

Switching in the Rain: Predictive Wireless x-haul Network Reconfiguration

Igor Kadota^{*†}, Dror Jacoby^{*‡}, Hagit Messer[§], Gil Zussman[†], and Jonatan Ostrometzky[‡]

Abstract

Wireless x-haul networks rely on microwave and millimeter-wave links between 4G and/or 5G base-stations to support ultra-high data rate and ultra-low latency. A major challenge associated with these high frequency links is their susceptibility to weather conditions. In particular, precipitation may cause severe signal attenuation, which significantly degrades the network performance. In this paper, we develop a Predictive Network Reconfiguration (PNR) framework that uses historical data to predict the future condition of each link and then prepares the network ahead of time for imminent disturbances. The PNR framework has two components: (i) an Attenuation Prediction (AP) mechanism; and (ii) a Multi-Step Network Reconfiguration (MSNR) algorithm. The AP mechanism employs an encoder-decoder Long Short-Term Memory (LSTM) model to predict the sequence of future attenuation levels of each link. The MSNR algorithm leverages these predictions to dynamically optimize routing and admission control decisions aiming to maximize network utilization, while preserving max-min fairness among the base-stations sharing the network and preventing transient congestion that may be caused by re-routing. We train, validate, and evaluate the PNR framework using a dataset containing over 2 million measurements collected from a real-world city-scale backhaul network. The results show that the framework: (i) predicts attenuation with high accuracy, with an RMSE of less than 0.4 dB for a prediction horizon of 50 seconds; and (ii) can improve the instantaneous network utilization by more than 200% when compared to reactive network reconfiguration algorithms that cannot leverage information about future disturbances

Index Terms

^{*}Both authors contributed equally to this research.

[†]I. Kadota and G. Zussman are with the Department of Electrical Engineering, Columbia University, New York, USA.

[‡]D. Jacoby and J. Ostrometzky are with the Faculty of Engineering, Tel Aviv University, Tel Aviv, Israel.

[§]H. Messer is with the School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel.




Fig. 1. (a) A wireless backhaul network in Gothenburg, Sweden (the map area is of approximately $10 \times 10 \text{ km}^2$). The data utilized in this paper was collected from this network by *Ericsson AB*. (b) An abstraction of the network topology (described in Sec. II).

Wireless Networks, Millimeter-Wave, Backhaul, 5G, Routing, Machine Learning, Rain Attenuation

I. INTRODUCTION

4G and 5G networks often use high bandwidth microwave and millimeter-wave (mmWave) links in their fronthaul, midhaul, and backhaul (x-haul) networks [1] for supporting applications that require high data rate and ultra-low latency. These wireless x-haul networks can connect a large number of base-stations, covering entire cities, as depicted in Fig. 1(a). A main challenge of using microwave and mmWave links is their high susceptibility to weather conditions. The signal attenuation due to different atmospheric and weather phenomena is described by the International Telecommunication Union (ITU) in [2]–[5] and depicted in Fig. 2. It can be seen that, apart from the oxygen resonance frequency at 60 GHz, the dominant factor affecting link attenuation is precipitation. This implies that signal attenuation may vary significantly over time and over geographic locations. *Hence, the need for a high bandwidth wireless x-haul that is robust to variations in the network conditions calls for the development of a predictive network reconfiguration framework that can dynamically allocate resources based on current and future estimated network conditions.*

Until recently, only local Physical/Link layer mechanisms were employed to alleviate the impact of the time-varying conditions of the links on the network performance. For example, the Automatic Transmit Power Control is a commonly used mechanism that adjusts the transmitter power based on measurements of the link attenuation [6]. However, with the emergence of Software-Defined Networking (SDN) [7]–[9], it is now possible to develop global Network




Fig. 2. Signal attenuation (in dB/km) for various atmospheric phenomena as a function of frequency [2]. The commonly used bands of 6–40 GHz (K-band) and 60–90 GHz (E-band) are highlighted.

layer mechanisms (such as NEC’s backhaul solution in [10]) that monitor the entire network and react to performance drops caused by weather-induced disturbances. A main drawback of *reactive reconfiguration* mechanisms is their delay in recovering from performance drops, which may severely affect time-sensitive applications. To overcome this challenge, *predictive reconfiguration* mechanisms can be employed.

Prior work on predictive network reconfiguration algorithms (see [11] for a survey) focused mainly on alleviating the effects of node mobility [12]–[17], traffic demand variability [18]–[25], and link quality degradation due to multi-path reflection, line-of-sight occlusion, and interference [26]–[33]. Weather effects pose fundamentally different challenges. In particular, weather-induced attenuation can be severe, affect large contiguous geographic areas, and last for extended periods of time. The literature on the prediction of microwave and mmWave signal attenuation due to weather conditions uses meteorological data (e.g., weather-radar echo measurements) to predict the current/future attenuation levels [34], [35] or uses past attenuation measurements to predict future attenuation levels [36]–[40]. Most relevant to this paper is our prior work in [37] which employs an encoder-decoder LSTM model to predict future link attenuation levels. The main drawbacks of [37] are that: (i) its prediction mechanism does not capture the significant spatial correlation of the rain-induced attenuation; and (ii) its prediction mechanism is not employed to inform any algorithm or protocol.

The literature on predictive weather-aware reconfiguration algorithms contains only a few works [34]–[36], [41]. Most of these works, in particular [34], [35], [41], develop modifications to standard *distributed* routing protocols such as Open Shortest Path First (OSPF) which due to the

lack of centralized coordination may converge slowly, making them unsuitable for networks that support time-sensitive applications. Only [36] leverages SDN to perform centralized predictive network-wide reconfiguration. The framework proposed in [36] predicts future link attenuation levels using a model *specific to rain fading* and then computes current and future routing decisions aiming to maximize throughput. Some limitations of the solution proposed in [36] are that: (i) its attenuation prediction mechanism does not capture the spatial correlation of the weather-effects; (ii) its prediction mechanism can *only* be employed during periods of rain; (iii) its network reconfiguration mechanism *allows transient link congestion* (i.e., it allows flows to temporarily exceed the link capacity); (iv) its network reconfiguration mechanism does not take fairness into account; and (v) its re-routing mechanism does not support flow splitting.

Our contributions: In this paper, we develop and evaluate, based on a real dataset, a Predictive Network Reconfiguration (PNR) framework that leverages existing local Physical/Link layer mechanisms and adds two new components: an Attenuation Prediction (AP) mechanism; and a Multi-Step Network Reconfiguration (MSNR) algorithm.

The AP mechanism employs an encoder-decoder LSTM model to predict the sequence of future attenuation levels based on past measurements, capturing both time and spatial correlation that are typical of weather-effects *without incorporating weather-related models*, which allows it to be used both in dry and rain periods, and *without relying on meteorological data from external sources* such as weather radars. To train, validate, and evaluate the AP mechanism, we use *a unique dataset obtained from the real-world city-scale backhaul network in Gothenborg, Sweden* (see Fig. 1(a)) *collected by Ericsson AB*. The dataset contains 2,295,000 measurements of link attenuation. In Fig. 3(a), we display the evolution of the measured attenuation for every link in the backhaul network over a period of 1.9 hours. Notice that in the interval between $t = 300$ and 600 time-steps there is an increased attenuation due to rain. The spatio-temporal correlation is evident. *The AP mechanism leverages this correlation to achieve high prediction accuracy*. In particular, the AP mechanism achieves a Root Mean Square Error (RMSE) of less than 0.4 dB for a prediction horizon of 50 seconds. We evaluate two benchmark time series prediction methods that do not capture the spatial correlation of the weather-effects and show that both of them can perform 30% worse than the AP mechanism in terms of RMSE.

The MSNR algorithm leverages the predictions from the AP mechanism and uses Model Predictive Control (MPC) [42] to compute the sequence of current and future routing and admission control decisions that: (i) maximize network utilization, while (ii) achieving max-min

fairness among the base-stations sharing the network and (iii) preventing transient congestion that may be caused by re-routing. This sequence of routing and admission control decisions are employed by the centralized SDN controller to reconfigure the network over time. For example, based on a prediction that a set of links will become unavailable in 30 seconds, the MSNR algorithm can determine when it is optimal for the SDN controller to redirect flows in order to avoid potential interruptions to service and can decide whether or not it is necessary to revoke network slices from low priority services. An important challenge associated with the MSNR algorithm is computational complexity. In Sec. IV-C, we proposed a principled implementation of the MSNR algorithm which has a computational complexity that grows polynomially with the prediction horizon, as opposed to a naive implementation that can have exponential complexity.

We evaluate the PNR framework using the data collected from the backhaul network. Our results show that the PNR framework can improve the instantaneous network utilization by more than 200% when compared to reactive network reconfiguration algorithms that do not prepare the network for future disturbances. To the best of our knowledge, this is the first attempt to propose and evaluate, based on a real dataset, an integrated framework for x-haul network reconfiguration that leverages the spatio-temporal correlation of the weather-effects to jointly optimize routing and admission control decisions. A patent including some of the results is pending [43].

This paper is organized as follows. Section II describes the network model and the dataset. In Sec. III, we develop the AP mechanism. In Sec. IV, we develop the MSNR algorithm. In Sec. V, we evaluate the performance of the PNR framework. Section VI concludes the paper and discussed future work.

II. PROBLEM FORMULATION AND DATASET

In this section, we present the network model used to develop the PNR framework. We first describe the model in general and then establish the connection between the model and the real-world backhaul network. Let $G = (V, E)$ be the directed graph that represents an x-haul communication network with base-stations, also called nodes, $n \in V = \{1, 2, \dots, N\}$, connected by wireless links $(k, l) \in E$ where $k, l \in V$ and (k, l) represents the link $k \rightarrow l$. Time is divided into time-steps with index $t \in \{1, 2, \dots, T\}$, where T is the time-horizon and the time interval between t and $t + 1$ is $\Delta = 10$ seconds. Let $d_n > 0$ be the demand associated with commodity $n \in V$. The demand d_n represents the uplink traffic that base-station n aggregates from its

associated users. Let $z_{n,t} \in [0, 1]$ be the fraction of the demand d_n admitted during time-step t . It follows that the *admitted demand* from base-station n during time t is given by $z_{n,t}d_n$. For simplicity, we assume that demands d_n remain fixed over time and that node N is the common destination for all commodities $n \in V \setminus N$. Naturally, for the common destination N , we have $d_N = 0$ and $z_{N,t} = 0, \forall t$. Let $f_{n,t}^{(k,l)} \in [0, 1]$ be the fraction of the admitted demand $z_{n,t}d_n$ that flows through link (k, l) during time t . By definition

$$f_{n,t}^{(k,n)} = 0, \forall (k, n) \in E, \forall t ; \quad (1)$$

$$f_{n,t}^{(N,l)} = 0, \forall n \in V, \forall (N, l) \in E, \forall t ; \quad (2)$$

$$f_{n,t}^{(k,l)} = 0, \forall n \in V, \forall (k, l) \notin E, \forall t , \quad (3)$$

where (1) is a constraint on the incoming flows at the source nodes, (2) is a constraint on the outgoing flows at the destination node N , and (3) enforces zero flow on non-existing links. It follows that the *total flow* in link $(k, l) \in E$ during time t is given by $\sum_{n=1}^{N-1} z_{n,t}d_n f_{n,t}^{(k,l)}$.

