Learning to Harness Bandwidth with Multipath Congestion Control and Scheduling

Shiva Raj Pokhrel, *Senior Member*, IEEE and Anwar Walid, *Fellow*, IEEE

Abstract—**Multipath TCP (MPTCP) has emerged as a facilitator for harnessing and pooling available bandwidth in wireless/wireline communication networks and in data centers. Existing implementations of MPTCP such as, Linked Increase Algorithm (LIA), Opportunistic LIA (OLIA) and BAlanced LInked Adaptation (BALIA) include separate algorithms for congestion control and packet scheduling, with pre-selected control parameters. We propose a Deep Q-Learning (DQL) based framework for joint congestion control and packet scheduling for MPTCP. At the heart of the solution is an intelligent agent for interface, learning and actuation, which learns from experience optimal congestion control and scheduling mechanism using DQL techniques with policy gradients. We provide a rigorous stability analysis of system dynamics which provides important practical design insights. In addition, the proposed DQL-MPTCP algorithm utilizes the 'recurrent neural network' and integrates it with 'long short-term memory' for continuously i) learning dynamic behavior of subflows (paths) and ii) responding promptly to their behavior using prioritized experience replay. With extensive emulations, we show that the proposed DQL-based MPTCP algorithm outperforms MPTCP LIA, OLIA and BALIA algorithms. Moreover, the DQL-MPTCP algorithm is robust to time-varying network characteristics, and provides dynamic exploration and exploitation of paths.**

Index Terms—**Multipath TCP, Deep Q-Learning, Stability Analysis.**

I. INTRODUCTION

CONGESTION Control design of both TCP and multipath TCP (MPTCP) is a fundamental problem in networking and has been widely investigated in the literature (see [2]–[9] and references therein). Most of the seminal single-path congestion control protocols proposed are either *delay-based* (e.g., TCP Vegas), or *loss-based* (e.g., TCP NewReno), otherwise hybrid (both) (e.g., Compound TCP). All of these congestion control protocols use some packet-based events (loss, delay, etc.) as an indication for the congestion and perform their window adjustment (based on some fixed control mechanism) to control the number of outstanding packets in the network.

To take advantage of the increasingly available wireless and wireline access technologies, and mutlipath and multi-homing capabilities in the Internet and data centers, Multipath TCP (MPTCP) has emerged as an enabler for harnessing and aggregating available bandwidth and improve applications' performance. Raiciu *et al.* [10] developed the MPTCP *Linked Increase Algorithm* (LIA) congestion control that elegantly uses

S. R. Pokhrel is with the School of Information Technology, Deakin University, Geelong, Australia. A. Walid is with Nokia Bell Labs, Murray Hill, NJ, USA. The revised version of this paper will appear in IEEE Transactions in Mobile Computing [1]. Email: shiva.pokhrel@deakin.edu.au, anwar.walid@nokia-bell-labs.com
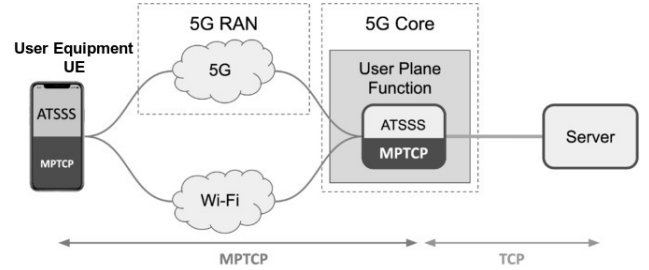

Fig. 1. 5G Network-assisted MPTCP: An abstract view of ATSSS [16]

their specific window increase function to couple the congestion windows running on different subflows; however, it suffers from throughput inefficiency problem which was addressed by *Opportunistic* LIA (OLIA) [11]. Packet scheduling in MPTCP determines which path, among available paths with open windows, to send a packet so that the need for packet re-sequencing at the receiver is minimized. The design of MPTCP packet scheduling has also been equally challenging [12]–[15].

As shown in Fig. 1, MPTCP has a vital role to play in 5G and 6G networks [16]; see the 3rd Generation Partnership Project (3GPP) 5G Release 17 [16] for further details. Since Release 15, the 3GPP introduced a new Access Traffic Steering, Switching, and Splitting (ATSSS) feature [17].[1]

ATSSS is one of the prevalent use case of network-assisted multipath data transport. It has two main components (see Fig. 1): User Equipment (UE) and 5G Core. The UE could natively use MPTCP while the User Plane Function may have a TCP to MPTCP proxy. The 3GPP has already defined the use of the 0-RTT TCP Convert Protocol for the MPTCP proxy, which addresses the proxy directly (IETF RFC-8803, July 2020) [18]. Another option relying on the activation of a dedicated function developed MPTCP *concentrator* [19], where the identity is explicitly considered on the hosts.

One early very prevalent use of MPTCP is the SIRI digital assistant of Apple iPhone, which has been using MPTCP since 2013 (iOS 7) to allow Siri to use both the Wi-Fi or Cellular whatever is available and attain the lowest delay and high reliability. In 2019, MPTCP was also introduced by Apple in their Maps and Music. Samsung is another key player in the MPTCP framework that has been introduced in phones used by Korean Telecom, SK Telecom and LG+ since 2015 in South

---

[1]*Steering*: choosing the best network depending on the user's location and network conditions; *Switching*:enabling seamless handovers from 4G/5G to Wi-Fi (or vice versa) *Splitting*: data transport over multiple paths for higher speeds.

Korea. The goal here was to merge Wi-Fi and Cellular and use them to attain Gigabit speeds at the same time. From 2016, on customer premises devices such as DSL home gateways and/or LTE/4G modems, Tessares and Swisscom adopted MPTCP, mainly to aggregate the capacities of the fixed and mobile networks.

By using fluid-flow model and dynamical system analysis, Peng *et al.* [20] designed *BAlanced LInked Adaptataion* (BALIA) to achieve a better balanced performance. However, all of the aforementioned congestion control approaches [6], [10], [11], [20] designed MPTCP on the basis of *model-based* frameworks, which are challenged in highly dynamic wireless networks (such as 5G[2] [21] and beyond[3] [22]–[24]) as the underlying models do not apply equally well in various scenarios and the solutions can be in suboptimal regimes and unstable states.

To elaborate, the current MPTCP congestion control design mainly consists of two steps: i) developing mathematical models for resource allocation problem (e.g., fluid and/or queuing models [6], [20]), ii) designing algorithms to solve the models (e.g., convex optimization based on its properties [6], [20]). Such models, which are based on the Network Utility Maximization (NUM) or control-theoretic frameworks, typically lead to algorithms with fixed parameters. Such algorithms, however, may be unable to handle highly fluctuating and time-varying wireless channels and data paths.

Furthermore, in current model-based MPTCP design, a user's utility is normally a function of its instantaneous rate, yielding 'bandwidth sharing' networks [25] which fits the case of infinite length flows. All the main single path TCP algorithms proposed in the literature have strictly concave utility functions, implying uniqueness and existence of a stable equilibrium and their convergence. However, the case of MPTCP is much more delicate: whether an underlying utility function exists depends on the design choice. It is worth noting that the utility function does not exist for some well-known MPTCP congestion control algorithms [6], [10], [11], [20] (see [20, Tab. 1]). Therefore, utility maximization with the standard MPTCP congestion control algorithms has been typically challenging and nontrivial.

Moreover, the current MPTCP design does not provide applications with full access to the transport layer features (cross-layer)[4], making it difficult to configure the transport layer behavior to take advantage of the available context and network information. Given that the file size is known, a natural alternative is to use, for a path, the optimal *Shortest Remaining Processing Time* (SRPT) scheduling. However, naively applying SRPT to networks creates starvation for multipath flows. In contrast, one may plan to develop an optimization framework and/or cross-layer approach with i) precise prediction of future values and ii) accurate mathematical model to quantify the network behaviour precisely. But both i) and ii) are extremely challenging and unscalable because of large state space (combinatorial in the number of packets

in flights) over future wireless networks. Therefore, there is a need for robust experience-driven algorithms that can adapt intelligently to dynamic environments without escaping from the TCP/IP layering fundamentals.

A. Literature

Recently, real momentum has been building up towards learning-based TCP designs to exploit and handle dynamic network environments [3], [4], [8], [26]–[33]. Winstein *et al.* [27] proposed *Remy*, a single path congestion control, which is capable of generating control rules for several different network settings. Authors in [26], [34] developed a (single path) performance-based congestion control protocol, such that the source can continuously track and adopts actions timely for a higher network performance [35]. Zaki *et al.* [28] proposed *Verus*, a congestion control protocol that learns a correlation between end-to-end delivery delay of packets to update the congestion window [4], [33]. However, these works [4], [26]–[28], [33], [34], are basically for designing single path congestion control protocols and are not applicable directly for MPTCP.

More recently, advances in machine learning have been applied to the design of MPTCP [3], [8], [15], [29]–[32], [36]. Mai *et al.* [29] applied deterministic policy gradient for learning the optimal congestion control strategies. Liao *et al.* [30] proposed a DRL-based multipath scheduler aiming at decreasing the out-of-order queue sizing under heterogeneous paths. Silva *et al.* [32] implemented an adaptive virtual reality with content-aware prioritisation to improve MPTCP's performance.

Of particular relevance to this work are the two multipath congestion control designs [3], [8] using *deep reinforcement learning* (DRL). Both of these designs [3], [8] focus on congestion control and do not address packet scheduling. Indeed, all of the aforementioned MPTCP designs [3], [8], [15], [29]–[32], [36] considered either congestion control or packet scheduling separately (instead of jointly). It is known that the performance of MPTCP is significantly impacted by packet scheduling which decides what packet to send on which path among the available paths that have open window. Such a decision impacts the goodput (rate of packets delivered to the receiver in the right order) and the required resequencing delay. The advantage of the MPTCP controller in [8], however, is shown only when it is applied to a group of sources, rather than the usual and versatile case of a controller per source. In addition, the stability of the proposed algorithms [3], [8], [29]–[32] are not addressed, whereas stability is a key desirable goal in the original MPTCP design philosophy [10]. Authors in [15], [36] introduced Deep Q Network (DQN) framework to enhance the MPTCP congestion control [36] and packet scheduling [15] performance in asymmetric paths.

In this paper, we develop DQL-MPTCP, a Deep Q-Learning (DQL) based framework for joint congestion control and packet scheduling of MPTCP. At the heart of the solution is an intelligent agent, which learns from experience optimal congestion control and scheduling mechanism using DQL techniques with policy gradients. DQL-MPTCP is a situation-aware learning driven by reward maximization, where the

---

[2]'Considerations for MPTCP operation in 5G' https://tools.ietf.org

[3]'5G ATSSS solution', https://www.tessares.net

[4]Cross-layer approach refers to sharing information among TCP/IP layers for efficacious use of network resources and achieving high degree of adaptivity. It is an escape from the established TCP/IP layering of protocols.

reward is a controllable function representing application performance. Our framework, unlike prior works, employs a DQL agent for joint learning of scheduling and congestion control for each DQL-MPTCP connection. The two controls (congestion and scheduling) jointly and effectively work together towards achieving the target reward. Furthermore, our implementation setup is versatile to use, i.e. whether the sender is a device or a server, readily inter-works with the current Linux kernel implementation (and makes comparisons with standard MPTCP algorithms fair).

