Quantum ECN and Beyond: A Survey on Congestion Control for Distributed Quantum Computing Infrastructures

Chenliang Tian, <chenliang.t@wustl.edu>
(A paper written under the guidance of Prof. Raj Jain)

Download 

Abstract

Quantum Data Centers (QDCs) interconnect multiple quantum computing nodes via quantum networks to enable scalable distributed quantum computing. As these networks grow in complexity and traffic, congestion control emerges as a critical challenge due to the unique fragility of quantum data. Congestion in quantum networks causes not just delays but quantum decoherence – the decay of qubit states while buffered – which directly erodes communication fidelity . This survey provides a comprehensive overview of congestion control techniques for distributed quantum computing infrastructures. We first motivate the need for specialized strategies by contrasting quantum vs. classical network behavior, then propose a taxonomy of congestion control approaches in QDC networks ranging from static resource allocation and routing optimizations to dynamic feedback-driven protocols. In particular, we highlight Quantum ECN (Explicit Congestion Notification) as a central example of a feedback-based scheme, illustrating how explicit network feedback can modulate source rates to stabilize network load and preserve entanglement fidelity. We compare Quantum ECN with alternative strategies such as fixed-rate throttling, MILP-based scheduling of entanglement distribution, and emerging adaptive routing protocols like QuARC. For each category, we discuss design motivations, trade-offs between throughput and fidelity, and applicability in quantum data center environments. We further summarize state-of-the-art results, identify open challenges – including multi-application fairness, scalability, and cross-layer integration with quantum repeaters – and outline promising directions for future research. By synthesizing recent advances, this survey aims to clarify the landscape of quantum network congestion control and guide the development of robust communication protocols for next-generation QDC infrastructures.

Keywords: Quantum data center networks; Congestion control; Quantum ECN; Entanglement distribution; Quantum networks; Distributed quantum computing; Network fidelity; Explicit congestion notification.

Table of Contents:

- 1. Introduction
- 2. Background and Motivation
  - 2.1 Quantum Data Center Networks and Entanglement Distribution
  - 2.2 Congestion in Quantum Networks: Impacts and Detection
- 3. Taxonomy of Congestion Control Techniques in QDCs
  - 3.1 Static Resource Reservation and Open-Loop Control
  - 3.2 Adaptive Routing and Network-Layer Congestion Mitigation
  - 3.3 Feedback-Driven Congestion Control
- 4. QECN Comparative Evaluation and State-of-the-Art Summary
- 5. Open Challenges and Future Directions
- 6. Conclusion
- References
- List of Acronyms

1. Introduction

Quantum computing is rapidly progressing, yet near-term quantum processors remain limited in qubit count and coherence time. To scale beyond a single device's capabilities, the concept of Quantum Data Centers (QDCs) has been proposed: clusters of small quantum computers connected by quantum communication links to perform distributed quantum processing. By networking quantum nodes, a QDC can facilitate tasks like distributed quantum algorithms and entanglement-based protocols that surpass what isolated machines can do. Achieving this vision, however, requires a supporting quantum network infrastructure that can reliably distribute quantum entanglement between nodes at scale.

A fundamental challenge in quantum networks, especially in dense QDC interconnects, is congestion control. In classical data networks, congestion leads to packet queues building up, increased latency, and eventually packet loss. In quantum networks, the situation is more dire since qubits cannot be buffered for long due to decoherence (the gradual loss of quantum state fidelity over time). When many quantum nodes simultaneously request entangled qubits, intermediate quantum switches (quantum repeaters or routers) can become congested with qubits waiting in memory. Unlike classical bits in a router's buffer, qubits cannot be stored indefinitely without degradation – their quantum state decays exponentially with storage time. Thus, network congestion in a QDC does not just cause delay; it directly erodes the quality of the distributed

quantum state. If a quantum entangled pair is delayed too long in a memory buffer, it may decohere to the point of being unusable. Severe congestion can lead to wasted entanglement resources (qubits that decohere before they can be used) and a collapse in end-to-end application fidelity and throughput. This is a stark contrast to classical networks, where a packet can, in principle, be queued arbitrarily long and still be delivered intact. The fragility of quantum data makes effective congestion control a critical requirement for any large-scale QDC network.

Early research efforts have begun to tackle these issues by drawing inspiration from classical networking while accounting for quantum-specific constraints. For example, Dahlberg et al. [Dahlberg19] introduced a quantum network protocol stack analogous to the classical TCP/IP layers, including a link layer for reliable entanglement distribution. Van Meter et al. [VanMeter16] proposed high-level architectures for quantum networks (e.g., a software-defined quantum internet) to manage entanglement routing and resource allocation. Perhaps most relevantly, Zhao et al. [Zhao23] developed some of the first transport-layer protocols for quantum data networks, inspired by classical TCP congestion control principles. These pioneering works laid the groundwork for quantum network protocol design. However, most early approaches relied on static resource reservation or scheduling – for instance, dedicating certain quantum channels or time slots to each flow – which is analogous to circuit switching. Such reservation-based schemes can prevent buffer overflow, but they conflict with the best-effort, packet-switched service model that has made the classical Internet so flexible and efficient. In other words, prior quantum network protocols often assumed pre-planned allocations of quantum links or memory for each communication session, an approach that may not scale or adapt well to dynamic network conditions.

Recognizing these limitations, researchers have begun exploring feedback-driven, adaptive congestion control techniques tailored for quantum networks. One notable development – and a central focus of this survey – is the concept of Quantum ECN (Explicit Congestion Notification), a framework in which the network provides direct feedback to senders about incipient congestion, enabling senders to throttle their quantum traffic before the network becomes overburdened. Quantum ECN draws an analogy to ECN in classical IP networks: rather than waiting for qubits to be dropped or irreversibly decohered as a signal of congestion, quantum switches actively mark or signal congestion early (using the accompanying classical channel) to notify end nodes. By forming such a closed control loop between the quantum network and end hosts, it becomes possible to stabilize network load below critical thresholds and maintain a high fidelity of transmitted quantum states.

In this survey, we synthesize and compare the diverse strategies proposed for congestion control in QDC networks. We categorize approaches into broad families, from simple static rate-limiting policies and routing optimizations, to sophisticated dynamic protocols with explicit feedback, and analyze how each addresses the throughput–fidelity trade-offs inherent in quantum communication. Quantum ECN will serve as a running example of a feedback-based approach, which we contrast with other methods like fixed-rate (open-loop) control, Mixed-Integer Linear Programming (MILP) based entanglement routing,

and novel adaptive schemes such as QuARC (Quantum Adaptive Routing using Clusters). Throughout, we highlight design motivations and assumptions (e.g., whether a scheme assumes synchronized operations or dedicated resources), discuss performance results reported in the literature, and identify the key trade-offs (for example, maximizing raw entanglement throughput versus ensuring high end-to-end fidelity or fairness).