Feasibility and Fairness. We assume that $G = (V, E)$ and d_n are given and remain fixed over time. We assume that routing and admission control decisions implemented by the centralized SDN controller at time-step t , namely $f_{n,t}^{(k,l)}$ and $z_{n,t}$, respectively, remain fixed in the interval between t and $t + 1$. Routing and admission control decisions at each time-step t are *feasible* when they satisfy flow conservation and capacity constraints. The *flow conservation* associated with commodity $n \in V$ and node $l \in V$ at time t is given by

$$\sum_{k=1}^N f_{n,t}^{(k,l)} - \sum_{m=1}^N f_{n,t}^{(l,m)} = \begin{cases} -1, & l = n \\ +1, & l = N \\ 0, & \text{otherwise} \end{cases} , \quad (4)$$

where $l = n$ indicates that node l is the source of commodity n and $l = N$ indicates that node l is the destination of commodity n . Let $c_t^{(k,l)} \geq 0$ be the capacity of link (k, l) at time t and let $\hat{c}_{t+1}^{(k,l)} \geq 0$ be the *predicted* capacity of link (k, l) at time $t + 1$. Since the exact moment between t and $t + 1$ in which the capacity changes from $c_t^{(k,l)}$ to $\hat{c}_{t+1}^{(k,l)}$ is unknown, we assume the worst-case and represent the capacity in this interval by $\min\{c_t^{(k,l)}, \hat{c}_{t+1}^{(k,l)}\}$. Hence, the *capacity constraint* associated with link $(k, l) \in E$ at time t is given by

$$\sum_{n=1}^{N-1} z_{n,t}d_n f_{n,t}^{(k,l)} \leq \min\{c_t^{(k,l)}, \hat{c}_{t+1}^{(k,l)}\} . \quad (5)$$

Definition 1 (Feasibility): The set of routing and admission control decisions at time t , namely $\{f_{n,t}^{(k,l)}, z_{n,t}\}$, $\forall n \in V$, $\forall (k,l) \in E$, is *feasible* when it satisfies the flow constraints in (1)-(3), the flow conservation in (4) and the capacity constraints in (5).

Definition 2 (Max-Min Fairness): The feasible set $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ at time-step t has admission rates $z_{n,t}$ that are max-min fair if, in order to *maintain feasibility*, an increase of any $z_{n,t}$ necessarily results in the decrease of $z_{m,t}$ of another source m for which $z_{m,t} \leq z_{n,t}$.

The goal of the PNR framework is to dynamically optimize routing and admission control decisions over time, taking into account future predicted network conditions, aiming to maximize the cumulative sum of admission rates $\sum_{t=1}^T \sum_{n=1}^{N-1} z_{n,t}$, while ensuring that, in *each and every time-step* t , the selected feasible set $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ is max-min fair and can be implemented without inducing transient congestion. Recall that transient congestion can cause increased delay which can severely affect time-sensitive applications. This challenging optimization problem and its computational complexity are addressed in Sec. IV.

Real-World Network and Dataset. Consider the backhaul network in Fig. 1(a) composed of 17 wireless links whose lengths vary from 0.6 to 5.9 km and that operate between 18 and 40 GHz. The directed graph $G = (V, E)$ with $N = 13$ nodes in Fig. 1(b) is generated by assuming that *link endpoints* in Fig. 1(a) that are in close proximity (up to 300 m apart) are connected by fiber which is not capacity-limited. Under this assumption¹, a node in $G = (V, E)$ represents one or more neighboring link endpoints in Fig. 1(a).

The backhaul network in Sweden contains a centralized data collection system (described in detail in [44]) that periodically gathers measurements from each link $(k,l) \in E$ in intervals of $\Delta = 10$ seconds. Each measurement in time-step t includes the transmitted and received signal levels (in dB) represented by $P_{Tx,t}^{(k,l)}$ and $P_{Rx,t}^{(k,l)}$, respectively. According to [44], the extra load associated with the transmission of measurements via the backhaul network is insignificant.

In this paper, we consider that, in each time-step t , the following events occur: (i) the data collection system shares the latest measurements with the centralized PNR framework; (ii) the AP mechanism predicts the future attenuation levels $x_t^{(k,l)} = P_{Tx,t}^{(k,l)} - P_{Rx,t}^{(k,l)}$ of every link and the MSNR algorithm generates new routing $f_{n,t}^{(k,l)}$ and admission control $z_{n,t}$ decisions; and (iii) the SDN controller implements the new network configuration by propagating $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ to the corresponding base-stations. We assume that both the transmission of measurements

¹Notice that other assumptions could have been made but they should not affect the generality of the results.

$\{P_{Tx,t}^{(k,l)}, P_{Rx,t}^{(k,l)}\}$ and the propagation of routes and admission control updates $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ utilize a negligible amount of resources from the backhaul network. It is important to emphasize that these control packets are transmitted at most once in every 10 seconds.

To train, validate, and evaluate the PNR framework we use the directed graph $G = (V, E)$ together with a dataset containing 2,295,000 measurements (i.e., 135,000 per link) and a *train-validation-test split of 80-10-10*. The *test* data utilized to evaluate the PNR framework in Sec. V is composed of three sequences of measurements, each containing a period of rain: *Test Seq. I* with 87,890 measurements collected over a period of 14.3 hours on 2015-06-02, *Test Seq. II* with 11,900 measurements collected over a period of 1.5 hours on 2015-05-19, and *Test Seq. III* with 94,690 measurements collected over a period of 15.5 hours on 2015-06-17.

III. ATTENUATION PREDICTION MECHANISM

In this section, we present the AP mechanism which predicts future link attenuation levels based on historical data, capturing both *time* and *spatial* correlation that are typical of weather-effects. Next, we describe the encoder-decoder LSTM model and the training process. In Sec. V-A, we compare the performance of the AP mechanism with two benchmark time series prediction methods.

A. Encoder-Decoder LSTM Model

The encoder-decoder LSTM model is a Recurrent Neural Network designed to address sequence-to-sequence prediction problems such as machine translation, natural language generation, and speech recognition [45]–[47]. The model is composed of two main parts: the *encoder*, which maps the input sequence into a state vector and the *decoder*, which maps the state vector into a sequence of predictions.

The AP mechanism employs the sliding-window method and the encoder-decoder LSTM model illustrated in Figs. 3(a) and 3(b), respectively, to predict the next H attenuation levels based on the previous W measurements. In particular, let $x_t^{(k,l)} = P_{Tx,t}^{(k,l)} - P_{Rx,t}^{(k,l)}$ be the *attenuation measurement* for link $(k, l) \in E$ at time t , and let $\mathbf{x}_t = (x_t^{(k,l)})$ and $\hat{\mathbf{x}}_{t+h} = (\hat{x}_{t+h}^{(k,l)})$ be the vector of *attenuation measurements* and the vector of *h-steps-ahead attenuation predictions* for all links at time t , respectively. In each time-step t , the encoder-decoder LSTM model employs the sequence of measurements in the *input window* $\{\mathbf{x}_{t-W+1}, \mathbf{x}_{t-W+2}, \dots, \mathbf{x}_t\}$ to predict the sequence of attenuation levels in the *prediction window* $\{\hat{\mathbf{x}}_{t+1}, \hat{\mathbf{x}}_{t+2}, \dots, \hat{\mathbf{x}}_{t+H}\}$. Notice that the




Fig. 3. (a) Measured attenuation for all the 17 links in the network in Fig. 1 with time-steps t separated by 10 seconds (adding up to 1.9 hours) together with an illustration of the sliding-window method with *input window* size of W time-steps and *prediction window* size of H time-steps. The different baseline attenuation levels are due to the different characteristics of the links in terms of distance between base-stations and transmission frequency. An increased attenuation due to rain is observed in the interval between $t = 300$ and 600 time-steps. (b) Encoder-decoder LSTM model that employs the last W measurements $\{x_{t-W+1}, x_{t-W+2}, \dots, x_t\}$ from every link in the network (i.e., the input window) to predict the future H attenuation levels $\{\hat{x}_{t+1}, \hat{x}_{t+2}, \dots, \hat{x}_{t+H}\}$ in each link (i.e., the prediction window).

measurements contained in the *input window* allow the encoder-decoder LSTM model to capture the spatio-temporal correlation that is typical of weather-induced attenuation. We employ an input window size of $W = 12$ and a prediction window size of $H = 5$, which corresponds to 120 seconds and 50 seconds, respectively.

We train the encoder-decoder LSTM model to minimize the prediction error. In particular, consider a dataset with a sequence of attenuation measurements in the interval $t \in \{1, \dots, T\}$. The encoder and decoder are jointly trained to minimize the objective function:

$$\mathcal{L}(\Theta) = \sum_{t=1}^{T-H} \sum_{h=1}^H \|x_{t+h} - \hat{x}_{t+h}\|^2 \quad (6)$$

where $\|\cdot\|$ represents the Euclidean norm, and Θ represents the parameters of the encoder-decoder LSTM model, i.e., weights and biases. We implement the encoder and the decoder LSTM with one hidden layer containing 128 units. We use the dataset collected from the backhaul network to train, tune, and evaluate the AP mechanism. We train the AP mechanism using Backpropagation Through Time [48] and Adaptive Moment Estimation (Adam) [49] with a batch size of 150. The prediction accuracy of the AP mechanism is evaluated in Sec. V-A.

IV. MULTI-STEP NETWORK RECONFIGURATION ALGORITHM

SDN enables the design of algorithms that dynamically reconfigure the entire network. Building on that, in this section, we develop the MSNR algorithm, which leverages information about

links' future conditions to compute the sequence of current and future routing and admission control decisions that attempt to maximize network utilization, while achieving max-min fairness (in every time-step t) among the base-stations sharing the network and preventing transient congestion that may be caused by re-routing. Hereafter, we denote this sequence of routing and admission control decisions as the *optimal sequence of network configurations*.

