Scaling 5G-TSN Bridges: Operating Regimes, Scheduling, and Time Synchronisation Under Heterogeneous Industrial Traffic

Mohamed Seliem
School of CS and IT
University College Cork
Cork, Ireland
MSeliem@ucc.ie

Utz Roedig
School of CS and IT
University College Cork
Cork, Ireland
U.Roedig@ucc.ie

Cormac Sreenan
School of CS and IT
University College Cork
Cork, Ireland
Cormac.Sreenan@ucc.ie

Dirk Pesch
School of CS and IT
University College Cork
Cork, Ireland
Dirk.pesch@ucc.ie

***Abstract*—3GPP Release 16 enables a 5G system to operate as a transparent IEEE 802.1 TSN bridge, but its scalability under heterogeneous industrial workloads remains insufficiently characterised. This paper uses the nascTime framework on OMNeT++/Simu5G to evaluate how many TSN endpoints a single 5G NR cell can bridge before per-flow QoS degrades. We model closed-loop control, machine vision, bulk telemetry, and IEEE 802.1AS traffic over a four-bearer SDAP architecture, varying the number of endpoints from 1 to 40, MAC scheduler, radio bandwidth (10 MHz and 20 MHz), and channel model. Results show three operating regimes. Below saturation, non-DRR schedulers perform similarly; near saturation, QoS-aware PF reduces critical-flow P99 latency by up to two orders of magnitude relative to channel-aware and fairness-based schedulers; and under overload, only QoS-PF maintains near-complete delivery for the highest-priority traffic. Across the two evaluated bandwidths, the saturation threshold approximately doubles when bandwidth doubles. We also show that isolating IEEE 802.1AS/gPTP traffic on a dedicated high-priority radio bearer reduces clock-servo instability, although endpoints carrying lower-priority data still experience elevated synchronisation delay under saturation because of reduced MAC scheduling frequency. Finally, the evaluated sub-6 GHz, 30 kHz SCS configuration exhibits an effective latency floor of approximately 2.25 ms, indicating that sub-3 ms TSN deadlines may require radio-configuration changes such as configured grants or higher numerology.**

***Index Terms*—5G-TSN integration; MAC scheduling; scalability; IEEE 802.1AS; QoS mapping; radio bearer; time synchronisation; industrial automation; OMNeT++**

I. Introduction

3GPP Release 16 specifies how a 5G system can operate as a transparent IEEE 802.1 Time-Sensitive Networking (TSN) bridge, mapping TSN priority classes through 5G QoS flows and radio bearers while supporting IEEE 802.1AS transparent-clock behaviour [1]. Although early testbeds demonstrate feasibility at small scale, a key deployment question remains unresolved: *how does per-flow QoS degrade as more TSN endpoints share a single gNB, and what determines the practical capacity limit?*

This question is strongly affected by implementation choices that the 3GPP specifications leave open, especially the gNB MAC scheduler, which allocates shared radio resources among competing traffic flows. Existing work has compared schedulers for generic URLLC traffic [2] and studied time synchronisation over 5G in isolation [3], but no prior study jointly characterises endpoint scaling, scheduler impact, bandwidth sensitivity, fading behaviour, and gPTP stability under a heterogeneous industrial TSN workload.

We address this gap using the nascTime simulation framework [14] on OMNeT++ 6.3, INET 4.6, and Simu5G. The model includes closed-loop control, machine vision, bulk telemetry, and IEEE 802.1AS/gPTP traffic mapped through a four-bearer Service Data Adaptation Protocol (SDAP) architecture. We evaluate 1 to 40 endpoints, four MAC schedulers—Maximum C/I, Proportional Fair, Deficit Round Robin, and QoS-aware Proportional Fairness—two bandwidths (10 MHz and 20 MHz), and both ideal and Jakes fading channels.

The results show three operating regimes. Below saturation, non-DRR schedulers provide similar performance and scheduler choice has limited impact. Near saturation, QoS-aware scheduling reduces critical-flow P99 latency by one to two orders of magnitude relative to channel-aware or fair scheduling. Under overload, only QoS-aware scheduling maintains near-complete delivery for the highest-priority traffic class, while other schedulers exhibit spatial unfairness or class starvation. Across the two evaluated bandwidths, the saturation threshold approximately doubles when bandwidth increases from 10 MHz to 20 MHz.

The study also exposes two architectural risks. First, the evaluated per-UE DRR implementation is unsuitable for heterogeneous TSN bridging: although it preserves closed-loop-control delivery, it starves machine-vision and telemetry flows because its fixed per-user quantum is not flow-class aware. DRR's marginally lower CLC P99 delay should not be read as better scheduling. By suppressing MV-HP to 2.1% and BLK to 7.6%, DRR effectively reduces the workload to a CLC-only stream, so its CLC latency is measured under far lighter contention than the other schedulers experience. Weighting the per-endpoint delivery counts by the $N = 20$ profile mix, QoS-PF delivers approximately 16% more packets in total

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

than DRR and 52% more than PF, while also protecting the highest-priority class. DRR's apparent CLC advantage is therefore purchased at the cost of overall bridge utility. Second, IEEE 802.1AS/gPTP traffic requires explicit bearer treatment. A dedicated high-priority gPTP bearer avoids the clock-servo divergence observed when synchronisation traffic competes on the default bearer, but it does not fully bound residence time under saturation because endpoints carrying lower-priority data receive fewer MAC scheduling opportunities.

This paper makes three contributions:

1) It characterises the scalability of 5G-TSN bridges under a heterogeneous industrial traffic mix, identifying three operating regimes and showing that, for the evaluated workload, the saturation threshold approximately doubles when bandwidth doubles.
2) It evaluates a four-bearer SDAP architecture with a dedicated high-priority IEEE 802.1AS/gPTP bearer, showing that bearer isolation reduces clock-servo instability but does not by itself guarantee bounded gPTP residence time under saturation.
3) It compares scheduler behaviour under ideal and Jakes fading channels, showing that the relative scheduler ranking is preserved in the tested configurations and motivating broader study of fading effects in 5G-TSN bridge deployments.

The remainder of the paper is organised as follows. Section II reviews 5G-TSN bridge architecture and related work. Section III presents the four-bearer SDAP design. Section IV describes the simulation methodology. Section V reports scaling, bandwidth, fading, and synchronisation results. Section VI discusses deployment implications and limitations, and Section VII concludes.

II. Background and Related Work

A. 5G-TSN Bridge Architecture

3GPP Release 16 specifies a 5G system as a logical IEEE 802.1 TSN bridge between wired TSN segments [1]. The bridge is formed by a Network-side TSN Translator (NW-TT), co-located with the UPF, and a Device-side TSN Translator (DS-TT), co-located with the UE. The TSN Application Function exposes bridge capabilities such as port delay, supported rates, and traffic-class support to the TSN Centralised Network Controller.

QoS mapping follows a multi-stage pipeline. At the NW-TT, the IEEE 802.1Q Priority Code Point (PCP) is mapped to an IPv4 DSCP value; the UPF Traffic Flow Filter maps DSCP to a 5G QoS Flow Identifier (QFI); and the SDAP layer maps each QFI to a Data Radio Bearer (DRB). Each DRB carries QoS attributes, including GBR flag, packet delay budget, packet error rate, and priority level, which the gNB scheduler may use but is not required to use. At the DS-TT, the reverse mapping reconstructs the original TSN priority.