Unlike earlier model-driven MPCTP designs [6], [10], [11], [20], our DQL-MPTCP utilizes runtime statistics of the underlying network environment. We combine the Deep Q-Learning (DQL) based control with policy gradients, as it is capable of effectively handling complex and dynamic state spaces. DQL-MPTCP learns dynamic system state and makes decisions without relying on any mathematical formulas or fixed control policies. DQL-MPTCP has to deal with mapping of traffic to multiple paths (subflows) efficiently. One may think of a simple approach and attach a DQL agent for each subflow independently, so as to separately conduct congestion control across all of the available paths. However, such approach is not fair to traditional single path TCP [10], [20] and fails to capture the necessary and essential coupling of paths. DQL-MPTCP maximizes total performance reward by performing coupled congestion control and packet scheduling over all active paths/subflows of a connection. As will be discussed, we adopt a prioritized experience replaying concept to attain desired performance even with a limited data sets.

B. Novelty and Contributions

In this paper, we exploit a setting where a DQL agent access datasets from the history of the system [37]; using prioritized experience replay [38] we store experience tuple (*current state, current action, reward, next state*) in the memory, and then sample most important transitions from the memory more frequently.[5] As the users could be mobile and the number of available network paths (subflows joining and leaving) changes over time, existing DRL and convolution neural network approximation framework. To achieve an MPTCP algorithm that leverages few samples of data to accelerate the online learning even from relatively small datasets, we exploit a *recurrent neural network* integrated with the *long short-term memory* (LSTM) as a replay buffer, and continuously capture packet flow dynamics. More specifically, our design is capable of automatically assessing and selecting the essential data sets from replay during learning.

Our main contributions in this article are as follows:

C$_1$ We propose a new approach for designing coupled multipath congestion control and packet scheduling using Deep Q-Learning with policy gradients.

In contrast to standard MPTCP designs, this new direction in designing multipath TCP performs both packet scheduling and congestion control. The two controls (congestion and scheduling) jointly and effectively work together towards

---

[5]For scalable design (i.e. memory requirements), we consider deleting the tuples after sampling and maintain only the samples.

---

maximizing the performance. Moreover, the aggregate reward can be selected as a function of packet goodput (rate of in-order received packets) and overall delay.

C$_2$ We develop an integrated '*Policy Gradient with Deep Q-Learning*' framework for multipath TCP, that uses an intelligent agent with the objective of maximizing a desired performance reward and performing dynamic congestion control and scheduling across all available network paths.

Accommodation of accurate state and fine-grained actions results in a large state-action space; therefore, we integrate policy gradient algorithm with an actor-critic mechanism. To respond to highly dynamic networks, we exploit the agent's policy distribution to average the value function over a set of actions, rather than only sampled action values.

To elaborate, when a connection is first initiated the end hosts have limited updated knowledge of the characteristics of the path. Typically, information is confined to that gleaned from the initial connection handshake and also from cached historical information. As the connection proceeds, feedback is obtained from packet transmissions but this feedback is delayed. The reason being 25-50 ms typical Round-Trip Times (RTTs) over local wireless internet paths, which can correspond to hundreds of packets 'in flight' (around 150ms for international paths). Forward error correction has been used [24], [39] to ameliorate such circumstances. Short connections (i.e. small TCP flows), thus have limited information as to the path characteristics. Longer connections, e.g., bulk TCP flows, need to learn the path characteristics on the fly while transmitting packets. Our new design is to ameliorate the delay in learning path characteristics.

C$_3$ In our design of DQL-MPTCP, we utilize long short-term memory (LSTM) *recurrent neural network* for tracking and learning of dynamic behaviour of wireless/wireline paths, and appropriately adjusting congestion control and packet scheduling.

We investigate the design of intelligent schedulers that assign packets to paths in such a way as to minimise the reordering delay at the receiver. Key aspects that we address include a delay-aware scheduling, the impact of connection length of the optimal scheduler strategy and the need for joint scheduling of transmissions and learning of path characteristics (see Sec. II-B for details).

C$_4$ We conduct stability analysis of DQL-MPTCP, which provides important practical insights for the protocol design (Sec. IV).

C$_5$ We evaluate DQL-MPTCP and compare its performance with the standard MPTCP algorithms. Our comparison demonstrates significant performance improvements with DQL-MPTCP.

II. PRELIMINARIES AND MOTIVATIONS

Multipath congestion control and packet scheduling mechanisms require continuous control action. However, it is impossible to apply Q-learning to continuous control action directly, because in continuous spaces, finding a greedy policy requires an optimization of control action a_t at every time stamp. In fact, this type of optimization is sluggish in nature and

hard to implement in practice with unconstrained approximates and large (control action) state spaces. Therefore, we use a deterministic policy gradient algorithm with an actor-critic approach. Such gradient-based algorithm maintains an actor (function) which can specify the policy by mapping states to a (specific) control action directly. The critic can be learned by using the Bellman equation (similar to Q-learning). By applying the chain rule to the expected utility, the actor can be modified from an initial distribution with respect to the parameters.

It is well-known that using a nonlinear approximation such as (deep) neural network for reinforcement learning is not desirable [3], [8]. In fact, such a nonlinear approximation is unstable and may lead to a diverging state [40]–[42]. The pioneering work on the Deep Q-networks by Mnih *et al.* [43] laid the foundations for solving complex decision problems; their idea of combining experience replay with Q-learning and convolutional neural network enabled the framework to learn and perform intelligently. This approach not only avoids divergence and oscillations, but also is comparable to humans in making multilevel decisions. We exploit the ideas underlying the success of DQL [35] in our continuous multipath congestion control and scheduling by combing it with policy gradients.

Recurrent neural networks exhibit a high capability of modeling nonlinear time series problems in an effective way; however, there are some issues to be addressed. In particular, they are unable to train time varying network lags that are common in future wireless networks, and rely on a predetermined time duration for learning making them inflexible. To overcome the aforementioned shortcomings of the recurrent neural networks, we use LSTM, that act as a prioritized experience replay buffer for continuously tracking the active subflows and their interactions with the networks.

A. Integrated Learning Framework

We start by employing a Markov Decision Process (MDP), a 5-tuple $(\mathbb{S}(t), \mathbb{A}(t), \boldsymbol{R}(t), \boldsymbol{T}(t), \gamma)$, where $\mathbb{S}(t)$ is a finite set of states, $\mathbb{A}(t)$ is a finite set of actions, $\boldsymbol{R}(t)$ is a reward function, $\boldsymbol{T}(t)$ is the transition function ($\boldsymbol{T}(t)$ determined from the probability state distribution $\boldsymbol{P}(t)$) and γ is a discount factor. In every state $s(t) \in \mathbb{S}(t)$, the agent takes an action $a(t) \in \mathbb{A}(t)$, receives a reward $\boldsymbol{R}(a(t), s(t))$, and attains new state $s(t+1)$ obtained from the probability distribution $\boldsymbol{P}(s(t+1)|s(t), a(t))$. The important notations used in our design and algorithms are summarized in Table I.

We have a policy that specifies for every state which action the agent can take, where the aim of the agent is to discover the policy that potentially maximizes the expected reward. For tractability [27], we consider 6-tuple (*sending rate, throughput, RTT, change in window, schedule, difference in RTT*) as the state of an MPTCP connection.[6]

Our DQL-MPTCP takes action on each MPTCP subflow and provides what change (window increase, decrease and/or

[6]In our experiments we have observed that adding more parameters, increasing data sampling complexity with no noticeable improvement in the performance of the agent.

TABLE I
SUMMARY OF NOTATIONS.

Notation	Description
$\mathbb{N}(.), \mathbb{E}(.), \boldsymbol{Q}(.)$	Representation $\mathbb{N}(.)$, Actor $\mathbb{E}(.)$ and Critic $\boldsymbol{Q}(.)$
$\hat{\mathbb{N}}(.), \hat{\mathbb{E}}(.), \hat{\boldsymbol{Q}}(.)$	Target Representation, Actor and Critic Networks; Replicate the structure of $\mathbb{N}(.), \mathbb{E}(.)$ and $\boldsymbol{Q}(.)$
$\mathbb{S}(t)$	Finite set of states, $s(t) \in \mathbb{S}(t)$, new state $s(t+1)$, $\mathbb{S}(t) = \{s(t,n)\}$ indicates subflow on link n
$\mathbb{A}(t)$	Finite set of actions, $a(t) \in \mathbb{A}(t)$
$\boldsymbol{R}(t)$	Reward function, a reward $\boldsymbol{R}(a(t), s(t))$ from distribution $\boldsymbol{P}(s(t+1)|s(t), a(t))$
$\boldsymbol{T}(t)$	Transition function determined from the probability state distribution $\boldsymbol{P}(t)$
γ	Discount factor.
$Q^\star(s, a)$	Estimate of the expected future reward
$\Pi^\star$	Optimal policy (greedy policy for $Q^\star$)
$\mathbb{F}(t)$	Final state for all subflows of the connection, ultimate final state $\boldsymbol{f}^N(t)$
$\bar{\boldsymbol{\tau}} = \{\bar{\tau}_n\}$	RTTs of all active subflows
$\mathbb{W}(t) = \{W_n\}$	Size of the congestion windows
$U(t, n)$	Utility of the subflow n at time t
$\hat{\boldsymbol{a}}(t)$	Control action for the target subflow
y_i	Target for the Critic $\boldsymbol{Q}(.)$

packet schedule change) needs to be made concurrently to the congestion windows and packet scheduling for the subflows.

An estimate of the expected future reward that can be obtained from $(s(t), a(t))$ is given by the value (using Bellman equation) [37]

$$Q^\star(s, a) = \mathop{\mathbb{E}}_{s' \sim P} \big[R(s, a, s') + \gamma \max_{a'} Q^\star(s', a') \big], \quad (A1)$$

The optimal policy, denoted by $\Pi^\star$ can be obtained as a greedy policy with respect to $Q^\star$:

$$\Pi^\star(s) = \arg\max_{a'} Q^\star(s, a). \quad (A2)$$

It has been known that with respect to $Q^\star$ approximator, algorithm approximates greedy policy as the optimal policy. Our prioritized experience replay technique mandates the DQL agent to sample more frequently the most important state transitions from the memory. For example, we consider the likelihood of sampling a transition proportion to its priority in the spirit of those underlying *learning from demonstrations* [35], [37]. Therefore, with the prioritized experience replay [38], we develop a mechanism to attain comparatively better performance even with a limited data sets and learning time.

