DESiRED - Dynamic, Enhanced, and Smart iRED: A P4-AQM with Deep Reinforcement Learning and In-band Network Telemetry

Leandro C. de Almeida, Washington Rodrigo Dias da Silva,
Thiago C. Tavares, Rafael Pasquini, Chrysa Papagianni and Fábio L. Verdi

October 30, 2023

Abstract

Active Queue Management (AQM) is a mechanism employed to alleviate transient congestion in network device buffers, such as routers and switches. Traditional AQM algorithms use fixed thresholds, like target delay or queue occupancy, to compute random packet drop probabilities. A very small target delay can increase packet losses and reduce link utilization, while a large target delay may increase queueing delays while lowering drop probability. Due to dynamic network traffic characteristics, where traffic fluctuations can lead to significant queue variations, maintaining a fixed threshold AQM may not suit all applications. Consequently, we explore the question: *What is the ideal threshold (target delay) for AQMs?* In this work, we introduce DESiRED (Dynamic, Enhanced, and Smart iRED), a P4-based AQM that leverages precise network feedback from In-band Network Telemetry (INT) to feed a Deep Reinforcement Learning (DRL) model. This model dynamically adjusts the target delay based on rewards that maximize application Quality of Service (QoS). We evaluate DESiRED in a realistic P4-based test environment running an MPEG-DASH service. Our findings demonstrate up to a 90x reduction in video stall and a 42x increase in high-resolution video playback quality when the target delay is adjusted dynamically by DESiRED.

1 Introduction

In the modern domain of computer networks, the necessity to meet progressively rigorous service requirements, including ultra-reliable low-latency communications and high bandwidth, has resulted in an unparalleled upsurge in network traffic, amplifying the intricacies associated with traffic management. Subsequently, approaches aimed at assisting congestion control mechanisms, such as Active Queue Management (AQM), are consistently embraced.

In scenarios where incoming packet rates exceed a network device's processing capacity, a transient queuing of packets occurs in the appropriate output queue, often causing transmission delays. To mitigate this bottleneck, an effective strategy involves

1

notification congestion status to the packet sender, allowing the congestion control algorithm to reduce transmission rates. The primary methods for conveying congestion conditions to senders include packet marking using Explicit Congestion Notification (ECN) bits and selective packet dropping. These approaches are the predominant means of communicating congestion information in network environments.

Traditionally, AQM mechanisms have been primarily focused on draining packets directly from queues, with the overarching objective of mitigating transient congestion occurrences and reducing the queuing delay. Prominent examples of these traditional AQM algorithms include Random Early Detection (RED) [1], Blue [2], CoDel [3], CAKE [4], and PIE [5]. More recently, owing to the inherent flexibility of the programmable data plane (PDP), the prevailing state-of-the-art AQM solutions designed to operate within PDP hardware environments and made publicly accessible comprise iRED [6], P4-CoDel [7], and the (dual) PI2 [8]. These AQM implementations exemplify the synergy between novel programmable data plane capabilities and the evolving demands of congestion control within modern network infrastructures.

An integral aspect of AQM algorithms pertains to the selection of an appropriate threshold value, often determined based on considerations of either queue delay (referred to as the target delay) or queue depth. Opting for an excessively small threshold value can lead to an increased occurrence of packet losses, resulting in a higher drop probability while reducing overall link utilization. Conversely, employing a high threshold value can lead to extended queuing delays but a lower likelihood of packet drops, characterized by a reduced drop probability. Additionally, the dynamic nature of network traffic necessitates the avoidance of static threshold values for specific applications. In this context, we explore this issue as the **fixed target delay problem**, as illustrated in Fig. 1, delving into the intricate dynamics of threshold determination in AQM algorithms

Figure 1: The trade-off: If the Target Delay is small, can increase packet losses and decrease link utilization. If is high, increase queueing delays and decrease packet drops.

At the core of this matter lies a fundamental trade-off, giving rise to a pivotal question: *What is the ideal target delay for AQM?* Estimating this value presents a challenging task. However, recent advancements in the field of artificial intelligence as applied to computer networks [9] introduce a potential avenue, leveraging the capabilities of Deep Reinforcement Learning (DRL) as a powerful tool to enhance decision-making in addressing this challenge.

Although DRL models are known for their appetite for data, the provision of real-

time data at the requisite granularity has posed an obstacle within the realm of computer networks. However, recent advances in the domain of PDP, in tandem with the integration of In-band Network Telemetry (INT) [10], have conducted a paradigm shift. These advancements have presented us with the capability to attain granular visibility, discernible on a per-packet basis, effectively altering the scenario of the challenges associated with data availability in the context of DRL applications within computer networks.

The hypothesis of this study posits that INT measurements can serve as valuable input features for a DRL model. This DRL model is intended to dynamically adjust the target delay, departing from our prior work with a fixed target delay in iRED [6]. The overarching goal is to utilize this DRL model for real-time optimization of QoS, thereby introducing a novel approach aimed at enhancing network performance and adaptability.

iRED represents a pioneering P4-based algorithm that introduced the concept of disaggregated AQM in PDP hardware. Disaggregated AQM involves the segmentation of AQM operations into distinct blocks, specifically Ingress and Egress, within the PDP architecture. Addittionally, iRED achieves full compliance with the L4S framework (Low Latency, Low Loss, and Scalable Throughput) [11]. It accomplishes this by categorizing traffic as either Classic (subject to dropping) or Scalable [1] (marked with the ECN bit), thus ensuring fairness among various flows through a combined packet dropping and marking mechanism.

Through the integration of INT, DRL, and the iRED framework, we introduce the innovative paradigm of DESiRED (Dynamic, Enhanced, and Smart iRED). To our knowledge, DESiRED serves as the leading implementation of a dynamic AQM system based on P4 architecture. This advancement combines the cutting-edge capabilities of fine-grained network measurements enabled by INT with the cognitive capabilities provided by the Deep Q-Network (DQN), thereby representing an integrated embodiment of state-of-the-art progress in the field of AQM.

We undertake a comprehensive evaluation of DESiRED within a realistic testbed environment, focusing on the delivery of an MPEG-DASH (Dynamic Adaptive Streaming over HTTP) service [12]. Our experiments involve the provision of diverse video catalogs to video clients traversing a programmable network. Fine-grained INT measurements, collected at line rate in the data plane, are utilized to inform the DRL mechanism in the control plane. The DRL mechanism guides the agent's actions, dynamically adjusting the target delay to optimize the QoS for the DASH service. This forms a Smart Control Closed Loop, as depicted in Fig. 2

Our empirical findings elucidate that DESiRED wields an impact, with the potential to alleviate video stall occurrences by a factor of up to 90x. Moreover, the enhancement in the QoS within the MPEG-DASH framework is evident, as measured by an augmentation of up to 42x in terms of Frames per Second (FPS), underscoring the considerable efficacy of DESiRED in elevating the video streaming experience. In summary, the main contributions of this work are:

1. We design and implement DESiRED, which is a smart Closed Control Loop

---

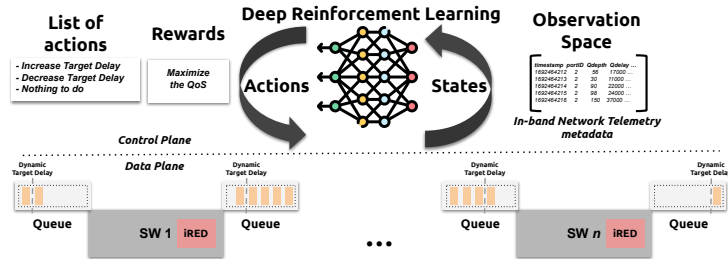[1]TCP Prague in L4S framework.

Figure 2: The Smart Control Closed Loop with DESiRED: The fine-grained INT measurements provide wide visibility into the state of the network in each observation space. The DRL mechanism guides the agent to dynamically adjust the target delay to maximize the QoS.

that unifies the state of the art in network telemetry (INT), Deep Reinforcement Learning (DQN), and congestion control in-network (AQM).

2. We conduct an evaluation of the DESiRED algorithm within the context of a DASH service. This entails the practical implementation of DESiRED within a real-setup DASH environment, followed by a systematic evaluation of its performance and effectiveness.

3. We have created and made publicly available datasets used throughout our experiments that encompass network and application data, collectively characterizing the complexities of an adaptive video service.

The remainder of the paper is organized as follows. In Section 2 we describe INT and DRL fundamental concepts. Additionally, we detail DESiRED, describing the main components implemented in the P4 language (data plane) and the DRL integration (control plane) in Section 3. In Section 4, the experiments and evaluation are presented, including a brief view of the testbed and the workloads used. Results and discussion are detailed in Section 5. Some Lessons learned are given in Section 6. Finally, the conclusions are depicted in Section 7.

2 Background

In this section, we expound upon the foundational principles that underpin the functionality of DESiRED. Sub-section 2.1 provides a concise elucidation of the programmable data plane and In-band Network Telemetry. Furthermore, Section 2.2 delves into the principal facets of Deep Reinforcement Learning.

2.1 In-band Network Telemetry

Recent progress in programmable hardware and the utilization of the P4 language [13] have enabled network devices to autonomously report the network's state, eliminating

the need for direct control plane intervention [14]. In this scenario, packets incorporate telemetry instructions within their header fields, facilitating the fine-grained collection and recording of network data.