The problem of finding the optimal sequence of network configurations is a generalization of the well-known Maximum Concurrent Flow (MCF) problem [50], [51] for the more challenging setting where: (i) a sequence of predictions of future network conditions are available and (ii) transient congestion due to re-routing is taken into account. The MSNR algorithm employs MPC to address this generalized MCF optimization problem. In particular, in each time-step t , the MSNR algorithm uses its knowledge of future (predicted) network conditions to evaluate and compare the performance of different congestion-free sequences of network configurations $\{f_{n,t+h}^{(k,l)}, z_{n,t+h}\}, \forall n \in V, \forall (k,l) \in E, \forall h \in \{0, 1, \dots, H-1\}$ and then it selects the max-min fair sequence that maximizes the cumulative sum of admission rates $\sum_{h=0}^{H-1} \sum_{n=1}^{N-1} z_{n,t+h}$. The SDN controller implements the first configuration in the selected sequence, i.e., the configuration $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ associated with the current time t . This iterative process allows the SDN controller to account for future predicted network conditions when optimizing the current network configuration.

An important challenge associated with the MSNR algorithm is computational complexity. A naive implementation of the MSNR algorithm computes and compares the performance of all possible sequences of network configurations within the prediction window $\{t, \dots, t+H\}$. The number of such sequences grows exponentially with H , as we will discuss in Sec. IV-C, which could render the MSNR algorithm impractical. To overcome this challenge, we develop a principled implementation of the MSNR algorithm which employs the structure of the optimization problem to recursively explore the space of all possible sequences of network configurations. This recursive method reduces the complexity from exponential $O(2^H)$ to polynomial $O(H^4)$.

Prior to introducing the MSNR algorithm, we describe: (i) the adaptive modulation mechanism in [52], which is a Physical layer mechanism employed by the backhaul network in Sweden to maximize link capacity over time; and (ii) the SWAN mechanism developed in [53], which is a Network layer mechanisms that eliminates transient congestion that may be caused by re-routing. The MSNR algorithm builds upon both these existing solutions to address the generalized MCF problem, enabling the optimization of routing and admission control decisions over time in a

setting where predictions of future network conditions are available.

A. Adaptive Modulation Mechanism

Three parameters that can be dynamically adjusted to compensate for high attenuation levels in microwave and mmWave links are: the transmission power, the coding rate, and the modulation scheme. The dataset utilized in this paper was collected for a backhaul network that uses radios similar to the ones described in [44], [52] which: (i) employ a constant transmit power $P_{Tx,t}^{(k,l)}$ and a constant coding rate over time; (ii) use Quadrature Amplitude Modulation (QAM) with *adaptive* constellation size M ; and (iii) use a fixed channel bandwidth of 28 MHz that achieves a capacity of 45 Mbps when $M = 4$. Recall that when M is increased by a factor of k , the capacity $c_t^{(k,l)}$ increases by a factor of $\log_2 k$ and the Bit Error Rate (BER) decreases according to [54, Eq. (18)].

The adaptive modulation (AM) mechanism adjusts the constellation size M over time, aiming to maximize link capacity $c_t^{(k,l)}$ while keeping the BER above a given threshold. For complying with the description of the radios in [52, Sec. II.B], hereafter in this paper, we consider a wireless x-haul network that employs the AM mechanism with hysteresis represented in Table I. In particular, we consider that every link $(k, l) \in E$ uses radios that adapt their constellation size M at each time-step t based on Table I and on their measured received signal level $P_{Rx,t}^{(k,l)}$. The *limit up* in Table I represents the received signal level in which the adopted M should increase. The *limit down* represents the received signal level in which the adopted M should decrease to keep the BER above the set threshold. Notice that Table I represents a mapping from the evolution of the received signal levels $P_{Rx,t}^{(k,l)}$ over time to the evolution of the link capacities $c_t^{(k,l)}$ over time.

B. The Cost of Re-routing

One possible approach to dynamically optimizing the network configuration *without resorting to predictions of links' future conditions* is for the SDN controller to carry out, in each time t , the following procedure: (i) gather information about the current link capacities $c_t^{(k,l)}$; (ii) employ existing solutions to the MCF optimization problem (e.g., [50], [51]) to find the configuration $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ that maximizes the *current* network utilization; and (iii) implement the new routing decisions $f_{n,t}^{(k,l)}$ and admission rates $z_{n,t}$ by sending control packets to the base-stations in the x-haul network. Upon reception of these control packets, the base-stations add/remove entries from

TABLE I
PARAMETERS ASSOCIATED WITH THE ADAPTIVE MODULATION MECHANISM WITH HYSTERESIS FOR A BER THRESHOLD OF 10^{-9} .

M	Bitrate (Mbps)	Limit up (dBm)	Limit down (dBm)
4	45	-72	N/A
16	90	-66	-74
64	135	-62.5	-68
128	157	-61	-64
256	180	-57	-62
512	202.5	-53	-58
1024	225	N/A	-54

their routing tables and adjust their network slice admission and provisioning accordingly. Two important drawbacks of this approach are the delay to recover from performance drops, which is characteristic of reactive reconfiguration mechanisms, and that it does not take into account the transient congestion that may be caused by re-routing. Both drawbacks may severely affect time-sensitive traffic. The MSNR algorithm proposed in Sec. IV-C addresses both drawbacks. In this section, we discuss the negative effects that re-routing may have on the network performance.

To update routes from $f_{n,t-1}^{(k,l)}$ to $f_{n,t}^{(k,l)}$, the SDN controller may have to send control packets to multiple base-stations. Due to communication and processing delays, some base-stations may apply the new routes $f_{n,t}^{(k,l)}$ while others still employ old routes $f_{n,t-1}^{(k,l)}$, which may cause significant transient congestion and over-utilization of communication links, namely violation of the capacity constraints in (5). Depending on the duration and magnitude of the congestion, data packets may be severely delayed or even lost. In this case, the re-routing process is clearly imposing a performance cost that should be taken into account when the SDN controller decides whether or not to re-route.

In order to reduce the transient congestion associated with re-routing, a common approach (e.g., [36], [53], [55]) is to subdivide the re-routing process into multiple stages. In each stage, the SDN controller updates a small number of base-stations, instead of updating all of them at the same time. Each stage is designed to generate zero (or little) transient congestion and the complete sequence of stages is designed to lead to the desired final routing configuration.

An important constraint is that the time for completing the re-routing process should be shorter than the interval between two consecutive time-steps, e.g., t and $t + 1$, which in this paper is of $\Delta = 10$ seconds. In [36], [53], [55], the authors propose different route implementation systems that attempt to minimize the transient congestion. Yet, these route implementation systems can only guarantee that re-routing is performed with zero congestion when a portion of the network capacity is vacant before the update. Naturally, when all links are fully utilized, the first update to take effect will always congest at least one link.

In this paper, we consider an SDN controller that implements any given set of new routes $f_{n,t}^{(k,l)}$ by employing the *SWAN mechanism* developed in [53]. The SWAN mechanism leverages *scratch capacity* in *every link* to perform congestion-free re-routing. In particular, the authors of [53] show that SWAN can update routes, i.e., change from $f_{n,t-1}^{(k,l)}$ to any given $f_{n,t}^{(k,l)}$, with zero transient congestion in *at most* $\lceil 1/s_t \rceil - 1$ stages, where $s_t \in (0, 1]$ represents the *scratch capacity of the network* at time t . Formally, s_t is given by

$$s_t = \operatorname{argmax} \left\{ s \in (0, 1] \left| \sum_{n=1}^{N-1} z_{n,t-1} d_n f_{n,t-1}^{(k,l)} \leq (1-s)c_t^{(k,l)}, \forall (k,l) \in E \right. \right\}. \quad (7)$$

For details on how the SWAN route implementation mechanism works, we refer the reader to [53]. Notice that when the network has no *scratch capacity*, i.e., $s_t \rightarrow 0$, the SWAN mechanism needs $\lceil 1/s_t \rceil - 1 \rightarrow \infty$ stages to complete a single congestion-free re-routing process. To limit the re-routing time, we impose a lower bound of $s_{min} = 0.05$ on the *scratch capacity*, s_t , needed for a re-route. Hereafter in this paper, we assume that the SDN controller is allowed to re-route at time t if and only if $s_t \geq s_{min} = 0.05$.

The SDN controller employs the MSNR algorithm to compute the optimal sequence of network configurations over time and, when necessary, it employs the SWAN mechanism to implement new routes. In particular, in each time-step t , given the prior routing and admission control decisions, $\{f_{n,t-1}^{(k,l)}, z_{n,t-1}\}$, the SDN controller employs (7) to calculate the current scratch capacity s_t . If $s_t \geq s_{min}$, the SDN controller employs the MSNR algorithm to compute the optimal sequence of network configurations and then it employs SWAN to implement the optimal configuration $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ at the current time t . Alternatively, if $s_t < s_{min}$, the SDN controller is not allowed to re-route at time t , but it can still optimize the admission rates $z_{n,t}$. In this case, the SDN controller employs the MSNR algorithm *with fixed routing parameters* $f_{n,t}^{(k,l)} = f_{n,t-1}^{(k,l)}$ to compute the optimal sequence of network configurations and then it implements the optimal admission rates $z_{n,t}$ at time t . It is easy to see that admission rates can be updated from $z_{n,t-1}$

to $z_{n,t}$ with zero transient congestion in at most two stages, irrespective of the value of s_t . In the first stage, the SDN controller updates all base-stations in which $z_{n,t-1} > z_{n,t}$ and, in the last stage, the SDN controller updates all base-stations in which $z_{n,t-1} < z_{n,t}$.

The routing and admission control decisions *at time t* determine the scratch capacity s_{t+1} *at time $t + 1$* , which determines whether or not the SDN controller will be allowed to re-route at time $t + 1$. Hence, if the SDN controller plans to re-route at time $t + 1$, it should select a network configuration $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ that will lead to $s_{t+1} \geq s_{min}$. This can be achieved by employing, at time t , the following capacity constraint for every link $(k, l) \in E$

$$\sum_{n=1}^{N-1} z_{n,t} d_n f_{n,t}^{(k,l)} \leq \min\{c_t^{(k,l)}, (1 - s_{min})\hat{c}_{t+1}^{(k,l)}\}. \quad (8)$$

Alternatively, if the SDN controller plans to keep the same routes in the next time-step, i.e., $f_{n,t+1}^{(k,l)} = f_{n,t}^{(k,l)}$, it should attempt to fully utilize the links, leaving no scratch capacity. This can be achieved by employing the capacity constraints in (5). Intuitively, this means that, in order to re-route in the next time-step $t + 1$, the SDN controller may need to reduce the admission rates $z_{n,t}$ in the current time-step t . *This potential reduction of $z_{n,t}$ represents the cost of re-routing*, as illustrated in the following example.

