



Full length article

Assisting mission-critical traffic flows with Active Queue Management in Industrial Internet of Things

Shuo Wang^a, Jonathan Kua^{b,*}, Jiong Jin^a, Yew Wee Wong^a, Prem Prakash Jayaraman^c, Zhibo Pang^d^a School of Engineering, Swinburne University of Technology, Melbourne, VIC, 3122, Australia^b School of Information Technology, Deakin University, Geelong, VIC, 3220, Australia^c School of Science, Computing and Emerging Technologies, Swinburne University of Technology, Melbourne, VIC, 3122, Australia^d School of Advanced Manufacturing and Robotics, Peking University, Beijing, China

ARTICLE INFO

Keywords:

Industrial Internet of Things
Active Queue Management
Mission-critical flows
Cloud-Fog Automation

ABSTRACT

Mission-critical Industrial Internet of Things (IIoT) traffic flows require bounded network latency and jitter guarantees to ensure the safe functioning of critical industrial infrastructure. These flows are typically communicated via commodity network routers with conventional First-In-First-Out (FIFO) buffers. FIFO has proven to be the culprit of the well-known bufferbloat phenomenon, and the deployment of Active Queue Management (AQM) schemes have demonstrated significant performance improvements for latency-sensitive applications over the Internet in the IT domain. However, the bufferbloat phenomenon and the efficacy of AQM schemes have not been studied in IIoT-based OT domain. In this paper, we propose the use of AQM as a lightweight and non-intrusive mechanism for assisting mission-critical traffic flows in IIoT networks. Our experimental results demonstrated that multi-queue AQM schemes provide substantial flow isolation and capacity sharing benefits, and significantly improve the performance of mission-critical traffic flows under network pressure. We further provide deployment recommendations based on our experimental insights.

1. Introduction

The rapid technological advancements and innovations in Industrial Internet of Things (IIoT) and Industrial Cyber-Physical Systems (ICPS) have propelled industrial automation systems towards a fully-digitized and Artificial Intelligence (AI)-driven Industry 4.0 paradigm [1,2]. In recent years, research efforts from academic and industrial communities have focused on shifting factory automation systems towards a more modular Cloud-based architecture, such as the newly proposed “Cloud-Fog Automation” architecture [3–5]. This architecture deploys lower layer applications such as supervisory controllers, machine controllers, I/O modules over wireless connections, whereas upper layers such as Enterprise Resource Planning (ERP) and Manufacturing Execution Systems (MES) are deployed in the Cloud.

Cloud-Fog Automation represents an ambitious vision and major paradigm shift from current industrial architectures. Although Cloud-Fog Automation provides significant technical performance and CapEx/OpEx business benefits, it requires considerable amount of time and effort to be fully realized, primarily due to its disruptive nature to

legacy systems. Many businesses could not afford the substantial downtime and cost-inducing upgrades required for their current automation infrastructure. The return on investments from such business decisions and deployments may not be immediately evident. However, businesses typically value and welcome non-intrusive technical decisions that have been experimentally-validated into their existing infrastructure [6].

To this end, we propose the use of Active Queue Management (AQM) schemes in an IIoT-enabled industrial automation systems to drive the transition towards Cloud-Fog Automation. The deployment of AQM is a backward-compatible, lightweight, and non-intrusive mechanism for improving network performance in IIoT, as opposed to new deterministic networking technologies which incur extensive developmental and deployment costs. Typical industrial automation systems use commodity network routers to provide connectivity for a broad range of operational IIoT devices and applications. Network routers (with their proprietary vendor software or open-source software) typically (by default) use the conventional *drop-tail* or First-In-First-Out (FIFO) queue management scheme due to its protocol and performance stability.

* Corresponding author.

E-mail addresses: shuowang@swin.edu.au (S. Wang), jonathan.kua@deakin.edu.au (J. Kua), jiongjin@swin.edu.au (J. Jin), yewweewong@swin.edu.au (Y.W. Wong), pjayaraman@swin.edu.au (P.P. Jayaraman), zhibo.pang@pku.edu.cn (Z. Pang).<https://doi.org/10.1016/j.jii.2026.101170>

Received 23 March 2026; Received in revised form 30 June 2026; Accepted 15 July 2026

Available online 22 July 2026

2452-414X/© 2026 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. (a) Experimental testbed emulating (b) Cloud-Fog Automation in a hybrid wired/wireless multi-cell smart manufacturing factory. The Cyber-Physical Factory cell consists of two AIVs, a FESTO® CP factory, IP camera and traffic generator, all connected to the MES via the router. The Additive Manufacturing cell consists of the Gantry system and Siemens smart meter, both connected to BaSyx via a wireless link.

However, FIFO has proven to be the culprit of the *bufferbloat* phenomenon over the Internet (*IT domain*), where unmanaged router buffers at either home gateways or Internet Service Providers (ISPs) lead to excessive delays and performance degradation in latency-sensitive and transactional traffic flows [7]. However, this phenomenon has not been studied in IIoT environments (*OT domain*), partly due to the common perception that high-bandwidth local area network connectivity are able to support the relatively low-rate industrial control applications in IIoT. The latency and reliability requirements of mission-critical applications (such as process control data, machine-to-machine communication, safety data, and control system signals, and so forth) is paramount.

Unlike conventional Internet environments, IIoT systems in the OT domain contain heterogeneous traffic flows that require not only throughput performance, but also predictable timing behavior, bounded latency and stable jitter characteristics. Many industrial traffic flows (e.g., OPC UA telemetry, machine-control signaling, robotic coordination and supervisory process communication) operate at relatively low data rates yet remain highly sensitive to transient queue buildup and timing perturbations. In addition, modern industrial environments increasingly integrate hybrid wired/wireless connectivity, edge-cloud coordination and heterogeneous industrial subsystems through commodity IT-based networking infrastructure. In such environments, unmanaged queueing behavior may introduce excessive delays and timing instability despite apparently sufficient bandwidth provisioning.

In traditional Internet applications, network congestion primarily degrades the Quality of Experience (QoE). However, queue-induced latency inflation in industrial environments may directly impact operational stability, synchronization and fail-safe system behavior. These characteristics motivate an experimental investigation into the behavior and deployment implications of modern AQM schemes in realistic IIoT environments. Our work focuses on studying the use of state-of-the-art AQM schemes in IIoT, and investigating how they can be deployed to assist mission-critical flows and alleviate network performance bottlenecks (paralleling the benefits in the IT domain, as recommended in RFC7567 [8]).

In this work, we emulate the shift towards Cloud-Fog Automation by constructing a sophisticated smart manufacturing experimental testbed using real-world industrial equipment with realistic hybrid multi-cell factory scenarios, and evaluated the performance of traffic flows generated by IIoT applications with or without the presence of AQM schemes. We compare the default FIFO queue management scheme with three representative AQM schemes: *Controlled Delay* (CoDel, defined in RFC8289 [9], representing single-queue AQM), *FlowQueue-CoDel* (FQ-CoDel, defined in RFC8290 [10], representing multi-queue AQM), *Common Applications Kept Enhanced* (CAKE, representing set-associative-based AQM [11]). Our experimental results

demonstrated the significant flow isolation and capacity sharing benefits of deploying AQM in IIoT networks.

In summary, our work offers the first real-world experimental validation of how AQM schemes can assist mission-critical IIoT traffic flows. While existing AQM schemes are not new, their adoption in industrial settings has been limited. Hence, we provide experimental evidence that multi-queue AQM schemes significantly outperform FIFO and single-queue mechanisms in protecting latency-sensitive flows. Our key contributions are as follows:

- (i) Identify unmanaged queue buildup and bufferbloat as the cause of performance bottleneck in IIoT, which motivates the use of AQM as a lightweight, backward-compatible and non-intrusive mechanism for assisting mission-critical industrial traffic flows.
- (ii) Experimentally characterize the operational behavior of representative AQM schemes in a real-world smart manufacturing IIoT testbed comprising heterogeneous wired/wireless industrial subsystems, OPC UA-based communication, robotic coordination traffic, real-time monitoring applications and competing elastic traffic flows.
- (iii) Analyze the impact of AQM schemes on mission-critical industrial applications' throughput stability, latency performance, jitter characteristics, flow isolation and operational responsiveness under congested and uncongested conditions.
- (iv) Provide practical deployment recommendations and operational insights into the deployment of AQM in industrial automation systems based on our experimental observations and in-depth analysis from realistic industrial network scenarios.

The rest of the paper is organized as follows. Section 2 provides background information and related work, Section 3 presents our experimental testbed and evaluation methodology, and Section 4 discusses our experimental results and provides deployment recommendations. Section 5 concludes the paper and outlines potential future work.