For time synchronisation, the 5G system implements IEEE 802.1AS transparent-clock behaviour [6]. The NW-TT timestamps each gPTP Sync message at ingress, while the DS-TT computes the 5G residence time and adds it to the gPTP correction field. This mechanism is effective only if residence time and its variance remain bounded. Bounding delay and jitter over a wireless link is a long-standing difficulty across wireless technologies, arising from variable propagation delay, channel asymmetry, interference, and clock drift, and TSN-5G/6G synchronisation interoperability remains an identified open challenge [7]. Thus, although the 3GPP bridge specifies the QoS and synchronisation interfaces, the MAC scheduler ultimately determines how shared radio resources are allocated among DRBs and endpoints. This implementation choice is the source of the performance variability studied here.

B. MAC Scheduling Disciplines

We evaluate four Simu5G scheduling disciplines [8], chosen to represent distinct radio-resource allocation strategies.

a) Maximum C/I (MaxCI): MaxCI assigns resource blocks to the UE with the highest instantaneous CQI. It maximises instantaneous throughput but can create spatial unfairness because UEs with persistently weaker channels receive fewer resources. MaxCI is unaware of DRB QoS attributes and treats all traffic within a UE identically.

b) Proportional Fair (PF): PF schedules the UE that maximises the ratio of instantaneous achievable rate to exponentially averaged past rate [16]. This balances throughput and fairness, but like MaxCI, it does not use DRB-level QoS attributes and therefore cannot distinguish traffic classes within a UE.

c) Deficit Round Robin (DRR): DRR allocates a fixed quantum of resource blocks to each UE in cyclic order, carrying unused deficit to later rounds [17]. It provides strict per-UE fairness and is channel-unaware. In the evaluated implementation, DRR does not distinguish among DRBs within a UE, so control, vision, telemetry, and gPTP traffic compete within the same per-user allocation.

d) QoS-aware Proportional Fair (QoS-PF): QoS-PF extends PF by weighting each connection according to DRB attributes. In the Simu5G implementation, the weight for DRB d is

$$w_d = (\mathrm{GBR}_d \ ?\ 2 : 1) \times \frac{10}{p_d + 1} \times b(D_d), \tag{1}$$

where p_d is the priority level, D_d is the delay budget, and $b(D_d) \in \{5.0, 3.0, 1.5\}$ for delay budgets ≤ 10, ≤ 50, and ≤ 100 ms, respectively. The scheduling score is $s = w_d r_{\text{inst}}/r_{\text{avg}}$, using the same instantaneous and averaged rates as PF. QoS-PF is therefore the only evaluated scheduler that can explicitly prioritise high-priority DRBs with tight delay budgets. Further details are provided in [13].

C. Related Work

Several simulation frameworks study 5G-TSN integration, but they differ substantially in radio fidelity, QoS mapping, scaling scope, and synchronisation support. 5GTQ [2] implements NW-TT and DS-TT modules in OMNeT++/Simu5G with QoS-aware scheduling, but maps traffic priority directly at the MAC using 5QI values rather than SDAP-based per-flow DRB selection; it also does not model gPTP residence time or

TABLE I
COMPARISON WITH EXISTING 5G-TSN SIMULATION STUDIES.

Study	Radio[a]	DRB[b]	Sched.[c]	Scale[d]	gPTP[e]	Fade[f]
5GTQ [2]	✓	×	○	×	×	×
Da Silva [3]	✓	×	×	×	✓	×
6GDetCom [4]	×	×	×	×	×	×
P5G-TSN [5]	✓	×	○	×	×	○
Zanbouri [18]	✓	×	×	✓	×	✓
CQDRR [19]	○	×	✓	×	×	✓
nascTime [14]	✓	✓	×	×	✓	✓
This work	✓	✓	✓	✓	✓	✓

✓ = supported, × = not supported, ○ = partial. [a]5G radio with MAC scheduler, HARQ, channel model, RB allocation. [b]Per-flow QFI-to-DRB mapping via SDAP. [c]Systematic comparison of ≥3 scheduling disciplines. [d]Endpoint count swept from single-UE to cell-capacity limit. [e]gPTP transport through simulated radio with residence-time measurement. [f]Small-scale fading model evaluated.

scaling beyond a small number of traffic classes. Da Silva et al. [3] extend INET's IEEE 802.1AS model to transport gPTP over 5G and demonstrate sub-microsecond synchronisation accuracy, but do not include data-plane forwarding and therefore cannot capture contention between TSN data traffic and synchronisation traffic. 6GDetCom [4] models deterministic communication across heterogeneous domains, but represents the wireless segment as a statistical delay element rather than a 5G RAN with MAC scheduling, HARQ, channel models, and resource-block allocation. P5G-TSN [5] adds TDD-pattern and resource-allocation analysis for private 5G, but does not provide SDAP-based per-flow DRB mapping.

The nascTime framework [14], used in this work, implements the full PCP→DSCP→QFI→DRB pipeline and transports gPTP through the simulated 5G radio path using L2-in-GTP-U encapsulation with per-message residence-time correction. Its initial evaluation, however, considers only a three-endpoint topology and does not study endpoint scaling or scheduler choice.

Prior scalability and scheduling studies are complementary but narrower than the present work. Zanbouri et al. [18] evaluate bounded delay as the number of factory devices increases, but do not compare MAC schedulers or study synchronisation under data-plane contention. Ginthör et al. [20] analyse multi-user scheduling for industrial 5G, but consider a single traffic class and do not use the SDAP/DRB QoS mapping pipeline. Zhang et al. [19] propose a channel- and queue-aware DRR variant for hybrid TSN-5G traffic; their focus is scheduler design, whereas we characterise existing schedulers under heterogeneous scaling. Larrañaga et al. [21] coordinate configured grants with TSN scheduling, addressing a different resource allocation mechanism from the dynamic schedulers evaluated here.

Table I summarises the gap. To our knowledge, no prior work jointly provides actual 5G radio simulation with MAC scheduling and channel models, SDAP-based per-flow DRB mapping, comparison of multiple schedulers, endpoint scaling to cell-capacity limits, gPTP analysis under data-plane contention, and both bandwidth and fading evaluation.

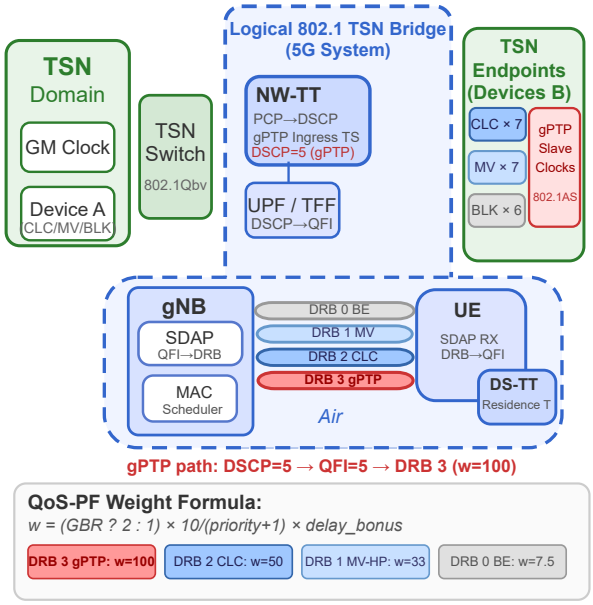


Fig. 1. 5G-TSN bridge architecture with four-bearer SDAP configuration. DRB 3 carries gPTP traffic at priority 0. QoS-PF weights are shown for each DRB.