Example: consider the network in Fig. 4 with $N = 3$ nodes, three links $\{(1, 2), (2, 3), (1, 3)\}$, and fixed demands $d_1 = 1$ and $d_2 = 0.5$. Assume that this network has capacities $c_t^{(k,l)} = 0.5$ for all links and predicted capacities $\hat{c}_{t+h}^{(k,l)} = 0.5$ for all links and prediction horizons h . Moreover, assume that $s_t \geq s_{min} = 0.05$, meaning that the SDN controller is allowed to re-route at the current decision time t .

Plan to not re-route. If the SDN controller plans to keep the same routes in future time-steps, then it adopts the capacity constraints in (5). In this toy example, it is easy to see that the corresponding max-min fair admission rates are $z_{1,t} = z_{2,t} = 2/3$. Notice that there exists feasible configurations $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ with higher sum $\sum_{n=1}^2 z_{n,t}$, but their admission rates are not max-min fair. An example of such unfair feasible admission rates are $z_{1,t} = 0.5$, $z_{2,t} = 1$.

Plan to re-route. Alternatively, if the SDN controller plans to re-route in the next time-step, then it adopts the capacity constraints in (8) with $s_{min} = 0.05$. It is easy to see that the corresponding max-min fair admission rates are $z_{1,t} = z_{2,t} = 2/3 * (1 - 0.05)$.

Two important observations are: (i) Planning to re-route at time $t + 1$ does not guarantee that the SDN controller will be able to re-route at time $t + 1$. In particular, if the capacity prediction is inaccurate and (by chance) $\hat{c}_{t+1}^{(k,l)} > c_{t+1}^{(k,l)}$, the SDN controller may not have enough scratch




Fig. 4. Illustration of a network with $N = 3$ nodes (two commodities and a destination) and three links. The admitted demands $z_{n,t}d_n$ at time-step t are shown within the corresponding nodes. The total flows and capacities at time-step t are shown next to the corresponding links.

capacity at time $t + 1$ to re-route. (ii) Planning to re-route at time $t + 1$, can only hurt the network performance at the current time t due to the provision of the scratch capacity, as illustrated in the example. The potential benefits of planning to re-route at time $t + 1$ can only be assessed by computing the performance of the network at future time-steps.

C. Optimal Sequence of Network Configurations

In this section, we develop the MSNR algorithm which leverages information about current and future predicted link capacities $\{c_t^{(k,l)}, \hat{c}_{t+1}^{(k,l)}, \dots, \hat{c}_{t+H}^{(k,l)}\}$ to dynamically optimize routing and admission control decisions aiming to maximize the cumulative sum of admission rates $\sum_{t=1}^T \sum_{n=1}^{N-1} z_{n,t}$, while ensuring that, *in every time-step t* , the selected feasible set $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ is max-min fair and can be implemented by the SDN controller without inducing transient congestion. The MSNR algorithm addresses a generalization of the MCF problem [50], [51] for the more challenging setting where: (i) a sequence of predictions of future network conditions are available and (ii) transient congestion due to re-routing is taken into account.

Prior to describing the MSNR algorithm, we introduce the concept of a re-routing plan. For a given time t and a prediction window size H , let $r_{t,h}$ be an indicator function that is equal to 1, if the plan is to re-route in time-step $t + h$, $\forall h \in \{0, 1, \dots, H\}$, and $r_{t,h} = 0$, otherwise. The *re-routing plan* at time t is given by the vector $\mathbf{r}_t = (r_{t,0}, r_{t,1}, \dots, r_{t,H})$. Notice that if $s_t < s_{min}$, then $r_{t,0} = 0$ and if $s_t \geq s_{min}$, then $r_{t,0} \in \{0, 1\}$. Next, we use an example to show how the re-routing plan $\mathbf{r}_t$ can be utilized to separate the problem of finding the optimal sequence of network configurations into simpler sub-problems.

Example: consider a network with a prediction window size of $H = 5$ and a plan $\mathbf{r}_t = (0, 1, 0, 0, 1, 0)$ to re-route only at times $t + 1$ and $t + 4$. The network parameters associated

TABLE II
EVOLUTION OF NETWORK PARAMETERS ASSOCIATED WITH THE RE-ROUTING PLAN $\mathbf{r}_t = (0, 1, 0, 0, 1, 0)$.

plan	capacity constraints in (9)	admis.	routing
$r_{t,0} = 0$	$\min\{c_t^{(k,l)}, (1 - s_{min})\hat{c}_{t+1}^{(k,l)}\}$	$z_{n,t}$	$f_{n,t-1}^{(k,l)}$
$r_{t,1} = 1$	$\min\{\hat{c}_{t+1}^{(k,l)}, \hat{c}_{t+2}^{(k,l)}\}$	$z_{n,t+1}$	$f_{n,t+1}^{(k,l)}$
$r_{t,2} = 0$	$\min\{\hat{c}_{t+2}^{(k,l)}, \hat{c}_{t+3}^{(k,l)}\}$	$z_{n,t+2}$	$f_{n,t+1}^{(k,l)}$
$r_{t,3} = 0$	$\min\{\hat{c}_{t+3}^{(k,l)}, (1 - s_{min})\hat{c}_{t+4}^{(k,l)}\}$	$z_{n,t+3}$	$f_{n,t+1}^{(k,l)}$
$r_{t,4} = 1$	$\min\{\hat{c}_{t+4}^{(k,l)}, \hat{c}_{t+5}^{(k,l)}\}$	$z_{n,t+4}$	$f_{n,t+4}^{(k,l)}$
$r_{t,5} = 0$	$\hat{c}_{t+6}^{(k,l)}$ is unknown	N/A	N/A

with this particular plan $\mathbf{r}_t$ are displayed in Table II. The capacity constraints in time-step $t + h$ depend on whether the plan is to re-route in the next time-step $t + h + 1$ or not, according to the following expression

$$\sum_{n=1}^{N-1} z_{n,t+h} d_n f_{n,t+h}^{(k,l)} \leq \min\{\hat{c}_{t+h}^{(k,l)}, (1 - s_{min} r_{t,h+1}) \hat{c}_{t+h+1}^{(k,l)}\}, \forall (k, l) \in E. \quad (9)$$

Equation (9) is a generalization of (5) and (8). The second column of Table II represents the RHS of the capacity constraint in (9). Notice from Table II that $f_{n,t+h}^{(k,l)}$ can be updated only at the re-routing times $t + 1$ and $t + 4$ while $z_{n,t+h}$ can be updated at every time-step. Hence, the routing decisions at time $t + 1$, namely $f_{n,t+1}^{(k,l)}$, affect not only $z_{n,t+1}$, but also $z_{n,t+2}$ and $z_{n,t+3}$. It follows that the optimization problem associated with $\mathbf{r}_t = (0, 1, 0, 0, 1, 0)$ can be subdivided at the re-routing times, resulting in three simpler sub-problems, each of which jointly optimizes: (i) $f_{n,t}^{(k,l)}$ and $z_{n,t}$; (ii) $f_{n,t+1}^{(k,l)}$, $z_{n,t+1}$, $z_{n,t+2}$, and $z_{n,t+3}$; and (iii) $f_{n,t+4}^{(k,l)}$ and $z_{n,t+4}$.

In general, the optimization problem associated with any re-routing plan $\mathbf{r}_t$ can be subdivided at the re-routing times (i.e., times $t+h$ in which $r_{t,h} = 1$) *without loss of optimality*. Let $\{t+h_1, t+h_1+1, \dots, t+h_2\}$ represent a subdivision of a re-routing plan $\mathbf{r}_t$. The Generalized-MCF (G-MCF) algorithm described in Algorithm 1 jointly optimizes the routing decisions $f_{n,t+h_1}^{(k,l)}, \forall n, \forall (k, l)$ at the initial time $t + h_1$ and the admission rates $z_{n,t+h}, \forall n, \forall h \in \{h_1, \dots, h_2\}$. To address this joint optimization, G-MCF solves a sequence of MCF problems with increasing admission rates $z_{n,t+h}$ until all commodities in the network become saturated. In particular, let k be the iteration index of the algorithm, let U be the set of unsaturated commodities at the beginning of iteration k , let (n, h) be the tuple that represents the index of the $(N - 1)(h_2 - h_1 + 1)$ different commodities in this subdivision of $\mathbf{r}_t$, and let $z_{(n,h)}^S$ be the admission rate $z_{n,t+h}$ that

saturates commodity (n, h) . Initially, we have $U = \{(n, h)\}, \forall n, h$. In each iteration k , the algorithm solves the MCF problem associated with the *unsaturated commodities*, i.e., it assigns $z_{n,t+h} \leftarrow \bar{z}, \forall (n, h) \in U$, and $z_{n,t+h} \leftarrow z_{(n,h)}^S, \forall (n, h) \notin U$, and finds the feasible configuration (Output 1) with maximum value of $\bar{z} \in [0, 1]$ which we denote by $\bar{z}^*$. Then, the algorithm identifies the commodities that become saturated² in the current iteration, stores their saturation values $z_{(n,h)}^S \leftarrow \bar{z}^*$, updates the set U accordingly, and proceeds to the next iteration $k + 1$. The algorithm terminates when $U = \emptyset$ and the admission rates $z_{(n,h)}^S$ that saturate every commodity in the network have been determined. The G-MCF algorithm finds the optimal routing and admission control decisions within a subdivision $\{t + h_1, t + h_1 + 1, \dots, t + h_2\}$ of the re-routing plan $\mathbf{r}_t$. To compute the optimal routing and admission control decisions associated with an entire re-routing plan $\mathbf{r}_t = (r_{t,0}, r_{t,1}, \dots, r_{t,H})$, the G-MCF algorithm is utilized in each of its subdivisions.

MSNR algorithm. To find the optimal sequence of network configurations at time-step t , the MSNR algorithm selects the plan $\mathbf{r}_t^*$ with highest cumulative sum of admission rates $\sum_{h=0}^{H-1} \sum_{n=1}^{N-1} z_{n,t+h}$. A naive implementation of the MSNR algorithm computes and compares the performance of the (at least) 2^H admissible re-routing plans. To reduce the computational complexity from exponential $O(2^H)$ to polynomial $O(H^4)$, we propose a principled implementation of the MSNR algorithm based on backward induction which leverages the fact that the optimization problem can be subdivided at the re-routing times without loss of optimality. Specifically, the algorithm separates re-routing plans $\mathbf{r}_t$ into $H + 1$ disjoint sets and then finds the best plan within each set. The first set contains plans that re-route *for the first time* at step $t + H - 1$, the second set contains plans that re-route *for the first time* at step $t + H - 2$, and so on, until the last set which contains a plan that never re-routes. A key observation is that computations for earlier sets can be used to simplify computations for later sets. A description of this computation is provided below.

