
Public Review for
A DualPI2 Module for Mahimahi:
Behavioral Characterization and
Cross-Platform Analysis

Nawel Alioua, Linghe Zhang, Aneesh Garg, Francis Y. Yan,
Elizabeth Belding

The authors of this paper present a module for the Mahimahi network emulator that allows to run latency experiments on the Low Latency, Low Loss, and Scalable Throughput (L4S) mechanism. The module is based on the DualPI2 active queue management and runs in user-space. The authors analyze the statistical characteristics of their implementation using several metrics and explore to which degree emulated and real execution in the kernel produce aligned outcomes. The paper received two reviews, both of which recommended to accept the paper after adding some further explanations and limitations, including a clear description of the conditions under which the Mahimahi model is truly useful for other researchers. The result is a paper that should be of interest to researchers in computer networking interested in the L4S paradigm when they are contemplating whether an emulator can be helpful in their work.

Public review written by
Ralph Holz
University of Twente, Netherlands

A DualPI2 Module for Mahimahi: Behavioral Characterization and Cross-Platform Analysis

Nawel Alioua
UC Santa Barbara
nawel@ucsb.edu

Linghe Zhang
UC Santa Barbara
linghe@ucsb.edu

Aneesh Garg
UC Santa Barbara
aneeshgarg@ucsb.edu

Francis Y. Yan
University of Illinois
Urbana-Champaign
fyy@illinois.edu

Elizabeth Belding
UC Santa Barbara
ebelding@ucsb.edu

ABSTRACT

Low Latency, Low Loss, and Scalable Throughput (L4S) is an emerging paradigm for latency control based on DualPI2 active queue management and scalable congestion control. While a Linux kernel implementation of DualPI2 is available, controlled and reproducible experimentation on L4S mechanisms can be facilitated by a modular, user-space alternative. In this paper, we present a DualPI2 module for the Mahimahi network emulator, designed to support extensible, component-level experimentation without kernel modification. We conduct a statistical behavioral characterization of the Mahimahi implementation by examining key metrics across diverse traffic patterns and network conditions, using the Linux kernel implementation as a reference baseline. Our analysis shows that behavioral alignment across execution environments is not automatic: identical DualPI2 parameterization does not guarantee identical dynamics. Instead, key control parameters exhibit environment-dependent sensitivity, leading to regime-dependent discrepancies across bandwidth-delay product (BDP) conditions. Through targeted parameter exploration, we identify configurations that improve cross-platform alignment in low BDP regimes, while revealing structural differences that persist under higher load. This work provides both a practical tool for experimental L4S research and empirical insight into cross-platform behavioral differences, highlighting the importance of systematic characterization and environment-aware parameter selection in emulation-based AQM studies.

CCS CONCEPTS

• **Networks** → **Network experimentation; Network performance analysis; Network architectures; Transport protocols;**

KEYWORDS

L4S, DualPI2, low-latency, network emulation, active queue management, ECN, statistical analysis

1 INTRODUCTION

The Low Latency Low Loss Scalable Throughput (L4S) architecture is a modern approach to network-wide congestion control that aims to reduce queuing delays while maintaining high throughput. L4S challenges the long-held belief that these two objectives form a zero-sum tradeoff [10]. This architectural shift is increasingly gaining traction amid two seemingly contradictory trends in modern

networking: (i) the growing availability of high-bandwidth access technologies (e.g., fiber, 5G), which continually push the frontier of achievable throughput [22]; and (ii) the growing recognition among ISPs and network operators that increasing bandwidth alone does not solve the core problem of user-perceived performance, namely insufficient interactive responsiveness and high latency [56]. As a result, the focus of innovation has shifted from raw throughput to end-to-end latency and, by extension, to application-level quality of experience (QoE). This shift is especially important given the rise of latency-sensitive services like AR/VR, cloud gaming, and teleconferencing.

The persistent latency problem emerges as a direct consequence of bufferbloat, characterized by “overly large, unmanaged, and uncoordinated buffers” [25]. While early fixes focused on in-network active queue management (AQM) and shallow queue designs [49], L4S builds on these principles and introduces a new paradigm: latency control through sender-side congestion responsiveness enabled by explicit and fine-grained congestion signaling (ECN), and supported by specialized AQMs like DualPI2 [4] at bottleneck routers.

Ten years after its early conceptualization, L4S is now being actively standardized at the IETF, with working group discussions, multiple experimental RFCs [10, 74, 75], and early deployments from ISPs. While industry and standards activity have accelerated in recent years, academic exploration of L4S is only beginning to gain momentum.

A key enabler of broader academic engagement will be the availability of accessible experimentation tools. Most L4S studies rely on the Linux kernel implementation of DualPI2, which, while functional, requires elevated system privileges for empirical evaluation, and non-trivial kernel reconfiguration and recompilation for design exploration of alternative AQM components. This limits exploratory testing, slows iteration, and complicates reproducibility. With L4S still evolving, the broader networking community is well positioned not only to assess its performance, but also to contribute to its refinement. Supporting this collaborative development will require tooling that is modular and accessible to researchers and developers across platforms.

To address this need, we present a modular implementation of the DualPI2 AQM [4] integrated into Mahimahi [46, 47], a lightweight, user-space emulator already widely used for mobile and transport-level network evaluation. Unlike the monolithic nature of the kernel implementation—dictated by the internal structure and integration constraints of the Linux networking stack—our version

separates the components of DualPI2 in alignment with the RFC architecture. This modularization encourages design experimentation, such as swapping in alternative schedulers or tuning marking strategies. We show that DualPI2 exhibits environment-dependent parameter sensitivity, such that default configurations do not ensure cross-platform behavioral consistency. Through targeted tuning, we identify configurations that improve cross-platform alignment under specific BDP regimes.

Our key contributions are as follows:

- We present the first L4S-capable DualPI2 AQM implementation integrated into the Mahimahi emulator.
- We release the implementation of the DualPI2 Mahimahi module¹ and experimental artifacts² to enable repeatable L4S experimentation without kernel modification.
- We statistically quantify the achieved throughput, queueing dynamics, ECN marking rates, and drop behavior of our implementation across controlled scenarios, highlighting similarities and differences relative to the Linux DualPI2 implementation.
- We analyze the sensitivity of DualPI2 behavior to its configuration across execution environments, identifying key parameters that govern alignment between Mahimahi DualPI2 and the Linux implementation. We show that targeted tuning improves behavioral consistency under specific network conditions and provide practical parameter recommendations for reproducible experimentation.
- We release the statistical testing framework as an independent project to enable similar validation and characterization work using any metric.³

2 BACKGROUND AND MOTIVATION

In this section, we first outline the core design of L4S to ground our work in its technical context. We then examine recent industry deployment efforts, which signal growing momentum but also highlight the need for broader collaboration. Finally, we articulate the current limitations preventing open, reproducible experimentation with L4S, particularly within the academic community, and how our work seeks to address them.

2.1 Overview of the L4S Architecture

The L4S architecture, as standardized collectively across RFCs 9330, 9331 and 9332 [10, 74, 75], is comprised of two main components: a compliant congestion controller on the end hosts that responds to the explicit congestion signals, and in-network AQM support at the bottleneck for managing the coexistence of classic traffic and L4S traffic, in the form of a Dual Queue Coupled AQM.

The end-host congestion controller plays a central role in the L4S architecture, as it is responsible for translating explicit ECN-based congestion signals into rate adaptations that jointly achieve low latency, low loss, and scalable throughput. Unlike classic congestion control, which primarily relies on rare packet loss events and deep queues for stability, L4S operates with frequent, low-amplitude

ECN marking and deliberately shallow queues at the bottleneck. These design choices fundamentally change the nature of the congestion signal exposed to the sender, requiring continuous and fine-grained rate adaptation in the absence of buffering-induced smoothing. Traditionally, ECN signals in “classic ECN” were treated as equivalent to packet drops, triggering the same abrupt congestion responses—such as halving the congestion window in New Reno. In contrast, L4S redefines the semantics of the ECN Congestion Experienced (CE) mark to indicate early, low-level congestion, intended to elicit smooth and responsive rate control rather than drastic backoff. This shift enables the coexistence of low latency and high throughput, but also demands that both endpoints and AQMs interpret ECN signaling in fundamentally new ways.

While the concept of scalable congestion control predates L4S [38], the signaling and queueing model introduced by L4S makes scalable congestion control a necessity rather than an optimization: congestion controllers that depend on infrequent loss events or queue buildup for stability are ill-suited to achieve L4S’ goals. In addition, loss-based controllers typically incur long convergence times following a loss event, particularly at high sending rates, which is incompatible with the L4S objective of maintaining both low latency and high utilization.

A congestion controller is *scalable* if it fulfills the following conditions regardless of the scale of the network bandwidth [38, 73]:

- i) the time it takes to double the transmission rate must be independent of the current rate. This ensures the scalable throughput aspect of L4S, even in high-bandwidth settings.
- ii) the number of congestion signals per RTT remains constant, allowing the protocol to remain sensitive to these signals.

Together, these properties enable scalable congestion controllers to respond proportionally to frequent ECN marks while maintaining stable operation at high rates and low queueing delay. In contrast, classic congestion control algorithms such as Reno or Cubic are optimized for loss-driven operation and implicitly rely on queue buildup and infrequent congestion events, leading to slower convergence after loss and larger delay oscillations.

On the deployability front, for a scalable congestion control protocol to be usable with L4S over the public Internet, it must also comply with the Prague L4S Requirements [8, 74]. By satisfying these requirements, a transport can safely use L4S signals on the Internet; its congestion control can exploit frequent and fine-grained feedback provided through AccECN [7]; and it is able to coexist with classic flows fairly at shared bottlenecks. In addition, it can fall-back to classic (Reno-friendly) behavior on packet loss. Examples of L4S-compatible congestion controllers include TCP Prague [8], UDP Prague [40] and SReAM [35, 36].

The design of the Dual Queue Coupled (DQC) AQM, specified in RFC 9332 [75] and depicted in Fig. 1, isolates traffic by way of two queues: a classic (C) queue serving legacy congestion-controlled flows, and an L4S (L) queue for flows employing scalable congestion control. Each queue is managed by its own AQM, which independently regulates queue delay around a target operating point by computing a base control probability, namely P' for the C queue, and P'_L for the L queue. Then, these probabilities respectively determine packet drops (P_C) for classic traffic, and ECN marking (P_L) for L4S traffic.

¹<https://github.com/5G-VCA-CC/mahimahi/tree/dualpi2>

²<https://github.com/5G-VCA-CC/Experimentation-Data-Scripts/tree/ccr26-submission>

³https://github.com/5G-VCA-CC/mahimahi_validation.git




Figure 1: Dual Queue Coupled AQM [4, 72, 73]

The L4S queue is given latency priority to minimize queuing delay for scalable flows; however, throughput fairness between classic and L4S traffic is enforced through a coupling mechanism between the two AQMs; this is realized by setting the final L4S marking probability as the maximum of its native marking probability P'_L and a scaled version of the classic base drop probability $P_C = k * P'$.

The coupling ensures that L4S traffic remains responsive to congestion driven by classic flows, while still allowing the L4S AQM to independently maintain low queueing delay. Finally, a conditional priority scheduler arbitrates service between the C and L queues, providing latency priority to the L queue while using preset queue weights to control long-term throughput sharing with the C queue.

DualPI2 AQM [4, 75] is the de facto Linux kernel reference implementation of the DQC AQM. It is derived from the PI^2 AQM [18], which itself is an evolution of the PIE AQM [58]. The DualPI2 traffic classifier directs packets based on the value of the ECN field in the IP header as shown in Table 1. Packets marked with ECT(1) or CE are directed to the L queue, while packets marked Not-ECT or ECT(0) are sent to the C queue.