III. SYSTEM ARCHITECTURE

The 3GPP bridge specification defines the QoS mapping pipeline and transparent-clock mechanism, but not how many Data Radio Bearers (DRBs) to provision, how traffic classes should be assigned to them, or whether the synchronisation plane should receive dedicated radio resources. These choices determine how the MAC scheduler observes and prioritises competing TSN flows. This section describes the bearer architecture and gPTP transport, shown in Fig. 1, which are used in our evaluation.

A. Four-Bearer SDAP Configuration

We configure four DRBs per UE, each carrying a distinct traffic class with QoS attributes matched to its TSN requirements (Table II).

Closed-loop control (CLC) has the tightest deadline (2 ms) and is assigned to DRB 2 with GBR treatment and priority 1. Machine-vision high-priority traffic (MV-HP) uses DRB 1 with a 10 ms delay budget. Machine-vision best-effort traffic (MV-BE) and bulk telemetry (BLK) share DRB 0 because their deadlines are looser and they do not require GBR service. gPTP traffic uses the dedicated DRB 3.

The TSN-to-5G QoS mapping is a pass-through pipeline: PCP 7/6/0 are mapped at the NW-TT to DSCP 7/6/0 for CLC, MV-HP, and best-effort traffic, respectively; the UPF Traffic Flow Filter maps DSCP i to QFI i; and SDAP routes each QFI to its configured DRB. This preserves TSN priority semantics without complex UPF policy rules.

Under QoS-PF, Eq. (1) gives a priority hierarchy: DRB 3 (gPTP) receives weight 100, followed by DRB 2 (CLC) at 50, DRB 1 (MV-HP) at 33.3, and DRB 0 (best effort)

TABLE II
SDAP BEARER CONFIGURATION.

DRB	Traffic class	QFI	GBR	Delay (ms)	Priority (0 = high)	w_d
0	Best effort (MV-BE, BLK)	0	No	50	3	7.5
1	MV high-priority	6	Yes	10	2	33.3
2	Closed-loop ctrl	7	Yes	2	1	50.0
3	gPTP (dedicated)	5	Yes	1	0	100.0

at 7.5. Under MaxCI, PF, and DRR, these weights are ignored: all DRBs within a UE are treated identically, with allocation driven by channel quality or cyclic order. Thus, the bearer configuration provides the mechanism for QoS differentiation, but only QoS-PF uses it. The experiments quantify the consequence of this scheduler–bearer asymmetry as endpoint count increases.

B. Dedicated gPTP Bearer

In the baseline 3GPP bridge path, gPTP frames share the user plane with data traffic. When encapsulated in GTP-U for Ethernet-type PDU sessions, the resulting IP packets carry no QoS marking by default and are routed to the best-effort bearer (DRB 0). In initial scaling experiments without a dedicated bearer, gPTP frames at $N = 20$ UEs and $10\,\mathrm{MHz}$ were delayed by data-plane contention; the IEEE 802.1AS slave clock servo then overcorrected, and in several runs the INET SettableClock oscillator compensation exceeded its valid range ($[-500{,}000,\ +1{,}000{,}000]$ ppm), terminating the simulation. This occurred under both PF and QoS-PF, indicating an architectural rather than scheduler-specific vulnerability.

We therefore assign gPTP to DRB 3. The NW-TT marks GTP-U-encapsulated gPTP packets with DSCP 5, the UPF maps DSCP 5 to QFI 5, and SDAP routes QFI 5 to DRB 3, configured with GBR service, a $1\,\mathrm{ms}$ delay budget, packet error rate 10^{-5}, and priority 0. Under QoS-PF this yields weight 100, so gPTP is scheduled ahead of data traffic whenever its UE receives a scheduling opportunity.

To retain failed cases for analysis, we modified INET SettableClock to clamp out-of-range oscillator-compensation values rather than terminate the simulation. This instrumentation does not affect normal operation; it records clock-servo stress in cases that would otherwise diverge. In the primary sweep with the dedicated bearer, no clamping events occur at $N \leq 10$; events appear only under deep saturation ($N = 20$, $10\,\mathrm{MHz}$) for endpoints with reduced scheduling frequency, as analysed in Section V-E.

The dedicated bearer provides *intra-UE* priority but not *inter-UE* scheduling guarantees. gPTP is highest priority within a scheduled UE, but a UE carrying only low-priority data may receive fewer MAC scheduling turns, delaying its gPTP frames despite DRB 3 priority. This coupling is evaluated in Section V. The capacity cost is small: gPTP generates one Sync/FollowUp pair every $125\,\mathrm{ms}$, each approximately 100 bytes, and the extra DRB uses one of the 29 DRBs permitted per PDU session [23].

TABLE III
5G NR RADIO CONFIGURATION.

Parameter	Value
Carrier frequency	$3.5\,\mathrm{GHz}$ (band n78)
Numerology (μ)	1 ($30\,\mathrm{kHz}$ SCS)
Slot duration	$0.5\,\mathrm{ms}$
Bandwidth A	25 RBs ($\approx 10\,\mathrm{MHz}$)
Bandwidth B	51 RBs ($\approx 20\,\mathrm{MHz}$)
Channel model	NrChannelModel (Indoor Hotspot)
Fading	Disabled; Jakes, 6 paths in fading sweep
Shadowing	Disabled
HARQ	Enabled (default Simu5G config.)

TABLE IV
TRAFFIC PROFILES AND TSN FLOW PARAMETERS.

Profile	Flow	Size (B)	Period	Deadline (ms)	DRB
CLC	clc_hp	100	$1\,\mathrm{ms}$ CBR	2	2
MV	mv_hp	1500	$5\,\mathrm{ms}$ CBR	10	1
	mv_be	500	exp($2\,\mathrm{ms}$)	50	0
BLK	blk_be	1200	exp($10\,\mathrm{ms}$)	100	0

IV. EXPERIMENTAL METHODOLOGY

A. Simulation Platform

Experiments use the nascTime framework [14], built on OMNeT++ 6.3, INET 4.6.0, and Simu5G v1.4.1 [12]. nascTime implements NW-TT and DS-TT modules integrated with INET's `LayeredEthernetInterface` and transports gPTP over the simulated 5G path using L2-in-GTP-U encapsulation with per-message residence-time measurement, as described in Section III-B.

The topology contains a TSN grandmaster, TSN switch, NW-TT connected to the 5G UPF, one gNB, N UEs each co-located with a DS-TT, and N downstream TSN endpoints. An upstream TSN device (Device A) hosts all traffic sources; each UE maps one-to-one to a downstream endpoint (Device B).

B. Radio Configuration

Table III lists the 5G NR parameters. We evaluate $10\,\mathrm{MHz}$ and $20\,\mathrm{MHz}$ bandwidths, corresponding to 25 and 51 RBs, respectively, in band n78 at $3.5\,\mathrm{GHz}$. Both use numerology $\mu = 1$ ($30\,\mathrm{kHz}$ SCS), giving a $0.5\,\mathrm{ms}$ slot duration. The primary experiments disable fading and shadowing to isolate scheduling and contention effects; the fading sweep enables six-path Jakes fading.