The telemetry instructions are defined in the INT data plane specification [10]. Figure 3 illustrates the operation of INT within an arbitrary network. The network comprises four end systems, namely *H1, H2, H3*, and *H4*, along with four nodes equipped with P4 and INT support, denoted as *S1, S2, S3*, and *S4*. Each network node possesses a set of metadata, represented by orange (S1), magenta (S2), green (S3), and blue (S4) rectangles. This metadata contains information specific to each node, such as Node ID, Ingress Port, Egress Spec, Egress Port, Ingress Global Timestamp, Egress Global Timestamp, Enqueue Timestamp, Enqueue Queue Depth, Dequeue Timedelta, and Dequeue Queue Depth, as specified in the V1Model architecture.
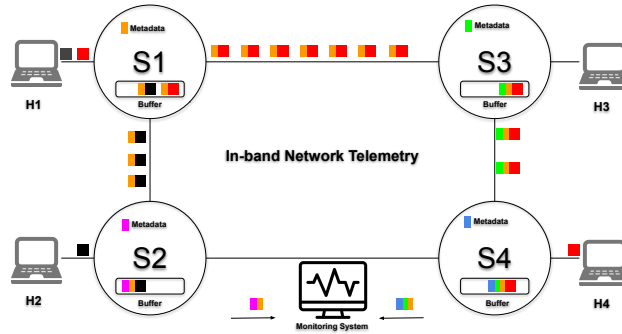

Figure 3: In-band Network Telemetry operation. INT metadata is appended on the packets in each hop. In the specific collection point, the monitoring system receives INT metadata.

In Figure 3, there are two distinct flows depicted: one represented by red packets and the other by black packets. The red flow is required to adhere to the prescribed network path *f1=H1, S1, S3, S4, H4*, while the black flow must traverse the designated path *f2=H1, S1, S2, H2*.

At each network hop along these paths, the data plane of the network devices employs telemetry instructions to facilitate the collection and inclusion of metadata within the packets as they traverse each node. This process is iteratively performed throughout the journey, starting from the first node after the source and concluding at the last node before reaching the destination. Upon reaching the destination node, the metadata is extracted from the packet and subsequently relayed to the monitoring system. The original packet is then directed to its final destination.

In addition to the modes delineated in the INT specification, alternative approaches exist for collecting metadata within programmable networks. One such approach involves the utilization of an "exclusive telemetry flow" to monitor the network's state, which, in this work, is referred to as "Out-of-band Network Telemetry" (ONT).

In the ONT scenario, dedicated probe packets are employed to gather metadata, eliminating the need for any modifications to the data packets associated with the ser-

vices operating within the network. The primary advantage of this approach lies in its ability to maintain the integrity of application traffic, as it traverses the programmable network without undergoing alterations, thereby mitigating issues related to packet growth, such as fragmentation.

Conversely, the use of an exclusive telemetry flow introduces additional overhead to the overall network traffic. This is due to the necessity of having a dedicated monitoring flow ONT for each service running within the network.

One of the primary advantages of employing telemetry lies in the exceptional level of granularity it offers. Every individual packet traversing the network carries pertinent information directly to the monitoring system at the line rate. This level of granularity aligns with the perspective presented in [15], wherein it is recognized that a substantial volume of data can prove immensely valuable for Deep Reinforcement Learning (DRL) algorithms, which have a voracious appetite for information.

2.2 Deep Reinforcement Learning

Reinforcement Learning (RL) is an Artificial Intelligence (AI) learning paradigm centered on actions and rewards. Unlike the conventional supervised and unsupervised learning approaches, where models learn from predefined dataset features, an RL learner, also known as an agent, interacts with an environment and receives rewards or penalties based on the actions it takes.

The model depicted in Figure 4 illustrates the formalization of a sequential decision-making strategy known as a Markov Decision Process (MDP). In this framework, the agent continually interacts with the environment by executing actions (A) at specific time steps (t) and observing new states (S_{t+1}) resulting from these actions. After each interaction, a reward value (R_{t+1}) is generated to assess the correctness of the action, with the aim of maximizing cumulative rewards throughout the agent's training process [16, 17].
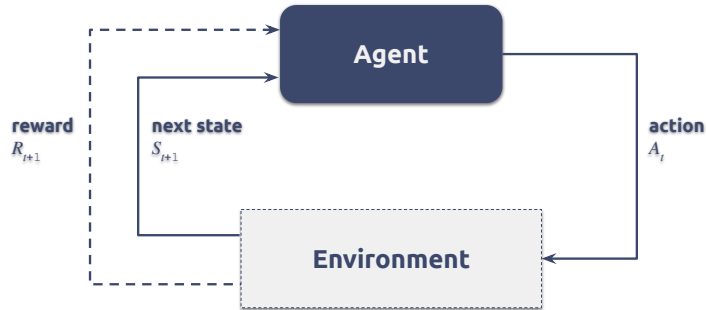

Figure 4: Interaction between the agent and environment in a Markov Decision Process. Adapted from [17].

In this context, the agent learns to maximize its cumulative rewards by determining a policy[2] that optimizes an action-value function, denoted as Q. This function estimates

---

[2]A policy defines the agent's strategy for associating actions with states. Such strategies can be stochastic,

the quality of actions taken by the agent in specific states.

Formally, an optimal action-value function, denoted as q_*, can be defined using the Bellman optimality equation [17]:

$$q_*(s, a) = \mathbb{E}[R_{t+1} + \gamma \max_{a'} q_*(S_{t+1}, a') \mid S_t = s, A_t = a] \tag{1}$$

Intuitively, the Bellman optimality equation suggests that the optimal Q-value for any state-action pair ($q_*(s, a)$) is expected to be the immediate reward obtained after taking an action (a) in a given state (s) at time step (t), augmented by the maximum expected return achievable by adhering to an optimal policy for subsequent state-action pairs, which are discounted[3] by γ [17].

Hence, the resolution of the Bellman optimality equation provides a pathway to ascertain an optimal policy, offering a potential solution to a RL problem. Nevertheless, it is imperative to acknowledge that, in practice, this solution is seldom feasible. It resembles an exhaustive search that requires consideration of all possible scenarios, involving the computation of occurrence probabilities and expected reward returns. Additionally, it relies on three assumptions that are often challenged when implementing solutions for real-world problems:

a) The accurate knowledge of environmental dynamics.

b) Sufficient computational resources to complete the computational requirements of the solution.

c) The adherence to the Markov property.

In light of these challenges, the only pragmatic approach to tackle the Bellman optimality equation is to seek a policy approximation derived from actual experiences, where transitions involving state s, action a, and reward r are considered, as opposed to relying solely on the expected outcomes [17].

When dealing with scenarios characterized by a well-defined set of finite states, it becomes feasible to model an approximation of the Bellman optimality equation using tabular data structures. Each entry within these structures corresponds to a state-action pair. In this context, the Q-Learning algorithm, introduced by Watkins [18], marked a significant milestone in the early stages of the RL paradigm.

Q-Learning is noteworthy for its direct approximation of the Bellman optimality equation, irrespective of the policy in use. It simplified the analysis of agent algorithms and facilitated early convergence proofs [17]. The Watkins Q-Learning algorithm is formally defined as follows:

$$Q(S_t, A_t) = (1 - \alpha)Q(S_t, A_t) + \alpha[R_{t+1} + \gamma \max_a Q(S_{t+1}, a)] \tag{2}$$

where the approximated optimal Q-value is calculated by blending the current Q-value, denoted as $Q(S_t, A_t)$, with the target temporal difference. The target temporal difference represents the reward R_{t+1} obtained when transitioning to the subsequent state

---

specifying probabilities for each action that could be taken when a particular state is observed [17].

[3]This discounting approach allows the agent to prioritize actions that maximize the cumulative rewards it receives in the future. The discount rate, denoted as γ, determines the present value of future rewards. For instance, a reward received k time steps in the future is worth only γ^{k-1} times what it would be worth if it were received immediately [17].

S_{t+1} after taking action a. This value is then weighted by the discount factor γ and modulated by a learning rate α ($0 \leq \alpha < 1$) [18, 17].

However, it is essential to acknowledge that Q-learning operates under the assumption of a tabular representation for state-action pairs and approximates the optimal Q-value in a linear fashion. In practice, real-world applications often exhibit complexity, characterized by non-linear relationships and encompassing high-dimensional state spaces. Such complexities render the storage of comprehensive tables unfeasible [17].

Networking management serves as a compelling example of such scenarios, where modern Tofino switches can process INT packets at a nanosecond timescale.

To address these limitations, Mnih et al. leveraged the Q-Learning algorithm by integrating it with a Deep Neural Network (DNN) to approximate the optimal Q-value, a methodology known as Deep Q-Network (DQN). In their seminal work [19], the authors showcased the effectiveness of this approach by training and evaluating the DQN on an Atari 2600 emulator. Impressively, the DQN-based agents achieved performance levels surpassing those of human players in 49 distinct games, relying solely on pixel inputs and game scores for guidance.

Of note, the authors maintained a consistent algorithm, DNN architecture, and hyperparameters across all games, eliminating the need for game-specific feature engineering. Thus, DQN not only outperformed agents employing linear function approximation but also demonstrated the capacity to attain or exceed human-competitive skills across diverse gaming environments. This pioneering work exemplified the synergy between RL and contemporary Deep Learning (DL) techniques, signifying a significant advancement in the state of artificial intelligence. It underscored the potential of RL when combined with modern DL methods, yielding remarkable outcomes [19, 17].

In line with this, we present an RL-based approach designed to dynamically fine-tune the iRED target delay to an optimal value during video streaming, named DE-SiRED. This process is facilitated by an agent built on the foundation of DQN. In the subsequent Subsection, we will delve into the constituent elements that constitute this innovative approach.

2.2.1 Deep Q-Network workflow

The DQN architecture, as proposed by Mnih et al. [19], consists of a Deep Convolutional Neural Network (CNN) designed to receive emulated game frames as input and subsequently generate predicted Q-values for each potential action within the given input state. To facilitate such predictions, Mnih et al. introduced two critical modifications to the conventional Q-Learning algorithm. These alterations were essential to mitigate instabilities inherent in using Deep Neural Networks (DNNs) for Q-value approximation [17].

The first modification entails the incorporation of a biologically inspired mechanism referred to as 'experience replay.' In this approach, the agent's experiences are stored as tuples containing the current state (S_t), the action taken (A_t), the reward received (R_{t+1}), and the subsequent state (S_{t+1}). Periodically, after reaching a predefined replay memory limit, a mini-batch of these experiences is uniformly sampled for training the DNN [19, 17].