B. Joint Scheduling and Congestion Control

A real momentum in recent years has been in the development of packet scheduling and congestion control mechanisms which aim to maximize the performance of MPTCP. See [5], [6], [11], [20], [24] for multipath congestion control [5], [6], [11], [15], [20], [24], [36] and [12]–[14] for the packet scheduling. Perhaps the most known scheduling policy is the one adopted by MPTCP scheduler [10], the *min–RTT*, which assigns packets to the smallest RTT path and fills its window, and so on to the other smallest RTT paths.

The time taken for a packet to traverse a mobile network path is always stochastic and time-varying (*variable path delay*) as a result of queueing, other flows sharing the path, wireless link layer retransmissions etc. When packets are sent via multiple paths they therefore can easily arrive at the destination

reordered in which case they need to be buffered until they can be delivered in order to higher layers, leading to head-of-line blocking. Moreover, the resulting buffering delay in MPTCP can be substantial [6], [24], [39], to the point where it largely undermines the throughput gain from use of multiple paths.[7]

In this work, we propose a new intelligent *Packet Scheduling* policy coupled tightly with multipath congestion control, with a notion that for controlled performance there is a potential to efficiently utilize the time-varying wireline and wireless links even when per-packet delays are highly fluctuating. To attain this, we have to move from the usual notion of considering packets individually to a new consideration of jointly scheduling collections of packets subject to an overall minimal delivery delay from all active subflows.

This new notion of scheduling is driven by our observation that the QoE (quality of experience) requirement is usually to transmit application layer entities such as video clips, icons, web pages, micro-blogs, etc. with the lowest attainable delay, and *it is the overall aggregate delay rather than the per-packet delay which is critical* in such context. This observation has fundamental implications in our packet scheduler design. We are scheduling packets in groups; therefore, such scheduling is capable of considering application layer contexts and their relationship to path uncertainties.

The packet scheduling policy in this work is coupled tightly with the multipath congestion control mechanism. Our scheduler assigns a number of packages to each path proportional to their short term average goodput (*number of successful packets per RTT*).[8] Given the availability in the windows (specified by the multipath congestion control process), our scheduler always assigns packets to the paths proportionally, ranging from the largest goodput path to the small goodput path. Such a scheduling policy has the effect of minimizing packet re-sequencing at the receiver [44].

III. THE PROPOSED MPTCP SCHEME

In this section, we present the design of our *Deep Q-Learning enabled MPTCP with policy gradients* for joint congestion control and packet scheduling across paths of different characteristics. At a high level, the interactions between modules of the proposed DQL-MPTCP are as illustrated in Figure 2. A detailed explanation of actor critic training, representation network and subflow state analysis is provided with algorithms in Section 3.1-3.3. As discussed in Section 3.3, the DQL agent (submodules inside the dashed portion in Figure 2) of the proposed MPTCP interacts with the links and the user to collect the freshest information of the state $\mathbb{S}(t) = \{s(t,n)\}$ here, n indicates an MPTCP subflow on the link n. At the beginning of time slot t, the agent computes the reward, $\boldsymbol{R}(t)$ (the sum of utilities of all active subflows from previous actions), by
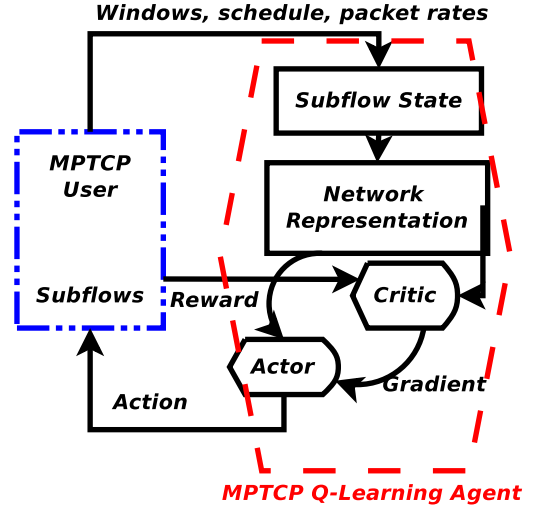

Fig. 2. A abstract view of interrelationship between the DQL-MPTCP modules. The modules inside the dashed region (form a DQL Agent) take packet rates, reward, congestion windows as input and generate control actions as outputs. A module on the left rectangle (is a MPTCP user with source program at the OS kernel), which provides new packet rates, schedule and congestion windows.

using the actor-critic network (Sec. III-B). The computation is based on the representation learned (Sec. III-A) by the LSTM and the current state of the subflows $\mathbb{S}(t)$ (final states $\mathbb{F}(t)$ and $\boldsymbol{a}(t)$ are used to compute $\boldsymbol{R}(t)$).

The DQL-agent is queried periodically by the MPTCP source (one query per slot), in order to i) update the size of the windows ($\mathbb{W}(t+1)$), ii) update the RTTs ($\bar{\boldsymbol{\tau}} = \{\bar{\tau}_n\}$) of active subflows and iii) schedule packets to the subflows (*schedule*, for next slot, see Figure 2). Rather than using only the actions that were actually executed, we consider a policy gradient algorithm that uses the agent's explicit representation of all action values to estimate the gradient of the policy. Therefore, our framework implements the action (via packet scheduling and congestion control at the MPTCP source in the kernel) by observing the reward $R(t)$. To ensure stability and maintain replay buffer, we have used target Network Representation $\hat{\mathbb{N}}(.)$ Actor $\hat{\mathbb{E}}(.)$ and Critic $\hat{\boldsymbol{Q}}(.)$, which replicate the structure of their corresponding networks in the state analysis, i.e., Network Representation $\mathbb{N}(.)$, Actor $\mathbb{E}(.)$ and Critic $\boldsymbol{Q}(.)$ respectively.

A key feature of DQL is its use of a target network whose purpose is to stabilize the learning process. In traditional Q-learning, the value of executing an action in the current state is updated using the values of executing actions in the next state. This process can be unstable since the values on both sides of the update equation can change at the same time. The target network keeps a copy of the estimated value function to serve as a stable target for a few future steps [43]. The main building blocks in Figure 2 and our algorithms are explained next.

A. Network Representation

The Network Representation module creates a (representation) vector simply by observing the states of MPTCP subflows

---

[7]In existing works [6], [24], [39], the problem of *variable path delay* was partially addressed by employing delay adaptation [6] and forward error correction [24], [39] at the MPTCP source.

[8]The transient behaviour, i.e., the trajectories through which short term average goodputs ($\theta_n(t)$) approach their steady values, is essential for the computation of exact reordering delay. However, we are concerned with the difference in short term steady goodputs, and our interest lies in designing a solution that can alleviate the severe reordering delay perceived at the application layer.

($\boldsymbol{s}(t,n)$ as inputs). However, it is challenging to deal with the representation when the number of subflows are time-varying (joining and leaving the MPTCP connection), which will be common in future wireless networks. Rather than using fixed input size Deep Neural Network approach (e.g., feed forward), we use LSTM which has the capability to handle variable number of inputs over time so as to capture the dynamics of varying number of MPTCP subflows (corresponding to the conditions of the available network paths and access interfaces). The approach is to provide the states of the subflows one by one into the LSTM for learning the Network (and Delay) Representation in a sequential manner.

The ultimate final state $\boldsymbol{f}^N(t)$, the output of this module (we denote $\mathbb{F}(t)$ for all subflows of the connection), is then fed for Actor-Critic training. Furthermore, we propose to train both 'LSTM-based representation' and 'Actor-Critic Network' jointly by using backward propagation. We have an inner loop between three modules, viz. *subflow state analysis* → *network representation* → *actor-critic training*. It is worth noting that jointly training the interrelated modules (by observing the state of the subflows) yields improved performance rather than training them individually.

B. Actor-Critic Training

The actor network consists of completely connected LSTM with two hidden layers, that consists of 128 neurons in both layers. The critic is very similar to the actor except of the output layer with a single linear neuron. We use Rectified Linear function in hidden layers and hyperbolic tangent function in the output layer for activation.

The final state $\mathbb{F}(t)$ produced by the LSTM-based representation is combined with the state of the corresponding MPTCP subflow, which is then supplied to the Actor-Critic network. The agent specifies at each time t how to update the size of the congestion window ($\mathbb{W}(t)$) and changes schedule (if any) for each subflows of the MPTCP connection in order to maximize the reward.

The reward in our framework is the sum of the utilities of all active subflows belonging to the MPTCP connection, i.e., $R(t) := \sum_n^N U(t,n)$ where, $U(t,n)$ is the utility of the subflow n at time t. It is worth noting that our framework is flexible to define and handle desired utility function (e.g., fairness-based [25] or, loss, delay, and throughput-based [34]) and scheduling policy.

Resources can be allocated amongst the competing flows according to various policies. To perform MPTCP congestion control and scheduling, we aim to facilitate applications with efficiency and fairness. Although there is no global notion of fairness, proportional and max-min fair allocations [45], [46] are the two well-known schemes.

Proportional Fairness [47, Eqn.(2.10), pp. 37]: For a set of users $\mathcal{U} = \{1, 2, ..., n\}$ the goodput allocation $[\theta_1^*, \theta_2^*, ..., \theta_n^*]$ to the set of users $\mathcal{U}$ is proportionally fair if and only if for any other feasible scheme $[\theta_1, \theta_2, ..., \theta_n]$, such that

$$\sum_{u \in \mathcal{U}} \frac{\theta_u - \theta_u^*}{\theta_u^*} \leq 0. \tag{1}$$

Algorithm 1 Training Algorithm

1: **procedure** TRAINING
2: Sample transitions $(\mathbb{S}_i, \boldsymbol{a}_i, \boldsymbol{R}_i, \mathbb{S}_{i+1})$ from priority replay
3: Compute $\boldsymbol{f}_{i+1}^N$ using $\hat{\mathbb{N}}(\mathbb{S}_{i+1})$
4: Evaluate target for $\boldsymbol{Q}(.)$, $y_i = \boldsymbol{R}_i + \gamma \hat{\boldsymbol{Q}}(\boldsymbol{f}_{i+1}^N, \hat{\mathbb{E}}(\boldsymbol{f}_{i+1}^N))$
5: Critic parameters $\min(loss)$, $1/k \sum_i^k (y_i - \boldsymbol{Q}(\boldsymbol{a}_i, \boldsymbol{f}_i^N))^2$
6: Compute Policy Gradient from Critic: $\bigtriangledown_a \boldsymbol{Q}(\boldsymbol{a}_i, \boldsymbol{f}_i^N)$
7: Update actor parameters, $1/k \sum_i^k \bigtriangledown_a \boldsymbol{Q}(\boldsymbol{a}_i, \boldsymbol{f}_i^N) \cdot \bigtriangledown_k \mathbb{E}(\boldsymbol{f}_i^N)$
8: Compute Policy gradient from Actor: $\bigtriangledown_k \mathbb{E}(\boldsymbol{f}_i^N)$
9: Update network representation (policy gradient),
   $1/k \sum_i^k \bigtriangledown_a \boldsymbol{Q}(\boldsymbol{a}_i, \boldsymbol{f}_i^N) \cdot \bigtriangledown_k \mathbb{E}(\boldsymbol{f}_i^N) \cdot \bigtriangledown_{\boldsymbol{R}_i} \boldsymbol{R}(\mathbb{S}_{i+1})$
10: **end procedure**

In (1), when the proportional change in one user's packet rate is positive, at least another user for which the change in packet rate is negative.