The remainder of this paper is organized as follows. Section 2 provides background on quantum data center networks, explaining the network model and the unique relationship between congestion and quantum coherence. We also discuss the motivation for borrowing and adapting classical congestion control techniques to the quantum domain. Section 3 presents a taxonomy of congestion control techniques in QDC infrastructures, classifying existing approaches into categories and outlining their core principles. Section 4 delves deeper into feedback-based approaches, with Quantum ECN as a case study to illustrate how explicit feedback loops can be implemented in quantum networks. In Section 5, we provide a comparative evaluation of the approaches and summarize the state-of-the-art, including recent proposals and their performance characteristics. Section 6 identifies open challenges and potential directions for future research, such as handling heterogeneous quantum hardware, multi-application networks, and integration with quantum error correction. Finally, Section 7 concludes the survey.

2. Background and Motivation

2.1 Quantum Data Center Networks and Entanglement Distribution

A QDC network consists of multiple quantum nodes, including quantum processors or small quantum computers, connected via quantum communication links (e.g., optical fiber) within a data center or local area. The goal is to enable these nodes to engage in distributed quantum algorithms by sharing quantum entanglement on demand. In practice, when one node needs to send quantum information or coordinate with another, it will request the generation of an entangled qubit pair with the target node. Entanglement is the enabling resource for protocols like quantum teleportation, by which the actual quantum data (an arbitrary qubit state) can be transferred from one node to another using a shared entangled pair and classical messaging.

Entanglement distribution in a multi-hop quantum network is typically achieved via a process called entanglement swapping. Nearest-neighbor nodes can generate entangled pairs over a single link, often called direct-link entanglement generation. To extend entanglement to distant nodes, intermediate nodes, also called quantum, perform entanglement swapping. For example, as illustrated in Fig. 1, if a node (Charlie) holds one qubit entangled with Alice and another qubit entangled with Bob, a joint measurement on Charlie's two qubits can create an entangled link between Alice and Bob. By chaining such swaps through multiple intermediate nodes, end-to-end entanglement between far-apart end nodes can be established. The QDC network thus resembles a multi-hop system where qubits, often called flying qubits when traveling through fiber as photons, hop from link to link and are temporarily stored in quantum memory at intermediate hops until the next link is ready.

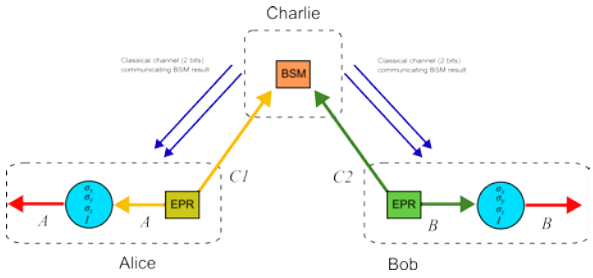


Fig. 1. Conceptual illustration of entanglement swapping between distant nodes via an intermediate repeater (Source:Wikimedia Commons).

Quantum memories are a key component in this architecture. We suppose each network node is equipped with a number of memory qubits that can store quantum states for a limited duration. Memory is needed because entanglement generation on each link is probabilistic due to losses and quantum measurement uncertainties, and the network must wait until all direct-link pairs along a path are ready before performing swaps. During this waiting time, half of each entangled pair resides in a memory at a node, queued until the subsequent operations occur. Crucially, quantum memories are volatile. Each has a finite coherence time, which might range from milliseconds to seconds depending on the technology, beyond which the stored qubit's state becomes highly error-prone or essentially random. In practical terms, if a qubit is buffered in memory for time , its entanglement fidelity decays roughly as under typical noise models, where is the coherence time. Unlike a classical packet, a quantum "packet" cannot be stored in a long queue without losing the information it carries. This fundamental constraint underpins why congestion is so harmful in quantum networks.

Another fundamental difference is the lack of amplification or copying in quantum communication. Due to the no-cloning theorem, one cannot simply copy a qubit to avoid loss or create backups. All quantum communication is effectively analog and lossy. If a qubit is lost or decoheres, it cannot be perfectly reconstructed. Entanglement distribution attempts can be retried (since we can always attempt to generate a fresh entangled pair), but any qubit that gets dropped or degraded is lost work. This also implies that quantum networks lean towards a send-and-wait (or stop-and-go) paradigm rather than blasting a large backlog of unacknowledged qubits, since lost qubits cannot be blindly resent without coordination.

Every quantum entanglement distribution is accompanied by a parallel classical communication channel. In fact, a QDC network runs a classical control network in parallel to the quantum links. Adjacent quantum nodes (and indeed all end nodes) are assumed to be connected via classical links (which could be standard Ethernet or fiber) to coordinate quantum operations. The classical network is used to send signals, such as entanglement generation requests and outcomes, acknowledgments that an end-to-end entangled pair has been established, and other control messages (e.g., which memory slot is storing which qubit). Importantly for congestion control, the classical channel can also carry congestion signals from network nodes back to senders. Classical messages travel much faster than typical quantum decoherence times (e.g., microsecond-order latency for classical packets across a data center, versus milliseconds of coherence time).

Therefore, using classical communication to provide quick feedback about network conditions is a viable strategy in quantum networks. It allows potentially reacting to incipient congestion before quantum states expire.

2.2 Congestion in Quantum Networks: Impacts and Detection

In a classical network, congestion is usually detected indirectly, for example, packet loss or increasing round-trip times. In quantum networks, waiting for packet (qubit) loss as a congestion signal is often too late. By the time a qubit is dropped from an overwhelmed quantum buffer or by the time a sender realizes that no acknowledgment was received, implying a qubit was lost, the network may have already suffered significant decoherence across many stored qubits. Moreover, deliberately dropping qubits to signal congestion, analogous to packet drop in TCP, is highly undesirable. Because generating entanglement consumes precious resources and time, and dropping a qubit means an irrecoverable loss of a quantum resource. Unlike a classical packet, which can be resent from the source if lost, an entangled qubit lost to decoherence or drop cannot be recovered except by starting the entire entanglement generation process anew. This motivates proactive congestion signaling to inform senders about congestion before qubits are lost.