This approach plays a pivotal role in mitigating the emergence of correlations

within the observed state space. By decoupling the dependence on successive experiences, it effectively reduces the variance in the parameters of the DNN. Fig. 5 illustrates the interaction between a DQN agent and an environment, taking into account the experience replay mechanism. Within this context, the agent selects actions following an ϵ-greedy rule.
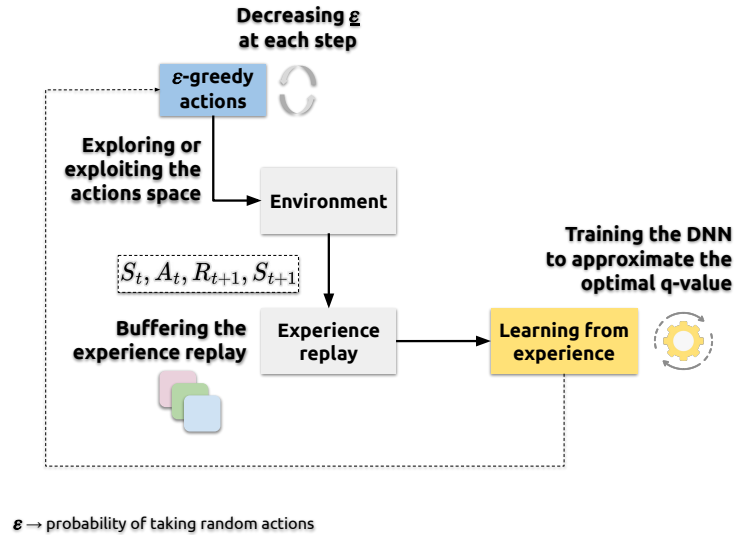

Figure 5: DQN high-level workflow [19, 17].

Specifically, when employing this rule, the agent chooses between two strategies: "exploitation" and "exploration". A "greedy action" involves selecting an action from the action space based on the maximum estimated Q-value. Conversely, a "non-greedy action" entails the random selection of an action. Exploitation, represented by the selection of a greedy action, aims to exploit the current knowledge to maximize immediate rewards. In contrast, exploration, represented by non-greedy actions, focuses on traversing the action space to maximize cumulative rewards in the long run [17].

In RL, achieving a balanced trade-off between exploration and exploitation is paramount. However, it's important to acknowledge that, at a single time step, it's not possible for an agent to simultaneously exploit and explore actions. To reconcile these opposing strategies, a solution is to allow the agent to primarily act greedily, favoring exploitation, while intermittently choosing an action from the action space at random, independent of the estimated Q-values. This random selection is determined by an exponentially decreasing probability parameter ϵ. Consequently, as the time steps progress, the probability of selecting an optimal action gradually converges to a value greater than $1-\epsilon$, approaching near certainty in favor of exploitation as the agent refines its strategy over time [17].

A second significant contribution introduced by Mnih et al. [19], relative to classical Q-Learning, pertains to the learning stage of the DQN. In this stage, a separate network, referred to as the 'target network,' is employed to estimate target values for the

Q-network, often referred to as the 'online network.' This modification enhances the algorithm's stability compared to using a single online network. The rationale behind this improvement lies in the fact that updating the parameters of the online network for the current state-action pair can inadvertently influence the Q-values of the next state, potentially leading to oscillations or even policy divergence.

To address this challenge, the online network's parameters are periodically cloned to the target network at intervals of every C time steps. Consequently, the target network's predictions serve as target values for the online network during the subsequent C time steps. This introduces a delay in updating the Q-values between the current and next states, effectively reducing the likelihood of policy oscillations or divergence [19].

Figure 6 illustrates the DQN learning workflow, incorporating the approach described above. A concise introduction to the functionality of DQN is presented in Subsection 3.2. Furthermore, for a detailed exposition on the implementation of DQN within the scope of this research, please refer to Subsection 4.4.
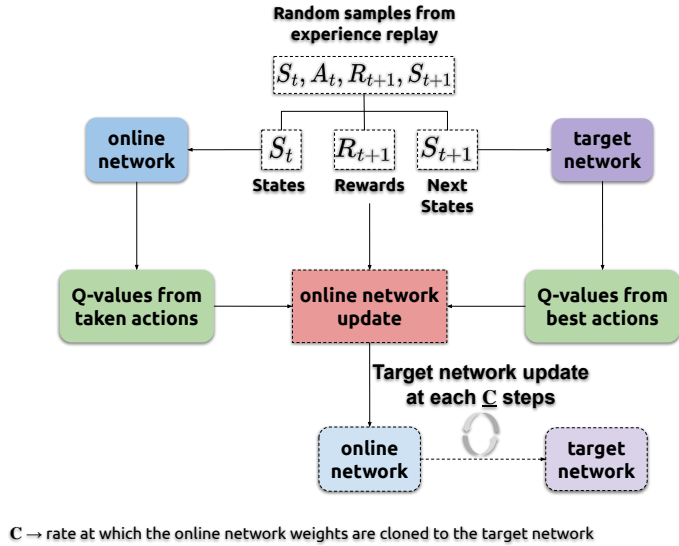

Figure 6: DQN learning stage workflow [19].

3 DESiRED - Dynamic, Enhanced, and Smart iRED.

DESiRED, herein referred to as an advanced iteration of iRED, which was initially introduced in the work of [6], constitutes a notable enhancement within the realm of intelligent network control systems. Specifically, it introduces a novel capability wherein the intelligent control plane harnesses the power of DRL to dynamically optimize and fine-tune the target delay parameters. In alignment with its predecessor, iRED, DESiRED remains faithful to the fundamental concept of disaggregated

AQM. In this paradigm, AQM operations are compartmentalized into discrete functional blocks within the architecture.

The concept of disaggregation emerges from the imperative to expedite packet discarding processes. In the pursuit of resource efficiency, we contend that the optimal location for packet discarding is the Ingress block. However, a noteworthy challenge arises as the vital metadata pertaining to queue delay (or queue depth), which constitutes the primary information utilized as input for the AQM algorithm to determine packet discarding decisions, is captured by the Traffic Manager and traditionally accessible within the Egress block. Within this context, DESiRED leverages a congestion notification mechanism, designed to incur minimal overhead, in order to relay the imperative to execute packet discarding actions to the Ingress block.
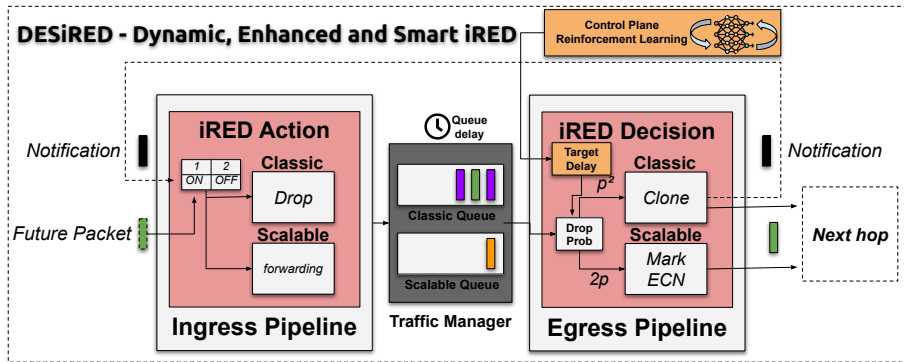

Figure 7: The Closed Control Loop overview with DESiRED. At the control plane side, DRL updates the target delay at the data plane.

As illustrated in Figure 7, the decision-making process within DESiRED takes place at the Egress block, while the corresponding actions are subsequently executed at the Ingress block. The following Subsections will elucidate the functioning of DESiRED, with a distinct focus on data plane and control plane operations.

3.1 Data plane operation (AQM)

To provide a more comprehensive understanding, we will commence our description of DESiRED's operation from the data plane perspective, focusing initially on the Egress block. Our exploration will initiate with the drop or marking decision process, a critical component housed within the decision module. At the Egress, iRED computes the Exponentially Weighted Mean Average (EWMA) of the queue delay (or queue depth[4]) for each individual packet, entirely within the data plane. The inherent absence of division and floating-point operations poses challenges in calculating average values within the data plane. To surmount this limitation, as applied in [20], we employ an approximation method following Eq. 3:

---

[4]The programmer can choose whether to use DESiRED's delay-based or depth-based approach.

$$S_t = \alpha \cdot Y_t + (1 - \alpha) \cdot S_{t-1} \qquad (3)$$

where S_t is the updated average queue delay, S_{t-1} is the previous average queue delay and Y_t is the current queue delay. The constant $\alpha \in [0, 1]$ determines how much the current value influences the average. We use $\alpha = 0.5$, such multiplication can be replaced by bit shifts operations. The output of the EWMA will represent the average queue delay over time. When the observed value, representing the average queue delay, falls within the range (minimum and maximum thresholds) configured by the DRL mechanism, DESiRED proceeds to calculate the drop probability in accordance with the RED approach. Simultaneously, it employs a coupling mechanism to generate various levels of congestion signals, which may entail either packet drops or packet marking (ECN bit).

Once the DESiRED decision module (Egress) has detected that a packet must be dropped, DESiRED must notify the action module (Ingress) to perform this action. The first challenge in the PDP context is to achieve communication between the Ingress and Egress blocks with minimum overhead. Obviously, DESiRED will not drop the packet that generated the discard decision, but a future packet [21]. Discarding future packets is one of the main features differentiating DESiRED from other state-of-the-art AQMs. For the congestion notification to reach the Ingress block, DESiRED creates a congestion notification packet (clone packet) and sends it through an internal recirculation port to reach the Ingress block.

The action module, situated in the Ingress block, maintains the congestion state table on a per-port/queue basis and activates the drop flag (ON) for the corresponding port/queue. The current packet is forwarded to the next hop without introducing any additional delay. Subsequently, future packets intended for the same output port/queue, where the drop flag is set to ON, will be dropped, and the drop flag will be reset to OFF. This mechanism, facilitated by DESiRED, ensures that the Ingress pipeline can proactively mitigate imminent queue congestion.

3.2 Control plane operation (DRL)

As mentioned earlier, DESiRED tackles the issue of fixed target delay through the implementation of an intelligent control plane, denoted by the orange box in Figure 7. This intelligent control mechanism is responsible for updating the register that maintains the dynamic target delay threshold, as determined by the DRL decision process. Now, let us provide a comprehensive account of the operational intricacies of the intelligent control plane, elucidating the inputs and outputs in detail.