First Set. Consider plans $\mathbf{r}_t$ that re-route *for the first time* at step $t + H - 1$, i.e., $\mathbf{r}_t \in \{(0, \dots, 0, 1, 0), (0, \dots, 0, 1, 1)\}$. The MSNR algorithm employs the G-MCF algorithm to compute the optimal routing and admission control decisions for these 2 re-routing plans and then selects the plan $\mathbf{r}_t^{(1)}$ with highest cumulative sum of admission rates at time $t + H - 1$, namely

²Identifying the commodities that become saturated in iteration k is not straightforward. The authors in [50] developed a saturation test which we adapt to our network setting in lines 10 - 23 of Algorithm 1.

Algorithm 1: Generalized-MCF (G-MCF) algorithm

```

1  % Let  $\{t + h_1, \dots, t + h_2\}$  be the subdivision of the re-routing plan  $\mathbf{r}_t$  under consideration and let  $U$  be
   the set of unsaturated commodities at the beginning of iteration  $k$ ;
2  Initialization:  $U = \{(n, h)\}, \forall n \in \{1, 2, \dots, N - 1\}, \forall h \in \{h_1, \dots, h_2\}$  and  $k = 0$ ;
3  while  $U \neq \emptyset$  do
4      % Find  $\bar{z}$  that solves the joint optimization;
5      for  $n \in \{1, \dots, N - 1\}$  and  $h \in \{h_1, \dots, h_2\}$  do
6          if  $(n, h) \in U$  then  $z_{n,t+h} \leftarrow \bar{z}$ ;
7          else  $z_{n,t+h} \leftarrow z_{(n,h)}^S$ ;
8      Solve:  $\max \bar{z}$ , s.t.  $\bar{z} \in [0, 1]$ , and (1)-(4), and capacity constraints in (9) for  $h \in \{h_1, \dots, h_2\}$ ;
9      Output 1: values of  $\bar{z}^*$  and  $f_{n,t+h}^{(k,l)}$ ;
10     % Identify the new saturated commodities  $(n, h)$ ;
11     Determine the set  $D$  of disconnected commodities in the residual graph associated with Output 1;
12     Saturation Flag  $= \emptyset$ ;
13     for  $(n, h) \in D$  do
14         % Find  $\bar{z}_{(n,h)}$  that solves the joint optimization;
15         Assign:  $z_{n,t+h} \leftarrow \bar{z}_{(n,h)}$ ;
16         for  $(m, j) \in U \setminus (n, h)$  do  $z_{m,t+j} \leftarrow \bar{z}^*$ ;
17         Solve:  $\max \bar{z}_{(n,h)}$ , s.t.  $\bar{z}_{(n,h)} \in [0, 1]$ , and (1)-(4), and constraints in (9) for  $h \in \{h_1, \dots, h_2\}$ ;
18         Output 2: values of  $\bar{z}_{(n,h)}^*$  and  $f_{n,t+h}^{(k,l)}$ ;
19         if  $\bar{z}_{(n,h)}^* = \bar{z}^*$  then
20             | Saturation Flag  $\leftarrow \textit{Saturation Flag} \cup (n, h)$ ;
21     for  $(n, h) \in \textit{Saturation Flag}$  do
22         | Assign:  $z_{n,t+h}^S \leftarrow \bar{z}^*$ ;
23         | Assign:  $U \leftarrow U \setminus (n, h)$ ;
24      $k \leftarrow k + 1$ 
25 % Find the max-min fair feasible configuration;
26 for  $n \in \{1, \dots, N - 1\}$  and  $h \in \{h_1, \dots, h_2\}$  do
27     | Assign:  $z_{n,t+h} \leftarrow z_{(n,h)}^S$ ;
28 Obtain:  $f_{n,t+h}^{(k,l)}$  that satisfy (1)-(4) and capacity constraints in (9) for  $h \in \{h_1, \dots, h_2\}$ ;
29 Output 3: values of  $z_{n,t+h} = z_{(n,h)}^S$  and  $f_{n,t+h}^{(k,l)}$ ;

```

$$\sum_{n=1}^{N-1} z_{n,t+H-1}.$$

Second Set. Consider plans that re-route for the first time at step $t + H - 2$, i.e., $\mathbf{r}_t \in \{(0, \dots, 0, 1, 0, 0), (0, \dots, 0, 1, 0, 1), (0, \dots, 0, 1, 1, 0), (0, \dots, 0, 1, 1, 1)\}$. Notice that in the *subset of plans* that re-route both at times $t + H - 2$ and $t + H - 1$, we know from the First Set that $\mathbf{r}_t = \mathbf{r}_t^{(1)} + (0, \dots, 0, 1, 0, 0)$ has the best performance and, hence, all other plans in this particular subset can be excluded from consideration. The MSNR algorithm computes the optimal routing and admission control decisions for the remaining 3 re-routing plans and selects the plan $\mathbf{r}_t^{(2)}$ with highest cumulative sum of admission rates at times $t + H - 2$ and $t + H - 1$, namely $\sum_{h=H-2}^{H-1} \sum_{n=1}^{N-1} z_{n,t+h}$.

Third Set. Consider plans that re-route for the first time at step $t + H - 3$, i.e., $\mathbf{r}_t \in \{(0, \dots, 0, 1, 0, 0, 0), \dots, (0, \dots, 0, 1, 1, 1, 0), (0, \dots, 0, 1, 1, 1, 1)\}$. Notice that in the *subset of plans* that re-route both at times $t + H - 3$ and $t + H - 2$, we know from the Second Set that $\mathbf{r}_t = \mathbf{r}_t^{(2)} + (0, \dots, 0, 1, 0, 0, 0)$ has the best performance and, hence, all other plans in this particular subset can be excluded from consideration. Similarly, in the *subset of plans* that re-route both at times $t + H - 3$ and $t + H - 1$, but do not re-route at time $t + H - 2$, we know from the First Set that $\mathbf{r}_t = \mathbf{r}_t^{(1)} + (0, \dots, 0, 1, 0, 0, 0)$ has the best performance and, hence, all other plans in this particular subset can be excluded from consideration. The MSNR algorithm computes the optimal routing and admission control decisions for the remaining 4 re-routing plans and selects the plan $\mathbf{r}_t^{(3)}$ with highest cumulative sum of admission rates from times $t + H - 3$ to $t + H - 1$, namely $\sum_{h=H-3}^{H-1} \sum_{n=1}^{N-1} z_{n,t+h}$.

Subsequent Sets. The MSNR algorithm considers the set of plans that re-route for the first time at steps $t + H - 4, t + H - 5, \dots, t$ and employs an analogous procedure in order to determine the best plans $\mathbf{r}_t^{(4)}, \mathbf{r}_t^{(5)}, \dots, \mathbf{r}_t^{(H)}$.

Last Set. The MSNR algorithm compares the performance of the best plans $\mathbf{r}_t^{(h)}, \forall h \in \{1, 2, \dots, H\}$ with the performance of the never re-route plan $(0, \dots, 0, 0)$ and then selects the plan $\mathbf{r}_t^*$ with highest cumulative sum of admission rates $\sum_{h=0}^{H-1} \sum_{n=1}^{N-1} z_{n,t+h}$ in the entire prediction window. The routing and admission control decisions associated with $\mathbf{r}_t^*$ are the optimal sequence of network configurations.

Remark 3 (Computational Complexity): To find the best plans $\mathbf{r}_t^{(1)}, \mathbf{r}_t^{(2)}, \dots, \mathbf{r}_t^{(H)}$ in each of the corresponding backward induction steps, the MSNR algorithm computes and compares the performance of $2, 3, \dots, H + 1$ re-routing plans, respectively. Then, in the last step of the induction, the MSNR algorithm computes and compares the performance of $H + 2$ re-

routing plans in order to find the plan $\mathbf{r}_t^*$ and the associated optimal sequence of network configurations $\{f_{n,t+h}^{(k,l)}, z_{n,t+h}\}, \forall h \in \{0, 1, \dots, H-1\}$, at time t . In total, the MSNR algorithm employing backward induction computes the performance of $(H+1)(H+4)/2$ re-routing plans, as opposed to the (at least) 2^H computations associated with the naive implementation. Notice from Algorithm 1 that to compute the performance of any given re-routing plan $\mathbf{r}_t$, the G-MCF algorithm solves $O(H^2 N^2)$ MCF optimization problems, each of which can be solved in polynomial time [50], [51]. It follows that the MSNR algorithm has polynomial computational complexity which grows as $O(H^4)$.

Proposition 4 (Max-Min Fairness of the MSNR algorithm): The optimal sequence of network configurations $\{f_{n,t+h}^{(k,l)}, z_{n,t+h}\}$ given by the MSNR algorithm has admission rates $\{z_{n,t+h}\}_{n=1}^{N-1}$ that are max-min fair in every time-step $t+h$ for any given $h \in \{0, 1, \dots, H-1\}$, irrespective of the network topology $G = (V, E)$, demands d_n , and current and predicted link capacities $\{c_t^{(k,l)}, \hat{c}_{t+1}^{(k,l)}, \dots, \hat{c}_{t+H}^{(k,l)}\}$.

Proof: Proposition 4 holds by the design of the MSNR algorithm. In the first iteration, Algorithm 1 finds the *lowest* admission rate $\bar{z}^*$ that saturates at least one commodity (n, h) , assigns $z_{n,t+h} \leftarrow \bar{z}^*$, and removes the new saturated commodities from the set of unsaturated commodities, i.e., $U \setminus (n, h)$. Similarly, in each subsequent iteration k , Algorithm 1 finds the *lowest* admission rate $\bar{z}^*$ that saturates at least one *unsaturated* commodity $(n, h) \in U$, assigns $z_{n,t+h} \leftarrow \bar{z}^*$, and performs $U \setminus (n, h)$. The algorithm terminates when all commodities are saturated, i.e. $U = \emptyset$.

Consider one of the commodities (n, h) that became saturated during iteration k . To increase its admission rate beyond saturation $z_{n,t+h}$, we would have to reduce the admission rate of at least one other commodity (n', h) that became saturated either in iteration k or in a previous iteration³. Notice that, by the design of Algorithm 1, the saturation admission rate of commodity (n', h) is lower or equal to $z_{n,t+h}$. This means that, in each iteration k , the set of saturated admission rates $\{z_{n,t+h}\}_{(n,h) \notin U}$ is max-min fair. It follows that, upon termination, Algorithm 1 yields admission rates $\{z_{n,t+h}\}_{n=1}^{N-1}$ that are max-min fair. ■

³Notice that if we could increase the admission rate of (n, h) beyond saturation $z_{n,t+h}$ without reducing the admission rates of another saturated commodity (n', h) , then (n, h) was not saturated.

V. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the PNR framework. In particular, in Sec. V-A we evaluate the prediction accuracy of the AP mechanism and compare it with two benchmark time series prediction methods. Then, in Sec. V-B, we evaluate the performance of the MSNR algorithm and compare it with two reactive algorithms using a small network with $N = 3$ nodes, synthetically generated attenuation levels $x_t^{(k,l)}$ and synthetically generated attenuation predictions $\hat{x}_{t+h}^{(k,l)}$ with different (adjustable) prediction accuracies. The goal is to draw insight from this small and controllable setting. Finally, in Sec. V-C, we evaluate the PNR framework (with both the AP mechanism and the MSNR algorithm) using the backhaul network with $N = 13$ nodes illustrated in Fig. 1 and the attenuation measurements from the dataset.

A. Evaluation of the AP mechanism

The prediction accuracy of the AP mechanism is evaluated using the test sequences of attenuation measurements described in Sec. II. In this section, we show the results associated with *Test Seq. I* and *Test Seq. II*, both of which include a period of rain. We first assess the prediction error of a given link, then we analyze the prediction RMSE of the entire network and, finally, we assess the empirical probability of large prediction errors.

Let $e_{t,h}^{(k,l)} = x_{t+h}^{(k,l)} - \hat{x}_{t+h}^{(k,l)}$ be the h -steps-ahead prediction error associated with link (k, l) at time t . In Fig. 5(a), we compare the evolution of the attenuation measurements $x_{t+3}^{(9,13)}$ from link $(9, 13)$ with the 3-steps-ahead attenuation predictions $\hat{x}_{t+3}^{(9,13)}$ generated by the AP mechanism during an interval of 300 time-steps from *Test Seq. I*. In Fig. 5(b), we display the relative frequency distribution of the 3-steps-ahead prediction error $e_{t,3}^{(9,13)}$ from link $(9, 13)$ associated with the entire *Test Seq. I*. The results in Fig. 5 suggest that: (i) the attenuation predictions accurately track the measurements and (ii) the distribution of the prediction error $e_{t,h}^{(k,l)}$ is similar to a normal distribution with zero mean.

Weather-induced attenuation varies over time and geographical location, and also depends on link's characteristics such as frequency, polarization, and length, meaning that prediction errors may differ considerably across different links. To capture the prediction error in the entire




Fig. 5. (a) Comparison of the attenuation measurements from link (9, 13) with the corresponding 3-steps-ahead predictions. (b) Relative frequency distribution of the 3-steps-ahead prediction error.

network, we employ

$$RMSE_h^{\text{avg}} = \sqrt{\frac{1}{T-H} \sum_{t=1}^{T-H} \frac{1}{|E|} \sum_{(k,l) \in E} \left(e_{t,h}^{(k,l)} \right)^2} \quad (10)$$

$$RMSE_h^{\text{max}} = \sqrt{\frac{1}{T-H} \sum_{t=1}^{T-H} \max_{(k,l) \in E} \left\{ \left(e_{t,h}^{(k,l)} \right)^2 \right\}} \quad (11)$$

which calculate the RMSE associated with the h -steps-ahead prediction errors of all links over the entire time-horizon and the RMSE associated with the *largest* h -steps-ahead prediction error among all the links in each time-step t , respectively. In Fig. 6, we display the $RMSE_h^{\text{avg}}$ and $RMSE_h^{\text{max}}$ (in dB) as a function of the prediction horizon $h \in \{1, \dots, H\}$ for *Test Seq. I* and *II* for three prediction mechanisms: (i) the AP mechanism; (ii) the naive AP method, also called random-walk method, which is a commonly used benchmark [56] that employs the latest measurement as future predictions, namely $\hat{x}_{t+h}^{(k,l)}(\text{naive}) = x_t^{(k,l)}, \forall h$; and (iii) the ARIMA model, which is a well-known time series prediction model. For an example of the ARIMA model being employed to predict rain-induced attenuation in Ku-band satellite links, we refer the reader to [39]. *It is important to emphasize that both benchmark methods (i.e., naive and ARIMA) consider each link in isolation when predicting future attenuation levels and, thus, they do not capture the spatial correlation that is typical of weather-induced attenuation.* The results in Fig. 6 suggest that the AP mechanism outperforms the benchmark methods in both *Test Seqs. I* and *II* and that this performance improvement increases as the prediction horizon h increases. In particular, when $h = 1$, the performance improvement (in terms of $RMSE_h^{\text{avg}}$) of employing the AP mechanism as opposed to any of the benchmark methods is between 0% and 12%, and when $h = 5$, the performance improvement is between 12% and 34%.




Fig. 6. $RMSE_h^{\text{avg}}$ and $RMSE_h^{\text{max}}$ of the prediction error for different prediction horizons h and for the AP mechanism, naive AP method, and ARIMA model.

To analyze the empirical probability of large prediction errors, we compute the percentile associated with the modulus of the h -steps-ahead predictions errors $|e_{t,h}^{(k,l)}|$. In particular, for a given test sequence with attenuation measurements $x_{t+h}^{(k,l)}$ and associated predictions $\hat{x}_{t+h}^{(k,l)}$ from the AP mechanism, the η th percentile value represents the lowest $|e_{t,h}^{(k,l)}|$ that is larger than or equal to $\eta\%$ of all the values of $|e_{t,h}^{(k,l)}|$ in the considered dataset. For example, if the 95th percentile value for *Test Seq. I* and $h = 3$ is 1 dB, it means that 95% of all the values of $|e_{t,h}^{(k,l)}|$ computed for the entire *Test Seq. I* are lower than or equal to 1 dB. In Fig. 7, we show the percentile values for different prediction horizons $h \in \{1, \dots, 5\}$ for *Test Seq. I* and *II*. The results in Fig. 7 suggest that, as expected, the percentile values increase with the prediction horizon h . Moreover, the results show that 95% of the one-step-ahead and 5-steps-ahead prediction errors are lower than 0.5 dB and 1.5 dB, respectively.

In summary, the results in Figs. 5, 6, and 7 show that *the AP mechanism predicts future link attenuation with high accuracy*. Next, we show that prediction accuracy has a significant impact on the performance of the MSNR algorithm.

B. Evaluation of reconfiguration algorithms in a small and controllable network

We compare the performance of the MSNR algorithm with two reactive network reconfiguration algorithms, namely NEVER RE-ROUTE and ALWAYS RE-ROUTE, in terms of their network utilization, which is captured by the evolution of the node-average admission rate $\sum_{n=1}^{N-1} z_{n,t}/(N-1)$ over time t . The considered network reconfiguration algorithms are:

- i) *MSNR algorithm*: leverages future predicted capacities to decide when to re-route. In particular, in each time-step t , it compares the performance of different re-routing plans




Fig. 7. Percentile for the modulus of the prediction error.

and selects the plan $\mathbf{r}_t^*$ with highest cumulative sum of admission rates, as described in Sec. IV-C.

- ii) *NEVER RE-ROUTE algorithm*: attempts to maximize the admission rates $z_{n,t}$ by never provisioning scratch capacity and, thus, fully utilizing links whenever possible. This reactive algorithm operates based on the MSNR algorithm. However, instead of selecting $\mathbf{r}_t^*$, it selects, in *every* time-step t , the re-routing plan $\mathbf{r}_t = (0, \dots, 0)$. Under this algorithm, the SDN controller is rarely⁴ allowed to re-route, but it is continually optimizing the admission rates.
- iii) *ALWAYS RE-ROUTE algorithm*: attempts to provision scratch capacity $s_t \geq s_{min} = 0.05$ at every time-step t , enabling the SDN controller to optimize routing decisions often. This reactive algorithm operates based on the MSNR algorithm. However, instead of selecting $\mathbf{r}_t^*$, it selects, in *every* time-step t with $s_t \geq s_{min}$, the re-routing plan $\mathbf{r}_t = (1, 1, \dots, 1)$, and in *every* time-step t with $s_t < s_{min}$, the re-routing plan $\mathbf{r}_t = (0, 1, \dots, 1)$.

Notice that all three network reconfiguration algorithms select max-min fair admission rates $z_{n,t}$ at *every time-step* t . The main difference between them is that only the MSNR algorithm employs the predictions of the links' future condition to decide when to re-route. Both the NEVER RE-ROUTE and ALWAYS RE-ROUTE algorithms simply react to the time-varying conditions of the network. The comparison with the predictive SDN-based routing framework developed in [36] is not possible due to the incompatible assumptions. Recall that the framework in [36] can *only* be employed during periods of rain, it allows flows to temporarily exceed the link

⁴Notice that if the predicted capacities $\hat{c}_{t+1}^{(k,l)}$ are inaccurate, in particular if $\hat{c}_{t+1}^{(k,l)} < c_{t+1}^{(k,l)}$, it may happen that $s_{t+1} \geq s_{min}$ and the NEVER RE-ROUTE algorithm is allowed to re-route at time $t+1$.

capacity, and it does not take fairness into account, thus, making the comparison unfit. Next, we evaluate the three reconfiguration algorithms in a small and controllable setting using synthetic data. In Sec. V-C, we evaluate the same three algorithms using measurements collected from the backhaul network.

The results in this section are associated with the network in Fig. 4 with $N = 3$ nodes and three links $\{(1, 2), (2, 3), (1, 3)\}$. The *normalized*⁵ demands associated with nodes 1 and 2 remain fixed at $d_1 = 1$ and $d_2 = 0.5$, respectively, during the time-horizon of 1,000 time-steps. The (actual) attenuation levels $x_t^{(k,l)}$ and predicted attenuation levels $\hat{x}_{t+h}^{(k,l)}$ are synthetically generated according to the following stochastic processes

$$x_t^{(k,l)} = \min\{\max\{x_{t-1}^{(k,l)} + \delta_t^{(k,l)}; -100\}; -50\}; \quad (12)$$

$$\hat{x}_{t+h}^{(k,l)} = \min\{\max\{x_{t+h}^{(k,l)} + \tilde{\delta}_{t,h}^{(k,l)}; -100\}; -50\}, \quad (13)$$

for all links $(k, l) \in E$, for all time-steps $t \in \{1, \dots, 1,000\}$, for all values of $h \in \{1, \dots, H\}$, and with $x_0^{(k,l)}$ sampled from a uniform distribution in the interval $(-100, -50)$. Notice that (12) establishes the variation of the attenuation $x_t^{(k,l)}$ over time, while (13) establishes the noise in the prediction $\hat{x}_{t+h}^{(k,l)}$ of the future attenuation $x_{t+h}^{(k,l)}$. The sequence of Gaussian random variables $\delta_t^{(k,l)}$ is i.i.d. over time t , independent across links, and sampled according to $\mathcal{N}(0, 6.25)$. Similarly, the sequence of random variables $\tilde{\delta}_{t,h}^{(k,l)}$ are Gaussian $\mathcal{N}(0, \tilde{\sigma}^2)$ with positive variance $\tilde{\sigma}^2$, i.i.d. over time, and independent across different links. Notice from (13) that, a high variance $\tilde{\sigma}^2$ represents an AP mechanism with poor accuracy, i.e. large prediction error. The choice of Gaussian distribution for $\tilde{\delta}_{t,h}^{(k,l)}$ was inspired by the relative frequency distribution of the prediction error shown in Fig. 5(b).