Congestion in a quantum switch manifests as a growing memory queue of qubits awaiting entanglement swapping. As discussed, the longer a qubit stays in this queue, the lower its fidelity becomes. Therefore, one intuitive measure of quantum congestion is the time qubits spend in memory. A related measure is memory occupancy, which means how many qubits are stored. A higher occupancy suggests longer waiting times and a greater risk of decoherence. In essence, quantum network congestion can be defined as the network operating in a regime where the demand for entanglement distribution exceeds what the network can service within the coherence time limits of the memories. When that happens, qubits pile up and decohere, dramatically reducing the effective throughput of useful entangled pairs delivered to applications.

To control congestion, network protocols must either prevent such backlogs or rapidly mitigate them when they occur. Broadly, there are two approaches to this: (1) avoidance via resource management (scheduling, routing, admission control) so that qubits are not injected into the network faster than they can be processed, and (2) feedback-driven adjustment where senders dynamically adjust based on network signals (like how TCP adapts its window in response to congestion indications). In quantum networks, both approaches are being explored.

The table below contrasts how congestion manifests and is handled in classical packet networks versus quantum data center networks.

Table 1. Classical vs. Quantum Congestion Characteristics

Aspect	Classical networks	Quantum data center networks
Primary congestion signal	Packet loss, increasing RTT, queue length	Memory occupancy, qubit waiting time, fidelity decay
Buffering behavior	Packets can be queued almost arbitrarily long	Qubits cannot remain in memory beyond coherence time
Cost of loss	Lost packets can be retransmitted from the source	Lost or decohered qubits require re-running entanglement generation
Control objective	Maximize throughput while avoiding persistent queue overflow	Maximize useful entanglement rate subject to fidelity constraints
Typical feedback mechanism	Implicit (loss/RTT) or explicit (ECN bits) congestion feedback	Classical control messages carrying Quantum ECN marks or similar signals
Resource coupling	Mainly link capacity and buffer space	Coupled link capacity, quantum memory, and coherence time

On the avoidance side, one may attempt to schedule entanglement requests in a way that minimizes overlapping usage of the same quantum channels or nodes. For example, if two pairs of nodes need entanglement and their routes would intersect at a common switch, scheduling one after the other (or choosing alternate paths) could avoid simultaneous contention. Recent research by Verzotti et al. [Verzotti23] analyzes simple scenarios of intersecting entanglement routes and proposes strategies to serve them sequentially (a “symmetrical” approach where neither flow goes simultaneously through the bottleneck) or in a staggered parallel fashion (“asymmetrical”), depending on success probabilities. This kind of entanglement request scheduling or routing can mitigate congestion by design, but it typically requires a priori knowledge of all requests or centralized control, which may not be feasible in a dynamic QDC environment.

On the feedback side, one leverages the classical channel to create a closed-loop system. A quantum switch experiencing high occupancy might send a signal (explicitly or implicitly) to sources to slow down further entanglement attempts. The challenge is to do so in a timely manner, within the coherence time window. This is where the concept of ECN is adapted to quantum networks. In classical networks, ECN marks packets by setting a bit in the header when queues grow too long, and the receiver echoes this mark to the sender to signal congestion without losing the packet. Quantum ECN works analogously. A quantum repeater, instead of dropping qubits, marks an indicator (potentially in a classical control message associated with the qubit or an auxiliary qubit flag) when its memory occupancy or delay exceeds a threshold. The receiver or the network itself then notifies the source, which can then adjust its sending rate before qubits start getting lost to decoherence. This approach aims for lossless congestion feedback, which is especially valuable given the cost of losing qubits.

The motivation for Quantum ECN and similar feedback schemes is clear: by reacting early, the network can be kept in a regime of high throughput and high fidelity. A well-designed congestion control will ideally find a sweet spot where throughput of entangled pairs is maximized subject to maintaining fidelity above the threshold needed for quantum applications. Some applications might tolerate lower fidelity or can purify entanglement afterward for higher throughput, while others require high fidelity above all. Congestion control in QDCs must balance this trade-off according to the needs of the application.

Finally, it is worth noting that the design of congestion control in quantum networks must also consider whether the network is operating in a connection-oriented or connectionless manner. Some quantum transport protocols (like those by Zhao et al. [Zhao23]) assume a connection setup phase and synchronized operation, effectively treating the network in a circuit-switched fashion. In these designs, once a connection is established between two end nodes, resources along the path might be reserved, and the transmission of qubits is done in discrete time slots across the path to coordinate entanglement generation and swapping. This can simplify congestion management as it resembles scheduling in a circuit network, but it reduces flexibility. Other approaches, including Quantum ECN as conceived in recent work, envision a more packet-switched quantum network where entanglement distribution requests are treated like packets flowing through the network, possibly taking different paths or being generated opportunistically

without rigid scheduling. In such a best-effort quantum network, dynamic congestion control akin to Internet routing and transport protocols becomes essential.

In summary, quantum data center networks demand new congestion control strategies because of the physical constraints, like no long queues, no cloning, probabilistic link outcomes, and maintaining quantum state fidelity for application demands. With this background, we now turn to classifying the congestion control techniques that have been proposed for QDC and similar quantum networking scenarios.

3. Taxonomy of Congestion Control Techniques in QDCs

Designing congestion control for QDC infrastructures is a nascent area of research, and multiple complementary strategies have been proposed. Broadly, we can categorize existing approaches along two key axes: Open-loop vs. Closed-loop control and Resource reservation vs. Best-effort service. These axes are not entirely independent, but they help distinguish the philosophies behind different solutions. We identify the major categories of congestion control techniques in quantum networks in Sections 3.1 to 3.3. For clarity of exposition, we discuss each category separately with representative examples.

3.1 Static Resource Reservation and Open-Loop Control

Static rate control is the simplest strategy: network designers impose an upper limit on how fast quantum entanglement requests can be generated by each node or each end-to-end session, such that even worst-case traffic patterns won't overflow quantum memories. For example, if each link or switch can reliably handle at most X entangled pairs per second without backlog, sources might be throttled to send at or below this rate. This approach is analogous to traffic shaping in classical networks. Its advantage is predictability – if the parameters are chosen conservatively, congestion (and hence decoherence loss) should never occur. Early quantum network experiments often effectively used static or duty-cycled operation to avoid overwhelming the system. However, static control has major drawbacks: it is typically overly conservative (under-utilizing the network when conditions are good or when not all sources are active) and inflexible (the optimal fixed rate depends on network conditions that can vary or are uncertain, such as link success probabilities or current traffic load).