Table 1: ECN codepoints for L4S [74]

ECN Codepoint	Binary Value	L4S Meaning
Not-ECT	00	Not ECN-capable Transport
ECT(0)	10	Classic ECN-capable Transport
ECT(1)	01	L4S-capable Transport
CE	11	Congestion Experienced

The L queue in DualPI2 employs a low-latency AQM, referred to as the *Step AQM* in the Linux implementation, with a very shallow ECN marking target (default `step_thresh` = 1 ms). In contrast, the C queue is managed by the PI^2 AQM and operates at a significantly higher target queuing delay (default `pi2_target` = 15 ms), which reduces the drop probability and allows loss-based congestion controls to achieve high link utilization.

PI^2 [18] is a proportional–integral controller parameterized by an integral gain (α) and a proportional gain (β). The integral component gradually eliminates persistent steady-state deviations from the target delay, while the proportional component reacts more rapidly to transient changes in excess load. Together, the proportional and integral components provide both fast response and convergent long-term behavior.

2.2 Industry Deployment Efforts

Historically, industry considerations of network performance have predominantly focused on increasing network capacity, also known as bandwidth. L4S emerged from an alternative strategy for improving quality of service, concentrating instead on the reduction of network latency, particularly loaded latency [55]. Nokia, through its Bell Labs research arm, has played a leading role in L4S research and standardization, which was initiated in the mid-2010s [50], and first introduced to the IETF community in 2016 [9]. Research activity gradually expanded, notably with CableLabs integrating L4S mechanisms into their Data Over Cable Service Interface Specification (DOCSIS) [13]. This, in turn, led to the publication of the first Low Latency DOCSIS Internet Draft [86] in 2019. The following years saw the technology mature into practical deployment. Early 2025, Comcast became the first ISP to roll out low latency DOCSIS to its customers [16, 31], recently reaching the milestone of 10 million L4S-enabled homes, representing 360 million devices [42].

Beyond fixed access networks, mobile network operators have also embraced L4S as a pathway to consistent low-latency services. In particular, Vodafone, in collaboration with Nokia Bell Labs, conducted one of the first end-to-end trials of L4S over passive optical broadband [81], yielding near-zero queuing delays under load. Another example of such partnerships was between Ericsson, Qualcomm and SoftBank, which demonstrated L4S-enabled ultra-low latency on commercial 5G Standalone networks, and integrated L4S with 5G scheduling to secure consistent latency for mission-critical services [20]. In the US, T-Mobile has begun deploying L4S capabilities across their 5G networks [71]. T-Mobile frames L4S as a flexible mechanism for supporting application-aware latency behavior rather than a single, uniform service tier. By combining L4S with a 5G standalone core and network slicing, T-Mobile is deploying L4S across multiple slices, where each slice may be tuned differently depending on the latency tolerance and traffic characteristics of the applications it supports [71].

Application-layer adoption is also gaining momentum. Apple has added L4S support to FaceTime, showing improved video quality and reduced stalls under congestion [5]. With native integration in iOS 17 and macOS Sonoma, apps using QUIC, HTTP/3, or TCP can benefit with little or no modification. Developers also have the option to implement scalable congestion control and ECN validation for fine-grained optimization. Meanwhile, Netflix is developing a receiver-driven rate adaptation algorithm designed to leverage signals like those provided by L4S [3]. Though still experimental, these efforts reflect growing alignment between modern applications and emerging low-latency network capabilities.

2.3 Obstacles to Open and Rapid L4S Experimentation

This section describes the challenges associated with reproducible, design-oriented experimentation within the L4S ecosystem, and discusses how an emulation-based DualPI2 implementation can complement system-level deployments by enabling broader access to controlled experimentation. In particular, it highlights how such an approach opens new opportunities for both network- and application-level evaluation across diverse access environments.

System-Level L4S Implementations. Multiple implementations of L4S are now available in Linux [75], FreeBSD [70], 5G programmable RAN such as OpenAirInterface [33] and srsRAN [82]. Among these, the Linux kernel DualPI2 queue discipline (qdisc), referenced in RFC 9332 [75], remains the most widely used implementation in the research community. While tailored for production deployments, the kernel implementation is not optimized for rapid experimentation or prototyping. In fact, the monolithic nature of kernel qdisc modules, particularly DualPI2, complicates the isolation of specific components of the architecture for experimentation, such as the scheduler, since it is tightly embedded within the dequeue logic. A more modular design would facilitate mapping between the code and conceptual components defined in the RFC, streamlining targeted architectural exploration and comparative studies. As a result, relying solely on the kernel module can raise the barrier to entry for new researchers or practitioners seeking to test modifications or evaluate design alternatives.

In contrast, a user-space DualPI2 implementation would support configuration via command-line arguments, enabling clean state resets between runs and simplifying test automation. Parameters can be easily modified and swept over large ranges without kernel interaction, supporting systematic exploration of algorithmic behaviors across operating conditions. Instrumentation is straightforward, and logging can be implemented directly in user space without specialized kernel tracing tools. Moreover, an implementation on Mahimahi would offer a version-stable experimental baseline. The kernel DualPI2 codebase remains under active development, and changes in the upstream Linux tree may affect module behavior or interface assumptions over time. The Mahimahi module serves as a frozen snapshot of the design at specific points, enabling consistent replication of experiments over long timescales. Together, these properties make the Mahimahi DualPI2 module a valuable complement to the kernel implementation, supporting a faster, more flexible design exploration, reproducible experimentation, promoting broader engagement with L4S development.

Limited Testing Opportunities on Mobile Networks. Testing congestion control mechanisms over real mobile infrastructure presents numerous challenges. Cellular networks exhibit highly variable channel conditions due to factors such as user mobility, environmental interference, and scheduling dynamics at the radio access layer. These fluctuations introduce significant noise into experimental results, requiring repeated trials under diverse conditions to reach statistically meaningful conclusions. Moreover, access to mobile infrastructure with sufficient control and visibility into the network stack is often limited, and conducting tests on commercial networks can be both time-consuming and costly.

As a result, mobile networking research frequently turns to emulation. Tools like Mahimahi (§8.1) are widely adopted for this purpose due to their ability to replay recorded mobile traces in controlled, repeatable environments. Mahimahi supports uplink and downlink bandwidth traces collected from real cellular networks, enabling researchers to evaluate protocols and applications under representative mobile conditions without the unpredictability of live deployments.

However, despite its widespread use in mobile networking research, Mahimahi currently lacks support for L4S mechanisms. As

discussed in §2.2, L4S deployments in commercial mobile networks remain limited, making direct experimentation on physical infrastructure impractical. This absence of emulator support represents not only a missed opportunity, but a significant gap in the current L4S research landscape, particularly for mobile scenarios. Without accessible, trace-driven emulation capabilities, researchers in both academia and industry are constrained in their ability to rapidly evaluate, refine, and innovate on L4S in mobile contexts, ultimately slowing progress toward optimization and broader adoption.

Challenges for Application-Level L4S Research. Of the prior L4S application research noted in this paper (§8.2), the majority is conducted on testbeds [17, 27, 44, 57, 78, 79] or using simulators or emulators [77], some of which are proprietary [11, 12], which limits access and reproducibility.

In contrast, Mahimahi has been instrumental to numerous application QoE-focused research efforts, such as VR [84], web browsing [45], telemetry on cellular networks [87], mobile application QoE [60], video conferencing [2, 32, 64], and most notably, video streaming [15, 24, 43, 48, 61, 83, 88–90, 92]. However, Mahimahi currently lacks an L4S-capable AQM, preventing controlled evaluation of application-level L4S mechanisms within this widely adopted framework. If Mahimahi had DualPI2, it would enable rigorous, accessible, and mobile-aware L4S experimentation across these diverse application domains.

Contribution. This work extends Mahimahi with a DualPI2 implementation, allowing L4S evaluations to be conducted without requiring specialized kernel deployments. The DualPI2 AQM-in-emulator thus provides the missing link for application-level L4S research, leveraging Mahimahi’s widespread adoption to bridge the gap between transport-layer L4S development and application-layer evaluation.

Concretely, with a Mahimahi DualPI2 module, researchers can evaluate L4S-enabled applications against a DQC bottleneck without relying on dedicated testbeds, which may be constrained by availability, root privileges, and operational complexity. Further, our implementation enables evaluation without the need to modify or rebuild kernels, deploy custom operating systems, or configure ECN on physical network interfaces. Finally, the modular design of our DualPI2 implementation (§3) enables experimentation with structural variations (e.g., alternative schedulers or queue coupling logics), allowing researchers to systematically and directly study their impact on application QoE and to share reproducible configurations alongside their findings.

3 DESIGN AND IMPLEMENTATION

The DualPI2 AQM module in Mahimahi ⁴, depicted in Fig. 2, consists of six additional classes. The main DualPI2 logic is contained in the DualQCoupledAQM class, which inherits from AbstractPacketQueue. This class is used in the LinkShell module to instantiate and return the AQM choice specified by the user.

The DualQCoupledAQM keeps two queues, the C queue and the L queue, which are instances of the ClassicPacketQueue and L4SPacketQueue, respectively. Both classes inherit from AbstractDualPI2PacketQueue. Although the initial intuition was to derive

⁴<https://github.com/5G-VCA-CC/mahimahi/tree/dualpi2>




Figure 2: Simplified UML class diagram of the DualPI2 module in Mahimahi

these two classes from DroppingPacketQueue, we decided to create a new class deriving directly from AbstractPacketQueue, because DroppingPacketQueue requires setting the maximum queue length at the queue level. However, in the DualPI2 conceptual model, the maximum queue length is intended to be set and maintained at the AQM level rather than at the individual queue level, since the coupling mechanism of the AQM, in tandem with the scheduler, determines the effective share of bandwidth across both queues as if they formed a single pool [10]. This prompted the need for a new base class for the two queues, AbstractDualPI2PacketQueue, which leaves queue capacity management to the AQM class.

The general structure of the DualPI2 AQM was inspired by the pseudocode available in RFC 9332 [75]. Any logic not explicitly specified in the RFC was carefully ported from the Linux kernel implementation⁵, ensuring behavioral consistency while adapting it to Mahimahi’s architecture. The weighted round robin (WRR) scheduler is an example of such a module adapted from the kernel code. In our implementation, we first created the AbstractDualQScheduler class to provide an interface for individual scheduler classes, of which the WRRScheduler inherits.

Overall, we deliberately decompose the AQM into well-defined abstractions at both the scheduler and queue levels. Importantly, this modular design facilitates both the addition of new schedulers and the introduction of alternative logics for the C and L queues through their respective base classes, moving away from the monolithic design of the kernel code and explicitly promoting flexible experimentation with new algorithms.

On the visualization side, Mahimahi provides live graphs displaying available bandwidth, achieved throughput, and queue delay. To clearly distinguish between L4S and classic traffic, we extend these visualizations with additional lines that separately represent L4S traffic (in blue) and classic traffic (in orange). Fig. 3 shows live graphs of an example dual-flow with separate lines for the L4S and classic flows.

⁵https://github.com/L4STeam/linux/blob/testing/net/sched/sch_dualpi2.c




Figure 3: Distinguishing L4S traffic (blue) from classic traffic (orange) in Mahimahi live graphs; the dark blue line is the cumulative throughput.

4 BEHAVIORAL CHARACTERIZATION AND COMPARATIVE ANALYSIS

To assess the fidelity our Mahimahi DualPI2 module with respect to the Linux kernel reference implementation, we employ a non-parametric permutation test to quantify similarities and differences in terms of four metrics and across multiple traffic and network scenarios. The following sections describe the metrics of interest and the distance measure, the statistical validation framework, and finally the traffic scenarios and network conditions considered.

4.1 Metrics and Distance Measure

The metrics of interest span the throughput achieved in each traffic pattern; the internal state of the AQM, captured by the queue occupancy trace; and the AQM’s outward signaling behavior, represented by ECN marks and packet drops.

Throughput. We use the iperf3 utility [21] as a traffic generator in our experiments and collect the average throughput value for each run directly from the final iperf3 report. In dual-flow traffic scenarios, we use *cumulative throughput*, which is the sum of the throughput values achieved by the L4S flow and the classic flow. Since we extract a single average throughput value per run, the distance measure between two runs is simply the difference between their respective average throughput values.