C. Traffic Model and Endpoint Scaling

The workload comprises closed-loop control (CLC), machine vision (MV), and bulk telemetry (BLK), derived from 3GPP TR 22.804 service requirements [22]. Table IV gives packet sizes, periods, deadlines, and DRB mappings.

CLC models a $1\,\mathrm{ms}$ periodic control loop and is mapped to DRB 2. MV contains a high-priority periodic stream on DRB 1 and a best-effort background stream on DRB 0. BLK models lower-priority telemetry on DRB 0. Each endpoint also

TABLE V
ENDPOINT PROFILE MIX AND AGGREGATE DOWNLINK OFFERED LOAD PER N. EACH MV ENDPOINT GENERATES BOTH MV-HP AND MV-BE FLOWS.

N	CLC	MV	BLK	Offered DL load (Mbps)
1	1	0	0	0.8
5	2	2	1	11.4
10	4	4	2	22.7
15	5	5	5	30.8
20	7	7	6	42.2
30	10	10	10	61.6
40	14	14	12	84.3

generates a reverse-direction CBR flow of 100 B every $10\,\mathrm{ms}$ ($0.08\,\mathrm{Mbit\,s^{-1}}$), excluded from the downlink load totals.

Endpoint counts range from $N = 1$ to 40 using the profile mix in Table V. Offered downlink load is computed as $L = 8S/T$, where S is packet size and T is the fixed period or mean inter-arrival time. Per-endpoint downlink loads are $0.8\,\mathrm{Mbit\,s^{-1}}$ for CLC, $4.4\,\mathrm{Mbit\,s^{-1}}$ for MV (MV-HP plus MV-BE), and $0.96\,\mathrm{Mbit\,s^{-1}}$ for BLK.

At $10\,\mathrm{MHz}$, effective downlink capacity is approximately $25\,\mathrm{Mbit\,s^{-1}}$; therefore $N = 10$ approaches capacity, while $N = 15$ and $N = 20$ exceed it by approximately 23% and 69%, respectively. At $20\,\mathrm{MHz}$, effective capacity is approximately $50\,\mathrm{Mbit\,s^{-1}}$; $N = 20$ remains below capacity, while $N = 30$ exceeds it by approximately 23%.

The protocol stack uses Simu5G v1.4.1 defaults unless stated otherwise: RLC UM, HARQ with up to three retransmissions, FDD duplexing, and $2\,\mathrm{MB}$ FIFO MAC buffers per logical channel (per DRB per UE). No application-layer deadline dropping is configured: packets remain queued until delivered, lost through RLC buffer overflow, or the simulation ends. Configuration files and analysis scripts are available in the nascTime repository [15].

D. Experiment Matrix

Each simulation runs for $30\,\mathrm{s}$ with a $2\,\mathrm{s}$ warmup; metrics are collected over $t = 2\,\mathrm{s}$ to $t = 30\,\mathrm{s}$. Each configuration uses three random seeds. Table VI summarises the four sweeps.

The primary sweeps compare MaxCI, PF, DRR, and QoS-PF across all endpoint counts and both bandwidths. The fading sweeps compare MaxCI and QoS-PF under ideal and Jakes channels to test whether fading alters the scheduler ranking. gPTP residence time and clock-servo data are collected in both primary and fading sweeps. A fifth Simu5G scheduler, MAXCI_COMP, was tested initially but produced results identical to MaxCI and is excluded.

E. Metrics and Packet Accounting

We report five metrics. Packet Delivery Ratio (PDR) is the number of unique application-layer packets received at Device B divided by the number generated at Device A, computed per flow class. End-to-end delay is measured from packet creation at Device A to reception at Device B using the receiver-side `endToEndDelay` vector; we report mean, P99, P99.9, and maximum. Deadline compliance is the fraction of (endpoint, seed) cells whose P99 delay does not exceed the flow's TSN deadline. Spatial fairness is the standard deviation of PDR across endpoints for a given $(N, \text{scheduler}, \text{flow class})$, averaged over seeds. gPTP residence time is extracted from the DS-TT `residenceTime` vector, and clock-servo stability from each Device B clock's `oscillatorCompensationChanged` vector.

TABLE VI
EXPERIMENT SWEEPS.

Sweep	BW	N values	Schedulers	Cells
Primary	10 MHz	1, 5, 10, 15, 20	4	60
Primary	20 MHz	1, 5, 10, 15, 20, 30, 40	4	84
Fading	10 MHz	1, 10, 20	2†	36
Fading	20 MHz	1, 10, 20	2†	36
Total simulation cells				216

†MaxCI and QoS-PF only, each with fading $\in$ {off, on}.

All metrics use the measurement window from $t = 2\,\mathrm{s}$ to $t = 30\,\mathrm{s}$. The PDR denominator is the number of unique application-layer packets generated in this window; the numerator is the number of matching unique packets received at the destination, whether or not they meet their TSN deadline. Delay statistics include all delivered packets, including late packets. Deadline compliance is therefore reported separately to distinguish "delivered eventually" from "delivered within deadline." Packets generated before the measurement window are excluded from both numerator and denominator.

V. RESULTS

We present results as a scaling narrative: no-contention baseline, saturation onset, overload, bandwidth sensitivity, fading resilience, and time synchronisation. The analysis focuses on MaxCI, PF, DRR, and QoS-PF; MAXCI_COMP is excluded as noted in Section IV-D. The comparison should be read as evaluating the benefit of exposing TSN class information to the MAC scheduler, not as a universal ranking of scheduling algorithms.

A. Baseline: The Latency Floor

At $N = 1$, a single CLC endpoint operates without contention. All four schedulers deliver 28,000 packets, i.e., 100% of the offered load, with mean and P99 delay of $2.25\,\mathrm{ms}$. This effective latency floor arises from the $0.5\,\mathrm{ms}$ slot duration at $\mu = 1$, scheduling-request/grant processing, and PDCP/RLC handling.

The $2\,\mathrm{ms}$ CLC deadline is an application requirement drawn from TR 22.804 motion-control profiles rather than a tunable parameter of the study. Retaining it while evaluating a standard mid-band, grant-based configuration is deliberate: the gap between the application requirement and the achievable floor is itself a deployment-relevant result. The floor is a property of the radio configuration and is unaffected by scheduler choice, so meeting deadlines below approximately $3\,\mathrm{ms}$ would

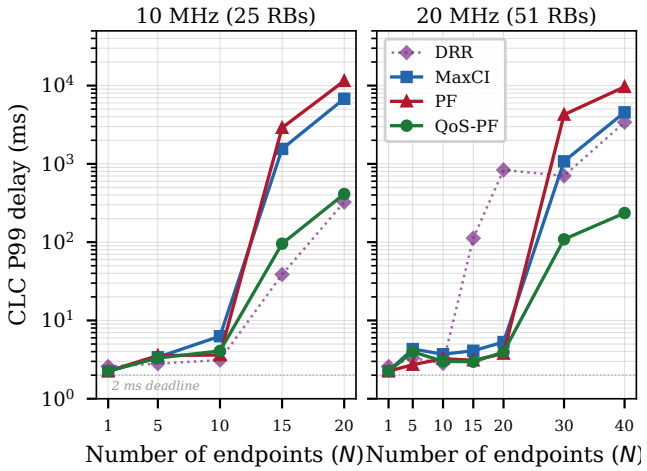


Fig. 2. CLC P99 end-to-end delay vs. endpoint count N at 10 MHz and 20 MHz. Logarithmic y-axis. The 2 ms deadline is shown as a dotted reference. DRR is omitted from the 10 MHz panel at $N \geq 15$ to preserve readability.