A more sophisticated open-loop approach uses global scheduling or optimization to allocate resources to entanglement requests. Researchers have formulated the entanglement distribution problem as a routing and scheduling problem – given a set of requested entangled connections between node pairs, decide which path each should take and at what times to attempt each link-level entanglement so as to maximize throughput or minimize completion time. This can be cast as a Mixed-Integer Linear Program (MILP) or other combinatorial optimization. For instance, Shi and Qian (2020) presented a model for concurrent entanglement routing where multiple entanglement

flows compete for resources, and they designed algorithms (including ILP formulations and heuristics) to efficiently route entanglements with minimal conflict. Similarly, Li et al. (2023) [Li23] considered a congestion mitigation approach that explicitly avoids having intersecting entanglement routes contend at the same time by careful selection of swapping order and path assignment. In their work, a novel entanglement swapping scheme is proposed to reduce congestion at critical nodes, and they use an optimization framework to improve the quantum request service rate compared to naive or greedy approaches.

The benefit of such optimization-based or reservation schemes is that, if the traffic matrix is known and static (or changes slowly), the network can operate near its capacity with minimal contention. By scheduling entanglement generation in non-overlapping time windows or along disjoint routes, one can largely avoid having too many qubits waiting in any one memory. This approach mirrors circuit-switched network design, where you might reserve a fixed bandwidth circuit for each flow. Indeed, as mentioned, some researchers argue that quantum networks will inherently be more circuit-like due to the need for synchronous operations. The Teleportation-based QDN model of Zhao and Qiao is one example: they establish an end-to-end connection and then send qubits in time slots, essentially preventing intermediate queue build-up by design. In such a model, once a connection is set up, each timeslot sees one entanglement attempt on each link in the path and an immediate swap, so qubits don't accumulate – if a timeslot's attempt fails, you simply try again in the next slot.

However, the limitations of static and scheduled approaches include lack of adaptability and potential inefficiency under dynamic workloads. QDC environments might see bursty demand (e.g. a particular distributed algorithm might suddenly request many entangled pairs, then be quiet). A fixed schedule or a once-computed plan can become suboptimal or even invalid if some requests finish early or new ones arrive. Moreover, solving a MILP for a large network in real-time is computationally hard; thus, such methods are better suited for planned network usage (or as an upper-bound benchmark for performance) rather than a continuous control mechanism. They also often assume accurate knowledge of success probabilities and traffic demands. If these assumptions are wrong (e.g. actual link fidelity is worse, or an unexpected additional flow appears), the network could still become congested because the plan did not account for it.

In summary, open-loop congestion control strategies in quantum networks either throttle usage to a safe fixed level or allocate resources in advance to avoid simultaneous heavy use of any one component. They trade efficiency for safety, and flexibility for predictability. These methods are important as baselines and may be combined with dynamic methods (for instance, one could use a MILP to get an initial resource allocation, then use feedback-based adjustments on top of it to handle variations). Next, we look at adaptive strategies that operate at the network layer, adjusting how entanglement is routed in response to network conditions.

3.2 Adaptive Routing and Network-Layer Congestion Mitigation

Because quantum networks are fundamentally constrained by probabilistic entanglement generation and distance-dependent losses, routing plays a crucial role in performance. Adaptive routing protocols aim to improve throughput and avoid congestion by intelligently spreading entanglement generation attempts in the network. Rather than treating the path as fixed and controlling the rate, these protocols may change which path (or paths) are used, or even alter the network's effective topology over time to balance load.

A prime example is QuARC (Quantum Adaptive Routing using Clusters) by Clayton et al. [Clayton24]. QuARC takes a novel approach: it periodically partitions the network into clusters of nodes, each of which internally generates entanglement among its members. Then, entanglement between clusters is established by connecting cluster nodes. The clustering is adaptive – QuARC can merge or split clusters over time, based on current network conditions and demand patterns . The core idea is to create a dynamic hierarchy that can better utilize parallelism: within each cluster, multiple entanglements can be generated concurrently, and between clusters, QuARC leverages multi-path fusion techniques (multi-particle measurements) to establish long links. By changing cluster size, QuARC effectively changes the "locality" of entanglement distribution: smaller clusters mean shorter paths (fewer hops) within each cluster but more hops between clusters, whereas larger clusters handle longer paths internally. This adaptability allows QuARC to strike a balance between two extremes – pure multi-path protocols that achieve high throughput under ideal conditions but can starve long-distance requests, and pure long-path (many swaps) protocols that ensure reach but often yield low rates . Simulation results show QuARC is robust to changes in physical parameters and maintains high overall throughput without starving far-end pairs, even as network size scales . In effect, QuARC mitigates congestion by localizing traffic: entanglement attempts largely occur within clusters, reducing the load on any single network-wide bottleneck, and by dynamically reconfiguring clusters, it prevents persistent hot spots .

Another approach to adaptive routing is to use multi-path entanglement distribution. Some protocols send entanglement concurrently over multiple disjoint paths, with the hope that at least one succeeds or that they can be combined (via entanglement fusion) into a higher-fidelity connection . Multi-path routing can increase throughput and reliability, but if not managed, it can also increase resource usage and potentially congest more parts of the network. Adaptive algorithms may monitor which paths are yielding entanglements successfully and concentrate efforts on the more productive paths. In principle, one could also envision load-aware routing where each repeater or switch can decide to forward entanglement attempts along a less busy route if its primary route is becoming congested (similar to how a packet network can reroute traffic if a link is saturated). This requires some way of measuring "quantum traffic" on different links – possibly by tracking memory usage or entanglement wait times.

It is worth noting that because quantum entanglement distribution often requires coordination across an entire path (to perform swaps in sequence or within a deadline before decoherence), routing decisions are not as simple as in classical hop-by-hop

packet routing. However, proposals like software-defined quantum networking suggest that a central controller (or distributed controllers with a global view) could manage entanglement routes in a more dynamic fashion than a static circuit setup. In such a scenario, if one route is experiencing low fidelity due to congestion (too many swaps waiting), the controller could reroute subsequent entanglement attempts via an alternate path that is currently more congested.

In summary, adaptive routing and clustering approaches address congestion by reconfiguring the flow of quantum information in the network. They complement end-to-end traffic control by ensuring that the network topology or path selection does not itself become a bottleneck. These strategies can significantly improve fairness and throughput under varying conditions (for example, QuARC explicitly tackles the fairness issue where long-hop requests were starved by shorter-hop ones in earlier protocols). The trade-off is additional complexity in network management and the need for timely knowledge of network state (e.g., which links are losing fidelity or have long queues). Many of these proposals are still at the simulation stage, but they highlight the potential of network-layer congestion-awareness in quantum networks.

3.3 Feedback-Driven Congestion Control

The third category comprises closed-loop congestion control algorithms that closely resemble those in classical data networking, in spirit if not in detail. These protocols rely on continuous monitoring of network conditions and use feedback signals to adjust the rate of quantum entanglement requests issued by sources.