In our implementation, we have used the well-known proportional fairness approach for all subflows (active paths) and maximize the function $U(t,n) = log\theta_n(t)$, where $\theta_n(t)$ is the short-term average goodput along path n in the previous slot perceived by the subflow. For delay minimization, our scheduling policy (recall Section II-B) implements batch scheduling of packets to the subflows with a rate proportional to their short-term average goodputs.

Remark III.1. Despite its empirical success and advancement, the theoretical perception behind the convergence of the actor-critic algorithm is lagging. Actor-Critic learning can be viewed as an alternating online bilevel optimisation process, the convergence of which is known to be fragile. In Sec. 4, we concentrate on applying DQL agent regulators, a basic but essential determination in Q-learning with gradients, to understand the learning dynamics' instability. Our aim in this work is to design a flexible framework, the stability of which is independent of the underlying utility function. Therefore, in this setting, we analyse a non-asymptotic convergence of the DQL-learning and demonstrate how it seeks a linear rate of convergence that is globally optimal. Our stability analyses in Sec. 4 can be the first step to thoroughly comprehend bilevel optimisation for the utility-based actor-critic problems, which will be investigated in future (in the worst case this can be NP-hard, non-convex and is mostly overcome with heuristics).

C. DQL-MPTCP with Policy Gradients

Our DQL-based multipath congestion control and scheduling mechanism is illustrated in Algorithms 1 and 2. Algorithm 1 illustrates the training process. Algorithm 2 initializes the parameters of Representation Network $\mathbb{N}(.)$, Critic networks $\boldsymbol{Q}(.)$ and Actor networks $\mathbb{E}(.)$. In order to ensure stable learning (see step 3, Algorithm 2), we have used target networks $\hat{\mathbb{N}}(.)$ $\hat{\mathbb{E}}(.)$ $\hat{\boldsymbol{Q}}(.)$, which replicate the structure of their corresponding networks, i.e., $\mathbb{N}(.)$, Actor $\mathbb{E}(.)$ and Critic $\boldsymbol{Q}(.)$ respectively.

For the convergence of the learning process, the target network parameters are updated using a smallish control parameter (0.001, see steps 17-19, Algorithm 2), such that the target parameters are adapted slowly in each iteration–recall that the target network, by design, needs to update slowly for

stability.[9] Our MPTCP DQL agent runs all the time, listening for periodic queries from the MPTCP source (implemented in the kernel).

To enable exploration, which is useful when training in-experienced DQL agent, we add correlated noise, using the Ornstein-Uhlenbeck stochastic process [35], [37] to control actions in run-time. Such coupling of prioritized experience replay with exploration helps the DQL agent improve its actions of adjusting subflow window sizes and packet scheduling.[10]

Recall that the actor and critic networks are completely connected LSTM with two hidden layers, that consists of 128 neurons in both layers. We use Rectified Linear function in the hidden layers and hyperbolic tangent function in the output layer for activation. The critic and actor networks are trained jointly using Adam optimizer (*we set learning rates to 0.001 and 0.01 respectively and discount factor $\gamma = 0.95$*).

The final state network representation of all active subflows $\mathbb{F}(.)$ is derived from the representation network $\mathbb{N}(.)$ (see step 8, Algorithm. 2), and the control action for the target MPTCP subflow $\hat{a}(t)$ is computed by using the actor $\mathbb{E}(.)$ (see step 9, Algorithm. 2). First, the generated transition samples are stored into kernel (memory), thereafter they are randomly sampled for training the tuple (*network representation $\mathbb{N}$, actor $\mathbb{E}$, critic Q*) jointly by using k prioritized samples (step 2, Algorithm. 1). The critic is a *Deep Q-Learning Network* and its parameters are updated by minimizing the squared error (step 5, Algorithm. 1), i.e., the target for critic y_i is evaluated by applying the Bellman equation (step 4, Algorithm. 1). The Q-function uses the Bellman equation and takes action ($a(t)$) and state ($\mathbb{S}(t)$) as inputs. For continuous control, the parameters of the network representation and actor networks are adapted together with the policy gradients using the chain rule [35, Eqn. (6)] by using k samples (steps 6-10, Algorithm. 1). Our DQL framework consisting of the representation, critic and actor networks has key ingredients for training and learning MPTCP packet scheduling and congestion control.

Remark III.2. The complexity analysis of DQL algorithm is another evolving research direction. To the best of our knowledge, despite the popularity of DQL with policy gradients, we are still investigating to accurately predict the computational complexity to train/learn a DQL network and solve a given problem. With the relevant insights from this work and [49], [50], an important future research direction would be a thorough complexity analysis of DQL-MPTCP.

IV. STABILITY ANALYSIS

Peng *et al.* [20] conducted a detailed study of the various variants of the MPTCP using a generalised fluid model and analyzed their stabilization aspects. In fact, analyzing stability has been an essential first step in designing a congestion control algorithm to guarantee it has a desirable equilibrium and convergence properties. Our DQL-MPTCP algorithm is

---

Algorithm 2 DQL-MPTCP Coupled Algorithm
1: **procedure** BEGIN
2:   Representation Network $\mathbb{N}(.)$, Actor $\mathbb{E}(.)$ and Critic $Q(.)$
3:   Target Networks $\hat{\mathbb{N}}(.)$ $\hat{\mathbb{E}}(.)$ $\hat{Q}(.)$
4:   Ornstein-Uhlenbeck process $\circlearrowright$ for exploration
5: **end procedure**
6: **procedure** SCHEDULING & CONGESTION CONTROL
7:   **While** (Schedule Change ‖ Window Adapt) **do**
8:     Compute final state $\mathbb{F}(.)$ and schedule by $\mathbb{N}(.)$
9:     Update target $\hat{a}(t)$ using Actor $\mathbb{E}(f^N(t))$
10:    Create action $a(t)$ applying $\hat{a}(t)$ and process $\circlearrowright$
11:    Take action $a(t)$, observe $R(.)$ and $\mathbb{S}(t+1)$
12:    Store transitions $(a(t), \mathbb{S}(t), R(t), \mathbb{S}(t+1))$
13:    **procedure** ON-LINE TRAINING
14:      Call TRAINING
15:    **end procedure**
16:    **procedure** UPDATE TARGET NETWORK
17:      $\hat{\mathbb{N}}(.) := .001\mathbb{N}(.) + .999\hat{\mathbb{N}}(.)$
18:      $\hat{\mathbb{E}}(.) := .001\mathbb{E}(.) + .999\hat{\mathbb{E}}(.)$
19:      $\hat{Q}(.) := .001\mathbb{C}(.) + .999\hat{Q}(.)$
20:    **end procedure**
21:   **endWhile**
22: **end procedure**

---

first trained offline (using Algorithm. 1) and then deployed, and goes along to learn online (step 14 in Algorithm. 2 calls Algorithm. 1). Therefore, we examine whether the learning process of DQL-MPTCP will indeed drive the network towards an equilibrium starting from an arbitrary initial state. Even though in reality a network is seldom in equilibrium, a stable online learning process ensures that it is always pursuing a desirable state, which also makes it easier to understand the global protocol behavior of the overall network.

For improving smoothness and responsiveness in the implementation, we indirectly guarantee the stability by imposing few constraints in the underlying Q-learning process of our MPTCP algorithm. Given the convergence of value iteration (the Q-learning process), the policy iteration is guaranteed to converge. Further, at the convergence point, the current policy and its value function are the optimal policy and the optimal value function. Recall (A2), the optimal policy $\Pi^\star$ is a greedy policy corresponding to $Q^\star$. In fact, any unconstrained policy iteration process leverages wider state space and region. Besides, the convergence of the policy iteration has been proven to be not slower than that of the value iteration, therefore, policy iteration has the potential to improve both smoothness and responsiveness.Based on above discussion, our focus in this section is mainly to quantify the convergence conditions of value iteration for the Q-learning process of our DQL-MPTCP.

Divergence is an important challenge in Deep Q Learning; it is not well understood and often arises in implementations. Therefore, it is important to analyze the stability of our proposed DQL-MPTCP algorithm [40], [42].

With relevant insights from [42], the DQL ideas implemented in Algorithm 1 is to learn an approximate to the optimal value function $Q^\star$, which satisfies (A1). With $\mathcal{T}^\star Q(s,a)$ given by the right hand side of (A1), after $\mathcal{T}^\star : \mathcal{Q} \rightarrow \mathcal{Q}$ operation on

---

[9]The rationale for selecting smallish control parameter values comes from our stability analysis in Section IV.

[10]An effective way to conduct retraining of the agent to adapt for the new network settings will benefits from transfer learning [48] and needs further investigations in future.

Q functions; then (A1) is given by

$$Q^\star = \mathcal{T}^\star Q^\star, \tag{A3}$$

where $\mathcal{T}^\star$ is the optimal Bellman operator, which with modulus γ is a contraction in the supremum norm. As a result, we have two cases.

CASE I. Given that the Q-function is represented by a finite table with (completely) known $\boldsymbol{R}_i$ (reward function) and transition kernel (steps 12 and 14 Algorithm 1), $\mathcal{T}^\star$ can be determined and therefore, $Q^\star$ can be estimated by using the successful method for computing an optimal Markov Decision Process policy (and its value), known as *value iteration*. For such cases, the value iteration starts at the end and then works backward, refining the estimate of $Q^\star$. Consider Q^i be the Q-function assuming there are i stages to go, then, these can be defined recursively; and the iteration starts with an arbitrary function Q_0 and uses the following equation to obtain the function for $i + 1$ stages (to move from the function for i stages to go):

$$Q_{i+1} = \mathcal{T}^\star Q_i. \tag{A4}$$

The convergence of the value iteration given by (A4) using Q_0 as an initial point is undertaken by the *Banach Fixed Point Theorem*.