The control plane operates by receiving data from two pivotal sources: the network state and the application state. In this particular implementation, fine-grained INT measurements constitute the input layer for the Deep Q-Network from the network state. The DQN's output layer is responsible for generating the agent's actions. Concurrently, the application state encompasses DASH metrics, including parameters such as FPS and the Local Buffer Occupancy (LBO) of the video player, which play a crucial role in computing the agent's reward. Fig. 8 illustrates this Control Loop.
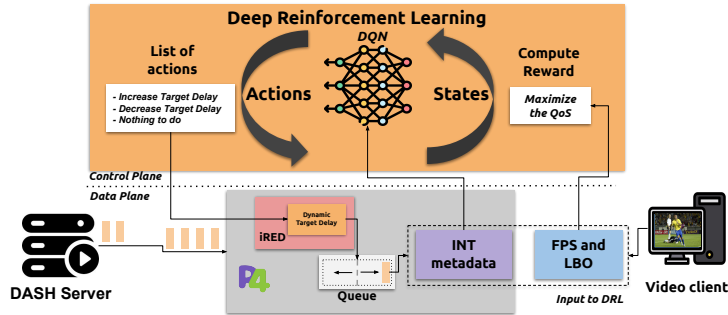
Figure 8: Operation of the Control Plane in DESiRED involves using fine-grained INT measurements as the input layer for the DQN. Additionally, DASH QoS measurements serve as the basis for calculating agent rewards.

INT measurements comprise observations that effectively depict the network's state with remarkable granularity, affording an unprecedented perspective on the extent of congestion. These measurements are acquired within the programmable data plane and subsequently routed to the intelligent control plane. Within the control plane, they are aggregated into compact dataframes, which collectively form what we term the "observation space." In the context of this study, the term observation space refers to the temporal window within which the intelligent control plane conducts an integrated analysis of both the network's state and the application's behavior.

For each received observation space, the DQN incorporates INT measurements as an input layer. Following neural network processing (refining its internal weights), the DQN generates an action, which is manifested as an activation in one of the neurons within the output layer. In this study, the possible actions include: 1) increasing the target delay; 2) decreasing the target delay; and 3) maintaining the current state (i.e., taking no action).

Subsequently, the control plane retains a record of the executed action and enters a state of anticipation for the forthcoming observation space. Upon the arrival of data from the subsequent observation space, the DRL mechanism evaluates whether the undertaken action has led to the optimization of DASH QoS, particularly with regard to enhancements in FPS and LBO metrics. In the event of a positive outcome, the agent is rewarded, whereas in cases of QoS deterioration, the agent incurs a penalty.

Leveraging insights from the dynamic network traffic patterns, DESiRED demonstrates a remarkable capability to adapt with precision to prevailing congestion conditions. This adaptability facilitates a continuous enhancement in the quality of video services offered.

It is imperative to elucidate that DESiRED is inherently application-agnostic, signifying its capacity to accommodate diverse reward policies tailored to evaluate a wide array of service metrics. This flexibility extends to metrics such as the response time of a web server or even the frame rate in video playback, underscoring its versatility across various service domains.

4 Evaluation

In this section, we provide a comprehensive overview of all the components utilized for the thorough evaluation of our proposal. This encompasses a detailed exposition of the research methodology, an in-depth portrayal of the experimental environment and its configuration, the load pattern employed, the DRL mechanism implemented, the metrics and measurements used for comprehensive analysis.

4.1 Research methodology

Our methodology is rooted in experimental research aimed at evaluating the effectiveness of the DRL mechanism within DESiRED. Specifically, our objective is to ascertain whether this mechanism can optimize the QoS for MPEG-DASH services by dynamically adapting the target delay under conditions characterized by both stationary and non-stationary loads within a Content Delivery Network (CDN) environment.

In this experiment, our aim is to conduct a comprehensive evaluation of DESiRED in comparison to iRED, where iRED employs fixed target delay settings of 5ms, 20ms, 50ms, and 100ms. We evaluate these approaches under both stationary (low and high) and non-stationary (sinusoidal) load conditions. To mitigate potential biases, each round of the investigation, spanning one hour, was repeated ten times for each approach, resulting in a cumulative duration of over fifty hours across independent runs. Furthermore, to gauge DESiRED's robustness, we aggregated the DRL agents derived from all preceding executions by employing an ensemble approach. This involved combining the model parameters through an exponentially decaying running average, as described by Eq. 4 [22, 23]:

$$\hat{\theta}^{(t)} = \alpha\hat{\theta}^{(t-1)} + (1 - \alpha)\theta^{(t)} \tag{4}$$

where θ represents a parameter from the Q-network; t the gradient descent iterations; $\hat{\theta}^{(t)}$ the average from such parameters ($\frac{1}{t}\Sigma_i\theta^{(i)}$); and α the exponential decaying factor (defined as 2.0).

We evaluate the application's performance from the client-side perspective, focusing on three key metrics: FPS LBO, and Rebuffering Rate (Starvation) as measured within the video player. Higher values for FPS and LBO correspond to improved QoS, while for Rebuffering Rate, a lower value signifies enhanced QoS.

In addition to evaluating application quality metrics, we also scrutinize the performance metrics of the DRL agent, encompassing Loss function and Rewards.

4.2 Environment description

The experiment was constructed within a realistic testbed, adopting an Infrastructure as Code (IaC) approach, and implemented using Vagrant, Virtualbox (version 6.1.28), and Ansible (version 2.10.8). In this setup, each infrastructure component is represented by an isolated virtual machine, interlinked through a P4 programmable data plane network, as visually depicted in Figure 9.
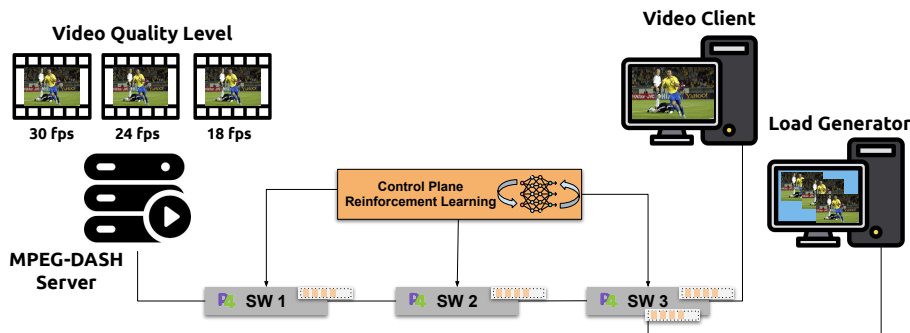
Figure 9: Setup Evaluation.

Each switch in the experiment was equipped with both the iRED and DESiRED approaches. On the control plane side, the DRL engine was implemented, comprising approximately 750 lines of code and utilizing Tensorflow as its backend framework. The CDN was deployed to facilitate an MPEG-DASH service, featuring live streaming of a soccer game and a playlist housing the ten most frequently accessed YouTube videos. Load management was executed using WAVE [24] [5], a versatile load generator that orchestrates instances of an application over time.

This infrastructure was hosted on a bare-metal server, namely the Dell EMC PowerEdge R720, equipped with 2 Intel Xeon processors (E5-2630 v2, 2.60GHz) boasting 6 cores per socket (amounting to 24 virtual CPUs), 48GB of RAM, a 2TB HDD, and running the Ubuntu 20.04.6 LTS operating system. All pertinent artifacts and resources can be accessed within the repository available at our GitHub[6].

The MPEG-DASH Server serves video content using the DASH standard to both the Video Client and the Load Generator. It offers various configurations, as detailed in Table 1, with each configuration having a chunk segment size of 4 seconds. The Video Client dynamically selects and transitions between these configurations based on network traffic conditions and the adaptation logic embedded within the video player.

The infrastructure is equipped with Apache version 2 as the web server, FFmpeg (version 2.8.17) for video encoding, and MP4box (version 0.5.2) for creating the MPEG-DASH manifest files, ensuring seamless video streaming.

Type	Resolution	FPS	Group of Pictures	Kbps	Buffer	Codec
vídeo	426x240	18	72	280	140	h264
vídeo	854x480	24	96	980	490	h264
vídeo	1280x720	30	120	2080	1040	h264
áudio	-	-	-	128	-	AAC
áudio	-	-	-	64	-	AAC

Table 1: Video parameters used in an MPEG-DASH Server.

The Video Client utilizes DASH.js, a contemporary DASH reference player equipped

---

[5]https://github.com/ifpb/wave

[6]https://github.com/dcomp-leris/DESiRED.

15

with an Adaptive Bitrate Streaming (ABR) algorithm. It employs this ABR algorithm to consume the video stream of the soccer game, with the TCP New Reno congestion control algorithm managing network congestion.

The Load Generator is responsible for introducing network noise, operating the WAVE framework with a variety of loads, including both stationary and non-stationary scenarios. It dynamically adjusts the number of video player instances over time to simulate changing network conditions. Further elaboration on this aspect can be found in Subsection 4.3.

All the switches utilized in this experiment were implemented within the BMv2 software switch environment, incorporating the respective P4 code for both iRED (fixed target delay) and DESiRED (dynamic target delay with DRL) approaches. Across all approaches, telemetry instructions were meticulously programmed to append telemetry metadata to all probe packets. Notably, this experiment follows the out-of-band (ONT) approach, wherein dedicated ONT probes are dispatched from the DASH server to the Video Client. Consequently, no modifications are made to data packets to accommodate telemetry metadata. The specifics of the telemetry metadata, consisting of 32 bytes, gathered at each node within this experiment, are elaborated upon in Table 2.

Name	bits	Description
Switch ID	31	the switch identification number
Ingress port	9	the port number that the packet entered in the switch
Egress port	9	the port number that the packet left of the switch
Egress spec	9	the port number (Ingress) in which the packet will leave the switch
Ingress Global Timestamp	48	the timestamp, in μs, of when the packet entered in the ingress
Egress Global Timestamp	48	the timestamp, in μs, of when the packet started processing in the egress
Enq Timestamp	32	the timestamp, in μs, of when the packet was enqueue
Enq Qdepth	19	the queue depth when the packet was queued
Deq Timedelta	32	the time, in μs, that the packet remained in the queue
Deq Qdepth	19	the queue depth when the packet was dequeued

Table 2: INT medatada.