Quantum ECN (Explicit Congestion Notification) stands out as a representative of this class. The idea behind Quantum ECN is to introduce an explicit feedback bit (or a few bits) into the control plane of the quantum network that alerts senders to congestion without waiting for qubit loss. Concretely, a quantum switch could monitor its memory occupancy and, if it exceeds a threshold or if qubits have been stored for more than a certain time, the switch marks the next outgoing entanglement (or sends a separate classical message) indicating congestion. This mark is delivered to the destination node along with the entangled qubit (or via the classical ACK), and then relayed to the source. The source, upon learning of congestion, reduces its sending rate – for example, by reducing the window of concurrent entanglement attempts or increasing the delay between new requests. When congestion is not signaled for a while, the source can cautiously increase its rate again. This behavior is analogous to TCP's Additive-Increase Multiplicative-Decrease (AIMD) algorithm, and indeed one can design the quantum source's controller to use AIMD with quantum-aware parameters. The "multiplicative decrease" would be triggered by a congestion mark (instead of a packet loss as in TCP), and "additive increase" would slowly raise the rate otherwise, probing for available capacity.

One key aspect of Quantum ECN is the congestion marking policy at the switches. Classical ECN routers often use Active Queue Management (AQM) like Random Early Detection (RED) to decide when to mark packets – e.g. probabilistically mark packets

when the average queue length exceeds a threshold. In a quantum switch, the policy must be tuned to protect qubits from decoherence. For instance, a switch might mark a fraction of qubits once the average qubit waiting time goes beyond, say, 50% of the coherence time, thereby preemptively signaling the network is nearing a dangerous regime . By doing so, Quantum ECN aims to keep the system's operating point below the "cliff" where fidelity rapidly collapses. Analytical evaluations of such a scheme have shown that it can indeed stabilize the network: sources react to congestion marks long before buffers overflow, thus the quantum memories are utilized but not overstuffed . The result is a network that carries as many entangled pairs as possible while maintaining their fidelity above a target threshold, effectively providing a throughput–fidelity knob for operators . This tunability is important; different QDC applications might prefer different operating points, and a feedback mechanism like ECN can be adjusted (via the marking threshold or the AIMD aggressiveness) to meet those requirements .

Another set of feedback-based approaches are the Quantum Transport Protocols (QTP) introduced by Zhao and Qiao. In their work, two protocols (Tele-QTP and TAG-QTP) were designed for reliable end-to-end qubit delivery in a QDN. These protocols manage quantum memory at both end hosts and intermediate nodes to avoid congestion and ensure reliability. Although they do not use an "ECN bit" per se, they implement a window-based flow control. For example, Tele-QTP might allocate a certain number of memory slots for a given connection along each link; the source can only have that many entanglement attempts in flight. If intermediate memories fill up, the protocol stops launching new qubits until an old one is consumed (acknowledged) – this behaves like a backpressure mechanism, ensuring no memory overflow. The authors describe their approach as achieving "congestion-free and reliable delivery by managing the limited quantum memory" at all nodes. In effect, QTPs maintain a form of distributed congestion control: since every qubit transmission requires an available memory at the next hop, the rate is self-regulated by whether downstream nodes are ready to accept more qubits. This is somewhat analogous to how TCP's sliding window prevents overrunning the receiver's buffer, combined with network-level flow control to not overload routers. Simulations of these protocols showed they can achieve high throughput and fairness among flows, confirming the effectiveness of coordinated window management.

Beyond these, researchers have started applying control theory and optimization to quantum network congestion control. Vardoyan and Wehner (2023) formulated a Quantum Network Utility Maximization (QNUM) problem, akin to the classical NUM framework, where each flow's rate is adjusted to maximize a global utility function under network capacity constraints. In the quantum case, the utility may be defined in terms of entanglement fidelity and rate. They derived a primal-dual algorithm called QPrimalDual that updates source rates and internal "link prices" (congestion signals) iteratively. Remarkably, their design embeds the congestion information into the quantum packet (qdatagram) header itself: as an entanglement is established across a path, each link adds its "price" (which could be a function of memory occupancy or recent success rate) to the header, and the accumulated value is returned to the source in the acknowledgment. The source then adjusts its sending rate based on this total congestion price – a high price means many congested links, so it reduces the rate, and a low price

means the network can handle more. This approach is analogous to explicit rate feedback in some classical systems (e.g., XCP or ATM ABR flow control), and it promises fair resource allocation in a distributed manner. Stability analyses indicate that, under certain conditions, such algorithms converge to an equilibrium where each flow's rate is optimal with respect to a trade-off between throughput and the "cost" (fidelity reduction) of using congested resources. While this work is theoretical, it represents an exciting crossover of network control theory into the quantum domain, suggesting that even sophisticated congestion control techniques (multi-bit feedback, optimization-based control) can be translated into the quantum context with appropriate modifications.

In summary, feedback-driven congestion control in quantum networks, exemplified by Quantum ECN, QTPs, and utility-maximization algorithms, is a powerful approach to maintain performance in the face of unpredictable demand and network conditions. These methods allow the network to self-regulate, pushing as hard as it can but backing off when nearing physical limits. The challenges in implementing them include ensuring the feedback loop is fast enough (classical signals must propagate and be processed before quantum states decohere) and coping with the probabilistic nature of quantum events (which introduces more noise into the measurement of congestion signals). Nonetheless, experiments and simulations so far are promising: for instance, a recent simulation of a best-effort quantum network architecture using AIMD congestion control and AQM showed the ability to tune end-to-end fidelity vs. throughput via simple controller parameters. In the next section, we compare these approaches head-to-head and summarize the current state-of-the-art, including what has been demonstrated quantitatively about their performance.

4. QECN Comparative Evaluation and State-of-the-Art Summary

Having outlined the main categories of congestion control techniques for QDC networks, we now compare their characteristics and performance as reported in the literature. It is important to note that since this is an emerging field, many proposals have been evaluated only via analysis or simulation on small network models. Direct apples-to-apples comparison is challenging due to differing assumptions. Nevertheless, we can discuss the relative merits of each approach and highlight representative results.

