***RESEARC PRORGAM**: Quantum Networks*

8. Quantum Sensing Protocol Optimization

For the general model of network of imperfect quantum sensors, where each sensor measures a single (or multiple) parameter, we analyze the impact of quantum state or measurement probing operator modeling imperfections on the results of measurement. Based on the results of previous work we suggest enhancements of the previous known methods and at the end suggest an original approach based on the joint optimization of the protocol stack in the sensing node consisting of the following operational composition (density-matrix / CPTP map language), which explicitly separates quantum state preparation, encoding, sensor noise, and measurement estimation error variance) but does not significantly alter optimal β, emphasizing the robustness of the chosen measurement imperfections namely: $\rho_0 \overset{\mathcal{E}_{prep}}{\rightarrow} \tilde{\rho} \overset{U(\theta)}{\rightarrow} \rho(\theta) \overset{\mathcal{E}_{sens}}{\rightarrow} \rho_{out}(\theta) \overset{\mathcal{M}_{imp}}{\rightarrow} p(x \mid \theta)$. *Here* $\mathcal{E}_{prep}$ *models state-preparation errors (imperfect entangling gates, mixedness).* $U(\theta)$ *is the ideal parameter encoding (local or global).* $\mathcal{E}_{sens}$ *captures decoherence during sensing (Markovian/non-Markovian, correlated spatial noise) and* $\mathcal{M}_{imp}$ *is an imperfect measurement channel (noisy POVMs, finite readout fidelity). The project presents general framework for optimization of the protocol stack of a node in quantum sensor networks. It enables the network designers to combine a number of possible choices of network parameters and find the optimum solutions for the remaining elements of protocol stack. Based on concrete simulation results for a subset of such cases we can summarize the results as follows: a) Single qubit: optimal measurement is a rotated Pauli measurement; POVMs (e.g., trine) do not surpass it. b) Two qubits: entanglement in both probe and measurement can be crucial for attaining maximum precision. c) Measurement families parameterized by β offer a powerful and experimentally realistic means to approach optimal metrology. d) Noise reduces Quantum Fisher Information (QFI) (increases parameter families. e) The results are consistent with quantum metrology theory and provide a concrete numerical demonstration accessible to experiments.*

8. A Quantum Sensing Protocol Optimization for Multiparameter Estimation in Sensor Network

We develop a systematic framework for multiparameter quantum sensing in a multinode network, by explicitly modeling the full operational stack: state preparation, parameter encoding, sensing noise, and measurement imperfections. Using CPTP maps and Kraus representations, we derive closed-form quantum Fisher information matrices (QFIMs) for GHZ, W, and product probe states, quantifying the effects of preparation mixing, dephasing, amplitude damping, and finite readout fidelity. Our analysis reveals that GHZ probes concentrate sensitivity into a single collective parameter, yielding a rank-deficient QFIM, whereas W and product probes produce full-rank QFIMs, enabling simultaneous estimation of multiple parameters. We provide explicit eigenvalue and covariance analyses, highlighting trade-offs between coherence, robustness, and multiparameter sensitivity, and discuss strategies to exploit ancillary nodes for noise mitigation or complementary parameter encoding. This operationally transparent approach extends prior studies that primarily focused on single-parameter or idealized probe models, offering design guidelines for realistic quantum sensor networks.

8.B Optimizing Quantum Multiparameter Sensing in Distributed Sensor Networks: A Unified Protocol-Stack and Performance-Bound Framework

We developed a complete theoretical and computational framework for end-to-end optimization of multiparameter quantum sensor networks, spanning state preparation, unitary encoding, sensing noise, and imperfect measurement. Starting from the covariance-based formulation of the Quantum Fisher Information Matrix (QFIM), we derived explicit symbolic expressions for two-node and three-node networks using local number-operator generators with general variances and covariances (v_i, c_{ij})*. Closed-form Quantum Cramér–Rao Bounds (QCRBs) were obtained, explicitly revealing how inter-node correlations modify identifiability and precision. We then embedded these results into a full protocol stack model* $\rho_0 \to \mathcal{E}_{prep} \to U(\theta) \to \mathcal{E}_{sens} \to \mathcal{M}_{imp} \to p(x \mid \theta)$*,and formulated optimization objectives based on classical Fisher information, QFIM, weighted mean-square error, and Bayesian risk. Gradient-based variational optimization strategies were presented for jointly tuning probe states and measurements under physical constraints and realistic noise. To address fundamental measurement incompatibility in the multiparameter regime, we derived a fully explicit Semidefinite Programming (SDP) formulation of the Holevo bound for two parameters, including the normally neglected imaginary covariance term* $\mathrm{Im}\, Z$*. Using operator-basis expansion and matrix lifting, the nonconvex Holevo optimization was converted into a convex SDP solvable with high-precision solvers such as MOSEK. A complete numerical implementation was supplied for two-qubit mixed states, yielding optimal measurement operators when the SDP relaxation is tight. Finally, we constructed a numerical two-qubit sensing example with loss and detector inefficiency and outlined a visualization framework comparing: the Quantum Cramér–Rao Bound (QCRB) from the QFIM, the classical Cramér–Rao Bound from a fixed local POVM, and the Holevo bound as the ultimate attainable precision. Together, these results provide a unified analytical–numerical toolbox for realistic multiparameter quantum sensing design, connecting abstract information-theoretic limits with experimentally implementable protocols.*

10.1 Imperfect Correlations and Information-Theoretic Limits of Quantum Measurement

We present an information-theoretic framework for modeling imperfections in quantum measurements, providing a principled, interpretation-neutral foundation for understanding measurement limitations. By treating measurement as a process of classical information transfer from a quantum system to an apparatus, we introduce the Imperfect Correlation Postulate, which asserts that perfect correlations are physically unattainable. Using the Holevo bound, we derive explicit upper limits on the mutual information between system observables and measurement outcomes, first for binary observables and then for arbitrary n*-outcome observables via a symmetric-overlap model. Analysis of the* n*-outcome case reveals that ideal projective measurements are singular limits in the space of physically realizable measurement processes, with non-commuting limits highlighting fundamental constraints on measurement sharpness and resolution. The framework naturally explains the emergence of classical outcomes, quantifies intrinsic limitations on information transfer, and provides a rigorous foundation for the use of POVMs in real experiments. Illustrative 2D and 3D plots visualize the decay of retrievable information with increasing pointer overlap and outcome number, highlighting the quantitative relationship between apparatus imperfection and measurement fidelity. Our results unify operational, mathematical, and foundational perspectives, offering new insight into the quantum measurement problem and the emergence of classicality.*

10.2 Measurement and Inference in Classical and Quantum Sensor Networks: A Comparative Framework

Classical sensor networks and quantum measurement systems are frequently modeled using similar estimation-theoretic concepts, such as likelihood functions, information measures, and distributed inference, despite being grounded in fundamentally different physical and mathematical frameworks. In this paper, we present a systematic, formal comparison between classical sensor-network measurement models and quantum measurement theory. We show that classical sensor networks can be described by commutative probability spaces in which measurements are represented as random variables conditioned on an underlying state, whereas quantum measurements are described by non-commutative operator algebras, with outcomes governed by positive operator-valued measures acting on density operators. By explicitly identifying the assumptions of passivity, state independence, and measurement commutativity implicit in classical models, we demonstrate why direct extensions of classical estimation and data-fusion techniques fail in general quantum settings. We further characterize the conditions under which classical tools admit consistent quantum analogues, including filtering, information metrics, and distributed inference under local operations and classical communication. Several representative case studies are used to illustrate the mathematical distinctions, including classical versus quantum Fisher information, Kalman versus quantum filtering, and consensus-based fusion versus LOCC-constrained protocols. The results establish a unified but assumption-aware framework for modeling measurement and inference in classical and quantum sensor networks, clarifying both their formal connections and their irreducible differences.