4.3 Load Pattern

The Load Generator, powered by WAVE, orchestrates the instances of video clients over time based on input parameters described by a mathematical function that defines the load pattern. In its current iteration, WAVE supports constant, sinusoidal, and flashcrowd load patterns. It initiates and concludes video player processes, generating network load through genuine video requests (real traffic) that flow from the video player to the MPEG-DASH Server.

In this study, our aim is to evaluate DESiRED under various load conditions, aiming to simulate diverse network state scenarios. To achieve this, we employ two distinct categories of load patterns: stationary and non-stationary. For stationary loads, which remain constant throughout the experiment, we classify them into two types: low and high. In this context, a low load is characterized by the presence of ten video client instances operating concurrently throughout the duration of the experiment, as depicted in Figure 10(a). Conversely, a high load is characterized by the simultaneous opera-

tion of forty video player instances, representing a high-intensity load, as illustrated in Figure 10(b).

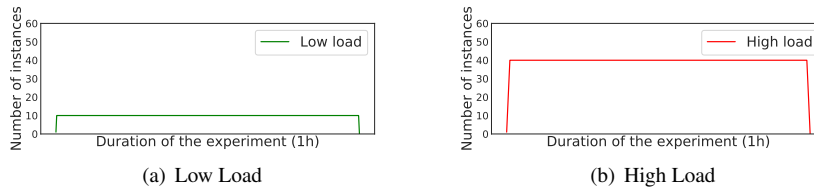(a) Low Load   (b) High Load

Figure 10: Stationary Loads.

Under low load conditions, it is anticipated that the target delay will be attained relatively infrequently, given the shorter queuing delays that typically prevail. In this scenario, both AQM strategies, whether employing a fixed or dynamic target delay, are likely to yield comparable results in terms of QoS.

However, when the network experiences predominantly high load, the surge in traffic volume can lead to an increase in queue delay, thereby prompting AQM strategies to respond in accordance with the specified target delay, whether fixed or dynamic. In such instances, the dynamic adaptability of DESiRED's target delay is expected to confer advantages in terms of QoS compared to the rigid, fixed target delay approach employed by iRED. This dynamicity enables DESiRED to better accommodate and optimize QoS in the face of fluctuating and demanding network conditions.

It is indeed unrealistic to assume that network loads will always remain stationary or static. Consequently, in the second phase of our evaluation, we undertook a more comprehensive evaluation under a realistic load scenario, one that mirrors the dynamic nature of real-world network environments. Our objective was to evaluate non-stationary load patterns, encompassing both peak (high load) and off-peak (low load) periods within a single experiment. To achieve this, we employed a sinusoidal periodic load pattern characterized by the sinusoidal function detailed in Equation 5, where A represents the amplitude, F denotes the frequency, and λ signifies the phase in radians. The specific input parameters utilized for this evaluation were: $A = 15$, $F = 1$, and $\lambda = 25$, culminating in the load pattern illustrated in Figure 11. This approach captures the fluctuations in network load more realistically, offering a dynamic and challenging environment for our evaluation.

$$f(y) = A \sin(F + \lambda) \tag{5}$$

4.4 Deep Reinforcement Learning mechanism

To accomplish the objectives outlined in this paper, we tailored the DQN architecture and agent-environment workflow to align with the distinctive characteristics of the DESiRED environment, as elucidated in Subsection 4.2. In doing so, we designed the DQN using a Multi-layer Perceptron (MLP) architecture, which is well-suited for handling the tabular nature of network telemetry metadata. The MLP network adopted
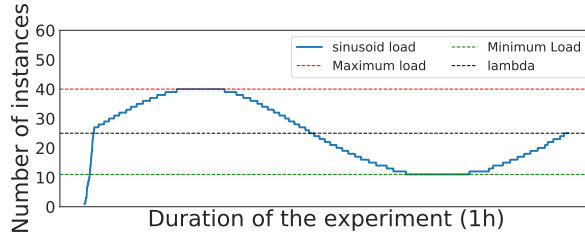
Figure 11: Non-stationary (Sinusoid) Load.

in our approach consists of an input layer featuring units corresponding to each INT feature, two hidden layers each comprising 24 neurons, and an output layer containing units for each possible action that the agent can undertake, as depicted in Figure 12. Importantly, both the online and target networks share this identical architecture. Table 3 provides a detailed breakdown of the hyperparameters utilized for training DESiRED.
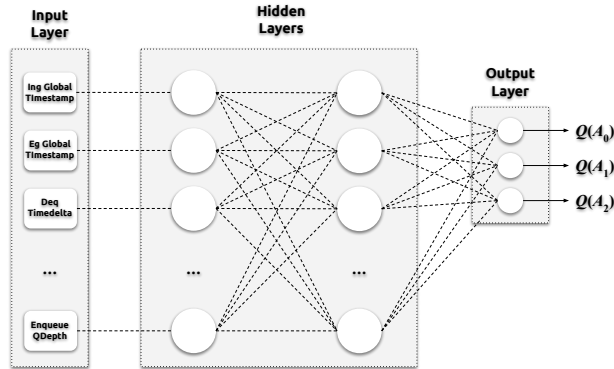

Figure 12: DESiRED DQN architecture. The input layer is a network of fine-grained measurements, provided by INT. Hidden Layers make up the DQN. The actions are defined in the Output Layer.

To facilitate the desired agent-environment interaction, we formulated the agent's behavior as an MDP with the video chunk size serving as the discrete time steps. In this framework, DESiRED operates within the environment, dynamically adjusting the target delay in all switches at 4-second intervals, synchronized with the video chunk size. A comprehensive discussion regarding the strategy of simultaneous actuation in all switches versus individual actuation in each switch is presented in Section 6. The agent's action space is delineated in Table 4, where it is evident that the action to increase the target delay brings about a modification that is proportionally twice as substantial as the decrease action. This choice was made to prompt DESiRED to respond promptly to transient congestion while retaining the flexibility to decrease the

18

Hyperparameter	Value	Description
Q-network input layer dimension	19	A scalar defining the state input shape.
Q-network hidden layers	2	A scalar defining the Q-network depth.
Q-network hidden units	24	A scalar defining the Q-network non-linear computing units.
Q-network output layer dimension	3	A scalar defining the Q-network predictions output shape.
Hidden units activation function	ReLU[1]	The non-linear activation function computed by hidden units.
Output units activation function	Linear	The activation function computed by the output layer.
Optimization function	SGD[2]	The function used to adjust the Q-network weights in order to minimize the predictions error in relation to the expected output.
SGD momentum	0.9	A scalar defining the momentum included in the optimization equation to accelerate the gradient descent.
Learning rate	1e-3	A scalar determining the pace at which the weights are updated.
Loss function	MSE[3]	The function used to compute the Bellman equation error.
Gamma	0.99	A scalar determining the discount factor in the Q-Learning update.
Tau	1e4	A scalar value determining how many updates the online network should perform before updating the target network (it corresponds to the parameter C depicted in the Fig. 6).
Experience replay capacity	1e6	A scalar defining the size of the list in which the agent's experience will be stored.
Experience replay minimum memory	100	A scalar defining the minimum experiences that should be stored before updating the online network.
Mini batch size	32	A scalar defining the number of experience samples over which the Q-network will be updated.
Starting epsilon	1.0	A scalar defining the initial probability to take random actions in the ϵ-greedy exploration.
Ending epsilon	0.01	A scalar defining the final probability to take random actions in the ϵ-greedy exploration.
Epsilon decay steps	250	A scalar determining how many steps the probability to take random actions in the ϵ-greedy exploration should decrease linearly before the exponential decay.
Epsilon exponential decay	0.99	A scalar determining exponential decay of the probability to take random actions in the ϵ-greedy exploration.

[1] Rectified Linear Unit.

[2] Stochastic Gradient Descent with Nesterov Momentum.

[3] Mean Squared Error.

Table 3: DQN hyperparameters.

target delay when necessary, mirroring the rationale discussed in [25].

It's important to highlight that the calculation of rewards does not occur immediately after an action is taken in the current state. This delay in reward calculation is attributed to the fact that the effects of the agent's action do not manifest instantly, primarily due to the inherent control mechanisms incorporated within TCP and ABR systems, as detailed in [26]. Consequently, the computation of rewards is deferred until the subsequent state's observation. In this context, the agent relies on network status data derived from INT measurements to form its states, selects actions, and is subsequently rewarded based on its ability to optimize the video's QoS, which is characterized by metrics such as FPS and LBO.

Indeed, the intrinsic correlation between metrics such as LBO and FPS presents a challenge when devising a reward policy. As the LBO increases, there is a tendency for the FPS to also increase. However, this relationship is not always straightforward due to the complex dynamics of network congestion and video streaming.

To calculate a reward (R_{t+1}) for a specific action (A_t), we adopt a strategy that first evaluates whether the LBO in the next state (LBO_{t+1}) improves compared to the LBO observed when the action was executed (LBO_t). Subsequently, a reward score is assigned based on the effects of this action on both the next state's LBO and FPS

Action Number	Value	Description
0	+ 2ms	increase target delay in all switches until 70ms (upper limit)
1	- 1ms	decrease target delay in all switches until 20ms (lower limit)
2	-	do nothing

Table 4: DESiRED actions space.

(FPS_{t+1}). Consequently, the agent receives maximum reward whenever the action taken leads to the maximization of LBO_{t+1}, and is penalized in an inversely proportional manner if the video experiences stalls. The algorithmic logic for calculating rewards is detailed in Algorithm 1. This approach ensures that the agent's reward is contingent on its capacity to optimize both LBO and FPS, balancing the trade-offs inherent to video streaming in dynamic network conditions.