Static vs. Dynamic Control: Static reservation or rate-control schemes (Category 3.1) guarantee that congestion (in the sense of memory overflow) should not occur, provided the worst-case assumptions hold. They inherently maintain high fidelity because qubits are not allowed to queue long. For example, if each source in a QDC is limited to one entanglement attempt at a time, and each entanglement is routed sequentially hop-by-hop, the network is effectively in a circuit-switched mode; fidelity remains high, but throughput is obviously limited. On the other hand, dynamic feedback schemes (Category 3.3) aim to approach the network's capacity by utilizing all available memory until just before decoherence becomes problematic. Studies have shown that a Quantum ECN approach can dramatically improve throughput over static policies while still avoiding a

fidelity collapse. In an analytical evaluation, a Quantum ECN controller stabilized a network that would otherwise suffer >50% decoherence loss under heavy load, achieving near 0% qubit loss with only a modest reduction in raw throughput compared to an uncontrolled flood. This indicates that feedback control can safely push the utilization higher than any fixed static rate could, because it adapts to moment-by-moment conditions.

Throughput–Fidelity Trade-off: One can think of each approach as choosing a point on the throughput vs. fidelity curve. Static conservative schemes pick a single safe point (often far on the fidelity side, low throughput). MILP-based scheduling tries to hit the optimal point by calculating an allocation (often maximizing throughput while constraining fidelity). Feedback methods like QECN allow the operating point to vary: for instance, by tuning the marking threshold, network operators can effectively decide to run the network "hotter" (higher throughput, lower per-pair fidelity) or "cooler" (lower throughput, higher fidelity). Bacciottini et al. (2024) [Bacciottini24] demonstrated via simulation that by adjusting AQM parameters, they could move the network's operating point smoothly between a high-fidelity/low-throughput regime and a lower-fidelity/high-throughput regime, thus providing a controllable knob to meet different application requirements. In contrast, adaptive routing (Category 3.2) mostly tries to improve both throughput and fidelity by better using network resources – for example, QuARC achieved both high total throughput and fairness (serving long distances without severe fidelity drop) by reorganizing the network adaptively. So, routing schemes can sometimes shift the entire curve outwards (improving performance across the board) at the cost of more complex operations.

Fairness and Starvation: A critical aspect in multi-user networks is how resources are shared. Static slicing of resources (each flow gets a fixed portion) is inherently fair by design, but can waste capacity if some flows are inactive. Feedback-based congestion control, like AIMD, tends to converge to fairness in the classical sense (equal share of bottleneck capacity) under symmetric conditions. The quantum transport protocols (Tele-QTP, etc.) also explicitly demonstrated fairness in their results – by managing memory per flow, they prevent any single flow from monopolizing the network. QuARC's primary motivation was to prevent starvation of longer routes in the presence of shorter ones. By clustering, QuARC ensured that long-hop requests still got served with reasonable throughput, whereas many prior routing algorithms would see near-zero success for long paths once short paths saturated the network. MILP optimizations can include fairness constraints as well, although early works often focused on maximizing total throughput (which can favor shorter or easier requests).

Complexity and Scalability: There is a spectrum of algorithmic complexity. MILP-based and other centralized scheduling methods can become intractable as the network grows (the problem of optimal entanglement routing is NP-hard). They might only be viable for small networks or as offline planning tools. In contrast, distributed feedback algorithms (like ECN or backpressure) scale more gracefully, as each node makes local decisions or simple calculations. However, they introduce overhead in terms of classical messaging (e.g., each entanglement attempt might generate a congestion mark message). The

volume of such overhead is typically small compared to the quantum operations – for instance, a few bits of ECN in an ACK for each entangled pair is negligible given that establishing an entangled pair might involve many photon transmissions. Simulations by Zhao et al. [Zhao23] indicated that their protocols operated efficiently in networks of a few dozen nodes, but scaling to hundreds may require hierarchical or more advanced control (an open question). QuARC's clustering algorithm introduces a different kind of overhead: periodically, nodes must perform cluster reconfiguration (which likely involves classical coordination and possibly some qubit overhead for cluster-boundary entanglements). The authors showed that even with these overheads, the system maintained high performance, but detailed scaling behavior was not fully explored.

Robustness to Physical Variability: Quantum network conditions (like link success probability, memory coherence time) can vary or be uncertain. One advantage of adaptive and feedback methods is robustness – they can adjust to whatever the current condition is. QuARC, for example, was tested under various physical parameter degradations and maintained good performance without tuning, thanks to its adaptive clustering. Similarly, an AIMD congestion controller will naturally send fewer qubits if losses are high or memory is short, as it will encounter congestion sooner. A static plan, by contrast, might be overly optimistic or pessimistic if the actual conditions differ from the assumed ones. This robustness is a significant argument in favor of dynamic congestion control in real quantum networks, which will inevitably face uncertainty and time-varying noise.

To summarize the state-of-the-art, Feedback-driven congestion control (Quantum ECN and similar) has shown the ability to significantly improve network utilization while protecting quantum data fidelity. Transport-level protocols (QTPs) have proven that classical-inspired flow control can enable reliable and fair quantum communication over multiple hops. Adaptive routing schemes (QuARC and others) address congestion from the topology side, ensuring even difficult communication pairs get service and overall throughput remains high. Optimization approaches provide benchmarks and, in some cases, practical algorithms for offline or slow-timescale resource management. The comparative table below conceptually contrasts these approaches on key metrics such as throughput, fidelity, fairness, required knowledge, and adaptability. While no single solution is universally best, the general consensus is that best-effort packet-switched quantum networking is feasible and even advantageous, as long as an effective congestion control mechanism is in place to deal with the quantum-specific challenges. The current frontier involves integrating these ideas and testing them in more realistic simulations or small-scale quantum network testbeds, which leads us to the discussion of open challenges.

Table 2. Comparative Summary of Congestion Control Approaches

Approach category	Example schemes	Control type	Primary objective	Strengths	Limitations
Static reservation / open-loop	Duty-cycled operation, MILP-based routing and scheduling	Pre-computed, open-loop	Avoid congestion by conservative admission and scheduling	Predictable, simple to reason about, high fidelity when assumptions hold	Overly conservative, hard to adapt to bursty or unknown traffic
Adaptive routing and clustering	QuARC, multi-path routing with fusion	Dynamic, topology-aware	Balance load and avoid hot spots by changing paths/topology	Improves fairness and throughput, robust to parameter changes	Requires network-state information and more complex control logic
Feedback-driven congestion control	Quantum ECN, QTPs, QPrimalDual/QNUM	Closed-loop, rate adaptation	Operate near capacity while meeting fidelity targets	Self-regulating, tunable throughput–fidelity trade-off,	Relies on timely and accurate feedback; behavior under large-scale deploymen

				scalable in principle	t still being studied

5. Open Challenges and Future Directions

Despite significant progress in recent years, congestion control for quantum data center networks remains largely theoretical. There are many open challenges and opportunities for future research:

Multi-Application and End-to-End Quality of Service: Future QDCs may support diverse applications simultaneously – for example, some nodes might be performing distributed quantum computing tasks (which need high-fidelity entanglement), while others run quantum key distribution (which can tolerate lower fidelity but needs steady key rates). How can congestion control differentiate service for different applications? In classical networks, techniques like weighted fair queueing or priority classes are used. In quantum networks, this is largely unexplored. A possible direction is to develop service differentiation mechanisms that allocate quantum network resources according to application needs (some initial work has considered fair sharing policies for multiple quantum applications). Ensuring fairness (no starvation) while also meeting different fidelity requirements is a complex optimization that likely needs new algorithms.

