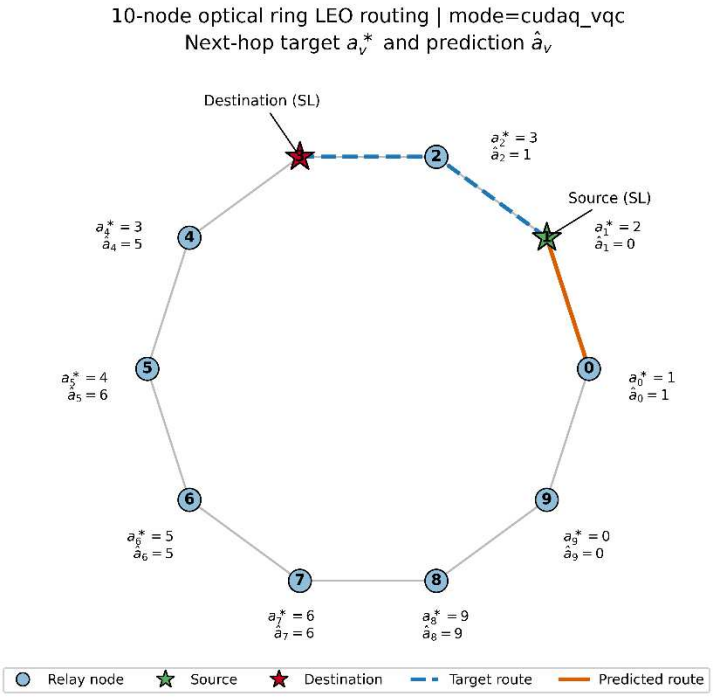


Fig. 1. Clustered optical satellite-ring topology used in the routing experiments. The figure shows a clustered satellite ring hosting datacenters (DCs), shown in blue, where vertices represent satellites hosting the DCs and gray edges represent optical inter-satellite links. Qualitative next-hop rollout visualization is shown for routing models on held-out 10-node optical ring LEO graphs for `cudaq_vqc`. The corresponding `classical_identity` and `photonic_sampler` graphs used in the benchmarking exercise are shown in **Fig. 16** in the Appendix. Blue circles denote relay nodes, green and red stars denote the source and destination, the dashed blue line denotes the target route, and the orange line denotes the predicted rollout. Node annotations show the target next-hop label a_v^* and the predicted next hop $\hat{a}_v$. The figures illustrate that exact node-wise agreement and end-to-end route success can diverge. In our simulations, the topology is modeled as a *clustered ring* rather than a full orbital ring, with a nominal nearest-neighbor spacing of approximately 1000 km.

The use of optical links is motivated by the role of free-space optical communications in next-generation LEO satellite networks [21], where narrow beams, large available bandwidth, and high-rate operation make optical inter-satellite links an appropriate physical-layer abstraction, making an aggressive optical data-plane abstraction physically plausible even though the topology here remains a simulation construct [4, 14].

In **Fig. 1**, we treat the *node count* ($|V| = 10$), *orbital regime* (LEO), and *neighbor spacing* (~ 1000 km) as explicit experimental parameters. This modeling choice is consistent with the role of optical/laser communications as the relevant high-throughput physical layer for future space networks: NASA describes optical communications [30, 31] as laser-based links intended to deliver substantially higher data return than conventional RF systems, and the TBIRD mission [49] demonstrated 100–200 Gbps laser

communication from LEO, which makes our use of an aggressive, forward-looking optical data-plane regime physically motivated even though the exact topology here is a simulation construct.

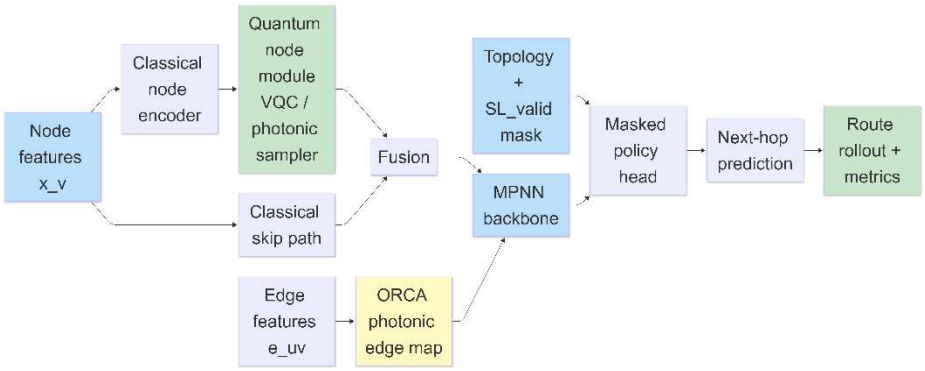


Fig. 2. Overview of the *hybrid routing training model*. Node features are encoded by a classical backbone augmented with an optional quantum node module (VQC or photonic sampler), while edge features can be enhanced by an ORCA photonic edge map. The fused representation is processed by an MPNN, combined with a topology-validity mask, and passed to a masked policy head that predicts the next hop. Predictions are then rolled out into full routes and evaluated with routing metrics. The *full model pipeline* depicted in **Fig. 22** in Section 6.10 of the Appendix.

A central challenge in such a hybrid model is that photonic samplers are inherently stochastic, so the downstream routing policy should be informative yet stable under repeated draws from the same node-local configuration. We address this by training the routing model as follows: combine supervised next-hop learning with a consistency regularization term that encourages agreement between policies induced by independent photonic draws. This follows the general semi-supervised view that predictions should remain stable under stochastic perturbations or alternative views of the same input, as in Temporal Ensembling and Mean Teacher [27, 53]. The resulting question is therefore not whether the photonic layer replaces classical graph learning, but whether it provides a useful stochastic representation that improves routing robustness and expressivity in a clustered optical ring.

3.3 Hybrid routing on clustered optical satellite rings

We model the routing policy as a supervised graph-learning problem in which node and edge states are processed by a message-passing backbone and then mapped to next-hop decisions. This is naturally expressed in the *message passing neural network* (MPNN) framework [21], where node representations are updated iteratively through neighborhood aggregation before a task-specific readout head is applied [14]. The `cudaq_vqc` branch augments this classical encoder with a variational quantum circuit (VQC), following the standard hybrid paradigm in which an Adam optimizer updates parameters of a trainable quantum circuit embedded inside a larger learning loop [7]. The `photonic_sampler` branch uses a boson-sampling-inspired feature map motivated by

linear-optical sampling over interferometric networks, which provides a physically grounded stochastic encoding rather than a deterministic classical projection [3].

To augment the classical node and link features, we attach a *node-local photonic sampling module* to each satellite and treat it as a *photonic sampling feature map*. Concretely, for node v, a local neighborhood representation is encoded into photonic parameters Θ_v, samples $s_v^{(k)} \sim \mathcal{P}_{\text{phot}}(\Theta_v)$ are drawn, and the resulting shot patterns are converted into a robust feature vector $z_v = \phi(\{s_v^{(k)}\}_{k=1}^{K})$, which is then fused with classical state variables before message passing. We deliberately keep this layer *general*, so that the pipeline is compatible with Gaussian boson sampling as well as related bosonic photonic primitives; this is motivated by prior work showing that boson-sampling distributions can encode graph structure and support graph-similarity or kernel-style representations. In the implementation studied here, the policy is trained in a *hybrid* manner: the photonic sampler supplies stochastic node-local features, while the routing backbone and optimization remain classical, using supervised cross-entropy on labeled next-hop actions together with a consistency penalty across independent photonic draws to encourage stability under sampling noise.

The MPNN then propagates information across the ring by iteratively aggregating neighborhood messages. At layer ℓ, node v receives an aggregated message $m_v^{(\ell)}$ formed from the states of its neighbors and the incident edge features, and its hidden representation is updated by the function U_ℓ. After L layers, a policy head g maps the final node state to destination-conditioned next-hop logits, yielding a routing policy $\pi_v(a|D)$. We use U_ℓ for the message-passing update and ϕ_{phot} exclusively for the photonic feature map to make the separation between the photonic and graph-learning components explicit. (a full notation table and associated theory is shared in **Table 3** in the Appendix.

Unless otherwise stated, experiments use the configuration shown in **Fig. 1**, destination-conditioned routing queries, and optical inter-satellite links as the communication substrate. The clustered-ring topology is not intended as a literal deployed constellation design; rather, it provides a compact graph in which routing decisions depend on both local node state and destination-conditioned topology. For each node, a local neighborhood representation is encoded into photonic parameters, samples are drawn from the resulting photonic distribution, and the sampled outcomes are converted into quantum-enhanced node features before message passing.

Training Dynamics of the Hybrid MPNN with VQC. For VQC based learning [16], we partition nodes into a supervised set V_{SL} and an unlabeled set V_{UL}. Supervised nodes carry next-hop labels a_v^*, while unlabeled nodes contribute only through consistency regularization. The training objective is

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda \mathcal{L}_{\text{cons}},$$

where $\mathcal{L}_{\text{sup}}$ is the cross-entropy loss over the subset $V_{\text{SL}}^{\text{valid}} \subseteq V_{\text{SL}}$ for which a next-hop label is defined, and $\mathcal{L}_{\text{cons}}$ is a symmetrized KL (Kullback–Leibler) [25] penalty between two policy outputs obtained from *independent photonic draws* of the same node-

local configuration. Intuitively, the supervised term teaches the model to reproduce labeled routing actions, while the consistency term encourages the policy to remain stable under stochasticity in the photonic layer. The resulting architecture is therefore *hybrid*: *photonic sampling* provides stochastic quantum-enhanced node features, whereas message passing and optimization remain classical. The detailed theory is shared Section 6.1 in Appendix.

3.4 vqc Expectation fidelity matrix

For each node v, we compute the final quantum state produced by the trained encoder after classical preprocessing maps the node features to circuit input angles. We then evaluate the expectation value of the Pauli-Z observable on each qubit and arrange the resulting values into a node-by-qubit matrix, where row $v$corresponds to node $v$and column q corresponds to qubit q. In variational quantum algorithms, optimization and prediction are often mediated through expectation values of selected observables, making such maps a natural diagnostic of what information the quantum circuit is actually exposing to the classical outer loop [7, 47].

3.5 Node-state fidelity extraction

To characterize the geometry of the learned quantum node representation, we extract a *node-state fidelity matrix* from the trained VQC encoder. After selecting a trained checkpoint, we pass each node feature vector x_v through the classical preprocessing block and the VQC kernel to obtain a pure output state $|\psi_v\rangle$ for node v, using CUDA-Q state retrieval on the trained kernel. For every pair of nodes i, j, we then compute the pairwise fidelity

$$F_{ij} = |\langle\psi_i|\psi_j\rangle|^2,$$

which is the standard pure-state specialization of quantum fidelity, as described in prior works [22, 32]. The resulting matrix $F \in \mathbb{R}^{N\times N}$ is visualized as a heatmap and summarized by an off-diagonal diversity score

$$\text{diversity}(F) = 1 - \frac{1}{N(N-1)}\sum_{i\neq j} F_{ij},$$

so that larger values indicate less overlap and, therefore, a less collapsed quantum embedding. In our setting, this diagnostic is used as a *representation analysis* rather than as a direct routing metric: it does not by itself imply better route success, but it reveals whether the VQC maps different nodes to distinguishable quantum states, which is important in hybrid variational models where the learned quantum representation is subsequently consumed by a downstream classical network [7]. Full theory is shared Section 6.2 in Appendix.

We use the fidelity matrix and its diversity summary as a qualitative and semi-quantitative diagnostic of whether the VQC backend is generating *node-dependent latent*

quantum states. This complements route-level metrics such as routing accuracy and route success and helps distinguish between backends that produce genuinely separated quantum embeddings and backends whose node states collapse to a narrow region of Hilbert space.

3.6 Node-by-Qubit Expectation Extraction

To characterize the *observable-level structure* of the learned quantum node representation, we extract a *node-by-qubit expectation matrix* from the trained VQC encoder. After selecting a trained checkpoint, we pass each node feature vector x_v through the classical preprocessing block and the VQC kernel to obtain a pure output state $|\psi_v\rangle$ for node v. For each node $v$and qubit $q \in \{1, \dots, Q\}$, we then compute the Pauli-Z expectation

$$E_{vq} = \langle \psi_v | Z_q | \psi_v \rangle,$$

where Z_q denotes the Pauli-Z observable acting on qubit q. The resulting matrix $E \in \mathbb{R}^{N \times Q}$ is visualized as a heatmap, where each row corresponds to a node and each column corresponds to a qubit. In variational quantum models, expectation values are the primary interface between the quantum circuit and the classical outer loop, so this matrix provides a direct view of what node-dependent information the quantum encoder is actually exposing to the downstream routing model, as shown in previous works [1, 7]. To summarize the heatmap numerically, we define an observable heterogeneity score

$$\text{heterogeneity}(E) = \frac{1}{Q} \sum_{q=1}^{Q} \text{Var}_v \left(E_{vq} \right),$$

where Var_v denotes the variance across nodes for a fixed qubit q. Larger values indicate that the measured qubit responses vary more across nodes, while smaller values indicate that the circuit exposes nearly constant observables and is therefore closer to observable-level collapse. As with the fidelity metric, this quantity is used as a *representation diagnostic* rather than as a direct routing objective: it does not by itself imply improved route success, but it helps determine whether the trained VQC produces node-dependent observables that can plausibly support different downstream routing decisions. Full theory is shared Section 6.3 in Appendix.

We use the node-by-qubit expectation matrix as a companion diagnostic to the node-state fidelity matrix. Fidelity reveals whether the quantum encoder separates nodes in Hilbert space, while the expectation matrix reveals whether that separation is actually visible through the selected observables passed to the classical routing stack. In practice, the two diagnostics play complementary roles: a model may exhibit distinct states but weak observable separation, or conversely show strong observable heterogeneity even when full-state diagnostics are harder to interpret.

3.7 Quantum Circuit

Fig. 3 shows the quantum circuit implemented. The circuit consisted of the node-local quantum encoder as a CUDA-Q quantum kernel over 4-qubits, with a circuit depth of 28, constructed with $12 \times R_y$, $8 \times C_x$ and $8 \times R_z$ gates. Each node is first classically preprocessed into a vector of input angles and encoded through R_y rotations, after which trainable R_y and R_z layers are applied and coupled by a ring of controlled-X entangling operations. The resulting circuit is executed through CUDA-Q's kernel interface, and the routing model consumes expectation values or state-derived diagnostics rather than raw measurement counts. NVIDIA's documentation describes CUDA-Q kernels as the fundamental representation of quantum circuits in the framework, while the visualization API provides a direct way to inspect the execution path of a kernel [7, 34, 35].

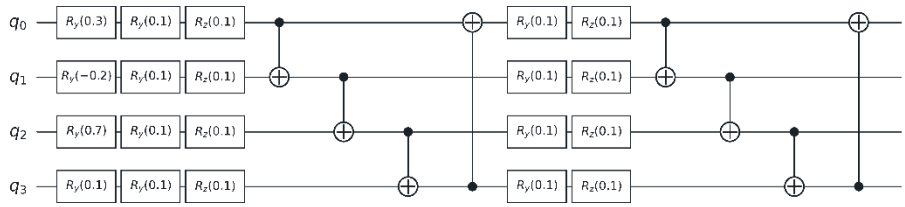


Fig. 3. The node-local variational quantum circuit (VQC) used in the `cudaq_vqc` routing backend. The raw circuit produced by CUDA-Q is depicted in **Fig. 20** of the Appendix.

3.8 Quantum Encoder

We use the quantum encoder as a structured representation layer rather than as a standalone controller, as a *node-local representation module* within the hybrid routing architecture. For each node v, the classical feature vector x_v is first mapped to a lower-dimensional latent representation and then encoded into quantum circuit parameters. With the node feature vector x_v, the encoder produces a quantum-enhanced embedding through a parameterized quantum circuit (for VQC) or photonic sampling map (for the ORCA photonic), and exposes that embedding to the classical graph network via observables or sampled outputs respectively. This design treats the quantum component as a feature-generating subroutine whose purpose is to reshape the geometry of the node representation before message passing, while correctness-critical routing logic remains in the classical control plane. The role of the encoder is therefore not to make routing decisions directly, but to provide a richer node representation before message passing and masked policy prediction. Hence, the quantum encoder functions as a *trainable or stochastic feature generator* inside a classically supervised graph-learning pipeline, allowing the downstream routing network to exploit nonclassical structure while preserving classical control over validity constraints, next-hop selection, and route rollout.

3.9 Boson-Output Diagnostics

To characterize the **distributional behavior** of the photonic edge module, we extract an empirical boson-output distribution from the photonic circuit after training or after fixing the edge-encoding parameters. For a given edge-conditioned photonic kernel, we execute repeated sampling on the linear-optical circuit and collect output occupation patterns $y \in \mathcal{Y}$, where each pattern records the measured photon occupation across modes. This follows the standard boson-sampling picture [3], in which identical photons are propagated through a linear-optical network and then non-adaptively measured at the output, and is compatible with CUDA-Q's orca-photonics [40] backend, which provides photonic circuit simulation and sampling primitives for photonic kernels. Theory behind our boson diagnostics is shared in Section 6.4 of the Appendix.

We use boson-output diagnostics as a *companion to route-level and node-level analyses*. Route success, latency, and congestion reduction measure whether the overall hybrid router is useful as a networked decision system, while the boson-output diagnostics measure whether the photonic subroutine is operating in a meaningful sampling regime. In practice, the top-k histogram of output patterns is used for qualitative visualization, while P_{coll}, P_{nocoll}, H_{norm}, and D_{TV} provide compact summaries that can be compared across backends, perturbation settings, or training checkpoints. CUDA-Q's orca-photonics backend is directly suited to this workflow because it supports photonic circuit simulation and repeated sampling of photonic kernels [34].

3.10 ORCA Photonic Configuration

The schematic of **Fig. 4** shows a 4-mode linear-optical circuit with two injected photons, mode-wise phase shifts $(\phi_0, \ldots, \phi_3)$, a sequence of beam splitters $(\theta_0, \theta_1, \theta_2)$ and photon-counting measurement at the output. Input edge-conditioned parameters are mapped to the interferometer, sampled to produce an output occupation distribution, and then converted into photonic features for the downstream routing model. The figure illustrates the role of the ORCA photonic block as a *structured stochastic feature generator* rather than as a standalone routing controller.