11.1 Fundamental Limits of Entanglement-Free Synchronization in Open Quantum Networks

We study synchronization in networks of n driven–dissipative quantum Van der Pol oscillators, focusing on the role of quantum correlations in collective phase locking. Using a Lindblad master equation with coherent coupling defined on an arbitrary graph, we introduce pairwise and global synchronization measures based on relative phase distributions. We show that global synchronization can emerge in fully separable steady states, generalizing initial two-oscillator results to networks. We derive rigorous bounds on the global synchronization order parameter in terms of pairwise phase coherence and network connectivity, establishing limits on achievable synchronization in open quantum systems. In contrast, strong coupling can enhance multipartite entanglement while degrading global synchronization. Our results identify discord-type correlations as the minimal quantum resource supporting network synchronization and provide a scalable framework applicable to quantum technologies and complex open systems.

11.2 Definition and Quantification of Quantum Synchronization: An Operational and Engineering Perspective

Quantum synchronization is a fundamental phenomenon with crucial implications for quantum networks, metrology, and communication, yet its precise definition, quantification, and robustness under realistic imperfections remain open challenges. In this work, we present a rigorous, theorem-driven framework for quantum synchronization across multiple operational levels, including expectation-value phase locking, trajectory-level coordination, state-level correlations, and task-oriented witnesses. We formally establish the inequivalence of different synchronization notions, and introduce operational witness operators capable of detecting synchronization even when semiclassical measures fail. Extending the framework to imperfect systems subject to detuning, decoherence, and parameter disorder, we analytically derive closed-form expressions for synchronization degradation and quantify tolerance thresholds that guide practical design and control. Numerical simulations and illustrative figures validate the theory, demonstrating that witness-based measures remain robust under realistic imperfections. Our results provide a systematic, quantitative foundation for understanding, engineering, and exploiting quantum synchronization in emerging quantum technologies.

11.3 Topology-Dependent Nonlinear Quantum Synchronization

We present a comprehensive theoretical framework for quantum synchronization in networks of weakly nonlinear oscillators, highlighting the interplay between network topology, nonlinearity, and dissipation. Starting from the full Hamiltonian and Lindblad dynamics, we systematically reduce the system to a slow-manifold of phase-bearing modes, making all approximations explicit and reproducible. Our analysis yields analytical expressions for the critical coupling $g_c(\lambda_2, \beta, \gamma, \Delta\omega)$ required for phase locking, explicitly connecting the spectral gap λ_2 of the network adjacency matrix, the oscillator nonlinearity β, the dissipation rate γ, and the intrinsic frequency detuning $\Delta\omega$ to synchronization thresholds. We construct nonlinear witness operators whose decay rates quantify collective phase coherence, providing a direct link between topology, nonlinearity, and observable measures. Illustrative examples for ring, random, and small-world networks demonstrate how topology influences synchronization speed, threshold, and robustness. The framework further accounts for finite temperature and imperfections, revealing that phase coherence persists across realistic experimental conditions. The methodology clarifies how weak nonlinearity and network connectivity govern synchronization, providing a predictive blueprint for designing and controlling quantum networks in superconducting, optomechanical, and trapped-ion platforms.

11.4 Topology-Dependent Quantum Network Stability

Synchronization in nonlinear quantum networks underpins collective quantum dynamics, with applications ranging from optomechanical arrays to superconducting circuits. Despite recent advances, a rigorous understanding of how network topology, nonlinearity, and dissipation jointly govern stability and robustness remains incomplete. Here, we develop a topology-dependent stability framework for coupled nonlinear quantum oscillators, combining operator-level slow manifold reduction (Dynamics Reduction DR) with Lyapunov-based contraction analysis. Our approach yields analytical stability thresholds that explicitly relate the network's spectral gap, Kerr nonlinearities, damping, and detuning. We demonstrate that network topology critically shapes stability: fully connected and small-world networks exhibit enhanced contraction rates, while ring networks are more susceptible to nonlinear destabilization. Robustness under bounded perturbations and thermal noise is established via input-to-state and mean-square stability bounds. Extensive numerical simulations confirm the analytical predictions, revealing topology-dependent transverse mode decay and synchronization witness operator dynamics. Our results provide quantitative design rules for robust synchronization in large quantum networks, bridging theoretical insight with experimentally relevant conditions. This framework extends classical spectral and Lyapunov techniques to the quantum domain and offers a transparent method to predict and control stability in complex quantum systems.

11.5 Structural Stability and Robustness in Open Quantum Networks

We present a general framework for structural stability and robustness in open quantum networks, unifying concepts from infinite-dimensional semigroup theory, operator-norm perturbation, normally hyperbolic invariant manifolds (NHIMs), and network topology. For a broad class of bosonic oscillator arrays, we identify a single spectral invariant—the residual gap $\kappa = (\gamma + g\lambda_2) - C_1\beta - C_2 \parallel \Delta\omega \parallel$ *that simultaneously governs exponential transverse stability, input-to-state robustness, thermal excitation limits, and NHIM persistence. We derive dimensionless small-gain conditions for nonlinearity and detuning that guarantee stability and illustrate how algebraic connectivity* λ_2 *provides a direct design knob for network robustness. Analytical results are supported by numerical simulations for small networks, showing topology-dependent decay rates, thermal floors, and the effect of nonlinear and detuning perturbations. Our framework bridges mathematical physics and quantum control, offering transparent design principles for scalable, robust open quantum networks.*

11.6 Laplacian-Guided Actuator Design for Quantum Network Synchronization

We develop a complete spectral control framework for synchronization in networked quantum systems with actuator-driven dynamics. Starting from an open quantum network model, we derive an exact Laplacian eigenmode decomposition that transforms the synchronization problem into a set of decoupled transverse modes. This yields closed-form exponential decay laws in which convergence rates are explicitly governed by the interplay between network eigenvalues and actuator strengths. Using Lyapunov methods, we establish global exponential stability of the synchronized manifold and derive topology-dependent scaling laws showing that convergence is fundamentally limited by the algebraic connectivity of the underlying graph. We prove that optimal actuator allocation follows a universal inverse-eigenvalue scaling law, whereby soft spectral modes require proportionally stronger control. Synchronization design is further formulated as a convex optimization problem via linear matrix inequalities, enabling systematic actuator synthesis under energy constraints. Robustness is quantified using both H_2 *and* $H\infty$ *norms, providing explicit disturbance amplification bounds and worst-case attenuation conditions. We additionally derive spectral controllability and observability criteria for distributed, node-level actuation and establish delay-dependent stability margins characterized by explicit analytical bounds. The resulting theory unifies spectral graph structure, optimal control design, and robustness analysis within a single mathematically transparent framework. Extensive simulations validate the analytical predictions across regular, random, and small-world networks. This work provides a rigorous foundation for actuator-driven synchronization in large-scale quantum networks and establishes spectral control principles applicable to emerging distributed quantum technologies.*

11.7 Stochastic Spectral Stability and Robust Actuator Synthesis in Imperfect Quantum Networks

Synchronization of quantum networks is increasingly relevant for distributed quantum computing, sensing, and communication architectures. Existing spectral actuator frameworks typically assume idealized network models and neglect the impact of parameter uncertainty, decoherence, stochastic noise, and time delays inherent to realistic quantum platforms. In this work, we develop a robust spectral actuator design framework for synchronization in imperfect quantum networks governed by stochastic Lindblad dynamics. We extend Laplacian-guided control to uncertain and noisy settings by deriving mode-wise stochastic decay bounds using Lyapunov and martingale techniques. Explicit analytical conditions are obtained for mean-square and almost-sure synchronization in terms of spectral gaps, actuator gains, and noise intensities. We formulate a min–max actuator allocation strategy under bounded parameter uncertainty and show that inverse-eigenvalue scaling remains optimal within provable robustness margins. Furthermore, we propose adaptive and distributed spectral-informed gain update laws that guarantee convergence without full network knowledge. The robust synthesis problem is recast as a convex semidefinite program enabling $H_2/H\infty$*-type performance optimization. Extensive numerical studies across ring, small-world, random, and scale-free topologies demonstrate finite-size scaling laws, robustness–performance trade-offs, and resilience against decoherence and communication delays. The results establish a unified spectral robustness theory for synchronization in realistic quantum networks and provide actionable design principles for experimental implementations.*

12.1 Adaptive Entanglement-Aware Routing in Multi-Hop Quantum Networks

Quantum networks promise secure communication and distributed quantum computing, yet efficient multi-hop routing remains a critical bottleneck due to probabilistic entanglement generation, decoherence, and limited quantum memory. Existing routing protocols often rely on classical heuristics or static metrics, failing to adapt to dynamic network conditions and link variability. In this work, we propose an adaptive entanglement-aware routing framework that integrates real-time network state information to optimize end-to-end entanglement fidelity while minimizing resource overhead. Our approach introduces a novel routing metric that balances path reliability, latency, and quantum memory utilization, enabling dynamic path selection in heterogeneous network topologies. Simulation results on representative quantum network scenarios demonstrate up to 45% improvement in entanglement

distribution success and significant reductions in network congestion compared to state-of-the-art methods. These findings highlight the potential of adaptive, fidelity-driven routing for scalable and resilient quantum communication, laying the groundwork for practical deployment in near-term quantum networks.