To determine the (actual) capacities $c_t^{(k,l)}$ and the predicted capacities $\hat{c}_{t+h}^{(k,l)}$ associated with the synthetic values of $x_t^{(k,l)}$ and $\hat{x}_{t+h}^{(k,l)}$, respectively, we adopt a constant transmission signal level of $P_{Tx,t}^{(k,l)} = 0$ dBm and use the AM mechanism described in Sec. IV-A. In Fig. 8(a), we display the evolution of the *normalized* values of $c_t^{(k,l)}$ employed to obtain the results in this section. Notice that this is a network with highly dynamic link capacities $c_t^{(k,l)}$.

In Fig. 8(b), we compare the evolution of the node-average admission rate $(z_{1,t} + z_{2,t})/2$ over time t for different reconfiguration algorithms operating with ideal attenuation predictions, i.e., with $\hat{x}_{t+h}^{(k,l)} = x_{t+h}^{(k,l)}$ and, as a result, $\hat{c}_{t+h}^{(k,l)} = c_{t+h}^{(k,l)}$. In Fig. 8(c) and in Table III, we show the time-

⁵Both demands and capacities are normalized with respect to the maximum achievable bitrate of 225 Mbps from Table I.




Fig. 8. Performance of the MSNR algorithm for the network in Fig. 4 with $N = 3$ nodes. (a) Evolution of the normalized link capacity $c_t^{(k,l)}$ over time. (b) Node-average admission rate $(z_{1,t} + z_{2,t})/2$ for different reconfiguration algorithms with ideal attenuation prediction ($\tilde{\sigma}^2 = 0$). (c) Time-average admission rate $\sum_{t=1}^T (z_{1,t} + z_{2,t})/2T$ for MSNR with different prediction window sizes $H \in \{2, 3, 4, 5\}$ and attenuation prediction accuracies $\tilde{\sigma}^2 \in \{0, 0.0025, \dots, 9\}$.

average admission rates $\sum_{t=1}^T (z_{1,t} + z_{2,t})/2T$ for different reconfiguration algorithms operating with attenuation predictions with different accuracies $\tilde{\sigma}^2 \in \{0, 0.0025, 0.25, 1, 4, 9, 25\}$ and different prediction window sizes $H \in \{2, 3, 4, 5\}$.

The results in Fig. 8(b) show that, as expected, NEVER RE-ROUTE has the worse performance, while MSNR with prediction window size $H = 5$ has the best performance in terms of network utilization. The poor performance of NEVER RE-ROUTE, especially between time-steps 500 and 800, results from the SDN controller not being allowed to re-route. The lower performance of ALWAYS RE-ROUTE when compared to MSNR is due to the frequent provisioning of scratch capacity $s_{min} = 0.05$. By leveraging the prediction of links' future conditions, MSNR can assess the potential future benefits of re-routing⁶, which allows it to choose when is the best time to re-route. Throughout the 1,000 time-steps, the SDN controller re-routes 31, 30, 28, and 29 times when employing MSNR with prediction window sizes $H \in \{2, 3, 4, 5\}$, respectively.

The results in Fig 8(c) and Table III suggest that: (i) the performance of MSNR improves as the prediction accuracy improves and as the window size H increases and (ii) the performance gain of improving the prediction accuracy is more significant than the performance gain of increasing the prediction window size H , which *highlights the importance of developing an accurate AP mechanism*.

⁶Recall from the discussion in Sec. IV-B that planning to re-route at the next time-step $t + 1$, can only hurt the network performance at the current time t due to the provision of the scratch capacity.




Fig. 9. Performance of the PNR framework using data collected from the backhaul network in Fig. 1. (a) Evolution of the normalized measured link capacity $c_t^{(k,l)}$ over time. (b)-(c): Evolution of the node-average admission rate $\sum_{n=1}^{N-1} z_{n,t}/(N-1)$ over time for different network reconfiguration algorithms.

TABLE III

TIME-AVERAGE ADMISSION RATES $\sum_{t=1}^T (z_{1,t} + z_{2,t})/2T$ FOR DIFFERENT NETWORK RECONFIGURATION ALGORITHMS AND FOR ATTENUATION PREDICTIONS WITH DIFFERENT ACCURACIES.

Prediction accuracy	Ideal	$\bar{\sigma}^2 = 1$	$\bar{\sigma}^2 = 25$
NEVER RE-ROUTE	0.362	0.359	0.343
ALWAYS RE-ROUTE	0.453	0.449	0.427
MSNR for $H = 2$	0.477	0.473	0.447
MSNR for $H = 3$	0.478	0.474	0.449
MSNR for $H = 4$	0.479	0.476	0.449
MSNR for $H = 5$	0.479	0.476	0.450

C. Evaluation of the PNR Framework with data from a real-world network

We now evaluate the performance of the PNR framework using the data collected from the backhaul network in Fig. 1 with $N = 13$ base-stations (12 commodities and one destination) and 17 links. The normalized demands assigned to the commodities are chosen according to a uniform distribution in the interval $(0, 2)$. In particular, the twelve demand values⁷ are $\mathbf{d} = [1.111, 0.557, 1.124, 1.266, 0.174, 1.485, 0.947, 0.067, 0.140, 0.596, 1.413, 0.999]$. The values of the (actual) capacities $c_t^{(k,l)}$ and future predicted capacities $\hat{c}_{t+h}^{(k,l)}$ are determined by the link attenuation measurements in the dataset, by the AM mechanism described in Sec. IV-A, and by the AP mechanism. To train, tune, and test the AP mechanism, we use a train-validation-test split of 80-10-10. To assess the performance of the PNR framework in a *challenging scenario*, we choose a sequence of more than 400 measurements (from *Test Seq. I* described in Sec. II) that includes a

⁷Notice that similar results can be obtained for different vectors of demands.

TABLE IV

PERFORMANCE GAIN OF MSNR WITH $H \in \{2, 5\}$ WHEN COMPARED TO A REACTIVE ALGORITHM: NEVER RE-ROUTE OR ALWAYS RE-ROUTE. THE PERFORMANCE GAIN METRICS IN COLUMNS 3, 4, AND 5 ARE DEFINED IN (14).

MSNR	Reactive	Time-average	Node-average	Instantaneous
$H = 5$	ALWAYS	7.74%	18.00%	170.19%
$H = 2$	ALWAYS	15.49%	26.84%	263.58%
$H = 5$	NEVER	1.67%	10.37%	68.37%
$H = 2$	NEVER	8.98%	22.04%	208.01%

period with high attenuation variability due to a rain event. Moreover, we consider transmission signal levels $P_{Tx,t}^{(k,l)}$ that are 10 dBm lower than the dataset measurements. In Fig. 9(a), we display the evolution of the normalized capacities $c_t^{(k,l)}$ from three selected links. Notice that the variation is significant. In Fig. 9(b), we show the evolution of the node-average admission rates $\sum_{n=1}^{N-1} z_{n,t}/(N-1)$ for different reconfiguration algorithms employing ideal attenuation prediction, i.e., $\hat{c}_{t+h}^{(k,l)} = c_{t+h}^{(k,l)}$. In Fig. 9(c), we display the node-average admission rates for algorithms employing the AP mechanism to predict $\hat{c}_{t+h}^{(k,l)}$ over time. The results in Figs. 9(b) and 9(c) show that MSNR outperforms both NEVER RE-ROUTE and ALWAYS RE-ROUTE.

In Table IV, we display the performance gain of MSNR with $H \in \{2, 5\}$ employing the AP mechanism when compared to reactive algorithms: NEVER RE-ROUTE or ALWAYS RE-ROUTE. Let $z_{n,t}^{(M)}$ and $z_{n,t}^{(R)}$ be the admission rates associated with MSNR and the reactive algorithm, respectively. The third, fourth, and fifth columns of Table IV are associated with

$$\frac{\sum_{t=1}^T \sum_{n=1}^{N-1} (z_{n,t}^{(M)} - z_{n,t}^{(R)})}{\sum_{t=1}^T \sum_{n=1}^{N-1} z_{n,t}^{(R)}}, \quad \max_t \left\{ \frac{\sum_{n=1}^{N-1} (z_{n,t}^{(M)} - z_{n,t}^{(R)})}{\sum_{n=1}^{N-1} z_{n,t}^{(R)}} \right\}, \quad \max_{n,t} \left\{ \frac{z_{n,t}^{(M)} - z_{n,t}^{(R)}}{z_{n,t}^{(R)}} \right\}; \quad (14)$$

which represent the *time-average performance gain*, the *maximum node-average performance gain*, and the *maximum instantaneous performance gain*, respectively. The results in Table IV show that the MSNR algorithm can improve the time-average admission rate $\sum_{t=1}^T (z_{1,t} + z_{2,t})/2T$ by more than 7% when compared to either ALWAYS RE-ROUTE or NEVER RE-ROUTE and, more importantly, they also show that the gain in terms of the instantaneous per commodity admission rate $z_{n,t}$ can exceed 200%. These significant instantaneous gains occur when severe rain-induced attenuation occurs, showing that the PNR framework is able to prepare the network ahead of time and alleviate the impact of these severe disturbances on the network performance, which can be paramount to time-sensitive applications.

An important observation from the results in Secs. V-B and V-C is that, when the AP mechanism has high accuracy, the performance gap between MSNR with $H = 2$ and reactive algorithms is significantly larger than the performance gain obtained from increasing the prediction window size H . Adding to this observation the fact that the computational complexity of MSNR grows with H , as discussed in Sec. IV-C, makes the PNR framework with $H = 2$ an attractive choice both in terms of performance and complexity.

VI. CONCLUSION

We developed the PNR framework, that includes: (i) the AP mechanism that uses historical data to predict the sequence of future attenuation levels, *without incorporating any specific weather-related models*; and (ii) the MSNR algorithm that dynamically optimize routing $f_{n,t}^{(k,l)}$ and admission control $z_{n,t}$ decisions over time aiming to maximize the cumulative sum of admission rates $\sum_{t=1}^T \sum_{n=1}^{N-1} z_{n,t}$, while ensuring that, in every time-step t , the selected feasible set $\{f_{n,t}^{(k,l)}, z_{n,t}\}$ is max-min fair in every time-step t and can be implemented without inducing transient congestion. We use a real-world dataset to thoroughly evaluate the PNR framework and to show that it allows the SDN controllers to prepare the x-haul for imminent (and possibly severe) weather-induced disturbances. There are several open problems that will be considered in our future work, including consideration of time-varying traffic demands d_n , consideration of downlink/uplink traffic, application to 5G slice admission and provisioning, and experimental evaluation in city-scale testbeds.