The photonic edge module is implemented as a small linear-optical circuit parameterized by edge-derived values and executed through CUDA-Q's orca-photonics backend. In the boson-sampling model, identical photons are injected into selected modes, propagated through a linear-optical network, and then measured by counting photons in the output modes. This model is adapted from previous works [3] and motivate its significance as a structured sampling process over linear optics. In CUDA-Q, orca-photonics provides a Q++-based photonic state-vector simulator and supports photonic operations including `create`, `annihilate`, `phase_shift`, `beam_splitter`, and `mz`, which makes it a natural implementation path for this kind of edge-level photonic subroutine. The histogram is then obtained by repeated sampling and retaining the most probable observed occupation patterns, while the diagnostics summarize collision probability, no-collision probability, and normalized entropy of the sampled law [3, 34].

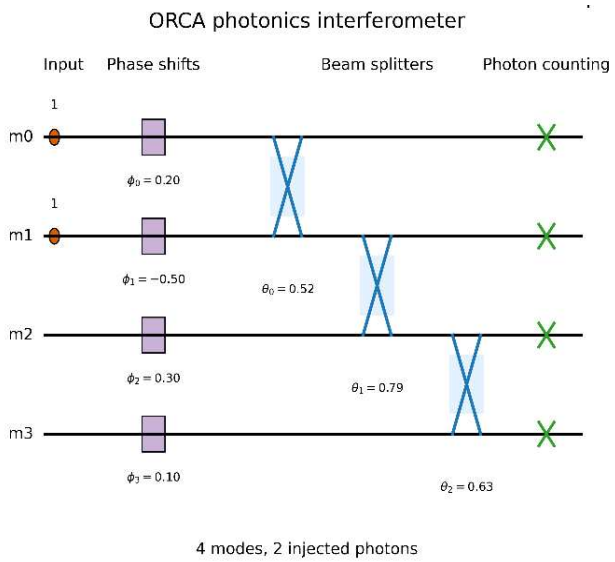


Fig. 4. Linear-optical edge circuit used by the ORCA photonic [40] edge map. The circuit consists of four modes, two injected photons, mode-dependent phase shifts, three beam splitters, and output photon-counting measurements. Samples from the resulting bosonic output distribution are used as edge-conditioned stochastic features within the hybrid routing architecture.

Training MPNN with ORCA. The routing policy is trained as a *hybrid photonic-classical graph model*. At each node, a local neighborhood representation is encoded into a photonic sampling program, and the resulting samples are transformed into node-local features that are concatenated with classical link and node-state variables before message passing. The routing backbone is an MPNN, which is appropriate because next-hop selection is inherently relational and depends on neighboring satellites and links rather than on node-local features alone [14]. The model is optimized with a combined objective

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda \mathcal{L}_{\text{cons}},$$

where $\mathcal{L}_{\text{sup}}$ is cross-entropy on labeled next-hop actions and $\mathcal{L}_{\text{cons}}$ is a symmetrized consistency penalty between policies obtained from two independent photonic draws of the same node-local configuration. This training rule follows the general logic of consistency regularization in semi-supervised learning, where the model is encouraged to produce stable predictions under stochastic perturbations or alternative views of the same input [27, 53]. In implementation, the photonic layer is executed through NVIDIA CUDA-Q's ORCA photonic stack: CUDA-Q provides an ORCA hardware target for boson-sampling-style photonic processors and an orca-photonics simulator backend for local photonic simulation, both within a unified hybrid programming model that couples classical and quantum resources.

3.11 Empirical node-output fidelity matrix

To quantify how strongly the node-local photonic layer distinguishes different satellites, we compute an empirical pairwise fidelity matrix over sampled node-output distributions. For each node i, we first encode its local neighborhood into photonic parameters Θ_i and execute the corresponding photonic program using the CUDA-Q ORCA stack. CUDA-Q supports both an *ORCA boson-sampling target*, which exposes ORCA's time-bin boson-sampling API, and the orca-photonics backend, which provides a state-vector photonics simulator; our procedure applies to either execution path because both ultimately return sampled output statistics over photonic measurement patterns. Let $n_i(s)$ denote the number of times output pattern $s$is observed for node $i$over Kshots, and define the empirical distribution,

$$\hat{p}_i(s) = \frac{n_i(s)}{K}.$$

We then compute the pairwise empirical fidelity

$$F_{ij} = \sum_s \sqrt{\hat{p}_i(s)\hat{p}_j(s)}\,.$$

This quantity is the Bhattacharyya coefficient [5, 6] also commonly interpreted as the classical fidelity between two probability distributions. Historically, this overlap-based measure originates with Bhattacharyya in 1943 and is later discussed in the broader context of distinguishability and fidelity measures [11]. Under this definition, $F_{ij} \in [0,1]$, $F_{ii} = 1$, and larger off-diagonal values indicate greater overlap between the node-output distributions of nodes $i$and j.

We use the resulting matrix $F \in [0,1]^{N\times N}$as a diagnostic for the *discriminative strength* of the photonic feature layer. If off-diagonal fidelities are low, then the photonic sampler yields strongly separated node signatures; if they remain high, then the photonic layer is stable but only weakly discriminative. This use of overlap-based comparisons is consistent with prior bosonic-sampling work on graph similarity, where similarities are derived from relationships between output distributions rather than from deterministic symbolic encodings alone [50].

3.12 LEO radiation model and device-health metric

To capture radiation-driven non-stationarity in the node-local photonic layer, we associate each node v with a scalar sampler-efficiency state $\eta_v(t) \in (0,1]$. The model is intentionally stylized, but it follows the basic distinction emphasized in NASA tutorial [42]: total ionizing dose (TID) is the absorbed dose from ionizing radiation and leads to *cumulative parametric degradation*, whereas a single-event effect (SEE) is a disturbance caused by the passage of a single proton or heavy ion through or near a sensitive circuit node. Additionally, in proton-dominated environments such as LEO, proton SEE

can constitute a substantial fraction of the overall SEE rate [42]. We therefore model device health as the combination of a slowly varying cumulative term and a transient upset term. Let $D_v(t)$ denote the accumulated ionizing dose at node v. A further theoretical insight and associated plot (**Fig. 17**) is given in Section 6.6 of the Appendix.

3.13 Metrics

We report three classes of metrics. First, we use *routing metrics*: next-hop accuracy, route success rate, and mean path length. Second, we use *systems metrics*: average time per graph, effective latency per 30-second refresh, successful-route queue reduction, successful-route congestion reduction, and stress-set accuracy under radiation/SEU-like perturbations. Third, we use *quantum and photonic diagnostics*: node-state fidelity, node-by-qubit expectation maps, noisy-versus-ideal expectation drift, boson-output histograms, collision and no-collision probability, normalized entropy, total variation drift under photonic noise, and logical circuit resource counts. This multi-level evaluation is important because route-level prediction quality alone is insufficient for judging deployability in LEO systems, where latency, congestion, and robustness matter alongside decision accuracy, as shown in previous works [54].

3.14 Backend Comparison Analysis

We compare backends within the same graph-learning pipeline, keeping the routing backbone fixed and changing only the representation module. The underlying routing learner follows the message-passing paradigm for graph-structured inputs, which is a standard choice for relational decision problems [14]. The `photonic_sampler` backend uses a boson-sampling-inspired stochastic map motivated by linear-optical sampling, whereas the `classical_identity` backend uses a purely classical representation path [3]. Validation accuracy is evaluated at each epoch and used as the primary model-selection signal; in light of the observed post-peak decline, future experiments should report best-checkpoint results under explicit early stopping rather than only final-epoch values [44].

To *characterize* the hybrid quantum branches beyond routing accuracy alone, we compute four backend-aware diagnostics. *First*, for the VQC node encoder, we extract node states using CUDA-Q get_state and compute a *state-diversity score* as $1 -$ the mean off-diagonal entry of the node-state fidelity matrix, so larger values indicate less overlap between node embeddings. *Second*, we compute *expectation drift* by comparing noisy and ideal VQC observable outputs and averaging the absolute difference in $\langle Z_q \rangle$ across nodes and qubits. *Third*, for the photonic edge module, we use CUDA-Q sample on the orca-photonics backend and measure *distribution drift* by the total variation distance between ideal and noisy boson-output distributions. Finally, we record *logical resource counts* for the VQC template, including total, one-qubit, and two-qubit gate counts, together with qubit count and logical depth. These diagnostics are motivated by the hybrid VQA setting, where expectation values and state geometry are often the effective interface between quantum and classical computation, and by the boson-

sampling view of photonic computation, where output distributions encode the relevant interference structure.

3.15 Robustness-Latency Comparison under Radiation

To evaluate robustness, we construct a *stress validation set* by perturbing held-out routing graphs with radiation-motivated changes, including elevated radiation level, injected SEU-like node events, and worsened queue/utilization proxies, and then measure routing accuracy under those perturbations. We plot the resulting stress-set accuracy against end-to-end refresh latency per graph so that resilience is interpreted together with runtime cost rather than in isolation. This design is motivated by the space-radiation setting: candidate spacecraft electronics are known to be vulnerable to single-event effects such as SEU, SEL, and SET [42] under proton and heavy-ion exposure, and SEUs have been reported during in-orbit operation of LEO missions such as *Flying Laptop*. We therefore use the stress-set evaluation as a controlled proxy for whether a routing backend remains useful under fault-like conditions without violating latency requirements.

3.16 Routing Accuracy vs end-to-end refresh latency

To produce this comparison, we evaluate each backend within the same graph-routing pipeline and vary only the representation module or execution path. The underlying learner follows the message-passing paradigm for graph-structured inputs, while the photonic branch uses a boson-sampling-inspired stochastic feature map motivated by linear optics. This design isolates the practical question of whether a backend improves predictive quality enough to justify its runtime cost under periodic LEO routing refresh, which is a natural criterion in ISL routing studies where graph learning, low latency, and optical networking all matter simultaneously [3, 14, 55]. This is shown in **Fig. 18** in the Appendix.

Notations: Notations and interpretations are shown in **Table 3** of Section 6 in the Appendix.

4 Results and Discussions

This section focuses on some of the key results obtained and associated discussions. The supporting data, such as additional performance results from the quantum and classical backends, extended training-validation plots for each backend and tables of performance, boson sampling data histograms from ORCA backend etc. are presented in the Appendix.

4.1 Benchmarks: Quantum Diagnostics across CUDA-Q Backends

We first present our benchmarking of the quantum backend of CUDA-Q [9] as shown in **Fig. 5**. The quantum diagnostics reveal a clear trade-off between representation

richness and noise sensitivity in the VQC branch. The GPU-backed VQC shows substantially higher state diversity than the CPU-backed variant (approximately 0.442 versus 0.101), indicating that its node embeddings are less collapsed and more separated at the fidelity level.

In the language of variational quantum algorithms, this is consistent with a more expressive observable representation, which can be beneficial if the downstream classical network can exploit the added structure. At the same time, the GPU configuration also exhibits substantially larger expectation drift (approximately 0.316 versus 0.123), meaning that the measured $\langle Z \rangle$ responses are more sensitive to the injected noise model. This suggests that the GPU-executed VQC is operating in a regime that is more expressive but also less stable, a tension that is well aligned with the broader VQA literature on expressibility, trainability, and robustness. The figure shows that the GPU VQC configuration yields higher state diversity but also larger expectation drift, whereas the ORCA photonic module exhibits a modest but nonzero distribution shift under noise.

In **Fig. 5** and **Table 1**, the ORCA photonic branch shows a total variation distance of about 0.123 between ideal and noisy boson-sampling distributions, indicating a moderate but not catastrophic shift in the sampled output law. This implies that the photonic edge map is neither perfectly noise-invariant nor completely destabilized. The resource panel further clarifies that the CPU and GPU VQC experiments are not using different logical ansätze; both correspond to the same 4-qubit, depth-9 template with 28 total logical gates, split into 20×1-qubit and 8×2-qubit gates. Thus, the differences in diversity and drift should be interpreted as *backend-/execution-regime-dependent behavior of the same logical circuit*, not as consequences of different circuit sizes. Across the current benchmark, the *photonic branch* is the most attractive hybrid configuration in the low-latency regime: it typically matches or improves upon the classical baseline on held-out routing accuracy while remaining close to the lowest end-to-end refresh latency.

The *ORCA-edge* variant is also competitive in latency and can yield useful route-quality behavior on successful routes, although its gains are more mixed. By contrast, the *VQC variants* tend to incur substantially higher runtime cost, with the GPU-backed VQC often producing richer quantum-state structure than the CPU-backed VQC but not yet translating that advantage into superior route-level performance. The quantum diagnostics make this trade-off visible: the VQC branch can show greater node-state diversity and more pronounced expectation structure, but also greater sensitivity to noise and backend-dependent drift. The photonic diagnostics show that the ORCA/photonic module generates a structured boson-output distribution rather than unstructured noise, which supports its interpretation as a meaningful feature mechanism even when routing gains remain modest.

Since CUDA-Q explicitly supports multiple execution modes for the same kernel and separately supports orca-photonics for photonic simulation, the panel of **Fig. 5** is a compact way to show that *backend choice affects the qualitative behavior of the learned quantum representation even before route-level metrics are considered.* Following on from this, next, we present our results on routing.

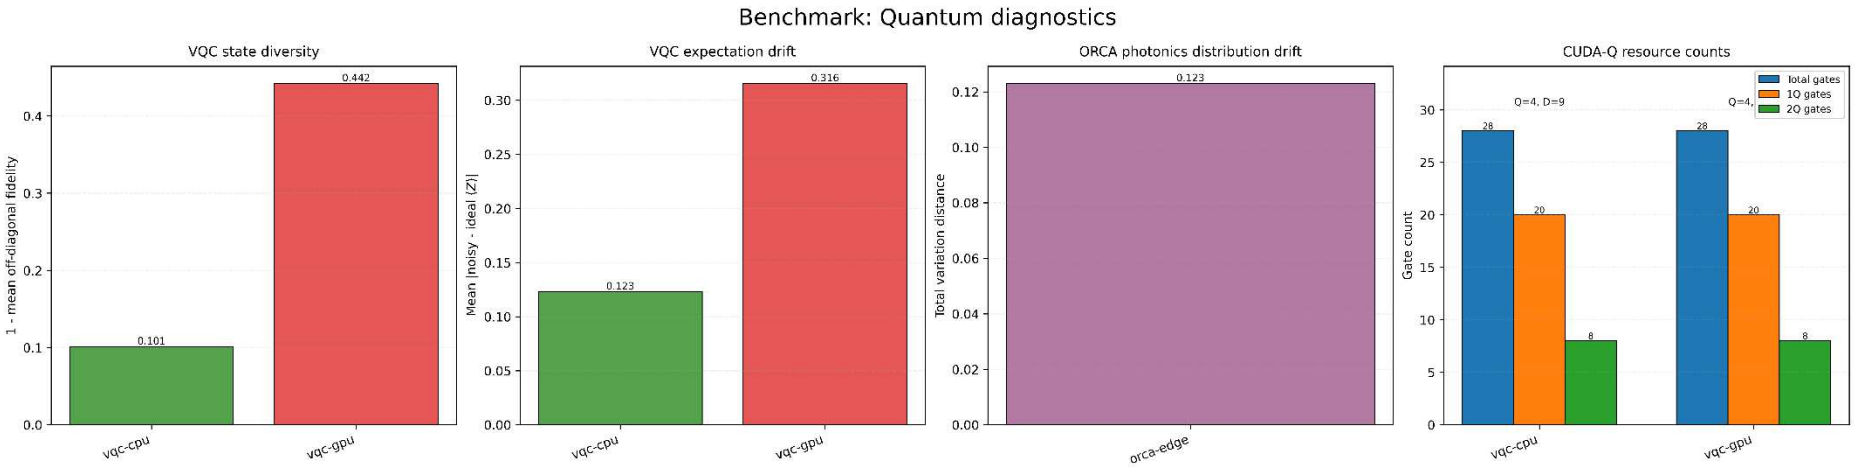


Fig. 5. Benchmark quantum diagnostics across CUDA-Q backends and the ORCA photonic edge module. The *left two* panels compare the VQC encoder under CPU (`qpp-cpu`) and GPU (`nvidia`) execution using two representation-level diagnostics: *state diversity*, computed as $1-$mean off-diagonal fidelity, and *expectation drift*, computed as the mean absolute difference between noisy and ideal $\langle Z \rangle$outputs. The *third* panel reports the *ORCA photonics distribution drift* as the total variation distance between ideal and noisy sampled boson-output distributions. The *right panel* shows logical resource counts for the VQC template, including total gates, one-qubit gates, and two-qubit gates; these remain identical across CPU and GPU because the logical circuit is unchanged, while the execution backend differs. **Table 1.** Backend configuration and ORCA precomputation. depicts the values of the backend configuration and ORCA precomputation for the quantum diagnostics benchmark.

Table 1. Backend configuration and ORCA precomputation.

Backend	Node mode	CUDA-Q target	ORCA edge	Status	ORCA total (s)	Avg / graph (s)	Avg/edge (s)	ms / shot /edge
Classical	classical_identity	—	No	ok	0.000	0.000	0.000	—
Photonic sampler	photonic_sampler	—	No	ok	0.000	0.000	0.000	—
CUDA-Q VQC (CPU)	cudaq_vqc	qpp-cpu	No	ok	0.000	0.000	0.000	—
CUDA-Q VQC (GPU)	cudaq_vqc	nvidia	No	ok	0.000	0.000	0.000	—
ORCA edge augmentation	classical_identity	—	Yes	ok	4.344	0.010	0.001	0.004

Note. Values are rounded to three decimals; — denotes not applicable.

4.2 Routing

We evaluate whether photonic sampling at the network edge improves routing quality over classical graph-based baselines in clustered optical satellite rings. Our experiments compare heuristic routing, a learned classical MPNN, and hybrid models that augment node representations with photonic sampling features under varying link conditions and radiation-induced device degradation. Across the main regimes of interest, the hybrid model improves routing robustness and preserves performance under moderate noise, while its gains diminish in the high-disorder regime where both classical and photonic features become less reliable.

Fig. 6 shows a representative destination-conditioned routing instance, where unlike the heuristic, which follows a fixed structural rule, the learned model conditions its decisions on graph topology, link context, and node-local photonic sampling features, enabling more adaptive routing behavior across representative clustered-ring instances.

