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

Igor Kadota*
Columbia University
New York, USA
igor.kadota@columbia.edu

Dror Jacoby*
Tel Aviv University
Tel Aviv, Israel
drorjacoby@mail.tau.ac.il

Hagit Messer
Tel Aviv University
Tel Aviv, Israel
messer@eng.tau.ac.il

Gil Zussman
Columbia University
New York, USA
gil.zussman@columbia.edu

Jonatan Ostrometzky
Tel Aviv University
Tel Aviv, Israel
jonatano@tauex.tau.ac.il

ABSTRACT

4G, 5G, and smart city networks often rely on microwave and millimeter-wave x-haul links. 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 nodes using the network (e.g., base-stations) and preventing transient congestion that may be caused by switching routes. 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. The full paper associated with this abstract can be found at https://doi.org/10.1145/3570616.

ACM Reference Format:
Igor Kadota, Dror Jacoby, Hagit Messer, Gil Zussman, and Jonatan Ostrometzky. 2023. Switching in the Rain: Predictive Wireless x-haul Network Reconfiguration. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '23 Abstracts), June 19–23, 2023, Orlando, FL, USA*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3578338.3593574

*Both authors contributed equally to this research.

SIGMETRICS '23 Abstracts, June 19–23, 2023, Orlando, FL, USA

ACM ISBN 979-8-4007-0074-3/23/06.
https://doi.org/10.1145/3578338.3593574

1 MOTIVATION

Microwave and millimeter-wave (mmWave) fronthaul, midhaul, and backhaul (x-haul) networks can connect a large number of Base-Stations (BSs), covering entire cities, as depicted in Fig. 1. However, a major challenge is the high susceptibility of microwave and mmWave links to weather conditions [2, 3]. In particular, precipitation may cause severe signal attenuation, which can reduce the capacity of the wireless links, leading to a sharp degradation of the network performance. Fig. 2(a) shows attenuation measurements from five links of the network in Fig. 1 collected during a rainy period on 2015-06-02. The impact of the attenuation on the capacity of the wireless links is displayed in Fig. 2(b). It is evident from Fig. 2 that weather-induced attenuation and the resulting capacity degradation vary over time, geographical location, rain intensity, and can be severe. *The need for a high capacity wireless x-haul that is robust to variations in the links' conditions calls for the development of a predictive network reconfiguration framework that can dynamically allocate resources based on current and future estimated 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. However, with the emergence of Software-Defined Networking (SDN), it is now possible to develop global Network layer mechanisms that monitor the entire network and react to weather-induced disturbances. An important drawback of *reactive reconfiguration* mechanisms is their delay in mitigating performance issues, which may severely affect time-sensitive applications. To overcome this challenge, *predictive reconfiguration* mechanisms should be developed.

2 CONTRIBUTIONS

In our full paper [4], we develop and evaluate (based on a real dataset) a *Predictive Network Reconfiguration (PNR) framework*, illustrated in Fig 3, that prepares the network ahead of time for

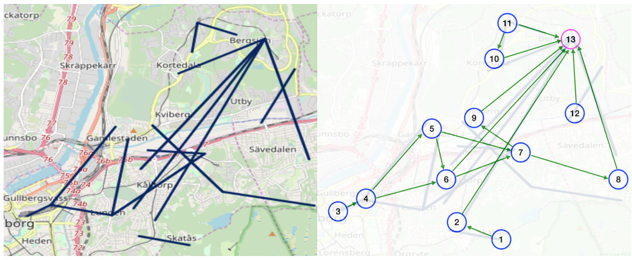

(a) (b)

Figure 1: (a) A wireless backhaul network in Gothenborg, Sweden (the map area is of approximately 10x10 km^2). The data utilized in this paper was collected from this network by *Ericsson AB*. (b) An abstraction of the network topology.