2. Background and related work

In this section, we provide background information on mission-critical IIoT traffic flows and AQM schemes, and discuss related work.

2.1. Traffic management for mission-critical IIoT flows

In IIoT environments, mission-critical traffic refers to communication flows where timely and reliable delivery is necessary to maintain correct operational behavior, process coordination, safety functions or control responsiveness. Examples include machine-control signaling, robotic coordination traffic, supervisory process communication, OPC

Table 1

Queue configurations (Baseline FIFO and AQM schemes).

Queue	Configuration
FIFO (baseline, no AQM)	Queue size = 500 packets
CoDel AQM	Queue size = 1000 packets, $T_{interval} = 100$ ms $T_{target} = 5$ ms
FQ-CoDel AQM	Queue size = 1000, $T_{interval} = 100$ ms, Flows = 1024, Quantum = 1514 bytes
CAKE AQM	RTT scheme = LAN, Diffserv mode = 3, Flow mode = triple isolate

UA telemetry and safety-related monitoring data. Unlike conventional best-effort Internet traffic, these traffic flows are often highly sensitive to transient delay inflation, jitter variation and communication instability (even at relatively low data rates).

Latency guarantees extend beyond achieving low average delay in industrial environments. Rather, it emphasize bounded and predictable timing behavior. Many industrial applications rely on periodic communication cycles, synchronized operations and deterministic control responsiveness. Excessive delay variation or transient queue buildup may negatively affect operational stability, synchronization accuracy and fail-safe behavior (even when overall throughput remains sufficient).

A crucial characteristic of IIoT is the need to ensure the timeliness of data transmission, with bounded latency and jitter guarantees, particularly for mission-critical applications. Control and data packets need to be transmitted and received within a strict deadline with minimum jitter (little variation in packet arrival). This represents ‘real-time’ and ‘deterministic’ network communication for mission-critical IIoT traffic flows. The Industrial Internet Consortium (IIC) categorizes criticality into low, medium and high [12]. Traffic types with *high* criticality necessitate the fulfillment of their Quality of Service (QoS) guarantees (such as latency, jitter, and packet loss) satisfied to ensure a functioning system. If QoS guarantees are not met, it will lead to critical system malfunction. Applications with *medium* criticality do not require their QoS standards to be consistently fulfilled; while unmet requirements may lead to performance degradation, it will not result in system failure. *Low* criticality refers to applications that are not directly relevant to the critical part of the system. These applications tolerate packet losses and retransmissions, and do not require latency and jitter guarantees.

2.2. Standards and protocols for deterministic networking

Current approaches for ensuring determinism primarily rely on wired connections (e.g., Industrial Ethernet), modified protocols (e.g., MQTT [13], OPC UA [14], EtherCAT [15], PROFINET [16]), network computing architectures [17] and integrated industrial networks (e.g., as proposed by the Internet Engineering Task Force (IETF) [18, 19]). Recent industrial developments include IETF’s Deterministic Networking (DetNet) as defined in RFC8578 [20,21], and IEEE 802.1’s Time-Sensitive Networking (TSN) [22,23], which focuses on high precision traffic scheduling and prioritization to ensure determinism and end-to-end guarantees over wired and wireless links.

2.3. State-of-the-art AQM schemes

Conventional FIFO queueing passively buffers packets until congestion becomes severe. AQM schemes are designed to proactively regulate/manage queue occupancy to reduce excessive delay buildup and improve responsiveness under congestion (i.e., minimizing overall system latency and jitter). Single-queue approaches such as CoDel primarily control queue delay within a shared bottleneck queue, whereas multi-queue mechanisms such as FQ-CoDel and CAKE additionally provide flow isolation and fairer capacity sharing among competing

traffic flows. These characteristics are particularly relevant in industrial environments where heterogeneous latency-sensitive and elastic traffic co-exist over shared communication infrastructure. AQM mechanisms dynamically drop or mark packets before the queue becomes full, hence signaling congestion to the endpoints.¹

PIE [24] estimates queuing delay from buffer occupancy and egress rate. It admits packets freely for the first 150 ms of an empty queue, then probabilistically drops packets based on a periodically adjusted drop probability (every 30 ms by default), targeting a delay of 15 ms. When the drop probability is below 10%, ECN-enabled flows are marked instead of dropped. PIE bases its decisions on both the current queuing delay and its trend. CoDel [9] tracks the minimum queuing delay over fixed intervals (initially 100 ms) using packet timestamps. If the minimum delay exceeds a target value throughout an interval, CoDel begins dropping packets based on a control law and continues until the delay falls below the target. ECN marking is also used when supported.

FQ-CoDel [10] classifies traffic flows into one of 1024 sub-queues by hashing the flow’s five-tuple. CoDel is used to manage each sub-queue independently, and the modified Deficit Round Robin (DRR) scheduler is used to dequeue a ‘quantum’ of bytes (1500 bytes by default), allowing FQ-CoDel to prioritize ‘sparse’ (latency-sensitive) traffic flows. It provides fair sharing of network resources and bandwidth sharing while controlling the flows’ queuing delay.

CAKE [11] is a shaping capable AQM that uses a combination of COBALT (CoDel+BLUE AQM), a traffic shaper, and a variant of DRR++ scheduler for flow isolation. It uses an 8-way associative hashing to eliminate hash collisions for placing flow packets into the independently-created sub queues, with each sub-queue running its own AQM state, and prioritizes latency-sensitive flows. CAKE provides several extra options, such as priority queuing with a simplified Diff-Serv implementation, overhead compensation, and Round Trip Time (RTT) tuning.

2.4. Related work and research gaps

Existing studies in the literature focused on evaluating and improving AQM schemes in consumer and commercial broadband networks in the wider Internet [25–27]. The work in [28] studied both the TCP incast and congested bottleneck problems in IIoT using network simulators. Other recent studies have focused on deep mathematical formulation and reasoning for designing cooperative TCP and AQM systems in the Internet [29].

New deterministic networking standards such as TSN are gaining traction in the community. TSN standards outline the data scheduling, synchronization and traffic shaping mechanism to achieve the specified latency requirements. TSN operates directly at Layer 2 (data link/MAC layer) and bypasses the TCP/IP layer [30,31]. It requires dedicated TSN switches (and infrastructure) that implements TSN stream scheduling/prioritization mechanisms to guarantee the deterministic latency requirements. In contrast to TSN, Active Queue Management (AQM) reduces intra-network latency at the network and transport layers without requiring specialized switches or infrastructure. AQM is supported by most commodity routers via the operating system kernel, and it manages internal buffer queues by creating logical sub-queues. It monitors queuing delays and proactively drops or marks packets to restrain aggressive flows and prioritize latency-sensitive traffic. While TSN offers deterministic latency and reliability, its adoption is limited by high deployment costs and poor interoperability with existing protocols such as EtherCAT and PROFINET. We present a low-cost, backward-compatible solution through the use of AQM in current IIoT

¹ In our work, FIFO is used strictly as a baseline for comparison and is not considered a form of AQM.

Table 2
Traffic flows' characteristics and functionalities.

Name	Application-layer Bitrate	Criticality	Traffic Properties and Functionalities
Autonomous Intelligent Vehicle (AIV) 1	9.3 Kbps, WiFi	High	Periodic, real-time monitoring of the movement and location with robot arm operation control, taking 3D printed sensor back housing to scan cell for quality inspection.
Autonomous Intelligent Vehicle (AIV) 2	9 Kbps, WiFi	High	Periodic, real-time monitoring of movement and location, exchanging PCB boxes for assembly robot to assist with part assembly.
IP Camera	1 Mbps, WiFi	Medium	Cyclic, real-time IP-based video streaming for surveillance, monitoring workplace environment and detecting intrusions.
FESTO® Cyber-Physical (CP) Factory	0.5 Kbps, Ethernet	Low	Sporadic, alarm and events data from CP Factory, modular production line testbed and programmed for the assembly of sensor products.
Gantry and BaSyx®	2 Kbps, Ethernet and WiFi	High	OPC UA flow from Omron PLC to BaSyx, which is a Digital Twin implementation for real-time control and monitoring of Gantry's operation.
Siemens® Smart Meter	25 Kbps, Ethernet and WiFi	High	OPC UA over TCP traffic flow for real-time energy monitoring of the 3D printing process.
Background Traffic	Elastic rate, Ethernet	Low	Elastic (<i>i.e.</i> , consuming all available bandwidth) TCP CUBIC traffic flows, representing background traffic with high-bandwidth utilization in congested networks.

infrastructure to enhance the performance of mission-critical traffic across both wired and wireless links.