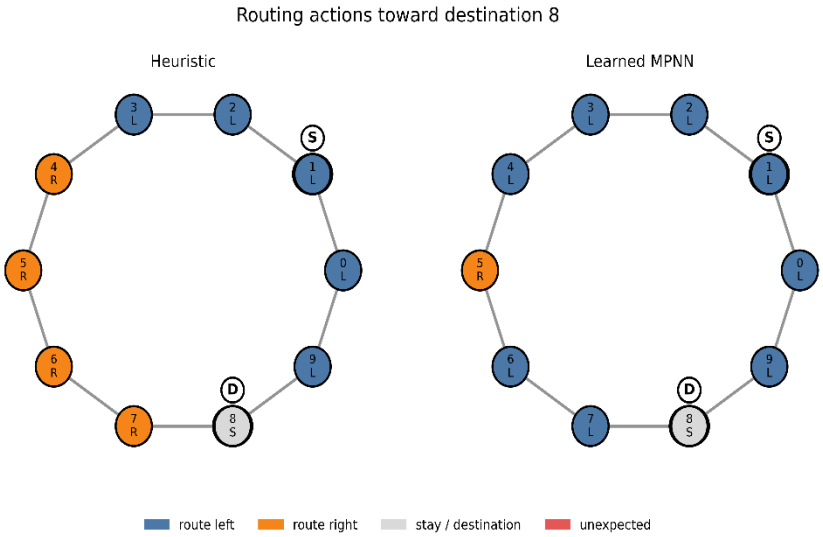


Fig. 6. Heuristic versus learned MPNN routing actions toward destination 8 on a clustered optical satellite ring. Each node is colored by the action selected for forwarding toward the fixed destination $D = 8$, with the source denoted as S: route left, route right, or stay at the destination. Relative to the heuristic baseline, the learned message passing neural network (MPNN) alters the routing pattern in a structured manner, changing the preferred forwarding direction for a subset of intermediate nodes while preserving coherent destination-directed behavior. This figure provides a qualitative view of how message passing modifies the induced routing policy on the ring.

We next evaluate the proposed hybrid photonic-classical routing model on clustered 10-node optical satellite rings with 30-second routing-table updates and radiation-aware node dynamics (**Fig. 7**). We compare heuristic routing, a learned classical MPNN, and hybrid variants that augment node representations with photonic sampling features. The main findings are that the hybrid model produces more topology-aware routing policies than heuristic baselines, remains stable under moderate loss and dark-count noise, and exhibits clear degradation only when radiation-induced device impairment becomes severe.

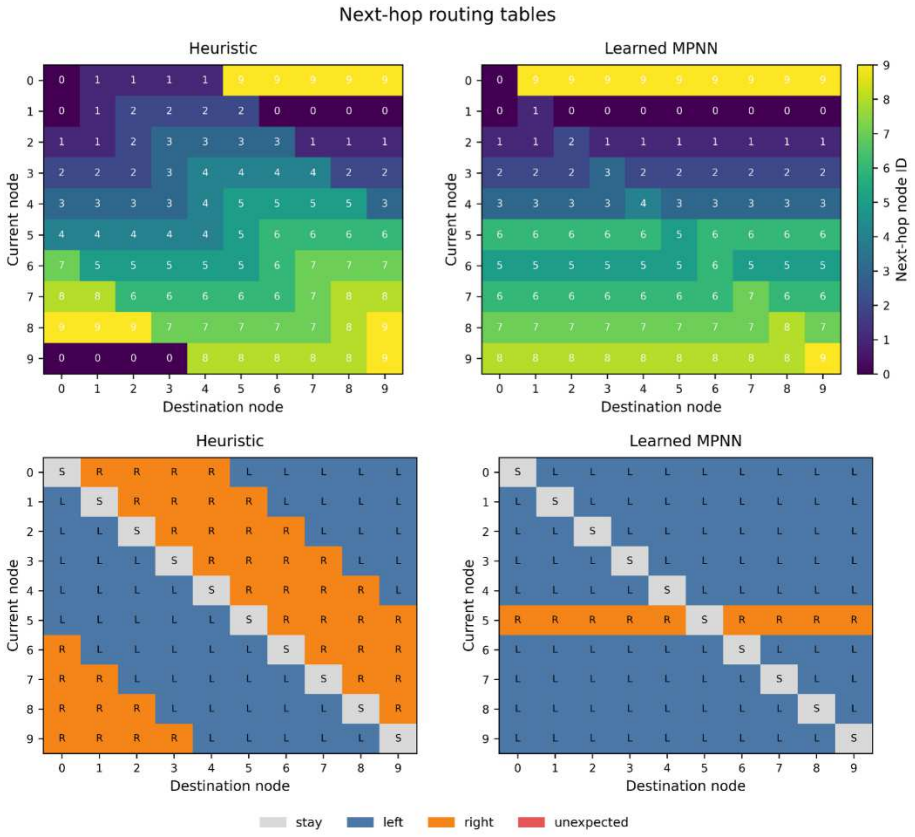


Fig. 7. Heuristic versus learned MPNN routing policies on clustered optical satellite rings. *Top*: next-hop routing tables for the heuristic baseline and the learned message passing neural network (MPNN), where each entry gives the selected next-hop node for a given current-node/destination pair. *Bottom*: the same policies expressed as action maps over the ring, where each entry indicates whether the selected action routes left, routes right, or stays at the destination. *Interpretation*: top line = node index; bottom line: L = send left; R = send right; S = stay; (destination) ; ? = unexpected next hop. *blue* = route left; *orange* = route right; *gray* = destination / stay; *red* = unexpected next hop.

Fig. 7 compares the routing policies induced by the heuristic baseline and the learned MPNN. The top row shows next-hop routing tables for the heuristic baseline and the learned message passing neural network (MPNN), where each entry gives the selected next-hop node for a given current-node/destination pair. The bottom row shows the corresponding action maps, indicating whether the policy routes left, routes right, or stays at the destination. Relative to the heuristic baseline, the learned MPNN induces a structured, destination-dependent routing policy, illustrating how message passing alters the resulting forwarding behavior. **Fig. 19** in Section 6.6 of the Appendix shows example of a disagreement map.

Route Latency–Utility Trade-offs. Fig. 8 highlights a latency–utility trade-off: the photonic and ORCA-edge variants achieve low-latency operation with positive congestion reductions, whereas the VQC variants incur substantially higher latency and mixed utility, with the CPU VQC underperforming the classical baseline on both metrics. Queue- and congestion-based comparisons are natural in dynamic multi-hop routing because queue backlog and link stress are standard proxies for network pressure and stability, while latency remains central in LEO ISL routing settings [54, 55].

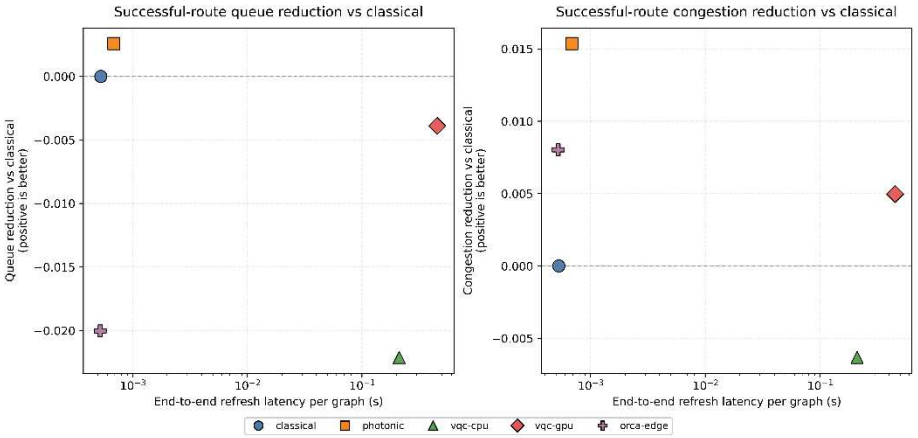


Fig. 8. Successful-route queue and congestion reduction versus classical baseline. Each point shows a routing backend positioned by end-to-end refresh latency per graph on the x-axis and by *successful-route-only* reduction relative to the classical baseline on the y-axis. Positive values indicate improvement over the classical route on the subset of rollouts that successfully reached the destination. The left panel reports queue reduction, while the right panel reports congestion-proxy reduction.

Fig. 8 compares routing backends by *successful-route-only* queue and congestion reduction relative to the classical baseline, plotted against end-to-end refresh latency per graph. This conditioning is important because failed or looping routes can distort path-quality averages and make short but invalid paths appear artificially favorable. Under this stricter view, the photonic and ORCA-edge variants occupy the most attractive low-latency region: the photonic branch achieves the strongest positive successful-route congestion reduction together with a small positive queue reduction, while the ORCA-edge branch also improves congestion at similarly low latency, albeit with worse queue behavior. By contrast, the VQC variants lie in a much higher-latency regime, with the GPU VQC offering only a modest congestion gain and slightly worse queue reduction, and the CPU VQC underperforming the classical baseline on both axes. Interpreted conservatively, the figure does not show that quantum backends uniformly improve routing, but rather that different hybrid subroutines induce different *latency–utility trade-offs*. This is a meaningful comparison in multi-hop routing, where queue backlog and congestion are standard proxies for network pressure and stability,

while latency remains central in LEO ISL operation, as emphasized in previous works [54, 55].

From a modeling perspective, the figure is significant because it links backend choice to network-level consequences rather than to prediction accuracy alone. The routing stack is still built around message passing on a graph, so the comparison isolates how the backend affects the quality of routes that actually succeed, while holding the graph-processing structure fixed [14]. Under that lens, the photonic backend is the clearest positive result in the current study: it remains in the lowest-latency regime while improving successful-route congestion and slightly improving queue reduction. The ORCA-edge backend is competitive but mixed, and the VQC branch currently appears to pay a substantial runtime cost for limited utility gains. The strongest conclusion is that *successful-route utility must be read jointly with latency*, and that the photonic branch currently offers the most favorable trade-off among the hybrid variants considered here [14, 54, 55].

Robustness vs Latency.

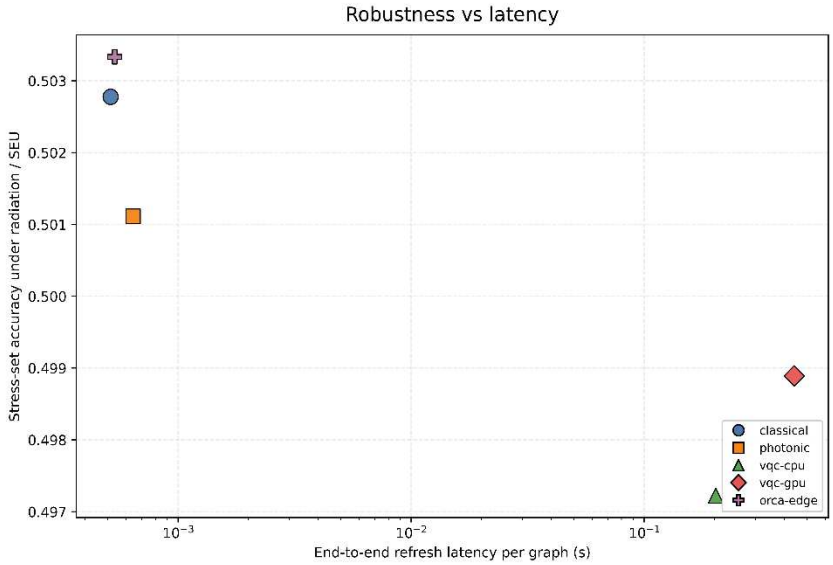


Fig. 9. Robustness–latency comparison under radiation/SEU stress. Each point shows a routing backend positioned by end-to-end refresh latency per graph on the x-axis and by stress-set accuracy on the y-axis, where the stress set applies radiation- and SEU-motivated perturbations to held-out routing graphs. Higher values indicate better predictive stability under fault-like conditions, while lower latency remains preferable for periodic routing refresh in LEO networking. This comparison is operationally meaningful because LEO routing must balance low-latency decision making with resilience to radiation-induced single-event effects, including SEUs and related transients that are well documented in spacecraft electronics and have been observed in LEO missions. **Table 2** shows a qualitative presentation.

Fig. 9 suggest that the most attractive robustness–latency trade-off in the current benchmark lies with the classical and ORCA-edge variants, which combine the highest stress-set accuracy with the lowest-latency regime. The photonic backend remains similarly fast but shows a modest drop in stressed accuracy, while the GPU VQC recovers competitive robustness only at substantially higher latency. The CPU VQC is dominated on both axes, combining the lowest stressed accuracy with the highest latency burden. The strongest interpretation is therefore not that one quantum backend uniformly outperforms the classical baseline, but that backend choice changes the *robustness frontier* in a measurable way: some hybrid variants preserve low-latency deployability better than others under radiation/SEU-like perturbations. This is the relevant lens for LEO routing, where low-latency networking and adaptive control are central concerns, and where radiation-driven single-event effects remain a real operational constraint.

Another interpretation is that this figure reflects *systems robustness*, not just prediction quality. The stress set is intended to proxy conditions motivated by the known susceptibility of spacecraft electronics to heavy-ion- and proton-induced single-event effects such as SEU, SEL (Single Event Latch-up) [13], and by in-orbit reports of SEUs in LEO spacecraft memory systems. Under that interpretation, the photonic and ORCA-edge branches are interesting because they stay in the low-latency region while remaining reasonably robust, whereas the VQC branches currently pay a substantial runtime cost for limited robustness gains. **Table 2** reflects the qualitative nature of the plots of **Fig. 9**.

Table 2. Stress-set robustness metrics.

Backend	Stress acc.	Stress-Route success	Stress hops	Acc. drop	Route-success drop
Classical	0.503	0.340	2.730	0.007	-0.050
Photonic sampler	0.501	0.275	2.700	0.005	0.010
CUDA-Q VQC (CPU)	0.497	0.275	2.670	0.012	0.040
CUDA-Q VQC (GPU)	0.499	0.240	2.635	0.016	0.025
ORCA edge augmentation	0.503	0.240	2.615	0.002	0.030

The table reports accuracy, route success, and mean hops on a stress validation set constructed with radiation/SEU-like perturbations. The final two columns report degradation relative to the corresponding nominal evaluation: Acc. drop is the reduction in routing accuracy under stress, and Route success drop is the reduction in end-to-end route success. Smaller drops indicate greater robustness, while lower stress hops should be interpreted together with route success rather than in isolation.

Table 2 shows that the strongest robustness profile under the current stress setting is achieved by the classical and ORCA edge augmentation backends, both of which attain the highest stress accuracy (0.503). The two methods differ, however, in route-level

behavior: the classical baseline retains the highest stress route success (0.340) with a moderate accuracy drop (0.007), while ORCA edge augmentation has the smallest accuracy drop (0.002) and the shortest stress-set routes (2.615 hops) but a lower stress route success (0.240). The photonic sampler remains competitive, with stress accuracy 0.501, route success 0.275, and a small accuracy drop (0.005), indicating that it preserves much of the classical robustness while remaining slightly weaker at the end-to-end route level. By contrast, the CUDA-Q VQC variants are the least robust in this comparison: the CPU and GPU versions have lower stress accuracy (0.497 and 0.499), larger accuracy drops (0.012 and 0.016), and reduced stress route success relative to the classical baseline. The GPU VQC is particularly notable for the largest accuracy drop, suggesting that its richer learned representation does not translate into greater stability under stress in the present configuration.

*Routing Accuracy vs End-to-End Refresh latency***:** Among the low-latency backends, the photonic branch achieved the highest routing accuracy, while the VQC backends incurred substantially higher latency for limited accuracy gains. The full results plot, associated results and discussions can be found in Section 6.6 of the Appendix showing the backend comparison in implementation trade-offs [55].

4.3 Fidelity Matrix

The first left-hand side matrices of **Fig. 10** and **Fig. 11** correspond to the same routing domain and the second right-hand matrices correspond to another routing domain.

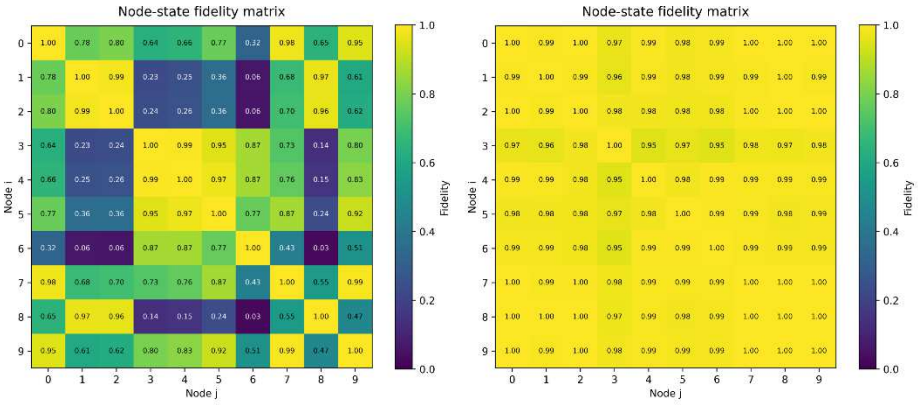


Fig. 10. Comparison *of node-state fidelity matrices for two quantum-routing configurations.* Each entry F_{ij} measures the pairwise fidelity between the quantum states associated with nodes $i$and j, where fidelity for pure states is the squared overlap between state vectors. The *left* matrix exhibits a broad spread of off-diagonal values, indicating that the learned quantum embedding separates several node pairs and preserves nontrivial state diversity. The *right* matrix is dominated by near-unity off-diagonal fidelities, indicating that most node states are nearly identical and that the quantum embedding has largely collapsed to a narrow region of state space. This contrast suggests that the *left configuration provides a richer node representation* for downstream routing than the right configuration.