State Similarity. We define *state similarity* as the degree to which queue occupancy traces are comparable between the two implementations at any given time. Inspired by the work of Ware et al. [85], which demonstrates that a queue occupancy time series can be used to classify congestion control protocols based on their buffer-filling behavior (under a fixed AQM), we adopt the converse hypothesis: given identical traffic, differences in queue occupancy time series can be attributed to differences in AQM behavior and/or system-specific deviations. Therefore, if we can reject the hypothesis that the distance between queue occupancy traces produced by the Mahimahi and Linux DualPI2 implementations is beyond a certain threshold, we can conclude that cross-system deviations can reasonably be attributed to internal stochastic noise rather than structural divergence of the AQM logic. The statistical test is described in more detail in §4.2.

Signal Similarity. Relying solely on state similarity for a comprehensive behavioral comparison is insufficient, as queue occupancy is an aggregate outcome influenced by traffic characteristics (e.g.

rate and burstiness) as well as internal AQM behaviors that may obscure finer differences between implementations. Although we control for traffic characteristics, ensuring that they are identical across both the kernel and Mahimahi experiments, any differences in AQM signaling behavior need to be further investigated. As an active queue management technique, DualPI2 influences congestion response by choosing whether to drop packets or mark them with ECN CE. These actions constitute the *signals* that the AQM sends through the network, influencing receiver feedback and, ultimately, sender adaptation. To evaluate whether the Mahimahi DualPI2 module produces equivalent signaling to the reference implementation, we collect and compare time series traces of ECN markings and packet drops, where each data point is the number of ECN marks or drops that occurred during the last time interval. Our goal is not to test for exact superposition or pointwise equality of traces from Mahimahi and the kernel reference, which is unrealistic in systems subject to unavoidable timing noise. Rather, we seek to assess practical differences between the two implementations.

Dynamic Time Warping as the Distance Measure for Time Series Data. The most fundamental building block of our evaluation methodology is the distance measure that compares two time series traces. Similar to [85], our approach for computing individual pairwise trace differences is based on Dynamic Time Warping (DTW) [66], which takes two time series as input and returns a scalar distance that quantifies how dissimilar the two traces are. DTW is used to account for temporal dilation and contraction between traces and has been traditionally used in pattern-matching domains such as automatic speech recognition [34], data mining [39, 59], financial time series analysis [65], gesture recognition [30], and bioinformatics [76].

To build intuition for DTW, we first consider a baseline approach using Euclidean Distance (ED). Let $X = (x_1, \dots, x_n)$ and $Y = (y_1, \dots, y_m)$ be two time series of length n and m , where x_i and y_i represent the value of the considered metric at time step i in traces X and Y , respectively. ED is computed as the sum of the squared differences between corresponding elements in the two traces. However, this one-to-one comparison often fails for network traces. Fig. 4 illustrates such a case: two queue occupancy traces exhibit similar patterns but with temporal misalignments.




Figure 4: Comparing two time series using the Euclidean distance vs. DTW distance. Adapted from [85].

Fig. 4a shows the inflated ED due to these misalignments, while Fig. 4b shows the DTW distance, which captures structural similarity even when traces are temporally misaligned, yielding a smaller and more meaningful distance than the ED. Network traces may experience temporal dilation or contraction due to dynamic conditions like congestion, rate adaptation and system limitations. In our case, distortions may be caused by differences in scheduling behavior and timing granularity between Mahimahi and the Linux kernel. These effects cause otherwise similar traces to appear stretched or compressed in time.

Unlike ED, which requires a strict point-to-point comparison, DTW allows one-to-many alignments between trace points. At a high level, DTW computes an optimal alignment (or *warping path*) between the two sequences that minimizes a total cost. This cost is defined as the sum of absolute differences between aligned sample values along the computed alignment path.

Furthermore, trace collection intervals are not always exactly aligned to millisecond boundaries, which means that different flow traces may contain different numbers of sampled points. Without correction, DTW distances would be artificially inflated for longer traces simply because more points must be matched, and therefore more pointwise errors are accumulated along the warping path.

To account for this effect, we normalize the raw DTW distance by the length of the optimal warping path π (with $|\pi|$ denoting the number of alignment steps). The normalized DTW distance, defined in (1), yields a length-invariant measure of per-alignment-point deviation, ensuring that DTW distances are comparable across traces of differing durations and sampling densities.

$$\text{DTW}_{\text{norm}}(X, Y) = \frac{1}{|\pi|} \sum_{(i,j) \in \pi} |x_i - y_j| \quad (1)$$

4.2 Statistical Equivalence Testing

To evaluate whether Mahimahi behaves similarly to the Linux kernel in terms of our four metrics (throughput, queue occupancy, ECN marks and packet drops), we employ a *nonparametric exceedance-based validation test*. This test assesses whether cross-system differences are extreme relative to natural within-system variability. Specifically, we evaluate how frequently kernel–Mahimahi differences exceed a predefined tolerance threshold. The approach makes no assumptions regarding equality of variances or the underlying distributional shape of the populations. The only assumption is independence between experimental runs. We enforce this by fully resetting the experimental environment between runs, including termination of *iperf* processes, teardown of each environment, and a fixed quiescence period prior to the next trial to prevent residual state carryover. The sections below outline the key components of our testing methodology, which is publicly available.⁶

Defining the Distance Measure. Let $\mathcal{M} = \{M_1, \dots, M_{100}\}$ and $\mathcal{K} = \{K_1, \dots, K_{100}\}$ denote the sets of traces collected from Mahimahi and the Linux kernel, respectively.

⁶https://github.com/5G-VCA-CC/mahimahi_validation.git

• **Time-series metrics (Queue occupancy, ECN marks, packet drops).** For temporal observations, we compute the normalized DTW distance between two traces X and Y :

$$d(X, Y) = \text{DTW}_{\text{norm}}(X, Y) \quad (2)$$

Equation (1) defines the normalized DTW distance between the two traces X and Y . We compute all inter-system pairwise distances for each group, as well as all cross-system distances between traces in $\mathcal{M}$ and $\mathcal{K}$. These collections form the empirical distance distributions used in our validation procedure.

• **Scalar metrics (average throughput).** For non-time-series metrics, the distance measure we use is the absolute difference between two scalar measurements X and Y :

$$d(X, Y) = |X - Y| \quad (3)$$

Defining the Test Statistic and Exceedance-Based Validation.

Having constructed the inter-system and cross-system distance distributions, we assess whether Mahimahi behaves similarly to the Linux kernel by comparing cross-system discrepancies against natural inter-system variability. Let $\epsilon_{\max}$ denote a tolerance threshold defined as the maximum of the 95th percentiles of the two inter-system distance distributions.

$$\epsilon_{\max} = \max\left(\text{Quantile}_{0.95}(D_{\text{within}}^{\mathcal{M}}), \text{Quantile}_{0.95}(D_{\text{within}}^{\mathcal{K}})\right) \quad (4)$$

Where $D_{\text{within}}^{\mathcal{M}} = \{d(X_i^{\mathcal{M}}, X_j^{\mathcal{M}}) : i < j\}$ denotes the set of all pairwise distances within the Mahimahi group, and $D_{\text{within}}^{\mathcal{K}} = \{d(Y_i^{\mathcal{K}}, Y_j^{\mathcal{K}}) : i < j\}$ denotes the set of all pairwise distances within the kernel group. $d(X, Y)$ is the distance measure defined in Equation (2) for time series inputs, and in Equation (3) for scalar inputs.

Consequently, a cross-system difference is considered excessive only if it exceeds the greatest natural fluctuation exhibited by either system. This avoids unfairly evaluating cross-system discrepancies against the tighter variability of the less noisy implementation and ensures that only genuinely extreme deviations are flagged.

We then estimate our test statistic, which is the true cross-system exceedance probability, denoted as $\hat{p}_{\max}$, and defined as the probability that a randomly selected cross-system distance exceeds $\epsilon_{\max}$.

$$\hat{p}_{\max} = \frac{1}{|D_{\text{cross}}|} \sum_{d \in D_{\text{cross}}} \mathbf{1}\{d > \epsilon_{\max}\} \quad (5)$$

D_{cross} denotes the set of all cross-system distances and is defined as $D_{\text{cross}} = \{d(X, Y)\}$, where $d(X, Y)$ is the distance measure defined in Equation (2) for time series inputs, and in Equation (3) for scalar inputs. The $\hat{p}_{\max}$ quantity is estimated empirically as the proportion of observed cross-system distances that satisfy $d > \epsilon_{\max}$. The resulting estimator quantifies how often Mahimahi–kernel discrepancies are more extreme than variability observed internally within either system.

Assessing statistical equivalence. We interpret $\hat{p}_{\max}$ under a practical behavioral similarity framework. Specifically, we define

the null hypothesis H_0 as follows:

H_0 : Mahimahi does not behave similarly to the Linux kernel under a 5% exceedance criterion

H_1 : Mahimahi behaves similarly to the Linux kernel under a 5% exceedance criterion

Operationally, this corresponds to:

$$H_0 : p_{\max} \geq 0.05 \quad \text{and} \quad H_1 : p_{\max} < 0.05$$

Here, $p_{\max}$ denotes the true (population-level) cross-system exceedance probability, whereas $\hat{p}_{\max}$ is its empirical estimate computed from the observed data.

If $\hat{p}_{\max} \geq 0.05$, cross-system differences exceed natural inter-system variability beyond the acceptable threshold of 5%, the cross-system distribution is significantly misaligned with the inter-system distribution, and we fail to reject H_0 . In this case, the two implementations cannot be deemed behaviorally similar under this criterion. Conversely, if cross-system differences are within the acceptable threshold of variability of 5%, the kernel DualPI2 AQM dynamics are preserved in the Mahimahi implementation, which rejects H_0 and concludes practical behavioral similarity when $\hat{p}_{\max} < 0.05$.

The statistical significance threshold is somewhat arbitrary and depends largely on the field of application. The value of 5% was originally proposed by Fisher in his textbook “Statistical Methods for Research Workers” [23]. Although some fields such as genomics and high-energy physics use far stricter thresholds, a value of 5% is still considered a conventional significance threshold and is used in many experimental disciplines [6].

4.3 Traffic Scenarios and Network Conditions

We collected data from identical experiments on both Mahimahi and Linux DualPI2 using the `iperf3` [21] bandwidth probing tool, following three different traffic patterns:

- **Single L4S flow**, using TCP Prague congestion controller to isolate the behavior of the L queue.
- **Single classic flow**, using a TCP Cubic flow to isolate the behavior of the C queue.
- **Dual flow**, one TCP Prague and one TCP Cubic, to validate the coupling mechanism between the L queue and the C queue.

We explore all combinations of the above three flow types with three representative Bandwidth–Delay Product (BDP) classes summarized in Table 2, resulting in nine distinct experimental scenarios.

Table 2: Representative BDP classes used in our evaluation.

Regime	RTT	Bandwidth	Example Networks
Low-BDP	20 ms	12 Mbps	Legacy DSL
Medium-BDP	40 ms	50 Mbps	4G LTE/5G; Starlink
High-BDP	100 ms	200 Mbps	Inter-continental fiber

We collect 100 `iperf3` flow traces from Mahimahi DualPI2 and 100 from the Linux kernel DualPI2 for each of the BDP regimes, for each traffic pattern. We provide a justification for the sample size in Appendix B.5. Each `iperf3` run is 30 seconds long. For every

test scenario, we collect average per-run throughput and extract time series samples of queue occupancy, ECN marks, and packet drops every tupdate (default 16 ms [75]) for both implementations. We chose tupdate as our sampling rate because it is the frequency at which the key probabilities of the DualPI2 AQM are updated, effectively representing the “internal clock” of DualPI2. Our data collection scripts are publicly available.⁷

5 VALIDATION SETUP

This section outlines the measures taken to ensure that the validation of the Mahimahi AQM against its kernel counterpart is conducted fairly, and describes the experimental settings used to enable controlled comparisons of the AQM dynamics.