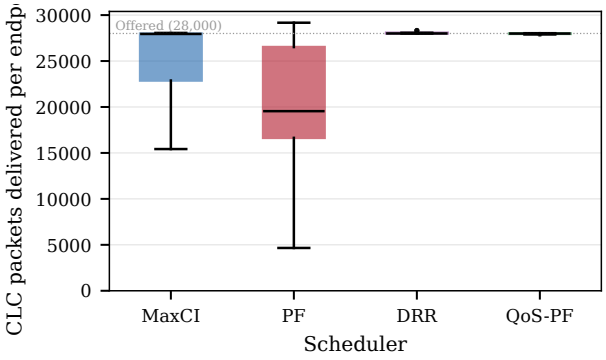


Fig. 3. Distribution of per-endpoint CLC delivery at $N = 20$, 10 MHz. Each box spans the interquartile range across endpoints and repetitions. MaxCI and PF exhibit high variance; QoS-PF and DRR maintain tight distributions.

TABLE VII
PERFORMANCE AT $N = 20$, 10 MHz. ALL DELIVERY COUNTS ARE MEAN PER-ENDPOINT VALUES AVERAGED ACROSS REPETITIONS. PDR IS COMPUTED PER ENDPOINT AND AVERAGED. FAIRNESS IS THE STANDARD DEVIATION OF PER-ENDPOINT DELIVERY ACROSS ENDPOINTS.

Metric	MaxCI	PF	DRR†	QoS-PF
CLC delivery	21,698	16,538	28,084	27,955
CLC PDR (%)	77.5	59.1	100.0	99.8
CLC mean (ms)	3,634	4,923	78	54
CLC P99 (ms)	6,790	11,472	324	412
CLC fairness (std)	5,756	6,961	54	54
MV-HP delivery	3,142	3,549	120	3,074
MV-HP PDR (%)	56.1	63.4	2.1	54.9
BLK delivery	2,797	1,929	214	2,382
BLK PDR (%)	99.9	68.9	7.6	85.1

†DRR CLC delivery count exceeds the nominal offered count due to measurement-window boundary effects; PDR is capped at 100.0%.

require configured grants, pre-emptive scheduling, or higher numerology rather than a different scheduling discipline.

Consequently, the evaluated sub-6 GHz, 30 kHz-SCS configuration does not satisfy strict CLC deadline compliance even under zero contention, and CLC deadline compliance is zero for every scheduler at every N. Subsequent scheduler comparisons therefore measure *relative robustness under load*, quantified through delivery ratio, P99 delay, and spatial fairness, rather than strict TSN deadline satisfaction.

B. Scaling at 10 MHz

Fig. 2 shows CLC P99 delay versus endpoint count. At 10 MHz, three regimes emerge.

1) Below Saturation ($N \leq 10$): At $N = 5$, all non-DRR schedulers deliver 100% of all traffic classes. CLC P99 ranges from 3.31 ms (QoS-PF) to 3.56 ms (PF), and MV-HP delivery is 100%. BLK delivery ranges from 2,818 (PF) to 2,895 (MaxCI) out of approximately 2,800 offered, reflecting the variance of the exponential arrival process. Scheduler choice has limited impact in this regime.

At $N = 10$, CLC delivery remains 100% for all non-DRR schedulers, with P99 ranging from 3.62 ms (PF) to 6.29 ms (MaxCI). PF achieves the lowest CLC latency at this operating point, while MV-HP delivery remains 100%.

DRR already shows its characteristic pathology. At $N = 5$, it delivers only 121 of 5,600 MV-HP packets (2.2%) and 215 of approximately 2,800 BLK packets (7.7%), despite the cell operating below capacity. MV-HP and BLK mean delays reach 4.4 s and 9.4 s, respectively, while CLC remains at 100% delivery with 2.33 ms mean delay. This starvation is structural: the fixed per-UE quantum accommodates CLC's 800 kbit s^{-1} CBR flow but cannot serve burstier MV and BLK traffic within a single round-robin turn.

2) Saturation Onset ($N = 15$): Between $N = 10$ and $N = 15$, offered load increases from 22.7 Mbit s^{-1} to 30.8 Mbit s^{-1}, exceeding the effective 25 Mbit s^{-1} cell capacity by approximately 23%. Scheduler rankings then diverge sharply.

QoS-PF maintains CLC delivery at 99.9% (27,984 of 28,000), with mean delay of 16 ms and P99 of 96 ms. MaxCI delivers 95.2% of CLC (26,658) with P99 of 1.55 s; PF delivers 89.3% (25,010) with P99 of 2.89 s. Thus, at saturation onset, QoS-PF reduces CLC P99 delay by one to two orders of magnitude relative to MaxCI and PF.

Spatial fairness explains this divergence. The standard deviation of CLC delivery across endpoints is 31 packets for QoS-PF, compared with 1,873 for MaxCI and 2,233 for PF. MaxCI favours endpoints with better channel quality, while PF distributes resources without flow-class awareness and therefore does not prioritise CLC. MV-HP delivery also begins to separate: QoS-PF delivers 98.0% (5,489 of 5,600), MaxCI 86.6% (4,851), and PF 87.4% (4,894).

3) Overload ($N = 20$): At $N = 20$, offered load substantially exceeds capacity and scheduler differences become stark. Table VII summarises the main metrics.

QoS-PF delivers 99.8% of CLC traffic with P99 delay of 412 ms and delivery standard deviation of 54 packets. PF delivers only 59.1% of CLC with P99 of 11.5 s and fairness

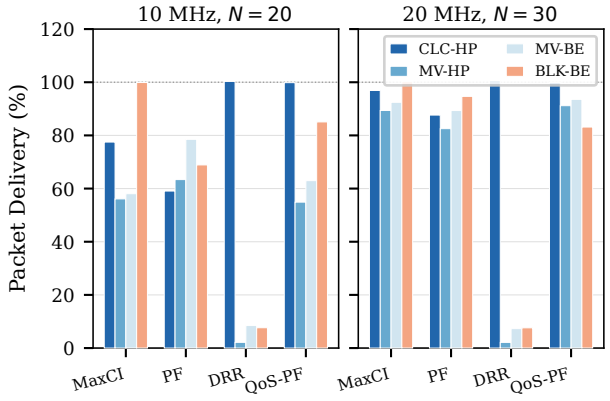


Fig. 4. Packet delivery ratio per flow class under saturation. Left: 10 MHz, $N = 20$. Right: 20 MHz, $N = 30$. DRR delivers CLC at 100% but starves MV-HP and BLK.

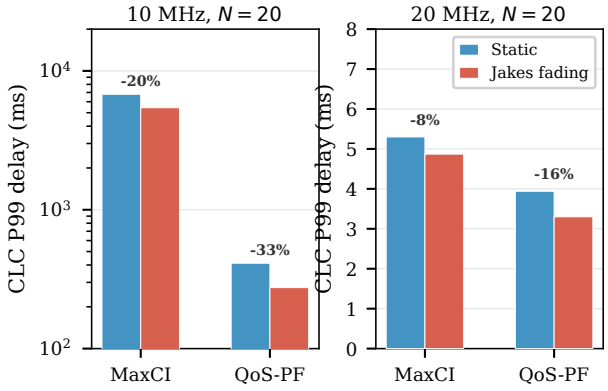


Fig. 5. Effect of Jakes fading on CLC P99 delay at $N = 20$. Left: 10 MHz saturated case. Right: 20 MHz below-saturation case. Percentage labels show change relative to the static channel.