CASE II. Given that $\boldsymbol{R}_i$ and the transition kernel are partially known (not completely known), it is viable to use Q-learning and learn $Q^\star$ in such settings [41], [42]. In fact, Watkins *et al.* [41] demonstrated that Q-learning converges to the optimum $\boldsymbol{a}^\star$ (action-values) with probability 1 so long as all actions are repeatedly sampled in all stages and the action-values are represented discretely (steps 10, Algorithm. 2 and 4-10, Algorithm. 1). The Q-values of the (s, a) pairs are updated using reward and next state for estimating $\mathcal{T}^\star Q_i(s, a)$ given by

$$\mathcal{T}^\star Q_i(s, a) = \boldsymbol{R}_i + \gamma \max_{\hat{a}} Q_i(\hat{s}, \hat{a})$$

as:

$$Q_{i+1}(s, a) = Q_i(s, a) + \alpha_i \big(\mathcal{T}^\star Q_i(s, a) - Q_i(s, a) \big). \tag{A5}$$

(A5) under much relaxed condition, viz., *the learning rates $\alpha_i \in [0, 1]$ must approaches zero and all state action pairs must be visited often enough*, converges to $Q^\star$. One can observe that the Q-learning given by (A5) works on time-varying difference as the updates in each stages are based on the temporal difference $(\mathcal{T}^\star Q(s_t, a_t) - Q(s_t, a_t))$ given by:

$$\Delta_t = \boldsymbol{R}_t + \gamma \max_{\hat{a}} Q(\hat{s}_t, \hat{a}_t) - Q(s_t, a_t). \tag{A6}$$

Motivated by the analysis in [42], our DQL-MPTCP is based on the generalization of (A5) to the approximation setting:

$$g' = g + \alpha \big(\mathcal{T}^\star Q_g(s, a) - Q_g(s, a) \big) \nabla_g Q_g(s, a) \big], \tag{A7}$$

where Q_g is a continuous function with parameters g. Observe that (A7) becomes (A5) given Q_g is a table. On the whole, such DQL representation uses gradient descent and experience replay, maintaining the expected update as:

$$g' = g + \alpha \mathop{\mathbb{E}}_{s, a \sim \rho} \big[\big(\mathcal{T}^\star Q_g(s, a) - Q_g(s, a) \big) \nabla_g Q_g(s, a) \big], \tag{A8}$$

where ρ, at the time of updates, is the distribution of the prioritized experience in the replay. Typically, to ensure stability, it is acceptable to replace, $\mathcal{T}^\star Q_g$ with the one based on slowly-updating target network, $\mathcal{T}^\star Q_h$, where h is obtained by Polyak averaging [35], [40]–[42].[11]

Using Taylor expansion of Q around g for the pair $(\hat{s}, \hat{a})$ we can observe the new Q-values based on (A8) as:

$$Q_g'(\hat{s}, \hat{a}) = Q_g(\hat{s}, \hat{a}) + \nabla_g Q_g(\hat{s}, \hat{a})^T (g' - g) + \mathrm{O}\big(\|g' - g\|^2 \big), \tag{A9}$$

Over a finite state action space, we consider matrix vector form in $\mathbb{R}^{|\mathbb{S}||\mathbb{A}|}$, therefore, combining (A8) and (A9) provides:

$$Q_g' = Q_g + \alpha \mathcal{K}_g D_\rho \big(\mathcal{T}^\star Q_g(s, a) - Q_g(s, a) \big) + \mathrm{O}\big(\|g' - g\|^2 \big), \tag{A10}$$

where D_ρ is a diagonal matrix obtained from the distribution from replay, $\rho(s, a)$ and $\mathcal{K}_g$ is $|\mathbb{S}||\mathbb{A}| \times |\mathbb{S}||\mathbb{A}|$ a matrix given by

$$\mathcal{K}_g(\hat{s}, \hat{a}, s, a) = \nabla_g Q_g(\hat{s}, \hat{a})^T \nabla_g Q_g(s, a).$$

Finally, our formulation (A10) provides us important insights to analyse the stability of DQL-MPTCP. In particular, the only condition essential now for convergence of DQL-MPTCP is to guarantee that the update operator $\mathcal{U} : \mathcal{Q} \to \mathcal{Q}$ with

$$\mathcal{U} Q_g = Q_g + \alpha \mathcal{K}_g D_\rho \big(\mathcal{T}^\star Q_g(s, a) - Q_g(s, a) \big) \tag{A11}$$

is a contraction on $\mathcal{Q}$. With relevant insights from [40]–[42], this will be considered next.

A. Practical Design Insights

In this subsection, we investigate the potential conditions and study how the update $\mathcal{U} : \mathcal{Q} \to \mathcal{Q}$ given by (A11) may give rise to instability in DQL-MPTCP and how to repair such instabilities. Next, motivated by the analysis in [42], we decompose our analysis into following cases.

CASE I. We consider $\mathcal{U}^i$, a special case of $\mathcal{U}$ where $\mathcal{K}_g = 1$ and $D_\rho = 1$, therefore, (A11) becomes

$$\mathcal{U}^i Q_g = Q_g + \alpha \big(\mathcal{T}^\star Q_g - Q_g \big). \tag{A11.1}$$

Theorem IV.1. *$\mathcal{U}^i$ defined by (A11.1) is a contraction on $\mathcal{Q}$ and $Q^\star$ is the fixed-point.*

Proof. Observe that $\mathcal{U}^i$ given by (A11.1) satisfies

$$
\begin{aligned}
& \big\| \mathcal{U}^i Q_1 - \mathcal{U}^i Q_2 \big\|_\infty \\
= \ & \big\| \alpha (\mathcal{T}^\star Q_1 - \mathcal{T}^\star Q_2) + (1 - \alpha)(Q_1 - Q_2) \big\|_\infty, \\
\leq \ & \alpha \| \mathcal{T}^\star Q_1 - \mathcal{T}^\star Q_2 \|_\infty + (1 - \alpha) \big\| (Q_1 - Q_2) \big\|_\infty, \\
\leq \ & \alpha \gamma \| Q_1 - Q_2 \|_\infty + (1 - \alpha) \big\| (Q_1 - Q_2) \big\|_\infty, \\
= \ & \| Q_1 - Q_2 \|_\infty (1 - \alpha(1 - \gamma)),
\end{aligned}
$$

and since $1 > (1 - \alpha(1 - \gamma))$, the update $\mathcal{U}^i$ contracts and $Q^\star$ is its fixed-point (follows straightforward using (A3)). $\square$

CASE 1 and Theorem IV.1 provide us the following important design insight for stable MPTCP.

---

[11] For simplicity and tractability [35], we omit the target network in our stability analysis.

DESIGN INSIGHT 1. Given $\mathcal{U}$ of the proposed DQL-MPTCP gets increasingly closer to $\mathcal{U}^i$, one can anticipate progressively more stable DQL-MPTCP behavior.

CASE II. We consider $\mathcal{U}^{ii}$, a special case of $\mathcal{U}$ where $\mathcal{K}_g = 1$, therefore, (A11) becomes

$$\mathcal{U}^{ii} Q_g = Q_g + \alpha D_\rho(\mathcal{T}^\star Q_g - Q_g). \quad \text{(A11.2)}$$

Theorem IV.2. $\mathcal{U}^{ii}$ defined by (A11.2) is a contraction on $\mathcal{Q}$ and $Q^\star$ is the fixed-point, if $\rho(s,a) > 0, \forall (s,a)$ and $\alpha \in (0, 1/\hat{\rho})$ where $\hat{\rho} = \max_{s,a} \rho(s,a)$.

Proof. One can observe that for any (s,a),

$$
\begin{aligned}
& [\mathcal{U}^{ii} Q_1 - \mathcal{U}^{ii} Q_2](s,a) \\
=\ & \alpha\rho(s,a)\big([\mathcal{T}^\star Q_1 - \mathcal{T}^\star Q_2](s,a)\big) \\
& + (1 - \alpha\rho(s,a))\big(Q_1(s,a) - Q_2(s,a)\big), \\
\leq\ & \alpha\gamma\rho(s,a)\|Q_1 - Q_2\|_\infty + (1 - \alpha\rho(s,a))\|(Q_1 - Q_2)\|_\infty \\
=\ & \|Q_1 - Q_2\|_\infty (1 - \alpha\rho(s,a)(1 - \gamma)).
\end{aligned}
$$

By taking the $\max_{s,a}$ on both sides,

$$
\begin{aligned}
& \left\|\mathcal{U}^{ii} Q_1 - \mathcal{U}^{ii} Q_2\right\|_\infty \\
\leq\ & \max_{s,a}\|Q_1 - Q_2\|_\infty (1 - \alpha\rho(s,a)(1 - \gamma)). \\
=\ & \|Q_1 - Q_2\|_\infty (1 - \alpha\tilde{\rho}(1 - \gamma)),
\end{aligned}
$$

where $\tilde{\rho} = \min_{s,a} \rho(s,a)$. Note that the condition, $\rho(s,a) > 0, \forall (s,a)$, is equivalent to $\tilde{\rho} > 0$. Therefore, $1 > (1 - \alpha\tilde{\rho}(1 - \gamma))$, which mandates that the update $\mathcal{U}^{ii}$ contracts and $Q^\star$ is its fixed-point (follows straightforward using (A3)). However, observe that when $\tilde{\rho} = 0$, we simply have an upper bound on $\left\|\mathcal{U}^{ii} Q_1 - \mathcal{U}^{ii} Q_2\right\|_\infty$. See [42] for more background. $\square$

Considering $\mathcal{U}^{ii}$, we observe that the missing data in the distribution has adverse impact on the convergence of the learning process. Given that the (exploration) policy explores all state action pairs enough, $\mathcal{U}^{ii}$ behaves as expected, however, missing data may cause a problem. This observation provides us another important design insight.

DESIGN INSIGHT 2. DQL-MPTCP may struggle to converge, when data is scarce at the beginning of the training, where initial conditions matter a lot.

The importance of DESIGN INSIGHT 2 will be discussed in Section IV-B.

CASE 3. We consider $\mathcal{U}^{iii}$, a special case of $\mathcal{U}$ where $\mathcal{K}_g = \mathcal{K}$ is positive-definite constant symmetric matrix, therefore, (A11) becomes

$$\mathcal{U}^{iii} Q_g = Q_g + \alpha\mathcal{K} D_\rho(\mathcal{T}^\star Q_g - Q_g). \quad \text{(A11.3)}$$

Therefore, $\mathcal{U}^{iii}$ is the case of linear function approximation, $\mathcal{K}$ is constant with respect to g and the Q-values before and after updates are given by

$$\mathcal{Q}_{g'} = \mathcal{U}^{iii} Q_g; \text{ where}$$

$$\mathcal{K}(\hat{s}, \hat{a}, s, a) = \phi(\hat{s}, \hat{a})^T \phi(s, a).$$

Theorem IV.3. $\mathcal{U}^{iii}$ defined by (A11.3) is a contraction on $\mathcal{Q}$ and $Q^\star$ is the fixed-point, if and only if i) $\alpha\rho_x\mathcal{K}_{xx} > 1$, $\forall x$ and

$$ii)\ (1 - \gamma)\rho_x\mathcal{K}_{xx} \geq (1 + \gamma)\sum_{x \neq y} \rho_y|\mathcal{K}_{xy}|, \forall x,$$

where x, y are the indices of state action pairs.