We installed a precompiled kernel⁸ that contains the DualPI2 AQM and TCP Prague modules. TCP Prague, with AccECN [7] enabled, is used for both the Mahimahi and the kernel experiments as the congestion control protocol of the L4S-enabled iperf3 flow. The DualPI2 kernel module is the reference against which we validate our Mahimahi DualPI2 AQM. We set the same, default values for the DualPI2 parameters [75] for both implementations, namely $target = 15$ ms, $step_thresh = 1$ ms, $limit = max_link_rate * 250$ ms, $tupdate = 16$ ms, $alpha = 0.16$, and $beta = 3.20$.

Mahimahi–kernel Matching Setup. To realize the various configurations discussed in §4.3 on Mahimahi, we use two of Mahimahi’s shells [47]: LinkShell and DelayShell. LinkShell emulates link dynamics using constant-rate packet delivery traces, while DelayShell enables the introduction of a fixed per-packet delay. For the Linux kernel runs, we use two network namespaces, sender and receiver, connected by a virtual link consisting of a veth pair (one veth interface in each namespace). At the sender egress we layer a Hierarchical Token Bucket (HTB) over the DualPI2 AQM to enforce the rate limit. To control end-to-end RTT in the kernel experiments, we emulate propagation delay using the Linux traffic control (tc) NetEm [28] queuing discipline by applying a one-way delay on each of the sender and receiver namespaces.

For rigorous validation, we augment the kernel DualPI2 implementation with additional counters exposed via the struct named `tc_dualpi2_xstats`, enabling direct observation of marking and drop behavior using tc tooling, which was also modified for this purpose. These changes reside in our Linux fork⁹ (which is derived from the DualPI2 testing branch). The parsing changes to tc are in our iproute2 fork.¹⁰ This setup provides a reproducible, Mahimahi-aligned bottleneck with a strict rate cap enforced by HTB, queuing/markings governed by DualPI2, and a configurable RTT introduced by NetEm. We also provide scripts to automate kernel module recompilation and loading, as well as namespace reinstallation.¹¹

⁷<https://github.com/5G-VCA-CC/Experimentation-Data-Scripts/tree/ccr26-submission>

⁸kernel version 5.15.72-48b3db6b4-prague-111 available at <https://github.com/L4STeam/linux?tab=readme-ov-file>. More recent kernel versions may have been released since the completion of this work.

⁹<https://github.com/gargAneesh/linux/tree/add-counters>

¹⁰https://github.com/gargAneesh/iproute2/tree/add_counters

¹¹https://github.com/gargAneesh/l4s_scripts

6 EVALUATION RESULTS & DISCUSSION

This section presents a comparative behavioral evaluation of Mahimahi DualPI2 and the Linux kernel implementation using our statistical testing framework. We begin with dual-flow experiments under the default DualPI2 configuration to establish a baseline comparison. Guided by the observed discrepancies, we then investigate the sensitivity of key parameters in single-flow experiments within Mahimahi. Finally, we evaluate dual-flow performance under refined parameter settings informed by this analysis.

Dual-flow experiments with default DualPI2 parameters. In the first part of our evaluation, we collect traces from 100 runs for the dual flows traffic pattern (TCP Prague $\times$ TCP Cubic) in all three BDP regimes on each of the Mahimahi module and the Linux kernel qdisc. In this set of experiments, we use the default DualPI2 parameters described in §5. The histograms in Fig. 5 show the empirical distribution of each of our test statistics in all three BDP regimes. Each distribution is accompanied by the corresponding ε_{max} and $\hat{p}_{max}$. The figure shows that the queue occupancy test statistic fails to reject the null hypothesis H_0 across all BDPs, indicating a considerable mismatch in the instantaneous state of DualPI2 between the Mahimahi and the kernel setting. This result is due to system differences between the kernel environment and the emulated Mahimahi environment. In particular, Mahimahi’s network traces encode the bandwidth as a succession of delivery opportunities at every millisecond, effectively implementing a token bucket that refills at 1 ms intervals, creating small traffic bursts, while packet delivery in the kernel execution is more smooth and evenly-spaced. We explore this observation in more detail in Appendix A. Meanwhile, two test statistics reject the null hypothesis H_0 in all BDPs within the acceptable threshold of 5%: packet drops and ECN marks. This indicates practical behavioral similarity with regard to these signal metrics. Looking at the traces more closely, we can see that the packet drops are close to zero across all experiments. This is expected even for classic traffic, since TCP Cubic supports classic ECN, which means the congestion signal is reinterpreted from packet drops to ECN marks. Considering the throughput test statistic, the test only rejects the null hypothesis for the low BDP regime (12 Mbps, RTT = 20 ms), indicating the achieved throughput in Mahimahi is very different from that of the kernel setting in both the medium and high BDP regimes.

To more deeply analyze the throughput discrepancy, we plot the CDF of achieved throughput in both systems across all BDPs, shown in Fig. 6. Mahimahi clearly fails to saturate the bottleneck capacity in the medium and high BDP scenarios, under-utilizing the link compared to the kernel qdisc.

The achieved throughput through DualPI2 is the result of queue dynamics, particularly the ECN marking behavior. Although the result of the hypothesis testing with ECN marks does not a priori indicate a mismatch in ECN marking behavior, we know that excessive ECN signaling causes the sender to throttle transmission, leading to reduced throughput. This is a direct consequence of the queue delay (sojourn time) exceeding $step_thresh$ for L4S traffic, and $target$ for classic traffic.




Figure 5: Validation of the TCP Prague $\times$ TCP Cubic runs in High, Medium, and Low BDP scenarios - Default DualPI2 parameters.

These observations reveal that DualPI2’s default parameterization, while appropriate for kernel deployment, does not transfer directly to an emulated execution environment. The interaction between Mahimahi’s timing and queueing mechanisms and DualPI2’s control thresholds alters AQM dynamics, thereby influencing throughput. This highlights the environment-dependent sensitivity of DualPI2 parameters and demonstrates that careful parameter selection is essential for achieving comparable behavior across platforms.

Single-flow experiments with varying DualPI2 parameters.

To understand the impact of these parameters, we conduct single-flow experiments in Mahimahi DualPI2, isolating L4S and classic traffic to assess how `step_thresh` affects L4S throughput and

target affects classic throughput. For a fair comparison, we also run a series of single flow experiments on the kernel, using the default parameters of DualPI2. The objective is to determine whether adjusting these parameters enhances behavioral alignment with the Linux kernel implementation.

In the L4S-only experiment, we sequentially run 100 `iperf3` flows using TCP Prague for two additional values of `step_thresh`: 5 and 10 ms. Similarly, we run 100 classic-only `iperf3` flows sequentially using TCP Cubic for two additional values of `target`: 30 and 45 ms. We plot the variation of the achieved throughput in each one of these scenarios as shown in Fig. 7 for L4S traffic and Fig. 8 for classic traffic. In each scenario, we also plot the corresponding kernel experiment results that use the default parameters for comparison.

We run the statistical test comparing each modified DualPI2 run on Mahimahi with its kernel counterpart that uses the default parameters. We compile the results of the L4S-only and classic-only experiments in Tables 3 and 4, respectively, where we show for each BDP regime and value of the considered DualPI2 parameter, the statistical test results for each metric in terms of $\hat{p}_{\max}$, denoted as p , and $\epsilon_{\max}$, denoted as ϵ_{ps} for brevity. The “OK” column refers to the result of the statistical test, which either fails to reject H_0 (x), implying a significant behavioral misalignment in terms of the corresponding metric, or rejects H_0 (✓), implying similarity under the 5% exceedance criterion.




Figure 6: Dual-flow cumulative throughput with default DualPI2 parameters in all BDP regimes




Figure 7: L4S-only throughput varying step_thresh in all BDP regimes




Figure 8: Classic-only throughput varying target in all BDP regimes

The results in Table 3 show that, in the L4S-only case, setting step_thresh to 5 ms improves the statistical test outcome of the throughput for both the low and medium BDP scenarios compared to the default value of 1 ms. As for the high BDP scenario of 200 Mbps, increasing step_thresh to 5 ms does not yield any improvement. At 10 ms, while the outcome for the ECN mark metric improves, there is no change with regard to achieved throughput. Concerning classic-only traffic scenarios, Table 4 shows that setting target to 30 ms improves the statistical test outcome for the medium BDP scenario, while it conserves the same outcome for the low BDP scenario. In the high BDP case, a target value of 45 ms improves the outcome for all considered metrics with the exception of queue occupancy.

Dual-flow experiments with refined DualPI2 parameters. Bringing together the results of these experiments, we rerun the dual-flow scenarios, combining the values of step_thresh and

target that optimize the statistical results for each BDP regime prioritizing throughput, namely:

Low BDP: step_thresh = 5 ms, target = 30 ms;

Medium BDP: step_thresh = 5 ms, target = 30 ms; and

High BDP: step_thresh = 10 ms, target = 45 ms.

For the medium BDP scenario, although the value of step_thresh that achieves the greatest metric alignment in the L4S-only experiments is 10 ms, we empirically find that the 5 ms value yields the best results in terms of throughput for the dual-flow traffic pattern. The empirical distribution of each of our test statistics using the above optimal parameters for each BDP regime is summarized in Fig. 9. We observe that, except for queue occupancy, the tuned parameters improve statistical alignment across most metrics in the medium and high BDP regimes. However, while throughput alignment improves at 50 and 200 Mbps, the resulting $\hat{p}_{\max}$ remains above the rejection threshold, indicating persistent behavioral differences. Furthermore, the parameter adjustments introduce a tradeoff in the high BDP scenario: ECN marking behavior, previously aligned with the kernel, exhibits increased divergence under the tuned configuration. These results motivate further investigation, not only into the interaction with emulation timing granularity, but also in its coupling mechanism to better understand the sources of persistent divergence under high load.

To begin to understand the systemic differences between the Mahimahi and the kernel namespace-based experiments, we investigate the packet dequeue cadence in both systems and report results in Appendix A. Briefly, the discrepancies we observed are partly due to the relatively coarse network channel emulation in Mahimahi compared to the kernel.

In Appendix B, we introduce a bootstrap-based confidence interval framework to more precisely characterize this divergence and to formally assess, with statistical significance, whether one configuration demonstrates improvement over another. Our results show that, at moderate and high bandwidths, the refined configuration

Table 3: Validation results of single L4S flows varying step_thresh across BDP scenarios

Step_thresh	Low BDP (12 Mbps, 20 ms)									Medium BDP (50 Mbps, 40 ms)									High BDP (200 Mbps, 100 ms)								
	1 ms			5 ms			10 ms			1 ms			5 ms			10 ms			1 ms			5 ms			10 ms		
Metric	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK
Throughput	0.10	1.00	×	0	0.010	✓	0	0.980	×	0.20	1.00	×	1.90	0.0001	✓	1.70	0.0001	✓	55.80	1.00	×	55.80	0.280	×	55.80	0.280	×
Queue occup.	0.232	1.00	×	0.368	1	×	0.363	1	×	0.218	0.886	×	0.617	1	×	0.710	1	×	0.144	0.949	×	0.925	1	×	0.865	1	×
ECN Marks	3.40	0.040	✓	1.088	1	×	2.822	0.006	✓	0.004	0.998	×	0.957	0.932	×	1.583	0.040	✓	0.163	1.00	×	0.198	0.984	×	1.154	0.003	✓
Drops	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓	0	0.00005	✓

Table 4: Validation results of single classic flows varying target across BDP scenarios

Target	Low BDP (12 Mbps, 20 ms)									Medium BDP (50 Mbps, 40 ms)									High BDP (200 Mbps, 100 ms)								
	15 ms			30 ms			45 ms			15 ms			30 ms			45 ms			15 ms			30 ms			45 ms		
Metric	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK	eps	p	OK
Throughput	0.30	0.013	✓	0.20	0.0101	✓	0.10	0.128	×	13.90	0.670	×	9.00	0.030	✓	9.50	0.140	×	46.00	0.049	✓	12.80	0.130	×	9.00	0.030	✓
Queue occup.	7.583	1	×	7.522	1	×	9.096	1	×	22.73	1	×	23.79	1	×	75.52	0.786	×	67.93	0.993	×	55.7	0.895	×	23.73	1	×
ECN Marks	0.0084	0.999	×	0.007	0.999	×	0.007	0.984	×	0.004	0.998	×	0.0005	0.018	✓	0.003	0.997	×	0.001	0.617	×	0.004	0.976	×	0.0005	0.002	✓
Drops	0.009	0.066	✓	0.009	0.062	✓	0.013	0.092	×	0.051	0.024	✓	0	0.00005	✓	0.042	0.040	✓	0	0.00005	✓	0.034	0.020	✓	0	0.00005	✓




Figure 9: Validation of the TCP Prague $\times$ TCP Cubic runs across BDP scenarios - refined Mahimahi DualPI2 parameters.

delivers clear, statistically significant improvements in throughput consistency over the default parameter setting. Nonetheless, the persistence of discrepancies at high BDP suggests that structural differences between Mahimahi’s emulation model and the kernel’s packet scheduling and timing mechanisms become more pronounced under higher traffic intensity.