TABLE VIII
CROSS-BANDWIDTH COMPARISON FOR CLC FLOW. CORRESPONDING OPERATING POINTS ARE PAIRED BY APPROXIMATE LOAD-TO-CAPACITY RATIO.

		Saturation onset		Overload	
		10 MHz N=15	20 MHz N=30	10 MHz N=20	20 MHz N=40
PDR (%)	MaxCI	95.2	96.9	77.5	85.5
	PF	89.3	87.6	59.1	76.8
	QoS-PF	99.9	99.9	99.8	100.0
P99 (ms)	MaxCI	1,552	1,078	6,790	4,553
	PF	2,889	4,256	11,472	9,646
	QoS-PF	96	109	412	236
Fairness (std)	MaxCI	1,873	1,965	5,756	4,942
	PF	2,233	3,282	6,961	6,412
	QoS-PF	31	45	54	75

standard deviation of 6,961, indicating severe endpoint-level imbalance. Fig. 3 shows the same pattern: MaxCI and PF have wide delivery distributions, while QoS-PF and DRR remain tight.

DRR preserves the CLC stream almost completely in the measured window, but only by starving other classes: MV-HP delivery is 2.1% and BLK delivery is 7.6%. Fig. 4 visualises this per-class delivery trade-off at the saturation points for both bandwidths. Thus, the evaluated per-UE DRR implementation is unsuitable for heterogeneous TSN bridging.

QoS-PF protects the highest-priority class through overload at the cost of lower service to less critical flows. Its MV-HP delivery (54.9%) is below PF's (63.4%), reflecting deliberate resource reallocation from MV to CLC. This trade-off is appropriate only when the deployment priority structure favours CLC preservation.

C. Bandwidth Sensitivity

Fig. 2 (right panel) shows CLC P99 delay at 20 MHz. The same three-regime structure appears, but the saturation threshold shifts to higher N.

At $N = 20$, all non-DRR schedulers deliver 100% of all traffic classes. CLC P99 ranges from 3.77 ms (PF) to 5.30 ms (MaxCI), comparable to the 10 MHz, $N = 5$ case. PF again marginally outperforms QoS-PF below saturation (3.77 ms versus 3.94 ms), suggesting that QoS weighting introduces a small cost that is only justified under contention. This cost is modest (typically $<$1.5 ms in P99) and disappears near saturation.

At $N = 30$, saturation onset is visible. QoS-PF maintains 99.9% CLC delivery with P99 of 109 ms; MaxCI drops to 96.9% with P99 of 1.078 ms; and PF drops to 87.6% with P99 of 4.256 ms. CLC fairness standard deviation is 45 packets for QoS-PF, compared with 1,965 for MaxCI and 3,282 for PF. At $N = 40$, QoS-PF still delivers 99.96% of CLC with P99 of 236 ms, while MaxCI and PF fall to 85.5% and 76.8% delivery with P99 delays of 4.553 ms and 9.646 ms, respectively.

Across the two evaluated bandwidths, the saturation threshold approximately doubles when bandwidth doubles. Table VIII shows that $N = 15$ at 10 MHz and $N = 30$ at 20 MHz produce comparable delivery, P99, and fairness values; similarly, $N = 20$ at 10 MHz mirrors $N = 40$ at 20 MHz. This proportional relationship can guide capacity estimation for the evaluated workload, although more bandwidth points would be needed to establish a general scaling law.

D. Fading Resilience

Table IX compares CLC P99 delay under static and Jakes fading channels for MaxCI and QoS-PF.

Two observations emerge. First, QoS-PF achieves lower CLC P99 than MaxCI under every tested channel condition, so the scheduler ranking is not contingent on ideal channels. Second, fading is associated with lower latency in these experiments. At 10 MHz, $N = 20$, QoS-PF P99 decreases from 412 ms to 275 ms (33%), and MaxCI decreases from 6.790 ms to 5.431 ms (20%). At 20 MHz, where the cell is below saturation, the effect is smaller than 15%.

This behaviour is consistent with multi-user diversity: under fading, CQI varies across TTIs, allowing the scheduler to exploit favourable channel instances. The effect is most visible

TABLE IX
CLC P99 DELAY (MS) UNDER STATIC AND JAKES FADING CHANNELS.

		MaxCI		QoS-PF	
BW	N	Static	Fading	Static	Fading
10 MHz	1	2.25	2.25	2.25	2.25
	10	6.29	5.29	4.08	3.58
	20	6,790	5,431	412	275
20 MHz	1	2.25	2.25	2.25	2.25
	10	3.71	3.42	3.00	2.75
	20	5.30	4.87	3.94	3.30

under saturation, where resource efficiency directly affects queueing. Fig. 5 visualises this contrast at $N = 20$: the reduction is substantial in the saturated $10\,\mathrm{MHz}$ case and modest in the below-saturation $20\,\mathrm{MHz}$ case. However, this observation is based on one fading model (Jakes, six paths), two bandwidths, and three endpoint counts. Broader channel, mobility, and geometry studies are required before generalising the result. CLC, MV-HP, and BLK delivery are unaffected by fading in the tested configurations.

E. Time Synchronisation

The dedicated gPTP bearer (DRB 3, Section III-B) prevents the clock-servo divergence observed in initial experiments: all simulation cells complete without oscillator-compensation errors. However, gPTP residence time still degrades under saturation.

1) Residence Time Under Scaling: At $N = 1$, 5G residence time is $2{,}772 \pm 249$ μs (mean $\pm$ std), with maximum $3.000\,\mu\mathrm{s}$, and is stable across schedulers.

At $N = 10$ without fading, residence time increases modestly. Under MaxCI, mean per-DS-TT residence time ranges from $3.500\,\mu\mathrm{s}$ to $5.000\,\mu\mathrm{s}$, with standard deviations of $1.200\,\mu\mathrm{s}$ to $2.500\,\mu\mathrm{s}$. Under QoS-PF, the range narrows to $3.000\,\mu\mathrm{s}$–$3.500\,\mu\mathrm{s}$, with standard deviations of $230\,\mu\mathrm{s}$–$1.080\,\mu\mathrm{s}$. QoS-PF therefore produces 2–5$\times$ lower residence-time variance, consistent with preferential scheduling of DRB 3.

At $N = 20$ under saturation at $10\,\mathrm{MHz}$, residence times diverge by endpoint type. CLC endpoints, which receive frequent scheduling turns because of high-weight DRB 2 traffic, show residence times of $15\,\mathrm{ms}$–$135\,\mathrm{ms}$. MV endpoints, which carry lower-priority data and receive fewer scheduling turns, show residence times of $394\,\mathrm{ms}$–$4.439\,\mathrm{ms}$, three orders of magnitude higher.

2) The Inter-UE Scheduling Coupling: This disparity shows that a dedicated bearer provides intra-UE priority but not inter-UE scheduling guarantees. DRB 3 makes gPTP the highest-priority traffic *within* a scheduled UE. Under QoS-PF, however, UEs carrying CLC on DRB 2 (weight 50) receive more scheduling opportunities than UEs carrying MV or best-effort traffic on lower-weight DRBs. As a result, gPTP residence time depends on the endpoint's data traffic class, even though all endpoints use the same gPTP bearer.