The fidelity diagnostics of **Fig. 10** *top row* suggest that the two configurations occupy different representational regimes. In the first matrix, low-fidelity pockets and heterogeneous off-diagonal structure indicate that the quantum encoder is assigning different latent states to different node contexts. This is precisely the behavior one would hope for from a hybrid VQC-based node encoder, since distinct node states can then be propagated and combined by the classical message-passing network. In the second matrix, by contrast, almost all off-diagonal entries are very close to 1, which means that the node states are nearly identical and the quantum encoder is providing very little discriminative structure. Since fidelity for pure states reduces to squared state overlap, a matrix with many low off-diagonal values signals greater representational diversity, whereas a matrix with uniformly high off-diagonal values signals state concentration and weak separation [22, 32]. Under that condition, the quantum block contributes little effective discrimination. Hence, the first matrix is consistent with a *useful quantum embedding*, whereas the second is more consistent with *state collapse or underutilization of the circuit.*

From a modeling perspective, these results are more informative than a single scalar fidelity summary because they reveal *where* discrimination is occurring. The first matrix shows that some nodes remain highly similar while others are clearly separated, suggesting structured clustering rather than random dispersion. That is an encouraging pattern for routing, since the task does not require every node to be maximally distinct; it requires the quantum representation to preserve distinctions that matter for downstream decisions. The second matrix, however, indicates that almost all pairwise relationships are treated as nearly equivalent, which is unlikely to help routing beyond what a purely classical encoder could already provide. This qualitative difference supports a stronger paper claim than "the fidelities are different": namely, that one configuration is actually producing a more informative geometric organization of node states for the hybrid routing pipeline.

The *node-state fidelity matrices* provide a compact view of *how strongly the quantum encoder separates different nodes in Hilbert space* [59]. In the first matrix, the off-diagonal fidelities span a wide range, including several low-similarity pairs alongside near-unity entries, which indicates that the encoder is producing *structured but non-collapsed* node embeddings. In the second matrix, by contrast, almost all off-diagonal values are close to one, implying that most nodes are mapped to nearly the same quantum state. Since pure-state fidelity is the squared overlap between state vectors, the first matrix is consistent with a more discriminative representation, while the second is more consistent with *state concentration* or effective collapse of the quantum embedding [22, 32, 59]. In practical routing terms, the first matrix is more favorable because it suggests that the downstream classical message-passing and policy layers receive distinguishable quantum features for different nodes, while the second matrix suggests that the quantum block is contributing only a weakly varying signal.

This diagnostic should nevertheless be interpreted as a *representation analysis*, not as a standalone routing metric. Greater state diversity does not automatically imply better route success, but it does indicate that the variational circuit is making fuller use of its available state space. This is aligned with the broader literature on variational quantum algorithms, where expressibility and trainability must be balanced rather than

maximized independently: a circuit should separate task-relevant inputs without becoming so random or unstable that optimization deteriorates [7, 19, 51]. Accordingly, we use the fidelity matrix as qualitative evidence about the geometry of the learned quantum representation. Under that interpretation, the first matrix provides stronger support for a useful quantum node encoder, whereas the second suggests that the VQC is contributing little discriminative structure beyond what the classical routing stack could already exploit.

The expectation-map comparison for **Fig. 11** supports the same qualitative conclusion suggested by the fidelity analysis of **Fig. 10**: one configuration is using the quantum block in a more informative way than the other. In the left panel, different nodes excite different qubit channels, with some nodes showing weak or negative $\langle Z \rangle$ responses and others showing strong positive values. In the right panel, most nodes activate nearly the same pattern, with qubit 2 especially close to a common high-positive response. From a routing perspective, the left panel is therefore more encouraging, because it suggests that the quantum encoder is producing node-dependent observables that can, in principle, support different downstream decisions. The right panel instead suggests that much of the quantum output is shared across nodes, which reduces its value as a discriminative latent representation. This does not prove better routing on its own, but it does show that the two configurations differ materially in the geometry of the observable features they expose to the classical model [1, 7].

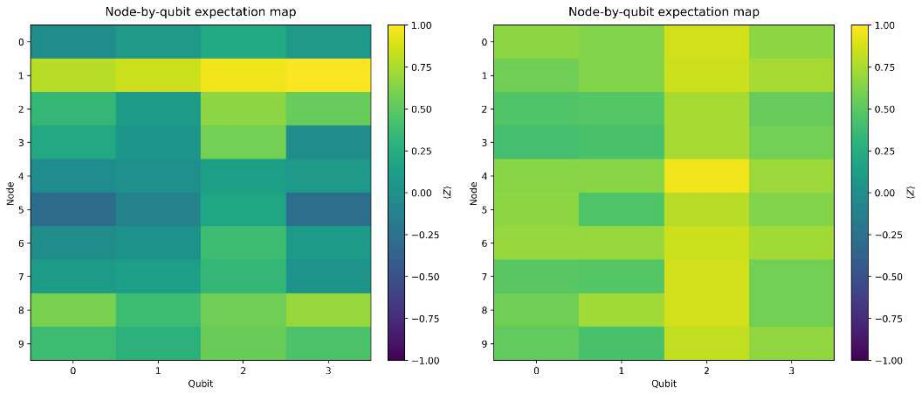


Fig. 11. Comparison *of Node-by-qubit expectation maps for two quantum-routing configurations.* Each entry shows the expectation value $\langle Z_q \rangle$ of qubit $q$for the node-specific quantum state produced by the trained encoder. The left panel exhibits stronger heterogeneity across nodes and qubits, indicating a more varied observable signature. The right panel is dominated by uniformly positive responses, especially along one qubit channel, suggesting a more concentrated observable representation. Because expectation values are the quantities passed from the parameterized quantum circuit to the classical learning pipeline in many variational quantum algorithms, these maps provide a compact qualitative diagnostic of how much node-dependent structure the quantum encoder is exposing to the downstream routing model [7].

The *node-by-qubit expectation maps* (**Fig. 11**) provide a compact view of the *observable signatures* produced by the trained quantum encoder. In one regime, the map is

visibly heterogeneous across both nodes and qubits, with mixed magnitudes and more node-specific structure; in the other, most entries are uniformly positive, and one qubit channel is close to saturation across many nodes. Interpreting these entries as per-node $\langle Z_q \rangle$ readouts, the first regime suggests that the encoder is exposing a richer set of measurable features to the downstream classical routing model, whereas the second is more consistent with an *observable-level concentration* in which many nodes look similar through the selected measurements. This distinction matters because, in variational quantum models, expectation values are the main classical interface used for learning and inference, so their diversity directly affects how much task-relevant information reaches the classical message-passing and policy layers [7].

At the same time, these maps should be interpreted as a *representation diagnostic*, not as a standalone routing metric. Greater variation in $\langle Z_q \rangle$ does not by itself guarantee better route success, but it does indicate that the circuit is using its measurement channels in a more discriminative way. This is aligned with prior work showing that parameterized quantum circuits should provide enough expressibility to separate inputs meaningfully, without collapsing to nearly constant outputs or becoming so unstructured that trainability deteriorates [1, 19, 51]. Under that interpretation, the more heterogeneous expectation map is the stronger qualitative result, because it suggests a node-dependent quantum representation that is more likely to benefit the hybrid routing stack than a map dominated by nearly uniform responses.

Fig. 10 and **Fig. 11** show that the interpretation should remain nuanced. High state diversity is not automatically better in every setting, and a fidelity matrix alone does not prove improved trainability or better routing decisions. Prior work on parameterized quantum circuits shows that expressibility, trainability, and gradient behavior are coupled in nontrivial ways: highly expressive ansätze can improve state coverage, but they can also become harder to optimize or approach barren-plateau-like regimes if the circuit becomes too random or too close to a 2-design. The ideal outcome is therefore not "minimum fidelity everywhere," but rather *task-relevant separation*: enough diversity to distinguish routing contexts, without driving the circuit into a regime where learning becomes unstable or uninformative [1, 7, 19, 51]. On that criterion, the first matrices in both **Fig. 10** and **Fig. 11**, which corresponds to the same routing domain, is the stronger qualitative result because it demonstrates selective separation rather than near-complete collapse

4.4 Fidelity Output of ORCA Photonics

Fidelity output in ORCA is shown in **Fig. 15** of Appendix. The behavior is consistent with the similarity of local neighborhoods across nodes and suggests that downstream routing gains, if present, arise from relatively subtle distributional differences rather than from strongly separated node signatures. In our setting, this behavior is plausible because many satellites occupy structurally similar local neighborhoods, so the photonic sampler is operating in a low-contrast regime: it produces stable distributions, but only modest node-to-node separation. Consequently, any downstream gain in routing quality must arise from relatively subtle, destination-relevant differences in these sampled distributions rather than from large pairwise distinctions. This interpretation is

consistent with prior work on bosonic-sampling-based graph similarity, where useful structure is often encoded in distributional overlaps rather than in mutually orthogonal node signatures [50]. The *photonic edge map* produced a structured, highly nonuniform boson-output distribution (Appendix, **Fig. 15**), confirming that the sampling block operated in a nontrivial interference regime [3].

The photonic edge module (**Fig. 4** and **Fig. 15**) is motivated by the standard boson-sampling model of linear optics. For an interferometer with unitary transfer matrix U, the probability of observing an output occupation pattern is governed by the squared magnitude of a matrix permanent of the corresponding submatrix of U, with factorial normalization for repeated occupations. In the collision-free case this reduces to the squared permanent of an $n \times n$ submatrix determined by the chosen input and output modes [3]. We therefore treat the photonic block as a *structured stochastic feature generator*: its sampled statistics are not arbitrary random features, but are determined by interference patterns induced by the programmed phases and beam splitters of the underlying linear-optical network.

4.5 GNN Policy Training with CUDA-Q ORCA Photonics

Fig. 12 shows the routing backbone which, as defined before, is an MPNN [14], while the quantum-enhanced node features are generated by a photonic layer executed through CUDA-Q's ORCA photonic stack, which supports both ORCA hardware backends and the orca-photonics simulator. The early transient in L_{cons} reflects mismatch between policies induced by independent photonic draws, whereas the subsequent decay of all losses indicates convergence to a routing policy that is both label-consistent and stable under sampling noise.

The plot in **Fig. 12** shows rapid convergence of the hybrid routing model. The supervised term L_{sup} starts high, then drops sharply within roughly the first 15–20 epochs, indicating that the model quickly fits the labeled next-hop actions. The consistency term L_{cons} exhibits a transient spike early in training before collapsing toward zero; this is a reasonable pattern when the policy is trained to agree across *independent photonic draws*, since the model initially sees the two stochastic views as inconsistent and only later learns a stable representation. After about 25–35 epochs, all three curves are essentially flat near zero, suggesting that the learned policy has reached a stable fixed point.

An interpretation is that the model is not only fitting the supervised routing labels, but is also learning to make *draw-invariant decisions* despite stochasticity in the photonic layer. In other words, the message-passing network is becoming robust to photonic sampling noise. At the same time, the very fast collapse of all losses suggests that the task may be relatively easy in the present 10-node clustered ring, so this figure should be read as evidence of *optimization stability* rather than, by itself, evidence of superior generalization.

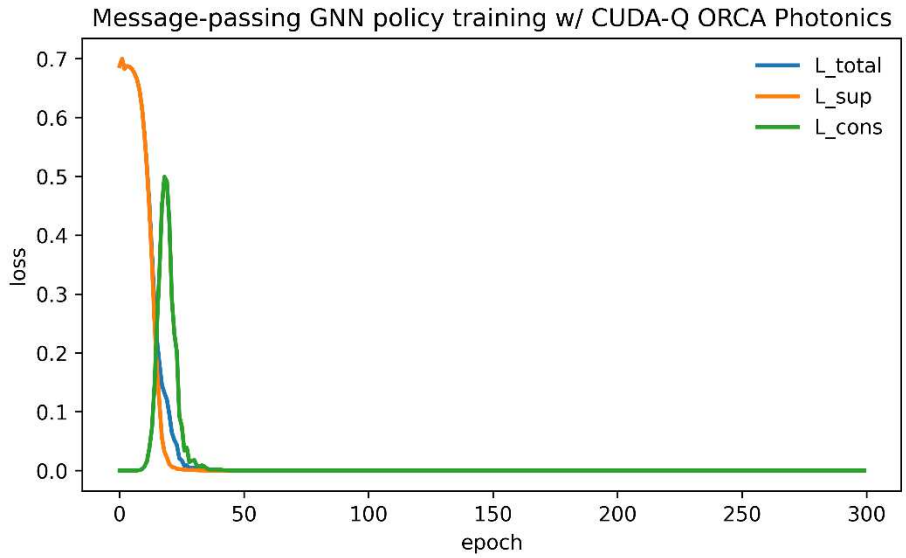


Fig. 12. Training dynamics of the hybrid message-passing routing policy with CUDA-Q ORCA photonics**.** The total loss L_{total}, supervised cross-entropy term L_{sup}, and consistency regularization term L_{cons} are shown as functions of training epoch.

Training dynamics and induced routing behavior: **Fig. 12** shows that the hybrid photonic-classical policy converges rapidly: the supervised term L_{sup} falls sharply within the first few tens of epochs, while the consistency term L_{cons} exhibits a brief transient before collapsing toward zero. This pattern is consistent with the role of consistency regularization in forcing agreement between policies induced by independent stochastic photonic draws [27, 54]. Crucially, the optimization dynamics are matched by the qualitative policy visualizations in **Fig. 5** and **Fig. 6**.

The learned route table does not differ from the heuristic baseline through isolated entry-wise perturbations; instead, it departs in coherent, destination-dependent blocks, and the corresponding action maps reveal a topologically consistent reorganization of left/right forwarding decisions on the ring. Such behavior is precisely what one would expect from a message-passing architecture, whose inductive bias couples routing decisions across neighboring satellites rather than treating each table entry independently [4, 14]. Taken together, the loss curve and the policy visualizations indicate that the photonic-augmented MPNN is not merely fitting labels, but is converging to a stable graph-aware routing policy whose learned deviations from the heuristic are structured and interpretable.

4.6 Backend Comparison

Fig. 13 compares validation accuracy over training for the `classical_identity` and `photonic_sampler` backends. The classical baseline reaches a clearly higher best validation accuracy and maintains an advantage over the photonic model for most

epochs. The photonic sampler exhibits a smoother but consistently lower validation trajectory, suggesting that its boson-sampling-inspired feature map is changing the optimization profile without yet improving held-out predictive performance. In the present configuration, the strongest claim is therefore not that the photonic backend outperforms the classical one, but that the hybrid feature map induces a meaningfully different learning dynamic. Because the classical curve also declines after its best epoch, the comparison further suggests that model selection should be based on validation checkpoints rather than final-epoch performance, consistent with standard early-stopping practice [44].

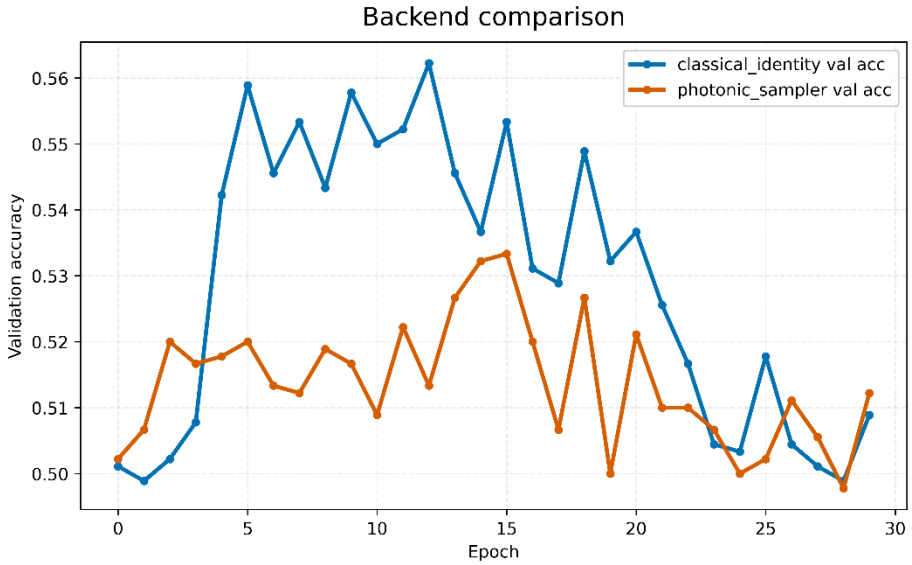


Fig. 13. Validation-accuracy comparison across training epochs for the classical_identity and photonic_sampler routing backends. The classical baseline achieves the higher peak validation accuracy and dominates for most of training, while the photonic sampler follows a flatter, lower-variance trajectory and remains below the best classical checkpoint.

Fig. 13 suggests that, in the current configuration, the photonic stochastic feature map changes optimization behavior but does not yet outperform the classical message-passing baseline on validation accuracy. Validation-based checkpointing or early stopping would be more appropriate than selecting the final epoch by default [44]. The post-peak decline of the classical curve also shows that final-epoch accuracy is not the right summary statistic here. A more defensible protocol is to report the best validation checkpoint or to use explicit early stopping, since validation performance begins to degrade after the strongest middle-epoch results. Under that reading, the main conclusion is that the classical backend is currently the best-performing choice on validation accuracy, while the photonic backend remains valuable as a hybrid feasibility result whose main contribution is to alter representation and learning behavior rather than to maximize peak score.

4.7 ORCA Photonics and Boson Sampling

Fig. 14 supports a mechanistic validation claim for the photonic branch, with the interferometer schematic depicted in **Fig. 4** of the methods section. The top-k histogram shows that the output mass is concentrated on a small number of occupation patterns, with one dominant pattern and a long tail of lower-probability alternatives, while the diagnostics panel shows nontrivial collision and no-collision probabilities together with high normalized entropy. This indicates that the photonic module is operating in a meaningful interference regime: it is neither flat nor purely deterministic but instead exposes a structured stochastic signature that can be used as an edge-level feature source downstream. That interpretation is exactly the one suggested by the standard boson-sampling picture, where the output law is determined by the interferometer rather than by arbitrary randomization [3]. This is consistent with standard boson-sampling model, in which identical photons are sent through a linear-optical network and then non-adaptively measured by counting photons in each mode, and with CUDA-Q's orca-photonics backend, which provides a photonic state-vector simulator with photonic operations such as create, `phase_shift`, `beam_splitter`, and `mz` measurement.