Overall, our findings demonstrate that while cross-platform alignment can be improved through targeted parameter tuning, structural differences between Mahimahi and kernel execution environments introduce regime-dependent effects that must be carefully considered in reproducible L4S experimentation.

7 LIMITATIONS

In this section, we discuss the known limitations of our behavioral comparison of the Mahimahi DualPI2 module and the Linux implementation.

In terms of the traffic types considered, our experiments only test a single congestion control protocol per traffic type, namely TCP Cubic for classic flows and TCP Prague for L4S flows. This selection may not be representative of the congestion control distribution on the Internet, and does not provide insight into the behavioral equivalence of the two systems under other congestion controllers. In particular, for L4S traffic, we do not test Google’s BBRv2 or BBRv3, which offer ECN/L4S support and are growing in prevalence [14, 91]. Furthermore, prior work has described fairness

issues using the DualPI2 reference implementation under certain conditions, especially when a BBR flow is involved [67, 68, 91]. In our experiments, the Cubic and Prague dual-flow scenarios exhibit relatively fair link sharing, although we do not formally quantify fairness or statistically validate it against the Linux implementation.

Regarding the network configurations tested, we only evaluate static scenarios across the different BDPs, with constant throughput and RTT values, because these settings are easier to configure and reproduce in both Mahimahi and the Linux namespace-based experiments than dynamic scenarios with time-varying throughput or RTT. Moreover, the Mahimahi emulator only allows for a constant RTT value per execution, and does not define dynamic RTT trace files that can be replayed in the same way as throughput traces.

8 RELATED WORK

8.1 Network Emulation & Simulation Tools

Our work builds on Mahimahi, a widely used user-space network emulator that supports the record-and-replay of packet-level network traces, originally introduced by Netravali et al. [46, 47]. While it was initially developed to emulate cellular conditions using real-world mobile bandwidth traces, Mahimahi is not limited to mobile scenarios. Users can define custom network traces, including constant-rate links or synthetic profiles, enabling its use across a broad range of network conditions. Its extensibility, support for time-varying network behavior and ease of use have made it the

go-to emulator for reproducible protocol experimentation, and has enabled a plethora of important work in networked systems and protocol design and evaluation [1, 15, 24, 26, 37, 43, 45, 48, 60, 61, 83, 84, 87–90, 92]. Other emulation tools include NetEm [28] and dummynet [63], which allow bandwidth and delay control but rely on static configurations that do not capture the time-varying nature of mobile networks. Mininet [41] emulates complete network topologies by creating virtual hosts and links using Linux namespaces. These systems are valuable for repeatable testing, but none currently support L4S-compatible AQMs.

Aside from emulators, it is worth mentioning that the ns-3 simulator [62] has started to integrate L4S-related support [51]. The DualPI2 AQM, although available as an external module [29], has not yet been included in the mainline release. Recent work has implemented the L4S end-host stack in ns-3 [80] including TCP Prague, and validated it against a Linux L4S testbed. However, testing real-world applications through ns-3 poses additional challenges not found in Mahimahi-based experimentation. First, although Direct Code Execution (DCE) enables existing userspace applications to run within ns-3, applications must be rebuilt in a DCE-compatible form [53], and compatibility depends on whether the required POSIX/library/system calls are supported by DCE. In particular, missing function symbols or system calls may require adding the corresponding support manually [52]. This adds uncertainty when integrating complex applications such as real multimedia systems, which are especially relevant for evaluating L4S. Moreover, DCE only supports C/C++ applications [53], limiting the range of applications that can be integrated into the ns-3 ecosystem. By contrast, Mahimahi enables applications to run without modifying or rebuilding their binaries, and our integrated DualPI2 module allows low-friction end-to-end testing of complex multimedia applications under L4S. Second, trace-driven experiments are simpler in Mahimahi: each trace line represents a packet-delivery opportunity, making cellular traces straightforward to replay [46]. In ns-3, variable-link behavior can be emulated by dynamically changing link attributes such as the point-to-point device DataRate, or by using more detailed wireless/channel models, adding to the complexity of the experiment preparation task [54].

Lastly, proprietary tools have also been used for L4S experimentation. For instance, Ericsson employed an internal simulator in early evaluations of SCReAM over L4S networks [11, 12]. While such setups enable advanced testing, they are not publicly available, limiting their usefulness for broader validation or follow-up work.

To our knowledge, our work is the first to introduce modular L4S capabilities into an open-source, record-and-replay emulator like Mahimahi, allowing the community to both evaluate existing protocols and experiment with new algorithmic designs under controlled, trace-driven network conditions.

8.2 Prior Work on L4S

One of the earlier research in the L4S applications sphere is work by Brunello et al [11, 12] on L4S in 5G networks. In particular, this work explores the adoption of L4S for latency-critical Augmented Reality (AR) video gaming traffic using the Self-Clocked Rate Adaptation for Multimedia (SCReAM) congestion controller [35, 36]. The study shows that enabling L4S significantly reduces end-to-end delay

compared to non-L4S operation, while maintaining application-layer throughput above the minimum requirements of high-rate, latency-sensitive applications, even under high system load. Although a fundamental throughput–latency trade-off remains, L4S consistently achieves lower packet loss. This study also investigates practical deployment challenges in 5G networks, such as encryption at the Radio Link Control (RLC) layer, which prevents packet marking within certain network nodes. To address this issue, the authors implement the L4S marking strategy at the Packet Data Convergence Protocol (PDCP) layer within the gNB, treating it as the effective network marking point. Still in the 5G context, the L4Span Radio Access Network (RAN) architecture [82] was proposed to enable end-to-end low-latency signaling on 5G’s last-mile link. It exposes RAN queue state to predict millisecond-scale queue occupancy and applies ECN marking for both low-latency and classic flows. The design aims to minimize required RAN modifications while remaining 3GPP and O-RAN compliant. A prototype is implemented in C++ on the open-source srsRAN platform. In the authors’ reported evaluation, their design reduces one-way delay of both low-latency and classic flows by up to 98% while maintaining near line-rate throughput. A number of other work have shown similar performance gains for various applications: Extended Reality (XR) [78, 79], cloud gaming [27], volumetric video conferencing [17], real-time streaming [44, 57, 77].

Some other research explored the effect of a partial deployment of the L4S ecosystem, pointing to L4S benefits not being guaranteed under partial deployment. Prior work [68] looks into the deployment of L4S-compatible congestion controllers amid partial L4S deployment on bottleneck nodes. The paper concludes that an L4S sender (using TCP Prague or BBRv2-style scalable congestion controller) cannot be assured it will not harm or be harmed by other flows when the sender does not know the bottleneck type and other cross-traffic conditions. In a preceding study [69], findings highlight cases where TCP Prague’s throughput or fairness can degrade in single-queue or incompatible coexistence settings (e.g., with non-ECN-responsive cross traffic).

Besides, there have been several attempts at devising novel congestion control protocols that leverage L4S’ signaling. Among these works, we can cite aL4S-CC [77]. The evaluation against GCC and SCReAMv2 reveals an improved average link utilization by 4% compared to GCC and 17.9% compared to SCReAMv2. Another such work by Pan et al. proposes L4S-GCC [57] for RTC/WebRTC: it combines ECN-derived queue information with delay-gradient logic to improve responsiveness and recovery; the authors report a reduction in stalling rate by 1.51–2.80% vs. GCC and an improvement in bandwidth utilization by 11.4–31.4% under delay jitter.

While these efforts have shown promising results, they remain hard to reproduce and extend. Our goal is to lower the barrier to entry for L4S research by embedding DualPI2 into a flexible user-space emulator.

9 CONCLUSION

Although a Linux implementation of DualPI2 is available, controlled and reproducible experimentation on L4S mechanisms benefits from a modular user-space implementation. By adding the DualPI2 AQM to the Mahimahi emulator, we provide an extensible module that

enables rapid iteration on control parameters, facilitates component-level experimentation, and supports application-focused studies without kernel modification and across diverse network conditions.

Through statistical behavioral characterization and cross platform comparison, we show that alignment between emulated and kernel executions is neither automatic nor invariant across BDP regimes. Instead, DualPI2's effective dynamics depend on the interaction between control parameters and the underlying execution environment, leading to regime-dependent behavior. Our analysis demonstrates that targeted parameter tuning can improve alignment in specific BDP conditions, while also revealing structural differences that persist under higher load. Table 5 summarizes the DualPI2 parameters that showed, with a 5% statistical significance threshold, the best alignment with the Linux kernel reference implementation for each BDP regime.

Table 5: Mahimahi DualPI2 parameters for each BDP regime

Regime	RTT	Bandwidth	step_thresh	target
Low-BDP	20 ms	12 Mbps	5 ms	30 ms
Medium-BDP	40 ms	50 Mbps	5 ms	30 ms
High-BDP	100 ms	200 Mbps	10 ms	45 ms

Collectively, this work lowers the barrier to experimental L4S research while providing an empirical understanding of cross-platform behavioral discrepancies. We hope it serves both as a practical tool for the community and as a foundation for deeper investigation into the parameter sensitivity of DualPI2 and the interaction between AQM control logic and emulation environments.

ACKNOWLEDGMENTS

We are grateful to Dave Täht for insightful discussions in the early stages of this work. We especially thank Dr. Sadjad Fouladi for his technical guidance and encouragement as we navigated the Mahimahi emulator project. We also acknowledge Dr. Prateesh Goyal for his advice on AQM development in Mahimahi. Generative AI tools were used to support scripting for experiments and statistical validation, plot generation, and language polishing. This work was partly funded by the Comcast Innovation Fund, grant 2024-1030.