12.2 Distributed Probabilistic Multi-Commodity Adaptive Routing for Quantum Networks: Capacity Characterization, Stability Guarantees, and Heavy-Traffic Optimality

Quantum networks must route entanglement across unreliable, decohering, and probabilistic links while supporting multiple simultaneous communication demands. Existing routing approaches are predominantly heuristic or static, and lack formal guarantees on stability, optimality, or performance under stochastic link failures. This project develops a unified stochastic control framework for distributed quantum routing under uncertainty. We first characterize the capacity region of probabilistic multi-commodity quantum networks via a flow-based linear program incorporating stochastic link success and fidelity constraints. We then propose Distributed Probabilistic Multi-Commodity Adaptive Routing (DP-MCAR), a fully distributed queue-aware policy that requires only local state and reliability beliefs. Using Lyapunov drift analysis, we prove DP-MCAR is throughput-optimal for all arrival rates within the interior of the capacity region. We further derive explicit finite-time regret bounds relative to an oracle routing policy. To establish performance in congestion regimes, we analyze the heavy-traffic limit and prove a state-space collapse result showing that scaled queue dynamics converge to a one-dimensional reflected Brownian motion. This implies asymptotic diffusion optimality and minimal workload growth near the stability boundary. Extensive Monte-Carlo simulations under stochastic link failures and multi-commodity traffic validate the theoretical results, demonstrating superior throughput, delay performance, and robustness compared to existing routing strategies. The proposed framework provides the first integrated capacity, stability, and heavy-traffic theory for distributed probabilistic quantum routing.

12.3 Robust Learning-Augmented Routing in Time-Varying Quantum Networks: Stability, Regret, and Heavy-Traffic Optimality

Quantum networks operate under severe stochastic uncertainty, hardware drift, and partial state observability, making classical routing and control policies fundamentally inadequate. In practical near-term systems, link-level entanglement success probabilities are time-varying, unknown, and subject to environmental perturbations, requiring routing algorithms that jointly perform learning and queue stabilization. This project develops a rigorous theoretical framework for learning-augmented routing in time-varying quantum networks. We introduce an adversarial drift model for link reliabilities with bounded variation and propose a Thompson-sampling–based adaptive routing policy that integrates Bayesian exploration with queue-aware scheduling. The proposed policy simultaneously estimates link success probabilities and performs congestion control. We establish four main results. First, we prove robust throughput optimality under adversarial link drift, showing stability for all arrival rates within a drift-adjusted capacity region. Second, we derive dynamic regret bounds of order $O\left(\sqrt{T(1+V_T)}\right)$ relative to an oracle with full trajectory knowledge, matching optimal non-stationary bandit rates while preserving queue stability. Third, under heavy-traffic scaling, we prove a diffusion limit and derive an explicit closed-form asymptotic optimality gap constant relative to the oracle controller, quantifying the precise second-order performance penalty induced by learning. Fourth, we establish a diffusion control equivalence theorem showing asymptotic optimality up to a learning-induced covariance perturbation. The results unify stochastic network control, non-stationary bandit learning, and heavy-traffic diffusion theory within a single mathematical framework tailored to quantum networking. Numerical experiments calibrated to realistic quantum hardware parameters validate the theoretical predictions and confirm the tightness of the heavy-traffic covariance characterization

12.4 Congestion Control for Entanglement Distribution in Quantum Networks.

Future quantum networks will support large-scale entanglement distribution among multiple users for applications such as distributed quantum computing, quantum sensing, and quantum cryptography. Unlike classical packet networks, quantum networks must manage fragile entangled states whose fidelity degrades due to decoherence and limited quantum memory lifetimes. As a result, uncontrolled entanglement generation and swapping may lead to severe congestion, memory overflow, and fidelity collapse. In this project, we develop a formal congestion control framework for quantum entanglement distribution in multi-hop quantum networks. We model the network as a stochastic queueing system where entanglement pairs are generated probabilistically over elementary links, stored in finite quantum memories, and consumed through entanglement swapping operations to create end-to-end connections. We formulate congestion control as a stochastic network control problem that jointly regulates entanglement generation rates and path utilization to maximize long-term throughput while maintaining stability of quantum memory queues. Using Lyapunov drift techniques and heavy-traffic analysis, we derive a dynamic congestion control policy that adapts generation rates to network congestion levels and decoherence constraints. We prove that the proposed policy achieves throughput optimality and stabilizes the network whenever feasible rates exist. Furthermore, we characterize the heavy-traffic diffusion limit, revealing a reflected Brownian motion structure similar to classical communication networks but modified by quantum decoherence effects. Simulation results on multi-node quantum network topologies demonstrate that the proposed scheme significantly improves throughput, reduces memory overflow probability, and maintains higher end-to-end fidelity compared to uncontrolled entanglement generation. These results provide a first step toward a principled congestion control architecture for large-scale quantum networks.

12.5 Age of Entanglement: A Freshness Metric for Quantum Networks

The performance of quantum communication networks is traditionally evaluated using metrics such as throughput, fidelity, and entanglement generation rate. However, these metrics do not capture an important temporal property of entangled resources: their freshness. Due to decoherence and storage imperfections, entangled quantum states degrade over time, making the age of available entanglement a critical factor for reliable

quantum applications. In this paper we introduce Age of Entanglement (AoE), a new metric that quantifies the freshness of entangled states in quantum networks. AoE measures the elapsed time since the generation of the most recently usable entangled pair between communicating nodes. We develop a stochastic framework for analyzing AoE under probabilistic entanglement generation and limited quantum memory. Analytical expressions are derived for the evolution and average value of AoE in single-link and multi-hop quantum communication scenarios. The analysis reveals fundamental trade-offs between entanglement generation rate, storage capacity, and entanglement freshness. We further investigate scheduling and generation policies designed to minimize AoE, including policies that prioritize fresh entanglement generation and adaptive resource allocation. Numerical experiments illustrate the behavior of AoE under different network conditions, including varying link reliability, decoherence rates, and memory constraints. The results demonstrate that optimizing entanglement freshness can significantly improve the usability of distributed quantum resources. The proposed AoE metric provides a complementary performance measure to traditional metrics and offers new insights for the design, control, and optimization of large-scale quantum networks.

12.6 Multi-Flow Quantum Networks: Capacity, Freshness, and Scheduling of Entanglement Flows

Quantum networks aim to distribute entanglement between distant nodes to enable applications such as distributed quantum computing, secure communication, and networked quantum sensing. While many existing models focus on single source–destination pairs, practical quantum networks must support multiple simultaneous entanglement flows that compete for shared resources including quantum links, memory, and entanglement swapping operations. The stochastic nature of entanglement generation and the decoherence of stored quantum states introduce significant challenges for resource allocation and network performance. This project develops a stochastic analytical framework for multi-flow quantum networks that captures the dynamics of entanglement generation, storage, and routing across a network supporting multiple concurrent flows. The network is modeled as a graph in which links generate entangled pairs according to probabilistic processes and intermediate nodes perform entanglement swapping with limited success probability and memory capacity. We derive analytical expressions for the expected completion time of end-to-end entanglement generation along multi-hop paths and extend the analysis to systems with multiple competing flows. To quantify the timeliness of entanglement resources, we introduce the Age of Entanglement (AoE) metric, which measures the staleness of the most recently available entangled pair for a given flow. Using this metric, we analyze the impact of stochastic generation processes and scheduling policies on the freshness of entanglement across the network. We further characterize bounds on the multi-flow capacity region, describing achievable entanglement distribution rates under link-level generation constraints. The proposed framework provides theoretical insights into the trade-offs among throughput, entanglement freshness, and fairness in quantum networks. The results suggest that scheduling and routing policies designed for classical networks must be adapted to account for probabilistic entanglement generation and decoherence effects. Numerical experiments on simulated networks illustrate the behavior of the proposed models and validate the analytical results.