Recent research efforts in industrial networking have increasingly focused on latency-sensitive communication [32], deterministic networking and QoS provisioning for IIoT/cyber-physical systems. Emerging approaches include TSN-based deterministic communication [33], edge-assisted industrial networking architectures, OPC UA real-time extensions [34–36] and adaptive traffic-management mechanisms for heterogeneous industrial workloads [37]. Modern AQM schemes such as CoDel, FQ-CoDel and CAKE have demonstrated significant latency-reduction benefits in broadband Internet environments by mitigating excessive queue buildup and improving flow isolation under network congestion. However, research on the operational impact of these queue management mechanisms in realistic OT-domain/IIoT environments remain limited. In particular, there are limited studies on the impact of AQM schemes on heterogeneous mission-critical traffic flows operating across hybrid wired/wireless industrial communication infrastructure, which require stringent operational timing and tight coupling with cyber-physical processes [38]. On the other hand, existing IIoT networking studies typically assume sufficiently provisioned local connectivity, rely on deterministic networking mechanisms such as TSN, or evaluate traffic behavior primarily through simulation-based models.

2.5. Novelty and contributions

In summary, most existing studies on IIoT networking assume high-bandwidth local connectivity or rely on deterministic mechanisms (*e.g.*, TSN, EtherCAT) without accounting for the effects of unmanaged queue buildup. As a result, the bufferbloat phenomenon (extensively studied in the IT domain) has received limited attention in the OT domain. For instance, the aforementioned prior works either use simulation-based models or pre-configured deterministic settings that do not reveal the queuing delay inflation caused by FIFO [25–28]. In contrast, our work highlights this performance bottleneck through real-world experiments and evaluates how AQM schemes can alleviate it across hybrid wired/wireless IIoT environments. This fills an important gap in the literature by showing that even in seemingly over-provisioned industrial networks, queuing effects can severely degrade mission-critical traffic unless explicitly managed.

To the best of our knowledge, this work represents the *first in-depth real-world experimental study* on investigating and evaluating the use of AQM schemes in the OT-domain/IIoT environments. It is our hope that this foundational work can motivate industrial network researchers and practitioners to adopt AQM in IIoT and advance IIoT-specific AQM innovations.

3. Experimental methodology

In this section, we present our experimental testbed setup and evaluation methodology for studying the benefits and impact of FIFO and AQM on mission-critical IIoT flows.

We note that this work does not propose a new AQM scheme or modify the internal mechanisms of existing queue management schemes. The primary objective of our work is to experimentally characterize and compare the operational behavior of representative state-of-the-art AQM schemes under realistic industrial IIoT environments. The evaluated AQM schemes are selected to represent distinct queue management design approaches and operational characteristics. CoDel represents single-queue delay-based AQM, FQ-CoDel combines queue-delay management with per-flow isolation and fair scheduling, while CAKE extends flow-isolated queue management with integrated shaping and enhanced traffic-handling mechanisms. These schemes are widely deployed (enabled by default in many ISPs' network routers/home gateways) and representative of modern low-latency queue management approaches supported by commodity networking infrastructure. The evaluated network conditions are designed to emulate realistic industrial communication scenarios in which IIoT traffic flows co-exist with competing elastic traffic over shared communication infrastructure.

3.1. System design and architecture

The primary goal of our experimental setup is to emulate a realistic Cloud-Fog Automation network in which mission-critical traffic coexists with background or elastic traffic flows. A central edge router is used to reflect common industrial deployment patterns where traffic from multiple sub-systems converges through a bottleneck queue before reaching upstream services (*e.g.*, cloud dashboards, digital twins). This router becomes the enforcement point for queue management policies.

To evaluate AQM under typical network load conditions, we introduce a background traffic generator that simulates bandwidth-hungry flows such as firmware updates, data logging, or external system backups. This configuration allows us to analyze how different queue management schemes (FIFO, CoDel, FQ-CoDel, CAKE) affect the latency and reliability of time-sensitive IIoT communications.

3.2. Experimental testbed setup and implementation

We built a real-world experimental testbed in our University's smart manufacturing testlab to investigate the benefits and impact of AQM on mission-critical traffic flows in IIoT environments. Fig. 1 illustrates

our physical experimental testbed setup and the corresponding network topology/configuration. Technical details and specifications of each key component are described in the following.

Routers. We use the MikroTik® hAP ac2 as our bottleneck router which connects the Cyber-Physical Factory cell to the Edge Cloud. An additional MikroTik router is connected to the bottleneck router via a wireless link to support the Additive Manufacturing cell, which is located in a farther physical location. Both internal routing software run Linux-based Router OS V7 that implements FIFO and CoDel, FQ-CoDel, CAKE AQM schemes,² and supports traffic shaping with Hierarchical Token Bucket (HTB) to provide a configurable bottleneck between the MES/BaSyX at the Edge Cloud level and field equipment/end-devices.

Background traffic generator. To emulate background industrial traffic and stress-test AQM's impact on evaluated IIoT traffic flows, we use *iperf3* to generate elastic TCP flows that maximize available bandwidth usage. The underlying TCP algorithm runs TCP CUBIC, which is known for its effective bandwidth utilization, and represents the default TCP algorithm of most industrial operating systems.

FESTO® Cyber-Physical (CP) Factory³ and Manufacturing Execution System (MES). CP Factory incorporates a suite of industrial functions into a modular and expandable factory. It consists of standard transportation modules for linear conveyor belts and branch modules. Our testbed features mobile modules with electrical, pneumatic, and mechanical interfaces, allowing for easy topological adjustments, utilizing Siemens® SIMATIC ET200SP PLC modules. Additionally, a symbiotic cell extended from the CP Factory enables the concurrent operation of collaborative robots (Cobots) and operators. Our MES operates on an Industrial PC (IPC), which serves as the central configuration, monitoring, and execution point for all machines and devices within the production cell. All control and process data to and from the MES are communicated via the MikroTik router.

Omron® Autonomous Intelligent Vehicles (AIVs). In our system, AIVs are deployed to assist human operators in performing quality inspection tasks on 3D-printed components. Two Omron LD-90 AIVs are utilized: AIV1 is equipped with a TM5-900 collaborative robotic arm and a PLC NX102-9020, while AIV2 operates without an onboard manipulator. The data traffic associated with AIV1 is categorized as mission-critical, as it encompasses time-sensitive control, coordination, localization, and situational awareness data streams between the robotic system and the MES. All communication is routed through the AIVs and a centralized bottleneck router. The job sequence proceeds as follows: upon receiving a task from the MES, AIV1 autonomously navigates to the designated workstation, retrieves the target components, and transports them to the 3D scanning cell. Once scanning is complete, a human operator performs a final verification, after which AIV1 returns to its docking location. Continuous and timely data exchange is essential to maintain operational integrity and avoid disruptions in the end-to-end production workflow.

Gantry and BaSyX® systems. A key component of our Additive Manufacturing cell is the Gantry and BaSyX system. Gantry is a crane-like system with industrial robots that operate on a mechanical framework that uses a mobile trolley over a linear bridge. Our setup uses Gantry as a smart inventory system for storing raw materials for 3D printing. It consists of a step motor-driven robot arm with a vacuum tip for picking and placing boxes and an Omron PLC NX102-9020 (from the AIVs) for logic control and data transmission. At the Edge Cloud level (in parallel with the MES), BaSyX®⁴ is a Digital Twin middleware that we implemented for monitoring and controlling Gantry's operations. The

traffic flows generated by Gantry are considered mission-critical as real-time inventory information is required to control and coordinate the 3D printing process.

Smart energy meter. We use Siemens® ET 200SP Analog input module AI energy meter 480VAC ST as our smart meter. It serves as the monitoring and reporting unit for the operation of the 3D printers in the Additive Manufacturing cell. Traffic flows generated by the smart meter to the BaSyX system are classified as mission-critical since the real-time energy consumption information of the printing process are required to optimize and improve the quality control of the 3D-printed parts while reducing material and energy wastage.

In the Additive Manufacturing cell, the PLCs driving the Gantry/BaSyX gantry subsystem and the smart-energy monitoring subsystem are linked via Ethernet to a local wireless access point. Both PLCs communicate using OPC UA. Process variables are published in real time, enabling client applications to subscribe to live telemetry streams. A UA-Expert OPC UA client continuously records the instantaneous energy consumption of the fused-deposition-modeling (FDM) printer. Although these OPC UA flows have a low data rate, their timeliness is critical; therefore, the proposed AQM schemes must prioritize and protect them.