REFERENCES

- [1] Soheil Abbasloo, Chen-Yu Yen, and H. Jonathan Chao. 2021. Wanna Make Your TCP Scheme Great for Cellular Networks? Let Machines Do It for You! *IEEE Journal on Selected Areas in Communications* 39, 1 (2021), 265–279.
- [2] Neil Agarwal, Rui Pan, Francis Y. Yan, and Ravi Netravali. 2025. Mowgli: Passively Learned Rate Control for Real-Time Video. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*. USENIX Association, Philadelphia, PA, 579–594. <https://www.usenix.org/conference/nsdi25/presentation/agarwal>
- [3] Paul-Louis Agneau, Grenville Armitage, and Scott Danahy. 2025. *Network Delivery Time Control*. Internet-Draft draft-agneau-ccwg-ndtc-00. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-agneau-ccwg-ndtc-00/> Work in Progress.
- [4] Olga Albisser, Koen De Schepper, Bob Briscoe, Olivier Tilmans, and Henrik Steen. 2019. DUALPI2: Low Latency, Low Loss and Scalable (L4S) AQM. In *Proceedings of the Netdev 0x13 Conference*.
- [5] Apple, Inc. 2023. Reduce network delays with L4S. (2023). <https://developer.apple.com/videos/play/wwdc2023/10004/> Apple WWDC 2023 video; accessed 2026-02-10.
- [6] Daniel J. Benjamin, James O. Berger, Magnus Johannesson, et al. 2018. Redefine statistical significance. *Nature Human Behaviour* 2 (2018), 6–10.
- [7] Bob Briscoe, Mirja Kühlewind, and Richard Scheffenecker. 2026. More Accurate Explicit Congestion Notification (AccECN) Feedback in TCP. RFC 9768. (April 2026).
- [8] Bob Briscoe, Koen De Schepper, Olga Albisser, Joakim Misund, Olivier Tilmans, Mirja Kühlewind, Richard Scheffenecker, Marcelo Bagnulo, and Asad Ahmed. 2019. Implementing the “TCP Prague” Requirements for L4S. In *Proceedings of the Netdev 0x13 Conference*.
- [9] Bob Briscoe, Koen De Schepper, and Marcelo Bagnulo. 2016. *Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service: Problem Statement*. Internet-Draft draft-briscoe-tsvwg-aqm-tcpm-rmcat-l4s-problem-00. Internet Engineering Task Force. Became RFC 9330 (2023).
- [10] Bob Briscoe, Koen De Schepper, Marcelo Bagnulo, and Greg White. 2023. Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture. RFC 9330. (Jan. 2023).
- [11] Davide Brunello. 2020. *L4S in 5G networks*. Master's thesis. KTH, School of Electrical Engineering and Computer Science (EECS).
- [12] Davide Brunello, Ingemar Johansson S, Mustafa Ozger, and Cicek Cavdar. 2021. Low Latency Low Loss Scalable Throughput in 5G Networks. In *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*. 1–7.
- [13] CableLabs. 2020. DOCSIS Technology: What's Changed in the Past Year and Why It Matters. (14 Oct. 2020). <https://www.cablelabs.com/blog/docsis-technology-whats-changed-in-the-past-year-and-why-it-matters> CableLabs blog post; accessed 2026-01-27.
- [14] Neal Cardwell, Ian Swett, and Joseph Beshay. 2026. *BBR Congestion Control*. Internet-Draft draft-ietf-ccwg-bbr-05. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-ccwg-bbr-05/> Work in Progress.
- [15] Shubham Chaudhary, Navneet Mishra, Keshav Gambhir, Tanmay Rajore, Arani Bhattacharya, and Mukulika Maity. 2025. COMPACT: Content-aware Multi-path Live Video Streaming for Online Classes using Video Tiles. In *Proceedings of the 16th ACM Multimedia Systems Conference (MMSys '25)*. Association for Computing Machinery, New York, NY, USA, 201–213.
- [16] Comcast. 2025. Comcast Introduces Nation's First Ultra-Low Lag Xfinity Internet Experience with Meta, NVIDIA, and Valve. (29 Jan. 2025). <https://corporate.comcast.com/press/releases/comcast-introduces-nations-first-ultra-low-lag-xfinity-internet-experience-with-meta-nvidia-and-valve> Corporate press release, Jan. 29, 2025. Accessed Jan. 27, 2026.
- [17] De Fré, Matthias and van der Hooft, Jeroen and Chang, Chia-Yu and De Schepper, Koen and Alfacc, Patrice Rondao and De Vleeschauwer, Danny and Wauters, Tim and Steenkiste, Peter and De Turck, Filip. 2025. Low-latency volumetric video conferencing in congested networks through L4S. In *PROCEEDINGS OF THE 16TH ACM MULTIMEDIA SYSTEMS CONFERENCE, MMSYS 2025*. ACM, 113–123.
- [18] Koen De Schepper, Olga Bondarenko, Ing-Jyh Tsang, and Bob Briscoe. 2016. PI2: A Linearized AQM for both Classic and Scalable TCP. In *Proceedings of the 12th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '16)*. Association for Computing Machinery, New York, NY, USA, 105–119.
- [19] Bradley Efron. 1979. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* 7, 1 (1979), 1–26.
- [20] Ericsson. 2025. SoftBank Corp., Ericsson and Qualcomm Technologies trial 5G and 5G Advanced capabilities including L4S on 5G SA commercial network. (23 Dec. 2025). <https://www.ericsson.com/en/press-releases/2025/12/softbank-corp--ericsson-qualcomm-technologies-trial-5g-and-5g-advanced-capabilities-including-l4s-on-5g-sa-commercial-network> Ericsson press release; accessed 2026-02-10.
- [21] ESnet. [n. d.]. iPerf - The TCP, UDP and SCTP network bandwidth measurement tool. <https://iperf.fr/>. ([n. d.]). Accessed: 2026-02-10.
- [22] Federal Communications Commission. 2020. *Appendix G-2: International Broadband Data Report Appendices. G-1–G-4*. Appendix to the 2020 Communications Marketplace Report FCC 20-188A8. Federal Communications Commission. <https://docs.fcc.gov/public/attachments/FCC-20-188A8.pdf>
- [23] Ronald A. Fisher. 1925. *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh.
- [24] Sadjad Fouladi, John Emmons, Emre Orbay, Catherine Wu, Riad S. Wahby, and Keith Winstein. 2018. Salsify: Low-Latency Network Video through Tighter Integration between a Video Codec and a Transport Protocol. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*. USENIX Association, Renton, WA, 267–282. <https://www.usenix.org/conference/nsdi18/presentation/fouladi>
- [25] Jim Gettys and Kathleen Nichols. 2012. Bufferbloat: dark buffers in the internet. *Commun. ACM* 55, 1 (Jan. 2012), 57–65.
- [26] Prateesh Goyal, Anup Agarwal, Ravi Netravali, Mohammad Alizadeh, and Hari Balakrishnan. 2020. ABC: A Simple Explicit Congestion Controller for Wireless Networks. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. USENIX Association, Santa Clara, CA, 353–372.
- [27] Philippe Graff, Xavier Marchal, Thibault Cholez, Bertrand Mathieu, Stéphane Tuffin, and Olivier Festor. 2024. Improving Cloud Gaming Traffic QoS: A Comparison Between Class-Based Queuing Policy and L4S. In *2024 8th Network Traffic*

- Measurement and Analysis Conference (TMA)*. 1–10. <https://doi.org/10.23919/TMA62044.2024.10558920>
- [28] Stephen Hemminger. 2005. Network Emulation with NetEm. In *Proceedings of the 6th Australian National Linux Conference (Linux.conf.au)*.
- [29] Tom Henderson. [n. d.]. DualQ Coupled PI2 Queue Disc implementation for ns-3. <https://gitlab.com/tomhenderson/ns-3-dev/-/blob/dual-queue/src/traffic-control/model/dualq-coupled-pi2-queue-disc.cc>. ([n. d.]). Source code file, dual-queue branch. Accessed: 2026-06-07.
- [30] Antonio Hernández-Vela, Miguel Ángel Bautista, Xavier Perez-Sala, Victor Ponce-López, Sergio Escalera, Xavier Baró, Oriol Pujol, and Cecilio Angulo. 2014. Probability-based Dynamic Time Warping and Bag-of-Visual-and-Depth-Words for Human Gesture Recognition in RGB-D. *Pattern Recogn. Lett.* 50, C (Dec. 2014), 112–121.
- [31] Kelly Hill. 2025. Test and Measurement: Comcast touts ‘ultra-low lag connectivity’ using L4S. (29 Jan. 2025). <https://www.rcrwireless.com/20250129/uncategorized/comcast-l4s-rcr-wireless-news-article>; accessed 2026-01-27.
- [32] Xiangjie Huang, Jiayang Xu, Haiping Wang, Hebin Yu, Sandesh Dhawaskar Sathyanarayana, Shu Shi, and Zili Meng. 2025. ACE: Sending Burstiness Control for High-Quality Real-time Communication. In *Proceedings of the ACM SIGCOMM 2025 Conference (SIGCOMM ’25)*. Association for Computing Machinery, New York, NY, USA, 1182–1198.
- [33] Mikel Irazabal and Navid Nikaein. 2024. TC-RAN: A Programmable Traffic Control Service Model for 5G/6G SD-RAN. *IEEE J.Sel. A. Commun.* 42, 2 (Feb. 2024), 406–419.
- [34] Shaohua Jiang and Zheng Chen. 2023. Application of dynamic time warping optimization algorithm in speech recognition of machine translation. *Heliyon* 9, 11 (2023), e21625. <https://doi.org/10.1016/j.heliyon.2023.e21625>
- [35] Ingemar Johansson. 2014. Self-clocked rate adaptation for conversational video in LTE. In *Proceedings of the 2014 ACM SIGCOMM Workshop on Capacity Sharing Workshop (CSWS ’14)*. Association for Computing Machinery, New York, NY, USA, 51–56.
- [36] Ingemar Johansson and Zaheduzzaman Sarker. 2017. Self-Clocked Rate Adaptation for Multimedia. RFC 8298. (Dec. 2017).
- [37] Nikhil Kansal, Murali Ramanujam, and Ravi Netravali. 2021. Alohamora: Reviving HTTP/2 Push and Preload by Adapting Policies On the Fly. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. USENIX Association, 269–287.
- [38] Tom Kelly. 2003. Scalable TCP: improving performance in highspeed wide area networks. *SIGCOMM Comput. Commun. Rev.* 33, 2 (April 2003), 83–91.
- [39] Eamonn J. Keogh and Michael J. Pazzani. 2000. Scaling up dynamic time warping for datamining applications. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’00)*. Association for Computing Machinery, New York, NY, USA, 285–289.
- [40] L4S Team. 2020. udp_prague: UDP Prague Congestion Control Implementation. https://github.com/L4steam/udp_prague. (2020). Accessed: 2026-01-17.
- [41] Bob Lantz, Brandon Heller, and Nick McKeown. 2010. A network in a laptop: rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (Hotnets-IX)*. Association for Computing Machinery, New York, NY, USA, Article 19, 6 pages.
- [42] Jason Livingood. 2024. LowLatency: discussion of bufferbloat, latency, and low latency networking. (15 Feb. 2024). https://www.linkedin.com/posts/jlivingood_lowlatency-latency-bufferbloat-activity-7420536190052802560-w3x2/ LinkedIn post; accessed 2026-01-27.
- [43] Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. 2017. Neural Adaptive Video Streaming with Pensieve. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM ’17)*. Association for Computing Machinery, New York, NY, USA, 197–210.
- [44] Lucas V. Monteiro, Vinícius S. Simão, Rodrigo de B. Lira, Leandro C. de Almeida, Ruan D. Gomes, and Paulo Ditarso Maciel. 2024. L4S in Private 5G Industrial Networks: A Case Study for Real-Time Video Transmission in Programmable Networks. In *2024 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*. 1–4. <https://doi.org/10.1109/NFV-SDN61811.2024.10807467>
- [45] Ravi Netravali, Vikram Nathan, James Mickens, and Hari Balakrishnan. 2018. Vesper: Measuring Time-to-Interactivity for Web Pages. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*. USENIX Association, Renton, WA, 217–231.
- [46] Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. 2015. Mahimahi: accurate record-and-replay for HTTP. In *Proceedings of the 2015 USENIX Conference on Usenix Annual Technical Conference (USENIX ATC ’15)*. USENIX Association, USA, 417–429.
- [47] Ravi Netravali, Anirudh Sivaraman, Keith Winstein, Somak Das, Ameesh Goyal, and Hari Balakrishnan. 2014. Mahimahi: a lightweight toolkit for reproducible web measurement. In *Proceedings of the 2014 ACM Conference on SIGCOMM (SIGCOMM ’14)*. Association for Computing Machinery, New York, NY, USA, 129–130.
- [48] Yunzhe Ni, Zhilong Zheng, Xianshang Lin, Fengyu Gao, Xuan Zeng, Yirui Liu, Tao Xu, Hua Wang, Zhidong Zhang, Senlang Du, Guang Yang, Yuanhao Su, Dennis Cai, Hongqiang Harry Liu, Chenren Xu, Ennan Zhai, and Yunfei Ma. 2023. CellFusion: Multipath Vehicle-to-Cloud Video Streaming with Network Coding in the Wild. In *Proceedings of the ACM SIGCOMM 2023 Conference (ACM SIGCOMM ’23)*. Association for Computing Machinery, New York, NY, USA, 668–683.
- [49] Kathleen Nichols and Van Jacobson. 2012. Controlling queue delay. *Commun. ACM* 55, 7 (July 2012), 42–50.
- [50] Nokia Bell Labs. 2024. L4S – Nokia Bell Labs pioneers L4S, the crucial enabler for large-scale deployments of real-time applications. (2024). <https://www.nokia.com/bell-labs/research/l4s/> Nokia Bell Labs research page; accessed 2026-02-9.
- [51] ns-3 Consortium. 2020. ns-3 L4S Support. <https://www.nsnam.org/wiki/L4S-support>. (2020). ns-3 Wiki page, last updated November 19, 2020; accessed 2026-01-23.
- [52] ns-3 DCE Maintainers. [n. d.]. DCE User Syscalls. <https://github.com/direct-code-execution/ns-3-dce/blob/master/doc/source/dce-user-syscalls.rst>. ([n. d.]). Documentation source file. Accessed: 2026-06-07.
- [53] ns-3 DCE Maintainers. [n. d.]. Direct Code Execution: Introduction. <https://ns-3-dce.readthedocs.io/en/latest/intro.html>. ([n. d.]). Documentation page. Accessed: 2026-06-07.
- [54] ns-3 Project. [n. d.]. ns-3 Doxygen: PointToPointNetDevice Class Reference. https://www.nsnam.org/docs/doxygen/dc/d89/classes3_1_1_point_to_point_net_device.html. ([n. d.]). API documentation page. Accessed: 2026-06-07.
- [55] Ookla. 2024. Introducing Loaded Latency. (2024). <https://www.ookla.com/articles/introducing-loaded-latency> Ookla article; accessed 2026-01-27.
- [56] Ookla. 2025. Loaded Latency and L4S: Measuring and Improving Network Responsiveness. Webinar recording, video hosted on Wistia. (2025). <https://ookla.wistia.com/medias/3fkfud7q7i> Accessed: 2026-02-26.
- [57] Guangjin Pan, Shugong Xu, and Pin Jiang. 2025. Optimizing 5G-Advanced Networks for Time-Critical Applications: The Role of L4S. *IEEE Wireless Communications* 32, 2 (2025), 230–237. <https://doi.org/10.1109/MWC.004.2400094>
- [58] Rong Pan, Preethi Natarajan, Chiara Pignone, Mythili Suryanarayana Prabhu, Vijay Subramanian, Fred Baker, and Bill VerSteeg. 2013. PIE: A lightweight control scheme to address the bufferbloat problem. In *2013 IEEE 14th International Conference on High Performance Switching and Routing (HPSR)*. 148–155.
- [59] Thanawin Rakthanmanon, Bilson Campana, Abdullah Mueen, Gustavo Batista, Brandon Westover, Qiang Zhu, Jesin Zakaria, and Eamonn Keogh. 2012. Searching and mining trillions of time series subsequences under dynamic time warping. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’12)*. Association for Computing Machinery, New York, NY, USA, 262–270.
- [60] Murali Ramanujam, Helen Chen, Shaghayegh Mardani, and Ravi Netravali. 2022. Floo: automatic, lightweight memoization for faster mobile apps. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys ’22)*. Association for Computing Machinery, New York, NY, USA, 168–182.
- [61] Devdeep Ray, Jack Kosaian, K. V. Rashmi, and Srinivasan Seshan. 2019. Vantage: optimizing video upload for time-shifted viewing of social live streams. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM ’19)*. Association for Computing Machinery, New York, NY, USA, 380–393.
- [62] George F. Riley and Thomas R. Henderson. 2010. *The ns-3 Network Simulator*. Springer Berlin Heidelberg, Berlin, Heidelberg, 15–34.
- [63] Luigi Rizzo. 1997. Dummynet: a simple approach to the evaluation of network protocols. *SIGCOMM Comput. Commun. Rev.* 27, 1 (Jan. 1997), 31–41.
- [64] Michael Rudow, Francis Y. Yan, Abhishek Kumar, Ganesh Ananthanarayanan, Martin Ellis, and K.V. Rashmi. 2023. Tambur: Efficient loss recovery for videoconferencing via streaming codes. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 953–971. <https://www.usenix.org/conference/nsdi23/presentation/rudow>
- [65] Aleksandra Rutkowska and Magdalena Szyszko. 2022. New DTW Windows Type for Forward- and Backward-Lookingness Examination. Application for Inflation Expectation. *Computational Economics* 59, 2 (February 2022), 701–718.
- [66] H. Sakoe and S. Chiba. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26, 1 (1978), 43–49.
- [67] Fatih Berkay Sarpkaya, Fraida Fund, and Shivendra Panwar. [n. d.]. To Adopt or Not to Adopt L4S-Compatible Congestion Control? Understanding Performance in a Partial L4S Deployment. Presented at the IETF 122 MAPRG meeting. ([n. d.]). <https://datatracker.ietf.org/meeting/122/materials/slides-122-maprg-to-adopt-or-not-to-adopt-l4s-compatible-congestion-control-understanding-performance-in-a-partial-l4s-deployment-00> Presentation slides. Accessed: 2026-06-13.
- [68] Fatih Berkay Sarpkaya, Fraida Fund, and Shivendra Panwar. 2025. To Adopt or Not to Adopt L4S-Compatible Congestion Control? Understanding Performance in a Partial L4S Deployment. In *Passive and Active Measurement: 26th International Conference, PAM 2025, Virtual Event, March 10–12, 2025, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 217–246.
- [69] Fatih Berkay Sarpkaya, Ashutosh Srivastava, Fraida Fund, and Shivendra Panwar. 2024. To switch or not to switch to TCP Prague? Incentives for adoption in a partial L4S deployment. In *Proceedings of the 2024 Applied Networking Research*