12.7 A Layered Framework for Entanglement Access Control in Quantum Networks

The emerging quantum internet will enable distributed quantum applications through the large-scale distribution of entanglement among networked quantum nodes. A fundamental challenge in such networks is the support of multiple simultaneous entanglement flows competing for limited quantum resources, including probabilistic link generation attempts, finite quantum memories, and time-sensitive entangled states subject to decoherence. While prior research has investigated entanglement routing and scheduling, the problem of systematic resource arbitration among competing flows has not yet been addressed through a dedicated control layer. In this project, we introduce Entanglement Access Control (EAC), a new architectural layer for quantum networks responsible for regulating access to entanglement resources in multi-flow environments. Analogous to the Medium Access Control (MAC) layer in classical communication networks, EAC manages contention among flows by coordinating entanglement generation, memory allocation, and scheduling decisions while accounting for the freshness and fidelity of entangled pairs. We develop an analytical framework to study the performance of multi-flow quantum networks under EAC, incorporating stochastic entanglement generation, memory constraints, and entanglement decay. Key performance metrics include throughput, Age of Entanglement (AoE), and fidelity. Several scheduling policies are examined within the EAC framework, including first-come-first-served, last-come-first-served, freshness-aware, and fairness-oriented strategies. Analytical insights are complemented by numerical simulations demonstrating how scheduling policies, network parameters, and hardware constraints jointly influence system performance. The results show that introducing an EAC layer provides a systematic abstraction for managing entanglement resources in multi-flow quantum networks and offers a principled foundation for scalable quantum network protocol design.

12.8 QTCP — Quantum Transport Control Protocol

Quantum networking requires reliable end-to-end delivery of entangled states despite the stochastic nature of entanglement generation, limited quantum memory, and decoherence effects. While recent work has addressed quantum routing and resource scheduling, the design of a dedicated quantum transport layer remains largely unexplored. This project introduces Quantum Transmission Control Protocol (QTCP), a transport-layer framework for reliable multi-flow entanglement distribution in large-scale quantum networks. QTCP provides end-to-end flow control, congestion management, and reliability mechanisms adapted to the unique constraints of quantum communication, including probabilistic entanglement generation, finite memory capacity, and fidelity decay. The proposed framework integrates with the previously developed Entanglement Access Control (EAC) layer and quantum routing mechanisms, forming a layered architecture for scalable quantum networking. To address realistic network conditions, the project further considers time-varying quantum networks and proposes a machine-learning-assisted QTCP (ML-QTCP) that predicts link availability and adapts transmission decisions accordingly. An analytical framework is developed to characterize throughput, Age of Entanglement (AoE), and protocol stability under dynamic network conditions. Numerical evaluations demonstrate that ML-QTCP improves

throughput and resource utilization compared to conventional scheduling approaches. The proposed transport architecture provides a foundation for future research on joint cross-layer optimization of quantum networking protocols.

12.9 Distributed and Hierarchical Quantum Routing for Large-Scale Networks.

The realization of large-scale quantum networks will require routing mechanisms capable of operating efficiently across geographically distributed and dynamically evolving infrastructures. Unlike classical networks, quantum networks must manage probabilistic entanglement generation, limited quantum memory, decoherence, and stochastic link availability, making scalable routing particularly challenging. In this project, we propose a distributed and hierarchical quantum routing architecture designed for large-scale multi-domain quantum networks. The proposed framework decomposes the network into routing domains connected through gateway nodes and employs distributed routing protocols that operate using local network state information while coordinating global path establishment across domains. To address the time-varying nature of quantum links, we further integrate a federated learning–assisted prediction layer that enables network nodes to collaboratively estimate link generation rates, entanglement success probabilities, and resource availability without sharing raw operational data. These predictions are incorporated into the routing metric to support predictive path selection in dynamic environments. We develop analytical models that characterize routing delay, scalability properties, and the impact of prediction accuracy on end-to-end entanglement establishment time. The proposed architecture enables efficient routing decisions while maintaining low signaling overhead and preserving domain autonomy. Extensive numerical evaluations demonstrate that hierarchical routing significantly improves scalability compared with flat routing architectures, while the federated learning–assisted prediction mechanism further reduces routing delay and improves entanglement throughput in time-varying networks. The results highlight the potential of combining distributed routing architectures with collaborative machine learning techniques to enable scalable and adaptive quantum network control. This work establishes a foundation for intelligent large-scale quantum networking and provides a key building block for future distributed quantum transport protocols

12.10 Latency–Decoherence Tradeoffs in Proactive Entanglement Distribution for Large-Scale Quantum Networks

The efficient distribution of entanglement is a fundamental requirement for the operation of large-scale quantum networks and the emerging quantum internet. A central architectural question concerns whether entanglement should be generated reactively, in response to connection requests, or proactively, by continuously generating and storing entangled pairs in quantum memories. Reactive entanglement distribution minimizes decoherence effects but introduces significant path establishment latency due to sequential link generation and entanglement swapping. In contrast, proactive entanglement buffering enables rapid path establishment by retrieving pre-generated link entanglement from memory, but suffers from fidelity degradation caused by decoherence during storage. This project presents a comprehensive analytical framework for studying the latency–decoherence tradeoff between proactive and reactive entanglement distribution strategies. We first develop stochastic models for serial and parallel reactive entanglement generation along multi-hop repeater paths and derive closed-form expressions for path establishment latency and end-to-end success probability. We then introduce a proactive entanglement distribution model in which link-level entangled pairs are continuously generated and stored in quantum memories. Using an exponential decoherence model, we derive analytical expressions for the probability of usable entanglement as a function of storage time and memory coherence properties. Building on these models, we establish threshold conditions under which proactive entanglement buffering becomes advantageous relative to reactive generation. We further derive the optimal entanglement refresh rate that maximizes usable entanglement throughput in the presence of decoherence and show that the optimal generation policy is closely related to the quantum memory coherence time. The analysis also reveals distinct operational regimes in which reactive, proactive, or hybrid entanglement distribution strategies are optimal. Numerical results illustrate the tradeoffs between latency, decoherence, and network scale, demonstrating that proactive strategies can significantly reduce path establishment delays in long multi-hop quantum networks when supported by sufficiently coherent quantum memories. The results provide new insights into the design of entanglement management policies and highlight the importance of coordinated entanglement generation and storage mechanisms for scalable quantum networking architectures.

12.11 Distributed Entanglement Scheduling in Multi-Domain Quantum Networks

The realization of a large-scale quantum internet requires efficient coordination of entanglement resources across multiple administrative domains with heterogeneous capabilities and limited information sharing. While prior work has addressed routing and entanglement generation, the problem of distributed entanglement scheduling remains largely unexplored, particularly under the combined constraints of decoherence, probabilistic entanglement generation, and domain autonomy. In this project, we develop a rigorous theoretical framework for distributed entanglement scheduling in multi-domain quantum networks. We introduce the notion of entanglement utility, which captures the trade-off between entanglement rate and decoherence-induced loss, and formulate the scheduling problem as a network-wide utility maximization under quantum-specific constraints. Building on this formulation, we propose a distributed entanglement scheduling algorithm based on local queue dynamics and congestion pricing, enabling scalable operation without centralized coordination. To address the challenges of inter-domain resource allocation, we further develop a quantum network resource market model, in which domains advertise entanglement prices reflecting both capacity and decoherence risk. We prove that the resulting distributed pricing mechanism converges to a global equilibrium that maximizes aggregate entanglement utility. In addition, we establish a decoherence-aware queueing model and derive stability conditions using Lyapunov analysis, demonstrating that the proposed scheduler ensures bounded entanglement queues under feasible load conditions. The paper also introduces a multi-domain scaling law, quantifying the impact of domain fragmentation on achievable entanglement throughput, and shows that hierarchical coordination is essential for maintaining scalability. Numerical evaluations validate the analytical results and illustrate the trade-offs between throughput, delay, and entanglement survivability across a wide range of network parameters.
The proposed framework provides a unified theoretical foundation for entanglement scheduling, bridging quantum network physics with distributed systems design, and constitutes a key building block for future quantum transport protocols and learning-assisted network control.