Proof. By using index notation

$$
\begin{aligned}
& [\mathcal{U}^{iii} Q_1 - \mathcal{U}^{iii} Q_2]_x \\
=\ & \alpha\sum_y \rho_y\mathcal{K}_{xy}\big([(\mathcal{T}^\star Q_1 - Q_1) - (\mathcal{T}^\star Q_2 \\
& - Q_2)]_y\big) + [Q_1 - Q_2]_x, \\
=\ & \alpha\sum_y \rho_y\mathcal{K}_{xy}[\mathcal{T}^\star Q_1 - \mathcal{T}^\star Q_2]_y \\
& + \sum_y (\Delta xy - \alpha\rho_y\mathcal{K}_{xy})[Q_1 - Q_2]_y. \\
\leq\ & \sum_y \big(\alpha\gamma\rho_y|\mathcal{K}_{xy}| + |\Delta xy - \alpha\rho_y\mathcal{K}_{xy}|\big)\|Q_1 - Q_2\|_\infty.
\end{aligned}
$$

Let

$$G(\mathcal{K}) = \max_x \sum_y \big(\alpha\gamma\rho_y|\mathcal{K}_{xy}| + |\Delta xy - \alpha\rho_y\mathcal{K}_{xy}|\big).$$

Then, assuming $\alpha\rho_x\mathcal{K}_{xx} > 1$, condition i),

$$
\begin{aligned}
G(\mathcal{K}) =\ & \max_x \Big(\alpha(1 + \gamma)\sum_{x \neq y} \rho_y|\mathcal{K}_{xy}| + |1 - \alpha\rho_y\mathcal{K}_{xy}| \\
& + \gamma\alpha\rho_y\mathcal{K}_{xy}\Big). \\
=\ & \max_x \Big(1 + (1 + \gamma)\sum_{x \neq y} \rho_y|\mathcal{K}_{xy}| - (1 - \gamma)\rho_x\mathcal{K}_{xx}\Big).
\end{aligned}
$$

Observe that $G(\mathcal{K}) < 1$ if an only if

$$\forall x,\ (1 - \gamma)\rho_x\mathcal{K}_{xx} \geq (1 + \gamma)\sum_{x \neq y} \rho_y|\mathcal{K}_{xy}|.$$

See [42] for other details. $\square$

Considering $\mathcal{U}^{iii}$, we observe that the conditions in **Theorem IV.1** are quite confining, it not only requires $\rho > 0$ everywhere but also for the choices of γ (eg. $\gamma = 0.999$) the (diagonal terms of $\mathcal{K}$)$\gg$(off-diagonal terms of $\mathcal{K}$). This observation provides us another important design insight.

DESIGN INSIGHT 3. The stability of DQL-MPTCP depends on the properties of the Q-approximator; in the sense that approximators which are more aggressive (larger off-diagonal terms in $\mathcal{K}$) may not demonstrate stable learning.

Generally, in our DQL-MPTCP learning, both the $\mathcal{K}$ and the ρ change between update stages, thus, each stage can be perceived as applying different updates. On this end, in general, if we sequentially apply different contraction maps with the same fixed point, then we will atain that fixed point. As a result, motivated by the findings in [42], we have the following theorem.

Theorem IV.4. *Assume a sequence of updates $\{\mathcal{U}^0, \mathcal{U}^1, \mathcal{U}^2, \dots\}$ with each $\mathcal{U}^i : \mathcal{Q} \to \mathcal{Q}$ being Lipschitz continuous with constant δ_i and the existence and uniqueness of the fixed point $Q^\star = \mathcal{T}^\star Q^\star$ is guaranteed, where all $\mathcal{U}^i$ share a common fixed point $Q^\star$. Then, starting from any initial point Q^0, the trajectory $\{Q^0, Q^1, Q^2, \dots\}$ generated by the DQL-MPTCP algorithm converges to a unique and globally stable fixed*

point $Q^\star$, if and only if the stages produced by $Q_{i+1} = \mathcal{U}^i Q_i$ satisfies (for an iterate k):

$$\left\| Q^\star - Q^0 \right\| \prod_{j=0}^{i-1} \delta_j \geq \left\| Q^\star - Q^i \right\| \text{ such that } \forall j \geq k, \delta_j \in [0,1).$$
(A11.4)

Proof. Using fixed-point assumption and iterative sequence of updates, $\left\| Q^\star - Q^i \right\| = \left\| \mathcal{U}^{i-1} Q^\star - \mathcal{U}^{i-1} Q^{i-1} \right\|$. From the definition (and property) of Lipschitz continuity $\left\| Q^\star - Q^i \right\| \leq \delta_{i-1} \left\| Q^\star - Q^{i-1} \right\|$. Therefore, $\left\| Q^\star - Q^0 \right\| \prod_{j=0}^{i-1} \delta_j \geq \left\| Q^\star - Q^i \right\|$. Finally, sequence $\{Q^0, Q^1, \dots\}$ for $\forall j \geq k, \delta_j \in [0,1)$ converges to $Q^\star$ as $\lim_{m \to \infty} \prod_{j=k}^{m} \delta_j = 0$. $\square$

B. Stability by Design

Even though DQL-MPTCP updates vary between stages, the aforementioned insights from different settings provide useful guidance for understanding and addressing possible instability issues in DQL-MPTCP. With the relevant intuitions from the aforementioned analysis, we next discuss four different possible causes of instability and how we address them by design.

CAUSE 1. Consider a scenario with aggressive learning rate (α is very high), then the term $O\left(\left\| g' - g \right\|^2\right)$ in (A9) is quite large, as a result DQL-MPTCP updates may not correlate well with the Bellman updates thus leading to instability.

CAUSE 2. A scenario when α is small enough for linearization, however, quite large that the $\mathcal{U}$ in (A11.3) start expanding rather then contracting, hence, leading to instability (Theorem IV.3).

CAUSE 3. Too aggressive generalization of Q function due to large off-diagonal matrix of $\mathcal{K}$ also may cause $\mathcal{U}$ to expand (Theorem IV.3).

CAUSE 4. If the distribution for updates are inadequate, the Q-values for missing state action pairs which are computed by general extrapolation, often incur errors. Such errors, which propagate through the Q-values to all other state-action pairs, may lead to unstable learning (Theorem IV.2).

Observe that the three insights from Section IV-A have been exploited for the convergence and stability in the DQL-MPTCP design (Section III-C). Earlier in Section III-C, we adopted smooth learning rates in relation to the Bellman updates for guaranteed contraction and linearization. We have also applied constructive generalization of the Q function and utilized learning from demonstrations to tackle the case of missing values.

V. PERFORMANCE EVALUATION

We have executed an extensive set of experiments to evaluate DQL-MPTCP under several network scenarios. The network settings used in our testbed are discussed next. Thereafter, we present our findings and reflections from the obtained results.

We evaluate and compare the performance of our DQL-MPTCP algorithm with the standard MPTCP algorithms whose implementation code is accessible: LIA [10], BALIA [20], OLIA [11]. We have used MPTCP v0.93 implemented in Linux.[12] We explain the findings from four different experi-
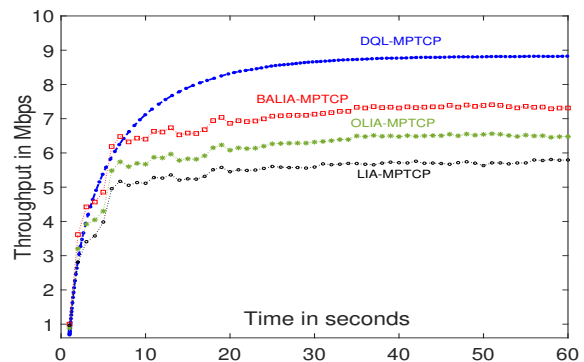
[12]*https://www.multipath-tcp.org*

Fig. 3. Experiment I: For both paths, we set equal delay (50ms), equal 3% path loss and equal bandwidth (10 Mbps each) while downloading replicas of 600MB file. We perform different runs for *DQL, BALIA, OLIA, LIA* with same setup.

ments in the following. Table II describes the performance metrics under consideration for the four experiments.

Our testbed consists of five laptops, four of them are users and one acts as a Server (running MPTCP Source). All four user laptops are connected with a Switch (Gigabit) consisting of two separate interfaces (Gigabit Ethernet) creating two different connections, always used for downloading files from the server laptop. For flexibility, quite similar to [20], every MPTCP connection handles two different subflows in our experiments. We employed and enhanced Linux traffic control facilities[13] and conduct experiments in the controlled environment not only by varying bandwidth, but also under different asymmetric path loss and delays. Each of our experiments consists of eight MPTCP connections (with two subflows) connecting the server with users. We analyze the results based on several data traffic dynamics upon downloading files of variable size (3MB to 500MB) from the server. The aggregate throughputs are extracted from tcpdumps, averaged approximately 900 different experiment runs and numbers in the figures represent the average for a connection (unless stated otherwise).

A. Results and Findings

Our experiments and findings are discussed as follows. Table II describes the organization of four different experiments.

1) Experiment I: Our testbed is set at equal delay of 50 ms, equal bandwidths 20 Mbps and the same 3% packet loss probability to quantify the impact on moving-average throughputs for 60 seconds upon downloads. We use the same file of size (30MB) for all MPTCP algorithms (several experimental runs with *DQL, BALIA, OLIA, LIA* separately using the same setup). Figure 3 shows the effect on the running average throughputs.

FINDING 1: In terms of perceived throughputs, for small size file transfers and short flows, our DQL-MPTCP outperforms the standard MPTCP algorithms under realistic channel conditions.

Remark V.1. FINDING 1 is the joint effect of coupled congestion control and scheduling policy adopted in this study. Note that the standard algorithms, LIA, OLIA, and BALIA inherent separately designed congestion control and scheduling.