3.3. Performance metrics and network conditions

We measure and analyze the flows' throughput, RTT, jitter and packet-size distributions to characterize the operational communication behavior of the industrial traffic flows. Full-length network packets are captured using *tshark* at the MES' and BaSyX's Network Interface Card (NIC), for deep packet inspection and traffic analysis. Throughput is calculated with a sliding window to determine the number of bits received per second (bps) per traffic flow. RTT and delay measurements are derived from the MES and BaSyX's perspective (systems at the Edge Cloud level perspective). It is important to note that our RTT measurements differs from the connotation of RTT in conventional data networks. The RTT measurements presented in this work includes path delays, induced queuing delays and the computational time of the tasks. Jitter is calculated as the variation of packet arrival times and interpreted in the context of operational timing stability for mission-critical industrial applications, and the packet size distribution is analyzed accordingly.

We emulate a Cloud-Fog Automation paradigm (Fig. 1(b)) where the MES is situated in the Edge Cloud, and data to/from the MES traverses through a bottleneck path in the operational factory environment. We set the bottleneck bandwidth at a symmetrical 10 Mbps for both uplink and downlink. The queue size for FIFO is configured at 500 packets (as per the Bandwidth Delay Product rule), and 1000 packets for CoDel, FQ-CoDel, and CAKE AQM (as per IETF's recommendations). Explicit Congestion Notification (ECN) is disabled at the end-hosts and router.

Table 1 summarizes the queue configurations and Table 2 presents the flow characteristics of the diverse traffic types in our test environment. The queue management parameters are selected based on commonly recommended operational configurations from the corresponding IETF RFCs and vendor-supported implementations. FIFO queue sizing followed the Bandwidth Delay Product (BDP) guideline, while CoDel and FQ-CoDel parameters use their default operational interval and target settings recommended in RFC8289 and RFC8290. CAKE is configured using its LAN RTT profile and triple-isolate flow mode to reflect a representative low-latency industrial deployment scenario.

All experiments are conducted using fixed queue management configurations and consistent traffic-generation settings across the evaluated scenarios to ensure comparability of the observed operational behavior among FIFO, CoDel, FQ-CoDel and CAKE. Note that the objective of these settings is not parameter optimization, but rather an evaluation of representative real-world deployment behavior under realistic industrial traffic conditions.

² <https://help.mikrotik.com/docs/display/ROS/Queues>.

³ https://www.festo.com/au/en/e/technical-education/training-concepts/highlights/training-factories/cp-systems-all-round-i4-0-training-factories-id_32122/.

⁴ <https://eclipse.dev/basyx/>.




Fig. 2. Time series and boxplots of throughput and RTT performance of AIV1's traffic flow in uncongested (a, d) and congested (b, e) network conditions in Scenario I.




Fig. 3. CDF plots of throughput, RTT, jitter and packet size distribution of AIV1's traffic flow in uncongested (a)–(d) and congested (e)–(h) network conditions in Scenario I.

4. Performance evaluation and analysis

In this section, we present our experimental findings and analysis of three representative industrial automation network scenarios. We then discuss deployment considerations and recommendations of AQM in industrial automation systems. The presented experimental scenarios are intentionally designed to represent an IIoT environment where heterogeneous mission-critical and background traffic flows co-exist over shared wired/wireless commodity networking infrastructure. The evaluated deployment reflects a transitional Cloud-Fog Automation setting in which industrial devices, PLCs, robotic systems, monitoring applications and Edge-Cloud services communicate through conventional IP-based routing infrastructure rather than a fully deterministic networking architecture. Our work focused on examining the scalability and operational robustness of AQM schemes under diverse industrial communication scenarios.

4.1. Scenario I: Typical industrial automation network

Scenario I focuses on how different queue disciplines impact mission-critical IIoT flows (AIV1 flow over wireless) in a typical industrial automation network (within the Cyber-Physical Factory cell as

illustrated in Fig. 1(b)), under congested/uncongested environments (with/without elastic TCP (*iperf*) flows). In this scenario, the field devices/equipment (comprising of two AIVs: AIV1 and AIV2, FESTO CP Factory and *iperf* traffic generator) are connected to the MES IPC via the router. The performance of AIV1 under uncongested environment over FIFO bottleneck represents the baseline and ground truth for the IIoT flows. We examine how AQM (either benefit or impede) the performance of these flows under congested/uncongested conditions.

4.1.1. The impact on throughput

Figs. 2(a) and 2(b) present AIV1's throughput across time (250 secs) in both uncongested and congested scenarios. As a ground truth, AIV1 generates 15 Kbps in an uncongested environment across all queue disciplines, i.e., FIFO, CoDel, FQ-CoDel and CAKE. This performance indicates that AQM has no negative impact on AIV1's flow, which provides us with a baseline. The impact of FQ-CoDel and CAKE is similar, where AIV1's throughput maintains the same 15 Kbps in FQ-CoDel and CAKE, while the throughput drops to 5 Kbps in CoDel. This performance is expected, as FQ-CoDel and CAKE both implements hashing and multiple sub-queue functions for even-capacity sharing, and protect latency-sensitive flows (which is reflected in AIV1's




Fig. 4. Time series and boxplots of throughput and RTT performance of AIV1's traffic flow in uncongested (a, d) and congested (b, e) network conditions in Scenario II.




Fig. 5. CDF plots of throughput, RTT, jitter and packet size distribution of AIV1's traffic flow in uncongested (a)–(d) and congested (e)–(h) network conditions in Scenario II.

flow in our experiments), whereas CoDel manages the packets within a single queue. The conventional FIFO queue shows the expected bufferbloat phenomenon, with a drastic 80% drop in throughput to 1 Kbps, which renders AIV1/robotic arm unusable. Figs. 2(c), 3(a) and 3(e) compare the throughput performance across queue disciplines under congested/uncongested conditions. The benefit of deploying AQM (specifically FQ-CoDel and CAKE) for preserving AIV1's flow throughput, and by extension, all latency-sensitive industrial mission critical flows is pronounced.

4.1.2. The impact on RTT and jitter

Figs. 2(d) and 2(e) show AIV1's induced RTT values across the experiment duration. In conjunction with the throughput performance, all queue disciplines demonstrate similar RTT values in the uncongested scenario, providing us with a baseline value of 10 ms (on average). Deploying AQM in an uncongested (low bandwidth utilization) scenario does not provide any noticeable difference. However, when faced with the congestion (when bandwidth is highly utilized by other traffic), the benefit of AQM is evident. As demonstrated in Fig. 2(e), the RTT values in FIFO is inflated to 150–200 ms, which is the typical bufferbloat phenomenon. When AQM is deployed, the RTT values are significantly

lowered, with FQ-CoDel and CAKE providing the same RTT values (10 ms) for the AIV (protecting AIV's flow) as the baseline. RTT values in CoDel is slightly higher (20 ms) due to its single-queue management algorithm. Figs. 2(f), 3(b) and 3(f) show the comparative boxplot and CDF plots for the induced RTT of AIV1's flow in both scenarios.

We note an important observation in the variation of packet jitter across all queue disciplines. As shown in Figs. 3(c) and 3(g), we observe that FQ-CoDel generally preserves the jitter characteristics most closely to the uncongested operational baseline across the evaluated scenarios. While the modified DRR scheduler (which is designed to serve a quantum of packets comprising 1500 bytes from each sub-queue in every iteration and prioritizes latency-sensitive flows, which is AIV1's flow in this case) likely contributes to this behavior by more evenly spacing packet transmissions across active flows, the observed jitter characteristics are also influenced by the interaction among per-flow queue isolation, queue management dynamics, traffic prioritization and congestion behaviors. We further note that CAKE also employs DRR-family scheduling and flow-isolation mechanisms, although differences in shaping behavior, queue management integration and traffic handling may contribute to the differences observed between CAKE and FQ-CoDel under congestion.

Although lower absolute jitter values are generally desirable from a networking perspective in the IT-domain, preserving timing characteristics close to the uncongested operational baseline is also important in the OT-domain. Lower absolute jitter values do not necessarily imply superior operational performance in industrial environments. For example, CAKE exhibits slightly lower absolute jitter distributions than FQ-CoDel in several congested scenarios. Nevertheless, preserving timing characteristics close to the normal operational baseline may be preferable for certain industrial applications. Abrupt changes in packet timing behavior under congestion may affect synchronization stability, control responsiveness or fail-safe operation (even when the absolute jitter values are lower). The observed jitter behavior is influenced not only by DRR scheduling, but also by the interaction among queue isolation, shaping behavior, traffic prioritization and flow dynamics under congestion.