12.12 Federated Learning for Quantum Network State Prediction: A Discrete Event Systems Approach with Stability–Utility Trade-offs

We consider the problem of predictive modeling in quantum networks, where distributed nodes must estimate future network states under stochastic entanglement dynamics, partial observability, and strict communication constraints. Centralized learning approaches are often impractical due to bandwidth limitations and the sensitivity of locally observed quantum states, while conventional federated learning methods fail to capture the event-driven nature of quantum network evolution. In this project, we propose a novel federated learning framework grounded in a discrete event systems (DES) formulation of quantum network dynamics. The proposed approach introduces event-triggered model aggregation, where communication and updates are activated based on prediction error and changes in quantum resource conditions rather than periodic schedules. We further define an effective quantum resource metric, Q_{eff}, capturing fidelity, availability, and coherence lifetime, and incorporate it into both the learning objective and aggregation weighting. We formalize the resulting stability–utility trade-off through a composite loss function that balances prediction accuracy, model stability, and quantum resource quality. Under mild assumptions on update frequency and communication delay, we establish conditions for bounded prediction error and convergence of the federated process. Simulation results demonstrate that the proposed method significantly reduces communication overhead while maintaining or improving prediction accuracy compared to periodic federated and centralized baselines.

12.13 Distributed QTCP with Federated Learning over Time-Varying Topology-Aware Multidomain Quantum Networks

We propose a distributed control and learning framework for Quantum TCP (QTCP) in time-varying, topology-aware multi-domain quantum networks. The approach integrates queue-aware entanglement flow control with federated learning, enabling network nodes to collaboratively adapt transmission policies under dynamic connectivity and quantum resource constraints. Each node maintains quantum buffer states and entanglement distribution rates, which evolve according to coupled queue dynamics and learning-driven control updates. The framework explicitly incorporates quantum network characteristics, including entanglement fidelity and decoherence, through a topology-aware resource weighting mechanism that guides the aggregation of local models. To reduce coordination overhead, we introduce an event-triggered communication scheme that selectively activates federated updates based on local state deviations. We establish stability conditions for the coupled quantum queue–learning system and characterize the trade-off between communication cost and network performance, including entanglement throughput, delay, and fidelity degradation. Numerical evaluations across multiple time-varying network topologies demonstrate that the proposed method achieves improved stability–utility trade-offs compared to topology-agnostic and non-learning-based approaches. The results highlight the effectiveness of integrating federated learning, event-triggered communication, and topology-aware design for scalable quantum network control. The proposed framework further provides a foundation for subsequent joint optimization of the quantum protocol stack.

12.14 On the Structure of Optimal Full-Stack Policies in Decoherence-Aware Quantum Network Protocol Stacks

We develop a rigorous framework for cross-layer optimization in quantum network protocol stacks under intrinsic decoherence constraints. Unlike classical approaches that treat layers in isolation, we consider a full-stack joint control policy in which physical resources, entanglement access control, network routing, transport-layer scheduling, and learning mechanisms are optimized simultaneously. Decoherence is explicitly incorporated into the system dynamics, ensuring that all results reflect fundamentally quantum limitations. We formulate the problem as an infinite-horizon stochastic control problem over a multi-domain, topology-aware quantum network with time-evolving states and delayed information exchange. Although the optimization problem is formulated at the network level, the resulting policy admits a distributed implementation, with each node executing its own protocol stack based on local information and limited interaction with neighbors. We establish well-posedness of the resulting dynamical system and derive sufficient conditions for stability and contraction in the presence of decoherence-induced decay. A key contribution is the characterization of the structure of optimal policies, showing that threshold-based control, effectiveness-weighted aggregation, and topology-aware routing emerge naturally from the joint optimization. Importantly, we prove that layer dominance is not imposed a priori but arises endogenously from the optimization, depending on network topology, decoherence rates, and communication constraints. We derive explicit bounds on the steady-state global error and quantify trade-offs between accuracy, communication cost, and physical resource consumption. Numerical results over multi-domain networks with heterogeneous topologies corroborate the theoretical findings and illustrate regime-dependent dominance of control, network, and physical layers.

This work provides a unified, mathematically rigorous foundation for full-stack cross-layer optimization in quantum networks and establishes a basis for subsequent studies under more specialized operational regimes.

12.15 Distributed Learning-Based Joint Optimization of Quantum Network Protocol Stacks with Decoherence-Aware Dynamics

Cross-layer optimization of quantum network protocol stacks is essential for achieving scalable and reliable operation in multi-domain environments subject to decoherence and dynamic topology. While existing approaches rely primarily on analytically designed control policies, they are inherently limited in their ability to adapt to complex and time-varying network conditions. In this project, we develop a learning-driven framework for joint optimization of all layers of the quantum network protocol stack, including the physical, entanglement access control, network, QTCP, and learning layers. The proposed approach formulates a global optimization problem over the network state, while the corresponding control policy is parameterized and learned through distributed, event-triggered updates across nodes. We introduce a class of decentralized learning policies that

operate on node-level protocol stacks and collectively implement a network-level objective. The framework explicitly incorporates decoherence dynamics, topology-aware interactions, and communication constraints, enabling adaptive decision-making under realistic operating conditions. We establish theoretical guarantees on the stability and convergence of the learned policies, including bounds on the global error in the presence of stochastic disturbances and asynchronous updates. In addition, we characterize the trade-off between communication cost and learning performance induced by event-triggered coordination mechanisms. Extensive simulations on multi-domain networks with heterogeneous topologies demonstrate that the proposed learning-driven approach significantly outperforms fixed-policy baselines in terms of convergence speed, steady-state error, and communication efficiency. The results highlight the critical role of data-driven cross-layer coordination in achieving near-optimal performance in large-scale quantum networks. The proposed framework provides a principled foundation for integrating learning and control in the design of next-generation quantum network protocol stacks.

12.16 Learning-Driven Predictive Cross-Layer Optimization for Time-Varying Quantum Networks

Quantum networks operate in inherently dynamic environments characterized by time-varying channel conditions, stochastic entanglement generation, and rapid degradation of quantum states due to decoherence. While recent cross-layer optimization frameworks have demonstrated the benefits of learning-driven control, existing approaches are predominantly reactive and rely on instantaneous or historical observations. Such formulations are fundamentally limited in their ability to anticipate future network states, particularly in the presence of propagation delays and fast temporal variations. This project introduces a predictive cross-layer optimization framework for time-varying quantum networks, extending prior learning-driven formulations by incorporating explicit prediction of future system states into the decision-making process. The proposed approach models the network as a stochastic dynamical system and integrates data-driven prediction with finite-horizon optimization, resulting in a learning-driven predictive control architecture. Within this framework, three formulations are developed, including a baseline Model capturing reliability-aware cross-layer interactions, a throughput-centric S-Model enabling direct comparison with classical network optimization, and a generalized multi-objective G-Model that jointly accounts for throughput, fidelity, and resource consumption over a prediction horizon. Theoretical analysis establishes that predictive control strictly improves performance relative to reactive policies under bounded prediction error and characterizes the impact of prediction accuracy on achievable performance. In particular, explicit bounds are derived linking prediction error to effective throughput degradation, revealing a fundamental trade-off between prediction horizon and robustness. Numerical results based on multi-curve evaluations demonstrate significant performance gains in dynamic regimes and illustrate how predictive policies adapt to varying network conditions. The proposed framework provides a unified methodology for anticipatory cross-layer optimization in quantum networks, bridging learning, prediction, and control. These results offer new insights into the role of foresight in mitigating decoherence and advancing the design of scalable and adaptive quantum communication systems.