- Workshop (ANRW '24). Association for Computing Machinery, New York, NY, USA, 45–52.
- [70] Deol Satish, Jonathan Kua, and Shiva Raj Pokhrel. 2024. Active Queue Management in L4S with Asynchronous Advantage Actor-Critic: A FreeBSD Networking Stack Perspective. *Future Internet* 16, 8 (2024).
 - [71] John Saw. 2025. Unlocking L4S on 5G Advanced. (21 July 2025). <https://www.t-mobile.com/news/network/unlock-l4s-5g-advanced> T-Mobile press release; accessed 2026-01-27.
 - [72] Koen De Schepper, Olga Albisser, Olivier Tilman, and Bob Briscoe. 2022. *Dual Queue Coupled AQM: Deployable Very Low Queuing Delay for All*. Technical Report. Technical Report. https://www.bobbriscoe.net/projects/latency/dctth_ton_tr.pdf
 - [73] Koen De Schepper, Olga Albisser, Ing-Jyh Tsang, and Bob Briscoe. 2019. “Data Center to the Home”: Ultra-Low Latency for All. Technical Report. Technical Report. http://www.bobbriscoe.net/projects/latency/dctth_journal_draft20190726.pdf
 - [74] Koen De Schepper and Bob Briscoe. 2023. The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S). RFC 9331. (Jan. 2023).
 - [75] Koen De Schepper, Bob Briscoe, and Greg White. 2023. Dual-Queue Coupled Active Queue Management (AQM) for Low Latency, Low Loss, and Scalable Throughput (L4S). RFC 9332. (Jan. 2023).
 - [76] Helena Skutkova, Martin Vitek, Petr Babula, Rene Kizek, and Ivo Provaznik. 2013. Classification of Genomic Signals Using Dynamic Time Warping. *BMC Bioinformatics* 14, Suppl 10 (2013), S1.
 - [77] Jangwoo Son, Yago Sanchez, Cornelius Hellge, and Thomas Schierl. 2024. Adaptable L4S Congestion Control for Cloud-Based Real-Time Streaming Over 5G. *IEEE Open Journal of Signal Processing* 5 (2024), 841–849. <https://doi.org/10.1109/OJSP.2024.3405719>
 - [78] Jangwoo Son, Yago Sanchez, Cornelius Hellge, and Thomas Schierl. 2024. Split Rendering with L4S Over 5G for Latency Critical Interactive XR Applications. *IEEE Communications Magazine* 62, 8 (2024), 46–52.
 - [79] Philipp Steininger, Rastin Pries, Yash Deshpande, Kaan Aykurt, Chia-Yu Chang, Koen De Schepper, and Wolfgang Kellerer. 2025. Performance Evaluation of L4S in XR Scenarios. In *Proceedings of the IFIP Networking Conference 2025*. Limassol, Cyprus. IFIP Networking 2025, May 26–30, 2025.
 - [80] Maria Eduarda Veras, Eduardo Freitas, Assis T. de Oliveira Filho, Djamel Sadok, and Judith Kelter. 2026. Implementing the L4S Architecture in the ns-3 Simulator. (2026). arXiv:cs.NI/2603.20166 <https://arxiv.org/abs/2603.20166>
 - [81] Vodafone. 2025. Vodafone and Nokia achieve world-first trial of lag-busting technology on fibre broadband. (10 Feb. 2025). <https://www.vodafone.com/news/newsroom/technology/vodafone-and-nokia-achieve-world-first-trial-of-lag-busting-technology-on-fibre-broadband> Vodafone press release; accessed 2026-01-27.
 - [82] Haoran Wan and Kyle Jamieson. 2025. L4Span: Spanning Congestion Signaling over NextG Networks for Interactive Applications. *Proc. ACM Netw.* 3, CoNEXT4, Article 25 (Nov. 2025), 25 pages. <https://doi.org/10.1145/3768972>
 - [83] Bo Wang, Mingwei Xu, Fengyuan Ren, and Jianping Wu. 2022. Improving Robustness of DASH Against Unpredictable Network Variations. *Trans. Multi.* 24 (Jan. 2022), 323–337. <https://doi.org/10.1109/TMM.2021.3050086>
 - [84] Shibo Wang, Shusen Yang, Hailiang Li, Xiaodan Zhang, Chen Zhou, Chenren Xu, Feng Qian, Nanbin Wang, and Zongben Xu. 2022. SalientVR: saliency-driven mobile 360-degree video streaming with gaze information. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking (MobiCom '22)*. Association for Computing Machinery, New York, NY, USA, 542–555. <https://doi.org/10.1145/3495243.3517018>
 - [85] Ranysha Ware, Adithya Abraham Philip, Nicholas Hungria, Yash Kothari, Justine Sherry, and Srinivasan Seshan. 2024. CCAnalyzer: An Efficient and Nearly-Passive Congestion Control Classifier. In *Proceedings of the ACM SIGCOMM 2024 Conference (ACM SIGCOMM '24)*. Association for Computing Machinery, New York, NY, USA, 181–196. <https://doi.org/10.1145/3651890.3672255>
 - [86] Greg White, Karthik Sundaresan, and Bob Briscoe. 2019. *Low Latency DOTS - Technology Overview*. Internet-Draft draft-white-tsvwg-ll-d-00. Internet Engineering Task Force. Expired Internet-Draft.
 - [87] Yaxiong Xie and Kyle Jamieson. 2022. NG-Scope: Fine-Grained Telemetry for NextG Cellular Networks. *Proc. ACM Meas. Anal. Comput. Syst.* 6, 1, Article 12 (Feb. 2022), 26 pages. <https://doi.org/10.1145/3508032>
 - [88] Francis Y. Yan, Hudson Ayers, Chenzhi Zhu, Sadjad Fouladi, James Hong, Keyi Zhang, Philip Levis, and Keith Winstein. 2020. Learning in situ: a randomized experiment in video streaming. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. USENIX Association, Santa Clara, CA, 495–511. <https://www.usenix.org/conference/nsdi20/presentation/yan>
 - [89] Hyunho Yeo, Youngmok Jung, Jaehong Kim, Jinwoo Shin, and Dongsu Han. 2018. Neural adaptive content-aware internet video delivery. In *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation (OSDI'18)*. USENIX Association, USA, 645–661.
 - [90] Danfu Yuan, Yuanhong Zhang, Weizhan Zhang, Xuncheng Liu, Haipeng Du, and Qinghua Zheng. 2022. PRIOR: deep reinforced adaptive video streaming with attention-based throughput prediction. In *Proceedings of the 32nd Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV '22)*. Association for Computing Machinery, New York, NY, USA, 36–42. <https://doi.org/10.1145/3534088.3534348>
 - [91] Danesh Zeynali, Emilia N. Weyulu, Seifeddine Fathalli, Balakrishnan Chandrasekaran, and Anja Feldmann. 2024. BBRv3 in the public Internet: a boon or a bane?. In *Proceedings of the 2024 Applied Networking Research Workshop (ANRW '24)*. Association for Computing Machinery, New York, NY, USA, 97–99.
 - [92] Bo Zhang, Thiago Teixeira, and Yuriy Reznik. 2021. Performance of Low-Latency HTTP-based Streaming Players. In *Proceedings of the 12th ACM Multimedia Systems Conference (MMSys '21)*. Association for Computing Machinery, New York, NY, USA, 356–362. <https://doi.org/10.1145/3458305.3478442>

APPENDIX

A MAHIMAHİ VS. KERNEL DEQUEUE CADENCE

To investigate the statistical misalignment described in §6 between the DualPI2 Mahimahi module and the kernel reference implementation, we collect dequeue traces, which consist of a series of timestamps corresponding to packet dequeue events from both systems. Fig. 10 shows a snippet of a dequeue trace obtained by running an `iperf3` flow using TCP Prague through the DualPI2 AQM with the default parameters outlined in §5, in the medium BDP scenario (50 Mbps with an RTT of 40 ms).