These actions were executed according to the ϵ-greedy strategy as elucidated in Subsection 2.2.1. To implement this strategy, we established initial and final probabilities for taking random actions, specified the number of decaying steps, and defined an exponential decay factor (as outlined in Table 3). In this scheme, ϵ commences its linear decrease over a span of 250 time steps to facilitate exploration. Subsequently, the probability of selecting random actions is exponentially reduced, gradually transitioning to a minimal value to emphasize exploitation over exploration. This strategy allows the agent to strike a balance between exploring new actions and exploiting its existing knowledge as it interacts with the environment.

Taking into consideration the agent's action frequency of once every 4 seconds and the requirement for 250 iterations to initiate the exponential decay of ϵ, the exploration phase is expected to persist for approximately 17 minutes (equivalent to 1000 seconds). In tandem, the experience replay memory buffer necessitates a minimum of 100 samples to facilitate the online network parameter updates (as indicated in Table 3. Since experiences resulting from the agent-environment interaction are stored every 8 seconds, it would take approximately 13 minutes (or 800 seconds) for this condition to be met. Consequently, the online network undergoes an update each time a new experience is stored, as illustrated in Figures 5 and 6.

In the case of the non-stationary load, it follows a trajectory of 15 minutes to reach its peak, maintains a plateau for an additional 15 minutes, and subsequently begins to decline. During this period, the agent explores the action space during the ascending phase of the sinusoidal curve and exploits these actions during the plateau and descending phases. Consequently, when the exploitation stage commences, the agent should have already gleaned insights from past experiences, encompassing both low and high load scenarios. This enables the agent to adapt and respond effectively to the fluctuating network conditions.

4.5 Metrics and Measurements

On the video client side, we evaluate the QoS by monitoring key metrics, including:

- FPS (Frames Per Second): This metric quantifies the number of frames displayed per second on the screen, reflecting the smoothness of the video playback.

Algorithm 1 DESiRED reward policy algorithm.

1: **function** $calculate\_reward(LBO_t, LBO_{t+1}, FPS_t, FPS_{t+1})$
2:     **if** $LBO_{t+1} > LBO_t$ **then**
3:         **if** $LBO_{t+1} > 30$ **then**
4:             $R_{t+1} \leftarrow 2$
5:         **else if** $LBO_{t+1} < 30$ **then**
6:             **if** $FPS_{t+1} == 30$ **then**
7:                 $R_{t+1} \leftarrow 1$
8:             **else if** $FPS_{t+1} == 24$ **then**
9:                 $R_{t+1} \leftarrow 0.5$
10:            **else**
11:                $R_{t+1} \leftarrow 0.1$
12:            **end if**
13:        **end if**
14:    **end if**
15:
16:    **if** $LBO_{t+1} < LBO_t$ **then**
17:        **if** $LBO_{t+1} > 30$ **then**
18:            $R_{t+1} \leftarrow 2$
19:        **else if** $LBO_{t+1} < 30$ **then**
20:            **if** $FPS_{t+1} == 30$ **then**
21:                $R_{t+1} \leftarrow 1$
22:            **else if** $FPS_{t+1} == 24$ **then**
23:                $R_{t+1} \leftarrow 0.5$
24:            **else**
25:                $R_{t+1} \leftarrow -2$
26:            **end if**
27:        **end if**
28:    **end if**
29: **end function**

- LBO (Local Buffer Occupancy): LBO measures the remaining time, in seconds, for frames stored in the player's local buffer. It provides insights into the buffer's capacity to absorb network fluctuations and maintain continuous playback.

From these primary metrics, we derive additional insights, including:

- Resolution Distribution: We analyze the percentage of video content played at different resolutions (Maximum, Medium, and Minimum) to assess the adaptive streaming capabilities.

- Rebuffering Rate: This metric represents the percentage of time during which the video experiences stalls or freezes on the screen, indicating interruptions in playback.

To facilitate these measurements, we configure the DASH.js player to log these metrics on a per-second basis. Within the DRL mechanism, we focus on evaluating the performance metrics of the DQN:

- Loss: This metric is calculated as the Mean Squared Error (MSE) between the predicted q-values for the current and next states. It reflects the convergence and accuracy of the DQN's predictions.

- Reward: Reward represents the cumulative rewards and penalties acquired throughout the experiment. It offers insights into the agent's performance in maximizing QoS.

Additionally, we capture the action history for each experiment, documenting the agent's selected actions at each observation space (every 4 seconds). These metrics provide a comprehensive view of the agent's learning and adaptation throughout the experiment.

5 Results

In this section, we will present the outcomes of our experiments, where we evaluate how DESiRED enhances the QoS of MPEG-DASH. We offer an in-depth analysis from the client-side perspective, showcasing the results and delving into instances where video QoS has benefited from the dynamic adjustments facilitated by DESiRED. Furthermore, we scrutinize the performance of the DRL model, presenting evidence that the agent has successfully learned the designated policy and has been able to identify an optimal target delay value that maximizes QoS across the range of experiments conducted.

5.1 Stationary Loads

The motivation behind evaluating performance under stationary loads stemmed from the necessity to ascertain whether the DRL agent would exhibit distinct learning behaviors during moments of low load (ample resources) and high load (congested resources) across separate executions.

When the network load predominantly remains low, as illustrated in Figure 10(a), network resources are readily available. In such scenarios, there is minimal contention for the use of the queue, resulting in limited or no intervention from auxiliary congestion control mechanisms like AQM. This phenomenon can be observed from the perspective of the video client, as depicted in Figure 13.
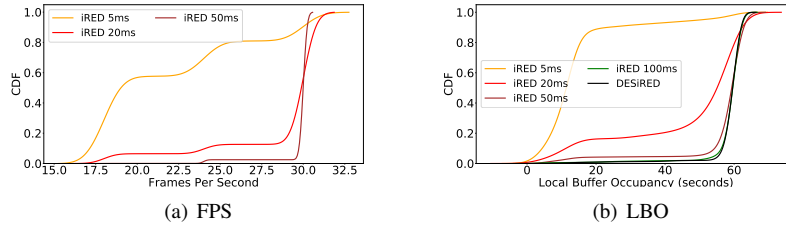

(a) FPS   (b) LBO

Figure 13: Low Load - Characterized by only ten video player instances managed by WAVE (Load Generator).

Figures 13(a) and 13(b) illustrate the Cumulative Distribution Function (CDF) of FPS and LBO under low load conditions. In Figure 13(a), we observe some variation in FPS for iRED with a fixed target delay of 5ms and 20ms/50ms. Conversely, in the cases of iRED with a fixed target delay of 100ms and DESiRED, the video client consistently played the video at 30 FPS throughout all experiments.

Concerning LBO, as depicted in Figure 13(b), the results exhibit similar behavior across approaches, with the local buffer maintaining a near-full state for most of the evaluations, approximately 60 seconds. The only exception is the iRED with a 5ms fixed target delay. In this specific scenario, the use of such a small threshold value appears to have triggered a higher frequency of AQM actions. This, in turn, might have led to more frequent drops within a time interval of less than one Round-Trip Time (RTT), as discussed in [27]. Paradoxically, this increased AQM activity, rather than alleviating congestion, may have exacerbated the situation, demonstrating the potential for unintended side effects when setting overly aggressive congestion control thresholds.

Conversely, when the network experiences predominantly high load conditions, as illustrated in Figure 10(b), the dynamics shift significantly. In such scenarios, all approaches employing fixed target delay mechanisms encounter challenges in maintaining acceptable MPEG-DASH QoS. DESiRED, on the other hand, manages to distinguish itself from the fixed target delay approaches, as evident in Figure 14.

To gain a deeper understanding of these results, it's important to clarify some aspects of the Adaptive Bitrate (ABR) adaptation logic employed by the DASH.js player, as described in [26]. The adaptation logic used in DASH.js, known as DYNAMIC, employs two different algorithms at different stages of video playback. During instances when buffer levels (LBO) are low, such as startup and seek events, a straightforward THROUGHPUT algorithm (based on throughput) is utilized. Conversely, when buffer levels are high, the player switches to the BOLA algorithm [28]. This dynamic adaptation approach aims to optimize video streaming under varying network conditions,
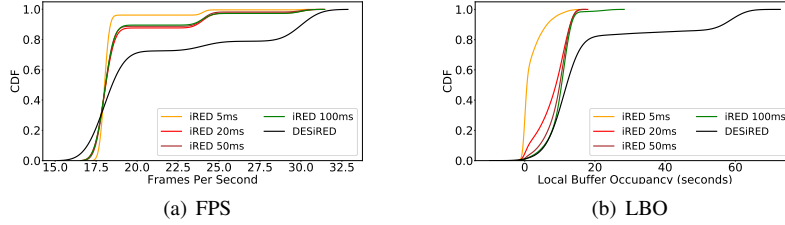
(a) FPS   (b) LBO

Figure 14: High Load. Characterized by forty video player instances managed by WAVE (Load Generator).

aligning the bitrate selection algorithm with the network's congestion state.

DYNAMIC starts with THROUGHPUT until the buffer level reaches 10s or more. From this point on, DYNAMIC switches to BOLA which chooses a bitrate at least as high as the bitrate chosen by THROUGHPUT. DYNAMIC switches back to THROUGH-PUT when the buffer level falls below 10s and BOLA chooses a bitrate lower than THROUGHPUT [26].

Indeed, from the perspective of the video player's adaptation logic, the LBO metric proves to be far more sensitive to variations in network buffer levels compared to FPS. It's important to note that changes in bitrate and FPS should only occur when the LBO drops below 10 seconds. Consequently, it is logical to aim for maintaining an LBO greater than 10 seconds for the majority of the time, as this instructs the ABR algorithm to select the highest-quality video levels.

Figure 14(b), which pertains to LBO, contributes significantly to understanding why DESiRED achieves superior FPS levels, as indicated in Figure 14(a). In this context, it is plausible to surmise that fine-tuning the target delay has provided an advantage in terms of preserving a sufficient LBO during periods of severe network conges-tion. This, in turn, aids the ABR algorithm in making optimal bitrate and quality level selections, ultimately leading to improved video QoS.