Therefore, jitter performance among FIFO, CoDel, FQ-CoDel and CAKE should be interpreted in the broader context of operational timing stability rather than solely on absolute jitter values. Due to this reason, our analysis considers not only absolute jitter values, but also the extent to which the AQM schemes preserve stable and predictable timing behavior relative to the uncongested baseline.

4.1.3. The impact on packet size distribution

Figs. 3(d) and 3(h) investigate the packet size generated by AIV1 across all queue disciplines in both uncongested and congested scenarios. As the ground truth, all queue disciplines show similar packet size distribution (60 bytes, median), demonstrating the regular operation of AIV1 in the presence of video stream traffic. In the presence of competing elastic TCP flows, we observe the packet sizes generated in the FIFO scenario is unusually large (150 bytes, median). A closer packet inspection notes that the underlying transport protocol of AIV1 has clustered multiple payloads due to the extremely long delays incurred by bufferbloat in the bottleneck router (cf. Fig. 3(f)). The internal mechanism of the transport protocol is designed conserve bandwidth utilization in low-bandwidth scenario, hence combining packet payloads while waiting for acknowledgment. However, this mechanism is detrimental to the performance of AIVs (and all IIoT flows, by extension) due to their requirement for timely data transmission for real-time operation. The adverse impacts are observed in our experiments where AIV1 enters the ‘fail-safe’/system malfunction mode.

4.2. Scenario II: Industrial automation network with real-time video streaming

In Scenario II, we added an IP camera for real-time video streaming (used for surveillance and monitoring) to the network in the Cyber-Physical Factory cell. The IP camera is variable bitrate (VBR)-encoded, and generates a 7 Mbps data rate on average [39]. We study the impact of such bandwidth-intensive application on industrial traffic flows with FIFO and AQM schemes, with/without the presence of elastic TCP flows.

4.2.1. The impact on throughput

Figs. 4(a) and 4(b) demonstrate AIV1’s throughput with the presence of IP camera traffic in both uncongested and congested scenarios. Since the bottleneck bandwidth is set to 10 Mbps, the VBR-encoded video stream generates 7 Mbps of throughput, and the IIoT flows are in the order of Kbps, all queue disciplines manage to handle the capacity sharing relatively well without the presence of elastic TCP flows. Mission-critical flows are not severely impacted. However, the performance degradation with FIFO and the benefits of AQM schemes are clearly evident with the presence of elastic TCP flows. Fig. 4(b) shows that AIV1’s throughput drastically drops to <2 Kbps, resulting in AIV1 stalling and going into ‘fail-safe’ mode. Throughput with CoDel is much better at 6 Kbps (due to single-queue AQM), however it still falls short of the bitrate requirement for a functional AIV. FQ-CoDel

and CAKE both provide the best throughput performance, preserving AIV1’s throughput to the baseline in the presence of competing elastic flows. Figs. 4(c), 5(a) and 5(e) present the comparison plots of AIV1’s throughput in both scenarios, and clearly demonstrate the superiority of FQ-CoDel and CAKE.

4.2.2. The impact on RTT and jitter

Figs. 4(d) and 4(e) presents AIV1’s induced RTT values with the presence of video streams, in uncongested and congested scenarios. As with the throughput performance, RTT values are similar across all queue disciplines in the uncongested scenario due to the sufficiently large bandwidth. However, the bufferbloat phenomenon is once again evident in the congested scenario, with RTT in FIFO spiking to 150 ms on average (i.e., the RTT induced by full buffer capacity), rendering AIV1 unusable (cf. Fig. 4(b)). CoDel, FQ-CoDel and CAKE provide significant latency improvements, with FQ-CoDel and CAKE performing the best (due to their set hashing-enabled capacity sharing abilities), hence allowing AIV1’s induced RTT to remain at 10 ms (as the baseline) and function as normal. Figs. 4(f), 5(e) and 5(f) compare the induced RTT results and demonstrate the benefits of FQ-CoDel and CAKE. Similar to the observations in Scenario I, we note that FQ-CoDel’s per-flow queue isolation reduces interference between elastic and latency-sensitive traffic, hence preserving operational timing stability under congestion. It preserves packet jitter values (as demonstrated in Figs. 5(c) and 5(g)) due to its implementation of the modified DRR scheduler. The modified DRR scheduler evenly spaces out packet transmission on the router’s egress ports with an independently-managed sub-queue AQM algorithm. In-depth analysis and insights into reasons (and broader context implications) behind the jitter variations in FQ-CoDel and CAKE are presented in 4.1.2.

4.2.3. The impact on packet size distribution

Figs. 5(d) and 5(h) provide insights into the AIV1’s packet size distribution as impacted by the queue disciplines in both uncongested and congested scenarios. As the baseline, all queue disciplines result in nearly identical (median) value of <60 bytes, representing the periodic and timely transmission of packet payloads. Under congestion, FQ-CoDel and CAKE preserve the distribution most closely to the baseline value, whereas CoDel’s values skews slightly to the right, and FIFO shows the worst performance in a heavily right-skewed CDF plot. Upon inspection, it is noted at the transport stack of AIV1 combines multiple payloads into a single data packet due to the extremely high delay in waiting for the packet acknowledgment (rationale as discussed in Scenario I’s analysis).

4.3. Scenario III: A hybrid wired/wireless multi-cell industrial automation network with real-time smart inventory and smart energy monitoring systems

Scenario III represents a sophisticated smart manufacturing scenario which implements a separate *Additive Manufacturing cell* (located in a different physical location on the factory floor), in addition to the *Cyber-Physical Factory cell* represented in Scenarios I and II. This scenario integrates all end-nodes, middleware and industrial network elements depicted in Fig. 1. As described in Section 3.2, the separate Additive Manufacturing cell provides real-time inventory of raw materials with Gantry/BaSyX and smart energy monitoring, which informs the operation and optimization of the manufacturing process in the Cyber-Physical Factory cell (with AIVs handling the parts during the manufacturing process autonomously) via the BaSyX system and MES.

The real-time smart inventory system monitors and manage the availability of raw materials for 3D printing. The live visualization of data storage capacity is used to record past designs and predict material usage. When aligned with MES, the inventory system is able to optimize production schedules, optimize material usage and minimize waste. The real-time energy monitoring system tracks the energy usage of 3D printers, which enables the system to identify processes that

consume large amounts of energy and subsequently reduce operational costs. Both smart inventory and energy monitoring systems represents real-time systems that are crucial for modern smart manufacturing to self-regulate and cooperate with different factory cells for an optimized production process.

Traffic flows from both smart inventory and smart meter systems are generated by the OPC UA protocol to enable system-wide subscription of the live data, and to enable the interoperability of data transmission across industrial devices manufactured by different vendors. Both flows have different traffic characteristics. The application bitrate of the smart energy meter flow is 25 Kbps, whereas the bitrate of Gantry's operation (generated from the OPC UA server enabled within the Omron PLC) is 2 Kbps. As illustrated in Fig. 1(b), the router connecting the entire Additive Manufacturing cell to the bottleneck router is implemented across a wireless link. Although the bitrates of both flows are relatively low, they are classified as mission-critical traffic flows due to their stringent latency requirements for real-time information exchange to/from the BaSyX system (located at the Edge Cloud level) via a hybrid wired/wireless communication infrastructure.

In summary, this scenario specifically (i) investigates the impact of FIFO and AQM schemes on OPC UA-based mission-critical flows (for smart inventory and real-time energy monitoring); (ii) present the benefits of AQM in a hybrid wired/wireless multi-cell industrial automation network; and (iii) demonstrate the scalability of deploying AQM schemes across multi-cell factory floors, while accommodating the heterogeneity of low-rate mission-critical IIoT traffic flows.

4.3.1. The impact on throughput

Figs. 8(a), 8(e), 9(a) and 9(e) present the throughput CDF distribution of Gantry and smart meter traffic flows in uncongested and congested environments, respectively. It is noted that the throughput performance of the traffic flows within the Cyber-Physical Factory cell are consistent with experiments conducted in Scenarios I and II. Here we analyze the throughput of the OPC UA flows within the Additive Manufacturing cell. Fig. 7 shows the comparison throughput and RTT boxplots of all three mission-critical traffic flows — AIV, Gantry and smart meter traffic flows. The throughput boxplot of Gantry's flows are omitted due to their low values, but presented as a CDF plot in Figs. 8(a) and 8(e). Both Gantry and smart meter flows perform consistently well without congestion across FIFO, CoDel, FQ-CoDel and CAKE AQM. Gantry's throughput are not severely impacted under congestion due to its sporadic low-bitrate traffic of 2 Kbps.