VII. ACKNOWLEDGEMENT

This work was supported in part by the NSF-BSF grant CNS-1910757 and the NSF grant OAC-2029295.

REFERENCES

- [1] J. Edstam, J. Hansryd, S. Carpenter, T. Emanuelsson, Y. Li, and H. Zirath, "Microwave backhaul revolution - reaching beyond 100 GHz," in *Ericsson Technology Review*, 2017.
- [2] ITU-R P.530, "Propagation data and prediction methods required for the design of terrestrial line-of-sight systems," *ITU*, vol. 530-17, 2017.
- [3] ITU-R P.840, "Attenuation due to clouds and fog," *ITU*, vol. 840-7, 2017.
- [4] ITU-R P.838, "Specific attenuation model for rain for use in prediction methods," *ITU*, vol. 838-3, 1992-1999-2003-2005, 2005.
- [5] ITU-R P.676, "Attenuation by atmospheric gasses," *ITU*, vol. 676-11, 2016.

- [6] S. T. Chung and A. J. Goldsmith, "Degrees of freedom in adaptive modulation: A unified view," *IEEE Transactions on Communications*, vol. 49, no. 9, pp. 1561–1571, 2001.
- [7] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe, "Design and implementation of a routing control platform," in *Proc. NSDI*, 2005.
- [8] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, 2008.
- [9] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: Experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM*, 2013.
- [10] NEC iPASOLINK Series, "Intelligent and high capacity wireless transport solution," 2016. [Online]. Available: <https://pdf.aeroexpo.online/pdf/nec-corporation/intelligent-high-capacity-wireless-transport-solution/171108-4848.html>
- [11] N. Bui, M. Cesana, S. A. Hosseini, Q. Liao, I. Malanchini, and J. Widmer, "A survey of anticipatory mobile networking: Context-based classification, prediction methodologies, and optimization techniques," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 3, pp. 1790–1821, 2017.
- [12] S. Naimi, A. Busson, V. Vèque, L. B. H. Slama, and R. Bouallegue, "Anticipation of ETX metric to manage mobility in ad hoc wireless networks," in *Proc. Ad-hoc, Mobile, and Wireless Networks*, 2014.
- [13] Z. Lu and G. de Veciana, "Optimizing stored video delivery for mobile networks: The value of knowing the future," in *Proc. IEEE INFOCOM*, 2013.
- [14] R. Margolies, A. Sridharan, V. Aggarwal, R. Jana, N. K. Shankaranarayanan, V. Vaishampayan, and G. Zussman, "Exploiting mobility in proportional fair cellular scheduling: Measurements and algorithms," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 355–367, 2016.
- [15] A. Chattopadhyay, B. Blaszczyszyn, and E. Altman, "Location aware opportunistic bandwidth sharing between static and mobile users with stochastic learning in cellular networks," *IEEE Transactions on Mobile Computing*, pp. 1802–1815, 2018.
- [16] Z. Yang, P. H. Pathak, Y. Zeng, and P. Mohapatra, "Sensor-assisted codebook-based beamforming for mobility management in 60 GHz WLANs," in *Proc. IEEE MASS*, 2015.
- [17] A. W. Doff, K. Chandra, and R. V. Prasad, "Sensor assisted movement identification and prediction for beamformed 60 GHz links," in *Proc. IEEE CCNC*, 2015.
- [18] L. Huang, S. Zhang, M. Chen, and X. Liu, "When backpressure meets predictive scheduling," in *Proc. ACM MobiHoc*, 2014.
- [19] N. Abedini and S. Shakkottai, "Content caching and scheduling in wireless networks with elastic and inelastic traffic," *IEEE/ACM Transactions on Networking*, vol. 22, no. 3, pp. 864–874, 2014.
- [20] A. Balachandran, V. Sekar, A. Akella, S. Seshan, I. Stoica, and H. Zhang, "Developing a predictive model of quality of experience for internet video," in *Proc. ACM SIGCOMM*, 2013.
- [21] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot, "Long-term forecasting of internet backbone traffic: Observations and initial models," in *Proc. IEEE INFOCOM*, 2003.
- [22] Y. Sun, X. Yin, J. Jiang, V. Sekar, F. Lin, N. Wang, T. Liu, and B. Sinopoli, "CS2P: Improving video bitrate selection and adaptation with data-driven throughput prediction," in *Proc. ACM SIGCOMM*, 2016.
- [23] X. Wang, Z. Zhou, F. Xiao, K. Xing, Z. Yang, Y. Liu, and C. Peng, "Spatio-temporal analysis and prediction of cellular traffic in metropolis," *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2190–2202, 2018.
- [24] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "Deepcog: Optimizing resource provisioning in network slicing with ai-based capacity forecasting," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 2, pp. 361–376, 2019.

- [25] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang, "A first look at cellular network performance during crowded events," *SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 1, p. 17–28, 2013.
- [26] X. Tie, A. Seetharam, A. Venkataramani, D. Ganesan, and D. L. Goeckel, "Anticipatory wireless bitrate control for blocks," in *Proc. ACM CoNEXT*, 2011.
- [27] S. Yin, D. Chen, Q. Zhang, and S. Li, "Prediction-based throughput optimization for dynamic spectrum access," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 3, pp. 1284–1289, 2011.
- [28] N. Bui and J. Widmer, "Mobile network resource optimization under imperfect prediction," in *Proc. IEEE WoWMoM*, 2015.
- [29] T. Mangla, N. Theera-Ampornpunt, M. Ammar, E. Zegura, and S. Bagchi, "Video through a crystal ball: Effect of bandwidth prediction quality on adaptive streaming in mobile environments," in *Proc. ACM MoVid*, 2016.
- [30] S. J. Tarsa, M. Comiter, M. B. Crouse, B. McDanel, and H. Kung, "Taming wireless fluctuations by predictive queuing using a sparse-coding link-state model," in *Proc. ACM MobiHoc*, 2015.
- [31] C. Yue, R. Jin, K. Suh, Y. Qin, B. Wang, and W. Wei, "Linkforecast: Cellular link bandwidth prediction in LTE networks," *IEEE Transactions on Mobile Computing*, vol. 17, no. 7, pp. 1582–1594, 2018.
- [32] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling robust 60 GHz links under blockage," in *Proc. USENIX NSDI*, 2016.
- [33] X. Xu, I. Broustis, Z. Ge, R. Govindan, A. Mahimkar, N. K. Shankaranarayanan, and J. Wang, "Magus: Minimizing cellular service disruption during network upgrades," in *Proc. ACM CoNEXT*, 2015.
- [34] A. Jabbar, J. P. Rohrer, A. Oberthaler, E. K. Cetinkaya, V. Frost, and J. P. Sterbenz, "Performance comparison of weather disruption-tolerant cross-layer routing algorithms," in *Proc. IEEE INFOCOM*, 2009.
- [35] J. Rak, "A new approach to design of weather disruption-tolerant wireless mesh networks," *Telecommunication Systems*, vol. 61, no. 2, pp. 311–323, 2016.
- [36] F. Yaghoubi, M. Furdek, A. Rostami, P. Öhlén, and L. Wosinska, "Consistency-aware weather disruption-tolerant routing in SDN-based wireless mesh networks," *IEEE Transactions on Network and Service Management*, vol. 15, no. 2, pp. 582–595, 2018.
- [37] D. Jacoby, J. Ostrometzky, and H. Messer, "Short-term prediction of the attenuation in a commercial microwave link using LSTM-based RNN," in *Proc. EUSIPCO*, 2020.
- [38] —, "Adaptive fuzzy-based models for attenuation time series forecasting," in *Proc. COMCAS*, 2021.
- [39] D. P. Patel, M. M. Patel, and D. R. Patel, "Implementation of ARIMA model to predict rain attenuation for Ku-band 12 GHz frequency," *IOSR Journal of Electronics and Communication Engineering*, vol. 9, no. 1, pp. 83–87, 2014.
- [40] J. Sosa, C. Sosa, and B. Paz, "ARIMA models in the rain attenuation prediction in a mexican tropical area," in *Proc. IEEE Antennas and Propagation Society International Symposium.*, 2000.
- [41] N. Javed, E. Lyons, M. Zink, and T. Wolf, "Adaptive wireless mesh networks: Surviving weather without sensing it," in *Proc. IEEE ICCCN*, 2013.
- [42] J. Maciejowski, *Predictive Control With Constraints*. Prentice Hall, 2002.
- [43] J. Ostrometzky, G. Zussman, H. Messer-Yaron, D. Jacoby, and I. Kadota, "Predictive weather-aware communication network management," Patent pending, 2021.
- [44] L. Bao, C. Larsson, M. Mustafa, J. Selin, J. Andersson, J. Hansryd, M. Riedel, and H. Andersson, "A brief description on measurement data from an operational microwave network in Gothenburg, Sweden," in *Proc. CEST*, 2017.
- [45] U. Ehsan, B. Harrison, L. Chan, and M. O. Riedl, "Rationalization: A neural machine translation approach to generating natural language explanations," in *Proc. AAAI/ACM Conference on AI, Ethics, and Society*, 2018.

- [46] H. Xie, S. Fang, Z.-J. Zha, Y. Yang, Y. Li, and Y. Zhang, "Convolutional attention networks for scene text recognition," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 15, no. 1s, pp. 1–17, 2019.
- [47] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proc. ICML*, 2014.
- [48] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [49] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [50] M. Allalouf and Y. Shavitt, "Centralized and distributed algorithms for routing and weighted max-min fair bandwidth allocation," *IEEE Transactions on Network and Service Management*, vol. 16, no. 5, pp. 1015–1024, 2008.
- [51] F. Shahrokhi and D. W. Matula, "The maximum concurrent flow problem," *Journal of the ACM*, vol. 37, no. 2, pp. 318–334, 1990.
- [52] L. Bao, J. Hansryd, T. Danielson, G. Sandin, and U. Noser, "Field trial on adaptive modulation of microwave communication link at 6.8 GHz," in *Proc. IEEE EuCAP*, 2015.
- [53] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," in *Proc. ACM SIGCOMM*, 2013.
- [54] K. Cho and D. Yoon, "On the general BER expression of one- and two-dimensional amplitude modulations," *IEEE Transactions on Communications*, vol. 50, no. 7, pp. 1074–1080, 2002.
- [55] H. H. Liu, X. Wu, M. Zhang, L. Yuan, R. Wattenhofer, and D. Maltz, "zUpdate: Updating data center networks with zero loss," in *Proc. ACM SIGCOMM*, 2013.
- [56] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International Journal of Forecasting*, vol. 22, no. 4, pp. 679–688, 2006.