5.2 Non-stationary Load

Recognizing the dynamic nature of network traffic, we embarked on an evaluation under non-stationary load conditions. To achieve this, we leveraged the WAVE frame-work, which effectively managed the execution of video client instances over time, adhering to a mathematical model of sinusoidal periodic load as detailed in Subsection 4.3.

The choice of a sinusoidal periodic load model holds significance because it encap-sulates moments of congestion and resource relief in the network, particularly within router buffers, within a single execution. This approach allows us to evaluate our agent's performance in situations of both high congestion, where rapid adaptation is crucial, and congestion-free states where shared resources are not overwhelmed. In essence, our expectation is that the agent will learn distinct patterns that differentiate between these varying states.

24

This evaluation under non-stationary load conditions provides valuable insights into how the agent responds to fluctuations in network congestion, thereby contributing to a more comprehensive understanding of its adaptability and effectiveness.
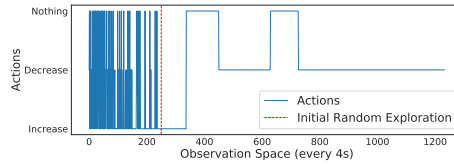

Figure 15: Actions performed by the agent in the environment. After the initial random exploration, the agent finds the best target delay value to maximize the QoS of MPEG-DASH.

The initial result we would like to present pertains to the actions taken by the agent (DESiRED) within the network environment. Figure 15 provides an overview of the agent's actions throughout the experiment. Notably, there is an initial phase of random exploration (indicated by the vertical dashed red line) extending up to the first 250 observations. During this exploratory phase, the agent gathers data about the network state, which is used to populate the experience replay buffer (as outlined in Subsection 2.2.1).

Subsequent to this initial exploration phase, the agent commences taking actions based on its learned knowledge, drawing from the experiences stored in the experience replay buffer. It's important to highlight that this buffer is continually updated, enabling the agent to learn from new states. Consequently, the agent can adapt to previously unseen states, a capability that proves particularly valuable in scenarios with non-stationary loads.

This analysis of the agent's actions provides insights into its learning process and the transition from exploration to exploitation as it becomes more knowledgeable about the environment.

Analyzing the agent's actions, it becomes apparent that during the initial phase of the experiment, characterized by an increase in network load, the agent frequently opted to increase the value of the target delay. Subsequently, as the load stabilized, the agent chose to take no action, potentially reducing the overhead of control plane operations in the data plane. Towards the end of the experiment, as the network load decreased, the agent shifted its strategy towards reducing the target delay.

Having observed how these actions mirror the agent's interactions with the environment, we can now delve deeper into the model's performance. Figure 16 provides an overview of the model's behavior, illustrated by the curves representing key performance metrics such as Loss and Reward.

Figure 16(a) illustrates the trajectory of Loss throughout the experiment. A decline in Loss signifies a lower MSE in predicting q-values. In essence, a low Loss value suggests that the model is effectively learning the policy by selecting actions that maximize rewards (QoS). During the initial phase of filling the experience replay buffer, Loss tends to be higher as actions are taken without the benefit of learning, effectively representing random actions. However, as the experience replay buffer becomes pop-
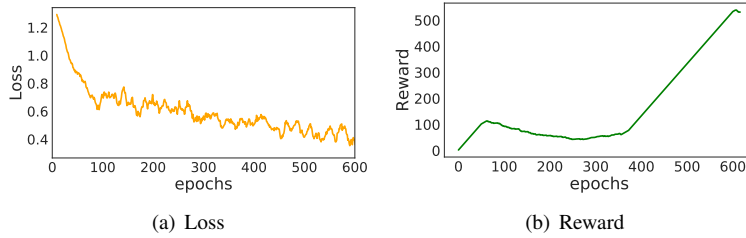
(a) Loss

(b) Reward

Figure 16: Model performance results - Decreasing Loss and increasing Reward indicate model convergence.

ulated and the Q-network is updated based on these experiences, the agent begins to make more informed and assertive decisions. This shift towards lower Loss values reflects the agent's ability to learn and improve its policy.

Turning our attention to Rewards (Figure 16(b)), we observe that the model incurs some penalties during the initial phase of the experiment. This corresponds to the period when the agent transitions from an initial stationary state with no charges to reaching the peak of the sinusoidal load curve, marked by the presence of 40 instances of the video player simultaneously. Subsequently, as the agent refines its decision-making, it starts receiving rewards consistently. These rewards indicate that the agent effectively maximizes the QoS of MPEG-DASH, further underscoring the model's learning and adaptive capabilities.

The insights gleaned from the agent's performance analysis are supported by the LBO and FPS metrics observed by the video client in response to DESiRED's actions, as outlined in Table 5 and depicted in Figure 17. At this conjuncture, we aim to provide an interpretation of the results from the video client's perspective, highlighting how DESiRED outperformed other approaches considered in this study.

An essential piece of data when evaluating the QoS of a video service is the resolution displayed on the screen by the video player. In this context, video consumers were offered three distinct quality levels:

1. Minimum Resolution: 426x240 pixels at 18 FPS.

2. Medium Resolution: 854x480 pixels at 24 FPS.

3. Maximum Resolution: 1280x720 pixels at 30 FPS.

AQM	Min. Resolution	Med. Resolution	Max. Resolution
iRED 5ms	91.83%	5.49%	1.36%
iRED 20ms	68.77%	12.96%	17.69%
iRED 50ms	46.74%	11.73%	41.27%
iRED 100ms	43.81%	9.91%	46.18%
DESiRED	31.71%	10.15%	58.07%

Table 5: Execution percentage at each video quality level.

Even under challenging conditions, Table 5 clearly demonstrates that DESiRED exhibits the highest percentage of video playback at the maximum resolution (58.07%) and the lowest rate of playback at the minimum resolution (31.71%). This finding aligns with the data presented in Figure 17.

The discussion initiated in Subsection 5.2 remains pertinent in this context as well. To reiterate, during periods of intense competition for shared resources, probabilistic drops facilitated by a target delay that adjusts in response to network load fluctuations have proven instrumental in maximizing the QoS of the video service. Once again, DESiRED effectively maintains a higher level of LBO filling, as depicted in Figure 17(b), ultimately contributing to superior FPS performance, as evidenced in Figure 17(a).
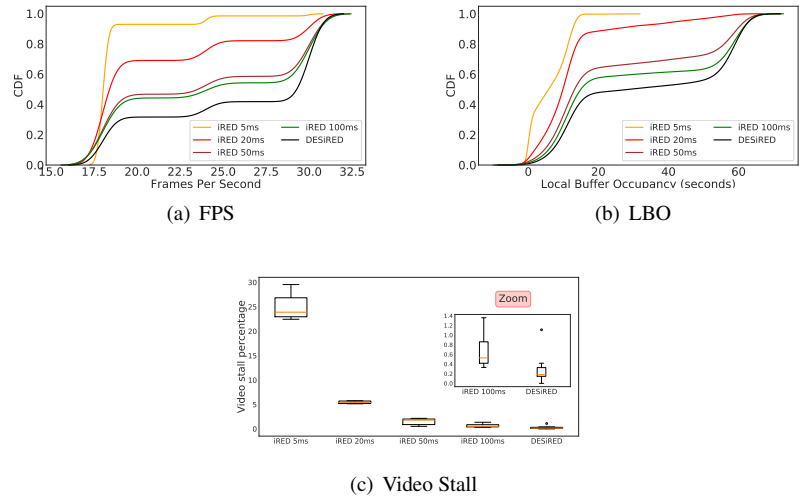

(a) FPS

(b) LBO

(c) Video Stall

Figure 17: QoS measurements of the MPEG-DASH video service. DESiRED improves FPS and LBO while minimizing video stall.

Figure 17(c) presents a boxplot representing the percentage of video stalls, which signifies moments when the video remains frozen without any frames being displayed. A cursory glance at this figure might lead to the incorrect assumption that a longer delay at a fixed target would yield better results. However, it's important to note that DESiRED imposes an upper limit of 70ms, which is lower than the value employed by iRED100ms, thereby dispelling this theory. In this context, we believe that DESiRED's fine-tuned approach enables it to determine the optimal target delay value for each network state during the sinusoidal load.

6 Lessons Learned

In this section, we will provide insights and lessons learned from our research on applying RL to computer network problems. These insights may be valuable to the scientific

community interested in using RL for similar applications.

1) The network has an intrinsic dynamism in its behavior: In the realm of RL, the challenges posed by computer networks present an intriguing and multifaceted problem. In essence, an RL problem can be likened to a strategic game where an agent interacts with an environment, making decisions and receiving rewards, all within the framework of a MDP. In each of these interactions, often referred to as episodes, the agent engages in a continuous process of trial and error, striving to acquire a policy that maximizes its cumulative rewards.

However, the application of RL models to computer network-related predicaments introduces a unique set of challenges. Contemporary networks, characterized by their dynamic nature and intricate traffic dynamics, necessitate a novel approach to the integration of RL. One of the central predicaments lies in adapting an RL agent to an environment that is in perpetual flux, a paradigm well-embodied by the ever-changing states of queues within network routers.

Of notable significance is the realization that RL agents draw their learning from the experiences accumulated through their interactions with the environment. This very dependence on real-time experiences, further compounded by the interdependence between video player metrics and network conditions—themselves subject to the agent's actions—renders the use of static datasets for agent training impractical. In situations where a physical network infrastructure is not readily available, a promising alternative entails the utilization of a model capable of simulating authentic network behaviors, such as a Generative Adversarial Network (GAN) as proposed by Navidan et al. [29].

An additional layer of complexity is introduced through the modulation of network load patterns, a deliberate endeavor aimed at inducing the RL agent to adapt dynamically to both peak (high load) and trough (low load) network scenarios. In this pursuit, an array of network load settings was meticulously explored, encompassing flashcrowd and sinusoid patterns. Notably, the most compelling outcomes were achieved when employing sinusoidal patterns, characterized by single instances of peak and trough conditions within the duration of video streaming.