12.17 Belief-Aware Predictive Cross-Layer Optimization in Quantum Networks

This predictive cross-layer optimization has recently emerged as a fundamental mechanism for improving throughput in quantum networks by leveraging time-varying system dynamics. Existing formulations, however, rely on the assumption of fully observable network states and treat prediction error as an external perturbation. In practical quantum networking scenarios, the underlying system state is inherently uncertain due to decoherence, stochastic entanglement generation, and limited measurement capabilities, necessitating a unified treatment of inference and control. In this project, we develop a belief-aware predictive cross-layer optimization framework that integrates estimation, prediction, and control within a single formulation. The network is modeled as a partially observable stochastic system in which control decisions, including entanglement generation rate, routing, and quantum transport window, are functions of a belief state representing the posterior distribution over latent network conditions. We derive a generalized throughput functional that incorporates both predictive gain and uncertainty-induced degradation, revealing a fundamental dependence on belief entropy in addition to prediction horizon. The proposed framework establishes new structural results characterizing the interplay between information and performance. In particular, we show that achievable throughput is bounded by an entropy-dependent gap relative to the fully observable case, and that predictive gain decomposes into temporal and informational components. Furthermore, we demonstrate that optimal operation requires joint optimization of sensing and control, yielding a strictly expanded Pareto region compared to passive prediction-based schemes. The analysis also identifies regimes in which uncertainty collapses the optimal prediction horizon, recovering reactive control as a limiting case. The developed Model, S-Model, and G-Model formulations provide a consistent hierarchy spanning tractable approximations to full belief-space optimization, and reduce to prior predictive results under perfect observability. These results establish belief-aware optimization as a necessary extension of predictive quantum networking and provide a foundation for information-constrained control and distributed inference in large-scale quantum systems.

12.17A Belief-Aware Predictive Cross-Layer Optimization in Quantum Networks

We propose a predictive cross-layer optimization framework for quantum networks that captures the fundamental interaction between prediction horizon selection, transport-layer control, uncertainty, and congestion. The model integrates diminishing-return predictive gains with horizon-dependent penalties arising from entropy, prediction error, and network load, resulting in a tractable yet expressive analytical formulation. A joint optimization problem over the prediction horizon and transport window is studied, and closed-form expressions for the optimal control variables are derived. Despite the non-convex structure of the objective function, the analysis establishes the existence and uniqueness of a global optimum under mild conditions. The results reveal a fundamental structural property: the optimal prediction horizon depends only on transport efficiency and network load, while uncertainty influences the feasible transport window without directly affecting the horizon. This decoupling provides both theoretical insight and practical advantages for system design. The framework is further extended to a generalized model with dynamic belief evolution, as well as to an alternative horizon-dependent error formulation that exhibits invariance of the optimal control variables with respect to constant prediction error components. These extensions unify multiple modeling approaches and highlight the role of information structure in predictive optimization. Scaling laws and feasibility conditions are derived, identifying distinct operating regimes and establishing that the optimal

prediction horizon decreases inversely with network load. A multi-user extension is also developed, revealing how congestion-induced coupling affects prediction and transport decisions in shared environments. Numerical results validate all theoretical findings and illustrate the structural properties of the optimal solution. The proposed framework provides a principled foundation for predictive control in quantum networks, bridging information-theoretic modeling and cross-layer system design.

12.17B Topology-Aware Predictive Control in Networked Quantum Systems: Spectral Scaling, Phase Transitions, and Graph-Structured Optimization

This project develops a topology-aware framework for predictive cross-layer optimization in networked systems, extending classical homogeneous models to graph-structured environments with spatially heterogeneous congestion. The proposed formulation introduces node-dependent load dynamics through the network adjacency structure, leading to a coupled optimization problem in which prediction horizons and transport control variables interact across the graph. We first derive closed-form expressions for node-level optimal prediction horizons and show that they depend explicitly on local congestion, which is determined by network topology. The resulting system is characterized as a fixed-point problem, and existence and uniqueness of equilibrium are established under mild conditions. A global variational formulation is then introduced, demonstrating that the equilibrium solution coincides with the optimum of a network-wide utility maximization problem, thereby connecting predictive control with classical congestion pricing interpretations. A central contribution of the paper is the derivation of spectral scaling laws that characterize the impact of topology through graph eigenvalues. In particular, the optimal prediction horizon is shown to scale inversely with the spectral radius of the adjacency matrix, while the heterogeneity of node-level horizons is governed by the spectral gap of the graph Laplacian. These results provide a unified framework linking predictive control performance to fundamental graph properties. The analysis further reveals a topology-dependent phase transition between predictive and congestion-dominated regimes, establishing conditions under which non-trivial prediction horizons exist. The role of canonical network topologies, including small-world, random, scale-free, and hierarchical graphs, is systematically examined, highlighting how structural properties influence both performance and fairness. Numerical results validate the theoretical findings and illustrate the interplay between topology, congestion, and prediction. The proposed framework provides a principled foundation for topology-aware design of networked systems and offers new insights into the fundamental limits of predictive control in complex networks.

12.18 Learning-Augmented Predictive Control in Time-Varying Quantum Networks

Quantum networks rely on efficient entanglement distribution under uncertainty, stochastic generation, and stringent resource constraints. In this work, we propose a learning-augmented model predictive control (MPC) framework for distributed management of entanglement flows in dynamic quantum network environments. The approach integrates predictive optimization with online parameter adaptation, enabling control policies that account for both network topology and uncertainty in link success probabilities. A distributed implementation is developed, establishing a connection between MPC solutions and diffusion-type control laws, thereby enabling scalable deployment. From a theoretical standpoint, we characterize the closed-loop system through time-varying network operators and establish explicit stability conditions in terms of spectral properties. In particular, the spectral radius of the induced dynamics is shown to govern convergence and tracking performance, and finite-time bounds are derived under bounded disturbances and parameter mismatch. These results provide a principled foundation for learning-enabled control in large-scale quantum networks. Extensive simulations across multiple network topologies, including ring, grid, random, and scale-free graphs, demonstrate that the proposed framework yields substantial system-level improvements. In particular, throughput gains of 30%–70% are consistently achieved relative to diffusion-based control, while discarded entanglement is reduced by up to 80%. The proposed method also significantly improves resource utilization efficiency and maintains robust performance under high-load conditions, where reactive strategies degrade due to congestion and mismatch effects. Overall, the results show that integrating prediction and learning into the control loop provides a scalable and robust solution for quantum network operation, bridging theoretical guarantees with practical performance gains.

12.19 Fundamental Limits of Information-Constrained Quantum Networking with Predictive Control

Quantum networks rely on the efficient generation, distribution, and utilization of entanglement under stringent resource and information constraints. In this work, we investigate the fundamental performance limits of information-constrained quantum networking and the extent to which practical control strategies can approach these limits. We develop a predictive control framework based on model predictive control (MPC) and compare it against diffusion-based routing across multiple network topologies and operating regimes. Performance is evaluated using throughput, resource efficiency, and discarded entanglement as key metrics. The results demonstrate that MPC consistently improves performance by proactively directing entanglement flows, reducing resource wastage, and mitigating congestion effects. The gains are particularly significant in heterogeneous network topologies, where structural asymmetries and bottlenecks limit the effectiveness of purely local strategies. In contrast, diffusion-based methods exhibit saturation and increased discard under high load conditions due to their reactive nature. To assess proximity to fundamental limits, we introduce an oracle benchmark representing an upper performance envelope under ideal coordination. The proposed MPC approach operates close to this benchmark across all scenarios, capturing a large fraction of the achievable throughput while maintaining high efficiency. These findings highlight the critical role of predictive, information-aware control in overcoming coordination constraints and improving performance in realistic quantum networking environments. The results provide new insights into the interplay between network structure, information availability, and achievable performance, and establish a practical pathway toward approaching fundamental limits in quantum networks.