[13]NetEm:https://wiki.linuxfoundation.org/networking/netem

TABLE II
METRICS USED FOR NETWORK EMULATION EXPERIMENTS

| Experiment ID	Metrics used in the emulated Network Settings and Scenario	Remarks
I	Moving Average Throughputs: Arithmetic mean of a set of previous throughputs until t seconds	Compare DQL-MPTCP with LIA, OLIA, BALIA
II	Throughputs vs. Fluctuating Delay: Variation in throughputs with increasing one path delays	How DQL-MPTCP reacts to fluctuating delays in one of the path
III	Throughputs vs. Varying Bandwidth: Average throughputs with increasing one path capacities	How DQL-MPTCP performs with respect to varying capacities
IV	TCP Friendliness : Ability of a new protocol to behave under congestion like the TCP protocol	Degradation of DQL-MPTCP when TCP competes for bandwidth
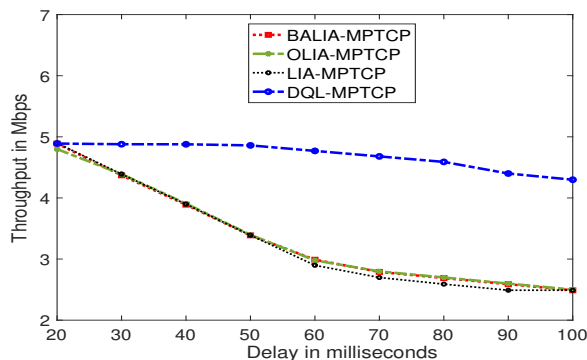
Fig. 4. Experiment II: Change the delay of one path gradually after every 20 seconds (increasing by 10 ms) starting from 20 ms to 100 ms; while downloading a large file. Delay over another path is set constant (20 ms). Both paths are set at same bandwidths (10 Mbps each) and suffers from the same path loss of 5%.
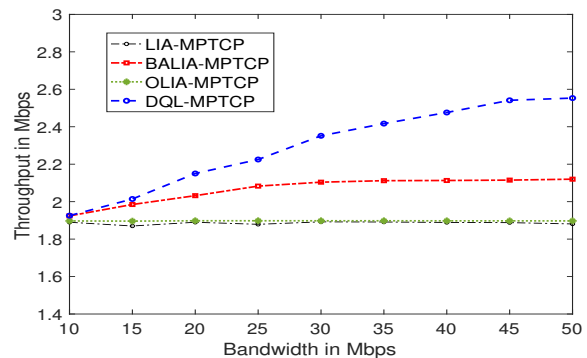
Fig. 5. Experiment III: We vary bandwidths of one path from 10 Mbps to 50 Mbps dynamically fixed another path at 10 Mbps). We set equal delay (150 ms) and same path loss (4%) for both paths to observe the impact of time-varying capacities on performance as perceived by all MPTCP algorithms.

2) Experiment II: For both paths, we set equal bandwidth (20 Mbps) and equal packet loss probabilities of 5% to see the impact on the aggregate throughput. One path delay increases from 20 ms to 100 ms after every 20 seconds, but the delay of another path is fixed at 20 ms.

Note that the users are downloading the same file of size 600 MB but using different MPTCP algorithms (different experimental runs with different versions of MPTCP: *DQL, BALIA, OLIA, LIA*) while the network setup is exactly the same. Figure 4 depicts the aggregate impact on throughputs by varying the delay of one of the path dynamically, which provides us another interesting finding.

FINDING 2: By increasing one of the path delay while downloading, the aggregate throughputs obtained by all MPTCP algorithms decreases. This behavior is as expected, but, the decreasing rate of throughputs is quite dissimilar with different MPTCP algorithms. In particular, the slope is quite high with all other MPTCP algorithms but is smallest with DQL-MPTCP.

Remark V.2. FINDING 2 is due to the delayed feedback of standard MPTCP algorithms and time to discover and adapt with the temporal changes in path delay. As a consequence, most MPTCP algorithms need retransmission and recovery of lost packets. This may also require re-sequencing of received packets at the user end (all packets need to be delivered in order). However, DQL-MPTCP is robust to delay variations as the DQL agent is learning packet scheduling and congestion control continuously. Further, the improved responsiveness is also due to the use of prioritized experience replay buffer,

which accelerates learning and adapts packet rates and schedule across paths quickly. As a result, with these time-varying delay changes, the reordering delay and waiting time for packets transmitted over the slow path, compared to standard MPTCPs, is minimal with DQL-MPTCP.

3) Experiment III: For both paths, we fixed the same delays of 150 ms and 4% packet loss probability, equal for both paths. Then, we vary bandwidths dynamically for only one of the paths (other path bandwidth is constant 20 Mbps) to see the impact on the average throughputs. All MPTCPs are downloading same size ($600MB$) file over the same testbed setup. Figure 5 illustrates the impact on throughputs.

FINDING 3: DQL-MPTCP obtains the highest throughput. All other state-of-the-art standard MPTCP algorithms fall behind in utilizing dynamically varying bandwidths.

Figure 5 shows that throughputs obtained by standard MPTCP algorithms are fairly low. DQL-MPTCP algorithm outperforms others in the same testbed setting. This is due to the fact that under realistic channel error, the standard MPTCP algorithms are guided by a predefined fixed loss and delay based policy.

Remark V.3. With the enhanced continuous learning using optimal packet scheduling action of DQL-MPTCP connections, they utilize increasing bandwidth pretty well even under significant magnitude of path loss. More importantly, the underlying LSTM inside the DQL agent of our MPTCP can quickly discover network bandwidth changes and intelligently perform balancing with joint packet scheduling and congestion control. Design of MPTCP algorithm requires balancing the
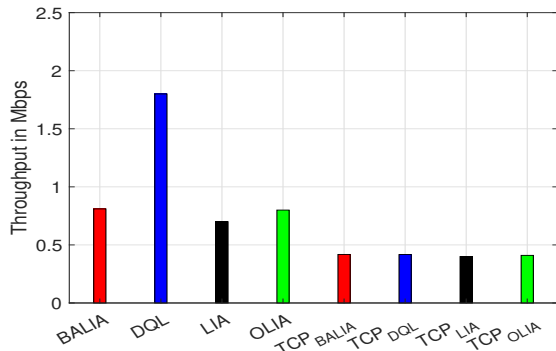
Fig. 6. Experiment IV setting: MPTCP flows contending with single path TCP. Using the setup of Experiment III, we conduct experiment with five MPTCP connections sharing one of the path with rival five single path TCP connections. We repeat same experimental runs using all four different versions of MPTCP: DQL, BALIA, OLIA and LIA algorithms.

tradeoff between fairness to single path TCP, when sharing a bottleneck, and responsiveness in the sense that MPTCP grabs available bandwidth when it becomes available in any of the available subflows [10], [20].

As noted in earlier experiments, we observed considerably high responsiveness of DQL-MPTCP. Our next experiment is to study the fairness of DQL-MPTCP to single path TCP.

4) Experiment IV: Figure 6 illustrates the throughputs comparisons for the TCP and MPTCP connections, when we set single path TCP flows competing with MPTCP connection in one of the paths.

Observe in Figure 6 that DQL-MPTCP is fair to rival single path TCP connections in terms of throughputs. It shows that the throughput perceived by single path TCP is quite close to that when competing with standard MPTCP algorithms; however, the DQL-MPTCP connection achieves higher throughputs. As compared to BALIA vs single path TCP, the throughput obtained by the single path TCP when competing with DQL-MPTCP subflows is negligibly smaller. The reason for this observation (negligibly smaller single path TCP throughput when competing with DQL-MPTCP) requires further investigations.

VI. FUTURE DIRECTIONS

With relevant insights from this work, we have identified the following three research problems for future directions.

A. Cost assessment Study

The cost assessment study of DQL-MPTCP with relevant insights from this work is another exciting research direction. At a higher layer, we can design a high-level policy framework i) to consider efficient energy consumption and economic data plans (e.g. Cellular vs. WiFi or number of interfaces in 5G and beyond) ii) to determine a set of suitable paths based on cost/energy constraints. Such high-level policy engine could determine the (cost/energy) optimal set of paths.

B. A hybrid (model- and DQL-based) design

Model-based designs such as LIA, OLIA, BALIA typically have very low computing overhead, as they adopt a fine-grain scale for their control loops and respond to every ack. In sharp contrast, DQL-based designs have potentially high computing overheads as their feedback loops are longer than the model-based designs.

To investigate the overhead cost of DQL-MPTCP, we send traffic from a server to the client (with two 10 Mbps links and 20ms delay) for 200 seconds. We measure the average CPU use on the MPTCP source side to examine learning overhead and compare DQL-MPTCP with other models. We repeat the same experiment and record iPerf's CPU utilization.

We remove the first few seconds of the experiment for all designs in order to have a reasonable comparison and limit the effect of the startup phases of different MPTCP designs. Figure 7 shows results. DQL-MPTCP, as expected, has a high overhead compared to other model-based designs (but lower than the DQN-MPTCP [36]). Although the current DQL-MPTCP user-space implementation has a high overhead, we think the final optimized kernel version would lower overhead.

This experiment provides new insights and an essential future direction towards a pragmatic and evolutionary strategy of rethinking a hybrid MPTCP (by blending model-based techniques with advanced DQL-based approaches).

C. Minimizing Training overhead

We trained the DQL agent for over twenty-five thousand transition samples using iPerf3 (*https://ipref.fr*), which generates packets continuously to keep the paths always busy in our network settings. This is useful to generate sufficient training samples. We perform experiments with different settings under variable packet loss, delays and bandwidths, hence our (offline) training time varies (2 to 10 hours). In one of the settings with two paths, the training runs for 3 hours when we set equal bandwidths 10Mbps each, equal RTTs 550ms each and equal packet loss probabilities 0.5% for each path. Nevertheless, the outcome is beneficial: the offline training makes our algorithm fully prepared for online deployment by minimizing the startup delay.
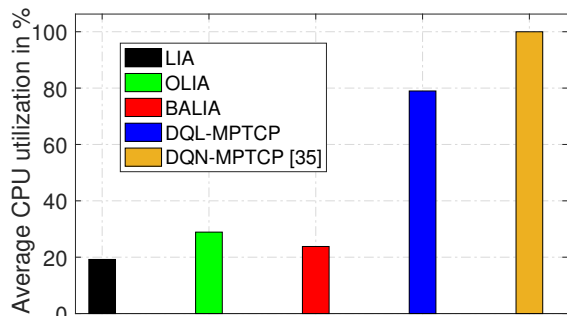

Fig. 7. We send traffic from a server to the client (with two 10 Mbps links and 20ms delay) for 200 seconds. We measure the average CPU use on the MPTCP source side to examine learning overhead and compare DQL-MPTCP with other models. We repeat the same experiment and record iPerf's CPU utilization. We repeat same experimental runs and measure the average % utilization of the five different versions of MPTCP [36].

It is worth noting that offline training is a one time job. Recall that we have used Deep Neural Networks (DNNs) for inference in our DQL-MPTCP implementation, each of which owns only two (hidden) layers. The online inference time observed during our experimentation is just about 0.8ms (which causes smallish overhead for real-time decision process). Retraining our DQL agent is needed only when the network settings changes drastically (e.g., from *high-bandwidth-delay* to *low-bandwidth-delay*). The cause for this is to collect sufficient transition samples needed to update (and learn) the DNNs and gain sufficient experience to make better decisions when comparable conditions occur in a new network environment. An effective way to conduct retraining of a trained DQL agent to adapt for the new network services, such as ultra reliable low latency communication, will benefits from transfer learning [48] and needs further investigations [51].

VII. CONCLUSIONS

We developed a novel DQL-based multipath congestion control and packet scheduling scheme using policy gradients and prioritized replay buffer for the future Internet and wireless technologies. It utilizes a DQL agent with policy gradients to jointly perform dynamic packet scheduling and congestion control, and conduct network reward maximization for all active subflows of an MPTCP connection. We provided stability analysis of DQL-MPTCP with relevant practical design insights. Our design consists of training intelligent agent for interface and actuation, which possesses a flexible LSTM-based representation network. The developed architecture appears to be capable of i) learning an effective representation of all MPTCP subflows, ii) jointly performing intelligible coupled congestion control and packet scheduling, and iii) dealing with dynamic network characteristics and time-varying flows.