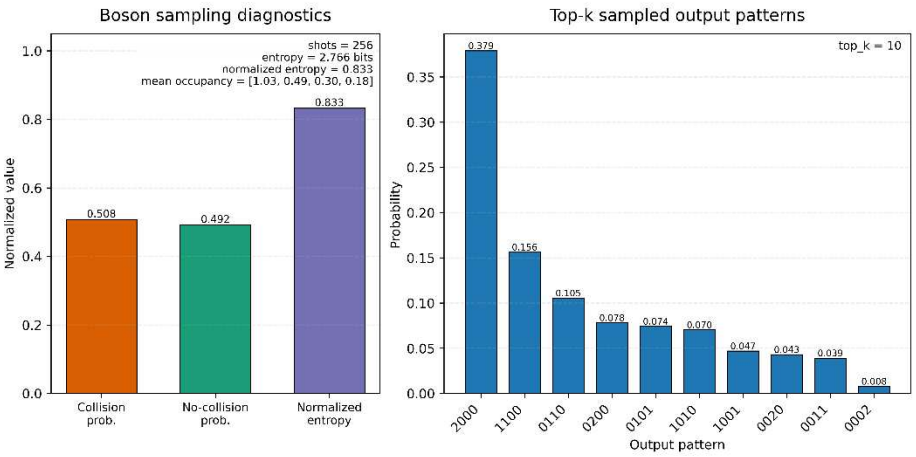


Fig. 14. ORCA photonics edge counts: (the interferometer schematic is shown in **Fig. 4**) boson-sampling diagnostics, and top-k sampled output patterns. The left panel summarizes three distributional diagnostics of the sampled output law: collision probability, no-collision probability, and normalized entropy. The right panel shows the top-k sampled output occupation patterns and their empirical probabilities. Together, the panels show that the photonic edge module produces a structured, highly nonuniform output distribution characteristic of boson-sampling-style linear optics rather than an unstructured random feature source.

The right-hand histogram of the Top-k Sampled Output patterns shows that the photonic module produces a *highly concentrated but structured* output distribution. The pattern 2000 dominates the sampled mass, followed by a smaller group of patterns such as 1100, 0200, and 0110, while the remaining top-ranked outputs have substantially

lower probability. This matters because the figure demonstrates that the photonic edge map is not behaving like an undifferentiated random feature generator. Instead, it is producing a stable interference-driven distribution over output occupations, which is the expected qualitative behavior of a boson-sampling-inspired linear-optical block [3]. In the context of the routing model, the histogram therefore serves as evidence that the photonic branch is injecting a *structured stochastic signature* into the pipeline rather than simple noise.

At the same time, this figure does not by itself establish improved routing quality; rather, it shows that the photonic subroutine is operating in a meaningful sampling regime and exposing nontrivial distributional information that can be used downstream. That makes the figure valuable as a qualitative or mechanistic diagnostic: it supports the claim that the photonic edge component is genuinely boson-sampling-like in behavior, even if route-level utility must still be judged by routing accuracy, route success, latency, and congestion metrics. CUDA-Q's photonics documentation makes this especially relevant methodologically, because orca-photonics is explicitly designed to simulate photonic circuits and sample their output statistics within the same hybrid framework used for the rest of the benchmark [34]. Route-level utility still needs to be judged separately using routing accuracy, route success, latency, and congestion metrics [34, 35].

4.8 Radiation Model output

The radiation plot and associated explanations for the toy LEO model is shown in **Fig. 17** in Section 6.6 of the Appendix.

4.9 Training Performances

All training loss, validation and accuracy plots, associated values and explanations (**Table 4** and **Table 5**) are shown in Section 6.10 of the Appendix. In summary, the photonic sampler provides the strongest *late validation* behavior, whereas the VQC branch reaches the strongest early validation accuracy but requires aggressive checkpointing because of later instability.

5 Sustainability, Limitations, Conclusions and Outlook

5.1 Sustainability

This study was designed as a small-scale hybrid-routing benchmark rather than a compute-intensive training campaign. Experiments were conducted on a 10-node optical LEO ring with shallow quantum circuits, modest shot budgets, and limited graph counts, and were implemented within a shared CUDA-Q workflow spanning classical, GPU-accelerated, and photonic simulation backends. To reduce computational cost, we reused common datasets across backend comparisons, kept circuit depth and problem size intentionally small, and reported latency- and route-level trade-offs rather than

scaling experiments aggressively. No large foundation-model training was performed. Hardware details, software versions, and benchmark runtimes are reported in the supplementary material to support reproducibility and downstream energy or carbon accounting.

5.2 Limitations

The main limitation of this study is that it is a small-scale simulation benchmark rather than a deployment-ready evaluation. The 10-node optical ring provides a controlled setting for comparing classical, VQC, and photonic backends, but it does not capture the scale or heterogeneity of real LEO constellations. The quantum components are evaluated only through CUDA-Q simulation backends, not on in-orbit hardware, and the learning objective is primarily next-hop supervision rather than a fully end-to-end routing objective. The radiation/SEU perturbations are also simplified stress proxies rather than hardware-validated fault models. For these reasons, the results should be interpreted as demonstrating backend-aware hybrid-routing feasibility and representation effects, not as establishing a deployable quantum routing advantage.

5.3 Conclusions

We conclude that the most appropriate architecture for the present problem is classical control for routing, with the quantum components used to enrich state representations rather than to replace the routing controller itself. The routing plane must enforce graph masks, destination admissibility, refresh timing, and other protocol-level constraints that are naturally represented and audited in a classical loop. By contrast, boson sampling is fundamentally a sampling model over linear-optical networks, and current programmable photonic platforms are built around reconfigurable, classically tuned elements such as phase shifters and interferometers. For this reason, a classically supervised control plane is the most credible systems design for the current study.

Within that classical outer loop, the VQC and photonic modules are best interpreted as controlled subroutines. This is consistent with the standard hybrid variational paradigm, in which a quantum circuit evaluates quantities that are expensive or structured in a useful way, while a classical optimizer or controller updates parameters and makes the final decision. The node-local VQC serves as a trainable embedding block, while the photonic boson-sampling module serves as a stochastic edge-feature generator whose outputs are consumed by the classical routing policy.

Outlook and Future Work. A natural direction for future work is to move beyond the current small-scale ring benchmark toward larger, time-varying LEO topologies with more realistic orbital dynamics, heterogeneous traffic, and longer routing horizons. On the learning side, the present next-hop supervision could be replaced or complemented by route-level objectives that directly optimize end-to-end success, loop avoidance, congestion reduction, and queue stability. On the quantum side, a key next step is to improve the efficiency and stability of the hybrid modules, for example by using more scalable photonic feature maps, better-regularized VQC encoders, and measurement-

adaptive updates across refresh intervals. It would also be valuable to extend the backend study to stronger multi-GPU and multi-QPU settings, and to test whether backend-aware training can better translate representational gains into route-level utility. Finally, a longer-term systems direction is to connect these hybrid routing methods to emerging space compute infrastructures, including optical networking for AI workloads and future orbital datacenter platforms, where backend-accelerated quantum simulation and photonic inference may become increasingly relevant.

Funding Declaration: The authors did not receive any funding support from any organization for the submitted work.

References

1. Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., Woerner, S.: The power of quantum neural networks. Nature Computational Science 1(6), 403–409 (2021)
2. Aharonovich, I., Crozier, K.B., Neshev, D.: Programmable integrated quantum photonics. Nature Photonics 20, 254–265 (2026)
3. Aaronson, S., Arkhipov, A.: The computational complexity of linear optics. Theory of Computing 9(4), 143–252 (2013)
4. Battaglia, P.W., Hamrick, J.B., Bapst, V., et al.: Relational inductive biases, deep learning, and graph networks. arXiv:1806.01261 (2018)
5. Bhattacharyya, A.: On a measure of divergence between two statistical populations defined by their probability distributions. Bulletin of the Calcutta Mathematical Society 35, 99–109 (1943)
6. Bhattacharyya, A.: On a measure of divergence between two multinomial populations. Sankhyā 7(4), 401–406 (1946)
7. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S.C., Endo, S., Fujii, K., McClean, J.R., Mitarai, K., Yuan, X., Cincio, L., Coles, P.J.: Variational quantum algorithms. Nature Reviews Physics 3(9), 625–644 (2021)
8. Cimini, V., Valeri, M., Piacentini, S., Ceccarelli, F., Corrielli, G., Osellame, R., Spagnolo, N., Sciarrino, F.: Variational quantum algorithm for experimental photonic multiparameter estimation. npj Quantum Information 10, 26 (2024)
9. CUDA-Q development team: CUDA-Q. https://github.com/NVIDIA/cuda-quantum (2024)
10. Edwards, B.L.: An overview of NASA's optical communications efforts. Presentation at OFC 2022, NASA Goddard Space Flight Center (2022)
11. Fuchs, C.A., van de Graaf, J.: Cryptographic distinguishability measures for quantum mechanical states. arXiv:quant-ph/9712042 (1998)
12. Ganguly, S.: Quantum Machine Learning: An Applied Approach: The Theory and Application of Quantum Machine Learning in Science and Industry. Apress (2021). https://doi.org/10.1007/978-1-4842-7098-1
13. Gaza, R., Ladbury, R.L., Allen, G.R., Ochs, L., Hodson, R.F.: Single-event latch-up in commercial electronics: Risk assessment and mitigation. NASA presentation for JEDEC JC13.4 / SAE CE12 (2024)
14. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: Proceedings of the 34th International Conference on Machine Learning, pp. 1263–1272 (2017)
15. Hamilton, C.S., Kruse, R., Sansoni, L., Barkhofen, S., Silberhorn, C., Jex, I.: Gaussian boson sampling. Physical Review Letters 119(17), 170501 (2017)

16. Heyraud, V., Li, Z., Donatella, K., Le Boité, A., Ciuti, C.: Efficient estimation of trainability for variational quantum circuits. PRX Quantum 4, 040335 (2023)
17. Hoch, F., Caruccio, E., Rodari, G., Francalanci, T., Suprano, A., Giordani, T., Carvacho, G., Spagnolo, N., Koudia, S., Proietti, M., et al.: Quantum machine learning with Adaptive Boson Sampling via post-selection. Nature Communications 16, 902 (2025)
18. Hoch, F., Rodari, G., Giordani, T., Spagnolo, N., Carvacho, G., Sciarrino, F.: Variational approach to photonic quantum circuits via the parameter shift rule. Physical Review Research 7, 023227 (2025)
19. Holmes, Z., Sharma, K., Cerezo, M., Coles, P.J.: Connecting ansatz expressibility to gradient magnitudes and barren plateaus. PRX Quantum 3, 010313 (2022)
20. Jia, Y., Zhang, H., Tian, G., et al.: High-throughput data routing technology for low-Earth-orbit mega-constellations all-optical networks. Aerotecnica Missili & Spazio (2025). https://doi.org/10.1007/s42496-025-00288-y
21. Jiang, W., Han, H., Zhang, Y., Wang, J., He, M., Gu, W., Mu, J., Cheng, X.: Graph neural networks for routing optimization: Challenges and opportunities. Sustainability 16(21), 9239 (2024)
22. Jozsa, R.: Fidelity for mixed quantum states. Journal of Modern Optics 41(12), 2315–2323 (1994)
23. Jun, I., Garrett, H., Kim, W., Zheng, Y., Fung, S.F., Corti, C., Ganushkina, N., Guo, J.: A review on radiation environment pathways to impacts: Radiation effects, relevant empirical environment models, and future needs. Advances in Space Research (2024)
24. Kedrowitsch, A., Black, J., Yao, D.: Resilient routing for low Earth orbit mega-constellation networks. In: Proceedings of the NDSS Workshop on the Security of Space and Satellite Systems (SpaceSec 2024) (2024)
25. Kullback, S., Leibler, R.A.: On information and sufficiency. The Annals of Mathematical Statistics 22(1), 79–86 (1951)
26. Ladbury, R.L.: Radiation environment effects on spacecraft. NASA Goddard Space Flight Center presentation (2017)
27. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: Proceedings of the International Conference on Learning Representations (2017)
28. Li, J., Yang, X., Peng, X., Sun, C.-P.: Hybrid quantum-classical approach to quantum optimal control. Physical Review Letters 118(15), 150503 (2017)
29. Lozano-Cuadra, F., Soret, B., Leyva-Mayorga, I., Popovski, P.: Continual deep reinforcement learning for decentralized satellite routing. arXiv:2405.12308 (2024)
30. NASA: Optical communications overview. NASA Space Communications and Navigation. https://www.nasa.gov/technology/space-comms/optical-communications-overview/ (2023)
31. NASA: Laser communications. https://www.nasa.gov/communicating-with-missions/laser-comms/ (2026)
32. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press (2000)
33. NVIDIA: How Starcloud is bringing data centers to outer space. NVIDIA Blog. https://blogs.nvidia.com/blog/starcloud/ (2025)
34. NVIDIA: CUDA-Q documentation: Quick start. https://nvidia.github.io/cuda-quantum/latest/using/quick_start.html (2026a)
35. NVIDIA: CUDA-Q circuit simulation backends. https://nvidia.github.io/cuda-quantum/latest/using/backends/simulators.html (2026b)
36. NVIDIA: CUDA-Q multiple QPUs. https://nvidia.github.io/cuda-quantum/latest/using/backends/sims/mqpusims.html (2026c)

37. NVIDIA: CUDA-Q photonics simulators. https://nvidia.github.io/cuda-quantum/latest/using/backends/sims/photonics.html (2026d)
38. NVIDIA: NVIDIA launches space computing, rocketing AI into orbit. NVIDIA Newsroom. https://nvidianews.nvidia.com/news/space-computing (2026f)
39. Noeldeke, C., Boettcher, M., Mohr, U., Gaisser, S., Alvarez Rua, M., Eickhoff, J., Leslie, M., Von Thun, M., Klinkner, S., Varatharajoo, R.: Single event upset investigations on the "Flying Laptop" satellite mission. Advances in Space Research 67(6), 2000–2009 (2021)
40. Nunn, J.: ORCA computing: the route from NISQ processors to fault-tolerance. In: Proceedings of SPIE 12446, Quantum Computing, Communication, and Simulation III, 124460D (2023). https://doi.org/10.1117/12.2651450
41. O'Bryan, M.V., Roffe, S.S., Wilcox, E.P., Campola, M.J., Osheroff, J.M., Casey, M.C., Joplin, M.B., Carstens, T.A., Barth, J.D., Ryder, L.D., Ryder, K.L., Wood, A.M., Lauenstein, J.-M., Majewicz, P.J.: Compendium of NASA Goddard Space Flight Center's recent radiation effects test results. In: IEEE Radiation Effects Data Workshop (2024)
42. Pellish, J.A.: Tutorial: Radiation Effects in Electronic Systems. NASA Goddard Space Flight Center presentation at the Applied Space Environments Conference (ASEC), Huntsville, AL, USA (2017)
43. Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P.J., Aspuru-Guzik, A., O'Brien, J.L.: A variational eigenvalue solver on a photonic quantum processor. Nature Communications 5, 4213 (2014)
44. Prechelt, L.: Automatic early stopping using cross validation: quantifying the criteria. Neural Networks 11(4), 761–767 (1998)
45. Qi, H., Xiao, S., Liu, Z., Gong, C., Gani, A.: Variational quantum algorithms: fundamental concepts, applications and challenges. Quantum Information Processing 23, 224 (2024)
46. Quadros de Aguiar, Y., Wrobel, F., Autran, J.-L., García Alía, R.: Radiation environment and their effects on electronics. In: Single-Event Effects, from Space to Accelerator Environments, pp. 1–28. Springer (2024)
47. Raisuddin, O.M., De, S.: Expectation value estimation. In: Quantum Computing for Engineers, chap. 27. Springer (2025)
48. Riesing, K., Schieler, C., Bilyeu, B., et al.: Operations and results from the 200 Gbps TBIRD laser communication mission. In: Proceedings of the 37th Annual Small Satellite Conference (2023). https://doi.org/10.26077/aggw-e480
49. Schieler, C., et al.: On-orbit demonstration of 200-Gbps laser communication downlink from the TBIRD CubeSat. In: Proceedings of SPIE 12413, Free-Space Laser Communications XXXV, 1241302 (2023)
50. Schuld, M., Brádler, K., Israel, R., Su, D., Gupt, B.: Measuring the similarity of graphs with a Gaussian boson sampler. Physical Review A 101(3), 032314 (2020)
51. Sim, S., Johnson, P.D., Aspuru-Guzik, A.: Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. Advanced Quantum Technologies 2, 1900070 (2019)
52. Starcloud: Data centers in space: carries the first NVIDIA H100 GPU into orbit. https://www.starcloud.com/starcloud-1 (2026)
53. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in Neural Information Processing Systems 30 (2017)
54. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. IEEE Transactions on Automatic Control 37(12), 1936–1948 (1992)

55. Westphal, C., Han, L., Li, R.: LEO satellite networking relaunched: Survey and current research challenges. ITU Journal on Future and Evolving Technologies 4(4) (2023)
56. Wiseman, H.M., Milburn, G.J.: Quantum Measurement and Control. Cambridge University Press (2010)
57. Xapsos, M.A.: The single event effects environment of space. NASA Technical Reports Server presentation (2021)
58. Zhu, Z., Rao, Z., Xiao, S., Yao, Y., Xu, Y., Meng, W.: Intelligent routing methods for low-Earth orbit satellite networks based on machine learning: A comprehensive survey. Ad Hoc Networks 173, 103995 (2025)
59. Schuld, M., Killoran, N.: Quantum machine learning in feature Hilbert spaces. Physical Review Letters 122, 040504 (2019). DOI 10.1103/PhysRevLett.122.040504.
60. Feilden, E., Oltean, A., Johnston, P.: *Why we should train AI in space*. White Paper v1.03, Lumen Orbit (now Starcloud). https://www.starcloud.com/wp (2024).
61. ASCEND: *Data Centres in Space*. https://ascend-horizon.eu/data-centres-in-space/

6 Appendix

Table 3. Notations and definitions

Symbol	Meaning	Type
$G_t = (V, E_t)$	Graph at decision time t	graph
V	Set of satellite nodes	set
E_t	Set of active optical links at time t	set
$\mathcal{N}(v)$	Neighborhood of node v	set
$x_v \in \mathbb{R}^{d_x}$	Classical feature vector of node v	vector
$e_{uv} \in \mathbb{R}^{d_e}$	Edge feature vector on link (u,v)	vector
Θ_v	Parameters of the node-local photonic sampler	parameter set
$\mathcal{P}_{\text{phot}}(\Theta_v)$	Photonic sampling distribution induced by Θ_v	distribution
$s_v^{(k)}$	k-th photonic sample drawn at node v	sample
K	Number of photonic shots per node	integer
$\phi_{\text{phot}}(\cdot)$	Photonic feature-map function	function
z_v	Photonic feature vector extracted from samples at node v	vector
$\mathbf{1}_D$	Destination indicator / destination encoding	vector
$h_v^{(\ell)}$	Hidden state of node v at message-passing layer ℓ	vector