However, smart meter flow across FIFO suffered significantly under congestion, with a drop of 40% in throughput. While CoDel performs better in maintaining the median throughput values, it has a large variation due to its single-queue AQM characteristics. FQ-CoDel and CAKE demonstrates the best performance in preserving the throughput of the smart meter flow under congestion due to their flow-isolation and bandwidth sharing capabilities. The ability to preserve OPC UA flows' throughput is important in ensuring both Gantry and smart meter transmit intact information with the BaSyX system in real-time.

4.3.2. The impact on RTT and jitter

Figs. 8(b), 8(f), 9(b) and 9(f) present the RTT distribution of Gantry and smart meter flows in uncongested and congested networks respectively, while Fig. 7(b) presents the comparison boxplots of RTT measurements for AIV, Gantry and smart meter flows. Figs. 8(c), 8(g), 9(c) and 9(g) demonstrate the corresponding jitter variations. It is clearly evident that both FQ-CoDel and CAKE provides the best latency-reducing performance for all mission-critical flows with their multi-queue AQM mechanisms, including the two low-rate OPC UA-based traffic flows. CoDel demonstrates RTT values lower than FIFO but still exhibits large variation due to its single-queue characteristics. The jitter distributions in Figs. 8(g) and 9(g) exhibit overlapping trends across the evaluated AQM schemes. While FQ-CoDel and CAKE generally demonstrate more stable timing behavior under congestion, the relative




Fig. 6. Boxplot of jitter performance of AIV and Gantry's traffic flows in uncongested and congested network conditions in Scenario III.

ordering of the jitter distributions varies across different portions of the CDFs. This behavior reflects the complex interaction among queue buildup dynamics, flow isolation, scheduling behavior, shaping mechanisms and traffic characteristics under congested conditions. Hence, the results should not be interpreted independently (as discussed in Section 4.1.2, but rather as demonstrating the difference in operational timing characteristics among the AQM schemes.

Fig. 6 presents the boxplot of the jitter performance of AIV and Gantry's traffic flows. This complements the RTT analysis by showing the timing stability of the AIV and Gantry flows across different queue disciplines. In the uncongested case, jitter remains relatively stable for all schemes. Under congestion, FIFO exhibits the widest distribution and the largest timing variation, while CoDel provides only partial improvement. FQ-CoDel and CAKE consistently achieve the most compact jitter ranges, which indicates that their multi-queue flow-isolation mechanisms are more effective in protecting IIoT traffic from queue buildup and contention. Although a left-shifted CDF plot can suggest lower absolute jitter, the boxplot shows if it comes with wider spread, longer whiskers or more variation. This observation supports the argument that stability and predictability are more important in IIoT traffic than lowering jitter to the smallest absolute value.

Another observation is that although Gantry's throughput across FIFO remains at 2 Kbps, its RTT values are inflated to around 200 ms (300% increase) which renders the flows unusable; whereas FQ-CoDel and CAKE preserve the RTT values of both Gantry and smart meter flows at 50 ms and 150 ms respectively (similar to when the network is uncongested), which satisfy their benchmarking requirements and translate into optimal performance in real-time data transmission to the BaSyX system. Therefore, the industrial significance of the observed jitter values should be interpreted in the context of timing stability and operational predictability rather than solely on absolute jitter values. Mission-critical industrial applications may tolerate small variations in instantaneous jitter, provided that the overall data communication behavior remains stable and predictable relative to the expected operational timing characteristics.

4.3.3. The impact on packet size distribution

Figs. 8(d), 8(h), 9(d) and 9(h) show the packet distribution of Gantry and smart meter flows in uncongested and congested environments, respectively. We observe a behavior in packet sizes in this scenario, which differs to Scenarios I and II. As demonstrated in the figures, the packet size distribution of Gantry and smart meter traffic flows are neither significantly impacted by network congestion nor the queue disciplines. Upon in-depth inspection of network and packet traces, we observed that the OPC server configuration within both PLCs has dedicated functional safety designs that ensure data packets of a specified length is consistently generated without relying on external acknowledgments, as opposed to TCP's behaviors (as observed and discussed in Scenarios I and II's analysis). Our results consistently demonstrate the operational distinction between single-queue and multi-queue queue




Fig. 7. Boxplots of throughput and RTT performance of AIV, smart energy meter and Gantry's traffic flows in uncongested and congested network conditions in Scenario III.

management approaches in IIoT. While CoDel effectively reduces excessive queue buildup compared to FIFO, the absence of explicit flow isolation limits its ability to protect latency-sensitive industrial traffic under sustained network congestion. In contrast, FQ-CoDel and CAKE are more effective in preserving timing stability and operational responsiveness by isolating competing traffic flows and reducing cross-traffic interference within the shared bottleneck/network infrastructure.

4.4. Limitations

While our work provide useful operational insights into the impact of AQM schemes in the OT-domain/IIoT environments, we acknowledge several limitations and scope boundaries below:

- (i) The evaluated scenarios reflect representative Cloud-Fog Automation communication conditions using commodity networking infrastructure and heterogeneous industrial traffic flows. However, they do not exhaustively capture all industrial deployment scales, topologies or traffic characteristics. Larger deployments may exhibit additional communication heterogeneity, synchronization requirements, wireless interference dynamics and distributed edge-cloud coordination complexity.
- (ii) The generated traffic patterns emulate representative mission-critical industrial communication behavior, including OPC UA telemetry, robotic coordination traffic, supervisory communication and competing elastic traffic flows. Nevertheless, real industrial environments may involve additional proprietary protocols, deterministic communication schedules and application-specific traffic dynamics that can influence the performance of AQM schemes.
- (iii) Our study focuses on AQM behavior within conventional IP-based networking infrastructure and does not evaluate deterministic networking technologies such as TSN or DetNet. Hence, the results should not be interpreted as a replacement for deterministic industrial networking architectures, but rather as

an investigation of lightweight and backward-compatible AQM approaches suitable for existing IIoT systems.

- (iv) Although our experiments demonstrate consistent operational trends across the evaluated scenarios, broader evaluation across larger industrial deployments, new wireless connectivity conditions and more diverse operational workloads remains important future work.

Despite these limitations, the work presented in this paper demonstrate consistent operational trends across heterogeneous industrial communication scenarios and provide practical insights into the deployment implications of modern AQM schemes in for current/legacy IIoT systems transitioning towards a new Cloud-Fog Automation-based paradigm.

4.5. Practical deployment considerations and recommendations

In this section, we provide several practical deployment considerations and recommendations for industry practitioners and researchers in the field.

4.5.1. Hardware and infrastructure

A practical advantage of deploying modern AQM mechanisms in industrial OT/IIoT environments is their compatibility with commodity IP-based networking infrastructure. Unlike deterministic networking approaches that may require specialized switching hardware and tightly coordinated scheduling mechanisms, AQM schemes such as CoDel, FQ-CoDel, and CAKE are widely supported in modern Linux-based networking stacks and commodity routing platforms. This enables incremental deployment within existing industrial communication infrastructure with relatively low hardware replacement cost.

4.5.2. Interoperability

Interoperability remains an important consideration in industrial networking environments due to the coexistence of heterogeneous protocols, legacy automation systems, vendor-specific equipment, and mixed communication technologies. In this context, the evaluated AQM mechanisms operate transparently at the network and queue-management level without requiring modification to application-layer industrial protocols such as OPC UA or changes to the underlying industrial devices and PLCs. This backward-compatible behavior is particularly important in transitional Cloud-Fog Automation environments where complete infrastructure replacement may not be operationally or economically feasible.

4.5.3. Deployment recommendations and caution

Despite the observed benefits, operational deployment of AQM in industrial environments should be approached carefully. Industrial systems often contain tightly coupled control processes and timing-sensitive applications where their behavior may depend on stable communication characteristics. AQM configurations should be validated under representative operational traffic conditions prior to deployment, and conservative parameter selection should be preferred unless extensive scenario-specific evaluation has been conducted. In summary, we provide the following deployment considerations/recommendations based on our experimental insights:

- (i) Modern IIoT facilities (e.g., smart factories) use IT-based routers as gateways/proxies for OT environments. Although AQM schemes such as PIE, CoDel, FQ-CoDel, CAKE are provided, FIFO remains the default queue discipline in most industrial routers. The trepidation to enable AQM in OT environment is the concern that AQM can cause unpredictable packet losses to OT flows. However, we have experimentally proven in this work that AQM does not adversely impact the performance of IIoT traffic. Rather, AQM schemes significantly improve throughput and RTT performance (especially mission-critical flows with stringent real-time requirements) in low-bandwidth and congested network environments.




Fig. 8. CDF plots of throughput, RTT, jitter and packet size distribution of Gantry's OPC UA traffic flow in uncongested (a)–(d) and congested (e)–(h) experimental trials in Scenario III.