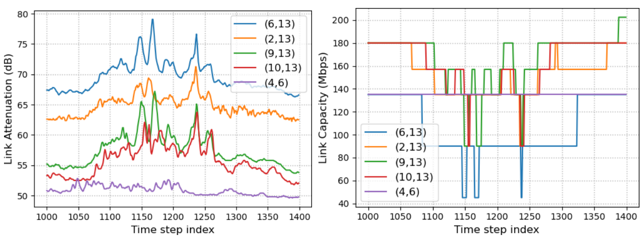

(a) (b)

Figure 2: (a) Measured attenuation over time for five links from the network in Fig. 1. Time-steps are separated by 10 seconds. Between time-steps $1,100$ and $1,300$ there was an increased attenuation due to rain. The spatio-temporal correlation is evident. (b) Link capacities over time.

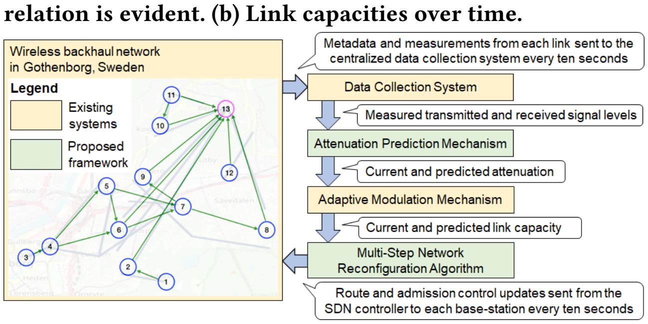

Figure 3: Overview of the main components and information flow in the proposed PNR framework.

imminent disturbances, significantly enhancing the network's robustness to variations in the links' conditions. The PNR leverages existing local Physical/Link layer mechanisms and adds 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*. This 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 dataset containing 2,295,000 measurements of link attenuation levels from the real-world city-scale backhaul network in Gothenborg, Sweden* (see Fig. 1(a)), *collected by Ericsson AB. The AP mechanism leverages the spatio-temporal correlation of the weather-effects to achieve high attenuation prediction accuracy.*

The MSNR algorithm leverages the predictions from the AP mechanism to prepare the network for future disturbances. Specifically, it generalizes the Maximum Concurrent Flow [1] problem and uses Model Predictive Control to compute the sequence of current and future routing and admission control decisions that: (i) maximize network utilization, while (ii) guaranteeing max-min fairness among the BSs sharing the network, and (iii) preventing transient congestion that may be caused by switching routes. These 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 decrease traffic admission from low priority services. An important challenge associated with the MSNR algorithm is computational complexity. We present 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. Specifically, we show that the PNR framework can achieve high attenuation prediction accuracy with a 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 PNR framework in terms of RMSE. In addition, we 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 [5].

ACKNOWLEDGMENTS

We thank our shepherd Prof. Suzan Bayhan and the anonymous referees for their helpful comments. This research was supported in part by NSF-BSF grant CNS-1910757 and NSF grants CNS-2148128, EEC-2133516.

REFERENCES

[1] M. Allalouf and Y. Shavitt. 2008. Centralized and distributed algorithms for routing and weighted max-min fair bandwidth allocation. *IEEE Trans. Netw. Service Manag.* 16, 5 (2008), 1015–1024.
[2] ITU-R P.530. 2017. Propagation data and prediction methods required for the design of terrestrial line-of-sight systems. *ITU* 530-17 (2017).
[3] ITU-R P.838. 2005. Specific attenuation model for rain for use in prediction methods. *ITU* 838-3, 1992-1999-2003-2005 (2005).
[4] I. Kadota, D. Jacoby, H. Messer, G. Zussman, and J. Ostrometzky. 2022. Switching in the rain: predictive wireless x-haul network reconfiguration. *Proc. ACM Meas. Anal. Comput. Syst.* 6, 3 (2022), 1–26.
[5] J. Ostrometzky, G. Zussman, H. Messer, D. Jacoby, and I. Kadota. US Patent Application No. 17/551,643. December, 2021. Predictive weather-aware communication network management. https://patents.google.com/patent/US20220110012A1/en