Symbol	Meaning	Type
$m_v^{(\ell)}$	Aggregated neighborhood message at node v and layer ℓ	vector
$M_\ell(\cdot)$	Message function at layer ℓ	function
$U_\ell(\cdot)$	Node-update function at layer ℓ	function
L	Number of message-passing layers	integer
$g(\cdot)$	Policy head mapping final node state to action logits	function
$\pi_v(a \mid D)$	Routing policy at node v conditioned on destination D	distribution
a_v^*	Ground-truth next-hop action label for node v	label
V_{SL}	Supervised node set	set
V_{UL}	Unlabeled node set	set
$V_{\mathrm{SL}}^{\mathrm{valid}}$	Supervised nodes with a defined next-hop label	set
$\mathcal{L}$	Total training loss	scalar
$\mathcal{L}_{\mathrm{sup}}$	Supervised cross-entropy loss	scalar
$\mathcal{L}_{\mathrm{cons}}$	Consistency regularization loss	scalar
λ	Weight on the consistency term	hyperparameter
$\mathrm{CE}(\cdot,\cdot)$	Cross-entropy loss	loss
$\mathrm{KL}(p \parallel q)$	Kullback–Leibler divergence between distributions p and q	divergence
$\pi_v^{(1)}, \pi_v^{(2)}$	Policy outputs from two independent photonic draws	distributions

6.1 Theory: Training dynamics of MPNN via VQC

Let $G_t = (V, E_t)$denote the ring graph at decision time t, with $|V| = 10$. For node $v \in V$, let $x_v \in \mathbb{R}^{d_x}$ denote its classical state (e.g., link-health, queue, radiation, or device-state variables), and let $e_{uv} \in \mathbb{R}^{d_e}$ denote edge features on $(u, v) \in E_t$. We write $\mathcal{N}(v)$for the local neighborhood of v. A node-local neighborhood representation is encoded into photonic parameters $\Theta_v = \mathrm{Enc}(\mathcal{N}(v), x_{\mathcal{N}(v)}, e_{\mathcal{N}(v)})$, from which we draw samples

$$s_v^{(k)} \sim \mathcal{P}_{\mathrm{phot}}(\Theta_v), k = 1, \dots, K.$$

These samples are mapped to a robust photonic feature vector

$$z_v = \phi_{phot}\left(\{s_v^{(k)}\}_{k=1}^K\right),$$

which is then fused with the classical state to form the initial node representation

$$h_v^{(0)} = [x_v \parallel z_v \parallel \mathbf{1}_D],$$

where $\mathbf{1}_D$ denotes a destination indicator. The MPNN updates node states through neighborhood aggregation,

$$m_v^{(\ell)} = \sum_{u \in \mathcal{N}(v)} M_\ell\left(h_u^{(\ell)}, e_{uv}\right), h_v^{(\ell+1)} = \varphi_\ell\left(h_v^{(\ell)}, m_v^{(\ell)}\right),$$

where,

$$h_v^{(\ell+1)} = U_\ell\left(h_v^{(\ell)}, m_v^{(\ell)}\right)$$

and a policy head outputs a destination-conditioned next-hop distribution,

$$\pi_v(a|D) = \mathrm{softmax}(g(h_v^{(L)}, D)).$$

We deliberately keep $\mathcal{P}_{\mathrm{phot}}$ *general*, so the node-local quantum layer can be instantiated by bosonic or related photonic samplers rather than by a single Gaussian-specific construction. This choice is motivated by prior work showing that bosonic sampling distributions, and in particular Gaussian boson sampling, naturally encode graph structure and can be used to define graph-feature or graph-similarity representations [15, 50].

For *learning*, we partition nodes into a supervised set V_{SL} and an unlabeled set V_{UL}. Supervised nodes carry next-hop labels a_v^*, while unlabeled nodes contribute only through consistency regularization. The routing objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{sup}} + \lambda\, \mathcal{L}_{\mathrm{cons}},$$

with

$$\mathcal{L}_{\mathrm{sup}} = \frac{1}{|\, V_{\mathrm{SL}}^{\mathrm{valid}} \,|} \sum_{v \in V_{\mathrm{SL}}} \mathrm{CE}\,(\pi_v, a_v^*),$$

and

$$\mathcal{L}_{\mathrm{cons}} = \frac{1}{|\, V \,|} \sum_{v \in V} \frac{1}{2}\left[\mathrm{KL}\left(\pi_v^{(1)} \,\|\, \pi_v^{(2)}\right) + \mathrm{KL}\left(\pi_v^{(2)} \,\|\, \pi_v^{(1)}\right)\right].$$

Here, $h_v^{(\ell)}$ denotes the hidden state of node v at message-passing layer ℓ, $m_v^{(\ell)}$ denotes the aggregated neighborhood message, ϕ_ℓ and ψ_ℓ are the node-update and message functions, λ is the consistency-loss weight, CE is the supervised cross-entropy loss, and KL is the Kullback–Leibler divergence [25] used to enforce agreement between policies

induced by independent photonic draws. $\mathcal{L}_{\text{sup}}$ is the cross-entropy loss over the subset $V_{\text{SL}}^{\text{valid}} \subseteq V_{\text{SL}}$ for which a next-hop label is defined, and $\mathcal{L}_{\text{cons}}$ is a symmetrized KL penalty between two policy outputs obtained from *independent photonic draws* of the same node-local configuration. Intuitively, the supervised term teaches the model to reproduce labeled routing actions, while the consistency term encourages the policy to remain stable under stochasticity in the photonic layer. The resulting architecture is therefore *hybrid*: photonic sampling provides stochastic quantum-enhanced node features, whereas message passing and optimization remain classical.

Here $\pi_v^{(1)}$ and $\pi_v^{(2)}$ are policy outputs obtained from *two independent photonic draws* for the same node-local configuration. The supervised term trains the model to reproduce labeled next-hop actions, while the consistency term encourages the policy to remain stable under stochastic photonic feature generation, device noise, and other local perturbations. In practice, the destination node is excluded from $\mathcal{L}_{\text{sup}}$ because no next-hop action is defined once the packet has arrived. The resulting architecture is therefore hybrid: the node-local photonic layer supplies stochastic quantum-enhanced features, while message passing and optimization remain classical.

6.2 Node-State Fidelity Metric

Proposition 1 (properties of the node-state fidelity metric). Let $\{|\psi_v\rangle\}_{v=1}^{N}$be normalized pure states extracted from the trained node-local quantum encoder, and let F be the pairwise fidelity matrix with entries

$$F_{ij} = |\langle\psi_i|\psi_j\rangle|^2.$$

Then the following properties hold:

1. Normalization and symmetry: $F_{ii} = 1$for all i, and $F_{ij} = F_{ji}$.
2. Boundedness: $0 \le F_{ij} \le 1$for all i, j.
3. Diversity range: the off-diagonal diversity score

$$\text{diversity}(F) = 1 - \frac{1}{N(N-1)}\sum_{i \neq j} F_{ij}$$

satisfies $0 \le \text{diversity}(F) \le 1$.

4. Extremal cases:
   - if all node states are identical, then $F_{ij} = 1$for all $i, j$and $\text{diversity}(F) = 0$;
   - if all node states are pairwise orthogonal, then $F_{ij} = 0$for all $i \neq j$and $\text{diversity}(F) = 1$.

The diversity score therefore measures the complement of the mean pairwise state overlap: larger values indicate a more separated quantum embedding, while smaller values indicate a more collapsed representation.

The identity $F_{ii} = 1$ follows from state normalization, and symmetry follows from $|\langle\psi_i|\psi_j\rangle|^2 = |\langle\psi_j|\psi_i\rangle|^2$. The bound $0 \leq F_{ij} \leq 1$ follows from the Cauchy–Schwarz inequality for normalized states. Since the diversity score is one minus the mean of bounded off-diagonal fidelities, it is itself bounded in $[0, 1]$. The extremal cases follow directly from identical-state and orthogonality conditions. This is a direct consequence of the standard pure-state fidelity formalism [22, 32].

For each node i, we extract the final pure state $|\psi_i\rangle$ produced by the trained quantum encoder after classical preprocessing maps node features to circuit input angles. We then compute the node-state fidelity matrix

$$F_{ij} = |\langle\psi_i|\psi_j\rangle|^2,$$

which is the standard pure-state specialization of quantum fidelity. This construction is appropriate because the diagnostics are taken from state-vector simulations of the trained circuit rather than from mixed-state reconstructions. Fidelity is a standard similarity measure in quantum information, and for pure states it directly quantifies overlap between two state vectors [22, 32]. In the context of hybrid variational models, such pairwise fidelity matrices provide a useful qualitative diagnostic of whether the circuit is producing diverse embeddings or collapsing many inputs into the same effective state.

We emphasize that this diagnostic should be interpreted as a *representation analysis*, not as a standalone routing metric. A more diverse fidelity matrix does not by itself imply higher route success or better latency, but it does indicate that the VQC is using its available state space in a more discriminative manner. This connects naturally to the literature on variational quantum algorithms and ansatz design, where representational richness, expressibility, and trainability must be balanced rather than optimized in isolation. For that reason, include this metric as evidence about the geometry of the learned quantum representation rather than as a replacement for end-to-end network evaluation [7, 19, 51].

6.3 Node-by-qubit Expectation Metric

Proposition 2 (properties of the node-by-qubit expectation metric). Let $\{|\psi_v\rangle\}_{v=1}^{N}$ be normalized pure states extracted from the trained node-local quantum encoder, and let $E \in \mathbb{R}^{N\times Q}$ be the expectation matrix with entries

$$E_{vq} = \langle\psi_v|Z_q|\psi_v\rangle.$$

Then the following properties hold:

1. *Boundedness:* $-1 \leq E_{vq} \leq 1$ for all nodes $v$and qubits q.
2. *Identical-state consistency*: if $|\psi_i\rangle = |\psi_j\rangle$, then the corresponding rows of $E$are identical, i.e.,
$$E_{iq} = E_{jq} \text{ for all } q.$$

3. *Observable collapse condition*: if all rows of $E$are equal, then the selected Pauli-Z measurements do not distinguish the nodes at the observable level, even though the underlying quantum states may still differ.
4. *Heterogeneity range***:** if $\text{heterogeneity}(E)$ is defined as above, then

$$0 \leq \text{heterogeneity}(E) \leq 1.$$

The lower bound is attained when all node rows are identical, and the upper bound is attained only in the extremal case where, for every qubit, node expectations split maximally between $+1$ and -1.

The intuition is that expectation matrix captures the *observable signature* of the quantum encoder. Larger variation across rows indicates that different nodes induce measurably different qubit responses, while nearly identical rows indicate observable-level concentration or collapse with respect to the chosen measurement basis. Since downstream classical layers consume expectation values rather than full wavefunctions, this diagnostic is especially relevant for hybrid routing models.

*Intuition***:** The bound $-1 \leq E_{vq} \leq 1$ follows because Z_q has eigenvalues ± 1, so the expectation value of Z_q in any normalized state lies in that interval. Identical-state consistency follows immediately from substituting the same state into the same observable. The observable-collapse condition is definitional: if all node rows are identical, the chosen observables do not distinguish node embeddings in measurement space. Finally, since each E_{vq} lies in $[-1,1]$, the variance across nodes for a fixed qubit is bounded by 1, and averaging over qubits preserves this bound. These are standard consequences of Pauli observables and expectation-value measurement in quantum information [32].

6.4 Theory and supporting Results: Photonic Boson Sampling

Boson-sampling output law. Let $U \in \mathbb{C}^{m \times m}$ be the unitary matrix implemented by an m-mode linear-optical interferometer, and let $s = (s_1, \ldots, s_m)$ and $t = (t_1, \ldots, t_m)$ be input and output occupation vectors with $\sum_i s_i = \sum_i t_i = n$. Then the transition amplitude from input Fock state $|s\rangle$to output Fock state $|t\rangle$is proportional to the matrix permanent of the submatrix of Uselected by the input and output occupations:

$$\langle t|U|s\rangle = \frac{\text{Per}(U_{t,s})}{\sqrt{\prod_i s_i ! \prod_i t_i !}},$$

and the corresponding output probability is

$$\Pr(t|s,U) = \frac{|\text{Per}(U_{t,s})|^2}{\prod_i s_i ! \prod_i t_i !}.$$

In the collision-free boson-sampling regime with one photon injected into each selected input mode and one photon measured in each selected output mode, this simplifies to the familiar rule that the output probability is the squared modulus of a permanent of

the corresponding $n \times n$ submatrix. This is the core mathematical reason boson sampling is relevant here: the sampled output statistics encode multi-path interference through a structured, nontrivial function of the interferometer parameters, making the photonic module a principled stochastic feature map rather than an arbitrary noise source [3].

Previous works [3] further argue that approximate classical simulation of boson-sampling output distributions would have implausible consequences for complexity theory, which is why boson sampling is often used as a canonical example of a photonic computation that is simple to state physically but hard to reproduce classically.

Theory: Boson-output Diagnostics. We convert the sampled counts into an empirical distribution

$$\hat{p}(y) = \frac{c_y}{K},$$

where c_y is the number of times pattern y is observed in K total shots. From $\hat{p}$, we derive three summary diagnostics: collision probability, no-collision probability, and normalized entropy. The collision probability is defined as

$$P_{\text{coll}} = \sum_{y \in \mathcal{Y}_{\text{coll}}} \hat{p}\,(y),$$

where $\mathcal{Y}_{\text{coll}}$ is the set of output patterns containing at least one mode with occupancy greater than one. The no-collision probability is

$$P_{\text{nocoll}} = \sum_{y \in \mathcal{Y}_{\text{nocoll}}} \hat{p}\,(y),$$

where all occupied modes have unit occupancy. Next, we compute the Shannon entropy [12],

$$H(\hat{p}) = -\sum_{y \in \mathcal{Y}} \hat{p}\,(y) \log \hat{p}(y),$$

and normalize it as

$$H_{\text{norm}}(\hat{p}) = \frac{H(\hat{p})}{\log|\text{supp}(\hat{p})|},$$

where $\text{supp}(\hat{p}) = \{y : \hat{p}(y) > 0\}$ is the observed support. Finally, to quantify sensitivity to photonic perturbations, we compare ideal and noisy boson-output distributions using the total variation distance

$$D_{\mathrm{TV}}(\hat{p},\hat{q}) = \frac{1}{2}\sum_{y\in\mathcal{Y}} |\,\hat{p}(y) - \hat{q}(y)|.$$

These diagnostics are used as *mechanistic measures* of the photonic feature map rather than as direct routing metrics: they indicate whether the photonic module is producing a structured sampling law and how stable that law is under perturbation, but they do not by themselves imply better route success [3, 34].

Proposition 3 (Basic properties of boson-output diagnostics). Let $\hat{p}$ be the empirical distribution over sampled output occupation patterns obtained from a photonic kernel with K shots, and let $\mathcal{Y}_{\mathrm{coll}}$ and $\mathcal{Y}_{\mathrm{nocoll}}$ denote the collision and no-collision subsets of the output space, respectively. Then the following properties hold:

1. *Normalization*: $\hat{p}(y) \geq 0$for all $y \in \mathcal{Y}$, and $\sum_{y\in\mathcal{Y}} \hat{p}\,(y) = 1$.
2. *Collision partition*: if $\mathcal{Y}_{\mathrm{coll}} \cup \mathcal{Y}_{\mathrm{nocoll}} = \mathcal{Y}$ and the two sets are disjoint, then

$$P_{\mathrm{coll}} + P_{\mathrm{nocoll}} = 1.$$

3. *Entropy bounds*: the Shannon entropy satisfies

$$0 \leq H(\hat{p}) \leq \log|\mathrm{supp}(\hat{p})|,$$

and therefore the normalized entropy satisfies

$$0 \leq H_{\mathrm{norm}}(\hat{p}) \leq 1.$$

The lower bound is attained when all observed mass is concentrated on a single output pattern, and the upper bound is attained when $\hat{p}$is uniform over its observed support.

4. *Total variation bounds*: for any two empirical distributions $\hat{p}$and $\hat{q}$over the same output space,

$$0 \leq D_{\mathrm{TV}}(\hat{p},\hat{q}) \leq 1.$$

The lower bound is attained if and only if $\hat{p} = \hat{q}$, while the upper bound is attained when the supports are disjoint. Normalization follows immediately from the definition of empirical relative frequencies. The collision-partition identity follows when the collision and no-collision subsets form a disjoint partition of the sampled output space. The entropy bounds are standard Shannon-entropy inequalities: entropy is minimized by a point mass and maximized by a uniform distribution over the support. The total variation bounds follow directly from the ℓ_1-distance form of D_{TV}. These are general properties of discrete probability distributions, while the relevance of the output law to photonic interference follows from established boson-sampling framework [3] and, their Gaussian photonic variants [15].

Hence, the boson-output diagnostics summarize complementary aspects of the photonic feature map. P_{coll} and P_{nocoll} quantify whether sampled mass concentrates on

bunching-like versus collision-free events; H_{norm} quantifies how concentrated or diffuse the output law is; and D_{TV} quantifies the sensitivity of the photonic distribution to perturbation or noise. Together, these quantities distinguish a structured photonic sampling regime from either near-deterministic collapse or unstructured noise. This is especially relevant in a hybrid routing setting, where the photonic module is used as a *stochastic feature generator* whose usefulness depends on both structure and stability [3, 15].

Fidelity-output in ORCA. Fig. 15 shows pairwise empirical fidelity between node-output distributions generated by the CUDA-Q ORCA photonic sampler.

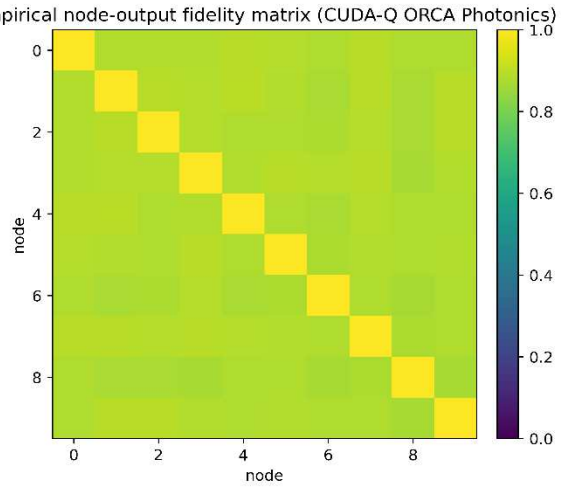


Fig. 15. Empirical node-output fidelity matrix for the CUDA-Q ORCA photonic sampler. Entry (i, j) reports the empirical fidelity $F_{ij} = \sum_s \sqrt{\hat{p}_i(s)\hat{p}_j(s)}$ between the sampled output distributions associated with nodes $i$and j. The unit diagonal is expected, while the uniformly high off-diagonal values indicate that node-output distributions remain broadly similar across the clustered ring. This suggests that the photonic feature layer operates in a relatively low-contrast regime, reflecting the structural regularity of the underlying topology.