Furthermore, the intricate calibration of parameters pertaining to the reward policy emerged as an arena of paramount importance within the implementation of the DE-SiRED system. It was during this phase that some of the most noteworthy findings and developments transpired. Remarkably, the strategic revision of the reward policy wielded a disproportionate influence over the observed outcomes, eclipsing the impact of various other elements intrinsic to the proposed approach. As such, it underscores the pivotal significance of meticulous and judicious reward policy design tailored to the specific problem domain.

In summation, the application of RL methodologies to the domain of computer networks is an enthralling endeavor replete with challenges and opportunities. It necessitates an astute orchestration of dynamic simulations, judicious load modulation, and the nuanced refinement of reward policies—a multifaceted tapestry of considerations aimed at navigating the intricate terrain of modern network optimization.

2) The core of solution design lies in the rewards policy: As previously mentioned, a significant portion of our modeling effort was dedicated to defining a rewards policy that aligns with our goal of maximizing the QoS in MPEG-DASH. Initially, we considered focusing solely on the FPS values during video playback. However, this approach

proved insufficient due to the dynamic nature of the video player's adaptation logic, which considers factors like throughput and buffer level. As the agent's actions influence the target delay on network devices, we anticipated that FPS values would only exhibit noticeable changes following alterations in the LBO, as LBO is more responsive to network variations. Consequently, we opted to construct our rewards policy, with a primary emphasis on evaluating LBO levels, and secondary consideration given to FPS.

3) Why actuate in all devices at the same time: As actions taken by our agent are intrinsically intertwined with the rewards policy, this study delves into several approaches, including the independent execution of actions on individual switches (each switch having its specific action) or the simultaneous execution of identical actions across all switches within the network. Initially, we contemplated that employing independent actions for each switch could be an appealing strategy. However, this approach did not align seamlessly with the scope of our problem.

Firstly, modifying the target delay for a single switch might not suffice to effectively assist the TCP congestion control algorithm, potentially yielding inconspicuous improvements in application-level QoS. Secondly, the adoption of such an approach would entail a proliferation of actions, scaling exponentially with the number of switches in the network (i.e., 2^n actions, with n representing the count of switches). This increase in action space complexity could substantially augment the neural network architecture's intricacy.

4) When we need to think in Transfer Learning (TL): Given the vast diversity of services, applications, topological configurations, and network loads encountered, it is imperative to acknowledge that an agent trained within a specific network environment cannot be expected to replicate its performance in other heterogeneous settings. In response to this challenge, TL has emerged as a promising approach, aiming to address several intricacies not typically encountered in the realm of RL. However, the application of TL within an RL framework is a non-trivial undertaking, necessitating numerous adaptations to enable the agent to effectively leverage knowledge acquired in a source domain for application in a target domain.

Amidst the inherent complexities of this context, numerous questions naturally arise, including but not limited to: a) What types of knowledge are amenable to successful transfer? b) Which RL structures are best suited for integration into a TL framework? c) What truly distinguishes a source domain from a target domain? These inquiries, among many others, prompt a comprehensive in exploration. While extant literature, such as previous work by [30], has endeavored to shed light on these considerations, we posit that a dedicated examination of these issues within the specific context of Transfer Learning in RL, particularly within computer network problem domains, is required.

7   Conclusions and Future Directions

In summary, this study introduces DESiRED (Dynamic, Enhanced, and Smart iRED) as an innovative solution to tackle the long-standing issue of fixed target delay in AQM systems. By harnessing advanced network telemetry within programmable data planes

and leveraging the capabilities of deep reinforcement learning, DESiRED emerges as a formidable tool to augment TCP congestion control mechanisms. In this novel framework, DESiRED utilizes high-resolution router buffer measurements, collected at line rate within the data plane, as inputs to deep reinforcement learning models residing on the control plane. Empowered by these synergistic components, the agent undertakes dynamic adjustments to the AQM's target delay in real-time, with the overarching goal of optimizing QoS for networked applications.

The comprehensive evaluation conducted within a realistic testbed, featuring the contemporary adaptive bitrate schemes for HTTP-based streaming (MPEG-DASH), reaffirms the viability of DESiRED. Throughout a diverse range of scenarios, encompassing various real-world traffic loads, our results consistently indicate the efficacy of dynamic target delay adjustments in enhancing the QoS of DASH video services for end users.

Considering the inherent dynamism of computer network environments, the prospect of transitioning toward TL has surfaced as a compelling avenue for future exploration. Nevertheless, the intricate challenges associated with this paradigm necessitate dedicated research endeavors to delve into these complexities in greater depth. As such, we recommend that this critical topic be addressed in forthcoming investigations.

References

[1] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, 1993.

[2] W.-c. Feng, K. G. Shin, D. D. Kandlur, and D. Saha, "The blue active queue management algorithms," *IEEE/ACM Trans. Netw.*, vol. 10, p. 513–528, aug 2002.

[3] K. Nichols and V. Jacobson, "Controlling queue delay: A modern aqm is just one piece of the solution to bufferbloat.," *Queue*, vol. 10, p. 20–34, may 2012.

[4] T. Høiland-Jørgensen, D. Täht, and J. Morton, "Piece of cake: A comprehensive queue management solution for home gateways," in *2018 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*, pp. 37–42, 2018.

[5] R. Pan, P. Natarajan, and F. Baker, "PIE: A Lightweight Control Scheme To Address the Bufferbloat Problem," Internet-Draft draft-ietf-aqm-pie-03, Internet Engineering Task Force, 2015. Work in Progress.

[6] L. C. de Almeida, G. Matos, R. Pasquini, C. Papagianni, and F. L. Verdi, "ired: Improving the dash qos by dropping packets in programmable data planes," in *2022 18th International Conference on Network and Service Management (CNSM)*, pp. 136–144, 2022.

[7] R. Kundel, A. Rizk, J. Blendin, B. Koldehofe, R. Hark, and R. Steinmetz, "P4-codel: Experiences on programmable data plane hardware," 2021.

[8] G. Gombos, M. Mouw, S. Laki, C. Papagianni, and K. De Schepper, "Active queue management on the tofino programmable switch: The (dual)pi2 case," in *ICC 2022 - IEEE International Conference on Communications*, pp. 1685–1691, 2022.

[9] R. Boutaba, M. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F. Estrada-Solano, and M. Caicedo, "A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities," *Journal of Internet Services and Applications*, vol. 9, 05 2018.

[10] P4, "In-band network telemetry (int) dataplane specification," tech. rep., P4 Consortium, 2021.

[11] B. Briscoe, K. De Schepper, M. Bagnulo, and G. White, "Rfc 9330: Low latency, low loss, and scalable throughput (l4s) internet service: Architecture," 2023.

[12] ISO, "Dynamic adaptive streaming over http (dash)-part 1: Media presentation description and segment formats," *ISO/IEC*, pp. 23009–1, 2014.

[13] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming protocol-independent packet processors," *SIGCOMM Comput. Commun. Rev.*, vol. 44, p. 87–95, jul 2014.

[14] S. Arslan and N. McKeown, "Switches know the exact amount of congestion," in *Proceedings of the 2019 Workshop on Buffer Sizing*, BS '19, (New York, NY, USA), p. 6, Association for Computing Machinery, 2019.

[15] A. Góes de Castro, F. Rossi, A. Lorenzon, M. Caggiani Luizelli, F. Vogt, R. Hohemberger, and R. Mansilha, "Orchestrating in-band data plane telemetry with machine learning," *IEEE Communications Letters*, vol. 23, p. 20, 10 2019.

[16] R. Boutaba, M. A. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F. Estrada-Solano, and O. M. Caicedo, "A comprehensive survey on machine learning for networking: evolution, applications and research opportunities," *Journal of Internet Services and Applications*, vol. 9, no. 1, pp. 1–99, 2018.

[17] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[18] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, pp. 279–292, 1992.

[19] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[20] C. Busse-Grawitz, R. Meier, A. Dietmüller, T. Bühler, and L. Vanbever, "pforest: In-network inference with random forests," 2022.

[21] X. Chen, S. L. Feibish, Y. Koral, J. Rexford, O. Rottenstreich, S. A. Monetti, and T.-Y. Wang, "Fine-grained queue measurement in the data plane," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, CoNEXT '19, (New York, NY, USA), p. 15–29, Association for Computing Machinery, 2019.

[22] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, 1992.

[23] I. Goodfellow, Y. Bengio, and A. Courville, *Optimization for Training Deep Models*, ch. 8, pp. 271–325. MIT Press, 2016.

[24] L. Almeida, J. Silva, R. Lins, P. M. Jr., R. Pasquini, and F. Verdi, "Wave - um gerador de cargas múltiplas para experimentação em redes de computadores," in *Anais Estendidos do XLI Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, (Porto Alegre, RS, Brasil), pp. 9–16, SBC, 2023.

[25] W. Li, F. Zhou, K. R. Chowdhury, and W. Meleis, "Qtcp: Adaptive congestion control with reinforcement learning," *IEEE Transactions on Network Science and Engineering*, vol. 6, no. 3, pp. 445–458, 2019.

[26] K. Spiteri, R. Sitaraman, and D. Sparacio, "From theory to practice: improving bitrate adaptation in the dash reference player," in *ACM Transactions on Multimedia Computing, Communications, and ApplicationsVolume 15Issue 2s*, pp. 123–137, 06 2018.

[27] J. Gettys, "Bufferbloat: Dark buffers in the internet," *IEEE Internet Computing*, vol. 15, no. 3, pp. 96–96, 2011.

[28] K. Spiteri, R. Urgaonkar, and R. K. Sitaraman, "Bola: Near-optimal bitrate adaptation for online videos," *IEEE/ACM Transactions on Networking*, vol. 28, no. 4, pp. 1698–1711, 2020.

[29] H. Navidan, P. F. Moshiri, M. Nabati, R. Shahbazian, S. A. Ghorashi, V. Shah-Mansouri, and D. Windridge, "Generative adversarial networks (gans) in networking: A comprehensive survey & evaluation," *Computer Networks*, vol. 194, pp. 1–21, 2021.

[30] Z. Zhu, K. Lin, and J. Zhou, "Transfer learning in deep reinforcement learning: A survey," *CoRR*, vol. abs/2009.07888, 2020.