Fig. 9. CDF plots of throughput, RTT, jitter and packet size distribution of Siemens smart energy meter OPC UA traffic flow in uncongested (a)–(d) and congested (e)–(h) experimental trials in Scenario III.

- (ii) Network performance is sensitive to the internal AQM parameters, as they dictate how the bottleneck queue signals and interacts with the end-to-end transport protocols for regulating data transmission. Global AQM parameters (which are extensively tested and intricately tuned) should be left as default and ECN should be disabled, unless extensive experiments have been conducted on scenario-specific network configurations.
- (iii) Practitioners should be guided by RFC7567 [8] and RFC7928 [40] when deploying AQM schemes in industrial settings. AQM behaviors should be first assessed in a controlled environment (with the recommended traffic mix), to ensure a safe operational deployment. Practitioners should consider the potential impact of the central aspects of an AQM algorithm, such as their burst absorption capacity, RTT fairness, resilience to fluctuating network conditions, on the intended use cases.

Although our experimental scenarios do not represent exhaustive large-scale industrial deployments, the observed behavior suggests that multi-queue AQM mechanisms remain beneficial in heterogeneous IIoT environments where mission-critical and elastic traffic co-exist over shared communication infrastructure.

As industrial deployments scale in terms of traffic diversity, subsystem heterogeneity and edge-cloud coordination complexity, flow-isolation and queue management behavior are expected to become increasingly important for preserving operational timing stability under network congestion.

5. Conclusions and future work

In this work, we proposed and experimentally studied the use of AQM schemes to assist mission-critical traffic flows in a real-world IIoT environment with three realistic network scenarios. There is currently a building momentum towards a more network-centric and collaborative paradigm, known as Cloud-Fog Automation. To progressively realize this paradigm with state-of-the-art networking technologies, operational networks in IIoT first require a backward-compatible and non-intrusive technique to alleviate existing network pain points.

Our work demonstrates that the commonly-observed bufferbloat phenomenon in the IT-domain is also prevalent in the OT-domain. We implemented emerging AQM schemes (CoDel, FQ-CoDel and CAKE) in our network, and showed that AQM schemes benefit all OT traffic flows, especially in assisting mission-critical traffic flows to be unperturbed (and protected) in heavily congested environments. Both

FQ-CoDel and CAKE are superior due to their sophisticated flow isolation and capacity sharing abilities. We observed that FQ-CoDel and CAKE enable mission-critical flows to achieve up to 80% performance improvement when compared with FIFO — allowing the production process to continue seamlessly in spite of heavy competing traffic. Furthermore, we have demonstrated that AQM schemes can assist OPC UA flows in a sophisticated hybrid wired/wireless environment, which encourages the wider deployment of AQM for supporting protocol interoperability in IIoT networks.

The core contribution of this work is the experimental characterization of the operational benefits and impact of modern AQM schemes on the mission-critical industrial traffic performance in the OT-domain/IIoT environments (and not in proposing a new IIoT-specific AQM scheme, which is the next step in our future research). It is our sincere hope that this foundational work provides the motivation, experimental characterization and validation for operators to deploy AQM in modern industrial networks.

Future work include investigating the limitations outlined in Section 4.4, and studying AQM parameter-tuning for IIoT environments, implementing AI/ML-driven decision-making into IIoT-based AQM schemes, validating protocol correctness, integrating AQM with in-network computing technologies to achieve network-wide deterministic low-latency networking. A rethink and redesign of queue management strategies is required to cater for non-TCP/IP, ultra-Reliable Low-latency Communication (uRLLC) traffic flows in emerging 5G/6G networks [41,42].

CRedit authorship contribution statement

Shuo Wang: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation. **Jonathan Kua:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Data curation, Conceptualization. **Jiong Jin:** Writing – review & editing, Writing – original draft, Supervision, Software, Resources, Project administration, Funding acquisition, Conceptualization. **Yew Wee Wong:** Writing – original draft, Supervision, Software, Methodology, Investigation, Formal analysis. **Prem Prakash Jayaraman:** Validation, Software, Resources, Project administration, Funding acquisition. **Zhibo Pang:** Supervision, Resources, Methodology, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.


Data availability

Data will be made available on request.

References

- [1] Y. Wu, J. He, K. Zhang, X. Huang, F. Wu, Y. Zhang, Embodied intelligence-based hybrid edge computing networks for scalable task execution in industrial IoT, *J. Ind. Inf. Integr.* (2025) 100916.
- [2] D. Xia, P. Xu, G. Han, J. Jiang, Emerging perspectives on embodied intelligence in future smart manufacturing, *J. Ind. Inf. Integr.* (2025) 101020.
- [3] J. Jin, K. Yu, J. Kua, N. Zhang, Z. Pang, Q.-L. Han, Cloud-fog automation: Vision, enabling technologies, and future research directions, *IEEE J. IINF* (2023) 1–16.
- [4] J. Jin, Z. Pang, J. Kua, Q. Zhu, K.H. Johansson, N. Marchenko, D. Cavalcanti, Cloud-fog automation: The new paradigm towards autonomous industrial cyber-physical systems, *IEEE J. Sel. Areas Commun.* (2025).
- [5] R. Mahmud, J. Jin, J. Kua, M. Afrin, S. Mistry, A. Krishna, Trusted microservice orchestration for secure edge computing in industrial cyber-physical systems, *IEEE Netw.* (2025).
- [6] M. Lezzi, M. Lazoi, A. Corallo, Cybersecurity for industry 4.0 in the current literature: A reference framework, *Comput. Ind.* 103 (2018) 97–110.
- [7] J. Gettys, K. Nichols, Bufferbloat: Dark Buffers in the Internet, *Queue* 9 (11) (2011) 40:40–40:54.
- [8] F. Baker, G. Fairhurst, IETF recommendations regarding active queue management, 2015, RFC 7567.
- [9] K.M. Nichols, V. Jacobson, A. McGregor, J. Iyengar, Controlled delay active queue management, 2018, RFC 8289.
- [10] T. Høiland-Jørgensen, P. McKenney, D. Taht, J. Gettys, E. Dumazet, The flow queue CoDel packet scheduler and active queue management algorithm, 2018, RFC 8290.
- [11] T. Høiland-Jørgensen, D. Täht, J. Morton, Piece of CAKE: A comprehensive queue management solution for home gateways, in: 2018 IEEE International Symposium on Local and Metropolitan Area Networks, LANMAN, 2018, pp. 37–42.
- [12] Industrial Internet Consortium, Time sensitive networks for flexible manufacturing testbed characterization and mapping of converged traffic types, 2019.
- [13] D.M. Anna, M. Vijayalakshmi, K.S. Raju, Enabling lightweight device authentication in message queuing telemetry transport protocol, *IEEE J. IOT* 11 (9) (2024) 15792–15807.
- [14] W. Dai, Y. Zhang, Y. Zhang, J. Kang, D. Huang, Automatic information model generation for industrial edge applications based on IEC 61499 and OPC UA, *IEEE J. IINF* 19 (4) (2023) 6093–6104.
- [15] G. Cena, I.C. Bertolotti, S. Scanzio, A. Valenzano, C. Zunino, Evaluation of etherCAT distributed clock performance, *IEEE J. IINF* 8 (1) (2012) 20–29.
- [16] G.S. Sestito, A.C. Turcato, A.L. Dias, M.S. Rocha, M.M. da Silva, P. Ferrari, D. Brandao, A method for anomalies detection in real-time ethernet data traffic applied to PROFINET, *IEEE J. IINF* 14 (5) (2018) 2171–2180.
- [17] C. Tang, C. Zhu, N. Zhang, M. Guizani, J. Rodrigues, SDN-assisted mobile edge computing for collaborative computation offloading in Industrial Internet of Things, *IEEE J. IOT* 9 (23) (2022) 24253–24263.
- [18] K. Makhijani, L. Dong, Framework for integrated industrial networks, 2022, draft-iiotops-km-iiot-frwk-00.
- [19] C. Tang, H. Wen, S. Ruan, B. Huang, X. Feng, IPv6 and 5G based architecture for IIoT, in: Informational, 2020, draft-tang-iiot-architecture-00.
- [20] E. Grossman, Deterministic networking use cases, 2019, <http://dx.doi.org/10.17487/RFC8578>, RFC 8578.
- [21] H. Yu, T. Taleb, J. Zhang, Deep reinforcement learning-based deterministic routing and scheduling for mixed-criticality flows, *IEEE J. IINF* 19 (8) (2023) 8806–8816.
- [22] Y. Li, J. Jiang, S.H. Hong, Joint traffic routing and scheduling algorithm eliminating the nondeterministic interruption for TSN networks used in IIoT, *IEEE J. IOT* 9 (19) (2022) 18663–18680.
- [23] L. Ji, S. He, C. Gu, Z. Shi, J. Chen, Routing and scheduling for low latency and reliability in time-sensitive software-defined IIoT, *IEEE J. IOT* 11 (7) (2024) 12929–12940.
- [24] G. White, R. Pan, Active queue management (AQM) based on proportional integral controller enhanced (PIE) for data-over-cable service interface specifications (DOCSIS) cable modems, 2017, RFC 8034.
- [25] C.A. Grazia, N. Patriciello, T. Høiland-Jørgensen, M. Klapez, M. Casoni, Aggregating without bloating: Hard times for TCP on WiFi, *IEEE J. NET* 30 (5) (2022) 2359–2373.
- [26] J. Kua, S.H. Nguyen, G. Armitage, P. Branch, Using active queue management to assist IIoT application flows in home broadband networks, *IEEE J. IOT* 4 (5) (2017) 1399–1407.
- [27] T. Song, Y. Kyung, Deep-reinforcement-learning-based age-of-information-aware low-power active queue management for IIoT sensor networks, *IEEE J. IOT* 11 (9) (2024) 16700–16709.
- [28] D. Morato, C. Pérez-Gómara, E. Magaña, M. Izal, Network simulation in a TCP-enabled industrial internet of things environment-reproducibility issues for performance evaluation, *IEEE J. IINF* 18 (2) (2021) 807–815.
- [29] A.M. Abdelmoniem, B. Bensaou, Enhancing TCP via hysteresis switching: Theoretical analysis and empirical evaluation, *IEEE J. NET* (2023).
- [30] M. Wollschlaeger, T. Sauter, J. Jasperneite, The future of industrial communication: Automation networks in the era of the internet of things and industry 4.0, *IEEE Ind. Electron. Mag.* 11 (1) (2017) 17–27.
- [31] T. Tran, Q. Nguyen, T. Luu, M. Tran, J. Kua, T. Hoang, M. Dien, Empowering robotic training with kinesthetic learning and digital twins in human-centric industrial systems, *J. Ind. Inf. Integr.* 43 (2025) 100743.
- [32] J. Cheng, W. Chen, F. Tao, C.-L. Lin, Industrial IIoT in 5G environment towards smart manufacturing, *J. Ind. Inf. Integr.* 10 (2018) 10–19.
- [33] J. Jiang, Y. Li, X. Zhang, M. Yu, C.D. Lee, S.H. Hong, Assessing the traffic scheduling method for time-sensitive networking (TSN) by practical implementation, *J. Ind. Inf. Integr.* 33 (2023) 100464.
- [34] C. Lee, H. Kwon, Y.I. Lee, OPC UA-based three-layer architecture for aggregated microgrids integrating edge cloud computing and IEC 62264, *J. Ind. Inf. Integr.* (2025) 100965.
- [35] M. Ladegourd, J. Kua, Performance analysis of opc ua for industrial interoperability towards industry 4.0, *IIoT* 3 (4) (2022) 507–525.
- [36] A. Martins, J. Lucas, H. Costelha, C. Neves, CNC machines integration in smart factories using OPC UA, *J. Ind. Inf. Integr.* 34 (2023) 100482.