Most off-diagonal entries in **Fig. 15** are high, indicating substantial overlap between the sampled photonic output distributions of different nodes, indicating a stable quantum feature regime. In the clustered-ring setting, this behavior is consistent with the similarity of local neighborhoods across nodes and suggests that downstream routing gains, if present, arise from relatively subtle distributional differences rather than from strongly separated node signatures.

This is the empirical node-output fidelity matrix for the CUDA-Q ORCA photonic layer. By construction, the diagonal entries equal one, but the more informative pattern is the uniformly high off-diagonal overlap across most node pairs. This indicates that the sampled node-output distributions are broadly similar across the 10-node clustered

ring, rather than sharply separated. In our setting, this behavior is plausible because many satellites occupy structurally similar local neighborhoods, so the photonic sampler is operating in a low-contrast regime: it produces stable distributions, but only modest node-to-node separation. Consequently, any downstream gain in routing quality must arise from relatively subtle, destination-relevant differences in these sampled distributions rather than from large pairwise distinctions. This interpretation is consistent with prior work on bosonic-sampling-based graph similarity, where useful structure is often encoded in distributional overlaps rather than in mutually orthogonal node signatures [50].

6.5 Qualitative Routing Behavior

Fig. 16 below and **Fig. 1** provide a qualitative comparison of rollout behavior for the `classical_identity`, `cudaq_vqc`, and `photonic_sampler` models on held-out 10-node optical ring LEO graphs. All three methods use the same graph-learning framing, with node and edge states processed by a message-passing backbone before a next-hop prediction is produced; the quantum variants differ only in the feature-generation subroutine, using either a variational quantum circuit or a photonic sampling-inspired map [3, 7, 14, 43]. In the displayed examples, the classical and photonic models recover the active source-to-destination route, although both still exhibit off-route node-label errors. By contrast, the VQC example predicts the first hop correctly but then reverses direction and forms a short loop, showing that local next-hop agreement is not sufficient for end-to-end route success. These examples therefore illustrate that the quantum-enhanced branches are not merely reproducing the classical baseline, but are changing the structure of routing errors in qualitatively distinct ways [3, 7, 43]. The `classical_identity` panel suggests a policy that is adequate for the active path but not globally consistent across the full ring. The `cudaq_vqc` panel shows a more brittle representation, where one incorrect intermediate decision invalidates the route despite an initially correct choice. The `photonic_sampler` panel is qualitatively closer to the successful classical rollout, but still displays residual off-path inconsistencies. Accordingly, the strongest interpretation is that the hybrid quantum modules affect the learned representation in a nontrivial way, but the present results remain better characterized as a feasibility demonstration of hybrid routing behavior than as a finalized claim of improved routing efficiency. This reading is consistent with the intended role of message passing for graph reasoning, variational quantum circuits as trainable hybrid embeddings, and photonic sampling as a structured stochastic feature source rather than a fully autonomous routing controller [3, 7, 14, 43].

The routing model can be described as a *message-passing graph learner* operating on the satellite ring, where node and edge states are aggregated before a next-hop readout is produced. This is naturally expressed in the Message Passing Neural Network (MPNN) framework, which provides a standard way to learn representations on graph-structured inputs [14]. The `cudaq_vqc` branch then augments this classical pipeline with a *variational quantum circuit* that acts as a trainable embedding subroutine inside a classical optimization loop, matching the standard hybrid variational quantum algorithm formulation [7]. The `photonic_sampler` branch is motivated by

linear-optical boson sampling, using a photonic stochastic map as a feature-generating component rather than as a standalone routing controller [3]. This interpretation is also consistent with photonic variational implementations in which a reconfigurable quantum processor is coupled to a conventional classical optimizer [43].

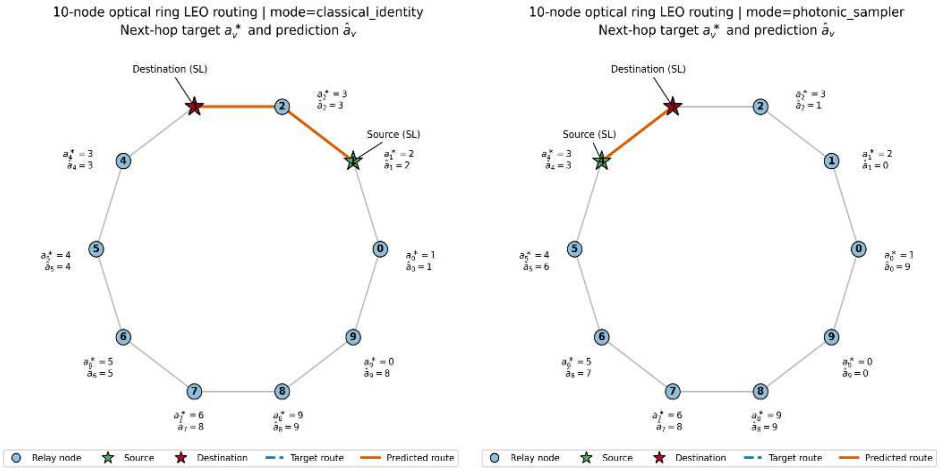


Fig. 16. Qualitative next-hop rollout visualizations for routing models on held-out 10-node optical ring LEO graphs: `classical_identity` and `photonic_sampler` with `cudaq_vqc` shown in **Fig. 1**. Blue circles denote relay nodes, green and red stars denote the source and destination, the dashed blue line denotes the target route, and the orange line denotes the predicted rollout. Node annotations show the target next-hop label a_v^* and the predicted next hop $\hat{a}_v$. The figures illustrate that exact node-wise agreement and end-to-end route success can diverge: the classical and photonic examples complete the displayed active route despite several off-route errors, whereas the VQC example predicts the correct first hop but then forms a local loop and fails to reach the destination.

The qualitative routing behavior differs across the three examples. In the `classical_identity` panel, the model correctly recovers the active source-to-destination path, but several noncritical nodes remain mislabeled, indicating that the model has learned a route-serving policy that is locally adequate for the active flow without being globally consistent over the full ring. In the `cudaq_vqc` panel, the model predicts the first hop correctly but then sends the packet back toward the source instead of onward to the destination, producing a short 2-cycle. This suggests that the learned quantum-enhanced embedding is capturing some local structure but has not yet stabilized the global routing rule. In the `photonic_sampler` panel, the active route is again correct on the displayed sample, and the surrounding node labels appear qualitatively more coherent than in the VQC example, although several inactive nodes are still incorrect. Overall, these figures support a cautious interpretation: the quantum-enhanced variants alter the error pattern of the learned policy, but the present evidence is stronger as a demonstration of hybrid routing behavior than as a claim of uniformly improved routing quality.

If these visualizations are discussed alongside the training curves and benchmark tables, the safest conclusion is that the current hybrid models are already rich enough to change the policy structure, but not yet regularized or constrained enough to guarantee route-consistent behavior on every held-out graph. In particular, the VQC example shows why end-to-end rollout success should remain a primary routing metric: even a visually plausible local policy can fail if one intermediate node violates the ring's shortest-path structure.

6.6 Radiation Model and Output

To capture radiation-driven non-stationarity in the node-local photonic layer, we associate each node v with a scalar sampler-efficiency state $\eta_v(t) \in (0,1]$. The model is intentionally stylized, but it follows the basic distinction emphasized in NASA tutorial [42]: total ionizing dose (TID) is the absorbed dose from ionizing radiation and leads to *cumulative parametric degradation*, whereas a single-event effect (SEE) is a disturbance caused by the passage of a single proton or heavy ion through or near a sensitive circuit node. Pellish also notes that, in proton-dominated environments such as LEO, proton SEE can constitute a substantial fraction of the overall SEE rate. We therefore model device health as the combination of a slowly varying cumulative term and a transient upset term. Let $D_v(t)$ denote the accumulated ionizing dose at node v. We evolve dose according to

$$D_v(t + \Delta t) = D_v(t) + \dot{D}\,\Delta t,$$

where $\dot{D}$ is a mission-level dose-rate parameter. The cumulative component of device health is then modeled by an exponential dose law

$$\eta_v^{\mathrm{TID}}(t) = \eta_{0,v}\exp(-k_{c(v)}D_v(t)),$$

where $\eta_{0,v}$ is the nominal pre-radiation efficiency, $c(v) \in$ {"rad-hard" ,"rad-tolerant"} denotes the device class, and $k_{c(v)}$ is a class-dependent degradation coefficient. This form is not intended as a calibrated part-level hardness law; rather, it is a compact simulation mechanism for representing the cumulative degradation behavior associated with TID. The distinction between cumulative degradation and upset-like disturbances is the key qualitative feature inherited from the NASA radiation-effects framework [42].

To represent upset-like behavior, we use4 a transient multiplicative SEU (single-event upset) into our model [42], also known as single-event error [SEE] factor,

$$u_v(t) = \begin{cases} \beta_v, & t_v^{\mathrm{SEU}} \le t < t_v^{\mathrm{SEU}} + \tau_v^{\mathrm{SEU}} \\ 1, & \text{otherwise,} \end{cases}, 0 < \beta_v < 1,$$

and define the effective sampler efficiency as

$$\eta_v(t) = \eta_v^{\mathrm{TID}}(t)\, u_v(t).$$

Here t_v^{SEU} denotes the onset time of an SEU-like event, τ_v^{SEU} its recovery interval, and β_v the temporary efficiency penalty. We report the class-averaged device-health metric

$$\bar{\eta}_c(t) = \frac{1}{|\,V_c\,|} \sum_{v \in V_c} \eta_v\,(t),$$

where V_c is the set of nodes assigned to class c. In this way, $\bar{\eta}_c(t)$ acts as a compact proxy for radiation-driven variation in the photonic feature generator: slow changes in $\bar{\eta}_c(t)$ reflect cumulative dose, while short excursions reflect SEU-like disturbances, matching the qualitative SEU or TID (total ionizing dose) separation [42].

Radiation model output. The radiation model plot of **Fig. 17** shows that, under the present toy LEO radiation model, the mean sampler efficiency remains close to unity for both device classes over the short simulation horizon, with the main visible deviation appearing as a brief transient drop rather than as steady cumulative decay. This is qualitatively consistent with the distinction between cumulative total ionizing dose (TID) effects and single-event effects (SEE)**:** TID produces gradual parametric degradation, whereas SEE-like behavior appears as abrupt disturbances to normal circuit operation. The figure should therefore be *interpreted as evidence that the simulator* can represent *time-dependent radiation-driven nonstationarity* in node health, not as a calibrated ranking of radiation-hardness classes. In particular, because proton-induced SEE can be a major contributor in LEO environments, short-lived excursions in η are plausible even when the mean health remains otherwise near nominal.
The trace is dominated by *near-unity mean sampler efficiency* for both device classes, with only a brief transient dip in the curve labeled “rad-hard” around 7 minutes and a rapid rise of the “rad-tolerant” curve from about 0.95 to nearly 1.0 early in the run. Read literally, this means the toy model is producing a regime in which both classes operate close to ideal most of the time, and the main visible distinction is a short-lived upset rather than steady cumulative degradation. That is consistent with a simulation in which SEEs are modeled as transient events superposed on a weak cumulative total ionizing dose (TID) background, since TID is cumulative while SEE-like behavior appears abruptly and locally in time. In LEO, both TID-like cumulative exposure and SEE-like transient upsets are standard design concerns because spacecraft electronics are exposed to trapped-belt particles, solar particle events, and galactic cosmic rays [30, 42].

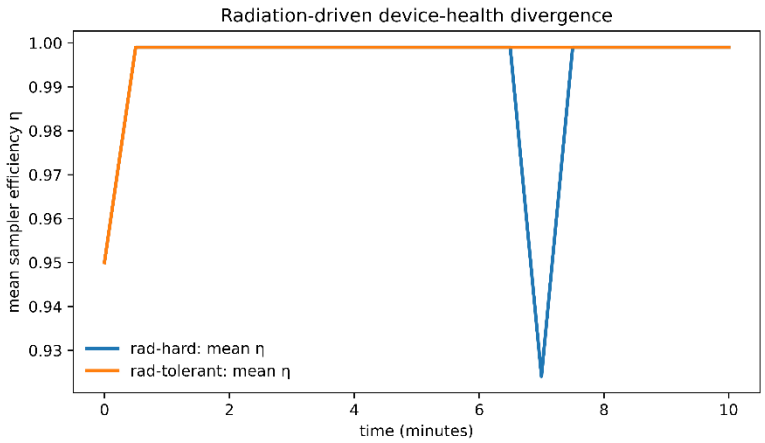


Fig. 17. Radiation-driven device-health divergence under the toy LEO sampler model. The curves report the mean sampler efficiency η for the two device classes over time. Both traces remain near unity for most of the run, indicating that cumulative degradation is weak on this short horizon, while the transient dip reflects an upset-like perturbation rather than sustained failure. Because TID and SEE are both relevant for electronics in LEO, but the present simulation is intentionally stylized, this figure should be interpreted as illustrating the model's ability to represent radiation-driven nonstationarity in node health rather than as a calibrated comparison of radiation hardness classes.

An important point is that the ordering of the two curves in **Fig. 17** is not what one would normally expect from the labels alone: in the plot, the "rad-tolerant" curve remains at or near the top of the axis while the "rad-hard" curve exhibits the only visible drop. For a long-duration radiation-hardness comparison, one would usually expect the more strongly hardened class to be at least as stable as the weaker tolerant class. In practice, the plot of the toy model demonstrates the model's ability to represent *transient divergence in device health under radiation-driven perturbations*, not that rad-tolerant devices outperform rad-hard devices in LEO. In our case, the divergence is heavily influenced by the simulation assumptions—for example, the chosen SEU trigger pattern, transient penalty, or hormetic uplift—not by a calibrated radiation-qualification model [30, 42].

This plot supports a modest but useful claim: the radiation layer introduces *time-dependent nonstationarity* into the photonic feature generator, and the downstream routing policy therefore has to remain stable not only under stochastic photonic draws but also under slowly evolving device health.

6.7 Routing Accuracy vs Latency

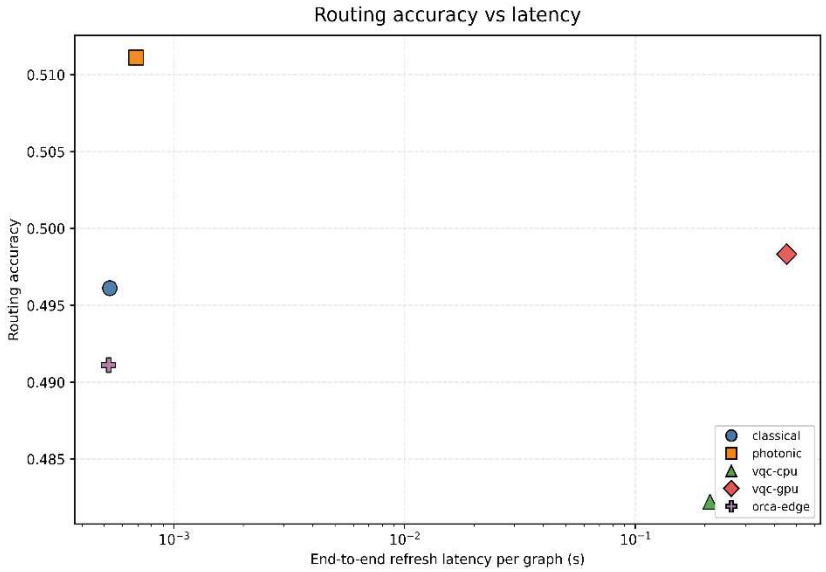


Fig. 18. Routing accuracy versus end-to-end refresh latency across routing backends. Each point shows a backend positioned by held-out routing accuracy on the y-axis and by end-to-end refresh latency per graph on the x-axis. The photonic backend attains the highest accuracy in the lowest-latency regime, the classical baseline remains competitive at similarly low latency, the ORCA-edge variant is slightly less accurate at comparable latency, and the VQC backends incur substantially larger latency costs, with the GPU VQC recovering only modest accuracy while the CPU VQC is dominated on both axes. This type of comparison is operationally relevant in LEO ISL routing because low-latency refresh remains a central requirement, while graph-structured learning provides a natural framework for backend comparisons inside a common routing pipeline, and photonic sampling modules are naturally motivated by linear-optical computation [3, 14, 55].

Fig. 18 shows a clear latency–accuracy split between the low-latency classical/photonic family and the slower VQC family. In this run, the *photonic backend* is the strongest point on the plot because it delivers the highest routing accuracy while remaining in the same sub-millisecond regime as the classical and ORCA-edge variants. The *classical baseline* is slightly less accurate but still competitive, whereas the ORCA-edge backend sacrifices a small amount of accuracy at roughly the same latency. By contrast, the GPU VQC regains only a modest amount of accuracy while moving two to three orders of magnitude to the right on the latency axis, and the CPU VQC is the weakest point overall. The strongest interpretation is therefore not that quantum backends improve routing accuracy in the present configuration, but that backend choice produces a measurable trade-off between representational style and deployable refresh-time performance in a graph-based LEO routing pipeline [14, 55].

A second useful reading is that the absolute spread in accuracy is fairly small, so this figure is best treated as a backend-comparison summary rather than as a headline performance result. The plot is still valuable because it shows that the photonic branch changes the learned representation in a way that is beneficial on this metric without paying the latency cost of the VQC backends, while the VQC backends currently do not justify their runtime overhead on routing accuracy alone. Since the photonic branch is motivated by linear-optical sampling rather than by a conventional deterministic feature map, this is still a meaningful hybrid result even if it is not yet the strongest route-level result in the paper [3, 55].