12.20 Fundamental Limits of Topology-Dependent Holevo Capacity in Imperfect Quantum Networks

We investigate the fundamental performance limits of quantum networks under realistic constraints, including decoherence, finite resources, and imperfect entanglement operations. In contrast to existing studies that rely primarily on simulation or idealized assumptions, we develop a unified analytical and computational framework that explicitly captures the interplay between network topology, loss mechanisms, and control strategies. We show that achievable throughput is intrinsically topology-dependent, with heterogeneous networks—such as scale-free graphs—exhibiting

significantly higher capacity due to reduced path lengths and enhanced connectivity. Building on these insights, we formulate a model predictive control (MPC) approach for adaptive routing and resource allocation, and benchmark its performance against diffusion-based schemes and an oracle upper bound. Our results demonstrate that MPC consistently outperforms diffusion across all considered topologies, achieving substantial gains in throughput and resource efficiency while significantly reducing discard rates. Moreover, MPC operates close to the oracle benchmark, effectively approaching topology-dependent capacity limits, particularly in well-connected networks. We further quantify the impact of loss and congestion, showing that performance degradation is governed by a nonlinear interaction between path length and load, and that MPC mitigates these effects by exploiting global network structure. These findings establish a clear connection between topology, control, and fundamental limits in quantum networks, providing both theoretical insight and practical guidelines for the design of scalable and efficient quantum communication infrastructures.

12.21 Fundamental Limits of Topology-Dependent Holevo Capacity in Imperfect Quantum Networks with Error Correction

We investigate the fundamental performance limits of time-varying quantum networks under realistic constraints, including decoherence, finite resources, and imperfect entanglement operations. Building on a unified analytical and computational framework, we quantify how network topology, path length, and quantum error correction (QEC) interact to determine achievable end-to-end throughput. In contrast to fixed-QEC schemes, we introduce a path-length-dependent adaptive QEC strategy, where code distances are selected based on network paths and target logical error probabilities. Our simulations reveal that adaptive QEC significantly improves throughput and resource efficiency across random, scale-free, and lattice networks, while mitigating logical errors on long paths. Multi-curve analysis demonstrates that lattice networks, previously insensitive to code distance, benefit substantially under adaptive QEC. These findings establish a direct link between topology, adaptive control, and fundamental communication limits, providing design guidelines for scalable and efficient quantum communication infrastructures.

12.22 Topology-Aware Adaptive Quantum Error Correction for Time-Varying Quantum Networks: Theory and Performance Limits

Quantum networks face fundamental challenges due to decoherence, limited quantum resources, and dynamically evolving topologies. While quantum error correction (QEC) is essential for reliable communication, existing approaches largely rely on static code configurations and neglect the impact of network structure and temporal variability. In this project, we propose a topology-aware, dimension-adaptive quantum error correction (TA-DA-QEC) framework that jointly optimizes code distance and code dimensionality based on network conditions. The proposed approach introduces a path-length–dependent effective error model and dynamically selects between two-dimensional (2D) and three-dimensional (3D) QEC schemes to balance reliability and resource overhead. We develop analytical scaling laws showing that logical error probability decays exponentially with code distance in 2D codes and super-exponentially in 3D codes, leading to distinct operational regimes. Based on these results, we derive a dimension-switching criterion and construct a phase diagram that identifies when higher-dimensional QEC provides net throughput gains. Extensive simulations across random, scale-free, and lattice network topologies demonstrate that the proposed adaptive framework achieves significant improvements in throughput, logical error suppression, and resource efficiency compared to static QEC strategies. In particular, topology-dependent transitions to 3D encoding yield substantial performance gains in high-noise and long-path regimes. These findings establish a new paradigm for network-aware quantum error correction, highlighting the importance of jointly optimizing coding structure and network dynamics for scalable quantum communication systems.

G1 Entanglement Graph Reconfiguration Under Noise: A Channel-Theoretic Framework for Quantum Network Topology Control

We develop a rigorous analytical framework for quantum state graph reconfiguration under imperfect operations, where entanglement links between nodes are dynamically established, removed, or rerouted in the presence of decoherence, noise, and non-ideal control. The system is modeled as a controlled stochastic process over an exponential graph state space, with reconfiguration operations inducing a path measure on graph trajectories. The main result establishes a unified exponential performance law, showing that both fidelity and success probability decay exponentially with the reconfiguration length, with exponents characterized by large deviations rate functions. We prove that achievable performance is fundamentally limited by the relative entropy between the implemented and ideal path measures, yielding tight information-theoretic bounds on fidelity and reliability.

G2 Topology Aware Entanglement Graph Reconfiguration in Time Varying Quantum Networks

We develop a channel-theoretic framework for entanglement graph reconfiguration under noise, unifying topology, control, and learning. Each entangling operation induces a multiplicative fidelity decay, yielding a log-additive model amenable to large deviations analysis. We show that topology influences performance through an effective exponent capturing interaction density and structural amplification. This leads to (i) a phase transition in optimal topology as a function of noise, (ii) finite-length corrections via a central limit refinement, and (iii) exponential regret for

suboptimal topology selection. We formulate topology control as a sequential decision problem and derive a threshold-based optimal policy. Under uncertainty, we prove a minimax sparsity principle, and propose an adaptive estimator that learns the critical threshold online with vanishing regret. Numerical results demonstrate distinct roles of noise sensitivity, topology-induced scaling, and cost–fidelity trade-offs, and reveal discrete resource-induced transitions in dense graphs. The framework provides a principled foundation for AI/ML-assisted topology control in noisy quantum networks.

G3 Entanglement Graph Reconfiguration Under Noise: Foundations of Stochastic Topology Control in Quantum Networks

Quantum networks rely on dynamically evolving entanglement topologies whose structure must be continuously adapted in response to noise, decoherence, link failures, and changing communication demands. While existing quantum routing and entanglement distribution frameworks largely assume static or idealized topologies, the fundamental principles governing noisy entanglement graph reconfiguration remain poorly understood. In this project, we develop a rigorous analytical framework for entanglement graph reconfiguration under imperfect operations, where entanglement links are sequentially established, removed, or rerouted in the presence of stochastic noise and decoherence. The evolution of the network topology is modeled as a controlled stochastic process over an exponential graph-state space, with reconfiguration operations inducing probability measures over graph trajectories. Using a channel-theoretic formulation, we show that reconfiguration fidelity and success probability obey a universal exponential scaling law characterized by large deviations rate functions. We establish information-theoretic performance bounds demonstrating that achievable reliability is fundamentally limited by the relative entropy between implemented and ideal graph trajectory measures. The reconfiguration problem is further formulated as a divergence-regularized stochastic control problem, yielding Bellman optimality equations and structural properties of optimal policies, including monotonicity and threshold behavior. We additionally prove computational hardness results showing that optimal reconfiguration is EXPTIME-hard in general network settings, thereby motivating structured and distributed protocol design. The framework further enables the analysis of practical protocol classes, including greedy, shortest-path, multipath, and entanglement-preserving strategies, whose performance is characterized through induced large deviations exponents and cost–fidelity trade-offs. Numerical results validate the theoretical predictions, confirming exponential scaling laws, topology-dependent phase transitions, and finite-length effects. The results establish a unified foundation connecting stochastic control, information theory, and quantum network topology evolution, and provide a principled basis for scalable and learning-assisted quantum network control.

G4 Cognitive Control of Time-Varying Topology for Adaptive Quantum Networks

Future quantum Internet systems will operate in highly dynamic environments characterized by stochastic entanglement generation, topology fluctuations, congestion variability, decoherence, and adaptive routing requirements. Classical static routing and fixed-topology formulations are insufficient for such large-scale adaptive quantum-network environments. This project develops a unified stochastic-control framework for time-varying topology control in adaptive quantum Internet systems. The network is modeled as a dynamic entanglement graph $G_t = (V, E_t, W_t)$ *whose topology, link fidelities, capacities, and congestion states evolve according to coupled stochastic graph dynamics. The framework integrates stochastic topology evolution, entropy-regularized routing, adaptive congestion control, amplification-aware optimization, and cognitive (learning-aided) adaptation. An entropy-regularized stochastic-control formulation is developed using the soft Bellman framework* $V(\mathcal{X}) = -\lambda log \sum_a exp\,(-Q(\mathcal{X}, a)/\lambda)$ *which provides robust adaptive routing under uncertain and time-varying network conditions. Stability and convergence properties are established through Lyapunov drift analysis, ergodicity theory, and contraction arguments for entropy-regularized routing dynamics. To characterize large-scale structural effects, the paper introduces the topology amplification observables* $\gamma(G_t)$ *and* $\kappa(G_t)$ *which quantify topology-induced amplification of routing inefficiency, congestion growth, fluctuation propagation, and stochastic instability. The framework further incorporates cognitive solutions based on Upper Confidence Bound learning, Thompson sampling, and reinforcement-learning-based adaptive routing, thereby enabling online topology adaptation under partially unknown network dynamics. Numerical experiments demonstrate robust adaptive rerouting under stochastic link failures, congestion stabilization through entropy regularization, topology phase-transition behavior, and favorable scalability properties for small-world hybrid architectures. The proposed framework establishes a unified theoretical foundation connecting adaptive quantum networking, stochastic graph dynamics, reinforcement learning, statistical mechanics, and large-scale autonomous cognitive topology control. The results further provide a roadmap toward future developments involving heavy-traffic quantum-network theory, topology phase transitions, mean-field scaling laws, and large-scale reinforcement-learning quantum control.*