- [37] R. Kumar, N. Agrawal, Analysis of multi-dimensional industrial IoT (IIoT) data in edge-fog-cloud based architectural frameworks: A survey on current state and research challenges, *J. Ind. Inf. Integr.* 35 (2023) 100504.
- [38] K.S. Hawaou, V.C. Kamla, S. Yassa, O. Romain, J.E.N. Mboula, L. Bitjoka, Industry 4.0 and industrial workflow scheduling: A survey, *J. Ind. Inf. Integr.* 38 (2024) 100546.
- [39] J. Kua, G. Armitage, P. Branch, A survey of rate adaptation techniques for dynamic adaptive streaming over HTTP, *IEEE O CSTO* 19 (3) (2017) 1842–1866.
- [40] N. Kuhn, P. Natarajan, N. Khademi, D. Ros, Characterization Guidelines for Active Queue Management (AQM), 2016, RFC 7928.
- [41] C.I. Bădoi, B.K. Çetin, K. Çetin, Information integration for industry 5.0: A review of IIoT-edge-cloud architectures with 6G/ISAC capabilities, *J. Ind. Inf. Integr.* (2026) 101117.
- [42] E. Zeydan, S. Arslan, Y. Turk, 6G wireless communications for industrial automation: Scenarios, requirements and challenges, *J. Ind. Inf. Integr.* 42 (2024) 100732.



Shuo Wang is a Ph.D. candidate whose research focuses on low-latency networking for industrial automation. He also works full-time as a senior technical officer at the Factory of the Future at Swinburne University of Technology, supporting technical facilities, equipment and daily operations. This role provides practical experience and technical insights that directly inform and strengthen his research.



Jonathan Kua received the B.Eng. (First Class Hons.) degree in telecommunications and network engineering and the Ph.D. degree in telecommunications engineering from Swinburne University of Technology, Australia, in 2014 and 2019, respectively. He is currently a Senior Lecturer and Academic Director (Cloud Computing) within the School of Information Technology, Deakin University, Australia. His research interests include low-latency data transport protocols, adaptive streaming, content delivery, Internet of Things and industrial cyber-physical systems, with an overarching focus on improving their network performance and quality of service. He was the recipient of the Netflix Ph.D. scholarship award (2015–2019) and the second runner-up of the international DASH-IF Best Ph.D. Dissertation Award on “Algorithms and Protocols for Adaptive Content Delivery over the Internet” (2019). He previously served as a Guest Editor for the IEEE Journal on Selected Areas in Communications (JSAC) Special Issue on “Co-Design of Communication, Computing, and Control in Industrial Cyber-Physical Systems”.



Jiong Jin received the B.E. degree with First Class Honours in Computer Engineering from Nanyang Technological University, Singapore, in 2006, and the Ph.D. degree in Electrical and Electronic Engineering from the University of Melbourne, Australia, in 2011. He is presently a Full Professor in the School of Engineering at Swinburne University of Technology, Melbourne, Australia. His research interests include network design and optimization, edge computing and intelligence, robotics and automation, and cyber-physical systems and Internet of Things as well as their applications in smart manufacturing, smart transportation and smart



Yew Wee Wong received the Bachelor of Engineering and Ph.D. degrees from Swinburne University of Technology, Melbourne, Australia in 2016 and 2021, respectively. He is a certified project and systems engineering manager with an expertise in robotics and mechatronics. He specialises in leading complex, high-value projects at the intersection of engineering, ICT and innovation, translating sophisticated technical systems into practical, value-driven solutions. His research interests include robotics, control, and the integration of AI and Internet of Things.



Prem Prakash Jayaraman is a globally recognised researcher and academic leader in the areas of digital manufacturing, distributed systems and emerging technologies. He is a Professor at Swinburne University of Technology, Melbourne, Australia. He currently serves as the Director of Swinburne's Factory of the Future and Digital Innovation Lab, and Deputy Director of the ARC Industrial Research Transformation Hub for Future Digital Manufacturing. He leads Australia's only university-based Industry 5.0 program for small-to-medium enterprise, delivered by industry-professionals to help manufacturing businesses navigate the transition to advanced digital manufacturing technologies.

His research spans across distributed systems, Internet of Things, and applied AI, attractive over \$20 million in competitive external research funding from industry, Government and Australian Research Council. He has (co)-authored over 170 papers in several high ranked journals and has been ranked in the top 2% of the world's most cited researchers by Stanford University every year since 2021. Prof. Jayaraman has been awarded three Vice Chancellor awards for excellence in translating research into industry outcomes and impact.



Zhibo Pang received MBA in Innovation and Growth from University of Turku in 2012 and PhD in Electronic and Computer Systems from KTH Royal Institute of Technology in 2013. He is currently a Full Professor with Tenure at Peking University, China. Before that, he was Senior Principal Scientist at ABB Corporate Research Sweden, Adjunct Professor at KTH Royal Institute of Technology, and Adjunct Professor at University of Sydney. He is a Standing Committee Member of IEEE Industrial Electronics Society, Steering Committee Member of IEEE IoT Technical Community, Chair of IEEE Technical Committee on Cloud and Wireless Systems for Industrial Applications, and Co-Chair of IEEE TC on Industrial Informatics. He is Associate Editor of IEEE TII, IEEE JBHI, IEEE TCE, IEEE JESTIE, and IEEE IoTM. He works on embodied AI, robotics, control, computing, communication, and electronics for autonomous systems.