6.8 Routing: Disagreement Map

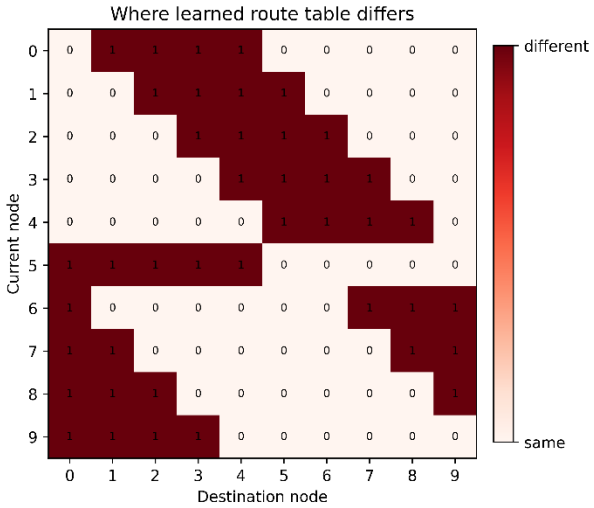


Fig. 19. Disagreement map between the heuristic baseline and the learned MPNN routing policy. Each entry indicates whether the learned message passing neural network (MPNN) selects the same next hop as the heuristic baseline for a given current-node/destination pair. Light cells denote agreement, and dark cells denote disagreement. The differences are structured rather than isolated, indicating that the learned model modifies routing decisions in coherent destination-dependent regions of the policy table rather than through random local deviations.

Fig. 19 shows *where the learned MPNN policy disagrees with the heuristic baseline*, so it is not describing the method itself; it is visualizing an outcome of training: 0 / light cells = learned MPNN chooses the same next hop as the heuristic; 1 / dark red cells = learned MPNN chooses a different next hop; rows = current node; columns = destination node. The main value of this plot is that the differences are structured, not random. The learned policy is not just making scattered changes; it is reorganizing routing decisions in a destination-dependent way across coherent regions of the table.

Routing data: Analysis. These figures are best read as *qualitative rollout diagnostics* rather than aggregate performance summaries. The `classical_identity` panel shows a successful active route from the displayed source to destination: the predicted next hops along the active path match the target route, even though several off-route nodes still have incorrect next-hop labels. The cudaq_`vqc` panel is more brittle: the first hop is correct, but the next decision reverses direction and creates a short loop, so the rollout fails despite one locally correct decision. The `photonic_sampler` panel again shows a successful active route on the displayed sample, but—as in the classical case—several inactive nodes remain misclassified. The key qualitative message is that *node-wise correctness and route-wise correctness are not the same*: a model can complete the displayed path while still making many global labeling errors, and a single error on the active path can invalidate the route.

Another point is that the panels should be compared cautiously across modes. The displayed source and destination locations are not identical in all three visualizations, which strongly suggests that these panels were generated from different held-out validation graphs rather than from one common rollout state. For that reason, the figures are most useful for showing *error structure*—for example, correct first-hop prediction, loop formation, and off-route label noise—rather than for making strict apples-to-apples quantitative comparisons between methods.

6.9 Discussion: VQC Circuit Architecture

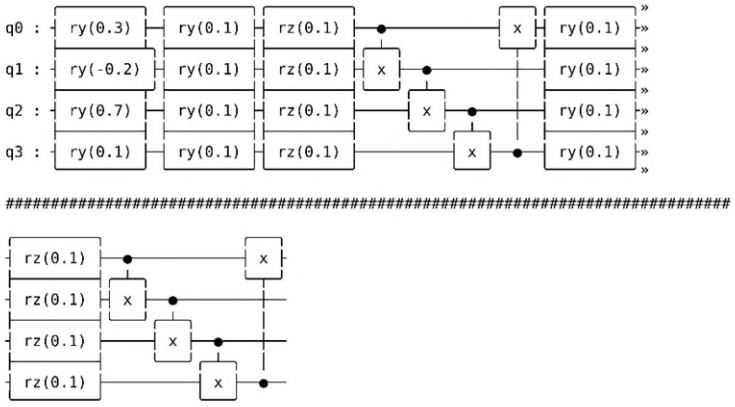


Fig. 20. CUDA-Q produced node-local variational quantum circuit (VQC) used in the cudaq_vqc routing backend. The diagram shows a 4-qubit CUDA-Q kernel with angle encoding on each qubit, followed by trainable single-qubit rotation layers and a ring-style pattern of controlled-X entangling gates. The rendered layout is split across lines by the circuit visualizer, so the lower portion should be read as a continuation of the same kernel rather than as a second circuit. In CUDA-Q, such circuits are defined as quantum kernels and can be visualized with cudaq.draw, while expectation values and state vectors are obtained through observe and get_state during hybrid execution. This circuit is therefore best interpreted as the quantum embedding subroutine inside the routing model, not as a standalone algorithmic output.

The main role for the quantum circuit of **Fig. 3** in Section 3.6 is to make the quantum encoder *concrete*: the routed node state is first classically encoded, then transformed by a shallow trainable circuit, and finally exposed back to the classical model through observables or state-derived diagnostics. Because CUDA-Q also supports direct visualization of kernel execution paths, the rendered diagram provides a transparent description of what the `cudaq_vqc` branch computes at the circuit level. The circuit makes explicit that the `cudaq_vqc` branch is a shallow hybrid variational model: node features are first mapped to input rotation angles, then processed by trainable single-qubit rotations and a ring entangling pattern before the downstream routing stack consumes observable outputs. This presentation is useful because it shows that the quantum branch is *not an opaque black box*; it is a concrete, parameterized kernel of the kind typically used in variational quantum algorithms, where a classical outer loop optimizes circuit parameters through repeated evaluation of observables, aligning with the standard hybrid VQA perspective and with CUDA-Q's kernel-based execution model [7, 35].

Fig. 8: Route Latency–Utility Trade-offs - Explanation. The negative y-values mean the backend is worse than the classical baseline on that successful-route metric**.** For these plots, the y-axis is defined as a reduction vs classical, i.e., reduction vs classical $=$ classical metric $-$ backend metric.

Hence, the intuition is, *positive value* $\rightarrow$ backend has lower queue / congestion than classical $\rightarrow$ better ; *zero* $\rightarrow$ same as classical; *negative value* $\rightarrow$ backend has higher queue / congestion than classical$\rightarrow$ worse.

In the two panels, for Successful-route queue reduction vs classical: negative means that, among routes that successfully reached the destination, that backend's routes had higher average queue load than the classical baseline. For Successful-route congestion reduction vs classical: negative means that, among successful routes, that backend's routes had higher congestion proxy than the classical baseline. *Example*: If classical successful-route queue mean is 0.40 and a backend's is 0.42, then: $0.40 - 0.42 = -0.02$. Therefore, the plotted value is -0.02, meaning the backend is *0.02 worse* than classical on that metric. Important nuance: These are successful-route-only values, so they do not mean the backend is worse overall in every sense. They mean: conditional on the route succeeding, the chosen path had worse queue or congestion than the successful classical path average. Next, **Table 4** and **Table 5** in Section 6.10 shows the training metrics.

6.10 Training Performance

Table 4. Training-curve summary by backend

Backend	Training fit	Validation behavior	Approx. best / late validation accuracy
classical_identity	Train loss falls steadily (~0.69 -> 0.36) while train accuracy rises strongly (~0.50 -> 0.86).	Validation loss rises almost monotonically after early epochs (~0.69 -> 1.17). Validation accuracy peaks early and then drifts back toward chance-like performance.	Peak val. acc. ~0.56 at epochs 5–7; late ~0.50–0.51.
cudaq_vqc	Train loss decreases strongly (~0.70 -> 0.20) and train accuracy rises to ~0.95.	Validation loss is flat early and then climbs sharply after mid training (~0.67 -> 1.15). Validation accuracy is highest early-to-mid and then becomes volatile.	Peak val. acc. ~0.73 at epochs 6–13; late ~0.53–0.58.
photonic_sampler	Train loss decreases strongly (~0.69 -> 0.17) and train accuracy rises to ~0.94.	Validation loss still rises overall (~0.69 -> 1.10), but validation accuracy gradually improves and remains comparatively stable late in training.	Late-stage / peak val. acc. ~0.66–0.68 at epochs 23–30; final ~0.66.

Table 5. Stability, checkpointing, and appendix-ready interpretation

Backend	Overfitting / stability	Suggested checkpoint window	Interpretation
classical_identity	Early overfitting. The train/validation gap widens quickly after the first third of training.	Prefer early stopping near epochs 5–7.	Strong training fit but weak generalization; best treated as the classical baseline, not the preferred final checkpoint.
cudaq_vqc	Best early validation signal, but later training is unstable and strongly overfit.	Use a best-checkpoint rule in the early-to-mid window, roughly epochs 6–13.	Promising early validation behavior, but not stable enough late in training to claim robust generalization without stricter checkpointing.
photonic_sampler	Still overfits in loss, but the validation-accuracy trajectory is the most stable late in training.	Use a late checkpoint near the validation-accuracy plateau, roughly epochs 23–27, or the exact best epoch by validation accuracy.	Best late validation behavior among the three curves; strongest candidate when the paper emphasizes deployable validation performance.

The metrics show that across all three variants, training loss decreases steadily and training accuracy rises, indicating that each model can fit the training set. The validation curves, however, separate clearly. The `classical_identity` model overfits earliest: validation loss rises after the first part of training while validation accuracy stays nearly flat. The `cudaq_vqc` model shows the strongest early-to-mid validation behavior, but later training becomes unstable and the train-validation gap widens substantially. The `photonic_sampler` model reaches the strongest final validation accuracy, although its validation loss continues to increase, suggesting that the model is improving top-class decisions while becoming increasingly overconfident on errors.

From the quantum perspective, the current behavior is consistent with the intended role of hybrid variational and photonic modules as representation-enhancing, classically supervised subroutines rather than standalone routing controllers [3, 7, 43]. Tables 4 and 5, as summary of the training/validation/acc plots for **Fig. 21**, show that the plots support an interpretation of the current system as a feasible hybrid-routing demonstration rather than a fully optimized routing learner, with the photonic branch showing the most favorable late-stage validation behavior and the VQC branch showing the most promising early-stage representation gains [7, 23, 43, 44]. **Table 6** and **Table 7** show the values extracted from these plots.

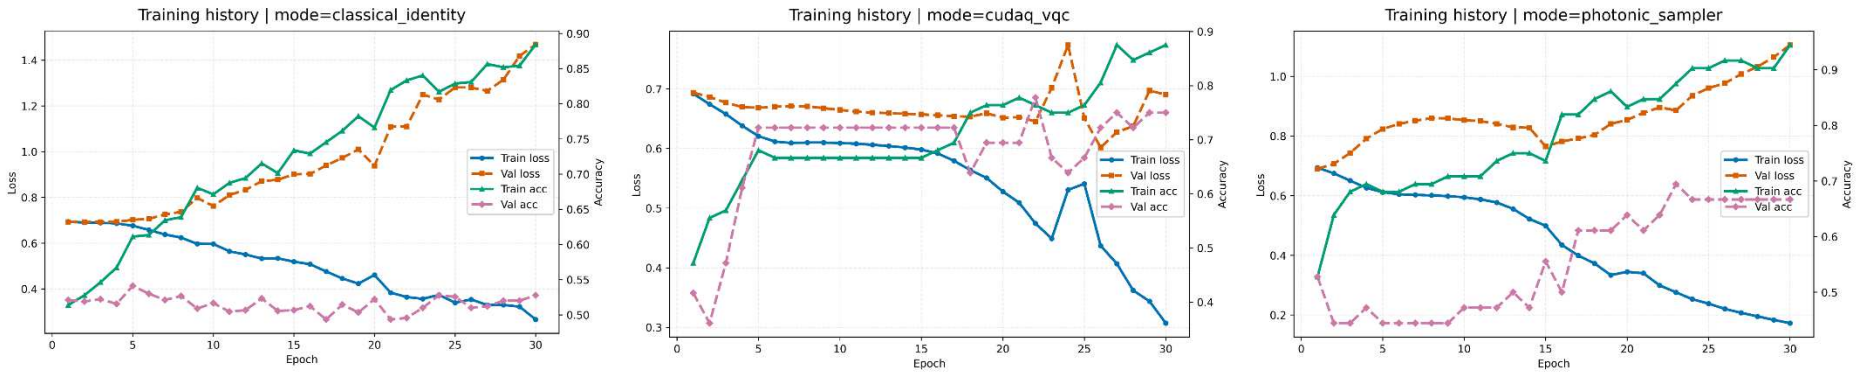


Fig. 21. Training histories for three routing models: `classical_identity`, `cudaq_vqc`, and `photonic_sampler`. In all cases, training loss decreases and training accuracy increases, indicating that each model can fit the training data. The validation curves, however, reveal substantial differences in generalization. The classical model overfits earliest, with rising validation loss and flat validation accuracy after an early peak. The VQC model reaches strong early validation performance but becomes unstable later in training, producing the largest train-validation gap. The photonic sampler achieves the strongest late validation accuracy, although its rising validation loss indicates increasing confidence on incorrect predictions. Overall, the figure suggests that the quantum-enhanced variants alter the generalization profile of the routing learner, with the photonic model offering the most favorable final validation behavior. **Table 6** and **Table 7** in Section 6.11 show the values tabulated from these plots.

Summary. All three models fit the training set well. The main difference is validation behavior: `classical_identity` overfits earliest, `cudaq_vqc` reaches the strongest early validation accuracy but becomes unstable later, and `photonic_sampler` gives the best late-stage validation accuracy and the smoothest late validation trajectory.

6.11 Measured Supporting Metrics for Training, Routing and Quantum Benchmarks

Table 6. Training and inference latency metrics.

Backend	Train loss	Train acc.	Train total (s)	Per epoch (s)	Avg / graph (s)	P95 / graph (s)	Refresh (s)	Shots	Latency/shot (ms)
Classical	0.414	0.823	7.455	0.086	0.001	0.001	0.001	128	—
Photonic sampler	0.246	0.913	8.542	0.095	0.001	0.001	0.001	128	0.005
CUDA-Q VQC (CPU)	0.327	0.844	9165.321	264.902	0.203	0.235	0.203	128	1.586
CUDA-Q VQC (GPU)	0.248	0.889	19776.105	571.788	0.441	0.539	0.441	128	3.443
ORCA edge augmentation	0.234	0.903	7.680	0.089	0.001	0.001	0.011	128	0.083

Table 7. Routing and path-level metrics.

Backend	Routing acc.	Route success	Mean hops	Path queue mean	Path util. mean	Congestion proxy	Queue red. vs network	Cong. red. vs network	Queue red. vs classical	Cong. red. vs classical
Classical	0.510	0.290	2.735	0.522	0.505	0.514	-0.017	-0.009	0.000	0.000
Photonic sampler	0.506	0.285	2.570	0.524	0.511	0.517	-0.019	-0.013	-0.002	-0.004
CUDA-Q VQC (CPU)	0.509	0.315	2.665	0.523	0.513	0.518	-0.018	-0.014	-0.001	-0.005
CUDA-Q VQC (GPU)	0.514	0.265	2.580	0.524	0.508	0.516	-0.019	-0.012	-0.002	-0.002
ORCA edge augmentation	0.506	0.270	2.615	0.523	0.510	0.517	-0.018	-0.012	-0.001	-0.003

Note. Reductions are reported relative to the classical baseline or the network average as named in the column.

6.12 Full Model Pipeline for the CUDA-Q Assisted Space Datacenter Routing

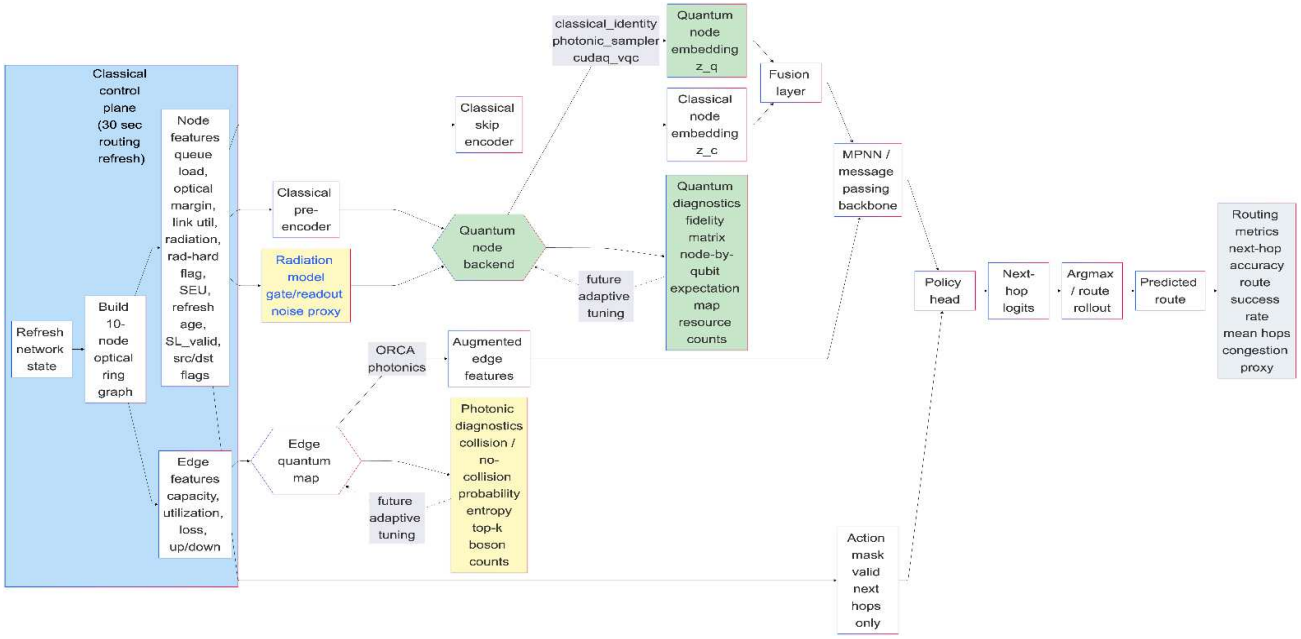


Fig. 22. Hybrid classical–quantum routing model for the space based datacenter [60, 61] in LEO orbit**.** Node features x_v are first processed by a classical node encoder and then passed to an optional quantum node module, instantiated as either a VQC or a photonic sampler, while a classical skip path preserves the original node representation. Edge features e_{uv} may be further augmented by an ORCA photonic edge map. The resulting node and edge representations are fused and propagated through an MPNN backbone, after which a topology-aware mask (SL_valid) is applied in the masked policy head to enforce valid next-hop decisions. The model outputs a next-hop prediction that is rolled out into a route and evaluated with routing metrics. The design reflects the paper's central principle of classical control for routing with quantum subroutines for representation enhancement.