G5 Heavy-Traffic Limits for Adaptive Quantum Networks with Dynamic Topology Control

Future quantum Internet infrastructures are expected to operate under highly dynamic conditions characterized by stochastic entanglement generation, topology fluctuations, congestion bursts, and near-capacity traffic loads. Although recent work has investigated routing and topology optimization for adaptive quantum networks, the behavior of such systems under heavy traffic remains poorly understood. This project develops a heavy-traffic framework for adaptive quantum networks with dynamic topology control. The network is modeled as a time-varying entanglement graph with stochastic queue evolution, adaptive topology reconfiguration, fidelity-aware routing, and congestion-dependent resource allocation. The analysis focuses on the heavy-traffic regime in which network utilization approaches saturation, i.e., traffic intensity approaches unity. Under this regime, fluid limits are derived to characterize deterministic large-scale dynamics, while diffusion approximations capture stochastic fluctuations around congested operating points. The proposed framework establishes reflected diffusion limits for congestion processes, derives state-space collapse results that reduce high-dimensional queue dynamics to lower-dimensional workload processes, and introduces heavy-traffic amplification measures that quantify topology-dependent congestion growth under overload conditions. The analysis further investigates adaptive routing

mechanisms including learning-based control, entropy-regularized routing, and bandit-based exploration under saturation. Simulation studies demonstrate queue blow-up phenomena, delay amplification, fidelity degradation, and strong topology dependence across multiple network architectures. The results provide a mathematical foundation for understanding congestion behavior in adaptive quantum networks and establish connections between heavy-traffic theory, dynamic topology control, and scalable quantum Internet architectures.

G6 Fundamental Limits of Cognitive Congestion Control: Learning–Congestion Tradeoffs in Dynamic Quantum Networks

Cognitive congestion control has emerged as a promising approach for adaptive resource management in dynamic quantum networks. However, learning-driven decision processes introduce exploration, delayed adaptation, and information uncertainty that may significantly increase congestion and degrade network stability. Existing approaches typically optimize either learning efficiency or congestion performance independently, providing limited understanding of the fundamental tradeoffs between adaptation and queue stability. This project develops a theoretical framework for cognitive congestion control based on learning–congestion tradeoffs in dynamic quantum networks. We introduce a joint queue–regret formulation that captures the interaction between learning efficiency and congestion accumulation and derive analytical conditions under which adaptive control policies achieve both logarithmic learning performance and bounded average congestion. The proposed framework establishes stability conditions under controlled exploration and characterizes how topology, delayed feedback, and incomplete information affect achievable performance. The analysis reveals fundamental limits governing simultaneous learning and congestion control and demonstrates that adaptive network optimization is constrained by intrinsic queue–regret tradeoffs. These results provide theoretical foundations for the design of scalable congestion-aware cognitive control mechanisms in future dynamic quantum infrastructures.

G7 Network Control Experience Management: A State-Space Framework for Intelligent Communication Networks

Intelligent communication systems increasingly rely on accumulated experience derived from observations, learning processes, and past operational outcomes to support network management and decision making. While Age of Information (AoI) has emerged as a fundamental framework for quantifying information freshness, it does not capture the freshness and usefulness of accumulated experience that drives adaptive system behavior. As communication networks evolve toward AI-native, autonomous, and experience-driven operation, managing the freshness of experience becomes increasingly important. This project introduces the concept of Network Control Experience Management (NCEM)), a general framework for experience-aware management of intelligent systems. The proposed framework represents experience through three complementary state variables: experience content, experience freshness, and experience relevance. Building upon these concepts, we develop models for experience evolution, relevance-aware refresh mechanisms, and adaptive experience management under resource constraints. To establish a direct connection between experience freshness and system performance, we introduce the notion of Experience Utility, which quantifies the effective value of accumulated experience by jointly accounting for experience content, freshness, and relevance. The resulting framework links experience management to decision support and provides a quantitative basis for analyzing how aging experience influences operational effectiveness. The paper further develops optimal experience-updating formulations, relevance-aware resource-allocation mechanisms, and stochastic Age of Experience (AoE) models that relate refresh probability to expected experience utility. A quantum-network case study is used to demonstrate the applicability of the proposed framework. Numerical results show that relevance-aware and adaptive refresh strategies improve Experience Utility and reduce AoE compared with conventional refresh mechanisms. The analysis also reveals a fundamental sampling-delay tradeoff, demonstrating the existence of an optimal operating point that balances experience freshness against communication overhead. The proposed NCEM framework extends freshness analysis beyond information-centric models and provides a foundation for experience-aware management, optimization, and decision support in future intelligent communication systems.

G7a Network Control Experience State and Management for Intelligent Communication Networks

Freshness-aware communication systems have traditionally relied on metrics such as Age of Information (AoI) and its extensions to characterize information staleness. While these approaches provide valuable insight into temporal freshness, they do not fully capture the multidimensional nature of accumulated experience required for intelligent network operation. This project introduces a Network Control Experience Management (NCEM) framework for intelligent communication networks. The framework is built upon a novel Network Control Experience State that jointly characterizes accumulated experience, experience age, relevance, evolution rate, and utility. Using this representation, we develop a state-space model for experience evolution and introduce the concepts of Experience Equilibria, Experience Operating Regions, and Experience Policies. These concepts provide a systematic foundation for managing network experience as an operational resource. The proposed framework is evaluated through representative numerical examples, including utility-aware management, disturbance recovery, and quantum-network operation. The results demonstrate that state-aware experience management provides improved adaptability and robustness compared with approaches based on individual freshness or utility metrics. The presented work establishes a foundation for experience-centric management of intelligent communication networks and offers a new perspective on the role of accumulated experience in future autonomous and adaptive network systems.

G7b Experience-State Management: A Mathematical Framework for Adaptive Information Acquisition in Intelligent Communication Networks

Autonomous management of future quantum networks requires timely and accurate knowledge of dynamically evolving network conditions. Unlike conventional communication networks, quantum networks are characterized by stochastic entanglement generation, decoherence, quantum memory limitations, and rapidly changing resource availability, making complete and continuous environmental awareness both costly and impractical. Existing quantum network management approaches generally assume that sufficiently accurate network state information is available prior to decision making and therefore do not explicitly address the acquisition of environmental information as part of the management process. This project introduces the Experience-State Management (ESM) framework, a mathematical foundation for adaptive quantum network control that explicitly distinguishes the physical Environment State from the controller's internal Experience State. Environmental information is acquired through adaptive probing, polling, telemetry, and state-request mechanisms, collectively represented by an Information Acquisition Operator governed by an Environment–Experience Coupling coefficient. Rather than assuming complete network knowledge, the proposed framework treats environmental information acquisition itself as an optimization variable, enabling the controller to balance management benefit against acquisition delay, communication overhead, and computational cost. The paper develops the mathematical foundations of the Experience-State framework, formulates the adaptive information acquisition problem, and establishes the basis for topology-aware optimization of environmental information collection in autonomous quantum networks. The proposed formulation provides a unified analytical framework for adaptive quantum network control and establishes a foundation for the design of scalable, self-adaptive quantum network management systems.

G7c Designing Network Intelligence: A State-Space Framework for Intelligent Communication Networks
G7d Modeling the Network Thinking Process for Intelligent Communication Systems
G7e Distributed Experience-State Management for Multi-Domain Quantum Communication Networks