Clock-servo data confirm this coupling. At $N = 10$ under fading, oscillator-compensation standard deviation is approximately 4,200 ppm, with corrections from $-24{,}000$ to $+24{,}000$ ppm. At $N = 20$ under saturation, MV-type endpoints exceed 190,000 ppm standard deviation, with peaks approaching the $1{,}000{,}000$ ppm SettableClock clamping bound.

Thus, complete gPTP protection under saturation requires more than a dedicated bearer: it also requires a minimum scheduling frequency per UE, for example via semi-persistent scheduling or configured grants. We return to this issue in Section VI.

VI. DISCUSSION

A. Synthesis of Findings

The results show that 5G-TSN bridge performance is governed by the interaction between endpoint load, bearer architecture, and MAC scheduling at a load-dependent operating point, rather than by endpoint count or scheduler choice alone.

a) Scaling behaviour: The transition from adequate to degraded performance is abrupt. At $10\,\mathrm{MHz}$, CLC P99 latency increases by two orders of magnitude between $N = 10$ ($4\,\mathrm{ms}$) and $N = 15$ ($96\,\mathrm{ms}$ under QoS-PF; $1.5\,\mathrm{s}$ under MaxCI). Thus, a deployment that performs acceptably below saturation may degrade qualitatively after only a small increase in endpoints or offered load. This sharp transition follows from fixed radio capacity: once offered load approaches the scheduling limit, queues grow rapidly and non-prioritised traffic is affected first.

b) Scheduler selection: In the evaluated configuration, QoS-aware scheduling offers the best risk trade-off. Below saturation, QoS-PF's cost relative to PF is small ($<1.5\,\mathrm{ms}$ in CLC P99); above saturation, it is the only tested scheduler that maintains near-complete delivery for the highest-priority traffic class. This supports using QoS-aware scheduling for heterogeneous TSN traffic, while recognising that the result applies to the tested Simu5G implementations.

The evaluated per-UE DRR implementation should not be used for heterogeneous TSN bridging. Its per-UE fairness model cannot distinguish a $1\,\mathrm{ms}$ control loop from bulk telemetry within the same UE, causing starvation of bursty flows independently of utilisation. Channel- or queue-aware DRR variants such as CQDRR [19] may mitigate this limitation but are outside this evaluation.

c) Bandwidth planning: Across the two evaluated bandwidths, the saturation threshold approximately doubles when bandwidth doubles. For the workload studied here, this corresponds to roughly $N = 12$ per $10\,\mathrm{MHz}$. This provides an initial planning heuristic, not a general scaling law: additional bandwidths and traffic mixes would be needed for broader validation.

d) Time synchronisation: Isolating gPTP on a dedicated high-priority bearer prevented the clock-servo divergence observed without it: at $N = 20$, 4 of 15 cells terminated without the bearer, while all cells completed with it. However, bearer isolation alone does not bound gPTP residence time under saturation. Endpoints carrying lower-priority data receive fewer MAC scheduling turns, delaying their gPTP traffic despite its intra-UE priority. This inter-UE scheduling coupling suggests that robust synchronisation may require a

minimum scheduling frequency per UE, implemented through mechanisms such as semi-persistent scheduling, configured grants, or priority handling before MAC scheduling. It is an instance of the broader resource and deployment constraints that shape wireless synchronisation across technologies [7], and our results indicate that bearer-level prioritisation alone is insufficient to satisfy them.

B. Deployment Guidelines

The following guidelines are scoped to the evaluated simulation platform, traffic mix, and scheduler implementations, and should be validated for the target deployment:

1) **Use or evaluate QoS-aware scheduling.** QoS-PF provides the strongest protection for critical traffic under saturation with negligible cost below saturation.
2) **Provision a dedicated gPTP bearer.** A high-priority DRB for gPTP prevents the observed clock-servo divergence, although residence-time inflation can persist for scheduling-starved endpoints.
3) **Plan around the saturation threshold.** The transition to overload is abrupt; deployments should characterise the threshold for their traffic mix and maintain margin below it.
4) **Account for the radio latency floor.** In the evaluated grant-based $30\,\mathrm{kHz}$-SCS Simu5G configuration, the effective floor is approximately $2.25\,\mathrm{ms}$; tighter deadlines may require configured grants, pre-emptive scheduling, or higher numerology.
5) **Monitor synchronisation per endpoint class.** Lower-priority endpoints can experience elevated gPTP residence time under saturation even with a dedicated bearer, so per-endpoint synchronisation monitoring is advisable.

C. Limitations

Several limitations scope the applicability of these findings.

a) Single cell and no mobility: All experiments use one gNB and stationary UEs. Multi-cell interference, handover, coordinated scheduling, and mobility may change both the saturation threshold and scheduler ranking.

b) Simulation fidelity: Results reflect Simu5G's scheduler implementations, INET's protocol stack, and the NrChannelModel channel model. Commercial gNBs may differ in scheduling, buffering, HARQ, and queue management. The $2.25\,\mathrm{ms}$ latency floor is therefore a property of this configuration, not a universal 5G limit.

c) Traffic model: The workload is a specific control/vision/telemetry mix derived from 3GPP TR 22.804. Correlated bursts, event-driven traffic, or different uplink/downlink ratios may shift the operating regimes.

d) Fading scope: The fading analysis uses one Jakes six-path model without mobility. The observed latency reduction under fading may not generalise to channels with deep fades, shadowing, mobility, or frequency selectivity.

e) No TSN-side scheduling: The TSN domain uses FIFO forwarding without IEEE 802.1Qbv Time-Aware Shaping or IEEE 802.1Qav Credit-Based Shaping. Coordinating TSN-side and 5G scheduling remains outside the scope of this study.

f) Statistical scope: Each configuration uses three seeds. Each delay percentile is nonetheless estimated from tens of thousands of packets per run, and the scheduler differences reported here span one to two orders of magnitude. Three seeds therefore support the ordinal claims (scheduler ranking, regime boundaries) but not cardinal claims about P99.9 or tail shape, which we do not make.

g) Scaling confounds: The scaling sweep varies endpoint count, offered load, and traffic-class composition together according to Table V. The operating regimes should therefore be interpreted for this workload; isolating endpoint count from load and class mix requires additional controlled experiments.

VII. Conclusion

This paper presented a scalability study of 5G-TSN bridges under heterogeneous industrial traffic, evaluating four MAC schedulers across up to 40 endpoints, two bandwidths, and ideal and Jakes fading channels. The results show three operating regimes separated by a saturation threshold: below saturation, non-DRR schedulers behave similarly; near saturation, QoS-aware scheduling sharply reduces critical-flow latency; and under overload, QoS-PF is the only evaluated scheduler that maintains near-complete delivery for the highest-priority traffic class. Across the two evaluated bandwidths, the saturation threshold approximately doubles when bandwidth doubles.

The study also shows that IEEE 802.1AS/gPTP traffic benefits from a dedicated high-priority radio bearer, which prevents the clock-servo divergence observed when synchronisation traffic shares the default bearer. However, bearer isolation alone does not bound gPTP residence time under saturation: endpoints carrying lower-priority data receive fewer MAC scheduling opportunities, coupling synchronisation quality to data-plane priority. Finally, the relative scheduler ranking is preserved under the tested Jakes fading configuration, where fading is associated with reduced critical-flow latency under saturation, consistent with multi-user diversity.