We observed better performance of DQL-MPTCP when comparing with standard MPTCP algorithms under similar network settings. Essentially, what the DQL policy may amount to is that each MPTCP agent either sends or abstains from sending based on a set of observables and the policy attempts to maximize goodput and maintain low delays over the whole network and for all contending agents. Moreover, we anticipate that the DQL should have accelerated MPTCP in discovering the best policy for the underlying partially observable Markov process (of the network dynamics), which requires further investigations and is left for future work.

REFERENCES

[1] S. R. Pokhrel and A. Walid, "Learning to harness bandwidth with multipath congestion control and scheduling," *IEEE Transactions on Mobile Computing*, vol. 0, no. 0, p. early access, 2021.

[2] A. Walid, Q. Peng, J. Hwang, and S. Low, "Balanced Linked Adaptation Congestion Control Algorithm for MPTCP," *Internet Engineering Task Force, Internet-Draft draft-walid-mptcp-congestion-control-04*, 2016.

[3] W. Li *et al.*, "SmartCC: A Reinforcement Learning Approach for Multipath TCP Congestion Control in Heterogeneous Networks," *IEEE Journal on Selected Areas in Communications*, 2019.

[4] P. Goyal, A. Agarwal, R. Netravali, M. Alizadeh, and H. Balakrishnan, "ABC: A Simple Explicit Congestion Control Protocol for Wireless Networks," *arXiv preprint arXiv:1905.03429*, 2019.

[5] F. Chiariotti *et al.*, "Analysis and Design of a Latency Control Protocol for Multi-Path Data Delivery With Pre-Defined QoS Guarantees," *IEEE/ACM Transactions on Networking*, 2019.

[6] S. R. Pokhrel and M. Mandjes, "Improving Multipath TCP Performance over WiFi and Cellular Networks: an Analytical Approach," *IEEE Transactions on Mobile Computing*, 2018.

[7] X. Nie *et al.*, "Dynamic TCP Initial Windows and Congestion Control Schemes through Reinforcement Learning," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1231–1247, 2019.

[8] Z. Xu *et al.*, "Experience-driven Congestion Control: When Multi-Path TCP Meets Deep Reinforcement Learning," *IEEE Journal on Selected Areas in Communications*, 2019.

[9] S. R. Pokhrel and C. Williamson, "A rent-seeking framework for multipath TCP," *ACM SIGMETRICS Performance Evaluation Review*, vol. 48, no. 3, pp. 63–70, 2021.

[10] C. Raiciu, M. Handley, and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols, IETF RFC 6356," 2011.

[11] R. Khalili *et al.*, "MPTCP is Not Pareto-optimal: Performance Issues and a Possible Solution," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, pp. 1651–1665, 2013.

[12] P. Hurtig *et al.*, "Low-latency Scheduling in MPTCP," *IEEE/ACM Transactions on Networking*, vol. 27, no. 1, pp. 302–315, 2018.

[13] A. Garcia-Saavedra, M. Karzand, and D. J. Leith, "Low Delay Random Linear Coding and Scheduling over Multiple Interfaces," *IEEE Transactions on Mobile Computing*, vol. 16, no. 11, pp. 3100–3114, 2017.

[14] Y.-S. Lim, E. M. Nahum, D. Towsley, and R. J. Gibbens, "ECF: An MPTCP Path Scheduler to Manage Heterogeneous Paths," in *Proceedings of the 13th International Conference on emerging Networking Experiments and Technologies*, 2017, pp. 147–159.

[15] J. Luo, X. Su, and B. Liu, "A reinforcement learning approach for multipath TCP data scheduling," in *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, 2019, pp. 0276–0280.

[16] 3GPP, "Study on access traffic steering, switch and splitting support in the 5G system architecture phase 2 (release 17)," *3GPP TS23.700-93, V17.0.0*, 2021.

[17] M. Simon, E. Kofi, L. Libin, and M. Aitken, "ATSC 3.0 broadcast 5G unicast heterogeneous network converged services starting release 16," *IEEE Transactions on Broadcasting*, vol. 66, no. 2, pp. 449–458, 2020.

[18] O. Bonaventure, M. Boucadair, B. Peirens, S. Seo, and A. Nandugudi, "0-RTT TCP converters," Internet-Draft draft-bonaventure-mptcp-converters-01. IETF Secretariat, Tech. Rep., 2017.

[19] M. Boucadair *et al.*, "An MPTCP option for network-assisted MPTCP deployments: Plain transport mode," *IETF*, 2016.

[20] Q. Peng, A. Walid, J. Hwang, and S. H. Low, "Multipath TCP: Analysis, Design, and Implementation," *IEEE/ACM Transactions on Networking*, pp. 596–609, 2016.

[21] H. Haile, K.-J. Grinnemo, S. Ferlin, P. Hurtig, and A. Brunstrom, "End-to-End Congestion Control Approaches for High Throughput and Low Delay in 4G/5G Cellular Networks," *Computer Networks*, p. 107692, 2020.

[22] C. Lee, J. Jung, and J.-M. Chung, "DEFT: Multipath TCP for High Speed Low Latency Communications in 5G Networks," *IEEE Transactions on Mobile Computing*, 2020.

[23] M. Polese, R. Jana, and M. Zorzi, "TCP in 5G mmWave Networks: Link Level Retransmissions and MP-TCP," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2017, pp. 343–348.

[24] S. R. Pokhrel and J. Choi, "Low-Delay Scheduling for Internet of Vehicles: Load-Balanced Multipath Communication With FEC," *IEEE Transactions on Communications*, vol. 67, no. 12, pp. 8489–8501, 2019.

[25] S. R. Pokhrel *et al.*, "Fair Coexistence of Regular and Multipath TCP over Wireless Last-Miles," *IEEE Transactions on Mobile Computing*, vol. 18, no. 3, pp. 574–587, 2019.

[26] M. Dong *et al.*, "PCC: Re-architecting Congestion Control for Consistent High Performance," in *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, 2015, pp. 395–408.

[27] K. Winstein and H. Balakrishnan, "TCP ex machina: Computer-generated Congestion Control," *ACM SIGCOMM'13*, pp. 123–134, 2013.

[28] Y. Zaki *et al.*, "Adaptive Congestion Control for Unpredictable Cellular Networks," in *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4. ACM, 2015, pp. 509–522.

[29] T. Mai, H. Yao, Y. Jing, X. Xu, X. Wang, and Z. Ji, "Self-learning Congestion Control of MPTCP in Satellites Communications," in *2019 15th International Wireless Communications & Mobile Computing Conference (IWCMC)*. IEEE, 2019, pp. 775–780.

[30] B. Liao, G. Zhang, Z. Diao, and G. Xie, "Precise and Adaptable: Leveraging Deep Reinforcement Learning for GAP-based Multipath

Scheduler," in *2020 IFIP Networking Conference (Networking)*. IEEE, 2020, pp. 154–162.

[31] J. Chung, D. Han, J. Kim, and C.-k. Kim, "Machine Learning Based Path Management for Mobile Devices Over MPTCP," in *2017 IEEE International Conference on Big Data and Smart Computing (BigComp)*. IEEE, 2017, pp. 206–209.

[32] F. Silva, M. A. Togou, and G.-M. Muntean, "AVIRA: Enhanced Multipath for Content-aware Adaptive Virtual Reality," in *2020 International Wireless Communications and Mobile Computing (IWCMC)*. IEEE, 2020, pp. 917–922.

[33] Y. Li *et al.*, "HPCC: High Precision Congestion Control," in *Proceedings of the ACM Special Interest Group on Data Communication*. ACM, 2019, pp. 44–58.

[34] M. Dong *et al.*, "PCC Vivace: Online-Learning Congestion Control," in *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, 2018, pp. 343–356.

[35] T. P. Lillicrap *et al.*, "Continuous Control With Deep Reinforcement Learning," Jan. 26 2017, US Patent App. 15/217,758.

[36] S. R. Pokhrel and S. Garg, "Multipath Communication With Deep Q-Network For Industry 4.0 Automation and Orchestration," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2020.

[37] T. Hester *et al.*, "Deep Q-learning From Demonstrations," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[38] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized Experience Replay," *arXiv preprint arXiv:1511.05952*, 2015.

[39] S. Ferlin, H. Kucera, Claussen, and O. Alay, "MPTCP meets FEC: Supporting latency-sensitive applications over heterogeneous networks," *IEEE/ACM Transactions on Networking*, vol. 26, no. 5, pp. 2005–2018, 2018.

[40] T. Haarnoja, A. Zhou, S. Ha, J. Tan, G. Tucker, and S. Levine, "Learning to Walk via Deep Reinforcement Learning," *arXiv preprint arXiv:1812.11103*, 2018.

[41] C. J. Watkins and P. Dayan, "Q-Learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.

[42] J. Achiam, E. Knight, and P. Abbeel, "Towards Characterizing Divergence in Deep Q-Learning," *arXiv preprint arXiv:1903.08894*, 2019.

[43] V. Mnih *et al.*, "Human-level Control Through Deep Reinforcement Learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.

[44] S. K. Saha *et al.*, "MuSher: An Agile Multipath-TCP Scheduler for Dual-Band 802.11 ad/ac Wireless LANs," in *The 25th Annual International Conference on Mobile Computing and Networking*, 2019, pp. 1–16.

[45] F. P. Kelly, A. K. Maulloo, and D. K. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research society*, vol. 49, no. 3, pp. 237–252, 1998.

[46] R. Mazumdar, L. G. Mason, and C. Douligeris, "Fairness in Network Optimal Flow Control: Optimality of Product Forms," *IEEE Transactions on communications*, vol. 39, no. 5, pp. 775–782, 1991.

[47] S. R. Pokhrel, "Modeling and performance evaluation of TCP over last-mile wireless networks," Ph.D. dissertation, 2017.

[48] S. R. Pokhrel *et al.*, "Multipath TCP meets transfer learning: A novel edge-based learning for industrial IoT," *IEEE Internet of Things Journal*, 2021.

[49] H. Kumar, A. Koppel, and A. Ribeiro, "On the Sample Complexity of Actor-Critic Method for Reinforcement Learning with Function Approximation," *arXiv preprint arXiv:1910.08412*, 2019.

[50] T. Degris, M. White, and R. S. Sutton, "Off-Policy Actor-Critic," *arXiv preprint arXiv:1205.4839*, 2012.

[51] S. R. Pokhrel *et al.*, "Towards ultra reliable low latency multipath TCP for connected autonomous vehicles," *IEEE Transactions on Vehicular Technology*, 2021.