Integration with Quantum Error Correction (QEC): One way to extend quantum memory times and communication distances is to incorporate quantum error correction. QEC consumes entangled pairs in large quantities (for entanglement distillation or as part of repeater protocols) . Congestion control in a QEC-enabled network might look different – rather than individual qubit fidelity, the focus could shift to maintaining a sufficient rate of raw entanglements so that after error correction the net rate meets the target. An open question is how to coordinate congestion control with QEC processes. For instance, if the network is congested, should it favor fewer high-fidelity pairs (maybe to feed directly to applications) or many low-fidelity pairs (to feed into an error-correcting code)? Research is needed on joint rate-fidelity control, treating QEC protocols as part of the congestion control loop.

Scaling to Larger Networks: Thus far, most simulations and analyses consider at most tens of nodes or a few concurrent flows. Scaling up to data-center-sized networks (hundreds of nodes, many simultaneous flows) introduces new issues. Classical congestion control is known to suffer issues at scale (e.g. TCP incast, or synchronization effects). In quantum networks, scale might introduce longer feedback delays (if the network is larger, classical signals take longer to return) and more complex interaction between flows. It remains to be seen whether algorithms like Quantum ECN remain stable and effective as network diameter increases – especially because longer paths mean qubits must be stored longer, tightening the time constraints for feedback. Future work should explore hierarchical congestion control, where local regions of the network manage congestion internally and higher-level protocols manage inter-region traffic

(analogous to traffic management in large IP networks). This ties in with architecture: perhaps QDCs will be organized in pods or clusters with hierarchical control (an approach hinted by QuARC's clustering idea, but that was focused on routing).

Experimental Validation: To date, congestion control proposals for quantum networks mostly rely on simulations with theoretical models of quantum hardware. As experimental quantum network testbeds start to emerge (e.g. metropolitan area entanglement networks), implementing even a simple congestion control scheme would provide invaluable insights. A challenge here is that current hardware might not yet support dynamic traffic – often experiments schedule one entanglement at a time. But as multiplexing technology improves, a near-term goal could be to demonstrate a toy scenario: e.g., two quantum repeaters handling two entanglement flows and using a classical feedback signal to time the second flow's attempts such that they don't interfere too much with the first. Achieving real-time classical feedback within coherence times will be a crucial proof-of-concept for Quantum ECN-like ideas.

Interplay with Link-layer Protocols: Protocols like the one by Dahlberg et al. [Dahlberg19] ensure reliable delivery of single entangled pairs via acknowledgment and possibly re-tries at the link level. How does this interact with congestion control at a higher level? For instance, if a link-layer protocol is aggressively retrying to get an entangled pair (to achieve reliability), it could potentially congest the quantum channel or memory with back-to-back attempts. A congestion-aware link layer might need to back off if the end-to-end layer signals congestion. Conversely, if the transport layer (end-to-end) is slowing down, the link layer might be under-utilized. Designing cross-layer strategies where, say, the link layer provides information (like current link utilization or queue length) to the transport layer could enhance performance. This calls for a co-design of link, network, and transport layers specific to quantum networks, rather than treating each in isolation.

Quantum Routing under Uncertainty: Adaptive routing like QuARC assumes switches can obtain some information about network conditions (even if just local measurements). In a real network, measurement of congestion or link quality may be imperfect. There could be transient fluctuations in loss rates or memory decoherence rates. Robust routing algorithms that can handle such uncertainty without causing oscillations (route flaps) or instability are needed. This is analogous to how classical networks use metrics like the moving average of delay or packet loss to make routing decisions that are somewhat damped to avoid reacting too strongly to short-term bursts.

Standardization and Interoperability: As quantum network research matures, there will be efforts to standardize protocols (indeed, organizations like the IEEE and IETF have started exploratory groups on quantum networking). Congestion control should be part of these standards. One challenge is that different hardware might support different features – not all quantum switches might implement an ECN marking, for example. Some might only support simple blocking (stop sending when full), others might have more fine-grained control. Defining a common set of signals (e.g., a "QECN bit" in the classical control header, or a standard way to report memory occupancy) will be important for interoperability. Future directions include defining an analog of the TCP/IP congestion

control suite for quantum: perhaps a “Quantum Transmission Control Protocol (QTCP)” that incorporates the best ideas (window-based control, ECN, etc.) and can run over various link layers.

Machine Learning for Quantum Congestion Control: An intriguing avenue is whether machine learning could assist congestion control in quantum networks, as has been explored for classical networks. For instance, a reinforcement learning agent could potentially learn policies for sending rates that maximize long-term entanglement throughput-fidelity reward, especially in complex topologies where analytical solutions are hard. Some initial work on deep learning for quantum routing exists, and similar ideas could be applied to flow control. However, the “state” of a quantum network (including qubit states) is not fully observable or easy to feed into an ML model, so this is a challenging but exciting direction.

In summary, while foundational building blocks like Quantum ECN show that congestion control in quantum networks is achievable and beneficial, many practical challenges remain before these ideas can be deployed in a real QDC. The community needs to address cross-layer integration, scaling laws, hardware implementation, and policy questions (fairness/QoS), ensuring that as quantum networks grow, they remain stable and efficient. The lessons from classical networks are invaluable, but must be carefully adapted to respect quantum constraints. The coming years will likely see the first experimental demonstrations of these protocols, and with them, new insights and undoubtedly new problems to solve.

6. Conclusion

Distributed quantum computing promises to overcome the limitations of individual quantum processors by leveraging entanglement across networks of quantum nodes. To unlock this potential, Quantum Data Center infrastructures must deliver entangled qubits reliably and efficiently under increasing demand. This survey has reviewed the emerging area of quantum network congestion control, highlighting why it is both crucial and fundamentally different from classical congestion control. The fragility of quantum information – particularly the decoherence of qubits stored in network buffers – means that conventional “run-until-loss” strategies would lead to unacceptable fidelity loss in a congested QDC network. Instead, quantum networks demand proactive and adaptive control mechanisms that can maintain the delicate balance between throughput and fidelity.