For the evaluated configuration, these findings support QoS-aware MAC scheduling, dedicated gPTP bearer provisioning, capacity planning below the saturation threshold, and explicit consideration of the radio latency floor when setting TSN deadlines. Future work will extend the analysis to multi-cell deployments, mobility, additional channel models, higher numerology, and configured-grant or semi-persistent scheduling for gPTP protection under saturation.

Acknowledgment

This publication has emanated from research conducted with the financial support of Research Ireland under Grant number 13/RC/2077_P2. For the purpose of Open Access, the author has applied a CC-BY public copyright license to any Author Accepted Manuscript version arising from this submission

REFERENCES

[1] 3GPP, "System architecture for the 5G system (5GS)," 3GPP TS 23.501, Rel. 16, 2020.

[2] R. Debnath, M. S. Akinci, D. Ajith, and S. Steinhorst, "5GTQ: QoS-aware 5G-TSN simulation framework," in *Proc. IEEE 98th Veh. Technol. Conf. (VTC2023-Fall)*, Hong Kong, 2023, pp. 1–7, doi: 10.1109/VTC2023-Fall60731.2023.10333602.

[3] S. R. B. da Silva, F. G. Vogt, and C. E. Rothenberg, "Towards TSN-5G integration: Simulating time synchronization through 5G via OMNeT++," in *Proc. Braz. Symp. Comput. Netw. Distrib. Syst. (SBRC)*, 2024.

[4] L. Haug, F. Dürr, S. Egger, E. Mostovaya, J. Gross, G. Sharma, and J. Sachs, "A data-driven simulation framework for logical 5G-TSN bridges," in *Proc. Int. Conf. Networked Syst. (NetSys)*, Ilmenau, Germany, 2025, pp. 21–24, doi: 10.22032/dbt.67110.

[5] L. Becker and W. Kellerer, "P5G-TSN: A private 5G TSN simulation framework," in *Proc. KuVS Fachgespräch Würzburg Workshop 6G Netw. (WueWoWAS'24)*, 2024.

[6] *IEEE Standard for Local and Metropolitan Area Networks—Timing and Synchronization for Time-Sensitive Applications*, IEEE Std 802.1AS-2025, Dec. 2025.

[7] M. Seliem, U. Roedig, M. Ahmed, R. Muzaffar, D. Hamidovic, A. Hadziaganovic, C. J. Sreenan, and D. Pesch, "Wireless clock synchronization: A comprehensive survey and taxonomy," *Proc. IEEE*, vol. 114, no. 3, pp. 425–460, Mar. 2026, doi: 10.1109/JPROC.2026.3705899.

[8] G. Nardini, D. Sabella, G. Stea, P. Thakkar, and A. Virdis, "Simu5G—An OMNeT++ library for end-to-end performance evaluation of 5G networks," *IEEE Access*, vol. 8, pp. 181176–181191, 2020.

[9] A. Varga and R. Hornig, "An overview of the OMNeT++ simulation environment," in *Proc. 1st Int. Conf. Simulation Tools Techn. Commun., Netw. Syst. Workshops (SIMUTools)*, 2008, pp. 1–10, doi: 10.4108/ICST.SIMUTOOLS2008.3027.

[10] Z. Satka, M. Ashjaei, H. Fotouhi, M. Daneshtalab, M. Sjödin, and S. Mubeen, "A comprehensive systematic review of integration of time sensitive networking and 5G communication," *J. Syst. Archit.*, vol. 138, 2023, Art. no. 102852.

[11] Z. Satka, M. Ashjaei, H. Fotouhi, M. Daneshtalab, M. Sjödin, and S. Mubeen, "QoS-MAN: A novel QoS mapping algorithm for TSN-5G flows," in *Proc. IEEE 28th Int. Conf. Embedded Real-Time Comput. Syst. Appl. (RTCSA)*, Taipei, Taiwan, 2022, pp. 220–227.

[12] M. Seliem, U. Roedig, C. Sreenan, and D. Pesch, "SDAP-based QoS flow multiplexing support in Simu5G for 5G NR simulation," in *Proc. IEEE 30th Int. Workshop Comput. Aided Model. Design Commun. Links Netw. (CAMAD)*, Tempe, AZ, USA, 2025, pp. 1–8.

[13] M. Seliem, U. Roedig, C. Sreenan, and D. Pesch, "QoS-aware proportional fairness scheduling for multi-flow 5G UEs: A smart factory perspective," in *Proc. Int. Conf. Model., Anal. Simulation Wireless Mobile Syst. (MSWiM)*, Barcelona, Spain, 2025, pp. 20–27.

[14] M. Seliem, U. Roedig, C. Sreenan, and D. Pesch, "nascTime: A full-stack 5G-TSN bridge simulation framework with SDAP-based QoS mapping and IEEE 802.1AS transparent clock," in *Proc. IEEE 31st Int. Conf. Emerg. Technol. Factory Autom. (ETFA)*, 2026, to be published.

[15] M. Seliem, U. Roedig, C. Sreenan, and D. Pesch, "nascTime: 5G-TSN bridge simulation framework—scalability study configuration files and analysis scripts," 2026. [Online]. Available: https://github.com/MohamedSeliem/nascTime-5gtsn

[16] R. Kwan, C. Leung, and J. Zhang, "Proportional fair multiuser scheduling in LTE," *IEEE Signal Process. Lett.*, vol. 16, no. 6, pp. 461–464, Jun. 2009, doi: 10.1109/LSP.2009.2016449.

[17] M. Shreedhar and G. Varghese, "Efficient fair queuing using deficit round-robin," *IEEE/ACM Trans. Netw.*, vol. 4, no. 3, pp. 375–385, Jun. 1996, doi: 10.1109/90.502236.

[18] K. Zanbouri, M. Noor-A-Rahim, and D. Pesch, "Scalability analysis of 5G-TSN applications in indoor factory settings," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Milan, Italy, 2025, pp. 1–6, doi: 10.1109/WCNC61545.2025.10978486.

[19] X. Zhang and Y. Lu, "Asynchronous channel-aware and queue-aware deficit round robin scheduling strategy for different TSN traffics in the TSN-5G network," *Electron. Lett.*, vol. 61, no. 1, 2025, Art. no. e70418, doi: 10.1049/ell2.70418.

[20] D. Ginthör, J. von Hoyningen-Huene, R. Guillaume, and H. Schotten, "Analysis of multi-user scheduling in a TSN-enabled 5G system for industrial applications," in *Proc. IEEE Int. Conf. Ind. Internet (ICII)*, Orlando, FL, USA, 2019, pp. 190–199, doi: 10.1109/ICII.2019.00044.

[21] A. Larrañaga, M. C. Lucas-Estañ, I. Martinez, and J. Gozalvez, "5G configured grant scheduling for 5G-TSN integration for the support of Industry 4.0," in *Proc. 18th Wireless On-Demand Netw. Syst. Services Conf. (WONS)*, Madonna di Campiglio, Italy, 2023, pp. 72–79, doi: 10.23919/WONS57325.2023.10062219.

[22] 3GPP, "Study on communication for automation in vertical domains," 3GPP TR 22.804, V16.3.0, Mar. 2020.

[23] 3GPP, "NR; NR and NG-RAN overall description; Stage-2," 3GPP TS 38.300, V15.9.0, Apr. 2020.