Figure 10: Dequeue events from Mahimahi and kernel. Each vertical line represents the beginning of a burst, and each dot represents a packet. For Mahimahi, each visible dot is in fact multiple dots stacked on top of each other.

The figure shows that the packet dequeue in Mahimahi is visibly more sparse than the smooth and regular dequeue in the kernel setting. Mahimahi has a granularity of 1 ms, resulting from how the network traces are encoded. Indeed, the trace is a series of millisecond timestamps, where if a given millisecond appears in the trace n times, then a maximum of n MTU-size packets (1500 bytes by default) can be released during that one millisecond. This naturally creates a burst of packets every millisecond, and the higher the throughput, the higher the amplitude of the burst.

To quantify the dequeue period in both Mahimahi and the kernel settings, we run a period extraction heuristic on the dequeue traces and obtain a period of 998 μ s for Mahimahi, confirming a 1 ms granularity, and 242 μ s for the kernel, corresponding to dequeue events that are about four times more frequent than in Mahimahi.

Fig. 11 shows a sample dot plot of dequeued packets over an arbitrarily small interval (we choose 120 μ s), where each dot is a single packet. In the top graph, representing Mahimahi, frequent bursts of up to 12 packets can be observed. By contrast, bursts in the kernel setting are both less frequent and smaller, with a maximum of two packets per interval.

Together, these observations highlight the inherent difficulty of emulating a low-latency AQM such as DualPI2. One major consequence of the 1 ms default granularity of Mahimahi is that the L4S queue’s default `step_thresh` of 1 ms cannot be precisely enforced, as packets build up in the queue while waiting for the next millisecond mark to be released. This helps explain why the throughput results improved when we set `step_thresh` and target to higher values for higher BDP regimes. Adjusting the DualPI2 parameters, as we do in this work, therefore provides a practical path toward using an emulated DualPI2 module and opens up opportunities for rapid end-to-end L4S research and experimentation.




Figure 11: Packets dequeued in an interval of 120 μ s. The top figure represents Mahimahi, while the bottom represents the kernel.

Alternatively, efforts aimed at calibrating emulation tools such as Mahimahi for low-latency experiments could be considered, as low-latency research is only going to receive increasing attention.

B BOOTSTRAP CONFIDENCE INTERVALS FOR FURTHER STATISTICAL VALIDATION

While $\hat{p}_{\max}$ provides a point estimate of the exceedance rate, it is computed from a finite number of experimental runs and is therefore subject to sampling variability. To ensure that our conclusions are not driven by a particular realization of the data, we quantify the uncertainty in $\hat{p}_{\max}$ using nonparametric bootstrap resampling with replacement [19]. This procedure approximates the sampling distribution of $\hat{p}_{\max}$ under repeated experiments, allowing us to assess how the exceedance probability would vary under repeated sampling of runs.

B.1 Bootstrap Resampling Implementation

In our implementation, resampling is performed at the run level within each system. For each replicate, we sample runs with replacement from the kernel set and independently from the Mahimahi set, preserving the original sample size in each group. Using these resampled run sets, we recompute all pairwise within-system and cross-system differences, re-estimate the tolerance threshold $\epsilon_{\max}$, and recompute the exceedance statistic $\hat{p}_{\max}$. Repeating this procedure across B replicates yields an empirical bootstrap distribution of $\hat{p}_{\max}$ from which confidence intervals are constructed.

We use $B = 2000$ bootstrap replicates, which provides stable percentile confidence intervals: the Monte Carlo error of the bootstrap quantiles decreases as $O(B^{-1/2})$, and $B = 2000$ yields sufficiently fine resolution for the 2.5th and 97.5th percentiles (indices 50 and 1950), ensuring that our conclusions are not sensitive to the number of replicates.

B.2 Bootstrap Confidence Interval

Let $\hat{p}_{\max}^{(1)}, \dots, \hat{p}_{\max}^{(B)}$ denote the bootstrap replicates obtained from B resamples. We construct a 95% percentile bootstrap confidence interval as:

Table 6: Bootstrap 95% CIs for $\hat{p}_{\max}$: Default vs. Optimized across bandwidths and variables. ✓ = yes; × = no.

Bandwidth	Variable	Bootstrap 95% CI for $\hat{p}_{\max}$			Sig. ($\alpha = 0.05$)	
		Default	Optimized	Improved	Default	Optimized
12 Mbps	packets	[0.613, 0.954]	[1.000, 1.000]	×	×	×
	throughput	[0.000, 0.018]	[0.000, 0.280]	×	✓	×
	ecn_mark	[0.019, 0.418]	[0.023, 0.046]	×	×	✓
	packet_dropped_total	[0.000, 0.020]	[0.000, 0.039]	×	✓	✓
50 Mbps	packets	[0.975, 1.000]	[1.000, 1.000]	×	×	×
	throughput	[0.222, 0.470]	[0.050, 0.169]	✓	×	×
	ecn_mark	[0.022, 0.051]	[0.025, 0.324]	×	×	×
	packet_dropped_total	[0.000, 0.000]	[0.000, 0.000]	×	✓	✓
200 Mbps	packets	[0.814, 1.000]	[1.000, 1.000]	×	×	×
	throughput	[0.338, 0.779]	[0.000, 0.115]	✓	×	×
	ecn_mark	[0.015, 0.812]	[0.076, 0.158]	×	×	×
	packet_dropped_total	[0.000, 0.000]	[0.000, 0.000]	×	✓	✓

$$\left[\text{Quantile}_{0.025}(\hat{p}_{\max}^{(1:B)}), \text{Quantile}_{0.975}(\hat{p}_{\max}^{(1:B)}) \right]$$

This interval captures the central 95% of the empirical bootstrap distribution. The lower bound represents a conservative estimate of the smallest plausible exceedance probability under repeated sampling, while the upper bound represents the largest plausible exceedance probability consistent with the observed variability.

B.3 Decision Rule

We formalize the validation criterion as a one-sided test of the exceedance probability:

$$H_0 : p_{\max} \geq 0.05 \quad \text{versus} \quad H_1 : p_{\max} < 0.05.$$

We reject H_0 if the upper endpoint of the 95% bootstrap confidence interval for $\hat{p}_{\max}$ lies below 0.05. In this case, the data provides statistically significant evidence that the cross-system exceedance rate is below the prescribed 5% tolerance. Conversely, if the confidence interval overlaps or exceeds 0.05, we do not reject H_0 . This means there is not sufficient statistical evidence to conclude that the exceedance probability is below 5%.

To assess whether the Optimized configuration improves upon the Default configuration, we compare their bootstrap 95% confidence intervals for $p_{\max}$. We declare a statistically significant improvement only if the entire Optimized confidence interval lies strictly below the entire Default confidence interval. In this case, even accounting for sampling variability captured by the bootstrap, all plausible values of $p_{\max}$ under the Optimized configuration are smaller than those under Default, providing statistical evidence that the exceedance probability has been reduced.

B.4 Results

We summarize the validation results in Table 6, which reports, for each bandwidth and variable, the bootstrap 95% confidence intervals for $p_{\max}$ under both the Default and Optimized configurations. The table additionally indicates whether the Optimized configuration demonstrates a strict improvement over Default (non-overlapping confidence intervals) and whether statistical significance at the 5% level is achieved according to the decision rule.

Table 6 reports bootstrap 95% confidence intervals for $\hat{p}_{\max}$ under both configurations. Statistical significance, defined by the upper bound of the confidence interval lying strictly below the 5% tolerance threshold, is achieved only in isolated cases. Specifically, significance is observed for throughput and packet_dropped_total at 12 Mbps under the Default configuration; for ecn_mark and packet_dropped_total at 12 Mbps under the Optimized configuration; and for packet_dropped_total at 50 and 200 Mbps under both configurations. In all remaining cases, the confidence intervals overlap or exceed the 0.05 threshold, and we therefore fail to reject H_0 , indicating insufficient statistical evidence that the exceedance probability remains below the prescribed tolerance.

Although the Optimized configuration does not demonstrate a universal improvement across all variables, it achieves clear and strict gains in throughput at higher bandwidths. At 50 Mbps, the Optimized confidence interval [0.050, 0.169] lies entirely below the Default interval [0.222, 0.470], indicating a non-overlapping and strictly lower exceedance probability. An even stronger separation appears at 200 Mbps, where the Optimized interval [0.000, 0.115] is fully contained below the Default interval [0.338, 0.779]. In both cases, the entire Optimized confidence band is strictly smaller than the corresponding Default band, providing statistical evidence of improvement under our non-overlapping CI criterion. These results indicate that at moderate and high bandwidths, the Optimized configuration substantially reduces the probability of cross-system deviations exceeding within-system variability, yielding improved statistical consistency in throughput behavior.

B.5 Sample Size Justification for 100 Flows per Configuration

To justify the adequacy of 100 flows per configuration, we examine the bootstrap 95% confidence intervals computed from the collected data. In several cases, the intervals are already extremely narrow and lie entirely below the 0.05 tolerance threshold, including throughput at 12 Mbps under Default ([0.000, 0.018]) and packet_dropped_total at 12 Mbps under both configurations ([0.000, 0.020] and [0.000, 0.039]), as well as packet_dropped_total at 50 and 200 Mbps where the intervals collapse to [0.000, 0.000].




Figure 12: Bootstrap 95% CI width vs. number of flows n for ECN mark rate across all bandwidth configurations.

Because these intervals are both tight and fully contained below the tolerance level, additional flows would be unlikely to alter

the decision outcome. This indicates that, for these scenarios, the current sample size already provides stable and statistically decisive estimates of $p_{\max}$.

However, the relatively wide ECN-mark confidence intervals could be interpreted as indicating insufficient sampling and the potential need for additional runs. However, when plotted against the number of flows (see Fig. 12), the confidence intervals clearly shrink and stabilize as n approaches 100, demonstrating convergence rather than instability.

The slight dip observed in panel (f) at $n = 100$ is most likely due to bootstrap resampling variability, since it is not part of a consistent downward trend and neighboring sample sizes exhibit similar confidence interval widths.

Overall, this behavior justifies the choice of $n = 100$ flows for these configurations, as the interval widths have largely plateaued by this point. At the same time, it highlights an important characteristic of ECN marks: they exhibit inherently higher variability compared to other metrics, suggesting that their dispersion reflects natural stochastic behavior rather than insufficient sampling.