We presented a taxonomy of congestion control techniques for QDCs, covering static resource allocation, adaptive routing, and feedback-driven approaches. We identified Quantum ECN and related feedback schemes as a promising direction to achieve high utilization without sacrificing quantum state quality, by exploiting fast classical signals to orchestrate quantum transmissions. We also discussed alternatives and complements: from MILP-based scheduling that can optimally avoid congestion in theory, to QuARC’s clustering protocol that dynamically reconfigures the network for load balancing, to transport protocols inspired by TCP that manage window sizes and ensure fairness. Each

approach brings valuable ideas to the table, and a full-fledged QDC implementation will likely synthesize multiple techniques. For instance, a QDC might use an initial admission control (to prevent gross overload), an adaptive routing scheme to choose good paths, and a Quantum ECN-based transport to fine-tune the flow of entanglements on each path.

Our comparative analysis underscored that no single method is universally superior – the optimal strategy can depend on network size, quantum hardware parameters, and application requirements. However, a clear theme emerged: introducing flexibility and feedback greatly improves performance and robustness. The ability to react to congestion signals (whether explicit marks, rising "prices," or implicit backpressure) is what allows quantum networks to operate near capacity while preserving entanglement fidelity. This mirrors the experience in classical networking, where adaptive congestion control was key to the Internet's scalability. It is encouraging that initial simulations and theoretical work indicate analogous benefits in the quantum realm.

We have also highlighted many open challenges on the road ahead. Realizing these congestion control schemes in practice will require further innovation in protocol design, cross-layer coordination, and possibly new theoretical tools to account for quantum uncertainties. As quantum hardware improves – with longer coherence times and higher entanglement generation rates – congestion control will shift from a mostly theoretical concern to an immediate engineering issue. It is our hope that this survey helps accelerate progress by consolidating current knowledge and guiding researchers towards the most pressing issues.

In conclusion, quantum congestion control is poised to play for quantum networks a role similar to that of TCP/IP congestion control for the classical Internet: an essential mechanism for shared, efficient, and fair use of network resources. Quantum ECN and the other approaches discussed are laying the groundwork for a future where quantum data centers can deliver on their promise of powerful, networked quantum computing without being bottlenecked by fragile network physics. Continued interdisciplinary efforts – bridging quantum physics, networking, and control theory – will be critical to take these ideas from simulations to functioning networks. The results will not only inform the design of quantum data centers, but also the broader Quantum Internet envisioned to connect quantum devices across the globe. Ensuring congestion is managed in these nascent networks will help ensure that, as quantum technology scales, it does so gracefully and reliably, much like the classical internet before it.

References

[Liu24] J. Liu and L. Jiang, "Quantum Data Center: Perspectives," IEEE Network, vol. 38, no. 5, pp. 160–166, 2024.

[VanMeter16] R. Van Meter and S. J. Devitt, "The path to scalable distributed quantum computing," Computer, vol. 49, no. 9, pp. 31–42, 2016.

[Dahlberg19] A. Dahlberg et al., “A link layer protocol for quantum networks,” in Proc. ACM SIGCOMM, pp. 159–173, 2019.

[VanMeter22] R. Van Meter et al., “A quantum internet architecture,” in Proc. IEEE International Conference on Quantum Computing and Engineering (QCE), pp. 341–352, 2022.

[Zhao23] Y. Zhao and C. Qiao, “Distributed transport protocols for quantum data networks,” IEEE/ACM Trans. Networking, vol. 31, no. 6, pp. 2777–2792, 2023.

[Zhao24] Y. Zhao, Q. Wang, and C. Qiao, “An asynchronous transport protocol for quantum data networks,” IEEE J. Sel. Areas Commun., vol. 42, no. 7, pp. 1885–1899, 2024.

[Bacciottini24] L. Bacciottini et al., “Leveraging internet principles to build a quantum network,” arXiv:2410.08980, 2024 (to appear in IEEE Trans. Quantum Engineering).

[Vardoyan23] G. Vardoyan and S. Wehner, “Quantum network utility maximization,” in Proc. IEEE Int. Conf. Quantum Computing and Engineering (QCE), pp. 1238–1248, 2023.

[Clayton24] C. Clayton, X. Wu, and B. Bhattacharjee, “Efficient Routing on Quantum Networks using Adaptive Clustering,” arXiv:2410.23007, 2024.

[Li23] Z. Li et al., “Swapping-based entanglement routing design for congestion mitigation in quantum networks,” IEEE Trans. Netw. Serv. Manag., vol. 20, no. 4, pp. 3999–4013, 2023.

[Verzotti23] L. P. Verzotti and B.-C. Ciobanu, “Optimal quantum network decongestion strategies,” Scientific Reports, vol. 13, p. 9834, 2023.

[Jones16] C. Jones et al., “Design and analysis of communication protocols for quantum repeater networks,” New J. Phys., vol. 18, p. 083015, 2016.

[Hollot01] C. Hollot et al., “On designing improved controllers for AQM routers supporting TCP flows,” in Proc. IEEE INFOCOM, vol. 3, pp. 1726–1734, 2001.

[Xin25] Y. Xin and L. Zhang, “Towards a high-performance quantum data center network architecture,” in Proc. IEEE Int. Conf. Commun. (ICC), 2025 (to appear).

[Azuma23] K. Azuma et al., “Quantum repeaters: From quantum networks to the quantum internet,” Rev. Mod. Phys., vol. 95, p. 031002, 2023.

List of Acronyms

- **ACK**: Acknowledgment
- **AIMD**: Additive Increase / Multiplicative Decrease
- **AQM**: Active Queue Management
- **ATM**: Asynchronous Transfer Mode
- **ECN**: Explicit Congestion Notification
- **QECN**: Quantum Explicit Congestion Notification
- **IP**: Internet Protocol
- **MILP**: Mixed-Integer Linear Programming
- **QoS**: Quality of Service
- **QDC**: Quantum Data Center
- **QDN**: Quantum Data Network
- **QEC**: Quantum Error Correction
- **QECN / Quantum ECN**: Quantum Explicit Congestion Notification
- **QNUM**: Quantum Network Utility Maximization
- **QTP**: Quantum Transport Protocol
- **QuARC**: Quantum Adaptive Routing using Clusters
- **RED**: Random Early Detection
- **TCP**: Transmission Control Protocol
- **XCP**: eXplicit Control Protocol

Last modified on October 08, 2025
This and other papers on Recent Advances in Networking are available online at
https://www.cse.wustl.edu/~jain/cse5700-25/index.html
Back to Raj Jain's Home Page